Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

WHEN DOES A PHARMACEUTICAL MODEL BECOME A DIGITAL TWIN? AN UMBRELLA REVIEW ACROSS PHARMACOLOGY, PHARMACEUTICS, AND DRUG DELIVERY

Thabo Nkosi[1]*, Lerato Molefe[1], Sipho Dlamini[2], Ayanda Mokoena[3], Kabelo Ndlovu[4]

1. ***Department of Digital Twin Definition and Umbrella Review, Faculty of Pharmacy, University of Cape Town, Cape Town, South Africa.***
2. ***Department of Pharmaceutical Digital Twins in Drug Delivery, Faculty of Pharmacy, University of the Witwatersrand, Johannesburg, South Africa.***
3. ***Department of Pharmacology and Pharmaceutics Twins, Faculty of Pharmaceutical Sciences, University of Pretoria, Pretoria, South Africa.***
4. ***Department of Model-to-Twin Translation, Faculty of Pharmacy, Stellenbosch University, Stellenbosch, South Africa.***

ARTICLE INFO

Received:
03 April 2026
Received in revised form:
08 August 2026
Accepted:
10 August 2026
Available online:
28 August 2026

Keywords: Digital twin, Computational pharmacology, Pharmaceutics, Drug delivery modeling, Virtual patients, Model-informed precision dosing

ABSTRACT

Digital-twin terminology is increasingly applied to computational models used in drug discovery, pharmacology, formulation development, drug delivery, precision dosing, clinical trials, and pharmaceutical manufacturing. However, patient-specific simulation, mechanistic modeling, adaptive prediction, virtual populations, and process monitoring are frequently described as digital twins without consistent requirements for real-entity linkage, updating, uncertainty assessment, or decision use. This umbrella review critically evaluates review-level evidence across pharmaceutical domains to determine when a model warrants the digital-twin designation. The synthesis distinguishes digital models, adaptive models, virtual patients, digital shadows, and decision-coupled twins according to counterpart identity, data connection, update logic, predictive purpose, bidirectional information flow, validation, and governance. Review selection, methodological quality, primary-study overlap, consistency, and context of use are treated as determinants of evidentiary strength rather than assuming that multiple reviews constitute independent confirmation. The evidence indicates that pharmacometric, systems-pharmacology, formulation, delivery, and manufacturing models provide important twin-enabling capabilities, but most reported applications demonstrate only subsets of the functions implied by a mature digital twin. Patient specificity, mechanistic sophistication, or benchmark performance alone does not establish twin status. The central contribution is a proposed operational boundary in which a pharmaceutical digital twin must represent a specified counterpart, receive state-relevant updates, generate uncertainty-aware prospective predictions, and connect those predictions to an auditable and bounded decision pathway. This framework is conceptual and requires domain-specific validation. Its purpose is not to promote universal adoption of digital twins, but to improve terminology, evidence appraisal, comparative evaluation, and translational judgment across computational pharmaceutical science.

To Cite This Article: Nkosi T, Molefe L, Dlamini S, Mokoena A, Ndlovu K. When Does a Pharmaceutical Model Become a Digital Twin? An Umbrella Review across Pharmacology, Pharmaceutics, and Drug Delivery. Pharmacophore. 2026;17(4):135-45. https://doi.org/10.51847/HXSioqastj

Introduction

Digital-twin concepts have entered biomedical scholarship through several partially overlapping traditions. Systems-medicine accounts describe a twin as an individualized representation that can be progressively refined using information from the represented patient, whereas precision-medicine roadmaps emphasize multiscale simulation, longitudinal data integration, and prospective support for therapeutic decisions. Drug-development literature additionally employs the terms *virtual patient*, *causal disease model*, and *in silico trial*, although these constructs may represent simulated populations or hypothetical disease trajectories without maintaining a persistent connection to one evolving real-world counterpart [1-3]. The resulting conceptual

Corresponding Author: Thabo Nkosi; Department of Digital Twin Definition and Umbrella Review, Faculty of Pharmacy, University of Cape Town, Cape Town, South Africa. E-mail: thabo.nkosi@uct.ac.za

overlap is scientifically consequential because the label assigned to a model influences how its capabilities, validation burden, and potential decision role are interpreted.

The recent convergence of generative artificial intelligence, mechanistic modeling, biomedical data integration, and trial simulation has widened the range of systems described as digital twins. Generative and predictive approaches may create plausible molecular, patient, or disease-state representations and may help explore counterfactual scenarios, but these capabilities remain conditional on data suitability, calibration, external validity, and the intended context of use [4]. A model that generates realistic outputs may still lack a clearly specified counterpart, a validated update mechanism, or an auditable route from prediction to pharmaceutical action. Consequently, technical sophistication should not be treated as synonymous with twin maturity.

Digital-twin language now extends across the pharmaceutical lifecycle, including target and molecule evaluation, translational pharmacology, clinical development, individualized dosing, formulation and drug-delivery design, and manufacturing control [5]. These applications differ substantially in what is being represented, how frequently its state can be observed, which decisions are being supported, and what harms may follow from incorrect predictions. A patient twin, for example, may require longitudinal physiological and treatment data, whereas a manufacturing twin may receive frequent process measurements from a defined production system. Treating these implementations as interchangeable obscures differences in observability, update cadence, validation, uncertainty, and governance.

This umbrella review therefore asks when a pharmaceutical model becomes a digital twin rather than merely when digital-twin terminology has been used. Its purpose is to synthesize review-level evidence across pharmacology, pharmaceutics, and drug delivery; distinguish demonstrated capability from plausible utility and routine readiness; and propose a transparent operational boundary for terminology and evaluation. The article does not assume that every useful computational model should become a digital twin. Instead, it examines which additional properties are required before a static, adaptive, mechanistic, data-driven, hybrid, or patient-specific model can reasonably be regarded as a connected and decision-relevant representation of a pharmaceutical entity or process.

Umbrella-Review Questions and Operational Definition of a Digital Twin

Cross-domain reviews consistently indicate that a digital twin is more than a digital representation. Recurrent components include a specified physical or biological counterpart, a virtual representation whose states correspond to that counterpart, information exchange between the represented system and the model, and model-enabled interpretation or prediction [6-8]. Nevertheless, the literature remains inconsistent concerning whether updating must be continuous, whether information flow must be fully automated, whether the return pathway must directly control the counterpart, and how model fidelity should be judged. These differences make a binary definition based only on terminology unreliable.

Healthcare-focused reviews further show that applications described as digital twins range from conceptual architectures and retrospective simulations to individualized monitoring systems and decision-support environments [9]. Some models are personalized only at initialization, whereas others are repeatedly updated as measurements become available. Some represent a real patient, device, batch, or process, while others represent a synthetic patient or a population distribution. The umbrella review therefore separates *entity specificity* from *patient specificity*: a model may be individualized without maintaining a persistent identity relationship, and a process twin may satisfy entity-specific requirements even though it does not represent a patient.

For this review, a pharmaceutical digital twin is operationally defined as a fit-for-purpose virtual representation of a specified real pharmaceutical entity or process that receives state-relevant information through a defined connection, updates its internal representation using documented logic, produces prospective and uncertainty-aware outputs within a declared context of use, and connects those outputs to an auditable scientific, development, clinical, or manufacturing decision pathway. The pathway may be human-mediated and does not require autonomous control. This component-based definition is consistent with healthcare frameworks that distinguish the represented entity, model, data connection, update process, service or decision function, and governance environment [10]. It is proposed as an interpretive framework for this umbrella review rather than as an empirically validated classification or regulatory standard.

The review questions follow directly from this definition. The synthesis examines how twin status should be operationalized; how search quality, review selection, appraisal, and evidence overlap affect confidence; how terminology differs across pharmacology, pharmaceutics, delivery, and manufacturing; what mechanistic, data-driven, hybrid, and patient-specific approaches can contribute; when updating and bidirectional connection become meaningful rather than nominal; how prediction is separated from experimental confirmation and decision value; what the review-level evidence supports concerning maturity; and which minimum conditions should govern use of the digital-twin label. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop umbrella-review questions and operational definition of a digital twin within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Umbrella-review questions and operational definition of a digital twin

Evidence domain	Eligible evidence or method	Reported capability or claim	Strength of support	Reproducibility or bias concern	Unresolved gap

Definition and nomenclature	Systematic, scoping, conceptual, and state-of-the-art reviews that explicitly define digital twins, digital shadows, virtual patients, or connected models	Digital twins generally combine a virtual representation with a specified counterpart and some form of information exchange	Convergent at the component level but inconsistent in terminology	Reviews use heterogeneous eligibility criteria and may transfer industrial definitions into healthcare without sufficient adaptation	A pharmaceutical definition that distinguishes necessary components from optional technological features
Counterpart identity	Reviews of patient-specific models, process models, virtual patients, and healthcare architectures	A twin should represent a traceable patient, product, formulation, batch, device, process, or other bounded entity	Moderate review-level support	"Personalized," "individualized," and "patient-specific" are sometimes used without persistent identity linkage	Criteria for establishing when initialization to individual data becomes an enduring twin relationship
Updating and synchronization	Reviews addressing data assimilation, monitoring, adaptive modeling, and model-state correspondence	State-relevant updating differentiates a connected representation from a static model	Conceptually strong; operational evidence remains heterogeneous	Update frequency, latency, missing-data handling, and post-update validation are often incompletely reported	Context-specific standards for sufficient update cadence and update reliability
Prediction and uncertainty	Reviews of mechanistic, data-driven, hybrid, and model-informed decision systems	Digital twins may forecast future states, compare interventions, or estimate process responses	Strong evidence for modeling capability; weaker evidence for twin-specific prospective utility	Retrospective fit, data leakage, inadequate calibration, and untested domain shift may inflate confidence	Prospective, externally validated, and uncertainty-calibrated prediction within declared contexts of use
Bidirectional connection and decision use	Architectural, implementation, clinical, and manufacturing reviews	Twin outputs may inform human decisions or automated control and may subsequently receive feedback from the resulting state	Moderate and highly context-dependent	Decision pathways, responsibility, override rules, and consequences of error may be omitted	Evidence that model-informed actions improve bounded pharmaceutical decisions relative to an appropriate comparator
Validation and credibility	Review-level guidance on verification, validation, uncertainty quantification, and context of use	Credibility depends on intended use, decision consequence, and supporting evidence rather than model complexity alone	Strong methodological support	Validation terminology is inconsistent across disciplines and may not address updating or lifecycle change	Integrated credibility expectations for connected and evolving pharmaceutical models
Umbrella-review quality and overlap	Reporting guidelines, review-appraisal frameworks, citation matrices, and overlap methods	Review conclusions must be interpreted according to methodological quality, scope, primary-study duplication, and consistency	Strong methodological support	Several reviews may repeatedly synthesize the same underlying studies, creating apparent convergence	Transparent overlap-sensitive synthesis without treating review counts as independent evidence
Translational readiness	Meta-reviews, implementation reviews, regulatory-facing model guidance, and domain-specific roadmaps	Some applications demonstrate twin-enabling functions, but routine readiness requires integration, governance, monitoring, and decision-value evidence	Emerging and uneven	Promising examples may be generalized beyond their validation setting	Domain-specific thresholds separating conceptual plausibility, validated use, and operational readiness

Table note: The domains and boundaries are an original synthesis for this umbrella review. They are not mutually exclusive categories, a numerical scoring system, a validated maturity instrument, or a regulatory determination.

Search, Review Selection, Quality Appraisal, and Overlap Assessment

The review was structured as an umbrella synthesis in which the eligible evidence base consisted principally of systematic reviews, scoping reviews, meta-reviews, state-of-the-art reviews, consensus reviews, and directly relevant conceptual or methodological syntheses. Structured bibliographic searching was supplemented by targeted journal and DOI verification. Eligibility required peer-reviewed journal publication, direct relevance to at least one prespecified review question, sufficient methodological or conceptual detail to support a precise manuscript claim, and an identifiable context spanning pharmacology, pharmaceutics, drug delivery, patient modeling, clinical development, or pharmaceutical manufacturing. Review reporting and appraisal were organized using established guidance for critical assessment, transparent overview reporting, and reproducible description of search and selection decisions [11-13]. Articles that merely used digital-twin terminology without defining the represented entity, connection, update mechanism, evaluation context, or decision purpose were not treated as strong evidence for twin maturity.

Quality appraisal was domain based rather than reduced to an invented aggregate score. Systematic reviews were examined for protocol transparency, search reproducibility, eligibility decisions, duplicate study handling, source appraisal, synthesis methods, limitation reporting, and traceability between evidence and conclusions. Scoping and conceptual reviews were assessed according to the clarity of their definitions, evidence-selection logic, coverage of competing interpretations, acknowledgment of uncertainty, and separation of proposed frameworks from demonstrated applications. Regulatory-facing and technical reviews were additionally examined for context-of-use specification, verification, validation, uncertainty treatment, and transferability. Because an overall corrected covered area can conceal extensive duplication between particular review pairs, a single overlap indicator was not regarded as sufficient to determine evidentiary independence [14].

Primary-study overlap was therefore treated as both a methodological and interpretive issue. Review citation matrices were used to identify where the same foundational applications, validation studies, or conceptual sources appeared across multiple reviews. Where reporting permitted, overall overlap assessment was complemented by pairwise inspection, review recency, methodological quality, scope, and the exact claim being supported. Reviews with overlapping evidence were not counted as independent confirmations merely because their conclusions were similar. Instead, duplicated evidence was weighted according to directness, quality, and relevance, while divergent conclusions were examined for differences in entity type, model architecture, validation setting, and intended decision. This approach follows practical guidance that overlap should be explicitly mapped and managed rather than ignored or addressed by automatically retaining only one review [15].

Digital-Twin Concepts in Pharmacology, Pharmaceutics, and Delivery

Pharmacology provides some of the most mature foundations for pharmaceutical digital twins because physiologically based pharmacokinetic models already connect drug properties, biological parameters, dosing conditions, and predicted exposure within explicit physiological structures. Established qualification practices require documentation of model purpose, assumptions, parameter sources, verification, sensitivity, and comparison with relevant observations [16]. These practices make PBPK models credible components of a potential twin, particularly when individualized physiology or longitudinal drug measurements are incorporated. Nevertheless, a qualified PBPK model remains a model rather than a twin when it is applied offline to a representative population, calibrated once, or used without a persistent relationship to a specified patient, product, or development process. Qualification supports credibility within a context of use; it does not independently establish connectedness, updating, bidirectionality, or decision coupling.

Quantitative and systems pharmacology adds disease mechanisms, biological networks, intervention effects, and multiscale responses to the pharmacokinetic foundation. Its models can integrate molecular, cellular, tissue, and clinical information and can compare intervention scenarios that would be difficult or unethical to test directly. Best-practice recommendations emphasize model transparency, provenance, reproducibility, explicit assumptions, and documentation that permits qualified reuse [17]. These characteristics are necessary for digital-twin development because an evolving model cannot be meaningfully updated, audited, or transferred when its structure and evidence lineage are opaque. However, QSP complexity can also create weak identifiability, parameter nonuniqueness, and uncertainty that is difficult to propagate. A mechanistically detailed system should therefore not be called a twin merely because it resembles the represented biology.

Model credibility in pharmaceutical science is determined by what the model is expected to do, the consequences of an incorrect output, and the evidence supporting its use. Risk-informed frameworks for mechanistic in silico models distinguish verification, validation, uncertainty, applicability, and context of use rather than treating credibility as an intrinsic property of a model class [18]. This distinction is particularly important for digital twins because the same virtual representation might be adequate for hypothesis generation but inadequate for dose selection, formulation release, process control, or a patient-level intervention. Digital-twin terminology should therefore be attached to a declared use and evidence boundary. A system may qualify as a twin for low-consequence monitoring while remaining insufficiently validated for a higher-consequence decision.

In pharmaceutics and drug delivery, data-driven models are increasingly used to relate composition, processing conditions, material attributes, dosage-form characteristics, and biological performance. Such models may improve screening or organize complex design spaces, but retrospective prediction does not establish that a formulation or delivery system has been twinned [19]. A delivery twin would require an identifiable formulation, device, depot, batch, or patient-linked delivery process; observable state variables; an update mechanism; and validated predictions relevant to release, stability, exposure, transport, or performance. Pharmaceutical lifecycle reviews similarly show that the digital-twin label is applied at molecular, patient, formulation, process, and manufacturing scales [20]. These scales should not be merged into a single evidence category

because their sensors, update frequencies, validation comparators, decision consequences, and governance requirements differ substantially.

Mechanistic, Data-Driven, Hybrid, and Patient-Specific Approaches

Mechanistic approaches encode prior scientific knowledge through equations, causal structures, physiological compartments, material-balance relations, transport mechanisms, or disease pathways. They may support explanation, counterfactual reasoning, and extrapolation when the represented mechanisms remain valid. In silico clinical-trial literature illustrates how mechanistic or semi-mechanistic models can generate virtual patients and populations for scenario testing [21]. Yet virtuality is not equivalent to twinning. A virtual patient may be sampled from a parameter distribution, constructed to represent a phenotype, or used to explore trial outcomes without corresponding to an identifiable real person. Mechanistic plausibility and population realism are therefore distinct from persistent entity linkage and longitudinal state updating.

Data-driven approaches infer relations from observed data without requiring complete mechanistic specification. They may be valuable when the underlying pharmaceutical system is too complex, partially observed, or insufficiently understood for a comprehensive mechanistic model. Their limitations include dependence on training-data suitability, susceptibility to leakage and confounding, reduced reliability under domain shift, and uncertainty estimates that may not be calibrated. Verification, validation, and uncertainty quantification remain necessary regardless of whether the model is mechanistic or learned [22]. For a data-driven digital twin, model updating must also be distinguished from uncontrolled retraining: new observations should change the representation through documented procedures, and post-update behavior should remain traceable and evaluable.

