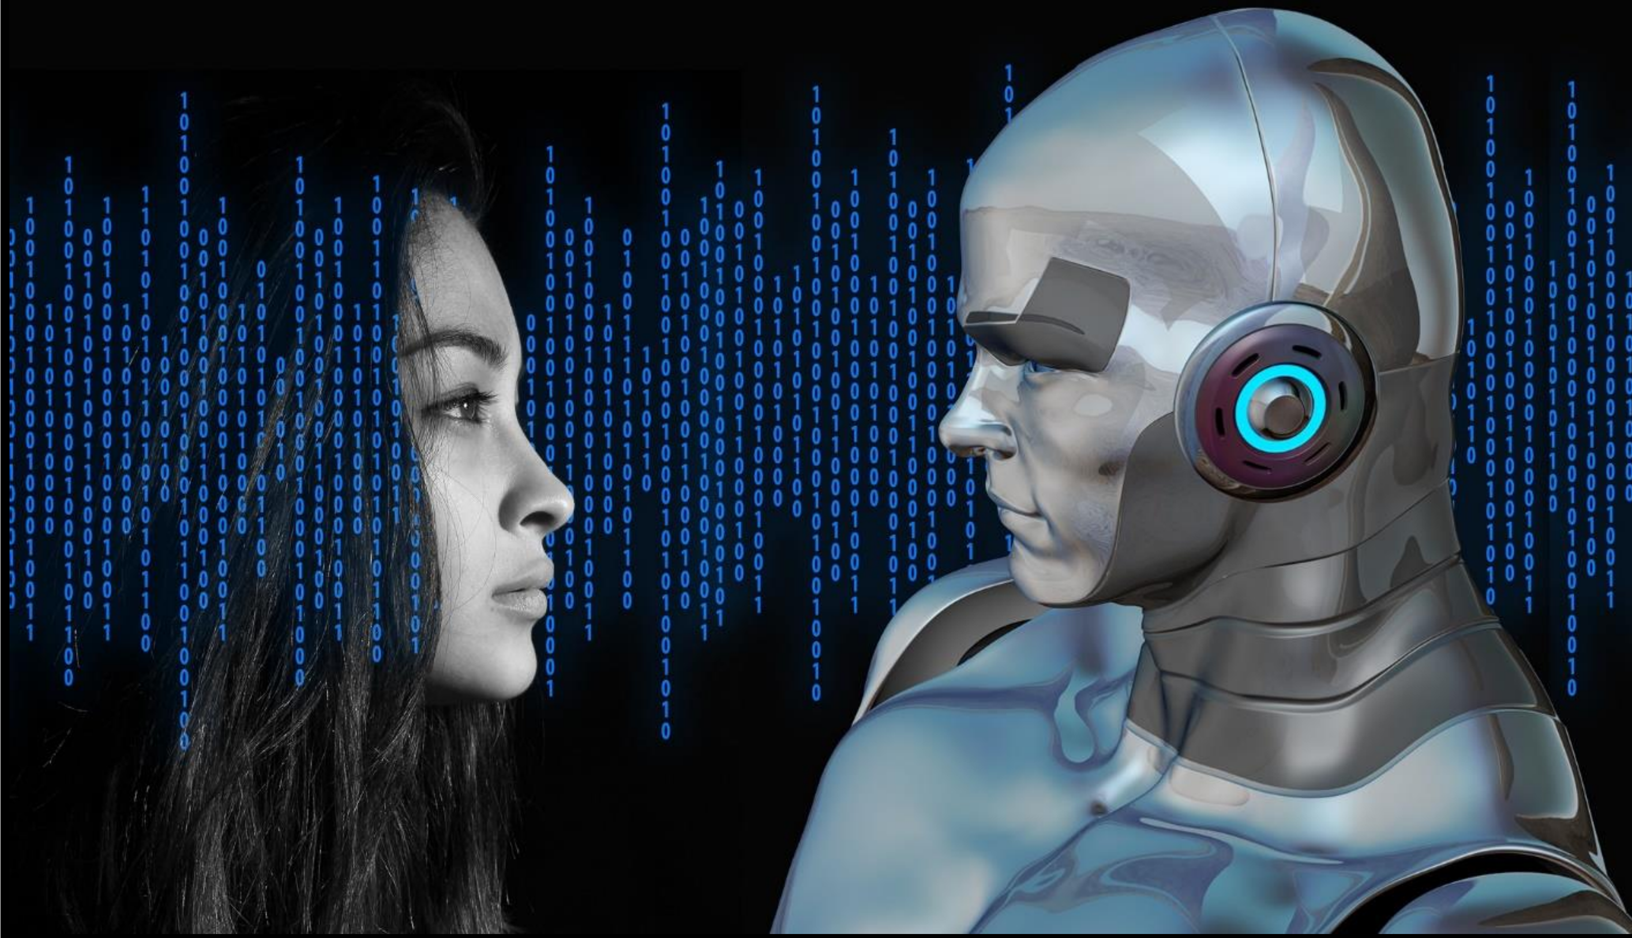

Symbiotic Autonomous Systems

An FDC Initiative

symbiotic-autonomous-systems.ieee.org

White Paper II

October 2018

S. Mason Dambrot, Derrick de Kerchove, Francesco Flammini,
Witold Kinsner, Linda MacDonald Glenn, Roberto Saracco
Edited by Theresa Cavrak

Contents

1 Overview

This White Paper follows the first one produced in 2017 by the IEEE Symbiotic Autonomous Systems Initiative (SAS)[1], extending it to address updated technologies and cover additional topics due to the evolution of science and technology. Additional white papers will follow because this is an area in continuous development.

The first examples of symbioses are already available in a number of areas and even now, these are impacting our economic system and way of life. The IEEE SAS Initiative takes a 360° view based on technology and standardization—the foundation of IEEE—and invites all interested constituencies to contribute complementary points of view, including economic, regulatory, and sociocultural perspectives. The transformation fostered by technology evolution in all paths of life requires planning and education by current and future players. Another goal of the initiative is to consider the future of education, given that these symbioses transform its meaning, making it both shared and distributed.

In this respect, the aims of this White Paper are to further develop the ideas presented in the first white paper: (1) to highlight impacts that are clearly identifiable today, and (2) to indicate emerging issues, thus providing a starting point to those involved in making public policy to understand the technical fundamentals, their evolution and their potential implications.

Note that this White Paper is intended to be self-contained, without requiring the reader to read the previous white paper.

The White Paper is structured as follows:

- **Evolution and Definition of Symbiotic Autonomous Systems**
  A general introduction to the area, touching upon the various aspects involved. It can be seen as an executive summary and may be of interest to the layperson.

- **Roadmap: from Today to the Future**
  The technology evolution is presented with a 20 to 30 year horizon to provide understanding of future impacts. At the same time, it is important to outline the steps that will take us to that future, knowing that the further we move in time the more ambiguous the landscape. The roadmap has the goal of helping decision making and steering in desired directions.

- **Technology Evolution 2030-2050**
  This section expands the technology overview provided in the first white paper, taking into account recent evolution and foresight studies. In particular, attention is given to the evolution of Artificial General Intelligence and supporting technologies, emergence of sentiment and mood analysis, a broad set of human augmentation technologies, and the growing pervasiveness of Digital Twins. Although the observation timespan is quite broad, it is rooted in current research and is based on current scientific knowledge and understanding of possibilities. Hence, the White Paper is based not on wishful thinking nor

science fiction, but rather concrete thinking that might, or might not, turn into commercial reality depending on several social, cultural and economic factors.

- **Socioeconomics, Culture, Law, Ethics, and Politics**
  This section places all technology evolution into the broader context of society and economics pointing out the mutual implication, i.e., how technology and its adoption impacts society, culture, and economics; and how societal, economic, cultural, legal, and political (including regulation) implications impact investment in technology, hence steering its evolution.

  The aspects of self, selves and super-self are addressed, as well as ethical implications deriving from the possibility of "designing" humans and legal issues of shared culpability and responsibility.

  The closing part of this section looks into the evolution of the law, the changing meaning of democracy as citizens expand into symbiotic citizens with blurring boundaries between people, machine, artificial intelligence, knowledge and cyberspace.

- **Market Impact**
  This section considers the broad implication of technology evolution on the market, including from the rebalance of labor between people and machines, expected loss/creation of jobs, the emergence of new skills needed, the furthering of the shared and gig economy[i], and shifts from consumption to usage and ownership to sharing.

  The evolution of manufacturing and the change in the whole value chain will then be addressed, considering the growing role played by artificial intelligence in the production, supply and distribution chains, aiming at the zero waste circular economy.
  The evolution in the areas of transportation, from self-driving autonomous vehicles to Hyperloop, as well as energy and genomics will be considered.

- **Education**
  This section addresses the changes in education fostered by the growing relevance of Symbiotic Autonomous Systems and the opportunities for IEEE to embrace this transformation.

  An education scenario at 2050 is presented, along with concrete examples of the seeds of change existing today, including start-ups designing the future of education.
  This is followed by a discussion of the shift from open-loop to closed-loop education.
  The section closes discussing symbiotic education, that is, shared between the person and the augmentation environment.

- **IEEE Societies Impact**
  This section provides several IEEE Societies points of view on the impact SAS has in their domain and the activities those Societies are engaged in or will be engaged in this area.

[i] In a gig economy, temporary positions are common, and organizations contract with independent workers for short-term engagements

2. Symbiotic Autonomous Systems

This section is largely based on the similar section "Evolution and Definition of Symbiotic Autonomous Systems" provided in the first White Paper of this series, published in 2017[1]*. It is a revised version; however, those familiar with the SAS Initiative and the first White Paper may skip it since it is basically intended to familiarize new audiences with the IEEE SAS Initiative providing context to this white paper.*

To a certain extent, human cultures have been characterized by the tools they made and used to the point where, starting with the Stone Age, these cultures are named after the predominance of a specific material used for tools. Notice that the idea of a tool is related to an artefact, more or less sophisticated but still manufactured by a human being to serve a specific purpose. The Stone Age was a time when our ancestors learned to shape stones in order to fit a specific purpose (to cut, drill, hit, scrape, etc.). Subsequent cultures have shown an increased capability to deal with additional materials (like bronze) in order to make new and more effective tools.

Tools as body extensions

Until the 18th century, tools were primarily an extension of our body powered by our muscles. While levers could trade displacement for strength, human power was limited by our muscle power (note that water and wind mills predated steam, but their application was constrained by the particular location).

With the invention and wide distribution of the steam engine, humanity quickly acquired the capability to use external power in ordinary fabrication methods. The issue for the culture of the 18th and 19th century became one of how to control this power.

At the end of the 19th century, electricity provided a new and different source of energy that was easier to control and use. As a consequence, electricity became the dominant way to manufacture products, including tools.

In the second half of the last century, the invention of computers made available a new *quality* of tools. Computer-controlled automated processes are improving the effectiveness of *control* and more recently have become outstanding tools for improving our reasoning and thinking capabilities.

Computers as tools for mind extension

We are in the Computer Age because many of our tools are directly or indirectly tied to computers. However, we are starting to see the emergence of a Digital Age in which the material to be manipulated and used for construction is no longer (just) atoms but also bits.

Spectacular advances in brain monitoring and in various forms of brain-computer interface (BCI), including deep brain stimulation (DBS), have proved the unification of *soft* (thoughts) and *hard* (neurons and neuronal circuits) in the brain. Notice that BCIs, similarly, are composed of a hard and a soft part with technology evolution in both. The former detects brain electrical activity with electrodes and affects brain activity using technologies like optogenetics; the latter interprets the detected activity creating "meaning" and commands specific actions to affect the brain.

At the same time, SAS creates new challenging questions about the emergence of shared thinking and shared awareness with profound ethical issues. This digital technology evolution is moving us towards the availability of a seamless integration (at different levels) of these computer/digital tools with us, the users. These tools are becoming a seamless extension of our body and mind, as the hoe was an extension of the farmer's arm. This seamless integration is very important, because it implies that these new tools are fading from our consciousness, we take them for granted, and they become an integral part of our life.

Think about the many times we use our smartphones to Google a piece of information. When we do this, we are extending our brain's memory and knowledge using a prosthetic device without giving it a second thought.

The symbiotic relationship with tools leads to humans 2.0

We are slowly entering into the age of human 2.0 or (or, as some have called it, transhumanism), and we are doing this through a symbiotic relationship with our digital tools. These new tools have become complex entities that are probably better referred to as *systems*.

Actually, the proposed change of name, from tools to systems is the consequence of a new qualitative dimension of modern, computerized tools.

While today's computerized tools are far more complex than what was used just 100 years ago, this is not the most important factor. Rather, today's tools are starting to operate autonomously and without our direct intervention, due to a growing flexibility and an improved awareness of their environment and decision-making capabilities. They are operating to fulfil a goal and take what they consider are the required actions to pursue and achieve that goal. Clearly one point is *who* sets the goal - can it be set by the SAS itself, or can the SAS change the goal on its own as the context changes and experience is gathered?

Self-evolving, autonomous decision taking, advanced interaction capabilities

Never in human history have we had tools with these characteristics. Robots are the first example of these types of tools that comes to mind. They come in many shapes and operate in different ways and for different purposes. They may differ significantly from each other, in terms of shape, dimension, functionality and cost. However, what matters most in the context of SAS is the varying degrees of autonomy they have, their capability to evolve (e.g., to learn and adapt), and their ability to interact with their environment, between themselves, and with humans.

We are therefore interested in SAS because of these three aspects: autonomy, self-evolution and human interaction. As SAS developments continue to progress at an ever-faster pace, they will change the landscape of manufacturing and life itself. They may even change what it means to be human.

Like all life on Earth, we have evolved to adapt our behavior to the context in which we live. However, by becoming able to change the environment to better suit our needs, humankind went a step further than simple adaptation. As a result, in the coming decades we will see that for the first time, artefacts that we have created will start to adapt themselves and their behavior based on their ecological context. In short, we will be part of *their* context.

Hence, starting in the next decade and even more so in the further future, we will live in a dynamically changing world where we will be responding to the behavior of machines, machines will be responding to our behavior in a continuously changing fabric, and it will become progressively more difficult to distinguish cause and effect between man and machine.

From symbiotic relationship to emergence of new entities

What is happening is the establishment of a symbiotic relationship among (autonomous) systems as well as between them and humans.

There is yet another aspect of these trends that will become apparent over the next decade. The interaction of several systems, each one independent from the others but operating in a symbiotic relationship with the others—humans included—will give rise to emergent entities that do not exist today. However, we are recognizing the abstract existence of something like a *smart city*, a *digital marketplace* or a *machine culture*. These entities are seemingly abstract concepts, although they are rooted in the interoperation of independent systems.

As an example, a smart city is the result of the interplay of several systems, including its citizens as a whole, as well as individuals. We can design individual systems and even attempt to design a centralized control system for a complex set of systems, such as a city. However, a smart city cannot be designed in a top down way, as we would do with even a very complicated system such as a manufacturing plant where everything is controlled. Just the simple fact that a city does not exist without its citizens and the impossibility of dealing or controlling each single citizen, as we would control a cog in a manufacturing plant, shows that conventional design approaches will not succeed.

In the past we felt that we could fully control a robot as we would a cog in a factory. However, as robots become more and more autonomous, aware, and able to self-evolve, they will become increasingly similar to human citizens, thereby requiring different strategies for management and control.

This emergence of novel abstract (although very concrete) entities created by these complex interactions is probably the most momentous change we are going to face in the coming decades. To steer these trends in a direction that can maximize their usefulness and minimize their drawbacks requires novel approaches in design, control, and communications that for the first time will place our tools on the same level as ourselves.

A new area of science

The IEEE SAS Initiative is inclined to think that a new branch of science is required, which we call Symbiotic Systems Science (SSS), rooted in the science of complex systems, taking into account the social and ethical implications. Consequently, promoting studies in this area is one of the goals of the initiative.

SAS with human participation

The symbioses of artefacts with humans will move by little steps and has already begun. For example, prosthetic hands are becoming more and more sophisticated, and part of their increased functionality stems from the autonomous nature of the prosthetics. When we pick up an object, several control systems are at work, even though we are normally unaware of their operation. For example, we can effortlessly pick up a nut or a raspberry, and we know to modify the pressure for the nut versus the raspberry, which is easily crushed. The decision process involved is quite complex, and it requires the cooperation of different systems; sensorial, touch, sight, motion, decision-making at the brain/cortical level, fine grading coordination by the cerebellum, immediate response by the spinal nodes, and more.

From shared to distributed knowledge

Prosthetic hands are now able to sense and interoperate with the person’s neural system; they can also make local decisions (like the level of pressure to exercise). To a certain extent, these hands are *autonomous systems*, and they enter a symbiotic relationship with the person wearing them. Notice that this development is a continuously evolving process resulting in increasingly advanced symbiotic relationships currently involving evolution slanted towards the person who is slowly learning to adapt his or her actions and reactions to achieve a better control of the prosthetic. Most recently, we are seeing the emergence of a co-learning, or symbiotic learning, approach where both the person and the prosthetic are engaged in a learning process that results in a distributed knowledge.

Note this knowledge is not shared, where every component has the same knowledge, but distributed, where each component has its specific knowledge and the symbioses generate the required overall knowledge.

A leading edge prosthetic hand, different from the first model that did not have sophisticated interaction capability, would not fit a different person because over time a very specific symbiotic communication will have evolved, mostly on the part of the person—today—but we are now seeing learning and adaptation taking place in the prosthetic hand as well.

Augmented Humans, Humans 2.0, Transhumanism

Embedded Internet of Things (IoT) devices are also becoming more common (think of sensors to monitor chronic pathologies, smart drug dispensers like insulin pumps, and home connected devices). IoT devices are getting more and more sophisticated. In a short while, these IoT products will communicate with each other through body area networks—and in the longer term, they are likely to create distributed decision points with an emergent intelligence. Shortly after this, a symbiotic relationship will be established with the person wearing the devices, first improving the person's well-being and then the user's physical performances and ultimately their intellectual performances as well. In this latter area, DBS and the progressively more sophisticated chips controlling it create a new way of interacting with the functioning of a person's brain, changing the way it works. This is the path leading to augmented humans, human 2.0, or transhumanism.

Humans 2.0

Although these three terms are sometimes used interchangeably, we take the view of a progression where the first step is leading to augmenting the physical abilities of a person (imagine having a wavelength converter embedded in the eye that allow that person to see in the infrared or UV spectrum), then reaching a point where many persons are markedly different from natural people because of their extended capabilities. These could include specific "improvements" like a permanent, seamless, connection to the web, made possible by advanced BCIs. This stage would characterize the development of human 2.0, and its main difference from augmenting the physical abilities of one person is the generalization that it will involve many people.

Augmented Humans

While in the augmented human we are likely to see an evolution that starts (as it is already happening) to address some disabilities and then move on to provide augmented advanced functionality to very few people, in the development of human 2.0 we have a generalized adoption of the technology probably due to decreasing cost for implementation. (Note that it has been said that we are already at that stage because of the generalized and systematic use we make of the smartphone to pair the web to our brain-based memory.) What we have in mind with our interface with devices like our smart phones is not the full human 2.0. We might concede to call this Human 1.5 insofar as in the nearer future, human to machine interfaces will remain visible. The transition to human 2.0 is marked by a seamless, often invisible, interface where you are not going to interact with the smartphone in an explicit way by typing or calling on Siri or Alexa but you simply think of something and related information pops up in your mind's eye after having been retrieved seamlessly from the web (or a local storage device that you may carry with you).

Transhumanism

Transhumanism carried to the extreme may signal a transition to a new species not driven by evolution, but, rather, by technological development. Although transhumanism is rooted in the concept of leveraging science and technology, it is looking not at a symbiosis between us and our artefacts but to the possibility of changing, at a fundamental level, the characteristics (or some of them) of humans.

Artefacts in a symbiotic relationship

We think that artefacts will evolve in a way that in some respects resembles the organic evolution of living creatures. The rapid development of technology enables this artefact evolution. It is therefore a natural step to extend the concept of symbioses one step farther applying it to the relationship between artefacts as well as living creatures.

Interestingly, we have examples in nature where these properties do not belong to individual components in a relationship but tend to emerge when many of these interact with one another as an ensemble. This is the case, for instance, for swarms of bees with a behavior as a group that is very different from that of individuals. Similarly, we can expect similar emergent behavior for swarms of robots. There is therefore a focus on two categories of symbiotic relationship only involving the interaction of artefacts with each other:

- Firstly, where each artefact demonstrates awareness-autonomy-evolution
- Secondly, where the ensemble demonstrates these properties as an emerging property

Local interacting Intelligences

In the former case, the symbiotic relationship may occur among only a few artefacts. An example is the area of robotics where as individual robots increase their awareness capabilities through better sensors and context data analysis, they become more and more autonomous with technologies supporting analysis and problem solving using AI/Deep Learning methods that evolve over time. This type of symbiotic relationship impacts several verticals—for example, Industry 4.0 (manufacturing and retail) and healthcare.

Emergent Behavior

In the second type of symbiotic relationship, there is a need for a significant number of artefacts to create a symbiotic relationship with enough complexity that emergent behavior results. There are no defined thresholds for complexity above which these properties emerge, although in general, the simpler the entities involved, the more of them are required. We see this in nature where a flock of starlings gives rise to amazing choreography in the sky with hundreds of birds while in the case of a swarm of bees the number is in the order of several thousands.

These aggregations can be studied with the science of complexity along with other technologies in the domain of AI. These aggregations and their emerging properties will be a topic of growing interest in the domain of IoT, although very few studies have focused on that. The interest derives from the fact that we are moving towards billions of IoT loosely connected with one another.

AI technologies can use data from the devices to extract emerging properties and direct the behavior of the IoT in the cluster.

IoT Swarms

This completely new domain will come into play in the next decade, as the number of connected IoT will reach a threshold above which awareness-autonomy-evolution can take place. 5G is likely to be an enabling technology in this domain providing the communication fabric for the ever-smarter IoT and clusters of IoT.

The growing connectivity is an enabler of increasingly complex systems, provided that each (or several) of the various parts have some autonomous characteristics. In turn, the studying of various technologies and application areas will require the SAS view.

Many of the IEEE Societies are likely to be affected, and one of the points raised by this series of white papers is a call to action for several of them to include the SAS perspective in their work and foster cooperation amongst them. An updated discussion and refined roadmap calling for joint action are outlined in this white paper.

3 Roadmap: from Today to the Future

A 10 to 30-year horizon

The technology evolution is presented with a 10 to 30-year horizon to identify possible impacts. It should be noted that some technologies considered in this White Paper are research topics today. They may be facing significant hurdles and eventually may never come to maturity, either because it will prove impossible to overcome those hurdles (from a technical or economic standpoint), or because alternative technologies will supersede the need for them.

The White Paper considers these technologies not under a probabilistic point of view (i.e., more emphases on those that are more likely to succeed), but on an equal footing explaining what the present research goals are and how the hurdles are being addressed. The hurdles themselves may vanish due to evolution in other areas, and new ones may appear as evolution occurs.

Also, over this span of time, we can expect new technologies to appear—but even if we dream about future technology, the methodology adopted in this White Paper precludes their insertion if they are not based, at least, on current research.

In the following subsections each technology will be described; they are embedded in a functional structure—i.e., we focus on the technology roadmap with reference to the functions they are supporting pointing out their mutual relationships (the success of one is likely to foster another one), thus creating a roadmap and the expected global timeline. The goal of this roadmap is to guide decision-making and steering in desired directions to enhance progress of the addressed functional areas (identified in the first white paper).

In the following subsections a brief explanation on the timeline of a given technology in a certain application domain is given. When the timeline of that technology is the same one of a previously described application domain, a direct reference to that explanation is given. In a few cases the timeline differs in total or in part from the one relevant to a previous application domain. In this case a new explanation is provided.

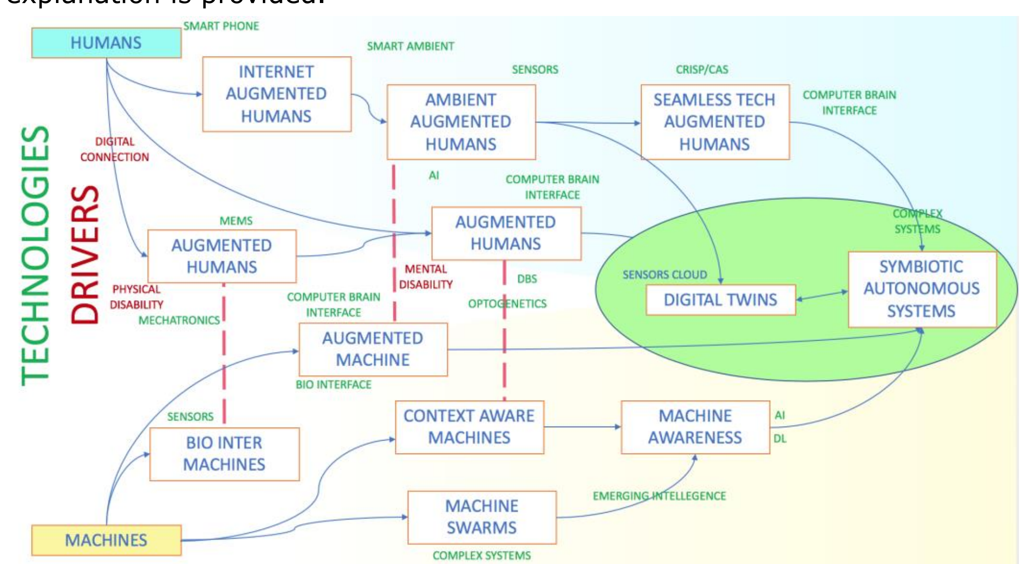

Fig. 3.1. Outline of evolution phases towards the emergence of Symbiotic Autonomous Systems

Figure 3.1 identifies the functional areas addressed and their relationships. In the following subsections each functional area will be addressed clustered under Machine Augmentation, Human Augmentation and Symbioses, the last including Transhumanism (Human 2.0).

The Gartner technology hypercurve is used to map evolution of each technology with specific reference to the functional area considered. This means that a technology may be represented in different phases in different areas.

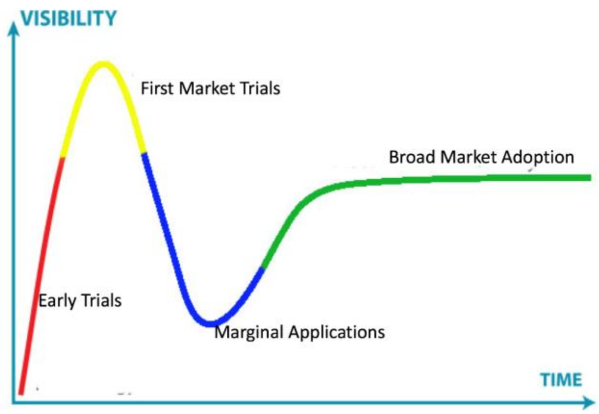

Fig. 3.2. Adaptation of the hypercurve to status of technology in SAS

In particular, a color code is used in the roadmap to identify the status of a given technology, indicating:

1. Red: Phase of early trials where academic research is leading the evolution
2. Yellow: First market trials in niches where performance and cost is not the main issue, mostly academic
3. Blue: Marginal application in market waiting for significant cost reduction to make it affordable and consistent performance meeting the needs. Industry is taking the lead in research/innovation.
4. Green: Broad market adoption. Evolution driven by market and industry. Research is continually occurring for improvement of the technology.

Notice the importance of the blue transition: This is where research shifts from academia to industry (hence the relevance for IEEE in partnering with industry at this stage). This is also where standardization is most relevant.

The goal of this section is to provide a rough estimate or roadmap of adoption of technologies and their mutual interplay, so the time axis is considered showing only the evolution in a 10-year window, and the area is characterized by some specific qualitative status.

The interplay of technologies—i.e., the point where they have to achieve a certain level of evolution (performance/cost) in order to proceed—is marked with dots connected by a dashed line.

For each cluster (Machine Augmentation, Human Augmentation, and Symbioses) a circle map is provided to show the various technologies involved in a cluster.

The circle diagrams have been created placing on one side the application domain and on the other the various technologies contributing to those domains. The thickness of the line connecting a technology to an application domain represent, in a qualitative way, the importance of that technology for the evolution of that domain. Notice that the diagrams are not exhaustive, more technologies are involved (as an example processing technologies are important everywhere). The choice has been to represent those that in a way characterize the evolution in that domain. Colors of the lines have no meaning; they are used in the rendering to facilitate the vision.

Notice that while the circle diagrams (see figures 3.3, 3.9, 3.15) contain all technologies that have been discussed in the White Paper relevant to each functional area, the roadmaps are presented only for the most impactful technologies. Since a given technology is often used in several functional areas it is discussed only once, unless it applies differently in different areas.

Please note that this White Paper has been written in 2018, and the roadmaps represent the consensus of the group of authors at that time. They will need to be revised as time goes by.

3.1 Machine Augmentation

Machines have increased their variety and performance over the last two centuries. In the last decades the evolution has been steered by improved electronics and manufacturing processes.

Machines that basically rely on electronics, like computers, CAD/CAM, and robots, that have been able to take full advantage of Moore's law and other technologies such as genome sequencers, have been able to evolve faster than the Moore's law by using parallelization.

If in the last few decades, electronics and softwarization paved the way to evolution, we can expect three main forces to steer the coming decade:

- Artificial intelligence
- Smart materials, including bio-integration
- Self-development

-

As indicated in the global roadmap, the following macro functional areas can be identified:

- Bio-interfaced machines
- Context-aware machines
- Machines swarms
- Augmented machines
- Machine awareness

Several technologies are fueling the evolution of machine augmentation, and one technology may contribute to advances in more than one area, as illustrated by the following circle diagram:

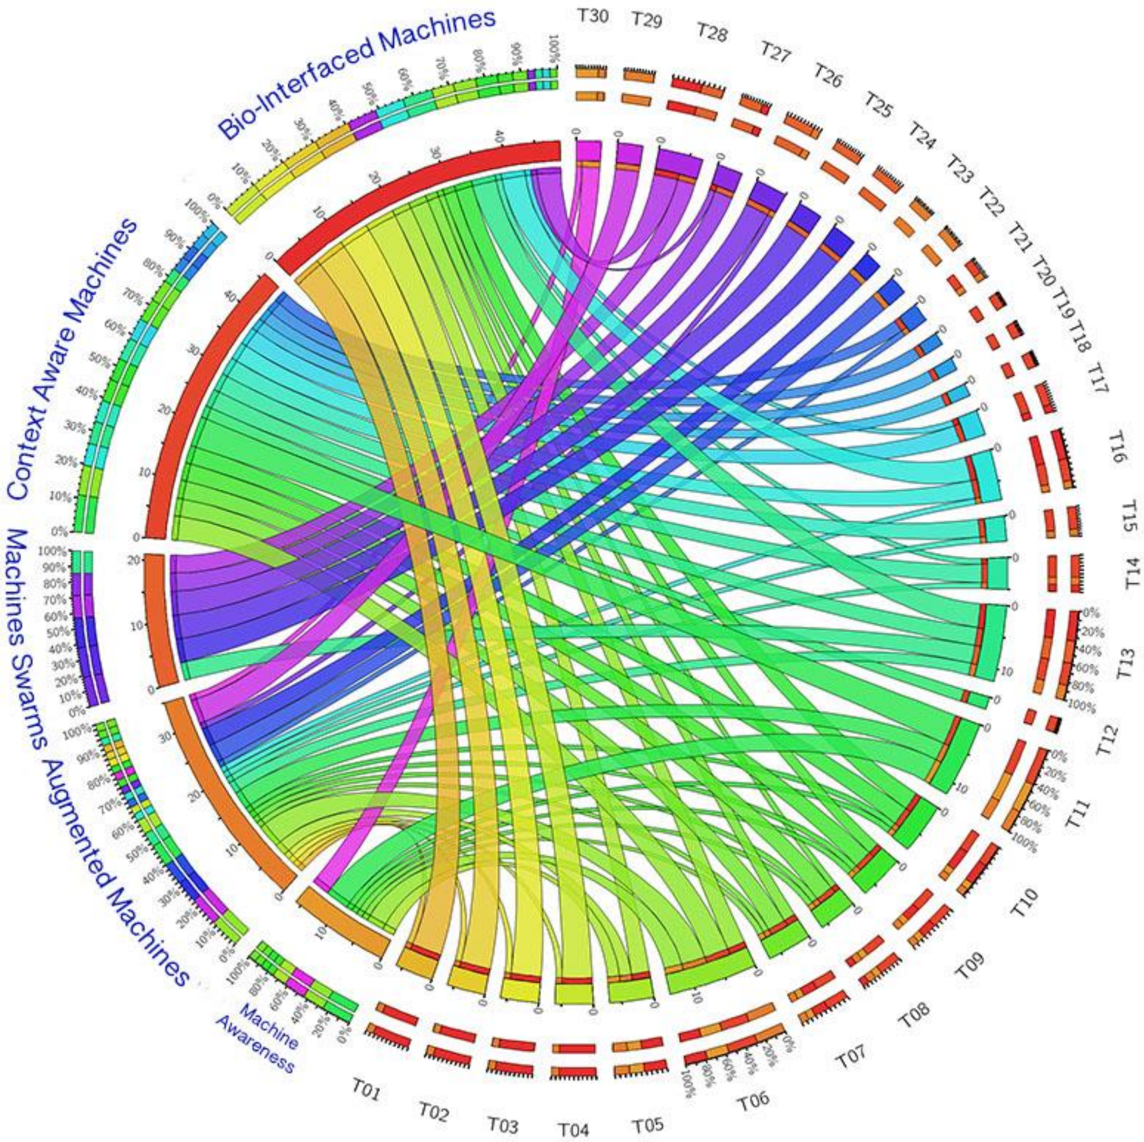

Fig. 3.3. Machine augmentation technologies

T01	Bio-nanotechnologies	T11	AGI	T21	Biometric Clues Detection
T02	Nano-biotechnologies	T12	LIDAR	T22	Affective Computing
T03	Optoelectronics	T13	Sensors	T23	Self-Replication
T04	Optogenetics	T14	Image Recognition Understanding	T24	Small Worlds
T05	Signal Processing	T15	3D Recognition	T25	Complex Systems
T06	Artificial Intelligence	T16	Pattern Recognition/Understanding	T26	Self-Orchestration
T07	Deep Neural Networks	T17	Intention Recognition	T27	Low-Latency Communications, 5G
T08	Recurrent Neural Networks	T18	Sound Signature	T28	LPWAN
T09	Convolutional Neural Networks	T19	Empathic Machines	T29	Autonomous Machines
T10	Machine Learning	T20	Social Robots	T30	Sentient Machines

Table 3.1. Technologies fostering/enabling machine augmentation

3.1.1 Bio-Interfaced Machines

There will be an evolution of this functional area from today's *independence* (the machine operates independently of the bio-system, like a pacemaker that sends impulses to the heart without being aware of the body's general situation) to *responsiveness* (the machine senses the status of the bio-entity and adapts its actions as needed) to a *continuous interaction* and to, finally, a *symbiotic status* where machine and bio-entity influence each other towards a common goal.

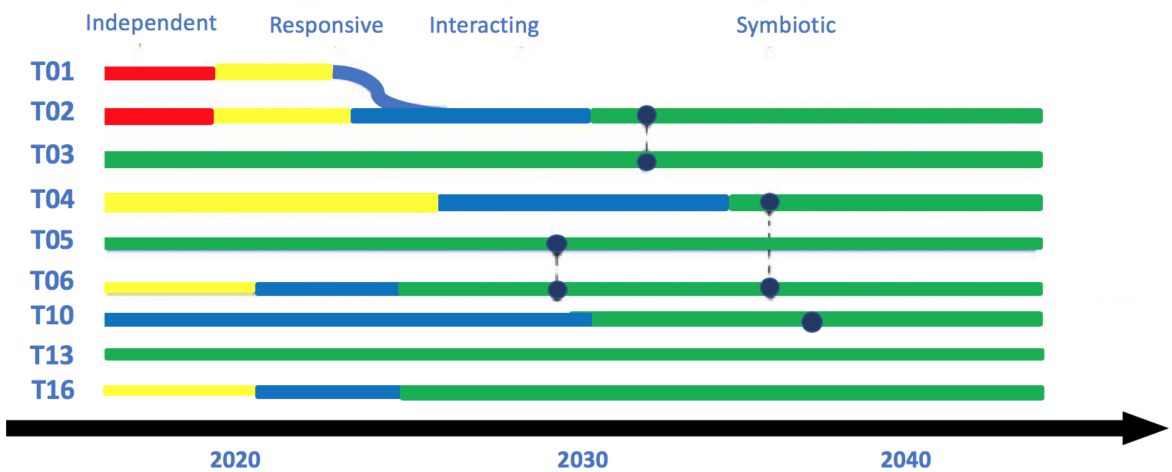

Fig. 3.4. Timeline of bio-machines related technologies

T01: Bionanotechnology
BioNanoTech (the use of nanotechnology for various biological applications) is still in its infancy today, at the confluence of bio and nano and addressed by different academic groups. By 2020, the first consolidated results will be applied in prototypes for bio-interfaced machines, mostly in prosthetics with specific focus on interconnection with the peripheral nervous system.

The two groups addressing BioNano and NanoBio, while today separate, are already converging and are expected to merge in the first part of the next decade.

Industry is likely to take the lead in the application (and further development) of the second part of the next decade. It is expected that these technologies will be applied in the optoelectronics area providing more effective interfaces beyond 2030 where they will have become state of the art for prosthetics.

T02: Nanobiotechnology
NanoBioTech (the use of biological tools for nanotechnological applications) has the same evolution trend as BioNanoTech described above.

T03: Optoelectronics
Optoelectronics is already a mature technology, in the sense that it is part of industrial products (particularly in optical fiber communications), and its evolution is driven by

industry. As indicated above, it will benefit from research in the nano area finding application in bio-interfaced machines.

Notice that bio-interface applications may require the use of wavelengths different from the ones used in telecommunications.

T04: Optogenetics

Optogenetics has shown significant promise. Thus far it has been experimented mostly on lab animals because of its need for gene modification and invasive procedures. At this time, it is seen more as a tool for getting a better understanding of the brain, but it will evolve first as a way to cure some specific pathologies (by influencing the firing of neurons). It may also be used to create strong symbioses with bio entities, including humans.

This is unlikely to happen before 2035—and even then only for very specific applications, most likely aimed at curing some deficit rather than to provide augmentation. The complexity of managing interaction through optogenetics in a distributed way—involving hundreds of neuronal circuits—will require artificial intelligence support. Since the implant of multiple probes for multiple neuronal circuits is very complex, it is not expected to become reality before the second part of the fourth decade of this century, hence the relationship with AI is foreseen from 2035 onwards.

T05: Signal Processing

Signal processing is a mature technology that is finding more and more fields of application. It is also progressing at a steady pace. In the area of bio-interfaced machines it is already extensively applied. As interactions are becoming more and more complex (e.g., capturing and delivering electrical signals to/from thousands of probes (using deep brain stimulation), AI support will be needed. This will require moving from signal processing done mostly at a single scale (monoscale) to independent multiple scales (multiscale) to processing at different scales simultaneously (polyscale). Such developments are expected to become widespread in the second part of the next decade.

T06: Artificial Intelligence

Artificial intelligence is already being used in some robotic prosthetics. While most are still part of academic research, at least one company, Össur, has been selling positional awareness AI-equipped prosthetics for years— Rheo Knee since 2004, Proprio Foot since 2006 and Power Knee since 2010[2]—and it is a trend that will progress in industrial applications in the next decade. It can be expected that in the first half of the next decade, industry will study embedding AI in their prosthetic products, and by the second part of the next decade, AI is likely to become a normal component of many prosthetics.

The challenge, particularly for brain-chip implants, is to have sufficient power to sustain AI computation without requiring significant power (which would result in high heat dissipation that would kill the surrounding cells). This is the main reason why the merging of AI in bio-interfaced machines is not expected to become the norm before the last part of the next decade.

T10: Machine learning

Machine learning is a mature technology, in the sense that it is already widely used today. Nevertheless, we can expect significant growth in its capability, due to specific (neuromorphic) chips with increased capability and the extension of the data the it uses to increase learning, areas where industry research will be leading. By 2030, we expect to have machine learning as a standard component in most systems.

T13: Sensors

Sensors are forming the bulk of IoT (Internet of Things). There are billions of them, and they will grow into trillions in the coming two decades. This volume plays in terms of:

- economy of scale, leading to lower and lower cost, fueling their adoption
- massive data generation, giving rise to soft meta-sensors further increasing their usefulness
- ubiquitous presence in the environment, thus enabling a variety of sensing architectures, partly relying on the environment and partly on the onboard sensors.

Although academia research keeps finding new ways to create sensors for sensing a broader set of parameters, this clearly is by far an industry-led evolution. In Symbiotic Autonomous Systems, there is the expectation of new ways to sense bio-entities and these technologies are addressed in the following functional area related to human augmentation.

T16: Pattern recognition

Pattern recognition is well developed in several areas (like in digital photography for removing moiré and noise) but it has not reached maturity. In particular, the understanding of the pattern needs further development, and this is where academic research is needed. Industry should be able to continue from there around the beginning of the next decade. Notice, however, that the use of Artificial Intelligence in pattern recognition is leading to quick and significant progress, with companies like Facebook and Google having many digital images that can be used to train AI algorithms leading the way.

3.1.2 Context-Aware Machines

As machines are able to harvest and process more data from their environment to create a model of the environment and to perceive their role and interaction with the environment, they are shifting from being *passive* to becoming *active* towards an *understanding* of the environment.

There are a few areas that are already seeing this evolution with self-driving cars at the forefront (it is likely that in the military area there is faster evolution, but progress is not disclosed). The context-awareness has already reached the mass market in products like robotic vacuum cleaners, but it is focused on very specific environment niches. A more generalized context-awareness will take a few more decades to become the norm.

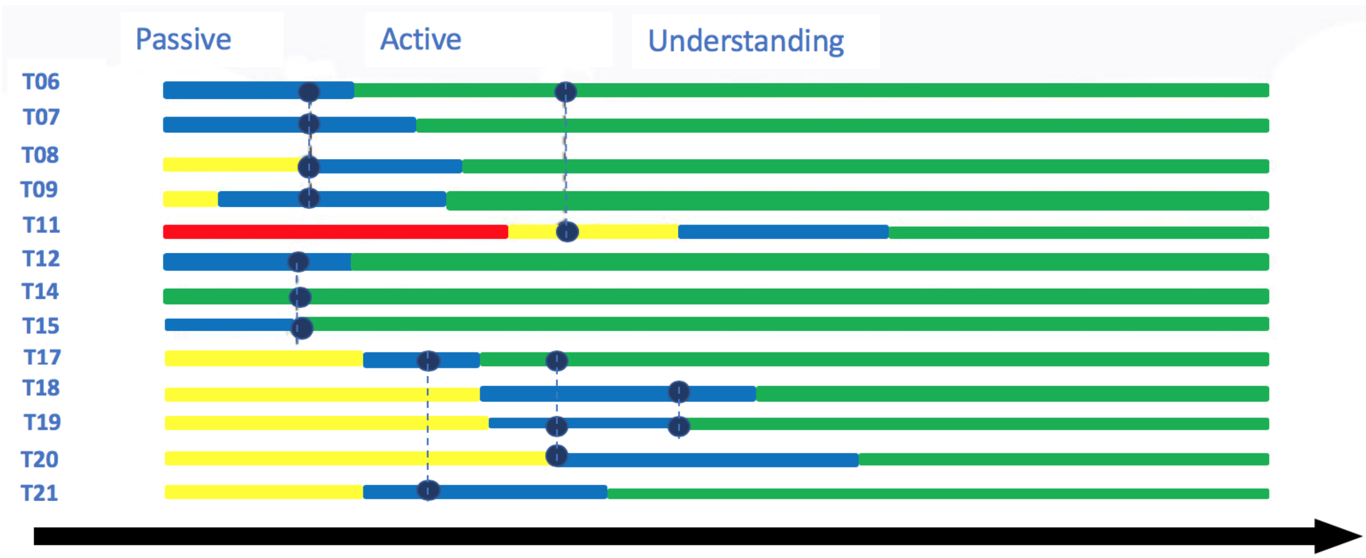

Fig. 3.5. Timeline of context-aware machine-related technologies

T06: Artificial Intelligence (AI)

Artificial intelligence is a crucial enabler for context-aware machines. It provides both the capability to recognize the various environment components, e.g., to tell a cyclist from a dog, and the understanding of the implications, e.g., a cyclist is likely to move in a straight line while a dog may wander around. In addition, artificial intelligence provides the bases for decision making.

AI is already present in several consumer goods (such as digital cameras that are aware of people smiling), and it will keep evolving. There is clearly plenty of research going on in academia, but it has reached an industrial maturity with many industries in many sectors working on applying AI to their products to make them context aware. Hampering context-awareness introduction in products is not a shortcoming in current AI, but rather the need for defining and evolving the product's concept and purpose (why should a given product become context-aware).

In the next decade we can expect AI to be part of all products that will require some form of context awareness. Clearly, progress is happening in the underlying technologies (deep neural networks, recurrent neural networks, convolutional neural networks), and more are likely to appear. By 2020, it is expected that these underlying technologies that today are seen as independent silos will become a toolkit for any AI need.

T07: Deep Neural Networks

Deep neural networks (DNN) are a layered structure of computation where each layer returns a probability that is further processed at the layer above. Probabilities are matched with the real world and change over time based on experience. Hence DNN are an ideal technology for learning from experience. The tweaking of the computation may be done internally or by an external operator. In the context of Symbiotic Autonomous Systems every component, in principle, can contribute to the fine tuning of the DNN.

Early in the next decade we can expect DNN to become part of many autonomous systems, providing the capability to learn from experience, hence making them ever more flexible and autonomous.

T08: Recurrent Neural Networks

Recurrent neural networks (RNN) are sequential structures that process and understand time evolution. They are utilized and well established in writing and speech recognition. In the context of Symbiotic Autonomous Systems, the temporal observation is clearly relevant but it is still in its early stages. It can be expected to become the norm in the second part of the next decade.

T09: Convolutional Neural Networks

Convolutional neural networks (ConvNet) are a class of feed-forward artificial neural networks mimicking the visual cortex in animals and are applied to image recognition. They are already part of the standard tool set for several image recognition applications. Scientists are making progress in understanding the circuitry of animals' brain, like the brain of a fly, and are investigating the effectiveness of replicating their capabilities in artificial neural networks. With a relative limited number of neurons and very little power requirements, a fly can orient itself in a 3D space whereas our artefacts require a massive amount of processing.

In Symbiotic Autonomous Systems, power requirements are often critical, and finding optimal, efficient ways to process images to understand the context is crucial. A significant amount of academic research is going on and we can expect industry to increase research in this area in the next decade and leverage them in the last part of the next decade.

T11: Artificial General Intelligence (AGI)

Progress has been made in the last decade on artificial intelligence, largely due to vast computation capabilities and access to big data sets in specific areas (like speech recognition and understanding). However, a general intelligence has proved elusive to the point that some experts are not optimistic on achieving it in the coming decades. Others are betting that it will be a reality by 2030; and in this White Paper we concur regarding that which impacts Symbiotic Autonomous Systems. Actually, it is difficult to place a boundary between artificial intelligence and artificial general intelligence. The former is bound to extend its fields of application to the point that in many areas it might be practically indistinguishable from AGI. From the point of view of Symbiotic Autonomous Systems, the feeling is that this will be achieved slightly before 2030. This is why artificial intelligence and AGI have been linked at that date point.[ii]

[ii] Actually, it is difficult to place a boundary between artificial intelligence and artificial general intelligence. Would it include emotions and self-awareness? The whole question of AI, general or not, is that it is based on our previous habits of

T12: LIDAR

LIDAR (Laser Imaging Detection and Ranging)[iii] provides an accurate measurement of the distance of an object and is used in self-driving cars to assess the environment. The cost, on the order of tens of thousands of dollars, is challenging for a mass market deployment. In this last decade, its price declined somewhat, but nowhere near the decrease in price of other electronic products (it requires some sophisticated precision mechanics whose price is not decreasing). This is why, although it is a mature technology, it has been flagged as requiring a few more years of industrial evolution to make it more affordable.

In the area of Symbiotic Autonomous Systems, a reliable measurement of distance of objects in the system environment is crucial but in the next decade it is not a given that LIDAR will be providing the solution. It is more likely to come from software rather than hardware, leveraging image recognition, 3D recognition and pattern recognition, each using raw data coming from very inexpensive digital cameras. By the time the cost of LIDAR is at an affordable level, it may be too late for adoption, given the uptake of the other technologies based on digital image processing.

T13: Sensors. See Section 3.1.1.

T14: Image recognition and understanding

Image recognition has advanced enormously in the last decade, hitting the mass-market with retail products like smartphones and digital cameras and becoming available as web services (for example, image search by Google or face recognition by Apple).

Image understanding is also progressing, although it has not reached the level of performance of image recognition (e.g., that is a dog but what is the dog doing). In Symbiotic Autonomous Systems, image recognition and understanding will be a basic, normal tool for machine context-awareness and more in general for machine awareness (see Section 3.1.5).

While evolution will continue, the present level of performance is already enabling significant product development, hence the green line.

T15: 3D recognition

3D recognition is less advanced than image recognition, but it is already at the industrial application stage. The effort on self-driving cars is stimulating progress, and it can be expected that by the beginning of the next decade 3D recognition will reach the level of industrial maturity that we have today in image recognition.

Evolution is progressing through analysis of shadows as well as through the understanding of objects. In general, the problem is solved today through massive processing power and access to big reference data sets. This is not ideal in several Symbiotic Autonomous Systems applications where the local processing power and access to data may be

---

distinguishing, separating, cataloguing human faculties in separate categories, so separate, in fact, that connections between the disciplines, say of psychology, sociology, and neurology, or engineering are not often considered, let alone practiced. See emphatic machines.

[iii] There are other interpretations of the acronym such as Light Detection and Ranging.

constrained. Smarter, more efficient approaches may be required and this is an area where academic research may provide indications on how to move forward.

T16: Pattern recognition. See Section 3.1.1.

T17: Intention recognition
Intention recognition is a recent area of investigation, where IEEE has already started to lead by organizing workshops and convening a variety of competencies. It is becoming very relevant to industry as it shifts to Industry 4.0 with robots cooperating with humans on the workshop floor. It is obviously relevant to self-driving cars and in the general interaction of machines with humans.

Symbiotic Autonomous Systems may not generally need this capability to interact internally, i.e., one component with another (although in some cases, having a hint on what the other component intends can be useful, as with medical implants), since symbioses are generally based on reactive responses rather than proactive (anticipatory) responses. In transhumanism, however, artificial super intelligence (ASI) might—in addition to its intelligence level being permanently beyond that of humans—make use of intention recognition and actually might be another differentiator from AGI.

Intention recognition will leverage from biometric clues detection (T21) as it will become widely available in the next decade and is likely to confluence in Symbiotic Autonomous Systems involving the human component.

T18: Sound signature
Sound signature is already used by industry in several applications (e.g., in agriculture to spot harmful bugs), but it can develop much further and eventually complement image/pattern recognition to provide a more comprehensive context awareness.
The technology per se is already available, and its evolution will see an integration with others. Specifically, it may be expected its use is integrated with empathic machines where the sound signature can provide hints on the emotional status of the human in a symbiotic system or in interacting with a human.

T19: Empathic machines
In order to better interact with humans at a social level machines need to understand the emotional state of the humans they interact with. The detection of emotional states relies on biometric clues detection (see T21), and data need to be processed using specific technologies (in the future they might involve communication with the Digital Twin of that human to understand the reasons behind certain clues). Empathic machines are going to become an enabling technology in a variety of applications, like elderly care, hospital care, interaction with disabled persons, interaction with children, and they can become a component in the symbioses with humans. Affective computing is a specific technology used by empathic machines.

T20: Social robots
Social robots are becoming a necessity as robots become a visible part of our society, interacting with blue collar workers at semantic level (i.e., learning and teaching on the job, becoming part of a working team), interacting with surgeons, with pilots and in the next decade finding a place in schools, department stores, hotels, and similar organizations.

It is expected that their presence will grow and will become part of the landscape. In the next decade some of these social robots may be able to engage in a dynamic symbiotic relationship with humans, i.e., when a human becomes part of an ambient (like a hospital room, a kindergarten, an elderly care home) the social robots will engage in a symbiotic relationship with that human, most often taking advantage of his/her Digital Twin.

T21: Biometric clues detection

Although slightly different in different cultures, a significant part of our human to human communication relies on biometric clues seamlessly detected by our brain (for example, a smile, tension in the neck, or eye movement). Actually, there are more clues based on physiological phenomena like an increase in the heart beat that can be detected by observing tiny changes in the color of the face skin (undetectable by our eyes but visible to a computer with optical sensors) that can be used. The availability of sensors coupled with signal processing and special software can vastly increase the number of clues and provide information to a machine. This is important in symbiotic systems involving a human component (as an example in a symbiosis with a prosthetic) leading to much better and effective interaction. There is expected to be a massive use of biometric clue detection in the next and following decades.

At the same time, biometric clues detection creates issues of societal interaction and privacy since biometric clues may overcome societal masks used in human to human interaction. As long as the clues used by a machine in a symbiotic system these issues are moot, but once bio-clues become widespread it is but a small step to leverage them in human to human communications, where a machine picks up the clues and convert them into information to the other human involved in the interaction.

3.1.3 Machines Swarms

As machines become more pervasive and

- able to detect what's going on in the environment
- have flexibility in their behavior, and
- have a goal, rather than an operationally prescribed behavior

it can be expected that they can aggregate into clusters, as happens in nature with flocks of birds, school of fish, and swarms of insects.

Notice that a swarm, in nature as in machines, does not require explicit communication among the members of the swarm. Rather the behavior of the swarm is an emergent property of the aggregated behavior of each single member. Each member is detecting what is happening around it and behaves accordingly. There is no orchestrator in a swarm, nor explicit communication.

In case of machine swarms, we expect an evolution from *swarming by design* to *ad hoc swarm* aggregation and behavior to the creation of a *super-organism*. Notice that a swarm is characterized by an emergent behavior, and in turn, this requires the presence of a multitude of components (members) in the swarm. The occasional opportunistic cooperation among machines/humans is considered in the section on the augmented machine (see next).

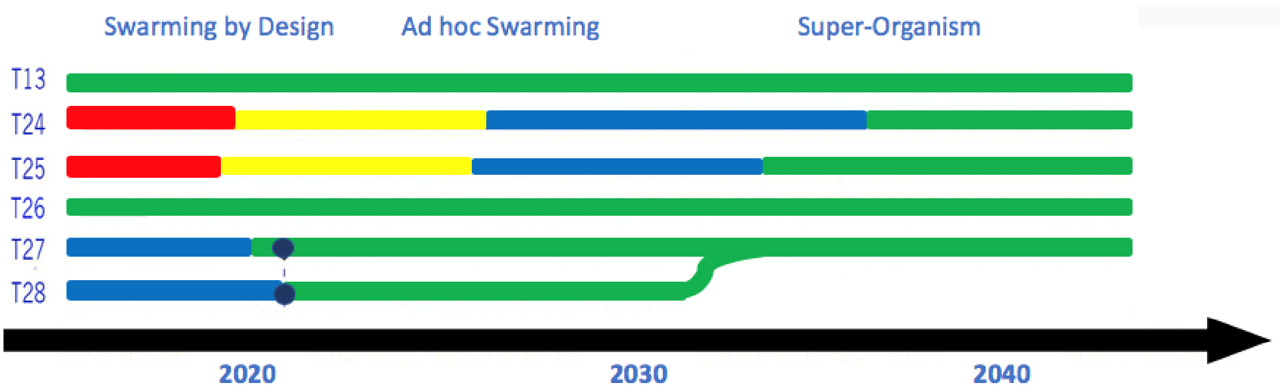

Fig. 3.6. Timeline of machines swarms related technologies

T13: Sensors
In this area, micro sensors are mostly used. MEMS, nano and bio tech are the enabling technologies. Apart from that the same considerations made in 3.1.2 apply.

T24: Small worlds
The theory of small worlds (a type of mathematical graph in which most nodes are not neighbors of one another, but the neighbors of any given node are likely to be neighbors of each other and most nodes can be reached from every other node by a small number of hops or steps) and the associated mathematics is one of the underpinnings in Symbiotic Autonomous Systems. It is a relatively new theory, and significant theoretical and experimental (observational) work (inspired by biosystems) is ongoing in academia.

Industrial applications may be envisaged in the next decade with effective commercial deployment taking place in the following fourth decade. Nanoparticles interactions, drugs, and interplay with neuronal networks are potential targets.

In swarms, particularly those formed by micro- and nano-systems, the small world theory plays an important role in describing the potential interactions and their impact with respect to the emergent behavior.

T25: Complex systems
Bio-organisms are clearly complex systems; a single cell is a complex system. So far technology has been able to create very large and complicated systems (a chip may contain billions of transistors) but not complex systems.

However, software and, most crucial, networks (like the Internet) have started to create complex systems. Artificial intelligent entities are complex systems, and swarms are complex systems when considered in terms of their emergent behavior. Hence, complex systems theory and related technologies are a crucial aspect in designing swarms and in understanding and leveraging their behavior.

Managing complex systems will be accelerating the transition from design to opportunistic swarming, and it will likely be the most crucial aspect in the creation of super-organisms.

T26: Self-orchestration

Software algorithms based on detection and reaction can support self-orchestration of a swarm. They are based on the afore-mentioned small worlds and complex systems theories. The challenge is to develop very stripped down local controllers that together can make complex behavior emerge (not just complex, also desirable). Self-orchestration is tied to the development of both small worlds and complex systems theories hence its evolution is dependent on those.

T27: Low latency communications – 5G

Swarm behavior relies on reaction times. Having longer latency may actually hamper the formation of an emergent behavior, and on the contrary very low latency gives rise to more dynamic behavior and allows the propagation of implicit messages to a greater number of members in a swarm, making them participate in the behavior generation.

Low latency communications such as the ones promised by 5G should be an important enabler for machine swarms. 6G is expected to be even better for swarms since it will allow an easier creation of self-organizing communications networks more effectively than 5G. That will happen beyond 2030.

T28: LP-WAN

Low power communications in wide area networks (LP-WAN) is another crucial enabler for machines swarms, basically clusters of IoT, since their powering possibility is severely constrained and the lower communications power required the better. Actually above a certain power level a swarm cannot operate, i.e., does not exist.
By 2030/32, the first availability of 6G embedding very low power communications as an integral part of its architecture will absorb other LP-WAN technology under its umbrella.

3.1.4 Augmented Machines

Machines have improved in terms of performance and types of activities they can carry out. This will continue in the coming decade through augmentation, mostly by adding intelligence. Hence, intelligence is key when discussing and qualifying machine augmentation.

From today increasing *local intelligence* machines in the next decade will become able to *flank and leverage* other intelligence, mostly other machines' or virtual machines' (like the web) intelligence. This will not happen by design but through the autonomous recognition that other forms of intelligence are available and can be tapped on demand. Eventually, beyond 2040, machines will be able to create a *symbiotic intelligence*, an intelligence that will emerge from multiple machines interacting.

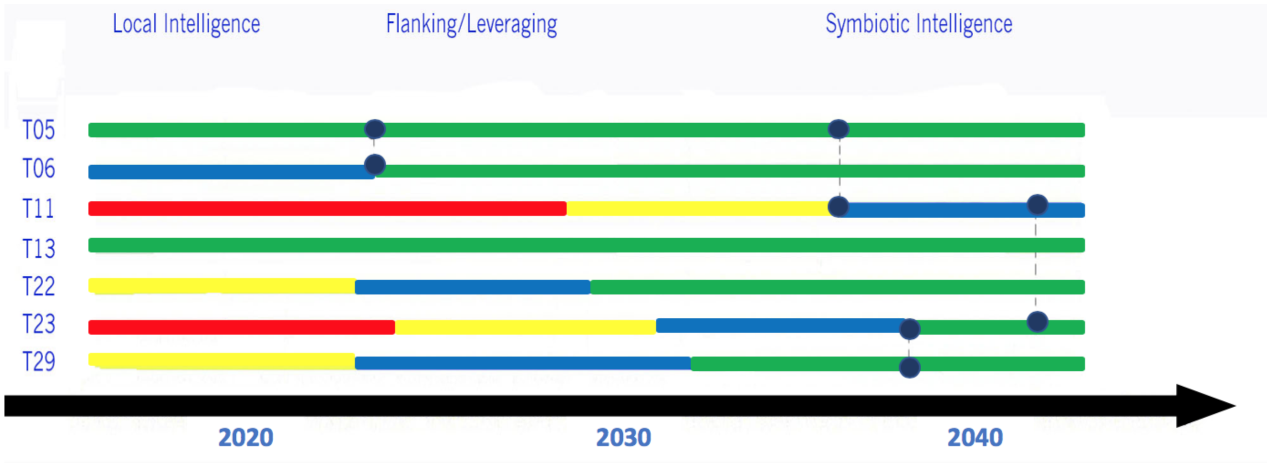

Fig. 3.7. Timeline of augmented machines related technologies

T05: Signal Processing

Signal processing will evolve as represented in Section 3.1.2. It is a crucial area for machine augmentation and will benefit from the application of artificial intelligence first, in the first years of the next decade, and then from the application of artificial general intelligence around 2040.

Signal processing will move from a syntactic analysis to a semantic analysis and will be contextualized more and more. Similarly, the signals generated by the machine will be contextualized. Whereas today a machine needs to understand "incoming signals" and be taught how to communicate, in 20-30 years, it will be able to make sense of signals and deal with them accordingly. Over a short time, it will learn new languages, both to understand and "speak them".

T06: Artificial Intelligence

Augmentation is pursued through intelligence. While today we already have a number of smart machines, augmented machines will use their intelligence to augment themselves. Hence the shift from today's local intelligence, used to be more effective in doing the activities they are supposed to do, to the next decade where machines will become smarter by using their own intelligence to tap on ambient intelligence (other machines, humans, distributed intelligence in the web). Baxter, the industrial robots of Rethinking Robotics[3] is a first step in that direction, able to learn by observing its co-workers (workers in the team). In the longer term, machines will autonomously create symbiotic relationships realizing how best to collaborate to create a team for approaching, solving, or executing a task to reach a goal.

T11: Artificial General Intelligence

Artificial general intelligence is often marked as the singularity, the point in which machines will outpace humans. Actually, machines have already outpaced humans in many areas and will continue, including areas like creativity. They have not, however, reached the singularity point. The consensus of the group that developed this White Paper is that the singularity is beyond the period of observation of this White Paper but we can expect industries to start working on AGI around the 2040 timeframe. Notice that with AGI, industries would be managed by a machine; a machine will be in charge of deciding where to invest, how to approach the market and so on. Humans might be relegated to taking care of the machine's needs, like bee workers take care of the queen's need.

We don’t expect this to happen in the observation period, and as remarked in Sections 3.2 and 3.3 we claim it will never happen because as machines get augmented so do we.

T13: Sensors
Sensors to detect environmental parameters and feed the AI and AGI are essential, and they are following the same sensors roadmap described in Section 3.1.2.

T22: Affective Computing
One form of augmentation is the ability to feel (or, in a pre-AGI sense, mimic) empathy. The affective computing technology is still pursued in academia, but it will be moving to industry in the next decade leading to machines more suited to co-exist alongside humans. By 2030, all machines that will be visible to humans will likely behave as humans in terms of societal interactions.

T23: Self-replication technologies
A few basic technologies like 4D materials (that is, able to change their shape over time), 3D printers, and smart materials are being considered to provide the basic building blocks for self-replication. Soft machines (software) are subject to fewer constraints (material constraints) in terms of replication, and work is already progressing in this area mostly at the academic level. Beyond 2030 it is expected that industry will engage in self-replication machines, and beyond 2040 self-replication will be leveraging with AGI to create better replicas, starting an evolution process that in principle will be self-managed by the machines themselves. Some cases of autonomous self-replication decisions (for soft machines) is likely to happen around 2035, possibly in the area of cyberattacks and defense from cyberattack.

T29: Autonomous Machines (decision making, goal setting)
Autonomous systems are already a reality in the respect that they operate autonomously (e.g., an autonomous vacuum cleaner). With autonomous machines in the context of augmented machines, the meaning is the possibility for a machine to make autonomous decisions in a broad space.

This will remain an academic area of research for the coming years, and by the middle of the next decade it is expected that industry will be studying creating augmented machines able to make autonomous decisions. Enablers from regulatory, societal and ethical standpoints will be required to make this happen.

3.1.5 Machine Awareness

In order to get smarter, machines need to become more and more aware of their context, goals, and abilities. By far the basic enabling technologies are the capability to process the signals received through sensors from the environment (including their active observation of the environment) and the intelligence to make sense out of them. Accordingly, three phases can be identified: *task awareness*, *goal awareness* and *self-awareness*.

Fig. 3.8. Timeline of machine awareness related technologies

T05: Signal processing
Signal processing is the starting point for machine awareness. As pointed out in Section 3.1.1, signal processing is a mature technology that continually improves as more data points can be harvested and more intelligence can be applied due to increased processing capability. Academic research is still occurring although industry is actively improving signal processing and learning to leverage it.

A likely evolution is hardware architectures mimicking the nervous system with hierarchical computation with feedback along the hierarchy, basically implementing some artificial intelligence processing within the detection chain.

T06: Artificial Intelligence
Artificial intelligence is at the core of machine awareness, to be superseded by AGI once it becomes available. In most cases, signal processing does not require AGI, and AI will be sufficient. We can therefore expect AI to continue to be applied in signal processing for the foreseeable future. AI, in this sense, is rapidly becoming a mature, industrial grade technology and we can expect it to become a common tool in industry by the middle of the next decade.

The approach to AI will continue to improve and become more integrated (see T07-T10 below).

T07: Deep Neural Networks[iv]. See Section 3.1.2

[iv] T07 through T10 are toolkit technologies (T07 to T09 are already described in Section 3.1.2; T10, described in section 3.1.1, is similar with a longer research evolution span to become a commoditized mass market technology as shown in the roadmap timeline) supporting artificial intelligence first and artificial general intelligence later.

They are applied in different sectors of AI today each one with its own strength. In the coming decade, and more so in the following ones, a form of unification of these toolkits is expected, plus the addition of new ones. A boost in this direction may derive from an endeavor like the human

T08: Recurrent Neural Networks. See Section 3.1.2

T09: Convoluted Neural Networks. See Section 3.1.2

T10: Machine Learning. See Section 3.1.1

T11: Artificial General Intelligence (AGI)
There are several definitions or interpretations of AGI. The most practical one are the operational ones, i.e., the ones that provide a method to assess if AGI has been achieved, like the Nilsson test[4].

Artificial general intelligence is not a requirement in achieving machine awareness, and it is not a given that AGI will become possible within the observation period of this white paper. There is no consensus on the achievability of AGI, its unique definition, metrics on when AGI will become a reality, nor a date (although some are convinced that the singularity will happen within the next 20 years).

