Digital Twins: Evolution in Manufacturing

By Roberto Saracco

Table of Contents

1. Introduction

Figure 1. Digital Twins used to be a straightforward, well defined concept. As they evolve it becomes more difficult to have a precise definition that is agreed by all. Image credit: IEEE DRI

Everything evolves, right? Why shouldn't Digital Twins evolve as well? Indeed, they have been evolving, and as I look at what is happening around the world, they will be evolving even more in many sectors. I guess one should take for granted a widespread knowledge of what a Digital Twin is, particularly with an audience coming to listen to the evolution of Digital Twins. As a matter of fact, the concept of a Digital Twin was straightforward just 5 years ago; a digital copy of a physical entity. Yet, when I discuss this with different people, I get a variety of nuances, and when I think about DTs today in different sectors and how they are evolving, the definition becomes fuzzier and fuzzier. In a way, this is proof that there is a significant evolution under way, but at the same time, the fuzziness supports further evolution (if something remains well defined, it is constrained by its definition, i.e. does not change, nor evolve!).

Hence, the very first point to address is to look at the definition of a DT (for the record, last year I participated in a discussion with the group of authors engaged in writing a book on DT (which should be out in a few months), and again, more recently, in a discussion within the EU expert group on personal digital twins. In both cases, different opinions on the definition of DTs have (and will) emerge, and it hasn't been possible to come to a single, unanimously agreed-upon definition).

2. Digital Twins Evolution

2.1 What is a Digital Twin?

To start, let's look at the "old" definition: a digital copy of a physical entity. The digital copy:

- Mirrors the physical entity, i.e. its digital model;
- Keeps track of its' real-time status and shadows the physical entity;
- Keeps a record, thread, of the evolution of the physical entity.

The first point – In this definition, it is "implicit" that a Digital Twin is not, and never was, a "copy" of the physical entity. First, the physical entity is "as is" at this precise moment, and because of the "thread", it does not have "memory" of its past (not necessarily at least). Second, even disregarding the thread, the digital model updated to the present status (through the shadow) is "always" a partial model of the physical entity (to be extreme, we would never be able to model each individual molecule making up the physical entity). This partial model is fine, as long as it represents what matters from the point of view of using the digital twin.

Now, we come to the second point: as the use of the DT changes, we need to change the digital model. I'll explain in subsequent post what this means from an evolution point of view.

When dealing with the digital model, one important aspect is how we can create (and re-create) it. Historically, the digital model of a product in manufacturing has been created (and by far still is) using the output of Computer Aided Design (CAD), a tool used in the design phase. Hence, most of the time, the Digital Model precedes the existence of the physical entity. In industries such as building construction, the digital model can be the result of the Building Information Modelling (BIM), a tool, and standard, used in that industry. Again, in this scenario, the digital model precedes the physical entity. In other cases, for example as in healthcare, the digital model is generated after the physical entity and can originate from the Electronic Health Record (EHR). The digital model can be "generic" or specific to a physical entity. In the end, we will always need a specific digital model that, in the case of a generic one (for example, the CAD model), requires simultaneous and constant shadowing of the physical entity

In other words, the manufacturing process produces many "pieces," each one similar, but each associated with a specific instance. All those instances will share the same digital model but will have different shadows and different threads.

However, there is another way to create a digital model – by observing interactions of the physical entity. This is what happens with Alexa, for example. Through its interaction with the user, Alexa has the potential to create a digital model of the user pertaining to their profile and behavior. Note, I am not saying that it does this today, just that it "might" and can. Voice-assistants can distinguish unique voices among users, and, consequently, have sufficiently and accurately developed a digital signature of a person's voice–different Digital Twins. The voice digital signature can, of course, show pattern alteration (still the same person speaking, but with a different intonation…), and Alexa can derive information about the mood of the user (and react accordingly). Of course, there is plenty of information in the interactions (what the user wants, when they want it…) enabling the creation of a model of the users' habits, interests, etc.

I will demonstrate how this manner of creating and expanding a digital twin may become crucial assets in manufacturing in the framework of Industry 4.0.

A digital model is fine in the design phase. Actually, we are hearing a new word: virtual twin. A Virtual Twin models a physical entity that does not yet exist, and may never exist, in the physical space. We create the "idea" of an entity, and we keep that entity in the cyberspace, ready to interact with other entities both in the cyberspace and in the physical space.

This provides industry with great flexibility – what used to be a step in the design phase (resulting in a physical product) becomes a soft product that can potentially be sold on the market.

Nevertheless, the physical dimension remains crucial, and the virtual twin derives its value in its capability to interact directly, or indirectly, with the physical world. The "mirroring" of the physical world implies the capability to remain updated on the real-time status of the physical world. In the digital twin, this is done through the shadowing of the physical entity. The updates may be generated by the physical entities themselves (through embedded IoT), or they can come from the environment (for example, video cameras on the shop floor assembly line that report what is going on in terms of video streams. These streams are analyzed by image-recognition AI that produces data "describing" the status in real-time). At a factory level, we are increasingly seeing a blending of data coming from the robots on the production line, and those coming from various types of cameras. Additionally, data may be derived from the interactions taking place amongst workers, and between workers and machines. The entire factory is becoming an aggregate of digital twins interacting with their physical counterparts and among one another.

Figure 2. The manufacturing process is quite complex. It involves many resources and players within the shop floor and through the supply chain. This generates a massive data flow, many GB per each single product, that can be collected and used through and after the manufacturing is completed. This is what happens in assembly lines, like at Tesla, where digital twins of equipment are working in sync, resulting in an extension of the product (in this scenario, a car) digital twin "instance". Image credit: Atria Innovations

Additionally, the assembly process may result in the assembly of the product digital twin that will be included as part of its data thread related to its manufacturing. The "construction" of the digital twin flanks the construction of its physical entity. This requires a new way of looking at the manufacturing process.

Digital twins can also support embedded IoT to provide the status of their physical entity or can be used by an external application to simulate a predicted outcome or status of the physical entity. This will need to be "confirmed" by data retrieved from the physical space. For example, the engine on a flying plane will be reporting data (pressure, fuel flow, thrust, etc.) at predetermined intervals, and the digital twin will be matching this data with additional data derived through simulations applied to the digital model. In case of discrepancy, the DT (or an external system) will be required to identify the issue and take the appropriate actions accordingly (these can also include a refinement of the digital model). As a matter of fact, notice that if it is the DT that carries out the analysis, this DT is a significant extension of the DT concept. All the shadowed data accumulates and results in the digital twin thread, becoming a source of “intelligence” for both that specific digital twin, and for “generic” instances of that digital twin. This is a very important possibility that opens the door to the provisioning of services flanking the product.

For example, Tesla does this through data analytics. Tesla monitors the behavior of approximately 2 million cars produced since 2009, and can assess both issues on a specific car, as well as issues derived during the production of a given batch of cars. Furthermore, the data retrieved is used to continuously refine the manufacturing process. Information derived from shadowing, and data analytics on threads, is used to provide customers with operation and maintenance support.

The thread includes both data derived from the physical twin, as well as data that can be acquired from the context of the physical twin. A growing part of the thread is formed by the analysis of the effects of interactions between the DT and the Physical Twin (PhT). In other words, the digital twin is evolving to include knowledge and understanding. This is relatively new, and it marks a departure from the original concept of Digital Twin.

What this means is that the Digital Twin that used to be a lesser version of an entity in comparison to its physical twin (because it represented a subset of the physical twin) is now becoming “larger” (in some respects) than its physical entity. In turn, this means that industry, and market, will increasingly start to use the DT to derive features that would not be available in the PhT.

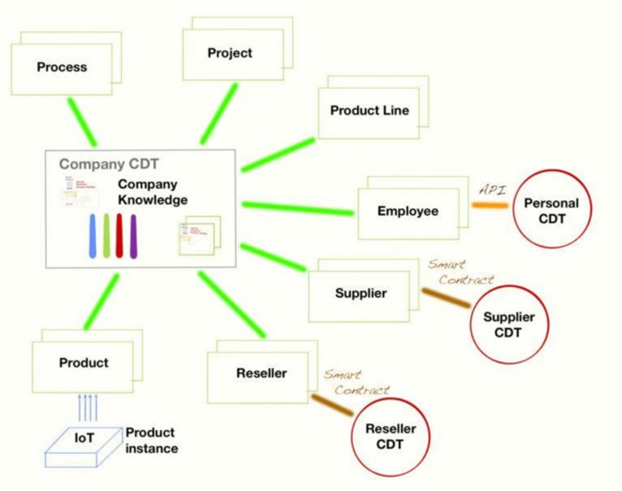

Figure 3. Knowledge is an infrastructure enabling business and operation within a Company and across the value chain. The Knowledge Infrastructure connects local knowledge, and creates an emerging system-wide knowledge, tied together through processes. Image credit: DRI IEEE

2.2 Picking up a Life of its Own

The continuous expansion of data sets accrued by the DT, and the embedding of software capable of providing analytics (more and more AI based) on this data is de-facto transforming the DT into a knowledge entity.

The knowledge is about the PhT—what it is, how it is performing, what are the interactions taking place with the environment. This knowledge is rapidly extending to the ambient in which the PhT also operates with knowledge derived from the analyses of the knowledge space of other instances of that DT. This latter knowledge is usually created outside the DT by an external function (most likely leveraging on AI and ML). Although this knowledge is created outside of the DT, the DT is designed to expand and refine its decision capability. Therefore, it becomes part of the DT itself.

DT knowledge is accrued to enable decision making and to record its interac
tion with the PhT: it is a knowledge “to take actions,” not a knowledge to know about things. This is usually referred to as “executable knowledge.”

Executable knowledge results in interactions among entities (autonomous players). As shown in figure 3, we find this knowledge in the workings of a company, manifesting itself in the ways activities are performed within the company, and in the interactions the company has across its value chain.

