Malays J Pathol 2025; 47(3) : 355 – 368

REVIEW ARTICLE

Digital twins: A new paradigm for innovation in clinical research and medical affairs

TORRESI, Giada[1,4] and VERNA, Roberto[1,2,3]*

[1]Research Center for Health and Social Wellness, Guglielmo Marconi University, via Plinio 44, Roma, Italy; [2]Academy for Health and Clinical Research, via Umbria 7, Roma, Italy 00187; [3]World Association of Societies of Pathology and Laboratory Medicine, IL, USA; [4]Department of Clinical and Molecular Sciences, Università Politecnica delle Marche, Ancona, Italy

Abstract

Digital Twin (DT) technology, originally conceptualised in engineering, has recently emerged as a transformative paradigm in healthcare, promising to redefine the generation, interpretation, and application of biomedical evidence. DTs enable real-time simulation, prediction, and optimisation of clinical outcomes. The review aims to elucidate how DTs may enhance methodological efficiency, ethical standards, and strategic innovation in biomedical science, while addressing their epistemological and regulatory challenges. A DT is a dynamic, data-driven virtual replica of a biological entity or clinical process, continuously updated through real-time data to simulate, predict, and optimise outcomes. Originating in engineering, DTs are now entering healthcare as enablers of predictive, preventive, and precision medicine. Supported by Internet of Things (IoT) technologies, cloud computing, and machine learning, DTs integrate heterogeneous data-genomic, physiological, behavioural, and environmental-into adaptive models capable of mirroring and anticipating patient trajectories. In clinical research, they enable synthetic control arms and in silico trials, reducing recruitment barriers, improving statistical power, and addressing ethical issues associated with placebo use. The recent qualification of DT-based methodologies such as PROCOVA™ by the EMA and FDA confirms their growing scientific and regulatory credibility. DTs are redefining Medical Affairs, strengthening its role as a bridge between data science and clinical practice. They enable patient-level insights and personalised scientific communication, transforming Medical Affairs into a predictive, data-driven discipline that supports evidence-based and patient-centered decisions.

Keywords: Digital Twins; Clinical Research; Medical Affairs

DIGITAL TWIN: THEORETICAL FOUNDATIONS, MODELS, AND APPLICATION ARCHITECTURE IN HEALTHCARE

Conceptual Origins and Operating Principles

The term *Digital Twin* was coined in 2010 by NASA engineer John Vickers, but its conceptual roots date back to the 1970s, when NASA used physical replicas of spacecraft for ground simulations, notably during the Apollo 13 mission.[1] These early efforts anticipated key Digital Twins (DTs) principles: controlled replicas, real-time data use, and enhanced management of complex systems under critical conditions.[2]

The foundational concept was theorised in the early 2000s by Michael Grieves, a University of Michigan researcher and Product Lifecycle Management (PLM) expert. Grieves *et al*. (2014) defined the DT as a structured set of virtual informational representations capable of fully describing a physical artifact, from the micro-atomic scale to its macroscopic geometry.[3] Ideally, all information obtainable via direct observation of a physical object could also be accessed through its digital counterpart, making the virtual twin an equivalent yet more accessible and manipulable tool. Grieves *et al.* (2014) conceptualised digital twinning as the synergy of three core components: the physical twin, the DT, and the digital

**Address for correspondence:* Roberto Verna. Research Center for Health and Social Wellness, Guglielmo Marconi University, via Plinio 44, Roma, Italy. Tel: +393711547775; E-mail: *r.verna@unimarconi.it*

thread—a three-tier architecture still widely recognised as the conceptual basis for advanced DT applications.[3]

- Physical Twin: A real entity, existing or planned, ranging from industrial products and humans to entire hospital structures. Essentially, any tangible object, system, or process can be associated with a physical twin.
- DT: Its virtual counterpart, constructed through computational models, data, and algorithms replicating its behaviour. The DT functions as an algorithm capable of producing responses analogous to the physical system under the same input conditions.
- Digital Thread: The continuous, bidirectional connection between the physical and DTs. Beyond passive data transfer, it enables dynamic interaction: the physical twin feeds the DT with updated data, which in turn provides insights, predictions, and operational guidance. While sometimes used interchangeably with DT, the term *Digital Thread* more accurately refers to the infrastructure enabling synchronous coexistence and interoperability between the two domains.

A further distinctive feature of the DT paradigm is that the life cycles of the physical and DTs need not coincide. DTs can precede the physical entity or persist beyond its lifespan, acting as dynamic repositories for analysis and optimisation.[4]

Evolution of DT Models

The DT has evolved into increasingly complex models, differing in autonomy, predictive capacity, and interaction with the physical counterpart.

- Static Twin: A basic digital replica capturing the state of a physical system at a specific moment using historical or rarely updated data. It is descriptive, suitable for documentation, analysis, or visualisation, but cannot process real-time inputs.
- Mirror Twin (Functional Twin): Maintains a unidirectional, real-time link with the physical system, updating the digital model continuously. It remains passive, unable to autonomously process data or generate actions.
- Shadow Twin (Self-Adaptive Twin): Extends the mirror twin by actively interpreting real-time data and adapting its behaviour in response to system changes. It is reactive and adaptive, capable of recalibrating based on observations.
- Intelligent Twin: The most advanced form, integrating AI and Machine Learning to learn from the environment, anticipate future scenarios, make autonomous decisions, and interact with other DTs or human agents. It functions as a cognitive avatar of the physical entity.[5,6]

From Industry to Healthcare: DT Adoption

DTs have become a central enabling technology across high-tech sectors, including advanced manufacturing, automotive, aerospace, civil engineering, and complex system management. They support smart design, lifecycle prediction, predictive maintenance, and real-time monitoring of critical infrastructure.

In healthcare, DT applications are emerging but remain at an early stage. Current solutions cannot yet integrate all individual patient characteristics-genetic, biochemical, anatomical, lifestyle, and clinical history-into a single sustainable digital model.[2,4,7] In healthcare, a DT is a dynamic digital replica of a physical entity or clinical process, designed to replicate its structural and functional characteristics virtually. Entities can include patients, organs, medical devices, or hospital infrastructure. Its key feature is real-time updating through continuous streams of clinical, biometric, environmental, and behavioural data.[2,8]

Enabling Technological Infrastructure

Implementing a healthcare DT requires integrating diverse technologies for data collection, management, analysis, and visualisation:

(IoT): networks of connected devices, including wearable, environmental, or medical sensors, collect real-time biometric and physiological data, continuously updating the virtual model.

- Cloud Computing: provides scalable storage and management of healthcare data with encryption and privacy safeguards.
- Artificial Intelligence and Machine Learning (AI/ML): enable advanced data analysis, identify complex clinical patterns, develop personalised predictive models, dynamically adapt DT behaviour, and support clinical decision-making.
- Modelling, Simulation, and Visualisation Systems: software tools generate virtual representations and allow interactive exploration of the DT.[2,8]

Data Sources

Healthcare DTs integrate multiple complementary data sources:

- Structured Clinical Data: electronic health records, lab reports, diagnostic imaging, insurance documentation, and pharmacological records.
- Digital Health Devices: smartwatches, connected glucometers, wearable/implantable/ingestible sensors capturing real-time physiological, biological, and behavioural parameters.
- Patient-Reported Data: self-assessments, questionnaires, and symptom reports.
- Real-World Data: observational studies and disease registries providing evidence outside controlled clinical trials.
- Non-Clinical Data**:** environmental and lifestyle factors, such as air quality, consumption habits, and social media activity, contextualizing patient behaviour.[2]

Logical Architecture and Representation Models of the Digital Human Twin

DTs can be modelled at varying levels of complexity, from the entire human organism to specific systems (e.g., digestive or respiratory), individual organs (e.g., the liver), or microscopic components such as tissues, cells, organelles, and even molecular structures. Disease-specific twins, such as a digital liver affected by non-alcoholic fatty liver disease, or models simulating interactions with external agents like viruses, also fall within this spectrum.

Composite DTs integrate multiple models to provide a systemic, multiscale representation of the biological subject. Instance Twins are identical digital copies of a single individual, used to test alternative clinical scenarios and compare therapeutic strategies. Aggregate Twins group in multiple instances across families, cohorts, or populations-supporting large-scale and epidemiological analyses.

The enabling technologies form the foundation of the DT architecture, which is structured around three core components: the physical entity, the virtual model, and the digital thread. The virtual twin is developed through advanced computational modelling tools, while continuous connection to the physical counterpart is ensured by the IoT. Interactions occur across multiple scales, integrating multimodal data-genetic, molecular, environmental, social, radiological, and clinical-throughout the individual's lifetime. The long-term vision for healthcare DTs envisions continuously updated, personalised models evolving dynamically with new measurement, test, or behavioural change. Such models could integrate genetic, physiological, environmental, and psychosocial data to chronically enable predictive, preventive, and personalised medicine, particularly in disease management.[9]

DT vs. Traditional Predictive Models

Traditional clinical prediction relies on statistical inference, deriving general conclusions from limited datasets. Regression models, linear or logistic, are commonly used to estimate disease risk, treatment efficacy, or clinical outcomes based on predefined mathematical assumptions such as linearity, normal error distribution, and independence among observations. These models are interpretable and robust under ideal conditions but struggle with big data environments characterised by high dimensionality and complex variable interactions.

Machine Learning (ML) offers a more flexible alternative. As a branch of Artificial Intelligence, ML algorithms autonomously learn patterns and generate predictions from complex, heterogeneous datasets, identifying relationships often undetectable by traditional methods. The DT synthesises two complementary paradigms:

- A deductive (mechanistic) approach, based on theoretical and mathematical models describing human physiology across scales, from molecular to systemic levels.
- An inductive (data-driven) approach, typical of ML, where models emerge directly from data and adapt to individual patient characteristics.

This hybrid framework allows the DT not only to mirror a patient's clinical state but also to predict health trajectories, simulate disease progression, and optimise therapeutic decisions. Thus, the DT functions as a predictive, adaptive system, supporting personalised clinical management throughout the entire care continuum, from early risk detection to treatment evaluation.[5]

DTS IN CLINICAL RESEARCH: ETHICS, EFFICIENCY, AND INNOVATION

Fundamentals of Clinical Research

Clinical research is essential for advancing biomedical knowledge, generating reliable evidence on the efficacy, safety, and therapeutic value of new interventions. It follows a structured, phased approach:

- Phase I primarily assess tolerability and pharmacokinetics in a small cohort, often healthy volunteers, to determine the maximum tolerated dose and identify acute adverse reactions.
- Phase II evaluates therapeutic activity and safety in a targeted patient population, optimising dosage for subsequent trials.
- Phase III tests clinical efficacy against standard therapy or placebo in large, multicenter, randomised controlled trials, producing data critical for regulatory approval.
- Phase IV, or post-marketing surveillance, monitors long-term real-world use, rare adverse events, and comparative effectiveness across broader populations.

Each phase adheres to internationally recognised ethical standards emphasising participant safety, protocol transparency, and scientific rigor. The integration of advanced biotechnologies and information systems has enabled the design of more targeted, efficient, and clinically representative trials.[10]

Introduction to the Use of DTs in Clinical Research

DTs hold significant potential for transforming clinical research. Approximately 80% of clinical trials experience enrolment delays, and 20% fail to meet recruitment targets, largely due to challenges in identifying suitable participants and the increasing focus on personalised medicine, which narrows eligible populations. These factors make traditional trials increasingly costly and time-consuming.

DTs offer an innovative solution by creating virtual replicas of real patients, enabling in silico experimentation with multiple therapeutic strategies. Such digital counterparts can function as control groups, allowing early drug testing in simulated environments while reducing both patient risk and study costs. Unlike traditional external controls based on historical or real-world data, DTs provide individualised predictions, estimating each patient's outcome had they been assigned to the control arm.

