



HAL
open science

Bridging the Worlds of Pharmacometrics and Machine Learning

Kamilė Stankevičiūtė, Jean-Baptiste Woillard, Richard Peck, Pierre Marquet,
Mihaela van der Schaar

► **To cite this version:**

Kamilė Stankevičiūtė, Jean-Baptiste Woillard, Richard Peck, Pierre Marquet, Mihaela van der Schaar. Bridging the Worlds of Pharmacometrics and Machine Learning. *Clinical Pharmacokinetics*, 2023, 62 (11), pp.1551-1565. 10.1007/s40262-023-01310-x . hal-04489693

HAL Id: hal-04489693

<https://unilim.hal.science/hal-04489693v1>

Submitted on 11 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Bridging the worlds of pharmacometrics and machine learning

Kamilė Stankevičiūtė^{1,*} & Jean-Baptiste Woillard^{2,3,*}, Richard W. Peck^{4,5}, Pierre Marquet^{2,3}, Mihaela van der Schaar^{6,7}

* These authors contributed equally to this work

¹ Department of Computer Science and Technology, University of Cambridge, 15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

² INSERM U1248 P&T, University of Limoges, F-87000 Limoges, France.

³ Department of Pharmacology and Toxicology, CHU Limoges, F-87000 Limoges, France.

⁴ Department of Pharmacology & Therapeutics, University of Liverpool, Liverpool, United Kingdom

⁵ Pharma Research & Development, Roche Innovation Center, Basel, Switzerland

⁶ Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom

⁷ The Alan Turing Institute, 96 Euston Road, London NW1 2DB, United Kingdom

Corresponding author:

Jean-Baptiste Woillard Title: PharmD, PhD ORCID: 0000-0003-1695-0695

Address: Univ. Limoges, INSERM U1248 P&T, 2 rue du Pr Descottes, F-87000 Limoges, France.

Phone: +33 5 55 05 61 40

Fax: +33 5 55 05 61 62

Email: jean-baptiste.woillard@unilim.fr

ORCID: 0000-0003-2489-9615, 0000-0003-1695-0695, 0000-0003-1018-9655, 0000-0001-7698-0760, 0000-0003-3933-6049

Running head: Bridging the worlds of pharmacometrics and machine learning

Abstract

Precision medicine requires individualized modeling of disease and drug dynamics, with machine-learning based computational techniques gaining increasing popularity. The complexity of either field, however, makes current pharmacological problems opaque to machine learning practitioners, and state-of-the-art machine learning methods inaccessible to pharmacometricians. To help bridge the two worlds, we provide an introduction to current problems and techniques in pharmacometrics—ranging from pharmacokinetic and pharmacodynamic modeling to pharmacometric simulations, model-informed precision dosing and systems pharmacology—and review some of the machine learning approaches to address them. We hope this would facilitate collaboration between experts, with complementary strengths of principled pharmacometric modeling and flexibility of machine learning leading to synergistic effects in pharmacological applications.

Keywords: pharmacometrics, artificial intelligence, machine learning, simulations, model informed precision dosing

Key findings/implications of the manuscript

Precision medicine requires individualized modeling of disease and drug dynamics, with machine learning-based computational techniques gaining increasing popularity.

The complexity of either field, however, makes current pharmacological problems opaque to machine learning practitioners, and state-of-the-art machine learning methods inaccessible to pharmacometricians.

To help bridge the worlds of pharmacometrics and machine learning, we provide an introduction to current problems and techniques in pharmacometrics, and review some of the machine learning approaches to address them.

1. Introduction

With the world moving towards precision medicine, physicians and clinical pharmacologists are moving their focus towards tailoring treatments to match the individual patient, since the same dose of a drug could result in different concentrations and responses from one patient to another. This variability in dose requirement can be linked to many factors, including varying disease states, differences in drug exposure and response of the target, or the downstream cascade initiated by the drug/target interaction. The challenge of finding the right dose for each patient is at the heart of the discipline of *pharmacometrics*.

Broadly, pharmacometrics investigate how a given dose relates to the *exposure* to the drug based on the changes in drug concentration (*pharmacokinetics*), and how the changes in exposure markers relate to molecular or clinical *effects* (*pharmacodynamics*) (Figure 1). A good understanding of these relationships is necessary for a reliable prediction of the dose with the best risk/benefit balance for an individual patient. While there have been developments in probabilistic and dynamical system-based techniques such as model-informed precision dosing, reliable prediction with current tools is still challenging.

First, the markers directly linked to the exposure or effect of the drug are often unreachable, since it would require being able to measure concentrations at the target tissue on a continuous basis, or being able to observe clinical effects far into the future. Although the development of biosensors for real time measurement of drug concentrations is promising, it is still far from routine care use [1]. Drug exposure and effect are therefore often estimated only through the more accessible but less accurate surrogate markers, for example blood plasma concentrations.

