International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com
Volume 3 Issue 8 , August 2026
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos

From Continuous Biological Sensing to Precision Pharmacotherapy: An Observability-Aware Dynamic Patient–Drug Modelling Framework

Satyam Tiwari[1], Swapnil Sharma[2]
[1,2] *Student, Apj Abdul Kalam Technical University.*

Article Info

Article History:

Published: 21 August 2026

Publication Issue:
Volume 3, Issue 8
August-2026

Page Number:
509-525

Corresponding Author:
Satyam Tiwari

Abstract:

Background. Precision medicine has historically relied on intermittent measurements, such as a laboratory panel, a genetic test, or a clinic visit, to characterize a patient who is, in reality, a continuously changing physiological system. Problem. Continuous biosensing, artificial intelligence (AI), digital twins, multi-omics, and model-informed precision dosing (MIPD) are each advancing rapidly, and several recent frameworks already combine them, including temporal causal inference on routine physiological data, state-space digital twins for critical-care precision dosing, and proposed multi-omics/AI/digital-twin architectures for predictive, preventive, personalized, and participatory (P4) medicine. This heterogeneity of partially overlapping efforts makes it easy to overstate novelty and difficult to see what remains structurally unresolved. Approach. This review maps the relevant literature against a common architecture and argues that the unresolved problem is not integration but biological observability: the gap between what is measured and what is physiologically true. Framework. We propose a conceptual, mathematically explicit synthesis, the Observability-Aware Dynamic Patient–Drug Model (OAD-PDM), which treats the patient as a partially observed dynamical system whose latent physiological state is estimated from heterogeneous, asynchronous observations together with an explicit uncertainty and observability annotation, and which couples that state bidirectionally to mechanistic pharmacokinetic/pharmacodynamic (PK/PD) models rather than to an unconstrained predictive layer. Evidence base. We illustrate the framework using three domains at different maturity levels: closed-loop insulin delivery (clinically established), hybrid mechanistic-machine-learning vancomycin dosing (actively converging, supported by recent comparative studies), and continuous gastrointestinal sensing for oral drug absorption (experimental, demonstrated only in animal studies to date). We position the framework against its closest existing counterparts rather than claiming to have invented latent-state modelling, digital twins, or AI-assisted dosing. Limitations. The OAD-PDM framework is conceptual. It has not been implemented or clinically validated, several of its proposed outputs lack established estimation procedures, and its translational value depends on sensor reliability, external validation, and formal uncertainty quantification that do not yet exist for most of the biological domains it spans. Conclusion. We close with a staged, deliberately conservative research roadmap and an explicit statement of limitations.

Keywords: Artificial intelligence; precision medicine; digital twins; continuous biosensing; physiological state estimation; pharmacokinetics/pharmacodynamics; model-informed precision dosing; uncertainty quantification

1. Introduction

1.1 From population medicine to a snapshot of the individual

Precision medicine reframed therapy around the individual rather than the population average, using genotype, phenotype, and clinical history to stratify patients and guide treatment choice. Pharmacogenomics and MIPD extended

this logic to dosing itself, combining population pharmacokinetic (PopPK) models with patient covariates and, where available, measured drug concentrations to individualize a regimen. These are real and productive research programs. But they share a structural feature: they treat the patient's relevant biological characteristics as comparatively fixed inputs to a model, updated only when a new measurement arrives.

1.2 The snapshot problem

Human physiology does not hold still between measurements. Renal and hepatic function, hydration state, inflammatory tone, and drug distribution all vary over the course of a single day, sometimes over minutes, especially in acute illness. Yet many of the variables precision medicine depends on, such as a metabolic panel, a hormone level, or a drug trough concentration, are sampled intermittently, sometimes only once. Between samples, the patient's true state is not observed; it is assumed to persist or is reconstructed only when the next data point arrives. In critical illness this mismatch has been raised explicitly in the literature: a proposed Dynamic Living Digital Twin (DLDT) framework argues that MIPD's reliance on intermittently updated data creates a temporal mismatch with rapidly evolving physiology, and proposes a state-space alternative to snapshot-based vancomycin dosing in critically ill patients [4].

1.3 The rise of continuous biological sensing, and its limits

Continuous glucose monitoring (CGM) demonstrated that dense, longitudinal physiological measurement is technically achievable and clinically useful. That success has motivated a much broader wave of continuous and semi-continuous sensing: wearable electrochemical biosensors capable of monitoring metabolites, electrolytes, and hormones beyond glucose in sweat, interstitial fluid, and saliva [15]; ingestible capsules that profile the gastrointestinal (GI) lumen directly [8]; and continuous or near-continuous drug monitoring aimed at supplementing sparse therapeutic drug monitoring (TDM) with denser pharmacokinetic trajectories. A 2026 Perspective on wearable molecular sensors in endocrinology and metabolism is explicit that this expansion is uneven: it identifies incomplete understanding of longitudinal biomarker physiology, technical barriers that still prevent continuous monitoring of most clinically relevant biomarkers, machine-learning difficulty with dense and noisy biomarker data, insufficient large-scale longitudinal clinical validation, and poor interoperability with clinical workflows as persistent, unresolved challenges, not solved problems [10]. A broader review of wearable electrochemical biosensors similarly frames biofouling, sensor drift, and long-term stability as open engineering problems rather than settled ones [15]. More sensors, in other words, do not automatically resolve the snapshot problem; they change its shape.

1.4 The observability problem

Even where continuous sensing exists, what is measured is rarely identical to the biological quantity of clinical interest. Peripheral oxygen saturation is not the complete state of gas exchange; sweat or interstitial sodium is not plasma sodium; skin temperature is not core temperature; a genotype is a static predisposition, not a real-time enzyme activity. Formally, if x_t denotes an individual's true physiological state at time t, sensors provide only:

$$y_t = \mathrm{h}(x_t) + \eta_t$$

Here y_t is a noisy, partial, sensor-specific projection of x_t, not x_t itself. We refer to this as the **biological observability problem**: the central obstacle is not simply acquiring more data, but converting heterogeneous, imperfect, asynchronously sampled observations into a defensible estimate of an unobserved biological state. We use the term “observability” here in a deliberately informal, clinical sense; Section 4.8 distinguishes this usage from the formal, technical notion of observability used in control theory.

1.5 Why AI alone does not resolve the problem

Multimodal, temporal machine learning can plausibly fuse heterogeneous streams and detect nonlinear relationships that simple thresholds miss. But an unconstrained predictive model has no obligation to respect mass balance, non-negativity of concentrations, or known pharmacological structure, and is vulnerable to distribution shift and spurious correlation in dense biomarker data, concerns raised directly in the wearable-sensing literature discussed above [10]. Recent hybrid pharmacokinetic modelling work makes the same point from the pharmacology side: purely data-driven models can produce physically implausible predictions unless mechanistic structure constrains them. The recurring argument across this literature is not "AI versus mechanistic models" but that each compensates for the other's blind spots.

1.6 Why pharmacology makes the problem harder, not easier

Adding a drug to this picture compounds the difficulty in two ways. First, drug response is not simply a function of the current state; effects are delayed and prolonged relative to dosing, so a system that only conditions on the immediately preceding state violates the Markov assumption that much sequential decision-making machinery, including standard reinforcement learning, depends on. Basu and colleagues formalized this directly, identifying delayed and prolonged medication effects as a structural obstacle to naïve reinforcement-learning-based dosing and proposing a Prolonged Action Effect–Partially Observable Markov Decision Process (PAE-POMDP) to restore a workable Markov structure [9]. Second, physiology and pharmacokinetics are coupled bidirectionally: renal and hepatic function shape drug clearance and volume of distribution, while the drug in turn perturbs the physiological state that determines those same parameters. A model of "the patient" that ignores this feedback is incomplete for pharmacotherapy specifically, even where it might be adequate for passive monitoring.

1.7 What has already been proposed, and what has not

Contrary to a "nobody has done this" framing, several recent efforts already occupy adjacent territory. A generalist framework using temporal causal inference, Temporal Causal Precision Medication (TCPM), learns a treatment-free physiological profile from routine clinical data to estimate unbiased treatment effects and recommend individualized strategies, validated across six clinical cohorts [3]. A separately proposed Dynamic Living Digital Twin (DLDT) framework reframes patient physiology in critical illness as a continuously evolving latent state, updated with high-frequency electronic health record (EHR) data through adaptive state-space approaches such as Kalman filtering, specifically to address the snapshot-dosing problem for vancomycin [4]. A 2026 perspective developed by Immorta Bio in collaboration with Theriome Inc. and the American Board of Precision Medicine proposes integrating multi-omics, AI, digital twins, and blockchain-secured data governance to move P4 medicine from aspiration toward clinical architecture [5]. A related, independently authored perspective by Emmert-Streib and colleagues similarly argues that combining digital twins with the P4 framework yields capabilities beyond either alone [6]. Disease- and population-specific mechanistic digital twins have also progressed substantially: Prunella and colleagues describe an evolutionary digital twin for aminoglycoside dosing in neonates with suspected sepsis that couples a physiologically based pharmacokinetic-pharmacodynamic (PBPK-PD) model with a data-driven, continuously updated renal-function predictor, calibrated on a cohort of 1,634 neonates [7].

Given this, this review does **not** claim that AI-enabled digital twins, latent physiological state estimation, continuous biosensing, counterfactual treatment prediction, or AI-assisted precision dosing are individually novel. Each already has a substantial and, in places, clinically evaluated literature. What we argue is narrower: that a broadly applicable, biologically interpretable account of *how* incomplete, asynchronous, multi-compartment biological observations should be converted into an uncertainty-aware latent state, and coupled bidirectionally to mechanistic PK/PD rather

than to an opaque predictive layer, remains underdeveloped in the frameworks examined in this review. We return to this comparison in more structural detail in Section 6.

1.8 Research questions

This review is organized around four questions:

- **RQ1 (Observation):** Which physiological and pharmacological variables can currently be measured continuously or longitudinally, and which remain latent or only indirectly inferable?
- **RQ2 (Modelling):** How are AI, state-space, and mechanistic approaches currently combined to estimate and forecast dynamic physiological and pharmacological states?
- **RQ3 (Patient–drug coupling):** How can a dynamic physiological state, static biological priors (genome, pharmacogenomics), and mechanistic PK/PD models be integrated to predict individualized treatment exposure and response?
- **RQ4 (Translation):** What validation, uncertainty, sensing, clinical, and regulatory barriers currently limit the translation of such frameworks into adaptive precision pharmacotherapy?

2. Review Approach and Scope

This is a **critical, purposively sampled conceptual review**, not a formal systematic scoping review. We searched general academic and publisher literature for work spanning six thematic areas: human/patient digital twins, continuous biosensing and molecular monitoring, AI and temporal/state-space modelling of physiology, MIPD/TDM/PK-PD/PBPK/QSP, multi-omics and pharmacogenomics, and closed-loop therapeutics, concentrating on 2023–2026 with earlier foundational work included where relevant. Because no fixed, pre-registered protocol was executed and no reproducible multi-database record count was generated, we do not report PRISMA-style flow statistics; doing so without an actual audit trail would misrepresent the rigor of the search. Every reference cited for a specific quantitative or empirical claim was individually located and read in primary or publisher-indexed form; we did not cite search-engine snippets as a substitute for the primary source, and quantitative statistics reported in Section 7 are quoted directly from the source's own reported results rather than estimated or recalculated. For two sources discussed in Section 1.7 (references 4 and 5), the paper's existence, title, and journal were confirmed through publisher-level indexing, but a complete, independently verified author byline, volume, and page range could not be obtained through the channels available during preparation of this manuscript; this is stated explicitly at first citation and in the reference list. Where a claim from earlier project working notes could not be independently verified, it has been omitted rather than reproduced. This is a genuine limitation of the review, discussed further in Section 10, and it means the evidence mapped here should be read as **illustrative of the field's structure and maturity gradient**, not as an exhaustive census.

