The digital twin is not an organ model

Genealogy of a contraction, epistemic asymmetry, and the distinction between peri-interventional, indication, and predictive regimes

Introduction

The term *digital twin* has imposed itself on the French health agenda with remarkable intensity. The launch of the MEDITWIN consortium in December 2023, backed by France 2030 and led by Dassault Systèmes with several University Hospital Institutes, the Nantes University Hospital and Inria, has installed in public debate a now dominant representation: a three-dimensional simulated organ, animated by multi-physics equations, queryable before a surgical or specialised procedure.

This representation is scientifically solid, clinically promising, and industrially structuring.

It does not, on its own, suffice to define the concept it embodies.

The term *digital twin* does not canonically designate a mechanistic organ model. It designates a broader class of numerical representations coupled to a real referent: object, system, organ, patient, trajectory, cohort, territory, or care pathway. The implicit reduction of the concept to the simulated organ therefore constitutes a *late semantic contraction*, made possible by the demonstrative power of certain use cases. This contraction is understandable. It nevertheless produces theoretical, industrial, and political effects that it is no longer possible to leave implicit.

The observation is not new. In *Do Buzzwords Dream of Clearer Substance?,* I already noted the constitutive polysemy of the term: *digital* may cover IoT signals, knowledge graphs, mechanistic models, or artificial intelligence systems; *twin* may designate a static visualisation, a data-fed model, or a feedback loop. The systematic review by Negri, Fumagalli and Macchi published in 2017 already showed that no univocal definition had stabilised in the literature, and that the dominant terms of the corpus, notably *physical* and *product*, signalled a structural bias toward the representation of a tangible object. This bias, observable very early in the industrial literature, precedes by several years the debate visible today in healthcare.

The thesis of this note is threefold.

1. First, the polysemy of *digital twin* is not an ambiguity to be resolved by arbitrary restriction; it is a structure to be preserved through terminological discipline.
2. Second, the current contraction of the concept proceeds from an implicit transfer from one epistemic regime, that of heavily instrumented industry, to an empirical regime, that of healthcare, which shares neither the data properties, nor the validity criteria, nor the modelling conditions of its source domain.
3. Third, the semantic contraction is compounded by a temporal contraction. What public communication today names *predictive medicine* covers, in a significant share of the cases put forward, optimised peri-interventional support or specialised indication aid. These objects are useful. They are not equivalent to predictive and preventive medicine in the populational sense.

Conflating these regimes is not a mere imprecision. It is a *silent transfer of agenda*.

Conceptual dilution does not lie in broadening. It lies in restriction.

1. Genealogy of a contraction

The concept of *digital twin* was not born in medicine. It is traditionally attributed to the work of Michael Grieves on Product Lifecycle Management in the early 2000s. In this first industrial matrix, the digital twin designates an informational representation of a physical object coupled to its real referent, mobilised to simulate its behaviour, track its evolution, and optimise its lifecycle.

A formal definition that became canonical in the aerospace field was published in 2012 by Glaessgen and Stargel in the NASA context: *an integrated multi-physics, multi-scale, probabilistic simulation of a vehicle or system*. The concept is then clearly anchored in a world where the referent is physical, heavily instrumented, and operated in an environment governed by modellable laws.

The most useful formalisation for architectural analysis, however, remains the one proposed by Kritzinger and colleagues in 2018. It distinguishes three regimes:

1. The *digital model* is a numerical representation without automated connection to the real referent.
2. The *digital shadow* is a representation automatically fed by data from the referent, but without automated feedback toward it.
3. The *digital twin* in the strict sense implies a bidirectional coupling: the referent feeds the representation, and the representation can influence the referent.

This distinction is not an academic refinement. It conditions infrastructure, validation, governance, and the system's degree of autonomy. It also implies that a significant number of systems qualified as *digital twins* are, in reality, *models* or *shadows*. Kritzinger already noted this in the manufacturing literature. The problem has not disappeared in healthcare; it has simply become more sensitive there.

