Pervasive and Connected Digital Twins – A Vision for Digital Health

Published Version:
Ricci, A., Croatti, A., Montagna, S. (2022). Pervasive and Connected Digital Twins – A Vision for Digital Health. IEEE INTERNET COMPUTING, 26(5), 26-32 [10.1109/MIC.2021.3052039].

Availability:
This version is available at: https://hdl.handle.net/11585/797735 since: 2023-06-06

Published:
DOI: http://doi.org/10.1109/MIC.2021.3052039

Pervasive and Connected Digital Twins – A Vision for Digital Health

Alessandro Ricci
Computer Science and Engineering Department (DISI), University of Bologna
via dell'Università, 50, Cesena, Italy

Angelo Croatti
Computer Science and Engineering Department (DISI), University of Bologna
via dell'Università, 50, Cesena, Italy

Sara Montagna
Computer Science and Engineering Department (DISI), University of Bologna
via dell'Università, 50, Cesena, Italy

***Abstract*—Healthcare is a main domain where digital twins are being explored and applied. Following a healthcare 4.0 perspective, in this paper we discuss a vision in which digital twins are pervasively used to digitalize any strategic assets of a health organization. We start from introducing a specific real-world case study – concerning trauma management – and then we broaden the discussion identifying some main points and challenges of a wider vision.**

IN THE LAST LUSTRUM, digital twins have been explored and applied in different domains besides manufacturing and aerospace – which is where the original idea was introduced more than two decades ago. Besides the specific use, there is a basic agreement about what digital twins are, i.e. a virtual/digital image (or replica) of some physical entity – either living or non-living – which is typically considered an *asset* for some kind of organization or context (such as a product, a service, a process) [7], [9]. Among the domains, a foremost important one is healthcare, where the introduction of digital twins can bring enormous advantages, as already pointed out by [17].

In the healthcare context, literature reports on different approaches that apply digital twins to model different entities, from medical devices and hospital processes to human organs or people. They all share the idea that digital twins have a role to play in collecting and integrating significant amount of data from the healthcare context, with the goal to improve data-driven decision making by possibly combining the predictive power of simulation and machine learning [12].

However, digital twin based medicine is far from being an established fact yet.

In this paper we focus specifically on the application of digital twins to health IT. We introduce and discuss a broader perspective in which digital twins are adopted as the reference paradigm for implementing a vision of pervasive and connected digital healthcare, thus enabling the integration of different aspects that concern digital transformation in healthcare and healthcare 4.0 [3]. In the case of manufacturing, the impact of Industry 4.0 (technologies, methods) is so profound that it is considered an industrial revolution. The same is going to occur in healthcare: healthcare 4.0 is accelerating an innovation process which is going to have a deep impact at all levels, from users/patients to professionals/practitioners. In this context, we argue that digital twins can be adopted as a paradigm for driving a pervasive digitalization, that will concern all the strategic assets of an health organization and will support the development of more interoperable, open and collaborative health IT systems.

In the remainder of the paper, we start from seeing this role in the context of a specific real-world case study, in which we applied digital twin to Trauma Tracking [14]—in a project supported by the local healthcare institution. Then, we broaden the perspective, introducing and discussion a more general vision in which digital twin is being pervasively exploited for shaping health IT systems.

A CASE STUDY: DIGITAL TWINS APPLIED TO TRAUMA TRACKING

Trauma management is one of the most challenging healthcare scenarios where physicians can be involved. The trauma team, a team with strong heterogeneous expertises, must promptly identify a diagnosis and quickly provide medical aid, since patient outcomes strongly depend on the first hour of treatment.

The whole trauma management process can be conceptually split in two macro-phases: the *pre-hospital phase* – when emergency services personnel administers first aid and basic life support at the accident site – and the *operative phase* – when the patient is treated by the trauma team within the hospital emergency department. In the following, we will refer to these two phases as *PreH* and *Trauma* respectively.

The TraumaTracker project [14], [5] has been developed in this context with the main goal to support the leader of the trauma team in trauma documentation: the trauma leader produces real-time a report containing all the possible information on the ongoing trauma (procedures, administered drugs, diagnostics reports, patient's vital signs trace, and so on). The report is useful not only for a posteriori analysis but also during trauma management, providing the trauma leader with a comprehensive picture of the whole trauma to make the best informed decision.