3. From Biomarkers to Biological State: The Observability Problem

3.1 State versus measurement

We distinguish the patient's true, largely unobserved physiological state x_t from the set of measurements y_t actually available at time t. A hierarchical decomposition is useful because different physiological domains have very different observability profiles:

$$x_t = \{ X_t^{metabolic}, X_t^{cardiovascular}, X_t^{respiratory}, X_t^{fluid/electrolyte}, X_t^{endocrine}, X_t^{immune}, X_t^{GI}, X_t^{organ}, X_t^{pharmacological} \}$$

This nine-domain grouping is a pragmatic, clinically motivated organization rather than a first-principles decomposition; it follows organ-system boundaries already used in clinical assessment and in the sensing and

pharmacology literature cited throughout this review, and is intended as a communication device rather than a claim that these are the objectively correct state coordinates. Systems-biology approaches to multiscale model composition (for example, quantitative systems pharmacology (QSP) model libraries and standardized model-exchange formats) already address the more general problem of principled state decomposition and cross-scale coupling; OAD-PDM does not propose a new multiscale-modelling formalism and should be read as sitting downstream of, and compatible with, that literature rather than in competition with it.

3.2 A maturity gradient, not a binary

Rather than a simple "measurable / not measurable" split, the literature supports a four-level gradient (Table 3):

Table 3. A four-level biological observability gradient

Level	Description	Illustrative examples
Directly observable, mature	Continuous, clinically validated	Heart rate, SpO_2, glucose (CGM), activity, some blood pressure measures
Directly observable, emerging	Continuous measurement technically demonstrated, limited clinical validation	Lactate, ketones, sweat/interstitial electrolytes, GI luminal pH and selected metabolites/hormones via ingestible capsule
Indirectly inferable	Estimated from correlated signals, not measured directly	Hydration state, systemic inflammatory tone, "cardiovascular stress," core temperature from skin temperature
Largely latent	No reliable continuous non-invasive proxy currently exists at scale	Tissue/organ-specific drug concentration, intracellular signalling, most cytokine dynamics, real-time hepatic or renal metabolic capacity

The gastrointestinal domain illustrates the "emerging" tier well. Commercial ingestible sensors have historically monitored only basic parameters such as pH and pressure; Min and colleagues reported an integrated smart capsule ("PillTrek") in 2025 capable of simultaneously detecting electrolytes, metabolites, and hormones within the GI tract [8]. This is a genuine expansion of what is measurable, but it remains, at the time of writing, an experimental technology validated in animal studies rather than a clinically deployed human monitoring tool, and it does not, by itself, resolve the harder problem of relating luminal measurements to systemic drug absorption.

3.3 Cross-compartment translation is not free

Even where a biomarker is measured continuously, the compartment sampled is frequently not the compartment of clinical interest: sweat sodium is not plasma sodium; interstitial glucose lags plasma glucose; skin temperature is not core temperature. A defensible state-estimation framework therefore needs an explicit, and where possible learned or calibrated, compartment-transformation step rather than treating peripheral measurements as interchangeable with the systemic quantity they are meant to proxy. This review does not propose a general solution to this problem; it names it as a first-class requirement that any implementation of the Layer 3 state estimator described in Section 5 would need to address.

3.4 Asynchrony

Different modalities arrive on entirely different clocks: cardiac signals at sub-second resolution, CGM every few minutes, drug concentrations intermittently, genomics essentially once per lifetime, laboratory panels daily to weekly. Any state-estimation approach for this problem must therefore handle irregularly and asynchronously sampled multimodal time series rather than assuming a synchronized observation grid, a requirement that favours continuous-time or explicitly asynchronous formulations (state-space models, neural ODEs, Bayesian filtering with irregular update times) over architectures that implicitly assume regular sampling.

4. Mathematical Framework for a Dynamic Patient–Drug System

The purpose of this section is not to introduce new estimation algorithms, several already exist and are cited where used, but to give the biological argument above a common, explicit notation, so that later sections (the proposed architecture, the comparison with existing frameworks, and the case studies) can refer to a shared formal object. Throughout, we distinguish established mathematical machinery (state-space models, Kalman filtering, do-calculus) from our own conceptual proposal (the explicit observability annotation introduced in Section 4.8), and we note this distinction again where it matters.

4.1 State transition

$x_{t+1} = \mathrm{f}(x_t, u_t, \mathrm{g}, e_t) + \varepsilon_t$

where u_t is the therapeutic intervention (drug, dose, timing), g is a relatively static biological prior (genome, pharmacogenomic variants), e_t is time-varying environmental/behavioural context, and ε_t is biological/process uncertainty. This is standard state-space notation; f is not specified further here and is understood to be approximated, in any implementation, by a mechanistic, statistical, or hybrid model appropriate to the physiological domain in question.

4.2 Observation model

$y_t^{(j)} = h_j(x_t) + \eta_t^{(j)}$

for modality j (wearable, laboratory panel, molecular sensor, imaging, EHR entry), each providing a different, imperfect projection of the same underlying state.

4.3 Individualized homeostatic deviation

Rather than comparing a patient's values against a population reference range, we define an individual, time-varying reference trajectory x_t^* (which may itself shift with circadian rhythm, age, or chronic disease) and the deviation:

$\Delta x_t = x_t - x_t^*$

The claim that homeostasis is better represented as an individualized trajectory than as a fixed population midpoint is a reasonable extension of long-standing physiological regulation theory rather than a new empirical finding, and we present it as such.