From an ethical standpoint, DTs may also address issues arising in comparative trials, especially when the experimental treatment is potentially lifesaving, and a placebo or standard therapy offers limited benefits. By replacing or supplementing control groups, DTs could preserve trial validity while avoiding patient exposure to ineffective or harmful treatments. Although still in its early stages, preliminary evidence suggests that DT-supported trials could effectively address the main challenges of clinical research, enabling the design of smaller-scale studies with greater statistical power, or recovering power in ongoing trials affected by recruitment difficulties or high dropout rates.[2,6,11,12]

Limitations of Public Documentation on DTs

To date, no publicly available case reports describe the use of individualised DTs as control arms supported by detailed quantitative data. This lack of transparency stems from multiple factors. Pharmaceutical companies have initiated advanced DT projects, but most results remain confidential due to industrial secrecy or ongoing pilot phases lacking full validation. In many cases, DTs are used internally for strategic purposes, such as dosage optimisation, adverse event prediction, or clinical planning, without formal publication. As a result, available documentation focuses primarily on general technological descriptions rather than quantitative evidence or direct comparisons with real-world cohorts. Given this context, the following section illustrates a representative study that, while not a fully individualised DT implementation, marks a concrete step toward their clinical adoption. Here, the term *DT* is used broadly to denote the statistical simulation of virtual cohorts derived from real-world data through predictive modelling. The study demonstrates the feasibility and adds value of integrating real-world evidence and artificial intelligence to construct robust synthetic control arms. While these results cannot yet be generalised to fully personalised DTs, they highlight a promising trajectory toward individualised, data-driven clinical trial design.[13]

Construction of a Synthetic Control Arm Using DTs: The Case of Chronic Graft-versus-Host Disease

Phesi, a company specialising in real-world data analytics and predictive modelling, developed a synthetic control arm (SCA) for patients with chronic Graft-versus-Host Disease (cGvHD), a long-term complication of allogeneic haematopoietic stem cell transplantation characterised by donor immune cell attacks on host tissues. Standard first-line therapy involves systemic corticosteroids, with 40–60% response rates and significant side effects. Given the

complexity of cGvHD, patient heterogeneity, and recruitment challenges, it represents an ideal setting for DT-based trial innovation.

The study aimed to simulate the efficacy of prednisone as standard therapy through virtual DT cohorts, thereby eliminating the need for new control-group recruitment. Using the *Trial Accelerator*™ platform-containing data on over 61 million patients across 232,909 cohorts-Phesi identified 2,042 patients with newly diagnosed cGvHD who received systemic first-line therapy. After data refinement, eight eligible cohorts totalling 438 patients were included. Efficacy was assessed via the six-month Overall Response Rate (ORR), representing complete plus partial responses. The mean ORR was 52.7%, ranging from 48% to 66% across cohorts, with no statistically significant differences between groups ($\chi^2 = 4.66$; $p = 0.70$). These findings aligned closely with published real-world data, which reports ORR values for steroid therapy in cGvHD generally ranging between 40% and 60%, confirming the validity of the synthetic approach.[13]

Methodological Validation of the Synthetic Arm
The study's primary objective was methodological: to assess whether a virtual cohort, aggregated from heterogeneous real-world sources, could replicate the reliability of a traditional randomised control group. Despite using real clinical data, this equivalence is not guaranteed, as real-world datasets often suffer from inconsistencies, missing information, and selection bias. Building a "clean" synthetic arm that mirrors the rigor of randomised trials requires standardised endpoints, harmonised inclusion criteria, and correction of temporal and geographic variability. Here, artificial intelligence proved essential, enabling automated patient identification, data harmonisation, and construction of a homogeneous, statistically valid synthetic population comparable to conventional trials.

The study thus demonstrates the feasibility of developing highly realistic synthetic control arms using high-quality historical clinical data. DTs emerge as ethical and efficient alternatives to traditional placebo groups, particularly in rare diseases where recruitment is challenging. The strong concordance between simulated and real-world outcomes suggests that, in the near future, DT-based synthetic controls could gain regulatory acceptance as valid tools for clinical trial design and evaluation.[13]

Unlearn.AI: DT Generators

Unlearn.AI is a leading company that applies artificial intelligence to clinical medicine, aiming to transform the design and conduct of clinical trials via DT technology. Its Digital Twin Generators (DTGs) are advanced predictive models designed to simulate the individualised clinical trajectory of each patient.

DTGs are trained on patient-level data from historical clinical trials and real-world evidence, encompassing a wide range of clinical variables, including biomarkers, demographics, vital signs, symptomatology, and disease progression. Essentially, DTGs are predictive models built on large datasets of prior control-group patients and comparable observational data.

Given the complexity of biological systems, no universal model exists; each DTG is disease-specific, tailored to accurately represent the clinical dynamics of conditions such as Alzheimer's disease, amyotrophic lateral sclerosis, Parkinson's disease, Crohn's disease, or type 2 diabetes. Upon enrolment in a trial, baseline clinical data of a new patient (history, imaging, physiological parameters) are collected. The DTG then predicts the patient's potential clinical course under standard-of-care or no intervention.

Experimentally, the patient's DT can serve as a synthetic control: while the real patient receives the investigational treatment, the twin simulates disease progression without intervention. Comparing actual and predicted outcomes allows assessment of treatment efficacy without a traditional placebo group.[14,15]

The PROCOVA™ Methodology
Unlearn.AI developed PROCOVA™ (Prognostic Covariate Adjustment), a statistical adjustment strategy enhancing traditional analyses by incorporating a powerful covariate: the prognostic score. First, baseline data are used to construct each patient's DT, simulating the hypothetical outcome under the control condition (placebo or standard therapy). This generates a prognostic score representing the expected outcome in the absence of active treatment.

During the statistical analysis phase, the prognostic score is included as a covariate within a traditional regression model. This allows the estimated treatment effect to be adjusted for each participant's individual prognosis. For instance, if a patient was already expected, based on their baseline characteristics, to achieve a

favourable outcome regardless of treatment, the inclusion of the prognostic score prevents this improvement from being incorrectly attributed to the experimental intervention. Conversely, if a patient was "predicted" to deteriorate but instead shows substantial improvement, the observed change can be attributed to the treatment with greater confidence. Prognostic scores explain a substantial portion of natural outcome variability, reducing background noise and isolating treatment effects more clearly. This improves statistical efficiency, enabling trials with fewer patients to achieve equivalent power or, at equal sample sizes, increasing the ability to detect genuine therapeutic effects.[14]

Regulatory Validation by EMA and FDA

In September 2022, the PROCOVA™ methodology received official qualification from the European Medicines Agency (EMA) for use as a primary analysis method in Phase II and III clinical trials with continuous clinical endpoints. In January 2024, Unlearn.AI also received positive feedback from the U.S. FDA's Center for Drug Evaluation and Research (CDER), confirming alignment with the EMA assessment. The FDA recognised PROCOVA™ as compliant with current guidelines, deeming it an accepted statistical methodology under both EMA and FDA regulations. These endorsements represent authoritative validation of PROCOVA™'s scientific and regulatory value, confirming it as one of the most advanced statistical solutions currently available to enhance clinical trial efficiency and accuracy.[14]

Operational Advantages of Integrating DTs in Clinical Trials

The implementation of DTs in clinical research offers substantial operational benefits across trial design and execution:

- Reduction in patient enrolment. By simulating individual patient trajectories without intervention, DTs reduce the need for recruiting patients into control arms, lowering logistical burden, operational costs, and ethical concerns related to placebo or suboptimal treatments.
- Increased statistical power. The comparison between the outcome observed in the real patient and that simulated by their DT allows for a reduction in inter-individual variability. This enhances the ability to detect clinically meaningful differences even in smaller sample sizes. Alternatively, it enables an increase in study power at a given sample size, thereby improving the efficiency of evidence generation.
- Optimisation of inclusion/exclusion criteria. The ability to simulate different population configurations allows for the early identification of the most relevant clinical characteristics, thereby improving participant selection. This targeted approach increases the likelihood of obtaining clinically and statistically significant results, while reducing heterogeneity and enhancing the quality of the study design.
- Improved Multiple Ascending Dose (MAD) studies. Early-phase studies are required to identify the maximum tolerated dose (MTD) and define the optimal therapeutic range for subsequent phases of development. DTs enable the virtual testing of individual responses to different doses, identifying tolerability thresholds and signalling potential adverse events before they occur in reality. Evidence generated from DTs can provide robust guidance on whether to proceed with drug development or to terminate the experimental program early, thereby limiting participants' exposure to unnecessary risks and containing overall study costs.
- Enhanced trial ethics. All patients receive the experimental treatment without being exposed to suboptimal therapies or placebo, enhancing the fairness of clinical trials. This increases study appeal, facilitates patient recruitment, and improves overall trial equity.
- Time and cost efficiency. Robust, validated virtual populations reduce the need for large real-world control arms, accelerating trial phases, lowering administrative burden, and bringing innovative treatments to patients sooner.
- Accelerated and adaptive clinical research. DTs facilitate timely interim analyses, disease progression modelling, and dynamic protocol adaptation based on emerging data.[6,11,12,14,16]

INTEGRATING DT INTO PERSONALISED MEDICINE AND MEDICAL AFFAIRS: A STRATEGIC AND SCIENTIFIC OPPORTUNITY

Personalised Medicine: Context and Potential

Personalised medicine is an emerging paradigm aimed at delivering the right treatment to the right patient at the right time, leveraging diagnostic and therapeutic tools tailored to an individual's genetic, phenotypic, biomolecular, physiological, and psychosocial characteristics. Unlike conventional medicine, which relies on generalised approaches, precision medicine emphasises interindividual variability, enhancing treatment efficacy, and overall healthcare efficiency. However, current healthcare systems often struggle to provide truly personalised care, particularly in complex conditions such as oncology, where diagnosis and treatment involve multiple layers and high clinical heterogeneity. A key challenge lies in the variability of therapeutic response among patients sharing the same diagnosis. This heterogeneity largely reflects the discrepancy between the underlying biological complexity, characterised by dysfunctional interactions among thousands of genes that vary significantly between individuals despite identical diagnoses, and the current diagnostic capacity, which remains limited to a small number of biomarkers. Consequently, a single diagnostic label may conceal diverse therapeutic needs.[8,17]

DTs in Precision Medicine

DTs can predict disease onset dynamically, considering not only patient history but also contextual factors such as environment, time, and activities, facilitating a predictive and personalised approach. They also serve as decision-support tools, allowing virtual simulation of multiple therapeutic options to optimise treatment selection. A comprehensive DT-based precision medicine model involves identifying a patient-specific pathological signal, generating multiple virtual replicas integrating thousands of clinically relevant variables, and testing various therapeutic scenarios. The optimal treatment identified virtually is then applied to the patient, enabling a scientifically guided, individualised intervention.[2,7,8,9]

Oncology: Optimising Therapy in Triple-Negative Breast Cancer

A study published in *NPJ Digital Medicine* demonstrated that MRI-based DTs can optimise neoadjuvant chemotherapy regimens in patients with triple-negative breast cancer (TNBC).[18] The study included 105 patients, each represented by an individual DT. Model validation involved comparing predicted treatment responses with actual clinical outcomes. The model's ability to discriminate between two clinical outcomes-complete pathological response (pCR) versus non-response following chemotherapy-was assessed using a ROC curve. Model performance was quantified by the Area Under the Curve (AUC), which reached 0.82 in this study, indicating high predictive accuracy and suggesting that the model can provide reliable and clinically meaningful predictions of individual therapeutic response. DTs were further used to simulate 128 alternative clinically plausible chemotherapy combinations (Adriamycin/Cyclophosphamide followed by Taxol, A/C-T). While the observed pCR was 60.95%, DT-guided simulations suggested alternative regimens could achieve pCR up to 85.71%, representing an absolute improvement of 24.76%. Notably, 26 patients who did not achieve pCR with standard therapy were identified as potentially benefiting from personalised regimens.[18]

Broader Clinical Applications

While oncology provides an ideal context to exploit the predictive power of DTs, their use is expanding across multiple clinical domains, highlighting their versatility and translational value.