At a high level of description, TraumaTracker can be seen as a (partially) proactive Personal Assistant of the trauma leader designed to observe and act accordingly with trauma team actions in order to produce a detailed trauma report. The system is in use for two years within the Trauma Center of an Italian hospital[1], and about 1000 trauma reports have been collected. During this time, a further important desiderata emerged, *i.e.*, the need of continuous monitoring of the complete state of the trauma, of the involved patient and care team, from the very beginning of the process (the pre-hospital phase) to the end of the trauma phase – possibly continuing the monitoring even beyond, during the patient hospitalization into the Intensive Care Unit (ICU).

Such a desiderata lead to a deep re-engineering of the TraumaTracker system, and the adoption of an agent-based digital twin, that models the two main phases of the process, seemed to us the most natural choice. This vision [4] could bring significant improvements. Among others, in particular:

1) offering a platform for an online and real-time monitoring of the ongoing operation of each trauma management session;
2) merge in an unique aggregate the whole data and information related to a particular trauma coming from different and heterogeneous sources;
3) allow for a more cohesive data collection enabling high level of data analysis;
4) providing a framework supporting and monitoring trauma management simula-

[1] Trauma Center of the "M. Bufalini" Hospital, Cesena, Italy.

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/MIC.2021.3052039

tions.

The basic idea behind a digital twin designed to support such a healthcare process is to consider the trauma management process as a physical asset which is suitably mirrored by two complementary digital twins, digitalizing both the involved phases related to the PreH phase (PreH digital twin) and the process related to the management of the trauma inside the hospital (Trauma digital twin). Figure 1 shows a conceptual representation of this design. This architectural choice follows the real evolution of trauma management. It is in the PreH phase that rescuers take in charge the case of the patient and decide if either the current situation is a severe trauma or not, and only in the former case the trauma team (and the second phase) is triggered. The physical assets and software agents in the two cases are different (refer to Figure 1).

The PreH Digital Twin

This digital twin represents the digital counterpart of the pre-hospital care process. Abstracting from details, the digital twin instance starts when the rescue central unit receives the call for assistance. The model of this digital twin involves:

- the *vehicle*, sent to the accident site with the rescuers;
- the EMT (emergency medical technician) and the *rescuer*;
- the *accident place* and, in particular, the *accident dynamic*;
- the *patient*.

This digital twin collects real-time information considering the information given by the central unit, the GPS System of the vehicle, and the smart devices held by rescuers to compile emergency forms. During its life-cycle, the most relevant moment is related to the transition to the state where rescuers decide the degree of severity of the ongoing trauma (considering patient's health state and its GCS–Glasgow Coma Scale value).

The Trauma Digital Twin

This digital twin represents the operative phase of trauma management and starts when, in the previous phase, the trauma is marked as severe. The fact that this digital twin starts before the arrival of the patient to the hospital is very important for this case study. In this way, the trauma team is pre-alerted about the incoming patient and start to collect and receive information directly from the accident site. Its internal state changes when the patient is delivered to the emergency department, that is, when the trauma team starts in taking care of the patient. The model of this digital twin involves the following assets:

- the *patient* and all information flow coming from *connected devices* (i.e., the vital signs monitor);
- the *trauma leader*, and the trauma team members;
- the *shock-room*, the room of the emergency department where generally each trauma is managed;
- other tools and equipment (e.g., rapid *diagnostics machinery*, *displays* to show real-time information of the ongoing trauma, and so on).

Some data composing the internal state of this digital twin come from the previous digital twin, others are collected by the coupled assets—we use the term *shadow* to indicate the coupling between physical assets and digital twin in order to capture the real-time synchronization between physical assets and their digital representations. The Trauma digital twin life-cycle contemplates the macro-steps related to trauma management – most of them occurring within the Emergency Department Shock-Room, but also within other hospital places related to the trauma path such as the diagnostic room – and terminates when the patient is ready for the hospitalization, most of the times into the Intensive Care Unit (ICU) department.