4.4 Pharmacokinetics coupled to physiological state

A minimal one-compartment oral absorption model,

$dA_g / dt = -k_a A_g , \quad dC_p / dt = (k_a / V)A_g - (CL / V)C_p$

is standard and is shown here only as the simplest illustrative case; real pharmacokinetic models for most drugs of clinical interest are more complex (multi-compartment, nonlinear, or physiologically based). The individualization

step of interest to this framework is to let clearance, volume of distribution, and absorption rate depend on the *current* physiological state and static biology rather than treating them as fixed patient covariates:

$\theta_i(t) = \{ k_{a,i}(t), V_i(t), CL_i(t) \} = \varphi(g_i, x_t, \text{disease}_i, e_t)$

Here φ denotes a covariate-mapping function in the same general sense already used in population pharmacokinetic practice, where clearance or volume terms are written as functions of covariates such as weight, renal function, or genotype; the only change proposed here is that the covariate set includes the dynamically estimated state x_t rather than only static or infrequently updated covariates. We use φ rather than f to avoid reusing the state-transition function symbol from Section 4.1 for a different mapping. This is the formal expression of the bidirectional coupling argued for in Section 1.6: physiological state shapes pharmacokinetic parameters, and the resulting drug exposure feeds back into x_{t+1} through Section 4.1.

4.5 Pharmacodynamics

$E(t) = \psi(C_p(t), x_t)$

for example an E_{max} -type relationship modulated by physiological state rather than concentration alone. We use ψ rather than g for this mapping to avoid reusing the symbol already assigned to the static biological prior in Section 4.1.

4.6 State estimation

The posterior belief over the patient's state given everything observed so far is:

$\pi_t(x_t) = P(x_t \mid y_{1:t}, g, \text{EHR}, u_{1:t})$

and $\hat{x}_t$ denotes a point summary of this distribution (for example its mean or mode), rather than the distribution itself. Multiple existing estimation approaches are compatible with this formulation, including Bayesian and Kalman/particle filtering, neural state-space models, and temporal transformers, and nothing in the framework requires committing to one; the DLDT proposal discussed in Section 1.7, for instance, instantiates this step with adaptive Kalman filtering on high-frequency EHR data for a specific clinical population [4].

4.7 Counterfactual prediction and treatment selection

$P(x_{t+k} \mid \text{do}(u_t = u)) , \quad u^* = \arg\min_u \; E[L(x_{t+k}, u)]$

using Pearl-style interventional notation; for simplicity this is written for a single intervention at time t rather than a full policy over the interval $[t, t+k]$, and L is left unspecified here but would combine, in any implementation, disease burden, deviation from the individualized homeostatic trajectory, toxicity risk, and estimation uncertainty. The TCPM framework discussed in Section 1.7 is the clearest existing instantiation of counterfactual treatment-effect estimation from routine physiological data at this level of generality, and our framework is explicitly not claiming to improve on its causal-inference machinery; it is asking a complementary question about how the *physiological state itself* should be constructed and constrained mechanistically before such optimization is performed [3].

4.8 Explicit observability and uncertainty as outputs, not afterthoughts

We propose, as this review's own conceptual contribution rather than as established methodology, that the state-estimation step (Section 4.6) should output three objects, not one:

$(\hat{x}_t , U_t , O_t)$

where $\hat{x}_t$ is the point state estimate, U_t is a measure of its uncertainty (for example the spread or entropy of π_t), and O_t is what we term an **operational observability, or confidence, profile**: a per-domain annotation indicating, for each

physiological block in Section 3.1, whether the current estimate rests on direct measurement, indirect inference, or an unvalidated extrapolation.

This use of "observability" should be explicitly distinguished from the term's established meaning in control theory, where observability is a formal, binary-in-the-linear-case property of a dynamical system, whether the internal state x_t can in principle be reconstructed exactly from a finite sequence of outputs y_t and inputs u_t, typically assessed through a rank condition on an observability matrix (for linear systems) or a local observability rank condition (for nonlinear systems). O_t, as proposed here, is not this. It is not a system-level structural property established once from the model equations; it is a proposed per-estimate, per-domain, time-varying annotation intended for clinical communication, analogous in spirit to a data-quality or provenance flag rather than to a control-theoretic guarantee. We are not aware of an established, generally accepted procedure for computing O_t in this sense; it is proposed here as a structural requirement, not a solved estimation problem, and formal control-theoretic observability analysis of any specific implementation of f and h_j would be a separate, complementary exercise. As a purely schematic illustration of the intended distinction, not a computed result: a state estimator might annotate the renal-function component of x_t with high confidence because it is anchored to a directly measured, frequently sampled creatinine trend, while annotating the systemic inflammatory-tone component with low confidence because it is inferred indirectly from heart-rate variability and temperature rather than measured. Formalizing how such confidence annotations should actually be computed and calibrated is left as an open problem in Section 9.

5. The OAD-PDM Conceptual Framework

Building on Sections 3–4, we propose organizing the problem into six layers plus a safety envelope, summarized in Table 1. Each layer individually draws on established or actively developing methodology; the contribution offered here is the explicit coupling between Layers 3 and 4 through the ($\hat{x}_t$, U_t, O_t) interface, and the insistence that Layer 4 remain mechanistically constrained rather than being replaced by an end-to-end predictive model.

Table 1. The OAD-PDM architecture

Layer	Function	Inputs	Output
1. Biological priors	Static individual biology	Genome, pharmacogenomic variants, prior phenotype/EHR history	g
2. Multimodal observation	Heterogeneous, asynchronous sensing	Wearables, molecular/GI biosensors, labs, imaging, TDM, EHR	$y_t^{(j)}$
3. State estimation	Convert observations into biological state, with confidence	Layer 2 + Layer 1	($\hat{x}_t$, U_t, O_t)
4. Mechanistic patient–drug model	Constrain prediction with pharmacological structure	($\hat{x}_t$, U_t, O_t), PK/PD/PBPK/QSP structure	$\theta_i(t)$, predicted exposure/effect

5. Prediction and counterfactual optimization	Forecast trajectories under candidate interventions	Layer 4 outputs	$P(x_{t+k} \mid do(u))$, u^*
6. Clinical decision support	Present recommendation, rationale, and uncertainty to a clinician	Layer 5 outputs	Human-reviewed action
Safety envelope	Verification, validation, uncertainty quantification (VVUQ), out-of-distribution detection	Cross-cutting	Gate on whether Layers 5–6 may act autonomously at all

The safety envelope is not decorative. Tudor and colleagues' 2025 scoping review of purported human digital twins found that of 149 studies claiming digital-twin status, only 18 (12.08%) fully met the National Academies of Sciences, Engineering, and Medicine (NASEM) criteria for a digital twin, personalization, dynamic updating, and predictive capability for clinical decision-making, and only two of the 149 studies mentioned verification, validation, and uncertainty quantification at all [1]. Any framework in this space, including the one proposed here, should be read against that finding: the label "digital twin" or "dynamic patient model" is not self-validating, and we use OAD-PDM as a conceptual architecture name precisely to avoid implying a level of clinical readiness the framework does not have.