In cardiology, DTs simulate cardiac electrical and mechanical activity, supporting ablation planning, selection for implantable devices, and arrhythmia risk assessment. In neurology, they predict the progression of neurodegenerative diseases such as Alzheimer's and multiple sclerosis and optimise antiepileptic treatments. For chronic conditions like diabetes, DTs enable continuous glucose monitoring and real-time insulin dose adjustment, enhancing patient self-management, and reducing complications. In rehabilitation, DTs personalise motor recovery protocols based on patient-specific musculoskeletal responses. In surgery, DTs allow the simulation of complex operative scenarios, improving procedural planning, and reducing complications.

DTs are also proving highly valuable in clinical pharmacology and pharmacovigilance, enabling the prediction of drug-related adverse events. A notable example is the liver DT, constructed using

mathematical models that integrate anatomical and physiological knowledge with clinical and pharmacological data. This model has been employed to assess the risk of drug-induced liver injury (DILI), providing an important tool to support both drug development and post-marketing safety monitoring.

Across these applications, DTs enable a shift from standardised to truly personalised medicine, anticipating, adapting, and optimising clinical decisions based on each patient's unique profile. Computationally validated models that identify potentially more effective treatments for specific patient subgroups also offer strategic insights for the pharmaceutical industry, particularly in Medical Affairs, bridging scientific innovation and clinical practice.[5,8,9]

Implementation of DTs in Medical Affairs: Impacts and Opportunities

Medical Affairs has emerged as one of the most dynamic and strategically relevant functions within pharmaceutical organisations, acting as a critical bridge between research, clinical practice, and patient needs. This evolution reflects a broader transformation toward evidence-based, personalised, and technology-integrated healthcare systems. In this context, Medical Affairs serves as a catalyst for scientific translation, converting data into actionable clinical insights.

In an era of data complexity and personalised medicine, its strategic influence is expected to grow further. Within this evolving landscape, DTs offer significant potential to enhance the impact of Medical Affairs, although their current application in this domain remains at an early developmental stage.[19,20]

Optimising Evidence-Based Strategies

The integration of DTs represents a methodological breakthrough in evidence-based medicine, enhancing the capacity to guide clinical, regulatory, and corporate decision-making through transparent, validated, and data-driven insights. These simulations enable the generation of an advanced form of real-world evidence (RWE), data derived from unselected, often more heterogeneous populations than those enrolled in registration trials, thereby reflecting the complexity of real clinical practice. The use of DTs allows for the integration and expansion of evidence obtained from traditional studies, extending analyses to patient groups frequently excluded from experimental protocols, such as the elderly, frail individuals, or those with comorbidities. Moreover, by digitally modelling specific subgroups, Medical Affairs can identify early on the patient profiles most likely to respond favourably, or to experience reduced tolerance, to a given therapy, thereby providing concrete support for treatment personalisation and clinical appropriateness.

Furthermore, DT-based simulations can explore alternative therapeutic scenarios, such as variations in drug sequencing, dosing, or combinations, that are ethically or logistically unfeasible in traditional trials. When aligned with scientific and corporate strategies, this predictive capability provides a competitive advantage, supporting the design of more targeted, sustainable, and patient-centred studies.

In essence, DT adoption transforms the notion of evidence-based strategy from a static, retrospective framework into a proactive, adaptive, and predictive model, capable of generating real-time, patient-specific clinical value.

Personalising Scientific Communication

In the context of precision medicine, personalising scientific communication has become a key strategic goal for Medical Affairs. DTs enable the virtual representation of patient-specific biological and clinical profiles, allowing the creation of dynamic, data-driven scenarios that illustrate treatment effects on individualised cases. These simulations enhance scientific materials and discussions by highlighting differences in efficacy, safety, or tolerability that are often obscured in aggregated trial data. As a result, Medical Affairs can enhance scientific dialogue with healthcare professionals, strengthen credibility as a trusted scientific partner, and convey complex data as intuitive, context-driven insights, particularly valuable in engagements with Key Opinion Leaders (KOLs) and institutional stakeholders.

This approach proves particularly effective in advisory boards and multidisciplinary scientific discussions, where the availability of interactive models fosters stakeholder engagement, facilitates debate on clinical rationales, and promotes shared reflection on potential therapeutic implications. DTs thus go beyond their role as analytical tools, becoming advanced communication mediators capable of

conveying scientific content with greater impact and clinical relevance.

In summary, integrating DTs into Medical Affairs communication strategies allows a shift from standardised information delivery to personalised, interactive, and patient-relevant scientific dialogue.

Clinical Decision Support

In an era marked by clinical uncertainty, interindividual variability, and the proliferation of heterogeneous data, personalised decision support has become a central challenge for healthcare professionals.

The application of DTs by Medical Affairs proves especially valuable in complex cases, such as frail, multimorbid patients or those underrepresented in clinical trials, where predictive models can:

- Identify the most appropriate therapeutic strategy when clinical evidence or guidelines are lacking
- Anticipate treatment efficacy and estimate the likelihood of adverse events or drug interactions, enabling comparative assessment of multiple therapeutic options
- Dynamically adapt therapeutic pathways based on real-time clinical evolution

Through this approach, Medical Affairs evolves from a passive repository of information to an active mediator between data science and clinical practice, translating DT-generated simulations into actionable, scientifically validated insights. This model establishes an advanced, evidence-based framework for precise and adaptive clinical decision-making.

Contribution to Corporate Strategy

Therapeutic innovation advances rapidly, demanding a sophisticated understanding of the clinical and regulatory landscape. Within this framework, Medical Affairs plays a pivotal role in guiding corporate strategy, and DTs emerge as high-value tools to support informed, evidence-driven decisions through predictive simulations, comparative analyses, and realistic modelling of therapeutic responses.

By integrating and analysing heterogeneous data sources, such as real-world evidence, genomics, imaging, and digital patient profiles, DTs enable:

1. Precise definition of target patient populations and identification of subgroups most likely to derive maximal therapeutic benefit.
2. Early prediction of differential efficacy across population segments, supporting evidence-based competitive positioning and strategic development planning.
3. Identification of new therapeutic opportunities, including off-label or adjacent disease areas emerging from simulated response patterns.
4. Exploration of potential indication extensions based on simulated evidence, guiding targeted clinical study design and providing a robust rationale for regulatory submissions.

