The Role of Personal Digital Twins in Control of Epidemics

An IEEE Digital Reality White Paper

April 2020
Roberto Saracco with Juuso Autiosalo, Derrick de Kerckhove, Francesco Flammini, and Louis Nisiotis

Edited by Theresa Cavrak
DigitalReality.ieee.org

Contents

Introduction

Since the shift from nomadic life to aggregation in clusters or cities humanity has faced epidemics. It is the cluster of people that provides the fertile environment for viruses to jump from one host to the next generating an epidemic. The geographical distance among clusters is a barrier to the spread of the epidemics; traveling from one locale to another was the only way to continue the spread. In the past, as shown in Figure 1, the epidemics spread along the commerce, maritime, and land pathways. Travel was slow and sporadic so an epidemic took years to become a pandemic.

Today we have both bigger clusters (megacities and cities that on average are much bigger than the ones of the past) and much faster and denser traffic among clusters. This fuels both epidemics and pandemics.

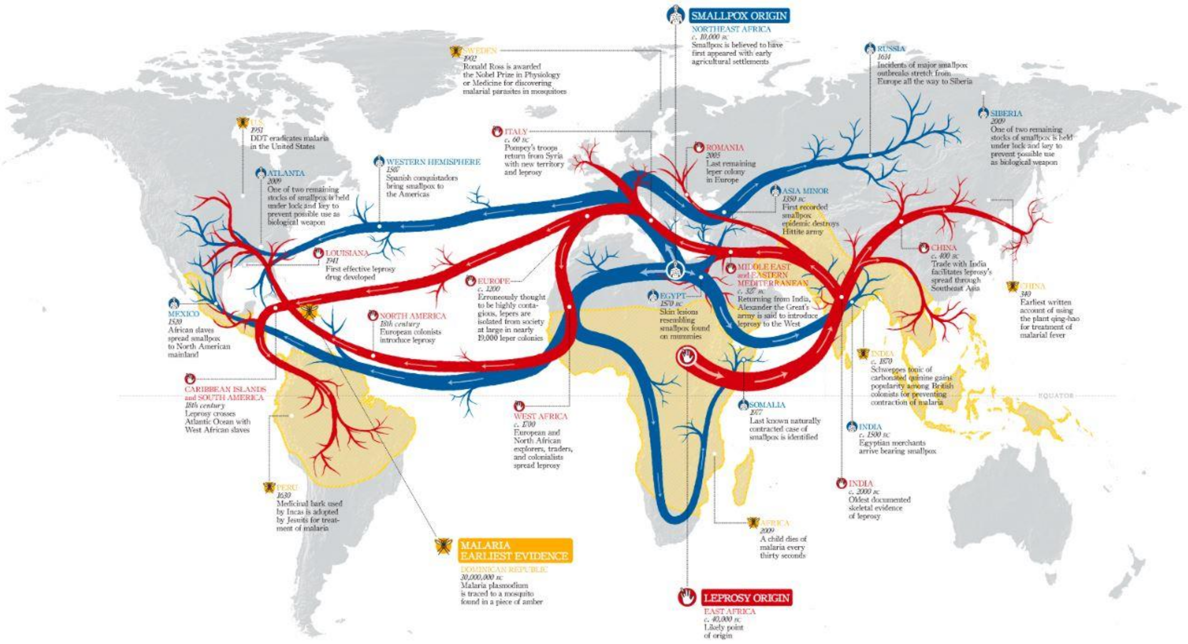
SMALLPOX ORIGIN NORTHEAST AFRICA c. 10,000 BC Smallpox is believed to have first appeared with early agricultural settlements
SWEDEN 1902 Ronald Ross is awarded the Nobel Prize in Physiology or Medicine for discovering malarial parasites in mosquitoes
RUSSIA 1614 Incidents of major smallpox outbreaks stretch from Europe all the way to Siberia
SIBERIA 2009 One of two remaining stocks of smallpox is held under lock and key to prevent possible use as biological weapon
U.S. 1951 DDT eradicates malaria in the United States
ITALY c. 60 BC Pompey's troops return from Syria with new territory and leprosy
ROMANIA 2005 Last remaining leper colony in Europe
WESTERN HEMISPHERE 1507 Spanish conquistadors bring smallpox to the Americas
ATLANTA 2009 One of two remaining stocks of smallpox is held under lock and key to prevent possible use as biological weapon
ASIA MINOR 1350 BC First recorded smallpox epidemic destroys Hittite army
CHINA c. 400 BC Trade with India facilitates leprosy's spread through Southeast Asia
LOUISIANA 1941 First effective leprosy drug developed
EUROPE c. 1200 Erroneously thought to be highly contagious, lepers are isolated from society at large in nearly 19,000 leper colonies
MIDDLE EAST and EASTERN MEDITERRANEAN c. 327 BC Returning from India, Alexander the Great's army is said to introduce leprosy to the West
EGYPT 1570 BC Skin lesions resembling smallpox found on mummies
CHINA 340 Earliest written account of using the plant qing-hao for treatment of malarial fever
MEXICO 1520 African slaves spread smallpox to North American mainland
NORTH AMERICA 18th century European colonists introduce leprosy
INDIA c. 1870 Schweppes tonic of carbonated quinine gains popularity among British colonists for preventing contraction of malaria
CARIBBEAN ISLANDS and SOUTH AMERICA 18th century Leprosy crosses Atlantic Ocean with West African slaves
WEST AFRICA c. 1700 European and North African explorers, traders, and colonialists spread leprosy
SOMALIA 1977 Last known naturally contracted case of smallpox is identified
INDIA c. 1500 BC Egyptian merchants arrive bearing smallpox
EQUATOR
INDIA c. 2000 BC Oldest documented skeletal evidence of leprosy
PERU 1630 Medicinal bark used by Incas is adopted by Jesuits for treatment of malaria
AFRICA 2009 A child dies of malaria every thirty seconds
MALARIA EARLIEST EVIDENCE DOMINICAN REPUBLIC 30,000,000 BC Malaria plasmodium is traced to a mosquito found in a piece of amber
LEPROSY ORIGIN EAST AFRICA c. 40,000 BC Likely point of origin

Figure 1 The origin and spread of smallpox, leprosy and malaria around the globe. Macro-representation. Image credit: Doug Belshaw blog

Luckily, today there are much better tools to fight the effect of viruses, however, there are some limitations, especially considering fast-moving pandemics such as the 2019-2020 Coronavirus. Particularly, fighting viruses requires vaccination which takes time to develop for novel viruses, such as the Coronavirus that jumped from animals, possibly bats, to humans. Additionally, these tools come in limited supply, and therefore it is of paramount importance to detect the potential insurgence of epidemics as soon as possible and delay the spreading.

The graphic presented was created based on historical information of actual contagion and would not have been useful to people at that time. There is a need to draw such a graphic in real time and to be able to predict its evolution. Today there are tools based on data harvested from various sources interpreted through epidemic models (taking into account the ways contagion happens and its diffusion speed) that look at the movement of people within a community and across different areas.

Social media are used to capture the manifestation of a virus and the habits of people that can lead to exposure. The current outbreak of coronavirus has been monitored and its growth predicted in various centers like the Network Science Institute at Northeastern University in Boston through big data analytics applied at social networks.

Social networks can indeed be used as sensors but their sensitivity and in particular their resolution is not optimal. This is where personal digital twins (PDT) may play a role. A PDT is a representation of various aspects of a person that might include the movement of the person, the interactions that person has in physical space with other people, and her health status (like presence of fever or coughing). These data can be accrued by the PDT using a few sensors already available: position and movement can easily be monitored by extracting data from her smartphone and health status can be monitored through wearable sensors, such as smartwatches measuring body temperature, heart rate, and other health markers.

There are not PDTs currently available. There are a number of companies that are starting to propose the creation of PDTs and specifically to use them in the healthcare domain, like General Electric, Siemens, and Philips, however there is no real adoption of digital twin technology to mirror individual persons.

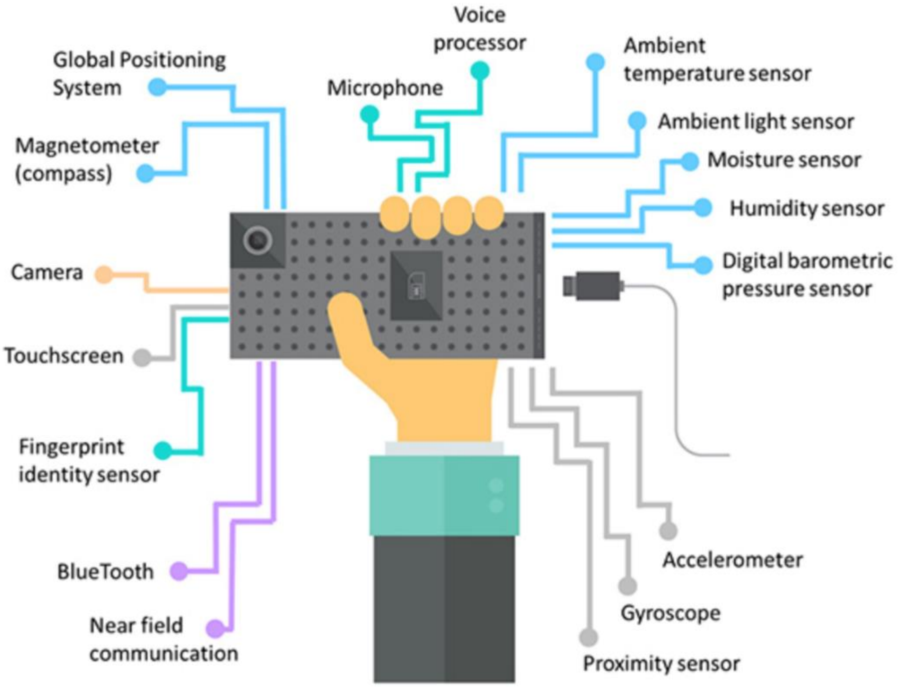

Figure 2 A smartphone is brimming with sensor,s and their number keeps growing. The data generated by these sensors can be used to identify and analyze the person's activity and health status with a growing sensitivity and precision. These data become useful in the monitoring of epidemics and in enacting control protocols. Image credit: B. Cheung

Evolution of Medical Cyber-Physical Systems and Introduction to PDTs in Epidemic Control

By Louis Nisiotis

Medical Cyber-Physical Systems (MCPS) are critical and distributed multilinked intelligent systems used in medicine, connecting physical and digital environments through connected devices, focusing on the Digital Twin as the representation and digitalization of a physical device (Jimenez et al, 2020). The PDT approach presented here goes beyond, and suggests the digitization of a physical person to enable anonymous and secure sensing, collection and analysis of data to inform strategic decisions that can disrupt the current ways the healthcare system works and manages pandemic situations. This approach has opportunities for providing accurate and consistent gathering and communication of data that can help mining actionable insights to enable decision making and help to emerge disruptive Cyber-Physical Systems. However, when technology is utilized in ways that require sensitive and private data to be shared and managed, there are significant privacy concerns that needs to be taken under deep consideration.

As a result of the COVID-19 outbreak, there is a significant increase in exploiting digital technologies to help in the battle of containing the virus. For instance, the European Commission issued an urgent call for technologies and innovation solutions that could help against the virus outbreak. The use of smart phones is one of the many technologies that have been utilized, offering opportunities that can contribute to the attempt to combat this pandemic crisis. Some countries have either already implemented or are considering measures using smart phone related technologies to fight COVID-19. Other examples include the efforts from researchers from Oxford University who are demonstrating feasibility evidence to European governments to explore the potential of developing smart phone apps for contact tracing, a recent development of a novel AI-enabled Framework to diagnose the virus using smartphones, and a few other scientific studies indicating the use of digital tracing to support containing the spread. Recent news revealed that a UK mobile operator is in talks with the British government in regards to using location and usage data to monitor the virus and determine if the government measures against the virus work, and in general, there is a widespread interest in the smart phone industry to explore the development of data sharing systems to support the efforts against COVID-19. Even Google and Facebook are now in talks with the USA government in regards to sharing anonymized location data, in an attempt to combat the spread of the virus.

However, for the successful implementation of any technological solution that collects, and processes personal data in an attempt to help against the virus, it is important to consider privacy. In any form of a digital group, the data protection, transparency, openness and honesty is key, and needs to be taken under deep consideration when emerging and utilizing such systems. It is also important to consider how willing people are to share their personal data and trust the 'trusted' source who will be managing and analyzing their personal data. Several mechanisms exist to support encryption and data protection. For instance the use of Blockchain methods for secure storage of data, encryption, sharing, and to enable the user/patient to manage their own health data may be an example to be considered.

