**Pharmacophore**

ISSN-2229-5402

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

# WHAT MACHINE LEARNING ADDS TO PHARMACOMETRICS—AND WHAT MECHANISTIC INFORMATION IT RISKS REMOVING

**Ying Li[1]*, Hui Wang[2], Xiaoming Chen[3], Jun Zhang[4]**

1. ***Department of Machine Learning in Pharmacometrics, Faculty of Pharmacy, Yunnan Agricultural University, Kunming, China.***
2. ***Department of Mechanistic Information and PK/PD Modeling, Faculty of Pharmaceutical Sciences, Zhejiang University, Hangzhou, China.***
3. ***Department of ML Contributions to Pharmacometrics, Faculty of Pharmacy, Nanjing Agricultural University, Nanjing, China.***
4. ***Department of Risk of Mechanistic Removal in ML Models, Faculty of Pharmacy, Ningxia University, Yinchuan, China.***

**ARTICLE INFO**

**Received:**
28 April 2025
**Received** in revised form:
10 August 2025
**Accepted:**
16 August 2025
**Available online:**
28 August 2025

***Keywords:*** Pharmacometrics, Machine learning, Pharmacokinetics, Pharmacodynamics, Scientific machine learning, Population pharmacokinetics

**ABSTRACT**

Machine learning is increasingly used alongside pharmacokinetic, pharmacodynamic, population, exposure–response, and systems-pharmacology models, but its contribution is frequently judged through predictive performance without equivalent attention to the scientific information altered by the modeling choice. This integrative critical review examines where machine learning adds genuine value to pharmacometrics, where it merely substitutes one representation for another, and where it may remove mechanistic constraints needed for explanation, extrapolation, or decision credibility. The selected literature was compared according to pharmaceutical task, data suitability, model structure, mechanistic information content, validation design, uncertainty treatment, interpretability, and intended context of use. The synthesis distinguishes mechanistic models, predominantly data-driven models, and hybrid or scientific machine-learning architectures. Machine learning can expand nonlinear function estimation, covariate discovery, concentration prediction, exposure–response characterization, and computational efficiency. These capabilities are most persuasive when the task is bounded, the prediction domain is adequately represented, and the consequences of error are limited or reversible. However, predictive flexibility can obscure compartmental meaning, physiological constraints, parameter interpretation, causal assumptions, and the basis for extrapolation. Hybridization may preserve selected mechanistic structures while learning unknown components, but the label "hybrid" does not itself establish mechanistic validity or regulatory credibility. The article proposes a comparative analytical framework based on an explicit mechanistic-information ledger and task-dependent evidence requirements. The available evidence remains limited by heterogeneous validation practices, restricted external evaluation, incomplete calibration assessment, and sparse prospective demonstration. Accountable integration therefore requires model choice to follow the pharmaceutical question, not technological preference, and requires claims to remain proportionate to the information retained and the validation actually performed.



**To Cite This Article:** Li Y, Wang H, Chen X, Zhang J. What Machine Learning Adds to Pharmacometrics—and What Mechanistic Information It Risks Removing. Pharmacophore. 2025;16(4):91-101. https://doi.org/10.51847/maxWWV33ys

## Introduction

Pharmacometrics and machine learning increasingly occupy overlapping territory in pharmaceutical science. Both can estimate drug concentrations, characterize variability, identify influential covariates, support dose selection, and connect exposure with response. Their apparent convergence, however, can conceal substantial differences in what the resulting models represent and how their predictions should be interpreted. Pharmacometric models typically encode assumptions about compartments, physiological processes, parameter distributions, dose timing, disease progression, or pharmacological response. Machine-learning models instead prioritize the estimation of predictive relationships from available data, often with fewer predefined biological constraints. Recent reviews therefore describe machine learning as a potentially valuable extension to

**Corresponding Author:** Ying Li; Department of Machine Learning in Pharmacometrics, Faculty of Pharmacy, Yunnan Agricultural University, Kunming, China. E-mail: ying.li@ynau.edu.cn.

pharmacometrics while also emphasizing that implementation remains heterogeneous, terminology is inconsistent, and methodological advantages are conditional on the task, dataset, and validation strategy [1-3].

The central difficulty is that predictive equivalence does not establish scientific equivalence. A population pharmacokinetic model and a machine-learning model may both reproduce observed concentration–time profiles, yet they may rely on fundamentally different sources of information. The pharmacometric model may attribute variation to clearance, volume, absorption, covariates, and interindividual random effects, whereas the machine-learning model may encode the same variation within a distributed function that lacks directly interpretable pharmacological parameters. Comparative work illustrates that similar performance on an observed-data task does not guarantee equivalent behavior under new doses, schedules, populations, or sampling conditions [4]. The practical question is therefore not simply whether machine learning predicts well, but whether the information discarded by the learned representation is required for the intended decision.

Scientific machine learning and other hybrid strategies have emerged partly in response to this tension. These approaches retain selected differential equations, compartmental structures, or population assumptions while assigning unknown, misspecified, or difficult-to-parameterize components to learned functions. Comparative evidence suggests that such architectures can combine mechanistic organization with data-driven flexibility, but their advantages are neither automatic nor universal [5]. A learned component may correct a genuine knowledge gap, yet it may also absorb structural misspecification, measurement artifacts, unrecognized confounding, or dataset-specific relationships. Consequently, the presence of a mechanistic scaffold does not guarantee that the complete model remains mechanistically interpretable.

This review addresses the unresolved question of how machine learning should be judged when used within or alongside pharmacometrics. Its purpose is not to defend one modeling tradition or to rank algorithms by benchmark performance. Instead, it evaluates what different approaches add, what mechanistic information they preserve or remove, and what forms of evidence are required before their outputs can support pharmaceutical decisions. The article introduces a comparative framework centered on task definition, data suitability, information retention, extrapolation basis, uncertainty, and context-of-use credibility. Within that framework, association is distinguished from mechanism, prediction from explanation, retrospective fit from transportability, and conceptual promise from routine pharmaceutical readiness.

*Review Scope and Comparative Analytical Framework*

The review was designed as an integrative critical synthesis of heterogeneous methodological, applied, translational, and regulatory scholarship. An integrative review is appropriate when the objective is to compare concepts, model classes, evidence traditions, and interpretation boundaries rather than to estimate a pooled effect. Such reviews require explicit definition of the problem, transparent literature identification, critical comparison of eligible evidence, and synthesis at the level of concepts rather than sequential description of publications [6]. Accordingly, the analytical unit used here is the relationship among the pharmaceutical task, the structure of the model, the information used for estimation, the information retained after fitting, the validation performed, and the consequence of an incorrect result. This structure allows evidence from population pharmacokinetics, exposure–response analysis, physiologically based modeling, neural differential equations, precision dosing, and broader pharmaceutical machine learning to be compared without implying that these domains are methodologically interchangeable.

The scope is intentionally critical but not falsely exhaustive. Review labels should correspond to the actual purpose and coverage of the evidence-identification process, and broad conceptual questions should not be presented as if they were answered by a narrowly defined meta-analysis or by a complete census of all published studies [7]. Eligible evidence was therefore assessed for direct relevance to at least one planned claim concerning pharmacokinetic or pharmacodynamic prediction, covariate modeling, population integration, hybrid architecture, mechanistic information, extrapolation, interpretability, uncertainty, model credibility, or task-dependent pharmaceutical use. Evidence was excluded from the conceptual synthesis when it supported only a narrow benchmark, lacked sufficient methodological detail, or encouraged equivalence between predictive accuracy and mechanistic, experimental, clinical, or regulatory validity. Where an article supported only a broader concept, the manuscript restricts the wording of the associated claim accordingly.

The comparative framework uses three model traditions. The first is mechanistic pharmacometrics, in which biological, physiological, pharmacological, or population assumptions are represented explicitly through states, parameters, structural equations, and variability models. The second is predominantly data-driven machine learning, in which the principal relationship is estimated from observed inputs and outcomes with limited predefined mechanistic structure. The third is hybrid or scientific machine learning, in which selected known relationships remain explicit and selected unknown relationships are learned. These traditions are compared across capability, information content, extrapolation basis, uncertainty, interpretability, and context of use. Model-informed drug-development practice further requires that evaluation be proportional to the model's role in a decision, because the same predictive error has different implications in exploratory screening, trial design, exposure–response interpretation, and patient-level dosing [8]. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop review scope and comparative analytical framework within the article's central argument.

**Table 1.** Evidence domains, core questions, scientific requirements, and interpretation boundaries for Review scope and comparative analytical framework

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

| | | | | | |
|---|---|---|---|---|---|
| **Integrative-review methodology** | Methodological guidance on integrative, systematic, and scoping review design | Heterogeneous evidence can be synthesized through explicit questions, eligibility logic, structured extraction, and critical comparison | Established methodological support | Search terminology, disciplinary boundaries, and reviewer judgment can affect coverage | No single review method resolves inconsistent terminology across pharmacometrics and machine learning |
| **Mechanistic pharmacometrics** | Population PK/PD, exposure–response, PBPK, QSP, and model-informed development literature | Explicit states, parameters, variability structures, and biological assumptions support interpretation and structured simulation | Strong for established pharmacometric applications | Structural assumptions may be misspecified; parameter identifiability may be weak | Mechanistic form does not guarantee mechanistic correctness |
| **Predominantly data-driven machine learning** | Supervised learning, neural networks, ensembles, representation learning, and direct prediction models | Flexible estimation of nonlinear, interacting, and high-dimensional relationships | Moderate to strong for bounded predictive tasks | Data leakage, unstable feature selection, opaque representations, and limited external testing | Conditions under which predictive gains remain reliable under distribution shift |
| **Hybrid and scientific machine learning** | Residual models, neural differential equations, learned covariate functions, parameter-estimation modules, and mechanistic surrogates | Known structure can be retained while uncertain or difficult relationships are learned | Emerging and context dependent | The learned component may absorb model misspecification or dominate system behavior | Standard methods for testing whether claimed mechanistic content is genuinely retained |
| **Data suitability** | Studies reporting data source, sampling design, variable quality, missingness, outcome definition, and representativeness | Appropriate data can support stable estimation within a defined task | Variable across application domains | Data availability may be mistaken for suitability; historical treatment patterns may become learned shortcuts | Prospective criteria for determining whether available data represent intended use |
| **Predictive evaluation** | Internal validation, external validation, calibration, error analysis, and sensitivity analysis | Predictive adequacy can be characterized for a specified population and endpoint | Strong for conventional internal metrics; weaker for broad transportability | Random train–test splits may overstate generalization when correlated observations are separated | Standard pharmacometric stress tests for new regimens, populations, compounds, and measurement processes |
| **Mechanistic-information assessment** | Parameter interpretation, state definition, conservation constraints, physiological plausibility, causal assumptions, and counterfactual behavior | Model classes differ in what biological and pharmacological information remains explicit | Conceptually strong but inconsistently operationalized | Interpretability may be asserted from model form rather than demonstrated | Validated procedures for auditing information preserved, obscured, or removed |
| **Uncertainty and credibility** | Parameter uncertainty, predictive uncertainty, structural uncertainty, calibration, applicability, and context-of-use assessment | Credibility depends on evidence appropriate to model influence and consequence of error | Strong as a general principle | Uncertainty sources are often combined or incompletely reported | Integrated credibility standards for models containing adaptive or opaque learned components |
| **Translational readiness** | Prospective evaluation, workflow integration, decision impact, maintenance, and governance | Pharmaceutical usefulness requires more than benchmark performance | Limited and task specific | Publication bias may favor positive retrospective demonstrations | Evidence linking model use to reproducible improvements in development or clinical decisions |
| **Review interpretation boundary** | Critical synthesis distinguishing capability, plausibility, validation, utility, and readiness | Evidence can support bounded claims without implying causality or routine deployment | Strong as an analytical discipline | Readers may interpret a proposed taxonomy as an empirically validated framework | Prospective testing of the proposed comparative and mechanistic-information framework |

*Machine Learning Contributions to Pharmacokinetic and Pharmacodynamic Modeling*

One of the clearest contributions of machine learning to pharmacometrics is the capacity to estimate complex mappings between patient characteristics, dosing histories, sampling information, and drug concentrations. Such models may accommodate nonlinear interactions, high-dimensional covariates, irregular measurements, or relationships that are difficult to specify through conventional parametric functions. Reviews of drug-concentration prediction nevertheless show that reported success is often closely tied to the dataset and evaluation design, with persistent concerns regarding missing data,

limited interpretability, inconsistent external validation, and uncertain use outside the development population [9]. Machine learning therefore adds representational capacity, but this capacity should be described as a method for estimating relationships present in the available data rather than as evidence that the underlying pharmacokinetic mechanism has been discovered.