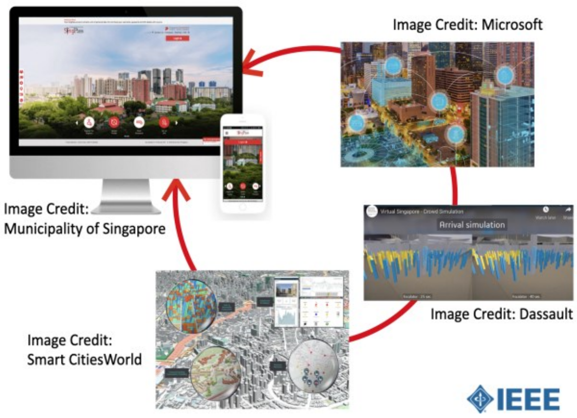

Figure 4. A Digital Twin can be a cluster of several other Digital Twins. This opens the door to abstraction and emerging intelligence. Image credit: DRI IEEE

This creates a knowledge infrastructure that, in turn, generates an emerging system-wide knowledge.

The DTs become knowledge hubs, and as they expand their capabilities, they become independent knowledge entities that can be used in other contexts. This is an interesting evolution in terms of manufacturing processes and business opportunities since they can be used "independently" of their PhT.

Furthermore, the possibility to share knowledge through interconnection of DTs leads to the creation of DT clusters.

This is the case in a smart cities. Singapore was the first city to leverage Digital Twins— independently created to mirror specific resources by clustering them, creating a "Singapore City" DT. This DT is an abstraction of the city modeling the interplay of its various infrastructures and components.

Likewise, in a manufacturing context, we can cluster the DTs of robots on a shop floor to create the DT of that floor. This is not just a static representation of the shop floor, it is a dynamic model of what is going on AND what can go on there. We are seeing more and more applications of these DT cluster concepts in manufacturing—used for monitoring and planning a new production line, determining how to restructure the shop floor, how to change/tune individual components (robots, teams, etc.), and/or a portential redesign of the whole factory[1].

2.3 Extending the Digital Twin

The evolution of Digital Twins, as schematically represented in figure 5 (on the next page), can be "read" from different perspectives. For examples, like the evolution of:

- the degree of representation of the physical entity.
- the interaction level among the physical entity and the DT.
- the relevance of the DT in the operation of the physical entity.
- the functionality offered by the DT.
- the autonomy level of the DT.
- …

Howver, one key component relevant to this discussin is of a different perspective:

★ the transformation of the DT in a product by itself

[1] Dassault. https://www.youtube.com/watch?v=AhNStk765DM

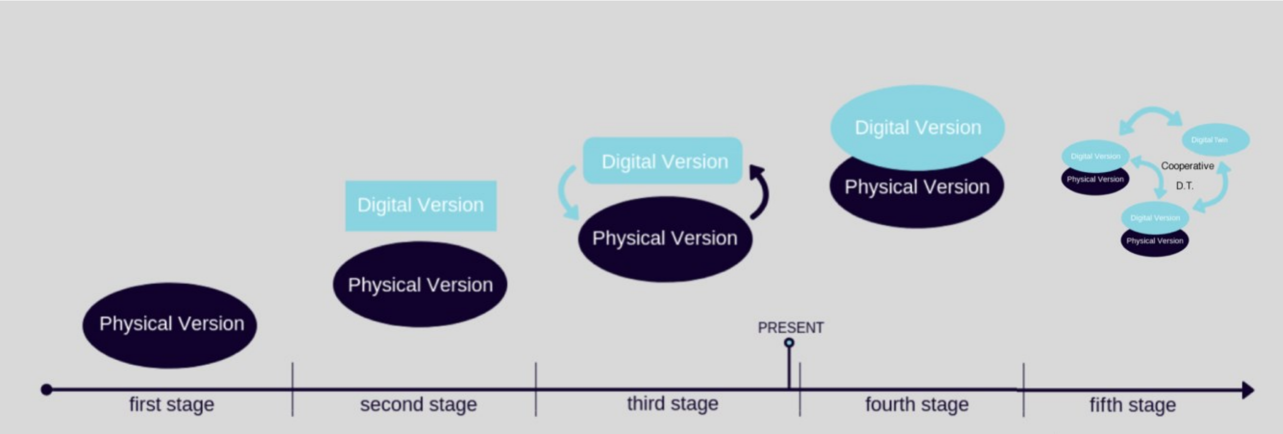

Figure 5. The Digital Twin evolution is described here in 5 stages for sake of clarity. However, there are many grey areas. It may be difficult to assign a DT to a specific stage, as it may be in between two of them. Image credit: DRI IEEE

Notice that this, in principle, applies to all stages represented in figure 5 (below).

Indeed, one could “sell” the digital model created in stage one as a blueprint that can be used by other companies, as well as “sell” (I am just providing examples here) a DT at stage 4 to a third party to embed additional functions that spice up the physical entity (actually this would be a good way to promote a value-added ecosystem on a product…).

This is a very significant change, and it is a “fall-out” of the Digital Transformation.

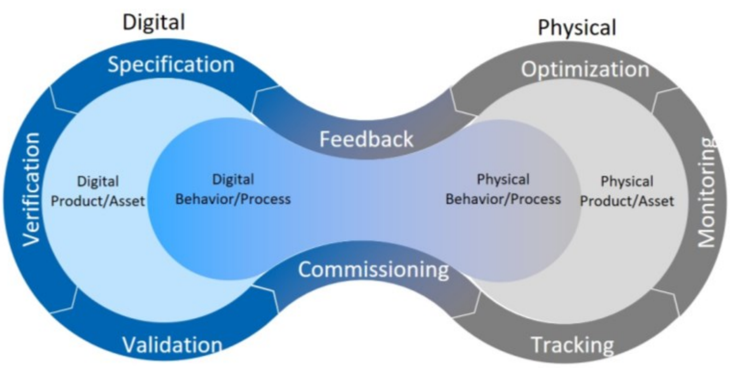

Figure 6. A schematic representation of the flow of activities across the physical and cyberspace resulting a a specular existence of products in the physical and in the cyberspace. Image credit: HCL

By shifting processes, assets, and part of the manufacturing “output” to cyberspace, the resulting products from the manufacturing process may be partly in the physical space (as before), and partly in cyberspace. In the latter case, it can be a Digital Twin—See figure 6.

Indeed, if we look at the manufacturing process, we have digital twins of the tools used (such as robots in the assembly line and mirroring the whole process/processes of the factory), and we have the digital twin (instance) of the product that is created and “manufactured,” along with the physical product itself.

This aspect becomes particularly relevant once we are dealing with stage 3 and onwards. A digital twin at stage 3 interacts with the physical product, and it might embed functionalities designed to enhance its operation and to monitor/provide maintenance. At stage 4, it can augment the functionality of the physical entity, and at stage 5, it can have functionality independent of the physical entity.

All these functionalities can be construed as “services” to be sold with the physical entity, after the sale of the physical entity, or even independently of the physical entity (DT at stage 5).

As a matter of fact, one can envisage in the (near) future a decoupling between the soft side of a product (a DT), and the hard component. This decoupling may result in some industries focusing on manufacturing the soft part, and others the hard ones. That is obviously the case for smartphones, tablets, and personal computers where you have a full decoupling between the application part and the “device”—each one being manufactured by different parties. A standardized platform (the OS in the examples mentioned) ensures that soft goes hand-in-hand with hard. The soft part, the DT, can be the real provider of features, using the hard part as an interface to deliver the features.

However, what I am pointing out here is that, as manufacturing is reshaped throughout the Digital Transformation, the Industry should look at ways to exploit the cyberspace. By requiring the creation of a Digital Twin of the Manufacturing process, and along with it of the product (since the product’s DT is used to steer the manufacturing), it makes more sense *not* to think of this DT as a tool. Instead, consider this DT a product in itself, and leverage the opportunities.

This, however, means that Manufacturing creates both products and services, and in turn, this requires a different set of business processes and procedures.

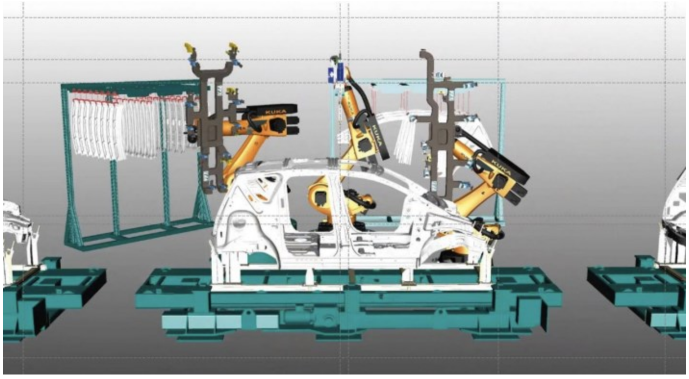

Figure 7. A digital twin can become a product itself as a company may create and sell it to a customer. It will be up to the customer to instantiate the Digital Twin to serve in the intended environment. This scenario portrays the idea of a digital twin, acquired by a third party, being used in the manufacturing process. Rather than buying a robot, a company will be able to buy the digital twin of a robot with the capability of instantiating it to match existing robots in a specific environment to take care of specific tasks. Image credit: Siemens

If we stretch the idea of a Digital Twin as a "product" able to deliver value to the customer, we can foresee an evolution (at stage 5 and beyond) where some companies will be creating and selling Digital Twins. These will operate as software applications that run on platforms such as smartphones, industrial platforms (such as Mindsphere), Government public platforms, and coming soon (by the end of this decade), a communication platform like 6G.