Secondly, plasma concentrations, in contrast to fixed doses, are variable and noisy. One reason is the inter-individual variability due to genetic, demographic and environmental factors—including age, sex, body weight, genetic polymorphisms in metabolism enzymes or membrane transporters and in drug targets, differences in physiological function, circadian rhythms, disease state, and many others. Measurements are noisy even for a single individual: while part of this intra-individual variability is quantifiable (e.g. analytical uncertainty in drug measurement), the other part is not (e.g. uncertainty in drug dosage and blood sampling times, food effect, drug-drug interactions, poor adherence to treatment, etc.) [2–6].

In this article, we provide an overview of pharmacometrics and explore the potential of machine learning applications across all areas of pharmacometric modeling, drug development stages, and real-time decision making for therapeutic drug monitoring (dose adjustment). Building upon previous perspectives on machine learning in clinical pharmacology [7] and collaborative works between machine learners and pharmacometricians so far [8–10], we discuss the ways to bridge these worlds further. We provide a summary of the pharmacometric challenges and opportunities in Table 1, accompanied by Figures 2 to 5 that outline our vision for their stagewise integration with machine learning.

2. Pharmacokinetic modeling

Pharmacokinetics (PK) quantify the *exposure* of a body to the drug using different *exposure markers* (such as peak concentration C_{\max} , trough concentration C_0 or C_{\min} , or total area under the concentration curve, AUC, which measures the average exposure between two doses) (Figure 1). For most drugs, C_0 is a good surrogate of the AUC due to their strong correlation. In other cases, C_{\max} or more accurate estimation of the AUC should be preferred. Concentration profiles depend on the dose and individual features, but exposure needs to be high enough to achieve the desired effect of the drug while avoiding toxic effects. Once modeled, the relationship between individual features and the concentration profiles can be used to propose the initial treatment regimen, and adjusting the dose based on subsequent feedback until the desired exposure is reached—a practice known as model informed precision dosing (MIPD). Pharmacokinetic models are therefore key for data-driven precision medicine, simulation studies, novel drug discovery, and target attainment hypothesis testing.

Population pharmacokinetic (popPK) modeling is used to fit pharmacokinetic parameters based on the time-concentration profiles of a set of individuals. The mean popPK parameters describe the *structural model* of a given population, and their variance the *statistical model* of their inter- and intra-individual variabilities [11]. In this section, we review current approaches to popPK modeling, challenges in modeling variability with small data samples and narrow study populations, and the potential of machine learning methods for addressing them.

2.1. Population pharmacokinetic model development

In pharmacokinetics, the human body is typically represented by a simplified compartmental structural model, where peripheral compartments correspond to different tissues in the body, and the central compartment includes the vascular system through which the drug is transported (Figure 1b). The compartments are interconnected—

usually through the central compartment—and the drug movements (absorption, distribution, elimination) between them are specified via rate constants either in their explicit forms or through ordinary differential equations (ODEs). Variability in rate constants can be linked to individual covariates (age, sex, genetics)—one task of popPK, then, is to estimate the relationships between individual covariates and the rate constants in a given population, while discovering the sources of and accounting for inter- and intra-individual variabilities.

One or two-compartment structural models are common choices made by human experts for ease of mathematical modeling, but they can lead to oversimplification or misspecification of the true pharmacokinetic mechanisms. To address this, data-driven machine learning methods for model selection have been proposed, including those based on genetic algorithms and neural networks [12–14] (Figure 3). In addition, a reinforcement learning algorithm has been developed to select the best structural non-parametric models [15].

Nonlinear mixed effect (NLME) modeling is the standard approach for fitting popPK models though other nonparametric methods have also been developed [11,16–19] (Figure 1c). As stated above, popPK parameters can depend on individual covariates; however, when the covariates are high-dimensional and have strong correlations structures, it is a challenge to select only those most relevant to the PK of a particular drug. This leads to a range of popPK models designed for the same drug and population, yet having distinct covariate sets [20]. Additionally, not all covariates may have a causal effect on pharmacokinetic parameters. Recent works have proposed using machine learning (random forests, support vector regression, and neural networks) to guide covariate screening, leading to performance improvements [21,22]. Machine learning could also be used to further improve the estimates of nonlinear relationships between individual covariates and rate constants typically modeled in popPK analyses.

2.2. A priori and a posteriori population pharmacokinetic modeling

Population pharmacokinetics could be categorized as *a priori*, when only individual covariates—but no drug concentration measurements—are available (initial dose prediction), or *a posteriori*, refining subsequent doses as more drug exposure information is obtained. Machine learning approaches for covariate-based first dose prediction, ranging from simple regression models to the more complex techniques, have so far been limited due to high inter-individual variability [23,24]. However, pure popPK-based approaches may also fall short of perfect target attainment [3,25–27]. We now briefly explore the potential synergies between *a priori* and *a*

posteriori popPK and machine learning-based predictions, combining the strengths of pharmacological domain knowledge captured by popPK and the powerful pattern-matching of machine learning.

One line of work combining popPK modeling with machine learning for *a priori* first dose prediction focuses on popPK as a way to generate more training data, since samples from real patients are often insufficient for training the machine learning model. For example, Monte Carlo simulations of one [24] or more [28] popPK models and training an XGBoost model on the resulting data has been effective at improving dose predictions for vancomycin. In relation to that, we have recently developed machine learning algorithms to predict tacrolimus exposure from Monte Carlo simulations based on a previously published popPK model, with performances as good as those of models developed using measured patient data [29]. However, we have also found that machine learning models tend to overfit with increasing simulated data, one example being everolimus exposure prediction in transplanted patients [30]. Alternatively, Tang et al. proposed using popPK models as a basis for training a machine learning predictor for individual PK parameter values, and observed improvements in performance [31].

In *a posteriori* popPK, exposure markers are predicted following observation of an individual response to the initial doses. Feedback from initial observations is used to improve individualized predictions of the next exposure and to propose a more accurate dose while updating the parameters of the popPK model. Bayesian methods (such as maximum a posteriori estimation) are typically used to obtain posterior distributions of popPK parameters [32,33]; however, their performance suffers when complex prior models cannot be complemented with sufficient individualized information, leading to parameter shrinkage to mean population values [19]. Hybrid popPK and machine learning algorithms have been developed to address these limitations: Hughes & Keizer propose a machine learning classifier for optionally ignoring popPK information by flattening the priors on an individual level[34]; Destere et al. use machine learning to predict the residual errors of popPK model outputs [35]. When larger datasets are available, popPK modeling may be skipped altogether in favor of machine learning-based PK parameter and exposure index prediction directly from individual data [36–38], leading to similar or improved performance and faster development [39].

3. Exposure/effect modeling

Ultimately, the goal of precision medicine is not only to accurately estimate exposure markers, but to improve the *outcomes* (clinically relevant endpoints) of a treatment. The relationship between exposure markers of a drug and the *outcome* (or *effect*) on the patient is at the core of *pharmacodynamics* (PD). The two are often analyzed together and jointly referred to as *exposure/effect* (or PK/PD) modeling.

At a microscopic level, PK/PD models describe relationships between drug concentrations and the activation or inhibition of (molecular) targets, typically using NLME methods (Figure 1d). At a macroscopic level, the effect could describe the *efficacy* or *toxicity* in terms of patient symptoms and the benefit/risk balance at a given exposure level. While the low-level molecular effects are relatively well-defined in PK/PD modeling, they are mostly based on *in vitro* experiments, e.g. by estimating the concentration required to reach 50% of the maximal or inhibitory effect (EC50 or IC50 respectively), or measuring receptor occupancy; however, some rare exceptions aside (such as effect of hypertension drugs observed through changes in blood pressure, or effect of statins observed through changes in LDL-cholesterol values), direct measurement of effect at target tissues is not possible *in vivo*. Similarly, high-level outcomes are often defined through binary efficacy and toxicity thresholds—a likely oversimplification in itself—which might be impossible to determine due to delayed outcomes such as survival; in this case, only imprecise surrogates (such as tumor growth in oncology or glomerular filtration rate in renal transplantation) are available. Machine learning could help determine the best exposure indices and their relationship to microscopic targets or macroscopic outcomes, taking into account non-linear relationships, delays in response or changes in dynamics due to cumulative exposure.

3.1. Pharmacokinetics/pharmacodynamics for chronic disease

In addition to challenges of inter- and intra-individual variability, chronic diseases tend to evolve over time, with the same treatment leading to a different exposure/effect relationship depending on the disease state. Longitudinal studies and disease progression modeling are therefore of crucial importance for determining correct PK/PD relationships. A standard statistical approach is to use joint modeling of longitudinal and time-to-event data [40–42], with the limitation that joint modeling relies on strict hypotheses about the choice of the hazard function for the survival model [43]. Machine learning models could be explored as an alternative. For example, Lee et al. develop a deep learning model for dynamic survival predictions for multiple competing risks

[44]. Combining this model with unsupervised learning to identify latent patient groups with their own submodels [45], they show promising results in predicting the progression of prostate cancer [46].

3.2. Pharmacokinetics/pharmacodynamics of drug interactions

A related challenge is that complex diseases are often multicausal and need to be treated by several drugs at the same time (e.g. in case of immunosuppressants, antibiotics or anticancer drugs). This could lead to noisy or incorrect exposure/effect estimates if the drugs are studied without accounting for interaction effects. For example, low exposure of one treatment could be misclassified as a non-predictor of a clinical outcome if the other treatment compensates for it due to pharmacodynamic (synergistic) interaction. The treatments may also interact pharmacokinetically by inducing or inhibiting each other's metabolism. A review of pure pharmacometric and machine learning-based approaches for compound combination prediction has been presented in Bulusu et al. [47]—categorizing the techniques into those based on gene expression and cell line sensitivity, pathways and networks, and mathematical models utilizing chemical information. In addition to standard approaches based on support vector machines, neural networks, and random forests, more recent works have been exploring more sophisticated techniques to detect drug-drug and drug-target interactions. This includes the line of work in graph neural networks [48–50], multi-armed bandits for efficient sequential exploration of safe drug dose combinations [51], as well as explainability for disentangling the different effects of each drug on the outcomes predicted by other black-box machine learning models [52]. Different computational approaches to drug-drug interaction prediction have been increasingly benchmarked through competitions (such as AstraZeneca-Sanger DREAM challenge [53]), encouraging the development of more innovative approaches while resulting in more publicly available data. This allows for systematic comparison of the merits of different techniques through the same training and blinded test datasets, but it also carries the danger of overfitting to a (still relatively small) competition dataset and the proxy performance metric, resulting in difficulties in generalizing to novel contexts outside of the dataset as well as *in vivo* applications [54].

3.3. Using more general pharmacokinetic/pharmacodynamic data

Since precise measurements of the effect—whether target activation or clinical outcome—may never be fully observed, it is important to determine appropriate surrogate measurements and how to leverage them for better outcome estimates. The relationships between the surrogate and the true outcome need to be estimated in order to develop a good model. However, often this may only be possible in animal models and *in vitro* studies that do

not directly correspond to human patients. Understanding how to identify good biomarkers from observational data, how to transfer the knowledge across different species (animal to human) and experimental scales (*in vitro* molecular studies to *in vivo* clinical outcomes), and how to incorporate the knowledge of the disease mechanisms could inform better estimates of the PK/PD relationships. In addition, it might be necessary to go beyond any single observational biomarker, and leverage the changes of patterns in multiple inter-related markers: for example, Barbieri et al. successfully develop and externally validate a deep learning approach combining previous observations and patient characteristics for individualized anemia treatments in chronic renal failure patients [55]. Another relevant line of work in machine learning for biomarker identification and knowledge transfer includes language models for large-scale medical text preprocessing [56] and knowledge graph development [57–61], where link prediction could be used to quantify the relationships of interest. The use of these techniques, as well as modeling of richer individual data (e.g. chronic conditions, drug-drug interactions), are visualized as part of a more advanced machine learning pipeline for pharmacometrics in Figure 4.

4. Pharmacometric simulations

Aside from the use of simulation and sampling of popPK models as a way of generating additional training data for machine learning models, simulations are useful in their own right. The key uses of simulation in PK and PK/PD modeling include (i) evaluation of popPK models using visual predictive checks [62], (ii) *a priori* dose evaluation based on the probability of PK or PD target attainment, and (iii) aiding the simulation of clinical trials to answer “what if” questions [19] (Figure 2).

4.1. Simulation of popPK models

Machine learning can aid with improving the quality of popPK models which in turn would improve the simulated data. Monte Carlo simulations from popPK models are typically drawn from normal or log-normal parameter distributions [16,17,19]; however, this requires that the popPK and PK/PD models are well-specified, that true parameters come from required distributions, and that the dynamics conform to strong assumptions of linearity that are required to extrapolate to unseen populations. Typically, a popPK model is estimated from a rather small number of subjects with heterogeneous profiles from a homogeneous population. Thus many different models have to be developed for each subpopulation (e.g. neonates and adults, healthy and critically ill subjects). Additionally, small sample sizes may not allow the model to generalize to the rest of the

subpopulation, and one subpopulation may not generalize to another [63]. Bayesian model averaging approaches have been proposed to address that by performing model selection or creating an ensemble or averaged popPK model either at a population [26] or individualized level [64]. In particular, the latter method uses the demographic information summaries reported alongside published popPK models to find those the most applicable to a particular individual, and using that to create individualized model ensembles [64]. Another approach is to use continual learning [65] to dynamically adapt and refine prior popPK models to fit populations of interest, thereby improving predictions. Future work could also investigate generative model-based simulation approaches; however, this is subject to the major limitation of scarce data. Overcoming this would likely involve finding new and better ways of incorporating domain expertise as prior knowledge.

4.2. