




MINISTÈRE
DU TRAVAIL, DE LA SANTÉ,
DES SOLIDARITÉS
ET DES FAMILLES

*Liberté
Égalité
Fraternité*

French
Digital Health Delegation



Digital Twins in Health Issues, Definitions and Ethical Challenges

BRIEFING NOTE FOR THE FUTURE WORKING
GROUP GT15 OF THE FRENCH DIGITAL HEALTH
ETHICS UNIT

SEPTEMBER 2025

Summary

The *digital twin* is a health innovation inspired by the industrial sector. It takes the form of a dynamic digital model of a patient, an organ, or a physiological process, built from medical, biological, environmental, and behavioural data. Unlike a static image, it continuously evolves and can simulate the progression of a disease or the potential effect of a treatment. Its promise is to make medicine more personalised, predictive, and precise. Yet this technology also raises profound questions.

Prepared by a group of experts: scientists, researchers, industry stakeholders, jurists, philosophers and sociologists, who are all actively involved in digital twin projects, this document underscores that while digital twins open promising perspectives, they also bring numerous ethical challenges.

The first challenge concerns the reduction of one person to data, and the risk of confusing the model with the human reality it represents. The very term “*twin*” fosters this ambiguity, suggesting equivalence, while it is only a partial construction, derived from technical choices and selected datasets. Another issue concerns the caring relationship. Traditionally based on the encounter between a patient and a physician, it is being reshaped by digital technologies: the dyad becomes a quadrangular relationship, now including the patient, the physician, the digital twin, and its designer. This new configuration redistributes trust and responsibilities, potentially weakening the autonomy of clinical judgment.

The document also highlights the danger that predictions produced by such models might become deterministic: statistical probabilities can be perceived as certainties, limiting the agency of patients and caregivers alike.

Beyond human and relational concerns, another dimension arises: sustainability. The power of digital twins relies on heavy, costly, and energy-intensive infrastructures. Their widespread adoption therefore presents a dual challenge: reducing the environmental impact of digital health and preventing such technologies (due to their ecological cost) from becoming exclusive tools, deepening inequalities in access to care.

In conclusion, the digital twin appears as a powerful and hopeful tool to enrich medicine and support clinical decision-making. However, it must be perceived for what it truly is: a representation, not a duplicate, of a person. Its value will depend on the way it is governed (i.e., with transparency, fairness, and responsibility), such that it remains a reflective ally serving the caregiver–patient relationship, and never a substitute for it.

Members of the expert working group

List of contributors to the work that made possible the development of this briefing note for a future working group, GT15, of the Digital Health Ethics Unit of the French Delegation for Digital Health (DNS), dedicated to the ethics of digital twins.

Last name	First name	Position
BARBIER	Michèle	Co-pilot of the Digital Twin Ethics Working Group Coordinator of Digital Twin Projects at INRIA, France
SEROUSSI	Brigitte	Co-pilot of the Digital Twin Ethics Working Group Director of the Ethics Unit, French digital health delegation
BOUNAUD-DEVILLERS	Gilbert	French association of Patients with Cardiac Implantable Electronic Devices (APODEC)
CHNEIWEISS	Hervé	President of the Ethics Committee of INSERM, eBRains, and EMBL
DESAIVE	Thomas	Professor of Biomedical Engineering, University of Liège, Belgium
DUVIGNAUD	Frédéric	Ethics Project Manager at the French Digital Health Agency
ELHADJ	Elisa	PhD Student, Life Sciences and Society Lab, Catholic University of Louvain, Belgium
GODART	Elsa	Philosopher and Psychoanalyst, Research Director at Gustave Eiffel University (LIPHA) & Associate Researcher at EHESS/CNRS (LAP), France
HIRSCH	François	INSERM, France
PECUCHET	Nicolas	Physician, 3DS, MediTwin
PEYRAT	Jean-Marc	Founder of inHEART
ROUCOUX	François	Physician and Computer Scientist, Medical Informatics Laboratory, École Polytechnique, Catholic University of Louvain, Belgium
SREEDHAR	Hari	Medical Imaging Researcher, Côte d'Azur University, INRIA, France
THEBAUT	Jean-François	French Federation of Diabetic Patients
Van GYSEGHEM	Jean-Marc	Center for Research on Information, Law, and Society, University of Namur, Belgium
WILLIATTE	Lina	Lawyer, Ethics Unit, French digital health delegation

Table des matières

Summary	2
Members of the expert working group	3
1. General Context	5
The Need for an Ethical Reflection on Digital Twins in Health	5
The Ethics Unit of the Delegation for Digital Health of the French Ministry of Health	6
2. The Conceptual Ambiguity of the Digital Twin: Between Representation and Replica	7
Origins and Industrial Heritage	7
Towards a Definition in Health	7
Distinctions with Other Concepts	10
Use Cases and Current Applications	11
Limits and Uncertainties	12
3. Ethical Issues of Digital Twins	13
Reducing the person to their digital twin: an ontological shift	13
Between Subject and Artifact: The Ambiguity of Ownership of Digital Twin Data	14
The Weight of Prediction: Between Determinism and the Capacity to Act	15
Knowledge or Wisdom: The Disappearance of Judgment in Simulated Clinical Practice	16
Consent and Algorithmic Heteronomy: Shared Decision-Making in the Era of the Digital Twin	17
The Digital Twin and the Care Relationship: From Dyad to Quadrangulation	18
Responsibility, Explainability, and Social Justice: Governance to Consider	19
Digital Twins and Ecological Footprint: Ethics Tested by Sustainability	20
4. Conclusion	22
5. Annexes	23
Presentation of Two French and European Projects on Digital Twins in Healthcare	23
EDITH	23
MEDITWIN	23

1. General Context

The Need for an Ethical Reflection on Digital Twins in Health

Over the past twenty years, digital technology has become a central tool in the field of health. Initially confined to administrative management and the digitisation of medical records, it has gradually established itself as a means to assist with diagnosis, therapeutic decision-making, and patient follow-up. The emergence of artificial intelligence, big data, and connected sensors has opened a new field: the possibility of **digitally modelling living beings**, to the point of producing a kind of virtual double.