What would be the difference between a software package we use today and the Digital Twin kind? Well, the Digital Twin (to be faithful to its name), is a software package that mimics an entity. For example, you might have a company that offers a Digital Twin mimicking a person. You and I will buy that Digital Twin (possibly running it via our smartphones), and we will instantiate it to mimic our person for the traits we pre-select. In another potential scenario, after buying this "person's digital twin" from a company, I will instantiate it by opening my EHR to it—providing access to my wearables (to get the stream of physiological data that these harvest), sharing it with my doctor, and providing its' identity in my EHR so that if I end up in an emergency room, any hospital can access to obtain and share data.

In the manufacturing industry, we could buy a Digital Twin (a model of a manufacturing process, a generic robot in an assembly line, a warehouse, etc.), and instantiate it to the factory environment. The Digital Twin "model" will be expanded/refined to match the current physical entities and will acquire the "thread" (or historical record) of those entities. Furthermore, it will be connected to the physical entities to shadow them. From that moment on it becomes an instance of the digital twin acquired, and the real digital twin of the associated entities.

This mechanism is based on the idea that we can create a generic Digital Twin, with an embedded model and a set of features, along with a tool (it can be part of a platform) that can support the client/user in the instantiation of the Digital Twin by adding specific knowledge.

The idea of a Digital Twin embedding knowledge derives from the work done by IBM to create a digital twin mirroring the newer models of robots used in manufacturing.

This idea takes the concept of static modeling (part of the Digital Twin Digital Model) to the next level. These newer models have a greater level of autonomy and can operate by taking autonomous actions and sharing them with the environment (such as other robots in the assembly line and in the supply/delivery chain). This autonomy requires a knowledge of the context and of the goals (plus a framework of do's and dont's). In 2018, IBM came up with the concept of Cognitive Digital Twins[2] [2] to match the evolution of robots in manufacturing, and this lead to an extension of the Digital Twin concept.

The ongoing shift in automation on the shop floor, involving smarter and smarter robots, is also known as Robotic Process Automation (RPA). The Cognitive Digital Twins (CDTs) are an integral part of this transformation.

Notice that knowledge is both embedded in a CDT, and shared across several CDTs, creating a knowledge infrastructure that characterizes the knowledge space of operation of CDTs and of DTs. In other words, the knowledge space of a CDT becomes the operation environment for all digital twins operating in that environment.

2.4 Personal Digital Twins

As should be clear from the previous discussion, Digital Twins are a powerful and flexible way to represent salient characteristics of a physical entity, and they have been evolving fast, extending their reach to represent an ever larger variety of physical entities.

One might wonder if they would also be suitable to represent a person. Indeed, this is not a hypothetical question since we already have a number of examples of digital twins used to represent partial characteristics of a person.

For example, Dassault has created a digital model of a human heart[3] and is looking into extending it into a Digital Twin by creating a shadow using data from wearables (measuring heart beat and monitoring the electrical activity of the heart), and keeping the thread. They are not alone. The pharma industry is routinely using organi simulation and fluidic chips, also referred to as organ on a chip, to experiment with drugs. These chips have an associated digital twin and there is interest in using these digital twins, through instantiation, as previously described, to monitor organs of living humans and their reaction to drug protocols. There is even a name for these types of Digital Twins: Deep Twins.

Through aggregation (already happening in Pharma with the shift from organ-on-a-chip to body-on-a-chip), we might expect to have a digital model that can mimic the physiology of the body and can be instantiated to create a Person Digital Twin (PDT), mirroring the physiology of a specific person, enriched with genomic data (DNA sequencing) and with a thread recording the healthcare history of that person. By

[2] https://www.ibm.com/blogs/internet-of-things/iot-evolution-of-a-cognitive-digital-twin/
[3] https://www.3ds.com/products-services/simulia/solutions/life-sciences-healthcare/the-living-heart-project/

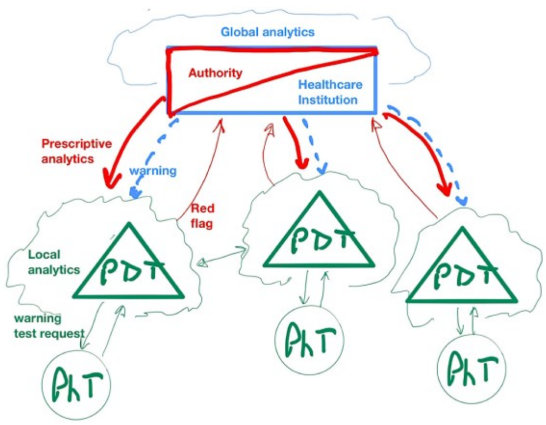

Figure 8. Example of use of a PDT in the context of epidemic monitoring and control. Notice the decoupling assured by the PDT between the physical person and the context. This is crucial in preserving privacy on the one side still ensuring societal benefit on the other. Image credit: IEEE DRI

connecting this PDT to the person's body using wearables and other types of ambient sensors, we would have a full blown PDT.

We are not there yet, but we already have some types of PDTs, limited in terms of mirrored characteristics, in the healthcare sector supporting very concrete and useful applications.

Obviously a "person" is much more than its physiology! The physical shape of a person is also another characteristic that may or may not be important. For example, if you are looking for apparel, such as a t-shirt or a pair of shoes, your physical shape is very important. On the other hand, if you are applying for a job as a data analyst, your body shape, your sex, and even your physiological characteristics are likely to be irrelevant. What matters, to you and your employer, is the type of skills, experience, and knowledge you can bring to the table.

Historically, particularly in the Western world, we have gotten accustomed to distinguishing the difference between the body and the mind (soul). This is not the place to enter into a discussion on this, but it is important to notice that the representation of the aspects related to the physical versus the cognitive sphere differ significantly.

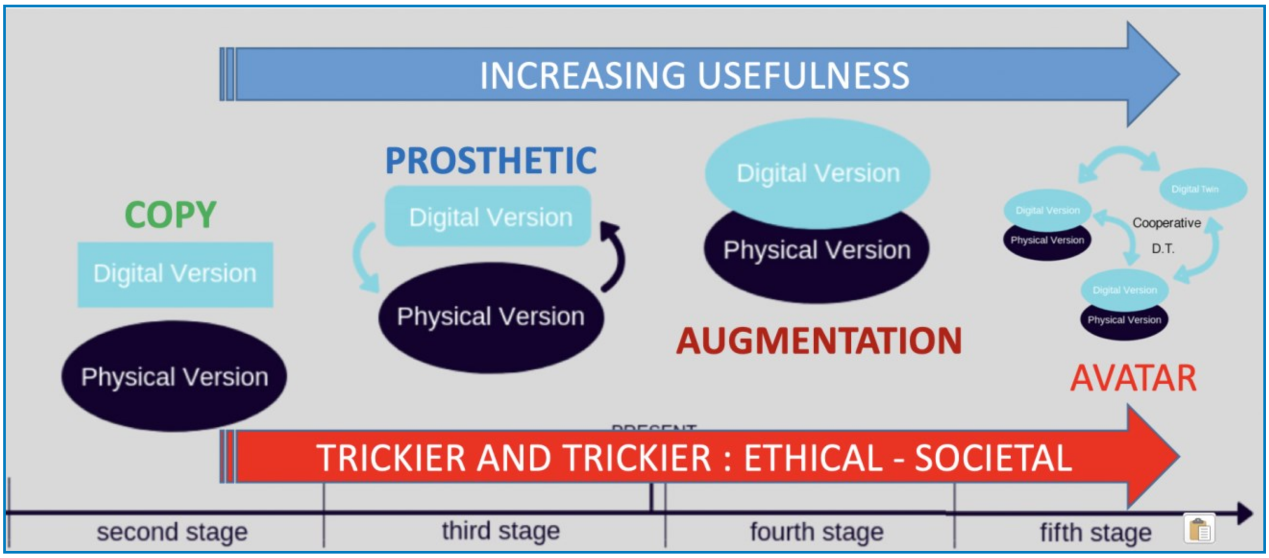

Figure 9. Personal Digital Twins are new on the Digital Twins landscape, but they are fast growing in capabilities and adoption. This image outlines four main stages of evolution, and their correspondence to application areas—from a pure copy of a person that can be used for study and simulation, to becoming an avatar of that person. Image credit: IEEE FDC DRI

Indeed, the work on extending the Digital Twin to humans has lead to the identification of the Cognitive Digital Twin of a person, which represents the knowledge, attitude, character, feelings, etc…, and a more general PDT (which may or may not include the soft aspects of a person). Therefore, with the concept of a PDT, we refer to digitally mimicking certain aspects of a person, and we need to specify what these aspects are. With CDT's, we are only referring to the cognitive aspects, and again, we need to specify the extent of mirroring being done.

In the case of PDTs, it makes sense to take a pragmatic approach and look at the way these may be used to outline the evolution roadmap. Although the roadmap looks similar to the evolution roadmap of Digital Twins (as reported in figure 4), here, the emphasis is on the application and the issues deriving from their application.

As shown in figure 9 (on page 13), we can compare PDT evolution with the evolution of DT's:

- At stage 2, we have a copy of certain characteristics of a specific person, such as the ones derived from genome sequencing, for example, to define a drug protocol for breast cancer (the first stage would be one where we only have a generic model of some person's characteristics, like the one used in Pharma for testing drugs on a chip);
- At stage 3, the PDT may become a sort of prosthetic flanking the physical person and interacting with it;
- At stage 4, the PDT can take over some aspects of the person. augmenting the person (such as a PDT that harvests information from the web and makes it available when needed);
- At stage 5, the PDT can behave as an avatar of that person, acting as a proxy in cyberspace (and possibly interacting with the physical space on behalf of the person).

This can be the case in manufacturing where the PDT of a technician can provide support in the shop floor (both to machines and to other workers) with no need of presence of the physical person and,, in principle, without the physical person being aware of the activity of her avatar.

It is obvious that the higher the stage, the trickier the management of the PDT, as well as newly emerging ethical issues.

2.5 Cognitive Digital Twin

As mentioned, the concept of Cognitive Digital Twin (first defined in the context of smart robots by IBM as a way to represent the knowledge of the robot and an ensemble of robots on the shop floor), has been applied to the representation of the knowledge of a person, becoming a "subset" of the Personal Digital Twin of that person. In fact, the CDT has the potential to express a given set of characteristics of a person that are relevant in a given context or situation. The following are some examples:

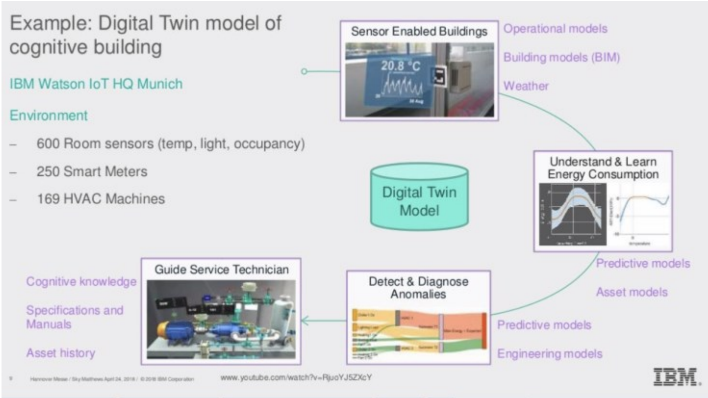

Figure 10. The slide used by IBM in the presentation of the Cognitive Digital Twin concept at the Hannover Messe in 2018. The Cognitive part is seen as an add-on to the digital twin model and is used to make smart and flexible decisions in understanding the environment, as perceived through IOT data, and interact with it. Image credit: IBM

- Knowledge management at a company level;
- Knowledge development in an educational environment, like a college, university, training program
- Knowledge asset management at a personal level (what do I know, what should I know?)
- As a trading asset in a business environment

It is important to take note and recognize that the management of "personal knowledge" is trickier than the management of a machine knowledge (be it a robot or an AI-based application). It is difficult to capture and portray, from the point of view of mirroring, what that person knows in terms of expoitation of thie knowledge—in other words, it is complicated to mirror a person's "executable knowledge."

A person may:

- Know something but be unable to apply that knowledge to the problem at hand
- know something but be unable to face a given situation (e.g. stress) and apply that knowledge
- Might have known something and then forgot all about it
- Might know something and be unwilling to apply, or share, that knowledge.

It should also be noted that, in the machine domain, there are also tricky issues in knowledge representation and management. For example:

- A set of knowledge is first embedded in the machine (application), both in terms of a static representation, models, data, procedures, and in terms of algorithms. How should we make sense out of existing data and interactions? The first set is fully controlled by the designer, and can be tested extensively. However, as more and more data becomes available, this first set of knowledge may prove to be difficult to be tested exhaustively (think about the millions of images used to train an image recognition application such as the one present in autonomous cars);
- The first set of knowledge is expanded through the lifetime of the machine/application operation, and it may become impossible to keep track AND to test the interpretation/implication of the new data accrued on the pre-programmed algorithms;

- The new wave of Artificial Intelligence (AI) is not “pre-designed” or “pre-programmed.” Rather it is emerging from algorithms that are competing with one another (like Generative Adversarial Networks, or GANs). Here, the designer teaches the AI how to learn by defining objectives and values, letting the AI identify algorithms more accurately approaching the goal, and maximizing values. The AI builds up both a knowledge and a reasoning on its own (this is what transforms knowledge into executable knowledge), and it becomes difficult to create a representation of that knowledge. The reality is that the only accurate representation is the AI itself. As in the case of human knowledge, the only accurate representation is the brain/mind itself, and this can only become visible as it is executed.

From this discussion it is clear that any CDT, both associated to a machine and to a person, is, at the very best, a limited and often imprecise model of the real executable knowledge of its physical entity. As in many other areas of our “understanding” of the world, we have to make do with what we have.

As long as the CDT proves to be useful, and we can control the potential shortcomings, it is fine. This is what is happening today. We have a tool that is not perfect, but can help in the management of knowledge as an asset.

The interest on CDTs is growing, and companies are starting to look at them as a tool to effectively manage knowledge assets. Digital Transformation (DX) is making knowledge a crucial component of business, and it is becoming even more important to manage the knowledge assets of a company. DX shifts atoms into data, but data as such are a commodity with very limited value. The value has to be leveraged through the “understanding of data” and its’ implication in a specific context, or at a specific time.

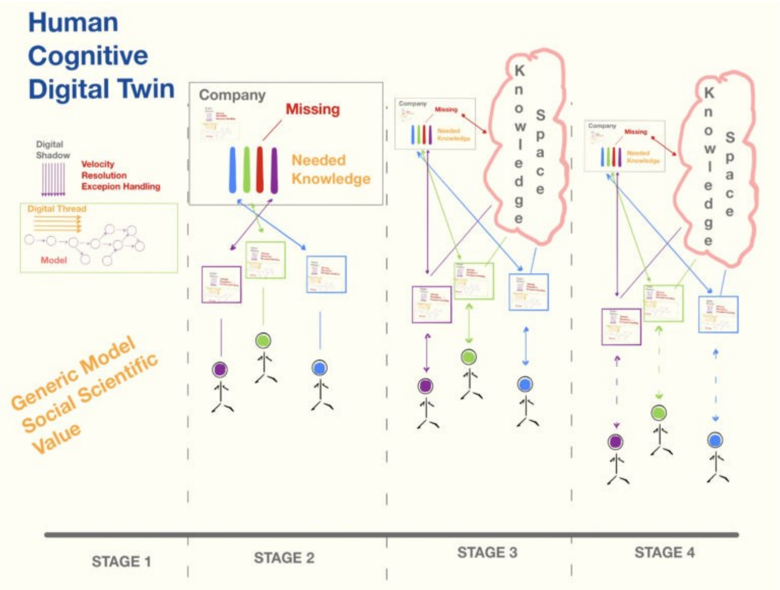

Figure 11. Knowledge has become a crucial asset for companies. Each company has knowledge embedded in terms of operation processes and tools supporting operation. This embedded knowledge reaches out to acquire the knowledge of its human resources. This knowledge is what makes the company competitive on the market. However, the overall knowledge, indicated as knowledge space in the graphic, is way larger, and part of it may be needed to keep the company competitive. Today, the challenge is to ensure knowledge needed is acquired by the human resources through training, hiring, consultancy—tomorrow, this additional knowledge may be used independently of the resource having it by interacting directly with the CDT associated to the resource. Image credit: DRI IEEE

A CDT can “capture” a knowledge asset and make it an active operational component of the company. In other words, the company may use the CDT in place of the physical entity that has that knowledge.

The first step, as shown in figure 11, is to use the CDT as a representation of a knowledge asset in the company. This can help in assessing what knowledge is available with respect to the ones needed. Notice that this is something that is already happening (even without the CDT): an HR department has a “map” of the companys’ knowledge space, i.e. who knows what. This is essential for completing and conquoring human resource tasks (technical departments have a map of the available tools and what they can be used for—for example, the flexibility of a robot and how it can be used in a given scenario). A CDT would provide, in a way, a standardized way to represent the knowledge of an asset. In addition, a CDT will have the capability to keep this representation up to date (through shadowing). It is also important, as mentioned, to identify gaps (usually this is a technical area that defines the needs, and the HR department looks for ways to meet those needs, identifying possible gaps).

The next step. shown in the graphic, is to identify the missing knowledge in the knowledge space outside the company (the IEEE knowledge ontology is a good reference point to navigate the knowledge space of technology, including the very latest of tech). Once this “missing” knowledge is identified it should be brought inside the company.

There are, of course, several ways to bring the needed knowledge "inside" the company:

- Train some employees to acquire that knowledge (in this case, one should also identify those employees that would be better suited for training—pre-existing competences, time availability, etc.);
- Hire a new employee with the desired knowledge;
- Hire a consultant to support the project with the needed knowledge (makes sense if a need is expected to be temporary…);
- Partner with another company that can provide that knowledge and take care of the parts of the project requiring such knowledge
- Buy a machine/application embedding that knowledge (add to or upgrade existing resources).

The added knowledge will be reflected in the related CDT, the one associated to the trained employee—to the newly hired one, … to the machine/application.

Further down the line, we can imagine that the acquisition of knowledge can happen at the CDT level (without involving the physical entity). Now, this might seem like science fiction, but as a matter of fact, that is what happens with robots and software applications where new sw versions can be installed, "adding" knowledge.

Could this be done for a person? we have clearly no way to download knowledge in a brain, it has to be acquired through "learning". However, if we consider CDT at stage 4 and 5, where the CDT is an augmented set of the associated entity knowledge, we can well add knowledge to it.

The crucial point here is that this CDT will in part mirror the existing knowledge of its associated person, and in part will augment them. Notice that with a CDT the knowledge owned constitutes a single "space"—Therefore, the (AI) functions that are transforming knowledge into executable knowledge take the entire knowledge space into account.

A new word has been coined to define this type of CDT—one that has an embedded augmented knowledge: a hybrid CDT (the same name applies to the compound CDT, including a machine CDT and a person CDT, cooperating in symbiosis).

In this "future" (but not science fiction) scenario, several challenges will arise as previously mentioned. We are also entering into a new business space as I will discuss later.

3. Use of Digital Twins as of 2022

After the previous discussion on the evolution of digital twins, basically, from an academics point of view with a separation into stages and types (DT, PDT, CDT, OPDT, Hybrid, etc.), and before looking at the further expected evolution, it makes sense to take a look at how today, in 2022, digital twins are used in various sectors—sort of a reality-check.

3.1 Manufacturing

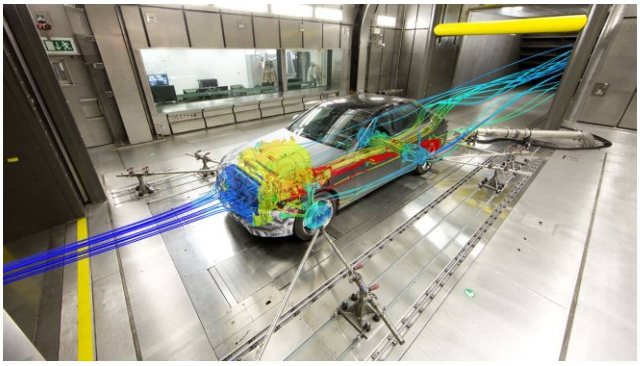

Figure 12. Product manufacturing relies on cyberspace. Digital Twins create a copy of tools, processes, and components in the cyberspace giving rise to the digital twin of the product as part of the PLM.

Manufacturing has been the first area to put the concept of Digital Twins to work, and they are now an integral part of manufacturing processes in many companies. In fact, they for sure have become one of the pillars of Industry 4.0.

Manufacturing is based on tools and processes—orchestrating the use of tools and resources throughout the entire Product Life Cycle Management (PLM). Digital Twins are derived from tools (CAD) and used in others (CAM). They themselves have become tools in supporting simulation and monitoring.

”Orchestration” is done through processes and tools supporting them. In the case of Digital Twins, the orchestration is achieved and supported via platforms. The major manufacturing tool providers have created their own platforms, such as Siemens Mindsphere[4].

Most DTs used in manufacturing are at stage 3, for example. When the DT interacts with its physical entity only for the sake of remaining in synch with it. It can also act as a gateway for other applications (analytics, simulation, etc.) to interact with the physical entity. An anomaly, detected by data analytics provided by the physical entity via its associated DT, can be processed by an external application and result in a command that will be handed over to the physical entity through the DT.

The same applies to the DTs associated to most products. They are created during the manufacturing process and remain in the ownership of the manufacturer to connect with the physical product throughout its life time.

A few of these DTs are starting to embed “intelligence” to perform data analyses and to assist the physical entity. In a way, this is a tiny step towards becoming autonomous. Some are also connecting to the cyberspace to get addtional data (autonomously) to be used internally. Self driving cars are a clear example of this kind of evolution (getting a better grasp of the context by communicating, *autonomously*, with other DTs).

Mevea[5] is possibly one of the most advanced users of DTs in the industry. They are basing business models and competitive advantages of the adoption of DTs throughout the life cycle and are using the DT of their products to deliver services. They use the shadowing to get insight on the use of the products. They compare shadowing of several DTs in a given product line to improve all of them, based on experiences derived from each of them. Their DTs are, in many cases, approaching

[4] https://siemens.mindsphere.io/content/dam/cloudcraze-mindsphere-assets/03-catalog-section/05-solution-packages/solution-packages/digitalize-and-transform/Siemens-MindSphere-Digitalize-and-Transform-sb-72224-A8.pdf
[5] https://mevea.com/solutions/digital-twin/

stage 4 since some of the product functionality is actually being delivered through the DT.

General Electric[6] is another company (as a matter of fact, one of the first) that heavily relies on Digital Twins to monitor the use of their products and provide proactive maintenance services (placing their DT somewhere between stage 3 and 4).

The Competence Industry Manufacturing 4.0[7], located in Turin, Italy (on the Turin Polytechnic Campus and clustering many companies in the manufacturing area), is developing a digital twin infrastructure[8] that can be used by their associated company to create a virtual lab[9], consisting of both physical and virtual objects that can be inspected and assembled in a hybrid mode (virtual +physical).

Here, DT's are present at all stages, from 1 to 5.

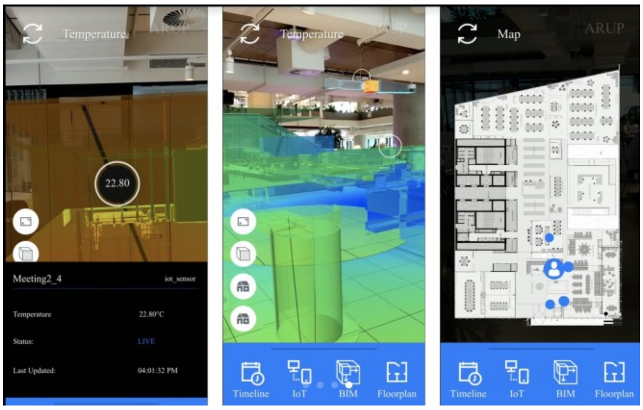

Figure 13. Building operation and maintenance benefit from the existence of an associated Digital Twin. Figure 13 demonstrates three screenshots from a workers' smartphone using an app that leverages digital twins to explore data associated to the building. In the first two screenshots, AR is used to highlight temperatures of different parts of the building. This helps in evaluating the level of insulation provided by the materials used. It can also show the presence of cracks in the structure (at the crack level, there is a clear temperature gradient). Image credit: ARUP

3.2 Construction

The Construction Industry has been working with digital models for quite a while, based on Building Information Modelling (BIM). In the last few years, more and more sensors (IoT) have been inserted into buildings during the construction phase (and in several cases, IoT has started to be retro-fitted in existing buildings). These sensors are generating streams of data that enable the constructor to create a shadow and a thread, i.e. the three components of a Digital Twin.

The development of software to manage and leverage these DTs is enabling operation and maintenance services, which has been an obvious next step.

ARUP, one of the largest construction companies operating worldwide in over 150 countries, is using Digital Twin technology through the construction phase that, after project completion, continues to monitor the buildings operations. Over time, a historical record grows, and this data can provide hints on the degradation of certain parts of the building that need to be fixed before any damage occurs. Proactive maintenance is cheaper (it can be scheduled) than repairing damage after the fact (reactive maintenance). In addition, the digital twin may interact with components in the building infrastructure to

[6] https://www.ge.com/digital/applications/digital-twin?utm_medium=Paid-Search&utm_source=Google&utm_campaign=HORZ-DigitalTwin-MoF-EU-Search&utm_content=%2Bdigital%20%2Btwin
[7] https://cim40.com
[8] https://cim40.com/projects/dtman/
[9] https://www.reply.com/brick-reply/en/content/brick-reply-is-part-of-the-competence-industry-manufacturing-4-0

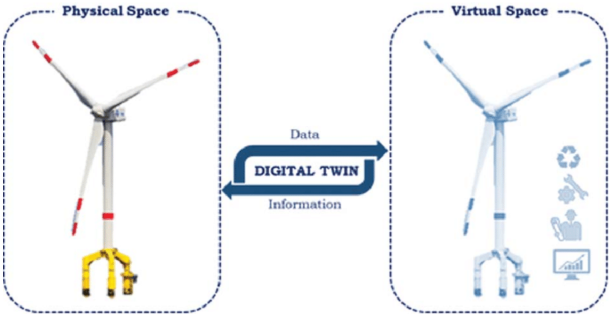

Figure 14. Digital twins are becoming a widespread tool in the operation and management of wind farms. Notice in this graphic that the digital twin does much more than mirror the wind mill—It connects to a variety of entities in the cyberspace, including processes, other equipment, and people. Image credit: Hooman Mohammadi Moghadam

tweak operation and decrease risk of damage (like decreasing pressure in pipes…). DTs operate at stage 3 in most buildings monitored and controlled by ARUP.

Interestingly, the DT can also be used to assist in maintenance work by providing digital information on the building. Use of AR allows the maintenance crew to inspect, virtually, the building infrastructures such as the pipes, inside walls and ducts,etc.

In addition, the data received from sensors is compared with data received from similar buildings. The data analyses take a variety of situations into account (for example, a building in Houston will be exposed to a different climate than one in Seattle) to derive meaningful comparison. This distributed knowledge (on single buildings) generates an emerging knowledge that helps in the operation of buildings (proactive maintenance), and in fine tuning the design of future buildings (what material demonstrated is best in the field in a given climate…).

3.3 Energy

General Electric (GE), as already mentioned, has been applying DTs to energy production for several years, specifically using them to monitor and control wind turbines. Wind farms are complex and costly systems, but by fine tweaking the angle of a blade and other components, efficiency can be increased, and this in turn alters the flow of air (wind blowing across the wind farm). Therefore, when fine tuning one must account for the impact on other nearby wind mills to achieve not just a local best, but global optimization. In addition, monitoring is important to enable proactive maintenance (rather than having to resort to recovery maintenance). Wind farms[10] all over the world equipped by GE have a digital twin for each wind mill, digital twins mimicking the processes, and a digital twin for the entire farm. These digital twins are “hosted” on Amazon AWS Cloud to provide both a local presence and a centralized hub (in cyberspace there are no distances).

Digital Twins mirroring equipment in a specific wind farm and other DTs across wind farms are communicating and collaborating. The GE Assets Performance Management Software[11] (APM) leverages Machine Learning to create knowledge and fine tune processes and operation/maintenance decisions. Interestingly, the creation of digital twins to mirror local conditions, processes, and equipment (in addition to the ones provided by GE that are already delivered with their associated DTs), can be accomplished using a Digital Twin library, provided by GE, which has reduced the time needed to create a customized DT by 75%.As shown in figure 14 above, digital twins play the role of connecting the physical entity and the wind mill to all the relevant components of the wind farm mirrored in the cyberspace[12].

[10] https://www.forbes.com/sites/amazonwebservices/2021/12/07/the-future-of-energy-using-digital-twins-as-a-strategic-asset-at-ge-digital/?sh=7acbfdac7d6c
[11] https://www.ge.com/digital/applications/asset-performance-management
[12] https://www.ge.com/digital/applications/asset-performance-management

Operation data from the wind farms demonstrated a 40% decrease in reactive maintenance, thanks to the use of Digital Twins.

Another interesting use case for GE DTs is as knowledge repositories. For example, when a staff turnover occurs, the DTs can be used for training the new staff and also to let the new staff get in touch with experienced ones located in other parts of the world. This illustrates a very smart use of DTs showing the convergence of product-service-knowledge.

This extended use of the Digital Twin has some aspects that would place it at stage 5.

3.4 Automotive

The automotive sector has adopted Digital Twins technology in manufacturing to mirror robots in the assembly line. In the last few years, it has started to create and use digital twins of the product—the vehicles produced. Automotive companies are increasingly equipping cars (and trucks) with IoT, and they receive a stream of data reporting the statuses for the various components. This stream of data often includes location data, and the issue of data privacy may emerge. Some manufacturers, to avoid these types of issues, keep the data record in the car, and this data is only harvested when the owner requires a car check up (some of the data ends up in the key fob). In this scenario, all data analytics take place in the car.

An entire new class of IoT and AI supporting chips, designed to provide intelligence at the edge (like the STM 32 series), now enables local intelligence and support for the local operation of a Digital Twin. This local intelligence would be able to both signal an emergent issue to the driver, as well as report the problem to a service center for proactive maintenance, possibly without disrupting the service. In other words, the DT may take action autonomously (or guided by the service center—software application), to alter the vehicle parameters, thus ensuring that it can defer the required maintenance to a later time.

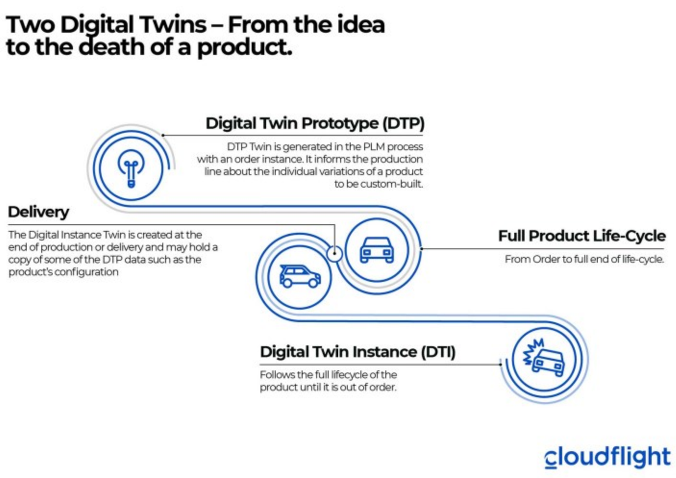

Figure 15. Graphic representation of the creation and instantiatied use of Digital Twins associated to each Tesla car. Image credit: Cloudflight

For the time being, however, (to my knowledge), only Tesla[13] has genuine digital twins associated to each one of their vehicles. Daimle [14] (trucks), Porsch[15] and Mercedes (DT used in production[16] and for Formula 1 cars[17]) are hinting at adopting DTs for their products.

[13] https://www.sas.com/en_us/insights/articles/big-data/modern-manufacturing-s-triple-play-digital-twins-analytics-iot.html

[14] https://www.digitalengineering247.com/article/daimler-truck-adopts-siemens-xcelerator/Digital-Twin

[15] https://grapeup.com/blog/how-porsche-developed-a-digital-twin-to-win-the-race-for-the-virtual-car-concept/#

[16] https://group.mercedes-benz.com/innovation/case/connectivity/industry-4-0.html

[17] https://www.businesswire.com/news/home/20220224005152/en/TIBCO-and-Mercedes-AMG-Petronas-Formula-One-Team-Continue-Data-Driven-Winning-Streak

Tesla has embraced Generative Design[18], an evolution of Computer Aided Design (CAD), which uses AI to optimize the designs studied by engineers based on the goal. By associating AI to the Digital Twins of the cars already in use (the data provided by those DTs), it is possible to receive and consider feedback from the "operation" field, in the true spirit of Industry 4.0. Each Tesla car[19] is paired with a DT, and that DT is reporting back to Tesla GB of data every day. The avalanche of data is analyzed through AI (and Machine Learning), resulting in monitoring, offering services, and fine tuning production. Since Tesla cars are basically computers with specialized software, it is possible to update the software whenever needed, both to fix glitches and to offer new features. Again, we are seeing the convergence of product, service, and knowledge. A massive amount of data is collected by Tesla from their cars--- an estimated equivalent of 3 billion miles of data are now on their servers, enabling unique (in the automotive market) data analytics, AI, and ML. Take a moment to consider that a car may generate a few TB of data every single day! No surprise that some analysts look at Tesla not as an automotive biz, but as a data company[20].

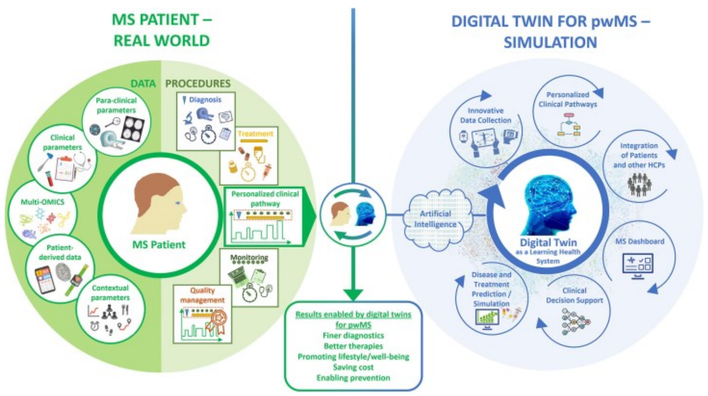

Figure 16. Use of a Personal Digital Twin to simulate situations involving patients and potential effects from healthcare procedures. On the left hand side, the patients physical space is generating data and executing procedures, and the right hand side provides analyses of harvested data and simulations of procedural impact. Image credit: First Hospital Pekin University

At this point, Tesla is using DTs at stage 3. However, there are a few nuances that show uses placing them at stage 4, and maybe even 5.

3.5 Healthcare

Healthcare is an industry with a complex infrastructure and plenty of equipment. Think about hospitals, medicine design and production, care centers … This is also an area that is very sensitive in terms of privacy, and crucial for the well being of single individuals and society. Last but not least, it is a major financial burden weighing on individual and geographic budgets. It is no surprise that Digital Transformation is seen as the key to a sustainable healthcare industry, and Digital Twins are used to increase effectiveness.

Big companies like GE[21] and Philips[22] have utilized DTs for monitoring and controlling the health of their equipment, and they have recently started to use Personal Digital Twins (PDTs) to monitor patients. PDTs are also used by pharmaceutical companies to support design and trials of drugs. Soon, medical insurance companies and healhcare service providers will likely use PDTs to support tele-consultations, diagnoses, prescriptions, etc.).

[18] https://www.industryweek.com/technology-and-iiot/article/21130033/how-digital-twins-are-raising-the-stakes-on-product-development
[19] https://www.cloudflight.io/en/blog/learnings-from-the-digital-twins-data-architecture-of-tesla/
[20] https://www.aidataanalytics.network/data-monetization/articles/tesla-automaker-or-data-company
[21] https://www.ge.com/news/reports/these-engineers-are-building-the-industrial-internet-for-the-body
[22] https://www.philips.com/a-w/about/news/archive/blogs/innovation-matters/20180830-the-rise-of-the-digital-twin-how-healthcare-can-benefit.html

The use of chatbots will be a driving force in the evolution. These will morph into PDTs that can provide personalized assistance.

DTs of human organs[23] have been designed and used in simulation. This is a small step towards these DTs instantiating a specific patient, transforming them into PDTs.

Each person has a growing set of personal data that can be used in medical evaluation and healthcare protocol decisions—from very basic data (related to the patient characteristics such as genome, proteome, metabolome, etc.), to patient health records (exams, diagnoses, drug prescriptions, etc.), patient behavior and environment (diet, habits, working environment, etc.), and potentially even inherited traits.

All this is augmented by monitoring data derived from wearables (smart watches, fit bands, ambient sensors, etc.) and medical devices which provide the "shadow".

The healthcare industry is using DTs and PDTs from stages 1 to 3.

4. Steering Evolution

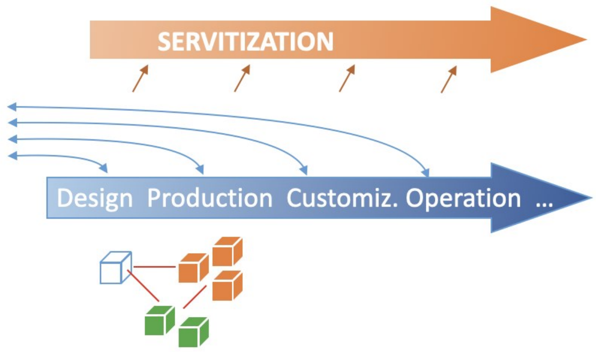

Figure 17. Representation of the 4 dimensions of the evolution of Digital Twins application.

Digital Twins are evolving in 4 major ways:

1. Extension through the product life cycle
2. Extension in depth
3. Extension over the value chain
4. Extension into the business area

Some industries have been using digital twins in the manufacturing phase, some in design, and others to monitor a product. It is quite natural that these industries will be looking at extending the use of Digital Twins throughout their entire project and product life cycle. In doing so, they will need to expand the data sets associated with the Digital Twin, and the set of interactions the Digital Twin engages with. At this time, almost no industry is using the Digital Twin to cover the product *and* lifecycle. The growing interest in sustainability and recycling is increasing interest in leveraging Digital Twins for that purpose.