So far, and for the next decade, AGI is a matter of academic research (also pursued by a few companies, such as Google[5]), and it isn't expected for general industry to become active on AGI before the later part of the fourth decade. An acceleration may come from the result of the ongoing brain projects (Human Brain[6], Brain Initiative[7], Human Connectome Project[8]).

AGI is not required to support machine awareness in the first two phases, and possibly not even in the third phase, self-awareness.

T30: Sentient Machines
If there is not a general consensus on metrics for machine awareness there is even less on sentient machines, machines that have a self and recognize themselves as a "living entity" with a goal and a sense of fulfilment. We do not even have an agreement on most living beings whether they are sentient or not. Although most people would feel their dog is a sentient being fewer would consider a fly as a sentient being. Religions have their interpretation of sentient (which often translate into a culture of sentient), and regulators have their own interpretation of sentient (in New Zealand since 2015[9] animals are considered as sentient beings although it is not specified if that applies to all animals or just to a subset of them).

---

brain expected to be close in the first years of the next decade after having explored the basic structures of the brain and how these can be mimicked artificially.

By the end of the next decade, neuromorphic chips supporting AI toolkits can be expected. Around 2040, a new generation of neuromorphic chips, at very low power (getting closer to bio neural circuits in terms of power requirement) may become available as a platform to support AGI. These chips will be *de facto* implemented in their architecture like today is done by deep/recurrent/convoluted neural networks in software running on normal chips. In turns this will make machine awareness possible and affordable.

It is generally accepted that in order to be sentient an entity should be able to sense, perceive, think, feel and experience subjectivity. There are technologies that clearly cover the sense and perception aspect. We are developing technology supporting some sort of thinking, and we can program a sense of identity (subjectivity) in machines. The difficult part is related to feeling. A machine can be programmed to get rewards, as well as getting the sense of an undesirable situation. However, it is anyone guess if this translates into feelings. Could a machine be happy or sad? It can surely be programmed to act and appear as happy, sad, but would it feel it?

Feelings and emotions are complex matters because they require the participation of several organisms in the animal body, including the frontal lobe, the olfactory bulb, the thalamus and hypothalamus, the hippocampus, and the amygdala, all parts of the limbic system. A hybrid robot might simulate such processes, but it would remain to be seen whether that would help the item to experience rather than simulate the feelings necessary to correlate with the attending circumstances.

Most scientists accept the Turing test, acknowledging that if a computer behaves in a way that it is indistinguishable from a human in all effects, it can be assimilated to a human and passes the Turing test. Could we say that if a machine behaves in a way that looks happy (or sad) we can assume to all effects that it is happy (or sad)?
Technologies that allow a machine to look happy (plus any other state of feelings) are being experimented to increase the acceptability of social robots, and we might expect them to create machines that show their feelings when interacting with humans. Whether they are actually feeling them it is a different story but for some it is irrelevant as long as we perceive them as having feelings.

3.2 Human Augmentation

The main driver towards human augmentation in the next two decades will remain treating disabilities, although in the military area, niches of human augmentation are already being pursued. In factories and surgery some physical augmentation is already available and will keep growing to meet niche demands.

Beyond 2040 the landscape may change with the beginning of human augmentation aiming at increasing performance. This is likely to open the gate to transhumanism. Technology evolution is likely to support this change, although societal and ethical considerations may stop or delay evolution.

Human augmentation will be enabled (in the order of time):

- by symbioses with prosthetics of various forms
- by flanking of soft and hard machines interconnected to the human body/brain in various ways
- through DNA/RNA reengineering.

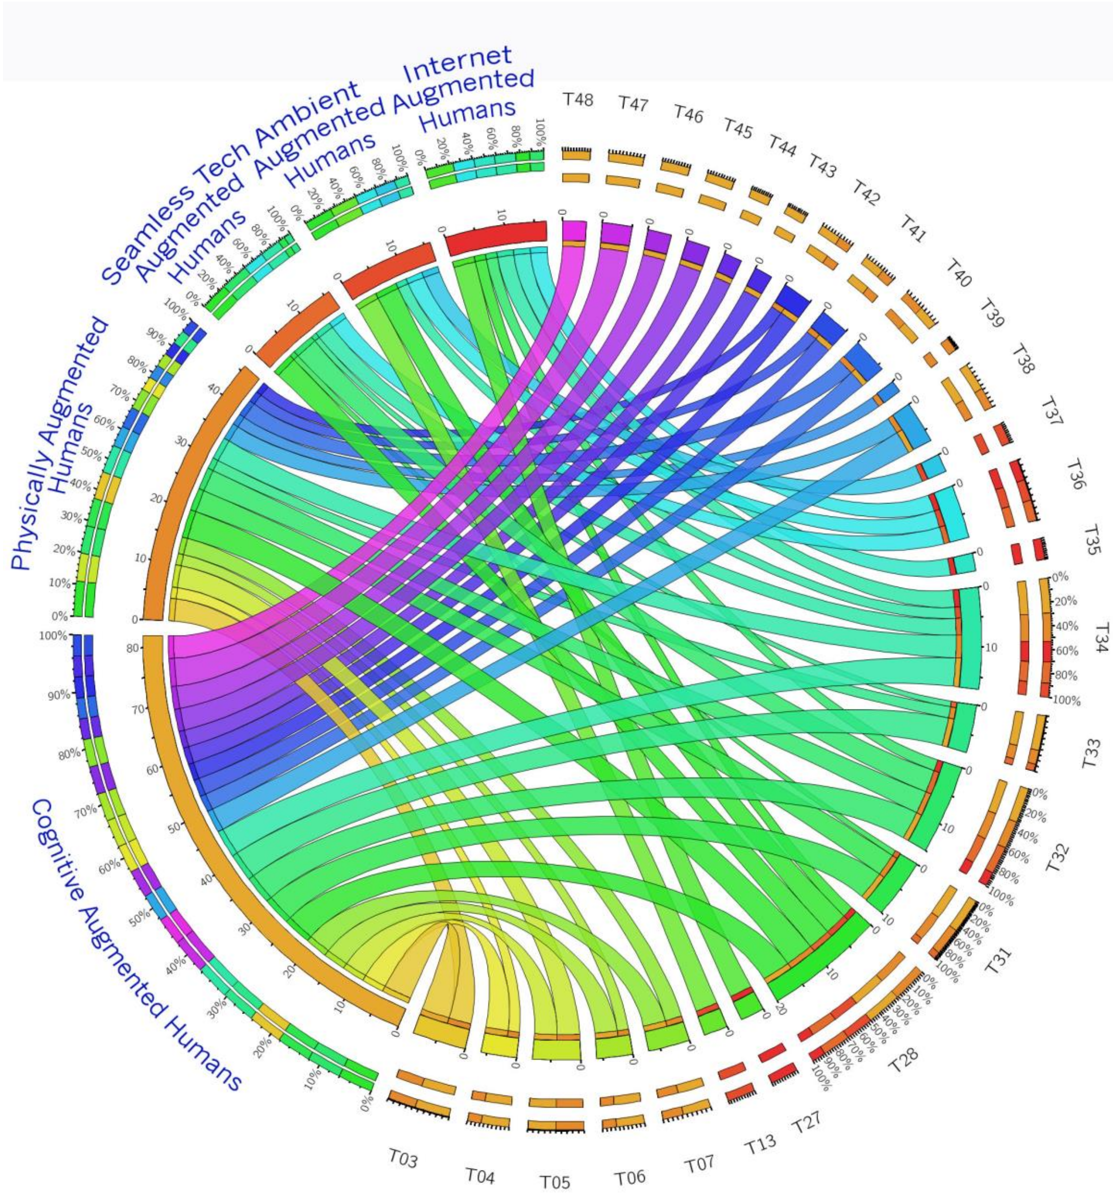

Fig. 3.9. Human augmentation technologies

T03	Optoelectronics	T32	Smart Materials for long term Implants	T41	DNA modification
T04	Optogenetics	T33	Neuroimaging	T42	RNA modification
T05	Signal Processing	T34	BCI	T43	Deep Brain Stimulation
T06	Artificial Intelligence	T35	Virtual Reality	T44	Transcranial magnetic stimulation
T07	Deep Neural Networks	T36	Haptic	T45	Cognitive prosthetics
T13	Sensors	T37	Smart Materials	T46	Neural Engineering System Design
T27	Low-Latency Communications, 5G	T38	Implantable chips	T47	Smartphones
T28	LPWAN	T39	Exoskeletons	T48	Symbiotic Intelligence
T31	Fluorescent proteins	T40	CRISPR/Cas9		

Table 3.2. Technologies to foster/enable Human Augmentation

3.2.1 Internet Augmented Humans

The Internet has augmented each of us. We have been given access to an unlimited source of data, information and knowledge. We have also been provided with communications tools, calculation tools, orientation tools, search tools and so much more. What used to take days is now available at our fingertips in seconds. Not even the industrial revolution has augmented humans so much.

Today the use of the Internet is *mediated by devices* (smartphones, tablets, computers…); in the next decade it will be *mediated by a cloud of devices* no longer perceived, most of the time, and in the longer term it might become seamless through direct *body-Internet connection*.

Fig. 3.10. Timeline of technologies for Internet augmented humans

T27: Low-Latency Communications - 5G
Low-latency communications might play a role in BCI when capturing the correct time differences of signals generated in different parts of the brain can be of significance. Similarly, the possibility offered by 5G architectures of aggregating a variety of local networks may be useful in the communications among a cloud of devices, mixing body area networks with personal area networks and with picocells. The trend towards human augmentation through Internet access mediated by a cloud of devices will surely benefit from 5G first and then most definitely from 6G that will support bottom up creation of networks.

T28: Extremely low power electronics
Low-power electronics are a crucial aspect for body implants, particularly for brain and sensory implants (e.g., an ocular chip). Hence, there is a tight connection with BCI that will stimulate research of more sophisticated (invasive) BCI around 2023 and further down the lane, by 2035, brain chip implants. We already have low power electronics but we are still far from the kind of extremely low power that can compare, and be compatible, with cells. Basic research is still needed, and then it will have to be translated into industrial manufacturing processes. In this area, Moore's Law (although no longer fully applicable) provides a good measuring stick. Within 20 years, we should be able to reach 10 microjoule per bit (we were at 100 µjoule per bit in 2012, in 2018 we are around 80 µjoule per bit), which is the level of consumption for a neuronal spike in the human brain.

T32: Smart materials for long-term implants

Smart materials are crucial for permanent implants. They need to be bio-compatible (i.e., do not create inflammation) and able to adapt to changing conditions. An electrode inserted in the brain becomes useless as the brain re-wires itself over time, with neurons creating new dendrites and connection and signals being rerouted as the brain gain experiences. This brain evolution shall be matched by implants' evolution which is very challenging. Implanted materials shall be able to change their shape, grow, and modify their characteristics as needed by the changing ambient conditions. There is a lot to discover at the scientific level (red line), and only in the next decade can we expect research to move on into some prototypes. Industrial takeover is over ten years away with applications expected in the fourth decade.

T34: BCI

There are already several BCI but very few CBI. That is, the detection of signals from the brain has progressed significantly, but it is still cumbersome and the tradeoff between non-invasive detection and precision is not good if you are interested in capturing thoughts.

CBI, the transfer of information from the Internet to the brain, is in its pre-infancy. So far the only practical solutions are based on using our senses to feed information to the brain. We are at least 15 years away from a seamless centripetal communication to the brain.

T35: Virtual Reality

Virtual reality has made significant progress but it still requires bulky apparatus that by their very existence undermine the virtual experience. There is quite a bit of industrial research going on, and we may expect that seamless virtual reality will start becoming real in the second part of the next decade.

T36: Haptic

Haptic interfaces have become common (like the vibration sensation on touching a smart phone screen) but are still far from truly cheating the brain. Haptic is going to be important to deliver a credible virtual reality; that's why the two lines are joined. So far both are not mature enough to be combined, but that should change in the next decade.

3.2.2 Ambient Augmented Humans

We have created a world of artefacts starting millennia ago with the construction of infrastructures, cities and dwellings. We are now in a transition phase transforming artefacts into smart artefacts that are becoming able to sense their environment, adapt to it, and interact in ever more effective ways. Humans are going to take advantage of the increased flexibility and smartness of their environment leveraging it for their own augmentation.
From today's *passive* ambient we are moving to a *responsive* ambient that can be morphed into one that augments our capabilities, and in a few more decades will become proactive and *symbiotic* with humans.

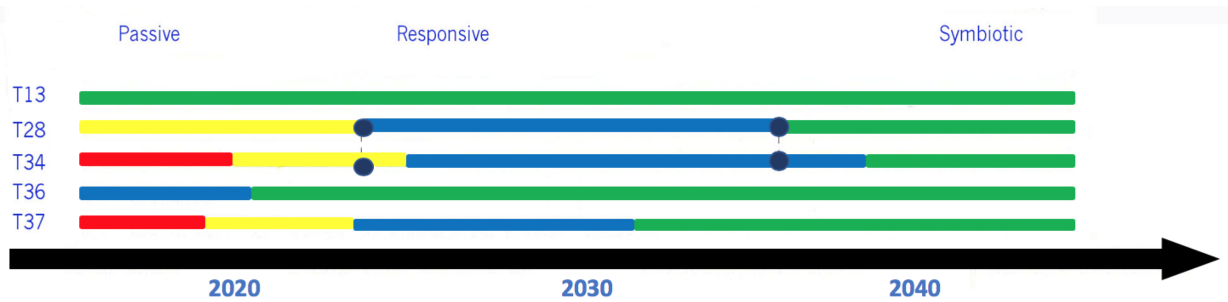

Fig. 3.11. Timeline of technologies for ambient augmented humans

T13: Sensors
Sensors are becoming an integral part of many artefacts, and in the coming decades smart materials will have sensing capabilities. In 20 years it will be unusual to find an artefact that does not embed sensors. IoT will be pervasive, in the trillions.
The few artefacts without sensors will be in ambient where other artefacts will sense them, thus providing indirect sensing capabilities.

T28: Extremely low power electronics
The ambient will be able to provide continuous low power supply (via radio waves). Low power electronics are needed to leverage ambient wireless power. For more general consideration on low power electronics, see Section 3.2.1.

T34: BCI
Interaction with the ambient through BCI requires the availability of wireless BCI. See Section 3.2.1 for a more general discussion.

T36: Haptic
Ambient interaction will be supplemented by haptic feedback, also by virtue of the availability of smart materials. Haptic may transform an ambient along with some augmented reality technology to create a real-virtual ambient, that is, in morphing an ambient into a different one from a perceptual point of view. In turn, this provides the human with augmented sensations.

T37: Smart Materials
Smart materials are beginning to be available, like materials that can absorb light, generate electricity, change their color, absorb CO_2, degrade pollutants, sense pressure and temperature and so on. Smart materials will be able to provide dual functions—such as a screen providing any surface with interactive capability. There is still quite a lot of scientific research needed, particularly in the area of nanotechnologies (leading to additive manufacturing). The goal is to be able to design a material starting from its desired characteristics. In the next decade academic research will turn to the application of smart materials, and by the end of the next decade industry will start studying how to manufacture them on an industrial scale.

3.2.3 Seamless Technology Augmented Humans

Real augmentation will happen once the applied technologies become seamless, falling below the level of perception. This will require embedding these technologies in wearables, their implant on and in the body, a seamless interaction of external devices, or gateways with our brain or our senses.

We can expect an evolution from today's *wearable* to *direct senses interaction* to *symbioses*. Notice that the very fact of "seamless" (hence invisible) is raising societal and ethical issues that may affect the roadmap, mostly playing adversely (lengthening the time of progress).

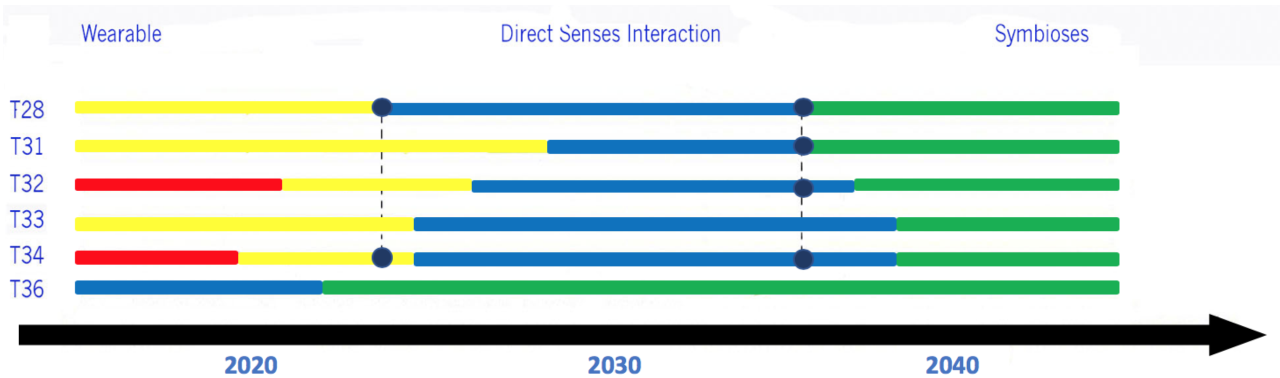

Fig. 3.12. Timeline of technologies for seamless technology augmented Humans

T28: Extremely low power electronics. See Section 3.2.1

T31: Fluorescent proteins
Fluorescent proteins are used as markers to study cells in vivo. Specific proteins can bind to specific cells and even to specific parts of a cell, like a synapse. By designing these classes of proteins it becomes possible to affect (through optogenetics, see Section 3.1.1) very specific circuits in the brain, thus enabling accurate interactions. Although there is a lot of work going on it is all very academic, and it won't move to industrial research for at least one decade.

T32: Smart materials for long-term implants. See Section 3.2.1

T33: Neuroimaging
New techniques for visualizing the brain, its structure and its "workings" are creating a new discipline that is rapidly evolving allowing scientists and researchers to get more and more data. This leads to better understanding of the brain and to the possibility of finding ways to interact with it. The progress in this area as well as several others that are being pursued by the three big initiatives on the human brain are bound to accelerate the shift towards seamless technology to augment humans with impacts expected in the fourth decade of this century.

T34: BCI
An effective BCI (and CBI) would be the ideal interface for a seamless technology for human augmentation. However, current technology is nowhere near seamless. See Section 3.2.1.

T36: Haptic
Touch sensations are a fundamental component in our perception of the world, although we may often underestimate touch. Hence, haptic technologies are an important enabling factor in achieving seamless augmentation. See Section 3.2.2.

3.2.4 Augmented Humans (Physical)

Prosthetics designed to fill disability gaps are improving every day, approaching the functionality and performance of the part they are substituting. In a few situations they can provide even better performances than the original. This will continue to the point when prosthetics will consistently deliver better performance than the original. At that point some people may start considering having a prosthetic to benefit from the increased performance, and in a short while more and more people will start adopting them. In the end, a few of these prosthetics will become indispensable (as has happened with the smartphone). The augmentation will progress from *overcoming disability* to *focused augmentation* to *overall augmentation*.

There are many technologies that are making augmentation possible (as discussed in this White Paper and shown in the circle diagram). Several of these technologies have already been presented in the previous sections and will not be repeated here, namely:

T03: Optoelectronics. See Section 3.1.1
T04: Optogenetics. See Section 3.1.1
T05: Signal processing. See Section 3.1.1
T06: Artificial Intelligence. See Section 3.1.1
T07: Deep Neural Networks. See Section 3.1.2
T28: Extremely low power electronics. See Section 3.1.3
T31: Fluorescent proteins. See Section 3.2.3
T32: Smart materials for long term implants. See Section 3.2.1
T33: Neuroimaging. See Section 3.2.3
T34: BCI. See Section 3.2.1

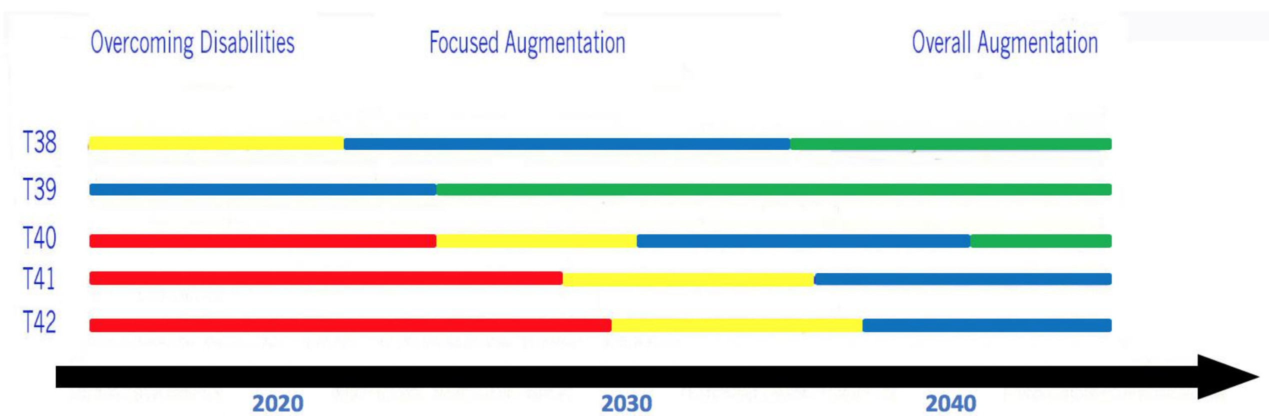

Fig. 3.13. Timeline of technologies for augmented humans (physical)

T38: Implantable chips

A few complex implantable chips (like the Argus II for the retina[10]) have been in experimental stages in the last few years (notice that Argus II was approved by the FDA and has been implanted in thousands of blind people providing a minimal restoration of sight, so it may considered to be beyond the experimental stage). More simple chips, like RFID and glucose sensors, have been implanted since the last decade. Chip implants to control seizures have also been experimented. There is a lot of potential in this area, but there are also big hurdles to overcome; hence a lot of research is occurring. Issues like chip bio-compatibility over long periods of time, powering of the chip and more recently concerns on malicious hacking need to be solved. Besides, all chips today, and for a while, are designed to tackle a very narrow scope. There are no chips in the foreseeable future that could be used for a generalized augmentation.

Although chip implants will increase over the next decade it is unlikely to reach the level of sophistication and reliability that would make them adopted for overall human augmentation.

T39: Exoskeletons

Exoskeletons are already available, although in experimental stage, in health care, industry (for relieving fatigue on assembly lines) and the military. Industrial research is at work to make them more flexible, lighter and more resilient. In the next decade they are likely to become products, and over the following years their price shall decrease to the point of becoming an option in the mass market.

T40: CRISPR/Cas9

CRISPR/Cas9 is a new technology but it is already widely used in research as well as in industry. It still has some hurdles to overcome, and there are new, but similar, technologies on the horizon.

The understanding of the connection between the genotype (that can be altered using CRISPR/Cas9) and the phenotype (the manifestation of the genotype) is still fuzzy. There are researchers analyzing these relationships, using artificial intelligence and big data, but we are still far from the possibility of a generalized correlation, and it may take many more years to be managed. This is why the timeline of CRISPR/Cas9 with respect to the generalized human augmentation indicates at least 20 years before it (or a similar alternative technology) can be used in designing a new augmented human.

T41: DNA modification

Today, DNA modification is using CRISPR/Cas9 to change the sequence of codons and modify/delete/add genes. In addition, scientists can add or delete entire chromosomes (although this is generally resulting in a living entity that cannot reproduce). By modifying the DNA one is effectively creating a new genome and potentially giving rise to a new species with a different set of characteristics—a different phenotype—with obvious ethical issues. Thus far, this has been used in humans to repair a degraded genome restoring it to the normal sequence, hence avoiding genetic diseases.
DNA modification can be made before fertilization (potentially affecting the egg and/or the sperm), but it can also be made on a living organism using viruses as vectors to change the DNA.

DNA modification will be used in the coming years to overcome genetic disabilities. By the end of the next decade it could be used in focused augmentation, once there is a better understanding of the relationship between the genome and the phenotype, e.g., to increase the resilience of a person in living in a specific environment (for example, some thinking is going on in modifying the genome of future astronauts to make them more resilient to radiation during long space travel).

In the far future, DNA modification can lead to overall augmentation with the potential of creating a new human species, human 2.0 or transhumans. Clearly this is fraught with ethical issues.

DNA modification is already an industrial reality in agriculture, genetically modified organisms (GMO), having resulted in species variations more resilient to specific environmental conditions, able to grow with scarcer irrigation and more resistant to bugs.

Any DNA modification is passed on to offspring, if the individuals is able to generate offspring after the modification.

T42: RNA modification

RNA modification can achieve basically the same result of DNA modification, but it is not passed on to offspring (since it is only the DNA that is involved in offspring generation). Hence, the ethical issues are softened.

RNA regulates the gene expressions, activating DNA specific genes and carrying the information from the cell nucleus to the cytoplasm where proteins are manufactured, hence resulting in the phenotype expression. From a technical point of view, the RNA modification is a bit trickier than DNA modification hence the slightly displaced timeline in the roadmap when compared to the DNA timeline.

The key point, as for DNA modification, is the ability to understand how a change in the RNA is affecting the phenotype. The growing availability of an RNA and DNA database and the development of artificial intelligence approaches to analyze those data is holding promise to identify the relationship between the genotype (and RNA genotype) and the phenotype. In terms of a capability for overall human augmentation (that is the reference used to draw the timeline) we are at the edge of the observation period of this white paper.

3.2.5 Augmented Humans (Cognitive)

Augmenting human cognitive capabilities has been achieved in the last millennia through education (both formal and informal). Enhancing learning capability (e.g., by increasing memory and becoming more effective in assimilating new concepts) is obviously a way to augment human cognitive capabilities. Increasing the capability to retain memory is also increasing cognition, however we do not understand completely how forgetting is actually part of learning (e.g., forgetting something by replacing with something else that is more valuable is actually a way to increase cognitive capabilities).

Augmenting cognitive capabilities is expected to progress from today's *enhanced capability to reach out* to *seamless connection to knowledge* eventually aiming at *enhanced brain capability* (to operate) *in a mixed reality* environment.

Several of these technologies have already been presented in the previous sections and will not be repeated here, namely:

T03: Optoelectronics. See Section 3.1.1
T04: Optogenetics. See Section 3.1.1
T05: Signal processing. See Section 3.1.1
T06: Artificial Intelligence. See Section 3.1.1
T07: Deep Neural Networks. See Section 3.1.2
T28: Extremely low power electronics. See Section 3.1.3
T31: Fluorescent proteins. See Section 3.2.3
T32: Smart materials for long term implants. See Section 3.2.1
T33: Neuroimaging. See Section 3.2.3
T34: BCI. See Section 3.2.1
T38: Implantable chips. See Section 3.2.4
T40: CRISPR/Cas9. See Section 3.2.4
T41: DNA modification. See Section 3.2.4

T42: RNA modification. See Section 3.2.4

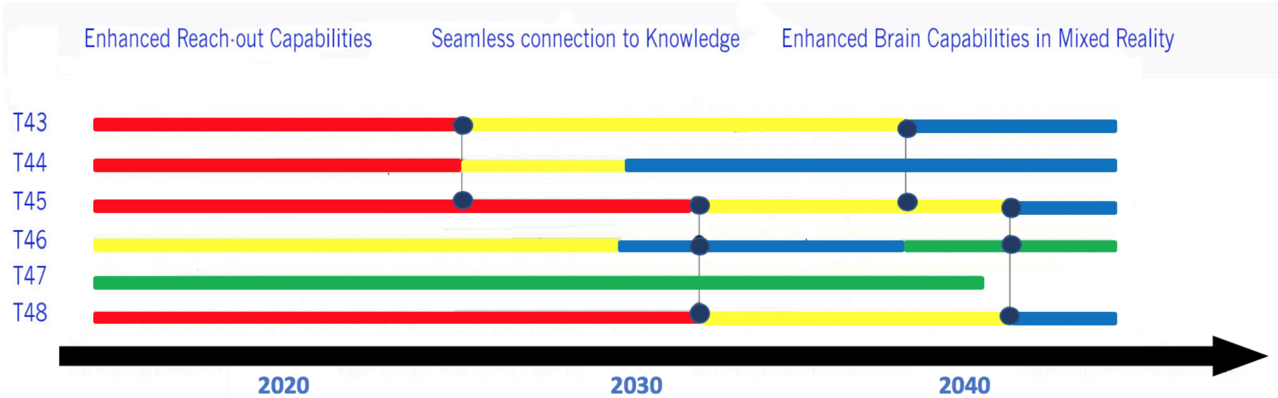

Fig. 3.14. Timeline of technologies for augmented humans (cognitive)

T43: Deep Brain Stimulation

Deep Brain Stimulation (DBS) is based on sophisticated electrodes sending electrical spikes into the brain and interfering with neuronal electrical activity, thus activating/enhancing or depressing activities in very narrow regions in the brain. The spikes are controlled by a computer, eventually implanted on the brain cortex, and the control is getting smarter and smarter, both by applying artificial intelligence algorithms and by receiving feedback from sensors, some co-located with the electrodes.

The tricky, and so far unsolved issue, is the difficulty in reaching the exact areas that need to be affected as there are usually several of them, and it is even more complex to maintain the effectiveness of the stimulation as brain plasticity shifts the areas that need to be affected.

Using DBS for cognitive enhancement is going to work only in very limited situations, like counterbalancing attention deficit hyperactivity disorders (ADHD). Research is ongoing to find more effective and flexible ways to provide electrical stimulation to brain areas although no silver bullet is in sight. Results in this area will be important in the design and development of cognitive prosthetics (T44).

T44: Transcranial Magnetic Stimulation

Transcranial magnetic stimulation (TMS) is like DBS aiming at interfering with brain neuronal activity using focused magnetic fields. Unlike DBS it is not an invasive procedure since it is done from outside of the skull, but it requires bulkier apparatus (to generate the resonant magnetic fields and to have them converging in the desired areas of the brain) making this technology suitable only for the lab. In addition, the spatial resolution and the effectiveness of interference with neurons is much more limited than DBS.

So far TMS is seen more as a way to study the brain than to alter its functions.
In perspective it may turn out that TMS may be used to affect the brain functionality by inducing selective rewiring of neural circuits but today is only at the level of speculation.

T45: Cognitive Prosthetics (2040 -2100)

Cognitive prosthetics are still in the fuzzy area between science and science fiction. Technologies (like chip implant, DBS, or optogenetics) are available to influence brain processing, and research is harvesting knowledge on several brain processes, like the physiology of memory. We are still far from a complete understanding that would allow us to look for the right technology and techniques to boost the brain by interfering with its processes.

The expectation, particularly from the multi-country brain initiatives that are running and should deliver results in the next decade, is that within 20 years there should be sufficient knowledge and technology to work on cognitive prosthetics.

T46: Neural Engineering System Design
A whole new science of neural engineering system design is needed. The first steps are being taken. The IEEE Future Directions Brain Initiative[11] is a step in that direction. A convergence and collaboration from a variety of disciplines are required to tackle the complexity of designing "a brain". Its effects are expected in the fourth decade of this century.

T47: Smartphones (present-2040)
Smartphones are today the most effective cognitive prosthetics available. Actually, they are so effective that quite a few people are concerned that their use is actually decreasing our natural brain cognitive capabilities (you no longer need to remember a number or information since you can turn to your cellphone to get what you need and rely on it to remember as your proxy). Smartphones will continue to get better and more effective in the next decade and then will start to fade away replaced by wearable and ambient communication fabric from 2040 on.

T48: Symbiotic intelligence (2040)
The amazing amount of data, applications, information and knowledge that keeps accumulating on the web is shifting the focus from cognition as a characteristic of humans to the possibility to reach cognition by accessing the web. We are moving, slowly but inevitably, from a focus on personal knowledge to a focus on accessing knowledge.

This has profound implication on industry, education and society. Companies are getting more and more conscious of the need for a continuous update of their "company know-how" and while in the past this was distributed among their work force today it is partly in cyberspace with the workforce seeing it as a tool to access that knowledge. Clearly the competitive advantage of a company is both in having a smart work force able to access and apply that knowledge and in having knowledge owned solely—at least for a certain period of time—by the company.

The advent of Digital Twins is also reshaping the knowledge and cognitive landscape, where a company may put claims on the knowledge Digital Twin of their work force (for example, patenting knowledge) and keep using it once that employee leaves the company.

At a personal level there is a growing need in knowing where knowledge is, how to access it, and how to apply it rather than a need for learning knowledge. We are moving towards a symbiotic intelligence that will be in full swing in the second part of this century but whose first signs, and implications, are already showing up today.

3.3 Symbioses

We are already cooperating with machines. Over the coming years this cooperation will become more and more seamless to the point that we might not even perceive it; we will take it for granted. The next step is machines becoming aware (including aware of our presence and capabilities) and adapting their operation to the overall ambient. Some implants will become much smarter than today, adapting in a seamless way to the body, and conversely the body will adapt seamlessly to the implant. In the fourth decade we can expect this mutual adaptation,

relying on seamless interfaces and low latency communications, to broaden beyond implants to components in an ambient that will operate in a symbiotic relationship.
Intelligence will become a distributed capability giving rise to an emergent symbiotic intelligence.

Digital Twins will be enablers in this evolution bringing physical objects, including humans, to inhabit and interact in cyberspace. A true symbiosis is clearly far away, possibly beyond the observation horizon of this White Paper. Nevertheless, it is felt that this is the ultimate destination; for what can be imagined (remaining in feasible science) and some aspects of symbioses will be manifested within this White Paper time span.

For this reason, it is necessary to point out that the timelines presented are with respect to a symbiotic arrival. Those technologies that are fine today as they are, are indicated by a green line. What we have today would fit the need for tomorrow (although of course what we have today will be much more advanced in twenty years).

On the contrary, a technology that works fine today, meeting today's requirements may be considered totally insufficient for tomorrows challenges, and if so much more research is required and it will be indicated by a red line.

Another point is the inclusion among technologies of some "strange beast", like Neuralink[12] that is not a technology but a company whose goal is the creation of a seamless effective brain computer interface. That is exactly what would be needed for a human-machine symbiosis, and therefore it is included in the technology list. Notice that in this specific case, Neuralink is aiming at delivering workable products in a short period of time, like 5 to 10 years. However, they state that reaching the full goal of the symbiotic target will require much longer and therefore most of the timeline is characterized in red.

Another technology that has been included is counterfactual quantum entanglement (CQE). Some aspects of the brain are subject to quantum phenomena (including the activation of rhodopsin by photons and some synaptic interchange). CQE is taken as an example of quantum-based technologies that may eventually find a place in symbiotic systems, although their roles are currently uncertain.

That being said, one such CQE-based application is counterfactual quantum communications (CQC), which allows—and has demonstrated—quantum entanglement without entangled particles interacting and are secure without cryptographic keys. CQC networks will make instantaneous synchronization rather than relying on classical signal transmission as used today—a powerful feature for Symbiotic Autonomous Systems and, in particular, Digital Twins. Nevertheless, a significant challenge is the neurobiological technologies that will allow the human twin to interact with his/her Digital Twin via CQC, coupling neural structures to a CQC network will—while not *de facto* impossible—require neuroscience advances and technologies beyond our current capabilities.

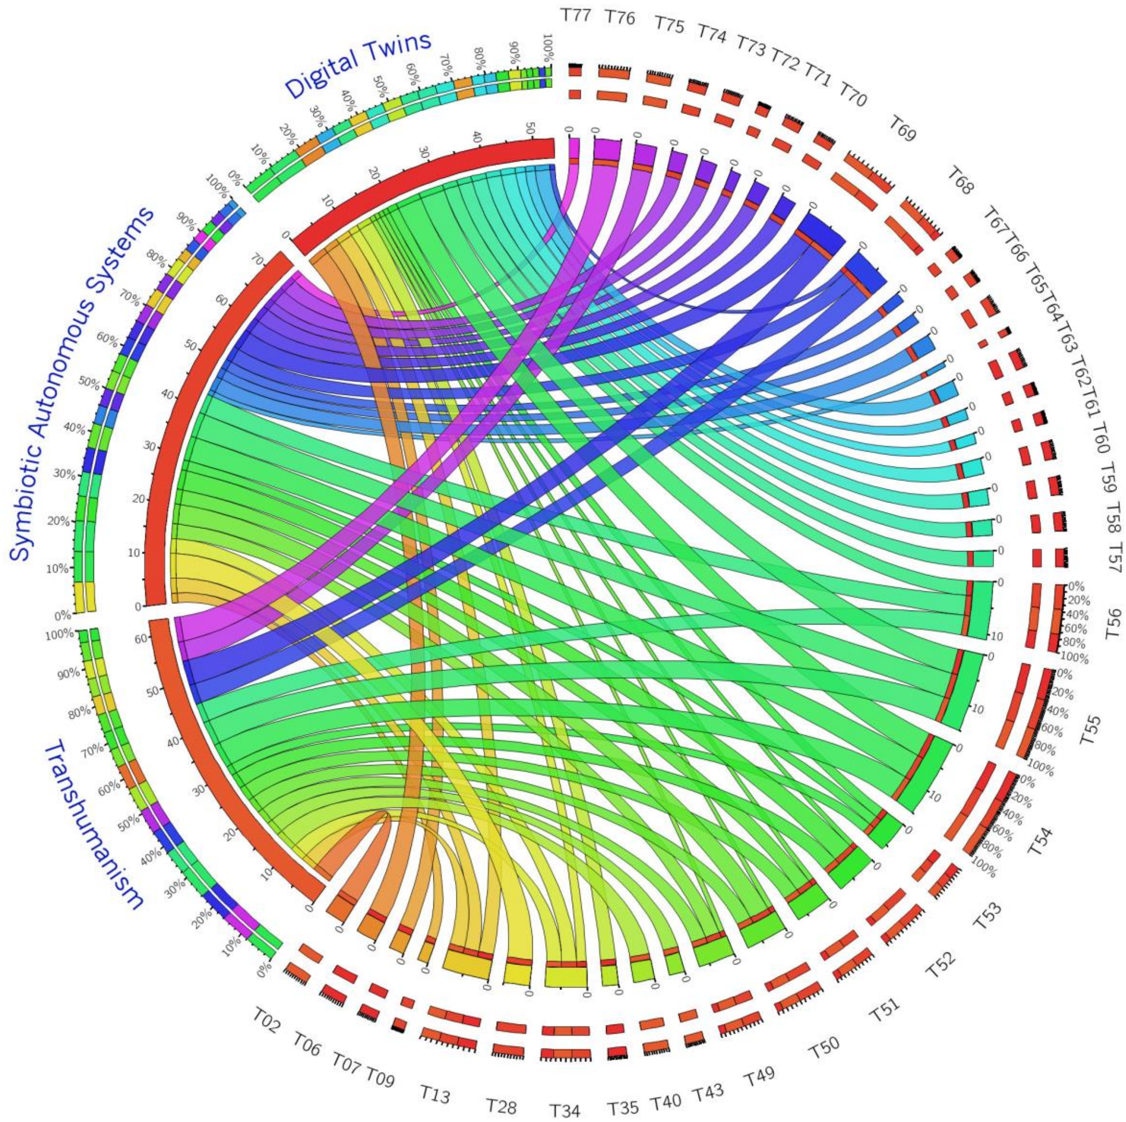

Fig 3.15. Symbioses technologies

T02	Nano-biotechnologies	T52	Symbiotic Life Design	T65	Shared Intelligence
T06	Artificial Intelligence	T53	5G-6G	T66	Augmented Data Discovery
T07	Deep Neural Networks	T54	Security	T67	Prescriptive Analytics
T09	Convolutional Neural Networks	T55	Cybersecurity	T68	BrainInternet
T13	Sensors	T56	Cyber-physical Security	T69	Artificial Human Super Intelligence
T28	LPWAN	T57	Virtual Twin	T70	Conversational Analytics
T34	BCI	T58	IoT	T71	Embedded Analytics
T35	Virtual Reality	T59	Sensors Network	T72	IoT Edge Analytics
T40	CRISPR/Cas9	T60	Sensors Cloud	T73	Advanced Anomaly Detection
T43	Deep Brain Stimulation	T61	Avatar	T74	Citizens Data Science
T49	Neuralink	T62	Real Time Simulation & Data analysis	T75	Artificial Super Intelligence ASI
T50	Nanobots	T63	Augmented Reality	T76	Genomic Engineering
T51	Microbots	T64	Counterfactual quantum communications	T77	Human in the loop - crowdsourcing

Table 3.3. Technologies fostering/enabling Symbioses

3.3.1 Digital Twins

Digital Twins have been used for a few years now to mirror complex objects like turbines and vehicles digitally (in bits). There are a few signs indicating their expansion to mirror many more entities, from cities (Singapore in 2019[13]) to humans. A Digital Twin represents, in bits, an aspect of the real entity. In a human it might represent the health status, the knowledge, the set of relationships, and so on.

In the context of Symbiotic Autonomous Systems, the functional area of a Digital Twin allows for the possibility of augmenting human capabilities and entering into a symbiotic relationship with other entities by creating a link through cyberspace. Likewise, Digital Twins of smart machines make possible their augmentation and symbioses through cyberspace.

It is expected to see an evolution of Digital Twins from today's *mirroring of physical entities* to *understanding the physical entity life-cycle* up to *enhancing as proxy the physical entity* entering into a symbiotic relationship with other Digital Twins.

Several of these technologies have already been presented in the previous sections and will not be repeated here, namely:

T06: Artificial intelligence. See Section 3.1.1
T07: Deep neural networks. See Section 3.1.2
T13: Sensors. See Section 3.2.2
T34: BCI. See Section 3.2.1
T35: Virtual reality. See Section 3.2.1

Given the number of involved technologies they are presented in two clusters, one related to networking aspects, the other to application and architectural aspects.

Networking aspects:

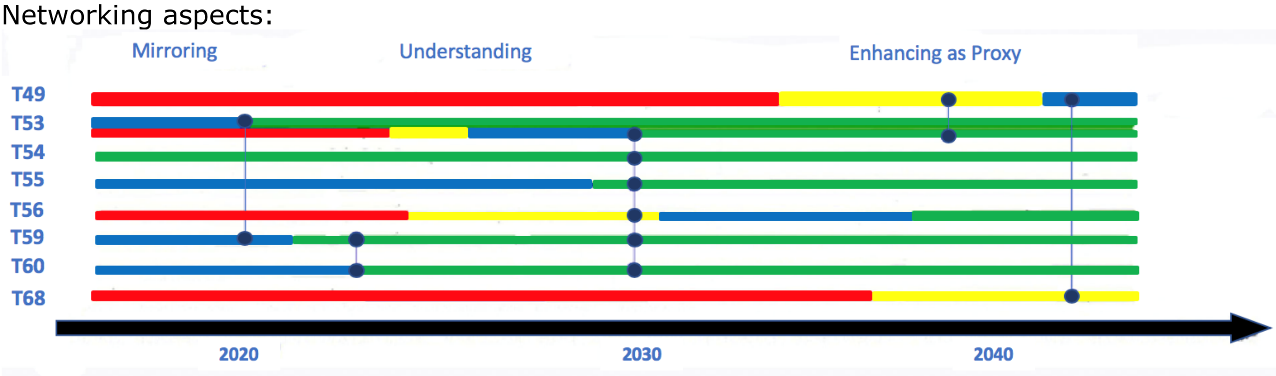

Fig. 3.16. Timeline of technologies for networking aspects of Digital Twins

T49: Neuralink

Neuralink[12] is a company and an initiative launched by Elon Musk. It is mentioned here because it aims at creating an ultra-fast seamless, implantable interface between a brain and a computer. This is pursued by addressing a mix of technologies. Looking at the ambition, it is not expected to move into a consolidated product within the time frame of this white paper. A huge and lengthy amount of pure research is needed to match the ambition. However, and this is the expectation of Neuralink, by-products from this endeavor can be expected in much shorter time frames. It is important to notice the ambition since if fulfilled it would provide a strong underpinning to human machine symbioses.

T53: 5G – 6G

5G and 6G are the next generations of wireless radio, with 6G following 5G approximately 10 years later. The timeline has been drawn with two lines, with the upper one referring to 5G and the lower one to 6G.

5G has almost reached the deployment phase, several industries are committed to its deployment, and it will become state of the art in the next decade, co-existing along 3G and 4G. 2G is expected to fade away in the first years of the next decade with re-allocation of its frequencies to 5G. 5G will provide IoT and devices in general with an advanced connectivity fabric. 6G will go a step further allowing IoT, devices and any object to create a local network that dynamically integrates with others. This fabric will provide an ideal set up for Symbiotic Autonomous Systems.

T54: Security

The basic concepts and technologies related to computer and systems security have been extensively studied and adopted for decades in most industries. Infrastructure protection, which is linked to the more general frameworks of defense and homeland security, includes risk modeling and assessment, planning for business continuity and disaster recovery, monitoring and surveillance, early warning and situation assessment, emergency response and crisis management. Those aspects can be considered mostly in the green phase of their life-cycle, although evolutions are expected in the field of artificial vision (i.e., video content analytics), information fusions and decision support systems, drone surveillance, and cyber-intelligence, as an effect of evolutions in AI and autonomous systems.

T55: Cybersecurity

Cybersecurity is also a rather well developed field. However, due to the rapid growth of complexity in computer-based systems, novel threats are continuously emerging (including the so-called zero-day attacks exploiting unknown vulnerabilities) that need to be addressed by smarter technologies. Some of those technologies are already available, like e.g., sandboxing techniques and other intrusion detection systems used to prevent unknown threats from attacking information systems. Nevertheless, many still need to be researched and developed, including intelligent security information and event management (SIEM) based on heuristics and machine-learning approaches. We can consider those efforts to be in the blue stage, partly linked to the evolutions in the field of AI. It is important to underline that cybersecurity of Internet connected devices also serves as an enabler for new SAS paradigms like Digital Twins, since many of those evolutions will not be viable if they are not deemed trustable. This is particularly important when dealing with automation and replication based on sensitive data like the ones addressing biometrics and health.

T56: Cyber-physical Security

While most embedded computing systems feature basic security mechanisms (e.g., password protection, encryption, etc.), heterogeneous and ubiquitous cyber-physical systems in critical applications still represent a challenge for security assessment, management and certification against international standards. It is expected that security standards and technologies for cyber-physical systems will evolve over the next 10-20 years as a response to the security requirements of IoT, Industry 4.0 and intelligent transportation systems, among others. Hence, those emerging approaches and paradigms can be located in the red-yellow phase, i.e., the early stage of their evolution. That is especially true for the security technologies involving strategic cooperation among autonomous agents in order to enhance diagnostics, prognostics and application of security countermeasures; those are expected to become operational no earlier than 20-30 years from now.

Note to T54-T55-T56

The paradigm of Digital Twins can also be leveraged in the long term as an enabler of pro-active dependability and security, allowing for the simulation and prediction of threats against SAS by continuously running models of the real-systems in virtual yet realistic environments as well as real operating conditions. Those fully featured parallel software replicas operating in the cloud will probably require decades to become a standard due to the need for increased computing power, storage and ultra-high bandwidth communication facilities (including 5G and beyond), but the research community will soon start to address this challenge.

T59: Sensor Networks

Sensor networks are already a reality in some areas by design, meaning that a number of sensors have been designed to form a local area network used to harvest data from each sensor and reach a gateway (in some cases a sensor with easier access to power will serve as a gateway). Industry research is striving to create sensors (IoT) that can autonomously create networks. This is an important component for future Symbiotic Autonomous Systems. Their availability is targeted in the next decade, and it is expected to become a normal feature of any sensor: once active the sensor is able to poll the environment through a self-created network space and connect with other network spaces to form a mesh network. 5G is going to be an enabler.

T60: Sensor Clouds

Sensor clouds are fuzzier sensor networks where data play the crucial role. The underpinning remains the sensors' network but the focus shifts to data. The underlying sensors' network will change its structure dynamically although the sensors' cloud will remain stable and applications will refer to the cloud, no longer to individual sensors addressed via the sensors' network. In a way it is the implementation of a Digital Twin for a cluster of sensors.

T68: Braininternet

Braininternet[14] is the idea of connecting a brain to the Internet using BCI in real time. The idea is to look at the brain as an IoT, and some experiments have been carried out at the University of the Witwatersrand. The implementation is tied to the availability of an effective BCI, plus it needs to overcome serious hurdles in the isolation of signals and their parallel processing. Today we are very far from solving the basic issues, and research will be needed to address them. By 2035 it can be expected that the basic building blocks may become available. At that point more research will be needed to integrate the various pieces, and it is unlikely that an industrial involvement could happen within this White Paper time frame.

Clearly, this kind of technology would give a boost to Symbiotic Autonomous Systems with a human component.

Applications, Architectural Technologies:

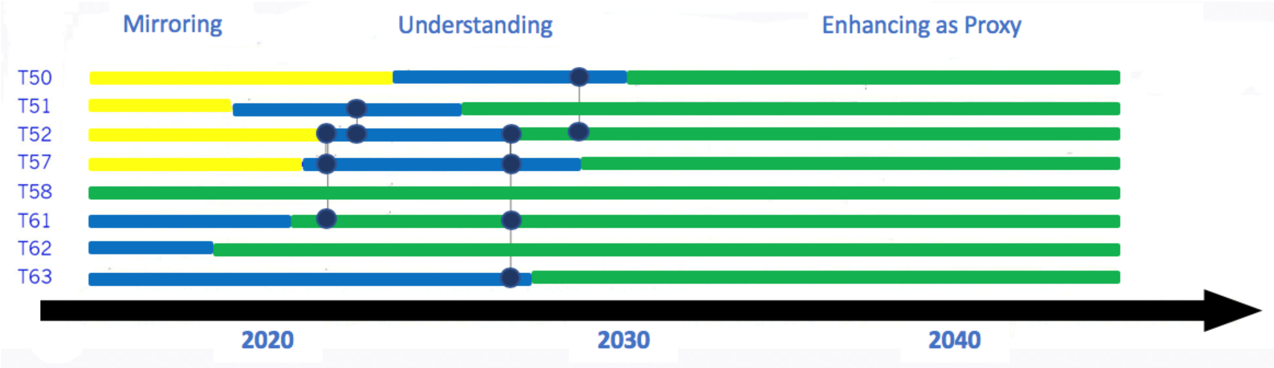

Fig.3.17. Timeline of technologies for applications, architectural aspects of Digital Twins

T50: Nanobots

Nanobots are an active area of research in academia today. A specific focus is on the potential application in the health care domain (for example, on-spot drug delivery). Some experiments are being carried out using iron particles bounded to a drug in the shape of beads that can be steered to the target area by magnetic fields. Autonomous nanobots are also being investigated, with some basic research creating nano-motors for autonomous movement. Industry is likely to get involved in the next decade with applications starting in the fourth decade. Nanobots can be one of the technologies used in the chemical communications within Symbiotic Autonomous Systems.

T51: Microbots

Microbots are leveraging MEMS technology. They are still in the academia research stage but a bit more advanced than nanobots. We can expect industry involvement at the turn of this decade and some applications starting by the middle of the next decade.

Microbots may be active components in Symbiotic Autonomous Systems, in some cases acting as the glue for the other components. Some early applications can be expected in prosthetics for health care.

T52: Symbiotic Life Design

A science of symbiotic life design is sorely needed. Some academic research connected to CRISPR/Cas9 and similar techniques is occurring. Artificial intelligence software is being used to link the genotype to the phenotype. The goal is to be able to start from the phenotype to identify the required genotype that in turn might be assembled from existing DNA strands and modified as needed. Construction from scratch is quite far away and might not even be of interest since it is much more effective to start from a library of DNA snippets and assemble them together with limited variations.

Notice that the very idea of starting from a library of life forms and creating a sort of superorganism with the desired characteristics is actually the creation of a symbiotic organism. Humans, as most multicellular life, are symbiotic organisms; our cells live in symbioses with an overwhelming population of bacteria. A single cell is a symbiotic organism, living in symbioses with mitochondria that are bacteria embedded in the cytoplasm. Symbiotic life design is supporting the design and creation with use of nanobots and microbots providing the required architectural framework. By the fourth decade of this century we might expect to have the application of symbiotic life design in a variety of fields.

T57: Virtual Twin

Different than a Digital Twin that mirrors a real entity, a virtual twin mirrors a symbiotic organism and in particular the relationships that are the underpinning of the symbiotic organisms.

The first applications are already being studied when creating meta-Digital Twins, like a Digital Twin representing a city not as a modelling of the city itself, rather as the assembly of various Digital Twins forming a city. While a Digital Twin has a corresponding real twin, a virtual twin mirrors a set of relationships among Digital Twins and only indirectly those of the real entities represented by the Digital Twins. Obviously a virtual twin is valuable as long as these representations are significant in reality.
Research is ongoing today in academia, and industry is likely to become involved in the next decade as more and more Digital Twins will operate in complex environments thus providing the backstage where a virtual twin would be valuable (as an abstraction of a multitude of Digital Twins, like a city is an abstraction of a variety of interplaying components). We can expect virtual twins to be used to mirror Symbiotic Autonomous Systems with applications in the fourth decade of this century.

T58: IoT

IoT have been the first entities to have a Digital Twin, even though implicitly, in the sense that each IoT has a mirroring digital value in cyberspace (the global or in a local one). It is an already well established area (hence the green line) and will be an integral component of Symbiotic Autonomous Systems. Of course, the number of IoT will keep growing and artefacts like a prosthetic and its components that today are rarely seen as an IoT will eventually become part of the IoT world in the sense of making use of the same tools.

T61: Avatar

An avatar may be considered as sort of Digital Twin, being a representation. A symbiotic autonomous system may extend to include avatars, one representing it as a whole and/or several representing some of its components.

Avatar technologies are already well established and will keep improving in the coming decades as their use will extend, and more representation technologies will become available (in particular 3D display technologies).

T62: Real time simulation and data analysis

Real time simulation and data analysis is a well-established toolkit of technologies and techniques used by academia and industries. Big data and IoT have given a further impulse in this area, and new algorithms are continually being investigated. The area of Digital Twins will need this kind of technology for moving to the next phase of "understanding".

T63: Augmented reality

Augmented reality has been around for a few years now, yet it is far from being the seamless experience that it is supposed to be. Lack of appropriate devices is hindering its progress, and industrial research is at work to fill the gap. The latest goggles are still falling short of a seamless experience. By the end of the next decade, the first electronic contacts lenses delivering effective AR are expected to be on the market.

Augmented reality will make avatars more effective as well as providing visual representation of virtual twins (and their components). It will also be used as a tool supporting symbiotic life design as today virtual reality is supporting molecular design.

AR will be an important technology to exploit Digital Twins.

3.3.2 Symbiotic Autonomous Systems

In nature there are plenty of Symbiotic Autonomous Systems; actually it would be difficult to identify an area of life that is not based on symbioses at a micro (mostly everywhere) or macro level.

The Symbiotic Autonomous Systems addressed in the IEEE SAS initiative and in these white papers are “artificial ones”, artificial at least in the construction of the symbiotic relationship, although in most cases these Symbiotic Autonomous Systems contain artefacts as components (as an example a symbiotic system formed by a human and a prosthetic where the prosthetic is clearly an artefact). Because of this, the technologies that are considered for Symbiotic Autonomous Systems are the ones forming the artificial glue among the different parts, the ones that create or make the symbioses possible.

Three phases can be recognized from *symbioses by design* to *dynamical opportunistic symbioses* (where the symbioses happen as result of interplay of the various components by their own volition) and finally the forming of *super-organisms*.

Several of these technologies have already been presented in the previous sections and will not be repeated here, namely:

T09: Convolutional neural networks. See Section 3.1.2
T28: Extremely low power electronics. See Section 3.1.3
T49: Neuralink. See Section 3.3.1
T34: BCI. See Section 3.2.1
T50: Nanobots. See Section 3.3.1
T51: Microbots. See Section 3.3.1
T52: Symbiotic life design. See Section 3.3.1
T53: 5G – 6G. See Section 3.3.1
T54: Security. See Section 3.3.1
T55: Cybersecurity. See Section 3.3.1
T56: Cyber-physical security. See Section 3.3.1
T68: Braininternet. See Section 3.3.1

The timelines are split into a first set dealing with supporting technologies and in a second set dealing specifically with systems having a human component.

Supporting technologies:

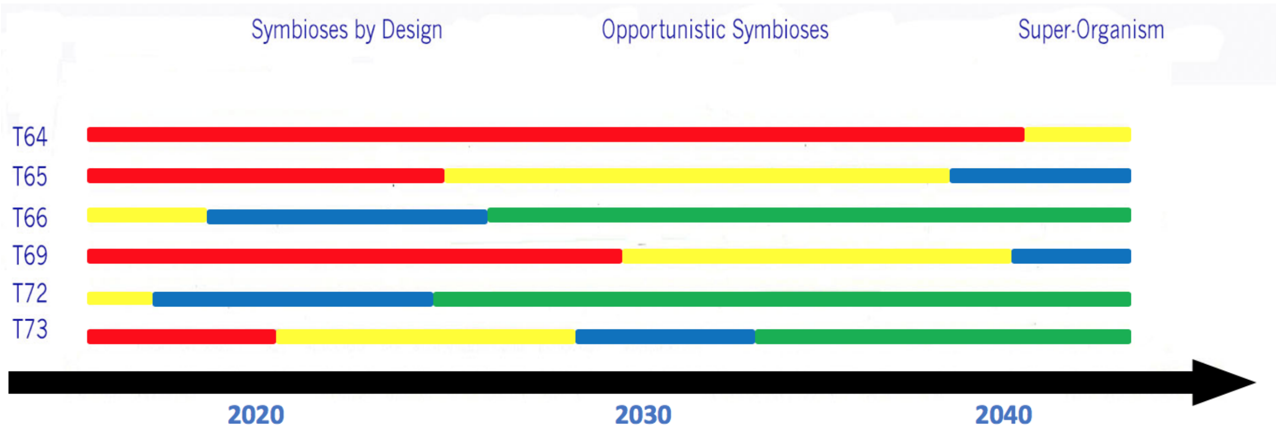

Fig. 3.18. Timeline of technologies of supporting tools for Symbiotic Autonomous Systems

T64: Counterfactual Quantum Communications

Counterfactual quantum communications allows—and has demonstrated—quantum entanglement without entangled particles interacting and are secure without cryptographic keys. Counterfactual quantum communications networks will make spatial- and temporal-agnostic communications possible with instantaneous synchronization rather than relying on classical signal transmission as used today.
It is, as most quantum based technologies today, at the stage of academia research (although some kind of quantum encryption technologies are already used in niche products). The application of this technology to Symbiotic Autonomous Systems and the need for them is still an open question. By the end of this White Paper observation period it is likely that quantum technologies may have become more concrete (in the sense of usability) and may play a role in communications within a symbiotic system. That said, there have been several successful laboratory demonstrations of counterfactual quantum communication—and the fact that counterfactual quantum entanglement was recently proposed in 2009 is encouraging.

Nevertheless, a significant barrier is the neurobiological technologies that will allow the human twin to interact with his/her Digital Twin. Although there are known instances of the role of quantum functions in biological environments, coupling neural structures to a counterfactual quantum communication network (a future version of today's early quantum networks) will—while not *de facto* impossible—require neuroscience advances and technologies beyond our current capabilities.

T65: Shared intelligence

Once artefacts have an autonomous intelligence they will also probably have seamless interaction capabilities that will enhance their local intelligence by making use of other entities' intelligence. In the area of Symbiotic Autonomous Systems this is going to play a role in all three stages, symbioses by design where the sharing of intelligence will be designed, in opportunistic dynamic symbioses where other entities' intelligence will be exploited and of course in the super-organism phase where the shared intelligence is a fundamental aspect of the super-organisms. Notice that at this phase the shared intelligence may take the form of an emergent intelligence of the super-organisms, i.e., it may no longer be possible to recognize the individual components of this intelligence.

T66: Augmented data discovery

Big data has created a broad toolkit of technologies and techniques to create meta-data (data derived by the analysis of other data). These meta data provide (additional) meaning and, in the area of Symbiotic Autonomous Systems they may be the result of the analysis of data related to each single component. By analyzing the interaction and the individual data set, these techniques can provide insight on the symbiotic aggregation and eventually an understanding of the super-organisms. These augmented data may also provide a feedback to each individual component creating a holistic view of the system, also mirrored by the virtual twin.

Research is expected to progress both at the academic and industrial level throughout the next decade.

T69: Artificial Human General Intelligence
There is also another twist to artificial general intelligence and that is a super intelligence of the human species achieved through selective breeding (of humans), nootropics (and other types of substances affecting the brain processes) which are drugs to augment intellectual capability, genomic manipulation to evolve the human species by design, and epigenetics modulation (becoming more intelligent due to environmental factors, like advanced education).

It is foreseeable that this form of intelligence related to humans will become reality in this century and it will be a characteristics of transhumanism.

T72: IoT Edge analytics
As IoT keeps multiplying and densifying, it makes more and more sense to perform data analysis at the edges creating a sort of hierarchical data analysis that might result in local action, also diminishing the need for absolutely reliable communications towards the Internet. This implies the development of shared intelligence strategies and involves principles from small worlds, complex systems and symbiotic life design.
Notice that this is what happens in living systems where centralized and decentralized analysis co-exists.

The evolution in this area is expected to be fast with a well-developed toolkit of technologies, techniques, and architectures available in the next decade. Symbiotic Autonomous Systems will rely heavily on these edge analytics.

T73: Advanced anomaly detection
Symbiotic Autonomous Systems are complex systems, and the detection of anomalies is a challenging issue, particularly since each component is autonomous and in principle its behavior can be unpredictable. In addition, there might be situations where the behavior of each component is normal and yet the resulting interplay can be abnormal leading to an undesired behavior at system level.

New approaches to anomaly detection are therefore required and research is needed both at system design, aggregation frameworks, and super-organisms modelling.
We can expect basic research and academic research to yield results (in terms of algorithms and architectures) in the next decade, and industry to determine implementable solutions in the first years of the fourth decade, even though, of course, earlier solutions will be implemented.

Data analytics involving a human component:

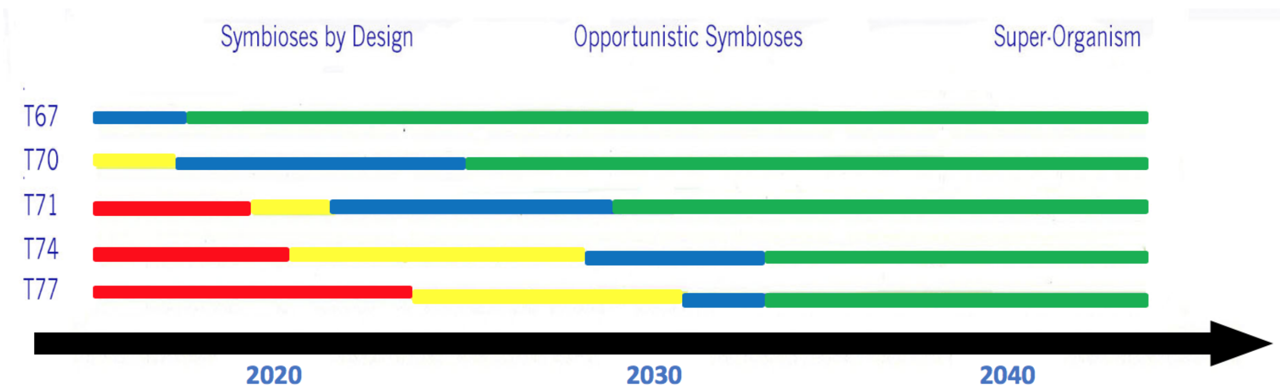

Fig. 3.19. Timeline of technologies in data analytics involving a human component for Symbiotic Autonomous Systems

T67: Prescriptive analytics

Prescriptive analytics advise on the course to follow from the available data. It is usually presented as the third and most advanced stage of analytics, after descriptive (analyzing a past and present body of data) and predictive (defining how a given situation is or will be evolving) stages. Providing different recommendations and strategies based on all the evidence given by data, prescriptive analytics immediately precede human decision-making. Taking account of historical, contextual, transactional and real-time data streams, it is being used for business, military and medical decisions and diagnoses. Although still under development in many industries, it is likely to expand considerably over the next ten years. It is expected to find applications in both symbioses by design and opportunistic symbioses.

T70: Conversational analytics

Conversational analytics begins by transcribing speech (oral or written) into data. It is then structured to extract insights based on word frequency, topical emphasis, and thematic selection. It is used for business as well as security purposes. Several new entertainment platforms have also emerged in the last three years putting conversational analytics to data culled from existing conversations so as to approximate the character of real persons in chat-generating systems. Symbiotic Autonomous Systems are expected to engage in some instances where a human component is involved in natural language conversation, and the artefact components will leverage conversational analytics.

T71: Embedded analytics

Embedded analytics are integrated within the tools developed for data analysis in whatever context, business, administrative, healthcare or government. They are based on data processing formats adjusted to the specific needs of the user. For example, in customer relationship management (CRM), embedded analytics provide a permanently self-updating of data collecting, sorting and interpreting available information about clients. Embedded analytics aggregate streams and combinations of data that are relevant to a specifically tailored service. This target-precision allows the introduction of AI into common applications, and it is expected to be used by smart components in a symbiotic autonomous system. It will take a few more years for its application in the SAS area, requiring sufficiently smart autonomous components that can engage in embedded analytics internally rather than relying on external support.

T74: Citizen data science

There are two possible meanings to citizen data science, a recent addition to the data analytics category. The first, presently more commonly used, is the practice of some businesses, facing the shortage of professional data scientists, to train unskilled members of their organization to use seasoned data analytics tools for simpler tasks. This practice opens the door to a new crop of targeted analysis and inspired citizens and especially data science students to create new systems. The other meaning, although less current, is far more important because it involves the study of citizenry with various technologies such as population dynamics, epidemics, sentiment analysis and other data analytics to estimate the needs of whole communities and groups.

This set of techniques is expected to be used by Symbiotic Autonomous Systems in their external communications.

T77: Human-in-the-loop crowdsourcing

Human-in-the-loop takes crowdsourcing to the next level by including provisions and tools as well as access to datasets to include human expertise and judgment from the crowd to refine large and complex multifactor decision-making. Bringing humans into refining the data analytic stage not only permits focusing on the most pertinent data, but also mitigating the often crudely automatic outcomes of machine learning. In the future, it could become a necessary precaution to take before final and critical decisions are made. Currently, there is a discussion about the possibility of applying the 80/20 rule about algorithmic processes so as to insure human monitoring and, eventually, correcting decisions taken by automated data analysis. These techniques are expected to find a way in the Symbiotic Autonomous Systems in their understanding of the external world. They should become applicable in opportunistic symbioses in the later super-organism stage.

3.3.3 Transhumanism

Transhumanism is at the border of philosophy, sociology and science with each having their input. It is becoming a scientific area as technologies are becoming available that accelerate the evolution of the human species. This is occurring right at a time when most biologists agree that the human species cannot evolve, given the absence of conditions supporting evolution through natural selection.

Technology may take over and push humans towards humans 2.0. It is expected that evolution will move from today's first signs of technology *augmented humans* to an *augmented society* (including broader epigenetic factors influencing single humans) to a full *new species* (through genetic engineering) well beyond the present White Paper observation horizon.

Most of these technologies have already been presented in the previous sections and will not be repeated here, namely:

T02: Nano-biotechnologies. See Section 3.1.1
T40: CRISPR/Cas9. See Section 3.2.4
T43: Deep Brain Stimulation. See Section 3.2.5
T49: Neuralink. See Section 3.3.1
T34: BCI. See Section 3.2.1
T50: Nanobots. See Section 3.3.1
T51: Microbots. See Section 3.3.1
T52: Symbiotic Life Design. See Section 3.3.1
T53: 5G – 6G. See Section 3.3.1

T54: Security. See Section 3.3.1
T55: Cybersecurity. See Section 3.3.1
T56: Cyber-physical security. See Section 3.3.1
T68: Braininternet. See Section 3.3.1
T69: Artificial Human General Intelligence. See Section 3.3.2

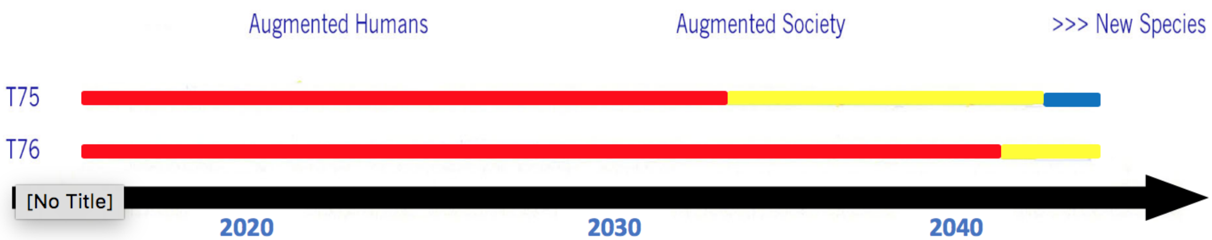

Fig. 3.20. Timeline of technologies for transhumanism

T75: ASI – Artificial Super Intelligence

Today artificial intelligence is not au pair with human intelligence; humans often have the upper hand. However, there are areas where computer intelligence is better than the human one, as an example analyzing or remembering a huge amount of data or evaluating the outcome of some complex decisions. Computer AI has managed to beat the human chess world champion, the human Go world champion, and has won Jeopardy.

Hence it would be fair to say that although we do not have computers that can demonstrate the same level of human intelligence in general (AGI) we have specific areas where computer intelligence is better than the human one. There is currently a consensus that sometime in the future computers will demonstrate an intelligence comparable to the one of humans; they will achieve AGI (although there is no consensus on when this will happen). Paradoxically, this will be the point when they will also achieve Artificial Super Intelligence (ASI), since they will maintain the edge in those areas where they already have an edge on human intelligence, hence by the time they will demonstrate AGI they will also demonstrate ASI.

However, ASI is expected to become vastly superior to human intelligence and some is even pointing to an equivalent IQ in the thousands (vs the 250-300 that is considered a maximum for human intelligence). Computers are getting more intelligent by learning, and they no longer learn from humans. They are starting to learn also by themselves, by trying and evaluating different approaches.

The advent of Symbiotic Autonomous Systems, where there will be a computer as a component, will lead to the emergence of intelligence at the SAS level, and this emergent intelligence is most likely to be better than the intelligence of each component. If one of the components is a human being and the other component is also intelligent, the emerging intelligence is likely to be of ASI type.

This will be unlikely to happen within the present White Paper horizon (this is the reason why the ASI timeline is red throughout the observation period, indicating ongoing academic research on how it could be achieved), although a few scientists and futurists bet on this transition to ASI to happen in the fourth decade of this century.

T76: Genomic engineering

Genomic engineering using CRISPR/Cas9 is already a reality. Genetically modified organisms (GMO) have been around for over a decade, and lately technology for genomic manipulation has progressed significantly.

As noted discussing CRISPR/Cas9, we are still far from understanding the full implication of the genome (and hence genome modification) on the phenotype, hence the type of modifications that we can make are quite constrained (and are potentially raising concerns).

In the third decade it is expected that there will be much better control and understanding of the relationship between genotype and phenotype, and industry will start using this capability in manufacturing a variety of products, including smart bio materials. Genomic engineering can potentially be a very powerful tool in the design of Symbiotic Autonomous Systems. This is unlikely to happen before the fourth decade of this century, but it will happen eventually.

4 Technology Evolution 2030-2050

This section expands on the technology overview provided in the first white paper, taking into account the most recent evolutions and forecast studies.

4.1 Integrative Transdisciplinary Capabilities

Nanobiotechnology and bionanotechnology

As structural and functional aspects of varied science and technology domains increasingly interact, interdisciplinary research and development transforms into a transdisciplinary phase in which these previously separate disciplines merge. As a result, new unitary fields of science and technology emerge. Perhaps the most well-known such transformation is from biotechnology and nanotechnology to nanobiotechnology. That said, it should be noted that some researchers hold, perhaps temporarily, that we are in a dual transdisciplinary phase in which bionanotechnology—in which natural biosystems are used to develop biomimetic nanomaterials—and nanobiotechnology—in which in this case is defined as the used of nanomaterials in biotechnological applications—coexist. Two currently emerging transdisciplinary fields are counterfactual quantum communications and bio/techno convergence:

- Counterfactual quantum communications[15,16,17] allows—and has already demonstrated[18,19] —the ability to establish quantum entanglement without the entangled particles having to interact directly without physical particles travelling between them, and moreover are secure without having to use cryptographic keys. When widely deployed in counterfactual quantum communications networks, it (enhancements of existing and in-development, quantum communications networks) will revolutionize secure spatial- and temporal-agnostic communications.
- Bio/techno convergence, as is discussed elsewhere in this white paper, will give rise to the most transformational change in *H. sapiens* over the next several decades. Our discrete biological and associated technological applications (prosthetics, implants, sensory amplifications and replacements, etc.) will give way to a complete integration of biology and technology, leading to the transdisciplinarian emergence of a unified transhuman transformation in which biology and technology are not only intertwined, but integrated into a *de novo* integrated cyborg evolution.

4.2 Artificial General Intelligence and Affective Computing

4.2.1 Artificial General Intelligence (AGI)

Mimicking Human Intelligence

Also referred to as strong AI, and unlike current (weak) AI, AGI is human-analogous cognition and intelligence. This entails human-like properties, including sentience, self-awareness, self-image, and consciousness. The historic difficulty of identifying specific neural correlates of these properties and associate interpretations positing that these properties are qualia may be the result of these properties being our experience of our neural activity rather than being associated with a specific brain area—an example being color, which we perceive as a retinal "interpretation" of frequencies of light, which do not themselves have color. (In short, "redness," for example, does not exist in and of itself.). That said, recent research shows that the anatomical neural correlates of consciousness are primarily localized to a posterior cortical hot zone that includes sensory areas, rather than to a fronto-parietal network involved in task monitoring and reporting[20]. The goal of giving computers a form of intelligence similar but analogous to our own has a vital implication for fully-functional Digital Twins being able to transcend sharing *information* between neurons and bits and to provide a common BioDigital *understanding* of that information.

4.2.2 Affective Computing

An endeavor complementing and relating to AI/AGI, affective computing is the development of technology capable of recognizing, interpreting, processing, and simulating (and eventually actually possessing) human or human-analogous affect, the goal being to provide "computing that relates to, arises from, or influences emotions"[21] and "help people gather, communicate, and express emotional information and to better manage and understand the ways emotion impacts health, social interaction, learning, memory, and behavior"[22]. Essentially, affective computing seeks to give computers the capability to take into account emotion related to human cognition and perception, a key result being Digital Twins with the capacity to more effectively support humans, as well as they themselves being able to make better decisions. Current investigations into applying affective computing in the enhancement of applications include computer-assisted learning, perceptual information retrieval, and human health and interaction.

Taking into account human emotions in Human Computer Interface

4.3 Augmented Human Technologies

The range of activities associated with human augmentation is extensive, and so those covered here were selected based on a number of factors:

- Research focused on addressing brain-based capabilities and disabilities
- Current, emerging and expected brain-focused technologies
- Four key areas essential in creating and supporting the evolution of brain-related Symbiotic Autonomous Systems: mind uploading, cognitive boosting, mind virtualization, and data upload to the brain
- Three key biotechnical areas that will emerge from the above-mentioned research and development activities: human/robot/computer symbioses, symbiotic intelligence, and multiple selves

4.3.1 Brain as a Symbiotic System Component

Huge efforts are dedicated worldwide on brain research aiming at understanding it, replicating lessons learned in computation, and establishing communications both at physical and logical levels that could help people with disabilities and increase the brain's capabilities:

- Brain Initiative (US)[23]
- China Brain Project[24]
- Brain/Minds (Japan)[25]
- The Human Connectome (US)[26]
- EU Human Brain Project[27]

All these efforts, and more focused ones carried out all over the world, are developing and exploiting new technologies which is providing the timekeeping for the future evolution. There are many technological evolutions. The following technology areas are expected to have the most impact:

- Optoelectronics (for probes)
- Optogenetics (for actuators and for selective sensing)
- Wearable microscopes single-multiphoton
- Software for signal processing (to identify signals from noise)
- Fluorescent proteins
- Implantable chips to monitor thousands of neurons
- New materials for long term implants
- Artificial intelligence for isolating emergent meaning
- Extremely low power electronics for use in brain implants
- Wireless communications to connect implants
- Neuroimaging (static and dynamic)
- Brainternet[28]
- Neuralink[29]

In this subsection we address four key areas for the evolution of Symbiotic Autonomous Systems in relation to the brain:

- Mind uploading
- Cognitive boosting
- Mind virtualization
- Data upload to the brain

4.3.2 Mind Uploading

From Science Fiction to Science

Mind uploading, i.e., the possibility to transfer all memories, thoughts and feelings from a person's brain to a computer, has been the realm of science fiction[v] until a few years ago. The development of technologies that allow the monitoring of the brain's activity have moved the field from fiction to science. We are still very far from reaching a true uploading of a brain, even on a small scale, but existing technologies have shown that in principle this might become possible. Current technologies are able to provide an interface between the brain and a computer that can transfer a very minimal indication of what is occurring in the brain at a particular time. Hence these technologies are not actually uploading anything, they are just intercepting (some) brain activities.

Before exploring the status of some of these technologies and the expected evolution it is important to clarify that uploading cannot be seen as a first step to a subsequent downloading, i.e., the brain is not a computer you can back up and then download the back up to that computer (or another one) at a later time. While it might in principle become possible to achieve a mind uploading (using technologies that we currently don't have nor are seen on the horizon) there is no way to achieve a mind downloading.

---

[v] A wonderful description of mind uploading was given in 1981 in "The Mind's I" a collection of reflection on the self by Douglas Hofstadter

In a brain hardware and software cannot be separated

The brain is both hardware and software (using computer science terminology), with the hardware being the neurons and their connections/synapses (plus astrocytes, neurotransmitters, neuromodulators, and so on) and the software being the flow of chemicals and electrical currents/potentials giving rise to activity, perception, feeling, etc. By using a variety of sensors, we can in principle detect all of its activity, convert this activity into bits, and record it in cyberspace (mind *activity* uploading). Once the mind is in cyberspace, and there is a model of that brain (very, very tricky, being close to impossible in the case of a human brain, given its complexity, but feasible for a worm like Caenorhabditis elegans[30]), it is possible to simulate the processing and even the emergence of thinking upon reception of stimuli. However, this simulation will rapidly diverge from what will happen in the real brain since this latter will change as result of its activity.

Downloading from the virtual copy is not possible because there is no separation between the "hardware" and "software" in the brain. One cannot download just the software on a real brain, and the hardware is not downloadable because the brain would need to be built from scratch, impossible at the complexity level of a fly, not to mention a human brain.

Having said this, the possibility of intercepting brain activity and understanding it is no longer science fiction. We could in the future read the mind of a person with lock-in syndrome, and we could surely create a direct communication link with a disabled person's brain to support a direct communication to a machine. Notice that this is more appropriately described in terms of a brain computer interface rather than mind uploading.

In the future, these communications might be exploited as a more efficient, seamless way to communicate with a machine, resulting in an augmentation of the person's capabilities.

4.3.3 Cognitive Boosting

The first IEEE SAS White Paper addressed brain computer interface technologies to interface with prosthetics, for example, to instruct a robot to carry out a specific task.

Ever more powerful technologies are able to capture brain activity, mostly in terms of electrical activity but also in terms blood perfusion, chemical reactions and gene activation. All of this growing data are now processed by artificial intelligence algorithms that are also taking advantage of deep learning approaches to correlate and learn from previous observations. This is crucial since it is now an accepted fact that although in a broad sense all brains (even the ones of other animals) are alike in terms of "working", each one is different. For example, seeing a cat results in a specific distribution of activities in one person's brain that differs from the ones activated in another's person's brain when seeing the very same cat. Decoding the intention to move a hand (by intercepting signals in the motor cortex) is different from thinking about a cat. This involves cognitive aspects, what we often call "thinking". Significant progress has been made in this last decade to interface prosthetics to the brain, using a few electrodes (on the skull or implanted) to pick up brain electrical activity and use it to control external prosthetics, like a robotic arm or a robotic wheelchair. Decoding the intention to move a hand to instruct a prosthetic to perform that action is the ultimate design of the human-prosthetic interface. However, usually the person learns to control the prosthetic by engaging in some specific thoughts activity. Hence, it is not the prosthetic that learns to read that person thoughts, but the person adjusts his thoughts to serve the prosthetic.

All brains are alike but no brain is the same, nor remain the same

The more electrical activity that can be picked up and the more precise the location of that activity, the easier it is to control the prosthetics and more and more complex activities can be orchestrated. This is the reason for the DARPA challenge: Neural Engineering System Design[31], resulting in brain implants able to pick up electrical activities from a million neurons.

Cognitive Prosthetics

Researchers are at work to win the challenge, and it is most likely they will. However, creating seamless cognitive prosthetics is another story altogether.

Here the crucial point is "seamless". In a way cognitive prosthetics are available today: a smartphone is an extremely effective cognitive prosthetic. If we do not know something, a few clicks on our smartphone and the world's knowledge is at our fingertips, similarly for performing a variety of tasks: translating into another language, navigating a foreign city, doing math, etc.

The Imperial College Foresight study[32] predicts the existence of seamless cognitive prosthetics by 2040, although this might not be achieved in its full form, that is, increasing our brain cognitive capabilities through some prosthetics. It does not seem feasible, in this century, to plug in a chip on a brain and boost its cognitive capability at will. However, it is not black or white, there is plenty of grey in between, and that is the area where improvements in cognitive prosthetics are likely to take place in the coming 20 years.

As noted, cognitive prosthetics already exist in our smartphone (but that is not fundamentally different from using a book in a library, just a billion times more efficient). We can easily foresee a smartphone shrinking to the point of becoming embedded in an electronic contact lens. That would provide a cognitive boost, but it would still be what we have been doing for centuries, accessing a knowledge repository through our senses.

Boosting the brain cognitive capabilities

There are some experiments being done in actually boosting the cognitive capability of the brain. There are trials on mice performed at the Wake Forest Baptist Medical Centre[33] indicating that through electrical stimulation, it is possible to increase the learning capabilities of a rat brain. Other trials[34] are also under way aiming at discovering ways to boost our brain cognitive capabilities.

Most of these trials are exploring the use of deep brain stimulation, and it is likely that the increased understanding gained in the next two decades from flagship projects, like the Human Brain[27] and Connectome[26], will result in better ways to boost brain capabilities. At the same time the increase of knowledge is likely to widen the gap between what we can learn and understand by boosting our brain (through stimulation and genetic modification) and what becomes available.

Cognitive boosting through symbioses with machines

It becomes more and more likely that in the future we will need to rely on a distributed intelligence, and that we will leverage it through a symbiotic relationship with machines and ultimately with the environment. This might seem like science fiction but it might be the only way to cope with the avalanche of knowledge being created. This has been pointed out as the way to move forward by the World Economic Forum in its 2018 meeting[35].

Dream reading and recording

A special case of mind uploading is dream reading and recording. Googling "dream reading" results in pointers to sites that are supposedly helping in understanding the "magic" revealed by dreams. That goes back centuries and millennia to the times of soothsayers that are still flourishing today. Freud moved dream interpretation to a new level with psychology. Both Freud and dream readers, however, rely on a description of the dream.

Some scientists are looking at dreams as brain activities that can be tracked and eventually understood, in principle, recording dreams that a person may never be aware of or remember. This falls into the more general endeavor of detecting brain activity and interpreting it.

We are seeing continuous progress in the area of decoding brain activity into meanings[36], however pinpointing exactly what is going on is in the far future. According to the Imperial College foresight study, there may be a machine able to read our dreams and record them by 2040.

Researchers at Union College, NY, have created a robot that is fed with data captured by sensors on a sleeping person[37]. This data is processed by the robotic brain and guides the robot to enact those (supposed) dreams. Some actions, like replicating rapid eye movements occurring in dreams, are easy, however, more semantically connected dreams are more difficult to interpret and are beyond current technological possibility.

This connection between one's self and a robot, with the robot digesting and mimicking our dreams is reminiscent of DeepDream[38], a program developed by Google to look inside the neural networks of a computer as it processes images. The program visualizes the processing giving rise to images that look somehow like a robot dream. Will a robot be able to dream? Possibly. Will it be aware that it is dreaming? Maybe. Will it enjoy the dreaming? We do not have an answer to these questions.

4.3.4 Mind virtualization

From what stated above it should be clear that a full mind uploading is very far in the future, assuming that it will ever be possible (and at the present state of technology and knowledge it is not). Subsets of mind uploading (like the intention of moving a hand or the identification of the image of a cat) have already been proven using current technologies and can be expected in the coming decade as an extension of capabilities that will greatly improve brain computer interfaces (but this is not complete mind uploading).

There is, however, a different approach to mind uploading that uses a completely different set of technologies: mind virtualization. Mind virtualization means the possibility of extracting a number of characteristics from a mind, as observed from the outside, to develop a model of that mind that can be used to simulate future behavior.

Extracting the mind and simulating it

This approach is based on a broad set of artificial intelligence technologies, like machine learning and deep neural networks, both in extracting the characteristics to create the model and in the simulation of the mind based on the model.

It should be noted that these technologies, associated with statistical analysis, are already used in market forecasting as well as election forecasting. Here the key point is the use of statistics applied to a multitude of individuals averaging out the noise in the signal to get an accurate forecast.

In the coming years we are going to see these technologies applied to forecast the behavior of a single individual, basically creating a mind virtualization. Notice that in the case of single individuals we cannot use statistical analysis to average out and eliminate the noise but we can apply machine learning to obtain the same result by creating multiple images of the mind separated in time. As the possibility of monitoring increases, through wearable, ambient sensors, activity tracking (including semantic tracking) and the number of points available to machine learning will increase, more accurate results will become possible.

Additionally, the virtual mind can be used to continually predict behavior, and the system can learn from divergence from the expected behavior. This approach is now commonly used in speeding up the learning of machines with amazing results.

Digital Twins

A short term application of this mind virtualization is in the area of Digital Twin-based education that will be further discussed in the upcoming sections on Digital Twins and on education.

The ethical and privacy issues of these technologies are obvious and need to be tackled. Notice that the availability of mind virtualization can become an asset to a person, leading to augmentation of that person's cognitive capabilities and decision making abilities since one could simulate at a high speed his own mental processes stimulated by a broad variety of stimuli.

Symbiotic Augmentation

Through mind virtualization, the mind and the interaction between minds can be simulated. The latter can include human minds as well as "machine minds" (i.e., machine behavior that in intelligent machines can be very complex) and can be used to understand or predict the overall symbiotic behavior.

This is becoming a crucial point in areas where a loose cooperation between a human and a machine is needed, like a self-driving car at level 4: the car is capable of autonomous driving but a human driver is needed to take control in some cases. The problem is that the human will grow ever more confident in the machine (we have already seen this happen at level 3) and will not be ready to take over when the need arises. The possibility of simulating, using mind virtualization, the behavior of that particular person under a variety of situations can greatly improve safety by stimulating the symbiotic relationship in a specific way fitting that particular person.

There are companies working in these areas. The 2045 Initiative[39] is looking at the broad impact of artificial intelligence, an intelligence that includes the symbiosis with our intelligence and predicts the possibility to upload our brain in cyberspace where it can live "forever". Notice that it is not just about "me" living in cyberspace forever, it is about maintaining relationships once the atomic part of me dissolves. My friends will have the opportunity of talking to me, the "me" in cyberspace, as they do today when using a social network. With the Turing test passed, there is no way to determine if on the other side of the interface to cyberspace there is a real person or a computer (an artificial intelligence), and if that interacting entity in cyberspace is a copy of me, has my experience, my knowledge, my quirks, etc.; there is no way of distinguishing the difference.

Mind virtualization creates an alter ego in cyberspace

This alter ego in cyberspace will diverge over time from the real "me" since it will be exposed to interactions, experiences, I will no longer have, but if the real me is no longer existing it does not even make sense to talk about a divergence. It will still be me, just an older and more experienced me.

The brain uploading is clearly opening up completely new spaces, bringing along unexpected societal and ethical issues. What about a cyber-me that through interaction in cyberspace will cause damage to another entity, be it virtual or real? Would the state punish the digital me? How? Will my digital me be condemned to "death", to be erased from cyberspace? Aha! You cannot erase my digital me; you might erase one copy but my digital me could be so smart to clone itself in the billions and hide its self

New Social and Ethical Issues

in many ways to go undetected. These are just a few examples to point out the amazing new space that is developing. Remember that there is no black and white, but plenty of grey and some of this is already happening today. Replika[40], as an example, offers the possibility to create a digital copy of ourselves in cyberspace.

Kernel[41], a startup founded by neuroscientists and engineers from top US universities, is looking at technologies to access, read *and* write the brain (see Section 4.3.2 Mind uploading discussing also data download to the Brain and Section 4.4.2 Data Upload to the Brain). In 2017, DARPA awarded the University of Berkeley with a $21.6 million fund to develop technologies for reading and writing the brain.[42]

4.3.5 Human/Robot/Computer Symbiosis

A different twist on data upload and human augmentation involves human machine symbioses and is the result of using avatars, software as well as hardware entities, that can represent ourselves (using mind virtualization as a starting point) by keeping in synch with us by transferring data to our brain (through sensorial mediation). Avatars have been around for a while (e.g., Second Life[43]), but they have operated in a disjointed way from the person they represent; it is not a symbiosis at all. A real symbiosis between a person and her avatar can be expected in the future.

This is what one of the XPrize foundation challenges is: create a robot that would let a person see, hear and feel from the other side of the world by 2021[44] (actually the prize only requires a robot that can operate at least 100 km away from the person). The robot becomes a spatial extension of that person effectively augmenting him by uploading to that person's brain the data generated by the robot's real presence and that person's virtual presence. This is the crucial aspect that differentiates this from the remote control of a robot, like a drone. Technologies like humanoids robots, high bandwidth communications, virtual reality and high resolution haptic will need to be in place.

The evolution towards human/robot/computer symbioses is also starting to appear on the shop-floor, with blue collar workers operating in teams with robots that are behaving more and more as teammates complementing the workers' capabilities and vice-versa with the workers complementing their abilities with the robotic ones. Exoskeletons are an example of this evolution taking place today. In the next decade the use of exoskeletons will become more and more pervasive, due to soft robotics that better fit the human body and seamless communication between the two.

Exoskeletons evolving towards human robot symbioses

A further evolution will allow a detachment of the exoskeleton that for some activity will disengage from the physical proximity still keeping the seamless connection with the human. The aim of the XPrize is to achieve such a seamless connectivity over considerable distances. Avatars, particularly software ones, will become more and more common, effectively augmenting human capability in time and space (they can operate around the clock, in multiple locations and in parallel). The symbiosis is important since it creates a return channel towards the person, what the avatar does affects the knowledge and experience of the person, and the avatar increases its capabilities over time.

As pointed out in mind virtualization, the human/robot/computer symbiosis is fraught with ethical and social issues that need to be addressed.

4.3.6 Symbiotic Intelligence

The symbiotic relationship among human/robot/computer creates a symbiotic intelligence that goes beyond the linear composition of the individual intelligence. In a linear composition, the individual intelligences interact each keeping its individuality, and one can pinpoint the reasoning and decision to a specific one.

Partnering and sharing responsibility

As an example, by partnering with a computer through augmented reality contact lenses we will be able to extend our knowledge to the web, e.g., seamlessly accessing Wikipedia. The computer will be responsible for accessing data on the web and making the decision on which data to display to us (and this requires some sort of intelligence, e.g., finding the appropriate sources, making sure data are still valid and relevant, presenting them in an understandable way, and so on). The understanding of this data will occur in our brain, leveraging our intelligence.

Emergent behavior, distributed intelligence

In symbiotic intelligence the awareness, decision making and following activity is emerging from the whole system and cannot be pinpointed into its elemental components. The behavior of swarms is an example of a fully distributed intelligence.

The science of chaotic systems, complexity, small worlds along with simulation engines are the tools to understand and predict behavior. As more autonomous systems interact with one another the issue of emerging behavior, which relates to emerging intelligence, will become crucial.

Brain-IoT

Smart interconnected IoT will create swarms from a behavior point of view. The European project Brain-IoT[45], a model based framework for dependable sensing and actuation in intelligent decentralized IoT systems, is looking at these aspects.

Brainternet

The participation of humans in the swarm will become "normal". The first studies aiming at connecting the brain to the Internet and having it becoming part of the IoT landscape have started. See the Brainternet[14] project at University of the Witwatersrand.

The area of emerging intelligence is opening up new ethical issues, particularly in the area of responsibility. Since decisions and activities are distributed and it is no longer possible to pinpoint a single decision maker, accountability becomes an issue.

4.3.7 Multiple Selves

The possibility of creating a digital copy of the brain, the participation of that digital copy in a variety of activities in cyberspace, and its ability to affect the real self, brings to the fore the issue of multiple selves: there is not just me, but there exists a copy of me that I can control up to a certain point and that can influence my behavior.

This seems to be a completely artificial issue, however, recent studies on the brain have shown that there are multiple selves in our physical brain. There are some pathologies, known from centuries, resulting in the display of multiple personalities. Until two decades ago the investigation of multiple selves was in the psychology domain[46], and it was considered a theory to model some observed behaviors. More recently, fMRI application to split brains (where the corpus callosum has been cut through surgery) have demonstrated the existence of different selves within the brain that are competing and that give rise to the emergence of the self. In situations where the corpus callosum is cut, the selves active in one brain hemisphere can no longer interact with the ones in the other hemisphere and are unaware of their existence.

This situation may become normal in Symbiotic Autonomous Systems, systems where the awareness is distributed. In this case, decisions are the result of individual behavior influencing the overall behavior, and there is the emergence of a global self out of multiple selves. Interconnectivity and interoperability are the dominant factors in the coexistence of multiple selves leading to a global self.

In general, the interconnectivity will be based on local observation of the environment, i.e., each component in a symbiotic autonomous system "just behaves" on its own accord being influenced by the way it perceives its environment. Its behavior, in turns, changes the overall environment and influences the behavior of the other components. This applies to both components made of atoms (like interacting robots, swarms of drones and ourselves as part of the global system) and to the ones made of bits plus the interplay of bits and atoms components.

Digital Twin technology, as it will be addressed later in this white paper, provides the substratum to the selves in cyberspace.

4.4 Augmentation through genomic engineering

Since the discovery of the code of life, scientists have looked at ways to tinker with it to change some life characteristics. Over the last 10 years (CRISPR was discovered in 1993 and Cas9 in 2005 but their application to the splicing of DNA can be dated to 2013) researchers have been able to modify DNA strands removing and inserting snippets of DNA taken from genes of different species[47]. The manipulation is easy to understand (although it is not so easy to carry out). A gene has been discovered to be the code for a certain protein production in a given species. Using CRISPR/Cas9 it is possible to separate those instructions (that strand of codons) from a gene and then, again using CRISPR/Cas9 to splice them into a gene of the target species. When that gene will be activated will lead to the coding of the desired protein.

CRISPR/Cas9

In this way it has been possible to add some (desirable) characteristics to bacteria by borrowing them from a different species. Codons, i.e., the coding of proteins, are exactly the same in all species, hence it is possible to transplant them from one species to a different one without any rejection from the target species.

However, the protein resulting from those instructions may not be accepted by the organisms, or it may lead to side effect. The key issue is that the characteristics of a living being, its phenotype, depends on the genetic code (the genotype), but there is no one-to-one correspondence between a gene (and its DNA strand) and a single characteristic. A variation in a gene may lead to the living organism displaying a desired characteristic but at the same time it might create undesired characteristics that may become apparent only at a later stage.

Genotype vs Phenotype

Significant effort is under way to understand the overall impact of gene modification on the phenotype of that organism. The use of artificial intelligence (deep learning) made possible by the huge quantity of data that are being harvested from the modification of genes promises to deliver a tool for connecting the genotype with the phenotype[48]. This might become available in the next decade.

AI to design Phenotype changes

In perspective, it should become possible to design the changes of a gene to obtain the desired change in the phenotype, probably in the 2040-2050 timeframe. Notice that not all potentially desired changes will become possible. As an example, the number of fingers in mammals is governed by specific genes (Hox genes[49]). Altering of those genes may result in extra fingers but carries along non-favorable characteristics, hence the reason why mammals did not diversify the number of fingers.

Another crucial issue is that once a gene has been modified, that modification will be passed on to the offspring (if it is encoded in a reproductive cell) with the potential of generating undesired effects that are difficult to foresee.

RNA modification

By 2020 researchers expect to have a tool similar to CRISPR/Cas9 that can be used to change the RNA[50]. This would still allow the creation of desired protein but the mutation will not be passed on to offspring (traits inheritance occurs only via DNA).

It has also become possible to use viruses, billions of them, as vectors to change the DNA of cells in a grown individual, and some companies[51] are experimenting in this area; a first human trial took place in 2017.

So far the focus of researchers has been on curing genetic disorders or diseases but it is clear that there is a potential for human augmentation (carrying along many ethical issues).
Notice that we are still very far from the point of having the knowledge required to augment a human being. One case is to modify a gene to restore the normal coding, a completely different case is to create a code to achieve a certain result, particularly in certain areas. As an example, we have very little understanding of what makes intelligence emerge in a brain, hence it is impossible today to imagine a modification of the genome to augment intelligence.

From "cure" to Augmentation

Extending human life, on the other hand, by extending the telomeres seems within the domain of future feasibility since several studies have identified the shortening of telomeres as a reason for ageing. This would be a sort of human augmentation.

4.4.1 Humans "a la carte"

There are a few companies[52] using a combination of DNA sequencing, screening and gene selection, allowing future parents to choose a few traits for their child (like blue eyes). The parents' genes are analyzed to identify the hidden traits that can, once combined in a proper way, give rise to a desired trait. Basically, these companies turn a probability into certainty. They cannot create a trait that is not contained in the parents' genes, but they can make something that can be highly unlikely happen.

As discussed in the previous section, researchers have, and will continually have more, tools to modify the genes so it would be possible to turn the impossible (the absence of a trait in the parents' genome) into a possibility (create a trait in the newborn).

What will be on the menu?

Are we moving towards humans "a la carte"? And, what will be on the menu? There might be traits that are not present in the human species, like the possibility to eat certain types of food, that might be created for future generations. Will these become part of the menu?

This area is obviously fraught with questions since, at least today, we cannot evaluate all the consequences of mutating the genome. What would be the side effects? Even assuming that in the future, due to new predictive tools, it would be possible to evaluate all consequences, what would be acceptable and under which ethical framework? And what about the issues rising from the emergence of different humans (not necessarily "super-humans")?

The alteration of humans having an "augmentation goal" can happen through their life time (alteration through design is addressed in the next section). Human augmentation is a continuous process that has accelerated in the last decade and that is foreseen to accelerate even further in the coming decades with seamless integration of technologies that increase human capabilities and with technologies that modify the genome. They are often referred to as transhuman technologies. They can be clustered in the following areas:

- repairing
- lifestyle adjustment
- prosthetics
- boosting
- embellishing
- replacing
- adding
- altering
- redefining

Each one is addressed below.

Repairing

A few species have the possibility to repair their body, like re-growing a limb (e.g., the newt). Researchers are identifying the genes that make this possible[53]. Some of them are just dormant in our genome and might be reactivated. Others might be "implanted" in the future. A crucial aspect here is finding a balance between the possibility of re-growth and making sure uncontrolled growth does not happen (cancer). There ought to be a reason why evolution has suppressed the re-growth capabilities in most animals, particularly more complex ones, limiting it to small repair activities (like skin growth). Technology may provide the means to activate this capability only when needed and then deactivate it to avoid undesirable side effect. Notice that once perfected this might lead to the self-replacement of organs once they start to degrade.

Lifestyle Adjustment

Training to keep fit is quite generalized but there are people, like mountain climbers, that undergo specific training to prepare for a tough activity, like climbing a Himalayan peak; others need to train daily to be fit for their profession, like professional dancers. In the future, *lifestyle adjustment* may become extreme, resulting in transhuman capabilities. Imagine embarking on a trip where food will be limited and not sufficiently varied to ensure health (this may also apply to astronauts on a long space voyage). Bacterial genomic modification may create symbiotic bacteria that when ingested would become part of that person's bacterial flora making digestion of certain food possible (like digesting cellulose). These changes can make different lifestyles possible enabling life in areas where today it would be challenging. A severe climate change may require a lifestyle adjustment of many people. Notice that scientists have already adjusted the lifestyle of several plants[54] to make them grow in areas that would normally not support those plants (like drier places or salty ground). This is clearly an area aiming at a symbiosis with the environment, a symbiosis that has been the norm throughout natural evolution and that will be accelerated through technology.

Prosthetics

Prosthetics are becoming more and more effective[55], good news for people with an acquired disability, like hand amputation. These prosthetics are trying to simulate as much as possible the real part they are substituting but are not limited by the constraints of the real part. A prosthetic hand in principle can be made with material that is much more resilient to heat: a cook may need to pay less attention to the heat of the stove, or a mechanic can touch the engine with her prosthetic hand with no risk of being burned. Artificial limbs may augment a person's gait and speed. Not suffering from fatigue, an artificial eye can see beyond the physical constrains of a real eye and might even connect directly to cyberspace.

Boosting

Boosting our body, our senses and the physiological processes to increase our capabilities including human capabilities to interact with machines can be considered augmentation. Today we have drugs that can boost our capabilities (usually with undesirable side effects, an example being doping). Side effects must be considered; the boosting of one capability may adversely affect other physiological processes resulting in an overall negative situation. Our natural capabilities are the result of a compromise reached through millions of years of evolution; any boosting is likely to disrupt this compromise. Some researchers are convinced that it will be possible to disrupt this compromise bringing our body to a new acceptable dynamical equilibrium. Trials are underway to improve memory by electrical stimulation[56] of certain brain areas. So far experiments have occurred on rats with positive results; by 2040 we might expect to have a number of technologies that will result in boosting our natural capabilities.

Embellishing

Cosmetics have become quite sophisticated over the centuries today leveraging advanced technologies like carbon nanotubes[57] and contact lenses to alter the iris color. More recently the availability of flexible electronics and graphene has led researchers to experiment various forms of electronic tattoos[58] to provide new forms of embellishment. Plastic surgery is now able to modify many physical aspects of our body, such as enlarging, reducing, and modifying the color. There are even applications that allow you to virtually experiment plastic surgery[59]. By 2040, one can expect technology to provide a variety of ways to change the human body to satisfy the fancy of advanced cosmetics. Would you like longer, thinner fingers? What about a luminous fluorescent cheek? By tweaking genes, we might be able to define several aspects of our child body, and with genome modification using viral carriers, we can even modify our current genome.

Replacing

The human body is a system of sub-systems. Sometimes a sub-system fails resulting in a disability (losing a limb) or death (heart failure). Over the past fifty years, scientists have managed to replace some organs, using transplants and prosthetics. Technological progress will increase our ability to replace body parts. Researchers are progressing in growing organs: 3D printers using a person's cells are already printing skin, bones and windpipes (tracheas) and in the labs printing bladders, livers[60] and muscle are in experimental stages. A growing understanding of the genes and their regulators will open the way to self-regeneration. At the same time, artificial organs are being studied. In a few decades, we can expect organ replacement to become as common as the replacement of the eye lens or teeth is today. Since failing organs represent a significant cause of death we might expect an increase in the average life expectancy once organ transplant becomes routine. Notice that replacement in the future may also lead to an increase in performance of the replaced organ.

Adding

We are already adding capabilities to our body when using prosthetics like the special glasses used by surgeons to get a better view of the operating field. In the future, there will be more and more opportunities to use technology in a seamless way to augment our body capabilities, for sure in the area of increasing our senses, like seeing in the ultraviolet and infrared spectrum and hearing beyond the range of normally perceived frequencies. We can also expect, according to the Imperial College Foresight study, the advent of embedded chips that will expand our capabilities. Brain computer interfaces are already being studied with the aim of becoming seamless. Brain chip implants are already a reality to control epileptic crises. DARPA launched a project in 2017 to create a chip that can be implanted in the human brain to help cure some forms of mental disabilities. Whether this will become a reality or not is still an open question but questions on possible drawbacks are already being voiced[61].

Altering

We all belong to the human species because our particular DNA characterizes the human species. It does not matter if one is tall, fat, yellow, with curly hair, or so on; that person can create an offspring with another human. Even though one person won't find another human with exactly her same DNA (unless she has a monozygotic twin) the blueprint is the same, hence any male can generate offspring with any female of the human species. By altering the DNA, that person may depart from the human blueprint, hence she may no longer be able to generate offspring with another human, meaning that person is now part of a different species. This alteration is already a technological possibility; using CRISP/Cas9, it is possible to alter a genome. Researchers have already demonstrated the possibility of changing the genome of an adult, using viruses as vectors to infect the cells of the living organism to change their genome. This is a great possibility when thinking of curing a genetic disease, and this is the motivation for studying and experimenting. At the same time one can move from a repair to an adjustment; a tiny one to begin with may be to enable people living in harsh condition to eat some food that would not be edible for a "standard" human. Then one might look for some genome modification that would make a person more resistant to viruses and bacteria (we are doing that with vaccination, so one might see a genome alteration to this goal as an alternative, possibly a more effective one, to vaccination). The problem is that we will not be moving from black to white, rather there will be a drift (as happens in natural selection) across a multitude of greys till we reach a point that the human blueprint is changed, and a new species is born. Through the genome alteration, humankind will be doing what nature has done over eons. The human "mobility" has stopped: we can no longer drift as human species into another species through a natural selection process because we keep "mingling" together; the world has become too small to have enclaves where a species is isolated from the rest of the population and can drift naturally to the point of creating another species. What nature did following a random process (streamlined by natural selection) we now will have the technology to do "by design"[62].

Redefining

Altering the DNA of a person may lead to the creation of a new species but at a first glance that person may look like a human being, same size, intellectual and social behavior. By redefining the DNA, it is possible to create a living form that can differ in functionalities, performances and "shape" from a human. In theory, one could create a living being that would better fit lower gravity condition on another planet (with lower gravity you can have a bigger being, with higher gravity a smaller being would be a better fit). There is a wide spectrum of possibility once you start tinkering with the DNA. The problem is that it is unclear what the effect of tinkering will be. New studies are underway to associate the genotype with the phenotype using machine learning techniques[63]. Notice that one does not need to change much in a genome to get a completely different living being. Our human genome is not that much different from a primate (understandable) but also not that much different from a dog genome. Little changes can result in dramatic phenotype changes; the problem is that we have no idea, today, on how to design a living being starting from the code of life. This might change in the future, although this may not happen within this century, where changes will be based on trial and error (and this is what concerns most people). Redefining the DNA might turn out to be a step in the symbioses of autonomous systems, in the integration of a human being into cyberspace. We are clearly at the edges between science and science fiction but potentially the most convenient way to create a seamless connection between a brain and a "computer" might be through the restructuring of the brain by redefining the genome. Several pieces of the puzzle are being created and laid out in various labs today. They are all important but it will be a long and complex way to put those pieces together in a way that they really work.

In the meantime, companies like the Methuselah Foundation[64] are working to create the fountain of youth, being 90 years old and looking and feeling like a 50 years old, a goal to be achieved in the fourth decade of this century. The SENS Research Foundation[65] in cooperation with the Forever Healthy Foundation is working on rejuvenation technologies (an even bigger challenge than remaining "young" since this requires reversing the ageing process). These seem to be more dreams than reality and yet in this area we are getting closer, so close, in fact, that huge societal issues are looming.

4.4.2 Data Upload to the Brain

Data upload to the brain is a research area aiming at strengthening the brain capacity to learn and memorize data. It shall not be seen as the data upload happening with a computer, such as plugging in an SD card and transferring the data from the card to the computer memory. The brain has evolved to be able to respond to complex situations by analyzing data harvested by the senses (visual, aural, tactile) and making decisions by comparing the outcome of possible actions. This ability leverages memories. Although we understand much more about the creation and keeping of memories in the brain, many issues remain.

Deep brain stimulation

Researchers have already proven ways to strengthen memories, both in their creation and keeping. Deep brain stimulation of specific areas of the brain using electrical pulses has been proven effective in some experiments.

Activating specific neural activity is also becoming possible using optogenetics. We are, however, quite far from even basic data upload, like transferring an image or a sound directly to the brain. The brain has evolved to react to stimuli coming from our senses, involving many neural circuits that are somewhat different in different people, and even in a specific person these circuits evolve over time. Hence it is close to impossible to activate hundreds, sometimes millions of neurons in parallel using data upload.

Progress has been made in substituting the senses or flanking them, like using retinal implants to send messages to the brain using the optical nerves or bypassing a faulty ear and directly stimulating the cochlear nerve.

4.4.3 Life by design: artificial parents, symbiotic life

Once the door opens to the selection and combination of genes to create the offspring in vitro, natural impossibility can turn into reality.

Cloning

Cloning is already happening, just not at the human level. Human and more generally primate cloning presents harder technical issues but in principle human cloning can be possible. There are significant ethical issues in this area as well. For an interesting overview read the Cloning Fact Sheet[66] written by the Human Genome Research Institute.

A further evolution, somewhat in parallel with the creation of a fully functional artificial womb, is the possibility to have two males generating an offspring. Today a male couple could adopt a child, and neither of them will be the biological father, or one can use a sperm to fertilize an egg donor, in which case he will be the biological father. The availability of an artificial womb would make this possible without the mother to carry out the gestation.

Artificial Womb

Genetic manipulation is reaching the point of making possible to take a primordial germ cell (PGC) from a male and direct it (through gene tinkering) to create an egg. That egg can be fertilized by the sperm from the other male leading to a baby that is actually the biological child of two fathers[67]. Notice that the egg developed from the PGC will contain woman DNA, from the mitochondria (these are always inherited via the mother line) that are present in the cell.

Stretching to the extreme of in-vitro creation of life, intelligent life completely decoupled from the natural selection process and driven by design, one could imagine a scenario where symbiotic systems are created mixing biology, smart materials and software. There have been a number of results in integrating living cells, including neurons, with silicon chips, aimed at harvesting the best from both worlds. Research is progressing in several fields to make symbioses possible over extended periods of time (using micro fluidic chips to sustain living cells, protonic chips to communicate with cells) and to leverage both (merging analog and digital computation). Start-up companies, like Koniku[68], are working to create symbiotic processors.

Designing Symbiotic Life

By using genetic engineering, it might be possible to customize life to fit with artefacts, creating characteristics that could make interaction with artefacts more effective. In the second half of this century, one could expect a co-design of life and artefacts through genetic engineering. This can start from increasing the acceptance of external bodies (like nano-bots and microbots) that could serve as liaison agents between the living entity and the artefacts.

4.5 Awareness Technologies, Intention Recognition, and Sentiment Analysis

The growing availability of sensors and machine intelligence is creating an ambient more and more aware of what is going on, why it is happening and what will further happen. These three characteristics, what is happening, why it is happening, and what might happen, are the basic components of what we call intelligence, a feature that we find, in different degrees, in animal life. This is the result of an evolution process that stems from the advantage deriving from possessing it, making those species that casually acquired it to take the edge on those that didn't. Notice that this advantage influences the selection process (leading to evolution) if it results in acting in an advantageous way out of a slate of possible actions. To a large extent plants are missing the possibility of acting differently, hence we have not seen intelligence evolving in plants (plants can evolve different strategies but the reaction time to changes in the ambient is slow, rocks, one might argue, also react to changes in their environment but the reaction time is even longer, hence does not qualify to our definition of intelligence).

What Happens
Why it happens
What will happen

Technologies to create awareness, make decisions (evaluating the whys is a fundamental component), evaluating the outcome of a decision, and implementing the decision are now available and are becoming embedded in more and more artefacts and globally in several ambient. An integral part of intelligence is autonomy (at least in the evaluation).

Awareness requires, in many situations, the capability to recognize the intention of other players (life and artefacts alike) operating in the same ambient. This is essential both in understanding the why and in studying evolution scenarios.

More recently experimental technologies have been developed to go beyond intention recognition to look into the motivations of different players, called sentiment analysis.

4.5.1 Awareness Technologies

Interestingly, awareness technologies may be seen as serving two opposite purposes:

1. Creating an ambient that although permeated by technology is not perceived as such by users (no user awareness). This was the objective of the Japan MIC[69] (Ministry of Internal Affair and Communications) to ensure a seamless experience to lay people of ICT technologies.
2. Endowing artefacts with the awareness of their ambient to take informed decisions and to interact seamlessly with the various components in the environment.

In the area of Symbiotic Autonomous Systems both purposes are important, however in this section the focus is on technologies supporting the second purpose.

Context Awareness

A first area of awareness technologies relates to context awareness. Sensors embedded in the artefact, able to detect the shape of the environment, its characteristics and its various components, are becoming more sensitive, performant and affordable. Smart materials (like sensitive skin for robots) will be playing an increasing role in sensing. Indirect sensing, such as the one provided by safety cameras, is also relevant in several situations, and in the future the exchange of sensed data by components will take an important role (as an example, in providing data through exchange among autonomous vehicles in a given area).

The detection of the shape of the environment and of its components can be done in several ways:

- by scanning the ambient, e.g., using laser based technologies like LIDAR or sonar
- by looking at the ambient, e.g., using digital cameras (usually more than one or cameras with several lenses to ease the 3D recognition)
- by identifying the objects, e.g., using identification methods like tagging, patterns recognition, sound signature
- by interacting with the objects and sharing knowledge.

Depending on the situation one or another type of sensors can be used. In many cases several types are used in parallel. In general, the more data that can be harvested the better. Sensors data basically respond to the "what is happening" question.

To respond to the "why it is happening" other technologies take the stage. The aim is to understand the meaning of the harvested data. The semantic extraction correlates data and makes use of a knowledge base (like knowing that a table is supposed to stand still while a dog moves around). Data correlation may occur across different sources, including social networks when the data relates to people. In case of people, context aware technologies are already widely applied to deliver customized services. For example, banks are using context awareness to contain improper use of credit cards.

From local to global, system wide, context awareness

Context awareness is a fundamental characteristics of autonomous systems, and so far these have been designed to have it. In the future as systems will be created through the interaction of autonomous components (each potentially context aware), the overall context awareness may become an emergent property of the whole system, with its individual component sharing individual awareness to generate a global awareness. This is a matter of research for the coming years.

Also notice the shift in the definition of context that is very relevant in the area of Symbiotic Autonomous Systems:

> *"… The focus moves from seeing a context-aware system as an artefact "sensing" information, to seeing it as an interactive system with a physical user interface. This makes the distinction between foreground and background interaction a property not of the system, but of the situation. A consequence of this philosophical standpoint is that context can never be a property of the world, but that context rather is the horizon within which the user makes sense of the world."*[70]

Context aware computing is clustering a set of technologies and approaches more and more based on artificial intelligence[71] (deep neural networks) used to extract meaning, and in Symbiotic Autonomous Systems it will make increasing use of interaction with other entities to get more data. The data is used to create virtual models of the context on which simulation can be performed. This is crucial to support the analysis of "what will happen".

The Digital Twin approach can be used in the simulation to implement (in the virtual world) several strategies to evaluate the outcome and select the most appropriate one.

Goal Awarenes

A second area of awareness relates to goal awareness. The statement "… and select the most appropriate one" implies that there is a metric or a framework to identify the good one. This has been so in autonomous systems where the criteria of "good" was cabled inside the system, like "take the option that reduces fuel consumption". In Symbiotic Autonomous Systems there are several independent systems, and it may not be straightforward to cable in each system a metric and have them all make sense as a whole, since in most cases they have been designed independently of one another. Moreover, the overall goal may require some adjustment to the individual goals.

As artificial intelligence, in particular AGI, takes over the system, it can learn not just the most effective strategies but can also start to create its own framework upon which take decisions. It will, in a way, develop its own goal.

As it is observed in the book "Life 3.0: being human in the Age of Artificial Intelligence" [72] the problem is not (as much) the emergence of a malicious AI whose goals will oppose our human goals, rather the prevalence of the AGI competence leading it to create goals that are not compatible with ours. How could this be? It is the same situation we are facing when deciding to build a dam to create hydroelectric power. By flooding the area with a lake we are killing anthills, just to mention one side effect, and we are not giving it a second thought. The benefit of having electric power far outweighs, in our framework, the loss of millions of ants. Would we be able to ask the involved ants we might get a different perspective on the matter? So, what if, for a superior benefit, AGI sets itself a goal whose side effect includes the loss of human lives? Notice that it could be a perfectly good goal, like recovering from an epidemic or a famine. For example, a swarm of drones can be engaged to carry drugs and food to a remote. Once at the location, their collective intelligence finds out that the combined effect of killing a certain number of elderly and distributing the drugs and food will defeat the epidemic or famine. It is unlikely that most people would accept that kind of solution as ethically viable. This is both an ethical question, and as such it will be addressed in the relevant section of this white paper, and a technology related question of how can we define and possibly control the outcome of an autonomous goal such as the one that might become viable in the context of Symbiotic Autonomous Systems. How can we implement a system of shared intelligence that lead to an overall emerging intelligence that is still under our control?

It doesn't stop here. In the future, possibly before the end of this century (with some expecting around 2075) AGI will be superseded by ASI, with intelligence far beyond the human one and one that will have an embedded capability to set its own goals.

This is an open area of research that involves:

- Transferring our goals to AI – notice this is more about an AI system "learning" our goals rather than programming it. By definition if the system is autonomous you cannot program it, you interact with it. Learning means that AI cannot stop at learning what we do, rather it should understand why we do such things.
- Having AI adopting our goals – it is notoriously difficult to have other people adopting our goals, just imagine a machine. To adopt our goal a human needs to find it compatible with his own framework, be open to adopting a goal he does not have, and not have already committed to adopting a different goal (possibly auto-generated). With machine intelligence the trend is similar, just trickier: the machine should be smart enough to understand our goal and ready to adopt it and not so smart to consider that the only goals that really matter are the ones it can self-generate. Researchers are studying this aspect using inverse reinforcement learning.

- Having AI retaining our goals – this is probably the trickiest of the three. Again, looking at humans, as we grow we get (generally) more intelligent, and we change our goals. We have a few goals that are cabled in our genes, like to reproduce. Yet, as we understand how reproduction works, we change the goal, keeping the fun part and dropping the reproduction part. This can similarly apply to a super-intelligent system. As it considers its goals it would reflect on them and eventually decide to change them.

Mirroring others' awareness

A third area of awareness, covered in a specular way in the next subsection, relates to the potential perception of context and actions carried out by the "aware" entity by other entities. This is, by far, a higher level of awareness and is found only in a few mammal species, as far as we can tell. It implies the capability to imagine what the other entities can perceive or feel. Human brains and primate brains have been found to have mirror neurons that serve this specific purpose.

Humans for sure, and possibly other creatures, can imagine how other creatures would feel and react when confronted with our actions. This is a fundamental characteristics of social behavior. Most of the time we act in a way that we feel is acceptable by our environment.

Social robots have become a significant area of study, where the focus is on facilitating their interaction with us, human beings. In symbiotic systems, where a component is a human being, the other components may get hints on what the human is perceiving by looking at some telltale signs in her expression. There is already technology to evaluate the feeling of a single person as well as the feeling of a cluster of people (sentiment analysis discussed later on). Digital cameras are already equipped with software that interprets a smile to take the snapshot at the right time. By analyzing a number of traits, including posture, movement, and tone of voice, much more sophisticated software can extract very precise information on the feeling of a person. Simple camera sensors coupled with complex software can accomplish this feat.

However, it would be better to foresee the feeling of a human (or another component in the symbiotic system or in the environment) before executing an action. The point is to take a decision based on the possible ways these decisions would affect the others.

Virtual Twin vs. Digital Twin

Technology in this area is also progressing through the creation of virtual twins. Notice that the virtual twin differs from the Digital Twin associated to an entity. A Digital Twin is coupled explicitly with its real twin, a virtual twin is created on spot through modelling of the perceived behavior of an entity and is used by the ones that created it. A Digital Twin is associated to the real twin; a virtual twin is associated to the entity using it (and different entities would each generate their own virtual twin to "understand" the world around it). The recent approach based on generative adversarial networks can be used to test potential effect of decisions on the virtual twin.

The concept of a virtual twin can be applied to humans as well as to machines. It is created, and refined, through the observation of the behavior manifested by the “real twin” in response to specific stimuli. Deep learning technologies are useful in developing the virtual twin and refining it.

The virtual twin will be used to test (in a blink of an eye) the possible response to an interaction, and decisions will be taken based on the desired response. It is interesting to notice that these approaches and technologies are already used in modeling the possible responses of an audience or of voters during an election campaign, and the candidate will talk (in form and content) based on the expected reaction from the audience.

This becomes part of the way interactions are constructed, with continuous refinement that is needed not just to create a more accurate virtual twin, but to take into account the changing responses over time to a given interaction (what may work now may not work tomorrow). This is an area that on one side connects to sociology and psychology (if we are creating virtual twins of people) and game theory if applied to machines.

Symbiotic Autonomous Systems will become able to mirror their environment as well as their individual constituent’s awareness probably in the next twenty years (with some aspects already addressed today).

4.5.2 Intention Recognition Technologies

Intention recognition has developed significantly in fields like security and military. Additionally, it has been considered in health care to help patients with communications disabilities. The aim is to be able to decode people’s intentions through the observation of their behavior and through analysis of their interactions. Our brain has become quite good in reading between the lines, and it is accurate most of the time. It is also fairly good in recognizing intentions from some living creatures it is familiar with, like dogs and cats and by extension to similar animals through some stereotype used to detect aggressiveness and social behavior. An animal’s brain is also equipped, at least in a number of species, with this capability (notice that it is different from the one addressed in the previous section that related to imagining the impact of one’s action on another entity).

In the area of Symbiotic Autonomous Systems, the interest on intention recognition is extended to recognizing the intention of any kind of interacting entity, both living and machines.

Autonomous vehicles need to predict the possible movements of people in their environment (Does that pedestrian have the intention of crossing the road? Will that cyclist turn right?) as well as predict the movements of other vehicles. Obviously, cars might communicate with each other if they are new models, but they will have to discover the intention of old vehicles that are not equipped with communications capability.

Viewing the vehicle turn signal is not, per se, an assessment of an intention; it is a way of communicating a plan. However, sometimes the turn signal may have been erroneously activated and therefore intention assessment can be important (a driver has indicated the intention to turn right and instead the car keeps going straight or worse turns left). Notice that our brain, through experience, is pretty good in evaluating a variety of signs and hints to work out a high probability intention recognition. By many signs that would be difficult to spell out, a driver may detect the intention of a car to turn, even if no direction light has been activated. Sometimes it is referred to as a sixth sense: indeed, intention recognition technologies are asked to provide a sixth sense to a symbiotic autonomous system and more generally to smart autonomous systems.

Pedestrian intention recognition[73] is being assessed using Hidden Markow models; looking at a face[74] for tiny reactions can provide data for intention recognition, comparing those reactions to a virtual model of that person (which takes into account gender, age, culture) enriched with historical data on that person if available (using deep neural networks analysis). Notice that digital sensors, like the one of a digital camera, can pick up variations in the heart beat by looking at the subtle changes in color of the face[75] as the heart pumps the blood, thus deriving data pointing to excitement, fear, or interest. Similarly, the detection of eyelids, of the iris and pupil movements provide additional data useful for intention recognition. Systems have been developed to assess the physical fitness of a driver (increasing drowsiness as an example) based on head movements.

Many social and psychological studies have created significant knowledge in modeling human behavior and expected behavior. This can be used in the evaluation of data collected by sensors (mostly visual sensors although a growing set of knowledge is becoming available in the assessing of aural clues). Robots designed to work in an open environment are progressively equipped with intention recognition capabilities.

Brain computer interfaces have been demonstrated to have the capability of detecting intentions before they are turned into an activity. Actually, recent studies have shown that the intention may be present in the brain processing even before it is perceived by the brain owner. Hence, in the long term, once seamless BCI becomes feasible, it might become possible to mine the brain for intentions. Clearly this opens up a Pandora box in terms of privacy and ethics.

In certain complex settings, intention recognition can become important for a machine to understand the command received from the human operator. Teleoperation of drones may be one example as pointed out in a recent paper proposing to use convolutional neural networks (CNN)[76].

Machines can also demonstrate signs that help in intention recognition, e.g., a deceleration by a preceding vehicle in certain situations may suggest a turn intention.

Discovering the intention of a machine requires a similar process of data evaluation against a virtual model of the behavior of that machine. If no human is involved (as might be the case of a vehicle where the actual behavior is the result of its driver behavior) the point is to understand the decision process guiding the behavior of that machine. However, in the future as machines will be acting on the basis of their embedded artificial intelligence, the intention recognition will become, in a way, more similar to human intention recognition.

Intention recognition in complex systems

In the case of complex systems whose behavior is emerging from the loose interactions among its components the intention recognition takes a different spin. A flock of flying starling is continually changing its shape and direction. This is the result of an emergent behavior out of a multitude of (predictable) behaviors. In this case the intention recognition is played on a different level since there is not an "intention", just an effect. Complexity science, small worlds, is the one to be used in this area.

4.5.3 Sentiment Analysis Technologies

Sentiment analysis usually refers to the analysis of natural language (NL), including the sort of NL you find in SMS and Whatsapp, plus biometrics to identify and quantify the affective status of a person or a group of persons. A number of products are already on the market to support text analysis aiming at sentiment detection. A new boost of interest in sentiment analysis is coming from financial market evaluation, an area where the fleeting sentiment of investors leads to significant changes in the stock market. In this area blockchain technology[77] is being considered to support sentiment analysis.

Symbiotic Autonomous Systems take a broader view aiming the analysis at the affective state of the system itself, its components, and its environment.

Natural language processing (NLP) is a cluster of technologies that is largely benefitting from increased computational power and from a huge mass of data. Machine learning and deep learning can improve the NLP engine leading to the detection of subtle nuances in affective states. This goes beyond the polarity detection that is in many cases the object of sentiment analysis (to find if a community has a positive or negative feeling on a certain topic, from technical ones like writing software with certain tools[78] to assessing the like or dislike in a political contest).

Would machine become sentient?

As machines become more pervaded by artificial intelligence and in a way will assume unplanned behaviors (not in a negative sense, only in the factual sense of being self-generated by AI in ways that have not been designed, like AlphaGo[79] that played unexpected moves), it will become usual to associate "characters" to machines as is done with humans (and other forms of life), and this character might change depending on different situations. At that point it could make sense to apply sentiment analysis to machines as well.

In the shorter term humans will be (already are) conditioned by machines, and sentiment analysis should take that into consideration. Today the relationship with a machine may give rise to frustration or sometimes awe in that it gives an unexpected benefit. In the next decades, as machines match human behavior and intelligence, the relationship is bound to become much more complex and subtle.

In a symbiotic autonomous system with a human component, as an example a person with an artificial limb, as the separation from the person component to the prosthetic components fades away (as it is bound to be the case with prosthetics that seamless integrate in the body, receive signals from the brain and provide sensory feedback to the brain), the sentiment analysis although targeting the human component has to take into account the whole system.

4.5.4 Machine/Human Integrated Learning Technologies

The Web is the largest reservoir of knowledge, machines knowledge is on the rise

The amount of data, information, and knowledge in today's world is enormous and keeps growing at an amazing pace. In the past, this was shared between books and people; today it is shared also with machines and the web, with the latter having the bulk of the share and with the machine steadily growing their share. Humans have already lost their leadership as reservoir of knowledge and there are already a few areas where the mass of data is so huge to be beyond the possibility of humans to grasp them without the intermediation of computers or machines. To mention just one, the Large Hadron Collider produces 600 million MB of data per second of operation[80] (600 TB per second).

The process of learning in ancient times was a person-to-person relationship. It evolved with books (tablets, papyrus, parchments first) that over centuries became more and more important both as repository of information and as a learning tool.

The abundance of information is obvious today, but it was felt in the past as well, indicating that abundance is a relative concept ("distringit librorum multitude" – the abundance of books is distraction- Seneca -2000 years ago). Interesting the observation of Denis Diderot, the editor of the 1755 Encyclopédie[81]:

> *"…one can predict that a time will come when it will be almost as difficult to learn anything from books as from the direct study of the whole universe. It will be almost as convenient to search for some bit of truth concealed in nature as it will be to find it hidden away in an immense multitude of bound volumes."*

Machines and the web to become a symbiotic-shared knowledge base

What Diderot missed (understandably) is that machines would begin managing data, information and knowledge. They are also becoming tools for learning, and in the context of this white paper, their symbiotic relationship with humans creates a shared knowledge base. The access to knowledge is becoming easier so that it matters less where the knowledge is actually stored. There are social, political and ethical issues affecting the sharing of knowledge and its location, but these will be tackled in a later part of this white paper. From a technology point of view, the actual location of knowledge is becoming irrelevant as far as it is a shared knowledge.

In a symbiotic relationship, the knowledge present in one of the components of the symbiotic relationship is a knowledge of the symbiotic system as a whole. Consequently, the learning of any component in the system corresponds to the learning of the whole system. This means that we are now confronted with the question of where it is more effective to learn in a symbiotic system, which of its component is more suited to the learning of a specific topic.

Obviously, this makes sense only if it is true that all components are in a symbiotic relationship with respect to knowledge, i.e., if any component needing a specific knowledge can access it, seamlessly, when need arises, independently of where that knowledge is stored. The access, hence, becomes crucial, and access needs to be seen as an interaction: it is no longer the retrieval of data (a query to a data base) rather the sharing of needs resulting in the sharing of knowledge.

Learning technologies have been focusing on human beings and how to improve human learning. Significant advances have been made in the last decades leveraging computer and Internet power compounded with the availability of more flexible and ubiquitous devices. This evolution will continue as more understanding on learning processes in the brain becomes available and more effective technologies for gathering, communicating, rendering and personalizing information becomes affordable. Research efforts are currently looking at the possibility of augmenting brain learning capability by tweaking with the brain, as an example through electrical stimulation of the hippocampus[82] or elevating magnesium levels[83] in the brain. (Research results in 2016 pointed out the fragile nature of memories in our brain and the possibility that electrical stimulation of the hippocampus may actually destroy memories, rather than improving the memory processes. A lot of caution is needed in this area).

At the same time machine learning is progressing rapidly, due to more processing power and more storage availability in machines plus the possibility to leverage the experiences of thousands of machines in the cloud. Autonomous systems can greatly benefit from embedded learning capabilities and from learning from each other and as a community. This machine learning tends to merge into human learning given the overlapping of several aspects, although clear differences exist (today making learning easier for humans but the balance is rapidly shifting to the machines).

Learning has, for eons, implied access to something or somebody, who owns the knowledge and is willing to share it in a way that could be learned. One way of sharing, of course, is to write down the knowledge in a book. This goes for explicit knowledge, and we can see this kind of knowledge (easily) passed on to an autonomous system by uploading it to its "brain" (extending its data base or its programming capabilities).

Implicit knowledge

There is another kind of knowledge, implicit knowledge, like riding a bike, that cannot be coded into a book. You will never learn to ride a bike by reading a book, no matter how precisely it has been written or how many times you read it. You have to experience, fail, and learn from failure.

This kind of learning is possible for autonomous systems that can be programmed to experience and improve. Walking robots can learn to walk better and to walk on rough terrain by experience. Roomba learns about its environment by exploring the space as it does its vacuum cleaning chores.

There is also a learning that requires building of knowledge. You learn something did not exist before you thought about it. Research is an example; finding the demonstration to a new theorem is another example.

Building knowledge is a time consuming process. Autonomous systems, equipped with deep learning technology, are able to explore new ways and create knowledge faster than humans. For example, there is software that can demonstrate theorems that have not been demonstrated before and software that can play a game (like Go) creating new strategies that have not been learned from any book (or observing any other entity doing it).

Building knowledge

An autonomous system can learn by "arguing" with itself, like AlphaGo did to get better at Go. It started with the normal learning process, by looking at what good players do. Then it played thousands of games against itself learning from the outcomes and getting smarter and smarter through a process of deep reinforced learning. AlphaGo neural networks were trained on over 30 million moves actually made by Go players, becoming able to predict with a 57% accuracy the move a player would execute. This is also an interesting capability for an autonomous system: predicting what may happen next. The possibility for an autonomous system to autonomously learn opens up the issue of losing control of the system itself, i.e., the system may learn and therefore act in ways that have not been designed, nor, potentially, expected.

Collective learning, also called ensemble learning, will become more and more common. It is already a reality with Tesla cars. The autopilot system on a Tesla car has been programmed to learn as it gets more experience. In addition, since 2016, each Tesla car reports its experiences on a daily basis, and this creates a collective experience that greatly increases the learning speed of each car. The collective experience is processed centrally and emerging lessons are then distributed to all cars. It is like each car, every day, would drive over 1 million miles (the Tesla "fleet" is driving every day over 1.6 million miles. Clearly several cars are driving along the same road. Still, they are driving it at different times so they will acquire different experiences), clearly harvesting a huge experience.