The Layer 6 principle, human-in-the-loop rather than autonomous prescribing, follows directly from Layer 3's uncertainty output: when U_t exceeds a pre-specified threshold, or O_t indicates the recommendation rests on a poorly observed domain, the system should defer to standard clinical assessment rather than act.

Figure specifications

Four figures are envisaged for this framework; specifications are provided here for the benefit of a scientific illustrator or diagram-generation tool, since no empirical plots are appropriate for a conceptual framework of this kind.

Figure 1 (OAD-PDM architecture). A six-box vertical or circular flow diagram corresponding exactly to the rows of Table 1, with the biological-prior box feeding into the observation box, both feeding into the state-estimation box, which outputs three explicitly labelled sub-arrows ($\hat{x}_t$, U_t, O_t) into the mechanistic patient–drug model box, which feeds prediction/optimization, which feeds clinical decision support, which feeds back (a dashed return arrow) to the observation box to represent the closed loop. The safety envelope should be rendered as a dashed boundary enclosing Layers 3–6 only, not the whole diagram, to reflect that Layers 1–2 are inputs rather than sites of active inference.

Figure 2 (biological observability gradient). A 4×9 grid or heat-map-style figure with the nine physiological domains from Section 3.1 as rows and the four maturity tiers from Table 3 (Section 3.2) as columns, with each cell populated by the representative examples already given in that table. This figure should not include any effect sizes, percentages, or performance numbers; it is a qualitative maturity map, not a quantitative result.

Figure 3 (mathematical patient–drug dynamical system). A schematic showing the observation equation (Section 4.2) feeding the state-estimation block (Section 4.6), which outputs to both the homeostatic-deviation comparison (Section 4.3) and the mechanistic PK/PD block (Sections 4.4–4.5), which together feed the

counterfactual/optimization block (Section 4.7), closing the loop back to a new observation. Equations should be reproduced verbatim from the corresponding subsections rather than re-derived for the figure.

Figure 4 (technology maturity and translational roadmap). A horizontal timeline or staged-funnel figure corresponding to the six stages in Section 9, annotated with the specific technologies from Section 7 placed at their appropriate stage (for example, closed-loop insulin delivery at "Stage V/VI," hybrid vancomycin dosing at "Stage III," GI capsule sensing at "Stage I"). No projected dates should be included, since none are supported by evidence in this review.

6. Positioning Relative to Existing Frameworks

Table 2 compares OAD-PDM against the closest work identified in Section 1.7, organized by the functional components each addresses. This table is intended to make the review's central claim falsifiable: if an existing framework already covers all rows for a given system, our proposed synthesis adds nothing beyond aggregation, and the review should be read that way.

Table 2. Functional comparison of existing frameworks and the proposed synthesis

Component	Conventional MIPD/TDM	TCPM [3]	DLDT [4]	P4/multi-omics digital-twin perspectives [5,6]	OAD-PDM (this review)
Multimodal continuous observation	Limited (mainly labs/TDM)	Routine physiological indicators	High-frequency EHR	Multi-omics plus digital-twin inputs (conceptual)	Explicit, multi-compartment, asynchronous
Latent state estimation	Implicit (population PK)	Yes (treatment-free physiological profile)	Yes (state-space/Kalman)	Conceptual	Yes, with explicit uncertainty
Explicit observability/confidence output	No	Not emphasized	Not emphasized	Not emphasized	Proposed as a first-class output (this review's contribution)
Mechanistic PK/PD/PBPK coupling	Yes	Not central	Yes (vancomycin-specific)	Not central	Yes, generalized and

					bidirectional (conceptual)
Counterfactual treatment prediction	No	Yes	Partial	Conceptual	Adopts existing causal-inference machinery
Domain scope	Single drug/indication	Six clinical cohorts, general	Critical illness/vancomycin	Broad, aspirational	Cross-domain, explicitly maturity-graded
Clinical validation status	Established	Multi-cohort retrospective validation	Proposed/conceptual	Perspective/conceptual	Conceptual, not validated

Read this way, the honest summary is: TCPM is the strongest existing instantiation of state-to-treatment causal inference from routine data; DLDT is the strongest existing instantiation of mechanistic, state-space precision dosing for a specific critical-care drug problem; and the P4/multi-omics perspectives are the strongest existing statements of the broad aspiration. In the frameworks examined in this review, none makes biological observability and cross-compartment confidence an explicit, first-class output of the state-estimation step, and none generalizes the bidirectional physiology-PK coupling beyond a single drug or a single causal-inference pipeline. That gap, not the general idea of combining AI, sensing, and pharmacology, is what OAD-PDM is proposed to fill conceptually; a broader or differently constructed literature sample might identify additional frameworks that close part of this gap, and the comparison in Table 2 should be read as bounded by the purposive sampling described in Section 2.

7. Illustrative Case Studies

We use three case studies at deliberately different maturity levels to show where the framework's components already have empirical grounding and where they remain aspirational. These are illustrative applications of individual layers or layer-pairs from Section 5, not new experiments, and none should be read as a test of the complete six-layer architecture.