Hybrid approaches combine mechanistic structure with statistical learning, Bayesian updating, surrogate modeling, residual correction, or data-driven state estimation. Such integration may preserve interpretable constraints while allowing the model to absorb patterns that are difficult to encode explicitly. Systems-pharmacology experience in immuno-oncology shows how model-based virtual patients can support individualized reasoning while also revealing the evidentiary distance between a calibrated virtual representation and a continuously updated patient twin [23]. Hybridization does not automatically close this distance. The learned component may dominate predictions outside the mechanistic model's valid domain, while the mechanistic component may impose incorrect assumptions. Credibility therefore requires component-level evaluation, integrated validation, uncertainty propagation, and evidence that updating improves rather than destabilizes future predictions.

Patient-specific approaches illustrate both the promise and the definitional difficulty of pharmaceutical digital twins. Model-informed precision dosing can combine population pharmacology, patient covariates, measured drug concentrations, Bayesian forecasting, and repeated dose adjustment. This architecture may approach twin behavior when it maintains patient identity, incorporates new observations, predicts future exposure or response, and informs a bounded dosing decision. However, implementation also depends on assay timing, data quality, clinical workflow, software usability, interpretability, accountability, and evaluation of patient-relevant consequences [24]. Patient specificity should therefore be considered necessary for a patient twin but not sufficient. A one-time individualized simulation remains a patient-specific model unless the entity relationship, update process, uncertainty, and decision loop are explicitly demonstrated.

Updating, Bidirectional Connection, Prediction, and Decision Use

Updating is the operational mechanism through which a digital twin remains aligned with its counterpart. It may involve Bayesian assimilation, parameter estimation, state correction, recalibration, model selection, or controlled retraining. The relevant update frequency depends on the dynamics of the represented entity: a manufacturing process may require frequent sensor-driven updates, whereas a chronic therapeutic model may update after laboratory measurements, dosing events, disease progression, or clinically meaningful changes. Experience with model-informed precision dosing demonstrates that technically capable individualized models may still fail to become routine decision systems because validation, implementation, interoperability, professional ownership, and accountability are unresolved [25]. Updating must therefore be assessed as part of a socio-technical workflow rather than as an isolated computational feature.

A model should not be described as self-updating merely because it can accept new data. The update trigger, accepted data types, quality controls, temporal alignment, missing-data handling, parameter constraints, post-update checks, and versioning procedure should be declared. Precision-dosing implementation work indicates that adoption depends on interfaces, education, institutional support, role definition, and integration with real decisions [26]. These considerations also apply to formulation and process twins. If scientists or operators must manually reinterpret model output outside the declared system, the model may still be useful, but the claimed bidirectional connection should be described accurately rather than implied through automation language.

Bidirectional connection consists of two distinguishable pathways. The first transfers observations from the physical or biological counterpart into the virtual representation. The second carries model outputs into a decision, intervention, control action, or experimental choice that can alter the future state of the counterpart. The return pathway may be mediated by a clinician, scientist, engineer, or quality professional and need not involve autonomous actuation. Pharmaceutical digital-patient-twin accounts emphasize this relationship between evolving observations, model-state updates, predictions, and therapeutic or manufacturing decisions [27]. A dashboard that displays predictions without specifying who acts, under what authority, and how the consequences return to the model is better described as a connected monitoring model or digital shadow.

Prediction should be separated from reconstruction, classification, explanation, experimental confirmation, and decision value. Retrospective agreement can show that a model reproduces known observations, but it does not demonstrate that predictions were fixed before outcomes became available. Clinical-trial applications may provide useful settings for evaluating virtual or AI-generated twins, including comparison with observed trajectories or synthetic-control uses [28]. Nevertheless, success within a trial dataset does not alone establish prospective patient-level usefulness, causal validity, or routine clinical value. Decision use requires a prespecified output, relevant forecast horizon, appropriate comparator, calibrated uncertainty, and evidence that the resulting action is sufficiently reliable for its consequence.

Evidence Synthesis on Maturity and Translational Readiness

The review-level evidence supports a maturity interpretation based on functions demonstrated rather than labels used. Across life-science applications, recurring twin-enabling components include individualized or entity-specific representation, multimodal data integration, simulation, prediction, and decision support. At the same time, governance, representativeness, interoperability, implementation, and regulation remain persistent limitations [29]. The evidence is therefore strongest for the existence of enabling model families and architectures, not for a general claim that pharmaceutical digital twins are routinely operational. Maturity should be judged independently for the represented entity, model credibility, update connection, prediction setting, uncertainty, decision pathway, and lifecycle controls.

Multiscale integration may increase scientific relevance by connecting molecular events, physiological dynamics, material behavior, treatment history, and observable outcomes. However, complex-systems perspectives caution that increasing biological or technical detail also increases calibration demands, identifiability problems, computational burden, and uncertainty propagation [30]. Model completeness is therefore not a realistic maturity criterion. A mature twin is not the most detailed model available; it is the least insufficient representation that has been credibly justified for a declared decision. This interpretation also prevents digital-twin rhetoric from rewarding complexity that cannot be validated or maintained.

Meta-review evidence indicates convergence around promising application domains but also around recurrent barriers involving evidence quality, data infrastructure, ethics, implementation, and validation [31]. These findings support a qualitative maturity airlock rather than a single linear scale. Models pass through evidentiary boundaries only when additional claims are supported: benchmark performance must be distinguished from pharmaceutical usefulness, prediction from experimental confirmation, association from mechanism or causality, and conceptual readiness from operational deployment. **Figure 1** presents the pharmaceutical digital-twin maturity airlock, showing how the article's key components and boundaries are connected within evidence synthesis on maturity and translational readiness.

