Federal Register Notice: 89 FR 51554, Federal Register :: Networking and Information Technology Research and Development Request for Information on Digital Twins Research and Development, June 18, 2024.

Request for Information on the National Digital Twins R&D Strategic Plan

Nature Computational Science

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.

nature computational science

Perspective

https://doi.org/10.1038/s43588-024-00607-6

Digital twins in medicine

Received: 9 October 2023

Accepted: 12 February 2024

Published online: 26 March 2024

Check for updates

R. Laubenbacher[1] ✉, B. Mehrad[1], I. Shmulevich[2] & N. Trayanova[3]

Medical digital twins, which are potentially vital for personalized medicine, have become a recent focus in medical research. Here we present an overview of the state of the art in medical digital twin development, especially in oncology and cardiology, where it is most advanced. We discuss major challenges, such as data integration and privacy, and provide an outlook on future advancements. Emphasizing the importance of this technology in healthcare, we highlight the potential for substantial improvements in patient-specific treatments and diagnostics.

A 52-year-old man was found confused in his prison cell by the staff and brought to a hospital emergency department (ED). He had a remote history of a stroke and consequent hemiplegia. In the ED, his initial blood pressure was low and responded to volume resuscitation. His laboratory studies showed a mildly elevated peripheral leukocyte count and a mild renal insufficiency, and his chest X-ray showed bilateral airspace disease. He was diagnosed with pneumonia and started on empiric antibiotics aimed at pathogens that cause severe community-acquired pneumonia. Six hours later, he became more confused, developed hypotension again (requiring initiation of intravenous pressor drugs) and required increasing supplemental oxygen. His chest X-ray also showed worsening airspace disease. Three hours later, he required endotracheal intubation and lung-protective mechanical ventilation for worsening hypoxic respiratory failure due to acute respiratory distress syndrome. Over the following day, he developed refractory septic shock, requiring multiple intravenous pressors, acute kidney injury, and escalating ventilator requirements. His blood cultures, obtained on admission, showed *Klebsiella pneumoniae*, an organism that was sensitive to the antibiotics he had received. The patient's condition continued to deteriorate, and he died of multi-organ failure 32 hours after his initial presentation.

In this case (encountered by B.M. in the medical intensive care unit (ICU)) and others like it, existing illness severity scoring systems provide quite accurate data on the likelihood of acute mortality, and are hence helpful in determining, for example, whether the patient can be treated at home, in the hospital or in the intensive care unit. Their main shortcoming, highlighted by this case, is that the information they provide is not otherwise actionable, for example, by identifying interventions that may have altered the course of the illness. This represents the biggest promise and challenge in applying computational modeling at the level of individual patients: given that biological heterogeneity leads to a wide range of responses to illness and treatments, can computational models, together with the right kinds of data, help the medical team intervene with more effective and better-timed interventions, tailored to an individual patient and resulting in better outcomes?

When President Obama announced his precision medicine initiative during the 2015 State of the Union address and established a working group to implement it[1], he called for an unprecedented effort to collect data on a million patients to bring the United States closer to an era where medical treatment can be tailored to each individual patient. The digital twin paradigm, as initially formulated and developed in industry and engineering, has a compelling analog in medicine as a powerful tool for truly personalizing medical care. The first attempts at medical digital twin (MDT) technology were made decades ago, including the very comprehensive pioneering Archimedes project on diabetes[2]. The literature has grown substantially over the past several years. A PubMed search for 'digital twin' retrieves over 1,400 citations, with an exponential increase since 2020. Three developments have probably contributed to this trend: an increased focus on precision (or personalized) medicine, the increasing availability of data characterizing human biology[3,4] and electronic health record (EHR) data for large patient cohorts, as well as the emergence of powerful modeling and simulation capabilities, such as machine learning (ML) algorithms and artificial intelligence (AI) models.

In this Perspective, we review the current state of development of MDTs and the challenges that need to be met to move forward with the development of an MDT industry, similar to what currently exists in the industrial sphere. These challenges are (1) technological, such as appropriate modeling technologies, (2) medical, such as our understanding of the biological determinants of health and disease, as well as the availability of appropriate data-generation technologies and (3) administrative, such as standards for regulatory approval of MDT-based devices and data-sharing standards. We will also highlight developments in two

[1]Department of Medicine, University of Florida, Gainesville, FL, USA. [2]Institute for Systems Biology, Seattle, WA, USA. [3]Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD, USA. ✉e-mail: [redacted]

medical fields in which MDTs have made substantial progress already. We describe some approaches to MDTs in oncology, and we give a detailed description of the use of MDTs in cardiology, in particular heart arrhythmias.

Here, we will focus primarily on mechanistic modeling as the basis for MDTs. A plethora of ML/AI modeling methods have been applied, such as causal AI and physic-informed neural networks. We refer the reader to two review articles that focus on the integration of mechanistic modeling and machine learning techniques[5,6], which we believe will be one of the most promising approaches to MDT technology.

Industrial and medical digital twins

Industrial digital twins are characterized by two features: (1) they are built on a mechanistic model of the physical system to be twinned and (2) they are dynamically calibrated to the system for the purpose of forecasting system performance and identifying interventions, such as preventive maintenance. In other words, the digital twin evolves with the physical system over time. This definition conforms to the vision of personalized medicine, either curative or preventive. However, there are only a few tools currently in use in medicine that completely fit this definition, and we give some examples later. Although the industrial paradigm represents the gold standard for personalized medicine, digital twins that fall short of this can still serve as valuable tools that improve on the standard of care in many cases.

There are three main challenges that distinguish medicine from industry when it comes to digital twins. First, for many medical applications, the relevant underlying biology is partially or completely unknown. For example, it is known that some diseases have an important microbiome component. However, in most cases, little is known about the mechanisms involved. Having said that, microbiome data are easy to collect and are abundant, so there are opportunities to apply data-driven approaches to patient stratification and potential actionable insights for targeting treatments to patient subpopulations identified through ML algorithms. Second, even if there are mechanistic dynamic models of the requisite human biology available that could be personalized, the needed data are often not available or are difficult to collect. Third, human biology is often not easily describable with deterministic models based on physical principles, such as systems of ordinary differential equations based on physical laws. In these cases, other modeling platforms need to be used, such as agent-based models. The resulting computational models can be multi-scale, hybrid and stochastic. However, the theoretical and computational infrastructure to analyze and control such models is not yet developed to a degree that is needed for medical applications.

It is worth mentioning that there is no broad consensus as to what constitutes a digital twin in medicine[7,8]. Candidates range from simply a computational model relevant to disease to a full digital replica of all or part of a patient that is continually or periodically updated with patient-derived measurements (this is the full analog of an industrial digital twin). For different applications, all of these can be effective. The comprehensive report *Foundational Research Gaps and Future Directions for Digital Twins* by the National Academies[9] proposes, as a general definition, a computational model of the system to be twinned (in our case all or part of a human patient) that is connected to the system in a bidirectional fashion over time, periodically recalibrated with patient data, and provides patient predictions over time. This definition is closely related to the one widely used in industry.

Applications for medical digital twins

There are three main types of application for MDTs (hypothetical scenarios are depicted in Fig. 1): keeping healthy patients healthy, restoring health in ill patients, and developing novel therapeutics, such as drugs and medical devices. We describe each of these applications later in this section.

It is worth noting that, in addition to these, there are also other important applications not discussed in depth here. For example, MDTs could be used to reduce the use of animals in drug and product development, a priority of the US Food and Drug Administration (FDA)[9]. In addition, MDTs could be used to address health inequalities. One of the important sources of inequality in healthcare arises from clinical trials that are not representative of certain parts of the population, such as women or people of color, or different geographic regions of the globe. Another source of health inequities is the scant attention paid to rare diseases, for which it is often a challenge to recruit large enough cohorts for clinical trials with sufficient statistical power. Computational models are used now in different contexts to create virtual patient cohorts or enlarge existing cohorts through virtual patients. If, for a given trial, a digital twin is available that can be customized to specific patient groups, then clinical trials could be run with more representative patient cohorts. Finally, a further potential impact of MDT technology is on the reduction of health disparities. Incorporating a model-driven decision support system into treatment decisions can help alleviate the healthcare disparities that patients face, where social sources of bias (race or ethnicity, sex and sexuality, body weight, socioeconomic status and so forth) can influence medical decisions.

Keeping healthy patients healthy

MDTs can be an important tool as we transition from curative to preventive medicine. Risk score calculators have been in use for some time, and they might use genetic data, data collected from wearables, such as heart rate and rhythm or sleep patterns, or exercise patterns from fitness trackers. A major obstacle is our lack of knowledge about how to define health in the presence of the great biological variability between patients and, consequently, our inability to build predictive models that can be used for this purpose. At the same time, this application of the digital twin concept is the most impactful one in the long run and comes closest to a major use of digital twins in industry, namely preventive maintenance. It is important to note that, for industrial applications, the use of AI/ML techniques face several challenges[10,11]. Mechanistic models are generally preferred because they provide the means to forecast the effect of specific interventions and can be used to identify optimal control interventions.

Restoring health in ill patients

The most progress in MDT technology has been made for the purpose of treating patients with a health condition. A successful example of this is the development of automated subcutaneous insulin delivery for patients with type I diabetes. There are now several US FDA-approved devices on the market for this purpose. One of these[12] has been approved for all age ranges, most recently for young children[13]. It is based on an ordinary differential equations model of human glucose metabolism, coupled with a closed-loop control algorithm. The model receives near-streaming glucose measurements from a subcutaneous sensor in the patient and calculates the appropriate amount of insulin required, and the control algorithm drives an insulin pump that automatically injects it under the patient's skin. The model is recalibrated to the patient every few minutes. Currently, the algorithm still requires some input from the patient about food intake and activity level.

Another application area where MDTs hold considerable promise is in critical care, such as ICUs. In a fast-paced environment where healthcare personnel are typically confronted with a continuous stream of large volumes of data, MDTs can be valuable as decision support tools or data integration devices. At the same time, an ICU is a comparatively data-rich environment where patient measurements are collected routinely. Many of the MDTs in this field are blackbox ML models. Some of these are scoring systems[14] that provide risk scores as output, such as mortality risk for a given patient as described above. An example of an alternate approach that is more likely to produce