7.1 Case Study I: Closed-loop insulin delivery (mature)

Automated insulin delivery (AID) and hybrid closed-loop systems couple CGM (Layer 2) with an onboard control algorithm (a simplified, disease-specific version of Layers 3–5) to titrate subcutaneous insulin delivery (Layer 6) continuously. This is the field's clearest existing proof that a Sense-Estimate-Predict-Intervene-Observe loop can operate safely and effectively in real patients over sustained periods. Wang and colleagues' 2025 systematic review and meta-analysis, pooling data from 1,156 subjects, reported that, compared with standard control, AI-driven closed-loop systems significantly reduced time outside target glucose range (standardized mean difference 0.90, 95% CI 0.69–1.10, $P < 0.001$) [11]. We do not claim OAD-PDM improves on these systems; we use closed-loop insulin

delivery as the field's existence proof that the general control loop is achievable when (a) the relevant physiological variable is measurable with adequate fidelity and latency and (b) the intervention is finely and safely titratable. Neither condition holds yet for most of the domains in Section 3.2's "emerging" or "latent" tiers, which is precisely why generalizing this success is nontrivial rather than a matter of swapping in a different sensor.

7.2 Case Study II: Vancomycin precision dosing (actively converging)

Vancomycin is a narrow-therapeutic-index antibiotic whose clearance depends heavily on renal function, itself dynamic in critical illness, making it a natural test case for hybrid mechanistic-AI dosing. Multiple recent studies converge on a similar conclusion using different populations and designs. Chen and colleagues developed a hybrid population pharmacokinetic-machine learning (PPK-ML) approach for predicting vancomycin exposure in adult sepsis patients, reporting that the hybrid approach improved on standalone population pharmacokinetic and machine-learning approaches under data-scarce conditions [12]. In critically ill children, a separate group led by Chen developed an ML-population-PK hybrid model specifically to predict pre-dose clearance and volume of distribution before the first dose, aiming to inform initial dosing under exactly the kind of data scarcity that motivates the OAD-PDM observability framing [13]. Shi and colleagues report a related PopPK-informed machine-learning model for vancomycin in a general TDM population, reporting that the model accurately predicted 24-hour area-under-the-curve exposure and proposing it as a point-of-care clinical decision-support tool to improve target attainment and reduce dose-adjustment iterations [14]. Together these studies support the review's core methodological claim in Section 1.5: hybrid mechanistic-ML approaches are not merely theoretically appealing but are actively being developed and reported as outperforming single-paradigm approaches for at least this one well-studied drug, which is consistent with, though does not by itself establish, the more general architecture proposed here.

7.3 Case Study III: Continuous GI sensing and oral drug absorption (experimental)

Absorption is the least observable step of oral pharmacokinetics: clinicians infer it from downstream plasma concentrations rather than measuring it directly. The PillTrek ingestible capsule described in Section 3.2, capable of continuous multiplexed biochemical profiling (electrolytes, metabolites, hormones) within the GI tract, is a concrete step toward direct observation of the absorptive compartment rather than only its consequences [8]. We present this case explicitly as **emerging and experimental**: the reported work was demonstrated in animal studies, not yet validated as a human pharmacokinetic monitoring tool, and no published work known to us has yet coupled such a device to a mechanistic absorption model of the kind sketched in Section 4.4 within a living patient. We include it precisely because it is the clearest current illustration of Layer 2 (observation) technology outpacing Layer 4 (mechanistic coupling), a gap the framework is designed to name rather than one it claims to have closed.

Taken together, these three case studies illustrate different layers and layer-pairs of the architecture in Section 5 at three different maturity levels; none constitutes evidence that the complete six-layer OAD-PDM system functions as an integrated whole, and no such evidence currently exists for any comparably general framework in this space, including the ones discussed in Section 6.

8. Discussion

8.1 What this review adds, stated plainly

We are not claiming to have invented digital twins, AI-based physiological state estimation, continuous biosensing, counterfactual treatment prediction, or hybrid PK/ML dosing. Section 1.7 and Table 2 name the strongest existing work in each of these areas explicitly. The contribution offered here is threefold and comparatively modest: (1) a conceptual argument, grounded in the observability gradient of Section 3, that biological observability and estimation

confidence should be explicit, structured outputs of any patient-state estimator used for pharmacotherapy, not an implicit property left to the model's internals; (2) a common mathematical notation (Section 4) that makes the physiology-pharmacology coupling explicit and bidirectional, rather than treating pharmacokinetic parameters as fixed covariates; and (3) an evidence-mapping exercise (Table 2) intended to make future novelty claims in this space easier to check against what has already been published, given how quickly this literature is moving.

8.2 Why hybrid mechanistic-AI modelling, specifically

The case for hybrid modelling made in Section 1.5 is not merely conceptual. The vancomycin literature reviewed in Section 7.2 reports hybrid PPK-ML models outperforming single-paradigm approaches under realistic data-availability constraints, in more than one independent study and population. The general argument mirrors a broader pattern described across pharmacometric and digital-twin literature: mechanistic models constrain predictions to what is biologically and pharmacologically plausible (non-negative concentrations, mass balance, known elimination pathways) but can struggle to capture patient-specific nonlinearities and heterogeneous data; unconstrained machine learning captures such nonlinearities readily but offers no guarantee of biological plausibility and is vulnerable to distribution shift. Neither failure mode is hypothetical in a pharmacotherapy context, where a physically implausible or poorly calibrated prediction can translate directly into patient harm.

8.3 Multiple timescales

Section 3.4's asynchrony problem has a second dimension beyond sampling rate: different physiological processes are meaningful on structurally different timescales, cardiac electrophysiology in seconds, glycaemic dynamics in minutes, hormonal and drug-effect dynamics in hours, inflammatory and organ-adaptive processes in days, and genomic or disease-trajectory information effectively static over the timescale of a single hospitalization. A state-estimation architecture for this problem needs to represent multiple coupled timescales explicitly rather than forcing every input onto a single update clock. This is a design requirement rather than a novel technical claim: timescale separation and multiscale coupling are already treated formally within quantitative systems pharmacology and multiscale systems-biology modelling, and existing continuous-time and irregularly sampled modelling approaches (state-space models, Bayesian filtering with variable update intervals, neural ODEs) are candidate implementations rather than open research questions in themselves. OAD-PDM's contribution here is limited to insisting that the clinical observation-to-state interface respect this structure, not to proposing a new multiscale-modelling method.

