University of Zurich[UZH]

Digital Society Initiative

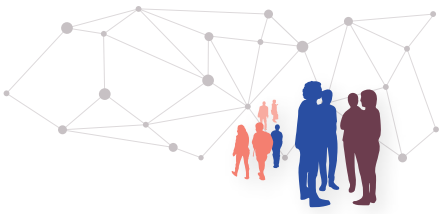

Position paper

Artificial Intelligence in Medicine – Objectives and Recommendations for the Responsible Use of Digital Twins

Editorial team:
Nikola Biller-Andorno, Markus Christen, Michael Krauthammer, Claudia Witt

Introduction

In the next few years, the use of artificial intelligence (AI) in medicine will influence the possibilities in prevention, diagnosis and therapy as well as the associated processes and responsibilities in healthcare to a large extent that is currently still difficult to assess. For this reason, experts from the University of Zurich and other organizations have addressed the question of which objectives and associated recommendations can lead to a responsible use of AI in medicine in the "Strategy Lab" of the Digital Society Initiative (DSI).

The possible range of applications of AI in medicine is vast. In order to facilitate a focused discussion, a form of application that is (currently) still in the future was selected to exemplify the impact of AI in medicine. The objectives presented here are based on the assumption that core issues of the future application of artificial intelligence (AI) in medicine can be defined by the concept of a "**digital twin**". A digital twin is a representation of an object from the real world in an information system. The functioning of a digital twin is dependent on the availability of sufficient data about this object (which may even be continuously collected) and the existence of computer models that use artificial intelligence methods to describe, predict and influence the properties or behavior of the represented object or to offer services about it. In the case of medical applications, data on health functions would then be collected and computer models would be used which, for example, could reproduce central vital functions of the original human such as respiration, circulatory function or digestion and, by means of simulations, make predictions of future health states.

This position paper was produced in a systematic and participatory process as part of the Digital Society Initiative's "Strategy Lab Artificial Intelligence in Medicine". The process extended from early 2022 to mid-2023 and included, among other things, a survey of medical subject matter experts in spring 2022, two workshops for the development of future scenarios in June and August 2022, a focus group workshop with representatives of the public in September 2022, and a workshop explicitly for the development of the objectives and recommendations in November 2022. The future scenarios and resulting objectives and recommendations were presented and discussed with stakeholders (including health professionals and citizens) at various events.

Participants in the workshop to develop the objectives and recommendations were, in addition to the editorial team: *Matthias Baumgartner, Abraham Bernstein, Susanne Gedamke, Janna Hastings, Manya Hendriks, Jeffrey David Iqbal, Christian Kauth, Lorena Kegel, Birgit Kleim, Viktor Kölzer, Tanja Krones, Titus Neupert, Verena Pfeiffer, Cristina Rossi, Luzia Rüdlinger, Reinhold Sojer, Bernd Stadlinger, Jade Sternberg, Reto Sutter, Viktor von Wyl and Andreas Wicki.*

We would like to thank all these people for their contribution to this process and their comments during the final drafting of the recommendations; in particular also Jeffrey David Iqbal for the operational support of the Strategy Lab process and Florent Thouvenin for feedback in legal matters.

Digital twins as a service

These objectives assume that digital twins will be part of everyday life in the future. However, the idea that every person has a single, fixed digital twin that is always

assigned to that person is rather unrealistic. Instead, the digital twin should be understood as a service offered by different providers using different data and AI technologies – analogous, for example, to today's common software subscriptions. Accordingly, people are likely to have several digital twins one day, which could also communicate with each other under certain circumstances (or exchange data). In the following, we therefore speak of "digital twin services" (DTS). Not all DTS will necessarily involve health aspects; they could also be used for pure entertainment purposes (e.g., new forms of computer games). Some DTS could also be in the health-related field and include, for example, topics such as wellness, psychological well-being or spirituality.

In this document, we will only talk about DTS that clearly have a **health-related function**, i.e., can perform at least one of the following four functions:

1) They support a medical **diagnosis** (e.g., by allowing the identification of parameters suitable for the diagnosis of a potentially developing disease).

2) They support the **prognosis of** a disease (e.g., by predicting the possible course of a disease matched to a person's individual characteristics and environmental factors).

3) They support the **therapy of** a disease (e.g., by using a digital twin to simulate the possible effect of different therapies on a person).

4) They support the **prevention of** disease and allow the identification of appropriate preventive measures (e.g., by alerting a person that certain behaviors may have unhealthy consequences).

Important distinctions

In the following, the term **"digital twin of person X"** shall only be used for the complex "data of X" and "AI model trained/updated with this data". Thus, the following in particular must be distinguished in technical terms:

A. In a first phase of the creation of a digital twin, presumably a lot of impersonal basic data (purely statistical quantities such as average life expectancies, data from basic biomedical research or pretrained deep *learning models*) will be used to create a kind of basic model of a certain function of the digital twin (for example: model of cardio-vascular function). We refer to this base model data as D_B. This base model may incorporate machine learning techniques (exactly what these will be in the future is difficult to determine); however, the base model is not yet a digital twin. The process of generating the basic model is therefore not the subject of these recommendations; this is a process of research and development to which the usual ethical and legal standards naturally apply, although adaptations and additions to the currently applicable law may be expected on the legal side.

B. To transform the basic model into the digital twin of a person X, personal medical data D_X of this person X is collected (presumably partly on a continuous basis, e.g. using implanted and environmental sensors). In this step, the digital twin is created. For the collection of these data D_X, which are usually the result of diagnostic or therapeutic interventions, person X has given consent (principle of informational self-determination). Since this consent must be informed, it is implicitly given that the design and type of use of the digital twin must also be clear (because that is what the data is collected for) and in a way that the person concerned understands. However, consent is not required for every type of data processing around the creation of a DTS, as long as the principles of data processing according to the Data Protection Act are adhered to. Questions may arise here, because in all likelihood the design and type of use of DTS may go beyond what one can know as a consenting person at the time of consent.

C. After completing step B while observing certain quality criteria (e.g., regarding the prediction accuracy of the digital twin), the digital twin is now an entity that is the subject of these recommendations.

Technically, the digital twin is a software service provided by a vendor. However, this entity is not fixed in the sense of a static software program, but is continuously changed (or ideally improved with regard to the specific function of the digital twin) by a continuous influx of data D_X of the person (this can then legally include further contracts with the providers of the data sources, e.g., sensors such as smartwatches, etc.) or also improved basic data D_B. In this process, new data D_{DTX} ("DTX" stands for "digital twin of person X") are now created as a result of the interaction of the health data D_X of person X and the personalized basic model that was created in step B. An example of such data is a time series of person X's blood pressure modeled into the future, which in turn is the basis for determining preventive measures. Legally, these data D_{DTX} are also to be regarded as personal health-related data of person X.

D. A large diffusion of DTS will lead to the availability, in principle, of enormous amounts of health-related data (D_X and D_{DTX}), which in turn could be the subject of population-based analyses for public health purposes (and presumably for other purposes as well). This potential is enabled by DTS, but would require various other technical innovations that could, for example, ensure data interoperability and privacy, or involve building and operating a comprehensive infrastructure for collecting data on individuals. However, the eventual exploitation of the potential is not a DTS, although machine learning methods can of course be used here as well. However, since DTS enable precisely this potential, corresponding recommendations for exploiting this potential form part of the following considerations.

This conceptual clarification will be followed by basic goals for the future considering the perspectives of four main groups of stakeholders. The DSI Strategy Lab believes that those goals will ensure that DTS are used in a way that is positive for promoting health from an individual and societal perspective.

Goals from the perspective of citizens or patients

From the perspective of citizens and patients – i.e. ultimately all people from whom digital twins will one day be created for health purposes – the DSI Strategy Lab believes that the following goals should be achieved in the future:

1. Citizens decide on the generation, data sources, design, type of use, and lifetime of their personal digital twin services.

2. The relationship of trust between patients and other healthcare stakeholders is preserved through the use of digital twin services.

3. Citizens are empowered through measures in education to understand personal digital twins in the best possible way and to be able to use them in accordance with their values and interests.

The first objective is to ensure that the development of DTS is oriented toward the needs of citizens; it is not the digital twin that is the "patient", but the individual person with his or her values, preferences and priorities. In particular, citizens should retain the rights of disposal over all data that is collected for the functioning of individual DTS and also arises through the use of DTS. The current level of data protection should be maintained; however, new questions arising with the use of DTS should also be taken into account – for example, the question of what should happen to a person's digital twin after his or her death. The implementation of this goal places specific demands on legislators, the state and manufacturers, which cannot be listed in detail here.

The second objective is to work toward conditions of use that not only maintain or promote patient trust, but also justify it. To this end, it must be transparent and comprehensible to lay-persons how data are collected, which data are combined, what data are used for, and under what conditions. In this context, the commercialization of data is a particularly sensitive issue. If citizens have rights of disposal over their data, they must also

be asked when their data is resold to third parties, for example. In this way, a justified relationship of trust can be established between the users and those who provide DTS (or the infrastructure and legal framework). This does not exclude that the status quo and developments are critically questioned again and again. Health professionals, manufacturers/operators of DTS and legislators in particular must contribute to achieving this goal.