In pharmacokinetic and pharmacodynamic modeling more broadly, machine learning can support direct outcome prediction, nonlinear function estimation, surrogate modeling, latent-state learning, classification of response, and extraction of interactions not readily represented by predefined equations [10]. These capabilities may be particularly useful when the scientific objective is predictive discrimination within a bounded domain or when the relationship of interest is too complex for a parsimonious structural model. However, flexible estimation also increases the number of relationships that can fit the observed data. Without adequate regularization, validation, and domain constraints, the model may learn sampling patterns, treatment-selection effects, assay differences, institutional practices, or other nonpharmacological signals. The relevant comparison is consequently not flexibility versus rigidity, but whether the selected degree of flexibility is justified by the question and whether the learned relationship remains reliable when the conditions of use change.

Covariate analysis illustrates both the promise and the limitation of this contribution. Conventional population-model procedures rely on predefined functional forms, staged hypothesis testing, diagnostics, and biological interpretation. Machine-learning methods can search a wider space of interactions and nonlinear effects and may therefore serve as screening tools or candidate function generators. A systematic comparison of covariate-selection approaches indicates, however, that algorithmic expansion of the search space does not remove the need to assess clinical plausibility, parameter stability, collinearity, missingness, and sensitivity to model specification [11]. In a concrete population-pharmacokinetic application, machine learning was used to support integration of covariate information, illustrating how data-driven tools can be positioned within rather than outside a pharmacometric workflow [12]. Such use is scientifically stronger when the algorithm proposes relationships that are subsequently evaluated within the structural model than when variable importance is interpreted as direct biological explanation.

Dose individualization represents a higher-consequence extension of the same logic. Machine learning may combine demographics, laboratory values, comorbidities, treatment history, concentrations, or other patient-level variables to predict a dose, exposure category, or risk state. The practical appeal is clear, particularly where relationships are nonlinear or information arrives sequentially. Yet dose recommendation is not merely a prediction task: it is an intervention whose consequences depend on calibration, treatment targets, uncertainty, clinical workflow, and the possibility of harm from incorrect action. Current evidence supports machine learning as a potentially useful component of individualized-dosing systems, but it does not justify treating retrospective predictive performance as sufficient evidence for autonomous dose selection [13]. The contribution of machine learning is therefore best characterized as conditional decision support whose acceptable role depends on validation depth, mechanistic requirements, reversibility of error, and human oversight.

*Integration with Population Models, QSP, and Exposure–Response Analysis*

Integration is most defensible when machine learning is assigned a clearly bounded role within a larger pharmacological argument. Across translational applications, three recurring configurations can be distinguished: machine learning operating beside a pharmacometric model as a comparator, machine learning supplying covariates or parameters to an established population model, and a learned function embedded directly within the model equations [14-16]. These configurations preserve different amounts of prior knowledge. A parallel model retains no structural dependence between the two approaches; a modular hybrid preserves the receiving model's states and equations but makes selected inputs data dependent; and an embedded hybrid permits the learned component to alter the system dynamics themselves. Calling all three configurations "integrated" therefore conceals important differences in interpretability, identifiability, and extrapolation.

Exposure–response analysis offers a particularly relevant use case because treatment effects may involve nonlinear exposure patterns, interacting patient characteristics, heterogeneous outcomes, and response thresholds that are difficult to capture with simple functional forms. Machine learning can identify predictive partitions, interactions, or response surfaces that complement conventional graphical and parametric analyses [17]. Nevertheless, an exposure–response association learned from observational or trial data does not alone establish pharmacological causation. Exposure may be related to dose modification, adherence, disease severity, treatment discontinuation, or time-varying clinical status. A flexible predictor may model these relationships accurately while providing an incorrect explanation of why the response occurred. Machine-learning findings are therefore most useful as exploratory structure, candidate effect modification, or predictive enrichment unless the study design and supporting evidence justify stronger interpretation.

QSP integration raises a related but broader problem. A QSP model commonly represents interacting biological pathways, disease processes, biomarkers, drug actions, and system-level feedback. Machine learning may support parameter estimation, emulation of computationally expensive modules, selection of plausible parameter regions, construction of surrogate functions, or learning of biological relationships that are not sufficiently characterized for explicit representation. Yet the scientific meaning of the resulting model depends on where the learned component enters. A surrogate placed outside the QSP system may accelerate simulation while preserving the underlying model for inspection, whereas an embedded neural component may change pathway behavior without offering a directly testable biological equation. Regulatory and development credibility consequently depend not on the presence of a mechanistic diagram but on documentation of model purpose, evidence, limitations, and decision influence [18]. **Figure 1** presents the mechanism–prediction pharmacometrics balance, showing how

the article's key components and boundaries are connected within integration with population models, qsp, and exposure-response analysis.