8.4 Translation and validation barriers

The gap between conceptual architecture and deployable system is large and is not primarily a modelling problem. Sensor-level barriers (biofouling, drift, calibration, long-term stability) documented in the wearable-sensing literature reviewed in Section 1.3 apply directly to Layer 2 [10,15]. Data-infrastructure barriers, including fragmented EHR systems, missing and non-standardized longitudinal data, and interoperability, apply to Layer 2's aggregation step. Model-level barriers include the near-total absence of external, prospective validation for most “digital twin” claims documented by Tudor and colleagues' 2025 scoping review discussed in Section 5 [1], and the well-documented risk that machine-learning models trained on one population or hospital system perform unpredictably on another. Each of these is an active, unresolved research area in its own right; none is solved by the architecture proposed here, which instead tries to make explicit where in the pipeline each barrier sits.

8.5 Safety, uncertainty, and the case against autonomy for now

We deliberately did not design OAD-PDM around autonomous prescribing. The Layer 6 human-in-the-loop requirement (Section 5) and the explicit uncertainty/observability output (Section 4.8) both follow from the same

premise: a system whose confidence in its own state estimate cannot be inspected should not be permitted to act on that estimate without review. Closed-loop insulin delivery (Section 7.1) remains, to date, the domain-specific exception where a narrow, well-instrumented, well-understood physiological loop has accumulated enough validated evidence to support a higher degree of automation, and even there, clinical oversight and fail-safes remain standard practice.

8.6 Ethical and governance considerations

Continuous, multimodal biological monitoring raises privacy and data-governance questions that are independent of the modelling questions addressed here: who owns continuously collected physiological and genomic data, how it is secured, and how patients retain meaningful control over its use. The Immorta Bio/Theriome/American Board of Precision Medicine perspective discussed in Section 1.7 addresses this directly by proposing blockchain-based data governance as part of a P4 medicine implementation [5]; we do not evaluate that specific mechanism here, but we note that any deployable version of the architecture in Section 5 would need an explicit governance layer of this kind, not merely a technical one.

9. Future Research Roadmap

We propose six non-simultaneous stages, deliberately ordered from least to most clinically consequential, consistent with the caution argued for in Section 8.5:

1. Reliable multimodal sensing. Longitudinal validation of continuous biosensors (beyond glucose) against reference-standard measurements, and quantification of drift, biofouling, and cross-compartment translation error.

2. Validated state estimation, including a workable definition of O_t. External, prospective validation of latent-state estimators against held-out populations and hospital systems, with uncertainty reported alongside point estimates as standard practice, and initial work toward an operational, calibratable definition of the observability annotation proposed in Section 4.8.

3. Hybrid mechanistic-AI PK/PD models. Extension of the vancomycin-style hybrid approach (Section 7.2) to additional narrow-therapeutic-index drugs and to non-antibiotic domains such as anticoagulation and immunosuppression.

4. Prospective, disease-specific digital twins. Building on existing mechanistic examples (Section 1.7's neonatal aminoglycoside twin, Section 7.2's vancomycin work) rather than attempting a single general-purpose twin first.

5. Clinician-supervised adaptive therapy. Decision-support deployment with mandatory human review, explicit uncertainty display, and prospective outcome tracking.

6. Highly constrained closed-loop systems. Automation restricted, as with insulin delivery today, to narrow, well-validated, well-instrumented physiological loops with demonstrated safety margins, not a general claim of readiness for autonomous pharmacotherapy.

10. Limitations

Of this review. This is a purposive, critical review rather than a formal systematic scoping review; no fixed multi-database protocol was pre-registered, and no PRISMA-style flow counts are reported, for the reasons given in Section 2. The literature evaluated is not exhaustive, and the pace of publication in this area (several of the papers discussed here appeared within the twelve months preceding this review) means the comparative landscape in Table 2 will likely shift quickly. Two references discussed in Section 1.7 and cited throughout, the DLDT paper [4] and the Immorta Bio/Theriome/American Board of Precision Medicine P4 perspective [5], could be confirmed as to title and venue,

and for [4] as to first author, but a complete, independently verified author byline, volume, and page range could not be obtained through the channels available during preparation of this manuscript; both should be checked against the primary journal record before submission. Several claims present in earlier drafts of this project's working notes, including specific participant counts, additional named studies, and several proposed mathematical refinements, could not be independently verified during preparation of this manuscript and have been deliberately omitted rather than reproduced on the basis of an unverified secondary source.

Of the proposed framework. OAD-PDM is conceptual. It has not been implemented as a working system, has not been validated against real patient data, and several of its components, most notably the explicit O_t observability annotation, are proposed structural requirements rather than established methods with a track record or an agreed estimation procedure. The mathematical formulation in Section 4 simplifies substantially: real PK/PD relationships for most drugs are more complex than the illustrative one-compartment model shown, multi-drug interactions are not addressed, and the framework does not resolve, it only names, the compartment-translation problem in Section 3.3.

Of clinical translation. Nothing in this review should be read as evidence that dynamic, AI-assisted, mechanistically constrained pharmacotherapy currently improves patient outcomes beyond what existing MIPD, TDM, and closed-loop systems already achieve in their specific, validated domains. That evidence, where it exists at all, is domain-specific (Section 7) and should not be generalized to the framework as a whole.

11. Conclusion

Continuous biological sensing, AI-based physiological inference, and mechanistic pharmacology are each advancing quickly, and, contrary to a naive reading of the field, they are already being combined in several serious, partially validated ways: temporal causal inference on routine physiological data, state-space digital twins for critical-care precision dosing, and hybrid mechanistic-machine-learning pharmacokinetic models for specific drugs. What remains underdeveloped in the frameworks examined in this review is not the aspiration to combine these tools but a common, biologically honest account of how incomplete, asynchronous, multi-compartment observations should be converted into a patient state whose confidence and observability are explicit, and how that state should be coupled, bidirectionally and mechanistically, to drug behaviour. The Observability-Aware Dynamic Patient–Drug Model proposed here is offered as a conceptual architecture for that problem, not as a validated clinical system. Its value, if any, will be established the way the closed-loop insulin literature establishes value for a narrower claim: through prospective, externally validated, uncertainty-reporting evidence, one well-instrumented physiological loop at a time.