According to the Chair of the European Data Protection Board (EDPB) "*data protection rules (such as GDPR) do not hinder measures taken in the fight against the coronavirus pandemic. However, I would like to underline that, even in these exceptional times, the data controller must ensure the protection of the personal data of the data subjects*". A recent statement from the UK Information Commissioner's Office (ICO), indicates that "*data protection and electronic communication laws do not stop Government, the NHS or any other health professionals from sending public health messages to people, either by phone, text or email as these messages are not direct marketing. Nor does it stop them using the latest technology to facilitate safe and speedy consultations and diagnoses. Public bodies may require additional collection and sharing of personal data to protect against serious threats to public health.*" The Director for Regulatory Assurance at the ICO points out some basic data protection suggestions on how to apply the law during the current outbreak of the virus, suggesting to keep clear how data is used, keep sharing, keep it lawful, keep it secure, keep it to a minimum, and also keep record of what has been done.

The European Commission identified the valuable role digital technologies can play in the current pandemic situation, and produced a series of key recommendations and measures (Toolbox) to develop a common approach for the use of mobile technology and mobile data to exit the Covid-19 crisis, considering security and the respect of EU fundamental rights such as data protection and privacy.

Considering the official data protection guidance, the EU Toolbox recommendations, laws and regulations, and other the relevant ethical considerations, the development of PDTs can provide opportunities that can contribute into the ongoing effort of monitoring and containing the virus. Using the data provided by the PDT, the science of outbreak analytics, an emerging data analysis science focusing on the collecting, analysis, modelling and reporting of outbreak data, among other data analysis techniques, can be implemented and help to inform real time and evidence based response decisions.

- The PDT approach can enable the evaluation of a person's condition through its Digital Twin by, for instance, by an AI agent that assesses the information provided by the user or thought wearable devices to the system, and flag infected Digital Twins.
- Using data provided by infected PDTs in anonymous ways to identify red zones and generate maps of infected areas. Further data analysis can enable identifying patterns of spreading and containing the disease, raise issues, and provide important insights to hospitals to allow better preparation for bed and equipment shortages, staff scheduling, and to manage patient flow.
- The development of social virtual spaces where Personal Digital Twins with infected status can connect to seek advice.
- Governments and healthcare organizations can disseminate announcements, general advice, dietary plans, and other information specific to infected PDTs directly on their digital device.
- Perform data analysis that can support decision making and further contribute into predictive forecasting to inform strategic decisions.

Again, the key for a successful implementation of such framework is to preserve privacy, ensure anonymity, transparency of data process and ethical commitment by everyone. The requirements for an approach like this to work are:

- Social awareness of the significance of the situation, and a conscious decision from people that their data will be shared, monitored, analyzed and evaluated in the attempt to contain the virus;
- Technology infrastructure is secure and efficient;
- Trusted parties commit to ethical usage, processing, and evaluation of data, must ensure the protection of the personal data, obey to the rules, and keep to these promises.

Personal Digital Twin Architecture for Epidemic Control

Let's suppose that the evolution of personalized healthcare in this decade will result in the adoption of PDTs. By the end of this decade one could imagine that every person will be flanked by a digital twin, able to raise a red flag in case of need. This red flag can be customized by the person, or more likely by that person's physician, to generate information when a certain situation emerges, or when there is a risk for the emergence of a certain situation. For an example, all personal digital twins can be designed with instructions to raise a red flag when a mix of data creates a specific pattern, like temperature at rest rising over 37.5 C and occurrence of rapid breathing—this is a sign of a possible COVID-19 infection.

Healthcare institutions at the government level can receive these red flags and in turn can analyze a variety of connected data (such as the occurrence of these red flags in a specific area, the movement of people in the previous months to correlate with the emergence of other red flags). Notice that in this scenario government and healthcare institutions may impose some kind of data analytics and red flag generation on all PDTs in order to harvest data. This would create an awareness of an incipient epidemic and make an accurate forecast of possible contagion based on the movement of people with red flags. It would provide a most timely and most accurate picture of the situation, worldwide.

This is a great benefit to epidemic control, however it is also raising big issues in terms of privacy and government/institution control. PDTs may take the Orwell 1984 vision a step closer as they would take all of our society into an uncharted territory.

This scenario is still a few years away, however Digital Twins are already considered in the simulation of some epidemics to support the creation of flu vaccines.

Although full-fledged epidemic control based on PDTs is still several years away, the South Korean government is using a similar approach. The Korean Center for Disease Control, KCDC, has organized an impressive collection of data from smartphone locations (tracked by local telecoms) and from public security cameras. This data is then used to create a contagion map focused on people's interactions at a micro level. Quick tests have been organized (taking samples from people in their own cars and processing them quickly), and based on the results, an updated contagion map led to the isolation of specific people rather than the generic lockdown of territories as it was done in China, Italy, Spain, Austria, and many more countries.

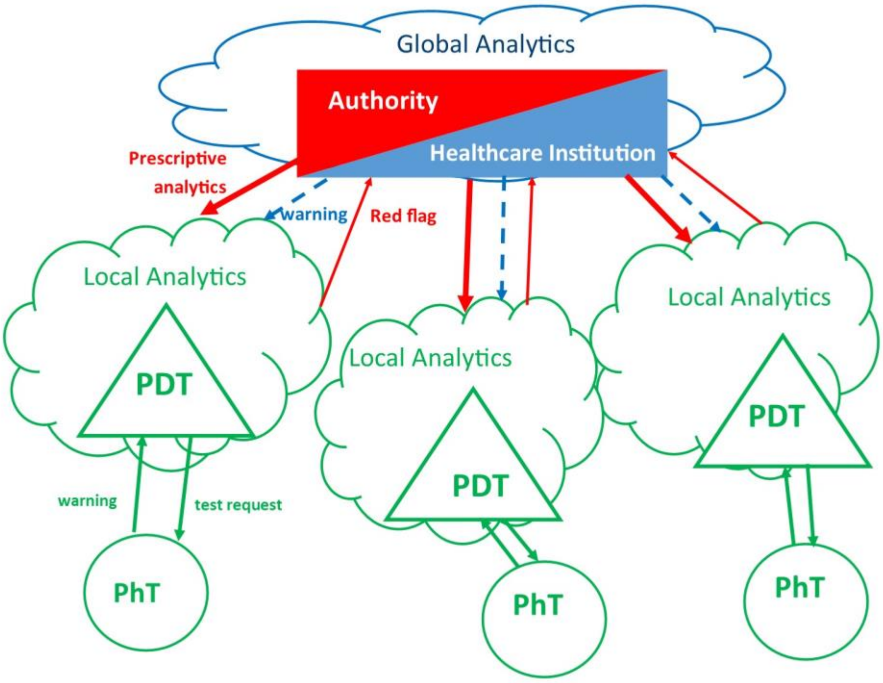

Figure 3 Schematic representation for using Personal Digital Twins (PDT) in epidemics control. The authority set the scene by requiring the PDT to monitor certain data patterns generated by its physical twin (PhT). Once a pattern is detected a red flag is generated and processed by the healthcare institution through global analytics taking into account several data streams. This might result in prescriptive actions, such as imposing testing or quarantine, and it contributes to increase the PDT/PhT local awareness. Nearby PDTs may interact based on this local awareness to ensure warning and appropriate behavior.

It is not just a different approach; it is one that has managed to get better results in healthcare terms (lower contagion rate and lower fatalities) and in economic terms (lower impact on business and lower expenditure in resources). The key to success was twofold: operating at a micro scale and taking immediate action. On the contrary, operating at macro scale becomes justifiable only when the epidemic is widespread. The Digital Transformation makes possible microscale operation in a most efficient way; the PDTs would be a natural way to implement this control.

To clarify the similarities and differences between the South Korean approach and a PDT approach, the two processes are outlined, noting that the aim is the same, to control the epidemic spread:

South Korea

1. A person shows symptoms that may point to a coronavirus infection. The person is tested and if positive starts the quarantine;
2. The person's contacts are discovered using technology: checking the smartphone movements, looking at presence of the person on security cameras feeds;
3. Data analytics are applied to the data on possible person's contacts to identify the probability of exposure of other persons and probability of contagion. These persons are tracked and tested. If positive they are put in quarantine and the sequence is repeated to find other possible exposures.

Personal Digital Twins

1. All PDTs are informed (*prescriptive analytics*) by the health institution of the need to report specific conditions, like presence of fever above 37.5 C, increased heartrate, and shortness of breath at rest as possible indicators of infection;
2. The health institution acquires the data by all PDTs and activates *global analytics* taking into account the location and the emergence of patterns. Based on this it signals to those PDTs presenting a suspicious pattern or with high probability of having suffered an exposure to request their physical twin (the person) to undergo a test;
3. The test result becomes both the trigger for action (quarantine) and a situational update that may lead to an updated reporting request and analytics to all PDTs affected;
4. In case of positivity the PDT is asked to report the history of contacts and movements of its physical twin;
5. Each PDT acquires an environmental awareness, dynamically updated both from the healthcare institution and through continuous communications to PDTs of nearby physical twins. This can result in warning signs like "DO NOT APPROACH/GET CLOSE" that can be generated in case of potential risk.

The advantages of the PDT approach are:

- Self generation of red flags, alerting the person of a possible critical situation;
- Extensive community and nationwide analytics with the possibility of detecting weak links, thus anticipating any significant spread;
- Lower cost;
- Increased focus, leading to less restrictions where there is very low risk and enforcement of stronger restrictions where risk is higher;
- Dynamic, just-in-time reactions to emerging situations;
- Increased personal awareness thus stimulating appropriate behavior.

It should also be noted that a PDT might be hijacked by the healthcare institution/government transforming it into a police enforcement agent. Any deviation of behavior by the physical twin from one instructed to the PDT can be immediately reported to the authority. This is clearly in the Orwell path and may raise more than one brow.

Smartphones have been used in China to trace movements, and an app was deployed in Wuhan to let people know if a person at high risk of being infected was getting close (in this case the smartphones in the vicinity started to ring). There are a number of initiatives, and more are likely to develop under the pressure of this ongoing epidemic, to develop apps that can track and create awareness, like CoEpi. The challenge is to balance societal benefits/needs with personal privacy and to be able to create an awareness that does not result in fear. All in all, it is a matter of creating trust.

As the epidemics in Italy and other countries continue to grow at an alarming rate, measures have been taken to limit the mobility of people, thus decreasing the chance of contact and spread. Yet, it seems that a few Italians keep moving around. How many?

Authorities in Milan requested the telecom operators to provide information on the mobility of cellphones (in an anonymous way), and it turned out that 40% of Milan cellphones moved around in a radius exceeding 300m. That is quite bigger than the average home where people have been requested to stay. On the bright side, 60% of people in Milan complied with the government order, but the problem is that 40% is way too large of a percentage to effectively contain the spread of the virus.

Notice the adjective "anonymous". European regulations (known as General Data Protection Regulation or GDPR) forbid the use of smartphone data to track people movement, unless there is an explicit consent. This was not an issue in China that actually leveraged smartphone software to impose and monitor movement restrictions on people (with a yellow/red code that was restricting/forbidding movement).

Figure 4 shows the timeline of contagion versus the awareness of contagion in the Hubei region of China. The availability of PDTs could have anticipated the detection of an incipient epidemic making containment actions much more effective.