It should also be noted that many industries are using digital twins, inherited from the tool providers, that are delivering the tools (robots) used in the production/assembly line, with associated digital twins that have to be used to operate the tools. In most cases they come along with a platform (like Mindsphere). It makes sense for these industries to adopt Digital Twins and extend their use to flank the product.

In other cases, the extension is the fall out of the industrial shift to cyberspace because it executes the Digital Transformation. All things considered, we can expect that DTs will become pervasive throughout the entire lifecycle, and more and more products will end up flanked by a DT.

[23] https://www.3ds.com/products-services/simulia/solutions/life-sciences-healthcare/the-living-heart-project/

The second extension is towards a more comprehensive mirroring of characteristics of the physical entities. This is fostered by the increased use of embedded IoTs. The increasing availability of data about an entity is, naturally, leading to an extension of the digital model of the entity—the shadow becomes more accurate and the thread significantly grows. Some industries might even see this as a step-by-step approach to digital twins. For example, you start with a minimal digital model with minimal shadowing, and then over time, you grow in-step with the use (exploitation). This is surely the case for PDTs in healthcare—it starts with a minimal set of mirrored characteristics of a specific person (like the Electronic Health Record, or EHR), and will likely evolve a more data becomes available (such as genome sequencing, data from wearables, etc.). As more data is collected, more intelligence can be derived from the DT (and embedded into the DT). This fosters more usage and, in turn, stimulates the quest for better mirroring.

The third extension involves the use of the DT outside of the company boundaries to span over the value chain. The component availability and value chain crises throughout the past two years have fostered industry interest on ways to have more flexibility in terms of sourcing and logistics. Value chains have pursued efficiency and optimization, achieving impressive results. However, this has led to much rigidity, and a glitch can disrupt the entire system.

The adoption of Digital Twins across the value chain, and the extension of DTs to interact, can provide data to feed AI software designed to explore work arounds and keep the value chain working at high efficiency despite issues. Machine learning can turn historic data into probability forecasts of future occurrences, and alternative strategies can be designed and implemented as needed.

Furthermore, the attention to the value chain is encouraging clusters of digital twins, or several DTs mirroring different “segments” of the value chain and owned by different parties. These can be clustered into a super DT mirroring the value chain as a whole (as an abstract entity). There are several issues we will face (such as ownership and data sharing). In addition, the trend towards the creation of a super DT goes in the opposite direction to the one of self orchestrating autonomous DTs (these are better in terms of ownership since these remain separate and no agreement is needed across the value chain). Logistic companies prefer the super DT approach (since it provides great control and it can be defined precisely), whereas industry on the value chain side (users of logistics) prefer the second approach which is more ecosystem orientated since it provides greater flexibility.

The fourth extension is possibly the most disrupting one. It brings the Digital Twin into the business space. A global mega trend in industry is the one towards servitization of products. This is sustained by the Digital Transformation that, by shifting assets and operation to the cyberspace, creates services out of assets and asset management. Additionally, we are seeing the softwarization of products, i.e. the embedding of software to provide features and functionality. Software can reside in the physical product, or can be in the cyberspace (cloud, edge cloud, and even devices/ambient cloud). The Digital Twin can be used to provide the features and functions (DT at stage 4 and beyond).

So far, and with a few exceptions, companies haven’t leveraged Digital Twins to deliver additional functionality. One reason is becauset DTs are born in industry, at the production level, whilst biz plans are developed at the sales and strategy level. In the future, we can expect a significant extension of DTs to that level.

4.1 Manufacturing Leads

The Manufacturing industry has been the first to adopt the technology of Digital Twins. No surprise since the use of CAD for product design and for shop floor design is rooted in the last century. It was a simple, natural step to take the digital model produced by CAD and use it for simulation, then for steering digital lathes, and then robots to manufacture the various components and assemble them into the finished product. First, IoTs became part of the toolkit and started to be disseminated on the shop floor to support monitoring, control, operation, and maintenance. Then it started to be embedded in the product. Merging all of this data to the digital model was a small step giving rise to the digital twin.

As the DT is being used in more extensive ways it is evolving, as previously discussed, and manufacturing is still in the lead. As shown in Figure 18, DTs are becoming autonomous (represented with a double circle) and smarter, both thanks to access to external intelligence and to embedded intelligence. One thing to note is the evolution of the "external intelligence." While it used to be an intelligence provided by a specific application running on a platform (in the factory or in the cloud), in the future, we are going to see the emergence of swarm intelligence, in other words, the intelligence that is emerging as result of the interaction of several entities (more or less intelligent on their own). Swarm intelligence does not require, per sé, intelligent entities as long as there are many of them, for example, think about bees and ants.

Robotic Process Automation will increasingly rely on this type of swarm intelligence (given the limited number of robots involved, each of them needs to have some degree of local intelligence). Logistic chains (including smart warehouses and smart receiving dispatching docks tied to internal factory processes) will also start to benefit from this kind of swarm intelligence.

On the shop floor, the advent of co-bots (robots that can cooperate with people), is also paving the way towards a cooperation among PDTs and robots' DTs. In perspective, we could foresee only the PDT of an expert technician being needed in the virtual space of the shop floor to interact with the robots DTs as needed.

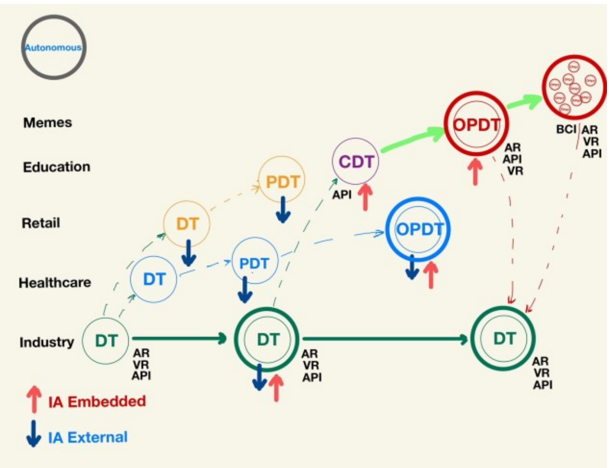

Figure 18. DT evolution and influence across different sectors

Connection with the DTs using Augmented and Virtual Reality is also being driven by the manufacturing world, and this will be another initial step into the industrial metaverse.

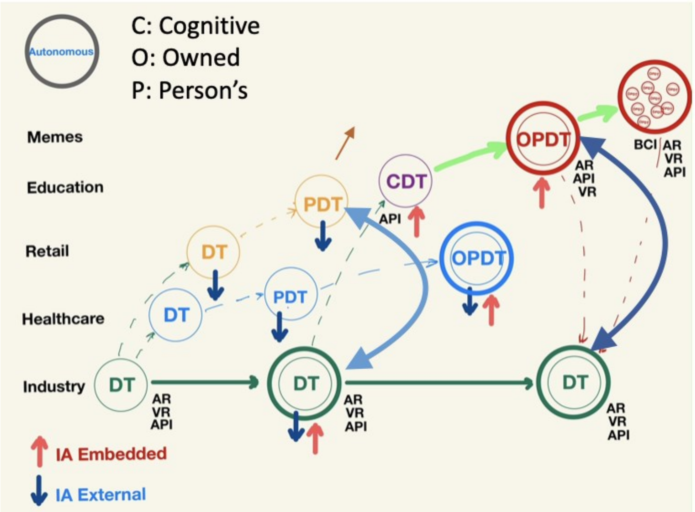

Figure 19. Evolution of Digital Twins, with the design of Personal and Cognitive DTs, on to the concept of Owned Personal Digital Twins. As shown, this evolution, while steered by the manufacturing industry, has contributions from other sectors, and there is now mutual influence across the various sectors. In the future, we are going to see the interplay of different types of DTs, also across different sectors.

4.2 Fostering and Inheriting

The widespread adoption in the manufacturing industry has fostered adoption in other segments, as shown in Figure 19.

Healthcare started adopting DTs soon after manufacturing, mostly thanks to GE as they had adopted DT technology in turbine manufacturing and they have a Health Division. More recently, GE created a GE Digital concept that is further leveraging on the DT. In the Health sector, the extension of a DT to a PDT to mirror a person's unique characteristics was a natural step. At industry level, health has probably been the first to extend DT to people, and it is still leading.

Interestingly, we are starting to see, studies as the evolution of the PDT proceeds on Owned PDTs, or a PDT that is owned by the person. It might seem an obvious point, actually one might even wonder why a PDT shouldn't be owned by the person that it mirrors digitally, but as a matter of fact, it requires a conceptual leap and a corresponding technology leap.

Today, DTs (and PDTs) are developed by industry (or companies in other sectors), and their development, operation, and management requires specific skills and tools that are beyond a single person. In the coming years, I expect to see new biz opportunities to democratize the creation and use of PDTs, but we are not there yet. I also expect, once we reach that point, that companies will be willing to create a PDT, based on the data relevant to them, and then offer its ownership to the person, as today companies may create my profile and then offer an interface to manage it. Regulatory framework (in line of GDPR) might actually impose this kind of opening.

Clearly, a company created PDT will mimic a very small "slice" of me, like my entertainment tastes, my travel record, my health record, my working experience (in a specific company), … and so on. Once the concept of OPDT is established and socialized, I expect to see tools supporting the aggregation of the multiple PDTs mirroring parts of me into a single one, MyPDT. At that point, I will be in command of my PDT, I will be able to nurture and grow it, to decide how it can interact with the world, and more…

As shown previously in figure 19, the Cognitive Digital Twin is derived from industry (IBM 2018), and it is now being considered in the frameworks for education and knowledge asset management. A CDT related to a person (we can have a CDT related to a company, organization, institution, etc.) is also a PDT (or it may considered the part of the PDT mirroring the knowledge space of that person).