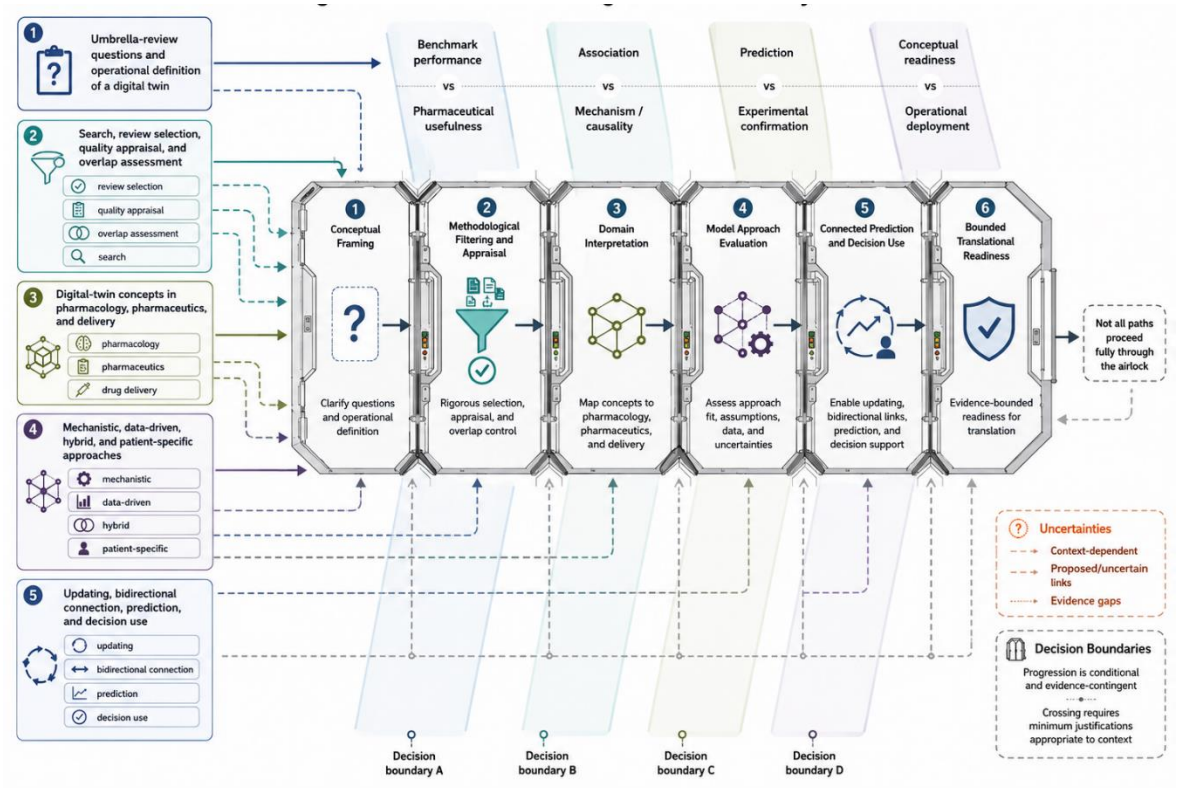

Figure 1. Pharmaceutical Digital-Twin Maturity Airlock

The figure is an original conceptual synthesis that organizes the article's central contribution across umbrella-review questions and operational definition of a digital twin, search, review selection, quality appraisal, and overlap assessment, review selection, quality appraisal, overlap assessment, digital-twin concepts in pharmacology, pharmaceutics, and delivery. Arrows

and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.
Broad technology reviews reinforce the finding that many reported systems remain enabling simulations, connected data environments, or proposed architectures rather than independently evaluated operational twins [32]. Translational readiness should therefore be reserved for applications that have passed prospective, context-specific evaluation and have documented data provenance, model versioning, applicability boundaries, drift monitoring, human oversight, and change control. A useful model need not satisfy these requirements when used for exploratory science. The stricter threshold applies only when the digital-twin label is used to imply an evolving counterpart relationship and a credible role in consequential pharmaceutical decisions.

Minimum Criteria for Calling a Pharmaceutical Model a Digital Twin
The first minimum criterion is a declared counterpart and context of use. The counterpart may be a patient, product, formulation, drug-delivery system, manufacturing batch, device, unit operation, or another bounded pharmaceutical entity, but its identity and lifecycle must be traceable. The intended decision must also be stated because credibility expectations depend on model risk and the consequences of error. Model-informed drug-development frameworks show that evidence should be scaled to context of use rather than inferred from model sophistication or publication status [33]. A model cannot be judged as a digital twin in the abstract; it can only be judged relative to what it represents, what it predicts, and how its output is intended to be used.
The second criterion is a defined twinning architecture. This includes the virtual representation, physical-to-virtual data pathway, state-estimation or update method, and model-to-decision pathway. Cross-domain digital-twin reviews identify modeling, sensing, communication, data management, and service layers as enabling elements while also distinguishing one-way digital shadows from more complete twinning relationships [34]. Pharmaceutical implementations need not reproduce an industrial architecture mechanically, but they should disclose equivalent functions. In particular, the existence of a return pathway should not be inferred from a model's ability to generate recommendations. The responsible actor, permissible action, latency, override mechanism, and feedback capture should be identifiable.
The third criterion is evidence that the connected model remains credible as it changes. Updating introduces risks that do not arise in a fixed model, including unstable parameter estimation, silent degradation, feedback-induced bias, loss of calibration, and uncontrolled change in decision behavior. Uncertainty quantification, sensitivity analysis, domain-of-validity monitoring, and bounded optimization are therefore central to twin credibility [35]. The twin should be able to communicate when evidence is insufficient, when inputs fall outside its supported domain, and when human escalation is required. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop minimum criteria for calling a pharmaceutical model a digital twin within the article's central argument.