Furthermore, DTs allow the simulation of alternative sequencing or combination strategies, optimising product use across its lifecycle. This capability supports launch planning, pipeline prioritisation, and the development of differentiated, scientifically grounded communication strategies—strengthening both clinical and corporate value creation.

Competitive Differentiation and Innovation Positioning

The strategic integration of DTs within Medical Affairs functions as a catalyst for competitive differentiation and innovation leadership in the biomedical sector. By leveraging DTs, pharmaceutical companies position themselves at the forefront of digital and personalised medicine, demonstrating a tangible commitment to predictive, patient-centred care. Moreover, DT-based simulations enhance the credibility of products and therapeutic positioning, even where traditional clinical evidence is limited or immature. Presenting simulated data on realistic patient profiles with individualised response predictions supports clinical rationales more authoritatively, differentiates the value of the molecule, and enables persuasive, scientifically grounded communication of treatment benefits. In sum, adopting DTs is not merely a technological choice but a strategic tool for scientific branding, reinforcing innovation identity, consolidating market position in digital health, and generating reputational value.

LIMITATIONS, CHALLENGES, AND ETHICAL IMPLICATIONS OF DT IMPLEMENTATION IN HEALTHCARE

Technological and Operational Challenges

Intrinsic Complexity of Human Biology

Unlike mechanical systems with well-defined components and standardised designs, the

human body is a highly complex and dynamic system. Everyone exhibits a unique health profile shaped by interactions among genetic, environmental, behavioural, and stochastic factors, combining in unpredictable ways. Developing high-quality predictive models requires advanced computational architectures capable of simultaneously representing multiple interdependent clinical features and dynamically adapting to physiological changes over time.

Data Acquisition and Integration

A major barrier to clinical translation and operational deployment of DTs in healthcare is the acquisition of accurate, synchronised, and multimodal data. Potential sources include electronic health records (EHRs), diagnostic imaging, wearable devices, biosensors, and genomic databases. However, these data are often stored across heterogeneous platforms with incompatible formats and structures. This IT heterogeneity complicates interoperability between systems, namely the ability to enable coherent and secure communication among diverse information sources. Challenges also extend to the temporal synchronisation of data streams, particularly when real-time integration of continuous physiological signals or dynamic clinical updates is required. The lack of adequate IT infrastructure and shared standards further exacerbates operational difficulties, hindering the development of reliable and timely updated digital models.

Data Quality and Accuracy

Beyond availability and integration, data quality is critical for the efficacy of healthcare DTs. Incomplete, fragmented, or inaccurate data undermine model validity, leading to misleading simulations and unreliable predictive analyses. Quality issues may arise from sensor malfunctions, background noise in signal acquisition, or semantic inconsistencies across datasets due to differing collection protocols or coding systems. Longitudinal patient data are especially important, as they enable the DT to adapt to individual clinical trajectories. However, such data are often sparse, heterogeneous, or temporally incomplete. Ensuring high standards of accuracy, consistency, and traceability is therefore essential to support robust, updatable, and clinically meaningful predictive models.

Lack of International Standardisation

The widespread adoption of DTs in healthcare is hindered by the absence of globally accepted standards for their design, implementation, validation, and interoperability. Most DTs are developed in isolated academic or industrial settings, using their own formats and criteria, limiting scalability, replicability, and cross-institutional transfer.

Semantic and terminological inconsistencies, such as heterogeneous clinical coding (ICD, SNOMED CT, LOINC) further impede system integration. International multi-stakeholder initiatives involving regulators, healthcare institutions, industry, academia, and civil society are urgently needed. Organisations such as *the Digital Twin Consortium*, *the Swedish Digital Twin Consortium*, and the European *DigiTwins project* are advancing in this direction, but a global, healthcare-specific standardisation framework remains to be established. Shared frameworks would facilitate DT development, ensure transparency, enhance safety, and promote equitable use.

Limitations in Methodological Transparency

Methodological transparency remains a critical limitation for DTs in healthcare. Many models rely on opaque machine learning or deep learning algorithms ("*black boxes*"), making it difficult for clinicians or researchers to interpret the rationale behind predictions. Documentation on algorithm structure, training datasets, validation metrics, and update procedures is often incomplete or absent, undermining reproducibility, verification, and external review. Moreover, scientific validation is frequently limited to small experimental cohorts, lacking robustness for real-world clinical applications. The absence of a standardised framework for predictive and clinical validation also impedes regulatory approval and safe clinical adoption.

Legal, Ethical, and Governance Considerations

Data Privacy and Security

Healthcare DTs rely on large volumes of highly sensitive personal data, including medical history, diagnostic results, real-time physiological parameters, genetic and behavioural information. Ensuring access to the health data required for DT functionality without compromising privacy is a major challenge. Healthcare DT infrastructures involve the collection, analysis, and storage of sensitive personal information, which exposes data to potential unauthorised access, security breaches,

or misuse. Informed consent is critical: patients must fully understand the implications, risks, and purposes of sharing their data. Transparent consent processes, compliant with regulations such as the EU *General Data Protection Regulation* (GDPR) and the US *Health Insurance Portability and Accountability Act* (HIPAA), are necessary to maintain trust. Users should retain control over their data, including the ability to limit sharing or withdraw consent. Advanced security measures, such as data encryption, secure storage and transmission, and controlled access systems, are indispensable. Establishing a secure, compliant digital environment is crucial to safeguard data integrity and ensure confidence in healthcare DT systems.

Data Bias and Health Equity

Current healthcare datasets often exhibit systematic biases reflecting historical and structural inequalities in medical research and healthcare access. Common sources include overrepresentation of specific demographic groups, such as adult Caucasian males, and underrepresentation of women, the elderly, ethnic minorities, children, or patients with complex comorbidities. Using non-representative datasets in constructing human DTs risks perpetuating these disparities, potentially producing suboptimal or discriminatory clinical recommendations.

Medical Skepticism and Legal Liability

Clinical adoption of DTs is hindered by physician scepticism toward algorithm-driven decisions, fuelled by concerns over diagnostic errors or inappropriate treatments. In the event of predictive or clinical errors, legal responsibility among developers, technology providers, clinicians, and institutions remains ambiguous. Clear legal frameworks are essential to define accountability, delineate roles, and establish obligations for all stakeholders involved in the development, validation, implementation, and clinical use of DTs.

Redefinition of Normality and Risk of Overmedicalisation

The adoption of DTs in healthcare enables a paradigm shift in defining normality and pathology by constructing high-resolution models of individual physiology, integrating molecular, phenotypic, behavioural, and environmental data across the lifespan. This longitudinal approach surpasses the limitations of sporadic measurements typical of traditional medicine, allowing a more precise characterisation of an individual's "normal" parameters.

Traditionally, normality has been population-based, defined through statistical averages and reference ranges. DTs enable a personalised model in which normality is relative to the individual's historical baseline. For example, continuous blood pressure monitoring via wearable devices may reveal circadian or lifestyle-related variations, defining an "*individualised normality*" that may deviate substantially from population standards. Such hyper-personalisation, while enhancing precision medicine, raises ethical and clinical concerns. Health may no longer be defined merely as the absence of clinically detectable disease but as an adherence to an individualised optimal functioning pattern.

This gives rise to the category of asymptomatic disease, in which an apparently healthy individual could be identified, based on data from their DT, as being at high risk of developing a pathological condition. In this context, what is considered "normal" may be perceived as improvable, opening the door to a medicalised interpretation of human enhancement. Within a DT model, health itself may no longer be regarded as the absence of disease, but rather as a configuration within a spectrum of possibilities, efficient or "optimisable".

In this scenario, health risks are conceptualised not as a state of equilibrium but as a process of continuous optimisation, potentially generating implicit or explicit social pressures to intervene on physiological states that are fully compatible with normal life. For example, an apparently healthy individual with a high genetic predisposition to a disease (e.g., Alzheimer's in ApoE-4 carriers) could be labelled as "*at-risk*" even in the absence of symptoms. This may lead to anxiety, overdiagnosis, and preventive interventions that are not always justified. From a psychological perspective, it can significantly affect self-perception, self-esteem, and mental health. It is therefore essential to provide support resources, such as psychological counselling, to help individuals manage the emotional impact of such detailed and evolving knowledge of their health status.[2,6,7,9,21]

DISCUSSION

Clinical Research and DTs: Between Simulated Efficiency and Complex Reality.

The integration of DTs into clinical research

promises to address structural inefficiencies such as recruitment delays, rigid protocols, high costs, and ethical challenges related to placebo use. However, this pursuit of efficiency risks oversimplifies experimental complexity. Predictive models, however advanced, remain in theoretical constructs rooted in historical and selective data. Introducing synthetic control arms into clinical settings effectively assigns evidentiary value to outcomes that have not occurred but are algorithmically generated. Consequently, the very notion of *evidence* shifts from the observation of clinical phenomena to their plausible projection. The central challenge is thus epistemological rather than technical: how can medicine validate what has not happened but holds predictive significance? How can fairness and transparency be maintained when comparison occurs not between real patients but between lived experiences and digital simulations? The future of clinical experimentation will hinge on this tension between the *verifiable* and the *credible*.

Toward a Medicine of Possibilities: Regulatory and Narrative Implications

DTs function not merely as predictive technologies but as generators of clinical possibilities, simulating outcomes that, though unrealised, can shape therapeutic and regulatory decisions. In this emerging paradigm, the boundary between observable evidence and hypothetical scenarios becomes blurred. Regulatory bodies and clinicians will increasingly face the challenge of distinguishing between *empirical evidence* and *simulated prediction*. The validation of predictive models and the regulatory acceptance of synthetic trial arms extend beyond statistical soundness, implying a paradigm shift: accepting algorithmic predictions as decision-grade evidence redefines the epistemic foundations of truth in medicine. This evolving conception of clinical truth also reshapes the physician-patient relationship. What is presented as personalisation may risk becoming adaptation to algorithmic logic rather than a genuine response to human complexity.

Beyond Digital Replication: Between Apparent Personalisation and the Risk of Over-Prediction.

Relying on DTs to continuously optimise clinical pathways may reduce diagnostic and therapeutic errors, yet it risks narrowing the space for unpredictability, ambiguity, and individual variability, core dimensions of human health.

A key paradox emerges although DTs aim to enable ultra-personalised medicine; they are inherently built upon standardised algorithmic structures. Patient data are quantified and compared against large databases, producing a form of "personalisation" that often reflects a recombination of known parameters rather than a genuine engagement with individual complexity. Consequently, *precision* may come to mean adherence to the model, rather than a deeper understanding of the person.

The challenge ahead is not merely to refine the predictive power of DTs, but to design models capable of incorporating the unquantifiable, the patient's narrative, clinical ambiguity, lived experience, and the essential uncertainty that has always characterised medical practice. Only by doing so can future medicine remain both technologically advanced and authentically human.

Medical Affairs and DTs: From Scientific Validation to Cultural Mediation

Traditionally, the custodian of post-approval evidence, Medical Affairs is now called to validate predictive models, interpret simulations, and guide the transition toward data-driven medicine. This expanded role requires a profound cultural shift: it is no longer sufficient to communicate therapeutic efficacy; professionals must contextualise predictive outputs, discerning what is clinically plausible from what is computationally optimal. Medical Affairs becomes the mediator between the language of algorithms and that of clinicians. Its authority will increasingly depend not only on data quality, but on its ability to question model assumptions, recognise bias, and preserve the diversity and complexity of real-world clinical experience.

CONCLUSION

The DT paradigm in healthcare emerges at a pivotal moment in contemporary medicine: the shift toward predictive, adaptive, and personalised models, where digital data no longer merely accompany clinical practice but can anticipate decisions, shape perspectives, and sometimes guide therapeutic trajectories. The aim of this work has not been to celebrate these technologies, but to critically examine their significance, applications, limitations, and implications.

