

Medical digital twins: enabling precision medicine and medical artificial intelligence

Christoph Sadée, Stefano Testa, Thomas Barba, Katherine Hartmann, Maximilian Schuessler, Alexander Thieme, George M Church, Ifeoma Okoye, Tina Hernandez-Boussard, Leroy Hood, Ilya Shmulevich, Ellen Kuhl, Olivier Gevaert



Lancet Digit Health 2025

Published Online
<https://doi.org/10.1016/j.landig.2025.02.004>

Stanford Center for Biomedical Informatics Research, Department of Medicine, Stanford, CA, USA (C Sadée MRes, T Barba MD PhD, A Thieme Dr med, Prof T Hernandez-Boussard PhD, Prof O Gevaert PhD); Division of Medical Oncology, Department of Medicine, Beth Israel Deaconess Medical Center, Boston, MA, USA (S Testa MD); Department of Medicine (S Testa), Department of Biomedical Data Science (M Schuessler MS MPP, Prof O Gevaert), and Departments of Mechanical Engineering and Bioengineering (Prof E Kuhl PhD), Stanford University, Stanford, CA, USA; Department of Internal Medicine, Édouard Herriot Hospital, Lyon, France (T Barba); Department of Radiology, Hospital of the University of Pennsylvania, Philadelphia, PA, USA (K Hartmann MD PhD); Department of Radiation Oncology, Charité-Universitätsmedizin Berlin, Berlin, Germany (A Thieme); Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA, USA (Prof G M Church PhD); University of Nigeria Teaching Hospital, Ituku-Ozalla, Enugu, Nigeria (Prof I Okoye MBBS, FWACS); University of Nigeria, Nsukka Center of Excellence for Clinical Trials, Enugu, Nigeria (Prof I Okoye); Institute for Systems Biology, Seattle, WA, USA (Prof L Hood MD PhD, Prof I Shmulevich* PhD)

*Prof Ilya Shmulevich died in April, 2024

Correspondence to: Prof Olivier Gevaert, Stanford Center for Biomedical Informatics Research, Department of Medicine, Stanford, CA 94304, USA ogevaert@stanford.edu

The notion of medical digital twins is gaining popularity both within the scientific community and among the general public; however, much of the recent enthusiasm has occurred in the absence of a consensus on their fundamental make-up. Digital twins originate in the field of engineering, in which a constantly updating virtual copy enables analysis, simulation, and prediction of a real-world object or process. In this Health Policy paper, we evaluate this concept in the context of medicine and outline five key components of the medical digital twin: the patient, data connection, patient-in-silico, interface, and twin synchronisation. We consider how various enabling technologies in multimodal data, artificial intelligence, and mechanistic modelling will pave the way for clinical adoption and provide examples pertaining to oncology and diabetes. We highlight the role of data fusion and the potential of merging artificial intelligence and mechanistic modelling to address the limitations of either the AI or the mechanistic modelling approach used independently. In particular, we highlight how the digital twin concept can support the performance of large language models applied in medicine and its potential to address health-care challenges. We believe that this Health Policy paper will help to guide scientists, clinicians, and policy makers in creating medical digital twins in the future and translating this promising new paradigm from theory into clinical practice.

Introduction

The development of personalised medicine—treatment planning based on data specific to an individual patient—has been a challenge in the medical community for decades. At present, treatment decisions rely largely on best practice guidelines and a physician's interpretation of the available data. With unprecedented clinical volumes and the ever-increasing number and complexity of data sources, this practice has become a difficult endeavour that results in a generalised treatment approach instead of personalised care. As summarised by the US National Academy of Medicine, “The accumulation of data has created a situation where health-care providers are responsible for interpreting, aggregating, and synthesising data far beyond human capacity”.¹ This issue is further exacerbated in low-income and middle-income countries where patient-to-doctor ratios far exceed levels recommended by WHO.² A particular application in engineering, the digital twin, translated to medicine, has been proposed as a potential solution—namely, the medical digital twin.^{3–6}

Medical digital twins combine diverse health data streams and disease modelling to produce a dynamic copy of a patient that guides the clinical team towards personalised treatment while alleviating workload. This process allows active sharing and interpretation of data to augment clinical decision making. Medical digital twins offer clear advantages over conventional care by generating a highly detailed and personalised disease model, the patient-in-silico, which can be used to visualise patient disease metrics, predict progression, and simulate different treatment outcomes. The highly personalised nature of this approach offers the potential to improve health care in which otherwise generalised treatment approaches are applied to allcomers. Critically, in contrast to conventional models, the

patient-in-silico evolves alongside the patient, a paradigm shift for treatment in the future.⁷

As a result of such promising features, an emerging body of literature recommends a digital twin approach for a widening array of diseases. However, no formal consensus exists regarding the characteristics that distinguish digital twins from conventional models because of the novelty of this approach in the medicine field. Standalone artificial intelligence (AI) and mechanistic disease models, already abundant in the literature, are often being rebranded as medical digital twins to increase their impact.⁸ This practice risks diluting the concept and creating a negative long-term perception as the technology falls short of its goals and suggested benefits.

To realise the potential of medical digital twins in providing personalised medicine and improving the health outcomes for all, the fundamental features of the medical digital twin and the nature of the technologies available and processes required to enable the use of digital twins in clinical practice need to be delineated. Therefore, in this Health Policy paper, we address the following question: what defines a medical digital twin and how can we build one? We outline the main components that define a digital twin and identify their medical equivalent. We consider key enabling technologies, proposing the symbiotic application of large language models and medical digital twins, and discuss strategies for testing and implementation.

From digital twin to medical digital twin

The term digital twin has sometimes been narrowed to a standalone model that replicates a physical object. However, while a detailed model is a central facet of a digital twin, the original concept incorporates additional components and properties. In engineering, the digital twin is defined as a

virtual representation connected to a physical object, implying that changes to the original are dynamically reflected in the copy. A digital twin can be broken down into the following five essential components: physical object, data connection, model, interface, and twin synchronisation.^{3,4,9,10} We outline each component and its respective medical equivalent of patient, data connection, patient-in-silico, interface, and twin synchronisation (figure 1, table).

The physical object is the focus of the digital twin. Within engineering, this physical object ranges from individual devices (such as engines) to entire processes including assembly lines and complex supply chains.^{11–13} The physical object is characterised by a stream of sensory data measuring its distinct properties. A medical digital twin focuses on the patient's health and environment, as described by multimodal data—including, but not limited to, laboratory tests, imaging, sequencing, and digital health assessments.^{14–16} The data are selected on the basis of medical necessity, which might result in initial sparse data acquisition that might later become extensive. The digital twin might rely on selective data; however, it should recommend the acquisition of additional data in cases of uncertainty. Although the patient as a whole receives treatment, not every localised medical issue necessitates a comprehensive whole-body digital twin; in such instances, the patient becomes the specific organ or system of interest until a broader context is needed.