References

1. Tudor BH, Shargo R, Gray GM, Fierstein JL, Kuo FH, Burton R, Johnson JT, Scully BB, Asante-Korang A, Rehman MA, Ahumada LM. A scoping review of human digital twins in healthcare applications and usage patterns. *npj Digital Medicine*. 2025;8(1):587. doi:10.1038/s41746-025-01910-w

2. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, Achenie L, Liu H, Jackson P, Xiao Y, Syeda-Mahmood T, Tuli R, Deng J. Digital twins for health: a scoping review. *npj Digital Medicine*. 2024;7(1):77.

3. Deng Z, Wu W, Zhang C, Zhao X, Pu M, Bu Y, Liao Y, Wang C, Yang J, Wang Y, Wang J. A generalist precision medication framework using temporal causal inference based on treatment-free physiological profiles. *Nature Communications*. 2026;17:6562. doi:10.1038/s41467-026-73238-2

4. Tezel Yalcin H, Yalcin N. Beyond snapshot dosing: a dynamic living digital twin framework for model-informed precision dosing in critical illness. *Expert Review of Clinical Pharmacology*. 2026. doi:10.1080/17512433.2026.2693120

5. Mohr AE, Bajnath A, King G, Ichim TE, Jasbi P. Aspiration to architecture: multi-omics, AI, digital twins, and blockchain for P4 medicine. *Journal of Translational Medicine*. 2026. doi:10.1186/s12967-026-08346-0

6. Emmert-Streib F, Parkkila S, Laubenbacher R, Mannermaa A, Hood L, Yli-Harja O. The role of digital twins in P4 medicine: A paradigm for modern healthcare. *npj Digital Medicine*. 2025;8:735. doi:10.1038/s41746-025-02115-x

7. Prunella M, Romano C, Borri A, Altini N, Di Benedetto MD, Annaert P, Allegaert K, Smits A, Bevilacqua V. Evolutionary digital twin framework for optimal aminoglycoside dosing in neonates with suspected sepsis. *npj Digital Medicine*. 2026;9(1):367. doi:10.1038/s41746-026-02558-w

8. Min J, Ahn H, Lukas H, Ma X, Bhansali R, Sunwoo SH, Wang C, Xu Y, Yao DR, Kim G, Li Z, Hsiai TK, Emami A, Jung HT, Gao W. Continuous biochemical profiling of the gastrointestinal tract using an integrated smart capsule. *Nature Electronics*. 2025;8(9):844–855. doi:10.1038/s41928-025-01407-0

9. Basu S, Legault MA, Romero-Soriano A, Precup D. On the Challenges of using Reinforcement Learning in Precision Drug Dosing: Delay and Prolongedness of Action Effects. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2023;37(12):14102–14109.

10. Güntner AT, Gerber PA, Dittrich PS, Serra N, Figalli A, Puhan MA, Beuschlein F. Challenges and opportunities of wearable molecular sensors in endocrinology and metabolism. *Nature Reviews Endocrinology*. 2026;22:50–60. doi:10.1038/s41574-025-01175-z

11. Wang X, Si J, Li Y, Tse P, Zhang G, Wang X, Ren J, Xu J, Sun J. Effectiveness and safety of AI-driven closed-loop systems in diabetes management: a systematic review and meta-analysis. *Diabetology & Metabolic Syndrome*. 2025;17(1):238. doi:10.1186/s13098-025-01819-0

12. Chen K, Wang C, Wei Y, Ma S, Huang W, Dong Y, Wang Y. Machine learning and population pharmacokinetics: a hybrid approach for optimizing vancomycin therapy in sepsis patients. *Microbiology Spectrum*. 2025;13(5):e00499-25. doi:10.1128/spectrum.00499-25

13. Chen J, Dai L, Xu H, Yao Y, Huang X, Ma J, Liu X, Yao H, Li J, Wu D, Sun J, Yu Z, Li Y, Yang H, Bu S. Machine-learning prediction of pre-dose pharmacokinetics optimizes initial vancomycin dosing in critically ill children. *Frontiers in Pharmacology*. 2026;17:1839423.

14. Shi L, Hu N, Ling J, Yang X, Chen R, et al. Model-informed vancomycin precision dosing by population pharmacokinetics combined with machine learning algorithms. *British Journal of Clinical Pharmacology*. 2026;1–11. doi:10.1002/bcp.70621

15. Duan H, Peng S, He S, Tang SY, Goda K, Wang CH, Li M. Wearable Electrochemical Biosensors for Advanced Healthcare Monitoring. *Advanced Science*. 2025;12(2).

16. Laubenbacher R, Mehrad B, Shmulevich I, Trayanova N. Digital twins in medicine. *Nature Computational Science*. 2024;4:184–191.

17. Moingeon P, Chenel M, Rousseau C, Voisin E, Guedj M. Virtual patients, digital twins and causal disease models: Paving the ground for in silico clinical trials. *Drug Discovery Today*. 2023;28:103605.

18. Jarab AS, Al-Qerem W, Jarab AS, Jrab R, Elsherif F, Al Meslamani AZ, Al Hamarneh YN, Aburuz S. Digital Twins in Precision Pharmacotherapy: Emerging Applications, Challenges, and Future Directions. *Frontiers in Digital Health*. 2026.

19. Silva A, Vale N. Digital Twins in Personalized Medicine: Bridging Innovation and Clinical Reality. *Journal of Personalized Medicine*. 2025;15(11):503.

20. De Domenico M, Allegri L, Caldarelli G, d'Andrea V, Di Camillo B, Rocha LM, Rozum J, Sbarbati R, Zambelli F. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *npj Digital Medicine*. 2025;8:1–11.