The third objective is to ensure that the use of DTS becomes part of the curricula in public schools as part of comprehensive health education and is also addressed in an appropriate form in vocational and adult education. This is less about teaching detailed technical knowledge about how digital twins work, and more about creating a basic understanding of the possibilities and limitations of using DTS. As a result, providers of DTS are also obliged to provide users with sufficient information about how DTS works and about the possibilities, risks, limits and biases of DTS – not only at the beginning when entering into a corresponding service contract, but also at regular intervals. Healthcare professionals will be trained more comprehensively on the specific functionality of DTS in appropriate education and training courses. Achieving this goal places particular demands on the education system and on the manufacturers/operators of DTS.

Goals from the perspective of health professionals

From the perspective of health professionals – i.e. representatives of all occupational groups with medical or medical-related training who are involved in the health care of citizens (including physicians, nurses, psychologists, pharmacists, occupational therapists, etc.) – the DSI Strategy Lab believes that the following goals should be achieved:

1. Digital twin services are integrated into interprofessional treatment teams, where the necessary competencies are available and responsibilities are clarified.
2. Healthcare continues to enable care for individuals who do not want to use digital twin services.
3. The setup and operation of a digital twin service infrastructure works internationally, so that location-independent use of digital twins is possible.

The first objective clarifies that DTS should not be understood as a conventional "tool", but that they can actively contribute to care and expand the interprofessional treatment team accordingly and change the interactions in the team. This affects necessary competencies and also responsibilities in the team, which are addressed equally. For example, the professions in the team need appropriate digital and communication skills, for example regarding the basic functioning of a DTS for the purpose of assessing the validity of a DTS prognosis. The current curricula must already be adapted with this goal in mind. It should be ensured that the assessments of patients and health professionals are given greater weight in medical decision-making than those of DTS. With regard to responsibility, the responsibility for the quality of DTS and the associated liability for quality defects of DTS must lie with the manufacturer or operator of the DTS and not with the health professionals. The added value of a DTS must be tested and proven in clinical studies, as is the case with medical devices. In this way, manufacturers and operators fulfill their duty of care. Providers of DTS should be able to be held liable for incorrect prognoses of their DTS, or it should be examined in more detail how exactly liability issues that can arise due to misdiagnoses or incorrect predictions of DTS should be regulated. This objective addresses in particular the training centers for medical professionals, the professional associations and the providers of DTS.

The second goal is to help prevent health care disparities with digital twins. In principle, access thresholds to DTS should be reduced. Incentives for use, e.g., lower health insurance costs, may also be provided. However, there should be no compulsion to use DTS. Therefore, care structures for people without digital twins must also persist. In this context, it is also important to ensure that the effects of the use of DTS on public health – particularly

with regard to psychological and social aspects – are systematically recorded. The state or health service providers with a state mandate are particularly called upon here.

The third goal is to ensure international compatibility and corresponding global cooperation of DTS, among other things so that this service is also available to people during travel and migration. Associated with this are requirements for data (protection) standards, but also other technical standards. Particularly in view of the current geopolitical situation, it should be prevented that health technologies such as DTS can only be offered in selected countries at all. Access to DTS from manufacturers in different countries should also be possible. For the implementation of this goal, international organizations in the field of standardization in particular, but also the international community, and finally also the healthcare institutions of the respective countries, which should implement the standards, are called upon.

Goals from the perspective of DTS manufacturers/providers

From the perspective of manufacturers and providers of digital twin services – these can be technology companies, startups, pharmaceutical companies, medical device companies, but also universities, etc. – the DSI Strategy Lab believes that the following goals should be achieved:

1. Digital twin service providers have access to as much anonymized health data as possible according to open data principles (open standards, interoperability).

2. Medical services are billed to reimburse for recommendations of high-quality and appropriately certified digital twin services.

3. Authorization procedures, certification and regulation of digital twin services, and information requirements for providers of such services are defined and agile.

The first objective aims to ensure that there is an innovation-friendly environment for the development of DTS. It clarifies that development and deployment of DTS requires a clear data strategy that enables sharing of health data at the national and international level, relies on open data principles, and promotes appropriate data standards. Since access to data is central to DTS innovations, the particular aim is to ensure that not only quasi-monopolists can build large data pools, but that these pools are in principle accessible to all innovators. Part of this data strategy is also the establishment of a government infrastructure for pooling data (see following section). In principle, all actors (government administration, research, companies) can access data from this pool after applying statistical anonymization techniques while complying with governance processes to be defined. Research in the field of anonymization techniques is also to be promoted. This is intended to promote innovation in the field of DTS, which can be driven by research and development by both public and private actors. Furthermore, healthcare providers (universities/hospitals; pharmaceutical and medical device companies, etc.) should have access to DTS for research and innovation. Both government institutions and the manufacturers and operators of DTS are called upon here in the joint development of such standards.

The second goal is to ensure the economic basis for the development and operation of DTS. Appropriate regulations and incentives are needed for the use of digital twins, including the possibility for DTS providers to make and implement decisions independently of medical specialists and to be reimbursed for these decisions. In this context, it would have to be clarified whether there is an obligation to treat or rather an obligation to reimburse on the basis of proposals made by a DTS. Adjustments are needed in the regulation of the distribution of competencies (for example, therapy decisions and the right to billing) in medicine with a focus on both new players such as big tech/pharma and patients who are "empowered" by digital tools. However, providers of DTS in the healthcare sector must also prove that their diagnostic/therapeutic recommendations are safe, that they function in accordance with the defined goals, and

that they contribute to a more cost-effective, but at the same time high-quality healthcare system that is accessible to all. In particular, the state and the health insurance companies are called upon to do this.

The third objective clarifies that innovation with DTS requires an agile regulatory framework. This includes both the approval procedures for DTS, including studies to assess benefits and risks, and education about potential side effects of DTS. This also includes adjustments in the regulation of drug development, specifically the consideration of data and simulations as evidence of drug efficacy. Technologies are changing much faster than it is possible to adapt corresponding rigid regulatory frameworks. Therefore, anticipatory change is needed now so that future regulatory frameworks allow for agile adaptations, which will be particularly challenging for regulators. The bodies of the regulatory authorities, in particular, are challenged here and should optimize their processes.

Goals from the perspective of regulators and payers

From the perspective of regulators and payers of digital twin services – this concerns both organs of state health administration and regulatory bodies, i.e., health insurance companies that implement state directives regarding cost efficiency of health care – the DSI Strategy Lab believes that the following goals should be achieved:

1. The state ensures the provision of a data infrastructure by means of which citizens can bring together data sources from all areas of life.

2. Certain health-related data generated or made available through digital twin services that are of greatest public health importance will be made available in anonymized form to third parties through this data infrastructure.

3. The state establishes benchmarks for quality and security of digital twin services.

The first objective is to help ensure that the functions "cross-domain merging of data" and "technical use of data" are separated in technical and regulatory terms. The aim is to ensure that data sets from the most diverse areas of life, which are indispensable for the use of DTS, can basically only be merged in one place. Instead of large quasi-monopolists attempting to bring together data from all areas of life, as is the case today, this bringing together should take place on an infrastructure created by the legislature and administered by the federal government, whose contents, however, can only be viewed and controlled by citizens. Every citizen has the right to store personal data in the individual "data account" and the individual then decides which third parties can use this data (subject to certain data falling under Objective 2). Data that DTS generate about the data subject should also be able to flow into that individual's data account. This data infrastructure is a prerequisite for DTS innovation in that DSZ providers can then access the data collected there in anonymized form. It also supports the creation of open data standards to promote the international exchange of data and to make it easier for citizens to choose between different providers of digital twins. This goal is addressed to legislators as well as technology companies and researchers, who would have to develop the necessary technologies.

The second objective is to clarify that access to certain health data is needed to enable quality testing of DTS, analysis of the public health impact of this technology, and generally identification of new health challenges. This is already the case today (for example, with mandatory reporting of certain dangerous diseases), but the comprehensive data infrastructure enabled by DTS raises the question of the minimum amount of health data that must be mandatorily collected. Analogous to the Scandinavian model, a democratic deliberative process should be used to determine what this scope is in order to meet the regulatory requirements for ensuring quality and capturing potential negative effects of DTS. Here, it would then have to be precisely defined (in accordance with the principle of proportionality) for which public health purposes an optout possibility is to be provided to citizens (or where consent is even required) and

where data disclosure is mandatory. This goal also requires a certain openness and solidarity of all citizens with regard to the use of their data by the community.

The third objective ensures that quality criteria for DTS exist and are enforced in the healthcare system. To this end, government actors must define benchmarks that DTS must not fall below. Linked to these benchmarks, a model for the approval of DTS must be developed. Security deserves a special focus. Technical standards for cybersecurity in the development (training) and operation of DTS must be developed and adhered to by the providers of this technology. A typology of attack opportunities as well as misuse opportunities must be established for the development of these standards. The development of such standards is a collaborative task between cybersecurity professionals, vendors, and the Federal Cybersecurity Agency. For the latter, the appropriate legal foundations should be created so that it can monitor the implementation of the standards.

Further information: dsi.uzh.ch/strategy-lab