The data connection processes the collected data from the physical object and forwards it to the model. This process poses a substantial challenge for medical applications because the data are largely unstructured and multimodal, warranting substantial harmonisation and data fusion capabilities for data collection. For example, deriving a diagnosis from the electronic health record often requires chart review by an individual with medical expertise. These challenges are exacerbated in instances such as searching for laboratory values with different names across various medical systems and deriving measurements from medical imaging. The data connection requires data fusion and harmonisation to combine data modalities in a meaningful way and obtain disease metrics for integration into a model. In recent years, a plethora of AI-based techniques have shown great capacity in diagnosing, identifying, and extracting this information.^{17–20}

The model generates an in-silico copy of the physical object based on the information supplied through the data connection. In terms of medical digital twins, this copy is termed the patient-in-silico, a detailed patient model that aims to accurately reproduce biological processes to predict their change over time and with treatment. One can use digital twins not just to model disease but also to consider wellness, working to maintain an individual's health and realigning it towards an ideal baseline. Mechanistic models have been the go-to tool to establish realistic models for a wide variety of purposes, from insulin receptors in diabetes and tumour growth in cancer, to the function of entire

organs.^{21–24} While mechanistic models can closely resemble the patient, it is often difficult to derive and accurately solve these models, requiring simplification and assumptions that lead to inaccuracies. Novel technologies that merge both mechanistic models and AI approaches offer enormous promise in generating a realistic and high-fidelity patient-in-silico model, while mitigating their respective shortcomings.^{25–27}

The engineering interface allows an operator to interpret the model, capture its uncertainty, and adjust the physical object or process accordingly. Translated directly to medical practice, direct application of the model would require the clinical team to face the challenge of interpreting a complex model instead of interpreting complex data. In this Health Policy paper, we propose uncertainty quantification alongside an AI-based interface, such as a ChatGPT-like large language model, to effectively act as an intermediary between the patient-in-silico model and the clinical team.²⁸ For example, such an interface might receive a request from the clinical team to establish the effect of different cancer therapies, query the patient-in-silico, and subsequently integrate its findings into the electronic health record. The findings are evaluated by uncertainty quantification, which tracks the accumulation of different errors from data acquisition to modelling and provides a measure of the deviation between the patient and the patient-in-silico model.²⁹ Furthermore, techniques such as parameter sensitivity analysis for mechanistic models can provide insights into the importance of specific data modalities for model prediction, with explainable AI techniques providing similar insights for machine learning approaches.^{30,31} This type of interface would allow flexibility in the diversity of tasks carried out by the clinical team, while basing its findings on the complex and rigorous models behind the patient-in-silico model.

The cycle is repeated when new data are collected, justifying a re-evaluation and a corresponding, near real-time update of the model, termed twin synchronisation.⁹ A major difference between a standalone model and a digital twin is the continuous updating of the digital twin when new data are available, synchronising the model within the digital twin to the object. For medical digital twins, synchronisation might occur less frequently than in real time. Instead, synchronisation takes place when recording a substantial deviation from the baseline or measuring the effects of treatment for a disease. For treatments of diseases such as cancer, twin synchronisation might involve updating the medical digital twin after each chemotherapy cycle. The choice of update frequency for a medical digital twin generally depends on the clinical context and the relevance of new data to ensure that it remains up to date.

Key enabling technologies

Medical digital twins rely on various technologies to enable incorporation of diverse multimodal data, accurate disease modelling, and interpretation. We describe key enabling

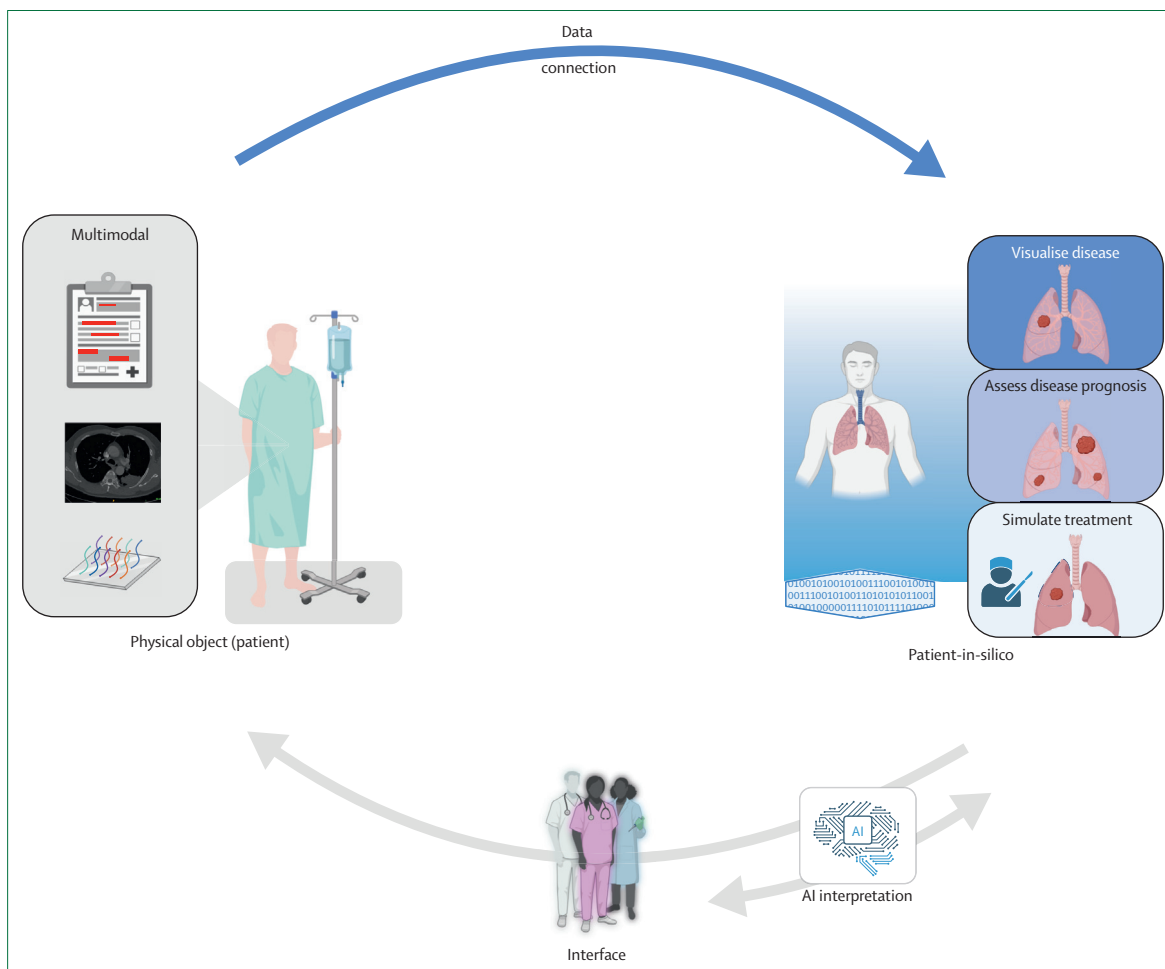


Figure 1: Outline of a medical digital twin

The physical object is described by a plethora of different data modalities (eg, electronic health records, imaging studies, and genetic data), which are processed and combined using data fusion approaches forming the data connection. The combined information is forwarded to the patient-in-silico model to visualise the disease, assess disease prognosis, and simulate treatments. The interface, with the aid of artificial intelligence (AI), allows the clinical team and patient to select an optimal treatment plan based on the patient-in-silico. The cycle is repeated as new patient data become available, synchronising the patient and the patient-in-silico (twin synchronisation). Figure created with BioRender.com.

technologies, as summarised in figure 2, that we foresee will facilitate medical digital twins.

Advances in continuous data acquisition

The fidelity of the patient-in-silico ultimately relies on the quality of the input. Repeated sampling or continuous measurements are essential for capturing dynamic interactions contributing to both wellness and disease. Ongoing advances in two fields in particular—sequencing and wearable technologies—are likely to have a substantial impact.

With advances in sequencing technologies, diseases are increasingly being characterised on a molecular basis with different subtypes revealed by specific genetic mutations. These same mutations have become targets for various therapies.³² This advance is perhaps most notable in the field of oncology. The development of liquid biopsies to detect

genetic material released by solid tumours in blood enables early and iterative detection and characterisation of heterogeneous tumours with a simple blood draw.³³ Moreover, further advances in sequencing technologies, including fourth-generation sequencing, will reduce time and cost constraints, offering the potential for swift and straightforward integration into the clinical workflow.³⁴ These newer technologies also offer the promise of direct RNA sequencing and sequencing of nucleic acid modifications that capture epigenetic changes and quantifiably reflect a patient's environmental exposures. However, these sequencing technologies only capture a snapshot at a time. Techniques have been developed that continuously record molecular events, although they have been performed only in mice so far.^{35,36} Using these technologies, it might be possible to continuously measure changes within an individual's genetic make-up for medical digital twin applications.

	Description	Enabling technology	Treatment example: oncology	Treatment example: diabetes
Patient	Organ systems or patient as a whole	Multimodal data: sequencing, medical imaging, electronic health record, wearables	Serial CT imaging, genotyping, methylation profiling, patient history, liquid biopsy testing	Measurement of blood sugar levels, nutritional assessment, monitoring of exercise and sleep habits
Data connection	Multimodal data harmonisation and data fusion	Convolutional neural networks, autoencoders, vision and language transformers	Derive tumour volumes, calculate genetic risk score, identify carcinogen exposures (eg, cigarettes, radon, and asbestos)	Extract nutrition information from food labels and calculate calorie deficit using exercise monitors
Patient-in-silico	High-fidelity organ, disease, or a whole-body patient model	Recurrent neural networks, ordinary and partial differential equations, physics-informed neural networks	Predict tumour growth, response to treatment, survival, quality-of-life measures	Predict glycated haemoglobin A1c levels and blood sugar response to food intake
Interface	Interpretation and querying of the patient-in-silico model by AI, clinical teams, and patients	Large language models, explainable AI, parameter sensitivity analysis, uncertainty quantification	Outline change in tumour size with confidence intervals, predict given different treatment options	Visualise the changes in blood sugar levels and recommend insulin dosing and food intake
Twin synchronisation	Patient-in-silico closely resembles patient and is updated when new data become available	Technologies that allow for frequent testing: wearables, liquid biopsy testing	Update the patient-in-silico after each chemotherapy cycle or equivalent treatment	Update the patient-in-silico after each insulin dose or food intake

AI=artificial intelligence.

Table: Outline of medical digital twin components, their description and enabling technologies, and examples for the treatment of cancer and diabetes

Wearable technologies offer a means to continuously measure a patient’s vital signs and environment. Unlike clinical tests that are designed and optimised to denote the onset of disease, wearables can enable detection of more subtle changes, providing insights into an individual’s wellness and early signs of an off-well state.^{37,38} Furthermore, wearables offer a means of direct communication with the patient, serving as an interface to relay information from the patient-in-silico model and clinical team to the patient—a powerful feature given the feedback loop of medical digital twins.

Artificial intelligence

AI has numerous applications within the digital twin framework, including within the data connection (feature extraction and data fusion), patient-in-silico (modelling), and interface between the clinical team and the patient-in-silico model.

Within the data connection framework, disparate and unstructured data formats must be harmonised. AI promises to achieve this harmonisation through feature extraction and data fusion. Feature extraction is the extraction of either physical or abstract features from unstructured data. Pertaining to physical features, AI algorithms can measure tumour volume using tumour segmentation from CT scans or identify laboratory values from the electronic health record among other applications.^{17,39} These metrics can then be directly integrated into the patient-in-silico model as structured data. Alternatively or supplementarily, each AI model can extract abstract feature vectors that distil key information from unstructured data. These feature vectors, while not directly mapped to physical features, can be used to predict a wide array of medically relevant metrics.^{19,40} These abstract feature vectors are then concatenated in a data fusion approach and passed to the patient-in-silico.¹⁶

Data fusion is an essential component of the medical digital twin paradigm that relies on multimodal data to make

predictions or derive disease metrics that are otherwise unattainable with a single-data modality. The field of radiogenomics is a success story that shows how combining imaging and genomic data via AI can yield valuable insights not achievable by either discipline alone.^{20,41,42} For example, a 2019 study merging transcriptomic and imaging data of medulloblastoma identified unique imaging features that predict molecular subgroups of the disease correlated with different outcomes.⁴³ A broader approach involves the fusion of the proteome, metabolome, microbiome, genome, and clinical laboratory values to predict transition from a state of wellness to that of early disease.^{44,45} Fusion of different data types will ultimately enable us to derive a more holistic overview of a patient.

In creating the patient-in-silico, AI can generate predictive models even if the underlying disease mechanisms are unknown, providing an important alternative and adjunct to mechanistic models. Recurrent neural networks are particularly well suited for the generation of predictions.⁴⁶ For example, recurrent neural networks have been used to predict changes in lung tumour shape throughout a radiation treatment cycle and mitigate exposure to healthy tissue.⁴⁷ Recurrent neural networks have also been used to predict the accumulation of somatic mutations in cancer cells, with implications for treatment strategy and prognosis.⁴⁶ While further testing is required, successful applications of AI allow for the diagnosis and prognosis of complex diseases. Importantly, AI can be used without a-priori knowledge of disease mechanisms, making it a powerful tool for medical digital twins.

True clinical application of a medical digital twin will require an interface that enables physicians and patients to interact with the model, interrogate its reasoning, and place it in the context of clinical guidelines. The rapidly advancing field of foundation models shows increasing potential.⁴⁸ Technologies such as ChatGPT by OpenAI have shown advanced reasoning skills, providing recommendations and outlines on everyday tasks. Steps are already underway to




Figure 2: Enabling technologies

The process of creating a medical digital twin begins with multimodal data. Artificial intelligence (AI) approaches are uniquely positioned to predict and extract parameters from complex data sources, such as tumour size and descriptors from test or genetic risk scores. In addition, AI can incorporate the different parameters to predict disease prognosis and treatment response, requiring only limited understanding of the disease. In contrast, mechanistic modelling can incorporate disease mechanisms and patient-specific parameters to predict disease prognosis and treatment response but is limited to simpler, well understood systems because mechanistic modelling relies on well defined mathematical relationships. The merging of AI and mechanistic modelling allows high-fidelity models, such as the patient-in-silico, that incorporate complex data types with a-priori knowledge about disease mechanics. ODE=ordinary differential equation. PDE=partial differential equation. Figure created with BioRender.com.

adapt this technology to the medical sector.^{28,49} However, ChatGPT, including its latest iteration GPTo4-mini-high, suffers from hallucinations—the making up of references and justifications for its reasoning.⁵⁰ These inaccuracies are most apparent when it fails in simple mathematical computations and logic statements, while presenting the results in a convincing manner.⁵¹ While current efforts are focused on additional training to alleviate these errors, it might only shift these hallucinations to more complex scenarios in which they are more likely to pass unnoticed, resulting in potentially more severe consequences. We propose the patient-in-silico model as a plug-in for foundation models, such as the ChatGPT-like models, to distil actionable treatment options for clinicians and patients. This strategy, although futuristic, has already proved to be effective in other fields. Through plug-ins, ChatGPT has gained the ability to query platforms such as Wolfram Alpha, which can accurately compute solutions to complex mathematical problems.⁵² Integrating large language models would enable the clinical team to query the verifiable and robust

knowledge of the patient-in-silico model while providing an easy-to-use and dynamic interface.

Mechanistic modelling

While AI has shown great predictive power for highly complex diseases, it often fails to extract biological insight, extrapolate disease evolution, and maintain predictive accuracy across different cohorts. In contrast, mechanistic models rely on understanding of the underlying disease processes, using mathematical frameworks to define relationships between disease parameters and clinical variables. With such a strong foundation in disease pathophysiology, these approaches excel in extrapolating disease evolution beyond direct observations and are robust across different cohorts. Mechanistic models can be used to model variables across a single dimension such as time (ordinary differential equations) or across multiple dimensions such as time and space (partial differential equations),^{23,53–55}

Mechanistic modelling is used to inform clinical guidelines and treatment at a population level; however, within medical

digital twins, these approaches become personalised to the individual for optimal care. For example, drug dosing guidelines are adjusted for age groups based on differences in renal function, receptor count, and other parameters as estimated by pharmacodynamics, described by an ordinary differential equation system.⁵⁶ While effective within an age group, this approach does not account for variability within each group. Within medical digital twins, drug dosing is individualised to the patient by identifying parameters of the ordinary differential equation system based on multimodal data inputs of the patient. For example, in a small clinical trial, Zhang and colleagues showed that use of personalised mathematical models to determine dosing of abiraterone for individual patients undergoing treatment for prostate cancer increased both the time to progression and overall survival.⁵⁷ Similar efforts are already being applied in the field of radiation oncology, in which radiation dose delivery can be individually estimated using partial differential equations and a patient's anatomy as obtained from imaging.^{58,59} Although these approaches show remarkable promise, they begin to fall short when the underlying biology is poorly understood or when the complexity exceeds simple modelling capabilities.

Fusion of AI and mechanistic modelling

Medical digital twins require interpretable and robust patient models that can describe complex disease phenomena. Mechanistic models, while highly interpretable, rely on a detailed understanding of the disease mechanisms and precise measurements. In contrast, AI can make predictions for complex diseases at the cost of explainability and the requirement for large-scale data. Fusion of mechanistic and AI approaches could enable each to circumvent their respective limitations. Specifically, by augmenting existing datasets, mechanistic models could help to mitigate the need for large-scale data, while the use of AI approaches could simplify parameter identification needed for mechanistic models.

Mechanistic models can augment existing datasets to supplement training data for AI and improve generalisation. Once a disease mechanism is partially understood and described by a mechanistic model, AI data can be generated by varying disease parameters within realistic bounds. For example, tumour growth can be modelled using an ordinary differential equation, and tumour oxygenation can be varied within physiological limits to generate multiple artificial growth curves. Each artificially generated curve can then be added to existing training data to improve performance of a recurrent neural network in predicting tumour growth. Alternatively, the appearance of a lung tumour within the thorax can be augmented using a partial differential equation that simulates respiration, displacing the tumour within the lung and providing additional training data.^{17,60} Therefore, mechanistic models can be used to augment pre-existing data in accordance with disease mechanics and physiology to help

with the generalisation of AI-based approaches beyond the given data.

Recently developed physics-informed neural networks (PINNs) and similar approaches offer an additional strategy to integrate mechanistic understanding in AI.^{27,61} PINNs constrain AI approaches to uphold physically accurate and representative models by using a neural network to analyse data and approximate the solution to a mechanistic model. The errors of the mechanistic model and neural network are then combined to optimise the neural network. This approach benefits from embedding of information in the neural network that aligns with the disease mechanisms of the mechanistic model in the neural network. Concurrently, the mechanistic model is solved, and its disease parameters are learned. This approach has been successfully deployed to estimate parameters of a glucose–insulin mechanistic model based on repeated glycaemic measurements.⁶² In this case, a neural network receives a time series as input and aims to predict model parameters as output. These approaches can be used to derive difficult-to-measure parameters for mechanistic models or, alternatively, to encode mechanistic understanding into AI.

Beyond using the fusion of mechanistic modelling and AI approaches to address their inherent limitations, the fusion of these two approaches in medical digital twins can also mitigate the inconsistencies in medical outcomes across patient groups. Medical treatments are commonly based on population cohorts, which might not adequately reflect the physiological or demographic variability of the wider patient population. This phenomenon is well illustrated by the use of tamoxifen in breast cancer treatment. The drug's effectiveness is influenced by specific genetic variants that affect the activity of the cytochrome P450 CYP2D6, which plays a crucial role in the activation of tamoxifen.⁶³ Different combinations of these variants result in varying levels of CYP2D6 activity, which can in turn result in dramatic changes in tamoxifen efficacy. The frequency of these genetic variants varies substantially across populations. For example, *CYP2D6*17*, associated with a 50% decrease in CYP2D6 function, has an allele frequency of nearly 20% in those genetically similar to a sub-Saharan African reference population, whereas in those genetically similar to a European population, the allele frequency is 0.4%.⁶⁴ The fusion of mechanistic models, which focus on understanding of biological processes, and AI to facilitate complex treatment predictions could provide a more comprehensive and individualised approach. Specifically, mechanistic modelling of tamoxifen levels using *CYP2D6* genotypes to derive CYP2D6 activity can enable predictions of drug toxicity and efficacy, in the absence of population-scale data coverage for AI training and, importantly, without using race as a proxy covariate. The predictions regarding drug pharmacokinetics, which are derived from such mechanistic modelling approaches, can then be integrated into AI models that predict treatment outcomes. We believe that these approaches have the potential to overcome data-driven distortions and improve treatment effectiveness

across different populations. However, further research is required to understand and eliminate the gaps in health care.

Medical digital twin implementation in research

A core tenet of medical digital twins and digital twins in general is the connection between the object and the model, such that both update alongside one another. This approach poses a particular challenge in medicine. Medical digital twins will require rigorous and in-depth testing before deployment; however, they also need continuous patient integration, creating a dilemma in which testing is required before deployment and deployment is needed for testing.

We highlight retrospective data staging as a potential research-and-test strategy to circumvent immediate testing in a patient-care setting while maintaining the core principles of a digital twin. For this purpose, retrospective patient data are staged at multiple timepoints (figure 3). Although comprehensive patient history is important for the formation of an accurate medical digital twin, its absence does not preclude formation of the twin. Instead, the absence of a comprehensive patient history would be reflected in model certainty. Instead of using the full scope of data from patient follow-ups, the digital twin model would be updated piecemeal, making predictions for the next timepoint alone. A similar approach for validation showed success in a study on the spread of COVID-19 in which model updates occurred in weekly increments based on all pre-existing data.⁶⁵ The staging of incremental updates underscores the dynamic component of a digital twin, which evolves in parallel with the patient being analysed retrospectively and approaches the concept of twin synchronisation. Importantly, for retrospective studies, the clinical decisions made can be compared against the medical digital twin to quantify the level of both uncertainty and accuracy. This approach will allow thorough testing of medical digital twins before their implementation in clinical practice.

Medical digital twins in oncology

Management of cancer depends on both patient characteristics and understanding of tumour dynamics. Medical digital twins that can integrate both have been proposed as the next step towards personalised medicine.⁷ In oncology, medical digital twins are undergoing rigorous research, with their clinical integration still in the early stages.⁶⁶ Novel treatment approaches such as adaptive therapy are showing early successes while reflecting simplified medical digital twin principles.

Adaptive therapy continuously integrates patient data into a cancer model to predict tumour resistance build-up and adapt therapy to forestall resistance. In a clinical trial by Zhang and colleagues,⁵⁷ in which such an approach improved outcomes related to time to progression and overall survival in patients receiving treatment for prostate cancer, the tumour was modelled as two cell populations, one sensitive to abiraterone treatment and the other resistant. In accordance with evolutionary theory, resistant cells are thought to experience a cost due to maintenance of

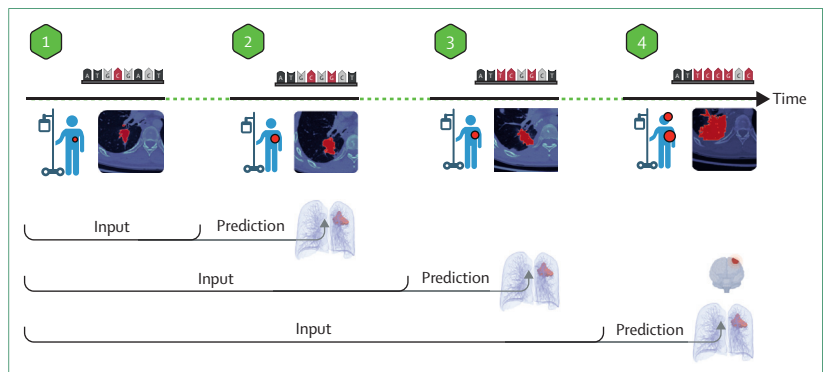


Figure 3: Retrospective data staging

Retrospective data staging allows the construction of medical digital twins using existing data. Data are separated into segments based on treatment regimens, which are then incrementally provided to the medical digital twin for updating and formulation of predictions. Retrospective studies are valuable for patient-in-silico validation and uncertainty quantification. Here, we illustrate a medical digital twin for the prediction of changes in tumour size and mutational load after every treatment step based on retrospective data. Figure created with BioRender.com.

the resistance mechanism.⁶⁷ During the clinical trial, instead of removing all sensitive cells with extended abiraterone treatment, treatment was halted once prostate-specific antigen levels fell below 50%; treatment was later restarted when this threshold was exceeded. The early cessation and reinitiation of treatment facilitate competition between resistant and sensitive populations and prevent the development of treatment resistance. In an additional retrospective patient-specific analysis by Silva and colleagues,⁶⁷ tumour growth patterns were simulated based on mechanistic modelling of the evolutionary dynamics of each patient. Notably, in this retrospective approach, similar to that outlined in the previous section, simulated growth patterns matched the observed prostate-specific antigen levels. Furthermore, the mechanistic model predicts an extended progression-free survival for even earlier cessation of abiraterone.

In the context of this application of a medical digital twin, the patient is the prostate cancer as described by blood tests and genetic markers. The data connection is simplified to the prostate-specific antigen biomarker indicating cancer levels. The patient-in-silico is built in more detail in the retrospective component of the study and uses the Lotka–Volterra equations, a type of competition mechanistic growth model, to visualise tumour growth, prognose future growth, and simulate growth based on different dosing intervals. Owing to the retrospective approach, no interface is implemented. The twin synchronisation is simplified to administration of abiraterone based on the changing and patient-specific prostate-specific antigen concentration. Overall, the study by Zhang and colleagues⁵⁷ highlights the potential of a medical digital twin, while also outlining required improvements such as personalising the model for the patient and creating an interface for final delivery. Similar approaches are underway in adaptive radiation therapy, in which treatments are tailored to adapt to the anatomical and physiological changes in a patient's body during treatment.^{68,69}

Search strategy and selection criteria

The topics in this Health Policy paper have been researched using the PubMed database and Google Scholar search engine. Articles published from Jan 1, 2015 to Dec 31, 2023 in the English language were considered, along with foundational articles such as the original study on digital twins by Grieves in 2014. An initial understanding of digital twins was established through review articles. Papers were then selected based on a combination of search terms such as “digital twin”, “mechanistic modeling”, “artificial intelligence”, and “deep learning”, and “health care”, “medicine”, “cancer”, “diabetes”, and “molecular”. Articles were excluded after an initial review if they did not highlight the core concepts of a digital twin. Key technologies required for the individual components of the digital twin, as highlighted in the selected papers, were individually searched for and reviewed. The final list of references was selected on the basis of relevance to the goal of translating the concept of the digital twin to medicine to reduce gaps in health care and improve health outcomes.