This double, known as a **digital twin**, has become one of today's most debated innovations. In industry, digital twins have proven their effectiveness by simulating the behaviour of aircraft, engines, or production lines. Translated into healthcare applications, this approach takes on an entirely different meaning: it is no longer about optimising a machine but about representing a living organism in all its complexity, or even an entire person.

This shift from the industrial to the living world profoundly transforms the concept's meaning. The human body is not just another machine: it carries history, subjectivity, vulnerability, and emotion. It evolves within a complex physical and social environment that influences it over time, across different life stages. Consequently, it belongs to a biographical and social narrative that resists complete modelling. Hence, the **ethical challenge arises immediately**. These virtual replicas of individuals, designed to improve the personalisation of care, rely on the massive use of personal and biometric data, requiring in-depth reflection on privacy protection and digital sovereignty. Questions of transparency and explainability of predictive models are also central, as they directly influence patients' and professionals' trust in these technologies.

Other challenges include **data bias**, the potential **widening of inequalities in access to care**, and the **temptation to reduce human singularity** to an object of simulation. Can we still speak of care when a virtual twin is entrusted with predicting or guiding medical decisions?

European and French public authorities quickly recognised the importance of these issues. The European Commission has integrated digital twins into its research and innovation strategies under the term *Virtual Human Twin*. In France, the Ethics Unit of the **Delegation for Digital Health (DNS) of the French Ministry of Health**, in collaboration with **INRIA**, launched in 2025 a cycle of reflection bringing together researchers, clinicians, engineers, jurists, and philosophers. These reflections build on the French **Framework for Digital Health Ethics (CENS)** published in December 2023¹, which sets

¹ <https://esante.gouv.fr/produits-services/referentiel-ethique>

out the principles of an inclusive, transparent, and rights-respecting digital health ecosystem. However, digital twins mark a new stage: they embody both the promise of profoundly improved medicine and the risk of major ethical drifts if ethics is not embedded at the heart of their design and use.

The purpose of this document is to provide a **framework for the work of the future GT15 working group** of the DNS Ethics Unit, dedicated specifically to the ethics of digital twins in healthcare. The aim is not to provide definitive answers, but to formulate the right questions, to map the ethical challenges, and to identify priorities for analysis. This document presents the adopted approach, the themes discussed, the key insights, and proposals regarding the scope and objectives of the future GT15 group.

The Ethics Unit of the Delegation for Digital Health of the French Ministry of Health

Since its creation in 2019, the Delegation for Digital Health (DNS) of the French Ministry of Health has included an Ethics Unit. Its mission is to make ethics a central pillar of the digital transformation in health, through practical tools for awareness and ethical evaluation aimed at industry, healthcare professionals, patients, and public authorities.

The Unit defines digital health ethics at the intersection of clinical ethics (as described in the Hippocratic Oath and bioethics principles) and digital ethics, which views technology as a tool serving human users. Digital health ethics currently rests on five well-established principles: beneficence, non-maleficence, respect for individual autonomy, justice or fairness, and environmental responsibility and sustainability.

The Ethics Unit operates within the CENS framework, whose goal is to make ethical principles operational, pragmatic, and human-centred, ensuring trust between patients, healthcare providers, and digital tools. The Unit is supported by the Committee for Digital Health Ethics (COMENS), which steers overall orientations and updates to the CENS framework, ensuring alignment with technological, social, and regulatory developments, as well as feedback from the Ethico-Vigilance Platform².

Currently, 16 working groups are active under the Ethics Unit. GT15, dedicated to Digital Twins, and GT16, focusing on Artificial Intelligence and Health Democracy, are the most recent. To prepare GT15's mission, the Unit convened a prefiguration group of multidisciplinary experts, ensuring the inclusion of perspectives from scientists, clinicians, legal specialists, engineers, philosophers, sociologists, and patient representatives.

² <https://esante.gouv.fr/produits-services/referentiel-ethique/plateforme-ethicovigilance>

2. The Conceptual Ambiguity of the Digital Twin: Between Representation and Replica

Origins and Industrial Heritage

The term *digital twin* was first coined in the field of engineering. In aeronautics, it refers to the virtual replica of an aircraft, continuously fed with sensor data, allowing simulation of its behaviour and anticipation of failures. In the automotive industry, it serves to optimise engine performance and vehicle safety. In all such cases, the digital twin has a predictive optimisation function.

The strength of the concept lies in its capacity to improve reliability, increase safety, and reduce costs. This model has inspired other sectors—from logistics to urban planning—but its potential appears most ambitious and controversial in healthcare.

Towards a Definition in Health

The expert group began by exploring a definition of the digital twin applied to healthcare, starting from a generic formulation proposed by INRIA:

“A digital twin is a virtual replica of a physical object, process or system, created from real-time data and advanced modelling. In health, it is a dynamic digital model of an organ, a patient, or a medical/physiological process, enabling simulation, analysis, and optimisation of care.”

This was complemented by the European EDITH programme’s definition of the *Virtual Human Twin (VHT)*:

“A Virtual Human Twin (VHT) is a patient-specific virtual representation of real-world systems or processes, built on data-driven or knowledge-driven (most often a combination of both) predictive computer models, and usable as a clinical decision-support system, a personal health forecasting tool, or a tool for the development and personalisation of medical products.” This can be summarised as a *“simplified mathematical representation of a specific individual, parameterised using the person’s data to help clinicians better understand their health status and evolution.”*

These definitions sparked rich discussions within the expert group, particularly about the terminology to be retained as the foundation for ethical reflection. The scope of this reflection was first clarified: it concerns only the digital twin of a human individual, in the context of care or research. Although increasingly used in healthcare, digital twins of populations, processes or organisations, pose ethical issues of a different order and will require separate analysis.