Software agents

From the agent perspective, the full digital twin system is observed and accessed by several agents acting as personal assistants of involved professional figures or as managers of involved rooms and places. For instance, the agent acting as a personal assistant of the trauma leader, assists this physician for the trauma documentation, updating the digital twin state considering the ongoing performed procedure and administered drugs. Agents behave according to the digital twin state, updating their belief upon it. Some of them, in particular the trauma leader personal assistant, can interact with a dedicated cognitive service [15] adapts alerts to the specific patient history and

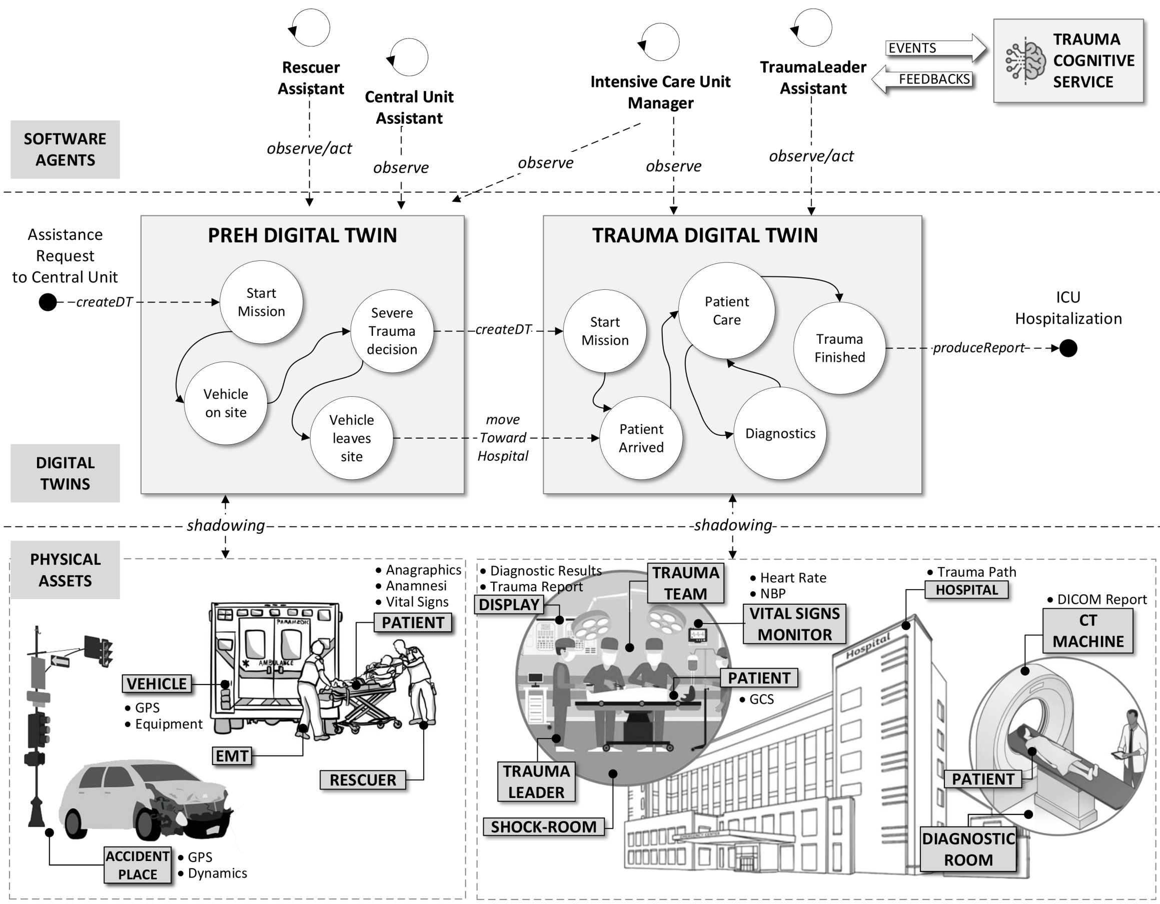

Figure 1. A conceptual representation of the involved digital twins for the trauma management process

context. Accordingly, the set of agent's beliefs, and possibly plans, is updated real-time during the process of trauma management. The Trauma Cognitive Service has been trained with data archived in the Trauma Store service, which is the knowledge base representing the shared data base of all trauma. Reasoning eventually provides a set of new feedback, e.g. further alerts or clinical indications, that are sent to the trauma leader assistant agent. Such feedback can be collected into the Trauma digital twin and proposed also to the trauma team exploiting, e.g., the digital shadow of the pervasive display in the shock-room.

DIGITAL TWIN FOR PERVASIVE DIGITAL HEALTH

The approach adopted in the trauma case study can be broaden towards a more general view, in which:

every strategic physical asset of an health organization has a corresponding digital twin, mirroring and augmenting its functionalities and services, eventually linked to related digital twins in an health IT ecosystem

The range of strategic assets is broad—from patients, devices, buildings to also processes physically occurring in the healthcare environment.

This view brings three main benefits, that we encountered in the trauma case study but that can be devised in general. A first one is about avoiding silos, and enabling a high-level of interoperability and collaboration among all the stakeholders that access and use assets, by a deeper exploitation of IT and ICT technologies. A second one concerns the effects of better supporting the connection and convergence of the IT and OT (Operational Technology) level in healthcare environments and organizations, by exploiting technologies such as the Internet of Medical Things [11], so as to strongly improve reactivity and efficiency in all processes. The third main benefit is about building the solid ground that can be leveraged by smart applications and systems running on top of it, including software agents like digital companions and personal assistant agents – as in the case of TraumaTracker.

In order to apply this view to any healthcare application, we identified three main points and challenges, that are discussed in the remainder of the section.

Digital Twins and Interoperability

Interoperability is a key issue in health IT, and specific standards have been developed, at different levels. A main example is FHIR (Fast Healthcare Interoperability Resources)[2], which is the most recent evolution of HL7[3] aimed at defining a standard for exchanging healthcare information electronically. Another main one is the HISA - Health Informatics Service Architecture (ISO 12967), which provides a blueprint at different viewpoints – enterprise, information and computational – to shape health service architectures. These standards then should be considered as a main reference in the design and implementation of digital twins, in order to concretely enable full interoperability among the different stakeholders sharing and collaboratively using them.

To that purpose, a clear direction that emerges for designing FHIR/HISA-compliant digital twin is the adoption of REST software architectural style and REST web-based service-oriented architectures, which clearly inspired by REST and most recent version of HISA[4]. In particular, FHIR is described as a RESTful specification, adopting the concept of linkable Resource as basic building block to define the information contents and structure for the information set shared by concrete implementations. FHIR servers allows to access and interact with resources by means of REST API.

Accordingly, the information model of a FHIR-compliant digital twin then can be developed on top of the basic types of resources available in FHIR. For instance, the Digital Twin of a patient may be based on the FHIR Patient resource[5] – which keeps track of demographics and other administrative information about a specific individual receiving care or other health-related services – and a collection of FHIR Observation resources[6], which can be used to define the current state of the patient according to measurements and assertions available, which are relevant for the purpose of the digital twin. Analogously, the digital twin of a device – which is a quite common and simple example – may be based on the FHIR Device resource[7] – used to track individual instances of a device and their location – as well as FHIR Observations concerning the last measurements performed by the device. The digital twin of a process like that an Ongoing Trauma – presented in previous section – could be based on the FHIR Encounter resource[8] including – among the others – FHIR Observations about the vital signs of the patient and FHIR Procedure[9] and Observation resources, organized following the FHIR Event pattern[10], to keep track of the relevant events occurred during the trauma management. Besides the current state, a digital twin meta-model typically includes also observable relevant *events* related to the physical asset as well as possible *actions* that can be requested on the digital twin having effect on the physical counterpart.

So a digital twin – in spite of the specific (meta)model to be adopted – could be conceived as a further resource dynamically aggregating (linking) a set of FHIR resources useful to represent the current state of the physical asset, in the shadowing process. A layer-based representation is shown in Figure 2—the digital twin layer is built on top of a FHIR-based Health Information System layer. Actually here it is important to given an explicit and rigorous semantics about *current state*. Generally speaking, all the FHIR Observation resources that are used to represent the current state in the digital twin of some physical asset should refer to a common temporal timeframe – that must be defined accordingly.