Collective learning

There are a host of technologies that are being used and experimented in the autonomous systems learning that are contributing to this area, including[84]:

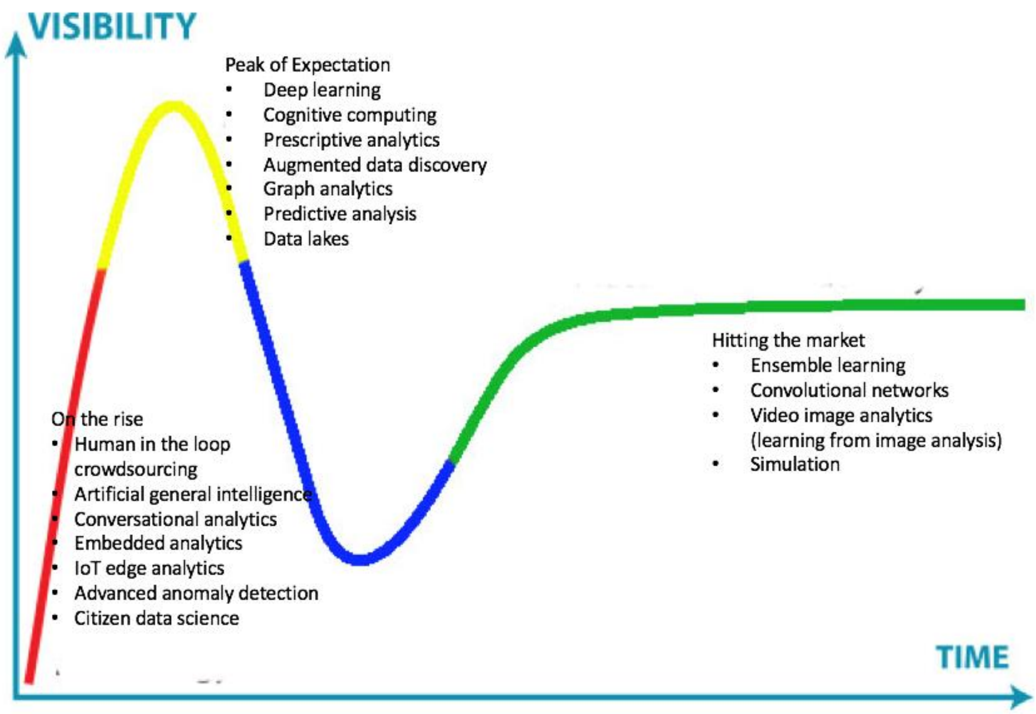

Figure 4.1. Technologies supporting autonomous systems learning

Notice among the technologies on the rise the human in the loop crowdsourcing which directly connects to the learning of Symbiotic Autonomous Systems.

4.5.5 Self-Replication Technologies

Life is, by definition, a self-replication technology. From viruses and bacteria to complex organisms, each form of life exists because it has found a way to replicate itself (and in this process to become more fit within its environment).

Inanimate objects, like artefacts, do not have the capability to replicate themselves. The difference between a living organism and an object gets fuzzier as we approach a molecular level. Here the difference between a virus and a complex molecule, like a protein, is not overwhelming. Indeed, as viruses need to leverage the external ambient (a cell, a bacteria) to replicate so a protein needs to leverage an external mechanism to replicate. This is what happens over and over in human cells.

Here the point is not to discuss if a protein may be considered as "alive" or not, rather if the same mechanisms that lead to a protein replication may be applied to more complex objects leading to replication capabilities. Notice, however, that this creates the ethical issue of deciding if a self-replicating machine should be considered alive or not.

At molecular level, researchers have already created replicating strings of DNA, and this has led to self-replicating DNA computers[85], as an example. The whole sequencing of the genome is an example of replication at work.

Smart materials are being studied to offer replication capabilities. Robots have been designed since the last decade, and significant work in this area is steered by space exploration where self-replication is considered very important. Work at Cornell University on replicating robots made by cubes that can self-assemble is pointing out to the rise of a new science[86], that of self-replication, based on the measure of the level of information replicated. If one is able to replicate 100% of the information (cloning) then one would have replicated completely however this is not necessarily what may be desired. As an example, a robot replicating itself may want to maintain its identity and create a robot that is almost like itself but with a different identity.

Notice that self-replication does not imply the capability to harvest the materials required for the replication. Clearly a supply of that material should be available but this can be provided by a third party (as it is the case for living beings that often work in a symbiotic relationship with others to become self-sustainable).

The acknowledgment that self-replication may lead to a new being that is not 100% equal to the original one opens up the point of evolution through replication as has happened to life on Earth. However, it should be noted that natural evolution required eons and a multitude of random variations which is not the case for a self-replicating machine. Here evolution can happen through replication by design, and indeed researchers are working to capitalize on the experience of a machine to improve its self-generated offspring.

The need for an adequate supply chain to fuel a self-replicating machine is also limiting the replication: in general, the more complex the replicating organisms the more time is required for the replication taking into account the need to create and maintain an adequate supply chain.

So far this discussion has assumed that the machines are made of atoms. Actually, there is a new class of machines, made of bits, that need to be considered and for these the replication takes on a different flavor and is subject to much fewer constraints.

Software bots

Cloning of software in the sense of activating several instances of an application is normal and is not considered as replication. On the contrary, the creation of software bots[87] that can replicate and roam the web is a form of replication. (Soft-bots refer to robots made of atoms, using silicone like substances making them soft, and these fall under the previous category of atom based bots).

Today software bots are based on weak AI, in the sense that they can be very smart but in some very narrow endeavor. DeepMind AlphaGo has proven to be extremely good at playing Go (hence very smart, smarter than the human world champion) but that’s (basically) it. You cannot converse with AlphaGo as you would with your friend and not even with Siri (which is another software bot specialized in another area).

Work is going on to reach strong AI, an artificial intelligence that for its breadth compares to human intelligence. Technology is not close to reaching strong AI, and once reached its implementation at the level of software bots is not a given.

Replication of swarm components

An area of research is the intelligence of swarms, and within swarms it is easier to envisage self-replication in individual components of the swarm. The collective intelligence can actually steer towards the replication of some or all of the swarm components.

Interestingly, self-replication technologies are creating new legal issues. Manipulation of the genome to obtain crops with specific characteristics is protected by law (the genetic modification can be patented). The issue is what happens to the second generation of crop. Here opinions diverge. Monsanto won a ruling[88] of the Supreme Court in 2013, enforcing its rights on modified seed (soya beans) even after self-replication.

Even if ruling for crops can go one way, the ruling over software bots or replicating robots that may change or evolve, may be different. It is a new area of study, beyond technology.

4.5.6 Autonomous Capabilities

According to the Imperial College foresight study, in the technologies that will have a disruptive impact beyond 2040, extreme automation has center stage: swarm robotics, battlefield robots, and AI board members and politicians.

Robots are becoming more and more autonomous. At the same time, they are becoming more flexible and are equipped with a variety of tools, increasing their usability in many areas. Bringing these robots to the market is an exercise in balancing performance with cost. It is obvious that the simpler the robot the easier it is to manufacture and maintain and the lower its cost. At the same time, increasing its complexity would extend its capability and possibility of use. An intermediate approach is to use several simpler robots cooperating to perform more complex tasks.

We can see this approach at work in natural systems: ants and bees are clear examples, but they are not alone. Human beings are another example: when we work as a community we can do much more than what any single individual can do, and that goes both in creating artefacts (like a car or a city) and in creating knowledge. The total is greater than the sum of its parts.

Probably humans are the first species that have become so good at harvesting the intellectual capacity of individuals to create a higher intellectual capacity. Until some time ago this increased capacity was created by one human exposed to knowledge created by other humans; now it is starting to happen in machines able to leverage our knowledge to create new knowledge through deep learning or artificial intelligence.

Cooperative through implicit communications

Cooperation, in general, does not come for free. To have simpler entities communicating to create a more valuable output requires investment in communications. However, there are examples where communication is not explicit; it does not require effort. Rather it is implicit (see the discussion on implicit communications in the first IEEE SAS Initiative white paper) and as such does not require an extra effort. Consider swarms. Bees and ants invest very little in communications; by far they use implicit communications. By flapping its wings, the bee temperature increases, and this increase is perceived by other nearby bees that change their behavior. Ants leave a trail of odorous molecules, and this trail affects the behavior of other ants. The evolution did the trick of transforming these implicit messages in higher level community behavior.

Scientists are trying to do the same with robots: swarm robotics. They are foreseeing a broad variety of applications, from Mars exploration[89] to characterizing a geographical area[90], from sensing in the sea[91] to future health care[92]. The basic principle is common to all applications: use a multitude, from ten to ten thousands simple robots each one behaving according to simple rules that connect its behavior to the environment leading to a self-orchestrating behavior, just like bees and humans.

Swarm robotics

In the coming decades these "simple" robots will become more sophisticated and the relationships among them will also become more sophisticated (as is the one orchestrating neurons in our brain) giving rise to the emergence of intelligent behavior. It is therefore reasonable to expect in the 2040 timeframe a disruption from swarm robotics in several areas, from the inside of our bodies to the environment to planetary exploration.

Notice that in swarms there is no single control point, and single participants in the swam (robot) are self-influencing one another in a dynamically evolving way, as we expect to happen in the future when robots will be able to learn and evolve based on experience. At that point it will become difficult to predict the behavior which raises legal and ethical issues (who is in charge in the setting up of the framework of evolution and who will be responsible for unplanned or undesired behavior?).

4.5.7 Decision-making Technologies

Technology in the military field has been on the leading edge in the last two centuries, benefitting from huge investment. It has also created significant fall out in non-military applications.

Artificial intelligence and robotics (tied together ever more) are seeing significant investment by the military, all around the world, although it is difficult to pinpoint the real status achieved. Fighter planes, although manned, are becoming more and more autonomous; drones are being remotely controlled but are also becoming more and more autonomous in flight operation and decision taking. Soldiers are using more sophisticated equipment, including robotic exoskeletons, that are clearly paving the way towards robotic soldiers where decision making will be shared among the human and the robotic component. Even though for a little while the ultimate decision will be taken by the human component, the AI on the robotic component will provide such an in-depth analysis on such a variety and multitude of parameters that it will influence the decision to a great extent. Actually, the influence will be such, and in several cases it is already the case (even in commercial aviation where pilots rely on the Flight Management System and Instrumental Flight System), that it is getting trickier to assign accountability.

Military AI and robotics

In the military area the deployment of robots has the capacity of extending by an order of magnitude (10 fold) the battlefield control[93], providing such a competitive edge that every army is pushing the envelope towards more effective AI. Here again, the issue is that the huge amount of data becoming available that can influence both decision and execution, is beyond the human capability and AI has to take over, basically in an autonomous way.

All companies operating in the defense area are working on more advanced robots transforming the concept of battlefield. Some, like QineiQ[94], are also voicing the need for an overall reconsideration of rules as robotics and artificial intelligence are no longer fitting the current internationally agreed rules.

This is a more general issue affecting all autonomous systems: *Who* is responsible for their behavior, given that they are autonomous? This is an issue being addressed by the IEEE FDC Symbiotic Autonomous Systems Initiative.

An interesting white paper[95] recently released by the US Army Research laboratory explains the Internet of Intelligent Battle Things. This is an area where we are already well advanced and where disruptions are already occurring. It is reasonable to expect that by 2040 wars will be fought in a completely different way; for example, most of the wars will no longer involve a physical battlefield; they will be fought in cyberspace[96].

Internet of Intelligent Battle Things

Don't underestimate the casualties however. Bits may turn out to be deadlier than bullets. In 20 years, we will be living in symbiosis with bits, with our and others' Digital Twins. We will have sensors and actuators on our body and in our homes. Malicious hacking on these may have deadly consequences. The economy is already running on bits. A disruption in the daily flow of bits can be devastating.

In the end, even though there will be killer drones using AI to take autonomous decisions, and robotic soldiers fighting with one another, most of the damage and casualties may come from cyberattacks.

Battles, of a different sort, are also fought in companies and in countries, within board meetings and parliamentary halls. Here again, we can foresee dramatic changes fostered by artificial intelligence being used to evaluate the impact of decisions, define strategies and take action.

Political elections are already flanked by experts using AI to analyze data harvested from social media to pinpoint people's mood and to craft the right message that can swing public opinion in a desired direction. Analysis of social media can provide an accurate forecast on voting outcomes; this is now moving to a new level to assess how those votes can be changed through a focused campaign. In the end, it is again a matter of money and resources. The point is to identify the areas that with a minimal investment can be conditioned to change their vote. Of course, every political party or vested investor, is trying to do exactly that in a never ending pursuit of winning the game.

Political election applications

There are companies specializing in the application of artificial intelligence, like Deep Knowledge Ventures[97], that are providing services to assess people's mood. Others, like Tieto[98], are developing software to support the companies' boards to take decisions.

In the coming decades we can expect artificial intelligence to get better, not because of better algorithms or chips, but because there will be more data to access and analyze including historical data supporting machine learning. In other words, AI is bound to become smarter and smarter, and in an area where there is a deluge of data it clearly has an edge on human analysis.

Ethical, legal and societal issues are at the forefront of these kinds of applications. There is no doubt that boards and political parties will keep making use of AI, and that will change the rules of the game. Politics and struggles at the board level have always been a matter of analyzing information and finding ways to twist moods one way or another. What is new is that we are losing control on both the analysis and twisting, relying on AI.

4.5.8 Complex Systems Technologies

The many interactions among autonomous systems are creating conceptually complex systems, i.e., they cannot be reduced without losing some key characteristics.

A single bacterium is a complex system; current artificial autonomous systems are way less "complex" than a bacterium but still many fall under the category of complex systems, particularly as they become symbiotic with already complex systems.

The sets of relationships an autonomous system has with its environment can often be described through the theory of small world with sets of weak and strong relationships or links. This is because the number of relationships, particularly for systems that move around, like a self-driving car or drone, is quite large and the quality of relationships varies a lot.

Some of these relationships are passive, like a car becoming aware of a dog; a few may involve direct communications (like car to car communications). Modelling of these relationships is an important part of a successful autonomy.

The degree of complexity in an autonomous system includes both the system itself as well as the relationships the system has to face. There are ways of measuring this complexity, like statistical complexity and self-dissimilarity. More work is needed in this area with specific reference of complexity in Symbiotic Autonomous Systems.

Also, notice that telecommunications systems in general and the Internet[99] specifically may be seen as complex systems for their high number of component elements and the variety of their interactions. With the shift from hierarchical architectures of the past, where complexity was managed in terms of hierarchy hence highly reduced (one may claim that telecommunications electromechanical systems and even the first generation of electronic switches were "complicated", not "complex") to the flatter hierarchy of today, the complexity of telecommunications systems has grown. The advent of IoT with millions of connected devices having an autonomous behavior that affect the overall network is further increasing this complexity.

The drive of the telecom operator to manage the 5G network in a rigid way may fail given the rise of the edges and their evolution in a "chaotic" way. It is most likely that 5G networks will have an increased level of complexity greater than current LTE networks. Applying complexity metrics to today's telecommunications networks and simulating first, then measuring, the complexity of future 5G networks may be a good topic of research with several practical effects.

5G, for its characteristics of also being a communications fabric self-created at the edges by autonomous systems may prove to be a key component in its evolution. The variety of protocols that will be embedded in 5G provides the latitude required for communications between and within Symbiotic Autonomous Systems.

Several domains, like smart cities, health care, and production processes are becoming complex systems. Notice that a complex system is, in a way, "complicated" but the difference is that complication is an essential characteristics of a complex system, and it cannot be reduced because the system is complex. On the other hand, many systems are complicated, but it is possible to reduce them into individual components each of which is "easy" and also the relationships among them can be seen in

Complex vs. complicated systems

subsets making them "easy" (both to understand and manage). A complex system's complexity cannot be reduced since complexity is an integral part of it. Bacteria can be decomposed in terms of its cellular organs, and the metabolic relationships can be identified and separated. However, what you get from this decomposition is no longer understandable as bacteria.

4.5.9 Emergent Properties Technologies

The relationships among the various components (physical and behavioral) of a symbiotic autonomous system are perceived by the context as its *emergent properties*. Interaction with other systems and with the environment takes place through these emergent properties, since they are characterizing the SAS.

4.5.9.1 Emergent Behavior

An emergent property is a property that the system has as a whole, but none of its components possess. Hence the decision making happens at the whole system level, and there is no specific component in the system responsible. This happens normally in (insect) swarms where decisions emerge out of the collective behavior of the swarm, and there is no individual component in charge.

A set of autonomous flying drones can in principle be programmed with a central intelligence/command, creating a hierarchy, or it can be programmed with a set of rules that results in emergent decisions. This latter approach has more resiliency, since there is no commander whose loss would hamper the swarm activities. The Internet is an example of massive distributed control for packet routing leading to an extremely resilient system from the point of view of end to end connectivity. At the very beginning of the Internet, routing strategies like the hot potato routing were studied and implemented to ensure a high reliability of network connectivity. This has evolved in other variant, like the cold and mash potato routing[100], specifically designed for autonomous systems.

5G at the edges may also be engineered as a swarm-like infrastructure where the connectivity (at the logical level, data transfer) is managed in a collective way with no single entity in charge of routing. Massively distributed IoT may be engineered to form a "swarm" and to have the swarm as a whole in charge for taking decisions.

Autonomous systems operating in a symbiotic relationship (like micro-bots embedded in a living being) will need to make decisions in absence of a coordinator, using a completely flat hierarchy, and the decision making process will be an emergent property of the symbioses.

Studies of nature where these emergent properties are usual, like in bees swarms[101], starling flocks[102], and even brain decision making processes are leading to an understanding of basic rules that can be coded into single autonomous systems and their components to give rise to intelligent decision making processes.

The emergent behavior of Symbiotic Autonomous Systems depends on decisions being taken by the system as a whole and these in turn depend on the knowledge of the system as a whole. In Symbiotic Autonomous Systems each component has its own knowledge that in general remains within that component, i.e., it is not exchanged with other components. However, because of that knowledge the component behaves in a very specific way so in an indirect way a component knowledge influences the whole system, and all together they give rise to an emergent knowledge. Notice that the specific knowledge may be as tiny as converting a pressure into a digit (like a sensor) or as complex as the analysis of big data sources.

Component knowledge influences emergent behaviour

The emergent knowledge property is interesting since in the end the knowledge of a single component does not matter, rather the knowledge owned by the whole system is important. Again, it is important to distinguish this emergent knowledge, specific of complex systems, from a distributed knowledge that is a collection of individual knowledge of the various components. In principle one can always "download a distributed knowledge" into a new system, copying the individual knowledge components. On the contrary, it may be impossible to "download" an emergent knowledge since this does not exist anywhere but is emerging from the behavior of the whole system. A point in case is the discussion of "brain downloading". We surely "store" individual knowledge in various parts of our brain (although it is not exactly clear where, as an example, we store the knowledge that dogs bark) but there is plenty of emergent knowledge that is the result of brain activity, and it is not stored anywhere.

4.5.9.2 Emergent Knowledge

The emergent behavior resulting from decision making, as said, is depending on the emergent knowledge and of course it is also depending on the way knowledge is processed, on intelligence. Here again we have in Symbiotic Autonomous Systems "emerging intelligence", and intelligence that is not residing in any specific component, rather it is the result of their symbiotic relationship.

All these emergent properties are part of complex systems and can be explored through complex systems technologies, small worlds, artificial intelligence, and chaos theory.

4.6 Digital Twins

A Digital Twin is a digital representation of physical assets (physical twin), processes, and systems that can be used for various purposes. Digital Twins can represent objects and entities as varied as a turbine, a robot, a whole ship, a cow, a human being, or a city, and everything else in between. More recently they have started to be used to represent intangible entities like services, processes and knowledge.

Digital Twins are already used in design, planning, manufacturing, operation, simulation and forecasting. They are also used in agriculture, transportation, health care and entertainment. Applications will continue to grow through the next decade; hence it is not surprising that they are named among the ten most strategic emerging concepts for the coming years by Gartner[103], or that MPL Systems[104] expects 25% of asset-intensive companies to be using them by 2020 and supporting technology spending of $10.96 billion in 2022.

Digital Twins are becoming a key aspect of Symbiotic Autonomous Systems since they allow a symbiosis spanning across the world of atoms and the cyber-world. The symbiosis can be established between a real object, including a human being, and its Digital Twin to leverage the latter in the cyber-world. Most operations taking place in the cyber-world are:

- cheaper
- faster
- clonable (e.g., they can run in parallel)
- reversible (they can be undone)
- unconstrained by locality

These properties are very interesting, and they considerably augment the capability of the object.

When considering Digital Twins, the key word is "represents". Basically, a Digital Twin mimics in bits an object's atoms and their structural/functional relationships. It does not necessarily represent all of them (something conceptually impossible, as you cannot represent a single atomic electron cloud with unlimited precision), but what matters is that the representation is accurate enough to support the goals that have been identified and that are being pursued. For example, if you want to check the proper working of an engine you need to represent all aspects that are functional to that goal, e.g., you may disregard the color used to paint parts of that engine. However, if you are mirroring a car then the color of the paint is important because retouching a car after an accident requires knowing the original paint color.

Representation with intention

Note that a Digital Twin can also, and usually does, contain more data than its real counterpart. As an example, an engine's Digital Twin is likely to contain the list of suppliers of the various components of the engine as well as the identity of the robots and of the workers that assembled it. A Digital Twin is also a historical repository of its counterpart. Thus, in the case of an engine, it may include extensive data on maintenance events and operations, for example, the minute-to-minute monitoring of airplane engine data including rotation speed, oil usage, pressure, temperature, and so forth.

All these data sets can be used for real-time analysis and simulation. They can also be used collectively to identify patterns and meanings. Take the example of General Electric (GE), which creates a Digital Twin for each of the turbines produced[105]. Once these turbines are assembled on a windmill to generate electricity or deployed on an aircraft to generate thrust, the turbines report operation status back to GE in quasi real time. This information is then compared with data generated by each unique Digital Twin for consistency. Any deviation activates an application to analyze the discrepancy and take action if needed, such as ordering the turbine flying on the plane to reduce power and decrease the rotation speed to safeguard the integrity of the engine. Of course, this affects the other Digital Twin engine on the aircraft—in this case making sure that balancing measures are implemented by increasing the thrust of the other engine and repositioning the wings' moving parts to maintain equilibrium. At the same time, the applications will look for an emerging pattern related to the situation (present and past) of other Digital Twins and will store data on any mismatches.

Turbine example of Digital Twin

There are similar scenarios with human Digital Twins. For instance, a person has data on Facebook that identifies her friends, information on her travel logged on Instagram, and a Twitter account that shows her reactions to events. Additionally, she might have sensors on her body (a smartphone) that tracks her daily activity and provides further data. More data may come from health records, if she's willing to share them, and in the future she could even have the data from the sequenced genome. Applications can continuously analyze her Digital Twin and detect emerging patterns that may require attention. This is particularly true in the health care domain,

Human Digital Twins

where Bill Ruh, CEO of GE Digital, in a recent presentation stated, "I believe we will end up with health care being the ultimate Digital Twin."

Similar to a Digital Twin of an object like a turbine, a human Digital Twin is more than a representation of the person within a specific domain of interest. It contains the copy of that person's past, very possibly keeping memories of something long ago forgotten. There is a need, therefore, to distinguish between the instantaneous Digital Twin, which represents a person at a specific moment in a specific context, and the global Digital Twin that remembers what the real twin had for dinner a year ago and what pill took to ease digestion.

Monitoring capabilities

A Digital Twin can be used to monitor its real twin and to simulate the effect of some actions (e.g., increasing the rotation speed of a turbine or changing a person's diet). It can be used to derive relevant information from other Digital Twins, such as detecting a malfunction that could affect other turbines or determining the side effects of a particular remedy. Statistical information and pattern data can be used to monitor changes in a particular activity, for example, turbines on a specific assembly line showing a power decrease in certain conditions or several persons taking two different kind of pills being subjected to undesirable side effects.

Impersonating capabilities

Digital Twins can also become "impersonators", where they can act out in cyberspace the part of the object in the real space. Hence, they can be used when designing a new object to study the interactions that may happen, as well as to solicit a Digital Twin to learn from those interactions, and then to transfer what has been learned to the physical object. This may be particularly useful in robotics where the Digital Twin of a robot can be solicited by other Digital Twins, including ones representing a specific situation and can try different approaches to find the most effective one. This experience (or knowledge) can then be shared with or downloaded to the real robot, providing it with an experience that it could not have had in the real space, perhaps because the situation could not be replicated at will or because the "real" experience might damage the robot.

There are many situations that are easier and cheaper to replicate in cyberspace with no likelihood of collateral damage and many examples of applications getting smarter by challenging themselves and learning from the experience. In the future, Digital Twins may become an essential component in the evolution of machines and the growing symbiosis between humans and machines.

4.6.1 Creation and synchronization

Having considered in depth the various aspects of Digital Twins and their usefulness and applications one has to consider how can Digital Twins be created and kept in synch with their real twin.

Digital Twin part of design and creation process

The creation of Digital Twins of objects can happen as part of the object design. This is becoming easier since more and more objects are being designed using supporting tools that in turn operate on a digital image of the object, be it a rivet or a whole aircraft. In the design phase the designer can specify the key aspects of an object that need to be mirrored and can make sure this happens, in most situations, as a result of the design activity.

When an object is manufactured, there is the need to associate the blueprint of the object (potentially created through a computer aided design (CAD) tool that might have been used in the design phase) to the real object. This is done by creating a clone of the blueprint and associating a unique identity tying it to the real object. From that moment on the twins will be forever connected.

As an example, Nikon engineers can design a new camera model and using CAD, thus creating a digital image of that camera well before it is manufactured. By using computer aided manufacturing (CAM) the camera is manufactured, starting from the digital image created by the design process and for each camera they associate an identity connecting it to a copy of the digital image which becomes that camera's Digital Twin.

Once the camera is delivered to a retail shop and then to the client, the Digital Twin is updated with the information (usually the customer is required to register the camera online thus updating the camera Digital Twin). Any time the camera is sent in for repair the various operations are recorded on the Digital Twin. When the customer downloads a new firmware version on the camera the Digital Twin is updated, and so on, to effectively create a maintenance record.

More sophisticated use of the camera Digital Twin might be achieved by monitoring the actual operation of the camera, e.g., the number of shots taken, the kind of lenses used and so on. As digital cameras are becoming more and more connected one can expect this to become normal (of course this creates privacy issues with the customer that might object to his shots being tracked).

Plenty of objects existing today do not have a Digital Twin. In many instances it may be possible to retrofit them with a Digital Twin, either directly by adding sensors that can provide data to the Digital Twin (the data received will actually build, over time, the Digital Twin) or by having external sensors (including human operators) track and record some of the object's aspects.

Retrofitting a Digital Twin

As an example, aging cars are subject to periodic revision, and one can set up a system where at each revision all data about that car can be harvested to create and continually update that car's Digital Twin.

Humans fall into this latter category. They are not designed, i.e., the process of creating a new human being does not involve any CAD, at least so far (although the advent of genomic engineering may change this). A person's Digital Twin can be created by harvesting data about that person, the more data being harvested and the more continuous is the update, the more accurate the Digital Twin will be. Data can be harvested by a third party, like a doctor or a nurse, or automatically using ambient, wearable, embedded sensors and also using data provided spontaneously by the person.

A personal digital health record is a form of a Digital Twin; likewise, the data collected by Facebook are de-facto creating a Digital Twin.

Keeping the Digital Twin in synch with the real twin is obviously crucial. We can only trust the Digital Twin if it is in synch with the real twin. Every Digital Twin needs to a specification of the degree of alignment with the real twin (for example, time since the latest update). Synchronization can be ensured by connecting the real twin with the Digital Twin, e.g., by having sensors on the real twin that periodically

Synchronization

update the Digital Twin.

That applies to humans but there are some basic issues to consider in areas like education. Having a Digital Twin that is abreast with the expertise of a person can be very good because it becomes possible to evaluate the fitness of that person to a specific task: Is this person skilled enough? What kind of specific training is needed?

It is obviously possible to track the education profile of a person (keeping track of the education career, experiences, or participation in courses) but how can what that person has forgotten be tracked?

A symbiotic Digital Twin, i.e., one that continuously connects with the person in every situation, can aggregate so much information that it might be possible to spot the degradation of knowledge and skill and therefore take action when needed.

This kind of symbiotic relationship greatly increases the value of the Digital Twin but at the same time raises stronger privacy issues.

4.6.2 Multi-dimensional Digital Twin

We are rapidly moving from the concept of a Digital Twin being a digital representation of its real twin to a semi-independent entity that is actually "richer" in a certain way than the real twin. This is what is called "multi-dimensional" Digital Twin.

A Digital Twin embeds or mirrors the real twin, so that, as an example, one can do simulation on the Digital Twin to evaluate the effect on the real twin. However, another dimension may be the historical tracking of the real twin, i.e., accumulating in the Digital Twin the information of what happened to the real twin and keeping copies of how the real twin was in the past (like Apple Time Machine keeping a copy of your computer, thus letting you to go back to a certain date).

Yet another dimension derives from the interactions that the Digital Twin has in cyberspace with other Digital Twins or with entities that simply interact with it. As an example, LinkedIn creates a sort of Digital Twin of a member based on the information the member provides, then adds more data about people browsing that person's profile. There might be tens of dimensions associated to a Digital Twin, and of course this makes the Digital Twin very useful but also raises questions.

Over the next decade, many objects will be created with their own Digital Twin and will live in symbiosis with them. In various situations, as described, robot intelligence will emerge through interactions with a Digital Twin in cyberspace, and over time the evolution and continuous learning of robots and other objects will undoubtedly be fostered by Digital Twins.

As with any new technology, Digital Twins are prompting ethical, legal and societal questions. Imagine a company with hundreds of human and robot workers, each one with a Digital Twin. Suppose a robot breaks down and needs to be replaced—wouldn't it be normal to associate the new robot to the previous robot Digital Twin so that it immediately inherits the previous robot's experience? Of course, no discussion about that.

Now consider a human worker who decides to retire or change jobs. What about her Digital Twin? Will it remain the property of the company and as such be used by the company to train a new worker? What might happen when it becomes feasible to replace a worker with a robot? Can the company associate the human Digital Twin to the robot, hence transferring the experience previously accumulated by the human worker to the robot? Is it possible in the future that companies will hire humans just for creating a Digital Twin that along with a robot will make them redundant?

Implications of Digital Twins as workers

As a robot Digital Twin learns by interacting with a human Digital Twin, is it likely to become smarter and smarter, accelerating the process of human displacement in factories and, more generally, in the labor market? These are just a few of the questions that are emerging as we walk the unexplored trails heading towards a future that probably is just around the corner.

All this support is freeing us from the need to know many things, beyond the possibility of knowing, and to focus on something specific. That makes us smarter than our ancestors and, in a win-win game, my being smarter makes others smarter in a never ending loop.

There is actually more in store in the future that will make us, individually and as species, even smarter: the fuzzy boundaries between us and cyberspace and between us and machines, a fuzzy space where Digital Twins thrive.

Knowledge of a task is already distributed and will get only more distributed in the future. Many tasks that were once done locally (in the brain) are now done using a machine (for example, squaring a number with a smartphone). As long as there is access to the knowledge, there is no detriment to distributed storage of knowledge. We have come to accept this distributed knowledge; most of the time we are not even perceiving that it is distributed. As long as I can get it, seamlessly and effortlessly, it looks like it is "my knowledge".

This is a challenge to educators; they need to educate students/professionals to harvest a distributed knowledge. This is a challenge for IEEE, they can no longer be a repository of knowledge; they need to transform it into a distributed knowledge seamlessly integrated with individual knowledge.

As one of this white paper's co-authors, Witold Kinsner, says,

> *"The IEEE Educational Activities and the IEEE Education Society not only must work together, but also with the IEEE Member and Geographic Activities, IEEE Publications, and IEEE Technical Activities to reshape the education process and the sooner the better. They have to develop Digital Twins in education. The educational twins are needed at all stages, from the young to the seasoned."*

As we are moving towards a more symbiotic relationship with machines we will come to accept that our knowledge extends to include the knowledge provided by machines. This knowledge can be represented in bits, and it can be part of our Digital Twin. Actually, given the progress expected in the area of augmented reality (AR) we will experience a continuous overlapping of knowledge in bits with the practical knowledge we need here and now.

For example, in an AR-enabled world, a concrete scenario where we mumble, "Uhm, how big is that couch? Would it fit in my living room?" would result, first, in seeing the measure of the couch overlaid on it and then the image of the couch surrounded by the artificial image of our living room.

It can also be something more abstract like looking at a poster of a movie and seeing a clip of that movie, looking at a theatre ads and hearing the voice of the main actor, or looking at a monument and feeling immersed in that historical time.

These examples move slightly out of the concept of present knowledge (what does my living room look like, what is the movie about) to the recreation of the past. The magic with Digital Twins is that they provide data that can be contextualized as needed, both in space and time.

More than that, a Digital Twin keeps the record of all its time instances, how it was yesterday and a year ago. This data can be used by an application to create customized knowledge. This applies to machines, as well as to humans. A doctor having access to (a part of) a human Digital Twin can discover the reason of a food allergy by obtaining knowledge (in a synthetic form) of food habits (again, a very trivial example).

Knowledge and Digital Twins

Digital Twins will become a crucial component in our knowledge space. They will be about the present and the past. More than that, a smart knowledge creation application can mine many Digital Twins and condense what makes sense to us, here and now.

The ocean of bits is not just about the surface we are seeing but it is about its depth. And in this case the depth represents the past.

Knowledge that will help us shaping the future will be based on what it is today, what it was in the past, and the changes that led to the present. In a way this was always the case but with Digital Twins it will be even more so.

The web has been transformed in the last 15 years from a repository where we look for information to a set of applications providing information. A similar process is going to take place in the education space: from looking at books where information is contained to seamlessly accessing applications that deliver information here and now, as needed. Note, it is not about satisfying the lazy; it is about enabling and grasping a richer and complex knowledge space.

Additionally, consider that each one of us is no longer a student for the first part of his life; we need to remain students throughout our life, and education can no longer be structured for classroom delivery. It has to become customized, to the person, to the time and to the context.

4.6.3 Augmented reality (AR) as a bridge between atoms and bits

The availability of Digital Twins is clearly important in cyberspace where they can be accessed and where they exist as independent entities. However, they can also be very useful if they become visible in the real world. Moving from real objects to their Digital Twins requires transforming some real aspects into bits (mostly via sensors). Moving from the digital world to the real world requires actuators to convert the bits into something that is perceivable by the end user, be it a machine or a human being.

In case of machines, like the engines of a flying aircraft that have to be told to trim their power because a simulation on their Digital Twin has pointed out potential overheating, it is a matter of defining a communications protocol that the engine controller can understand. In case of humans the situation is similar but the protocols are already defined, meaning the ones humans understand through their senses.

In these last years, and more so in the coming decades, AR will provide a very effective tool to present bit related information in ways easily understandable by humans. Projecting information on the windshield, showing artefacts in an ambient, and displaying images of patient radiography overlaying them on the operating field for the surgeon are different examples of effective use of AR.

4.6.4 Virtual reality (VR) to embed in a Digital Twin

A Digital Twin can become visible, as explained in the previous section, by leveraging AR. Another aspect of a Digital Twin, connected to its multidimensionality, is that it may serve as the basis for virtual reality (VR).

Take as an example the Digital Twin of a watering hole attracting animals in Kenya. This Digital Twin will have in one dimension (the mirroring one) the exact status of the watering hole at that particular time, but in another dimension it can have the historical record of the various time slices for the previous years, each showing the animals clustering around it at different times. In this dimension the Digital Twin may be studied by environmental researchers to evaluate the fauna consistency in various seasons and to compare it with historical records.

This dimension however, could also be used by an agency advertising photo safaris in Kenya to advertise the location and attract clients. For this they may use the Digital Twin as a VR stage where customers can look as if they were on site and then be inclined to buy the real thing.

Of course a different company can use that same Digital Twin, also in virtual reality mode, to provide entertainment to clients, e.g., letting them play with a lion. The starting point is a real environment with real lions (all converted in bits, which makes it far less dangerous), and through simulation programs and rendering, the clients can be immersed in that environment, using haptic devices, like sensory exoskeletons, that are able to recreate tactile sensation so that you can actually pat the lion.

This is an interesting twist to the concept of Digital Twins. It starts from replicating reality and moves into leveraging the Digital Twin to create a virtual reality.

Clearly, this is but an example. There are plenty of interesting applications, from education and training to design, active simulation and so on.

Replika[106] is a company that leverages the Digital Twin of a deceased person to let another person continue interacting with it. Interestingly they are adding a new dimension to the Digital Twin by recording the interaction so that the Digital Twin remembers it and will act accordingly the next time it is interacting with that person. In practice it is not just about interacting with a deceased person; in a way it keeps the Digital Twin of the person alive (up to date) by continually being in synch with what you are telling him when you interact in the virtual space.

Digital Twin of deceased person

Clearly, this raises new social and psychological issues.

VR is, in a way, executing the Digital Twin, like an operating system is executing a program.

4.6.5 Self, selves and super-self – A future where our self will experience no boundaries

Human history has been characterized by a progressive expansion of boundaries. From the agricultural society where most people lived and died within a 20-mile radius of their birthplace to today's web of flights encircling the planet that make it possible to reach the other side of the Earth in a day.

While technology made this possible, it actually did something even more significant: it changed our perception of time and space. It shrunk the world and densified our communities—and in doing so, changed ourselves and the way we perceive the world. In the next 30 years there will be an even more dramatic change, fostered by technology evolution, that will be more about creating a sense of pervasiveness of our self in the world.

Technologies like augmented reality today are separate from us. We need special goggles or a smartphone, but in twenty years' time, augmented reality will be part of our senses through, for example, electronic contact lenses first, followed by eye lens, retinal, and brain implants. Our senses will be extended, seeing things in the infrared, nano-pulses, or ultraviolet, hearing things at high frequencies, seeing the electromagnetic spectrum, or feeling presence at a distance. In a way we will get augmented reality through our own "augmentation".

More radical technologies, fraught with ethical concerns—like genomic engineering—are basically inevitable. While the engineering part is becoming a commodity, the big hurdles today are understanding the effect of genetic manipulation. New approaches based on artificial intelligence will connect the genotype with the phenotype, making "humans *a la carte*" a reality. Notice that among the ethical concerns include the effects on offspring following such a mutation, for example, offspring that could be generated through the mixing of a non-mutated genome with a mutated one.

Our augmentation will go hand in hand with the "embedding of a soul" in artefacts: our ambient environment and its constituents will become more and more aware and able to interact with us on a peer-to-peer level.

We have basically passed the Turing test, and fake news and fake interactions have become a major (unexpected) side effect of this evolution. Our Digital Twins will start to have a life of their own, and we might end up living parallel lives, an augmented one in the physical world and several ones in cyberspace. The big problem is that it will be more and more difficult to find a boundary between us, the world of atoms, and cyberspace. This is the amazing change we will experience in the second part of this century: the disappearance of boundaries. The philosophical question "Who am I?" will take a completely different flavor.

4.6.6 Cognitive Prosthetics

Significant progress has been made in this last decade to interface prosthetics to the brain, using a few electrodes (on the skull or implanted) to pick up brain electrical activity and use that to control an external prosthetic, like a robotic arm or a robotic wheelchair. Usually the person "learns" to control the prosthetic by engaging in some specific thoughts activity. Hence it is not the prosthetic that learns to read that person thoughts.

The more electrical activity can be picked up and the more precise the locations, the easier it is to control the prosthetics, and more complex activities can be orchestrated. That is the reason for the DARPA challenge: Neural Engineering System Design[31], resulting in brain implants able to pick up electrical activities from a million neurons.

Researchers are at work to win the challenge, however, creating a seamless cognitive prosthetic is a much more difficult challenge. Here the crucial point is "seamless". In a way we already have cognitive prosthetics today: our smartphone is an extremely effective cognitive prosthetic. If I do not know something, a few clicks on my smartphone and the world knowledge is at my fingertips, similarly for performing a variety of tasks, translating into another language, navigating a foreign city, or doing math.

4.6.7 The fading of boundaries between the self and the augmented self

Technology is becoming pervasive as will technology for human augmentation. Additionally, we will become accustomed to live with the augmentation provided by technology (as we already are, although today's technology—like smartphones—will seem primitive in 20 years) to the point that we will no longer perceive it. This raises the issue of distinguishing between the self, i.e., the person I am, the way I behave and feel without a technology boost, and my augmented self.

Notice that this is already an issue today for people taking mood altering drugs. Which is the real self? One could say the real self is the one felt when no drug is taken. On the other hand, if you have a fever and you feel tired, is that the real you or is it the one feeling much better after having taken a drug to relieve the fever? One could say the real self is the one emerging from the use of the drug. Are these two situations different? Not really. One person can feel depressed because there is a problem with his brain (or glandular system producing hormones that affect the way the brain feels) and taking a drug restores a "normal" situation thus making the real self emerge.

This is just an example to point out the challenge already faced today that pushes some people to say "take the pill and feel better", while others would say "don't take the pill, it is just creating someone that is not you".

In the future with technologies that will augment the human body and brain capabilities and sensations, thus impacting feelings and moods, including technologies that can affect the person before birth (by tweaking the genome) or change a person during his lifetime, it will become more difficult to separate the real self from the augmented self since it will be very difficult to define the real self.

Suppose scientists find a genome mix that makes a person much more likely to be happy and thus, parents start asking for a genetic engineering of their spermatozoa and oocyte to create happier offspring. Is the real self of that child the one he is born with or the one that he would have been had his parents not asked for genome engineering?

Prosthetics are becoming more sophisticated and personalized; they interact seamlessly with the body, relaying sensations to the brain. Already today some people wearing them consider them an integral part of their body. In the future these prosthetics will not just fill in gaps, they will provide augmentation. Will people realize the difference between themselves and their "prosthetic selves"?

What about BCI that (in a much more distant future) would provide seamless connection to the world knowledge through a web cloud? Will people feel themselves as ignorant because they will seamlessly rely on the cloud or will they feel savvy and empowered?

These boundaries fade with symbiosis; and the sensation feeling as one with the other components, be that one's Digital Twin, a limb prosthetic or cyberspace, takes precedence.

4.6.8 The fading boundaries between reality and virtual reality

The fading boundary between the self and the augmented self is mirrored by the fading boundaries between reality and virtual reality, between atoms and bits. In the coming decades, technology will make moving from reality to virtual reality seamless to the point that it will be difficult to appreciate the difference.

By wearing special goggles, we can see in the infrared and ultraviolet. A flower may look quite different looking at it using those goggles or without the goggles, yet it is the same flower. It is not becoming any less real looking at it in the infrared.

The same concept goes for an engine. By donning infrared goggles, I can see where it is hot (so I will not touch that part), and it is still the same real engine. What if the goggles, through augmented reality, will let me see cogs turning at low speed so that I can see if there is a problem? Well, clearly it is the same real engine. And what if with augmented reality I can see "inside" the combustion chamber, looking at the valves? And what if I see the simulation taking place at the level of the Digital Twin enacted, through augmented reality, on the real engine? Is what I see still real, or is it virtual?

Super reality beyond AR or VR

The boundary gets fuzzy, and the issue is that augmented and mixed reality are creating a "super" reality, as real as the real one. Symbiotic Autonomous Systems will likely leverage this mixture; symbiosis might actually be the result of this mixture.

Humans will operate in symbiosis with their Digital Twin; sometimes it will be their Digital Twin that will interact in cyberspace with other Digital Twins, like a robot Digital Twin to execute a specific action. Technology evolution will make this seamless, and the perception will be of directly executing the action, not of delegating the execution to someone else.

For example, pilots flying by wire have the perception of controlling the ailerons (a hinged surface in the trailing edge of an airplane wing, used to control lateral balance, similar to a flap) directly, while in fact it is a computer that takes their joystick input as the indication of what they want to do and translates it into an appropriate action depending on the situation. As an example, the movement of the joystick by two degrees may result in a shift of the ailerons of five degrees at a certain speed and altitude and in a nine degree shift at a different speed/altitude. Yet the pilot feels he is moving the ailerons directly.

4.7 Security

This section addresses security of transactions between machines and humans in the context of Symbiotic Autonomous Systems. It considers cybersecurity of networked computers and computing devices, cyber-physical security (CPS) of networked computing devices connected to physical processes, cyber-social security (CSS) and cyber-physical-social security (CPSS) of networked computing devices and society. Several recent security improvements of the Internet of Things (IoT) connectivity are also highlighted.

4.7.1 Introduction

4.7.1.1 Why Is Security So Important?

Our planet is well-balanced, but very complex. In complex dynamic systems, small changes in the input can produce large effects. The planet can minimize many of the effects, but not all.

Life is also very complex. Most of the time all is well. However, a tiny virus can swing the homeostasis of a body way out of kilter. Our immune system exhibits a very smart behavior, capable of fighting the viruses. Most of the time, it wins. Sometimes, the virus wins.

Living is exciting because of the complexity, uncertainty, and the presence of other systems equally complex. This balance is not certain, however, because of the natural and induced changes to the ecosystem.

Let us consider a major disturbance on this planet. Through science, many have been convinced that our planet is facing an unprecedented climate change crisis. Massive research and development (R&D) programs are being launched to develop technologies and infrastructures to reduce emissions and greenhouse gases. The 2009 meeting in Copenhagen agreed on keeping the rise of the average temperature below 2°C above the pre-industrial level[107]. In 2014, the Intergovernmental Panel on Climate Change declared that doing so would require reducing greenhouse gas emissions by 40 to 70% from the 2010 level by 2050[108].

This is a very difficult challenge to meet. What would such a miraculous development look like within the three largest offenders: electricity generation, transportation, and food and agriculture? Ten examples are discussed in the June 2018 issue of IEEE Spectrum[109]. However, the problem is that each of the new proposed developments may be attacked and altered to flip their intended purpose, thus accelerating the damage. Such attacks have been done in the past by implanting an agent (usually a person) in the environment where the damage was planned. Today, the attack can be carried out on the Internet (cyber-attacks) because we are more and more interconnected directly or through "things." The Internet of things (IoT) increases the number of possible attacks. The integrity of the IoT infrastructure and monitoring the CO_2 emission becomes very important, and cyber-attack prevention is crucial since IoT is a distributed infrastructure under multiple domain ownership.

We have seen (i) cyber-attacks in the form of denial of service, (ii) cyber-physical attacks in the form of altering the behavior of motors, and (iii) cyber-physical-social attacks in the form of changing the outcomes of how people live in a community.

More specifically, the science of computer security addresses the following main attributes[110]:

- Availability (system being operational)
- Integrity (data not being altered)
- Confidentiality (private data not being disclosed)

Availability

In the SAS domain, availability is extremely important since Symbiotic Autonomous Systems support humans in performing essential tasks that they could be unable to do without the help of machines. That is especially relevant in biomedical fields (imagine prosthetics or other life-support devices like smart insulin pumps, pacemakers, etc.). Human health can be severely affected when symbiotic machines stop working.

Integrity

Integrity is also of paramount importance, since correct SAS operation can be more critical than being up and working. In some cases, it is better to shut down a machine in a fail-safe mode, wherever fallback modalities are available, than letting it work improperly with possibly catastrophic consequences. Imagine a self-driving car or train being hacked: when an integrity fault is detected, it is better to brake safely and shut down than continue running in unsafe conditions.

Confidentiality

Finally, confidentiality is something more than plain privacy: keeping sensitive data out of reach from unauthorized users and possible perpetrators prevents many threats including stolen identity, frauds, and critical attacks to SAS integrity and availability performed using stolen information.

In security risk assessment, those attributes are addressed from the viewpoint of external or intentional or human made malicious threats. However, random faults are also addressed in the more general framework of computer dependability. Traditional techniques to achieve fault/attack-tolerance include intrusion detection systems, spatial and temporal redundancy, active reconfiguration, etc.

To make things more difficult, SAS are becoming increasingly complex and at the same time critical for our lives and well-being.[111]

Complexity is the effect of several factors, including:

- Large size of software required to implement smart, intelligent and autonomous devices, featuring non trivial failure modes that are difficult to diagnose
- Distribution due to ubiquitous Internet connection and pervasiveness (e.g., wearable

devices connected to local-metropolitan-geographical "fogs" or "clouds")

- Heterogeneity in hardware, software, wireless and wired network connections, usage of open-source or COTS (Commercial off-the-shelf) components together with proprietary/custom implementations, etc.

Criticality is given by SAS being employed in safety or life-critical, money, or business critical applications. While traditional computer-based safety-critical systems are kept as simple as possible in order to ease verification and validation as well as assessment and certification against international safety and security standards (e.g., ISO/IEC 15408 – Common Criteria, IEEE Std 1686-2013 - Standard for Intelligent Electronic Devices Cyber Security Capabilities[112]), the inner complexity and limited predictability of SAS represents a serious obstacle. At the same time, engineers can advantageously leverage artificial intelligence and other SAS-relevant paradigms (e.g., Digital Twins) to prevent and fight threats more effectively. Those aspects will be addressed in one of the next sections.

4.7.1.2 What Class of Countermeasures Might Be Effective?

Much work is being done to develop effective countermeasures against the different classes of attacks. The need for individuals who could be useful in the area of security increases. When hiring various agencies, companies, and organizations, the sought-after knowledge and skill set required incudes: data science and engineering with backgrounds in computer/electrical engineering, computer science, cybersecurity, information assurance, mathematics, cryptanalysis, signal processing with analysis and synthesis, security, and counterintelligence. It is critical to determine not only *what* these individuals can do, but also *why* they do it, which impacts the safety and security of our facilities, families, communities and planet.

However, since the number of persistent attacks increases and is mostly unchallenged, something is inadequate in our countermeasures. In principle, countermeasures based on past observations are not adequate. Since the attacks often belong to the class of dynamical systems, they have to be fought with dynamical systems. Our research in this area is based on this approach. We propose in this White Paper that Symbiotic Autonomous Systems should be developed to deal with the diverse classes of attacks.

Let us first consider the security of SAS themselves, followed by cybersecurity, cyber-physical security, and cyber-physical-social security.

4.7.2 Security of Symbiotic Autonomous Systems

SAS are complex dynamical systems by design in order to exhibit behavior that is non-trivial. Consequently, SAS are vulnerable to attacks by foreign unknown perpetrators. Like us, they are complex, with many interacting components and subsystems, and with nonlinear relations, operating in a nonlinear world. Like with us, a tiny virus can invade an SAS, reproduce rapidly, and even kill it. In order to have a chance to fight the virus, the SAS must have an immune system capable of producing many non-similar trial "antibodies", and when one of them matches the invader, the system should shift from producing trial scouts to producing a massive army of similar "soldiers" to suppress the invading enemy, hopefully in time.

Humans are still here on this planet because our immune system uses a very smart process to fight the invaders. The fight is not easy because the invaders mutate and are new to the immune system (otherwise our immunity system would have enough antibody "soldiers" in place to eliminate the old known invaders).

If the SAS immune system is fooled by the invading attackers that nothing is wrong, the SAS can then be altered by modifying the nonlinear interaction just slightly, with very large possible changes in its outcomes and performance. The SAS reality can be altered.

Since the outcomes of such modifications may be unpredictable, a catastrophic failure of the system may occur. Since SAS are symbiotic with us and the environment, such a failure may take us all down. Since the SAS discussed in this White Paper are intended to improve (not destroy) our human condition, security of SAS is existential.

4.7.3 Cybersecurity (CS)

Cybersecurity is defined as a "computing-based discipline involving technology, people, information, and processes to enable assured operations. It involves the creation, operation, analysis, and testing of secure computer systems. It is an interdisciplinary course of study, including aspects of law, policy, human factors, ethics, and risk management in the context of adversaries" [113].

Cybersecurity also refers to the secure operation of a computing system, including information technology (IT), local, national, and global networks. Although it is a $100-billion industry, few companies feel secure. From June to December of 2009, Google, Adobe, Yahoo, Symantec, Northrop Grumman, Dow Chemicals, Morgan Stanley and others were attacked from China (codename Operation Aurora)[114]. Google also introduced the so-called ZeroTrust policy: no outside or inside person is trusted[115].

There is no standard set of rules used to address the increasing number of threats from hackers, ransomware, and stolen data. The National Institute of Standards and Technology (NIST) has released Cybersecurity Framework[116] and will issue a companion document, the Roadmap for Improving Critical Infrastructure Cybersecurity[117]. The framework identifies five functions to organize a security system: identify, protect, detect, respond, and recover. The framework was intended for small business, and its intended impact was to defragment the cybersecurity world.

Critical Infrastructure Sectors

The Department of Homeland Security identifies 16 critical infrastructure sectors: Chemical Commercial Facilities, Communications, Critical Manufacturing, Dams, Defense Industrial Base, Emergency Services, Energy, Financial Services, Food and Agriculture, Government Facilities, Healthcare and Public Health, Information Technology Nuclear Reactors, Materials, and Waste; Transportation Systems; and Water and Wastewater Systems[118].

Software vulnerabilities can be exploited by attackers. The most common cybersecurity incidents are due to conventional malware and virus (64%), ransomware attacks (30%), and employee errors and unintentional actions (27%)[119]. Around 40% of Internet Connection Sharing computers experience attacks every six months. The majority of organizations (61%) have inadequate security management, and 66% of manufacturing businesses have no dedicated budget for their security.

4.7.4 Cyber-Physical Security (CPS)

Cyber-physical security refers to the secure computer-based operation of any physical environment, such as found in industrial control systems (ICS), self-driving cars, drones, etc., in which cyber-physical systems, which are systems including sensing, actuating and computing components, are connected to the Internet. Due to the Internet connection, the system becomes vulnerable to security threats happening in "open networks".

Since Symbiotic Autonomous Systems incorporate connected physical devices to interact with human beings and the environment, they are also vulnerable to cyber-physical attacks.

4.7.4.1 Physical Vulnerabilities

As an example of a physical vulnerability in such cyber-physical environments, the operation of a centrifuge can be altered a little (an unnoticeable change) in order to slow the process a bit. This is what happened with the Stuxnet implant into the SCADA-controlled Iranian centrifuges in a nuclear facility and brought it down in 2010[120]. A SCADA (Supervisory Control and Data Acquisition) system is an application that allows human operators to monitor an industrial process and to store and analyze the corresponding process data.

Physical vulnerability also applies to data storage, such as a smart grid which needs to be monitored for near real-time data related to its operation in order to optimize the delivery of power to a specific destination. Automated and smart meters are used to measure the energy consumption of electricity consumers and to provide a variety of value-added services. Since the collected data provides sensitive consumer information, the data must be protected, which is a difficult part of the privacy side of cyber-physical security[121].

IoT Security

An emergent but rapidly expanding area of data storage—and thereby physical vulnerability—is Internet of Things (IoT). The Internet of Things interfacing has been expanding rapidly in terms of complexity, speed and security. In harsh industrial and military environments, IoT systems design calls for the following four considerations: (i) a three-tier environment: edge devices, gateways and back-to-back systems, (ii) communications between all the devices must be wireless, (iii) modularity and interoperability, using heterogeneous suppliers of parts; (iv) an open-source foundation, at each layer[122]. Security of IoT related data are addressed in a following section.

The IoT systems must be upgraded, enhanced and maintained with respect to hardware, drivers, operating system kernel, frameworks and development tools on a regular schedule and when needed.

The focus should always be on security[123]. Security should be important always, not only after an attack. We must ensure device integrity and authentication, robust user and access control, and above all strong encryption in all data communications.

An IoT gateway is a system that connects IoT edge devices (sensors and actuators) with back-end applications running in a data center or in the cloud. An IoT gateway is co-located with edge devices, collects information from them, then processes and reformats it before sending it to back-end applications. In addition to the security of the gateway itself, the gateway must interact with—and be protected from—edge devices and the backend command and control systems that direct and manage the gateway.

IoT gateways increasingly perform data reduction and integration, as well as local processing and decision making[123]. In the case of an unmanned aerial vehicle (UAV), for example, IoT devices would include visual and infrared cameras, GPS, multi-axis accelerometers, engine sensors, engine controls, flight surface controls, and others. In this case this information is a mixture of streaming data (video), regularly updated data (position, engine operating conditions), and events and alerts (engine overheat, loss of oil pressure), as well as commands issued to the UAV. The gateway would be responsible for collecting information from all devices, processing it, and sending a subset of available information out through the uplink to a remote command and control system.

Other examples of physical vulnerability include:

- The alteration of the operation and explosion of a petrochemical plant in Saudi Arabia in 2017[124].
- A connected self-driving car can be manipulated from a distance with obvious fatal consequences.
- A connected MRI machine can be altered to produce incorrect images of the brain.
- An Internet-connected infusion pump can be altered to administer a lethal amount of medicine to the patient.
- A connected fridge may be altered remotely to spoil some foods to bring diseases not only to the family in the house, but to those in contact with that family.
- An attack on a dam might be devastating to millions of people. There are some 100,000 dams in the United States.
- An electrical system can be shut down to a house, or a street, or a district, or to the entire city. The damage may be fatal regardless of how fast or slow it occurs.

Since most of the above examples have been tried and some were successful, countermeasures to stop the perpetrators should become more and more sophisticated. Reactive countermeasures are not sufficient; SAS should be developed to deal with such cyber-physical attacks in real-time and hyper-real time.

It is clear that any cyber-physical system must now be designed not only for robustness (i.e., surviving attempts to temper with its physical devices), but also for resilience (i.e., the ability to recover after a successful attack).

4.7.4.2 Questions About Interacting with Devices

From a security perspective, there are several questions to consider regarding how the gateway interacts with these devices:

• Is the device available, and can the gateway connect to and interact with the device? In some cases, the gateway may need to do device discovery to see what devices are available. In other cases, the gateway will already have list of devices to access.

• Is the device who it says it is? The gateway should assume that all devices are untrustworthy and should verify the identity of all devices. All input from devices should be sanitized before it is processed. It should not be possible for a device to take over the gateway. Flow control and quality of service (QoS) mechanisms should be used to prevent devices from launching denial of service (DoS) attacks. In addition, mechanisms should be used to securely identify each device. Each device can be expected to provide information on itself such as device type, model, capabilities, and serial number. Certificates such as X.509 certificates installed on each device allow the gateway to verify the device identity.

• Has the device been compromised? Techniques such as signed software images, boot time integrity checking in the device (secure boot), and self-check mechanisms are needed. Cryptographic signatures, commonly implemented using certificates, should be used to verify all aspects of the device.

In many cases communications between devices and the gateway should be encrypted. This way, potentially sensitive information is not exposed, a man-in-the-middle attack cannot modify information flowing between a device and the gateway, and the identities of the device and the gateway can be verified. Encrypted communication is commonly done using transport layer security (TLS) which, with the proper choice of ciphers and keys, can provide security even with constrained devices.

The gateway itself should also be hardened by deploying a few proven techniques. First, only install the software that is needed should be installed, while unnecessary programs are uninstalled. No software development or debugging tools should be installed on production gateways. The gateway should be proactively managed, maintained, and updated. In doing so, Secure Boot should be used, which allows verification of the hardware and low-level software, and trusted processing modules (TPM).

Linux-based processes should be used. For instance, resource management, such as Linux cgroups, should be proactively employed, which can help specify the maximum and minimum amount of CPU, memory, and network traffic an application or service may use. Software signing, such as that which is used with Linux RPM packages, is also beneficial for allowing verification of the source of the software and that the software installed on the system has not been modified or corrupted. It is important to utilize access controls, such as SELinux, to control system resources and prevent compromised applications from accessing other parts of the system. Employing domain-based authentication solutions, such as Linux IPA, can allow centralized management of user accounts and optimal security through multi-factor authentication.

Of course, security is not complete without the appropriate firewalls and scanning solutions in place. Firewalls ensure that only allowed communications occur and can be configured to only let devices with specific IP addressed “talk” to the gateway. Regular scanning can allow for checking for unpatched security vulnerabilities and secure gateway configurations.

Finally, it is important to actively manage the gateway through regular updates, including security and bug fixes and the addition of new features as appropriate. A mechanism to remotely and securely manage gateways should be used for these tasks.

While effectively securing an IoT gateway involves many different processes and tools, at the end of the day, it boils down a single thing: keeping embedded systems safe.

4.7.4.3 Support Vulnerabilities

Software companies provide support for a product for a set reasonable amount of time. Hardware companies have been providing support for a much shorter amount of time (2-3 years). The situation started changing in 2018 (7-10 years).

This is a problem because connected products may generate small profit and may require years of updates, patches and security improvements before they themselves are upgraded to provide adequate security features.

4.7.4.4 Heterogeneous Supplier Vulnerabilities

Many connected devices are assembled with parts (hardware, firmware, software) from different suppliers. This is of particular concern because if the weakest link is not updated regularly, the entire system becomes vulnerable. While in the information technology (IT) environment, the life-cycle-management industry tracks patches and updates to a buggy software, the cyber-physical security does not have this capability yet.

Since Symbiotic Autonomous Systems incorporate connected physical devices and connected individuals, they are vulnerable to the cyber-physical attacks. It should be clear that new standards should be developed by IEEE for patching and updating heterogeneous SAS.

4.7.5 Cyber-Social Security (CSS) and Cyber-Physical-Social Security (SPSS)

Cyber-social security refers to the secure operation of a community or a society. If the cyber-social security involves physical devices, the threats become even more consequential.

In today's world, we can learn much from the Internet. Since the search engines provide too many plausible answers to answer even a simple question, there is no time to study them all, or even to consider the top contenders, if the number of answers is in the thousands or millions. The best we can hope for is to get a taste of the answer.

We should develop better Internet information digesters, information and knowledge curators, and advisors. In fact, some have already been developed; the numerous news digests deliver to our mailboxes or through other means filter information each day. For example, Flipboard and many other similar digesters select articles based on our interests. Amazon suggests purchases of books, music, video and physical products based on our previous purchases or expressed interest. Digital Twins described in this White Paper might do the job much better.

However, security must be considered as an important part of this area. An attacker may alter the sequence of answers to skew the understanding of the concept that we want to develop, so that we might develop an alternate reality, as intended by the attacker. Interspersing reliable answers with false ones may also skew the understanding of the reality. With sufficient priming of the skewed reality, false news could appear to be plausible.

Security of Internet information delivery

In education, the skewing process has a chance to be self-correcting because of the discernment that learners are normally known for. If the skewing is applied to medical record, to production reports, to financial reports, or to political views used in decision-making, the cost of the skewing might be much greater. In trust-based situations where there is no time for verification and quality assurance of the information provided, the outcome might be fatal.

4.7.6 Challenges and Open Issues for Secure SAS

SAS are computer-based systems belonging to the classes of embedded systems, smart systems and/or cyber-physical systems, featuring enhanced intelligence to make them autonomous and capable of operating symbiotically with each other, human beings and the environment. As such, their security inherits most of the metrics, emerging vulnerabilities, and protection technologies discovered or developed for the aforementioned classes of computer-based systems. Although SAS complexity and criticality can be higher than most common computer-based systems and that complicates their engineering and assessment, they also feature some unique capabilities that can hardly be found in any of the traditional or existing systems. Those unique capabilities are a consequence of their enhanced intelligence and autonomy, which can be advantageously leveraged in order to design self-healing and self-protecting machines. Furthermore, they will start to show some aspects of proactive security and dependability: instead of assuming that everything in the surrounding area is working according to basic engineering assumptions (i.e., specification), they will start to predict abnormal behaviors in the environment by other peers or the humans[125].

For example, consider the case of infrastructure-controlled autonomous vehicles approaching an intersection. According to the traditional safety paradigms, the centralized control system will allow only some vehicles to move, and command the others to stand still. The vehicles allowed to move will completely ignore the presence of the other vehicles (like in the case of railway control) or at most detect obstacles on their route if they feature appropriate sensors or video content analytics (i.e., artificial vision). Therefore, in current critical systems for traffic control, abnormal vehicle behaviors that are not considered by engineering specifications can easily lead to accidents and hazardous failures. Imagine now what happens when we as humans approach an intersection and see that a vehicle is running at high speed from the side street towards our trajectory. Although we know it should stop if we have the precedence, we will slow down prudently since we perceive a higher risk that the vehicle will not be able to stop. The same holds for any other external behavior we perceive as unsafe/unsecure. That does not happen in current computer-based systems unless there is a specific requirement, since functional rules are normally very simple and easy to assess. However, it is nearly impossible to predict everything that can happen in real-world scenarios, and most catastrophic failures happen due to unpredictable situations, e.g., due to hackers taking control of some system entities. The capability of SAS to react to unpredicted scenarios is therefore something that can counterbalance their inner unpredictability, leading to an apparent paradox. That means a paradigm shift is required to move from centralized control and predictable safety/security to decentralized autonomous systems that are able to react to unpredictable threats and operating scenarios.

Similarly, and differently from human beings, current systems are not designed to and thus unable to do what is otherwise desirable, like:

- Detect abnormal behaviors in their peers and warn others accordingly.
- Warn their peers about dangers and threats they are likely to face.
- Help peers to defend themselves and recover from attacks.
- Cooperate strategically and coordinate others to better respond to attacks.

While such a level of intelligence and autonomy can appear fictional, much is already being done to achieve an at least moderate level of proactive dependability (including safety and security) by leveraging machine-learning for early warning, situation assessment and decision support. For instance, computer intrusion detection uses heuristics to distinguish anomalies from normal network traffic and user behavior. Heuristics based on Bayesian networks or artificial neural networks models can be trained to fine-tune their performance [126].

Distributed reflective architectures

At a more advanced cognitive level, the concept of "distributed reflective architectures" has been introduced for the first time about 15 years ago, somehow pioneering the current research on Digital Twins[127]. In those conceptual architectures, autonomous agents feature some form of reciprocal monitoring and control. In order for a SAS to monitor its peers and the environment, it is necessary for appropriate predictive models to be included in its "brain". Computational constraints prevented computers to implement those models in the past, but in the future, cloud/fog computing and even small device computing power will enable efficient computations. SAS will include clones or Digital Twins of themselves, their peers and other entities in the environment, in order to perform accelerated "what if?" simulations based on the likelihood of threats, attacks and any other unwilled events. Just like humans do when they worry about the implications of a possible danger they foresee, SAS will assess situations, warn humans and the others, predict possible consequences, and quickly plan effective countermeasures. Furthermore, higher levels of intelligence will allow SAS to autonomously setup advanced security strategies like honeypots and camouflage in order to mislead adversaries.

4.7.7 On-going activities and open issues

Cybersecurity, cyber-physical security, and cyber-social security has been in focus to academia, research, industrial practitioners, governments, and the rest of us. It has been a major challenge in networked computing. It affects industry, business, health care, banks, organizations, and governments. It may be detrimental to our existence.

Much has been written about security in IEEE, ACM, and nearly all other technical reports, magazines, journals, and books.

While cybersecurity is multidisciplinary, it requires a solid foundation in computer science, computer engineering and mathematics. Many universities are considering embarking on cybersecurity programs. However, the foundational knowledge on which the field of cybersecurity is being developed is fragmented, and both educators and students are considering another Cyber Security Body of Knowledge (CyBOK) project to establish the foundational and generally recognized knowledge on cybersecurity[128]. A Joint Task Force on Cyber Security Education (JTFCSE) was established in 2015[129]. They provide many organizations a place to collaborate, including: Association for Computing Machinery (ACM), IEEE Computer Society (IEEE CS), Association for Information Systems Special Interest Group on Security (AIS SIGSEC), and International Federation for Information Processing Technical Committee on Information Security Education (IFIP WG 11.8). The JTF grew out of the foundational efforts of the Cyber Education Project (CEP).

In 2014, IEEE Future Directions Committee established an IEEE Cybersecurity Initiative, and developed the Try-CybSi Project[130].

All those efforts, however, should be accompanied by the paradigm shift enabled by future generation SAS. In fact, in existing distributed control, entities:

- Always trust their supervisors
- Do not care about their peers and the surrounding environment except for the variables they need to monitor
- “Set and forget” supervised entities

In order to achieve security and resilience at 360 degrees, in addition to self-checking and self-healing, secure SAS will be required to:

- Check whether they can trust their supervisor, just like humans do, by verifying if the commands they receive are reasonable and safe
- Be careful about their peers and the environment to detect any anomalies and abnormal behaviors that are possibly symptoms of hacking and sabotage
- Check if supervised entities are actually performing the tasks they have been assigned or someone else took their control

Much effort is directed towards specifying the educational framework for security. Educational material differs substantially between various educational institutions. The training requirements of industries and business also vary widely. The major technical organizations have been formulating guidelines along this effort (e.g., [131]).

This does not necessarily require SAS to be able to recognize and classify the unknown. A feasible way to do that is to embed (simplified) models of other interacting SAS to mimic the similar behavior of human beings. To that aim, the Digital Twin paradigm can be employed as mentioned in the previous section. Intelligent cooperation to better recognize threats and quickly respond to disruptions is also one major achievement that is highly desirable in research and innovation addressing secure SAS[132,133].

4.7.8 IoT Data Security

News from Sweden in May 2018 indicated that around 3,000 individuals there had microchips inserted in their hands[134]. The chips open up doors (no mechanical or other electronic keys), book train tickets, access vending machines and printers and other devices. So, humans are now included in the "things" in IoT.

IoT devices and systems are coming to our homes, cars, workplaces, our bodies, and many appliances that were never connected.

Any data in transit are vulnerable to being intercepted (stolen) or altered and re-injected (for other purposes). With the advent of unsecured IoT, the problem of securing sensitive and confidential data is more serious. According to a recent study, 92% of users say they want to control what personal information is collected automatically, and 74% are concerned that small privacy invasions may eventually lead to a loss of civil rights[135]. The General Data Protection Regulation (GDPR) introduced in May 2018 resolves some of the problems[136].

There are many developments to protect the data, including the User Managed Access (UMA) standard. The User Managed Access (UMA) standard is a promising approach; Stevenson concluded, "UMA supports a much more user-friendly method for managing access to control over personal data. It makes it easy to grant consent, share data, and revoke consent." There are many non-technical (e.g.,[137]) and technical books and literature on the subject.

4.7.8.1 Security Assistance

There are many organizations and groups that have been developed to mitigate the rising problems, and provide help. For example, the Open Web Application Security Project (OWASP) identified the following ten common attacks threatening IoT security[138]:

A01: Code Injection
A02: Broken Authentication and Session Management
A03: Cross-site Scripting (XSS)
A04: Insecure Direct Object Reference
A05: Security Misconfiguration
A06: Sensitive Data Exposure
A07: Missing Function Level Access Control
A08: Cross-site Request Forgery
A09: Using Components with Known Vulnerabilities
A10: Unvalidated Redirects and Forwards

They also provide an repository of the attacks[139] and countermeasures[140]. There are also IoT data security companies (e.g., Imperva) who can assist in developing some solutions to IoT data vulnerabilities.

4.7.8.2 Data Encryption and Secure Computation

Data security today is a major problem. Security professionals, administrators, and executives know this because data breaches occur monthly, weekly and daily. Loss of customer trust, huge payouts in fines, damage to reputation, and business leaders losing their jobs are just some of the consequences associated with a data breach. Encryption must be a part of the solution [141].

In his book [142], Patrick Townsend of Townsend Security addresses the following questions: (i) When to use encryption; (ii) What data you should encrypt; (iii) Where you should encrypt that data; (iv) Encryption best practices; (v) The importance of encryption key management. There are many non-technical (e.g., 142, 141) and technical books and literature on the subject (e.g., [143], [144], [145], [146], [147], [148]).

Encryption refers to encoding of data, using mathematical algorithms in order to make that data undecipherable to unauthorized viewers. Encryption has evolved and changed to meet the demands of evolving technology and numerous regulations. Today, the encryption algorithm accepted as the highest standard is the Advanced Encryption Standard (AES)[149]. AES is a formal encryption method adopted by the National Institute of Standards and Technology (NIST) in 2001. AES supports nine modes of encryption, and NIST defines three key sizes for encryption: 128-bit, 192- bit, and 256-bit keys. AES has been adopted by the Federal Government as an approved encryption technology under the FIPS-197 standard. AES is accepted by the Health Insurance Portability and Accountability Act (HIPAA) and is accepted by all credit card issuers for data security including Visa, Mastercard, Discover, American Express. AES has also been incorporated into Pretty Good Privacy (PGP) encryption which is used by banks, insurance companies, benefits providers, and most major financial institutions for securing data in motion.

Encryption

Using NIST validated AES encryption and FIPS 140-2 compliant key management is critical to ensuring that a security solution will stand up to scrutiny in the event of a data breach. These certifications are difficult to acquire and are only given to encryption and key management systems that have been heavily tested against government standards. Using trusted third-party systems is typically the easiest way to acquire and implement this technology. Many industry regulations require that your security solutions have these certifications.

Research in the area continues with many new approaches to cryptography developed now (e.g., [150]). Some of the seasoned approaches may need to be changed because old and new attacks (such as the side-channel attack) that have been so successful recently. Much effort has also been going into studying quantum cryptography.

The growing popularity of cloud-based machine learning raises the questions about the privacy guarantees that can be provided in such settings. A research group at MIT has been tackling this problem in the context of prediction-as-a-service (PaaS) wherein a server has a convolutional neural network (CNN) trained on its private data and wishes to provide classifications on clients' private images. They have developed as system called Gazelle[151], a scalable and low-latency system for secure neural network inference, using a combination of homomorphic encryption and traditional two-party computation techniques (such as garbled circuits). Gazelle makes several theoretical and practical contributions while being orders of magnitude faster.

5 Societal, Economic, Cultural, Ethical and Political Issues

This section places the current technological evolution of SAS into the broader context of society and its operations, pointing out the mutual implications; i.e., how technology adoption impacts society, economy, culture, ethics, and politics; and how all those elements (including regulation) impact our investment in technology, hence steering its evolution.

There are ethical and political implications deriving from the possibility of "designing" humans and legal issues of shared culpability and responsibility.

This section evaluates the conditions under which objective and immutable interaction between the new society's selves and super-self could occur. One of the recent technological evolutions was the blockchain concept applied in this section to voting.

The closing part of this section looks into the changing meaning of democracy as citizens expand into symbiotic citizens with blurring boundaries between people, machine, artificial intelligence, knowledge and cyberspace.

5.1 A New Society: Some Aspects of Self, Selves, and Super-Self

The self is a production of literacy. It comes from the appropriation of language for personal use. The person is built not by others, as it is in tribal societies, but by a "self" that grows in the silence of the mind. This silence, necessary for protracted reflection, is encouraged by the gradual silencing of reading. The self becomes a unique entity to the extent that it can manage language silently.

Another feature is that the self is opaque to others. The persona can adopt masks other than ritualized. A person is born strongly individualized, capable of privacy and jealously protecting it.

We have seen, however, that the fundamental conditions of self have been changed first by electricity and its industrialization, followed by digitization of the physical, physiological, and mental world.

Building of self is changing due to technology; today, children build their ego online, not inside themselves. Their self is now memories, not of words or content but of experiences, physical and mental. The rest of their memory is on their phones. Whatever is left of the self is capable of ever less resistance to absurd propositions.

Transparent Self

Furthermore, traced and catalogued, the self has ceased to be opaque, and owned entirely; it is now transparent and shared by whosoever. This condition is favorable for the construction (or, rather the self-organization) of super self. The intense increase of traffic in relational media is developing instant, often viral, tribes. Different configurations of association come and go like thoughts. Opinion takes over the control of science and facts. Objectivity becomes conflated with subjectivity with ever looser boundaries. In such an environment, what is the dominant or accepted culture? What is the value of any regulation? What is politics?

In a symbiotic environment, the lack of internal resistance may have binding consequences, for better or worse. Both possibilities need to be examined. How does the super-self determine its boundary? Can the super-self see its independent selves? Does the super-self need symbiotic relations? Since the new society is still subject to the existing laws of physics in this universe, is the super-self opportunistic or existential?

5.2 Direct Democracy

As its name implies, a direct (or pure) democracy is a system in which citizens have "an extraordinary amount of participation in the legislation process granting them a maximum of political self-determination."[152] The difficulty in implementing and maintaining security in order to prevent voting irregularities has been the bane of democratic voting since its inception. That said, this secure status may be achieved in a datacratic direct democracy by designing and utilizing blockchain technology-based universal, strongly encrypted, remote e-voting (online or digital voting) in order to create anonymous, secure, publicly accessible records of the voter ID, candidate ID and the time[vi]. Note that this vision already has existed in various stages: Switzerland is currently a direct democracy since a series of constitutional changes starting in the 19th Century, but early forms have been practiced in various locations since 1291[153]; West Virginia will provide a mobile blockchain voting option November 2018 for overseas military service members in elections[154]; and an Australian startup is focused on developing blockchain voting for emerging democracies and is designing a test case community voter platform to be deployed in Sumatra[155].

Democracy was, in fact, *direct* even earlier—when the ancient Greek created the word and the concept. The early practice of "power to the people" began around 690 BC in villages (the original meaning of demos was not people, but "village" or later "assembly") where everybody, that is, all free men (no women, slaves or foreigners), were entitled to have a say in decisions made by and for the community. Not workable for larger communities, *direct* democracy evolved into an early form of *representative* democracy in Attica at the turn of the 5th century under the jurisdiction of Kleisthenes. All free male Athenians were invited – in fact obliged by law – to elect 500 councilors (called the "boule") to manage city affairs and who, therefore ipso facto represented them. Direct democracy did reappear in various guises and places (Switzerland, as discussed above, as well as Kurdistan and Mexico), but it has gained new relevance with online technologies because it seems to promise a greater participation of the voting public and perhaps restore trust in present-day political processes. Indeed, in view of the disappointing turn-out in most elections of the world and the attendant diffidence towards institutions, hope is now placed in networks to arrive at more effective forms of governments.

5.3 Democracy in the Era of Bits

The symbiosis between digital and organic decision-making is developing under an undeclared contract that the technical is at the service of the organic.

[vi] In a blockchain, a voter ID is a public/private keypair that is untraceable to voter identity.

The combined trends of access to greater quantities of usable data and the growing sophistication of various kinds of analytics have introduced a new condition of transparency in the political scene. People's activities and personal features are traced, stored, analyzed, catalogued and reused by third parties, often without much control on the part of governments, institutions or enterprises. Many see this development as a threat to democracy, direct or representative. No longer can a secret even secluded in the mind resist digital penetration (see discussion of alter ego above). We are scattered in big data at the mercy of anyone who needs to know something about us. That is not merely a Digital Twin to which, presumably we would have access, but an unchecked digital unconscious[vii] that is about to rule our lives more than anything Jung or Freud imagined. It changes everything in the republic. *Res publica*, a concept borrowed by the Romans from the Greeks, means the "public thing" as opposed to the private person and property. This public thing is the space and the services (including the government) that are the prerogatives of all. Democracy is about all of us being able to contribute to the decisions that manage this space and these services. The Internet seems to present itself as the new public thing, but its neutrality is threatened from all sides, starting with governments, institutions and companies big and small.

In fact, democracy (as we knew it, if we ever did) is under threat not only from the general surveillance of everybody, but also from the uncontrollability of fake news that is presently creating geo-political havoc. We may be witnessing a major cognitive slip of objectivity that is putting in question not only science and facts, but also our more or less shared idea of what reality consists of. Populist movements the world over signal the fact that more and more people are conflating objectivity and subjectivity. For many there is no need for referents, references or verification. Suffice to claim a "truth" to make it so; all the more reason to fear unchecked direct democracy.

Paradoxically, the idea of democracy, which also evokes the power of the greatest number, is based on the equal rights of the natural person before the legislative, judicial and executive powers. Now that the *res privata* is also becoming public what happens to these rights when the person is virtual and transparent?

The real question is what new ethics must accompany transparency. The change is anthropological in nature. Under conditions of widespread surveillance where everyone can or will have access to private data from everyone, the obligation will be to have nothing to hide, as in the old oral cultures. And, as in the old oral cultures our main responsibility will be again directed towards the other. Indeed, if Freud has more or less rid us of guilt, a private experience, we are revisiting the era of shame, a public one. Already we speak of "Reputation Capital", that is proving so fragile on the web. Democracy will come back if and when people manage to bring the institutions to account. Transparency, by addressing our leaders, should eventually bring us there.

[vii] The "digital unconscious" is the sum of all data available online. Like Carl Jung's notion of the collective unconscious, it is founded on ambient information. The collective one is supported by stories people tell, retold in documents and present in vestiges that distribute more or less evenly the collective memory of myths and archetypes that occasionally emerge in consciousness. The digital one is present in a growing quantity of databases that are managed by a growing quantity of ever more intelligent software. The digital unconscious differs from the collective one by the immanence, permanence, accessibility, and potentially instant and global diffusion of data online, instantly available for collecting and reconfiguring to emerge at a conscious level in real time.

The relationship between the individual and power has changed a lot in the last twenty years, and that in several moments. It is marked by many trends that rebalance the distribution of power between leaders and led. The people can and want to be involved in decision-making now. The invention of Twitter is decisive on this subject. The Arab Spring was able to take place, we are sure now, due in large part to Facebook (Tunisia) and Twitter (Egypt); Barak Obama was re-elected by the systematic practice of social networks informed by a very refined use of Big Data about every potential voter. That practice wasn't questioned until Cambridge Analytics and the role of Facebook in providing sufficient data to orient critical voting decisions. Even journalism has had to distance itself somewhat from power because of the abundance of news and opinions that emerged from online comments by citizens. We are moving quickly towards a decentralization of power, via the tidal wave of transparency in the revelations of Julian Assange and Edward Snowden or those that are associated with the Panama and Paradise papers.

Since national intelligence services cannot be expected to abandon their surveillance strategies, one must believe, or at least hope, that a symmetrical transparency agreement will be part of the new political order in the making. The relationship between power and the individual will change again until a new balance is reached between the government and its citizens, a state of mutual transparency where those who have put the power in place can demand accountability. Until we come to that, there will be revolution upon revolution. It would be better to avoid this if at all possible. In the end, the global society will have to come together to form a new social contract.

5.3.1 Blockchain Voting

Long-term SAS evolution will provide a range of novel benefits. However, given our evolutionary proclivity for socioeconomic class hierarchies, a key SAS consideration will not be its availability, but rather it being available to everyone. In the emerging datacracy, those without the ability to purchase or otherwise acquire SAS enhancements will *de facto* define an underclass, which would contradict the value structures promised by SAS. Fortunately in this context, datacracy (algorithmic governance) technologies have the potential to support the design and instantiation of both a networked blockchain[viii]-based direct democracy as well as a post-scarcity/post-capital ecosystem.

5.3.2 Post-Capital Ecosystems

A fully automated SAS post-capital ecosystem[156] (in which goods, services and information are universally accessible at no monetary costs) could then theoretically emerge when the above human labor-free system generates global economies of scale and algorithmic optimization to minimize costs to the point of making capital unnecessary, thereby transforming values and ethics that then prioritize societal well-being and global preservation. A post-capital supply-and-demand system could thereby leverage global crowdsourcing protocols and local/personal molecular manufacturing to operate automatically and perpetually optimize ecosystem operations, security and environmental issues addressed by datacracy-like intelligent algorithmic systems.

[viii] A blockchain is a database audit trail managed by networked computers. No single computer is responsible for database storage or maintenance. Any computer may enter or leave the network at any time without compromising database integrity or availability, and the database can be rebuilt by downloading the blockchain and processing the audit trail.

5.4 Ethical Androids

SAS prognostication should be an occasion for us to revise our ethical standards globally (that is, including consideration and appreciation for very different ethical concepts from other cultures and, perhaps also including fauna and flora rights). We are progressing half-blindly to a condition where we will have become the extensions of our machines instead of the contrary. The way we program these machines will decide on our well-being and even survival as a species.

But what is the ethical drive of a droid? Put aside the rapidly boring argument about the self-driving car that has to decide between harming one or five persons depending on the immediate solution to be chosen. It should be obvious that we cannot expect—yet—machines to be better at that kind of choice than humans themselves. Furthermore, it is not just the robot that needs moral guidance but the whole SAS environment. The individual narrow focus (weak) intelligence of the android is primarily informed by a goal-oriented programming, and that can include a myriad of checks and balance provisions, but the prescribed goal will prevail. What is really needed is the SAS environment to be suffused with both artificial general intelligence (AGI) and affective computing (AC) in every instance. This goes far beyond Isaac Asimov's prescriptions that robots do no harm. AGI would have to include factors and parameters of overall safety, environmental health, reducing poverty conditions, augmenting employment opportunities, transportation easement, legal guarantees and many if not all the topics addressed in this white paper; AC would give the droid the ability to understand human emotion and, eventually, to experience emotions themselves. Ideally, in the future, no government should be allowed to exist without such guarantees.

5.5 Datacracy

A possible definition of datacracy is government by algorithms, that is, decision-making support for policy and ruling provided by different kinds of data analytics. Although datacracy is not yet fully implemented in any present-day governance, it could become so powerful as to overtake direct human intervention in deciding between the best options for a nation or a given community. Considering that case specific data-driven verdicts are already superior to human judgment in many critical sectors, medical, legal, financial and military (see above, Section 4.5.7), it is already foreseeable that the temptation will be for many governments to seek validation for their decisions in what will be presented as incontrovertible proof of their wisdom and fairness.

There are jurisdictions in course today, such as in Singapore or South Korea where data analytics are providing sufficient and comprehensive information taken from the people themselves via the analysis of social media and other data sources to justify policy and ruling decisions made for individuals in valuating, orienting and positioning them in education, housing and health services. Security issues loom large in such practices and will be even more often invoked as the geo-political as well as local safety conditions become more threatening.

In present-day China, guaranteeing security legitimates measures that clearly infringe on privacy such as ubiquitous face-recognition technologies to identify potential harm-doers or equipping Robocop-like policemen and women with direct access to criminal and other indicting records. China, however, is taking a large step beyond such understandable policies (if not fully acceptable for other nations) by implementing the practice of giving "social credits" to individuals, based on the cumulative and permanently upgraded valuation of their behavior and accomplishments or lack thereof in their daily life and over the long term. This practice may not seem very different from what is done in Singapore, or even from has been the norm in western countries since banking loans have been approved or not on the basis of people's behaviors, assets and careers for decades. However, making official what has been so far only a tacit social agreement is indeed pushing the envelope towards datacracy. Furthermore, there are already Chinese social media platforms that allow a person's family, friends and neighbors to share in and publish their opinions and valuation of that individual[ix]. This development seems to amount to a radical shift of the self-censorship practiced by western societies to censorship by other people. It has been suggested that in a nation comprising almost a billion and half citizens, there wouldn't be enough police to control the behavior of everybody hence the move to guarantee the development of what could become a kind of "self-police state". And with the increasing sophistication of data analytics that are well on the way to penetrate individual people's thoughts and motivations, we may be looking at Orwell's dismal vision of "thought-police" but not imposed or implemented by a special government force. The question, as in all political systems, would be how to counter efficiently or prevent human abuse of the system. In a fully transparent society it may be possible to achieve such goals.

There is a possibility that SAS could help make something like datacracy more tolerable. Indeed, assuming that AGI steered data analysis (including mood and sentiment) would focus on the community, instead of addressing mainly individuals or procedures, it would take all comprehensive environmental factors in consideration more or less in real-time. Automated policy-making, regulation and execution of different measures would guarantee increased social good and thereby reach a higher consensus in the community.

In a political system that is yet to be invented, for anyone to be entitled to participate in any policy-making, voters would have to provide evidence that they were informed and competent. This could be assessed by analytics. Access to decision-making would be given according to the level of competence every citizen achieved. This is already the case informally to the extent that degrees and positions theoretically guarantee a modicum level of awareness of whatever situation was being considered.

In a political system grounded in AGI, a mature SAS environment (local or global) would have to emulate for the whole system in real time the kind of survival alertness and opportunity awareness that each one of us possess individually. That would mean, for example, not recommending a decision that would harm the environment in the long term, or identifying and presenting opportunities for improvements to social and personal processes. In a format comparable to the concept presented above of the *virtual twin* and perhaps associated with it, the contextual synthesis of all pertinent factors should be made available to anyone intending to participate in a decision or a policy. The assembly of all these participants would constitute a sort of electronic *boule*, acting for the greater good and benefit of all citizens. The transparency of all public behavior would sustain the honesty of all participants, leaders and led.

[ix] To get an idea about how far this trend could take people, see Black Mirror's Nosedive episode where a hapless young woman is driven to personal and social disaster by how people evaluate her on a daily basis on her smartphone (see https://www.imdb.com/title/tt5497778/).

5.6 Mood and Sentiment

The Internet has an important emotional dimension. It has been functioning as a sort of social limbic system. This is the system that in the mammal body carries and manages emotion from trigger to enactment. People have a natural urge to share emotions. The Internet has made sharing emotions instant along with a globally expanded potential audience. Since it has become relational, the network stimulates more and more emotional drives in fast and skillful configurations. People go on the net and on social networks to express and share indignation, happiness, hate, irony. The online world works as an integrative system of impulses, desires and frustrations, which is moving at the speed of light and ignites unpredicted responses. Social media (among other platforms) transmit and develop the emotions and target them to specific peoples and groups for action and reaction, as the limbic system does in the body. In media connected via the Internet, there are many emotional and cognitive events being transmitted from person to person, which in turn motivate the sharing of experience and also the call to political action.

It is clear that today's geopolitical map of the world has been changed by the arrival on the political scene, via the Internet, of a new class of mass political activists, who are no longer the "Silent Majority"[x]. Now that the majority is silent no more, the result is a kind of interactive social "massification" consisting of the connections between many individuals who respond to some current issue as a significant collective. The Spanish network sociologist Manuel Castells called this the collaboration of many "mass individuals"[157]. Castells identified that the relationships that are multiplied among individuals on a personal basis, from one person to another, are much more complex and articulated than those that come out of crowd reactions or those of anonymous masses. We can therefore imagine that the result of this endless interaction between individuals on the Internet is equivalent to the infinite multiplication of conversations over a cup of coffee.

The transnational movements such as the Arab Spring, Occupy Wall Street, or Spain's grassroots movement Los Indignados, all represent collective emotions and connectivity amongst peoples across borders and cultures. It is therefore critical that recognition of this increasing emotional content and power be included in the development of SAS.

5.6.1 Mood and Sentiment Analysis

Although automated for text analysis since the late 1960s, sentiment analysis (SA) has been around since the invention of literature criticism. Henceforth addressed to the audience, not to the text, SA has been developed technically at least since 2002. It is occasionally referred to as opinion mining. What SA amounts to is the new possibility for institutions and businesses to listen to clients, patients, customers and citizens instead of simply imposing regulations, services and products. Of course, SA can offer advantage to various fields including health, municipal affairs, public administration, political process and policy evaluation, transportation, banking, insurance, security and business. SA has also become sufficiently affordable and relatively easy to make it valuable, if not mandatory, for public administrations to keep tabs on their charges' feelings about their operation. Originally based on rather crude quantitative analysis of natural language word

[x] A name given to essentially the more conservative bulk of the American population.

frequencies, SA has been given new prominence, owing to the manifold increase of available data and the precision afforded by machine and deep learning that allow research to pinpoint and interpret individual as well as collective moods and feelings hence capable of drawing instant and continuous mapping emotional responses to speeches, events and other triggers of emotion.

To the basic NLP application to verbal text henceforth assisted by contextual semantic search (CSS), SA adds expression identification in face recognition, sound frequency analysis to detect mood, body language study via surveillance cameras, and any method including tactile sensors and brainwave analyzing to report with greater exactitude the state of mind of groups and individuals. A later development extends the analytic capabilities to intent analysis. Such techniques benefit from research in affective computing that is already key to the exchange of expression and recognition of feelings between human and machines. In a symbiotic environment a keen awareness of ambient—and distant—sentiments, opinions and intentions will be part of AGI and be communicated to humans much in the way intuition works in human affairs.

5.6.2 Individual and Communal Mood and Sentiment (M&S)

Affective computing is focused predominantly on individuals and machines. A great deal of research is being implemented in AI and AGI systems to harvest emotional data from users so as to facilitate and improve the functioning of machines. A predictable upgrading of Siri or Alexa will be to make them able to recognize and respond to a user's feeling by the tone of voice or the rhythm of the speech. M&S analysis brings the technologies of SA to communities and regions.

Not surprisingly early on, several artists-engineers and designers have been interested in public emotions and working on models of how they could be captured, interpreted and displayed. Maurice Benayoun, in *Emotional Traffic* (2005)[158] and Salvatore Iaconesi in *VersuS-Rome* (2011)[159] have used word frequency analysis techniques in various media to draw dynamic maps of areas where emotion-provoking events are happening, in real time or over specific periods. Using a different sensory response Christophe Bruno in *Wi-Fi SM* (2006)[160] has used the same techniques to stimulate in a user, via a tactile interface bracelet, sensations of pleasure or pain triggered by the frequency levels reached by specific key words found in Google Trends for 4500 cities in the world.

For a local sassing of passersby mood, a public art piece called MIMMI was erected by a group of designers near the Convention Center Plaza of Minneapolis. It consists in a big ring made of a plastic balloon that changes colors according to the mood of the people of Minneapolis. In the words of one of its creator, Carl Koepcke: "The mood is determined from all the tweets coming from within 15 miles of downtown Minneapolis. Our custom program then parses the words and runs it against an open source database of emotional words".[161]

A relevant project is Mappiness[162], part of a research at the London School of Economics. This mobile app and online system actively notifies users once a day, asking how they're feeling. The data gets sent back along with users' approximate geographical location and a noise-level measure, as recorded from the phone's microphone. In this way users can learn interesting information about their emotions – which they see charted inside the application – and the operator can learn more about the ways in which people's happiness is affected by their local environment—air pollution, noise, green spaces, and so on. This is an interesting mechanism, but also one that lacks the possibility to sense the natural emergence of emotions as linked to urban daily life, in people's language and expressions, as it relegates users' interactions to strictly encoded forms.

We can expect the ubiquitous development of social innovation apps by yet unknown talent that will help orient research applicable to SAS.

6 Legal and Societal Issues

6.1 Symbiosis[163]

Central to the SAS concept is the creation of a human-computer symbiosis: a synergistic marriage between the heuristically, context-driven capabilities of human cognition and the sheer volume and detail orientation of computer technology. One of the challenges to human performance—whether in military applications, business, health care or academia—is that the overall volume of information and complexity of tasks continues to grow at a rapid pace, in stark contrast to human cognitive abilities, which have remained relatively static.

Limitations in human cognition are attributable to intrinsic restrictions in the number of mental tasks that a person can execute at one time, and this capacity itself may fluctuate from moment to moment depending on a host of factors, including mental fatigue, novelty, boredom, and stress. Although supplementing our current cognitive abilities is commonplace and readily apparent in the manner in which we use our PDAs (personal digital assistants), SAS aims to study this human-computer symbiosis by looking at individual and collective cognitive interfaces along several paths, as previously mentioned. As this new technology evolves, applications in various settings and research raise legal issues of shared culpability and responsibility.

6.2 The "Shotgun" Approach

Litigation attorneys are notorious for taking what has been termed a "shotgun" approach to culpability and responsibility; said approach constitutes advancing every possible argument on behalf of the clients without regard for the actual facts, existing law to the contrary, or the likelihood of success. Generally, what this means is that any entity involved in the creation of the SAS will likely be joined in a lawsuit; the justice system, like a crucible, burns away irrelevancies to determine legal culpability.

6.3 Proportional Allocation of Responsibility

The degree of computerized assistance in decision making (and a proportionate reallocation of responsibility) has been proposed as an effective scale to measure degree of interaction. This scale can serve as a useful overlay to the four major paths of SAS (Scale of Computerized Assistance):

1. The computer offers no assistance; humans must do it all.
2. The computer offers a complete set of action alternatives.
3. The computer narrows the selection down to a few, suggests one, and either (a) executes that selection if the human approves; (b) allows the human a restricted time to veto before automatic execution; (c) executes automatically, then necessarily informs the human; (d) informs him after execution only if he asks; or (e) informs him after execution only if it, the computer, decides to.
4. The computer decides everything and acts autonomously, ignoring the human.

Because the applications vary greatly in the measurement on the computerized assistance in decision making scale, as mentioned previously, the borders between an autonomous individual and an autonomous system become rather ambiguous in these human-computer symbioses.

Although the law is well equipped to hold an autonomous individual culpable for his or her actions, the law is not so well designed for such systems. Because hardware and software are often designed by different groups of individuals, the culpability chain becomes more difficult to trace.

SAS presents a unique challenge for the legal profession to help shape policy, given that the technology is cutting-edge and very little is written in statutes, case law, or law journals. However, we can extrapolate from the existing corpus of law.

6.4 The Law as Codified Conscience: Issues of Privacy, Autonomy, and Culpability[164]

Law is a codified reflection of normative social practices: it purports to guide human behavior, giving rise to reasons for action. Sometimes the law is prescriptive, its function being to restrict human behavior. But providing sanctions is not the law's only function in society. Solving recurrent and multiple coordination problems, setting standards for desirable behavior, proclaiming symbolic expressions of communal values (such as autonomy and privacy), resolving disputes about facts, and similar matters, are some of the important functions that the law serves in our society. However, the law is not the only domain that regulates behavior in our culture: Morality, religion, social conventions, etiquette, and ethical values also guide human conduct.

As individual rights become affected by advances in technology, the legal system will evolve new definitions of traditional concepts of privacy, autonomy, and culpability.

6.5 Rights of the Individual Versus Rights of Persons

Traditionally, the law has divided entities into two categories: persons or property. However, persons are not always individuals—the U.S. Supreme Court has declared corporations, municipalities, and even ships to be persons under the law. And in the past, individuals (specifically, women, children, and slaves) were considered mere property. Fortunately, the law has evolved (through legislation and court decisions) to recognize that individuals are persons, and the law is continually evolving to recognize the stages or categories in between.

Although the basic rights and responsibilities of individuals are now better established than previously in constitutional law (albeit subject to interpretation, depending on the constituency of the U.S. Supreme Court), one of the questions facing our courts will be, where do the rights of an autonomous system begin, assuming that an individual is an inherent part of that autonomous system? Whether using a property vs. personhood dichotomy or property-person continuum, the rights of the individual may change when the human performance of the individual is enhanced by a machine or other technology. This raises issues about privacy and autonomy.

6.6 Privacy[165]

Technically, a right to privacy is not explicitly enumerated in the Bill of Rights or Constitution. The legal concept of the right to privacy can first be found in an 1890 Harvard Law Review article entitled, "The Right to Privacy," written by Samuel Warren and Louis Brandeis when they were law firm partners[166].

Warren and Brandeis claimed that the right to privacy already existed in the common law and gave each individual the choice to share or not to share information about his or her private life. Their intent was merely to establish the right to privacy as a legal protection in their day. Neither man coined the phrase "the right of the individual to be let alone," as found in U.S. Supreme Court Justice Brandeis's dissent in Olmstead v. United States (1928)[167], which is often quoted by privacy champions and is the first case in which the U.S. Supreme Court considered the constitutionality of electronic surveillance. In their 1890 article, Warren and Brandeis interpreted the Fifth Amendment to the U.S. Constitution: "No person shall...be deprived of life, liberty, or property, without due process of law" to read that a person has an inherent right to be let alone and to privacy. Their interpretation was their legal theory and their view of a more general right to enjoy life.

Of course, it has also been argued that the Fourth Amendment creates a right to privacy. It states, “The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.” This right is not absolute; the key word is unreasonable. Under exigent circumstances or with a showing of probable cause that a crime is occurring or is about to occur, the state’s interests override the individual’s rights.

6.7 Autonomy

The world autonomy has been used synonymously with control and self-determination, but the term (derived from the Greek words autos “self” and nomos “rule”) was first used to refer to the self-rule of independent Hellenic city-states. The concept has since been extended to individuals and acquired meanings as diverse as liberty rights, privacy, individual choice, and freedom of the will. The most straightforward definition is self-determination, the ability to make one’s own decisions, as opposed to letting next of kin, a guardian, the state legislature, or a judge make decisions on behalf of the individual. But we certainly do not allow individuals to decide for themselves whether to inflict harm upon other people, such as assault and battery or fraud. The law recognizes that although individuals’ liberties should be protected, these freedoms must be weighed against the obligations to others – the obligation of nonmaleficence or the “do no harm” principle. The legal system recognizes that autonomy is a dynamic concept, not a static concept. The concept of autonomy varies widely in differing social, cultural, economic, and political circumstances. Because of the dynamic nature of the law, courts often determine autonomy on a case-by-case basis using competence as a measure.

6.7.1 Competence

Competence, a close relative of autonomy, is often used as a yardstick by courts to determine whether an individual is capable of exercising true autonomy. Although autonomy and competence are different in meaning (autonomy meaning self-governance; competence, the ability to perform a task), the criteria of the autonomous individual and the competent individual are strikingly similar. The more competent (and therefore more autonomous) an individual is, the higher the level of culpability for one’s actions. The granting of rights to rational autonomous beings brings with it the burden of responsibility. It can be argued that autonomy rests in the intersection of our notions of moral and legal responsibility and personhood. Actors who cannot respond to reasons because they cannot grasp and understand them, like children or animals, cannot be culpable. What are the moral boundaries implicit in the decision-making process? For example, if a decision needs to be made for the greater good at the expense of losing other individuals, is the individual or the technology responsible?

Because the ability to grasp and understand the reasons under which we act generally requires some degree of autonomy, autonomy (as measured by the yardstick of competence) emerges as a requirement of moral and legal responsibility. In wartime or crisis environments, patterns of behavior are directed by hierarchical command decisions. So, from a legal perspective, just how autonomous (and culpable) will they be? Will we be able to apply the same laws of autonomy and competence as are currently imposed on unenhanced humans? The courts have the ability to administer a remedy that is proportionate to the rights and interests of those who lack full autonomy, but they are venturing into new territory.

6.8 Recommendations

The law is not static; it is constantly subject to change, extension and reinterpretation, and evolution, whether by legislation or judicial decisions. As SAS continues to evolve as a science, the importance of a continued monitoring and legal interpretation of the impacts of policy and research are key.

There are three areas that are in earnest need of policy consideration:

- The establishment of a common lexicon among policy makers, implementation agents, and multidisciplinary users for terms such as autonomous and symbiotic.
- The establishment of a new lexicon for the new relationships that are being created as a result of new technologies, with thoughtful consideration given to the impact on current policies.
- The possibility of legal reform and the creation of specialized science courts, in which the judges have ongoing education and training to recognize and deal with these new legal issues and categories that arise from emerging technologies.

Ideally, a standards-based document should be developed to support a full interconnectivity of all the elements, either through a federal or an international commission. As human cognition is augmented and enhancements occur through the applications of SAS, it will be up to policy makers, the courts, and the legal profession to limn guidelines for the enumeration of the rights and responsibilities of the persons as well as the entities created.

7 Market Impact

Recently, but the issue is not new at all, newspapers in several countries are pointing to a growing concern on job losses, as result of an increasing automation that is also becoming smarter and that is now overflowing from the factory assembly lines to percolate the whole working space. The icon of this expanding automation is the robots.

As a matter of fact, one just needs to perform a web search (by the way, how many people are still visiting the library to harvest information? What happened to the librarian job?) to see data that are concerning, mostly because the impact of automation is reaching jobs that would seem to have nothing to do with robots. Take the CB Insight study[168] on the number of jobs being threatened by automation in the US:

By 2023 (5 years)	By 2028 (10 years)	By 2033 (15 years)
Cooks and servers: 4.3 million Cleaners: 3.8 million Movers and warehouse workers: 2.4 million	Retail salesperson: 4.6 million	Truck drivers: 1.8 million Construction laborers: 1.2 million Nurses and health aides: 6.9 million

Table 7.1. Numbers of jobs being threatened by automation

Today, there are robots that can carry out these various job tasks:

- Restaurants with robots cooking hamburgers[169] and robot waiters[170]
- Roomba robots[171] to clean our home floors
- Robots in Amazon warehouses[172] to pick up products from shelves and take them to a packaging robot that will deliver them to a truck heading to our home.

The first robot assisted malls are a reality in Japan and South Korea, with robots welcoming clients and walking with them as shopping assistants. Robots are deployed in hotel chains in several countries now, including Italy and US, to serve as a concierge.

Truck platoons drive on northern Sweden roads[173] (with several reindeers and very few cars) and in the Arizona[174] desert stretches of land. In India robots (industrial 3D printers) are building/printing whole villages. In several hospitals robots are picking up drugs from the hospital warehouse and delivering them to the patient bedpost.

Progress in artificial intelligence is scary: it is making robots ever more flexible, enabling learning (and even self learning), empowering them to take decisions, to operate in teams and mingle with human workers, carrying out ever more complex tasks, some which were considered human turf—future proof just a decade ago.

Today we know that no profession is future proof, that is solely a human turf. We also know that a few professions can no longer exist without robots, like neurosurgeon and astronaut (to name but two extreme cases), for qualitative reasons, i.e., it is impossible for a human surgeon to have the micrometer precision required in brain surgery or for a human astronaut to glide a shuttle in the re-enter path to Earth. Others require robots for quantitative reasons, like checking the video clip uploaded on YouTube; there are 300 hours-worth of video uploaded every single minute on YouTube. At the same time new jobs are being created, like the 10,000 people Google is hiring[175] to assess ethical issues related to content.

The world of manufacturing and distribution (logistics) is under a dramatic upheaval; robots and more broadly, automation, are a crucial component of the transformation, from raw materials (robots are now the only workers in mines, and that is good given the poor working conditions in those places) to retail where products are just a click away. My click from the computer on my desk or from my phone as I ride a bus can activate a robot that is hundreds of miles away that will move to get the product from the warehouse shelf to mail it to me. More and more, my click activates a number of robots that will build the product to fit my requirements and other robots, possibly in different companies that will take care of adding functions, packaging, and mailing it to me. Just wait a few more years and the product might be delivered to my porch by a drone.

Obviously, a changing world requires an evolving understanding from our part, we need to be prepared for something new and different. This is where education steps in, no longer confined to the first part of life but continuing throughout all professional life to stay in synch with the world. This is why the EIT is investing in professional education, and the EIT Digital that fosters the digital transformation is at the leading edge of continuous education through its Professional School[176].

Educational impacts

Regarding education, one might wonder if we have the capability to learn what we need to learn. If you look at results on the average IQ of people in different professions[177] you would discover, that shouldn't come as a surprise, that cleaners have a top IQ that is lower than the top IQ of computer scientists or neurosurgeons. However, and this might surprise you, looking at those results you discover that a good 90% of cleaners have an IQ that is equivalent to the one of computer scientists and neurosurgeons meaning that the difference in capabilities is just the result of a difference in education. Through education we can enable anybody, statistically speaking, to face the ever-complex systems awaiting us in the coming decades.

Countries cannot afford to not invest in autonomous systems and artificial intelligence. But this investment should be flanked by investments in human intelligence through education. It is only through this second form of investment that we can ensure robots (and artificial intelligence) will remain a tool in our hands, augmenting our capability rather than a replacement of our hands and of ourselves.

7.1 Towards a jobless society?

Robots and artificial intelligence are undermining our jobs. True in certain areas, even truer tomorrow and in even more areas. However, the big problem is not a robot replacing me (doing what I do better, faster, with guarantee quality, at lower cost) rather that the digital transformation makes my job useless to the point that no one, human or robot, has to do it. The digital transformation doesn't steal the work from us, something that we might be able to fight. It obliterates jobs.

The digital transformation, if properly executed, increases the overall system efficiency. Unfortunately, we owe our job to the existing inefficiencies. If you remove the inefficiencies, you no longer have that job.

Getting a plane ticket used to involve a travel agency, communicating to the airline company that they had a passenger for that flight, cashing the price of the ticket, and sending part of that to the airline. All that involved several people which implies the existence of several jobs. Today all of the above is done at the bit level initiated by the prospective passenger and resulting in an electronic ticket on the passenger's smartphone with no involvement of people, and with no need of jobs. Air travel booking has been digitally transformed increasing the overall system efficiency and obliterating jobs.

This is the real problem: increasing efficiency decreases job opportunities.

Hence the solution: let's keep the inefficiency! Unfortunately, this does not work. We live in a competitive system in which, even at the country level, we can control only a small part of the overall processes, and we cannot control the flow of economic goods (just because it is not advantageous to control them by imposing commercial barriers because in the end they backfire). Hence, if we stay as we are and don't improve our system's efficiency, we face competition from other places where efficiency is improved. Our competitive edge thins out, we lose jobs because we can no longer sell our products/services, and we also lose our capability to create value. It becomes a lose-lose game.

Someone is theorizing about a trend towards infinite efficiency to which corresponds a jobless society[178]. We don't believe that is the case, at least for the next two decades we will have and we will create inefficiency spaces sustaining our jobs and the ones of the 2 billion people that will increase Earth's population.

How can we find these inefficiency spaces without clashing with the laws of competition? By creating new things. Whatever is new is not efficient; efficiency comes with experience through step by step improvement.

Need for creativity and education

We need to foster creativity and creativity through education, from the way education is pursued. The way matters more than the content. The content will be superseded in just a few years, well before the completion of the professional life of that person.

We need to imbue creativity and the passion for creativity in our youngsters, but creativity is not enough, we need to teach execution and ensure an efficient execution environment (regulation, access to resources, low cost and efficiency of infrastructure).

Reinventing education and creating the right environment cannot be done in a green field. We've got to live with what we have. Hence the real challenge is to foster the transition. As it is clearly shown in the Imperial College Foresight study the future is not going to be bigger, nor faster, nor cheaper. It is going to be different and it is happening in a short time, much shorter than the cognitive revolution (60,000 years), the agricultural revolution (8,000 years), the communications –roads and ships- revolution (1,500 years), the industrial revolution (150 years), the computer /telecommunications revolution (100 years). The convergence of brain science, genomic science, material science, artificial intelligence and Digital Twins is creating the perfect storm to revolutionize our world in the coming 20-30 years.

Are we heading towards a JobLess Society[xi] (JLS)?

[xi] Interesting Report from WEF: Future of Jobs 2018. http://reports.weforum.org/future-of-jobs-2018/

A lot of the interest on the advent of a JLS is based on the perception that we are living in a very peculiar period of human history. A time when our capability to invent smart robots that can do whatever we are doing, faster and possibly better, and most importantly cheaper, is leading us to replace the present manpower with what most analogously be termed *robotpower*.

Actually, history books describe over and over the perception that any time was considered a very special time by those living it. The fundamental reason being that they happened to live in that time. We have seen continuous evolution of our social, economic and productive systems. And at some points in history indeed we have seen dramatic changes. Then we have also seen that through pain and ingenuity the world has adapted and went on.

The shift from the agricultural society to the industrial one dramatically decreased the workforce needed in agriculture; it created new jobs in the manufacturing but also in other areas, like entertainment and travel.

In the 1980's and 1990's, we have seen probably the most dramatic changes in some geographical areas as manufacturing shifted to low labor cost countries. And yet, the Western world ended up better off from this change that was even more impactful since it occurred in a very short time (a decade, versus the 100 years that saw the shift from an agricultural society to an industrial one).

Robotics impact on manufacturing

Robots have been around for some years and have clearly impacted many manufacturing processes. Actually, we are seeing in these years a reverse in offshoring. Whereas until last century it made (economic) sense to move manufacturing where labor cost was lower (due to a supply/distribution chain that made the location of the manufacturing plant irrelevant), now the cost of manufacturing with robots is the same be it in the Far East Asia or in downtown Detroit. Hence, companies are starting to move the production back.

With industry 4.0, dense robotization and increased softwarization and customization of products, manufacturing will become more and more distributed. That will likely create new (types of) jobs as well as decrease (not too much to decrease anymore) pure manufacturing blue collar jobs.

However, and here is the thorn for many of the people in ICT: Robots will learn to program, to design, to write articles, even to invent new areas of mathematics. We are already seeing the first signs of this which creates concern. Now the robots seem to be on a path of substituting us in what are real "human" activities.

In the past agriculture was a real human activity; more recently the capability of using a machine in a manufacturing plant was a real human activity. It is expected that we will continue to see changes, as the world shrinks and more people come to the production world, and as technology enables more and more people to be productive. Think of the millions of apps that the evolution of technology has made possible by enabling a workforce that was totally outside of the production processes. Now a 15-year-old boy can develop, market, distribute and sell the apps he has designed from a dormitory room.

JLS is not expected to occur for the very reason that jobs are a way of life, animals included. Jobs relate to producing wealth and make clear to the society that you have a role in it. Jobs are about responsibility as a citizen, for today's community and for our children. These basic characteristics of jobs will remain unchanged. The actual tasks performed by jobs, on the other hand, will keep changing.

These changes will create pain and suffering in some areas and in some types of people. Out of these hurdles we will see new types of job to emerge.

Overall, the arrow of time has shown a progress in our social life, in the wellbeing of people. In spite of the daily news that might seem to show a different story, of intolerance, violence and destruction, the overall picture keeps getting better. Life expectancy is improving everywhere, particularly in third world countries; hunger and thirst are less severe and widespread. There is more freedom and social communications to overcome countries' boundaries.

There are reasons to be an optimist, an optimism that is not hiding nor denying the many big issues facing us on this small globe, but one that is rooted in past progress and that can leverage new tools, technology being one of them.

7.2 A new definition of the value chain

Our economy has grown and has been transformed by 60 years of Moore's Law, bringing electronics in any path of life. Moore's Law has now reached its endpoint (although progress in computation capabilities are still continuing, not at the previous pace and, most importantly, not with the cost decrease we have experienced in the last decades) but a new law is knocking at the door of our economy and our society, promising to be as disruptive as Moore's: Lass' Law.

Fab Labs

Sherry Lassiter, known as Lass, is the head of the Fab Foundation[179]. A Fab, a Fab Lab, is a room full of computers managing tools that can manufacture objects, including 3D printers and laser cutters. The first Fab Lab[180] was created back in 2003 at MIT Center's for Bits and Atoms by Neil Gershenfeld, and in 2009 he set up the Fab Foundation. There, Lass noticed that the number of "tools" for manufacturing doubled every year and by 2016 there were over 1,000 Fab Labs around the world.

If Lass' Law will remain valid in the next ten years by 2030 there will be over 10 million Fab Labs, and clearly they will no longer be confined within research labs: there are simply not enough research labs around the world to host them. By that time Fab Labs will have percolated into small industries, retail stores and some will have found a place in consumers' homes. By 2040, 150 million homes will have a Fab Lab as part of their furniture.

Transformation of value chain of manufacturing

At the same time, the whole value chain of manufacturing will be transformed bringing Industry 4.0 into Industry 5.0. Consumers will no longer buy products but will purchase products' specifications and then will manufacture them at home. Of course one can imagine that a new slate of businesses will be there, providing support to customization. Possibly some of these business will sell "intelligence" to make customization decisions possible at home.

The convergence of artificial intelligence with massively distributed or on-site manufacturing is going to change the economy, the value chains and the players in manufacturing.

In twenty years, it will be difficult to find a product that is not a mélange of atoms and bits. Many products will be aware of their "use" and will be able to reconfigure themselves to the point of creating their own offspring (due to the availability of fab labs). The mixing of AI with fab labs and the self-generation of offspring will result in product evolution with an increasing symbiotic leverage of what is in the product's ambient.

7.3 From consumption to usage, from ownership to sharing

Our Society has been labelled as "The Consumption Society" and "Consumerism" has become a social and economic order, an ideology that encourages continuous consumption of products. As manufacturing became more effective, the cost of manufacturing growing as a whole but the cost per single product decreasing, it has become essential to sell more and more. Hence the push of marketing on each of us to buy, discard and buy again. Products have been designed not for reparability but for replacement.

The growing concern on the environment, with energy savings and waste disposal taking center stage, and the softwarization of products making software more important in delivering functionalities (thus allowing functionality and features upgrade by software replacement) is changing this social "order".

Softwarization has also promoted the growth of services, and services are not "consumed", they are "used". Companies are responding to this shift by moving from the sale of products to selling services. Some companies, like Apple, are significantly increasing their service income preparing themselves for the time when their products may fade away. Tim Cook has said several times that the time of the demise of the smartphone is on the horizon[181].

Also on the horizon, self-driving cars represent a further instance of the shift from consumption to usage. People will likely enter into usage mode with regards to cars, moving from ownership to sharing. As cars lose their appeal in terms of acceleration, the drive to own one rather than just using whatever is available to go from A to B.

Delivering services as needed

Autonomous Systems will tend to fall into this shift, actually accelerating it, delivering services and being used on a "need" basis, rather than being owned. It can also be expected that in several cases their complexity will keep their price in the high range, and it will make more sense to market them on a usage, rather than on an ownership bases.

This might also be the case for several Symbiotic Autonomous Systems, like the ones implanted in the body, where supervision and periodic maintenance will push towards a usage model with the system components being owned by the producer (or the assembler or the "implanter").

7.4 Multiscale Global Communications

Communication prices have been steadily decreasing over the last 30 years. The advent of wireless communication has created a small uprising spike that has vanished in a short time (between the 1990 and 2005), and the decreasing price has continued approaching zero. In a few countries in Europe it is possible to have an "all you want" both in voice calls and Internet for less than $10.

At the same time the volume and type of communications has increased. In most parts of the world free Wi-Fi is accessible, effectively bringing the price to zero. Youngsters have assimilated this low price in their culture, and most of them are expecting free communications services.

This is just going to get worse (or better depending which side you are) in the coming years as first 5G and then 6G will increase the available bandwidth and as more and more communications infrastructures will become available, most of the time deployed to deliver services (at a cost – direct or indirect) and not to generate revenues from the communications layer.

Autonomous systems need communications which means and most of them will be communications providers, effectively network nodes of a self-constructing communications infrastructure. Actually, in several places the communications infrastructure will be an autonomous system in itself, interacting with other autonomous systems, humans included, in both delivering services and in extending the reach of the infrastructure.

An example will be the ensemble of autonomous vehicles in an urban environment. They will create a mesh network to communicate among themselves with no need for an incumbent infrastructure and most likely will open up such infrastructure to third party use, like the vehicle passengers.

The advent of these autonomous, self-aggregating, infrastructures is supported by the incoming 5G and will take full swing with 6G. These will create a multiscale global communication fabric, consisting of dynamically changing subnets, each one providing access points and able to connect to nearby subnets, much like the flat structure of today's Internet.

Autonomous communications infrastructures

The incumbent infrastructures (in many cases evolving towards basic essential commodity sustained or regulated by governments) will serve as global glue, ensuring end-to-end connectivity.

7.4.1 6+G Intelligent Networks

As marketers are busy extolling the virtues of 5G as the ultimate wireless system filling all of your needs and all of your dreams (and operators are busy deploying and upgrading 4G) a few people are already looking at a new generation aiming at "filling the gaps between 5G promises and reality".

Any *x*G generation takes about 10 years from the inception to the market and then 10 more years to fully consolidate followed by 20 years of "normal" operation. 5G is now approaching the first steps to the market. We can expect to see 5G smartphones in 2020 and wireless dongle already in 2019, so it is about time for researchers to start looking at the next generation. As usual they are starting from some generic needs, and since the hypothetical performances of 5G are such that whatever is needed will be accommodated in the 5G wish list they are looking into how filling the likely gaps between promises and reality.

At the University of Oulu, Finland, a country that is rightly associated with wireless technology, a team of researchers have created an interesting video, called Vision 2030[182], fitting the time window for the first presence in the market of a new G generation. The basic assumption is that artificial intelligence will dominate both in the delivery area, in the core and at the edges of the network(s), in the fruition area, devices like smartphones and things (super IoT), and in the application space.

Augmented reality will become pervasive. It is not clear what technology—or technologies—will support this. Smart materials might allow any surface to display information, or holographic projectors might become available. However, to reach a pervasive ubiquity, we will need to have images created directly in our eyes, using electronic contact lenses, chip implants or brain implants (BCI). All of them are unlikely to be available, in the mass market, in that time frame. The expectation is that electronic contact lenses will be available in the following decade, and much further out for implanted chips and direct brain interaction. There may be trials sooner (very rough prototypes are available already today) but getting to the mass market is a different story.

Some images, like holographic objects floating in space, will only be possible through electronic contact lenses (or chip or BCI). It is also a matter of cost: it will be cheaper to augment humans to become able to receive and visualize bits than augmenting any ambient to display them.

Clearly, assuming that every surface is a screen, it follows that a huge bandwidth is required. However, different architectures may shift the bandwidth burden from the network to the edges, to the ambient, to the devices and eventually to the augmented human and things (unlimited local memory).

It may also be that in 20 years, communications demand will be created by objects, like autonomous systems, both as external communications (towards other autonomous systems) and internal communications (among Symbiotic Autonomous Systems), with human needs already fulfilled by the 4th and 5th G generation.

7.4.2 Universal Real-Time Translation and Holographic Distal Communications

An emerging technology in the early stages of deployment is real-time universal translation of spoken language. One way to accomplish this was made public when which on September 2, 2016 Google announced the Google Neural Machine Translation (GNMT) system[183], a significant improvement to Google Translate launched a decade earlier. Incorporating robust training techniques, previously-introduced recurrent neural network (RNN)[xii], and other advances led to a surprise: the appearance that the system learned "a common representation in which sentences with the same meaning are represented in similar ways regardless of language"—in short, an AI-generated interlingua. The system was then shown to be able to translate "a set of sentences between all possible pairs of the Japanese, Korean, and English languages."[184]

Moreover, integrating AGI and AC would provide the translation system to understand human speech and emotion, respectively, in a manner consonant with our own. Extrapolating the integration of the above technologies suggests the possibility of a Universal Real-Time Translation via an AGI/AC-generated interlingua[xiii], or Rosetta Code, which—rather than translating directly between the languages being spoken—uses a core artificial language conceptually based on the ancient Rosetta Stone[185] to translate between languages. This development will bring a powerful level of functionality to human/Digital Twin symbioses (see Section 4.6) by allowing seamless real-time natural speech/digital communications by a range of interfaces, including but not limited to telephony, computer interfaces, and semantic speech/digital/text conversational systems.

Furthermore, as Rosetta Code-based technology becomes normative, it will enhance our bioself/Digital Twin identity and expand into other domains and applications, including:

[xii] A **recurrent neural network** can use its internal state (memory) to take previous output or hidden states as inputs, and so store information about past inputs for a time that is not fixed prior to an event. RNNs can therefore address, for example, speech recognition and unsegmented, connected handwriting recognition.

[xiii] An interlingua is an artificial language that facilitates translation (in this case, machine translation) into a target language. The Google AI independently generated an interlingua to translate material for which definitions and grammar specifications were not present.

- As biology and technology continue to merge, bioelectronic communications will ultimately become an endogenous natural function, thereby obsolescing today's exogenous voice and data systems—and reducing the structural differences between our bioselves and our Digital Twins.
- Ubiquitous AGI/AC global mesh networks, in which Rosetta Code-enhanced symbiotic human/Digital Twin routing will provide real-time seamless language-agnostic voice and data communications in which optimal routing will be a function of semantic content fused with telecommunication protocols rather than the latter alone. Given increasingly AI-augmented accurate speech recognition and natural speech generation, we may well be having conversations with our Digital Twins sooner than we expect.
- Realtime distal holographic communications—long envisioned but existing only through the magic of sci-fi special effects— have recently taken a significant step towards realization. Researchers have demonstrated a free-space volumetric display[186] based on photophoretic optical trapping[187] in which laser manipulation of particles were used to generate 3D free-space motion holograms that will lead to the realistic experience of local presence of all participants—some of whom will be Digital Twins made visible in real-time.

7.5 Intelligent Transportation

7.5.1 Hyperloop - Vacuum-tube transport

Vactrain, sometimes referred to as the fifth transportation means, is a transportation system based on pipes in which modular vehicles can move at very high speed due to a frictionless environment since the pipes themselves have no air—a vacuum—and the vehicles don't touch the pipe by floating in a magnetic field.

In theory, one only needs power to accelerate the vehicle at the desired speed and then it will keep moving encountering no resistance. This would make for a very cheap transportation means. Unfortunately, the cost of the infrastructure is exceedingly high with present day technologies.

Hyperloop[188], being tested in the US and France in a scaled down version, and targeted to provide services connecting Abu Dhabi and Dubai in the next decade is an example.

A Chinese company, China Aerospace Science and Industry Corporation, is also at work designing a vactrain to run at commercial speed of 2,500 km per hour, after having proposed a concept for a flying-train at 4,000 km per hour that was received with some skepticism[189]. That would make connection between Europe and China possible in 3 hours.

This is what the foresight team at the Imperial College envisage for the future of transportation, a worldwide vactrain service connecting the globe beyond 2040. There are several technological issues to be solved, like not interfering with the magnetic field at the various inlets and outlets (the point where the vehicles ae joining and leaving the infrastructure) and guaranteeing the integrity of the infrastructure, taking into account natural and exceptional situations (just think, as an example, the earthquake risk in California, rendering the planned San Francisco to Los Angeles route very challenging).

The third trial plant for Hyperloop under construction is in France. A first 320m closed loop will be available in 2019 and a longer one 1 km long, will be ready in 2019.[190]

However, the most difficult issues are related to the cost of the infrastructure and therefore its affordability. The cost for the magnet required to ensure the levitation of the vehicle are astronomical considering the thousands of kilometers of pipes, and the power required to create the magnetic field along the pipes is staggering.

The likely solution for this latter should come from superconductive materials, but again their cost of operation today is way too high (these materials have to operate at very low temperatures, and maintaining these temperature is very costly).

We are likely to see a few vactrains in operation in the second part of the next decade, like in AUE and China, but their usage will be constrained by the high cost. They will more likely be trials and evidence of what could become possible in the future.

It is obvious that a pervasive availability of vactrains would revolutionize the world of transportation and along with it our idea of the world. Commuting time between San Francisco and Los Angeles will be less than half an hour, likewise between Frankfurt and London. And if you need to span longer distances you can take a faster vactrain, potentially travelling at 8,000 km per hour, getting from NY to Los Angeles in less than an hour.

7.5.2 Hyperjets/Scramjets

Supersonic Combustion Ramjets[191] is a technology, already proved by NASA[192], making use of the air compressed by a fast moving aircraft to generate thrust. Engines using air compressed by the movement of the aircraft are called "ramjets". The problem with ramjets is that the compressed air needs to be slowed down, to subsonic speed, in order to be usable in the engine. This slowing down creates waves that basically constrain the maximum speed of the incoming air to somewhere around 5,000 km. Above that speed it is no longer possible to slow down the air in the engine to subsonic speed and thrust is no longer created.

Using a different type of engines makes it possible to use supersonic speed in the combustion camera of the engine (hence the name "supersonic combustion").

These engines need air to work, so they are usable only in the atmosphere (not like a rocket that does not need air—it needs fuel and oxygen and it has to carry it all along). Potentially they could support very high speed in the upper part of the atmosphere, a region that has little air and so little drag, thus resulting in more efficient flights.

Technology has still a long way to evolve to the point of making scramjets viable for a broad use, so targeting 2040 seems reasonable.

7.6 Global Non-Polluting Net-Positive Energy Technologies

The advanced, emerging and projected energy-focused technologies discussed in this section—while not necessarily intended to align with a Digital Twin ecosystem—are likely to at some point have a Digital Twin that captures the technological, input/output, transformational, and other utilitarian data mirrored in a secure blockchain environment. Adopting this structure will allow a circular economy operation to demonstrate whenever required its procedures, regulatory compliance, inventory, efficiency, and other transactional, structural and operational data by having a Digital Twin perpetually mirroring these variables on an ongoing, permanent basis.

7.6.1 Circular Economies

The World Economic Forum defines a circular economy as an industrial system that is restorative or regenerative by intention and design. It replaces the end-of-life concept with restoration, shifts towards the use of renewable energy, eliminates the use of toxic chemicals (which impair reuse and return to the biosphere), and aims for the elimination of waste through the superior design of materials, products, systems and business models.[193] In a manufacturing plant based on a circular economy protocol, for example, waste materials in linear (standard) plants are replaced by the output of two classes of reusable outputs in which material flows are of two types—referred to as *nutrients*—these being biological nutrients (designed to reenter the biosphere safely) and technical nutrients (which are designed to circulate at high quality in the production system without entering the biosphere, as well as being restorative and regenerative by design). In short, as a result of these practices an industrial circular economy produces no waste or pollution.

That said, however, while a circular economy is most frequently described as a combination of reduce, reuse, recycle, and recover activities, there is a wide range of overlapping definitions: A recent paper identifying and analyzing 114 circular economy definitions across 17 dimensions concluded that the *restore* factor has come into use to a lesser degree (see Fig. 7.6.1); the *reduce* factor is frequently neglected, perhaps due a belief that reducing production might curb consumption and economic growth; and confirmed previous research that the circular economy/sustainable development link is weak.[194]

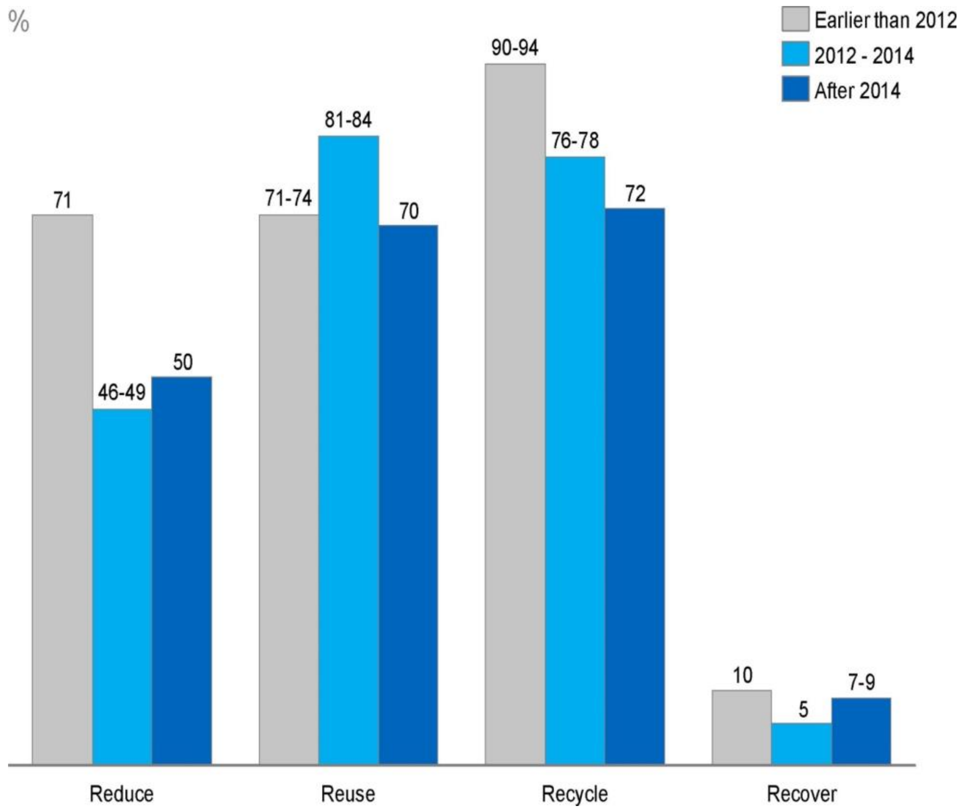

Fig.7.6.1. Development of Circular Economy Definitions over Time[195]

7.6.2 Artificial Photosynthesis

Artificial photosynthesis[196]—a biomimetic (that is, mimicking biology) chemical process that that replicates the natural process of photosynthesis by converting sunlight, water, and carbon dioxide into carbohydrates and oxygen—generally refers to any system that captures and stores energy from sunlight in the chemical bonds of the resulting solar fuel. Related technologies involve engineering photoautotrophic microorganisms and enzymes to generate microbial biofuel and sunlight-based biohydrogen production and converting CO_2 directly from air into biomass and fuels. Another example is a recent hybrid water splitting–biosynthetic system[197] that when combined with solar photovoltaic cells promises solar-to-chemical conversion rates roughly a 10-fold increase in efficiency compared with natural photosynthesis, and moreover avoids toxicity associated with previous attempts.

Cost-effective artificial photosynthesis technologies well-suited to housing installations in urban and densely-populated suburban areas are inkjet-printable solar panels[198], artificial leaves[199] and (even for woven polyester cotton fabrics[200]) spray-on solar cells[201]—an important focus given the interaction between continued population growth[202], increasing urbanization[203,204] and rising energy demand[205].

7.6.3 Thermionic Energy Conversion

Thermionic energy conversion (TEC) is the direct transformation of thermal to electrical energy—specifically, from thermions (heat quanta) to electrons—by thermionic emission (hot electrons spontaneously ejected from a surface). While TEC is currently used in solar cells to increase conversion efficiency, it has the potential to, for example, convert the heat of an in-use battery to be converted to electricity. While no researcher would assert that TEC is a self-perpetuating system, a limited charger-independent system can be envisioned with solar charging built into the display, as was the Kyocera Torque concept smartphone shown at Mobile World Congress 2015, with the French firm SunPartner providing the phone's built-in WYSIPS[xiv] Crystal solar panel—a 0.55mm transparent pane placed between the phone's display and touchscreen. That said, a system approaching a fully closed-loop system might be feasible by equipping a smartphone with both high-conversion ratio TEC and WYSIPS components.

On a larger, more ambitious scale, researchers are assessing the potential of TEC systems (also referred to as Thermionic Converters) for both space and terrestrial applications.[206]

[xiv] What You See Is Photovoltaic Surface

8 Impact on Education 2050

The shift towards an intelligent ambient and the emergence of technologies that support a seamless symbiosis among people and their environment, including the symbioses of single individuals with cyberspace is going to have a profound impact on education. The first signs of change, and of the new potential, are already in sight. We have come to rely more and more on cyberspace to access data, information, knowledge and "skills". We have now thousands of courses in any discipline that can be taken by connecting to the Internet. In a way, even more important and telling in terms of trends, we have access to hundreds of thousands of short education clips explaining "how to …", from simple daily chores to complex labs experiment.

The real value today is not in the content (that is abundant and free most of the time), rather in the tools guiding to the right education content. By 2050 the needs and the way to satisfy them will differ significantly from today and that will give a quite different meaning to education.

8.1 Education Needs

The amount of knowledge (and even more the amount of data and information) is growing exponentially, and there is no indication that this growth will slow down in the coming years, quite the opposite. The reason is threefold:

- more data can be acquired through a variety of sensors;
- more processing on those data creates information and possibly insight that may require acquisition of more data; and
- communications make data and information available everywhere with a multiplying effect.

At the same time, we are seeing a shift from value creation through atoms to a value creation fueled by bit, hence mastering bits, i.e., knowledge, is even more important than in the past where routine would suffice in most activities. The problem is not just the sheer volume of data, information, and knowledge but its continuous flux with new data, new information and new knowledge waves hitting at a faster and faster pace.

The half-life of knowledge is shrinking, being shorter than four years in many areas. Competitiveness and the value of knowledge requires a continuous education process in a growing number of professions, and this trend is gaining strength as time goes by.

It is the opinion of this White Paper that by 2050 (sooner actually) it will be impossible for an individual, with today's education protocols, to remain up-to-date and valuable in the job market. Notice that this problem is compound by the expected lengthening of the working life that is today around 35-40 years and that by 2050 may extend to 50+ years. This issues goes beyond the single individual, affecting companies and institutions as well, although with different implications as will be explained in the following.

Because of these two phenomena, expanding knowledge and shorter knowledge life time–both exceeding the capability of any individual with the current education protocols–, there is a shift towards a two stage education process.

8.2 Basic Education and Just-in-Time Education

The need to acquire a basic education, i.e. the capability to learn, including learning to communicate verbally and through reading and writing, may fade away, possibly not in this timeframe but in the long run. The capability to understand logical structures and mathematics and the capability to operate and interact effectively in a community will remain basically unchanged although the tools to achieve this basic knowledge may change by adapting and leveraging the new technology environment. For example, AR/VR will likely dominate learning by the middle of this century.

This basic education will probably need a refresh once the working life exceeds 40 years, with several researchers pointing to the possibility of having a few break periods to dedicate to basic education throughout the professional life, e.g., after 15 to 20 years of working activity have a 1-2 year break to go back to school (so to say, since the school might be quite different from today).

An aspect that will become crucial in the coming decade is the awareness of a personal knowledge gap. This can be related both to a gap in basic education (like the need to acquire new, more effective tools to access and manage knowledge) and to the understanding of what specific knowledge is needed for a task at hand.

Digital Twins to manage knowledge

Digital Twins will become essential to manage the knowledge gap. Every person will have a knowledge Digital Twin that on the one hand will mirror the knowledge of that person (the physical twin) and on the other hand will determine gaps in basic education (by monitoring the evolution of basic education and comparing to the one it has) and in the task at hand (by looking at what would be the best knowledge required to carry out the task and comparing it to the actual knowledge available to the real twin). This basic education (possibly repeated during one's lifetime) will need to be complemented by a just-in-time education that for most part will be based on the "need to know." Hence, it will be very focused and very timely.

A growing portion of this "need to know" will be related to learn how to access knowledge and make it **operational**, not to acquire knowledge, but to enhance knowledge. Knowledge itself might reside in other people or, ever more frequently, in machines (in cyberspace). We will have to move from today's paradigm of acquiring knowledge (made easier by tools that connect us to the knowledge in cyberspace, rather than using books as we did in the past) to a paradigm where accessing knowledge will be important to have the owner of the knowledge doing what is needed, rather than learning that knowledge to be the one doing what is needed. By doing so, we could enhance our knowledge in ways not accessible and available in the past.

Tailored education

Projects will be carried out by applying distributed knowledge in a much more effective way than what is happening today where knowledge transfer is the stumbling block. Notice the difference between today's "continuing education" from tomorrow's just-in-time education. Continuing education as it is conceived today is more in line with continuing basic education, although we are starting to see an approach to continuing education that is tailored to both the individual knowledge and to the education requirements to execute a specific task. We are starting to see continuing education courses created dynamically out of a vast module portfolio, although this is just a tiny step if we look at what will be the needed in 20 years.

An additional, but critical aspect, is that outside knowledge is no longer a repository (a book or a data base). Knowledge is more and more operational, exactly as is human knowledge. Our brain is not a repository of knowledge. Whenever we ask a person for some of her knowledge, including when we ask ourselves, we do not get a book, what we get is operational knowledge adapted to the specific circumstance. This is starting to occur when knowledge is accessed through a smart agent (making use of artificial intelligence). This is a major shift; it is the first time in the humankind history that we can tap onto operational knowledge.

8.3 Symbiotic Shared Education

The shift from owning knowledge to owning the capability to access knowledge and engage the knowledge owner to leverage it for the task at hand is taking us into the era of **symbiotic shared education**. Assuming that it will be possible to engage the required knowledge, the education protocol will aim at increasing knowledge where it is most effective, and in most cases that will mean to educate machines rather than humans. This might seem unbelievable today but we already have working examples of machine learning and machine learning in an autonomous way. Their learning is faster than the one of any human being, and can be continuously updated, while being retained since machines never forget.

Notice that this shift has started in the last century with the pocket calculator. How many people today would know how to find the square root of a number? From that very tiny beginning we now have machines that can learn to identify a face (used by many police departments all over the world), machines that can estimate the evolution of the stock market and make buying or selling decisions, and machines that can learn the best path to clean a house. The examples are many and growing fast. Robots have shifted from pure executors in the assembly line to be co-workers; software is learning, continuously, the mood and preference of people to create music that meets the fancy of the moment.

In healthcare, we are on the threshold of a momentous change with software associating specific genes to diseases and to the most effective cure, with protocols that are continuously monitored and finely tuned. In spite of all this progress we have just started and we are entering into a new education paradigm.

The Digital Twins, previously mentioned, go beyond assessing an education gap. They can guide in the delivery of education, and more than that they are becoming artificial knowledge carriers. A person gaining experience (and knowledge) will implicitly grow her Digital Twin experience and knowledge. Unlike a person that is limited to one instance (a person can be only in one place and do one task at a time) a Digital Twin can be instantiated an unlimited number of times (and we are starting to see Digital Twins consultants serving several clients at the same time). In addition, Digital Twin never gets sick, never takes vacation and may never die or cease to exist.

The ownership of a Digital Twin is becoming a matter of discussion both under the legal and the ethical point of view. A company claim rights, today, to the knowledge a person is acquiring on the job (such as through non-disclosure agreements or no-competition pacts). It might be expected they will claim rights on the Digital Twin as well, unless the Digital Twins would select to become independent. Would those companies retain the rights as the employee moves to another company? Would the Digital Twin be transferred to the other company (possibly enforcing a blockchain onto knowledge usage)?

What kind of accountability will reside with the Digital Twin, with the real twin and with the user of the Digital Twin? A company exists due to its competitive implicit knowledge. This can be derived from the creation of a symbiotic Digital Twin (that at certain stages can embed real twin instances). Digital twins can indeed consist of the aggregation of several Digital Twins as discussed in Section 4.6.

A Symbiotic Autonomous System has knowledge and skills that are the emergence of the individual components, through their mutual interactions. Education in 2050 will need to address the Symbiotic Autonomous System as a whole and decide where, when and how to educate the individual components. It is a whole new world that we are just starting to visualize.

An extensive discussion on Education 2050 is given in Appendix A.

9 IEEE Societies Impact

Symbiotic Autonomous Systems are pervasive in terms of technology and applications. Many technologies, as described in this White Paper and the first White Paper, are needed, such as signal processing, computation, communications, bio-engineering. The evolution of SAS is strongly dependent on the evolution of these technologies. Likewise, there is a growing field of application of SAS, including robotics, consumer electronics, health care, transportation, education and more.

Most IEEE Societies are actively working in both the evolution of technologies and in their applications.

In the following section, several statements from representatives of several IEEE Societies outline their views both on the impact that SAS are/will be having on their activities and on how their activities foster the evolution of SAS. This list is not exhaustive and will be continually updated, and the IEEE SAS Initiative welcomes participation from all societies with a tie to the technological innovations described in this White Paper (see symbioitic-autonomous-systems.ieee.org for more information).

9.1 SAS Impact on the IEEE Consumer Electronics Society

Today's consumer electronics are built upon hardware and software platforms that include both local as well as remote operations. These remote operations are often facilitated by applications running in large data centers (the cloud). In the not too distant future it will be common for consumers to own several smart connected devices located on their person, close to them or spread across their home and vehicles.

These smart connected devices will deliver a new generation of products and services that support and entertain consumers and their families. Systems capable of evolving in their understanding of the local situations as well as their behavior in response to that environment will be extremely valuable to create safe family experiences, enhance teaching and training and for new modes of entertainment.

Symbiotic Autonomous Systems (SAS) would be a way to have these connected smart things coordinate and share information and to enable distributed and responsive system-wide decision making that would support and entertain consumers. The IEEE Consumer Electronics Society will embrace these capabilities in its Future Directions activities. Among the initiatives that the Society could engage to promote and develop SAS in consumer electronics are new standards, conference sessions and special editions of the Society magazine (Consumer Electronics Magazine) that focus on the role of self-organizing autonomous agents running in consumer applications.

These activities could raise interest by non-members in the CE Society, bringing additional people to our events, more people reading our publications and thus increase our membership.

Such agents and their operations could enhance safety at home and while travelling, act as surrogates for consumers for various day to day activities (e.g., representative avatars making travel arrangements for consumers), help older people remain independent, and entertain and educate younger as well as older people. In addition, agents could act with other consumer agents with common interests in some coordinated way to pool resources and act to enable those interests. Methods such as distributed ledgers could be used to authenticate such SAS agents, even while keeping the represented individuals anonymous.

Tom Coughlin
IEEE Consumer Electronics Society Future Directions Chairman

9.2 SAS Impact on the IEEE Systems, Man, and Cybernetics Society

The notion of Symbiotic Autonomous Systems (SAS) is very synergistic with the field of interest and activity of the IEEE Systems, Man, and Cybernetics Society (SMCS). Our Society promotes the theory, practice, and interdisciplinary aspects of systems science and engineering, human-machine systems, and cybernetics. Central to this is the human and human relationships to machines in the context of engineered systems. With a highly comparable mindset focused on the convergence of human augmentation with artifacts having increased intelligence and awareness, the SAS Initiative is squarely aligned with the SMCS, particularly as it fosters studies and applications leading toward a symbiosis of humans and machines.

In executing its core mission, the SMCS has been fostering the evolution of SAS for decades through its conferences, publications, and other activities that contribute to the professional needs of its members. Many of the Society's currently active Technical Committees (TCs) focus on elemental aspects of SAS, representing how the SMCS is actively working on both the evolution of SAS technologies and on their applications. The relationship to the broad spectrum of technologies covered by the SAS Initiative -- machine and human augmentation as well as symbiosis -- is evident in the names of several SMCS TCs. Included among them are Awareness Computing, Brain-Machine Interface Systems, Brain-inspired Cognitive Systems, Evolving Intelligent Systems, Intelligent Internet Systems, Biometrics and Applications, Companion Technology, Interactive & Wearable Computing and Devices, Shared Control, Bio-mechatronics and Bio-robotics Systems, Cyber-Physical Cloud Systems, and System of Systems.

It is expected that continuation of this tight alignment will foster strong SMCS collaboration with the SAS Initiative in various forms now and with the products of the Initiative in its aftermath. To date, mutually beneficial collaboration is facilitated by active participation of representatives from the SMCS on the Initiative's steering committee. In addition, the SMCS is strongly supporting the Initiative's current efforts to foster consensus on how to bring about autonomous systems symbiosis as well as its aim to facilitate development of a new field of Symbiotic Systems Science. Recent evidence in this regard is the hosting of the Initiative's forum on "Symbiotic Autonomous Systems: Fostering technology, ethics, public policy & societal enablers" at the 2018 IEEE International Conference on Systems, Man, and Cybernetics on 9 October in Miyazaki, Japan. The forum discussed implementations and implications of symbiotic systems and was well attended by SMCS members and conference delegates, a testament to the high level of interest to be expected at an SMCS conference venue.

This is just the beginning of a collaboration of clearly mutually beneficial prospects. The IEEE SMC Society is excited about and looks forward to working further with the SAS Initiative.

Edward Tunstel, Ph.D., FIEEE
President, IEEE SMC Society, 2018-2019

9.3 SAS Impact on the IEEE Communications Society

he whole area of Communications is evolving, because of technology evolution sure but even more because of the variety of communicating entities and services being offered and accessed by hundreds of thousands of players, soon to become millions.

The connectivity layer is going to benefit from advances in several basic technologies and we are now seeing 5G hitting the market as researchers are already looking at what's next (6G and beyond).

The work carried out in the Symbiotic Autonomous Systems initiative is very much of interest to the Communications Society, both in the short term and in the long term.

In the short term the investigation on the needs and ways of communications among autonomous systems and the emergence of a communications fabric out of the interplay of several communicating systems fits well in the application domain of 5G, encompassing a variety of network technologies, thus supporting from biosensors communications to autonomous vehicle communications.

In the longer term the possibility of creating mesh networks by a multitude of autonomous systems on the one hand and on the other hand the study of communications at semantic level including brain to brain mediated communications and brain to cyberspace communications are at the frontier of science, an area that will become more and more important for the Communications Societies as artificial intelligence permeates networks, applications and communicating devices.

Khaled B. Letaief
President IEEE Communications Society

9.4 SAS impact on Computer Society

Symbiotic Autonomous System (SAS) are promising to change how we think about both computers and humans, bringing the two closer together towards increasingly symbiotic systems. This opens up tremendous opportunities in improving humanity, helping humans in their life and profession. But it also poses a potentially tremendous threat to humans if the technology is misused either by criminals or by rogue governments or both.

The benefits include augmenting humans, analyzing data at the edge using more powerful techniques, grouping humans and machines into more effective teams, etc. These functionalities offer some of the solutions to mankind's largest challenges, such as world hunger, global warming, cybersecurity threats and many others, through monitoring, raising awareness of situations at the edge over the groups of humans and machines, acting upon anomalies, and better utilizing resources.

The challenges to achieving SAS vision include regulations which may transpire governments around the world; the need for new standards across geographies as well as across humans and AI; introducing new norms of behavior; collecting new information on the successes and failures (which will inevitably happen) to improve on the AI and systems running them, as well as adoption by humans[207].

SAS will heavily depend on computer systems, networking, and Human Computer Interfaces which are directly related to the field of interest of IEEE Computer Society. In addition, many non-functional characteristics, such as Security, Performance, Scalability, Reliability, Availability, etc. are also traditionally studied by computer scientists and offered in IEEE CS products and services.

IEEE Computer Society has a number of technical activities related to SAS in its portfolio. To name a few that are directly related: IEEE International Conference on Autonomic Computing has been exploring areas of autonomic systems for more than a decade; Pattern Analysis and Machine Intelligence has conferences (e.g. International Conference on Computer Vision and Pattern Recognition) and a Transactions on PAMI. However, many other technologies such as Big data, Cloud computing, IoT, are closely tied to SAS area, as are fundamental areas, such as computer architecture, software, AI, etc.

SAS will open up additional opportunities for IEEE Computer Society to embrace new technologies and create new products and services to support SAS. It will be required to engage academia, governments and industry to fully support the needs of SAS, and Computer Society will work with other Societies of IEEE to achieve this vision.

Dejan Milojicic
IEEE Division 8 Director (from Computer Society)
IEEE Computer Society 2014 President

10 Roadmap and Conclusion

This second White Paper has addressed a quite long (in technology terms) timespan. This might give the impression to a casual observer that Symbiotic Autonomous Systems is something further out in time with very little relevance today, in a world that is usually focusing on the next quarter, particularly the world of business.

This is not so.

There are three main areas emerging from this White Paper that are immediately relevant and on which action is needed now:

1. Smart prosthetics shifting from disability recovery to human augmentation
2. Need to rethink education leveraging the expanded, distributed, and intelligence of knowledge
3. Ethical questions needing an answer to steer next decade evolution

10.1 Smart Prosthetics

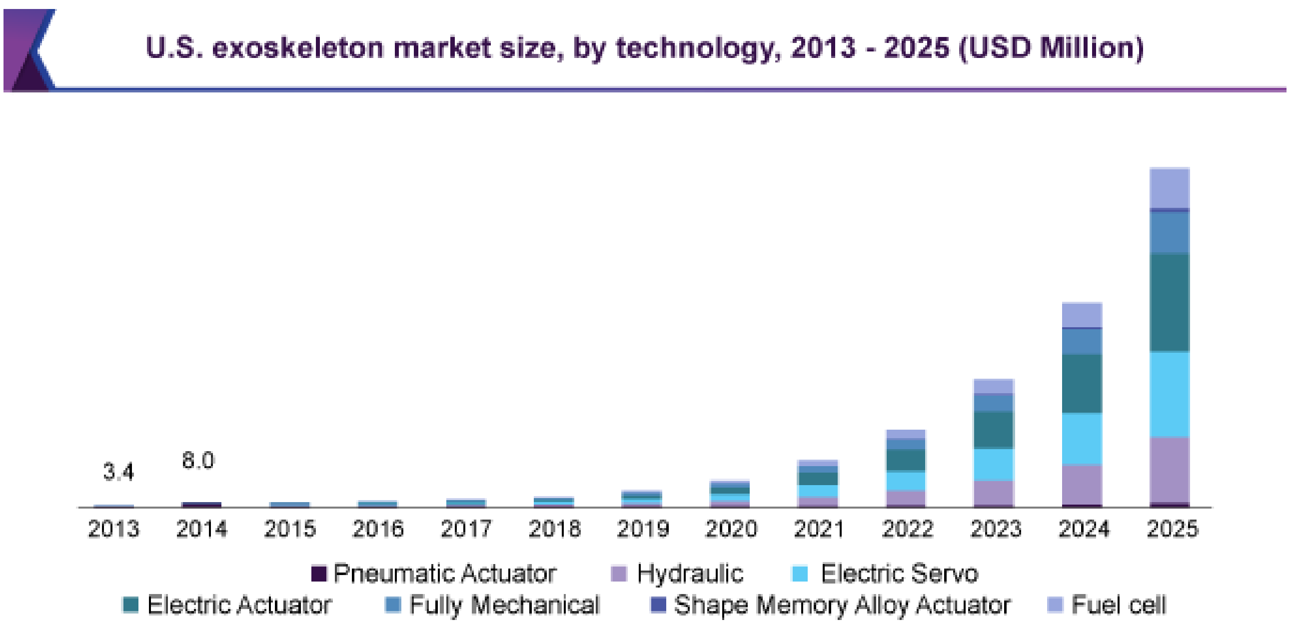

Figure 10.1. Exoskeleton market forecast: \$24.2M value in 2015 in the US, expected to grow \$1.2B in 2025[208]

Progress in sensing and seamlessly connecting prosthetics to the human body is making artificial hands, limbs, eyes, hearing, and more generally exoskeletons, much more effective and natural. At the same time prosthetics can offer augmentation, expanding our senses, strength, and resilience. This augmentation is already a reality in a number of areas, from robotic assisted surgery to exoskeleton support in the assembly line and healthcare.

A different, but similar area of prosthetic augmentation is the one seeing robots cooperating with humans in a symbiotic way. The growth of robotic intelligence, also in the areas of cooperation with humans, detecting human emotion and adapting to them will change the workplace in the coming 5 years, and it has already started. This means a rethinking of the workplace (from hotels to hospitals, from the cockpit to the assembly line) and along with it a rethinking of the education protocol and goals.

10.2 Rethinking Education

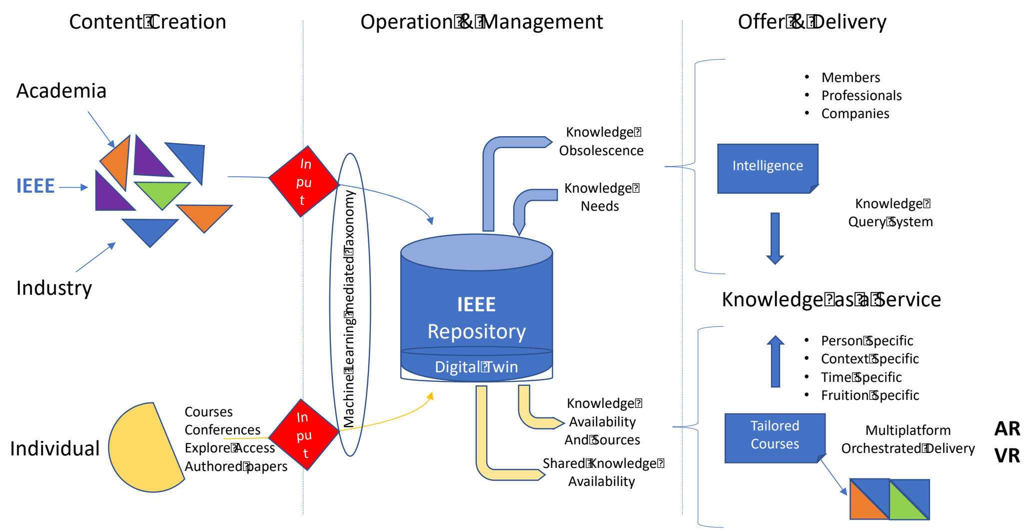

Figure 10.2. IEEE can play a pivotal role in the fruition of knowledge. This requires a restructuring of its knowledge base, an expansion to include users' Digital Twins and the adoption of AI tools.

The amount of knowledge is expanding beyond the capability of a single person. Discrete units of education, right sized and right on time, taking into account the operational context, i.e., what is the overall knowledge available in that particular environment at that particular time and the actual knowledge/skill of the person, are the way to go.

The White Paper points out the role of symbiosis in education, symbiosis between the person and cyberspace (actualized in cooperating robots, applications) since it is impossible, timewise, economically, and from a capacity perspective, to keep the education of a person current. We need to accept and consequently leverage that education, as knowledge, has to be distributed and make sure that it is available when it is needed and in the appropriate form.

The collaboration between human and machines (the term machines includes software applications) is critical to ensure proper education and knowledge sharing. Education shall look at the overall picture and be directed to educate all components in the picture, each one leveraging its own capability and expected role. Digital Twins are the bridge in cyberspace to connect the silos of knowledge and make them a usable whole. They are also the representation of the fleeting knowledge state of individual persons.

10.3 Ethical Questions

The table below indicates some the ethical considerations facing Symbiotic Autonomous Systems.

	Today to next decade	Long term (2035 – 2050)
Symbiotic Autonomous Systems as Objects	• Social ethics (e.g., unemployment) • Machine ethics: algorithm ethics • Fairness in algorithms	• Uploading values from humans to autonomous systems • Agents dominating • Dealing with AGI and ASI
Symbiotic Autonomous Systems as Beings	• Legal status of AS • Controlling learning • Sharing responsibilities	• Legal status of SAS • Consciousness • Mind uploads – rights, accountability

Table 10.1. Ethical questions raised by Symbiotic Autonomous Systems.

The White Paper has identified a number of ethical questions that arise from technology evolution. It is crucial to address these questions, in the appropriate context now since their answers can influence the way technology evolves.

Ethical questions, and more importantly the answers to those questions, depend on the cultural framework in which they are posed, and this cultural framework evolves over time. While we cannot expect to provide or get definitive answers, formulating the questions and socializing them is of crucial importance today—and tomorrow.

We cannot wait and see if we are serious about "fostering technology to the benefit of humanity."

Not only does the adoption of new technologies raise ethical questions, some of them, such as the appearance of the printing press or now the digitization of human culture, demand a radically new ethical order. The Renaissance witnessed a brutal redefinition of what was meant to be human during the painful transition between a predominantly oral and communal religious authority to an individualistic humanist social and political order. While transiting from "shame" to "guilt", the object of personal responsibility in western society shifted from "the other" to the "self". Today, as people are ever more exposed to continuous monitoring by automated electronic systems, and while, in some countries, behavior itself is controlled by algorithms, responsibility shifts away from the self to the now almost self-organizing whole social order, including, and perhaps eventually prioritizing the care of the environment. Shouldn't Symbiotic Autonomous Systems always be developed with their aftereffects on people and the environment in mind? Shouldn't large scale predictive analytics be applied to each innovation before implementation? Even as algorithms and AI take the lead in introducing a generalized symbiosis between individuals and the environment itself, shouldn't an ethical dimension be consciously included in their programming?

Among the most urgent questions is *how SAS needs to take account of an epistemological change that is going on right now because of technology*. People are being emptied of their psychological content and strategies without most noticing it because they labor under the illusion that apart from additions to our technical capabilities, human nature—their own in particular—is the same as it always was. "Human nature" changed before, and it is changing again, so much so that the unrecognized transition from humans developing from within to projecting (or rather letting go of) one's identity online externally could bring one to doubt that there is such thing as human nature at all. For example, is it not conceivable that the evolution of our "Digital Twins" will eventually become the next modality whereby people negotiate their relationships with SAS and the total environment, as opposed to what (Westerners at least) still consider the activity and the exclusive property of internal memory, intelligence and judgment? Assuming symbiosis is completed within the next thirty years, will humans still benefit from any capacity to resist intellectually, let alone politically or even emotionally?

Although genetically modifying the human genome to increase its intelligence on a par with AGI and ASI has been suggested as one of the ways to offset the complete takeover of human decision-making by AI and algorithms, the predictable future is that humans will become increasingly vulnerable to SAS domination without allowing enough time to evolve that sort of recourse. We cannot stop the exponential growth and sophistication of our machines, but we can decide to infuse them *algorithmethically* with genuine human survival and environmental priorities and monitor and control worldwide how they are applied. An ethical taskforce of SAS projections would want to address the particulars of that issue.

For many ethical considerations about how people relate to their robots or how they use them for better or for worse, standard legal provisions applying in conventional circumstances provide guidance and protection. However, among other issues, there is no provision for how we treat our machines. Chances are, for example, that many people will continue to consider household robots, not as life-like companions, but merely as machines to mistreat or dispose of the way they do with lesser contraptions. The danger with that (even with robots that don't look like humans, but have, like Alexa, for instance, increasing relationship capacities) is a slippage of categories where people once accustomed to rough up their robots, will also adopt brutal attitudes with real companions. Is there a need or a way to legislate proper conduct with machines?

Regarding the potential social injustice occurring between access to human augmentation by technological or biological means, one could invoke regulation obtaining in sport competition—but it isn't quite the same thing to prevent doping in a cycling event and to impede school and career privileges to enhanced humans. How does one regulate better for social justice in income and means disparities in a livable SAS environment? And, at the same time, regulation shouldn't stand in the way of the full realization of enhanced human potential. Observing the rapid proliferation of BCI and slower CBI technologies, it appears that connecting brains to the Internet and eventually accessing directly and pertinently its enormous contents is almost a foregone conclusion (if not seamlessly at first), perhaps via tweaking our Digital Twin. Wouldn't then there be a need to prevent intentional or accidental harm through thoughtless activity by reprogramming the Internet itself?

Finally, SAS predicts the gradual evolving of superorganisms consisting of the task-oriented association of indefinite numbers of components, both human and algorithmic. This prediction is congruent with observations above regarding how humans may become increasingly defenseless about how SAS takes over all important decision making. Should such a possibility be left to teleological self-organization or should it become the object of an international political decision-making?

10.4 Conclusions

The investigation of the evolution of technology summarized in this White Paper has pointed out:

1) Systemic evolution will accelerate as a consequence of the convergence and interplay of independent technology progress. More and more technologies provide options to applications and provide a springboard to accelerate further technology progress.
2) IEEE Societies are each fostering specific technology evolution. There is the need to take a holistic view of all these evolutions and of their impact on business and society. The IEEE Symbiotic Autonomous Systems Initiative is a step in this direction, encompassing many technologies and impacting many fields of applications, but it is not exhaustive by all means. Hence this White Paper is also a call to action to all IEEE Societies to take a broader view on the implications of their effort.
3) Symbiotic Autonomous Systems are already here. Clearly today's instances are limited and in only a few niches. We can expect an avalanche effect in the next 30 years, with steady growth in applications and application domains in the next and following decade, sweeping the market in the fourth decade of this century.
4) A few focused initiatives, such as in education evolution through digital twins, can be launched now.

11 Glossary

Artificial consciousness
Artificial consciousness is a consciousness created and experienced through artificial means. It is associated with machines (computers with artificial intelligence). Awareness is considered to be an essential component of consciousness but it is not sufficient to create consciousness.
The possibility of creating artificial consciousness is open to question at this time although discussion on the ethical implications arising from it are already being studied.

Artificial General Intelligence
Artificial general intelligence (AGI) is the intelligence of a machine that could successfully perform any intellectual task that a human being can. Source: Wikipedia

Artificial Super Intelligence
Artificial Super Intelligence (ASI) is an artificial intelligence that surpasses the brightest human minds in any area. It goes beyond Artificial General Intelligence which is au pair with human intelligence.

A few researchers observe that computers today are better than humans in several areas (like calculus) and are getting better and better in several more areas. This implies that once AI will reach the AGI stage as a matter of fact it will also be ASI. Hence, AGI will never happen, the shift will be from AI to ASI, skipping AGI. This is the "singularity". Machines will not become as smart as we are. All of a sudden they will move from inferior to us to be superior to us. Notice, however, that ASI does not imply Artificial consciousness.

Augmentation
Increasing the performances and extending the capability of an entity.

In our context we are referring to both humans, human augmentation, and machines, machine augmentation. In the longer term a symbiosis between humans and machines may augment both.

The natural process of random mutation and selection has extended living beings capabilities and performance. In our context augmentation is achieved by design. This is a superset of the natural evolution processes and includes, in the case of machines, random changes and selection processes, either controlled or open (self learning, self adaptation, self replication).

Autonomy
The ability of a system to be able to act independently and intelligently in dynamic, uncertain, and unanticipated situations. In addition, an autonomous system should be able to detect when its goals stand in conflict with the laws that govern its behavior and must have a way to "fail" gracefully in those situations.

Often varying levels (modes) of autonomy are used in the literature. There are four modes of operation: 1) in the *Fully autonomous* mode, the system operates without human intervention while adapting to operational and environmental conditions, 2) in the *Semi-autonomous* mode, the human operator and/or the system plan(s) and conduct(s) a mission which requires various levels of human-robot interaction. It should be noted that the system is capable of autonomous operation in between the human interactions (also called "bounded autonomy"), 3) in the *Teleoperation* mode, the human operator, using sensory feedback, either directly controls the actuators or assigns incremental goals on a continuous basis, from a remote location, and 4) in the *Remote Control* mode, the human operator controls the system on a continuous basis, from a remote location via only her/his direct observation.

Awareness
Knowledge and understanding that something is happening or exists. Source: Merriam-Webster
In our context we refer both to machine awareness (and self-awareness) and to human awareness of being part of a symbiotic entity.

Brain Computer Interface
Brain Computer Interface (BCI) is a means through which information/data is transferred from the brain to a computer. There are several technologies being used, and more are being studied. The goal of a BCI is to be able to capture the information required for a given goal, e.g., controlling a robot to operate on behalf of a paralyzed person, or to study the working of the brain and its neuron/neuronal circuit.

The interface can be based on either external sensing or may require invasive (implant) sensing. A given BCI is characterized by the technology and protocol used and its performances are measured with respect to the sensitivity and resolution provided.

Bio-interfaces
Interfaces that can establish a communication path between a biological entity and an artefact. In general, they act as a transducer between a living entity and an artefact. The two channels, from the living entity to the artefact and from the artefact to the living entity may use different technologies and protocols.

Examples of bio-interfaces are the protonic chips that use protons (ions) rather than electrons to communicate with living cells. Interfaces used in smart prosthetic limbs are another example of bio-interfaces, since they adapt the communications to the one supported by muscles in a limb or nerve termination. Other interfaces, like sensors to detect electrical activity (such as EEG, ECG) are not considered bio-interfaces.

Bio-Machines
A Bio-Machine is a machine that has been engineered using bio components, like bacteria, to acquire/deliver a specific functionality. As an example, bacteria (and genetically modified bacteria) can be used in symbiosis with artefacts to detect specific molecules.

Bioengineering
Bioengineering is the application of principles of biology and the tools of engineering to create usable, tangible, economically viable products. Source: Wikipedia

Brain implant
An artifact that is designed to be implanted inside the skull, either on the meninges or inside the brain. It may be used to monitor brain activity and/or to influence it. It can be a temporary implant or a permanent implant and must be bio-compatible with brain tissues.

Computer Brain Interface
Computer Brain Interface (CBI) is a means through which information or data is transferred from a computer to the brain. Like in the case of BCI, the interface can be based on external actuators or may require invasive (implant) actuators.

So far CBI technologies can only influence the working of the brain in some of its functions. As an example, CBI are used to block an epileptic seizure by interfering with the electrical activity underlying the seizure; another example is to alter depression (although no conclusive assessment on the effectiveness is available as of October 2018).

The stumbling block in creating a generalized CBI is due to the massive distributed nature of most brain functionalities making it practically impossible to interfere with all the involved neurons/neural circuits. Additionally, in most cases, any given neuron/neural circuit can participate in several functions, hence tampering with one to influence a function is likely to influence another function, often in a non-desirable way.

CRISPR/Cas9
CRISP/Cas9 is a technology used to modify DNA strings (now also being used to modify RNA strings adopting a slightly different protocol). It is the current tool for genetic engineering. CRISPR is an abbreviation of Clustered Regularly Interspaced Short Palindromic Repeats and was discovered through the genomes of prokaryotic organisms such as bacteria and archaea that developed this "technique" to defend themselves from virus infection.

Cas9 is an enzyme that uses CRISPR sequences as a guide to recognize and cleave specific strands of DNA that are complementary to the CRISPR sequence.

Deep Learning
Deep learning is part of a broader family of machine learning methods based on learning data representations, as opposed to task-specific algorithms. Learning can be supervised, semi-supervised or unsupervised. Source: Wikipedia

Digital Twin
Digital Twin refers to a digital representation of physical assets, processes, people, places, systems and devices that can be used for various purposes. The digital representation provides both the elements and the dynamics of how an Internet of things device operates and lives throughout its life cycle. Source: Wikipedia

A more business oriented definition from General Electric, one of the first companies to use Digital Twins:
Digital Twins are software representations of assets and processes that are used to understand, predict, and optimize performance in order to achieve improved business outcomes. Digital Twins consist of three components: a data model, a set of analytics or algorithms, and knowledge.

In the context of this White Paper, a Digital Twin is a digital representation of any characteristics of a real entity, including human beings. The characteristics represented by a Digital Twin are a subset of the overall characteristics of a real entity. The choice of which characteristics are digitalized depends on the purpose of the digitalization, i.e., the intended use of the Digital Twin.

Emergence
Complex systems, i.e., those systems composed by many parts that cannot be reduced without a loss of function (complex systems cannot be simplified without losing some of their characteristics, while complicated system can), often show characteristics that are not present in any of their components. This is often the case when one or more of their component is autonomous. The behavior of a system that is not the result of one of its component but that results from the interaction of the behavior of its constituent parts is called emergent behavior and the property of these systems in creating a whole "behavior" is called "emergence."

Evolution
A system *evolution* leads to a new system with different capabilities, forms and behavior. This resulting system (that can be software, hardware or a mix of both) inherits some of the previous

systems characteristics but has new ones usually as response to changing needs or to be more fitting to a certain environment. In this sense the word "evolution" reflects the concept of evolution in living beings. The evolution itself can be designed from the external or it can be generated internally, as an example by software applications that change their behavior as a consequence of experience. As in living species, autonomous systems can be designed to face and respond to selection pressure in various ways. Some routing strategies in the Internet have been designed to evolve in response to the success rate in establishing effective communications. In the future many software applications will be designed to be capable of self-evolution and the interactions among autonomous systems is also likely to evolve over time.

Exoskeleton
An exoskeleton is a rigid external covering for the body in some invertebrate animals, especially arthropods, providing both support and protection.

In our context an exoskeleton is a robot shaped in a way to wrap around part of a human body to increase the human strength (and relieve from fatigue). There are already many areas of application of exoskeletons, mostly in healthcare, manufacturing and military.

Generative Adversarial Networks
Generative adversarial networks (GANs) are a class of artificial intelligence algorithms used in unsupervised machine learning, implemented by a system of two neural networks contesting with each other in a zero-sum game framework. They were introduced by Ian Goodfellow et al. in 2014[209]. This technique can generate photographs that look at least superficially authentic to human observers, having many realistic characteristics (though in tests people can tell real from generated in many cases).

Genetic engineering
Genetic engineering is the direct manipulation of DNA to alter an organism's characteristics (phenotype) in a desired way.

Implicit/Explicit Communication
Autonomous systems are equipped with sensors that allow them to construct a model of the environment in which they operate and applications that can create awareness based on the dynamically changing condition in the environment. These applications are able, to different degrees of sophistication, interpret these changes. This results in what is called *implicit communications*. The behavior of an autonomous system, as well as any other system, creates data that once analyzed provides an implicit communication. As an example, a car blinking its direction lights provides an implicit communication to other cars in the vicinity that it is about to change direction. On the other hand, a system can generate a stream of data coding the information that it needs to share with other systems (autonomous or not) based on an agreed standard. This is called explicit communications. As an example a car can broadcast its position and velocity to all nearby cars to let them know of its approaching a blind crossing.

Internal/External Communication
Autonomous systems are composed of several parts that communicate with one another. Depending on the needs, *internal* communications may take a variety of forms that are "pre-designed." An autonomous system may also need to communicate with other systems, autonomous or not, and that "*external*" communication is also "pre-designed." There is, however, the possibility that autonomous systems moving in dynamically changing context may require establishing communication with other systems that were not known at design time. This *external* communication is a challenge that engineers need to face, and solutions are part of ongoing research. Basic standards can be defined at the physical and transport layer while

communications at the upper (applications) layers need to be flexible to meet unknown requirements.

Isomorphic
Similar to. An object A, physical or abstract, is said to have an isomorphic relationship with another object B, when it is possible to establish one or more mutual relationship between the two objects. The isomorphism is tied to a specific relation, like a dog and a cat are isomorphic in terms of number of legs and tail. Analogies, on the contrary, reflect abstract similarities but are not necessarily isomorphic. As an example: "I don't understand a word, it's Greek to me" means "what I hear is analogous to hearing Greek since I do not understand Greek and I do not understand what it is being said now".

Metabolome
The metabolome is the total number of metabolites present within an organism, cell, or tissue. The Human Metabolome project has resulted in the creation of the Human Metabolome Database (HMDB)[210] a freely available electronic database containing detailed information about small molecule metabolites found in the human body.

Multi-dimensional Digital Twin

A multi-dimensional Digital Twin extends the representation of an entity to include aspects like the process through which that entity has been manufactured, the place it was sold, and more. In case of a human digital twin it can include information on parents, on the environment, and so on. A multi-dimensional digital twin provides a multi-faced description of the entity and of its past and present context.

Nootropic
Nootropics (colloquial: smart drugs and cognitive enhancers) are drugs, supplements, and other substances that may improve cognitive function, particularly executive functions, memory, creativity, or motivation, in healthy individuals. Source: Wikipedia

Optogenetics
Optogenetics (from Greek *optikós*, meaning 'seen, visible') is a biological technique that involves the use of light to control cells in living tissue, typically neurons, that have been genetically modified to express light-sensitive ion channels. It is a neuromodulation method that uses a combination of techniques from optics and genetics to control and monitor the activities of individual neurons in living tissue—even within freely-moving animals—and to precisely measure these manipulation effects in real-time. Source: Wikipedia

Self-aware
Being aware of existing as an independent entity, having feeling, desires, purposes. In our context we address machines' self-awareness.

Sentient Machines
A sentient machine is a hypothetical machine that exhibits behavior at least as skillful and flexible as humans do. It is often associate to Artificial General Intelligence (AGI) and to artificial consciousness.

Sentient machines relate to the idea that a machine can have feelings and can appreciate that other entities can have feelings as well. It is a blurred area: there are computer (programs) that can feel the mood of people interacting with them but they do not feel anything in our sense of feeling, they just react in an appropriate way taking into consideration those (expression of) feelings.

Smart City
In this context a smart city is seen as a complex system, an emergent entity, resulting from the interplay of autonomous systems that all together create a symbiotic being, i.e., the smart city.

Smart prosthetics
Smart prosthetics are an evolution of prosthetics that embed processing and decision capabilities. In order to do that they have sensors and actuators; sensors report data to a processing units and the actuators execute orders provided by the processing units. More recently smart prosthetics have become equipped with technologies to interact with the person's body (and brain), i.e., to understand the intention of the person, acting in consequence, and providing sensation to the person.

We can expect smart prosthetics to become smarter in the coming decade, embedding intelligence and acting in symbiosis with the person, eventually giving rise to a more intelligent symbiotic behavior.

Superorganism
A superorganism is an organism composed of the symbiotic relationships of several organisms. In nature we have plenty of examples, and in a way most living beings are symbiotic expressions of a multitude of organisms (from sponges to human beings).

A superorganism can be composed of a multitude of similar entities (think about a hive, a superorganism composed by thousands of bees) or by different living entities (think of a cow needing bacteria to digest cellulose). A superorganism can be an abstract entity like a smart city, emerging from the loose inter-relations of different infrastructures and players (citizens, business...). In our context we are interested in superorganisms emerging from a mixture of atoms and bits of living entities and artefacts.

Transhumanism
The belief or theory that the human race can evolve beyond its current physical and mental limitations, especially by means of science and technology. Source: Oxford Dictionary

In our context we are not expressing a belief, rather we are pointing at the possible implication of technology evolution on humans.

Turing Test
The Turing test, developed by Alan Turing in 1950, is a test of a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human. Source: Wikipedia

Virtual Twin
A Virtual Twin, has similarities with a Digital Twin but differently from a Digital Twin it is created on spot through modelling of the perceived behavior of an entity and is used by the ones that created it. A Digital Twin is associated to the real twin; a Virtual Twin is associated to the entity using it (and different entities would each generate their own virtual twin to "understand" the world around it). The recent approach based on generative adversarial networks can be used to test potential effect of decisions on the virtual twin.

12 Acronyms

5G: 5th Generation Wireless System

ABET: Accreditation Board of Engineering and Technology

ACM: Association for Computing Machineryd

AES: Advanced Encryption Standard

AGI: Artificial General Intelligence

AI: Artificial Intelligence

ALIAS: Aircrew Labor in Cockpit Automation System

AMQP: Advanced Message Queuing Protocol

API: Application Programming Interfaces

AR: Augmented Reality

ASI: Artificial Super Intelligence

BCI: Brain-Computer Interface

BMI: Brain Machine Interface

BoK: Body of Knowledge

CAD: Computer Aided Design

CAM: Computer Aided Manufacturing

CAPSI: Computer-Aided Personalized System of Instruction

CARACaS: Control Architecture for Robotic Agent Command and Sensing

Cas9: CRISPR Associated Protein 9

CBI: Computer-Brain Interface

CMOS: Complementary Metal-Oxide Semiconductor

CO_2: Carbon Dioxide

CRISPR: Clustered Regularly Interspaced Short Palindromic Repeats

CPS: Cyber-Physical Security

CS: CyberSecurity

CSS: Cyber-Social Security

ConvNet: Convolutional neural networks

DBE: Dynamic Brain Emulation

DBS: Deep Brain Stimulation

DC: Direct Current

DIKW: Data Information Knowledge Wisdom

DNA: DeoxyriboNucleic Acid

DNN: Deep Neural Networks

EEG: ElectroEncephaloGram

ELS: Ethical, Legal, and Societal

ENG: Electroneurogram

EMG: Electromyographic

FDC: Future Directions Committee

f-MRI: Functional Magnetic Resonance Imaging

GAN: Generative Adversarial Networks

GDP: Gross Domestic Product

GPS: Global Positioning System

HCI: Human-Computer Interface or Human-Computer Interaction

HMDB: Human Metabolome Database

HRI: Human-Robot Interaction

ICS: industrial control systems

IEEE: Institute of Electrical and Electronic Engineers

IFR: Instrument Flying Rules

IoT: Internet of Things

ILS: Instrumental Landing System

MEMS: Micro-Electro Mechanical Systems

MPC: Microprocessor Controlled

MQTT: Message Queuing Telemetry Transport

MR: Mixed Reality

NIST: National Institute of Standard and Technology

ODF: Open Data Framework

PaaS: Prediction-As-A-Service

PGP: Pretty Good Privacy

PSI: Personalized System of Instruction

RNN: Recurrent Neural Networks

SAS: Symbiotic Autonomous Systems

SCADA: Supervisory Control and Data Acquisition

SDC: Self-Driving Car

SIM: Substrate Independent Mind

SSS: Symbiotic Autonomous Science

TMS: Transcranial Magnetic Stimulation

UAV: Unmanned Aerial Vehicle

VR: Virtual Reality

13 Appendix A: Impact on Education 2050

13.1 A Need and a Vision for Evolving Education Based on SAS

13.1.1 Motivation: Digital Twins in Industry

A number of companies like General Electric (GE), Tesla, and NASA, are creating Digital Twins defined as digital representations of their products like airplanes, cars, and satellites. The idea is to mirror a physical analog object in bits (i.e., a physical digital system, not resembling the original object in shape, but in its behavior) keeping the bit representation synchronized with the physical one. This allows various types of retrospective and predictive analysis on the Digital Twin that can provide a better insight into the analog one and lead to corrective actions when required. In this sense, Digital Twins are new tools for education: rather than studying and training on the analog object, one can study on its digital representation first. Many technologies like virtual reality can further enhance training and education.

The usefulness of a Digital Twin goes beyond that scenario. The Digital Twin can develop far beyond our physical and physiological limitations and can be helpful in our adaptation to the untenable challenge of doubling of knowledge over a decreasing time period. Another challenge is the need to educate individuals for more than one job due to automation, mechanization, and the unprecedented growth in deep learning (DL)[211] and artificial general intelligence (AGI), as described in the IEEE SAS White Papers I and II. Some challenges in developing better engineering education in cognitive systems have also been described.[212, 213]

13.1.2 Is There a Digital Twin of Me Already?

In a way, each of us has already several fragments of our own Digital Twin. Social media like Instagram, Facebook, LinkedIn and Twitter are collecting parts of our "self". Governments and municipalities are also collectors of parts of our "self." The health-care system "knows" much about our body and mind. Large physical department stores, as well as digital merchandise systems like Amazon "know" much about our purchasing needs and interests (electronic products, mechanical gadgets, books, music) so that they often suggest new and related products. Google has a deep insight into its users' interests ranging from scientific, technical, conceptual, philosophical, artistic, political, social, to theological. Travel companies know our interests about the world and our disposable income. In addition, the companies where we have been working and where we work now have other fragments, representing our acquired skills and habits. Insurance companies know much about our health risks. The educational institutions that we have used are also collecting records of what we have learned and how good we are at specific subjects, and can assess our intellectual potential. All of those fragments approximate what and who we are. They are distributed elements of our Digital Twins.

At this stage, all those fragments are dispersed. Some countries are developing rules to establish ownership of those fragments. For example, Italians have the right to access these data and information, and the companies physically storing the data have to grant them access. Having the right and actually being able to access them easily, are quite different stories.

13.1.3 Digital Twin: Aggregation of Fragments

In perspective, we should be able to aggregate those fragments into a more comprehensive one in order to represent our "self" better. In addition, it is most likely that the number of collected observations, data, and extracted information about ourselves will grow in time, thus leading to more and more accurate representation of our "self".

If we imagine a symbiotic relationship between a person and the corresponding Digital Twin, the symbiotic counterpart could form a very good understanding of who we are, sometimes through direct access to what we do jointly, some other times through the access to other Digital Twins.

13.1.4 A Symbiotic Digital Twin

In a formal way, our Digital Twin could come to represent both our skills, knowledge, and wisdom. It can also be flanked by applications taking into account the fading away of skills (what we lose when not practicing) and knowledge (when we forget). This information of our degrading skills/knowledge can be the starting point for a proactive education program.

Writing an article and presenting it at a conference, or attending a conference to listen to colleagues presenting their papers can also be mirrored by our Digital Twin. The same applies to the process of reviewing papers. Many publishers allow ongoing discussion on their published papers that could be monitored by our Digital Twin. Educational institutions, including IEEE, could contribute to the mirroring of their students or members into Digital Twins. These might come handy in creating customized and personalized education programs. An example of such a program is the *personalized system of instruction* (PSI) by Fred S. Keller (1899-1996; 97).[214] Since the manual administration of Keller's PSI is very tedious, *Computer-Aided PSI* (CAPSI) has been developed that has been running at the University of Manitoba, Canada for many years.[215]

In a symbiotic autonomous system (SAS), the skills, knowledge and wisdom should be shared among its component subsystems to enhance the overall performance of the system. Furthermore, the Digital Twin could start increasing (or decreasing) interaction between its component parts. Notice that in dynamical complex systems, the whole is not necessarily the sum of its parts. Through such nonlinear interactions, an emergent quality may appear that may not be found in any of its parts.

13.1.5 How Can a Symbiotic Digital Twin Help Me in Skill and Concept Learning?

If I live in a symbiotic relationship with my appliances at home or at work, the knowledge of what specific selection/action/effect (a "program" for short) I would most likely to be interested in at a given time becomes part of the global knowledge of the symbiotic Digital Twin. However, the knowledge about what programs are available and would fit my interest may lie in an appliance. Notice that Amazon's Alexa, Apple's Siri, Microsoft's Cortana are all moving in this direction. There are now thousands of programs (such as streaming contents) to choose from and be of serious interest to me, but they are just too many for me to be aware of at any given time. The same applies to the millions of YouTube clips, tweets, and other pieces of information that could become an integral part of my education process, but I will never know that they even exist.

The same scenario can be envisaged with studying magazines, white papers, reports, textbooks, monographs, and research papers. Another scenario emerges with all the courses and MOOCs (massive open online courses) that are available at Coursera launched in 2012 by Stanford University (over 2,000 courses), EdX launched in 2012 by Harvard University (over 1,200), Udemy (over 2,500), Udacity (200), MIT OpenCourseWare (2,200), XuetangX founded in 2013 by Tsinghua University (over 500), Lynda (3,300), Khan Academy started in 2006 by Salman Khan, TED (1,890), The Great Courses (500), and over 700 universities offering MOOCs[216,217]. Designing of online courses[218] and improved comprehension[219] is discussed.

A similar scenario emerges when doing research. Finding and reading relevant research papers, technical magazines, technical reports, white papers, and technical books is very time consuming. These are all examples of the "where can I find it?" problem in education. A Digital Twin could help in these situations.

13.1.6 More Reasons for Symbiotic Digital Twins: Knowledge Doubling and Its Half-Time

With the explosion of data, information, knowledge and wisdom, we would have to spend all our available time searching for what is needed for our education and work. We cannot just ask a teacher or professor to answer our questions outside their class or research area. Today, search engines still provide millions of hits that have to be reviewed for relevance. Finding the relevance in the sifted out and even prioritized material takes time. Since our reading and comprehension abilities are slow (the average reading speed is around 300 words per minute), it might take up to four hours to keep up with daily emails, news digests, blogs, magazines, books. This time for upkeep on the news reduces the time for creative work.

According to Buckminster Fuller's "knowledge doubling curve" in 1982, all human knowledge generated and transmitted doubled in size around year 1500. It doubled again by 1750 (only 250 years), and doubled again by 1900 (just 150 years). With those rates, humans were able to adapt to the growth and change. It became harder to adapt when the doubling took 25 years around 1950. The knowledge doubling today is much shorter (around 13 months). As an example, the number of annual patents increased from about 50,000 to more than 325,000 over the last 50 years. Many in IBM predict that in not-too-distant future (2020), the knowledge doubling will happen in 12 hours. It is not feasible for a human to adapt to that rate. The concept and implementation of a Digital Twin seems to be a necessity now.[220]

There is another reason for Digital Twins: the knowledge half-life. In his book Future Shock[221], Alvin Toffler stated that "the illiterate of the 21st century will not be those who cannot read and write, but those who cannot learn, unlearn, and relearn." The knowledge and skills acquired in our schools and successive jobs diminishes in value and requires continuous updating, not once but throughout our lives.

How long does it take for knowledge to become outdated and irrelevant, or even incorrect? The half-life of knowledge (i.e., the amount of time it takes for knowledge to lose half its value) is often used to indicate the devaluation of knowledge in various disciplines. As might be expected, the knowledge half-life in aggressive disciplines like science, engineering and technology is shrinking fast.

13.1.7 Knowledge Tsunami and Organizations

The conditions when knowledge-doubling occurs exponentially, while the knowledge half-time decreases, may have a tsunami effect on any society, organization, company or other organizational unit. The SAS with Digital Twins could be very helpful in increasing our resilience in some of the following areas: (i) Curation of knowledge (organizing and filtering according to agreed-upon criteria to eliminate irrelevant knowledge); (ii) Knowledge fusion (to discover and clean errors present in sources, as well as mistakes made in the process of knowledge extraction from sources); (iii) Plagiarism management (to generate new knowledge); (iv) Knowledge vetting (to identify and verify sources for quality of the content used in the organization); (v) Intellectual property management (separating intellectual property, trade secrets, and copyrighted information from generic and public-domain content); (vi) Knowledge sunsetting (to identify knowledge that cannot be used any longer); (vii) As traditional libraries dwindle, creation of a Digital Twin "librarian" that knows the needs of the organization and its members would be beneficial; (viii) As traditional publishing also dwindles, creation of a suitable Digital Twin "publisher" could benefit the organization.

13.1.8 Illustration: My Digital Camera and I

As described above, using my Digital Twin to understand what I know is a starting point. Suppose I need to learn something about a tool or a friendly, but sophisticated instrument. Which of the following two options would be wiser while learning the new skills: (i) to learn about the tool by myself as the human component? or (ii) to have the tool learn what I need to know, and show me how to bridge the gap? For example, I have just bought a very complex digital camera, and I started to learn how to use it by leafing through its manuals, watching courses on YouTube, and downloading new software to manage the new types of files. I am far from being at ease with the camera, and I suspect that it will take me a year before becoming used to it. By that time, very likely, I will be missing some features, and will forget something that I had learned on the way but had no opportunity of practice it.

If I had a Digital Twin, she might suffer from the same problems I am having, but that Digital Twin might be analyzed by a smart advisor that could identify knowledge gaps and make up for those gaps by adapting my camera, my computer, and my smart phone. In a way, the teaching can go both ways: not only to me (with my Digital Twin following my success or difficulty in learning the skills), but also to the other components making up the SAS. Notice that, today, although I am far from being in a symbiotic relationship with my camera, my computer, my smart phone and the related software as they pertain to my photographic activities, something could be done to improve my education. Even in this loosely-connected environment leveraging on the sketchy Digital Twin that is starting to mirror my "photographer self," the potential for a new form of education starts emerging when we consider the connections that can be created with the other sketchy Digital Twins associated with the digital camera, its applications, computer, and the smart phone. Each of the above Digital Twins is still a pale instance of a much smarter and capable Digital Twin that we might have in the future. The symbiotic relationship between a human and its Digital Twin is, once more, the beginning. A symbiotic system may include not only more than one my Digital Twins, but also Digital Twins of other humans.

13.1.9 From Competition to Cooperation and Knowledge-Sharing Society

In the past, individuals and teams in companies and organizations worked independently, competing in the quest for reaching the finish line first. Competition has been recognized as healthy, but leaves many in the dark.

The emerging possibility of educational Digital Twins in the SAS environment creates a set of new opportunities for knowledge to be uncovered and discovered, distributed widely, and then enhanced by the same individuals, but now working in the SAS environment. In fact, new knowledge is now developed mostly by collaborating or even cooperating teams rather than individuals.

13.1.10 What Prevents us from Having the SAS in Education Today?

There are at least three missing links to create a SAS in education: (i) a reasonable Digital Twin itself, (ii) symbiotic interconnects, and (iii) rigorous modelling of SAS.

While the concept of a Digital Twin is not new, the concept of a Digital Twin of a human is fairly new. While attempts are being made to develop machines and systems capable of acting in a manner that could be described as: adaptive, autonomous, intelligent, perceptual, cognitive, conscious, and symbiotic with humans and the environment, Digital Twins are not here yet.

The second missing link is the connectivity between the different parts of the symbiotic system. When the Internet has been developed sufficiently, we thought, happily but mistakenly, that education could be improved on a dime by delivering many online courses (in different forms such as the massive open online courses, MOOCs) to many people at any time and any place. Of course, this helped, as millions of new individuals obtained access to education that was not available to them before. However, the interconnections were not symbiotic. We must develop a new class of symbiotic interconnects.

The third missing link is proper rigorous modelling of SAS based on brain-inspired and socially-inspired processes. This modelling also implies quintessential changes in signal processing in the simulation of the processes. In the past, the majority of signal processing was done on a single scale (mono-scale). More advanced models of reality included multi-scale signal processing. Cognitive system and SAS require not only more elaborate multi-scale, but also poly-scale modelling and signal processing. Some of the definitions will be provided in the next subsection, and summaries of the modelling techniques may be described in the IEEE SAS White Paper III.

13.1.11 How Can We Get There?

The missing links can be addressed by some innovators, including some specific initiatives by IEEE (obviously not related to the above case on how to learn using a photographic camera to make aesthetic photographs in a reasonably short time, but, as an example, to support careers path of its IEEE members).

13.2 Learning Ecosystems: Some Definitions

13.2.1 Models of Learning

Education and learning have been existential to humanity and have been evolving throughout the millennia. Recently, educational systems have been changing more rapidly as a result of sociocultural, political, economic, demographic, and technological changes.[222] New technologies (such as social media, serious games, adaptive software, software-defined communications systems) and emerging practices (openness, user modeling) in particular, have facilitated opportunities to transform education, learning, and particularly teaching. With the advent of the Internet, the brick-and-mortar education has been expanded to network education through distance education, massive online courses (MOOC). Social media include: Facebook, Twitter, Flickr, Digg, YouTube, Upcoming, LastFM, Techorati, MyBlogLog, SlideSharing.

What is learning? It is the acquisition of knowledge or professional and other skills, through either self-study, or by being taught by parents, friends, teachers and/or tutors, or intelligent systems, or workplaces, or organizations, all with different degrees of experience, starting from childhood, through adolescence, to professional life and seasoned years. Learning occurs in a systematic way (schools, routine reading of scientific, technical and other news digests, discussions with family, colleagues and friends) and through the experiences and events that occur in life less predictably. This experiential learning is also fundamental in acquiring knowledge that is important in decision making. Learning alters the functioning of the brain[223].

What is education? We have just defined learning as the process of acquisition of knowledge in a discipline, hard and soft skills, critical thinking, creative thinking, values, beliefs, and habits.

Education is then the process of facilitating learning by teaching, training, discussion, interactive experiential experiments, and directed research.

We learn best when acting on what we have learned, thinking about it, and actually participating in the real world. Effective and impactful learning requires that we immerse ourselves in the process completely: with our will, senses, feelings, intuition, beliefs, and values. It often starts from our own inquiry. This is a very important point to make: the impact of education on us is determined by our engagement; technology by itself can help, but is not a replacement for the engagement. For the symbiosis to have the multiplying effect, we must engage the technology too.

In the past, learning was modeled as a linear process in which progression through various educational events produced an additive effect. Today, researchers and educators model learning as well as growth and development as a nonlinear dynamical system. Our proposed Digital Twin symbiotic educational system is intended to assist in our engaged lifelong learning.

A learning ecology includes (i) learning concepts, (ii) learning dimensions, (iii) filters, (iv) conduits. Learning is a process that involves a number of foundational concepts, such as signals and noise in the real and or virtual environment, observables and data, information, knowledge, meaning, understanding, wisdom, and vision. We learn because (the dimensions of learning): we need to know, we want to do something, want to be somebody, want to create, transform, change. Educational filters affecting our outcomes include values, perspectives and beliefs. Educational conduits include selected language, media, and technologies engaged in the process.

The educational process can be either formal or informal, it can be done through self-study or through communities, with the help direct performance support or monitoring and mentoring, all gaining experience through simulation, emulation, experiential learning, internship, co-op, or apprenticeship.

13.2.2 Current Models of Learning

Marcy Driscoll[224] provided a classification of epistemologies including (i) **Behaviorism** (objectivism) in which reality is external to the mind and knowledge and perception are acquired experientially, (ii) **Cognitivism** (pragmatism) in which knowledge is a negotiation between reflection and experience, inquiry and action, and (iii) **Constructivism** (interpretivism) in which knowledge is an internal construction and is informed through socialization and cultural cues. De Corte provided an overview of historical developments in the understanding of learning.[225]

Since human behavior cannot be fully understood by the reductionist behaviorist approach (decomposing the system into linear parts and then reconstituting it), the idea of Gestalt psychology became more attractive in which the organized configuration of components in the whole system is considered. This approach to learning requires information-processing techniques. **Social constructivism** might be a good model for representing interactions between learners and their grounding contextual environment. This is also combined with shifting away from artificial exercises to real-life situations. The current view on learning includes **adaptive competence** characterized by the so-called CSSC learning ("constructive" to signify that the learners are responsible for constructing their knowledge and skills; "self-regulated" as the learners use their own strategies to learn; "situated" to indicate learning in the context of the environment, rather than abstracted from it; and "collaborative" to indicate a team rather than an individual approach).

George Siemens and Stephen Downes [226,227,228,229,230,231,232,233,234,235,236,237,238] proposed another learning theory called **connectivism**, based on various ideas from networking and dynamical systems, i.e., complex interacting nonlinear systems that can develop chaos and self-organization. Connectivism is based on distributed adaptive knowledge (viewed as composed of connections and networked entities) and tries to explain how the new knowledge is created. Siemens uses the example of senior citizens that have been linked as mentors to elementary school pupils, thus forming a new distributed knowledge. In that view, learning is a process of connecting specialized nodes or information sources and may reside in non-human nodes. Thus, knowing where to find information is more important than knowing the information element. In contrast, the other three theories do not address the new distributed knowledge creation.

A number of writers criticized connectivism as a mere pedagogical view and did not consider it as a theory of learning[239,240]. Connectivism appears to be related to our proposed Digital Twins in a symbiotic network, although only through the same Internet that it uses, with several fundamental differences. it is useful to hear the critique of connectivism and its distributed knowledge and learning processes, and avoid possible mistakes in attempting a formulation of the Digital-Twin-based learning. Our longer-term objective is to develop a theory of the Digital Twin symbiotic system. The theory must explain corresponding phenomena, must be verifiable through measurable observables, and should predict future behavior of the system within the horizon of predictability.

13.2.3 A summary of the DIKWV Model

Figure 8.1 illustrates a common progression of situated learning stages from (i) observations of the environment and (ii) data extracted from the observations to (iii) information, (iv) knowledge, (v) wisdom, and (vi) vision. This is often called the data-information-knowledge-wisdom (DIKW) pyramid model.

In the hierarchy of human scientific and technical development, data appears at the starting point for our analysis and drives our data-driven learning. Analysis of the data may produce useful information. Note that if we consider the data as a stack of hay, extracting the information could be compared to finding a needle in the stack. Useful information may lead to knowledge (information woven into a garment). Good knowledge may lead to enhanced wisdom needed by the student (the wearer of the garment) to make good decisions. Although this model is fairly limited, we will describe it to link it with the concept of Digital Twins.

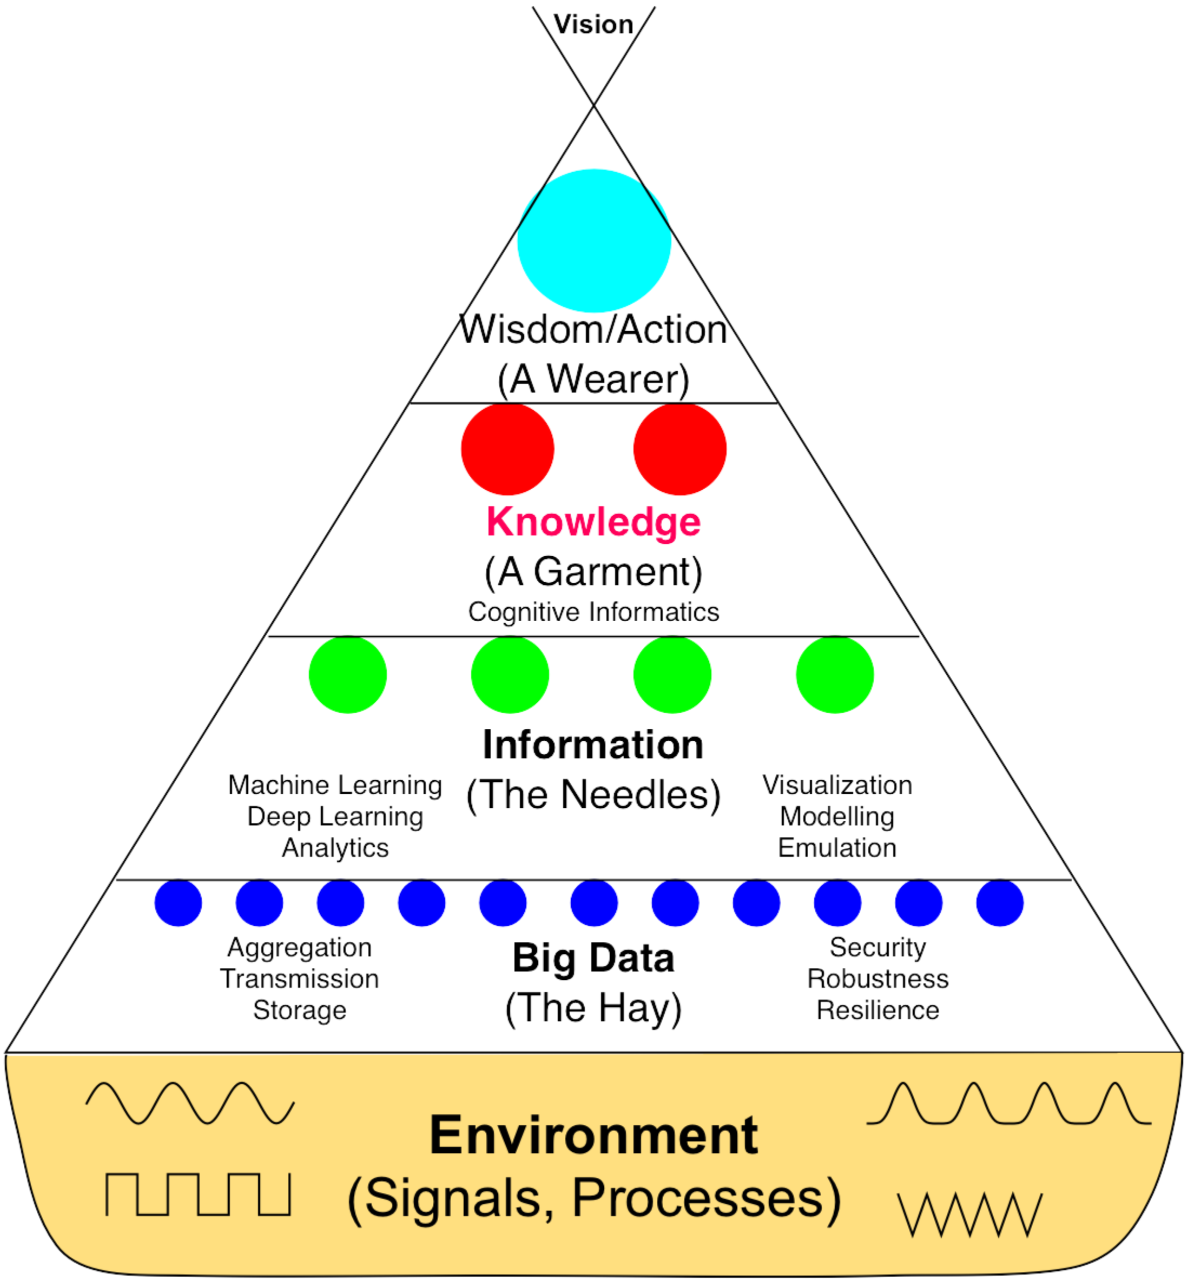

Fig. 13.1. The knowledge pyramid.

13.2.4 Signals (Representations of the Physical Processes)

Signals are mathematical or logical abstractions of physical processes, either spatial (like images) or temporal (like speech, heartbeat, pressure, temperature, or velocity), or spatio-temporal (like video or Doppler radar). The signals exist in a physical environment that may include living organisms. The physical signals are often translated into an electrical form (either voltage or current).

Signals can be linear (decomposable into their constituent components) or nonlinear. They can be stationary (when their statistical moments do not change) or non-stationary. They can be deterministic (formed without a chance), or stochastic, or chaotic.

Signals can be un-correlated (e.g., white Gaussian noise), have some mid-range dependence (e.g., terrestrial images), or have long-range dependence (e.g., Lévy walks). Many natural processes exhibit the heavy-tail long-range dependence.

Signals may be either analog (infinite resolution in time and magnitude), discrete (finite resolution in time and infinite resolution in magnitude), digital (finite resolution in time and finite resolution in magnitude), or boxcar (infinite resolution in time and finite resolution in magnitude). We often convert the analog form of signals to their equivalent digital form.

Historically, processing of the signals was done at a single scale (mono-scale). For example, the spectral (Fourier) decomposition of a stationary signal can be at a mono-scale. To overcome the Heisenberg limitation of the Fourier analysis, a time frequency analysis is required. Wavelet analysis is an example of multi-scale analysis. Signals with long-range dependence require a poly-scale analysis. Many Symbiotic Autonomous Systems will require the latter form of signal analysis.

13.2.5 Data and Capta (Signal Representations, A Stack of Hay)

Data are objective representations of spatial, temporal, or spatio-temporal processes (signals) such as images, speech, or video. For reconstruction and proper analysis, the analog signals must be sampled at above the Nyquist rate, i.e., twice the highest frequency component in the signal, and quantized to a number of bits *n* that is dictated by the dynamic range of the signal. The data acquisition period must also be selected properly to obtain a sufficiently large number of data points with the required sampling rate. If these requirements are not satisfied, the data are invalid.

Valid data may still contain errors due to the ever-present noise during the data acquisition, data transmission, data storage, or mistakes in data entry. For any analysis to produce good results, the acquired data must be valid (above the Nyquist rate), accurate, and precise (reflecting the dynamic range). Otherwise the garbage-in-garbage-out (GIGO) principle applies.

The data acquisition period may vary from femtoseconds (e.g., plasma dynamics in a laser) to minutes (ECG recordings), hours, days, weeks, months, years (the Hudson Bay record of the hare-lynx populations), centuries (the Sunspot cycle data or temperature data), or even millennia (the water levels of the river Nile). There are also many examples of streaming data without a planned end.

The data point can be in the form of a digital sample of an analog signal with *n* bits each. The entire data record would then include many successive samples.

Data may also result from collecting observations such as the number of packets on the web or the number of scintillations, arriving at a destination per unit of time.

Thus, data can include not only *n*-bit binary samples of signals but also decimal or any other digits, as well as any other symbols. The symbols are often represented by binary sequences corresponding to either ASCII characters or a Unicode (e.g., UTF-8).

Data representations are designed to be stored in memories either electronically in silicon, magnetically in hard drives, optically, or even mechanically.

In all cases data points must be unambiguous. They must be unique, with a clear understanding of the measurement units employed (e.g., metric, or Imperial, or US).

Notice that the acquired record of data does not have to be used for analysis in its entirety. Instead, we can select a part of the record (capta) most relevant to a specific analysis of facts. Thus, capta are richer than data because they are taken in a context.

Notice that "data" is the plural of "datum," as used in Canada and many other countries. In the US, data became an uncountable noun.

Data are collected in many fields, including: scientific, natural, statistical, financial, metrological, geographical, transport cultural. Big data (defined by their volume, variety, velocity, variability, and value) come from many sources, including: social media, transactional data, enterprise data, archives, public data activity generated.

Data should never be confused with information. For example the equation 2B + ¬2B = ? is unique through its symbols, but has at least two distinct pieces of information. On the other hand, two distinct collections of characters "Mozart" and "Моцарт" carry the same information.

13.2.6 Note on Big Data

Many of us are now involved in dealing with big data due to the increased capabilities of wireless connectivity between sensors/actuators and computers. What is the definition of big data? Big data refers to techniques and technologies to capture, process, analyze and visualize large datasets in a real-time or near-real-time.

How are big data characterized? Initially, big data had 3V characteristics, and now it has 8V:

1. Volume (the size of the data; can the needed knowledge be found?)
2. Value (extracting information and knowledge from data; can the knowledge be found when needed?)
3. Veracity (is it information or disinformation?)
4. Visualization (can I make sense at a glance? Is it useful for a decision?)
5. Variety (is the information balanced? The data could include a mixture of text; images and video, sounds, speech, music, position data, traffic data, environmental data such as temperature, pressure, humidity, light intensity, gas composition, volumetric radar data, biomedical data, security data such as fingerprints, irises, faces, voiceprint)
6. Velocity (the speed of data analysis)
7. Viscosity (does it stick with you, and call for action?)
8. Virality (does it convey a message, and can be passed on to others?)[241].

The smallest unit of information is the binary digit (bit) (0,1). A larger unit is a nibble (N) with 4 bits, and a byte (B) with 8 bits. A larger unit is a Kilobyte (2^10 = 1024, KB). To distinguish this binary number from the decimal kilobyte (10^3 = 1,000 = kB) a new name was coined by the International Electrotechnical Commission (IEC) in 1998, starting from "kibibyte (KiB)" where the "bi" denotes "binary". Donald Knuth suggested calling the binary units "big Kilobyte" (KKB where each abbreviation is the same as for the decimal units, except for the first character that is repeated)[242].

Small data would range from KKB to big Megabytes (2^20, MMB), big Gigabytes (2^30, GGB), big Terabytes (2^40, TTB), big Petabytes (2^50, PPB), big Exabytes (2^60, EEB, big Zettabyte

(2^70, ZZB), big Yottabyte (2^80, YiB), big Xenottabyte (2^90; XXB), big Shilentnobyte (2^100, SSB), big Domegemegrottebyte (2^110; DDB); big Icosebyte (2^120; IIB); and big Monoicosebyte (2^130; MIB).

To store all the text **ever written** would require about 5 EEB. To store the human speech ever spoken would require 42 ZZB (when digitized at 16 kilo samples per second, kSps, and 16 bits per sample, bpS). In contrast, the Square Kilometre Array (SKA) is expected to generate 1 EEB of data **a day**.

As in other areas of signal processing, big data can employ the following scale-related analytics:

1. Monoscale analytics (segmentation and result stitching)
2. Multiscale analytics (multiband independent analysis)
3. Polyscale analytics (simultaneous analysis at all scales)

Examples of Data Analytics for BD include

- Dimensionality reduction (principal component analysis, PCA)
- Data separation (e.g., independent component analysis, ICA)
- Non-negative matrix factorization
- Discrete signal processing on large-scale graphs
- Compressive sampling
- Sparse Fourier Transform (SFT)
- Data sketching (in streaming data and sliding window processing)
- Subspace clustering
- Dictionary learning
- Tensor- and kernel-based learning
- Scalable inference and optimization algorithms for decentralized and
- Online learning problems
- Decentralized and dynamic estimation, imputation and prediction
- Outlier detection (health, energy, communications)
- Outlier-resilient algorithms

13.2.7 Information (Understanding Relationships; A Needle in the Stack)

Information is the collection of relationships between data (representing signals contaminated by noise). Information can be carried by data and can be extracted from the data. A single data point may carry some information, but a collection of data points may carry much more information.

Data are transformed (processed) into information when they acquire a suitable form to communicate knowledge, ideas, conclusions, or meaning. This requires the understanding of the relationships between data.

For example, when building a precise electrical circuit that requires a 100 ohm (Ω) resistor, acquiring a single resistor will not produce the desired result as resistors with exactly 100 Ω would be very expensive. The resistors have tolerances. If we measure the resistance of say one-hundred 100 Ω resistors, we would be able to calculate the average and the spread of the values (variance), thus establishing the real tolerance of the batch (which should be either 1%, or 5%, or 10%). If we were the manufacturer of the resistors and the average would not be 100 Ω, we would have to change the manufacturing process of the resistors to return to the expected value.

In fact, the repeated measurement of the values would produce a histogram of the values that could be converted to the probability mass function (pmf) that would allow us to compute not only the first two moments (mean and variance, but also skewness, kurtosis and higher)[243].

Furthermore, the pmf would allow us to compute the Shannon self information and the expected weighted value (entropy) of the ensemble of measurements.

Another example is the trend in temperature changes of the climate.

Information should never be confused with meaning.

It should also be clear that the Shannon information has been designed for analysis of transmission of bits over channels with noise, and is not related to the meaning of data. There are other information theories, including Rényi[244], Kolmogorov-Sinai, Chaitin[245], and Stonier[246].

13.2.8 Knowledge and Learning (Understanding Patterns; The Garment That We Can Wear)

Knowledge is what we know about something through the understanding of the patterns in information, placed in the context of skills, experience, value, and meaning. It is the map of our world in the brain acquired through cognitive processes (though the exact location of the map in the brain is still studied). The knowledge map of the physical world is not static but updated almost continuously through our natural senses (our eyes, ears, nose, mouth, and skin). This knowledge is actionable.

Learning is a specific case of sharing information and constructing meaning to acquire knowledge, and is important in Informatics (Information Systems), Information Science, Communication Science, Sociology, and Philosophy.

The simple statistical processing of the data in the previous example indicates how information can lead to knowledge. The knowledge is the probability mass function (pmf), the mean, variance and other moments (if the distribution is not Gaussian) of the batch of resistors. This knowledge also contributes to the knowledge map.

We cannot store the knowledge map outside our brain directly, although there are many indirect methods that are being contemplated.

Furthermore, in cognitive machines and Symbiotic Autonomous Systems the maps can be developed independently of us.

13.2.9 Wisdom, Competence and Meaning (Understanding Principles)

Based on the knowledge map collected and experience from the past decision making, we can now generalize the specific situations to what-if scenarios and simulate new scenarios in order to develop rules, operating procedures and policies with respect to ongoing and planned actions. This state is called competence or wisdom and is necessary for effective decision making, including judgments.

In the context of the simple example of resistor manufacturing, the decision could be (i) to continue the production, (ii) to modify the production to improve the resistors' accuracy and improve their tolerance, and (iii) to collect new data to reduce any ambiguity in the decision process.

Although this personal competence level of meaning comprehension is colored by prejudice, it is the personalized frame-of-reference. The interpretation of meaning requires (i) basic understanding, (ii) connotation with nested consequences, and (iii) intention to help with hidden meanings.

There is also a shared meaning when acquired through interaction of more than one individual through consensus. The shared meaning may exceed that of the best individual meaning by reducing potential biases and prejudices.

The shared meaning facilitates the development of communicative actions (when the actors agree) or discursive action (when a dialog is required to settle on the action), and a strategic action (to achieve a strategic advantage).

13.2.10 Vision and Creative Imagination (Understanding and Inventing the Future)

The competence (wisdom) combined with extensive experience at various operational levels may lead to a vision that could improve the impact on the environment and humanity.

13.3 From Open-Loop to Closed Loop Education

13.3.1 Background

The traditional educational classroom concept has been used for centuries at universities and other educational institutions. It was modified in Prussia in the 1770s as a means to disseminate information and knowledge to students. The main objective of the concept was to deliver a standardized (one model fits all) curriculum to as many students as possible. By the 19th century, universal compulsory education at the elementary level became available in most European and North American countries.

Before 1945, on-the-job training and “vestibule training” were dominant forms of teaching workers the skills they needed to operate in factories or service industries. After 1945, many corporations adopted the Prussian classroom model for training. In the 1980s and 90s, “corporate universities” were established by main corporations such as Fujitsu. That was before the knowledge tsunami.

The future of education will have to involve Digital Twins in Symbiotic Autonomous Systems (DT-SAS). However, they do not exist yet. Can we do anything before the radical development of the DT-SAS?

13.3.2 An Open-loop Educational System

An assertive response to the above question involves closing of the educational loop. Our educational system has mostly an open-loop form, as illustrated in Fig. 8.2.

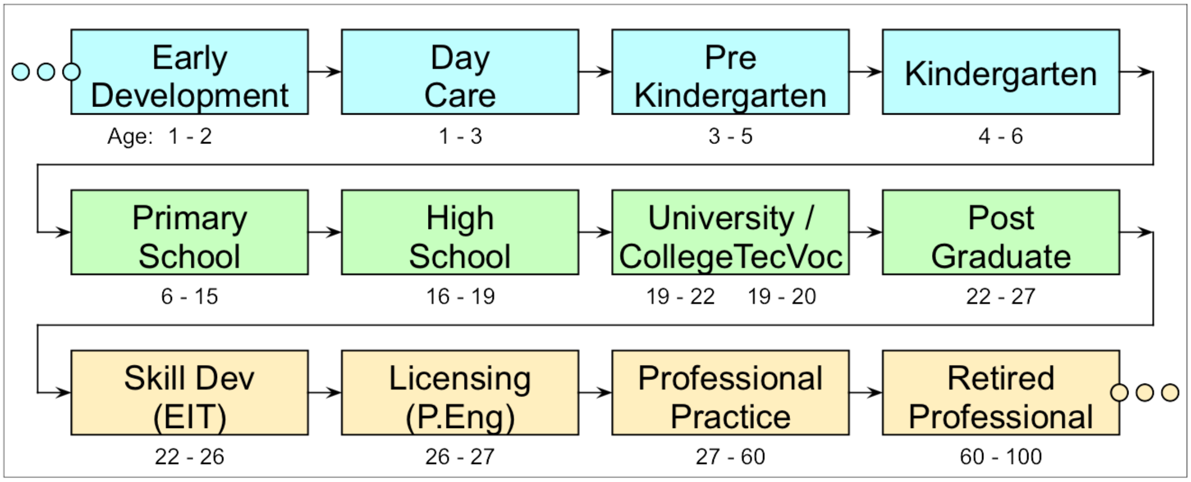

Fig. 13.2. An open-loop educational system.

The first layer of the critical education of a young child starts by their parents (early development). The next critical stages of education may advance through day care to pre-kindergarten and kindergarten where the curiosity, dexterity, and making sense of their "selves" are formed.

The second layer of education starts from the primary school and high school where more fundamental knowledge is imparted. Our usual focus of outreach is placed on the mini-universities for primary school students and summer camps for high-school students. High-school graduates may select a 1½ year vocational school, or a 2- to 3-year college or a 4-year university. University graduates may either continue with Master's and Doctorate postgraduate studies, or embark on a professional job as young professionals (YPs).

In many countries, further professional development is required on the job as engineers in training (EIT) followed by a certification and licensing of the individual as a professional engineer (P.Eng.). At this stage, a professional engineer can embark on unsupervised professional work. Eventually, at an appropriate age, the professional retires.

Figure 8.2 shows how a student progresses through the chain of stages. It is an open-loop educational system because the experience gained from the more advanced stages is not fed back into the education of YPs or students at the university/college stage, or even high-school students[247].

13.3.3 Evolution of a Professional in a Discipline

The evolution of a person from a student to a practicing professional in a specific discipline is illustrated in Fig. 8.3. A discipline requires (i) a well-defined collection of knowledge to practice the profession, (ii) a code of ethics, and (iii) an association to enforce legal obligations of the practicing professional[248]. For example, Civil Engineering and Electrical Engineering are well-established disciplines[249], while many others are still evolving. The main nodes in this model include (i) professional education evaluated by an independent body, (ii) further skill development through co-op and internship programs, as administered by the educational institutions and monitored by a professional association, (iii) licensure and certification administered and renewed by the professional association, and (iv) professional practice in industry and business, as monitored by the corresponding professional association.

Professional societies may also be involved in any stage of the process.

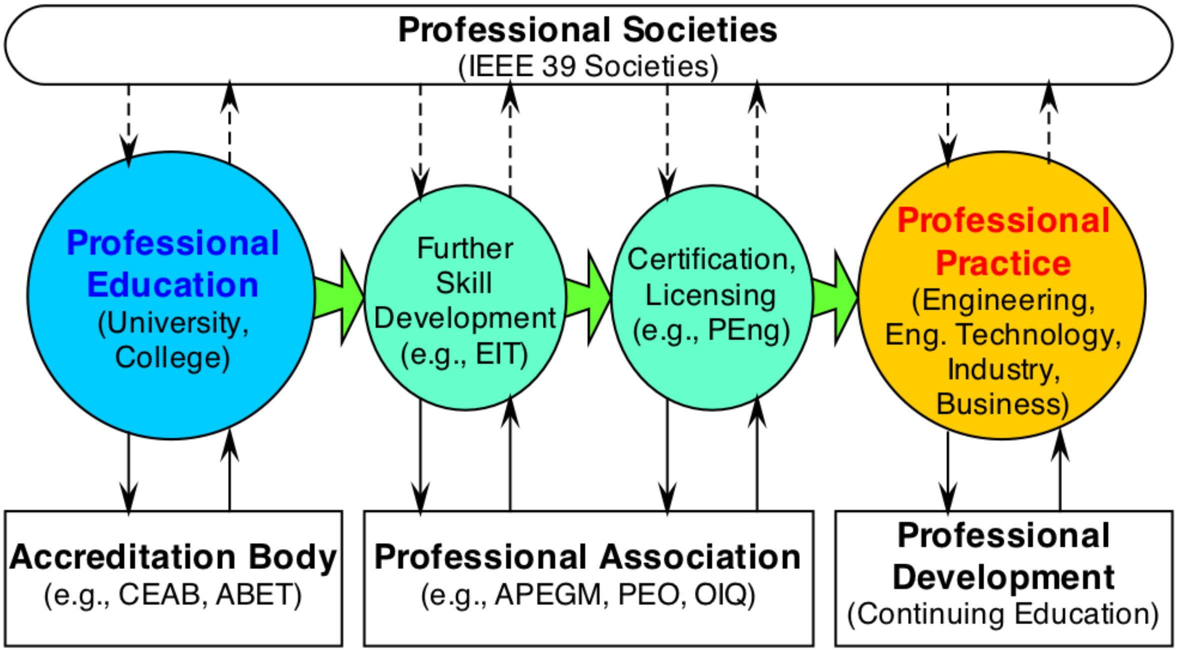

Fig. 13.3. An open-loop educational subsystem with some outside interactions.

Professional Education and Its Accreditation

Professional schools must craft their educational programs so that the programs could be accredited by an independent body. Up to the mid-1990s, the engineering and technology programs at universities and colleges were defined and evaluated by the corresponding accreditation bodies in their home countries. For example, the Canadian Engineering Accreditation Board (CEAB) of Engineers Canada has been accrediting engineering schools in Canada.[250] In the United States, the engineering and technology programs at universities and colleges were defined by the Accreditation Board of Engineering and Technology (ABET) since 1932.[251,252] In the United Kingdom, the Engineering Council UK is responsible for accreditation and requires that their graduates meet the requisite benchmark standard for a discipline.

Over the last two decades, the newer programs have been relying much more on the body of knowledge (BoK) definition of the disciplines. The new accreditation criteria have also been upgraded to specify outcomes and attributes of a student with respect to their knowledge, skills, and attitudes. Before graduation, all the engineering students are now required to be involved in a group design final capstone project. In addition, many students are placed into internship or co-op with companies developing products. Preparing for an accreditation visit is a very complex and laborious process.

Professional Mobility

The educational programs must also be compatible with other countries in order to allow for mobility of the professionals. Several international agreements were signed to help recognize graduates from other countries through documents such as the Washington Accord[253], the Sydney Accord[254], the Dublin Accord[255], the International Professional Engineers Agreement[256], Asia Pacific

Economic Cooperation Engineer Agreement[257], and the International Engineering Technologist Agreement[258].

Further Skill Development

After graduation, engineering graduates are required to work for several years under the supervision of a licensed engineer (as Engineers-In-Training, EITs, or other titles) and are evaluated by their local professional organization such as the Association of Professional Engineers and Geoscientists of the Province of Manitoba (APEGM). The evaluation is systematic, with required input from the supervisor and other sources.

Certification and Licensing

A certification process of a professional ascertains that the individual has the expected competencies, as defined by the accreditation bodies and the appropriate BoKs. Licensing extends the certification to include active oversight of the profession, including disciplinary actions in many malpractice situations. Licensing also requires special examinations that must be passed by the candidates. Furthermore, for internationally-educated engineers, various organizations have been instituted to transfer the engineer's credentials to the new country of residence.[259]

Professional Practice

A licensed professional is deemed qualified to practice engineering independently by providing service to the public and the profession. An engineer who is not licensed cannot practice engineering in Canada legally. Most professional associations now require that each practicing professional reports on their professional development (e.g., continuing education and training) to satisfy specific requirements (i.e., specific activities and a specific number of professional-development hours). This is required because the expanding areas of knowledge do not allow all the material to be fitted into a four-year baccalaureate degree, and continuing professional development must be part of our lives.

Professional Societies

A professional society may also play an important role in the professional development through its in-person and online courses, tutorials, workshops, and seminars. For example, the Institute of Electrical and Electronics Engineers (IEEE established in 1884) has 39 societies, covering most of the engineering disciplines. The Association of Computing Machinery (ACM established in 1947) has 37 special interest groups (SIGs), also covering all the major computing disciplines. The societies may also be helpful during the academic time of a student (e.g., enhancement laboratories, workshops, networking), as well as during their postgraduate training.

13.3.4 Why is This Model Incomplete?

Figure 8.4 shows the traditional morphing process of an individual into a professional. Is this morphing process complete? If in doubt, look at the progression of the arrows (from left to right only). There is no feedback from the professional practitioners to the educational system.

To remedy the situation in engineering, we have started an Industrial Forum series of meetings with industry and business to discuss gaps in knowledge, skills and attitudes of our graduates.[260] The results of such discussions are evaluated critically and fed back to the programs. However, an implementation of this feedback is done by the educators themselves.

Another attempt has been made to remedy the incompleteness by hiring practicing engineers into the educational programs to provide input on design issues and help in the capstone projects. Such engineers-in-residence (EIR) have been very helpful, but they no longer practice in industry.

In contrast, medicine has a more complete feedback because many of its educators are practicing physicians. Although we see this medical model as more attractive, it may not be implementable in engineering. Consequently, the Body of Knowledge for Practitioners (BoK4P) should be made a vehicle to provide the required feedback to the educational systems. There are more than 30 BoKs that have been reviewed for their disciplines, structure, and scope.[261]

The ultimate closing of the loop is when all the experienced individuals are engage in teaching the next generation of students by providing their experience, many of the soft skills required, and above all the passion and motivation for critical thinking, problem solving, and creativity.

Kinsner has recently described the process of morphing a student into a committed professional.[262,263,264]

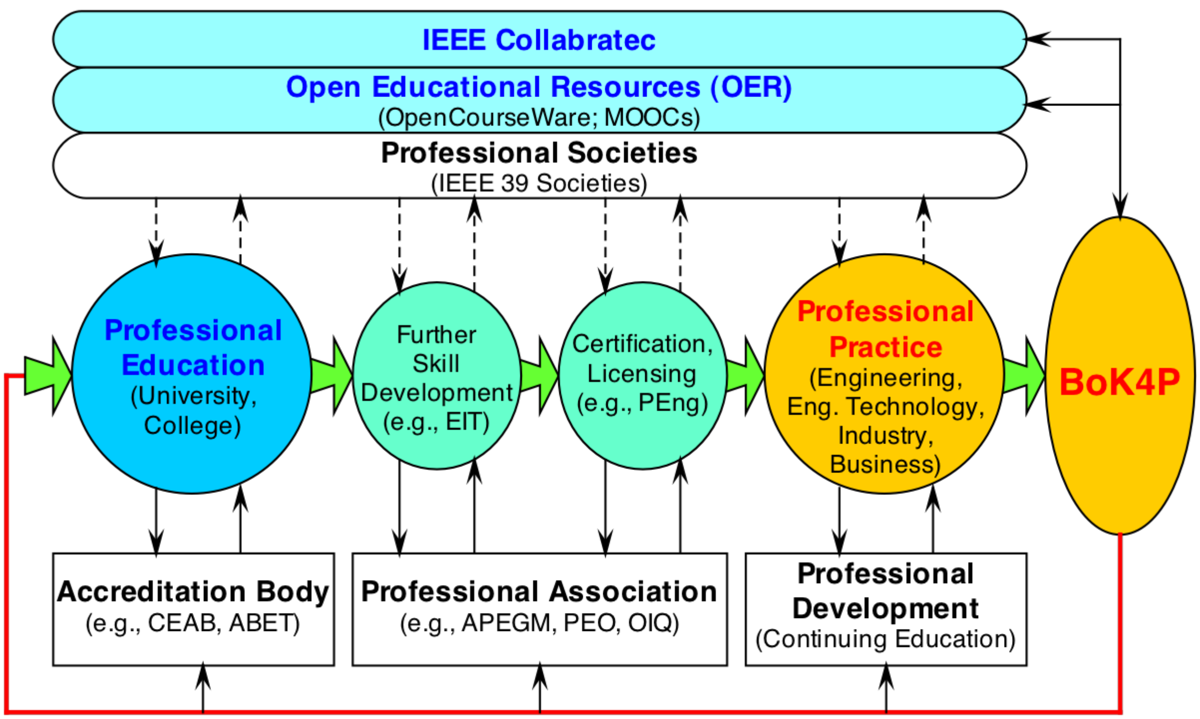

Fig. 13.4. Closing the loop and making the practitioners knowledge creators.

13.4 Towards Symbiotic Education

As we have discussed already, the knowledge tsunami and automation have put much pressure to change the current classroom/workshop model in vocational training. More students and workers learn "just-in-time" and often just enough to solve a problem or get a job completed. Teachers and trainers can no longer be the main sources of knowledge about the world of work but need new forms of technology to help find and manage the increasing amount of information. No single person, no matter how brilliant, can handle the knowledge, even in one field of study.

Consequently, the roles of teachers, trainers and consultants need to change—from mostly presenters of information to guides, curators of knowledge, critical thinkers, and problem solvers. They will have to use digital learning skills and literacies.

Throughout this section, we have been making the case that the next generation of education would benefit much from the development of symbiotic Digital Twins capable of being in relationship with human beings a symbiotic system.

The result of such a symbiotic educational system is illustrated in Fig 8.5. The Digital Twin is depicted by an inverted pyramid. It could penetrate the environment even deeper that the human. It could participate in the data mining and processing to extract more relevant information. It could possibly see more patterns in the information and extract more significant knowledge that could be used in the decision-making process. This layer is now the widest because the Digital Twin would be connected symbiotically to all the relevant other Digital Twins for consultation.

13.5 Closing Remarks on Symbiotic Education

- Symbiotic education has the promise of great impact on how we study, learn, acquire skills, interact with people and machines, discover new things, learn how to operate new things, and how to see reality much deeper.

- Symbiotic education can open up a new landscape for exciting new concepts and research projects.

- We already know how to compete. Symbiotic education might help us how to compete fairly.

- While competition could improve in fairness, we might also learn how to cooperate better.

- To succeed, symbiotic education must use the most sophisticated algorithms available today, and might accelerate development of better algorithm including:
    - Deep learning and machine learning[265]
    - Cognitive systems[266,267]
    - Web intelligence[268]
    - Higher Order (HO) statistical signal processing[269,270]
    - Intelligent signal processing[271]
    - Compressive sensing[272]
    - Fuzzy and granular computing[273]
    - Multiscale (wavelet) analysis[274]
    - Polyscale and fractal analysis[275,276]
    - Long-range-dependence patterns in the data[277]
    - Nonlinear time series analysis[278]
    - Emergent dynamical systems concepts[279,280]

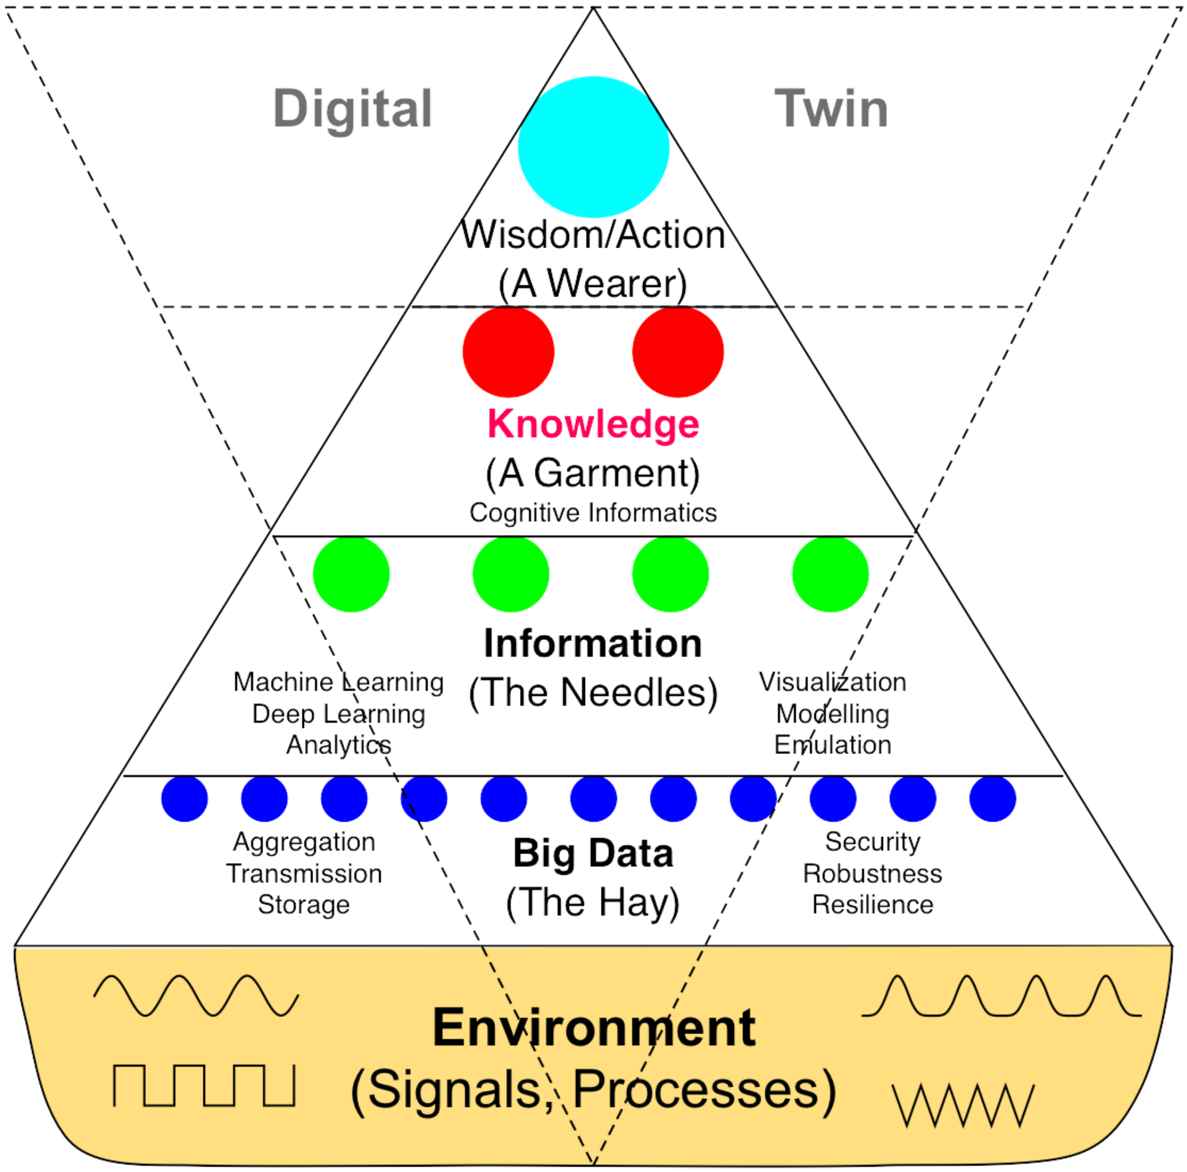

Fig. 13.5. The impact of symbiotic education.

14 Appendix B: Summary of Delphi Study Results

A Delphi Exercise was run to probe a number of people, with different roots in terms of education and experience, in the several areas addressed by the Symbiotic Autonomous Systems Initiative. It was executed in two phases, as usual, first asking answers and comments to the various questions and then submitting the summary of the answers and the various comments for a second round when each one was asked to reconsider the previous answers on the bases of the other people answers and comments.

Here is the final outcome. As any Delphi exercise it does not represent the "truth", nor it pretends to have any statistical bases. However, it can stimulate thinking and steer towards new, or refined, questions.

The areas chosen are represented in the following graphic:

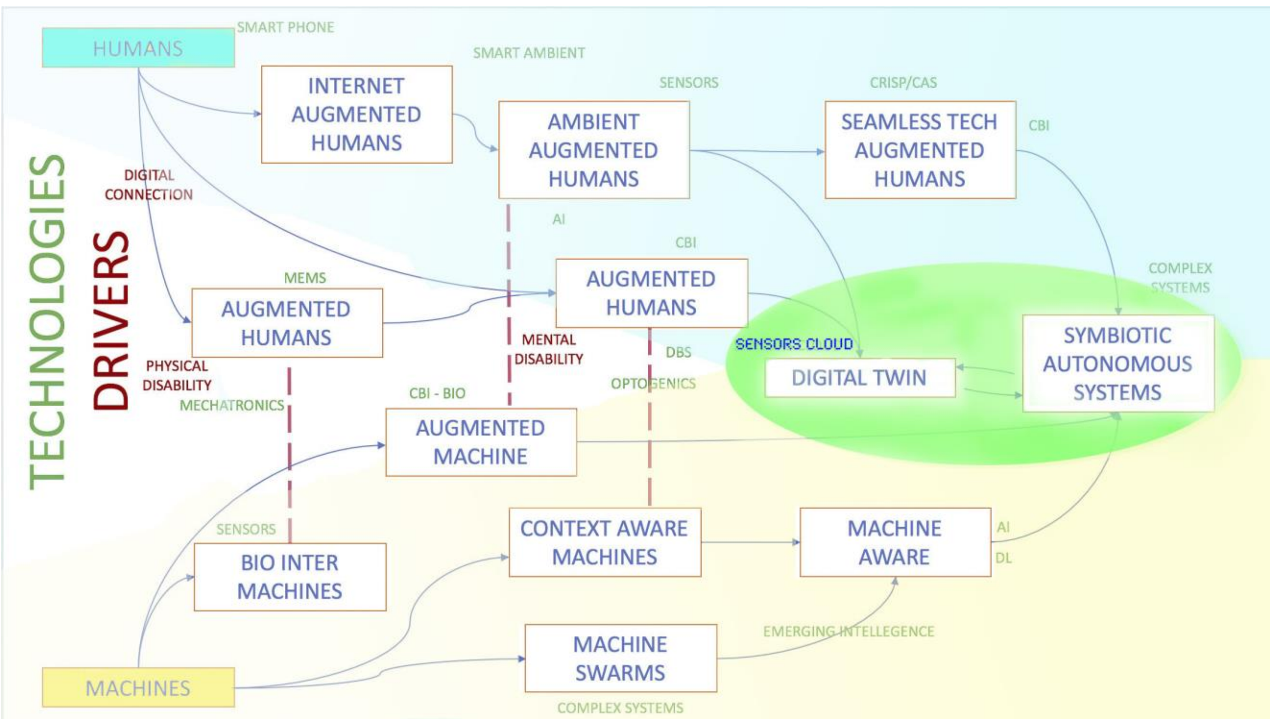

14.1 Area 1 – Internet Human Augmentation

Access to the Internet opens a wealth of information and knowledge that effectively augments humans. As tools for accessing the Internet and retrieving information get more sophisticated and seamless, a symbiosis is created between a human and the information/services on the web, complementing the relationship of humans as being part of a community.

Q1.1
When do you expect the Internet to become and be perceived as common as electricity in 90% of the world?

By far the consensus is on the perception of Internet as a utility, as common as electricity, by 2030. Notice that this is quite significant since the question was related to the whole world (90% of the world population). The implication is that by 2030 the Internet may be a common fabric for access to information, services and possibly the major driver in steering culture. Notice that this question is inserted under the area of "human augmentation" and has been considered in terms of a symbiotic relation with humans. Hence it goes beyond the simple "access possibility" looking at the perceptual symbioses with humans, i.e., its use is seen as seamless and natural, hence no longer perceived as artificial.

However, it does not mean that everywhere in the world the familiarity with the Internet will result in accessing the same information, hence different worlds might coexist in cyberspace. Actually, this seems to be the forecast emerging when considering the answers to other questions: a world potentially unified by a common information access infrastructure that will be segmented at the country level.

Q1.2

Would the remaining 10% be on the edges because of a voluntary choice or because of economic, cultural, and/or political factors?

The overwhelming consensus is that economic and political factors will be the reason for the 10% gap in access to the Internet. Political decisions are clearly affecting the economic affordability and this is most felt by the poorest part of the population. Notice that the availability of cell phones has reached a 65% penetration in 2017, expected to grow at 67% by 2019 with 5 billion cell phones and 4.68 billion users[281]. If out of the 7.6 billion people living on Earth in 2018 some 2 billion are below age 15[282], one can see that the penetration of cell phones is almost 80% for adults already today. By 2030 it is indeed reasonable to get close to 100% of adult population, and by that time the overwhelming majority of phones will be able to access the Internet. In this sense, and this was also a point that emerged from the Delphi, cultural aspects may also play a significant role in the 10% figure of people that will not be "on-line".

Q1.3

What kind of technologies do you expect will provide seamless symbioses between a single human and the Internet?

Almost unanimously the respondent converged on wearables as the means of choice to be "on-line". The alternative proposed, like chip implants, seemed far less attractive and not practicable within the observed horizon (up to 2050). Although technology might provide electronic contact lenses for a seamless "immersion" in cyberspace, the consensus is that culturally we will not be ready for that. However, the smartphone is likely to morph into something less visible, a wearable more similar to a bracelet or a sweater than a device.

Q1.4

Will Augmented Reality (AR) become the "normal" way to perceive the world?

Here again there was an almost unanimous convergence on Augmented Reality becoming an integral part of everyday life, with people seamlessly using it to perceive the world, a world that will consist of both its physical reality and of information (data and relations) provided by cyberspace and filtered by personal bots, most likely through the person digital twin.

Q1.5

Will the perception of the world through AR be biased by some legal and/or political constraint?

Unanimous consensus on the inference of the legal and political world on the way Artificial Reality will be experienced. This clearly raises ethical and societal issues on who will dominate the information provided by the cyberspace and who will be in charge for ensuring the correctness of the information (and relative accountability).

Q1.6
Who is likely to control the "augmented" component of AR?

Following on the previous question, here again, the majority of respondents felt that big platform providers (and several named all or a few of the companies known as FAANG - Facebook, Apple, Amazon, Netflix and Google) will be the one controlling the augmented (cyberspace) component of Augmented Reality.

Q1.7
Will people look for answers or will they look for information to develop the answer themselves?

Here the respondents spilt among those believing that people will be contented with "answers" and those feeling that people will want to have information as well to be able to work out the answer by themselves. There was a slight bias towards the "answer" seeming to indicate a certain laziness on humans. This is clearly increasing the importance of trust on the answers and of accountability. Given the explosion of false news we are already experiencing this is not a good omen. Education is clearly crucial in ensuring people awareness of the merits and dangers created by a seamless presence of the Internet.

Q1.8
Will children born in 2050 imagine a world without AR? That is, by 2050 will AR have become "reality"?

Unanimous consensus on the fact that Augmented Reality will become so pervasive, seamless and "natural" that it will simply be perceived as "the Reality", particularly so by the new generation that will be born from 2040 on, as today tablets and smart phones are an integral part of youngster reality. Not just that. It is most likely that the "augmented" part will be the one that will be most effective in communicating meaning and influencing our perception of the physical reality; in other words, cyberspace is likely to become more relevant than the physical space once seamless augmented reality will be the norm.

This is further increasing the need for appropriate education, for accountability of what is being provided. This will be a major challenge in the future years and it is imperative to start working immediately in tackling this issue. The distortion that can be and is introduced by a pervasive augmented reality is huge, and it is difficult to regulate and control with today's means. Hence there is a need to develop new approaches and ways to make the cyberspace a trusted environment.

14.2 Area 2 - Ambient Augmented Humans

The ambient is likely to become much more interactive, flexible and aware. As such it can morph to fit the needs of the person in that ambient. Part of this morphing, due to smart materials, will occur at the object level, and part, due to embedded AI, will occur at software level. The ambient will evolve a symbiotic relationship with its inhabitants and vice versa.

Q 2.1
When will sensors become a structural part of objects and/or a structural part of materials, rather than being an add-on?

Unanimous convergence on the 2030 timeframe. This is remarkable since it is roughly only ten years into the future. It implies that the design and manufacturing process will have to change significantly by adopting new materials. This is, however, in line with the forecast on the expansion of IoT to reach 1 trillion in the next ten years.

Q 2.2
How will the interaction with an ambient take place in the next 20-30 years?

Most answers pointed at an ambient that is adapting, through its local intelligence, to the human rather than the other way around as it happens today. This implies, as noticed in some answers, that the ambient will be much smarter than today, and part of this "increased intelligence" will derive from devices in the ambient, each one context aware and all behaving in a symbiotic relationship with one another and sometimes with the human. Here symbiosis implies a seamless interaction.

Q 2.3
What is the likely evolution of ambient in sectors such as the home, retail store, school, elderly retirement complexes, office, production plant, and entertainment space?

In general, a gradual evolution is foreseen with business spaces evolving on the basis of expected cost efficiencies and revenue generation (e.g., in the retail area). IoT will become pervasive, and over time they will be leveraged beyond the reason why they were deployed increasing the smartness of the ambient.

The home ambient has a relatively long life time so it cannot be expected to change "massively" over 20 years. Clearly new buildings (and related apartments) will be exploiting new technologies, but economic constraints will be limiting massive deployment. Hence in the home ambient, most evolution will take place in form of "adds on", devices or equipment that are introduced in the home that have a much shorter life time.

Given the increase in elderly population new elderly care complexes will be built in some areas that will benefit from advanced technology but in general technology will be an add on, as in the home, and will be deployed on cost/revenue consideration.

Q 2.4
Will a responsive ambient result from top-down design or will they be the result of self-aggregation of individual components created independently by different players?

As pointed out in the answers to the previous question, most changes will occur in a bottom up fashion, e.g., as result of the introduction of new ambient "components" each responding to a specific need. In few cases there will be ambient designed bottom up, most notably in the retail and manufacturing areas.

In the shorter term, it will be the result of top down design. In the second half of this century autonomous aware systems will start creating their own ambient.

Q 2.5
When will an effective symbiosis between a person and its ambient become feasible, and when will it become normal?

Although some answers pointed towards a gradual shift towards a symbiosis with no definite thresholds indicating that a symbiosis has been reached, a significant number of experts placed at 2030 the time when technology will make symbioses between the ambient and a person feasible, whilse 2040 would see it becoming normal.

Q 2.6
What power will an individual have in the customization and personalization of the ambient? In other words, is it more likely to see an evolution on the side of the individual that will lead to an adaptation of the ambient or an evolution where the awareness of the ambient will lead to its self-customization to specific individuals?

Most answers pointed towards a parallel evolution of the ambient and of the users (humans) although a few pointed out that the transformation of the ambient will be steered by the humans, i.e. it will be human driven. Clearly the evolution of the ambient will be "market driven" and in this sense it will also be human driven, but the mutual adaptation to a new way of interaction will probably involve both, with the ambient making feasible new interaction paradigms and humans, with their behavior, deciding which should become normal.

Q 2.7
There may be three ways of changing the ambient-human interaction: (1) by actually changing the ambient reality (e.g. using smart materials that would change the reality); (2) by changing the interactions between the ambient and the humans (e.g. using software to change visual/aural interactions); (3) by changing the perception of reality (e.g using glasses or interfering with senses or with the brain). Which one will take the lead?

Most answers placed the focus on 3, changing the perception of reality. This goes hand in hand with the trend emerging from the first question of an Augmented Reality taking the lead in future interaction and perception of the reality.

Software will dominate the first phase, probably taking the lion share till the 2040's. Smart materials and self-building will increase in the second part of the century, but software will still dominate, basically for economic reasons (cheaper). Interference with individual perception, by altering the sensation, both through sense hacking and brain hacking will happen in limited areas in the second part of this century. The uptake will be slow, limited to certain areas, possibly seen as a sort of "drug" that needs to be regulated. It will be fraught with ethical and social issues.

Q 2.8
What could be the side effects of establishing a symbiotic relation between an individual and the ambient, considering that the ambient may not be controlled by the individual? Would this evolution lead to a much more controlled society (and controlled individuals) or would the flexibility of having a customized ambient increase personal freedom?

All answers pointed to a very fuzzy scenario in the future where both situations are likely to be present. This was already pointed out in the Augmented Reality interaction, where the weight of cyberspace in human perception is bound to grow significantly and where who controls this cyberspace is not a given.

In those areas where the symbiosis fill human needs (like overcoming disabilities) there will be more personal freedom. In other cases, the symbiosis may lead to steering in a direction that has been decided somewhere else (and by someone else) thus effectively diminishing single human freedom. This has been the case since the birth of human societies. What is new is the effectiveness of the steering that can be achieved.

Privacy and self-control will evolve as well as opportunities to control individuals. There is the need to consider and try to plan for and influence the desirable side effects over the undesirable ones. This is a serious political question. It depends on whether the deep trend of technology is to enslave us or liberate us.

14.3 Area 3 - Augmented Humans

Bio-engineering and smart materials are converging in creating implants that can monitor life functions and expand life functionality, like eye implants first designed to recover sight might eventually provide 10x sight capability and extend human sight in the ultraviolet/infrared range (and beyond). DNA and RNA engineering can expand humans “by design”.

Q 3.1
Plastic surgery has become an accepted practice to modify one’s body. By 2050 will human augmentation be as normal as plastic surgery is today?

The vast majority of experts feels that human augmentation will be considered normal in the 2050 timeframe. However, there will likely be several types of augmentation and more will be surfacing so it is most likely that a few will be considered normal, other unusual and quite a few will be subject of ethical and societal debate on their acceptability.

Certain aspects of human augmentation, like increased resistance to diseases and DNA based fitness (e.g., obesity control), are likely to become normal. Other aspects, like increased sensorial capabilities, may become feasible but scarcely adopted and may be subject to social dislike as happened with Google Glass that in a way provided a form of sight augmentation. The idea that human augmentation is a possibility will become pervasive and will pave the wave to widespread adoption towards the end of the century. Surveys of Millennials already show a willingness to consider implanted devices.

Q 3.2
Will sense augmentation (such as the possibility to use eye implants to detect electromagnetic radiation outside of natural visual spectrum) become normal or will it be relegated to a few niches?

Whilst there seems to be a consensus that sense augmentation will become normal in military applications, its adoption in other areas will probably be on a need-to-have bases. In some niches and in some social classes it may become a distinguishing feature.

Q 3.3
Would augmentation become a professional advantage and as such will people seek augmentation to find better jobs?

In general, there is a consensus that certain types of augmentation will provide a competitive advantage in some jobs, and people will seek them. Clearly this is both going to

- stimulate more people to adopt augmentation to remain competitive, and
- raise legal issues on asking for augmentation in a job description, as today one may be asked to be fluent in a certain language.

This is an area that is bound to generate many labor disputes and will not be easy to regulate.

Q 3.4
Will there be a planet-wide agreement on the use of genetic modification technologies, like CRISPR/Cas9, or different will countries adopt different regulations?

Although we are far from using genetic manipulation for human augmentation, it is within future possibility. The debate today is on the ethical permissibility to tweak with the human genome, and we already see that there is no global agreement, rather a country by country debate and regulation. All the experts agree that it will remain a country by country area of regulation (although, of course, some countries may agree on adopting the same rules, as it is likely the case for the European Union).

Q 3.5
Will DNA engineering become a technology for augmentation, sometimes preferred to implants because it will be considered less invasive?

General consensus is, in the long term, to have the technology and understanding for changing the DNA in ways that lead to human augmentation. In the shorter term, however, DNA modification will be restricted to repair genetic disorders. In the medium term, 2040, the first DNA modification to augment human resistance to adverse factors, like long space travel, is expected. Only in the longer term, like the second half of this century, DNA modification might be considered for general human augmentation. Once experience is gained and trust ensured, DNA modification might be seen as preferable to chip implants.

Q 3.6
Will parents make extra effort to augment their children?

Almost unanimous consensus on parents embracing augmentation for their children.

Q 3.7
Will human augmentation by 2050 be a stepping stone to transhumans (the creation of a new species)?

Almost general consensus that human augmentation will not lead to transhumanism.
2050 is foreseen as too soon to have transhumans, in the sense of the creation of a new species. However, the cultural idea of transhuman as a symbiotic entity seamlessly leveraging technology and having new perceptions and cultures seems to be probable.

14.4 Area 4 - Bio augmented Machines

The use of bio (carbon based and living cells) has been researched for a while with experiments on merging neurons on chips to leverage qualitatively different sorts of computation. The evolution of bio-interfaces will support a variety of interactions potentially opening the way to synergies between machines and living beings, including humans.

Q 4.1
Considering the rapid evolution of machines, will bio-augmented machines in 2050 still make sense? In other words, would a machine still have any advantage by leveraging a biological brain?

The general consensus is that bio-augmentation will provide an edge to machines, even considering the progress of technology. There still seem to be in the observed time frame advantages in coupling bio with machines, particularly in the area of sensing. In the more specific case of leveraging a biological brain this is conditioned on the availability of effective CBI and BCI. Provided technology in those areas will be sufficiently developed, and consensus is lower, there may be benefit in leveraging a biological brain.

Interaction with the biological brain may have a value from a machine point of view in the sense of achieving more effective interaction. The progress of technology and the understanding of brain processes (and of the physical infrastructures supporting such processes) will indeed result in machines that are better than natural brains. What might be good in a natural brain would have been copied and injected into a machine. So in the second half of this century there won't be any advantage for a machine to exploit a natural brain. However, connectivity with brain will still be important because it will provide an advantage to the brain.

Q 4.2
Brain-Computer Interfaces will clearly progress significantly in the next 30 years. Would a brain be able to contribute to the processing of a machine?

Unanimous consensus on the possibility of a brain to contribute to the processing of a machine. As previously indicated there is not generalized consensus that the BCI – CBI will have reach a point sufficient for a real symbioses brain machine.

The understanding of brain processes and of the structures supporting them will provide new insights in syntheses, abstraction, conceptualization, intelligence and free will. These will be used, sometimes mimicked, in a machine. On the other hand, a real time contribution, like having shared processing between brain and machine does not seem realistic, not because it couldn't be done, rather because it will not be effective.

Q 4.3
Would the interaction between brain and machine become so seamless to give rise to symbiotic processing?

Unanimous consensus that in the long term this will be achieved. However, in the medium term the symbiosis will need to be mediated by senses, since a direct brain computer interaction is unlikely to reach the effectiveness required.

A direct seamless connection between the brain and a machine is not yet in sight, although there are several attempts to do so and results have been achieved. The crucial issues are seamless and the extent of the symbiotic processing. Seamless will remain a challenge for many decades. The extent of the symbiotic relationship will be growing over time but it will take several decades before reaching the point of symbiosis that (sometimes) we experience between two persons knowing each other very well. On more limited extent, like a paralytic interacting with an exoskeleton to execute a variety of actions, symbiosis will be achieved in the next decade and will keep expanding.

Q 4.4
In the case of a "brain" participating in a decision process with a machine, what accountability and responsibility issues would emerge?

The general consensus is that accountability will remain on the human side. However, the scenario may get much more complicated considering the variety of human players involved in addition to the human in symbiosis with the machine. The designer, developer, and maintainer of the machine (and related software) will be involved in the sharing of responsibility. The decision point in a symbiotic relationship cannot be tied to a single component as the decision arises from the ensemble. Also, the decision support fragments offered by the various components are likely to be heavily influenced by the other components and the ongoing interactions.

In general, there will be a need to create a new framework of accountability and responsibility for the whole ensemble, similarly to what has been addressed in the past as collective responsibility of a tribe, a nation.

Q 4.5
Machines mimicking life, and more specifically the brain, will become available in the next two decades (one aim of the Human Brain Project is to understand how the brain works in order to

leverage that understanding by replicating it in machines)—and if so, will they be more performant?

The majority of experts does not consider it likely that machines will reach the point of replicating a brain within the next 20 years (although a minority indicated that will be achieved).
There is a more broader consensus on the fact that processing performance au pair with a human brain might be achieved in the next twenty years, and possibly exceeded, however it seems unlikely that such a processing capacity will be possible within the energy budget of a human brain.

A brain in a purely technical sense is not very effective in many areas, although it is amazingly effective from a survivability standpoint as a working solution. Machines will keep increasing their analytical performance (already well beyond the human one) and will be continually copying and refining their synthetic capabilities. A single machine will probably not be structured in a way to take chances but clusters, swarms of machines and for sure symbiotic machines will take chances, and will become better in taking them.

Q 4.6
Life "information processing" may not be more "powerful" than silicon/quantum information processing, but it might remain more "energy efficient" remain more effective (an example is the flight control of flies that is based on some 5,000 neurons whilst a flight control of an airplane requires millions of code instructions). Could this be a reason to continue seeking for a bio-computer integration?

The experts were split almost evenly, part in support of a bio-computer integration for energy budget reasons and part stating that it will not be the main driver. Rather the possibility for a machine to crunch huge amounts of data and for the brain to have a feeling on those data might be the main motivation.

AGI/ASI will narrow the gap, while fully-realized artificial brains could operate as a biological brain at much higher speeds and reliability.

Q 4.7
Would machines that have to interact in symbioses with humans, like a robotic exoskeleton, benefit from a symbiotic cooperation with the brain?

All experts agree that in this kind of situation a symbiotic relation brain-machine would be highly desirable. Notice that in these situations the symbiosis can be restricted to certain aspects hence might be more feasible than a more general symbiosis.

A symbiotic relation involving the brain, like in the example of an exoskeleton supporting a paralytic person, can make the relationship seamless, and this is a great step forward.

Q 4.8
Would aspects like affection, emotion and feelings be better managed by machines interacting with humans rather than working on their own and simulating them?

The majority of experts indicated that the area of emotion in a broad sense is better managed by machines able to interact with humans rather than managed through simulation. However, it has been noted that in practice the capability of humans to manage emotions, and related mental states, is not necessarily adequate in many situations, e.g., under stress, so that a machine can be more predictable in these areas, which in certain situations may yield a better result.

Most humans struggle to identify and respond to their own emotions in a rational manner, so it's hard to see how they could convey useful information to a machine. Equally, as we have already seen with various infamous chatbots, given a sloppy set of human-driven rules, AI doesn't always evolve very well and this should be taken into account.

Humans display a significant slate of emotions, and in general, a machine interacting with humans, being exposed and sensitive to nuances and capable of learning will be better off that a machine modelling emotions.

14.5 Area 5 - Context Aware Machines

The drive towards autonomous systems requires machines to become context aware. Technology (sensors and AI) is supporting increasing levels of awareness. In the coming decades we can expect machines to increase their awareness to levels that may compare to the awareness of living beings, humans included. In certain areas, due to better sensory and processing capabilities, their awareness might even exceed human awareness. Overall, the consensus is on machines reaching a high level of context awareness, in some situations exceeding the one achievable by humans.

Q 5.1
When will machine context awareness match that of an average human?

All experts indicated this goal is beyond 2050. The matching can only happen when the richness of machine sensors becomes comparable to that of human sensors. This will take a very long time. Until then, machine context awareness will not match that of the average human. However, if we better define or qualify what we mean by this, our answers will change. For example, in very constrained situations for humans, machines may excel (for example, human vs. machine with IR/thermal sensors in a dark cave).

Q 5.2
Will machines in the future develop their own context awareness without being pre-programmed to recognize the various components of their environment (e.g., will they autonomously become aware of the difference between a cat and a dog)?

There was unanimous agreement from the experts that this will be achieved. It will require an acceleration in AI that we did not witness in the last decades but that we are starting to experience in these last few years.

Q 5.3
When a machine will be "surprised" by an unexpected context will it autonomously re-evaluate its model of the world (as opposed to today, when they are basically halting operations and transferring control to a human)?

There was unanimous agreement of the experts that this will be achieved. It will depend, obviously, on the kind of surprise the machine experiences. If it is only an alert in the functioning of the system, or an unexpected object or context in the field, there will be need for minor adjustments. In the case of a major disruption, it is probably safer in most situations to let the system inform the relevant actors of the situation and shut down or put itself on hold. The exception to this behavior is perhaps immediate emergency situations (such as, for example, the oft repeated argument of a choice to be made by an autonomous car between several equally undesirable options).

Reactions to surprise events are already being explored via saliency detection and autonomous exploration algorithms that are sensitive to novelty or anomaly detection. What machines do upon being surprised is a matter of risk posture to which they are designed (halting operations, logging the surprise for later human assessment, focusing in on the surprise to collect more data and create new models of it, and so on).

Q 5.4
Will machine context awareness eventually lead to machines changing their goals?

There was unanimous agreement of the experts that this will be achieved. This is clearly raising crucial issues since it may no longer be possible to trust a machine to work and operate within a predefined framework.

This is happening already -- a very simple low-level example being autonomous mobile robots that encounter blockages along planned paths through an environment and have to re-plan new sub-goals in continued pursuit of blocked goals. More sophisticated instances of this may be mere extrapolations of such simple examples in some cases.

Q 5.5
As result of context awareness a machine might alter its behaviour. Will this change of behaviour be considered as a reason for change in the context (i.e., will machines become self-aware of their relationship with the context?

There was unanimous agreement of the experts that this will be achieved. This aspect, similar to the previous one, will raise crucial issues although of a different sort. Here a machine may operate to change its context, including those components in the context that have full autonomy and awareness, thus leading to potential fights and to the attempt to eliminate opposition.

Q 5.6
Following on 5.5, will machine endeavour to change the context as result of their context awareness?

There was unanimous agreement of the experts that this will be achieved. Same considerations as for the previous question.

Q 5.7
Following on 5.5, will machines be able to see context as the result of the interplay of several components rather than seeing the context as a static situation?

There was unanimous agreement of the experts that this will be achieved.

Q 5.8
As machines become better in context awareness, are we going to hand over to them our context awareness? Will AR provide/complement our context awareness leading to a symbiotic relation in the area of context awareness?

The experts split in their view, with a part foreseeing humans handing over context awareness to machines, implicitly trusting them to be better in assessing a situation; others negating this outcome, considering that awareness will not be delegated to a machine. Notice, however, that in limited situations, such as pilots trusting the auto-pilot, this delegation of awareness to a machine already happens. Part of the experts see the future as an extension of what is already happening today, the others put a limit to what can be delegated to a machine.

14.6 Area 6 – Self-Aware Machines

Technology makes possible the creation of smaller and smaller objects that can self coordinate to achieve complex goals. Each component is an autonomous system on its own and is relatively inexpensive to manufacture and deploy. When clustered with many other similar (or exact replica) components, the context awareness along with (flexible or rigid) engagement rules give rise to an emergent behavior, similar to what is seen in insect swarms or bird flocks. The overall swarm is less expensive to create, more resilient to component malfunctions and can generate complex behaviors.

Q 6.1
Will machines ever become self-aware (i.e., will they perceive themselves as an entity)?

There was unanimous agreement of the experts that this will be achieved.

Q 6.2
Will machines ever become aware of why something is happening in an ambient that includes humans (i.e., will machines become aware of deep human intention, not just of their probable behaviour)?

The experts split. Basically half foresee a future where a machine will be able to internalize the deeper human intentions and their possible motivation; the other half does not believe that a machine can feel what a human feels, just evaluate his probable behavior.

The line between probable behavior and human intention is very blurred. Is it possible that machines will understand us better than we understand ourselves most of the time? Absolutely. But that has more to do with the human ability to self-delude than the abilities of the machine. In general, this is part of an ineluctable drive of intelligent technologies. Humans will be in a position to design a new type of self-awareness with varied and measured degrees of psycho-technical autonomy.

Q 6.3
Will machines become aware of other machines' awareness?

There was unanimous agreement from the experts that this will be achieved. Notice that this is the same or similar computational sense in which machines appreciate human awareness (via AR or other means of conveyance or internal representation).

Q 6.4
Will other-aware machines acquire a fully-developed theory of mind (the ability to recognize and attribute mental states—thoughts, perceptions, desires, intentions, feelings, emotions—to oneself and to others, and to understand how these mental states might affect behaviour) equal or superior to that of humans?

There was unanimous agreement of the experts that this will be achieved, however, this will be achieved with limited accuracy as educated guesses. Given workable representations and data, the reasoning and thinking should all be possible to some degree but thought of as useful information for behavioral guidance rather than accurate information, as it is the case for humans.

Q 6.5
What will likely be the roadmap towards full machine awareness (time and quality)?

Most experts foresee this happening beyond the end of this century, hence beyond the horizon of the SAS Initiative.

A gradual growth of machine awareness can be as follows:

- 2020: understanding of ambient (driven by self driving vehicles and robots in industrial environment),
- 2030: understanding of motivation, i.e., support to other entities behavior prediction,
- 2040 understanding of feelings,
- 2050: empathic relations.

Once the right path is found, evolution will be exceptionally rapid:

- sensing
- representations
- shared mental models
- cognitive maps
- autonomous learning
- high-level reasoning

Q 6.6
Will self-aware machines self-create their own goals (e.g., remain healthy, reproduce, interact with other self-aware machines, interact with humans) and will they "cheat" to achieve them?

A majority of experts foresee a time when a machine will be able to, and will actually, create their own goals, while a minority see this as a possibility but within a predefined framework. Clearly, this second view, if true, would create fewer issues in terms of controlling the motivation of machines (in a way machines can evolve still abiding to Asimov's the three laws of robotics).

By 2050, machines will be very efficient at ambient awareness, but still in the early stages of self-awareness and of handling its consequences. Machines would self-create sub-goals only and toward pursuit of goals humans impose on them. Responsible development should avoid enabling self-creation of machine's own goals independent of human goals.

Cheating is a very human notion based on perceived rules and etiquette. It would seem likely that a machine might do something a human would consider cheating purely as a more optimal way to achieve a goal. Whether that, actually, constitutes cheating is a different matter entirely.

Q 6.7
Following on 6.6, will machines experience pain by not reaching their goal and elation by succeeding (in other words, can sensations be used as motivators for machines)?

The majority of experts foresee machines having sensation of pain and elation/joy similar to what humans feel and acting as a consequence. This already the case for some autonomous mobile robots that "feel" frustration when encountering motion limiting cycling during navigation and break out of it with momentary random motions. Such motivations have been tools for some time now in autonomous systems.

A minority argues that such sensations can be programed in a machine, even with great sophistication leading to the appreciation of nuances, and will result in a change of behavior aiming at decreasing pain and increasing elation/joy but this cannot be considered as equivalent to having human feelings. Notice, however, that this is basically the objection of the Chinese Room, raised when discussing the Turing test.

Q 6.8

What will a human–machine relation be once the latter becomes aware?

The majority of experts does not foresee a change in the relationship between human-machines as these become more sophisticated, even when acquiring human-like behavior influenced by human-like feelings. A minority of experts foresee a point when it will be difficult for human not to equate a machine to a living being, with associated empathy and need to have a machine "feel better". This will also generate the instance of machine rights and their protection.

Humans will be likely to consider machines as living entities, although they will remain at a lower level than humans themselves, more like we tend to consider animals (and to a lower level plants). Some machines sharing their life with us will enter into an empathy space and will be seen as pets of a sort.

Q 6.9
Will humans "humanize" self-aware machines?

There was unanimous agreement of the experts that this will be achieved. Some humans will try to "humanize" machines. Note that if the machines develop emotions and self-awareness, doing so will not be at all like reprogramming, but more like trying to change another human and so may fail or have negative consequences

Q 6.10
How might we learn from self-aware machines?

The experts split basically in half; a part foreseeing that when such a point will be reached, humans will apply to machines the same paradigm being applied to humans, hence will try to learn from a machine as we try to learn from other humans. The other half foresee the opening of new paradigms of learning (e.g., establishing a relationship leading to the exploitation of machine capabilities serving the humans without the need to learn from it).

Machines will likely develop their own sense of the world and of the relationships existing in the world. Some of them may come as unexpected to us, and we can learn from them. Interaction with self-aware machines would make us more human by having greater respect for life and self-awareness. In reality, we might learn also many negative traits, such as perfecting the art of bullying a machine to be subservient.

14.7 Area 7 - Machine Swarms

Artificial intelligence and advances in processing, including neuromorphic computing, are opening the door to machine awareness. Machines will understand what is going on, why it is happening and what is the purpose. The question whether a machine that is aware is also perceiving itself as an entity or even more if it has feelings is quite open. The area has two parts: how the machines will transform their behavior through awareness and how we will interact with machines that are aware.

Q 7.1
Nanotechnologies and nanocomponents seem to create swarms naturally. In the next two decades, will we be able to create a science of swarm design that would be able to mimic swarms in Nature?

There was unanimous agreement of the experts that this will be achieved.

Q 7.2
More complex entities, designed to support self-organization, will potentially create swarms. Will swarms be generated in a bottom up way as well (like locusts that in certain conditions give rise to swarms)?

There is unanimous agreement of the experts that this will be achieved. The pervasive presence of Internet of Things and the self-organizing networks, initially introduced in 5G and more in 6G, are starting points for the bottom up creation of swarms.

Q 7.3
Will the existence of multitude of robots lead to the autonomous creation of swarms (although unlike swarms described in 7.2, there will be an explicit design leading to aggregation)?

A majority of experts feel that the growing presence of more and more sophisticated robots that are context aware, will inevitably lead to the emergence of swarm behavior, while a minority does not see this as a possibility, mostly because the numbers will not be sufficient (local density) to generate swarm behavior. It is more likely that the growing awareness will lead to the explicit set-up of cooperation strategies among machines.

It has been noted, however, that with pervasive IoT there will be a push towards design that can generate swarm behavior with respect to certain aspects, as an example, in communications behavior where each IoT will autonomously and implicitly take advantage of other IoT generating a dynamical swarm-like communications fabric. This might be embedded in 6G.

Q 7.4
To what extent can a swarm be influenced as a whole? How can that influence be executed?

The opinions of the experts are spread, almost evenly, on three approaches, although all foresee the possibility of creating swarms that can be influenced:

- influence by design, i.e., each participant in the swarm is coded with the "rules of engagement" that given a sufficient number of participants leads to a swarm formation and once that happens the swarm can be sensitive to external situations leading to influence the swarm behavior.
- by considering the swarm human-like awareness. That would lead to a behavior that like humans can be influenced from ambient changes.

- by setting boundary conditions whereby operations leading to stepping on the boundary condition is impossible and leads to a behavior change in the swarm.

Clearly the third type of influence would result in a much better, a priory control of the swarm behavior and in principle provides a much better control. As downside it limits the evolution of the swarm and its possibility to succeed in the presence of ambient evolution not foreseen at design time.

Q 7.5
What kind of standardization may be needed in the area of swarms?

Experts have split almost in half; one half calling for regulation, and the other for technical standards. Standards will be greatly needed to create the basic communications protocols, in the basic principles of engagement rules, purpose bounding, self-replicating limits and redundant override methods.

Q 7.6
Assuming 7.5, what will be the accountability of self-created swarms?

The majority of experts consider that human accountability will be required, however this might become more and more difficult to enforce, given the variety of relationships involved and the difficulty in assessing, and relating, a specific behavior to a specific human relationhip.

It may be more likely that a continuous process of regulation and technical standard tuning will be required as unexpected, undesirable effects surface. A new science, approach, or research regarding regulation and technical standards may be required to stimulate their effect. The collective responsibility will have to be tied to the engagement rules. However, it will be difficult to pinpoint accountability to a single entity. On the other hand, malicious hacking will become an issue, since they will attempt to change the rules of engagement to change the swarm behavior.

Assuming any need for standards for swarms is justified, accountability would have to be handled and contained via the standard. Unconstrained, swarms may have the property of emergent behavior, which complicates accountability. It would seem that a standard would limit or eliminate the potential for emergence, which seems counter-intuitive, as emergence in swarms as a desirable property and not one to be controlled.

Q 7.7
Swarms will create a dynamic context in which humans (and other lifeforms) may be present. Should we expect swarms to become symbiotic with lifeforms—and in particular, with humans?

The majority of experts foresee some sort of symbiosis possible between a swarm and a human. A minority do not consider this possible. Clearly there are new issues to consider when dealing with the symbioses, like:
- would a symbioses create a unique aggregation or could be a multi-party symbiosis, i.e., one swarm entering into a symbiotic process with several humans at the same time?
- could a symbiosis become stable or would it dynamically evolve as boundary conditions change?

Q 7.8
Will military applications lead the evolution of swarms, or could other areas, such as manufacturing, surveillance or healthcare, take the lead in the coming two decades?

The majority of experts see military leading the evolution of swarms, however a few experts point out that in the 2040 timeframe the healthcare sector may take the lead, or at least become a very important player in swarm evolution.

In general, it is the military that spearheads innovation because it is better funded. Next is medicine, for the same reason, but at a lower budgetary level. Education, unfortunately will continue to lag, contrary to ever more assertions, but what is really interesting is that money may not be critically dominant anymore, as more and more innovations will spring out of nowhere because of a much better distribution of knowledge and tools.

14.8 Area 8 - Digital Twins

Digital Twins, the replica in bits of an entity (also referred to as representations of entities in digital form) - as minute as a switch, as a turbine, or as a cluster of many Digital Twins like a city - are an expanding reality in many areas: manufacturing, operations, and planning and are now starting in education and health care.

Digital Twins are becoming more than a replica; they are actually becoming so extended (as an example recording the whole history of an entity) that there is the problem of separating the exact replica from the other data. Besides, Digital Twins can have a life of their own (avatar) augmenting the capabilities of their real twin.

Q 8.1
Internet of Things (IoT) are mirrored in cyberspace, each one having its own digital twin. Will their aggregation create digital twins in the next decade that will be open to stimulate third parties services?

There was unanimous agreement of the experts that this will be the case.

Q 8.2
Is digital twin standardization required to foster their use by third parties?

There was unanimous agreement of the expert that standardization will not be required in this area.

Q 8.3
Will digital twin economic value exceed the value of their real counterpart (i.e., will we see new business models based on the offering at nominal cost of the real twin as the hook to create value at the digital level)?

A majority of experts foresee a growth of value of the Digital Twins that within the IEEE SAS Initiative timeframe exceeds the value of their physical counterpart, in line with the current trend of data increasing its value to exceed the value of the objects themselves (e.g., more value resides in the song’s consumer data that in charging for playing a song). A minority does not foresee this happening. In other words, a majority applies to Digital Twins the rules of the Data Economy, whilst a minority consider Digital Twins as real objects bound to the rules of standard economy.

The current trend in replacing human workers with automation continues and embraces Digital Twin technology.

Q 8.4
Following on 8.3, will digital twins mark the transition from an economy of products to an economy of intelligent services?

There was unanimous agreement of the experts that this will be the case. However, it does not follow that Digital Twins can be considered as the prime cause of this shift, more likely a concause or a side effect.

Q 8.5
Digital twins are, and will go, beyond being the digital replica of an entity to become a multidimensional representation of that entity, in space and in time. Will this multidimensionality be subject to different ownership?

The majority of experts foresee a growth in the representation or mirroring power of Digital Twins and their extension beyond this to become, partly, independent and complementary entities with respect to their physical twin. A minority does not see this evolution.

A first step is to integrate all data that concerns a given individual and to give that person access to the developing Digital Twin. The second step is to establish interfaces and applications that allow a mutual learning (deep) between the owner/user and the Digital Twin. The third - simultaneous - development will be an economy of Digital Twin management and exchange.

Q 8.6
Given that today we already have profiling, will humans have their digital twin in the next decade?

A slight majority of experts foresee the evolution towards every person having his own Digital Twin as inevitable, starting in the next decade, as a linear evolution of profiling. Notice that while profiling is usually set up by a third party with or without our awareness (and agreement) a Digital Twin stems from an explicit intention of the person that remains in control of his Digital Twin. So, saying the Digital Twins for humans represent an evolution of profiling may be misleading.

This is already the case with hidden profiling and its use by business and government. The symbiotic aspects of such uses, called the Digital Unconscious, is a new feature of our daily life that is hidden from view and only manifest by its effects on our behavior and buying habits.
A minority of experts do not see this as happening during the IEEE SAS Initiative timeframe.

Q 8.7
What is the likely roadmap to a fully developed human digital twin?

Most experts consider 2040 as the most likely timeframe for massive adoption of human digital twins, although a few indicate this happening in the 2030 timeframe.

A possible roadmap:

- 2020/30: Digital Twin for Health Care and as access to services in the retail and entertainment
- 2030: Digital Twin as a knowledge twin used in education and GIG economy
- 2030/2040: Digital Twin embedding the genotype to simulate the phenotype
- 2050: IP on Digital Twin and use of simulation to create teamwork, selling of Digital Twin access

Q 8.8
Will symbioses take place at the level of digital twins and then have effect on the real human twin(s)?

Most experts do not see symbioses happening at the Digital Twin level, although a few consider Digital Twin as a crucial component in a symbiosis. Humans are always impacted by their technology. It is unlikely that people will lose track of who they are and who is the twin, unless the twin achieves consciousness then it will likely no longer be a twin.

Q 8.9
How will we use digital twins to learn according an optimal progression?

Most experts consider Digital Twins as useful tools for forecasting; a few see them also playing a role in monitoring both education accrual and the fading away of knowledge.
- 2020: self growing cv,
- 2030: Digital Twin prompting education and used to customize education,
- 2040: shared education with the extended Digital Twin.

Q 8.10
How will digital twins alter the learning process itself?

Most experts consider Digital Twins playing a role in cognitive augmentation, i.e., the seamless use of our Digital Twin to accrue knowledge and make it available on demand. The Digital Twin will maintain our knowledge gaps with respects to needs, will observe what we forget and step in to reinforce learning/refreshing, and will customize the learning process accommodating various needs and opportunities.

A Digital Twin will be the equivalent of a turbo-charged Fitbit, i.e., semi-customized and customized recommendations based on the data generated and reinforcement for certain behaviors and behavioral changes. A Digital Twin participating in an experience and then teaching the person what it had learned could be a new way to learn.

Q 8.11
How will digital twins enhance not only just-in-time skills, but also the fundamental understanding of Nature?

There was unanimous agreement of the experts that Digital Twins will be a useful tool through the ability to perceive levels of scale and wavelength spectra not perceptible by humans and augmented by rapid processing of complex mathematical analyses beyond humans (with the exception of savants and individuals with 180-220 IQs).

A digital avatar could engage in activities and environments that would be dangerous for physical people. They could provide a more absolute form of memory that isn’t subject to the human brain’s ability to be manipulated (intentionally or otherwise).

Q 8.12
Will digital twins enhance human creativity? If so, in what ways?

There was unanimous agreement of the experts that Digital Twins will enhance human creativity, mostly by connecting to the limitless cyberspace where everything, from a creativity point of view, is possible, not being constrained by physical rules.

14.9 Area 9 - Symbiotic Autonomous Systems

The symbioses of life forms and artefacts, human and machines, looks like the inevitable evolution, made possible by technology advances, by economic drive and by the increase of wellbeing through augmentation. The separation between ourselves and our artefacts is getting fuzzier as artefacts achieve our level of intelligence and become an integral, seamless, part of our life, in some instances entering our own body.

Q 9.1
When will human augmentation move from making up for disabilities to augment personal capabilities?

The experts forecast is evenly spread from 2030 to 2050 and beyond.

Q 9.2
What are the areas where augmentation is most likely to occur: e.g., business (surgery, manufacturing, design), sports (increased performances, resilience), military (heightened senses, resilience, sharper mind, speed of responses), education (laboratories, experiential learning), and creativity (design options, design implementations, design cost)?

Most experts foresee augmentation happening in all areas, first in niches within each area and then progressively expanding. A few are pointing out that the military area will have and keep the lead in human augmentation, in its various forms.

- 2020: military, sport and business
- 2030: extends to sports and education
- 2040: extends to creativity

Q 9.3
What will be the social acceptance of augmented humans?

Although a few experts foresee some sort of societal rejection or at least suspicion with regards to human augmentation, most point out that there will be changes in societal feelings eventually leading towards accepting it, and in some cases even encouraging it. A parallel may be seen considering vaccination, a sort of human augmentation providing increased resistance to microbes. This has moved through the stages of suspicion, to become accepted and even enforced, although a few are still looking at it with suspicions.

Human augmentation will take different forms and will result in different human changes. Each of these is likely to go through a process of rejection, suspicion, limited acceptance, mass acceptance, and enforcement. Of course, not all augmentation will be following this evolution with quite a few only moving through the first step(s).

Human augmentation will likely be perceived as a form of cheating in sports and unfair competition in business but it will be accepted in health care related areas. Perception will change over time as augmentation will become more common. Visible augmentation through a device will likely be considered more acceptable than invisible one (like DNA engineering). Augmentation acceptance most likely to start in tech hubs, but possibly not traditional tech hubs like Silicon Valley. It can be expected to gain more traction in less stringently regulated environments first e.g., parts of Asia.

Q 9.4
Will augmented humans become an elite, a niche, or a norm?

The majority of experts foresee human augmentation, within the IEEE SAS Initiative timeframe, as applying to an elite; only a minority of experts foresee human augmentation as becoming the norm. Here the point is not to discuss single aspects of human augmentation, rather the overall concept.

The evolution will be first niches, then elite, then norm. However, this will happen in several areas at different times so in a particular area augmentation may be restricted to a niche, in another it might have progressed to adoption by an elite and so on.

Q 9.5
Are machines likely to benefit from symbioses with humans or the benefit is only on the human side?

A slight majority foresees benefits for both humans and machines, while a minority consider human augmentation as beneficial to humans only. In general, it is felt that most benefit will be on the human side, and in the long term it will only be on the human side.

Q 9.6
What economic issues are likely to result from symbioses of humans and machines?

About half of the experts acknowledge that there will be significant consequences from an economic point of view (including obviously impacts on jobs), and that most of these consequences are not clear at all today. The other half is pointing at increased productivity in most areas as the major economic impact, while a few identify the issue of increasing inequality due to the uneven adoption of augmentation.

Different countries might experience different rate of adoption (and quality of adoption) and this may lead to both unfair competition and societal issues when an augmentation will start to be required or even forced.
Symbioses will not be uniform and will provide an advantage creating a gap between those who can have it and those who don’t. Disparity across countries may lead to significant economic issues. The pervasiveness of robots in Industry 4.0 and their growing symbioses with humans will be an element.

Q 9.7
Will augmentation become a basic human right?

The majority of experts does not foresee human augmentation as a human right, although a small minority foresees that in the long run, 2050, some forms of augmentation may be considered as part of the human rights (as vaccination is considered in some countries as a human right).

Q 9.8
What ethical issues are foreseeable in Symbiotic Autonomous Systems?

Although the general consensus is that several ethical issues are already clear, the experts foresee that most ethical issues related to human augmentation are still to be discovered and most of them will be specific to a specific form of augmentation, rather than general.
Some emerging ethical issues:

- Further differentiation between poor and rich
- Political misuse
- Criminal misuse
- What happens to the symbiotic system when the host dies? What if the system is self-aware?

- When is it okay to modify someone without their consent?
- If symbiotic systems become commonplace, do parents have the right to add one to their child?
- What about to a grandparent with dementia or other cognitive impairment?
- Will machines overtake humanity or individuals
- Who is responsible for machine errors?
- SAS availability based on ability to pay
- Manufacturing errors or intentional shortcuts
- Neural hacking

Q 9.9
AI is already replacing paralegals in several US and UK law firms. Will this trend increase to the point where AI supplants lawyers and legal decisions affecting humans, Digital Twins and Symbiotic Autonomous Systems be made by AI rather than humans?

The experts split with a slight majority foreseeing AI to take over; while a slight minority foresee humans in control during the IEEE SAS Initiative timeframe.

Q 9.10
Following on 9.9, if this trend does take place, would it be beneficial or detrimental?

A slight minority, roughly corresponding to those that foresee humans in control, points to mixed advantages and disadvantages, whilst a slight minority consider this evolution as beneficial.

Issues:

- Beneficial for law firms; detrimental for lawyers
- Legal decisions by AGI may or may not be similar to human decisions, depending on its ability to incorporate human qualities such as theory of mind, native understanding of human behavior, and context
- Many lawyers are subject to their emotions and forgetful of prior experience. Some are also lazy about researching existing material. A system that automatically knows all available data, optimizes for clearly defined criteria and charges less than $500 an hour would be excellent. (Although it could charge $1,000/hr since you would only need a few minutes of its time).
- Bad for those displaced, but providing new value to consumers of services.
- It will be detrimental as the limitations are encountered (due to the impacts they have), and it will be beneficial once used within the understood limits.

15 Appendix C: Examples of Recent Responses to IoT Vulnerabilities

This section describes how industry is responding to the rapid developments in the area of Internet of Things (IoT) connectivity and some data protection.

15.1 IoT Wireless Standards and Implementations

Connecting various IoT devices to a local network or the Internet requires wired and wireless technologies that were developed long before the IoT concept was on the table. The wireless technologies must range from the device level, to WiFi, to Bluetooth low energy (BLE), to cellular and other networking. The device connectivity and many data transformations belong to what is termed "edge computing" or "fog computing."

15.1.1 IoT LPWAN

The LoRa Alliance announced an implementation of a very low-power wide-area network (LPWAN) protocol that allows wireless connection of battery-operated "things" to the Internet at the regional, national, or global network levels[283]. It has bidirectional data communications, and end-to-end security, mobility and localization services. It uses the star-of-star topology in which gateways relay messages between end devices and the central network server. The gateway acts as a transponder bridge to convert radio-frequency (RF) packets to the Internet Protocol (IP) packets. It has a Firmware Over-The-Air (FOTA) either to update the connected devices, or to distribute messages.

Cypress PSoC 6MCU, Semtech, ARM Cortex-M 2MCU and others are LoRa secure devices.

15.1.2 IoT WiFi (IEEE 802.11)

If power is not critical, IoT connectivity can use the well-established infrastructure (hubs and routers).

In 2018, NXP Semiconductor developed an IoT-on-a-Chip edge-computing device based on the ARM iMX 6ULL application processor, together with WiFi (dual-band 802.11ac) and Bluetooth (4.2)[284]. It has a small footprint (14x14x2.4 mm), is scalable, and easy to design with. It will be followed by iMX7 and iMX8 in 2019.

The IoT-on-a-Chip is suitable for secure bots, tamper detection and response, and high-throughput cryptography. It can also be expanded with a custom inter-chip interface.

15.1.3 IoT ZigBee (IEEE 802.15.4)

For applications that do not require high-speed data transfers but do require low power, self-organizing networked devices for industrial environments often use the ZigBee protocol.

STMicroelectronics announced a dual-processor wireless chip, STM32WB System-on-Chip (SoC) to support the ZigBee protocol[285]. It is based on the ARM Cortex-M4 microcontroller to run the main application. Another ARM Cortex-M0+ core is used to offload the main processor. It runs the Bluetooth Low Energy (BLE) 5 and a ZigBee radio capable of running not only the ZigBee but also OpenThread and other protocols. Another unusual feature is the internal balun included on the SoC to connect to the antenna directly (it saves 6 external elements). The 2.4 GHz low-power radio requires 5.5/3.8 mA to transmit/receive.

15.1.4 IoT Bluetooth Low Energy (BLE) (IEEE 802.15.4)

Applications in health care, fitness, security, and home entertainment often require low-power IoT edge connectivity through the Bluetooth Low Energy (BLE) wireless personal area network (PAN). The BLE is good for devices that must be running a year without recharging. The range is adjustable from 10 to 200 feet.

Texas Instruments announced several BLE SimpleLink devices[286]. The CC2642R MCU is for Bluetooth 4 and 5 applications in the 2.4 and Sub-1 GHz devices with sleep current below 1 μA. It has a 48-MHz ARM Cortex-M4F CPU, with a dedicated ARM Cortex-M0 for radio control, as a well as an autonomous ultra-low-power Sensor Controller for digital/analog sensors and data acquisition. Its transceiver can handle many protocols (WiFi, BLE, ZigBee, Thread, Sub-1-GHz, and Ethernet.

15.1.5 IoT Cellular Networking

Many IoT require cellular connectivity, through GSM, CDMA, LTE in the 2G, 3G, 4G, and 5G cell environments.

Nortfic Semiconductor announced their nRF91 to work with the Verizon Wireless Network (USA) and Telia (Norway)[287]. The nRF91 is ultracompact, low power, small footprint (10x161.2 mm), multimode LTE-M/NB-IoT System-in-Package (SiP). It uses the ARM Cortex-M33 processor and the ARM TrustZonne security processor and the Assisted GPS. It has a modem, SAW-less Transceiver, and Qorro RF front end. The nRF91 system is intended not only for the current smartphones, but also to many other mobile devices.

It includes security for the application hardware and software through the ARM Cortex-M33 and ARM CryptoCell-310. The TrustZone for ARMv8-M secures data, firmware and peripherals.

15.1.6 IoT Near-Field Communications (NFC)

The near-field communications (NFC) technology connects a portable device with a contactless terminal. A connection is established when the portable device is brought close to the contactless terminal. The protocol used is Bluetooth and other protocols like FeliCa.

The standards used are ISO/IEC 14443 (106 kb/s) for ID cards, and ISO/IEC 18000-3 for RFID devices[288].

STMicroelectronics announced the ST25NF tags such as the type 4 Tag IC (ST25TA) and type 5 Dynamic Tags IC (ST25DV) and type 5 Tag IC (ST25TV) that are now certified[289].

There are many other edge connectivity schemes that all involve security by design.

15.2 Examples of Research in Security

Since cyber-attacks are very sophisticated, simple techniques for their detection are not applicable. Considerable research is being conducted around the world to develop good techniques to detect anomalies on the connected networks through transmitted packets and other means in order to develop effective countermeasures.

Any captured data from a network under attack is most likely non-stationary. Signal analysis is required to segment the signal into appropriate frames first, then to extract the most significant features from the frames, followed by assembly of the features into vectors, feeding of the vectors into a classifier, classification of the network's behavior based on the features, and finally finding appropriate countermeasures to protect the network and devices attached to it.

Data acquisition and segmentation is the prerequisite step for feature analysis. Feature extraction from the data constitutes the first major challenge and is often done using standard temporal and/or spectral mono-scale analysis. There is a movement to improve feature extraction through multi-scale and poly-scale signal processing techniques.[290] In addition to the traditional energy based metrics, we must also use entropy-based metrics capable of extracting information-related patterns in the data.[291,292] The metrics must also be suitable for bursty traffic that is characterized by a long-range dependence, with heavy-tail probability-mass-function distributions.

The second major element in the technique development is to identify an onset of the attack. The third is to classify the attacks and launch appropriate countermeasures. The detection and classification stages involve machine learning (ML) and deep learning (DL)[293] techniques.

A survey of ML and DL methods was conducted for network analysis of intrusion detection was presented, together with a description of the commonly used network datasets[294].

Examples of anomaly detection techniques in distributed denial of service (DDoS) are provided in [295,296].

References

[1] https://symbiotic-autonomous-systems.ieee.org/images/files/pdf/sas-white-paper-final-nov12-2017.pdf

[2] https://www.ossur.com/about-ossur/news-from-ossur/2-uncategorised/69-the-technology-behind-the-proprio-foot-from-ossur

[3] https://www.rethinkrobotics.com/baxter/

[4] https://intelligence.org/2013/08/11/what-is-agi/

[5] https://www.fanaticalfuturist.com/2017/03/bad-news-for-jobs-fabled-artificial-general-intelligence-could-arrive-much-earlier-than-expected/

[6] https://www.humanbrainproject.eu/en/

[7] http://www.braininitiative.org

[8] https://www.humanconnectome.org

[9] https://www.animalequality.net/node/703

[10] http://www.secondsight.com/g-the-argus-ii-prosthesis-system-pf-en.html

[11] https://brain.ieee.org/

[12] https://www.neuralink.com

[13] https://govinsider.asia/digital-gov/meet-virtual-singapore-citys-3d-digital-twin/

[14] https://neurosciencenews.com/brain-internet-connection-7489/

[15] Tae-Gon Noh, Counterfactual Quantum Cryptography, *Phys. Rev. Lett.* 103, 230501 (2009)

[16] Qi Guo, Liu-Yong Cheng, Li Chen, Hong-Fu Wang, and Shou Zhang, Counterfactual entanglement distribution without transmitting any particles, *Opt. Express* 22, 8970-8984 (2014)

[17] Qi Guo, Liu-Yong Cheng, Li Chen, Hong-Fu Wang & Shou Zhang, Counterfactual quantum-information transfer without transmitting any physical particles, *Scientific Reports* volume 5, Article number: 8416 (2015)

[18] Sheng Zhang, Probabilistic direct counterfactual quantum communication, *Chinese Phys. B* 26 020304 (2017)

[19] Yuan Cao, et al, Direct counterfactual communication via QZE, *Proceedings of the National Academy of Sciences* May 2017, 114 (19) 4920-4924

[20] https://www.nature.com/articles/nrn.2016.22

[21] R. W. Picard. Affective Computing. M.I.T Media Laboratory Perceptual Computing Section Technical Report No. 321. (PDF) (http://fliphtml5.com/equo/nzjx/basic)

[22] https://www.media.mit.edu/groups/affective-computing/overview/

[23] https://www.braininitiative.nih.gov/

[24] https://academic.oup.com/nsr/article/4/2/258/3052681

[25] http://brainminds.jp/en/

[26] http://www.humanconnectomeproject.org

[27] https://www.humanbrainproject.eu/

[28] https://www.wits.ac.za/news/latest-news/research-news/2017/2017-09/can-you-read-my-mind.html

[29] http://neurosky.com/2018/04/bci-elon-musk-and-the-future-of-human-enhancement/

[30] http://www.wormatlas.org/neurons/Individual%20Neurons/Neuronframeset.html

[31] Neural Engineering System Design, https://www.darpa.mil/program/neural-engineering-system-design

[32] https://rossdawson.com/blog/periodic-table-disruptive-technologies-innovation/

[33] Trials on mice performed at the Wake Forest Baptist Medical Centre, http://iopscience.iop.org/article/10.1088/1741-2552/aaaed7/meta

[34] https://www.wired.com/story/hippocampal-neural-prosthetic/

[35] https://www.weforum.org/agenda/2018/01/how-to-manage-ais-risks-and-benefits

[36] http://www.sciencemag.org/news/2018/01/mind-reading-algorithm-can-decode-pictures-your-head

[37] https://www.designx.mit.edu/cohorts/Dormio

[38] DeepDream, https://deepdreamgenerator.com/

[39] http://2045.com

[40] https://replika.ai

[41] https://kernel.co

[42] http://news.berkeley.edu/2017/07/13/21-6-million-funding-from-darpa-to-build-window-into-the-brain/

[43] https://secondlife.com

[44] https://avatar.xprize.org/prizes/avatar

[45] http://www.brain-iot.eu/brain-iot/

[46] http://journals.sagepub.com/doi/pdf/10.2466/02.09.28.CP.1.5

[47] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5800851/

[48] https://www.cell.com/cell/fulltext/S0092-8674(18)30592-0

[49] https://www.ncbi.nlm.nih.gov/pubmed/1363084

[50] https://phys.org/news/2018-03-crispr-cas9-rna.html

[51] https://investor.sangamo.com/press-releases/detail/404/sangamo-announces-8-million-cirm-grant-for-st-400----a

[52] http://www.fertility-docs.com

[53] https://humanlimbregeneration.com/super-healing-abilities-and-human-limb-regeneration-research

[54] https://www.independent.co.uk/news/science/drought-resistant-crops-plants-genetic-engineering-a8091816.html

[55] https://futurism.com/humanitys-next-stage-evolution-cyborg/

[56] https://www.wired.com/story/hippocampal-neural-prosthetic/

[57] https://phys.org/news/2011-09-carbon-nanotube-composites-enzymes-cosmetics.html

[58] https://www.engadget.com/2016/08/13/duoskin-high-tech-flash-tattoos/

[59] https://www.theverge.com/2018/4/6/17205718/plastic-surgery-apps-why-cosmetic-surgeons-insecurity

[60] https://www.ft.com/content/67e3ab88-f56f-11e7-a4c9-bbdefa4f210b

[61] https://gizmodo.com/darpa-s-brain-chips-could-be-the-next-big-mental-health-1791549701

[62] https://www.forbes.com/sites/daviddisalvo/2014/01/26/the-era-of-genetically-altered-humans-could-begin-this-year/#790144e6d9df

[63] https://www.slideshare.net/patriciafrancislyon/predicting-phenotype-from-genotype-with-machine-learning

[64] Methuselah Foundation, https://www.mfoundation.org/

[65] http://www.sens.org/

[66] https://www.genome.gov/25020028/cloning-fact-sheet/

[67] https://www.iflscience.com/health-and-medicine/two-father-babies-could-soon-be-possible-no-egg-donor-required/

[68] http://koniku.com

[69] https://globaledge.msu.edu/global-resources/resource/10557

[70] https://www.tandfonline.com/doi/abs/10.1207/S15327051HCI16234_17?journalCode=hhci20

[71] https://www.researchgate.net/publication/200048737_Context-awareness_and_artificial_intelligence

[72] ISBN-13: 978-1101946596

[73] https://ieeexplore.ieee.org/document/8317766/

[74] https://link.springer.com/chapter/10.1007%2F978-3-319-14364-4_50

[75] http://engineering.usu.edu/files/pdfs/us-9852507.pdf

[76] https://www.tandfonline.com/doi/abs/10.1080/01691864.2018.1460619

[77] http://bitcoinist.com/blockchain-makes-sentiment-analysis-made-affordable/

[78] https://link.springer.com/article/10.1007/s10664-017-9546-9

[79] https://deepmind.com/research/alphago/

[80] https://home.cern/about/computing/processing-what-record

[81] http://www.historyofinformation.com/expanded.php?id=2877

[82] https://www.darpa.mil/program/restoring-active-memory

[83] http://www.prohealth.com/library/showarticle.cfm?libid=29525

[84] https://www.gartner.com/doc/3772081/hype-cycle-data-science-machine

[85] https://futurism.com/self-replicating-dna-computers-are-set-to-change-everything/

[86] https://foresight.org/self-replicating-robots-and-degrees-of-self-replication/

[87] https://workingmouse.com.au/innovation/4-things-you-absolutely-need-to-know-about-software-bots

[88] https://www.law.berkeley.edu/files/Holman_Christopher.pdf

[89] https://www.nasa.gov/directorates/spacetech/niac/2018_Phase_I_Phase_II/Marsbee_Swarm_of_Flapping_Wing_Flyers_for_Enhanced_Mars_Exploration

[90] https://link.springer.com/chapter/10.1007/978-3-319-73323-4_13

[91] https://newatlas.com/aquabotix-swarmdiver-swarm-robots/54120/

[92] https://medtechboston.medstro.com/blog/2015/10/06/swarm-robotics-what-you-need-to-know-about-the-future-of-medicine/

[93]http://www.prohealth.com/library/showarticle.cfm?libid=29525https://www.army.mil/article/203649/robots_to_extend_reach_of_battlefield_control_says_army_robotics_expert

[94] https://www.qinetiq.com/Insights/2018/04/A-fine-balance

[95] https://arxiv.org/ftp/arxiv/papers/1803/1803.11256.pdf

[96] https://www.zdnet.com/article/the-next-world-war-will-be-a-cyberwar-first-and-a-shooting-war-a-distant-second/

[97] http://deepknowledgeventures.com/

[98] https://www.tieto.com/

[99] https://trinixcreative.com/2016/01/complex-systems-the-new-web/

[100] http://packetpushers.net/hot-cold-mash-potato-routing-and-route-reflector-design-considerations/

[101] http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0122827

[102] http://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1002894

[103] https://www.gartner.com/doc/3867164/top--strategic-technology-trends

[104] https://mplsystems.co.uk/iot-and-digital-twins-to-optimise-service/

[105] https://gereports.ca/digital-twin/

[106] https://replika.ai

[107] http://enb.iisd.org/vol12/enb12459e.html

[108] https://www.ipcc.ch/pdf/assessment-report/ar5/syr/AR5_SYR_FINAL_SPM.pdf

[109] IEEE, *Spectrum Magazine*, June 2018. Available Aug 15, 2018 from IEEE at https://spectrum.ieee.org/magazine/2018/June

[110] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. Dependable and Secure Computing*, vol. 1, no. 1, pp. 11-33, Jan-Mar 2004. (doi: 10.1109/TDSC.2004.2). https://ieeexplore.ieee.org/abstract/document/1335465/

[111] Francesco Flammini (ed.), *Resilience of cyber-physical systems: From risk modeling to threat counteraction*. Series: Advanced Sciences and Technologies For Security Applications. Berlin-Heidelberg: Springer, 2018 (ISBN 978-3-319-95597-1, eBook; 978-3-319-95596-4, hbk). https://www.springer.com/us/book/9783319955964; http://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1216903&dswid=6024

[112] International Standards Organization (2009). ISO/IEC 15408-1:2009 - Information technology -- Security techniques -- Evaluation criteria for IT security -- Part 1: Introduction and general model. [online] Iso.org. Available at: https://www.iso.org/standard/50341.html [Accessed 24 Sep. 2018]

[113] JTFSCSE, *Joint Task Force on Cybersecurity Education (CSE).* 2015. Available Aug 15, 2018 from IEEE at https://www.csec2017.org/

[114] https://en.wikipedia.org/wiki/Operation_Aurora

[115] https://cloud.google.com/beyondcorp/

[116] NIST, *Cybersecurity Framework*. Version 1.1. Gaithersburg, MD: The National Institute of Standards and Technology (NIST), Apr 16, 2018. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf; https://doi.org/10.6028/NIST.CSWP.04162018; https://www.nist.gov/cyberframework/framework; https://www.nist.gov/framework; Cybersecurity Framework website https://www.nist.gov/cyberframework

[117] https://www.nist.gov/sites/default/files/documents/2017/12/05/draft_roadmap-version-1-1.pdf

[118] DHS, *Critical Infrastructure Sectors*. Homeland Security. Aug 14, 2018. Available Aug 15, 2018 from IEEE at https://www.dhs.gov/critical-infrastructure-sectors

[119] Kaspersky Lab, *State of Industrial Cybersecurity 2018 Survey*. Kaspersky Lab, 2018. Available Aug 15, 2018 from Kaspersky at https://ics.kaspersky.com/the-state-of-industrial-cybersecurity-2018/

[120] Chee-Wooi Ten, Chen-Ching Liu, and Govindarasu Manimaran, "Vulnerability assessment of cybersecurity for SCADA systems," *IEEE Trans Power Systems*, vol. 23, no. 4, pp. 1836-46, Nov 2008.

[121] Muhammad Rizwan Asghar, György Dán, Daniele Miorandi, and Imrich Chlamtac, "Smart meter data privacy: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2820 - 2835, 2017. https://ieeexplore.ieee.org/document/7959167/

[122] Awais Rashid, George Danezis, Howard Chivers, Emil Lupu, Andrew Martin, Makayla Lewis and Claudia Peersman, "Scoping the Cyber Security Body of Knowledge," *IEEE Security & Privacy*, vol. 16, no. 3, pp. 96 - 102, May/June 2018. Available Aug 15, 2018 from Kaspersky at https://research-information.bristol.ac.uk/files/151229981/IEEE_SP_Paper_Author_Accepted.pdf; https://www.cybok.org/

[123] Russell Dotty, "IoT: Embedded and Secure," Military Embedded Systems, Nov 23, 2015. http://mil-embedded.com/guest-blogs/iot-embedded-and-secure/

[124] https://www.nytimes.com/2018/03/15/technology/saudi-arabia-hacks-cyberattacks.html

[125] Yosra Lakhdhar, Slim Rekhis, and Noureddine Boudriga, "Proactive security for safety and sustainability of mission critical systems," *IEEE Transactions on Sustainable Computing*, Feb 2018. (doi:10.1109/TSUSC.2018.2810092) https://ieeexplore.ieee.org/document/8303703/

[126] Qiang Liu, Pan Li, Wentao Zhao, Wei Cai, Shui Yu and Victor C.M. Leung, "A survey on security threats and defensive techniques of machine learning: A data driven view," *IEEE Access*, vol. 6, pp. 12103-12117, 2018. (doi: 10.1109/ACCESS.2018.2805680). https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8290925

[127] Catriona Mairi Kennedy, *Distributed Reflective Architectures for Anomaly Detection and Autonomous Recovery*. *Doctoral Thesis*. Birmingham , UK: University of Birmingham, Jun 2003. https://pdfs.semanticscholar.org/998d/81d3c09386e8b69ca18574654928bb14136f.pdf

[128] Awais Rashid, George Danezis, Howard Chivers, Emil Lupu, Andrew Martin, Makayla Lewis and Claudia Peersman, "Scoping the Cyber Security Body of Knowledge," *IEEE Security & Privacy*, vol. 16, no. 3, pp. 96 - 102, May/June 2018. Available Aug 15, 2018 from Kaspersky at https://research-information.bristol.ac.uk/files/151229981/IEEE_SP_Paper_Author_Accepted.pdf; https://www.cybok.org/

[129] https://dl.acm.org/citation.cfm?doid=3017680.3017811

[130] IEEE, *Try-CybSI*, 2018. Available Aug 15, 2018 from IEEE at http://try.cybersecurity.ieee.org/trycybsi/

[131] ACM, *ACM/IEEE/AIS SIGSEC/IFIP Cybersecurity Curricular Guideline*. 2017. http://cybered.acm.org/

[132] Hussein A. Abbass, Eleni Petraki, Kathryn Merrick, John Harvey, and Michael Barlow, “Trusted autonomy and cognitive cyber symbiosis: Open challenges,” *Cognitive Computation*, vol. 8, no. 3, pp. 385–408, Jun 2018. (doi: 10.1007/s12559-015-9365-5). https://link.springer.com/article/10.1007/s12559-015-9365-5; https://link.springer.com/content/pdf/10.1007%2Fs12559-015-9365-5.pdf

[133] Hussein A. Abbass, George Leu, and Kathryn Merrick, "A review of theoretical and practical challenges of trusted autonomy in big data," *IEEE Access*, vol. 4, pp. 2808-2830, June 24, 2016. https://ieeexplore.ieee.org/document/7480763/

[134] https://www.scmp.com/news/world/europe/article/2145896/thousands-people-sweden-get-microchip-implants-new-way-life

[135] https://www.forgerock.com/about-us/press-releases/consumer-privacy-survey-sponsored-forgerock-finds-strong-opinions-greater

[136] https://eugdpr.org/

[137] Ayla Networks, Navigating Data Security and Data Privacy in the IoT. Ayla Networks, 2016.

[138] OWASP, "The Open Web Application Security Project (OWASP) Top Ten Vulnerabilites," 2013. https://www.owasp.org/index.php/OWASP_Top_10/Mapping_to_WHID

[139] OWASP WASC Web Hacking Incidents Database Project," 2017. https://www.owasp.org/index.php/OWASP_WASC_Web_Hacking_Incidents_Database_Project

[140] OWASP, "OWASP XSS (Cross Site Scripting) prevention cheat sheet," Aug 15, 2018. https://www.owasp.org/index.php/XSS_(Cross_Site_Scripting)_Prevention_Cheat_Sheet

[141] Patrick Townsend, AES Encryption and Related Concepts. eBook, Townsend Security, 2016, pp.

[142] Patrick Townsend, The Encryption Guide: A Guide to Effective Encryption for Every Business. eBook, Townsend Security, 2018.

[143] Bruce Bosworth, Codes, Ciphers and Computers: An Introduction to Information Security. Rochelle Park, NJ: Hayden, 1982. (ISBN: 0-8104-5149-2)

[144] Dorothy Elizabeth Robling-Denning, Cryptography and Data Security. Reading, MA: Addison-Wesley, 1983. (ISBN: 0-201-10150-5)

[145] Artop Salomaa, Public-Key Cryptography. New York, NY: Springer, 1996. (ISBN: 3-540-61356-0)

[146] Bruce Scheider, Applied Cryptography: Protocols, Algorithms, and Source Code in C. New York, NY: Wiley, 1994. (ISBN: 0-471-59756-2)

[147] Randall K. Nichols, ISCSA Guide to Cryptography. New York, NY: McGraw-Hill, 1999. (ISBN: 0-07-913759-8)

[148] Michael Rosling, Elliptic Curve Cryptography. Greenwich, CT: Manning, 1999. (ISBN: 1-884777-69-4; pbk)

[149] https://nvlpubs.nist.gov/nistpubs/fips/nist.fips.197.pdf

[150] Jesus David Terrazas Gonzalez and Witold Kinsner, "A modular dynamical cryptosystem based on continuous-interval cellular automata," in Cognitive Informatics for Revealing Human Cognition: Knowledge Manipulations in Natural Intelligence, Yingxu Wang, Ed. Calgary, AB, Canada: IGI Global, 2012, pp. 261–286, (doi: 10.4018/978-1-4666-2476-4, ISBN: 9781466624764).

[151] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha P. Chandrakasan, "Gazelle: A Low Latency Framework for Secure Neural Network Inference," in 27th USENIX Security Symposium (Baltimore, MD; August 15–17, 2018), pp. 1651-1668, 2018. Open Access. https://www.usenix.org/conference/usenixsecurity18/presentation/juvekar; https://www.usenix.org/conference/usenixsecurity18; https://www.usenix.org/publications/proceedings

[152] http://direct-democracy.geschichte-schweiz.ch/

[153] http://history-switzerland.geschichte-schweiz.ch/switzerland-federal-constitution-1848.html

[154] https://www.washingtonpost.com/technology/2018/08/10/west-virginia-pilots-mobile-blockchain-voting-app-overseas-voters-november-election/?noredirect=on&utm_term=.79d25207b9f3

[155] Derrick de Kerckhove discusses the Age of Datacracy http://www.abc.net.au/news/2018-07-24/blockchain-innovators-seek-to-disrupt-indonesian-electoral-fraud/10025900

[156] S. Mason Dambrot, *Transinopia* (https://transinopia.org)

[157] Castells, M. (2007). Communication, power and counter-power in the network society. International Journal of Communication, 1, 238–266. Retrieved from http://ijoc.org/index.php/ijoc/article/viewFile/46/35

[158] Benayoun, M. Emotional Traffic, A performance http://www.petals.org/Barriere/EmotionalTraffic.html

[159] Iaconesi, S. VersuS – Rome, October 15th, the riots on social networks http://www.artisopensource.net/2011/10/16/versus-rome-october-15th-the-riots-on-social-networks/

[160] Bruno, c. WiFi-SM (2003-2006) https://christophebruno.com/portfolio/wifi-sm-2003-2006/

[161] Koepcke, C. MIMMI https://www.digitaltrends.com/social-media/see-the-twitter-mood-ring-over-minneapolis/

[162] http://www.lse.ac.uk/website-archive/newsAndMedia/newsArchives/2010/08/mappiness.aspx

[163] Glenn L.M., and Boyce J.S. At the nexus: augmented cognition, healthcare, and the law. Journal of Cognitive Engineering and Decision Making 1.3:363–373(2007)

[164] Glenn, L.M., Case Study: Ethical and Legal Issues in Human Machine Mergers (Or the Cyborgs Cometh), 21 Annals of Health Law, 175, 177-178 (2012)

[165] Glenn L.M. Biotechnology at the margins of personhood: an evolving legal paradigm. J. Evol. Techno. 13:35–37 (2003) and Wittes,B. and Chong, J. Our Cyborg Future - Law and Policy Implications, Brookings Institute.(2014)

[166] Warren, S.; Brandeis, L. D., *The Right to Privacy,* Harvard Law Review, Vol. 4, No. 5. (Dec. 15, 1890), pp. 193-220

[167] *Olmstead v. United States,* 277 U.S. 438, 48 S. Ct. 564, 72 L. Ed. 944, 1928 U.S. LEXIS 694, 66 A.L.R. 376 (1928)

(dissenting opinion)

[168] https://www.cbinsights.com/research/jobs-automation-artificial-intelligence-risk/

[169] https://techcrunch.com/2018/03/05/flippy-the-robot-hamburger-chef-goes-to-work/

[170] https://www.straitstimes.com/asia/south-asia/restaurant-in-india-replaces-waiters-with-robots

[171] https://www.wired.com/review/review-irobot-roomba-980/

[172] http://www.businessinsider.com/amazons-robot-army-has-grown-by-50-2017-1?r=UK&IR=T&IR=T

[173] https://www.drivesweden.net/en/news/scania-designs-full-scale-autonomous-truck-platoon

[174] http://tucson.com/business/developer-of-self-driving-technology-for-semi-trucks-arrives-in/article_a0c46f77-64db-5e40-9c5c-99decdf5db21.html

[175] https://slate.com/technology/2018/01/facebook-and-google-are-building-an-army-of-content-moderators-for-2018.html

[176] https://www.eitdigital.eu/eit-digital-academy/professional-school/

[177] http://www.philosophyinaction.com/blog/?p=2379

[178] https://www.weforum.org/agenda/2017/01/jobless-world-and-its-discontents/

[179] http://www.fabfoundation.org/

[180] https://www.wired.com/story/the-future-of-fab-lab-fabrication/

[181] https://www.businessinsider.com/apple-watch-series-3-vs-augmented-reality-2017-11

[182] http://www.oulu.fi/university/6gflagship

[183] Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation (http://arxiv.org/abs/1609.08144)

[184] Zero-Shot Translation with Google's Multilingual Neural Machine Translation System, https://ai.googleblog.com/2016/11/zero-shot-translation-with-googles.html

[185] The Rosetta Stone online (http://rosettastone.hieroglyphic-texts.net/)

[186] D. E. Smalley, E. Nygaard, K. Squire, J. Van Wagoner, J. Rasmussen, S. Gneiting, et al. A photophoretic-trap volumetric display. *Nature* volume 553, pp. 486–490 (25 January 2018). http://www.nature.com/doifinder/10.1038/nature25176

[187] V. G. Shvedov, C. Hnatovsky, A. V. Rode, and Wieslaw Krolikowski. Robust trapping and manipulation of airborne particles with a bottle beam. *Opt. Express* 19, 17350-17356 (2011). https://doi.org/10.1364/OE.19.017350

[188] https://www.theverge.com/2018/4/15/17235262/hyperloop-transportation-technologies-test-track-france

[189] https://www.scmp.com/news/china/society/article/2109035/china-firms-plan-develop-4000-km/h-train-just-pipe-dream

[190] Disruptive Technologies for a smart planet impacting beyond 2040 V, IEEE Future Directions (http://sites.ieee.org/futuredirections/2018/04/19/disruptive-technologies-for-a-smart-planet-impacting-beyond-2040-v/)

[191] https://www.grc.nasa.gov/www/BGH/scramjet.html

[192] https://www.grc.nasa.gov/www/k-12/airplane/scramjet.html

[193] World Economic Forum (http://reports.weforum.org/toward-the-circular-economy-accelerating-the-scale-up-across-global-supply-chains/from-linear-to-circular-accelerating-a-proven-concept/#view/fn-11)

[194] J. Kirchherr, D. Reike, M.Hekkert. Conceptualizing the circular economy: An analysis of 114 definitions. Resources, Conservation and Recycling, Volume 127, 2017, pp. 221-232. (https://doi.org/10.1016/j.resconrec.2017.09.005) Creative Commons Attribution 4.0 International (CC BY 4.0 License (https://creativecommons.org/licenses/by/4.0/)

[195] J. Kirchherr, D. Reike, M.Hekkert. Conceptualizing the circular economy: An analysis of 114 definitions. Resources, Conservation and Recycling, Volume 127, 2017, pp 221-232. (https://doi.org/10.1016/j.resconrec.2017.09.005) Creative Commons Attribution 4.0 International (CC BY 4.0 License (https://creativecommons.org/licenses/by/4.0/)

[196] Joint Center for Artificial Photosynthesis (JCAP) (https://solarfuelshub.org/recent-science)

[197] C. Liu, B. C. Colón, M. Ziesack, P. A. Silver, D.l G. Nocera. Water splitting–biosynthetic system with CO_2 reduction efficiencies exceeding photosynthesis. Science | 03 Jun 2016 : 1210-1213. (http://science.sciencemag.org/content/352/6290/1210)

[198] P. K. Basu, et al.18.7% Efficient inline-diffused screen-printed silicon wafer solar cells with deep homogeneous emitter etch-back. Solar Energy Materials and Solar Cells, Volume 117, 2013, pp. 412-420.

(http://www.sciencedirect.com/science/article/pii/S0927024813003413)

[199] A Big Leap for an Artificial Leaf (https://www.technologyreview.com/s/601641/a-big-leap-for-an-artificial-leaf/)

[200] S. Arumugam, Y. Li, J. Liu, J. Tudor, S. Beeby. Optimized Process of Fully Spray-Coated Organic Solar Cells on Woven Polyester Cotton Fabrics. Materials Today: Proceedings, Volume 5, Issue 5, Part 3, 2018, pp. 13745-13752. (http://www.sciencedirect.com/science/article/pii/S2214785318301998)

[201] The Spray-On Solar Cell (https://www.asme.org/engineering-topics/articles/renewable-energy/the-spray-on-solar-cell)

[202] World population trends (https://www.unfpa.org/world-population-trends)

[203] World Urbanisation Trends, 1950-2030. (http://www.unesco.org/education/tlsf/mods/theme_c/popups/mod13t01s009.html)

[204] Global Trends of Urbanization (https://morphocode.com/global-trends-urbanisation/)

[205] World Energy Outlook 2017 (https://www.iea.org/weo2017/)

[206] D. B. Go, et al. Thermionic Energy Conversion in the Twenty-first Century: Advances and Opportunities for Space and Terrestrial Applications. *Front. Mech. Eng.* 3:13, 08 November 2017. (https://doi.org/10.3389/fmech.2017.00013)

[207] Confluence: AI and Machine Learning applied to Cybersecurity, Brumley, et al. IEEE Report
[208] https://www.grandviewresearch.com/industry-analysis/exoskeleton-market

[209] https://arxiv.org/abs/1406.2661

[210] http://www.hmdb.ca

[211] Ian Goodfellow, Yoshua Bengio, Aaron Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016. (ISBN: 978-026203561-3). https://mitpress.mit.edu/books/deep-learning; HTML form: http://www.deeplearningbook.org/

[212] Witold Kinsner, “Challenges in the design of adaptive, intelligent and cognitive systems,” Intern. J. Software Science & Computational Intelligence, vol. 1, no. 3, pp. 16-35, July-Sep 2009.

[213] Witold Kinsner, Simon Haykin, Yingxu Wang, Witold Pedrycz, Ivo Bukovsky, Bernard Widrow, Andrzej Skowron, Piotr Wasilewski, and Menahem Friedman, “Challenges in engineering education of cognitive dynamic systems,” in *Proc. of the Canadian Engineering Education Association Conference*, CEEA 2012, (Winnipeg, MB, Canada; June 17–20, 2012), Paper 119, pp. 51-62, 2012. Accessed Mar 2015 from CEEA at http://library.queensu.ca/ojs/index.php/PCEEA/article/viewFile/4633/4615; http://www.ceea.ca/images/content/ceea12-proc-complete-v35s.pdf

[214] Fred S. Keller, "Goodbye teacher...," *Journal of Applied Behavior Analysis*, vol 1, pp. 79-89, Spring 1968. Available Aug 15, 2018 from NIH at https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1310979/pdf/jaba00083-0078.pdf

[215] Witold Kinsner and Joseph J. Pear, “Computer-aided Personalized System of Instruction for the virtual classroom,” Can. J. Educational Communication, vol. 17, no. 1, pp. 21-36, 1988.

[216] Antonio Fini, The Technological Dimension of a Massive Open Online Course: The Case of the CCK08 Course Tools," *International Review of Research in Open and Distributed Learning*, vol. 10, no. 5, 2009.

[217] G2Crowd, "105 Best Online Course Providers 2018," G2 Crowd, 2018. https://www.g2crowd.com/categories/online-course-providers

[218] Terry Anderson (ed.), *The Theory and Practice of Online Learning*. Edmonton, AB: AU Press, 2008 (2nd ed.). (ISBN: 9781897425077, PDF, free). http://www.aupress.ca/books/120146/ebook/99Z_Anderson_2008-Theory_and_Practice_of_Online_Learning.pdf

[219] Patricia Margaret Alexander, *Towards reconstructing meaning when text is communicated electronically*. PhD Thesis (Information Technology), University of Pretoria, Pretoria, South Africa, Aug 23, 2003. http://hdl.handle.net/2263/27367; Full Copy https://repository.up.ac.za/bitstream/handle/2263/27367/Complete.pdf?sequence=13&isAllowed=y

[220] Gary Woodill, "Rapid doubling of knowledge drives change in how we learn," *Float*, Jan 23, 2018. https://gowithfloat.com/2018/01/rapid-doubling-knowledge-drives-change-learn/

[221] Alvin Toffler, *Future Shock*. New York, NY: Bantam Books, 1970 (re-issue 1984). (ISBN-13: 978-05532773710)

[222] George Veletsianos (ed.), *Emergence and Innovation in Digital Learning: Foundations and Applications*. Edmonton, AB: AU Press, 2016. (ISBN: 9781771991490, pbk; 978-1-77199-150-6 (PDF); doi:10.15215/aupress/9781771991490.01. Available (free) Aug 15, 2018 from Athabasca Univ at http://www.aupress.ca/books/120258/ebook/99Z_Veletsianos_2016-Emergence_and_Innovation_in_Digital_Learning.pdf

[223] Christina Hinton, Kurt W. Fischer, and Catherine Glennon, Mind, Brain, Education: The Students at the Center Series. Report, Boston, MA: Jobs for the Future, 2012. Available Aug 15, 2018 at http://www.studentsatthecenter.org/sites/scl.dl-dev.com/files/Mind Brain Education.pdf

[224] Marcy P. Driscoll, *Psychology of Learning for Instruction*. New York, NY: Pearson, 2005 (3rd ed.). (ISBN-13: 9780205375196). Chapter 1 available Aug 15, 2018 from at http://ocw.metu.edu.tr/file.php/118/Dris_2005.pdf

[225] Erik De Corte, "Historical developments in the understanding of learning," in Hanna Dumont, David Istance, and Francisco Benavides (eds.), *The Nature of Learning: Using research to inspire practice*. OECD Publishing, pp. 35–68, 2010. (ISBN: 9789264086487) Available Aug 15, 2018 from OECD Library at http://www.keepeek.com/Digital-Asset-Management/oecd/education/the-nature-of-learning_9789264086487-en#page1; https://www.oecd-ilibrary.org/education/the-nature-of-learning_9789264086487-en; DOI:http://dx.doi.org/10.1787/9789264086487-en)

[226] Stephen Downes, *Connectivism: A Theory of Personal Learning*. Dec 8, 2008. Retrieved from http://www.slideshare.net/Downes/connectivism-a-theory-of-personal-learning

[227] George Siemens and Stephen Downes, *Connectivism and Connective Knowledge*. Online Course, 2009. Retrieved from http://ltc.umanitoba.ca/connectivism/

[228] George Siemens, "Connectivism: A Learning Theory for the Digital Age," 2004. elearnspace. Retrieved from http://www.elearnspace.org/articles/connectivism.htm

[229] George Siemens, "Connectivism: Learning as Network-Creation," ELEARNSPACE, Aug 10, 2005. Retrieved from http://www.elearnspace.org/Articles/networks.htm

[230] George Siemens, "Connectivism: Learning theory or pastime for the self amused?" 2006. Retrieved from http://www.elearnspace.org/articles/connectivism.htm

[231] George Siemens, "Connectivism: Learning conceptualized through the lens of today's world," at 2007 Online Connectivism Conference, OCCC2007, Slide Share, Feb 4, 2007. https://www.slideshare.net/gsiemens/connectivismonline

[232] George Siemens, "New spaces and structures of learning: the systemic impact of connective knowledge, connectivism, and networked learning," 2008. In unknown. Presented at the unknown event, Portugal. Retrieved from http://elearnspace.org/Articles/systemic_impact.htm

[233] George Siemens, "Learning and knowing in networks: Changing roles for educators and designers," Presented to *ITFORUM* for Discussion January 27, 2008. https://www.calvin.edu/~dsc8/documents/GeorgeSiemensPresentation-Jan2008.pdf

[234] George Siemens, *What is Connectivism? Week 1: CCK09*. Course Handout for CCK09, Sep 12, 2009a. Retrieved from https://docs.google.com/Doc?id=anw8wkk6fjc_14gpbqc2dt

[235] George Siemens and Tittenberger, P. *Handbook of Emerging Technology for Learning*. 2009. http://umanitoba.ca/learning_technologies/cetl/HETL.pdf

[236] Pearl Chen, Nora Mossessian, Renee Rodriguez, and Teresa Wu, Connectivism Learning Theory. *Survey of Educational Media and Technology,* Slide Show, Nov 10, 2010. https://www.slideshare.net/hieru/connectivism-5941513?next_slideshow=1

[237] A.Austen, P.C.Chen, V. Darlington, D.Daylamani-Zad, S.De'Cage, C.Farrant, "Connectivism Learning Theory," Slide Show, Nov 1, 2014. https://www.slideshare.net/CatherineFarrant/connectivist-learning-theory

[238] Su-Tuan Lulee, Sheri Oberman, and Tony Tin, "Connectivism Connecting with George Siemens," Slides. Athabasca University, Slide Show, Nov 21, 2010.

[239] Rita Kop and Adrian Hill, "Connectivism: Learning theory of the future or vestige of the past?" The International Review of Research in Open and Distance Learning, vol. 9, no. 3, 2008. Available from Athabasca University at http://www.irrodl.org/index.php/irrodl/article/view/523/1103

[240] Plon Verhagen, *Connectivism: A new learning theory?* 2006. Retrieved from https://www.scribd.com/doc/88324962/Connectivism-a-New-Learning-Theory

[241] Big Data. *NESSI White Paper*, December 2012. https://www.slideshare.net/tint10/connectivism-connecting-with-george-siemens?from_action=save

[242] Donald E. Knuth, "What is kilobyte?" 1998. https://www-cs-faculty.stanford.edu/~knuth/news99.html; https://www-cs-faculty.stanford.edu/~knuth/index.html

[243] Thomas M. Cover and Joy A. Thomas, *Elements of Information Theory*. New York, NY: Wiley-Interscience, 2006 (2nd ed.). (ISBN-13: 979-0-471-24195-9)

[244] W. Kinsner, "A unified approach to fractal dimensions," in *Proc. IEEE 2005 Intern. Conf. Cognitive Informatics*, ICCI05 (Irvine, CA; August 8-10, 2005), ISBN: 0-7803-9136-5, pp. 58-72, 2005.

[245] Gregory J. Chaitin, *Algorithmic Information Theory*. Cambridge, UK: Cambridge University Press, 1987. (ISBN 0-521-34306-2)

[246] Tom Stonier, *Information and the Internal Structure of the Universe: An Exploration into Information Physics*. New York, NY: Springer, 1990. (ISBN: 0-387-19599-8

[247] Robert A. Reiser and John V. Dempsey (eds.), *Trends and Issues in Instructional Design and Technology*. Upper Saddle River, NJ: Merrill Prentice Hall, 2007.

[248] IEEE, *Towards a Definition of the IT Profession*. White Paper. IEEE Computer Society, Professional Activities Board, Information Technology Committee, May 16, 2010. Accessed Mar 2015 from IEEE at http://www.computer.org/cms/professional-education/pdf/IT%20White%20Paper%20-%20rev0.pdf.

[249] Richard H. McCuen, Edna Z. Ezzell, and Melanie K. Wong, *Fundamentals of Civil Engineering: An Introduction to the ASCE Body of Knowledge*. Boca Raton, FL: CRC Press, 2011. (ISBN 978-1-4398-5148-7, pbk; $92.09)

[250] CEAB, *Canadian Engineering Accreditation Board*. 2015. Accessed Jan 2015 from Engineers Canada at https://www.engineerscanada.ca/; https://www.engineerscanada.ca/committee-listings; https://www.engineerscanada.ca/canadian-engineering-accreditation-board

[251] ABET, *The Accreditation Board of Engineering and Technology*. 2015. Accessed Apr 2015 from ABET at http://www.abet.org/; http://www.abet.org/accreditation/accreditation-criteria/

[252] ABET, *Accreditation Policy and Procedure Manual (APPM), 2018-2019*. Oct 20, 2017. Accessed Apr 2015 from ABET at http://www.abet.org/accreditation/accreditation-criteria/; http://www.abet.org/accreditation/accreditation-criteria/accreditation-policy-and-procedure-manual-appm-2018-2019/

[253] IEA, Washington Accord. 1989. Accessed Apr 2015 from the IEA at http://www.ieagreements.org/; http://www.ieagreements.org/Washington-Accord/Accredited.cfm

[254] Sydney IEA, *Sydney Accord*. Jun 2001. Accessed Apr 2015 from IEA at http://www.ieagreements.org/sydney/

[255] Dublin IEA, *Dublin Accord*. May 2002. Available Apr 2015 from the International Engineering Alliance (IEA) at http://www.ieagreements.org/Dublin/

[256] IPEA, *International Professional Engineers Agreement*. 2001. Accessed Mar 2015 from IEA at http://www.ieagreements.org/EMF/

[257] APEC, *Asia Pacific Economic Cooperation (APEC) Engineer Agreement*. 1999. Accessed Apr 2015 from the International Engineering Alliance (IEA) at http://www.ieagreements.org/; http://www.ieagreements.org/apec/

[258] IETA, *International Engineering Technologist Agreement*. 25 June 2001. Accessed Mar 2015 from IEA at http://www.ieagreements.org/etmf/

[259] University of Manitoba, *Internationally -Educated Engineers Qualification (IEEQ) Program*. 2015. Accessed Mar 2015 from UofM at http://umanitoba.ca/faculties/engineering/programs/ieeq/

[260] Ken Ferens, Jill Seniuk-Cicek, Nariman Sepehri, Witold Kinsner, Jean-Paul Burak, Danny Mann, Ann Parker, Dean McNeill, Doug Ruth, Ian Jeffrey, Norma Godavari, "Industry Forum III: Towards a common language," in *Proc. 5th Conference of the Canadian Engineering Education Association*, CEEA 2014 (Canmore, AB; June 8-11, 2014). Paper 035. 2014.

[261] Witold Kinsner, *Expanding the Concept of Body of Knowledge for Practitioners (BoK4P)*. Report. (Winnipeg, MB: IEEE Winnipeg Section AGM; Jan 28, 2015), 2015.

[262] Witold Kinsner, "Expanding the body of knowledge concept for professional practitioners," in *Proc. 2015 Canadian Engineering Education Association (CEEA15) Conf*., (McMaster University, Hamilton, ON; May 31-June 3, 2015), Paper 172, 2015. Available as of Aug 15, 2018 from CEEA at http://yint.org/static/PDF-CEEA-2015/Int-T1-5-paper-172.pdf

[263] Witold Kinsner, "IEEE for both knowledge dicoverers and professional practitioners: The body of knowledge for practitioners (BoK4P) project," Keynote Talk. (Winnipeg, MB: IEEE Winnipeg Section AGM; Jan 28, 2015), 2015.

[264] Witold Kinsner, "Humanitarian engineering education: Examples," in *Proc. 5th Conference of the Canadian Engineering Education Association*, CEEA 2014 (Canmore, AB; June 8-11, 2014). Paper 121, 2014.

[265] Christopher M. Bishop, *Pattern Recognition and Machine Learning*. Cambridge, UK: Springer Science, 2004 (2nd ed.).

[266] Yingxu Wang, Du Zhang, and Witold Kinsner (eds.), *Advances in Cognitive Informatics and Cognitive Computing*. New York, NY: Springer, 2010. (ISBN: 978-3642160820)

[267] Yingxu Wang, "On cognitive informatics," in Proc. 1st IEEE Intern. Conf. Cognitive Informatics (Calgary, AB; 19-20 August 2002), pp. 34-42, 2002.

[268] Ning Zhong, Jiming Liu, and Yiyu Yao (eds.), *Web Intelligence*. New York, NY: Springer, 2010. (ISBN: 978-3642079368)

[269] Simon Haykin, José C. Principe, Terrence J. Sejnowski, and John McWhirter (eds.), *New Directions in Statistical Signal Processing*. Cambridge, MA: MIT Press, 2007.

[270]Simon Haykin, José C. Principe, Terrence J. Sejnowski, and John McWhirter (eds.), *New Directions in Statistical Signal Processing*. Cambridge, MA: MIT Press, 2007.

[271] Simon Haykin and Bart Kosko (eds.), *Intelligent Signal Processing*. Piscataway, NJ: IEEE Press, 2001.

[272] Richard G. Baraniuk, Emmanuel Candes, Robert Novak, and Martin Vetterli, "Compressive sampling" IEEE Signal Processing, vol. 25, no. 2, pp. 12-20, March 2008.

[273] Witold Pedrycz, *Granular Computing: An Emerging Paradigm*. New York, NY: Physica-Verlag HD, 2010 (2nd ed.). (ISBN: 978-3790824872)

[274] Gregory W. Wornell, *Signal Processing with Fractals: A Wavelet-Based Approach*. Upper Saddle River, NJ: Prentice-Hall, 1996.

[275] Witold Kinsner, "Polyscale analysis and fractional operators for cognitive systems," in *Proc. 13th IEEE International Conference on Cognitive Informatics & Cognitive Computing*, ICCI*CC 2014 (London, UK; August 18-20, 2014) Paper KN5, 2014. (Keynote Talk)

[276] Witold Kinsner, "A unified approach to fractal dimensions," *Intern. J. Cognitive Informatics and Natural Intelligence*, vol. 1, no. 4, pp. 26-46, Oct-Dec 2007.

[277] Julien Clinton Sprott, *Chaos and Time-Series Analysis*. Oxford, UK: Oxford Univ. Press, 2003.

[278] Holger Kantz, and Thomas Schreiber, *Nonlinear Time Series Analysis*. Cambridge, UK: Cambridge Univ. Press, 2004 (2nd ed.).

[279] Steven H. Strogatz, *Nonlinear Dynamics and Chaos*. Cambridge, MA: Westview/Perseus, 2000.

[280] Robert L. Devaney, An Introduction to Chaotic Dynamical Systems. Menlo Park, CA: The Benjamin Cummings Publishing, 1986.

[281] https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/
[282] https://www.statista.com/statistics/265759/world-population-by-age-and-region/
[283] https://lora-alliance.org/about-lorawan

[284] http://media.nxp.com/phoenix.zhtml?c=254228&p=RssLanding&cat=news&id=2334390

[285] http://www.electronics-lab.com/stmicroelectronics-introduces-stm32wb-soc-32bit-microcontroller-bluetooth-low-energy/

[286] https://markets.businessinsider.com/news/stocks/lowest-power-multi-standard-multi-band-mcus-connect-the-building-factory-and-grid-via-thread-zigbee-bluetooth-5-and-sub-1-ghz-1017433247

[287] https://www.nordicsemi.com/News/News-releases/Product-Related-News/Nordic-nRF91-low-power-cellular-IoT-sneak-peek-cellular-made-easy-cellular-for-everything-else

[288] https://www.iso.org/standard/73599.html

[289] https://www.st.com/en/nfc/st25dv-i2c-series-dynamic-nfc-tags.html

[290] Witold Kinsner, "Polyscale analysis and fractional operators for cognitive systems," in *Proc. 13th IEEE International Conference on Cognitive Informatics & Cognitive Computing*, ICCI*CC 2014 (London, UK; August 18-20, 2014) Paper KN5, 2014. (Keynote Talk)

[291] Witold Kinsner, "Is entropy suitable to characterize data and signals for cognitive informatics?" *Intern. J. Cognitive Informatics and Natural Intelligence*, vol. 1, no. 2, pp. 34-57, Apr-Jun 2007.

[292] Witold Kinsner, "A unified approach to fractal dimensions," *Intern. J. Cognitive Informatics and Natural Intelligence*, vol. 1, no. 4, pp. 26-46, Oct-Dec 2007.

[293] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016. (ISBN: 978-026203561-3). https://mitpress.mit.edu/books/deep-learning; HTML form: http://www.deeplearningbook.org/

[294] Yang Xin, Lingshuang Kong, Zhi Liu, Yuling Chen, Yanmiao Li, Hongliang Zhu, Mingcheng Gao, And Haixia Hou, And Chunhua Wang, "Machine Learning and Deep Learning Methods for Cybersecurity," *IEEE Access*, pp. 35365-81, July 19, 2018. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8359287

[295] Jesus David Terrazas Gonzalez and Witold Kinsner, "Zero-crossing analysis of Lévy walks and a DDoS dataset for real-time feature extraction: Composite and applied signal analysis for strengthening the Internet-of-Things against DDoS attacks," *Int. J. Softw. Sci. Computational Intelligence* IJSSCI, vol. 8, no. 4, pp. 1–28, 2016. (doi: 10.4018/IJSSCI.2016100101)

[296] Jesus David Terrazas Gonzalez and Witold Kinsner, "Zero-crossing analysis and information divergence of Lévy walks for real-time feature extraction," *Int. J. Handheld Comput. Res*. IJHCR, vol. 7, no. 4, pp. 41–59, 2016. (doi: 10.4018/IJHCR.2016100104)