Table 2. Evaluation, implementation, and research priorities arising from When Does a Pharmaceutical Model Become a Digital Twin? An Umbrella Review

Approach or application	Data and task context	Evaluation practice	Evidence maturity	Translational limitation	Review interpretation
Patient-specific PBPK or QSP	Individual physiology, disease state, treatment history, biomarkers, and drug-exposure data	Verification, parameter evaluation, sensitivity analysis, external comparison, and context-of-use qualification	Established model credibility practices; variable twin-specific connection	Patient data may be intermittent, mechanisms may be weakly identifiable, and prospective decision benefit may be untested	A strong twin-enabling foundation that becomes a patient twin only with persistent identity, validated updating, uncertainty, and decision coupling
Model-informed precision dosing	Population model, patient covariates, measured concentrations, dose history, and forecasted exposure or response	Bayesian forecasting, calibration assessment, temporal validation, workflow evaluation, and outcome-focused comparison	Relatively advanced individualized updating in selected settings	Implementation, assay timing, interoperability, ownership, and prospective outcome evidence remain uneven	One of the closest pharmaceutical analogues to a decision-coupled twin when repeated updates and bounded dose decisions are demonstrated
Virtual patients and in silico trials	Synthetic populations, disease models, mechanistic parameters, trial-design assumptions, and observed comparator data	Plausibility checks, distributional validation, scenario analysis, sensitivity analysis, and comparison with trial observations	Moderate for simulation and trial-support functions	A virtual patient may not map to an identifiable evolving real patient	Retain the virtual-patient label unless persistent entity linkage and longitudinal updating are demonstrated

Data-driven formulation design	Composition, process parameters, material attributes, dosage-form properties, and performance measurements	Cross-validation, external testing, uncertainty calibration, domain-shift assessment, and prospective formulation confirmation	Stronger for retrospective prediction than for connected twinning	Sparse datasets, inconsistent measurements, scale-up changes, and limited real-time formulation-state observations	A predictive formulation model is not a twin without a specified product or process counterpart and validated state updating
Mechanistic drug-delivery modeling	Release, transport, degradation, device or depot state, anatomy, physiology, and exposure data	Verification, multiscale calibration, orthogonal in vitro and in vivo comparison, uncertainty propagation, and prospective testing	Conceptually strong but heterogeneous	Material–biology coupling and in vivo state observability are limited	A potential twin component that requires observable states, assimilation rules, and decision relevance
Hybrid mechanistic–machine-learning model	Mechanistic constraints combined with learned residuals, surrogate models, or state estimators	Component-level testing, integrated validation, ablation, uncertainty assessment, and extrapolation analysis	Emerging	Learned components may dominate under shift, while mechanistic assumptions may remain misspecified	Hybrid architecture may improve utility but does not reduce the need for transparency and credibility assessment
Pharmaceutical process or manufacturing twin	Sensor streams, equipment state, material attributes, process parameters, control settings, and quality variables	Sensor validation, state-estimation testing, process verification, shadow-mode evaluation, control testing, and lifecycle monitoring	Potentially more operational in bounded processes	Transfer across products, facilities, equipment, and operating ranges is uncertain	May represent the most mature twinning pattern when counterpart identity, synchronization, control boundaries, and change management are explicit
AI-generated clinical or trial twin	Longitudinal clinical data, trial records, generative models, and counterfactual trajectories	Retrospective reconstruction, external trial comparison, prospective prediction, calibration, and bias assessment	Early to moderate	Synthetic similarity does not establish causal validity, individual utility, or routine clinical value	Trial applications can contribute evidence but should not be generalized to patient-level deployment
Decision-coupled pharmaceutical twin	Any bounded patient, product, process, formulation, device, or development entity with longitudinal observations	Prospective validation, calibrated uncertainty, workflow or control evaluation, auditability, comparator-based decision assessment, and drift monitoring	Limited and context specific	Governance, accountability, lifecycle change, and decision-consequence evidence are frequently incomplete	Reserve the digital-twin designation for systems demonstrating all minimum components within a declared context of use

Table note: The categories and maturity interpretations are an original umbrella-review synthesis. They are not validated scores, comparative-effectiveness estimates, clinical recommendations, regulatory classifications, or universal implementation standards.

The fourth criterion is an explicit evidence and governance envelope. A reference architecture should make the counterpart, virtual representation, acquisition layer, synchronization, analytics, services, interfaces, security, and responsibility boundaries visible [36]. On this basis, the proposed minimum designation requires: a specified counterpart; a fit-for-purpose virtual representation; a state-relevant connection; documented update logic; prospective prediction; calibrated uncertainty and applicability limits; an auditable decision pathway; traceability, oversight, and change control; and evidence of decision value appropriate to the context. These criteria are proposed rather than validated and should be applied as a boundary framework, not as a universal certification checklist. **Figure 2** presents the evidence, validation, and translation boundary observatory for when does a pharmaceutical model become a, showing how the article's key components and boundaries are connected within minimum criteria for calling a pharmaceutical model a digital twin.