[2] https://www.hl7.org/fhir/overview.html
[3] http://www.hl7.org
[4] ISO 12967-1:2020, https://www.iso.org/standard/71037.html
[5] https://www.hl7.org/fhir/patient.html
[6] https://www.hl7.org/fhir/observation.html
[7] https://www.hl7.org/fhir/device.html
[8] https://www.hl7.org/fhir/encounter.html
[9] https://www.hl7.org/fhir/procedure.html
[10] https://www.hl7.org/fhir/event.html

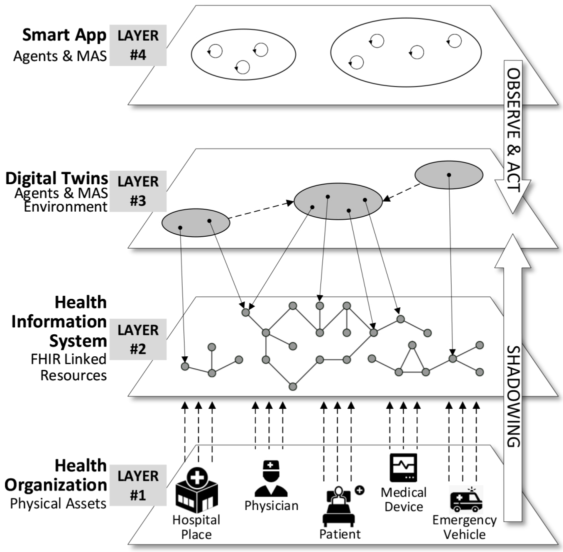

Figure 2. A layered representation integrating Digital Twins with FHIR-based HIS and Agents and MAS, as a high-level abstract blueprint for conceiving stage 4-digital twins in healthcare.

Towards a Web of Digital Twins

Generally speaking, the idea of pervasively introducing a digital twin for each strategic physical asset of an health organization leads us to consider the issue of managing a large set of digital twins, of different kinds, eventually *linked* by semantic relationships as defined by e.g. FHIR standards. This opens up issues concerning the design and availability of proper middleware and infrastructures providing services to create, execute, manage, monitor digital twins, featuring a proper level of scalability, reactivity, etc.

In this case, besides the general guidelines provided by standards such as the ISO 12967, we argue that an important lesson comes from the Web of Things W3C initiative[11], which is about exploiting the web as basic platforms for enabling interoperability at the application level for Internet of Things systems. Digital Twins appear among the architectural patterns that are considered in the WoT architecture proposal. A clear direction suggested in this case is to conceive and design digital twin platforms for executing and managing digital twins as hypermedia-based decentralized middleware based on the web and its technological stack.

[11] https://www.w3.org/WoT

Leveraging Extended Digital Twins by means of Multi-Agent Systems

The stage 4 digital twin envisioned in [17] is about extended digital twins delivering additional capabilities besides the physical asset ones, eventually including an autonomous part flanking the basic digital twin ones. In literature, agents and multi-agent systems (MAS) are a main reference paradigm for modeling and engineering autonomous, possibly distributed, systems, providing well-defined cognitive architectures for integrating AI techniques such as learning and planning [10], [18]. Therefore, the identification of effective approaches to exploit agents and MAS on top of digital twins appears a general relevant research direction to explore for shaping digital twin at stage 4.

The approach that we used in the trauma case study accounts for seeing digital twins as the basic blocks structuring the *environment* where agents and MAS are logically situated (layer 3 and 4 in Figure 2). That is: software agents can perceive the observable state exposed by digital twins – eventually modeled in terms of FHIR resources and delivered by means of web-based standard protocols – and act on them by means of their exposed API. Main examples of software agents in this case could be the personal assistant agents (also referred as digital companions) that are used to assist patients or support practitioners' individual and cooperative work. In this view, the agent layer and the environment (digital-twin based) are loosely coupled – allowing software agents to dynamically discover and access digital twins for their purposes, in an open system perspective.

RELATED WORKS

Scientific literature already identifies in healthcare one of the key domains for digital twins application. Moreover, different companies already developed products that adopt digital twins to improve the healthcare quality, efficiency and coverage, e.g. [8].

Without claiming to be exhaustive, some examples follow. GE Healthcare's understanding of digital twin focuses on hospital management and processes design, with the purpose of optimizing medical pathway planning and medical resource allocation [12]. Some attempts in industries have been made to create digital twins of human organs, such as Siemens Healthineer, that developed a digital twin of the heart simulating the physiological processes of a patient's organ with the goal to visualize responses to treatment before the intervention[13]. However few data are still available in literature reporting successful clinical applications. Digital twins of patients have been introduced as an *in-silico* representation of an individual. In literature they are adopted as a bag for accumulating, aggregating and representing medical data [13]: the twin stores medical and health data, laboratory results, medical imaging, individual genome sequences, vital signs, lifestyle of an individual person. Data usually come from different sources and are typically heterogeneous. The main goal is to feed analytical tools, that exploit simulation and AI algorithms, to enable personalized medicine [2]. The underlying vision is "Making mistakes on computer models instead of people[14]". An example is available in [1], where digital twins are adopted for identifying the best drug among the thousands possible to treat a certain disease. There, a digital twin of a patient with symptoms of a specific disease is developed in unlimited copies, based on network models of all molecular, phenotypic, and environmental factors relevant to disease mechanisms. Simulations with different drugs are then performed to identify the treatment. A combination of these applications is already reported in literature: for instance, DTH (Digital Twin Healthcare) is presented in [12] and models different physical objects: medical devices, wearable devices, patients, and external factors such as social behavior or government policy.

CONCLUDING REMARKS

In the context of healthcare, digital twins are being explored and applied for different specific purposes. In this paper we focused on the application of digital twins in health IT, injecting a vision where DTs are pervasively used to digitalize any strategic assets of health organizations, on top of existing standards such as FHIR and as the base ground to conceive and develop smart applications and systems on top. We see this vision and effort aligned with other even-broader initiatives, such as the National Digital Project proposed by the Centre for Digital Built Britain[15] and related principles (defined as the *Gemini Principles*[16]).

Besides the specific application domain, the perspective about DTs adopted in this paper is strongly based on the vision of agent-based mirror worlds discussed in [16], which is rooted on Gelernter's seminal mirror world idea introduced three decades ago [6]. An important aspect of that perspective is about *augmentation*, so that a mirror world – and correspondingly an extended digital twin – could be useful to augment a physical environment with functionalities to better support the individual and collaborative work of the persons situated in that environment—in a smart environments perspective. Such an augmentation includes also mixed reality extensions, e.g. holograms shared and used by the persons in the physical environments, but defined by a computational entities executing in the digital layer. Augmentation appears an interesting and relevant aspect to be further explored for digital twin as well, in particular for digital twins mirroring the environments where e.g. practitioners work and cooperate.

[12] https://www.gehccommandcenter.com/digital-twin
[13] https://www.siemens-healthineers.com/en-be/press-room/press-media-gallery/im-20190412002shs.html
[14] https://www.digitwins.org/
[15] https://www.cdbb.cam.ac.uk/DFTG
[16] https://www.cdbb.cam.ac.uk/DFTG/GeminiPrinciples

REFERENCES

1. Bergthor Björnsson, Carl Borrebaeck, Nils Elander, Thomas Gasslander, Danuta R. Gawel, Mika Gustafsson, Rebecka Jörnsten, Eun Jung Lee, Xinxiu Li, Sandra Lilja, David Martínez-Enguita, Andreas Matussek, Per Sandström, Samuel Schäfer, Margaretha Stenmarker, X. F. Sun, Oleg Sysoev, Huan Zhang, Mikael Benson, and on behalf of the Swedish Digital Twin Consortium. Digital twins to personalize medicine. *Genome Medicine*, 12(1):4, 2019.
2. Koen Bruynseels, Filippo Santoni de Sio, and Jeroen van den Hoven. Digital twins in health care: Ethical implications of an emerging engineering paradigm. *Frontiers in Genetics*, 9:31, 2018.
3. Janya Chanchaichujit, Albert Tan, Fanwen Meng, and Sarayoot Eaimkhong. *Healthcare 4.0: Next Generation Processes with the Latest Technologies*. Springer, 01 2019.