The concept's entry into medicine took place through industrial analogy. The Living Heart Project by Dassault Systèmes, initiated in 2014 with the FDA, installed an extremely visible prototype: an organ modelled mechanically, electrically, and physically, whose fidelity to the biological referent (a *generic* heart of a healthy male individual?) can be confronted with imaging and clinical measurements. The scientific, mediatic, and industrial success of this form progressively produced a synecdoche effect: the most spectacular part of the concept, the simulated organ, ended up designating the whole.

It is this synecdoche that structures a significant share of today's public debate.

2. An epistemic asymmetry

The contraction of the concept in healthcare does not stem only from a lexical drift. It proceeds from an implicit transfer of epistemic regime.

In the industrial sectors where the digital twin was structured (aeronautics, automotive, energy, discrete manufacturing), the referent is physical, geometric, heavily instrumented, and governed by explicitly modellable laws. The Industry 4.0 paradigm rests on massive and continuous data capture from distributed sensors, SCADA systems, and industrial IoT. The volume, regularity, and quality of these data enable the training and validation of strongly constrained models. A twin of a turbine, an aircraft, or a production line commonly disposes of high-frequency telemetry, over several usage cycles, in environments where the relevant variables are themselves instrumented.

Healthcare presents a radically different configuration.

Data there are fragmented across heterogeneous systems, produced in multiple clinical contexts, often incomplete, and characterised by limited volumes relative to the dimension of relevant variables. This is the classical regime of *High Dimension, Low Sample Size* (HDLSS). Measurements are collected at discrete moments, with observation biases that are difficult to control. The relevant determinants include behavioural, social, environmental, and longitudinal dimensions that are neither easily measurable, nor stationary, nor homogeneous across populations.

To this data asymmetry is added a computational asymmetry worth quantifying, because it silently structures the industrial perimeter of high-fidelity organ twins. The reference review published in December 2024 in *EP Europace* (Bhagirath, Strocchi, Bishop, Boyle, Plank, *From bits to bedside: entering the age of digital twins in cardiac electrophysiology*)

documents this point without ambiguity: with a discretisation of around 400 µm, already used in several recent clinical applications in VT stratification, VT ablation, or AF ablation, computational costs per patient are reported at the scale of several days, even on advanced HPC resources; doubling spatial resolution may multiply cost by a factor close to ten, consistent with the properties of reaction-diffusion solvers on fine tetrahedral meshes. This figure holds for reference bidomain models; reduced-order or differentiable surrogate approaches lower this cost, with a loss of fidelity whose evaluation remains an open construction site.

What this datum does not prove: that the mechanistic paradigm would be disqualified by its cost. It is not. What it does prove: that the generalisation of the paradigm to massive cohorts (a few tens of thousands of patients per year for a given indication, several hundreds of thousands at the scale of a secondary prevention policy) is not an incremental engineering problem, but an architectural one. A paradigm whose computational cost is measured in CPU-days per patient cannot, without substantial order reduction, be directly transposed to populational scale.

In this regime, the modelling exercise is structurally more uncertain, more dependent on hypotheses, and more sensitive to validation conditions than in a heavily instrumented industrial universe. This is not a methodological weakness. It is an *epistemic constraint*.

Directly transposing to healthcare a paradigm elaborated for abundantly instrumented physical systems therefore produces a framing effect. This framing privileges objects most compatible with available data and mastered models: a precisely imaged organ, in a peri-interventional context, on a patient already referred to a specialised pathway. It mechanically tends to render less visible the objects nevertheless central to the transformation of the healthcare system: longitudinal trajectories, territorial cohorts, care pathways, primary and secondary prevention.

This regime difference explains why the mechanistic paradigm works within MEDITWIN's perimeter, where medical imaging offers rich instrumentation on the organ, and why it cannot be extended without rearticulation to the families of twins that do not benefit from this instrumental fortune.

3. Restoring the taxonomy

A digital twin is not defined by its visual form. It is defined by a *quadruplet of independent dimensions*.