These examples of digital twin approximations in the field of oncology are indeed promising, allowing one to extrapolate and conceptualise medical digital twins that could be created at the time of cancer diagnosis. These models could simulate not only cells within a tumour, but also the patient’s overall health trajectory. Loaded with the patient’s medical history (such as demographics, comorbidities, allergies, and lifestyle), cancer symptomatology (performance status and symptoms), and tumour features (radiology, pathology, genetic profiling, transcriptomics, and metabolomics), this virtual representation could assist patients and providers in evaluating the risks and benefits of various treatment options more effectively. Moreover, the patient-in-silico might even predict potential complications of each treatment, thereby helping to facilitate expedited and optimised care. It is incumbent on physicians, scientists, and physician-scientists to translate these promising technologies into clinical practice.

Medical digital twins in diabetes

In addition to facilitating personalised medicine, digital twin technology has the potential to mitigate uneven health-care outcomes owing to a lack of access, a notable challenge both within and outside the USA. It is well documented in the literature, in reports of both the landmark Diabetes Control and Complications Trial and smaller studies, that regular, reliable medical care with frequent follow-ups improves glycaemic control and outcomes.^{70,71} Rising shortages of physicians make providing this level of care difficult if not impossible, a challenge that is not evenly distributed across world regions. The ADVICE4U clinical trial deployed a similar medical digital twin framework, demonstrating non-inferiority of an AI-based decision support system (AI-DSS) continuously guiding insulin therapy adjustments.^{72,73} Here, the patient is represented by blood glucose levels. In the data connection, multiple parameters are measured such as previous blood glucose levels, insulin dose, and carbohydrate intake. The proprietary AI-DSS software, akin to the patient-in-silico, provides personalised recommendations on insulin dosing. A simplified interface delivers physician-reviewed

recommendations via a mobile app; for a detailed example of what an interface entails, see extended data figure 7 (data management system population tracker with AI-DSS report) in the 2020 report by Nimri and colleagues.⁷² This process is then repeated on a 3-week basis, in line with twin synchronisation. To reach the full potential of a medical digital twin, future iterations will need to consider the mechanisms and effects of nutrition and exercise in addition to generating daily recommendations.^{74,75} A survey of physicians using the AI-DSS software revealed that the recommendations were reliable, helped to save time, and aided in communicating dosing decisions to patients, a factor of particular importance in communities with limited access to care.

Conclusion

Many existing computational models are being rebranded as digital twins. In this Health Policy paper, we sought to define the general concept of digital twins and to make the case for the five pillars of medical digital twins: patient, data connection, patient-in-silico, interface, and twin synchronisation. The real-time evolution across all components differentiates digital twins from conventional modelling approaches. Accordingly, we provide strategies on designing a medical digital twin, considering the recent advances in AI, and their potential role in addressing health-care challenges. To illustrate their use, we highlight examples of the use of medical digital twin frameworks for the treatment of cancer and diabetes. Our findings are also relevant for funders, health care institutions, and policy makers who will play a critical role in realising the potential of medical digital twins to improve outcomes while addressing long-term economic sustainability in health care. We believe that our Health Policy paper will guide researchers and users in creating medical digital twins for wellness and other disease applications in the future, and inspire discussions across disciplines that will help to translate this promising new paradigm into clinical practice.

Contributors

CS conceptualised the paper, established the search criteria, and wrote the original draft of the manuscript. ST, TB, KH, MS, AT, and IO provided clinical insights and where applicable, provided medically relevant examples and feedback. CS, GMC, LH, IS, EK, and OG provided the definitions of digital twins and key technologies. TH-B and IO gave insights regarding the application of digital twin technology in global health and the need to address global health-care gaps.

Declaration of interests

LH is the CEO of Phenome Health, dedicated to the creation of a second genome-like programme, aiming to link 1 million genomes and phenomes over the next 10 years. In this context, the latest computational approaches will be used, including the medical digital twin paradigm. IO is the founder of the Breast Without Spot initiative, dedicated to the early detection of breast and prostate cancers in Nigeria. Additionally, she serves on the board of Uburu, a company that aims to bridge the gap in biomedical research and industry by facilitating access to African medical data and expertise. OG is involved in unrelated research funded by the National Cancer Institute; AstraZeneca; US Food and Drug Administration; National Center for Artificial Intelligence,

Saudi Arabia; Owkin; Onc.AI; Union Chimique Belge; and Roche Molecular Systems. In addition, OG holds patents for Learning Gene Regulatory Networks Using Sparse Gaussian Mixture Models (patent S21-177, 11/22/2022), RNA to Image Synthetic Data Generator (provisional patent S22-425, 12/13/2022), and Explainable Computational Methods for Predicting Treatment Response to Immunotherapy from Histology Images of Non-Small-Cell Lung Cancer (provisional patent S24-079, 04/18/2024). GMC holds advisory roles in companies focusing on wearable health technology (LogicInk), liquid biopsy (4baseCare, Invitae, and Clinomics), and fourth-generation sequencing (INanoBio). An overview of all the advisory roles held by GMC is given at <https://arep.med.harvard.edu/gmc/tech.html>. CS and KH are married but maintain independent research careers at different research institutions and contributed separately to this work. All other authors declare no competing interests.

Acknowledgments

We thank Alex Arthur Dils, Serena Zhang, Marija Pizurica, Francisco Carrillo-Perez, Sandra Steyaert, and Yuanning Zheng for their comments and suggestions on the manuscript. We also thank Mary Kane Goldstein for her suggestions and edits to the clinical sections. There was no funding source for this piece of work.