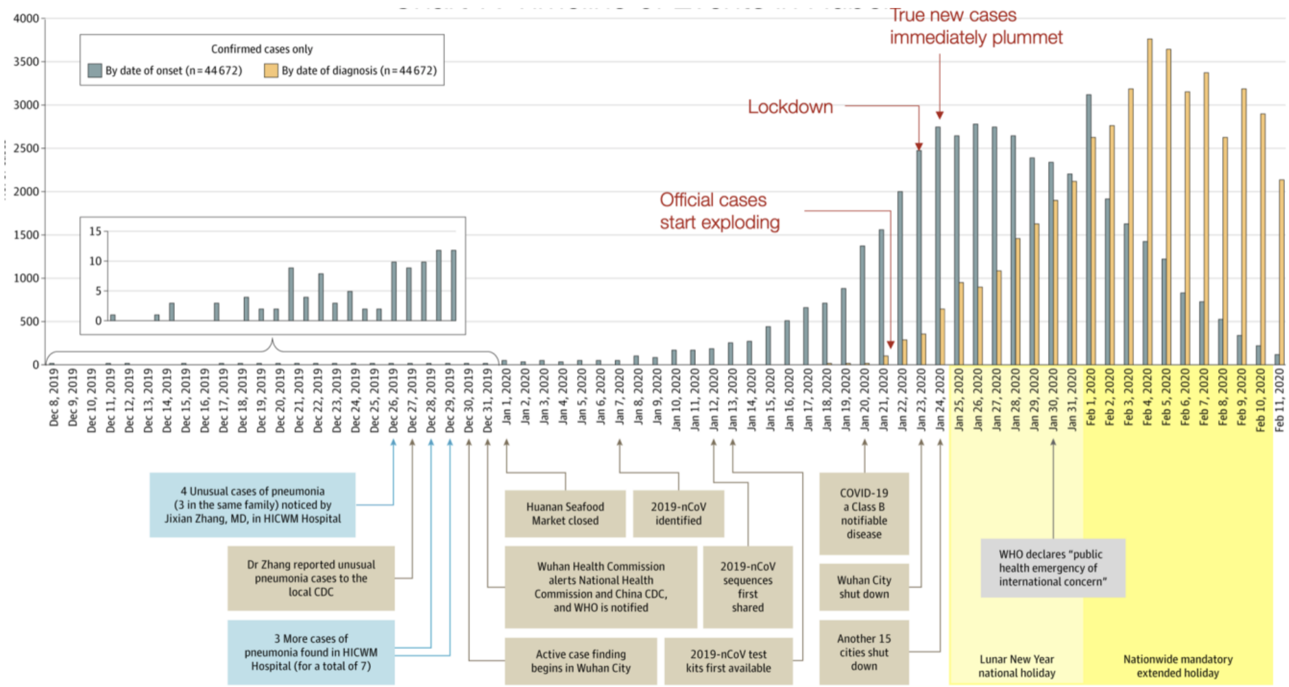

Figure 4 This graphic shows the timeline of contagion versus the awareness of contagion in the Hubei region of China. The availability of PDTs could have anticipated the detection of an incipient epidemic making containment actions much more effective. Image credit: Tomas Pueyo

PDTs are a must if we are serious in detecting and containing an epidemic because they offer the possibility of effective micro management, taking actions well before an epidemic can be detected by the authority using normal channels (i.e. by analyzing the growing number of cases reported through hospitals where people seek help). Evidence to this is provided in an analysis of the awareness growth in China.

The grey bars are based on interviews of sick patients, as reported to a hospital, on when their symptoms started. The orange bars show the growth of confirmed cases. The "delay" between the insurgence of symptoms and the awareness of contagion by the authority is obvious. The availability of a PDT, triggered by the healthcare institution to look for specific symptoms would have resulted in a much quicker detection of cases. More than that, a PDT can have a sensitivity that is far greater than the one of its physical twin. What can be disregarded initially by a person as an incipient cold or a mild flu can be identified by a PDT as a potential contagion, with the resulting raising of a red flag, demonstrating the importance of PDT in early awareness of a developing situation.

Privacy in the Presence of an Epidemic

The value of a PDT in early detection is not just for the specific COVID-19 case, rather it can be applied to many diseases and can be a very useful service to the person. This latitude of application can also provide the business incentive for their deployment.

The privacy aspect surfaces when the red flag is not directed to the physical twin, rather to the healthcare institution or (possibly more critical) to a third party (government, employer, insurance company). It is also a privacy issue when the red flag is communicated to other PDTs in the surrounding area so that they can inform their physical twin to take protective actions. In the Middle Ages, leprosy sufferers where required to ring a bell as they walked around to let other people know of their condition. The communication among PDTs of a potential health risk would be a modern way of warning and is likely to be rejected by many. This may well have a finger-pointing effect that can lead to harassment, bullying and further isolation of the infected person.

Several countries have legislation that requires a person suffering from a contagious disease to ensure other people are not put at harm but that is different than the obligation of declaring one's situation publicly. Also, there is quite a big latitude between black and white and many questions need to be considered. Where is the point when one has to inform another person or take evasive actions? If you have a cold should you behave in such a way to avoid spreading it around or otherwise you may get fined? Who is instructing the PDT on what to disclose, and then what is the entity receiving the data entitled to do? The general approach in Europe has been quite protective of personal data but there are now discussions, prompted by the current situation, that exceptions may be acceptable for the greater society good. A statement has been issued from the EDPD Chair on this very issue.

Privacy is an issue today. Even without a PDT there are ways to get personal data. Can the Municipality of Milan request the identity of the people moving around and follow up on them (possibly with fines and jail time)?

PDT Architecture 1: Personal Responsibility with No Sharing of Personal Data

One idea is to have the physical twin (the person) control the data and the decision on what to do with it and allow the healthcare institution only the permission to direct the PDT analysis in such a way to detect early signs, inform the person and provide the person the responsibility to act on them. It would then be a responsibility of the person to act according to the request of the institutions, failing that would be a crime. This could actually be an early detection service provided (and charged) to the person.

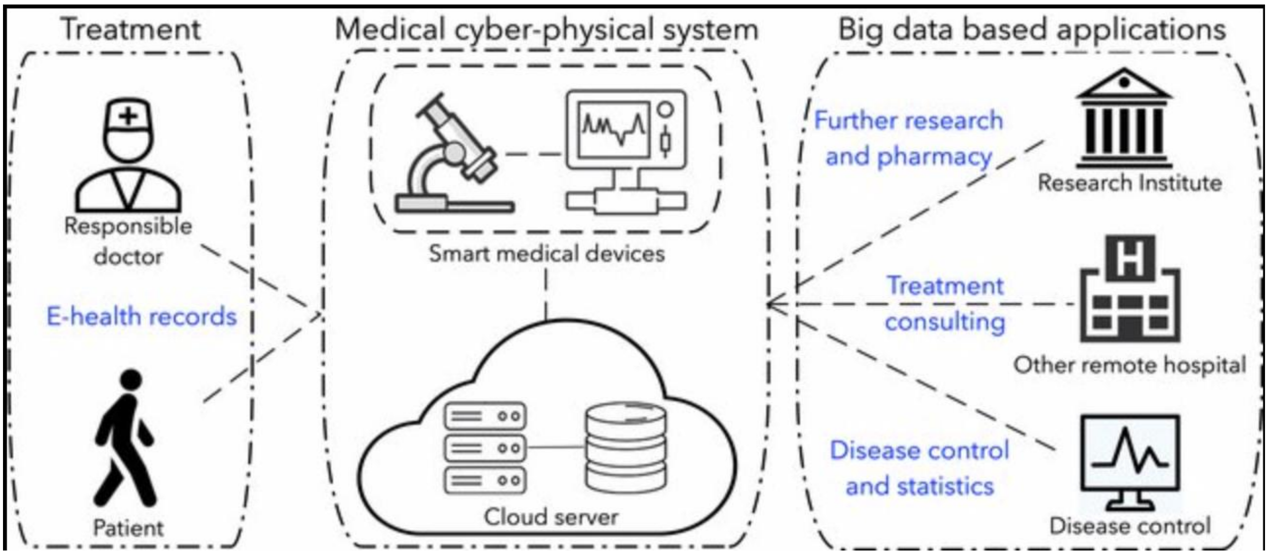

Figure 5 A schematic on medical cyber-physical systems showing the users, physicians and patients on the left, and the data set generated and shared on the right. A personal digital twin will provide an interface between the patient and the medical cyber-physical system (MCPS) or might be considered an integral part of the MCPS. Image credit: Han Qiu

There is a growth of Medical Cyber-Physical Systems (MCPS), basically resulting from the interconnection of several devices and systems in the medical landscape (see the section "Evolution of Medical Cyber-Physical Systems and Introduction to PDTs in Epidemic Control"). As these devices and systems become smarter, able to acquire, process and make sense of data (through artificial intelligence), they are playing a bigger role in healthcare decision making process. Part of these devices, by the way they deal with data, can be considered as being composed by a hard and a soft part, atoms and bits, and the soft part can be seen as a digital twin (perhaps in its infancy). The creation of a PDT would integrate well in this scheme, with the PDT acting as an interface between the person and the MCPS. An alternative way of seeing this would be to consider the PDT as an integral part of the MCPS. Whilst a first stage of PDT (simple mirroring of the person's health data) would better fit the interface role, more advanced stages (stage 3 and even more stage 4 where the interaction and convergence between the physical person and PDT is increased) would become part of the MCPS. In other words, the PDT may be one of the several components of a MCPS.

PDT Architecture 2: Macro Data Sharing Between PDT and Organizations

There are potential issues in the integration of PDTs at the societal level. While the first "architecture" provides for a clear separation between a PDT and the rest of the world (the personal space is fully preserved and it may be up to the PDT/rules of the game what to disclose), the second architecture no longer separates the PDT from the overall systems. On the contrary, the system is defined as the collection of all PDTs integrated with the healthcare system. The advantage of this architecture is the possibility to have each PDT contributing to the whole system; as a matter of fact the system status is the tuple including all PDT status among other things (like available healthcare resources). This would immediately allow detection of epidemic patterns as well as steer the counteraction in ways that balance resources with needs.

Independently of the architecture chosen, there are several advantages of creating PDTs and having them interacting with the healthcare system, such as:

- Effortless (or at least low cost) monitoring of people through their digital twins;

- Real time or almost real time feedback on actions taken on the physical persons through data collected and shared via their PDT;
- Red labelling or flagging those PDT that may be critical (either infected or getting too close to infected people) for specific monitoring;
- Development of services targeting infected people through their PDTs creating specific virtual communities;
- Providing direct access to support services, thus creating a triage in cyberspace;
- Supporting a more efficient use of resources based on the effective need dynamically monitored and balanced against competing needs and resources availability;

By having the PDTs part of the overall MCPS landscape (second architecture) it becomes possible to get the overall pulse of the healthcare system, segmented if needed in areas, and predict the possible evolution taking appropriate actions.

Actions in the healthcare industry are constrained by several factors, such as resource availability, cost, ethical and societal consideration. Visibility to the broad picture does not necessarily solve all issues but can provide a factual field for making informed decisions and, most important, to evaluate the result of those decisions and learn along the way.

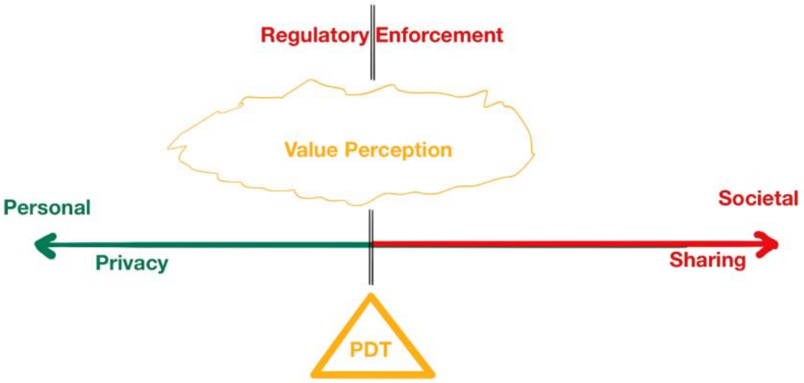

Figure 6 A very simple diagram to show the opposing interests of the personal sphere versus the societal one. On one side is privacy, on the other a full sharing of data and information. The perceived value differs over time and depends on the context (cultural and situational). Regulators impose a separation between personal and societal spheres that vary under different jurisdictions and geographical areas, as well as in response to societal challenges, as is the case with the current epidemics. It is essential to keep this enforced division within the range of value perceived by both society and people. Hence the importance of the lever of social communications to expand or modify the value perception. The PDT can act as the gateway between the personal and societal sphere.

During the coronavirus outbreak, there was ongoing discussion in Italian newspapers on the extent to which it is reasonable to protect privacy rather than monitoring people's movement to identify risks of contagion. Originally there was, at least in newspapers and in their reporting, a firm standing that "we are not like China or South Korea, we cannot, and we do not want to impose military force nor track single individual movements," however, the feeling has changed under the pressure of an epidemic that seems more and more difficult to contain with just friendly advices like "please stay home".

Other countries that were not in favor of locking down people and business are changing their policies. A macro approach was taken in the Western world to control the epidemics, however, South Korea took a micro approach, looking at individual people . So far it seems that the South Korean approach works better both in terms of containing the spread and in safeguarding business.

The trade-off is privacy and civil rights. However, our society is all about trade-offs between the personal rights and community rights, personal benefit/advantage and societal/community benefit/advantage. The big issue, on which there is not, and there cannot be, universal agreement is on where to draw the line between the personal and the societal space.

This is a matter of culture, and, of course, of political choice. It is, however, also a matter of situation, as it is becoming more and more clear in the present epidemic. Privacy and security are laying on the opposite vertex of a segment. The more we feel secure the less we are willing to give up our privacy; the more unsafe we feel, the more we are willing to compromise our privacy. Having security cameras looking at you might be annoying if you feel safe in a certain area. On the other hand, having security cameras in a place you perceive as dangerous will boost your confidence.

Technologies can impact these ideas significantly:

1. To start with, technology makes it easier to monitor individual people with an accuracy unheard of just in the last century. Our smartphone can provide plenty of data to pinpoint where we are and what we do. Through data analytics it is possible to extract meaning from different streams of data, correlating them. There is no need to deploy additional sensors. Smartphones are enough. Of course there are now so many security cameras around (there are an estimated 500,000 security cameras in London, in particular 408 CCTV in King's Cross St. Pancras monitoring 82 million people a year), whose feeds can be used, thanks to image recognition software, to provide further data on people, not just about where you are and where you go, also what you "feel". Face recognition is now able to detect moods. Add in credit card transactions, data generated from cars, information from social networks, and so on. Through data analytics and correlation, it becomes possible to create a map of close encounters, hence to pinpoint possible chains of contagion.
2. Technology can help to enforce policies, both in a soft way, by increasing awareness, and in a hard way by letting the authority identify and persecute those not complying with the rules. Notice that if people become aware that any single violation to a rule can be detected and persecuted the behavior will change. No one is willing to get fined or worse. In turns, this forced compliance generates a culture of compliance over time.
3. Technology can help in positioning the dividing line between privacy and societal interest in such a way that privacy can be protected, while up to a certain extent still meeting societal interest. This is where PDTs can and should play a role.

PDT Architecture 3: Macro and Micro Data Sharing Depending on Severity of Crisis

Two alternative architectures were previously discussed, one where data and interpretation of data are encapsulated in the PDT and another where the PDT is a transparent gateway conveying personal data to a "centralized" cloud where data analytics takes place. The following will explore in more detail how PDTs might be used for a flexible balance between privacy preservation and societal needs.

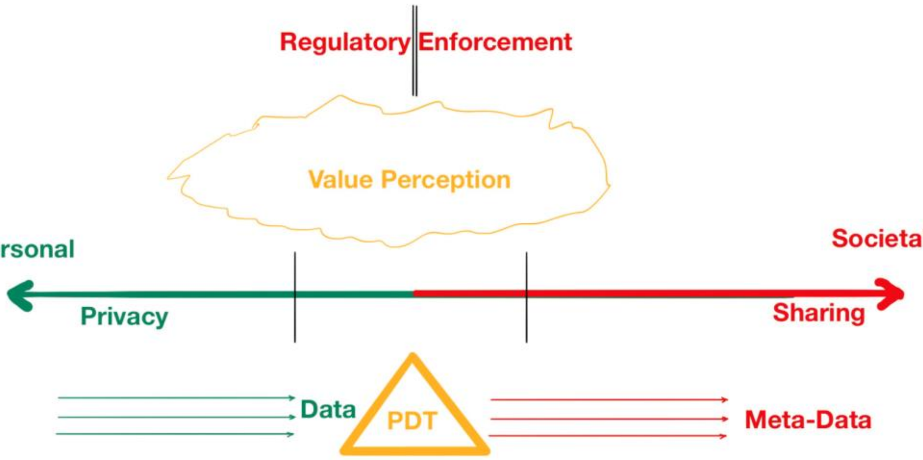

Figure 7 A PDT can act as a decoupler between the private and the societal spheres. Private data can be kept private, only releasing a set of metadata to the societal sphere. This creates a buffer area that can be dynamically defined and bounded by the regulator.

Earlier it was pointed out that there exists a fuzzy separation between the personal and the societal spheres, where personal privacy meets society's "need to know" and the fact that a PDT might serve as the gateway to the flow of information. The "fuzzyness" relates to the different perception of trade-offs but a regulatory body can determine a very precise boundary between the two. GDPR is an example of a legislator decision on where the societal/business sphere shall stop to preserve citizens/personal privacy. Yet, as pointed out in previously, exceptional situations cast a doubt on this boundary, claiming that broader disclosure would create a societal benefit that exceeds the value of private benefit.

The value of a PDT could be the possibility to create a privacy shield that while preserving privacy can disclose, at the same time, personal data in terms of information meeting societal needs. More than that, a PDT can become a tool for dynamically setting and modifying the boundary at a "personal" level. GDPR applies to all data of all people, and by using PDTs one could create a framework that applies, as it should, to all people but whose implementation is person and context specific.

As shown in Figure 7, the PDT is expanding the boundary between the personal and the societal spheres creating a sort of neutral space. Moreover, the "latitude" of this neutral space can be negotiated by the owner of the data (the physical twin) and the regulatory body (on behalf of societal needs). In other words, the two vertical black lines are no longer the ones imposed by the regulator (upper vertical black lines) but move apart to create a no-man land that is only subject to PDT algorithmic interpretation. Notice that in a way the personal sphere is reduced (left vertical line moves to restrict the personal privacy) and the societal sphere, the visibility/sharing of personal data with third parties (society, healthcare institutions, government) is also restricted. However, at the same time, metadata are generated and released by the PDT, meeting the needs of those third parties (more on this later).

Beginning with a normal situation where everything is fine, and there is no need to impose restrictions on people's privacy: the regulatory body has defined the boundary between private data, their disclosure and their usage. Part of this regulatory framework can also include the fact that each PDT shall detect some basic health parameters, and if some thresholds are reached an anonymous metadata shall be shared, such as "someone in this area has a spike of fever". These metadata are accrued by an authorized institution, like a healthcare national board, and analytics are performed taking into account several contextual data. This may lead, as an example, to the detection of an emerging flu and might result in some general warning to the population, or to a subset, like elderly in a specific affected geographical area, to consider vaccination and/or other measures. It may also result in a request to pharmaceutical companies to increase availability of the vaccine in that region.

Suppose now, that the analytics indicate an epidemic. Society, and in particular the healthcare management system, needs to know who might be affected, their whereabouts, and the encounters they have had and are having, in order to calculate the risk of contagion and take countermeasures. Anonymous data are no longer sufficient. On the other hand, just because there is an epidemic there is no need to track every single person, only the ones that might be affected directly or indirectly. Hence, upon the establishment that there is a surge of fever case that can be related to a specific virus whose effect on society is more dangerous than a normal flu virus, the healthcare board can suggest the regulatory body to release some privacy constraints, such as the identity of persons having specific symptoms and of those whom those persons have been in contact over the previous 3 days, as an example.

The regulatory body will take responsibility and will enforce a new framework on all PDTs within its jurisdiction to activate specific intelligence on the data owned (through algorithms provided by the healthcare institution, as an example). Based on the outcome of these algorithms, further data/metadata shall be disclosed to the healthcare institution, as well as, possibly, to other parties, like the city mayor or the police, to the extent of visibility required by each of those parties for carrying out the activities in their role. Notice that the physical twin, the person, will be immediately notified by her PDT of the new framework and of the amount of privacy restrictions imposed.

It is important to understand the flow and repartition of responsibility. The healthcare institution receives anonymized data so it cannot go back to the PDT generating a specific pattern of data. Rather, it must ask the regulatory body to update the global framework to which all PDTs have to comply so that those PDTs showing a specific pattern now will start reporting in a different way, possibly releasing some privacy constraints.

Notice that this decoupling of

- regulation (that is general and applies to all),
- implementation (that is specific and subject to conditions that are personalized, localized and contextualized),
- and data sharing (that is based on "need to know" principle)

maintains a high level of privacy where there is no need for personal data and restricts privacy where disclosure is needed but this restriction is limited in latitude and in to whom metadata are shared with (need-to-know). Besides, privacy is always subject to legislation and not to any third party wish/interest.

The use of blockchain can help in tracking the flow of metadata and protect privacy at a global level.

The healthcare institution can determine what should be done with the awareness generated by personal metadata and the related analytics. As an example, it can act to request the regulator to involve other PDTs that might have been at risk of contagion so that they also implement the new framework and disclose metadata to the healthcare organization, or may be on alert, and execute specific intelligence algorithms, to detect early signs of contagion. These PDTs may also be required to notify the potential danger to their physical twin and request quarantine. A PDT under quarantine operates within a different framework and accordingly might be required to report unlawful movements to the authority.

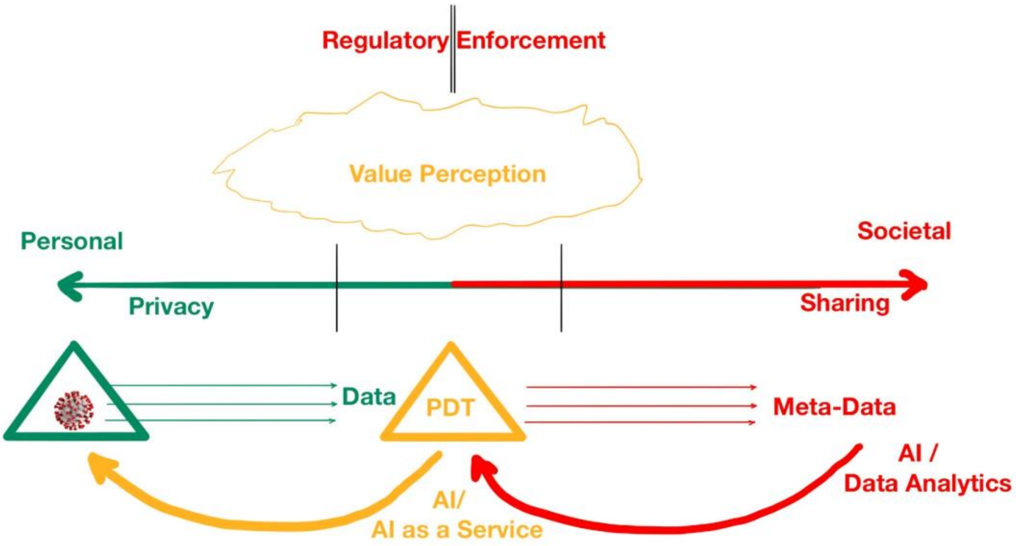

Figure 8 Metadata accrued from the PDT and data analytics/AI applied to them plus other streams of data can result in assessing the probability that a specific person has been infected and is becoming a contagion risk for others. This information can be sent to the PDT resulting in a change of state - like suspicious mode - and in the warning to the physical person of the risk situation. Notice that this "warning" can be delivered during all of the times that a potential risky situation arises, not just once. This can be much more effective since the warning can be situation related (an example might be, don't touch, use gloves, stay inside). The change of status of the PDT may activate a number of local data analyses invoking AI as a Service (not all AI computations may be feasible in a smartphone, for example, but the smartphone hosting the PDT can invoke AI as a Service from the network). This, in turn, will result in the generation of more meaningful metadata creating a continuous loop where the PDT can receive better information and advice from the healthcare institution.

A crucial aspect in epidemic control is the capability to involve a whole community and obtain compliance to the set of rules defined by the healthcare institutions and imposed by the government through a regulatory framework. The best way to achieve compliance is through awareness; if people understand what is at stake and their role in achieving a goal, they will be willing to comply and will also become agents of compliance. It should also be noted that different cultures show different levels of compliance but in general the better awareness the more compliance achieved.

Clearly, fines and other measures shall be used as a last straw, and a PDT may be required to detect the situations of non-compliance and report them back to the authority. This is clearly a very strong intrusion into privacy and may be objected in several cultural domains. This is probably a matter of the assessed danger determined at the government level and the acceptance by cultures that are not used to central control can be a function of the perceived danger by the population.

In Italy, as pointed out previously, there is a shift towards acceptance of restraining measures that affect individual privacy for societal benefit. As discussed previously, the adoption of PDTs could restrict these privacy limitations to those people that are at risk of spreading contagion and only for the time they represent a risk. This will be more acceptable, in general, since it is both restricted in application and in time. Furthermore, the release of non-compliance data will happen through a personal device not through an external –Big Brother- system.

