*Available online* **www.ijpras.com**

**International Journal of Pharmaceutical Research & Allied Sciences, 2026, 15(2):73-82**
https://doi.org/10.51847/KoWNdomZZK


**Original Article**

**ISSN : 2277-3657**
**CODEN(USA) : IJPRPM**

# *A Foundation Model for Pharmacokinetic and Pharmacodynamic Trajectories with Mechanistic Constraints, Missingness Awareness, and Calibrated Uncertainty*

**Nokuthula Dlamini[1]*, Sipho Zulu[1], Thabo Nkosi[2], Lerato Molefe[3]**

*[1]Department of Foundation Models for PK/PD Trajectories, Faculty of Pharmacy, University of KwaZulu-Natal, Durban, South Africa.*
*[2]Department of Mechanistic Constraints and Missingness Awareness, Faculty of Pharmacy, University of Cape Town, Cape Town, South Africa.*
*[3]Department of Calibrated Uncertainty and Model Validation, Faculty of Pharmaceutical Sciences, University of the Witwatersrand, Johannesburg, South Africa.*

***Email:** dlaminin@ukzn.ac.za*

## ABSTRACT

*Clinical pharmacokinetic and pharmacodynamic modelling supports the interpretation of dose, exposure, biological response, and patient variability, but most models remain tied to a particular drug, population, sampling design, endpoint, or clinical task. This task-specific orientation limits reuse across evolving treatment contexts and does not inherently address informative observation processes, distribution shift, or uncertainty under transfer. This Original Foundation-Model Architecture Article proposes a trajectory foundation model that combines reusable longitudinal representations with explicit pharmacological structure. Patient timelines, administered doses, concentration measurements, pharmacodynamic observations, time-varying covariates, drug and formulation characteristics, and observation-process information are represented through distinct provenance-preserving encoders. These representations inform continuous-time latent pharmacokinetic, exposure, effect-site, and pharmacodynamic states governed by modular hard or soft mechanistic constraints. Missing values, irregular sampling, and measurement intensity are represented as features of the observation process rather than being treated only as data-cleaning problems. The architecture separates aleatoric, epistemic, transfer, and out-of-domain uncertainty and permits selective prediction or abstention when evidence is insufficient. Evaluation is proposed against classical population models, mechanistic PK/PD models, task-specific machine-learning systems, and unconstrained pretrained models across predictive accuracy, calibration, transfer, mechanism consistency, missingness robustness, and decision relevance. Interpretability is grounded in state definitions, provenance, constraint auditing, and traceable adaptation rather than post hoc explanation alone. The contribution is an original conceptual synthesis, not an empirically validated system. Computational verification, independent external evaluation, prospective assessment, uncertainty validation, and context-specific clinical and implementation studies would be required before any decision-facing use.*

**Key words:** *Foundation model, Clinical pharmacokinetics, Pharmacodynamics, PK/PD, Dose–exposure–response, Population modelling*

**Received:** 15 February 2026; **Revised:** 22 May 2026; **Accepted:** 23 May 2026

## INTRODUCTION

Clinical pharmacokinetics and pharmacodynamics connect administered treatment with exposure, biological response, and patient-level variability. Model-informed precision dosing extends this relationship by combining

population knowledge, patient covariates, measured concentrations, and iterative updating, yet routine implementation remains limited by validation, workflow, software, and evidentiary requirements [1-3].

These established approaches should not be conflated. Therapeutic drug monitoring describes the measurement and interpretation of drug concentrations, whereas population pharmacokinetics estimates structural and variability parameters across individuals. PK/PD modelling relates exposure to response, and model-informed therapy uses such models within a defined decision process. None of these concepts is equivalent to a foundation model, a causal model, or a digital twin.

A foundation model is relevant because pharmacological trajectories repeatedly share temporal structures across drugs, populations, and tasks. Dose events precede exposure; exposure may drive delayed or indirect responses; covariates evolve; observations occur irregularly; and future tasks frequently reuse information embedded in earlier clinical histories. Nevertheless, general representation learning cannot by itself establish mass balance, parameter identifiability, causal treatment effects, calibrated transfer, or clinical decision validity.

This article develops an original, non-empirical trajectory foundation-model architecture for PK/PD applications. The proposed contribution integrates longitudinal pretraining, pharmacological representations, mechanistic constraints, observation-process modelling, transfer assessment, calibrated uncertainty, and decision boundaries. It is intended to define testable architectural requirements rather than to claim empirical superiority, clinical effectiveness, regulatory acceptance, or deployment readiness.