Another major question concerned the extent of ethical consideration: should it focus only on the current capabilities of patient-specific digital twins (mainly clinical modelling tools), or also anticipate an “ideal twin” — a complete digital double of an individual across all dimensions? The group agreed that, given the rapid pace of technological progress, it is essential to extend ethical reflection to the notion of a “perfect” digital twin, such as to foresee and prevent potential ethical drifts.

Anticipating the future evolution of the digital twin is complex, both technically and conceptually. The first patients participating in research projects express diverse expectations: some view the twin as an assistant, others as an interlocutor for the care team, or as an analytical mirror of their own body. Some even imagine it as an “other self” that could be tested, even “harmed,” without consequence.

Through the group’s discussions, a three-part working definition emerged and is still open to debate:

1. *“A Digital Twin Is a Dynamic Virtual Representation of Human Physiology and/or Pathology taking into account multiple scales (from tissue to organ, to bodily systems, up to the whole person), built from data (personal or not: clinical, biological, genetic, behavioural, environmental, etc.), updated via sensors and data collection, and never reducible to its physical counterpart.”*

The word “representation” was preferred to “replica” which implies unattainable fidelity. The idea of a “faithful representation” was discussed but discarded, as introducing a quality judgment would create unnecessary tension (especially from a legal standpoint) since the level of fidelity must always depend on the context of use (clinical decision support, research, therapeutic development, etc.).

Every digital representation is necessarily a simplification of living complexity. The group stressed that even the most sophisticated models remain partial abstractions of reality. They must never be perceived as equivalent to the person. Overly mimetic models risk erasing this distinction, reducing the patient to their data.

This terminological choice also allows a useful legal analogy: just as a photograph represents a person at a given moment, while still invoking personal rights (image, privacy, etc), the digital twin is a representational construct that can be legally framed without being identified with the individual.

The definition also emphasizes the dynamic nature of the digital twin. Unlike a static model, it is an evolving entity, continually enriched by new data from various sources: health records, biometric sensors, environmental and behavioural inputs. This enables broader applications beyond anatomical or biological description, including psychological or cognitive dimensions, especially within the emerging field of digital psychiatry.

Finally, the group chose not to restrict the definition to a specific entity (organ, system, or whole individual), but to encompass all physiological and pathological phenomena

across scales. Thus, the digital twin may represent part or all of a person, depending on clinical or research needs. This inclusive approach captures both current and future uses, while preserving the fundamental ethical boundary: no representation, however advanced, can replace a human being.

2. *“It Uses Digital Tools — Based on Mathematical Models and Algorithms — to Infer or Compute the Evolution of Its Physical Counterpart for the purpose of*
 - *representing and understanding complex physiological and/or pathological phenomena;*
 - *testing hypotheses; and anticipating responses to treatments.”*

This part highlights the technical foundations of the digital twin: mathematical and computational models (epidemiological, statistical, mechanistic, or probabilistic), associated with algorithms enabling complex data processing and inference. The group acknowledged the risk of “freezing” the definition in today’s technologies, which may evolve, but decided that referencing mathematical models and algorithms was still necessary for precision.

The term “physical counterpart” is used broadly: it may refer to the whole person (still an aspiration more than a reality) or a subsystem (organ, tissue, biological process). The twin does not claim to be a full replication, but rather a computational construct simulating specific biological dynamics.

The phrase “testing hypotheses and anticipating treatment responses” also encompasses simulation, understood as the virtual experimentation of clinical scenarios, allowing exploration of possible disease or treatment trajectories without directly exposing patients to risk.

The definition must also state what a digital twin cannot do, as well as the ethical limits associated with it, to prevent misuse or overinterpretation:

- The digital twin cannot predict the future with certainty. The output of its simulations yields probabilities and possible scenarios, not deterministic outcomes.
- The digital twin is never neutral: models and algorithms carry biases from the data on which they are built (historical, cultural, social, etc.); thus, it is an interpretation of reality, not reality itself.
- The digital twin cannot replace clinical judgment: medical decisions cannot rely solely on computational projections. The risk is to reduce patients to algorithmic profiles, dehumanizing care. The twin must remain a decision-support tool, not a substitute for human responsibility.

In short, the digital twin is an instrument of knowledge and projection: it sheds light but does not decide. It represents but does not embody. Its power lies in its ability to simulate biological complexity; its legitimacy depends on our ability to keep it in its proper ethical place within care.

3. *“Its aim is to provide decision support for the care of the physical person, with the goal of adapting care to patient needs while respecting their rights, improving care quality and health outcomes, and optimising the cost–benefit performance of healthcare delivery.”*

While industrial digital twins are designed to optimise performance and reduce costs, in health the goal must always be to improve patient care and well-being. The digital twin thus becomes a support for shared decision-making, enabling clinicians and patients to explore the best therapeutic options together.

The term “improve” was debated: though potentially subjective, it was retained because it expresses an ethical aim from the patient’s perspective improving care quality, health, and quality of life consistent with autonomy, dignity, and the right to fair and comprehensible treatment.

The reference to performance was also retained, as it reflects one of the motivations for developing digital twins. However, it must remain subordinate to the primary aim: improving care and quality of life, not mere technical or economic efficiency.

This definition of the digital twin includes several key points:

- Fidelity and dynamism: The digital twin is not a static snapshot but a living model, continuously updated.
- Data complexity: It integrates multiple types of information, from genetics to lifestyle behaviors, as well as environmental data.
- Non-reduction: It is emphasized that the physical person cannot be reduced to their digital counterpart.
- Medical purpose: The goal is to support care, not to replace human decision-making (whether prognostic, diagnostic, therapeutic, etc.). The digital twin is also a valuable tool for basic and clinical research.

This definition is proposed as an input for the future working group (GT15), which may review and refine it.

Distinctions with Other Concepts

It is essential to distinguish the digital twin from other technological objects that are sometimes perceived as similar:

- Synthetic data: These are artificially generated data from statistical models and anonymized real datasets, able to simulate a set of fictitious patients for research purposes or algorithm testing, without any connection to real individuals. The term “synthetic patient” is sometimes used in the context of cohorts. While their use poses minimal privacy risks, reusing them as a substitute for real data (which is useful for increasing the statistical power of models) requires ethical oversight.

- Patient avatars: Visual or interactive representations, sometimes used in therapeutic education or treatment (for example, through simulations in the metaverse to address physiological or psychological issues, behavioral problems, or phobias). They are generally educational, without direct predictive or decision-making functions. Unlike the digital twin, they are not grounded in real biomedical data and do not aim to model the dynamics of a living organism. However, they do represent a set of personal and behavioral data.
- Traditional simulation models: Already used in pharmacology or epidemiology, they differ from digital twins due to their lack of real-time personalization.

The digital twin is thus distinguished from these technological objects by its grounding in an individual's data (both personal and non-personal), with a modeling, predictive, and decision-making purpose, making it an ethically more sensitive object.

Furthermore, it should be noted that the use of the word “twin” for the digital twin could be problematic from an epistemological and ethical perspective. It suggests symmetry or identity with the real person it represents, which creates a misleading confusion: the digital twin is not a duplication of the patient, nor a clone, but a computational construct based on a reduction of human complexity to a set of observable data. Some group members have proposed the term “avatar” as an alternative to avoid undue personification. However, the term “digital twin” has already become established in practice, both in scientific publications and industry discourse, and it seems difficult to move away from it. The group acknowledges this linguistic reality while calling for constant vigilance regarding symbolic or interpretative misuses that this vocabulary could induce, particularly in representations of care and the person.

It is also important to note that the meaning of “digital twin” in medicine differs from its meaning in other sectors, where its applications are more mature, such as aerospace or manufacturing. Many other terms have also been used to describe what is now considered digital twins, such as the Virtual Physiological Human.

Use Cases and Current Applications

Among all the use cases mentioned during the discussions, only those concerning a human digital twin are retained within the scope of future work:

- The digital twin as a clinical decision support system: The ability to simulate the effects of a treatment or intervention on a given patient through their digital twin allows for informed medical decision-making and the selection of the most appropriate option among several alternatives. This decision support applies to treatment choices as well as personalized prevention.
- The digital twin as a facilitator of Patient Therapeutic Education (PTE): PTE is a regulated framework whose access is still limited for patients. The digital twin is seen here as an opportunity to expand the scope of PTE and enable a larger

number of patients to benefit. Being able to visualize on one's digital twin the effects of different treatment or intervention choices, and more broadly of behavioral decisions, is highly engaging for the patients, empowering them to take an active role in their care.

- The digital twin of an organ for transplantation (heart, skin, respiratory system, etc.): In the specific context of organ transplantation, the digital twin can help increase the chances of success by refining donor-recipient compatibility assessments and anticipating potential complications. Its contribution can be particularly valuable in this field, where numerous parameters must be considered in the context of rapid decision-making.
- The digital twin as a physiological and pathological representation of the patient: The digital twin of a patient in the broadest sense (not limited to a single organ or group of organs but representing the patient in all their dimensions) represents an aspirational model for ongoing developments in this field, opening the way to a new approach to individual patient care.

Limits and Uncertainties

Despite these potentialities, the definition of digital twins remains problematic. Several questions remain:

- Should they include the psychological and subjective dimension of the patient, or be limited to measurable physiology?
- What is the appropriate scale of modeling: organ, system, or the whole individual? Alongside, what temporal resolution should be used for data collection?
- How can we distinguish between healthcare uses (where uncertainty remains critical) and research uses (where abstraction may be acceptable)?

These debates show that digital twins are not a single, well-defined object, but rather a constellation of practices and representations still under development.

3. Ethical Issues of Digital Twins

Ethical reflection arises to question whether we should do something that we can do and want to do. Figure 1 presents a representation of the ethical dimensions of the social value of a health innovation: between technical feasibility (*Can we do this?*), social desirability (*Do we want this?*), and ethical viability (*Should we do this?*). Their intersection constitutes the true value for science and society, reminding us that the “best science in the world” must also be the “best science for the world.”

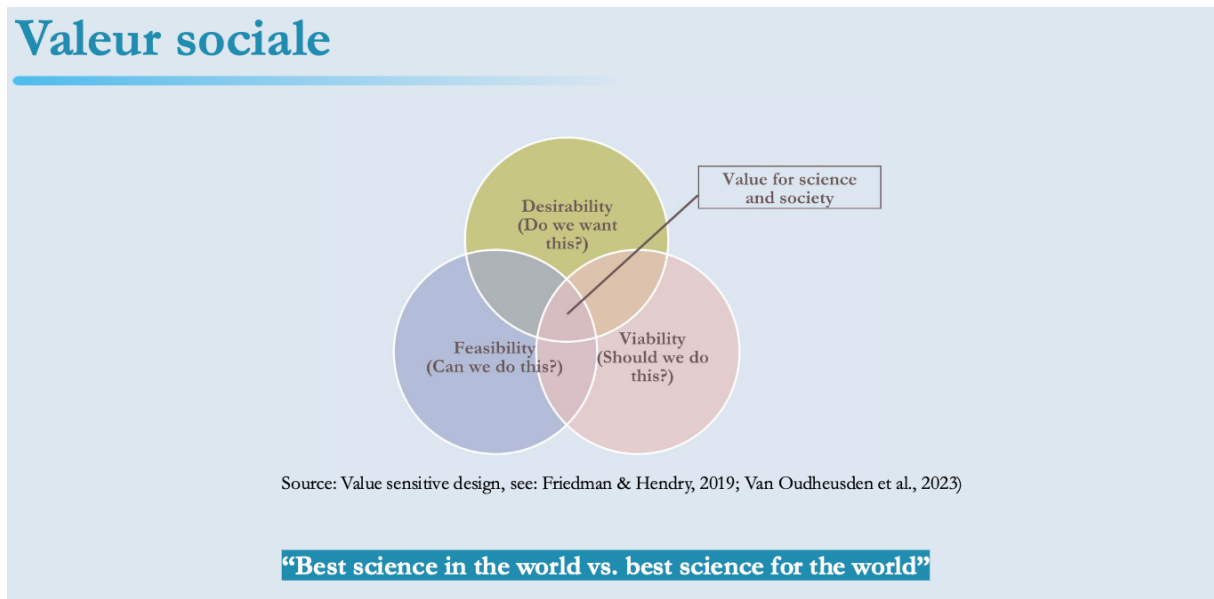


Figure 1: Representation of the social value of an innovation at the intersection of feasibility, desirability, and ethical viability.

In the case of digital twins, the question is whether we are developing digital twins because “we have the capability” of doing it or because “it addresses a genuine need in the field of healthcare.” During the expert group meetings, the questions proposed by Elisa Elhadj, a member of the group, were discussed (see Figure 2). The main lines of ethical reflection and questioning are listed below.

Reducing the person to their digital twin: an ontological shift

By its very principle, the digital twin performs a radical transformation: it transposes a human being, in all their complexity, into a digital space that is reducible, interpretable, and computable. Modeling a person does not simply consist of producing an abstract image; it involves relocating them into another space, where their uniqueness becomes a set of variables, their temporality is reduced to predictions, and their possibilities are reduced to probabilities. The digital twin is therefore not an extension of the person: it constitutes a calculable and necessarily simplified version of them.

- Are we developing digital twins because we have the technological capability to do so, or because they truly address a genuine need in the field of health?
- Who will have access to them? Who will benefit—and who will be left behind?
- Who profits from the data generated by digital twins, and how might these data be used for purposes other than patient care (e.g., malicious uses, commercialization of health data, surveillance, etc.)?
- Which bodies and lives are represented—or excluded—by digital twins? Who is missing in the data? At the same time, how can we use these models and simulations to make invisible bodies visible?
- How can we combine the strengths of humans (e.g., clinicians) and digital tools (e.g., digital twins)?




Illustration by James Steinberg for TIME
"Who Controls AI" (2023)

4

Elisa Elhadj | 02.04.2025

KU LEUVEN

Figure 2: Ethical questions regarding digital twins proposed by Elisa Elhadj.

While inevitable in itself, this simplification becomes problematic when it is forgotten or made invisible, or when the representation is confused with the people themselves. The risk then is the reification of the subject: the individual ceases to be a bearer of subjectivity, intention, and history, becoming instead an object for experimentation, simulation, and calculation.

In a medical context, this reduction can have harmful effects: a clinician might be tempted to rely on the model's projection rather than the unique experience of the patient; conversely, a patient might perceive themselves as "reduced" to their data, to the detriment of their lived experience, emotions, and capacity to defy prognoses.

Therefore, an ethics of the digital twin requires critical vigilance. It involves documenting modeling choices, making assumptions explicit, revealing margins of error, and preventing biases embedded in training data. More profoundly, it also requires emphasizing that the model must never obscure the person: the digital twin is only a tool that is useful and powerful, but incapable of capturing the narrative, relational, and existential richness of the subject. Ethics must ensure that this distance is maintained so that technology enriches medical practice without ever replacing the art of care.

Between Subject and Artifact: The Ambiguity of Ownership of Digital Twin Data

The ownership of digital twin data is a complex issue because it touches on a fundamental tension between personal identity and technical creation. On one hand, health data are inseparable from the person. They emanate from their body, biological history, and behaviors. Intuitively, one could say they constitute an "extension of self," a digital double that remains morally linked to the person. Considering them as "belonging" to a company or third party would reduce the individual to an exploitable object, which is ethically problematic. From a Kantian perspective, manipulating these data without consent would amount to instrumentalizing the person, treating them as a

means rather than as an end in themselves. From this angle, the data of a digital twin always belong (at least morally) to the person they represent, just like their image or voice.

However, there is another perspective. These data only take the form of a digital twin through technical work. They are collected, processed, used to create models, integrated into algorithms, and made visualizable. This process is not neutral: it involves methodological choices, interpretations, and intellectual creativity. The computer scientist, company, or laboratory creating the digital twin mobilizes expertise, tools, and know-how. From this point of view, the digital twin is not a mere extension of the person but also a technical “creation,” which could be subject to intellectual property rights. Moreover, because it is “a representation of me,” this digital twin potentially affects me, for better or worse, and should therefore be legally regulated in case of harm.

This leads to a dilemma: is the digital twin “mine,” because it contains my health data, history, and behaviors, or “theirs,” because it has been transformed through models, computations, and visualizations produced by others? Reducing its ownership to a single pole (i.e., either the person or the creator) ignores this dual nature. It therefore seems relevant to consider a principle of co-ownership: the individual retains an inalienable moral right over their data, while the computer scientist or company holds rights over the form and architecture of the model created.

Ethically, this co-ownership requires safeguards. The person must retain control over their data and its uses (principle of consent and autonomy). The company must ensure transparency and accountability in its technical choices. Finally, society must ensure that these models do not become instruments of alienation but remain tools for care and knowledge in the service of humanity.

In essence, asking “who owns digital twin data?” acknowledges that it is neither an exclusive “private good” nor an absolute “technical property.” It is a relational construction at the intersection of the human body and technology, whose governance must be conceived according to principles of sharing, responsibility, and respect for the individual.

The Weight of Prediction: Between Determinism and the Capacity to Act

The digital twin, designed to anticipate, introduces into the field of care a logic of the probable. It does not merely accompany; it calculates, projects, and traces trajectories. Yet care, as shown by the philosophy of care, resists calculation because it is primarily a space of attention to the unexpected, to surprise, to ambivalence. As in the famous aphorism attributed to Hippocrates regarding the art of medicine: *Vita brevis, ars longa, occasiō praeceps, experimentum periculōsum, iūdicium difficile*.

The patient is not the projection of their data; they are a being in becoming, always capable of diverging, disobeying models, or reinventing themselves. The temptation of the digital twin, however, is to freeze the person within a statistical norm, to assign their

future based on profiles derived from past data. The body is no longer primarily lived, but read through metrics, alerts, and risk scores. It becomes monitored rather than inhabited, observed rather than experienced.

This shift in the relationship to the body is profound: it induces a form of dispossession, where the individual becomes a spectator of themselves, dependent on what their digital double indicates, without always being able to influence this interpretation. The history of medicine reminds us that this logic is not entirely new: human genetics, with the belief in a “destiny inscribed in genes,” had already paved the way, until epigenetics highlighted its limits. But the digital twin, because it claims to be the equivalent of the individual themselves, radicalizes this determinism and makes its experience more immediate and constraining.

The danger lies not only in prediction but in normalization. The digital twin calculates deviations, compares models, also defines thresholds and margins. It locates the individual within a population, evaluating them against a statistical norm. Yet this norm is never neutral since it is constructed, localized, and normative. When a person is told, “you are at risk,” the risk is that this probability becomes an assignment: “you will get sick.” Prediction transforms into determinism, closing off the horizon of possibilities.

This shift raises a central ethical issue: preserving the subject’s capacity to act. Paul Ricoeur reminded us that medicine must protect the patient as a capable self, capable of narrative, choice, and reconfiguration. If the digital twin becomes prescriptive, it threatens this capacity. It risks substituting a possible future being open, unpredictable, reversible, with a probable future reduced to statistical inference from data that simplify the reality of the patient.

The patients are not some lines of code; they are a being capable of resistance, change, and defying expectations. Reducing care only to simulation, would risk neglecting this essential dimension and replacing deliberation with calculation, prudence with prediction, confusing the model’s projection with the singular trajectory of a life.

An ethics of the digital twin must therefore ensure that this openness remains alive: recognizing that the model illuminates but does not decide; that it describes but does not prescribe; that it calculates probabilities but must never abolish the freedom to act, nor the capacity of a subject to surprise their destiny.

Knowledge or Wisdom: The Disappearance of Judgment in Simulated Clinical Practice

Ancient philosophy, particularly Aristotle, distinguished between two forms of knowledge: *episteme* (i.e., demonstrative, rational, and generalizable knowledge), and *phronesis* which is practical wisdom, rooted in experience and attentive to the singularity of human situations. By its very logic, the digital twin embodies *episteme* as it calculates, infers, and predicts from data, situating the individual within a statistical norm. However, care, as it unfolds at the patient’s bedside, fundamentally belongs to *phronesis* as it

requires prudent discernment, the ability to interpret the unpredictable, and the recognition of what escapes the model and cannot be reduced to calculation.

Introducing algorithmic simulation as a tool for medical decision-making risks substituting predictive authority for clinical judgment-authority that is all the harder to challenge because it presents itself as objective, derived from a “science of data.” This reversal carries significant consequences. It delegitimizes doubt, limits discussion, and encourages caregivers and patients to conform to the model. The risk is that medical decisions become acts of adjustment to a prediction rather than shared deliberation. Simulation thus becomes prescriptive, not because it compels but because it imposes itself through apparent neutrality, assumed performance, and the scientific aura surrounding it.

Confusing the model’s projection with a person’s trajectory risks erasing the human dimension of care. It substitutes prediction for prudence, calculation for experience, and the algorithm for relational engagement. Aristotle’s distinction between *phronesis* (prudent judgment rooted in experience), and *epistēmê* (demonstrative and universal knowledge) illuminates this danger. Care is not merely a matter of knowledge and prediction; it requires the art of discernment, a meeting in which the patient’s subjectivity, narratives, and choices come into play.

The ethical challenge, then, is to preserve clinical judgment as a space for deliberation and prudence. Predictive models can provide guidance, but they must never decide on behalf of the clinician. The physician’s role is not to mechanically apply what the digital twin suggests but to integrate it within a broader reflection informed by their experience, knowledge, and the patient’s voice. Care is not reducible to a technical act: it is a practice of discernment, responsibility, and encounter, and a communicative practice in which the dignity of the person comes first.

Consent and Algorithmic Heteronomy: Shared Decision-Making in the Era of the Digital Twin

In medical tradition, informed consent is based on information that is understandable, proportionate, and freely accepted. Shared decision-making, on the other hand, requires a space for discussion between caregiver and patient, each contributing their knowledge and perspective to jointly shape a therapeutic direction.

The introduction of the digital twin into this relationship introduces a singular third party: an algorithmic simulation that claims to anticipate the future and whose “voice” appears difficult to contest. Faced with a model presented as “optimal,” how can a patient who is often technically inexperienced assert their preferences, values, or intuition? What freedom do they still have to refuse a trajectory that their digital twin indicates as preferable?

The risk is that consent becomes default adherence, driven by social pressure, the authority of technical expertise, or the inability to challenge predictions. Choice no longer originates with the patient but with the model that prefigures it. This is the emergence of

true algorithmic heteronomy: decisions are shaped by an opaque tool perceived as more competent than human judgment.

This situation undermines the very foundation of relational medicine. Dialogical authority (i.e., the authority exercised in exchange between caregiver and patient) risks being supplanted by technical authority, bolstered by the scientific aura of the data. As Hannah Arendt reminded us, the danger lies not only in the power of technology but in our critical disengagement from it. If patients passively adhere to simulations, and clinicians align unquestioningly with their recommendations, medicine loses its deliberative character and becomes a mechanical application of predictions.

In this context, informing the patient cannot be reduced to a one-time signed form. Consent must become evolving, reversible, and modular. It must incorporate genuine education on digital uses, rely on accessible resources, and recognize the cognitive and emotional vulnerability of patients confronted with tools they do not control. The ethical requirement is therefore twofold: to preserve individuals' freedom of choice while making models sufficiently explainable such that their predictions can be discussed and, if necessary, challenged.

Shared decision-making in the era of digital twins must not be a fiction. It requires keeping alive the dialogical dimension of care, where the algorithm illuminates without constraining, the model proposes without imposing, and clinical judgment (rooted in *phronesis*) retains its central place.