1. **The first dimension is the nature of the referent.** The referent may be an object, an organ, a physiological system, a longitudinal patient, a cohort, a territory, or a care pathway. These referents are not substitutable. A heart, a valve, a polypathological person, a population basin, and an oncology pathway do not pose the same questions.

2. **The second dimension is the coupling regime.** A system may be unconnected, unidirectionally fed, or inscribed in a feedback loop. In healthcare, this feedback is rarely a direct command on the biological referent. It is most often mediated by a human decision, a clinical recommendation, an alert, a programmable medical device, or an organisational adaptation. This mediation is not a detail. It conditions regulatory qualification, responsibility, traceability, and the level of human guarantee required.
3. **The third dimension is temporal scale.** A twin may be designed for a point intervention, for an individual trajectory over months or years, or for a populational dynamic over long horizons. These temporalities mobilise neither the same data architectures, nor the same performance expectations, nor the same regimes of promise.
4. **The fourth dimension is the validation regime.** A mechanistic organ twin is validated by physical and physiological fidelity. A patient trajectory twin is validated by predictive, counterfactual, and clinical utility. A territorial twin is validated by its capacity to improve a collective decision. A care pathway twin is validated by its fidelity to real-world data and by its capacity to explain ruptures of continuity.

These regimes are not hierarchisable in abstract terms of rigour. They correspond to *distinct finalities*.

Current confusion stems from the reduction of this quadruplet to a single dimension: the three-dimensional representation of an organ. This narrowing does not bring precision. It erases the other constitutive dimensions of the concept.

4. The temporal regime: peri-interventional, specialised indication, strong prediction

Temporal scale deserves separate treatment, because it silently redistributes public promises.

Three regimes are often conflated under the single word *prediction*.

1. **The first is prediction in the strong sense.** It assumes an open temporal horizon, a not-yet-declared event, a real intervention window, and the possibility of modifying the trajectory through a present decision. In the vocabulary of public health policy, this regime grounds *predictive and preventive medicine*: identify before declaration, intervene before the irreversible.
2. **The second is the specialised indication regime.** The patient is already identified, often already ill, sometimes already referred to an expert pathway, but

the intervention decision remains open. Should a defibrillator be implanted? Should an ablation be proposed? Which patient will truly benefit from a heavy procedure? This regime is not purely peri-interventional, because the procedure is not always decided. It is also not populational predictive, because the patient is already in a specialised trajectory. It is a *hybrid zone*, strategically important.

3. **The third is the peri-interventional regime.** The intervention is decided, or strongly framed, and the principal question becomes technical: how to execute the procedure with the best precision, the best dimensioning, the lowest risk, the best anatomical or electrophysiological anticipation? Temporality is then compressed: days, hours, sometimes minutes. The intervention window is not populational; it is operative or procedural.

To these three regimes is added *acute diagnosis*. Faced with chest pain in the emergency department, the question is neither predictive nor peri-interventional in the strict sense: it consists in classifying an event in progress, on a horizon of minutes. Here again, using the same word to designate all these uses produces a convenient confusion.

When the public communication of a digital twin consortium announces enabling *predictive medicine*, several objects may therefore be placed under the same term.

- The cardiometabolic strand, when it works on familial hypercholesterolaemia, metabolic risk, prevention, or longitudinal trajectories, indeed falls under strong prediction. But these objects then belong more naturally to the family of trajectory twins, even if they mobilise organ submodels.
- The peri-interventional strand (congenital surgery, TAVR planning, ablation, procedural simulation) belongs to another regime. It predicts in the sense that a solver predicts a deformation, a flow, a mechanical or electrophysiological risk conditional on a procedure. This is *computational prediction*. It is not necessarily predictive medicine in the preventive sense of the term.

Predicting the probable result of an imminent procedure is not the same operation as predicting a disease that has not yet declared.

The strategic consequence is clear-cut. The peri-interventional optimises a societal cost already engaged. Strong prediction aims to avoid that a societal cost be engaged. The specialised indication regime occupies an intermediate position: it may avoid unnecessary procedures or better select patients, but it intervenes already within an advanced clinical pathway.