References

- Matheny ME, Whicher D, Thadaneys Israni S. Artificial intelligence in health care: a report from the National Academy of Medicine. *JAMA* 2020; **323**: 509–10.
- WHO. The World Health Report 2006: working together for health. Oct 31, 2006. <https://www.who.int/publications/i/item/9241563176> (accessed Aug 23, 2023).
- Grievens M. Digital twin: manufacturing excellence through virtual factory replication. March 2, 2014. <https://www.3ds.com/fileadmin/PRODUCTS-SERVICES/DELMIA/PDF/Whitepaper/DELMIA-APRISO-Digital-Twin-Whitepaper.pdf> (accessed April 19, 2022).
- Glaessgen EH, Stargel DS. The digital twin paradigm for future NASA and U.S. Air Force vehicles. April 16, 2012. <https://ntrs.nasa.gov/search.jsp?R=20120008178> (accessed July 7, 2022).
- Björnsson B, Borrebaeck C, Elander N, et al. Digital twins to personalize medicine. *Genome Med* 2019; **12**: 4.
- Laubenbacher R, Mehrad B, Shmulevich I, Trayanova N. Digital twins in medicine. *Nat Comput Sci* 2024; **4**: 184–91.
- Hernandez-Boussard T, Macklin P, Greenspan EJ, et al. Digital twins for predictive oncology will be a paradigm shift for precision cancer care. *Nat Med* 2021; **27**: 2065–66.
- Wright L, Davidson S. How to tell the difference between a model and a digital twin. *Adv Model Simul Eng Sci* 2020; **7**: 13.
- Jones D, Snider C, Nassehi A, Yon J, Hicks B. Characterising the digital twin: a systematic literature review. *CIRP J Manuf Sci Technol* 2020; **29**: 36–52.
- National Academies of Sciences, Engineering, and Medicine. Foundational research gaps and future directions for digital twins. Washington, DC: The National Academies Press, 2024.
- Qi Q, Tao F, Zuo Y, Zhao D. Digital twin service towards smart manufacturing. *Procedia CIRP* 2018; **72**: 237–42.
- Lee D, Lee S. Digital twin for supply chain coordination in modular construction. *Appl Sci* 2021; **11**: 5909.
- Renaud G, Liao M, Bombardier Y. Demonstration of an airframe digital twin framework using a CF-188 full-scale component test. ICAF 2019-Structural Integrity in the Age of Additive Manufacturing. Proceedings of the 30th Symposium of the International Committee on Aeronautical Fatigue, June 2–7, 2019, Krakow, Poland. Springer, 2019: 176–186.
- Boehm KM, Aherne EA, Ellenson L, et al. Multimodal data integration using machine learning improves risk stratification of high-grade serous ovarian cancer. *Nat Cancer* 2022; **3**: 723–33.
- Levy R, Magis AT, Earls JC, et al. Longitudinal analysis reveals transition barriers between dominant ecological states in the gut microbiome. *Proc Natl Acad Sci U S A* 2020; **117**: 13839–45.
- Steyaert S, Pizurica M, Nagaraj D, et al. Multimodal data fusion for cancer biomarker discovery with deep learning. *Nat Mach Intell* 2023; **5**: 351–62.
- Primakov SP, Ibrahim A, van Timmeren JE, et al. Automated detection and segmentation of non-small cell lung cancer computed tomography images. *Nat Commun* 2022; **13**: 3423.
- Chaunzwa TL, Hosny A, Xu Y, et al. Deep learning classification of lung cancer histology using CT images. *Sci Rep* 2021; **11**: 5471.
- Wang SY, Tseng B, Hernandez-Boussard T. Deep learning approaches for predicting glaucoma progression using electronic health records and natural language processing. *Ophthalmol Sci* 2022; **2**: 100127.
- Zhou M, Leung A, Echeagaray S, et al. Non-small cell lung cancer radiogenomics map identifies relationships between molecular and imaging phenotypes with prognostic implications. *Radiology* 2018; **286**: 307–15.
- Tikenogullari OZ, Costabal FS, Yao J, Marsden A, Kuhl E. How viscous is the beating heart? Insights from a computational study. *Comp Mech* 2022; **70**: 565–79.
- Metzcar J, Wang Y, Heiland R, Macklin P. A review of cell-based computational modeling in cancer biology. *JCO Clin Cancer Inform* 2019; **3**: 1–13.
- Rathee S, Nilam. ODE models for the management of diabetes: a review. *Int J Diabetes Dev Ctries* 2017; **37**: 4–15.
- Baillargeon B, Rebelo N, Fox DD, Taylor RL, Kuhl E. The living heart project: a robust and integrative simulator for human heart function. *Eur J Mech A Solids* 2014; **48**: 38–47.
- Peng GCY, Alber M, Tepole AB, et al. Multiscale modeling meets machine learning: what can we learn? *Arch Comput Methods Eng* 2021; **28**: 1017–37.
- Alber M, Buganza Tepole A, Cannon WR, et al. Integrating machine learning and multiscale modeling-perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Digit Med* 2019; **2**: 115.
- Cai S, Mao Z, Wang Z, Yin M, Karniadakis GE. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mech Sin* 2021; **37**: 1727–38.
- Han T, Adams LC, Papaioannou JM, et al. MedAlpaca – an open-source collection of medical conversational AI models and training data. *arXiv* 2023; published online April 29. <http://arxiv.org/abs/2304.08247> (preprint).
- Chan S, Elsheikh AH. A machine learning approach for efficient uncertainty quantification using multiscale methods. *J Comp Phys* 2018; **354**: 493–511.
- Sumner T, Shephard E, Bogle IDL. A methodology for global-sensitivity analysis of time-dependent outputs in systems biology modelling. *J R Soc Interface* 2012; **9**: 2156–66.
- Yang G, Ye Q, Xia J. Unbox the black-box for the medical explainable AI via multi-modal and multi-centre data fusion: a mini-review, two showcases and beyond. *Inf Fusion* 2022; **77**: 29–52.
- Aggarwal S. Targeted cancer therapies. *Nat Rev Drug Discov* 2010; **9**: 427–28.
- Pereira B, Chen CT, Goyal L, et al. Cell-free DNA captures tumor heterogeneity and driver alterations in rapid autopsies with pre-treated metastatic cancer. *Nat Commun* 2021; **12**: 3199.
- Gorzynski JE, Goenka SD, Shafin K, et al. Ultrarapid nanopore genome sequencing in a critical care setting. *N Engl J Med* 2022; **386**: 700–02.
- Loveless TB, Grotts JH, Schechter MW, et al. Lineage tracing and analog recording in mammalian cells by single-site DNA writing. *Nat Chem Biol* 2021; **17**: 739–47.
- Bhattarai-Kline S, Lear SK, Fishman CB, et al. Recording gene expression order in DNA by CRISPR addition of retron barcodes. *Nature* 2022; **608**: 217–25.
- Sadee W, Hartmann K, Seweryn M, et al. Missing heritability of common diseases and treatments outside the protein-coding exome. *Hum Genet* 2014; **133**: 1199–215.
- Inui T, Kohno H, Kawasaki Y, et al. Use of a smart watch for early detection of paroxysmal atrial fibrillation: validation study. *JMIR Cardio* 2020; **4**: e14857.
- Lossio-Ventura JA, Sun R, Boussard S, Hernandez-Boussard T. Clinical concept recognition: evaluation of existing systems on EHRs. *Front Artif Intell* 2023; **5**: 1051724.
- Pai S, Bontempi D, Hadzic I, et al. Foundation model for cancer imaging biomarkers. *Nat Mach Intell* 2024; **6**: 354–67.