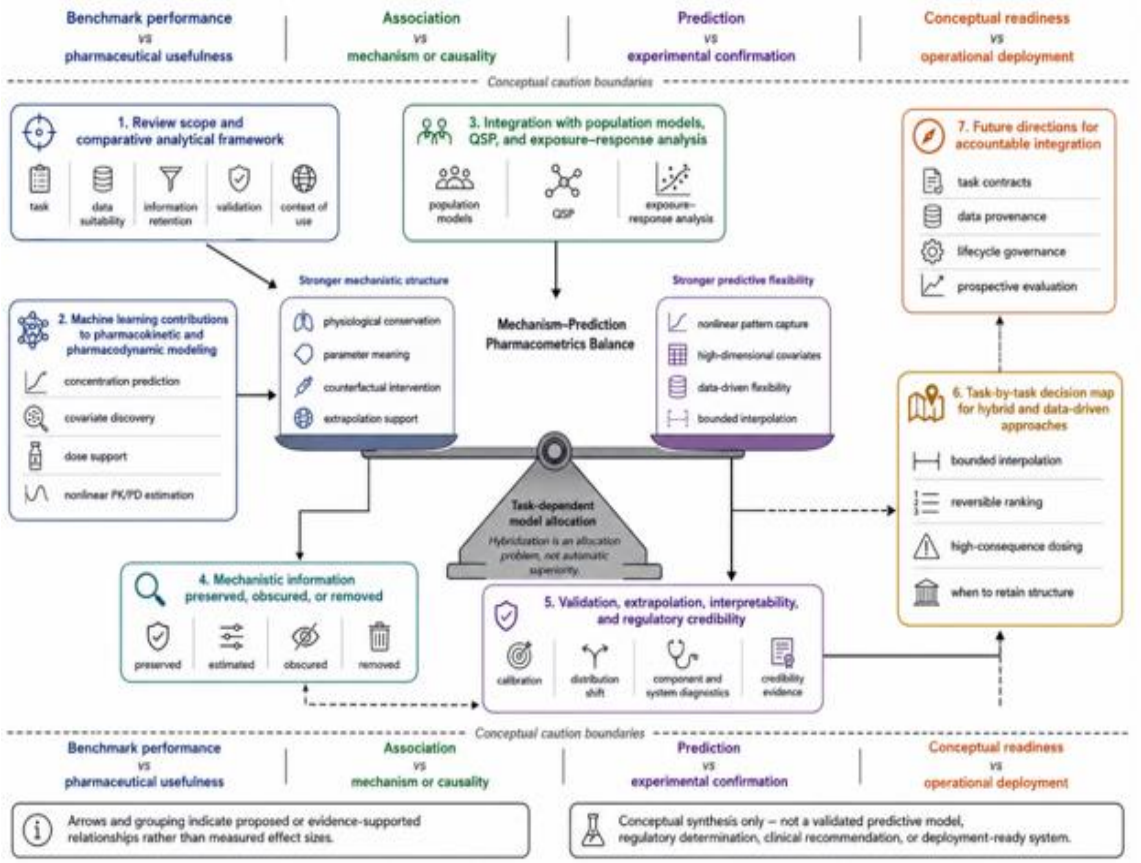


**Figure 1.** Mechanism–Prediction Pharmacometrics Balance.

The figure is an original conceptual synthesis that organizes the article's central contribution across review scope and comparative analytical framework, machine learning contributions to pharmacokinetic and pharmacodynamic modeling, integration with population models, QSP, and exposure-response analysis, integration with population models, exposure-response analysis, mechanistic information preserved, obscured, or removed. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

The proposed mechanism–prediction balance treats hybridization as an allocation problem rather than a compromise that is automatically superior to either parent tradition. Explicit structure should be retained where physiological conservation, parameter meaning, counterfactual intervention, or extrapolation is essential. Learned functions may be appropriate where relationships are poorly specified, highly nonlinear, or supported by sufficiently representative data. The balance shifts toward stronger mechanistic constraints as the prediction moves farther from observed conditions or as the consequence of error increases. Conversely, a predominantly data-driven approach may be sufficient for reversible ranking or bounded interpolation when biological explanation is not required. The central evaluative question is therefore whether every component contributes information needed for the intended task and whether the total architecture can be challenged through diagnostics appropriate to both its mechanistic and learned parts.

*Mechanistic Information Preserved, Obscured, or Removed*

Mechanistic information is not an all-or-none property. A model may preserve dose timing and longitudinal dynamics while removing conventional compartments; preserve compartments while replacing a clearance relationship with a neural function; or retain physiological labels while estimating parameter values from correlations that do not identify the proposed biological mechanism. Direct neural-network methods for concentration–time prediction illustrate this distinction. Such models can generate clinically recognizable trajectories without estimating clearance, volume, absorption rate, or interindividual variability in their conventional forms [19]. The resulting prediction may be useful, but the model cannot automatically answer mechanistic questions that depend on those parameters. A successful concentration prediction therefore does not demonstrate that the model has recovered the pharmacokinetic process represented by a compartmental model.

Neural ordinary differential equations retain an explicit temporal evolution law and may offer stronger structural discipline than static or sequence-based predictors. Their differential form can support continuous-time prediction and evaluation of

dosing schedules not represented identically in the training observations [20]. However, the derivative function remains learned, and its states do not necessarily correspond to identifiable organs, compartments, pathways, or biological quantities. Apparent extrapolation may also depend on interpolation within a learned dose–time region rather than recovery of a general pharmacological law. Differential-equation form should consequently be recognized as preserved mathematical structure, while biological interpretation of its states and functions must be established separately.

PBPK–machine-learning integration demonstrates an even wider range of information outcomes. Machine learning may estimate tissue-partition inputs, approximate computationally intensive simulations, support calibration, identify parameter regions, or replace the entire PBPK simulator with a surrogate [21]. These uses should not be treated as equivalent. Parameter-support applications can preserve the PBPK equations and physiological inventory, whereas full surrogates may reproduce outputs without maintaining an inspectable relationship to tissue flows, organ volumes, enzyme activity, or mass balance. Neural differential-equation perspectives similarly show that a model can occupy an intermediate position between explicit pharmacological equations and unconstrained prediction [22]. The amount of mechanism retained depends on the meaning and constraints of the states, the placement of the learned function, and whether the model reproduces pharmacological behavior under interventions rather than merely observed trajectories.

This review therefore proposes a four-part mechanistic-information ledger: preserved, when a relationship remains explicit and testable; estimated, when a mechanistic quantity remains defined but its value or functional form is learned; obscured, when relevant information may be encoded but cannot be reliably inspected or attributed; and removed, when the architecture no longer represents the quantity needed for the scientific question. The ledger is an interpretive device rather than a validated scoring system. It also prevents post-hoc explanation from being mistaken for recovered mechanism. Feature attribution, saliency, or local explanation may describe how a fitted predictor uses its inputs, but such procedures cannot recreate physiological constraints or causal structure absent from the model itself [23]. Interpretability should therefore be evaluated against the information required by the decision, not against the visual attractiveness of an explanation.

*Validation, Extrapolation, Interpretability, and Regulatory Credibility*

Validation of pharmacometric machine learning should begin with an explicit context of use. The evidence needed for exploratory covariate screening differs from that needed for trial simulation, exposure–response support, or individualized dosing. Credibility frameworks developed for model-informed development and mechanistic simulation converge on the principle that verification, validation, applicability, and uncertainty must be proportionate to the model's influence and the consequence of an incorrect result [24-26]. For hybrid systems, this principle should be applied at both component and system levels. The mechanistic module must be checked for implementation correctness and scientific plausibility, the learned module must be evaluated for data leakage, instability, and distribution shift, and the coupled model must be tested for emergent behaviors not evident from either component alone.

Calibration is central because predictive discrimination and numerical reliability are not interchangeable. A model may correctly rank patients by risk or generate low average error while systematically underpredicting high exposures, overstating response probability, or producing uncertainty intervals that fail to reflect observed error. Calibration assessment should therefore be aligned with the output used in the decision, whether that output is a concentration, parameter, probability, dose, trajectory, or treatment policy [27]. Internal random splits are rarely sufficient when observations from the same individual, site, regimen, compound, or development program share structure. Validation should instead challenge the dimensions along which the model is expected to operate, including new patients, new schedules, new populations, altered assays, changing clinical practice, or related compounds. **Figure 2** presents the evidence, validation, and translation boundary observatory for what machine learning adds to pharmacometrics and, showing how the article's key components and boundaries are connected within validation, extrapolation, interpretability, and regulatory credibility.

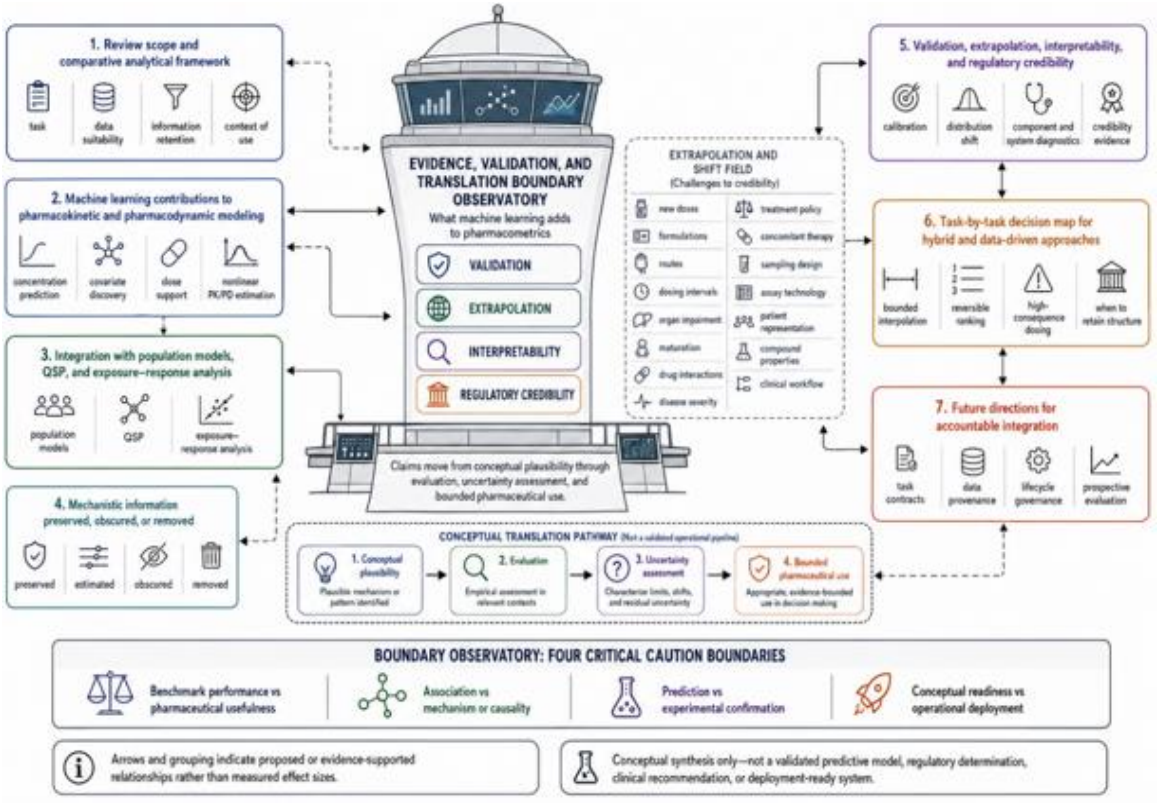


**Figure 2.** Evidence, Validation, and Translation Boundaries.

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

Extrapolation requires more than applying a fitted model to values numerically outside the training range. Relevant shifts include changes in causal structure, treatment policy, disease severity, concomitant therapy, sampling design, assay technology, patient representation, and clinical workflow. Healthcare machine-learning literature has accordingly criticized the assumption that performance obtained in one dataset establishes general applicability [28]. In pharmacometrics, a model may encounter additional shifts produced by new doses, formulations, routes, dosing intervals, organ impairment, maturation, drug interactions, or compound properties. Mechanistic models do not solve these problems automatically, but explicit assumptions make some extrapolations inspectable. Data-driven models may extrapolate successfully when the learned relationship is stable, yet the basis for that stability must be demonstrated rather than inferred from model complexity.

Regulatory credibility should therefore be understood as a structured evidence argument, not as a label attached to a model family. Regulatory experience with model-informed development emphasizes clear purpose, data provenance, model construction, evaluation, sensitivity analysis, uncertainty, limitations, and documentation of how the result influences a decision [18]. A hybrid model can potentially satisfy these expectations, but its learned component creates additional questions concerning version control, reproducibility, training-data relevance, update procedures, and performance under shift. Mechanistic-model credibility principles provide a valuable foundation, although they must be extended to address adaptive, opaque, or data-dependent components [25]. Interpretability contributes to credibility when it supports error detection, assumption checking, or decision justification; it is insufficient when it merely generates a plausible narrative after prediction.

*A Task-By-Task Decision Map for Hybrid and Data-Driven Approaches*

The appropriate modeling architecture depends first on the pharmaceutical task. Sequential dose adaptation, for example, differs fundamentally from static concentration prediction because each recommendation changes the subsequent state and data-generating process. Model-informed reinforcement learning has been proposed as a way to combine pharmacological simulation with sequential policy learning [29]. Its value is plausible where delayed responses, competing objectives, and repeated decisions are central. However, reward misspecification, unsafe exploration, incomplete representation of clinical constraints, and unreliable off-policy evaluation create substantial risks. Reinforcement learning should therefore be considered a high-evidence-burden approach rather than a natural replacement for established model-based dosing procedures.

For precision dosing, the starting point should usually be the least complex model capable of supporting the required decision. Established pharmacometric approaches may be preferred when concentration data are sparse, interpretable parameters are required, prior biological knowledge is strong, and extrapolation to new dosing conditions is central. Machine learning may add value when rich longitudinal covariates, complex interactions, waveform data, or high-dimensional clinical information

materially improve prediction. Hybrid models become attractive when an adequate pharmacometric scaffold exists but a specific component—such as a covariate relationship, biomarker function, residual process, or difficult parameter—is insufficiently described. Reviews of model-informed precision dosing nevertheless show that utility depends on workflow integration, updating logic, clinical feasibility, and prospective performance rather than model novelty [30].

The consequence and reversibility of error provide a second decision axis. Exploratory feature screening, compound prioritization, or hypothesis generation may tolerate lower interpretability and weaker external evidence because results can be rejected or tested before consequential action. Trial-design simulation, benefit–risk assessment, and individualized therapy require progressively stronger calibration, uncertainty characterization, transport evaluation, and traceability. Model-informed precision dosing also depends on validated models, relevant patient information, decision-support infrastructure, and evidence that recommendations can be implemented safely [31]. A task-by-task map should therefore classify not only the model but also the action it informs, the opportunity for human correction, the availability of confirmatory evidence, and the harm produced by failure.

A predominantly data-driven model is a reasonable starting point when the intended task is bounded interpolation, data are sufficiently representative, mechanistic interpretation is unnecessary, and prediction errors remain reversible. A mechanistic model is favored when the decision depends on mass balance, physiological counterfactuals, dose or regimen extrapolation, sparse-data borrowing, or interpretable sources of variability. Hybridization is justified when it solves a defined structural gap without surrendering information required by the decision. In PBPK applications, the role assigned to machine learning should therefore follow the task and the need for physiological traceability rather than a general preference for automation [32]. Early drug-discovery ranking may accommodate data-driven models more readily than patient-level dosing because candidates can undergo subsequent experimental evaluation before therapeutic use [33]. These distinctions constitute a proposed decision map, not a universally validated selection rule.

*Future Directions for Accountable Integration*

Accountable integration begins with narrower claims. Pharmaceutical machine learning should be developed around explicit task contracts defining the target, population, prediction horizon, acceptable error, intended intervention, excluded uses, and required evidence. Broad reviews of machine learning in drug discovery and development indicate that useful applications are most credible when the data and evaluation correspond closely to a well-specified pharmaceutical question [34]. Future studies should therefore report not only what a model predicts but why that output is needed, what competing approach it replaces or supplements, and what decision would change if the prediction were available. This task contract should precede algorithm selection.

Second, evaluation should separate component-level performance from development-level value. Improvement in a concentration metric, ranking function, surrogate computation, or covariate fit does not automatically shorten development, improve candidate quality, reduce attrition, or improve patient outcomes. Critical assessments of artificial intelligence in drug discovery emphasize that local technical success can be mistaken for evidence of end-to-end pharmaceutical impact [35]. Comparative studies should therefore identify the decision bottleneck being addressed and measure whether the model changes that decision reliably. Negative results, failed transport, calibration deterioration, and situations in which a simpler model performs adequately are particularly important for defining realistic value.

Third, the field requires stronger attention to data-generating processes. Chemical, biological, pharmacokinetic, and clinical datasets are not neutral collections of independent observations. They reflect experimental selection, assay limitations, treatment policies, missingness, institutional practice, censoring, compound families, and historical development decisions. These properties restrict the relationships that can be learned and may create shortcuts that remain invisible under random validation [36]. Future models should be accompanied by data-provenance maps, representativeness analyses, shift hypotheses, and explicit statements of which intended-use conditions are unsupported. Data suitability must be evaluated independently of data volume.

Finally, accountable integration requires lifecycle governance. Reproducible preprocessing, transparent model development, external validation, calibrated uncertainty, versioning, drift monitoring, and predefined revalidation triggers should be treated as scientific requirements rather than administrative additions. Pharmaceutical machine-learning guidance already emphasizes fit-for-purpose development, transparent evaluation, and reproducibility [37]. End-to-end visions will remain credible only when each linked model is independently and jointly validated and when failures can be traced to a specific component or assumption [38]. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop future directions for accountable integration within the article's central argument.

**Table 2.** Evaluation, implementation, and research priorities arising from What Machine Learning Adds to Pharmacometrics and What Mechanistic Information It Risks

| Approach or application | Data and task context | Evaluation practice | Evidence maturity | Translational limitation | Review interpretation |
|---|---|---|---|---|---|
| **ML-assisted covariate screening** | Population PK/PD datasets with multiple candidate covariates | Resampling stability, missing-data sensitivity, comparison with prespecified covariates, and pharmacometric diagnostics | Within-domain evaluated in selected applications | Variable selection may reflect correlation, treatment policy, or sampling structure rather than biology | Appropriate as a hypothesis-generation or model-building aid, not as automatic biological adjudication |