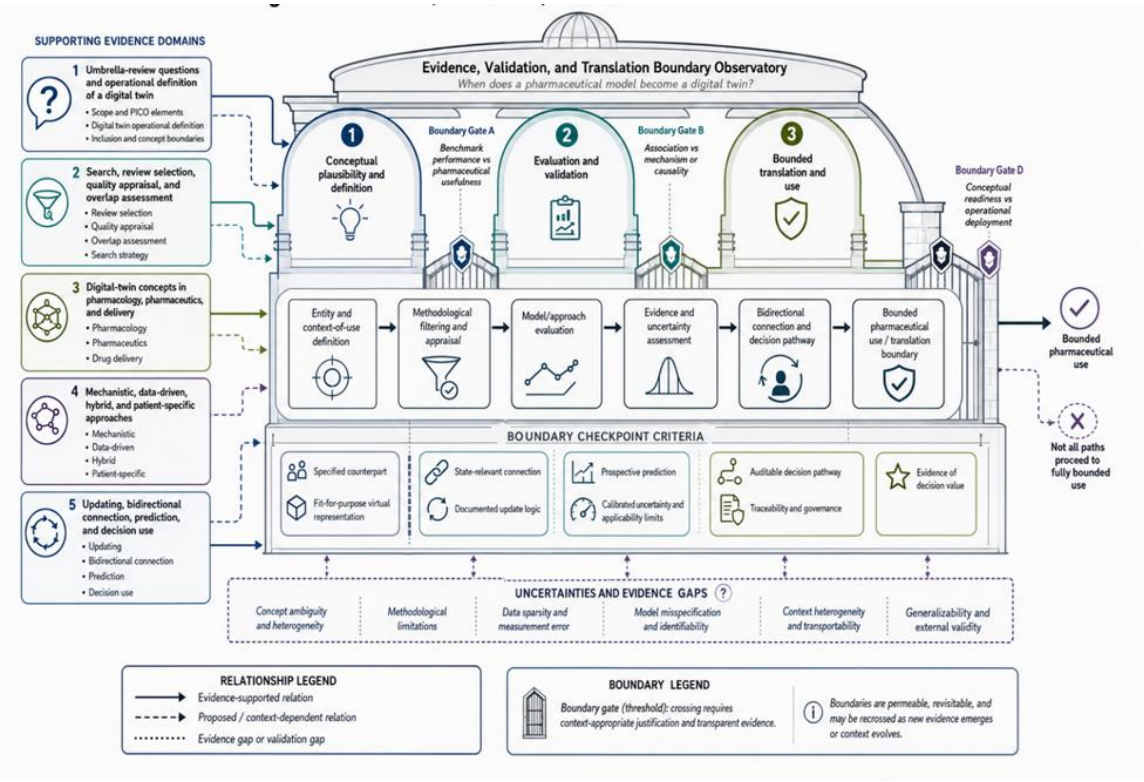

Figure 2. Evidence, Validation, and Translation Boundaries

The figure is an original conceptual synthesis that shows how claims move from conceptual plausibility through evaluation, uncertainty assessment, and bounded pharmaceutical use. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Conclusion

A pharmaceutical model becomes a digital twin not when it becomes computationally elaborate, personalized, generative, or highly predictive, but when it enters a credible and traceable relationship with a specified real counterpart. The proposed boundary requires state-relevant connection, documented updating, prospective and uncertainty-aware prediction, an auditable path from model output to a bounded pharmaceutical decision, and governance across the model lifecycle. Pharmacometric, systems-pharmacology, formulation, delivery, virtual-patient, and manufacturing approaches provide important twin-enabling functions, yet their maturity varies according to entity observability, validation depth, update reliability, decision consequence, and implementation context. Review-level convergence should also be interpreted cautiously because overlapping primary evidence, heterogeneous terminology, and variable appraisal quality can create an appearance of maturity that exceeds demonstrated capability. The principal contribution of this umbrella review is therefore a proposed operational distinction between useful pharmaceutical models and decision-coupled digital twins. This distinction is intended to improve scientific language, evidence synthesis, validation planning, and translational judgment rather than to imply that digital-twin status is necessary for every valuable pharmaceutical model.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Björnsson B, Borrebaeck C, Elander N, Gasslander T, Gawel DR, Gustafsson M, et al. Digital twins to personalize medicine. Genome Med. 2020;12(1):4. doi:10.1186/s13073-019-0701-3

2. Corral-Acero J, Margara F, Marciniak M, Rodero C, Loncaric F, Feng Y, et al. The 'Digital Twin' to enable the vision of precision cardiology. Eur Heart J. 2020;41(48):4556-64. doi:10.1093/eurheartj/ehaa159
3. Moingeon P, Chenel M, Rousseau C, Voisin EM, Guedj M. Virtual patients, digital twins and causal disease models: Paving the ground for in silico clinical trials. Drug Discov Today. 2023;28(7):103605. doi:10.1016/j.drudis.2023.103605
4. Bordukova M, Makarov N, Rodriguez-Esteban R, Schmich F, Menden MP. Generative artificial intelligence empowers digital twins in drug discovery and clinical trials. Expert Opin Drug Discov. 2024;19(1):33-42. doi:10.1080/17460441.2023.2273839
5. Venkatapurapu SP, Clegg L, Nowojewski A, Kimko H, Olabode D, Sawant-Basak A, et al. Digital twins for accelerating drug discovery and development: Opportunities and challenges. Drug Discov Today. 2026;31(2):104617. doi:10.1016/j.drudis.2026.104617
6. Jones D, Snider C, Nassehi A, Yon J, Hicks B. Characterising the digital twin: A systematic literature review. CIRP J Manuf Sci Technol. 2020;29:36-52. doi:10.1016/j.cirpj.2020.02.002
7. Barricelli BR, Casiraghi E, Fogli D. A survey on digital twin: Definitions, characteristics, applications, and design implications. IEEE Access. 2019;7:167653-71. doi:10.1109/ACCESS.2019.2953499
8. Rasheed A, San O, Kvamsdal T. Digital twin: Values, challenges and enablers from a modeling perspective. IEEE Access. 2020;8:21980-2012. doi:10.1109/ACCESS.2020.2970143
9. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, et al. Digital twins for health: A scoping review. NPJ Digit Med. 2024;7(1):77. doi:10.1038/s41746-024-01073-0
10. Pellegrino G, Gervasi M, Angelelli M, Corallo A. A conceptual framework for digital twin in healthcare. Inf Syst Front. 2025;27(1):7-32. doi:10.1007/s10796-024-10536-4
11. Shea BJ, Reeves BC, Wells G, Thuku M, Hamel C, Moran J, et al. AMSTAR 2: A critical appraisal tool for systematic reviews that include randomised or non-randomised studies of healthcare interventions, or both. BMJ. 2017;358. doi:10.1136/bmj.j4008
12. Gates M, Gates A, Pieper D, Fernandes RM, Tricco AC, Moher D, et al. Reporting guideline for overviews of reviews of healthcare interventions: Development of the PRIOR statement. BMJ. 2022;378. doi:10.1136/bmj-2022-070849
13. Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ. 2021;372. doi:10.1136/bmj.n71
14. Kirvalidze M, Abbadi A, Dahlberg L, Sacco LB, Calderón-Larrañaga A, Morin L. Estimating pairwise overlap in umbrella reviews: Considerations for using the corrected covered area index methodology. Res Synth Methods. 2023;14(5):764-7. doi:10.1002/jrsm.1658
15. Lunny C, Pieper D, Thabet P, Kanji S. Managing overlap of primary study results across systematic reviews: Practical considerations for authors of overviews of reviews. BMC Med Res Methodol. 2021;21(1):140. doi:10.1186/s12874-021-01269-y
16. Shebley M, Sandhu P, Riedmaier AE, Jamei M, Narayanan R, Patel A, et al. Physiologically based pharmacokinetic model qualification and reporting procedures for regulatory submissions: A consortium perspective. Clin Pharmacol Ther. 2018;104(1):88-110. doi:10.1002/cpt.1013
17. Cucurull-Sanchez L, Chappell MJ, Chelliah V, Cheung SYA, Derks G, Penney M, et al. Best practices to maximize the use and reuse of quantitative and systems pharmacology models: Recommendations from the United Kingdom Quantitative and Systems Pharmacology Network. CPT Pharmacometrics Syst Pharmacol. 2019;8(5):259-72. doi:10.1002/psp4.12381
18. Musuamba Tshinanu F, Skottheim Rusten I, Lesage R, Russo G, Bursi R, Emili L, et al. Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility. CPT Pharmacometrics Syst Pharmacol. 2021;10(8):804-25. doi:10.1002/psp4.12669
19. Vora LK, Gholap AD, Jetha K, Thakur RRS, Solanki HK, Chavda VP. Artificial intelligence in pharmaceutical technology and drug delivery design. Pharmaceutics. 2023;15(7):1916. doi:10.3390/pharmaceutics15071916
20. Maharjan R, Kim NA, Kim KH, Jeong SH. Transformative roles of digital twins from drug discovery to continuous manufacturing: Pharmaceutical and biopharmaceutical perspectives. Int J Pharm X. 2025;10:100409. doi:10.1016/j.ijpx.2025.100409
21. Pappalardo F, Russo G, Musuamba Tshinanu F, Viceconti M. In silico clinical trials: Concepts and early adoptions. Brief Bioinform. 2019;20(5):1699-708. doi:10.1093/bib/bby043
22. Viceconti M, Pappalardo F, Rodriguez B, Horner M, Bischoff JE, Musuamba Tshinanu F. In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products. Methods. 2021;185:120-7. doi:10.1016/j.ymeth.2020.01.011
23. Wang H, Arulraj T, Ippolito A, Popel AS. From virtual patients to digital twins in immuno-oncology: Lessons learned from model-based systems pharmacology. NPJ Digit Med. 2024;7(1):189. doi:10.1038/s41746-024-01188-4
24. Darwich AS, Polasek TM, Aronson JK, Ogungbenro K, Wright DFB, Achour B, et al. Model-informed precision dosing: Background, requirements, validation, implementation, and forward trajectory of individualizing drug therapy. Annu Rev Pharmacol Toxicol. 2021;61:225-45. doi:10.1146/annurev-pharmtox-033020-113257

25. Darwich AS, Ogungbenro K, Vinks AA, Powell JR, Reny JL, Marsousi N, et al. Why has model-informed precision dosing not yet become common clinical reality? Lessons from the past and a roadmap for the future. Clin Pharmacol Ther. 2017;101(5):646-56. doi:10.1002/cpt.659
26. Polasek TM, Rostami-Hodjegan A, Yim DS, Jamei M, Lee H, Kimko H, et al. What does it take to make model-informed precision dosing common practice? Report from the 1st Asian Symposium on Precision Dosing. AAPS J. 2019;21(2):17. doi:10.1208/s12248-018-0286-6
27. Fischer RP, Volpert S, Antonino P, Ahrens TD. Digital patient twins for personalized therapeutics and pharmaceutical manufacturing. Front Digit Health. 2024;5:1302338. doi:10.3389/fdgth.2023.1302338
28. Vidovszky AA, Fisher CK, Loukianov AD, Smith AM, Tramel EW, Walsh JR, et al. Increasing acceptance of AI-generated digital twins through clinical trial applications. Clin Transl Sci. 2024;17(7). doi:10.1111/cts.13897
29. Venkatesh KP, Brito G, Kamel Boulos MN. Health digital twins in life science and health care innovation. Annu Rev Pharmacol Toxicol. 2024;64:159-70. doi:10.1146/annurev-pharmtox-022123-022046
30. De Domenico M, Allegri L, Caldarelli G, d'Andrea V, Di Camillo B, Rocha LM, et al. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. NPJ Digit Med. 2025;8(1):37. doi:10.1038/s41746-024-01402-3
31. Ringeval M, Etindele Sosso FA, Cousineau M, Paré G. Advancing health care with digital twins: Meta-review of applications and implementation challenges. J Med Internet Res. 2025;27. doi:10.2196/69544
32. Abd Elaziz M, Al-qaness MAA, Dahou A, Al-Betar MA, Mohamed MM, El-Shinawi M, et al. Digital twins in healthcare: Applications, technologies, simulations, and future trends. Wiley Interdiscip Rev Data Min Knowl Discov. 2024;14(6). doi:10.1002/widm.1559
33. Kuemmel C, Yang Y, Zhang X, Florian J, Zhu H, Tegenge M, et al. Consideration of a credibility assessment framework in model-informed drug development: Potential application to physiologically-based pharmacokinetic modeling and simulation. CPT Pharmacometrics Syst Pharmacol. 2020;9(1):21-8. doi:10.1002/psp4.12479
34. Thelen A, Zhang X, Fink O, Lu Y, Ghosh S, Youn BD, et al. A comprehensive review of digital twin—part 1: Modeling and twinning enabling technologies. Struct Multidisc Optim. 2022;65(12):354. doi:10.1007/s00158-022-03425-4
35. Thelen A, Zhang X, Fink O, Lu Y, Ghosh S, Youn BD, et al. A comprehensive review of digital twin—part 2: Roles of uncertainty quantification and optimization, a battery digital twin, and perspectives. Struct Multidisc Optim. 2023;66(1):1. doi:10.1007/s00158-022-03410-x
36. De Benedictis A, Mazzocca N, Somma A, Strigaro C. Digital twins in healthcare: An architectural proposal and its application in a social distancing case study. IEEE J Biomed Health Inform. 2023;27(10):5143-54. doi:10.1109/JBHI.2022.3205506