*Limitations of task-specific PK/PD models*

Task-specific PK/PD models are usually developed for a defined drug, formulation, population, endpoint, sampling design, and intended use. Their parameters and priors therefore encode assumptions about the development context. When those assumptions no longer match the patient or setting, a fixed prior can constrain inference inappropriately, and selective data-driven adjustment may perform better than unquestioned prior retention [4]. Clinical populations and accumulated evidence also evolve, creating a rationale for controlled model updating rather than permanent reliance on an unchanged model [5].

Transportability cannot be inferred from satisfactory development performance. Published nonlinear mixed-effects models may differ in structural form, covariate relationships, residual-error assumptions, assay context, or population composition. Direct challenges using real-world clinical observations have shown that externally applied population models may not reproduce their original performance [6]. Such failure may reflect model misspecification, population shift, inconsistent dosing records, measurement error, altered treatment practice, or combinations of these factors.

Task-specific development also fragments pharmacological knowledge. Separate models may be built for concentration forecasting, biomarker response, treatment monitoring, adverse-effect prediction, or regimen simulation even when the tasks share dose histories, latent exposure states, covariates, and biological pathways. Repeatedly learning these relationships from limited task-specific datasets can reduce data efficiency and make inconsistencies between models difficult to detect.

A second limitation is that conventional prediction pipelines may treat observation as passive. In clinical care, measurements are often ordered because a patient is deteriorating, a concentration is suspected to be abnormal, a dose has changed, or a clinician requires additional information. The resulting trajectory contains both biological evolution and an observation process. Models that collapse these processes may learn institutional measurement behaviour without distinguishing it from pharmacological dynamics.

**Figure 1** presents the original conceptual synthesis for foundation-model architecture for pk/pd trajectories, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

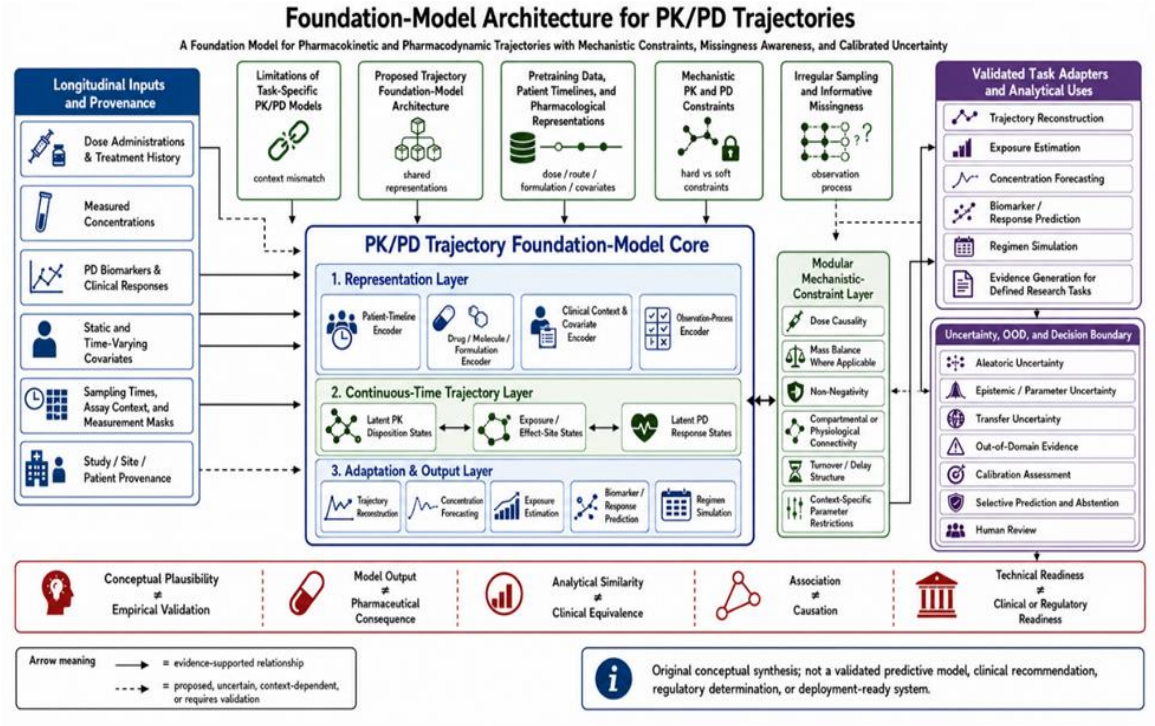


**Figure 1.** Foundation-Model Architecture for PK/PD Trajectories

The figure is an original conceptual synthesis developed for "A Foundation Model for Pharmacokinetic and Pharmacodynamic Trajectories with Mechanistic Constraints, Missingness Awareness, and Calibrated Uncertainty." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

The proposed response to these limitations is not to discard task-specific or mechanistic models. A foundation architecture should instead preserve their strengths while creating reusable representations and explicit interfaces for adaptation. It must permit a classical model to remain preferable when data are sparse, mechanisms are well characterized, transfer is unnecessary, or the intended decision requires transparent parameter interpretation.