The control approach deployed in several countries has leveraged smartphones. One example during the COVID-19 epidemic is an imposition in China to people living in the most affected area, Wuhan, to have an app that is connected to a system issuing permission to move about and tracking any movement. Another example of smartphone used in South Korea is the global monitoring of movements of cell phones through the cellular network, red flagging those belonging to a confirmed Covid-19 and tracing back other cell phones that got into vicinity, thus falling into a high contagion risk. The owners of those cell phones were notified and required to take a test. Both mechanisms proved, as far as we know so far, quite effective in containing the spread of the virus, once deployed.

During this time, the Italian Ministry of Innovation opened a call for proposal for:

- technologies and solutions for continuous tracing, alerting and immediate control of the level of exposure risk at personal level and
- evolution of the epidemic on the national territory with high sensitivity on specific location

Among the several requirements posed on the solicited proposal is the efficacy and friendliness in the identification procedure, authentication and care in the management of data ensuring personal awareness and compliance with personal data management.

Here one can see the attempt to kill two birds with one stone, getting data and preserving privacy. Clearly a compromise will be needed and the specific request to ensure awareness of data disclosure is a way to mitigate the loss of privacy.

Within the first 48 hours over 2,500 proposals[1] were submitted proving that technology is not really a stumbling block. Technology can currently trace and analyze data. Italy is not alone in looking at technology for tracing people's movement and assessing potential contagion risk.
There are a flurry of apps aiming at monitoring in real time, analyzing symptoms, providing guidance. On March 25th 2020, there were 20+ apps available on iPhone for these functions, when one month prior, there were none. It should also worth noting

PDT and AI-as-a-Service

A major issue in containing the epidemics is the identification of silent infections and super-spreaders. This is an area where PDT and AI as a Service can play a significant role. Many, if not most, of COVID-19 infected people perceive no tell-tale symptoms (like dyspnea) or very few symptoms easily ascribed to other mild conditions like a passing cold. These are carriers of a silent infection that can nevertheless spread the contagion. Other people are potential super-spreaders if they become infected because of the vast web of interaction they have as part of their working and/or social life. In both cases, it is essential to identify these potential moving bombs and make them aware of their possible infectivity. A PDT can detect subtle signs pointing to an infection and can, with the help of external data analytics, evaluate the risk level raising a red flag.

PDTs can help in the negotiation of the boundary between the personal sphere and the societal sphere creating an intermediate space where the overlapping of the two can be on the one hand regulated by an authority (under the advice of a healthcare institution) and on the other hand provide awareness on what is being disclosed and the benefit of disclosure in terms of feedback (like "you are likely at high risk of being infected/or infecting others). This assumed shared benefit derives from data correlation. The ones present in the PDT, detected by communication with the physical person (through sensors, mostly), are not sufficient, in general to derive a precise meaning. This is why the possibility to correlate with data coming from other PDTs, from the environment and the application of AI as a Service is so valuable at the personal level.

[1] As of April 8th, 371 proposals have reached the final examination stage, being evaluated under technical and societal parameters. Unfortunately, the intention of the Italian Government seems to make the adoption of the app "optional", i.e., each person will be free to download and enable it or not. This is likely to significantly limit its value since it is likely that it will only be used by those people that are already law abiding citizens.

However, the big issue is the credibility of such correlation. As shown in Figure 9 there are plenty of similarities or patterns that can be found in our life that have no interdependency whatsoever. The one shown in the graphic is clear to the point of being "laughable": clearly the trend on marriages in Kentucky cannot be related to the trend in people dropping out of a fishing boat and being killed. The two lines just happen to have a similar shape in the considered time window. There is no way to predict knowing the evolution of one what will be the evolution of the other.

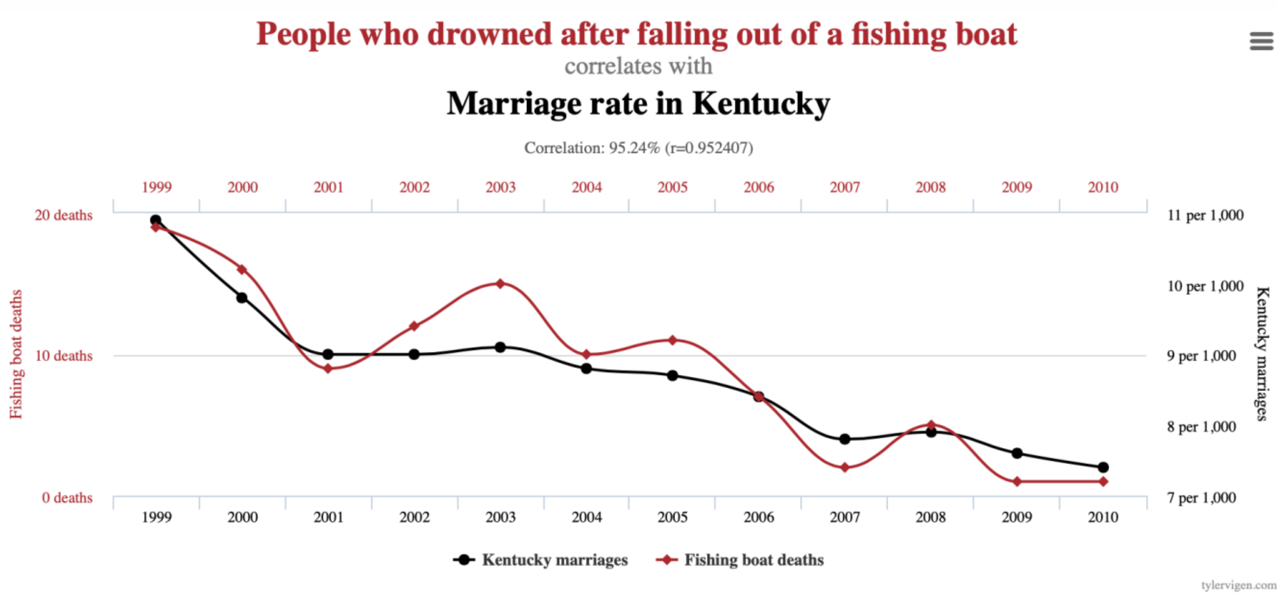

Figure 9 The two lines in the graphic represent (red) the number of people killed after falling out of a fishing boat and (black) the number of marriages (per thousand people) in Kentucky. The two lines seem to have similar shape, however, it would obviously be wrong to correlate the two. Several other examples of spurious correlation are presented in the book Spurious Correlations*.*

This same issue is faced today when scientists are trying to correlate epidemic patterns to possible causes. It is a fact that elderly people are experiencing more problems when infected. The question is why is that so? A first correlation is with the fact that as you get older you are more likely to have other pathologies, like diabetes, respiratory problems, or cancer, and indeed there are a significant number of deaths resulting from Multiple Organ Dysfunction Syndrome (MODS). In Italy the approach has been to count these deaths as COVID-19 fatalities. In Germany, on the contrary, it is considered that MODS is a general condition where COVID-19 increased the severity, but the resulting death is not attributed to COVID-19. That leads to quite different numbers in the statistics.

Recently, a study pointed out that for the virus to penetrate the cells, it attaches to a specific receptor, ACE2, whose numbers increase as result of hypertension antagonist drugs. Based on this, the conclusion should be that elderly people are more affected because most elderly people have been taking hypertension drugs for a longer period of time making them more susceptible to the virus attack. Others claim this is not necessarily true and anyhow discontinuing hypertension drugs will not decrease the ACE2 receptors in the short term but will expose those people to danger of hypertension; hence the advice: keep swallowing your hypertension pill.

These are just two examples of possible correlation emerging from data. The question is: are these causal correlations or are they just casual relations?

The adoption of PDT can combine the benefit of a macro view, where correlation takes place on multiple streams of data, with the local view, focusing on a specific person's data, both actual and historic. This can help in discriminating among true cause and effect relationships and simple coincidences. Indeed, the value of a PDT is based on the capability it has to retain the history of its physical twin. This history can lead, through machine learning, to a local understanding of that specific person. By coupling the always growing understanding of the specific person to contextual data and to metadata emerging from data correlation, it becomes possible to get a much more accurate understanding of what is going on "locally".

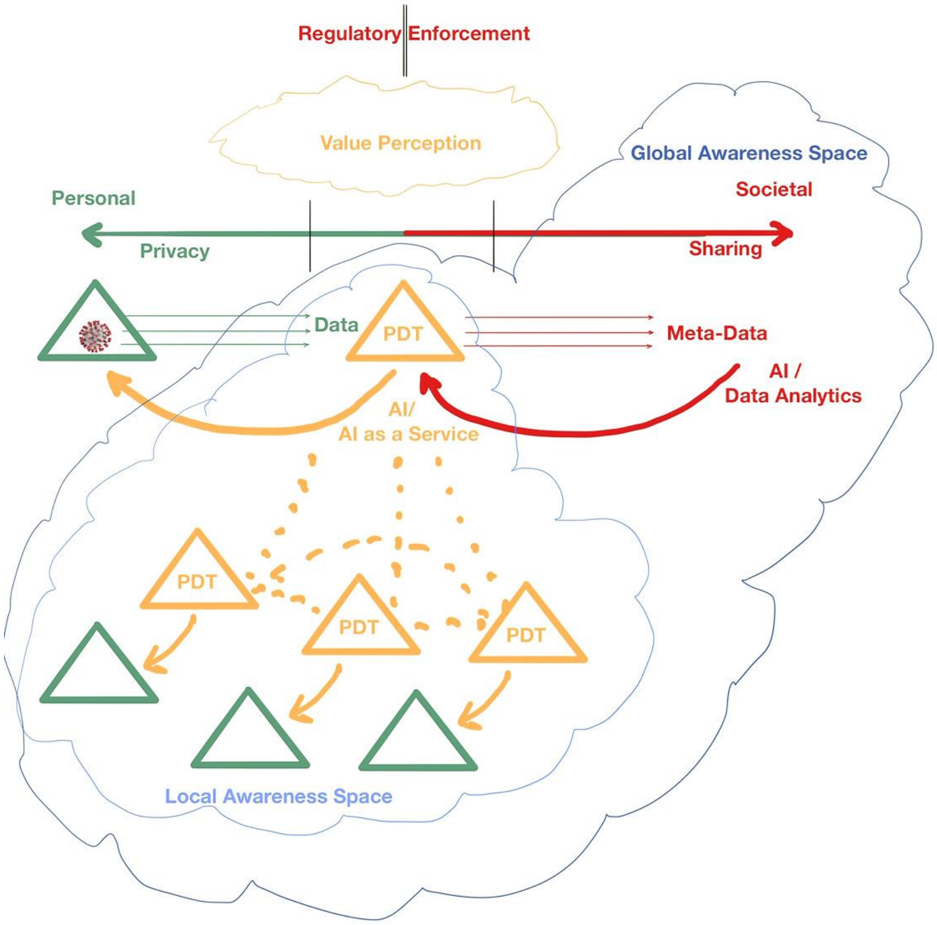

Figure 10 The PDT can play the role of an autonomous agent. As such it can exchange data and information with other PDTs in the geographical vicinity or in a functional vicinity (such as the PDTs of co-workers). This leads to the emergence of a local awareness space, versus the global awareness space. Notice that although the graphic represents the local awareness space as a subset of the global one, in reality the interactions among PDTs, being autonomous, may lead to a local awareness space that is only partially shared with the global one, thus enforcing visibility boundaries and preserving privacy. Whether this can be allowed by the authority is a different issue.

The evolution of Digital Twins is towards the acquisition of an autonomous role, i.e., a digital twin may start to act as an autonomous entity performing tasks on behalf of its physical twin. Although this is a natural evolution as digital twins shift from a static copy to a dynamic aggregation of services customized for its physical twin, there is also no doubt that this evolution changes the nature of the idea of a digital twin. New issues are coming up, like:

- the divergence of the digital twin from its physical twin. By taking up activities on behalf of its twin it accrues experiences that are unique to the digital twin and most likely unseen by its physical twin;
- taking autonomous actions, on behalf of its physical twin, raises the issue of accountability and responsibility;
- acting on behalf of the physical twin can play at several levels of responsibility. Consider the example of a turbine's digital twin. A digital twin detecting a potential issue in its physical turbine could autonomously decide that a replacement of

> lubricant is needed and can search the enterprise resources to locate the lubricant and the maintenance crew to execute the replacement. It can also interact with digital twins of other turbines to "understand" the impact of the lubricant used by different turbine versus the actual work and stress situation to make a decision on what could be the best choice for the replacement. This choice may not be available within the enterprise and the digital twin may "order" the appropriate lubricant from a seller. It is easy to see the complications of operating in an autonomous space.