The Digital Twin and the Care Relationship: From Dyad to Quadrangulation

The arrival of the digital twin introduces a non-human third party into the caregiver-patient relationship. This new actor can enrich the medical encounter by providing objective and predictive support, but it can also weaken it by creating technological dependence. The question, then, is whether the digital twin will serve as a reflective ally, stimulating clinical deliberation, or, on the contrary, as a factor of human disengagement, replacing discussion with computational authority.

Traditionally, the care relationship was conceived as a dyad: the patient's vulnerability and the physician's expertise. With the rise of digital health, a third actor entered this exchange: the technological tool, whether an electronic health record or a decision-support algorithm. The digital twin goes even further: it complicates the relationship to the point of transforming this triangulation into a true quadrangulation. In addition to the patient and the physician, there is now the patient's digital twin, an algorithmic avatar that can be consulted, interpreted, or even opposed to the patient's own testimony while behind it the engineer or design team, custodians of the technical knowledge necessary to understand and adjust the model. Care is therefore no longer conducted solely between two or three visible actors but within a four-voice network, in which trust and responsibility are redistributed.

This shift profoundly redistributes trust. Previously, the patient entrusted their body and data to the physician, a representative of an identifiable institution. Now, they must also place their trust in a model whose logic they do not control and, indirectly, in those who designed it. Do they trust their physician, the digital twin that “predicts” a trajectory, or the computer scientist who sets the model’s parameters? This dispersion of trust weakens the clarity of the therapeutic relationship and necessitates the invention of new ethical agreements.

The status of the digital twin remains deeply ambiguous. As a tool, it should remain at the service of care. As a double, it tends to acquire symbolic autonomy, potentially weighing more heavily than the real person. As a third party, it becomes an implicit interlocutor in clinical discussions: “What does the twin say? What does it simulate?” In all cases, it shifts the center of gravity in the relationship. The patient’s voice, marked by subjectivity, risks losing legitimacy in the face of the computational rationality of a model.

The computer scientist appears as a paradoxical actor. Invisible in the clinical relationship, they nevertheless act behind the scenes as a co-therapist, defining and adjusting the tool that guides medical decisions. Yet the computer scientist may also be perceived as a new “clerk,” holding technical knowledge on which physicians and patients depend without being able to control it. This reinforces a new asymmetry: the patient depends on the physician, but the physician now depends on a tool interpreted by an external expert.

This shift from a dyadic to a quadrangular relationship requires rethinking classical medical ethics. In this context, transparency becomes essential to make the model’s functioning as intelligible as possible, while mediation must recognize the role of the designers without erasing the autonomy of the clinician and the patient. The patient’s lived experience must remain primary, as the twin is not the person, but only a simplified computational representation.

Notably, the digital twin does not abolish the care relationship, but instead it enhances it. Like the stethoscope in its time, it adds a technical mediation between the body and the medical gaze, within a principle of an “augmented” physician. However, unlike the stethoscope, the digital twin carries symbolic autonomy and the power of simulation, making it almost a full-fledged actor. It therefore compels us to redefine what caring means: not managing digital doubles, but continuing to encounter flesh-and-blood people, for whom technology must remain a tool, never a substitute.

Responsibility, Explainability, and Social Justice: Governance to Consider

The challenges raised by the deployment of digital twins extend far beyond the anthropological sphere: they also touch on the political, legal, and social dimensions of contemporary medicine. The first question concerns responsibility. Today, healthcare professionals are accountable for their decisions; however, what can we say about tomorrow? For instance, if an error arises from a simulation outputting an incorrect

prediction leading to a misinterpreted recommendation, who will be held responsible: the physician who relied on the model, the designer who developed it, the technology provider who operated it, or the healthcare institution that used it? This indeterminacy creates a dilution of responsibility, which would be problematic in a context where medical decisions must remain clearly assumed and attributable.

This uncertainty is compounded by an issue of transparency. The models underpinning digital twins often rely on complex algorithms whose internal logic remains opaque. This opacity limits the ability of both caregivers and patients to understand, critique, or discuss the results produced. Explainability is therefore not an optional add-on, but rather constitutes a democratic requirement. Without understanding, there is no trust; and without trust, no legitimacy can be established. Making models auditable, providing interpretable interfaces, and developing educational tools to support their use, are all minimal conditions for an ethical deployment.

Finally, the question of social justice is unavoidable. Who will have access to these technologies? Which patients will benefit from this “augmented care”? There is a significant risk of widening a new divide between connected, educated, insured patients who could benefit from personalized medicine through the digital twin, and others such as those in vulnerable situations, underrepresented in data, are left behind. An innovation that exacerbates disparities in access to care betrays the fundamental principles of medical ethics, with the risk of creating a two-tiered healthcare system. Political philosophy, as John Rawls reminds us, emphasizes that any technological improvement should primarily benefit the most vulnerable. Should the digital twin become a privilege reserved for a few, it ceases to be a tool for empowerment and instead becomes a marker of inequality.

Ultimately, the future of the digital twin does not lie solely in its technical performance, but in how we organize its conditions of use. Clarifying responsibility, ensuring transparency, and safeguarding equitable access are the three pillars of an ethical framework capable of sustaining the digital twin legitimacy and its role in the medicine of tomorrow.

Digital Twins and Ecological Footprint: Ethics Tested by Sustainability

The rise of digital twins in healthcare generates legitimate enthusiasm. By enabling the simulation of disease progression, the virtual testing of treatments, or the anticipation of complications, they pave the way for more personalized, predictive, and hopefully more effective medicine. Yet behind the appealing image of this cutting-edge tool lies a dimension often overlooked, which is its environmental impact.

A digital twin is not an abstract entity floating in cyberspace. It relies on a particularly heavy hardware and software infrastructure: biomedical sensors, connected devices to continuously collect data, data storage centers, and artificial intelligence algorithms requiring intensive computation to process this information. Each of these components has an environmental impact. Sensors depend on the extraction of rare metals that are

difficult to recycle; servers consume enormous amounts of energy and water for cooling; while training complex models can generate carbon emissions comparable to hundreds of transatlantic flights.