- 41 Gevaert O, Echegaray S, Khuong A, et al. Predictive radiogenomics modeling of EGFR mutation status in lung cancer. *Sci Rep* 2017; 7: 41674.
- 42 Hartmann K, Sadée CY, Satwah I, Carrillo-Perez F, Gevaert O. Imaging genomics: data fusion in uncovering disease heritability. *Trends Mol Med* 2023; 29: 141–51.
- 43 Iv M, Zhou M, Shpanskaya K, et al. MR imaging-based radiomic signatures of distinct molecular subgroups of medulloblastoma. *AJNR Am J Neuroradiol* 2019; 40: 154–61.
- 44 Price ND, Magis AT, Earls JC, et al. A wellness study of 108 individuals using personal, dense, dynamic data clouds. *Nat Biotechnol* 2017; 35: 747–56.
- 45 Omenn GS, Magis AT, Price ND, Hood L. Personal dense dynamic data clouds connect systems biomedicine to scientific wellness. *Methods Mol Biol* 2022; 2486: 315–34.
- 46 Auslander N, Wolf YI, Koonin EV. In silico learning of tumor evolution through mutational time series. *Proc Natl Acad Sci U S A* 2019; 116: 9501–10.
- 47 Li R, Roy A, Bice N, et al. Managing tumor changes during radiotherapy using a deep learning model. *Med Phys* 2021; 48: 5152–64.
- 48 Moor M, Banerjee O, Abad ZSH, et al. Foundation models for generalist medical artificial intelligence. *Nature* 2023; 616: 259–65.
- 49 Singhal K, Tu T, Gottweis J, et al. Towards expert-level medical question answering with large language models. *arXiv* 2023; published online Dec 2. <http://arxiv.org/abs/2305.09617> (preprint).
- 50 Alkaissi H, McFarlane SI. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus* 2023; 15: e35179.
- 51 Plevris V, Papazafeiropoulos G, Rios AJ. Chatbots put to the test in math and logic problems: a preliminary comparison and assessment of ChatGPT-3.5, ChatGPT-4, and Google Bard. *AI* 2023; 4: 949–69.
- 52 Wolfram S. ChatGPT gets its “wolfram superpowers”! March 23, 2023. <https://writings.stephenwolfram.com/2023/03/chatgpt-gets-its-wolfram-superpowers/> (accessed Aug 26, 2023).
- 53 Baillargeon B, Costa I, Leach JR, et al. Human cardiac function simulator for the optimal design of a novel annuloplasty ring with a sub-valvular element for correction of ischemic mitral regurgitation. *Cardiovasc Eng Technol* 2015; 6: 105–16.
- 54 Sadée C, Kashdan E. A model of thermotherapy treatment for bladder cancer. *Math Biosci Eng* 2016; 13: 1169–83.
- 55 Salgia R, Mambetsariev I, Hewelt B, et al. Modeling small cell lung cancer (SCLC) biology through deterministic and stochastic mathematical models. *Oncotarget* 2018; 9: 26226–42.
- 56 Nathwani D. Antimicrobial stewardship: from principles to practice. British Society for Antimicrobial Chemotherapy, 2018.
- 57 Zhang J, Cunningham J, Brown J, Gatenby R. Evolution-based mathematical models significantly prolong response to abiraterone in metastatic castrate-resistant prostate cancer and identify strategies to further improve outcomes. *eLife* 2022; 11: e76284.
- 58 Fogliata A, Nicolini G, Clivio A, Vanetti E, Cozzi L. Dosimetric evaluation of Acuros XB advanced dose calculation algorithm in heterogeneous media. *Radiat Oncol* 2011; 6: 82.
- 59 Principi S, Wang A, Maslowski A, et al. Deterministic linear Boltzmann transport equation solver for patient-specific CT dose estimation: comparison against a Monte Carlo benchmark for realistic scanner configurations and patient models. *Med Phys* 2020; 47: 6470–83.
- 60 Werner R, Ehrhardt J, Schmidt R, Handels H. Patient-specific finite element modeling of respiratory lung motion using 4D CT image data. *Med Phys* 2009; 36: 1500–11.
- 61 Chen RTQ, Rubanova Y, Bettencourt J, Duvenaud D. Neural ordinary differential equations. *arXiv* 2019; published online Dec 14. <http://arxiv.org/abs/1806.07366> (preprint).
- 62 Yazdani A, Lu L, Raissi M, Karniadakis GE. Systems biology informed deep learning for inferring parameters and hidden dynamics. *PLOS Comp Biol* 2020; 16: e1007575.
- 63 Kanji CR, Nyabadza G, Nhachi C, Masimirembwa C. Pharmacokinetics of tamoxifen and its major metabolites and the effect of the African ancestry specific CYP2D6*17 variant on the formation of the active metabolite, endoxifen. *J Pers Med* 2023; 13: 272.
- 64 Clinical Pharmacogenetics Implementation Consortium. CPIC® Guideline for tamoxifen based on CYP2D6 genotype. January, 2018. <https://cpicpgx.org/guidelines/cpic-guideline-for-tamoxifen-based-on-cyp2d6-genotype/> (accessed Dec 4, 2024).
- 65 Area I, Fernández FJ, Nieto JJ, Tojo FAF. Concept and solution of digital twin based on a Stieltjes differential equation. *Math Methods Appl Sci* 2022; 45: 7451–65.
- 66 Ahmadian H, Mageswaran P, Walter BA, et al. A digital twin for simulating the vertebroplasty procedure and its impact on mechanical stability of vertebra in cancer patients. *Int J Numer Method Biomed Eng* 2022; 38: e3600.
- 67 Silva AS, Kam Y, Khin ZP, et al. Evolutionary approaches to prolong progression-free survival in breast cancer. *Cancer Res* 2012; 72: 6362–70.
- 68 Varian, a Siemens Healthineers Company. Daily adaptive external beam radiation therapy in the treatment of carcinoma of the cervix: a prospective trial of an individualized approach for intestinal toxicity reduction (ARTIA-Cervix). Jan 20, 2022. <https://clinicaltrials.gov/study/NCT05197881> (accessed Aug 29, 2023).
- 69 The University of Hong Kong-Shenzhen Hospital. Study on adaptive radiotherapy and multimodal information of cervical cancer assisted by artificial intelligence. March 27, 2020. <https://clinicaltrials.gov/study/NCT04022018> (accessed Aug 29, 2023).
- 70 Mehta PP, Santiago-Torres JE, Wisely CE, et al. Primary care continuity improves diabetic health outcomes: from free clinics to federally qualified health centers. *J Am Board Fam Med* 2016; 29: 318–24.
- 71 Nathan DM, DCCT/EDIC Research Group. The Diabetes Control and Complications Trial/Epidemiology of Diabetes Interventions and Complications study at 30 years: overview. *Diabetes Care* 2014; 37: 9–16.
- 72 Nimri R, Battelino T, Laffel LM, et al. Insulin dose optimization using an automated artificial intelligence-based decision support system in youths with type 1 diabetes. *Nat Med* 2020; 26: 1380–84.
- 73 Nimri R, Dassau E, Segall T, et al. Adjusting insulin doses in patients with type 1 diabetes who use insulin pump and continuous glucose monitoring: variations among countries and physicians. *Diabetes Obes Metab* 2018; 20: 2458–66.
- 74 Shamanna P, Saboo B, Damodharan S, et al. Reducing HbA1c in type 2 diabetes using digital twin technology-enabled precision nutrition: a retrospective analysis. *Diabetes Ther* 2020; 11: 2703–14.
- 75 Iacobucci G. NHS to trial “artificial pancreas” for patients with type 1 diabetes. *BMJ* 2021; 373: n1538.

© 2025 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).