Now, all of this gets even more awkward when we are dealing with a PDT. The example of the creation of a personal digital twin by UBS to clone[2] their Chief Economist is a point in case. The real Chief Economist, Daniel Kalt, can only talk to a very limited number of clients every day; by creating his PDT, UBS using IBM Watson, is able to share the knowledge of Daniel to a multitude of clients and in doing so, through clients' interactions, the PDT acquires information that is not available to Daniel. It can also analyze that information with thousands of data available on the Internet, learn from the correlation and quickly get "better" in some sense than Daniel. Now, this is a crucial point: what does it mean to get better than Daniel? In general, it would mean for the client to receive more actionable and accurate advice resulting in a gain (or limiting losses), and for the bank to make more money from consultancy services (the two may not necessarily go hand in hand). Would the real Daniel, if he was aware of all the info available to his PDT react in the same way? If not, how can we say that Daniel is accountable for his PDT? And if he is not accountable, who is? IBM Watson, UBS? (Don't worry on the specifics, UBS and IBM have plenty of lawyers that crafted documents for limiting their liability).

Getting back to the use of PDTs in the context of epidemics, it is natural that having an autonomous PDT could in principle increase its effectiveness. The PDT could interact with other PDTs creating a local awareness of what is going on, without having to rely on information coming from "above". This might also be a way to preserve a higher level of privacy, since data would only be shared among a limited number of PDTs and that information will fade away as it will no longer be needed. At the same time, security issues, that are present in all aspects of data use and processing in the frameworks previously discussed, become crucial. How can a PDT trust another PDT and to what extent can data be shared with confidence? Blockchain and reputation mechanisms may be of help.

Once a person has overcome COVID-19, been infected and recovered, she is no longer at risk of re-infection (at least this seems the general scientific agreement) nor can she be a source of contagion for others (for a certain period of time). Hence her PDT can advise other PDTs that may get, or plan to get, in her vicinity that there is no risk. Conversely, a person that has been in the vicinity of another person that turned out positive to Covid-19 can have his PDT signal the potential risk to others. Again, this approach can better preserve privacy than sharing this information at global level.

Decentralised Epidemic Alert Systems Leveraging on Digital Twins

By Juuso Autiosalo

Schematic showing the basic principles of decentralized epidemic alert system (DEAS). Decentralized systems run on multiple computers and leverage distributed ledger technologies (DLT), such as blockchains. Image credit: Juuso Autiosalo

Distributed ledger technologies (DLT), such as blockchains, offer a new kind of decentralized solution for preserving transparency and privacy in internet-based systems. Transparency is

[2] https://fortune.com/2018/07/05/ubs-digital-clone-chief-economist-daniel-kalt/

enabled by the public nature of transactions in a distributed ledger (an implementation of DLT), whereas privacy is ensured by using pseudonyms (usually public cryptographic keys) for the transactions. An apparent issue with privacy is that all transactions signed with single public-private key pair can be associated to each other.

The security of DLT based solutions can be considered superior to traditional centralized systems thanks to their innate use of cryptographic hashing with private-public key pairs and the distributed means of ensuring information.

Blockchain-based cryptocurrencies are the common example of how DLTs are leveraged currently, but distributed ledgers (implementations of DLT) can in fact be assembled in a wide variety of ways. Ledgers can trigger actions based on predefined rules (i.e., smart contracts), the mining capacity of participants can be used for meaningful computing instead of calculating trivial math problems, and a ledger can even be controlled by a trusted administrator. (It is negotiable if an administered ledger should be called a distributed ledger, but the technology can be used this way nevertheless.) Furthermore, a distributed ledger doesn't actually need a cryptocurrency; it is just a convenient way of ensuring computing and storage capacity in a truly distributed ledger. If computing and storage is provided some other way, a currency is not needed. For example, one or more governmental bodies can ensure the computing and storage capacity of a ledger, which might be the viable solution for the system presented in this post.

Decentralized epidemic alert system (DEAS) resides in a DLT-based cloud, catering the needs of the real world through via digital twins (DT). The DEAS cloud communicates with the personal digital twins (PDT) of people, and with the DTs of other stakeholder organizations, namely a hospital, a health organization and a governmental body. The PDT of a medical doctor has a special role in DEAS.

The main function of DEAS is to analyze the health information coming from the PDTs of the people. DEAS looks for possible signs of health problems and upcoming epidemics. The PDTs can analyze health data locally or send raw health data to the DEAS. (The health data will include information such as heartbeat readings from a smartwatch and/or body temperature from a tracker ring.) The analysis algorithms can detect an anomaly in health data, implying a possible virus infection. Accompanied with location data, DEAS can detect clusters of multiple infections and detect an upcoming epidemic before it even starts. Detection is performed by machine learning algorithms that actively monitor people that get sick around the same time, backtracking the location and searching for possible common locations. This way, the origin of the virus infection is tracked automatically, and the potentially contaminated people are notified via their PDTs. DEAS also generates reports on the data, providing crucial information to decision makers.

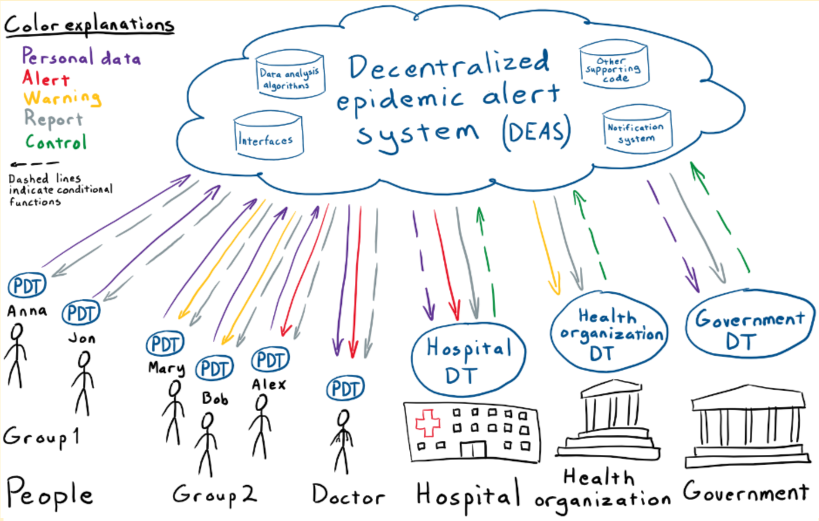

Figure 11 Schematic showing the basic principles of decentralized epidemic alert system (DEAS). Decentralized systems run on multiple computers and leverage distributed ledger technologies (DLT), such as blockchains. Image credit: Juuso Autiosalo

Error! Reference source not found.portrays a simple example case in which one person is potentially infected during a known pandemic outbreak, accompanied with consequential notifications and other information flow. The members of Group 1, Anna and Jon, have been practicing social distancing, and their PDTs are sending their personal data to the DEAS system. In return of sharing the data, Anna and Jon can rely on the knowledge that they will be notified by DEAS in case of nearby infection. If they wish, they can also review the public situation reports generated by DEAS.

Contrastingly, Group 2 is in the middle of an unfortunate situation. The DEAS algorithm has detected a potential infection for Alex, who is sent an alert about the detection. The other Group 2 members Mary and Bob receive a warning. The personal doctor of Alex also receives an alert, along with the personal heath data from the PDT of Alex that the doctor can use to pre-analyze the situation.

The alert is also sent to the nearby hospital that can access the personal data of Alex if needed and make the required preparations. The hospital is up to date on the underlying epidemic situation thanks to the reports provided by DEAS, which are also sent to a health organization and a governmental body. The health organization receives warnings of suspected epidemic outbreaks, and the government may receive personal data if enforcement procedures are required to keep the pandemic in control. Each of the three organizations may also be granted various types of control privileges to DEAS.

The reason why the organizations also have DTs is to simplify the information interfaces of the system; this way DEAS only needs to communicate with entities that follow the same DT data exchange standards (which are still to be defined).

The contents of DEAS consist of three main software components and other supporting code, which are all implemented as open source. Interfaces convey data in and out of the system, data analysis algorithms crunch the numbers to draw conclusions, and the notification system alerts the stakeholders. Open source enables transparency, creating trust for the system.

There are still open questions to be answered: Should people be able to limit their data sharing? How can the system be supported economically? What kind of DLT implementation should DEAS be based on? Who manages the updates of the system? Will DEAS be a national or global system? Some of these questions need decisions by leaders of the world, whereas others may only be solved on the go.

DEAS presents a novel way for decentralized pandemic prevention. There are known problems with DLTs, but they are being solved by numerous researchers and practitioners. The whole field is developing fast, enabling new kinds of implementations. Implementing the DEAS system described here is becoming increasingly trivial; we just need to make the initial decisions and start building it.

Economic Impact of Epidemics

Thus far, this white paper has not addressed the economic impact of epidemics, although it is and will remain, a crucial component. Notice that the means used to contain the epidemics can lead to very different impact on the economy, both in the short and in the medium term. A complete lock down of activities and business can dramatically affect the GDP but it will do so for a very short time, resulting in a V shape economy: quick decrease of GDP and rebound once the lock down is over (see Figure 12). A more restricted lock down can have a lower impact on GDP but extends the effect over a longer period of time, resulting in a U shape economy (the down period gets long). A quick lock down with a too early release may lead to the need of further lock down, as the contagion picks up again after the release of restrictions, resulting in a W shape economy where the decrease of GDP and rebound is followed by a further decrease. This can relieve a bit liquidity issues (since the business can restart after a brief period) but can generate a gloomy perception of the long term success in fighting the epidemic that in turn can have very negative impact on the market (since the market is all about expectation). Finally, a strong lock down or a prolonged weaker lock down can lead to stagflation where the production collapse leads to massive unemployment and is no longer able to meet demand, creating high inflation that makes situation even worse. This is the L shape economy where the difficulty is not liquidity but solvability of enterprises.

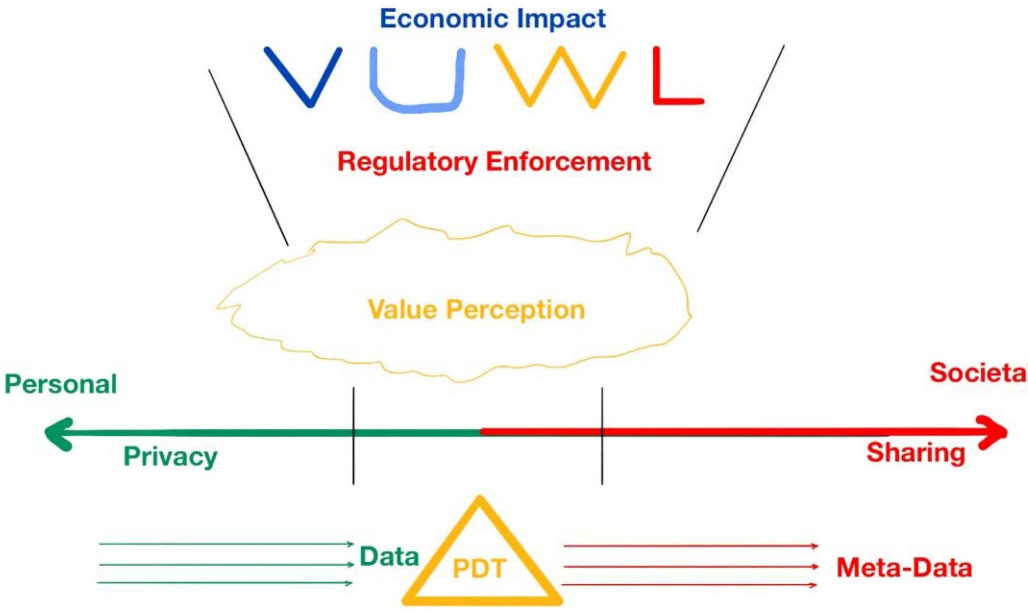

Figure 12 The epidemic impact is obviously significant onto the economy and the method chosen to contain the epidemic has different implications on the economy. In a time of economic crisis, it is useful to make reference to the four evolution trends of the crisis: V, U, W, and L. The use of PDTs can help in the implementation of the containment approach but also from the impact on the economy, providing better monitoring and flexibility in executing and correcting the execution by finely tuning the framework enforced by the authority.

Obviously, investment in technology and framework to get prepared is far less costly than having to manage a pandemic and, even more, its consequences.

The questions are: how far are we from the scenario presented and what does it take to turn it into reality? There are obviously several other possible solutions, and it is actually likely that several different implementations will take place in different parts of the world. This is acceptable as long as a minimum of interoperability can be ensured.

1. The basic technologies are already available, including smartphones, a pervasive communications network, data analytics, and artificial intelligence. More than that. These are not lab technologies; these are widespread and most people see them as part of their everyday life, especially smartphones and communications. Data analytics and AI are seen as features and services, like image recognition, speech recognition, language translation, stock market forecast and so on. Digital twins are a reality in several industries, particularly in manufacturing. They are just starting to be considered in the healthcare domain, and it will take a few more years to see PDTs adoption widespread.
   Data Analytics and Artificial Intelligence today are partly overlapping since AI requires huge data volume to perform. Having AI embedded in a PDT, not invoked as a service, requires a significant progress in research that possibly may be achieved by the ongoing brain research as well as synaptic chips evolution. The important point, however, is that we do not need to have a full implementation of a PDT to make it useful. Even a scaled down version, like the one that can be implemented today, could be useful.
2. Although current technology is not perfect it is sufficient to pursue a PDT based approach to epidemic management. Research will surely make better technology available with this decade but what is really needed is a regulatory framework supporting it. The use of data from smartphones, credit card transactions, and security cameras should not be seen as a temporary measure that can be taken just by releasing privacy rights. Rather it should be a starting point that has to lead to a framework where data accessibility and privacy are both valued. A first shell of PDT can be embedded in all smartphone as part of the OS, but data shall be managed not by the OS providers but by apps in the smartphone under the control of the smartphone owner. A back up copy of the PDT shall reside in the cloud, by a trusty provider, but the access to data shall be under the control of the data owner. Policy/authority may impose the (temporary) disclosure of data to specific authorized parties (like healthcare institutions or police), and these will need to be responsible on the use of such data and derived information. Each person should have the option to opt out, i.e., do not allow the access to his data but this can be open to prosecution by authority. It is like today where we are free to buy a Lamborghini, potentially pushing it to 300kmh but at the same time we can be prosecuted for doing that. This scheme may be thwarted in an authoritarian political framework leading to a Big Brother situation, however this should not be a reason to not foster a PDT approach since there are Big Brother regimes even without PDTs.

There are clearly several possible architectures supporting a PDT approach. Having the PDT embedded in a personal device, like the smartphone, can give a better perception of control than having it somewhere in cyberspace. Security aspects are obviously crucial to the acceptance of this data based approach, and there are already many hacking and malicious attack going on as outlined in the section "COVID-19, New "BlackNet" Cyber Threats Raises Worldwide Phishing Attacks to +350%".

Additionally, cultural aspects play a most important role in the acceptance and effectiveness of a PDT as discussed in the section "COVID-19 is a Perfect Storm Accelerating Digital Transformation Societal Waves".

There is another question to consider: should PDT be pursued in the healthcare domain and used for epidemics detection, monitoring, and containment? As pointed out earlier, there is a trade-off between the personal sphere and the social sphere. Privacy in many cultures is considered as a basic right of the person and PDTs, if used as gateways to access personal data by third parties, can be perceived as a threat. In Italy and other Western countries, under the pressure of the expanding epidemic, several constituencies have agreed that privacy rights can be suspended as long as an emergency situation remains. Tracking of personal movement using technology, mostly leveraging smartphones since they are so prevalent, has been allowed by several authorities. However, and here is the catch, this will no longer be allowed once the emergency disappears. This means that technology can be used to contain the spread once the situation has become very bad but not to detect and stop the spread when it is in the initial phases.

This is why the adoption of PDTs along the lines described would be important. PDTs can be used in normal situations to anonymously allow the detection of an incipient epidemic (data gathered from their physical twins can be analyzed and correlated) and if and when this happens the healthcare institutions can request the government authority to change the PDTs framework as needed to provide more data to pinpoint persons infected and those at risk. If needed the framework can be extended to monitor compliance.

This approach can manage and balance the personal rights with those of the community. The PDT technology has the potential of supporting this flexibility.

With regards to the economic impact, as briefly outlined, the approach to epidemic containment leads to different economic outcomes. Using PDTs, it is possible to finely tune the containment actions, making them both more localized and effective. More importantly, the use of PDTs may provide a much better pulse of the situation making adjustments possible in real time, thus decreasing the impact on the economy and steering those impacts towards the intended direction (sometimes a W shape may be better than a V or U shape, but an L shape economy is not desired at all). So far we have seen countries selecting one or another policy, all basically constrained by the reality of the healthcare resources available. China clearly went for a V shape, most Western Countries aimed for a U or W shape approach.

Impact of PDTs on Evolution of Epidemic

The adoption of PDTs will provide much more flexibility in decisions and much better visibility on trends. Data analytics used to define the PDTs framework can take these economic considerations and policies into account.

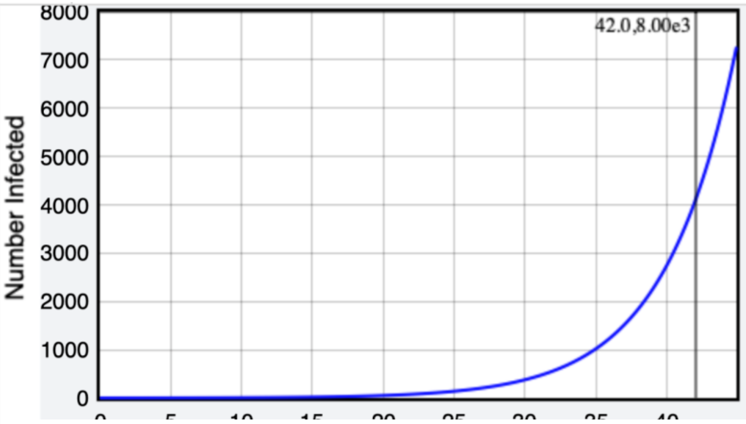

Figure 13 Exponential growth curves representing real events are limited in time. In their first phase, epidemics show an exponential curve whose derivative (the slop which is the ratio of progression) depends on the infectivity of the pathogen and the environmental limits to its infectiveness. Image credit: MIT Media Lab

In mathematical terms, an infection is characterised basically by four factors:

1. The basic Reproductive Number, R_0, indicating the number of people that would be infected by a single person, assuming they are all susceptible of infection (e.g., not vaccinated). The greater this number the steeper the curve, i.e., the faster the propagation of the infection (if the value is lower than 1 the epidemic will not

spread). For COVID-19 the consensus is on an R_0 between 2 and 2.5, with an accepted figure of 2.2 in Western countries; although in some regions we have seen the value of R_0 between 3 and 3.5. The formula governing this growth is

$$(\partial N/\partial t)=R_0N$$

where N is the number of people in the considered set and t is the time.
The spreading time depends on factors like the time from being infected to the time of becoming contagious.

2. The availability of people to be infected. It may seem obvious but it seems many miss this point. There are not, in the real world, an unlimited number of people to infect; there are about 7 billion on the planet and once they are all infected that's it. This means that the epidemic will stop once all have been infected. However, the infection curve is not like the one shown in the first diagram where at the top, once all people have got infected it simply plummet to zero. Epidemics are not following an exponential curve, rather a logistic curve (see Figure 14).

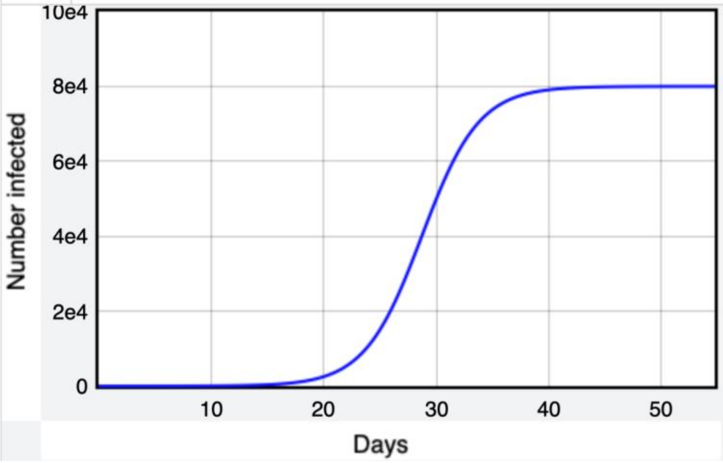

Figure 14 The real growth curve in an epidemic has an S shape. As more people get infected there are less people available to be infected, hence the R_0 value decreases. In other terms, a contagious person is less likely to find other people in the vicinity to infect. Image credit: MIT Media Lab

This formula governs the logistic curve:

$$(\partial N/\partial t)=R_0(1- N/N_{max})N$$

where N is the number of infected people and N_{max} is the total number of people. In the first phase of the infection N is small hence the ratio N/N_{max} is close to zero, and it follows an exponential curve with R0 as a factor. As the number of people infected grows (and these people can no longer be infected) the ratio between number of people that can no longer be infected versus total number of people approaches 1, hence the value of R_0 decreases, approaching zero. Thus it follows the S curve.

3. The generation time, that is, the time it takes from one person being infected to the time it infects another. For example, in measles the generation time is 12 days; for the flu, it is around 2 days. For COVID-19 current studies point to a 4-day generation time. Unfortunately, the incubation time (the time from infection to the appearance of symptoms) is longer, around 5 days, meaning that there is a period when a person feels healthy and yet can infect others (1 day). Notice that these numbers reflect a median value, with ranges much broader, up to 10 days for incubation time, meaning that the window of infectiousness without apparent symptoms can be large.

4. Influencing factors (like environmental ones) that can change the R_0 over time (leading to different *effective reproductive number*). This leads to a further expansion of the logistic formula. Several viruses, like the flu, are susceptible to climate and temperature, becoming more virulent during winter and relenting in summer. This leads to a sequence of infection waves that are still regulated by the logistic formula. Since the second wave is acting on a smaller N_{max} (those infected by the first wave will no longer be part of the target population), the number of affected people will be lower. Notice, however, that for a given virus, all conditions being equal, the total number of people infected in case of a single wave hitting the whole population or a subsequent series of waves hitting decreasing subsets of the

total population will result in the same number of infected. In other words, multiple waves just spread the infection over a longer period of time, they do not change the number of people affected.

The *effect* of an epidemic on people's health depends on the severity of the disease (often varying in different demographics, in COVID-19 the elderly are more compromised) and on the possibility of a cure. This, in turn, depends on the existence of curing protocols and their availability. As an example, in the current epidemic, most people presenting severe symptoms require the availability of ICU and respiratory support. Since these are in a limited number, healthcare institutions and governments are deploying measures to spread the number of people needing this kind of support over time, to make sure that for any new person requiring an ICU bed there is a person recovering that can be dismissed from the ICU. That is the peak of epidemics that can be sustained by healthcare resources. Going over this peak means some people cannot be assisted.

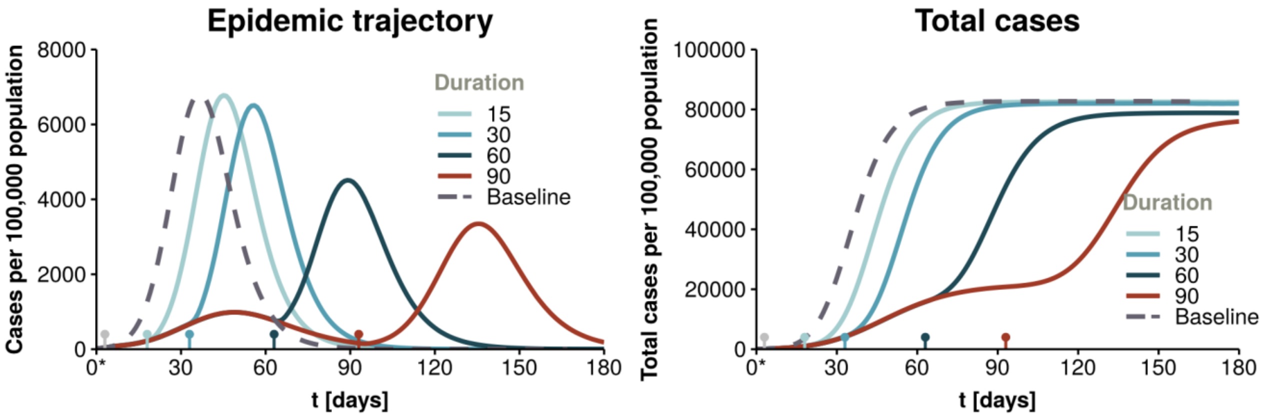

Figure 15 Epidemic trajectory when constrained by some external factors limiting its spread, its R_0. The bell shape remains but it gets lower and larger at the base as R_0 decreases. Notice how in the end they all lead to the same number of infected people. Image credit: MIT Media Lab

Controlling the spread is therefore crucial to decrease the strain on resources, as well as to provide time to find better ways to relieve symptoms, eventually hoping that a vaccine will be found which will decrease the N_{max}, the number of people that can be affected (leading as a side effect, logistic curve, to a herd immunity).

PDTs can affect the various parameters:

1. PDTs can increase the single individual awareness, increasing safe behavior. This will not hold true for all people; some people will consider themselves as superman, above the virus, and will keep an unsafe behavior. However, a significant portion of the people will take preventative actions to avoid being infected. Statistically this is equivalent to a reduction in R_0, and any reduction in R_0 leads to a slower infection rate (the steepness of the curve in the first graphic decreases) giving more time to the healthcare system to take countermeasures (increase the number of resources to support the infected with severe symptoms) and making resources available to more people over a longer period of time.
2. PDTs can provide timely information on the status of their physical twin, including the probability of being contagious, and data analytics on communities of individuals can identify individuals that might have been infected. This generates both individual awareness (R_0 - see previous point) and community awareness with imposition of lock down measures by the authority (following direction from the healthcare institutions) leading to a decrease of the N_{max} and thus to an

acceleration of the shape change in the S curve (logistic curve). In other words, data analytics made possible by PDTs can help in containment measures. The probability or risk of being infected can correlate to the probability of being infectious, and as noted before this would be crucial in containment. Locking down a community with possibly infectious members helps greatly in overall containment. Besides, data analytics can also point out the probability of infection in different communities providing the authority the means to locking down those sets of individuals that are not just most likely to infect others but also activate the lock down in a way the minimizes cross contagion in a community. This kind of information is also crucial at the time of relenting lock down measures. This can be done in a selective way to decrease risk of an uptake of the contagion, or at least to minimize its effect.

3. PDTs can release data on the locations of the physical twins that have been identified as potentially contagious (or definitely contagious having been found positive by a test) that are not allowed out of quarantine zones. This data can be used by the authority to enforce quarantine, getting police to pick up the person and persecuting him. Notice that the certainty of being caught and persecuted is a strong deterrent to avoid this kind of behavior in the first place. This contributes to decrease the R_0.

These three ways of using PDTs have been presented in a growing intensity of privacy reduction.

- The first is basically preserving all individual privacy. It is only the physical twin that uses his data, and the analytics are derived from the whole data space but without becoming aware of individual people's data, thus not infringing on people’s privacy.
- The second act is on a community level and does not disclose personal data although the fact of becoming part of a community implies the sharing of commonalities and therefore some personal data.
- The third provides full exposure of certain personal data to the authority, like health data and location or activity.

Clearly, the above assumes the existence of a security scheme that protects data from non-authorized parties.

There is a fourth contribution from PDTs: intelligence emergence. The data harvested at the micro scale by each PDT can be used by machine learning and data analytics to correlate into actions taken (like creation of locked down community, use of pharmaceutical protocols such as masks, social distancing, broadcasted warnings) and to assess their effectiveness. These data do not need to be visible in the micro scale; only the result of the analytics is of importance, hence privacy can be preserved.

The machine learning part is very important in the understanding of the ongoing epidemics and the effectiveness of countermeasures and in proactively preparing for new epidemics interception and management.

It is imperative to take the opportunity offered by this emergence to create a safety net to remain in place so that the impact of future possible epidemic can be minimized. The idea of implementing some patches under the strain of the current situation and then removing them once the emergency is over is not a good approach. As mentioned, a PDT approach can offer a way to balance private desire of privacy with societal need-to-know.

COVID-19 is a Perfect Storm Accelerating Digital Transformation Societal Waves

By Derrick de Kerchove

Regarding both the viral threat of COVID-19 and the haphazard regulatory reactions of the world's governments, the virus is being greeted by a 'perfect storm'. The still recent globalization of the planet, the almost suffocating network of communications available and the need for a reboot of human culture in the face of climate change present ideal conditions to a form of disease entirely predicated on communication. Coronavirus is a malady of communication, a meeting point – and collaboration – of biological and digital progression patterns. It is also a civilizational climax, marking a point of no return in a transition already begun.

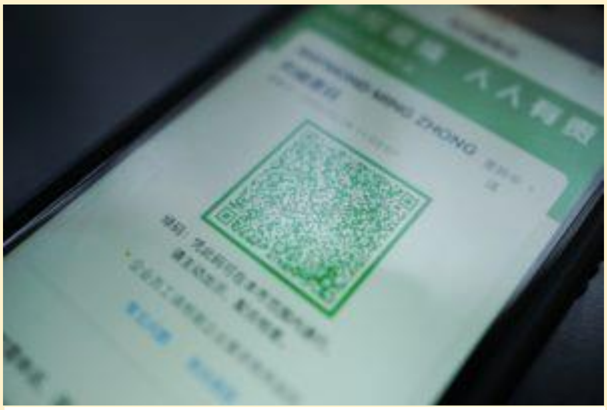

Figure 16 The green QR code shown on a smartphone in China by the Alipay app, certifying permission to move around. Yellow would restrict movement; red would forbid any movement outside the person's premises. The app automatically notifies authorities in case of non-compliance. Image credit: Raymond Zhong, NY Times

We are at a critical (epochal) moment of the so-called digital transformation (DT). When speaking about the digital transformation, it's typically all about business, production, delivery and management. Some modest considerations are given to social life and politics, but basically nothing close to its reality, namely a radical and profound reset of the individual human and society. We would really like to think that the DT is serving us, but it's becoming clear that the opposite is true: we are serving the DT. So the not-so-honorable behavior of a few Italians on the run is not only a selfish life-saving strategy but a typical and unconscious resistance to the incoming digital tsunami that eliminates all forms of autonomy, beginning with privacy. But it's no use. Long before we suffer the ignominious consequences of tracking and naming fleeing individuals and publicly punishing and shaming them, both inevitable if the tendency to escape continues, we have already abandoned the fight and conceded our privacy in myriad ways. For the government to threaten first and then act on the removal of privacy protection in the event of a national emergency is just a technicality. If we don't have the COVID-19 challenge under reliable control by midsummer 2020 (and perhaps earlier), starting with Italy, all Western governments including the Libertarian United States of America will put the entire population under surveillance in the way Chinese (and soon Indian) authorities have done. The police will most likely need to start tracking even offline all the smart people who have thrown away their phones thinking they can thus escape surveillance. It shouldn't quite be said that the coronavirus pandemic was part of some sort of spontaneous self-organizing DT strategy to accelerate its conquest of humanity, but it certainly came in handy to ensure that it happens. There's no escaping it. Moreover, the brazen official response of states to block people and force social distances is precisely the message of the DT: that is to dematerialize goods and services, yes, but at the same time remove autonomy, immobilize the population, and increase communication to unprecedented levels.

And, by the same occasion the social distancing program is reorganizing our sensory lives reducing the need for touch, for travel, for transport and turning us all symbolically legless. Jean Baudrillard said it in 1976 in his L'échange symbolique et la mort when he suggested that staying at home at watching TV would make most people physically and mentally crippled. At the time, readers just laughed. Today the situation is much more advanced; as Lev Manovich observed, the balance between physical and virtual life in front of some screen, already threatened since long ago, is lost in favor of the virtual where we are bound to spend more and more time porting our professional engagements on screens. There, of course, we will be tracked and catalogued in databases. Once the last defense of our private being has been removed, no government will return it. Goodbye to GDPR and other fantasies of democracy! The image of the post-coronavirus environment is predictable, a Kafkaesque metamorphosis, not into a cockroach, but into myriads of data and algorithm producers in the global hive.

COVID-19, New "BlackNet" Cyber Threats Raises Worldwide Phishing Attacks to +350%

By Francesco Flammini

While locked down at home to work, study and shop online, connection times increase and the chances of being a victim of cyber threats also increase. Recent reports highlight the danger of an explosion of phishing attacks.
Working and studying from home means connecting to the network every day, for several hours, sometimes for a period of time equal to office hours, with all the consequences one can easily imagine in terms of exposure to the current cyber threats.

In particular, cyber-investigators found a website - "antivirus-covid19 [.] Site" - on which a product called "Corona Antivirus" is sold as "a digital antivirus that promises to protect against the current epidemic of coronavirus ".

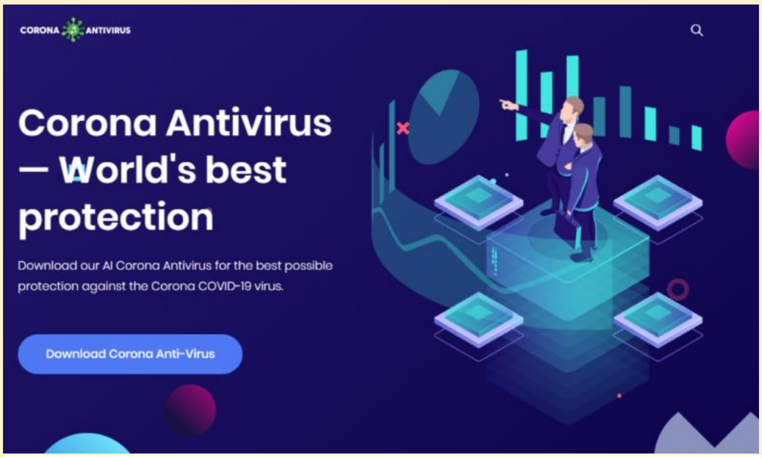
CORONA ANTIVIRUS
Corona Antivirus — World's best protection
Download our AI Corona Antivirus for the best possible protection against the Corona COVID-19 virus.
Download Corona Anti-Virus

Figure 17 The webpage of the fake "Corona Antivirus" used for one phishing attack.

In reality, by selecting the link offered on the site in question, the victim finds himself downloading malware capable of transforming the PC into a botnet, that is one unit of a large network of "infected" machines that are controlled remotely to perform several illegal operations without the knowledge of the unsuspecting user. Those illegal operations include launching DDoS attacks, making screenshots, downloading files, stealing saved passwords, acting as keyloggers, executing scripts and stealing from cryptocurrency wallets.

As usual in these cases, the suggestion is to be wary of these and similar messages, to avoid downloading applications or opening attachments, and to denounce any suspect cases to the appropriate governmental agencies and police authorities.

Another threat linked to the spread of the coronavirus epidemic in the world is the exponential increase in phishing attacks. According to the new Atlas VPN Report, in March 2020 the number of those threats grew by +350% compared to the month of February, bringing the number of dangerous websites to over 522 thousand. That is a significant increase on a monthly basis, if we consider that in February there were 293,235 and in January only about 150,000. According to experts, this is a phenomenon closely related to the COVID-19 pandemic and the odd fact is that these are phishing sites that arise in clusters of thousands from one day to the next.

Regarding phishing, the Italian postal police recently issued an alert in which they asked network users to not open the "CoronaVirusSafetyMeasures.pdf" file for any reason because it was a new online scam circulated by email and Whatsapp messages.

Many other similar cyber threats related to phishing scams and other malware have been denounced around the world, including free Netflix passes, "Coronavirus finders", promise of COVID-19 safety masks, fake goodwill payments, and fake fines for breaking COVID-19 quarantine. Coronavirus-themed phishing emails can also take different forms, including CDC (U.S. Centers for Disease Control) alerts, health advice emails (this has happened in China), workplace policy emails, etc.

"Many in the digital security community are coming together to combat malicious actors during the coronavirus disease 2019 (COVID-19) global outbreak. One of the most visible of these new efforts is the COVID-19 CTI League. Made up of approximately 400 volunteers living in approximately 40 countries, the COVID-19 CTI League is working to block attackers from health care organizations and other medical facilities at this juncture. Reuters reports that the threat intelligence collective is also looking to preventing lowly phishers and fraudsters from capitalizing on people's fears surrounding the outbreak." [3]

[3] https://www.tripwire.com/state-of-security/security-awareness/covid-19-scam-roundup-week-of-3-23-20/
https://www.tripwire.com/state-of-security/security-awareness/covid-19-scam-roundup-week-of-3-23-20/
https://www.trendmicro.com/vinfo/hk-en/security/news/cybercrime-and-digital-threats/coronavirus-used-in-spam-malware-file-names-and-malicious-domains
https://us.norton.com/internetsecurity-online-scams-coronavirus-phishing-scams.html