This reality poses a major ethical tension. Healthcare is, by definition, aimed at protecting and improving human life. However, what meaning does this have if the improvement of individual health comes at the cost of planetary health? Can we accept that a medical innovation, however sophisticated, indirectly contributes to worsening climate change, which already threatens global health? The question of sustainability becomes unavoidable.

Integrating environmental ethics into the design of digital twins requires a true paradigm shift. First, data collection must be limited to what is strictly necessary (principle of minimization), rather than succumbing to the temptation to “capture everything,” which multiplies unnecessary data flows. Next, massive and redundant duplication of databases should be avoided by favoring shared infrastructures among healthcare institutions and research organizations. Additionally, energy optimization of algorithms (eco-design) should become a design criterion alongside precision and robustness. Approaches such as “frugal artificial intelligence” already aim to reduce computational consumption; adopting these in the field of digital twins would signal responsible innovation.

Beyond the technical aspect, another requirement emerges: questioning the purpose of each simulation. Every calculation, update, and prediction carries an invisible environmental cost. It is therefore essential to further ask: is this simulation clinically relevant? Does it provide real added value for the patient or the healthcare system? Or is it merely an additional technological sophistication that is scientifically appealing but medically superfluous? A sustainability-oriented ethics requires prioritizing uses: favor those that prevent resource-intensive hospitalizations or invasive procedures, and critically evaluate the others. Otherwise, there is a risk of developing a paradoxical form of digital medicine: on one hand, presented as “innovative” and “sustainable” because it optimizes care pathways; on the other, reliant on energy-intensive infrastructures that contradict ecological transition goals. As numerous environmental ethicists have emphasized, the responsibility of care is not limited to the individual but rather extends to the community and the planet. A medical innovation that jeopardizes the very conditions of life on Earth cannot claim to be ethically neutral.

Ultimately, the future of digital twins will not be measured solely by their predictive power or clinical utility. It will also depend on our ability to manage their ecological footprint and embed these innovations within a logic of restraint. Care is inseparable from the ecosystem that makes it possible: protecting human health requires protecting the Earth on which it depends.

4. Conclusion

In the rapidly evolving landscape of healthcare, digital twins emerge as an innovation carrying the promise of more personalized, predictive, and effective medicine. Yet they also carry a risk: that of reducing the human being to a calculable abstraction, exacerbating inequalities, and undermining the fundamental principles of medical ethics.

This transformation raises fundamental questions: what does it mean to provide care when one can anticipate and simulate a patient's trajectory? What space remains for human judgment, relationships, and each individual's unique history? Ethics, far from being a mere overlay, becomes the very core of innovation as it guides how technology is conceived, developed, and integrated into care. It demands rigor, transparency, inclusion, and accountability, while defending a narrative and relational approach that recognizes the patient as a subject, not as a computational model.

The digital twin should not be considered a duplicate of the person, but as a reflective tool in the service of human relationships. This entails informed and dynamic consent, auditable and explainable models, and constant attention to equity, societal impact, and environmental footprint.

Yet the debate remains open and evolving: the ethical issues surrounding digital twins are never fixed. Every technological advance, new use, and feedback from experience calls for revisiting our frameworks and engaging in a multidisciplinary dialogue involving philosophers, legal experts, practitioners, engineers, and patients. The goal is not to slow down innovation, but to embed it within a humanistic logic where innovating goes hand in hand with caring, and where care remains primarily respectful, attentive, and shared. The challenge is not merely technical but profoundly human, aiming to make the digital twin a tool that illuminates care without ever replacing it.

5. Annexes

Presentation of Two French and European Projects on Digital Twins in Healthcare

EDITH

The European program EDITH (European Virtual Human Twin) aims to bring together an entire ecosystem around the digital twin. It serves as a coordination and support mission to pool the work of the ecosystem involved in the development of digital twins, assess legal barriers, address intellectual property issues, and more.

Since late 2022, EDITH and its partners have been fostering a collective dynamic through regular meetings and sharing the latest information on digital twins. To address all the issues, several working groups have been created: studying existing digital twins in Europe, mapping all stakeholders, initiatives, and obstacles, building a roadmap towards a European digital twin, and establishing a reference framework and simulation platform based on specific use cases.

<https://www.edith-csa.eu>

<https://www.inria.fr/fr/sante-partenariats-europeens-jumeau-numerique>

MEDITWIN

The MEDITWIN project brings together seven IHU (hospital-university institutes), Nantes University Hospital, Inria, several startups, and Dassault Systèmes. It aims to develop personalized virtual twins—organs, metabolism, cancer tumors—to improve patient diagnosis, monitoring, and treatment by allowing physicians to simulate future scenarios.

Seven new medical practices will be created in neurology, cardiology, and oncology, giving rise to seven “virtual health products” deployed on a sovereign industrial cloud platform. MEDITWIN will promote the industrialization, clinical validation, and standardization of these technologies to make them accessible to as many people as possible. Best practices will be codified as virtualized experiences, forming a global quality reference and a learning platform for scientific progress.

The project leverages the recognized expertise of its partners: Dassault Systèmes, as the industrial leader; the seven IHUs for medical and scientific excellence; Nantes University Hospital via the Thorax Institute; Inria with eleven project teams; and startups inHEART, Codoc, Qairnel, and Neurometers. Experience gained from digital twins such as the Living Heart and Living Brain initiatives, or the PEPR Santé Numérique, will serve as a foundation for healthcare.

The benefits of virtual twins will be assessed at multiple levels: for medical teams, by improving efficiency and the quality of multidisciplinary decisions; for patients, by enhancing safety and the effectiveness of interventions; and for the healthcare system, by optimizing practices and standards.

The MEDITWIN initiative will run for five years, from 2024 to 2029, with funding supported by the state under the France 2030 framework.

<https://www.3ds.com/fr/newsroom/press-releases/meditwin-launch>

<https://www.3ds.com/fr/insights/corporate-reports/transforming-medicine-virtual-twins-precision-healthcare>

<https://ihuican.org/meditwin-lutilisation-du-jumeau-numerique-pour-developper-la-medecine-personnalisee-de-demain/>