What we are starting to see, and it will become even more important in the coming years, is the interaction among PDTs and DTs, represented by the first arrow in figure 19 above. On

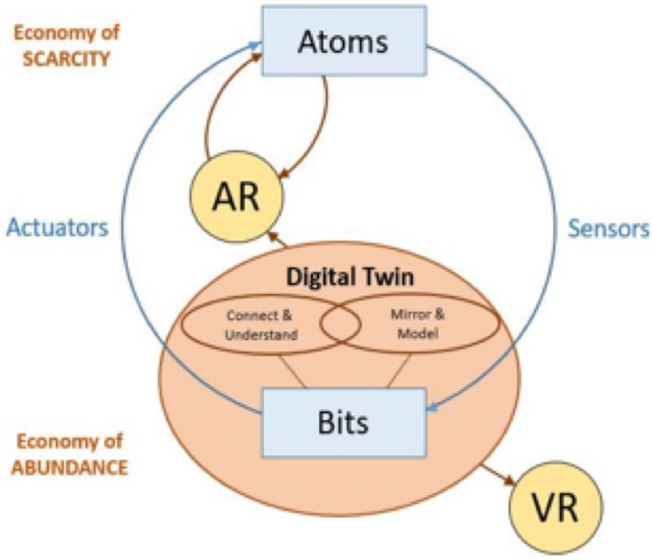

Figure 20. The Digital Transformation shifts a good portion enterprises into cyberspace. Digital Twins are a crucial component in the operation in cyberspace, and, more importantly, they can support interactions within the cyberspace through Virtual Reality, and overlay the cyberspace onto the physical space through Augmented Reality. In other words, they are both part of the operations in the cyberspace, and bridges linking cyber and physical spaces.

the shop floor, we are going to see that the interaction among workers and tools (robots) is mirrored in cyberspace, as are interactions among their respective (P)DTs.

A future step, represented in figure 19 by the second arrow (dark blue), is to have the knowledge of a person embedded and made accessible through that person's owned PDT. The accessibility can be asynchronous from the person, i.e. it can occur in the cyberspace in an autonomous way. A knowledge worker may "rent" her OPDT to a company to apply her knowledge to a situation. Clearly this would require monitoring of the interaction with a related valorization (i.e. that person will get some revenue out of it).

There is more, and this opens up a number of issues. Who is accountable for those interactions? One would tend to say that the person that harvests the revenue is also accountable for the interaction… However, it is more complex than it might seem. The executable knowledge that is provided (this is what we provide when we are involved in a work, not necessarily when we teach) depends on the owned knowledge, as well as on the context where it will be executed, on the knowledge that is harvested in the environment (and the environment in the cyberspace is the whole world!), and on the algorithms that are being applied locally and externally.

When we execute our knowledge, the algorithm guiding it is our brain, and execution capabilities are an important part of our assetts. When we are in a digital context, everything becomes fuzzier and complex. That is the reason why, for example, it is an open question to decide on accountability for a self driving car (the owner? the car manufacturer? the software provider? the data provider? the sensors provider? …)

We have to face difficult questions, but that shouldn't be surprising since the questions are brand new!

4.3 Digital Transformation

The Digital Transformation (DX) is ongoing and, actually, it has been accelerated by the pandemic as many companies have been forced to move as much of their activities as possible to cyberspace.

As shown in figure 20, Digital Twins are tools that support operation in the cyberspace for many entities and processes. Data resulting from the shift to the cyberspace creates a model of the entity/process in the physical space, and the stream of data provided by sensors supports the shadowing. In addition, all data continues forming the thread. Hence, the three components making up a Digital Twin are available as result of the DX.

Does the DX require Digital Twins? No. Would the DX benefit from Digital Twins? Absolutely.

Digital Twins provide entities with a structure in the cyberspace and connect them to their counterpart in the physical space. They provide a “method” and a standard of operation, through encapsulation of entity. This is particularly important when we are dealing with data and want to preserve their ownership while simultaneously sharing them. Rather than sharing the value a Digital Twin can make possible, to share its “meaning” through interactions. These interactions can be monitored and managed according to a framework. This applies to Digital Twins at stage 3 and beyond, particularly to those at stage 4 and beyond since at those stages Digital Twins offer functions and interactions with third parties (not those with the physical entity), and they are based on function activation.

One of the issues with shifting the physical to the “digital” is the loss of visibility. Seeing “bits” is not easy, and most of the times it is meaningless. Here is where DTs come in handy. They are models, and these models can be rendered and made visible in a meaningful way using Virtual Reality. This happens during the design phase, when the physical entity does not yet exist. Designers, even if they are scattered around the globe, can look and interact with the design—seeing what the physical entity will be like and even experimenting with it. Maintenance crews can use the models to see “inside” engines, to simulate ways of fixing a problem. Actually, an exciting application for Digital Twins and virtual reality is in the area of Surgery where surgeons can practice with the digital model of the patient by trying different approaches to a specified surgery.

Access to the cyberspace is not only supported by virtual reality, it also allows the possibility to use the Digital Twin to connect the cyberspace with the physical space through Augmented Reality. In this case, the Digital Twin can steer the rendering of data into the physical entity, or render a physical entity in a different physical ambient.

5. Bringing it all Together

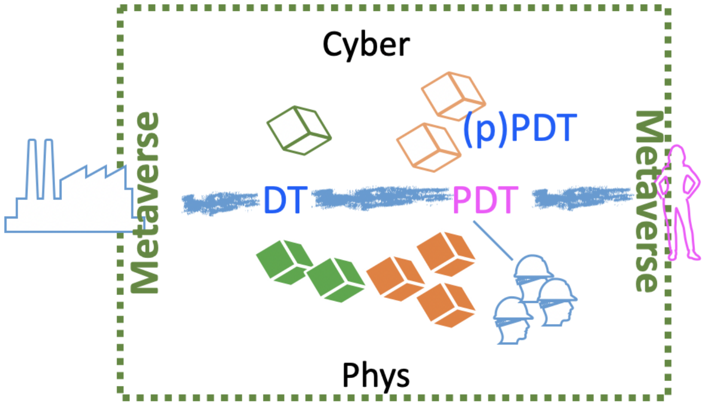

Figure 21. We can expect with reasonable confidence that the future will see a Digital Reality, both for Biz and each one of us—a reality where physical and cyber space will overlap, and where, from a perceptual point of view, the boundaries between the two will tend to fade away. Companies will have to operate in both spaces to reach customers that will "live" in both spaces.

The future of Digital Twins remains to be written, although what is going on now and what is going to happen in the coming 2-3 years is pretty straightforward.

DTs are a reality in manufacturing, and their use will further expand to cover more equipment and processes. The association of a DT to products will keep growing. While the DTs used "inside" factories are focussing on higher efficiency and flexibility, those associated to a product are considered tools to get data from operation. These are used to fine-tune production, enhance operation, and support pro-active maintenance. In the next few years, companies will learn to leverage DTs to offer services and generate additional revenue streams.

What is happening in the manufacturing area is spreading to several other areas, like construction, healthcare, insurance, finance and banking, retail, entertainment, and education.

Healthcare is steering the creation and the evolution of personal digital twins, and these will soon percolate in other areas as well.

The step from DT to PDT is not easy, mostly because the "P" brings along ethical and societal issues. Besides, Personal today refers not to the ownership, rather to what is mirrored, a persons set of characteristics. By the end of this decade, I am pretty sure that the "P" will be associated with the ownership. Whatever mirrors a person needs to be owned by that person.

The extension of DT and PDT to CDT (Cognitive Digital Twin) is in sync with the growing demand to manage knowledge. This is required fo the single person that is going to increasingly compete on the market on the executable knowledge available, and at the company level since knowledge is rapidly becoming tool of the trade.

Cognitive Digital Twins offer the hope of capturing and delivering executable knowledge. Here again, as with the "P" in PDTs, the "C" gives rise to new and yet to be explored ethical and societal issues.

The Digital Transformation is transforming much more than just the way of doing business. It is transforming the perception of values, and it has societal implications. It is a transformation into a life in the metaverse, a Digital Reality where people no longer perceive a separation between the physical and the digital cyberspace, and where businesses and companies needs to reach a market that is both in the cyber and in the physical space.

Advances in AR and VR will change the rules of the game once effective interfaces become available (seamless and affordable), and on one-hand can further blur the separation of the physical and cyber, while on the other hand making it possible to live in both at the same time.

As shown in figure 21 above, both companies and customers/users will operate in the metaverse, a space where the boundaries between the atoms and the bits are fuzzy. DTs and PDTs will bridge these two spaces. In addition, DTs and (p)PDTs (partial PDTs) will become products in themselves. Tools will be available to let people buy a DT and (p)PDT, and create their own PDT, a real alter ego in the cyberspace able to interact on their behalf with physical entities, people included, in the physical space.

These latter evolutions are still very much in the future, and may not happen any time soon as they will need to go hand in hand with evolution in many other areas, including AI, AR/VR transducers, cultural/societal acceptance and regulatory framework.

6. Acronyms

AI:	Artificial Intelligence
APM:	Assets Performance Management
AR:	Augmented Reality
AWS:	Amazon Web Services
BIM:	Building Information Modelling
CAD:	Computer Aided Design
CAM:	Computer Aided Manufacturing
CDT:	Cognitive Digital Twin
DT:	Digital Twin
DX:	Digital Transformation
EHR:	Electronic Data Record
GDPR:	General Data Protection Regulation
GE:	General Electrics
IA:	Intelligence Augmentation
ML:	Machine Learning
PPDT:	Owned Personal Digital Twin
PDT:	Personal Digital Twin
PhT:	Physical Twin
PLM:	Product Life Cycle Management
STM:	ST Microelectronics
VR:	Virtual Reality