| | | | | | |
|---|---|---|---|---|---|
| **Direct concentration prediction** | Repeated concentrations with patient and dosing descriptors | Patient-level separation, calibration, regimen-shift testing, subgroup error analysis, and comparison with pharmacometric baselines | Within-domain evaluated | Conventional PK parameters and physiological explanations may be absent | Useful for bounded prediction when interpretability and broad extrapolation are not required |
| **Neural ordinary differential equations** | Longitudinal PK or PK/PD trajectories with temporal structure | State and derivative diagnostics, dose-intervention tests, external regimen evaluation, and comparison with structural models | Proof-of-concept | Learned states may lack biological identifiability and may fail under structural shift | Preserves continuous-time form but does not automatically preserve pharmacological mechanism |
| **Residual-correction hybrid** | Adequate mechanistic scaffold with systematic unexplained error | Residual analysis, ablation of the learned correction, extrapolation tests, and assessment of whether structural misspecification is being hidden | Proof-of-concept to within-domain evaluated | Learned residuals may absorb assay artifacts, missing covariates, or incorrect mechanism | Appropriate when the correction is bounded, inspectable, and not used to disguise model inadequacy |
| **Embedded scientific machine learning** | Partially known dynamical system with one or more uncertain functions | Parameter recovery, conservation checks, component ablation, intervention testing, and uncertainty propagation | Proof-of-concept | A flexible embedded function may dominate system behavior despite a mechanistic exterior | Promising when unknown components are clearly identified and mechanistic claims remain component specific |
| **ML-supported PBPK or QSP parameterization** | Physiological model with uncertain or difficult-to-measure inputs | Plausibility constraints, sensitivity analysis, external compound or population testing, and propagation of input uncertainty | Emerging | Correlational parameter estimates may be mistaken for physiological confirmation | Useful when learned estimates remain constrained and the underlying model is independently credible |
| **PBPK or QSP surrogate modeling** | Computationally intensive simulations with a defined input domain | Emulation error, boundary stress tests, out-of-domain detection, and preservation of decision-relevant outputs | Proof-of-concept to within-domain evaluated | Surrogate accuracy may collapse outside the simulated design space and may remove physiological traceability | Appropriate for acceleration within a declared domain, not as an unrestricted replacement for the source model |
| **ML exposure–response analysis** | Trial or observational data with exposure and outcome measures | Confounding analysis, calibration, endpoint validation, temporal assessment, and comparison with conventional exposure–response models | Proof-of-concept | Predictive associations may not represent treatment-effect mechanisms | Valuable for exploration and effect-modification hypotheses; stronger causal claims require additional design and evidence |
| **Machine-learning precision dosing** | Rich patient-level longitudinal data and actionable dosing decisions | External calibration, silent prospective evaluation, workflow testing, uncertainty thresholds, and human-factor assessment | Early translational | Errors may directly affect treatment, and historical prescribing patterns may be learned as targets | Should remain supervised decision support until prospective safety and utility are demonstrated |
| **Model-informed reinforcement learning** | Sequential dosing or treatment policies with delayed outcomes | Safe simulation, off-policy evaluation, reward sensitivity, constraint testing, and prospective staged assessment | Conceptual to proof-of-concept | Unsafe exploration, incomplete rewards, and policy instability | Requires a substantially higher evidence burden than static prediction |
| **Pharmaceutical lifecycle governance** | Models updated across populations, sites, compounds, assays, or software versions | Version control, drift detection, audit trails, revalidation triggers, rollback procedures, and periodic context-of-use review | Emerging | Updating may invalidate prior calibration, explanation, or mechanistic assumptions | Governance must be integrated into model development rather than added after deployment |
| **Task-contract-based model selection** | Any pharmaceutical decision involving mechanistic, hybrid, or data-driven modeling | Explicit definition of intended use, consequence of error, mechanism requirement, extrapolation demand, and minimum evidence | Proposed synthesis | The decision map has not been prospectively validated across development programs | Provides a structured basis for choosing the least complex model capable of supporting the decision |

## Conclusion

Machine learning adds genuine value to pharmacometrics by expanding nonlinear estimation, covariate analysis, direct prediction, exposure–response characterization, surrogate modeling, and the representation of incompletely specified system components. Its contribution, however, cannot be judged by predictive performance alone. Every modeling choice changes the scientific information available for explanation, intervention, simulation, and extrapolation. The original contribution of this review is a proposed task-centered synthesis in which mechanistic information is classified as preserved, estimated,

obscured, or removed and in which the required evidence increases with distributional distance and the consequence of error. Mechanistic, data-driven, and hybrid approaches should therefore be treated as different allocations of assumptions and information rather than as stages in a universal hierarchy. Hybridization can be valuable when it retains the structures required by the decision while learning a clearly bounded knowledge gap, but a mechanistic exterior does not automatically make a learned model interpretable or credible. Accountable integration requires explicit task contracts, suitable data, component-level and system-level validation, calibrated uncertainty, transport testing, prospective evaluation where consequences are substantial, and lifecycle governance. These requirements do not diminish the role of machine learning; they define the conditions under which its flexibility can strengthen rather than silently weaken pharmacometric reasoning.

**Acknowledgments:** None

**Conflict of interest:** None

**Financial support:** None

**Ethics statement:** None

**References**

1. McComb M, Bies R, Ramanathan M. Machine learning in pharmacometrics: Opportunities and challenges. Br J Clin Pharmacol. 2022;88(4):1482-99. doi:10.1111/bcp.14801
2. Stankevičiūtė K, Woillard JB, Peck RW, Marquet P, van der Schaar M. Bridging the worlds of pharmacometrics and machine learning. Clin Pharmacokinet. 2023;62(11):1551-65. doi:10.1007/s40262-023-01310-x
3. Janssen A, Bennis FC, Mathôt RAA. Adoption of machine learning in pharmacometrics: An overview of recent implementations and their considerations. Pharmaceutics. 2022;14(9):1814. doi:10.3390/pharmaceutics14091814
4. Keutzer L, You H, Farnoud A, Nyberg J, Wicha SG, Maher-Edwards G, et al. Machine learning and pharmacometrics for prediction of pharmacokinetic data: Differences, similarities and challenges illustrated with rifampicin. Pharmaceutics. 2022;14(8):1530. doi:10.3390/pharmaceutics14081530
5. Valderrama D, Teplytska O, Koltermann LM, Trunz E, Schmulenson E, Fritsch A, et al. Comparing scientific machine learning with population pharmacokinetic and classical machine learning approaches for prediction of drug concentrations. CPT Pharmacometrics Syst Pharmacol. 2025;14(4):759-69. doi:10.1002/psp4.13313
6. Snyder H. Literature review as a research methodology: An overview and guidelines. J Bus Res. 2019;104:333-9. doi:10.1016/j.jbusres.2019.07.039
7. Munn Z, Peters MDJ, Stern C, Tufanaru C, McArthur A, Aromataris E. Systematic review or scoping review? Guidance for authors when choosing between a systematic or scoping review approach. BMC Med Res Methodol. 2018;18(1):143. doi:10.1186/s12874-018-0611-x
8. Marshall S, Madabushi R, Manolis E, Krudys K, Staab A, Dykstra K, et al. Model-informed drug discovery and development: Current industry good practice and regulatory expectations and future perspectives. CPT Pharmacometrics Syst Pharmacol. 2019;8(2):87-96. doi:10.1002/psp4.12372
9. Huang S, Xu Q, Yang G, Ding J, Pei Q. Machine learning for prediction of drug concentrations: Application and challenges. Clin Pharmacol Ther. 2025;117(5):1236-47. doi:10.1002/cpt.3577
10. Tang A. Machine learning for pharmacokinetic/pharmacodynamic modeling. J Pharm Sci. 2023;112(5):1460-75. doi:10.1016/j.xphs.2023.01.010
11. Karlsen M, Khier S, Fabre D, Marchionni D, Azé J, Bringay S, et al. Covariate model selection approaches for population pharmacokinetics: A systematic review of existing methods, from SCM to AI. CPT Pharmacometrics Syst Pharmacol. 2025;14(4):621-39. doi:10.1002/psp4.13306
12. Zhu X, Zhang M, Wen Y, Shang D. Machine learning advances the integration of covariates in population pharmacokinetic models: Valproic acid as an example. Front Pharmacol. 2022;13:994665. doi:10.3389/fphar.2022.994665
13. Li QY, Tang BH, Wu YE, Yao BF, Zhang W, Zheng Y, et al. Machine learning: A new approach for dose individualization. Clin Pharmacol Ther. 2024;115(4):727-44. doi:10.1002/cpt.3049
14. Terranova N, Venkatakrishnan K, Benincosa LJ. Application of machine learning in translational medicine: Current status and future opportunities. AAPS J. 2021;23(4):74. doi:10.1208/s12248-021-00593-x
15. Destere A, Marquet P, Salmon Gandonnière C, Åsberg A, Loustaud-Ratti V, Carrier P, et al. A hybrid model associating population pharmacokinetics with machine learning: A case study with iohexol clearance estimation. Clin Pharmacokinet. 2022;61(8):1157-65. doi:10.1007/s40262-022-01138-x
16. 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

17. Harun R, Yang E, Kassir N, Zhang W, Lu J. Machine learning for exposure-response analysis: Methodological overview and case study. Pharmaceutics. 2023;15(5):1381. doi:10.3390/pharmaceutics15051381
18. Wang Y, Zhu H, Madabushi R, Liu Q, Huang SM, Zineh I. Model-informed drug development: Current US regulatory practice and future considerations. Clin Pharmacol Ther. 2019;105(4):899-911. doi:10.1002/cpt.1363
19. Bräm DS, Parrott N, Hutchinson L, Steiert B. Introduction of an artificial neural network-based method for concentration-time predictions. CPT Pharmacometrics Syst Pharmacol. 2022;11(6):745-54. doi:10.1002/psp4.12786
20. 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
21. Chou WC, Lin Z. Machine learning and artificial intelligence in physiologically based pharmacokinetic modeling. Toxicol Sci. 2023;191(1):1-14. doi:10.1093/toxsci/kfac101
22. Losada IB, Terranova N. Bridging pharmacology and neural networks: A deep dive into neural ordinary differential equations. CPT Pharmacometrics Syst Pharmacol. 2024;13(8):1289-96. doi:10.1002/psp4.13149
23. Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat Mach Intell. 2019;1(5):206-15. doi:10.1038/s42256-019-0048-x
24. Kuemmel C, Yang Y, Zhang X, Florian J, Zhu H, Tegenge M, et al. Consideration of a credibility assessment framework in model-informed drug development: Potential application to physiologically-based pharmacokinetic modeling and simulation. CPT Pharmacometrics Syst Pharmacol. 2020;9(1):21-8. doi:10.1002/psp4.12479
25. Musuamba FT, Skottheim Rusten I, Lesage R, Russo G, Bursi R, Emili L, et al. Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility. CPT Pharmacometrics Syst Pharmacol. 2021;10(8):804-25. doi:10.1002/psp4.12669
26. Viceconti M, Pappalardo F, Rodriguez B, Horner M, Bischoff J, Musuamba Tshinanu F. In silico trials: Verification, validation and uncertainty quantification of predictive models used in the regulatory evaluation of biomedical products. Methods. 2021;185:120-7. doi:10.1016/j.ymeth.2020.01.011
27. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW, Topic Group 'Evaluating Diagnostic Tests and Prediction Models' of the STRATOS Initiative. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
28. Futoma J, Simons M, Panch T, Doshi-Velez F, Celi LA. The myth of generalisability in clinical research and machine learning in health care. Lancet Digit Health. 2020;2(9). doi:10.1016/S2589-7500(20)30186-2
29. Ribba B, Dudal S, Lavé T, Peck RW. Model-informed artificial intelligence: Reinforcement learning for precision dosing. Clin Pharmacol Ther. 2020;107(4):853-7. doi:10.1002/cpt.1777
30. Minichmayr IK, Dreesen E, Centanni M, Wang Z, Hoffert Y, Friberg LE, et al. Model-informed precision dosing: State of the art and future perspectives. Adv Drug Deliv Rev. 2024;215:115421. doi:10.1016/j.addr.2024.115421
31. 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
32. Talkington AM, Cao Y, Kearsley AJ, Lai SK. Opportunities for machine learning and artificial intelligence in physiologically-based pharmacokinetic (PBPK) modeling. Adv Drug Deliv Rev. 2025;227:115716. doi:10.1016/j.addr.2025.115716
33. Pillai N, Dasgupta A, Sudsakorn S, Fretland J, Mavroudis PD. Machine learning guided early drug discovery of small molecules. Drug Discov Today. 2022;27(8):2209-15. doi:10.1016/j.drudis.2022.03.017
34. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
35. Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 1: Ways to make an impact, and why we are not there yet. Drug Discov Today. 2021;26(2):511-24. doi:10.1016/j.drudis.2020.12.009
36. Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 2: A discussion of chemical and biological data used for AI in drug discovery. Drug Discov Today. 2021;26(4):1040-52. doi:10.1016/j.drudis.2020.11.037
37. Talevi A, Morales JF, Hather G, Podichetty JT, Kim S, Bloomingdale PC, et al. Machine learning in drug discovery and development Part 1: A primer. CPT Pharmacometrics Syst Pharmacol. 2020;9(3):129-42. doi:10.1002/psp4.12491
38. Ekins S, Puhl AC, Zorn KM, Lane TR, Russo DP, Klein JJ, et al. Exploiting machine learning for end-to-end drug discovery and development. Nat Mater. 2019;18(5):435-41. doi:10.1038/s41563-019-0338-z