https://doi.org/10.1038/s43588-024-00607-6

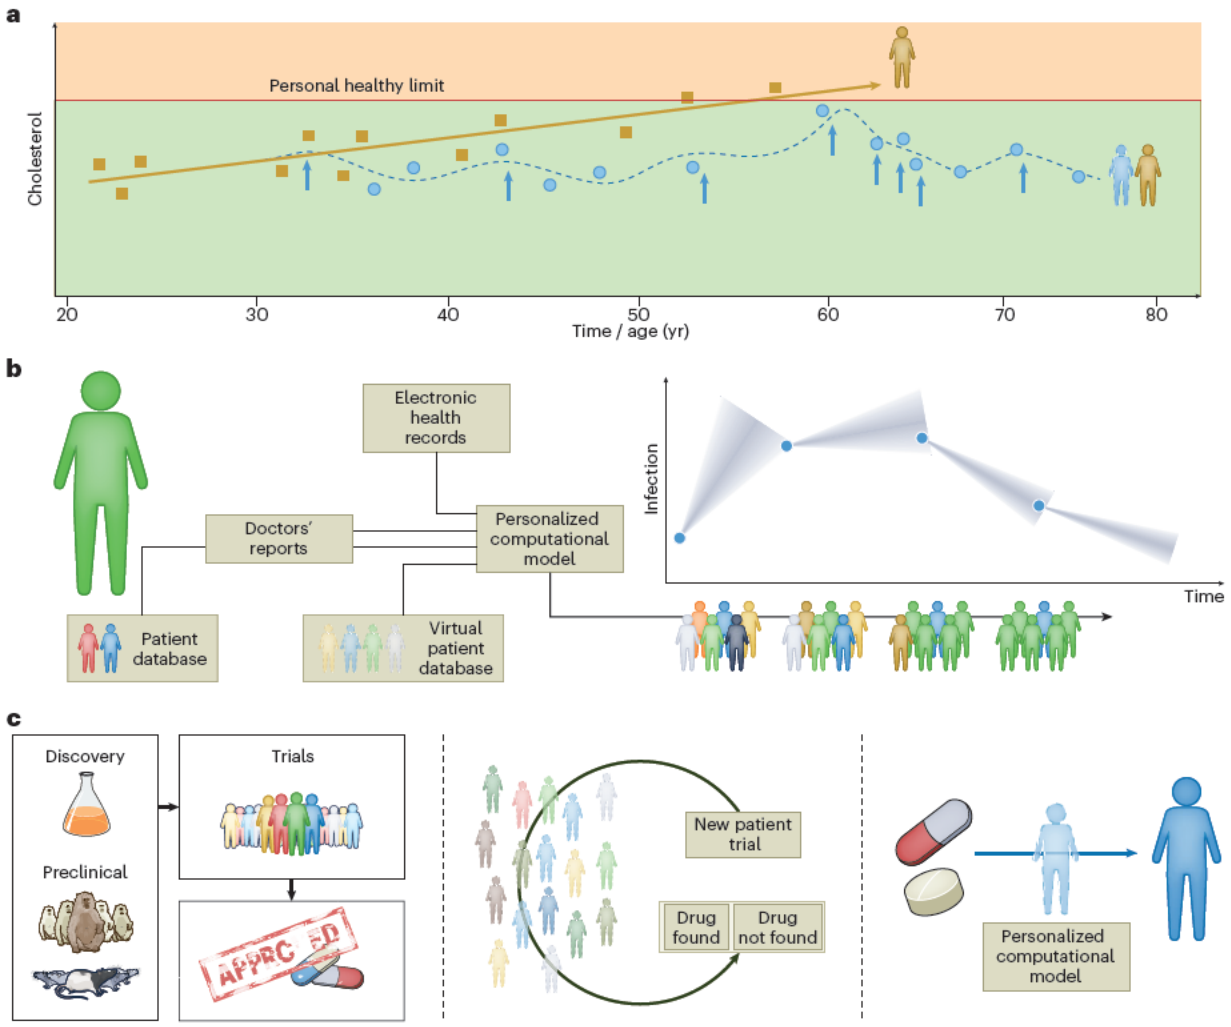

Fig. 1 | Applications for medical digital twins. a, Keeping healthy patients healthy. For a given patient (in yellow), a safe cholesterol level is determined, using genetic information, family history and other data. The yellow line indicates the trend of the patient's cholesterol levels over time, if untreated. Yellow boxes represent measurements. The patient's digital twin (blue), on the other hand, forecasts the trajectory and recommends periodic preventive interventions (blue arrows), resulting in cholesterol levels following the blue curve. **b**, Restoring health in ill patients. Upon admission to the ICU, the patient (green) is evaluated and receives initial treatment for an infection. A computer algorithm personalizes an appropriate computational disease model, together with information from a database of reference patients to recommend optimal interventions. As more repeated measurements are taken from the patient, the reference population is refined, the model is recalibrated to the patient at later time points, and the recommendations for optimal treatment are refined. The cones represent the likely trajectory of the infection as determined by the digital twin. With time and a larger number of patient data points, the uncertainty in the predictions decreases (the cone becomes narrower) and subsequent patient time points fall closer to the center of the previous prediction cone. The improvement in the parameter ensemble that describes the patient is reflected by the corresponding virtual cohort that describes the patient at each time point, which is depicted as increasingly containing more green subjects, like the patient being treated. **c**, Development of novel therapeutics. Currently, clinical trials typically involve the use of animals and patient cohorts (left panel). With the advent of MDTs, it will be possible to reduce the number of animals used in preclinical trials and to optimize patient trials using virtual patients. They can be used to screen large numbers of drug targets and drug candidates, and to perform initial optimization studies using large numbers of patient MDTs and virtual patients (middle panel). Optimal drug regimes, doses and combinations can also be inferred by MDT before administering drugs to patients, thus minimizing side effects (right panel).

such actionable information combines a causal AI model with expert rules for clinical decision making, in which a digital twin is introduced for acute stroke care[15]. The basic structure is a directed acyclic graph, built from expert-curated statements, such as 'ischemic stroke leads to cerebral edema' or 'timely administration of thrombolytics can lead to improvement of outcomes in ischemic stroke', together with likelihood scores obtained from consensus levels among a group of experts. The graph with associated probabilities is fashioned into a Bayesian network that is then further trained with patient data. Once sufficiently trained and validated, the model can be used for decision making for individual patients. A similar digital twin has been developed for the critical care of sepsis patients[16].

As mentioned earlier, one of the biggest challenges for MDTs is accounting for the biological heterogeneity of patients, which manifests itself in a highly variable response to therapeutic interventions. Higher-resolution models that take account of a patient's disease-relevant biology could help with more accurate predictions. A possible roadmap for such an approach in the case of sepsis[17,18] is to use a detailed mechanistic model of systemic inflammation and then create variability in the parametrization of this model through ML algorithms trained on large sets of patient data; a particular patient can then be matched to an appropriate parametrized model that can be used for forecasting of intervention outcomes and choice of optimal ones. Another early-stage project to develop a digital twin for pneumonia

also uses a mechanistic model of the early immune response to respiratory infections[19].

Developing novel therapeutics

Virtual clinical trials are a third application area for MDT technology[20–22]. The basic approach here is to begin with a computational model—often mechanistic, in the form of a system of ordinary differential equations or an agent-based model—that captures the human biology relevant to the compound or intervention to be tested. In the case of a mechanistic model, the model parameters will typically have biological meaning, such as immune cell counts, or relate to a patient's physiology, such as heart rate or glucose levels (as opposed to parameters in many ML models). Based on expert knowledge or published information, one can determine physiologically reasonable ranges for these parameters and create a virtual patient population by sampling the parameter space of the model within these intervals. Available clinical data might provide information about the distributions of parameter values across these ranges. Each specific parametrization represents an individual virtual patient. If there is already an existing patient population that one desires to enlarge through virtual patients, then one can determine ranges for the model parameters based on measurements from this existing population. This could be useful, for instance, for drug development focused on relatively rare conditions, where recruitment of a sufficiently large patient cohort is difficult (as described before). Finally, one might aim to create a digital replica of a real patient population by creating a digital twin of each individual patient and assembling them to an exact digital replica of the patient population. This would be required to create a digital twin that fits the industrial paradigm. Currently, we do not know of any published instantiations of this last approach.

In the following, we provide a more detailed discussion of two domain areas for which MDT technology development has been very encouraging: oncology and cardiology.

Medical digital twins in oncology

Digital twins for patients with cancer are emerging as a transformative tool in oncology, enabling a highly personalized and dynamic approach to cancer treatment and research. These digital replicas facilitate a comprehensive understanding of individual cancer cases, allowing for the simulation, analysis and prediction of cancer progression and treatment outcomes in a virtual environment. This technology is poised to make substantial contributions to clinical practice by enhancing the precision and effectiveness of cancer care[23,24]. The development of such a digital twin involves the meticulous integration of diverse patient data, such as genetic information about the patient and the tumor, clinical history and detailed imaging data. This rich dataset forms the foundation for the digital twin, enabling the simulation of tumor behavior and the assessment of potential treatment strategies.

Advanced ML algorithms and computational modeling techniques, such as multi-scale models that span molecular, multicellular and organismal scales, are integral to this process. These modeling techniques may include systems of ordinary differential equations or agent-based models, as well as other dynamical systems models. The latter are crucial for modeling molecular interactions within cancer cells that ultimately determine cellular phenotypes. Moreover, ML algorithms contribute by identifying patterns and correlations in large datasets, helping to predict tumor behavior and response to treatments with greater accuracy and efficiency.

In addition to enhancing treatment planning and assessment, digital twins for patients with cancer can play an important role in monitoring cancer progression and evaluating treatment responses. This continuous assessment allows healthcare professionals to make real-time, data-driven adjustments to the treatment plan, ensuring the delivery of the most effective and personalized care. This adaptability is pivotal in enhancing the likelihood of positive treatment outcomes and improving the overall quality of life for patients.

The multifaceted nature of these digital twins also enables the bridging of various biological scales, from molecular changes to physiological responses. Cancer digital twins incorporate data about the patient's pre-existing health, cancer type, size and location of tumors, their metabolic activity, and molecular markers expressed by the tumor. The model will then learn and adapt to the evolving patient data (for example, timing and type of chemotherapy, effect on tumor size, development of adverse effects, occurrence of metastases), ensuring that the models remain up to date and reflective of the patient's current condition. This periodic updating enhances the predictive accuracy of the digital twins, ensuring that healthcare professionals have access to the most relevant and current information for making clinical decisions. The integration of real-time data enhances the responsiveness and adaptability of cancer digital twins, ensuring that they remain a reliable and effective tool for guiding cancer treatment and management.

The application of digital twins also holds promise for addressing issues of equity in cancer care. By providing a platform for the exploration and assessment of diverse treatment options, digital twins enable the delivery of personalized and effective cancer care to a broader patient population. This inclusivity ensures that individuals from various demographic backgrounds have equitable access to advanced and innovative cancer treatments, promoting fairness and equality in healthcare delivery.

There are now several MDT projects under way that can realize this promise and become effective tools in the clinic. For example, Wu and colleagues have developed an MDT project[25] that is designed to predict the progression of breast tumors using a partial differential equations model of breast tissue, calibrated with patient-specific images from both magnetic resonance imaging (MRI) and quantitative positron emission tomography. The images are used to derive model parameters that capture the cell-migration and tumor-cell proliferation properties of the specific tumor. The model can be used to forecast the effect of drug treatments or predict the efficacy of immunotherapy[26].

A similar approach has proven effective when using a digital twin for patients with glioblastoma, a highly aggressive type of brain tumor with poor prognosis, which is part of the growing field of mathematical neuro-oncology[27]. This MDT, developed by Swanson and colleagues, also uses MRI images, in this case of the brains of glioblastoma patients[28]. A time course of such images allows the estimation of two parameters in the partial differential equation, one that captures the proliferation rate of the tumor and another that captures the rate of cell migration. These two parameters are independent of each other, as it has been observed that tumor cells either migrate through tissue or divide, but do not do both at the same time. With these patient-specific parameters, one can predict the actual extent of the tumor, informing surgery planning. The reason this cannot be done accurately simply from the MRI image is that tumor cells invade adjacent brain matter in a diffuse fashion that is not captured accurately by imaging, so the actual dimensions of the part of the brain containing tumor cells cannot be determined from MRI images alone. Modified versions of the digital twin can also predict the effect of different treatments[29].

Although there has been much progress in this field, there are still big challenges ahead, including a lack of sufficient patient data and a lack of sufficient mechanistic understanding underlying the many subtypes of different cancers and of the effectiveness of an ever-growing supply of cancer drugs, impeding model-based prediction of appropriate interventions at a given stage of the disease.

Medical digital twins in cardiology

MDTs are transforming the field of cardiology by providing tools for patient care, treatment planning and healthcare delivery. These virtual representations of the heart and its functions hold immense potential to improve the diagnosis, management and outcomes of cardiovascular

diseases. A cardiac digital twin is typically composed of three main components: (1) data acquisition (imaging, EHR, genetic data and wearables), (2) modeling and simulation (based on anatomy and physiology) and (3) clinical decision making. Here, we showcase the pivotal role that heart MDTs can play in decision support and patient care, and we focus on one of the most important aspects of cardiology: arrhythmia care management. We present two of the clinical applications of heart MDTs: the prediction of sudden cardiac death due to arrhythmias in various diseases and the use of MDTs to provide guidance in arrhythmia treatment by catheter ablation.

The incidence of sudden cardiac death due to arrhythmias is increasing globally, and accurate individualized risk assessment of death remains a major unmet clinical need. Heart MDTs have made major strides in predicting a patient's risk of sudden death, for patients with ischemic (that is, caused by coronary atherosclerosis) and non-ischemic cardiomyopathies. The study by Arevalo and colleagues[30] demonstrated the first utilization of MDTs created from contrast-enhanced MRIs of a cohort of patients ($n = 41$) after myocardial infarction (scarring in the heart) to determine the patients' likelihood of developing infarct-related ventricular arrhythmias and sudden death. The MDT prediction outperformed all current clinical risk assessment metrics, indicating that MDTs can be used to determine the need for a prophylactic implantation of defibrillator devices to prevent sudden death. A more complex approach to the assessment of the arrhythmia propensity of patients with previous infarcts using MDTs involves the additional incorporation of penetrating adipose tissue (fat)[31].

In relation to predicting the risk of sudden cardiac death in patients with non-ischemic cardiomyopathies, heart MDT studies have demonstrated the clinical utility of the approach in pediatric patients, such as those with acute myocarditis[32] and with repaired tetralogy of Fallot[33]. MDT technology has also been used[34] to predict arrhythmia risk in hypertrophic cardiomyopathy, a common genetic disease characterized by a thickening of heart muscle, substantially outperforming current clinical risk predictors. Another non-ischemic cardiomyopathy associated with high risk of sudden death and difficult risk prediction is cardiac sarcoidosis, an inflammatory disease. Shade and colleagues have developed a two-step prediction approach, combining MDT with ML in a study of 45 patients[35]. The results from MDT simulation were fed, together with a set of clinical biomarkers, into a supervised classifier. Finally, a genotype-specific heart MDT (Geno-DT) approach was recently developed to predict the arrhythmia circuits in patients with arrhythmogenic right ventricular cardiomyopathy (ARVC) of different genotypes[36]. This approach revealed that the underlying arrhythmia mechanisms differ among ARVC genotypes. The Geno-DT approach demonstrated the potential to augment therapeutic precision in the clinical setting, which can lead to more personalized treatment strategies in ARVC.

Catheter ablation plays a major role in the contemporary management of arrhythmias. This procedure involves the use of catheters that are maneuvered into the cardiac chambers and deliver radiofrequency energy to specific locations to terminate the perpetrator of arrhythmia. Identification of these specific locations in the heart is, however, difficult, and ablation targets are often inaccurate, with new (emergent) arrhythmias occurring post-ablation and necessitating repeat procedures and re-hospitalization. Personalized MDT technology has made major strides in improving ablation precision by providing non-invasive localization of ablation targets. Following a validation study[37], MDTs were used to predict the ablation targets and guide the ablation in post-infarction patients[38,39]. This work highlighted the enormous potential for MDT technology to impact the clinical management of ventricular arrhythmias. The MDTs predicted not only the targets for initial ablation, but also the ablation targets for re-do procedures several years later.

Another exciting aspect of personalized MDTs is the ability to plan different atrial fibrillation management strategies, and even predict a patient's risk of recurrence. Atrial fibrillation occurs in the upper chamber of the heart and is the most common human arrhythmia, affecting 1–2% of the population. Although not as dangerous as ventricular arrhythmias, it is associated with a high probability of stroke and a high burden of healthcare expenditures due predominantly to patient re-hospitalization. Several atrial MDT studies have tested the effectiveness of different ablation strategies in patients with a persistent form of the arrhythmia. The discovery of atrial fibrosis as a substrate for atrial arrhythmias resulted in the development of atrial MDTs reflecting the patient-specific atrial fibrosis distribution[40–44]. Boyle and colleagues pioneered a prospective ablation study for patients with persistent atrial fibrillation and fibrosis entirely with personalized atrial MDTs[45]. In that study, the MDT-proposed ablation targets were used to steer patient treatment. Finally, atrial MDTs have been used, often in combination with ML or other technologies, to predict atrial fibrillation recurrence[46,47].

The initial successes with heart MDTs constructed from imaging and other health data described above have opened new pathways for the development and application of MDTs in cardiology. Of particular importance will be the ability to incorporate continuous data from various streams, thus ensuring that the patients' MDT continuously reflects the state of the patient's heart.

Opportunities and challenges

MDT research has largely been scattered across individual laboratories and companies, often without explicitly using the MDT label. In recent years, there has been an emphasis on research funding for MDT projects in Europe, through the Horizon Europe grant program of the European Commission. Possibly the most ambitious project is the Ecosystem Digital Twins in Health (EDITH)[48], a comprehensive initiative to develop a roadmap for digital twin technology in healthcare in Europe, funded by the European Commission. The Virtual Physiological Human (VPH) is a European initiative to lay the groundwork for a collaborative framework to investigate the human body as a complex system[49]. Other examples of a large-scale MDT project include the Swedish Digital Twin Consortium[50,51], aiming to create MDTs for the entire Swedish population. In the United States, the National Institutes of Health and the US Department of Energy have partnered to support the development of digital twins for cancer patients[24]. These and other community- and agency-driven MDT efforts provide a wide range of opportunities for research collaborations, funding and community infrastructure.

One application area that is particularly promising and urgently needed is the response to infectious diseases. The SARS-CoV-2 pandemic was marked by the major challenge of the highly heterogeneous response to infection and treatment. The availability of MDTs that capture certain features of the immune system, even in rudimentary ways, could have provided additional decision support for healthcare providers[52]. Subsequent efforts to advance MDT technology focusing on the immune system were developed[53–55]. These and other efforts have begun to catalyze a research community focused on the immune system, as a major contributor to infectious-disease outcomes, as well as other diseases such as cancer, autoimmune diseases and diabetes.

As mentioned previously, the main three factors limiting MDT development are our incomplete knowledge of human biology, the availability of patient data in sufficient quantity and quality (and at all scales) and the lack of a well-developed modeling technology that can form the basis of an MDT industry comparable to that existing for industrial applications. Furthermore, to bring digital twin technology into the clinic, we need to solve a range of problems related to patient privacy, security, ethics, standards for models and data, and regulatory requirements. For the latter, we refer the reader to the strategic plan[56] released by the EDITH project. This plan addresses a comprehensive range of regulatory and business aspects for developing and implementing digital twin technology in healthcare at scale.

In particular, the plan addresses issues of data collection, privacy concerns and ethics guidelines. A plethora of different data types and sources are potentially valuable for MDT applications, collected from wearables and other types of mobile sensor, patient charts, a wide range of imaging data, as well as data collected from samples, such as blood or tissue samples from biopsies. Some of these are subject to patient protection regulations under the HIPAA law, whereas others are unprotected, such as data from fitness trackers or genetic sequence information generated by private companies. A particular concern, as with other personal data, is that an MDT is a vehicle for systematic data integration—one of its strengths as a medical tool—but potentially damaging to patient privacy. Issues such as who controls the MDT of a patient, who it belongs to, and what can be done with it, are not currently settled and will likely require regulatory actions.

Another challenge that must be addressed in the future is related to the models underlying MDTs, which in many cases will be multi-scale. Most drugs, for example, have mechanisms of action at the intracellular scale, but have organ- or organism-level effects. Models will probably be hybrid, combining for instance blood flow through an artery, modeled by a partial differential equation, with intracellular signaling in the endothelial cells lining the artery, modeled by a system of ordinary differential equations or a Boolean network. The models will probably be stochastic too, reflecting, for example, features of the immune system or gradient-based movement of cells in a tissue. These characteristics pose challenges to most of the established model building and analysis tools available, as well as mathematical control approaches. There are no formal methods available for this type of model, and tools that are standard for differential equations models, such as global sensitivity analysis, identifiability, forecasting and optimal control, are not directly applicable. New approaches to model validation are required too. Furthermore, MDT models will need to be updated and expanded over time, as our knowledge of biology or the application type changes. Standard model implementation methods do not result in models that are robust with respect to these operations. To provide the basis for a large-scale standardized MDT industry, extensive research in this area is required.

The future

Over the past two decades, digital twin technology has evolved to become an increasingly mature and rapidly growing industry, projected to reach US$183 billion by 2031[57]. Fitzgerald and colleagues[58] provide a discussion of the most important opportunities and challenges for the further development of digital twin technology for industrial applications, including the following central research questions:

1. What are the specifications that are necessary for a dependable digital twin?
2. What are the key specifications for usability and credibility that are required for a digital twin?
3. How accurate does a digital twin have to be to be useful?
4. What benefits does a digital twin have to provide to justify constructing one?

These same open questions need to be answered for MDTs. Industrial digital twin technology is benefiting from a more highly developed infrastructure, including standards for computational model specification, standard operating procedures and physics-based models for many of the systems to be twinned, to name a few. In biomedicine, in contrast, existing MDTs and MDT projects are still early in their development, without broadly available infrastructure, standards and templates for the successful development of commercial products. Successful strategies for commercialization remain largely unexplored, and the regulatory requirements and hurdles are formidable.

A special challenge MDT technology faces is a very complex regulatory environment for the use of computational models in medicine. These challenges need to be addressed before the technology can be broadly adopted. The FDA has encouraged the use of modeling and simulation in the development and approval process of drugs and medical devices, and has issued guidance on this topic[59,60]. A comprehensive set of standards in relation to the credibility of computational modeling for medical devices is also available[61], as well as a summary of all regulatory efforts and guidelines[62]. Unfortunately, there are currently no standards available for models that are not physics-based and use modeling platforms other than systems of differential equations. As mentioned earlier, in many cases, MDTs will have to rely on other model types, and additional standards and requirements will need to be developed for approval.

One of the biggest challenges to precision medicine and the use of MDTs is the biological heterogeneity of patients. To account for this, we will probably need to develop higher-resolution models than the ones that are currently available, given our knowledge of human biology and the data to capture it. By analogy, the accuracy of numerical weather prediction models over longer forecasting windows provides a good paradigm. Accuracy has increased dramatically over the past two decades, largely due to the higher-resolution models that have been made possible by higher-resolution data (and more powerful computation resources). This reflects the fact that both humans and the atmosphere are truly complex systems in which microscopic perturbations can result in macroscopic effects, the proverbial butterfly effect. Once we make progress on these challenges, medical digital twins will change healthcare fundamentally, helping us to transition from curative to preventive medicine.

References

1. *The Precision Medicine Initiative Cohort Program—Building a Research Foundation for 21st Century Medicine* (NIH, 2015); https://acd.od.nih.gov/documents/reports/PMI_WG_report_2015-09-17-Final.pdf
2. Eddy, D. M. & Schlessinger, L. Archimedes: a trial-validated model of diabetes. *Diabetes Care* **26**, 3093–3101 (2003).
3. Tomczak, K., Czerwińska, P. & Wiznerowicz, M. The Cancer Genome Atlas (TCGA): an immeasurable source of knowledge. *Contemp. Oncol. Pozn. Pol.* **19**, A68–A77 (2015).
4. Regev, A. et al. The Human Cell Atlas. *eLife* **6**, e27041 (2017).
5. Alber, M. et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical and behavioral sciences. *NPJ Digit. Med.* **2**, 115 (2019).
6. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
7. *Opportunities and Challenges for Digital Twins in Biomedical Research* (National Academies of Science, Engineering and Medicine, 2023); https://nap.nationalacademies.org/catalog/26922/opportunities-and-challenges-for-digital-twins-in-biomedical-research-proceedings
8. *Opportunities and Challenges for Digital Twins in Atmospheric and Climate Sciences: Proceedings of a Workshop—In Brief* 26921 (National Academies Press, 2023); https://doi.org/10.17226/26921
9. *Foundational Research Gaps and Future Directions for Digital Twins* (National Academies of Sciences, Engineering and Medicine, 2023); https://doi.org/10.17226/26894
10. Wright, L. & Davidson, S. How to tell the difference between a model and a digital twin. *Adv. Model. Simul. Eng. Sci.* **7**, 13 (2020).
11. Vogelsang, A. & Borg, M. Requirements engineering for machine learning: perspectives from data scientists. In *Proc. 2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)* (eds Damian, D., Perini, A. & Lee, S.-W.) 245–251 (IEEE, 2019); https://doi.org/10.1109/REW.2019.00050
12. Cobelli, C. & Kovatchev, B. Developing the UVA/Padova type 1 diabetes simulator: modeling, validation, refinements and utility. *J. Diabetes Sci. Technol*. https://doi.org/10.1177/19322968231195081 (2023).

13. Breton, M. D. et al. A randomized trial of closed-loop control in children with Type 1 diabetes. *N. Engl. J. Med.* **383**, 836–845 (2020).
14. Quintairos, A., Pilcher, D. & Salluh, J. I. F. ICU scoring systems. *Intensive Care Med.* **49**, 223–225 (2023).
15. Dang, J. et al. Developing DELPHI expert consensus rules for a digital twin model of acute stroke care in the neuro critical care unit. *BMC Neurol.* **23**, 161 (2023).
16. Lal, A. et al. Development and verification of a digital twin patient model to predict specific treatment response during the first 24 hours of sepsis. *Crit. Care Explor.* **2**, e0249 (2020).
17. Cockrell, C. & An, G. Sepsis reconsidered: identifying novel metrics for behavioral landscape characterization with a high-performance computing implementation of an agent-based model. *J. Theor. Biol.* **430**, 157–168 (2017).
18. Larie, D., An, G. & Cockrell, R. C. The use of artificial neural networks to forecast the behavior of agent-based models of pathophysiology: an example utilizing an agent-based model of sepsis. *Front. Physiol.* **12**, 716434 (2021).
19. Ribeiro, H. A. et al. Multi-scale mechanistic modelling of the host defence in invasive aspergillosis reveals leucocyte activation and iron acquisition as drivers of infection outcome. *J. R. Soc. Interface* **19**, 20210806 (2022).
20. Chasseloup, E., Hooker, A. C. & Karlsson, M. O. Generation and application of avatars in pharmacometric modelling. *J. Pharmacokinet. Pharmacodyn.* **50**, 411–423 (2023).
21. Moingeon, P., Chenel, M., Rousseau, C., Voisin, E. & Guedj, M. Virtual patients, digital twins and causal disease models: paving the ground for in silico clinical trials. *Drug Discov. Today* **28**, 103605 (2023).
22. Allen, R., Rieger, T. & Musante, C. Efficient generation and selection of virtual populations in quantitative systems pharmacology models. *CPT Pharmacomet. Syst. Pharmacol.* **5**, 140–146 (2016).
23. Hernandez-Boussard, T. et al. Digital twins for predictive oncology will be a paradigm shift for precision cancer care. *Nat. Med.* **27**, 2065–2066 (2021).
24. Stahlberg, E. A. et al. Exploring approaches for predictive cancer patient digital twins: opportunities for collaboration and innovation. *Front. Digit. Health* **4**, 1007784 (2022).
25. Wu, C. et al. Integrating mechanism-based modeling with biomedical imaging to build practical digital twins for clinical oncology. *Biophys. Rev.* **3**, 021304 (2022).
26. Jarrett, A. M. et al. Optimal control theory for personalized therapeutic regimens in oncology: background, history, challenges and opportunities. *J. Clin. Med.* **9**, 1314 (2020).
27. Baldock, A. L. et al. From patient-specific mathematical neuro-oncology to precision medicine. *Front. Oncol.* **3**, 62 (2013).
28. Jackson, P. R., Juliano, J., Hawkins-Daarud, A., Rockne, R. C. & Swanson, K. R. Patient-specific mathematical neuro-oncology: using a simple proliferation and invasion tumor model to inform clinical practice. *Bull. Math. Biol.* **77**, 846–856 (2015).
29. Hawkins-Daarud, A. et al. In silico analysis suggests differential response to bevacizumab and radiation combination therapy in newly diagnosed glioblastoma. *J. R. Soc. Interface* **12**, 20150388 (2015).
30. Arevalo, H. J. et al. Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models. *Nat. Commun.* **7**, 11437 (2016).
31. Sung, E. et al. Fat infiltration in the infarcted heart as a paradigm for ventricular arrhythmias. *Nat. Cardiovasc. Res.* **1**, 933–945 (2022).
32. Cartoski, M. J. et al. Computational identification of ventricular arrhythmia risk in pediatric myocarditis. *Pediatr. Cardiol.* **40**, 857–864 (2019).
33. Shade, J. K. et al. Ventricular arrhythmia risk prediction in repaired tetralogy of Fallot using personalized computational cardiac models. *Heart Rhythm* **17**, 408–414 (2020).
34. O'Hara, R. P. et al. Personalized computational heart models with T1-mapped fibrotic remodeling predict sudden death risk in patients with hypertrophic cardiomyopathy. *eLife* **11**, e73325 (2022).
35. Shade, J. K. et al. Predicting risk of sudden cardiac death in patients with cardiac sarcoidosis using multimodality imaging and personalized heart modeling in a multivariable classifier. *Sci. Adv.* **7**, eabi8020 (2021).
36. Zhang, Y. et al. Predicting ventricular tachycardia circuits in patients with arrhythmogenic right ventricular cardiomyopathy using genotype-specific heart digital twins. *eLife* 10.7554/eLife.88865.2 (2023).
37. Ashikaga, H. et al. Feasibility of image-based simulation to estimate ablation target in human ventricular arrhythmia. *Heart Rhythm* **10**, 1109–1116 (2013).
38. Prakosa, A. et al. Personalized virtual-heart technology for guiding the ablation of infarct-related ventricular tachycardia. *Nat. Biomed. Eng.* **2**, 732–740 (2018).
39. Sung, E. et al. Personalized digital-heart technology for ventricular tachycardia ablation targeting in hearts with infiltrating adiposity. *Circ. Arrhythm. Electrophysiol.* **13**, e008912 (2020).
40. McDowell, K. S. et al. Virtual electrophysiological study of atrial fibrillation in fibrotic remodeling. *PLoS ONE* **10**, e0117110 (2015).
41. Roney, C. H. et al. In silico comparison of left atrial ablation techniques that target the anatomical, structural and electrical substrates of atrial fibrillation. *Front. Physiol.* **11**, 1145 (2020).
42. Zahid, S. et al. Patient-derived models link re-entrant driver localization in atrial fibrillation to fibrosis spatial pattern. *Cardiovasc. Res.* **110**, 443–454 (2016).
43. Loewe, A. et al. Patient-specific identification of atrial flutter vulnerability—a computational approach to reveal latent reentry pathways. *Front. Physiol.* **9**, 1910 (2019).
44. Roney, C. H. et al. Predicting atrial fibrillation recurrence by combining population data and virtual cohorts of patient-specific left atrial models. *Circ. Arrhythm. Electrophysiol.* **15**, e010253 (2022).
45. Boyle, P. M. et al. Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nat. Biomed. Eng.* **3**, 870–879 (2019).
46. Ali, R. L. et al. Arrhythmogenic propensity of the fibrotic substrate after atrial fibrillation ablation: a longitudinal study using magnetic resonance imaging-based atrial models. *Cardiovasc. Res.* **115**, 1757–1765 (2019).
47. Shade, J. K. et al. Preprocedure application of machine learning and mechanistic simulations predicts likelihood of paroxysmal atrial fibrillation recurrence following pulmonary vein isolation. *Circ. Arrhythm. Electrophysiol.* **13**, e008213 (2020).
48. *EDITH: European Virtual Human Twin* (Virtual Physiological Human Institute); https://www.edith-csa.eu/
49. Viceconti, M. & Hunter, P. The virtual physiological human: ten years after. *Annu. Rev. Biomed. Eng.* **18**, 103–123 (2016).
50. Swedish Digital Twin Consortium (SDTC); https://www.sdtc.se/
51. Björnsson, B. et al. Digital twins to personalize medicine. *Genome Med.* **12**, 4 (2019).
52. Laubenbacher, R., Sluka, J. P. & Glazier, J. A. Using digital twins in viral infection. *Science* **371**, 1105–1106 (2021).
53. Laubenbacher, R. et al. Building digital twins of the human immune system: toward a roadmap. *NPJ Digit. Med.* **5**, 64 (2022).
54. *Forum on Precision Immunology: Immune Digital Twins* (UF Laboratory for Systems Medicine); https://systemsmedicine.pulmonary.medicine.ufl.edu/working-groups/forum-on-precision-immunology-immune-digital-twins/

55. *Building Immune Digital Twins* (Institut Pascal); https://www.institut-pascal.universite-paris-saclay.fr/en/scientific-programs/building-immune-digital-twins
56. *EDITH CSA Deliverable 3.2: First Draft of the VHT Roadmap* (EDITH Consortium, 2023); https://doi.org/10.5281/ZENODO.8200955
57. *Gartner 2018 Hype Cycle for IT in GCC Identifies Six Technologies That Will Reach Mainstream Adoption in Five to 10 Years* (Gartner, 2018); https://www.gartner.com/en/newsroom/press-releases/2018-12-13-gartner-2018-hype-cycle-for-it-in-gcc-identifies-six-technologies-that-will-reach-mainstream-adoption-in-five-to-10-years
58. Fitzgerald, J., Larsen, P. G., Margaria, T., Woodcock, J. & Gomes, C. in *Leveraging Applications of Formal Methods, Verification and Validation, Lecture Notes in Computer Science* (eds Margaria, T. & Steffen, B.) 13704 (Springer, 2022); https://doi.org/10.1007/978-3-031-19762-8
59. *Assessing the Credibility of Computational Modeling and Simulation in Medical Device Submissions; Draft Guidance for Industry and Food and Drug Administration Staff* (FDA, 2021); https://www.fda.gov/media/154985/download
60. Nuwer, R. US agency seeks to phase out animal testing. *Nature* https://doi.org/10.1038/d41586-022-03569-9 (2022).
61. *Assessing Credibility of Computational Modeling through Verification and Validation*: *Application to Medical Devices* (The American Society of Mechanical Engineers, 2018).
62. Ahmed, K. B. R., Pathmanathan, P., Kabadi, S. V., Drgon, T. & Morrison, T. M. Editorial on the FDA report on 'Successes and opportunities in modeling & simulation for FDA'. *Ann. Biomed. Eng.* **51**, 6–9 (2023).

Acknowledgements

R.L. acknowledges support from the US Army (ACC- APG- RTP W911NF), the National Institutes of Health (NIH) 1 R01 HL169974-01), the US Department of Defense DARPA (HR00112220038), NIH 1 R011 AI135128-01 and NIH 1 R01 HL169974-01. B.M. acknowledges support from NIH 1 R01 HL169974-01, NIH 1 R011 AI135128-01 and NIH 1 R01 HL169974-01. N.T. acknowledges support from NIH R01HL166759 and R01HL142496. I.S. acknowledges support from NIH NCI R01CA270210.

Author contributions

The authors are listed in alphabetical order in the author list. R.L. conceived the Perspective and drafted the first outline. N.T. wrote the section on digital twins in cardiology. B.M. and I.S. contributed to the other sections. All authors reviewed and edited the final version.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence should be addressed to R. Laubenbacher.

Peer review information *Nature Computational Science* thanks A. M. Alaa and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Primary Handling Editor: Fernando Chirigati, in collaboration with the *Nature Computational Science* team.