4. Angelo Croatti, Matteo Gabellini, Sara Montagna, and Alessandro Ricci. On the integration of agents and digital twins in healthcare. *Journal of Medical Systems*, 44(9):161, 2020.
5. Angelo Croatti, Sara Montagna, Alessandro Ricci, Emiliano Gamberini, Vittorio Albarello, and Vanni Agnoletti. Bdi personal medical assistant agents: The case of trauma tracking and alerting. *Artificial Intelligence in Medicine*, 96:187 – 197, May 2019.
6. David Gelernter. *Mirror Worlds or the Day Software Puts the Universe in a Shoebox: How Will It Happen and What It Will Mean*. Oxford University Press, Inc., New York, NY, USA, 1991.
7. Edward Glaessgen and David Stargel. The digital twin paradigm for future nasa and us air force vehicles. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 20th AIAA/ASME/AHS Adaptive Structures Conference 14th AIAA*, page 1818, 2012.
8. Siemens Healthcare GmbH. The value of digital twin technology transforming care delivery in a single department or an entire healthcare enterprise, 2019. White Paper.
9. Michael Grieves and John Vickers. *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*, pages 85–113. Springer International Publishing, Cham, 2017.
10. Nicholas R. Jennings. An agent-based approach for building complex software systems. *Commun. ACM*, 44(4):35–41, April 2001.
11. G. J. Joyia, M. Liaqat, R, A. Farooq, and S. Rehman. Internet of medical things (iomt): Applications, benefits and future challenges in healthcare domain. *Journal of Communications*, 12:240–247, 2017.
12. Ying Liu, Lin Zhang, Yuan Yang, Longfei Zhou, Lei Ren, Fei Wang, Rong Liu, Zhibo Pang, and M Jamal Deen. A novel cloud-based framework for the elderly healthcare services using digital twin. *IEEE Access*, 7:49088–49101, 2019.
13. Rainer Lutze. Digital twins in ehealth – : Prospects and challenges focussing on information management. In *2019 IEEE International Conference on Engineering, Technology and Innovation (ICE/ITMC)*, pages 1–9, 2019.
14. Sara Montagna, Angelo Croatti, Alessandro Ricci, Vanni Agnoletti, Vittorio Albarello, and Emiliano Gamberini. Real-time tracking and documentation in trauma management. *Health Informatics Journal*, 26(1):328–341, 2020. PMID: 30726161.
15. Sara Montagna, Stefano Mariani, Emiliano Gamberini, Alessandro Ricci, and Franco Zambonelli. Complementing agents with cognitive services: A case study in healthcare. *Journal of Medical Systems*, 44(10), September 2020.
16. Alessandro Ricci, Michele Piunti, Luca Tummolini, and Cristiano Castelfranchi. The mirror world: Preparing for mixed-reality living. *IEEE Pervasive Computing*, 14(2):60–63, 2015.
17. Roberto Saracco. Digital twins: Bridging physical space and cyberspace. *Computer*, 52(12):58–64, 2019.
18. Michael Wooldridge. *An Introduction to MultiAgent Systems*. Wiley Publishing, 2nd edition, 2009.

Alessandro Ricci is with the Department of Computer Science and Engineering (DISI) of *Alma Mater Studiorum* Università di Bologna. His general research line is about agents and multi-agent systems as a paradigm for engineering software systems. Specific investigations include the agent-based mirror worlds and digital twins. Contact him at a.ricci@unibo.it.

Angelo Croatti is with the Department of Computer Science and Engineering (DISI) of *Alma Mater Studiorum* Università di Bologna. His research is about the engineering of pervasive software systems, integrating multi-agent paradigm and augmented reality technologies. Healthcare is the primary application domain of his studies. He got his M.S. degree cum Laude in Computer Science and Engineering in 2014 and his Ph.D. in Computer Science and Engineering in 2019. He has co-authored several research publications, both on international journals and conference proceedings. Contact him at a.croatti@unibo.it.

Sara Montagna is with the Department of Computer Science and Engineering (DISI) of *Alma Mater Studiorum* Università di Bologna. She is working in between biology and computer science (BIO-ICT convergence) studying the application of techniques devised in computer science – such as agent-based modeling and simulation, Boolean Networks, Cognitive Computing algorithms – to analyzing complex bio/healthcare scenarios. She got her M.S. degree in biomedical engineering in January 2007 (110/110 cum Laude) and her Ph.D. degree in computer science in April 2011. She has published more than 50 papers, out of which 17 on renowned international journals. Contact her at sara.montagna@unibo.it.