A public policy that places these three regimes under the same banner without distinguishing them may finance essentially tertiary optimisation while believing it finances populational prevention.

5. Four non-competing families

Once the taxonomy is restored and the temporal regime disambiguated, four families of healthcare digital twins appear distinctly.

5.1 The organ twin

The organ twin has as referent an organ or a physiological system. Its coupling, in clinical use, is most often unidirectional or mediated: the simulation informs the decision, but does not act directly on the biological referent. Its dominant scale is peri-interventional or specialised. Its validation regime is mechanistic fidelity, evaluated by confrontation with imaging, physical measurements, and clinical outcomes. Its finality is specialised decision support, procedure optimisation, therapeutic dimensioning, training, and, in certain cases, risk stratification.

The cardiac case illustrates this family. The Living Heart Human Model by Dassault Systèmes is documented (Baillargeon et al., *European Journal of Mechanics A/Solids*, 2014) as a four-chamber finite-element model of a healthy reference human heart, capable of integrating mechanical, electrical, and fluidic dimensions, and calibrated to reproduce cardiac physiological behaviours under reference conditions. It constitutes a generic multi-physics model, personalisable through several regimes.

The first regime proceeds by geometric adaptation of a reference model: morphing, scaling, parameterisation, manual adjustment. The second proceeds by patient-specific reconstruction from imaging: segmentation of cardiac CT or MRI, late gadolinium enhancement MRI to map fibrosis, meshing, then adjustment of electrophysiological parameters to reproduce an observed behaviour. The literature now distinguishes more clearly these two strata: *anatomical twinning* and *functional twinning*. The second is more demanding, more costly, and more dependent on the quality of patient data.

Data feeding does not obey the calendar of an industrial sensor. The dominant mode is pre-procedural: static imaging serves to reconstruct a twin used to prepare or guide a decision. Per-procedural feeding exists in certain cases in the form of registration, augmented visualisation, or substitute submodels, but the high-fidelity multi-physics solver does not necessarily function in the real-time loop. Post-procedural longitudinal feeding already belongs to another regime, closer to the trajectory twin.

Feedback toward the referent must be formulated with rigour. The myocardium is not an addressable actuator. The heart does not receive a command from its twin. Three limit cases exist: programmable implantable devices, circulatory assistance systems, and certain guided procedures such as ablation. But even in these cases, the twin most often acts on a human decision or on an intermediate device. *It acts on the hand that acts on the heart*.

In Kritzinger's strict sense, a significant share of the clinical objects currently qualified as organ twins therefore fall under the *digital shadow* or the *mediated patient-specific model* rather than the fully bidirectional *digital twin*. This terminological precision is not a coquetry. It avoids promising a control loop where decision support is delivered.

Within its proper perimeter, the organ twin addresses clinically major questions: scar-related ventricular tachycardia ablation, persistent atrial fibrillation ablation, TAVR (Transcatheter Aortic Valve Replacement) or TMVR (Transcatheter Mitral Valve Replacement) planning, complex congenital cardiac surgery, virtual testing of therapeutic responses, post-infarction rhythm risk stratification. These questions share a common point: the decision is strongly structured by patient geometry and by modellable physiology.

The most advanced terrain in prospective validation today is scar-related post-infarction VT. Hwang and colleagues, in *Circulation: Arrhythmia and Electrophysiology* (2024), created digital twins of 18 patients from pre-procedural LGE-MRI, simulated re-entrant VT circuits through rapid pacing, derived virtual ablation targets, then confronted these targets with invasive ablation. The twins identified critical VT sites with a sensitivity of 81.3%, a specificity of 83.8%, and a negative predictive value of 98.8%. These figures do not establish that the twin replaces invasive mapping. They establish that, on a limited but prospective cohort, a model reconstructed from standard imaging can non-invasively designate targets whose agreement with critical sites mapped in the lab is high. The validity perimeter is explicit: 18 patients, ischaemic scar substrate, LGE-MRI of interventional quality. The result is actionable within this perimeter.

Which does not mean that all major cardiovascular questions fall under the organ twin. Predicting an NSTEMI, for example, brings into play plaque rupture or erosion, inflammation, thrombosis, vascular state, biological factors, and risk trajectory. It is not primarily a question of myocardial geometry. The relevant objects exist elsewhere: dedicated coronary models such as CT-FFR, risk scores, longitudinal models, EHR data, biological and behavioural trajectories. The simulated organ does not exhaust the clinic.

MEDITWIN constitutes in France the most visible and most structured expression of this family. There is no need to diminish its importance. One must only avoid that it absorb the definition of all the others.

5.2 The patient trajectory twin

The trajectory twin has as referent a patient in his or her longitudinal dimension. It aggregates clinical events, biomarkers, treatments, behaviours, exposures, follow-up data, social and environmental context when these data are available. Its scale is that of months or years. Its validation regime is the capacity to produce information actionable for anticipating, preventing, adjusting, or comparing trajectories.

This is the regime of longitudinal probabilistic models, counterfactual architectures, synthetic comparator arms, clinically-informed models, and more broadly of systems that treat the patient not as an instantaneous geometry, but as an *evolving trajectory*.

This regime is that of predictive medicine in the strong sense. It does not only ask: what will happen if this procedure is performed? It asks: what trajectory is this patient at risk of following, and what present decision may modify this trajectory?

The difference with the organ twin is not a difference of scientific nobility. It is a difference of referent, of temporality, of data, and of validation. A trajectory twin may be less spectacular visually than a 3D organ. It may also be more directly transformative for prevention, chronic follow-up, and territorial medicine. As often, what is least telegenic is sometimes the most structuring. A banal observation, but one that an era that readily confuses a proof with a beautiful image would do well to remember.

5.3 The territorial twin

The territorial twin has as referent a population basin. It mobilises epidemiological, environmental, demographic, organisational, and medico-social determinants. Its coupling is typically unidirectional or decisional: it feeds arbitrations of allocation, prevention, planning, or coordination. Its horizon is that of months, years, sometimes decades.

Its validation regime is not anatomical fidelity. It is *decisional fidelity*: does the system allow for better resource allocation, detection of weak signals, prioritisation of a public intervention, reduction of an access rupture, anticipation of hospital tension, or targeting of a prevention campaign?

The territorial twin is closer to a populational decision support system than to an organ simulator. Naming it a digital twin is not an abusive extension. It is a return to the initial idea: represent a complex real referent, couple it to data, simulate possible evolutions, and improve a decision.

5.4 The care pathway twin

The care pathway twin has as referent a pathology, a care episode, or a medico-organisational sequence traversed by a heterogeneous population of patients. Its objective is not to simulate an organ, but to understand how patients circulate, drop out, return, change therapeutic lines, undergo delays, ruptures, or losses to follow-up.

Its validation regime rests on fidelity to real-world data and on the explanatory power of observed variations. It allows the identification of frictions, delays, lost-chance situations, organisational dependencies, and territorial gaps.

It is particularly relevant in chronic pathologies, oncology, rare diseases, mental health, geriatrics, post-acute pathways, and complex care management. Here again, the referent is not an anatomical structure. It is *an organised collective trajectory*.

These four families are not competing. They do not cover the same needs, are not validated by the same criteria, and do not have the vocation to live within the same industrial architectures.

A coherent national strategy articulates them. A confused strategy lets one of them silently absorb the other three.

6. What the contraction renders invisible

The implicit reduction of the digital twin to the organ twin alone produces a direct consequence: it renders less visible the objects that potentially carry the broadest transformation of the healthcare system.

Take a first scenario.

An elderly, polypathological patient living in a nursing home or in a medically underserved area presents heart failure with preserved ejection fraction, diabetes, hypertension, atrial fibrillation, and chronic bronchopathy. No intervention is planned. The clinical question is not to plan a procedure, but to anticipate a decompensation. What changes, week after week, is not the myocardial geometry. It is volume status, respiratory rate, sleep quality, adherence, intercurrent infection, renal function, weight, care environment.

In this context, a high-fidelity organ twin has limited utility (especially if its validity domain has not been defined). The right object is a compact trajectory twin: a cardio-renal compartmental model coupled to home telemetry data, learning on comparable cohorts, early alert, therapeutic adaptation before hospitalisation.

The order of magnitude of accessible performance is now documented. Pavon and colleagues, in a non-randomised prospective study on 150 post-discharge HF patients over one year, trained an LSTM model exploiting telemonitored vital signs and patient profile to dynamically predict 30-day readmission risk. The model detects emerging readmissions with a sensitivity above 71%, a specificity above 75%, and an AUROC of around 80%, performance maintained above 78% AUROC even when reducing telemonitoring frequency to one day out of two. These figures do not prove that a trajectory model universally outperforms structured clinical follow-up; they prove that, at a defensible performance level, on a real clinical cohort, an object of the trajectory family exists, is validated, and is deployable on light infrastructure. What this literature does not resolve: the net effect on hospitalisations in routine practice, which depends as much on care organisation as on the predictive model itself. Recent meta-analyses on HF

telemonitoring remain contrasted on this point. The trajectory twin is an instrument; what it becomes clinically depends on the chain in which it is inserted.

The scale asymmetry with the organ twin is instructive. Where a high-fidelity electrophysiological solver mobilises several CPU-days per patient, a compact HF trajectory model executes in near-real time on light infrastructure, for cohorts potentially of several hundred thousand patients. The geometry of a heart does not change in seven days. The vital prognosis, sometimes, does.

A second scenario.

A post-infarction patient presents reduced ejection fraction and finds himself in the grey zone of an indication for an implantable cardioverter-defibrillator. Here, the organ twin may be extremely relevant. Scar geometry, electrophysiological substrate, and in silico inducibility may enrich the decision beyond ejection fraction alone. This is precisely what the Hwang cohort cited above documents: on 18 patients, sensitivity 81%, specificity 84%, negative predictive value 98%. This last metric deserves to be highlighted: an NPV of this order, if confirmed on larger cohorts, opens the way to an *informed de-indication* for patients whose twin does not generate clinically plausible VT inducibility.

The stake is not only scientific. It is industrial: can one move from a costly and specialised high-fidelity model to a version sufficiently robust, sufficiently fast, and sufficiently integrable to be used beyond a few expert centres? In France, several thousand patients per year fall under an ICD indication discussion, including a significant share in primary prevention; the perimeter of centres capable today of producing a high-fidelity electrophysiological twin covers only a fraction of this population. Reducing the gap between what the paradigm can prove scientifically and what it can industrialise is precisely the work of reduced-order models, differentiable surrogates, and hybrid architectures.

This scenario illustrates the specialised indication regime. The patient is already in a pathway. The procedure is not yet automatically decided. The twin may avoid an unnecessary intervention or better select patients who will benefit. It is neither pure peri-interventional nor populational prevention. It is precisely for this reason that a finer taxonomy is needed.

These two scenarios are not opposed to MEDITWIN. They trace another map.

The first shows a question for which no organ twin, whatever its fidelity, constitutes the principal object, because the question is not geometric.

The second shows a question for which the organ twin is conceptually relevant, but where the stake becomes order reduction, industrialisation, integration, and scaling.

This asymmetry is not intrinsic to the clinical value of the approaches, but to their relative industrial maturity and to their computational and data constraints. It may evolve.

In both cases, *architecture counts as much as the model*.

Populational preventive medicine operates further upstream.

Its question is not: how to optimise this procedure?

Its question is: who must be identified as at risk so that this procedure, this hospitalisation, or this lost chance does not occur?

This question is not resolved by the mechanistic modelling of an organ alone. It is resolved by the modelling of trajectories and populations, in other words by families of twins that do not have the same data, the same architectures, or the same regimes of proof.

Excellence in healthcare for all, everywhere on the territory, is not compatible with a definition of the digital twin that would implicitly assume a University Hospital Institute, advanced imaging, and a heavy industrial platform. These conditions of possibility are those of the high-fidelity organ twin. They are not those of the territorial twin, nor of many trajectory twins, nor of care pathway twins.

It is not a matter of opposing the families. It is a matter of recognising that the national strategy in digital health needs all four, and that this recognition supposes a vocabulary that does not pass one off as another.

7. Doctrinal articulation

This terminological thesis is not an isolated hypothesis. It inscribes itself within a more general position: in regulated healthcare, AI is not first an algorithm problem. It is a problem of architecture, of referent, of validation, and of decision.

The same reasoning holds for benchmarks, for agents, for clinical recommendation systems, for event-driven architectures, and for digital twins. A technical artefact can only be correctly evaluated if its regime of use is correctly named.

A peri-interventional organ twin does not become a populational predictive medicine instrument through mere maturity gains. It may move from specialised prototype to specialised clinical use; this passage remains *intra-regime*. To become a populational prevention tool, one would have to change referent, scale, data, validation, and deployment economics. This is no longer a promotion of the same artefact. It is the *constitution of another artefact*.

It is here that the distinction becomes operational. It prevents a success in one regime from being used as implicit proof in another.

A peri-interventional simulator may be excellent without constituting a prevention strategy. A trajectory model may be useful without producing a spectacular 3D anatomy.

A territorial twin may be decisive without resembling a medical device. A care pathway twin may transform an organisation without ever simulating an organ.

The common point is not form. The common point is *the disciplined coupling between a real referent, a numerical representation, a decisional finality, and a validation regime*.

8. Objections

First objection: broadening the term dilutes it.

The opposite. We do not dilute, we restore. The taxonomy proposed here is closer to the historical formulations of Grieves, to the aerospace definition by Glaessgen and Stargel, to Negri's observations, and to Kritzinger's categorisation than the recent contraction of the concept to the simulated organ alone. Dilution does not lie in restoring the field. It lies in the capture of a concept by its most visible instance.

Second objection: this distinction changes nothing in practice.

It changes, on the contrary, the criteria of financing, the evaluation grids, the validation regimes, and the industrial trajectories. A trajectory twin evaluated with the criteria of an organ twin structurally fails, not because it is mediocre, but because it is judged within a frame that is not its own. The reverse is also true.

Third objection: this note is a disguised critique of MEDITWIN.

No. It does not bear on the scientific quality of the consortium nor on the relevance of its use cases. It bears on the conceptual contraction that may install itself around it. Recognising the excellence of a project within its validity perimeter does not preclude recalling that this perimeter does not exhaust the conceptual category in which it inscribes itself. The two operations are jointly necessary.

The MEDITWIN organ twin is what it says it is, at a very high level. What it is not, is the single object under which the entire public promise of digital twins in healthcare can hold.

Fourth objection: the distinction between strong prediction, specialised indication, and peri-interventional is excessive.

It is on the contrary minimal. A solver predicts a deformation. A cardiovascular risk model predicts an event at ten years. An indication model predicts the expected benefit of a procedure in an already-identified patient. The three use the same verb. They do not belong to the same regime of decision.

The first inscribes itself within a procedural logic. The second within a preventive logic. The third within a specialised arbitration logic. They are not financed in the same way, are not validated with the same criteria, do not industrialise within the same architectures, and do not produce the same public value.

Conflating these regimes leads to judging an object with the promises of another. This produces structurally false arbitrations that nevertheless preserve an appearance of procedural rationality, that is, the most difficult form of framing defect to correct, because it presents itself as its opposite.

9. Limits

This note does not resolve the question of clinical validation proper to each family. Nor does it treat in detail the governance of data, which becomes central as soon as one leaves the strongly delimited tertiary frame. It does not prejudge, finally, the relative industrial maturity of the different families. The organ twin is today more advanced than the territorial twin, both in clinical validation and in industrial integration.

It does not exhaust either the questions of internal architecture: relevant modelling scale, hybrid models, multi-agent composition, operational condition maintenance, regulatory qualification, medico-economic validation, auditability, human supervision, reversibility. These questions fall under a complementary frame.

The rapid evolution of reduced-order models (hybrid eikonal, physics-informed networks, differentiable surrogates, compact models coupled to longitudinal data) could in the coming years reduce part of the computational asymmetry between high-fidelity organ twins and lighter twins. This convergence will not annul the taxonomic distinction. It will displace some of its industrial consequences.

The most plausible scenario is not the substitution of one family by another. It is *cohabitation*: high-fidelity organ twin in the expert tertiary sector, compact organ twin in certain specialised indications, trajectory twin in chronic follow-up and prevention, territorial twin in populational steering, care pathway twin in organisational optimisation.

Conclusion

The digital twin is not a simulated organ. It is a class of representations coupled to a real referent, defined by four independent dimensions: referent, coupling, temporal scale, validation regime.

In healthcare, this class declines into at least four families: organ twin, patient trajectory twin, territorial twin, care pathway twin. To this taxonomy is added a decisive temporal distinction between peri-interventional, specialised indication, and strong prediction.

This clarification is not a terminological correction of pure form. It conditions the legibility of complementarities between projects, the relevance of financing criteria, the rigour of evaluation grids, and the very possibility that populational preventive medicine may find its place within the digital health strategy beyond the circle of tertiary excellences.

Before asking how to fabricate a digital twin, one must determine of what type of reality it is a question, for what decision, at what temporal scale, with what validation regime, and under what form of human mediation. The choices of architecture, data, computation, deployment, and regulation follow from this.

In the absence of this distinction, arbitrations may remain rational in form while being structurally ill-posed.

The debate does not oppose approaches. It bears on the capacity not to confuse a visible instance of the concept with the concept itself, and on the manner of articulating technological excellence and effective transformation of the healthcare system.

The high-fidelity organ twin holds its promises within its perimeter. The public promise of populational predictive and preventive medicine is of another order. A national strategy that rigorously distinguishes them serves both. A strategy that conflates them compromises both.

Main references

Baillargeon, B., Rebelo, N., Fox, D. D., Taylor, R. L., & Kuhl, E. (2014). The Living Heart Project: A robust and integrative simulator for human heart function. *European Journal of Mechanics A/Solids*, 48, 38-47.

Bhagirath, P., Strocchi, M., Bishop, M. J., Boyle, P. M., & Plank, G. (2024). From bits to bedside: entering the age of digital twins in cardiac electrophysiology. *EP Europace*, 26(12), euae295.

Glaessgen, E., & Stargel, D. (2012). The Digital Twin Paradigm for Future NASA and U.S. Air Force Vehicles. *53rd AIAA Structures, Structural Dynamics, and Materials Conference*.

Grieves, M. (2002-2003). Product Lifecycle Management, lectures at the University of Michigan.

Hwang, M., Han, S., Park, M.-S., Yu, H. T., Kim, T.-H., et al. (2024). Heart Digital Twins Predict Features of Invasive Reentrant Circuits and Ablation Lesions in Scar-Dependent Ventricular Tachycardia. *Circulation: Arrhythmia and Electrophysiology*.

Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital Twin in Manufacturing: A Categorical Literature Review and Classification. *IFAC PapersOnLine*, 51(11), 1016-1022.

Negri, E., Fumagalli, L., & Macchi, M. (2017). A Review of the Roles of Digital Twin in CPS-based Production Systems. *Procedia Manufacturing*, 11, 939-948.

Sakata, K., Bradley, R. P., Prakosa, A., Yamamoto, C. A. P., Ali, S. Y., et al. (2024). Assessing the arrhythmogenic propensity of fibrotic substrate using digital twins to inform a mechanisms-based atrial fibrillation ablation strategy. *Nature Cardiovascular Research*.

Trayanova, N. A., & Prakosa, A. (2024). Up Digital and Personal: How Heart Digital Twins Can Transform Heart Patient Care. *Heart Rhythm O2*.

See also in this corpus

Vetillard, J. (2025). *Do Buzzwords Dream of Clearer Substance? The Journey of Digital Twins*. LinkedIn.

Vetillard, J. (2026). *Healthcare AI is not only an algorithm problem. It is above all an architecture problem*. Twingital Institute.