In clinical research, DTs offer solutions

to well-known challenges: methodological complexity, recruitment difficulties, and the ethical dilemmas associated with placebo use. The creation of synthetic control arms, the simulation of personalised therapeutic responses, and the reduced reliance on purely observational data open compelling opportunities, but also new scientific responsibilities. Predictive value does not equate to clinical value, and reproducible simulations cannot replace real patient experiences.

It is within this intermediate space, between algorithm and decision, that Medical Affairs assumes an increasingly central role, balancing innovation, evidence, and sustainability. Its evolution is both technical and cultural: Medical Affairs must interpret, integrate, and communicate predictive models while maintaining clinical coherence and methodological rigor. DTs position Medical Affairs as an active mediator between scientific complexity and clinical applicability, between digital models and therapeutic needs, between the language of data and the language of care.

Yet, like any transformative technology, DTs carry structural, methodological, and ethical vulnerabilities. They require large, harmonized datasets, interoperable infrastructures, and transparent, interpretable models. Challenges remain in data fragmentation, algorithmic opacity, and reproduction of biases in clinical datasets. Without critical and responsible management, simulation of risk perpetuating inequalities rather than addressing them.

Adopting DTs also entails confronting a new epistemological landscape, from lived reality to modelled projections. Here, the personalisation promised by predictive models can become an algorithmic illusion, interpreting patients through statistical configurations rather than their unique individuality. The notion of normality is also redefined: no longer a population average, but a continuously monitored, optimised, and potentially medicalised individual trajectory.

Ultimately, this work offers a clear-eyed perspective on innovative but non-neutral technology. DTs are powerful transformative tools, yet they demand new competencies, nuanced responsibilities, and constant critical oversight. True progress will arise only if these technologies are integrated into a model of medicine that embraces complexity, relationality, and uncertainty. In the end, care cannot be reduced to a model-derived prediction, it is foremost a human encounter, grounded in trust, shared decision-making, and respect for patients' values, fears, and preferences.

Authors' contributions: G.T. and R.V Conceptualisation; formal analysis; writing; review and editing; supervision; funding acquisition; final approval of the manuscript. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: Both authors declare no conflict of interest."

REFERENCES

1. Vickers J. The Promise and Applications of Digital Twins. NASA Advanced Manufacturing February 26 - 28, 2025 Baton Rouge, LA
2. Armeni P, Polat I, De Rossi LM, Diaferia L, Meregalli, S, Gatti A. Digital Twins in Healthcare: Is it the Beginning of a New Era of Evidence-Based Medicine? A Critical Review. J. Pers. Med. 2022; 12, 1255.
3. Grieves M. Digital twin: manufacturing excellence through virtual factory replication. White paper 1. 2014; 1-7
4. Pellegrino G, Gervasi M, Angelelli M, Corallo A. A Conceptual Framework for Digital Twin in Healthcare: Evidence from a Systematic Meta-Review. Information Systems Frontiers. 2025; 27:7–32
5. Tortora M, Pacchiano F, Ferraciolli SF et al. Medical Digital Twin: A Review on Technical Principles and Clinical Applications. J. Clin. Med. 2025; 14, 324.77
6. Katsoulakis E, Wang Q, Wu H, et al. Digital twins for health: a scoping review. NPJ Digit Med. 2024;7(1):77. Published 2024 Mar 22.
7. Cellina M, Cè M, Alì M et al. Digital Twins: The New Frontier for Personalised Medicine? Applied Sciences 2023; 13 (13): 7940.
8. Papachristou K, Katsakiori PF, Papadimitroulas P, Strigari L, Kagadis GC. Digital Twins Advancements and Applications in Healthcare, Towards Precision Medicine. J. Pers. Med. 2024; 14, 1101.
9. Kamel Boulos MN, Zhang P. Digital Twins: From Personalised Medicine to Precision Public Health. J. Pers. Med. 2021; 11, 745.
10. Wang T, Liu M, Peng B et al. From Bench to Bedside: A Review of Clinical Trials in Drug Discovery and Development. arXiv:2412.09378v2. 2024
11. Bordukova M, Makarov N, Rodriguez-Esteban R, Schmich F, Menden MP. Generative artificial intelligence empowers digital twins in drug discovery and clinical trials. Expert Opinion on Drug Discovery. 2024; 19:1, 33-42.
12. Mann DL. The Use of Digital Healthcare Twins in Early-Phase Clinical Trials: Opportunities, Challenges, and Applications. JACC Basic Transl Sci. 2024;9(9):1159-1161.
13. Li G, Chen YB, Peachey J. Construction of a digital twin of chronic graft vs. host disease patients

with standard of care. Bone Marrow Transplant. 2024;59(9):1280-1285.
14. Fisher C. How to Position Clinical Trials with Digital Twins for Regulatory Success. Unlearn AI, San Francisco, CA (no date)
15. Fisher C, D'Angelo H. Defining Digital Twins in Drug Discovery and Development. Unlearn AI, San Francisco CA (no date).
16. Vidovszky AA, Fisher CK, Loukianov AD, et al. Increasing acceptance of AI-generated digital twins through clinical trial applications. Clin Transl Sci. 2024;17:e13897.
17. Lajmi N, Alves-Vasconcelos S, Tsiachristas A, et al. Challenges and solutions to system-wide use of precision oncology as the standard of care paradigm. Camb Prism Precis Med. 2024;2:e4.
18. Wu C, Lima EABF, Stowers CE, et al. MRI-based digital twins to improve treatment response of breast cancer by optimising neoadjuvant chemotherapy regimens. NPJ Digit Med. 2025;8(1):195
19. Demuth S, De Sèze J, Edan G, Ziemssen T, Simon F, Gourraud PA. Digital representation of patients as medical digital twins: data-centric viewpoint. JMIR Med Inform. 2025;13(1):e53542.
20. Setia S, Ryan N, Nair PS, Ching E, Subramaniam K. Evolving Role of Pharmaceutical Physicians in Medical Evidence and Education. Adv Med Educ Pract. 2018;9: 777–790.
21. Bruynseels K, Santoni de Sio F, van den Hoven J. Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm. Front Genet. 2018;9:31.