*Proposed trajectory foundation-model architecture*

A foundation model is distinguished from an ordinary large model by reusable pretraining and adaptation across downstream tasks rather than by parameter count alone [7]. In the proposed architecture, the reusable object is a temporally indexed representation of patient, drug, dose, exposure, response, context, and observation processes. The architecture is therefore a conditional trajectory model rather than a general text generator or an undifferentiated clinical predictor.

The first layer contains provenance-preserving encoders. Structured clinical events can support pretrained contextual representations [8], while longitudinal transformers demonstrate that clinical concepts and temporal context can be learned jointly [9]. The proposed system extends this principle by assigning separate encoders to patient timelines, drug and formulation information, time-varying covariates, and the observation process.

The central layer is a continuous-time trajectory backbone. It contains distinguishable latent PK disposition states, exposure or effect-site states, and PD response states. Their separation is intended to prevent a generic latent vector from being interpreted automatically as concentration, target engagement, causal biological response, or clinical outcome. Each state must retain links to the observations and assumptions used to estimate it.

A modular constraint layer surrounds the trajectory backbone. It may impose dose causality, temporal ordering, non-negativity, mass balance where applicable, compartmental connectivity, turnover structure, or parameter restrictions. Constraints can be hard when violations are logically or physically inadmissible, or soft when knowledge is incomplete. Their activation must be drug-, formulation-, route-, population-, and task-specific.

Downstream adapters translate the shared representation into defined analytical tasks, including trajectory reconstruction, concentration forecasting, exposure estimation, biomarker-response prediction, or regimen simulation. Each adapter requires separate qualification. A successful pretraining objective does not authorize

every downstream task, and adaptation to one endpoint does not establish validity for another endpoint, population, or clinical consequence.

*Pretraining data, patient timelines, and pharmacological representations*

Longitudinal clinical records can support scalable learning from heterogeneous events without requiring every variable to be converted into a manually engineered feature [10]. Reusable sequence representations may then support multiple downstream tasks [11]. For PK/PD pretraining, however, the timeline must preserve administered dose, route, formulation, infusion duration, concentration time, biomarker time, interventions, covariates, assay context, and source provenance.

A patient timeline is defined here as an ordered sequence of time-stamped clinical and pharmacological events rather than a simple matrix of visits. Prescribed treatment must be distinguished from administered treatment, and a scheduled sample must be distinguished from an obtained result. Static characteristics, time-varying covariates, treatment changes, organ-function measures, concomitant therapy, and clinical-state transitions require explicit temporal alignment.

The drug representation should integrate identity with pharmacologically relevant context. Knowledge-guided molecular pretraining shows that domain knowledge can be introduced during representation learning rather than relying exclusively on unconstrained molecular tokens [12]. The proposed drug encoder therefore combines molecular or physicochemical information, route, formulation, dose units, target or pathway knowledge, and prior mechanistic structure where these data are available.

Joint patient–drug representations are proposed to condition the trajectory backbone, but they should not erase clinically important distinctions. Pretraining data may overrepresent common medicines, frequently monitored patients, large institutions, or dense measurement periods. Provenance, dataset partitioning, temporal separation, duplicate-patient control, and explicit documentation of unsupported drugs or populations are consequently architectural requirements rather than optional reporting details.

*Mechanistic PK and PD constraints*

Neural components can be embedded within continuous-time, pharmacokinetic, or mechanism-organized structures rather than replacing all established pharmacological structure [13-15]. This hybrid position is central to the proposed architecture because reusable representation and mechanistic discipline perform different functions.

A hard constraint prohibits an inadmissible state or transition, whereas a soft constraint penalizes departures from an expected relationship. Dose causality and non-negativity may justify strong restrictions; compartmental structure, turnover assumptions, or parameter relationships may require softer, context-dependent treatment.

Constraints should be applied through modular registries specifying their source, scope, mathematical form, affected states, and failure conditions. PK constraints may govern disposition and exposure, while PD constraints may govern effect-site delay, turnover, saturation, tolerance, or disease-linked response.

Mechanistic consistency is not equivalent to mechanistic truth. An incorrect compartment, omitted pathway, erroneous parameter restriction, or inappropriate transfer assumption can produce plausible but misleading trajectories. Every constrained model therefore requires comparison with unconstrained and classical alternatives and explicit challenge under mechanism misspecification.

*Irregular sampling and informative missingness*

Irregular observations should not be converted automatically into equally spaced, fully observed trajectories. Values, measurement masks, and elapsed times can be represented separately so that the model retains information about what was observed and when [16].

Clinical observation may be informative because measurement intensity depends on disease severity, treatment change, clinician concern, or expected instability [17]. A concentration obtained after suspected toxicity does not arise from the same observation process as a routine scheduled sample.

Missing laboratory patterns may themselves contain contextual information [18]. The proposed architecture therefore includes an observation-process encoder for scheduled, completed, delayed, cancelled, unavailable, censored, or below-quantification measurements, with reasons recorded when known.

This approach does not make missingness ignorable. Unobserved decisions, undocumented dose administration, assay limitations, and site-specific workflows may still bias learned representations. Robustness must be evaluated under altered masking, sampling frequency, measurement policy, and documentation quality.

**Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in a foundation model for pharmacokinetic and pharmacodynamic trajectories with.

**Table 1.** Definitions, components, relationships, and intended uses in A Foundation Model for Pharmacokinetic and Pharmacodynamic Trajectories with.

| Component or scientific dimension | Problem addressed | Inputs or determinants | Proposed mechanism or relationship | Expected contribution | Evidence required | Failure or uncertainty risk | Boundary statement |
|---|---|---|---|---|---|---|---|
| **Patient timeline** | Fragmented longitudinal information | Dose, concentration, biomarker, covariate, intervention, time and provenance | Proposed ordered event representation | Reusable temporal context | Temporal reconstruction and external testing | Misordered or duplicated events | A timeline is not a causal patient state |
| **Pharmacological representation** | Drug identity alone is insufficient | Molecule, route, formulation, dose and prior knowledge | Proposed alignment with patient context | Cross-drug representation | Ablation and transfer evaluation | Irrelevant molecular similarity | Representation does not establish in-vivo equivalence |
| **Continuous PK/PD state** | Discrete observations incompletely describe dynamics | Encoded timelines and drug context | Continuous-time latent evolution [13] | Trajectory reconstruction and forecasting | Numerical and predictive validation | Non-identifiability | Latent state is not a measured concentration |
| **Mechanistic constraints** | Unconstrained trajectories may be implausible | Dose history and justified PK/PD structure | Hard or soft restrictions [14] | Reduced inadmissible behaviour | Constraint and misspecification tests | Incorrect mechanism | Constraint satisfaction is not biological proof |
| **Mechanism-organized connectivity** | Generic networks obscure pathway structure | Known state relationships | Structured neural connectivity [15] | Greater traceability | Comparison with alternative organizations | False mechanistic reassurance | Network organization does not establish causation |
| **Observation-process representation** | Sampling and missingness may be informative | Mask, elapsed time, ordering and missingness reason | Parallel observation encoder [16] | Missingness-aware inference | Policy-shift and masking tests | Workflow shortcut learning | Missingness modelling does not remove selection bias |
| **Informative observation** | Measurements depend on clinical context | Severity, treatment change and clinician behaviour | Observation intensity linked to patient state [17] | Better distinction between biology and measurement | Cross-site validation | Unmeasured decision process | Association with observation is not treatment effect |
| **Missing laboratory pattern** | Absent values may carry information | Laboratory availability and measurement history | Context-dependent missingness representation [18] | Improved robustness | Sensitivity analysis | Institution-specific artefact | Missingness cannot substitute for measurement |
| **Task adapter** | One representation serves different analyses | Shared trajectory state and task specification | Proposed task-specific adaptation | Controlled reuse | Task-level qualification | Negative transfer | Validation for one task does not authorize another |
| **Decision boundary** | Prediction may be converted | Output, uncertainty | Proposed evidence and | Prevention of unsupported dosing use | Prospective and implementation evidence | Automation bias | Analytical output is not a |

| | | | | |
|---|---|---|---|---|
| | prematurely into action | and intended use | human-review gate | clinical recommendation |

*Transfer across drugs, populations, and clinical tasks*

Transfer must be defined by the dimension that changes. Temporal transfer concerns later practice periods; population transfer concerns patient characteristics; drug transfer concerns pharmacological entities; and task transfer concerns the prediction or simulation objective.

Pretraining may improve robustness when clinical distributions change over time, but such improvement does not establish transfer to unseen drugs, mechanisms, routes, or populations [19]. Each transfer claim must therefore specify its source and target domains.

External evaluation of population-PK models remains heterogeneous and must be matched to the intended precision-dosing use [20]. The proposed model should be tested through held-out drugs, studies, institutions, regimens, populations, and time periods rather than random record splitting alone.

Prospective validation provides stronger evidence than retrospective evaluation but remains conditional on its drug, population, workflow, and decision context [21]. Successful adaptation should therefore be interpreted locally unless broader transportability is demonstrated independently.

*Calibration, uncertainty, and out-of-domain detection*

Uncertainty should be communicated in forms that support interpretation rather than as an unexplained confidence score [22]. The architecture separates aleatoric uncertainty from epistemic, parameter, transfer, data-quality, and out-of-domain uncertainty.

Uncertainty quantification is necessary when model outputs may inform decisions under incomplete knowledge [23]. It is not sufficient, however, unless uncertainty estimates are calibrated for the relevant outcome, horizon, subgroup, and use context.

Out-of-domain detection should evaluate both obvious and near-distribution shifts. Medical tabular benchmarks show that difficult, clinically plausible shifts may not be identified reliably by generic detection methods [24].

The proposed system should permit abstention when uncertainty or domain mismatch exceeds the evidence supporting interpretation. Thresholds must be established empirically for a defined use and cannot be inferred from architecture alone.

**Table 2** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for a foundation model for pharmacokinetic and pharmacodynamic trajectories with.

**Table 2.** Testable propositions, evidence requirements, and validation criteria for A Foundation Model for Pharmacokinetic and Pharmacodynamic Trajectories with.

| Architecture layer or process stage | Core function | Information or material flow | Interaction with other components | Validation criterion | Potential failure mode | Human or experimental responsibility | Readiness boundary |
|---|---|---|---|---|---|---|---|
| **Multidomain pretraining** | Learn reusable trajectory structure | Patient and drug histories to shared representation | Supplies downstream adapters | Better data efficiency under controlled transfer | Memorization or leakage | Curate partitions and provenance | Pretraining alone is research evidence |
| **Patient–drug alignment** | Condition trajectories jointly | Patient context plus pharmacological representation | Feeds continuous state model | Added value over patient-only and drug-only models | Spurious similarity | Design ablations | Alignment does not prove equivalent exposure |
| **Mechanistic constraint layer** | Restrict selected dynamics | Assumptions to latent-state evolution | Modifies backbone behaviour | Fewer inadmissible trajectories without degraded validity | Misspecified mechanism | Pharmacometric review and challenge testing | Plausibility is not validation |
| **Observation-process layer** | Represent sampling and missingness | Masks, timing and measurement context | Informs state and uncertainty modules | Robustness under altered | Workflow shortcut learning | Audit collection processes | Missingness awareness does not |

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | | | | observation policies | | | identify unseen states |
| **Transfer evaluation** | Test adaptation across domains | Source representation to target task or population | Challenges backbone and adapters | Preserved calibration and performance under declared shift [19] | Negative transfer | Define source and target domains | No transfer claim beyond tested domains |
| **External validation** | Test independent data | Model to new site, population or study | Evaluates transportability [20] | Reproducible performance and calibration | Hidden compatibility or data leakage | Independent investigators | Internal validation is insufficient |
| **Prospective assessment** | Evaluate incoming data and workflow | Model output to silent or supervised use | Tests operational interaction [21] | Reliable operation without unsupported influence | Workflow mismatch | Clinical and implementation teams | Retrospective success is not clinical readiness |
| **Uncertainty module** | Quantify and communicate uncertainty | Predictions to intervals, risk or abstention | Interacts with OOD and decision gates [22] | Calibration and selective-risk performance | Miscalibrated confidence | Statistical validation and user testing | Confidence is not authorization |
| **OOD module** | Identify unsupported inputs | Incoming case to domain evidence | May trigger abstention [24] | Detection of predefined shifts | Near-domain failure | Construct challenge sets | OOD detection cannot guarantee safety |
| **Decision gate** | Restrict consequence of outputs | Evidence and uncertainty to human review | Final interface with use context | Context-specific benefit, safety and monitoring evidence | Automation bias | Authorized human decision-maker | No autonomous dosing without validation |

*Evaluation against classical and mechanistic models*

Evaluation should compare the proposed model with classical nonlinear mixed-effects models, mechanistic PK/PD models, task-specific machine-learning systems, and an unconstrained pretrained trajectory model. Comparator selection must reflect the actual analytical question rather than favouring one architecture.

Discrimination or average prediction error alone cannot establish medical decision usefulness. Calibration, clinically relevant error, classification consequences, and decision-analytic performance must be matched to the intended use [25].

Mechanistic evaluation should test dose causality, non-negativity, regimen extrapolation, parameter plausibility, and behaviour under deliberately misspecified constraints. Validation depth should be proportional to the model's context, risk, and decision consequence rather than governed by one universal checklist [26].

A staged programme should include computational verification, reproducibility, internal validation, external validation, prospective silent assessment, and clinical-impact evaluation when decision use is proposed. Translation into dosing software additionally requires validation of model selection, implementation, input handling, and software-facing operation [27].

Success must be multidimensional. Superior trajectory prediction cannot compensate for failed calibration, implausible dynamics, unsupported transfer, poor provenance, or unsafe interaction with clinical workflow. Conversely, a classical mechanistic model may remain preferable when interpretability, data efficiency, or established validation outweigh broader representation capacity.

*Interpretability, governance, and decision boundaries*

Interpretability should begin with state definitions, provenance, constraint visibility, adapter scope, and uncertainty decomposition. Post hoc explanations alone may be unstable or misleading and cannot substitute for model validity [28].

Every output should retain traceable links to dose history, observations, covariates, drug representation, active constraints, model version, adaptation data, and uncertainty state. Such traceability supports auditing but does not prove that the model's internal representation corresponds to the true biological mechanism.

Lifecycle governance should address fairness, universality, traceability, usability, robustness, and explainability [29]. For this architecture, those principles require drug- and population-specific performance review, version control, change assessment, monitoring, and documented human responsibility.

A strict boundary must separate research output from pharmaceutical or clinical consequence. Concentration forecasts, simulated responses, latent trajectories, and uncertainty estimates cannot be converted automatically into dose changes, treatment selection, or regulatory claims without context-specific validation and authorization.

*Limitations and research priorities*

Foundation-model research can be weakened by inconsistent datasets, poorly matched tasks, inadequate baselines, and non-comparable evaluation practices [30]. PK/PD development is especially vulnerable because leakage may occur across patients, studies, repeated dosing episodes, or temporally related records.

The architecture may also fail when drug representations omit determinants of in-vivo disposition, when mechanistic constraints are wrong, when latent states are non-identifiable, or when uncertainty is calibrated only within the development domain. Multidrug pretraining may produce negative transfer rather than reusable pharmacological structure.

A longitudinal trajectory model should not be described automatically as a digital twin. Digital-twin claims require validated dynamic correspondence, updating, multiscale representation, and a defensible relationship between the computational system and the biological system it represents [31].

**Figure 2** presents the original conceptual synthesis for evaluation framework for transfer, mechanism consistency, and uncertainty, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

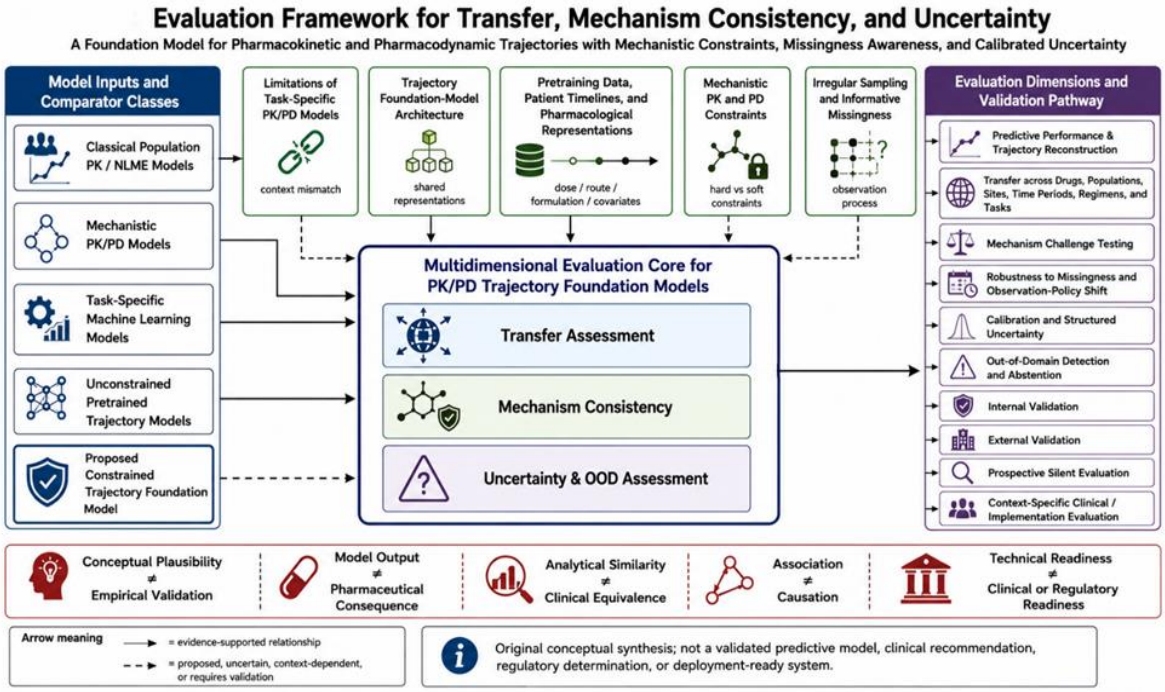


**Figure 2.** Evaluation Framework for Transfer, Mechanism Consistency, and Uncertainty

The figure is an original conceptual synthesis developed for "A Foundation Model for Pharmacokinetic and Pharmacodynamic Trajectories with Mechanistic Constraints, Missingness Awareness, and Calibrated Uncertainty." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

Research priorities include pharmacologically meaningful pretraining objectives, independent multidrug benchmarks, observation-policy challenge sets, constraint-misspecification experiments, uncertainty evaluation

under near-domain shift, prospective silent studies, and comparisons with strong classical models. These studies must identify when the proposed architecture fails, not merely when it appears advantageous.

## CONCLUSION

This Original Foundation-Model Architecture Article proposes a reusable framework for representing PK/PD trajectories while preserving patient timelines, pharmacological context, mechanistic constraints, observation processes, transfer conditions, and structured uncertainty. Its central contribution is the separation of shared representation from mechanism, task adaptation, evidence qualification, and decision authorization. The architecture remains a conceptual synthesis. It does not establish superior prediction, causal treatment response, clinical utility, regulatory acceptance, universal transportability, or deployment readiness. Its value must be determined through transparent comparator studies, mechanism challenges, external and prospective validation, calibrated uncertainty assessment, and context-specific evaluation of any proposed decision use.

**ACKNOWLEDGMENTS:** None

**CONFLICT OF INTEREST:** None

**FINANCIAL SUPPORT:** None

**ETHICS STATEMENT:** None

## REFERENCES

1. 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
2. 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
3. Wicha SG, Märtson AG, Nielsen EI, Koch BCP, Friberg LE, Alffenaar JWC, et al. From therapeutic drug monitoring to model-informed precision dosing for antibiotics. Clin Pharmacol Ther. 2021;109(4):928-41. doi:10.1002/cpt.2202
4. Hughes JH, Keizer RJ. A hybrid machine learning/pharmacokinetic approach outperforms maximum a posteriori Bayesian estimation by selectively flattening model priors. CPT Pharmacometrics Syst Pharmacol. 2021;10(10):1150-60. doi:10.1002/psp4.12684
5. Maier C, de Wiljes J, Hartung N, Kloft C, Huisinga W. A continued learning approach for model-informed precision dosing: Updating models in clinical practice. CPT Pharmacometrics Syst Pharmacol. 2022;11(2):185-98. doi:10.1002/psp4.12745
6. Kluwe F, Michelet R, Huisinga W, Zeitlinger M, Mikus G, Kloft C. Towards model-informed precision dosing of voriconazole: Challenging published voriconazole nonlinear mixed-effects models with real-world clinical data. Clin Pharmacokinet. 2023;62(10):1461-77. doi:10.1007/s40262-023-01274-y
7. Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65. doi:10.1038/s41586-023-05881-4
8. Rasmy L, Xiang Y, Xie Z, Tao C, Zhi D. Med-BERT: Pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. NPJ Digit Med. 2021;4(1):86. doi:10.1038/s41746-021-00455-y
9. Li Y, Rao S, Solares JRA, Hassaine A, Ramakrishnan R, Canoy D, et al. BEHRT: Transformer for electronic health records. Sci Rep. 2020;10(1):7155. doi:10.1038/s41598-020-62922-y
10. Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, et al. Scalable and accurate deep learning with electronic health records. NPJ Digit Med. 2018;1(1):18. doi:10.1038/s41746-018-0029-1
11. Steinberg E, Jung K, Fries JA, Corbin CK, Pfohl SR, Shah NH. Language models are an effective representation learning technique for electronic health record data. J Biomed Inform. 2021;113:103637. doi:10.1016/j.jbi.2020.103637

12. Li H, Zhang R, Min Y, Ma D, Zhao D, Zeng J. A knowledge-guided pre-training framework for improving molecular representation learning. Nat Commun. 2023;14(1):7568. doi:10.1038/s41467-023-43214-1
13. Lu J, Deng K, Zhang X, Liu G, Guan Y. Neural-ODE for pharmacokinetics modeling and its advantage to alternative machine learning models in predicting new dosing regimens. iScience. 2021;24(7):102804. doi:10.1016/j.isci.2021.102804
14. Valderrama D, Ponce-Bobadilla AV, Mensing S, Fröhlich H, Stodtmann S. Integrating machine learning with pharmacokinetic models: Benefits of scientific machine learning in adding neural network components to existing PK models. CPT Pharmacometrics Syst Pharmacol. 2024;13(1):41-53. doi:10.1002/psp4.13054
15. Mann J, Meshkin H, Zirkle J, Han X, Thrasher B, Chaturbedi A, et al. Mechanism-based organization of neural networks to emulate systems biology and pharmacology models. Sci Rep. 2024;14(1):12082. doi:10.1038/s41598-024-59378-9
16. Che Z, Purushotham S, Cho K, Sontag D, Liu Y. Recurrent neural networks for multivariate time series with missing values. Sci Rep. 2018;8(1):6085. doi:10.1038/s41598-018-24271-9
17. Sisk R, Lin L, Sperrin M, Barrett JK, Tom B, Diaz-Ordaz K, et al. Informative presence and observation in routine health data: A review of methodology for clinical risk prediction. J Am Med Inform Assoc. 2021;28(1):155-66. doi:10.1093/jamia/ocaa242
18. Tan ALM, Getzen EJ, Hutch MR, Strasser ZH, Gutiérrez-Sacristán A, Le TT, et al. Informative missingness: What can we learn from patterns in missing laboratory data in the electronic health record? J Biomed Inform. 2023;139:104306. doi:10.1016/j.jbi.2023.104306
19. Guo LL, Steinberg E, Fleming SL, Posada J, Lemmon J, Pfohl SR, et al. EHR foundation models improve robustness in the presence of temporal distribution shift. Sci Rep. 2023;13(1):3767. doi:10.1038/s41598-023-30820-8
20. El Hassani M, Marsot A. External evaluation of population pharmacokinetic models for precision dosing: Current state and knowledge gaps. Clin Pharmacokinet. 2023;62(4):533-40. doi:10.1007/s40262-023-01233-7
21. ter Heine R, Keizer RJ, van Steeg K, Smolders EJ, van Luin M, Derijks HJ, et al. Prospective validation of a model-informed precision dosing tool for vancomycin in intensive care patients. Br J Clin Pharmacol. 2020;86(12):2497-506. doi:10.1111/bcp.14360
22. Kompa B, Snoek J, Beam AL. Second opinion needed: Communicating uncertainty in medical machine learning. NPJ Digit Med. 2021;4(1):4. doi:10.1038/s41746-020-00367-3
23. Begoli E, Bhattacharya T, Kusnezov D. The need for uncertainty quantification in machine-assisted medical decision making. Nat Mach Intell. 2019;1(1):20-23. doi:10.1038/s42256-018-0004-1
24. Azizmalayeri M, Abu-Hanna A, Cinà G. Unmasking the chameleons: A benchmark for out-of-distribution detection in medical tabular data. Int J Med Inform. 2025;195:105762. doi:10.1016/j.ijmedinf.2024.105762
25. Van Calster B, Collins GS, Vickers AJ, Wynants L, Kerr KF, Barreñada L, et al. Evaluation of performance measures in predictive artificial intelligence models to support medical decisions: Overview and guidance. Lancet Digit Health. 2025;7(12):100916. doi:10.1016/j.landig.2025.100916
26. Lemaire V, Hu C, van der Graaf PH, Chang S, Wang W. No recipe for quantitative systems pharmacology model validation, but a balancing act between risk and cost. Clin Pharmacol Ther. 2024;115(1):25-28. doi:10.1002/cpt.3082
27. Taylor ZL, Poweleit EA, Paice K, Somers KM, Pavia K, Vinks AA, et al. Tutorial on model selection and validation of model input into precision dosing software for model-informed precision dosing. CPT Pharmacometrics Syst Pharmacol. 2023;12(12):1827-45. doi:10.1002/psp4.13056
28. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11). doi:10.1016/S2589-7500(21)00208-9
29. Lekadir K, Feragen A, Fofanah AJ, Frangi AF, Buyx A, Emelie A, et al. FUTURE-AI: International consensus guideline for trustworthy and deployable artificial intelligence in healthcare. BMJ. 2025;388. doi:10.1136/bmj-2024-081554
30. Wornow M, Xu Y, Thapa R, Patel B, Steinberg E, Fleming SL, et al. The shaky foundations of large language models and foundation models for electronic health records. NPJ Digit Med. 2023;6(1):135. doi:10.1038/s41746-023-00879-8
31. 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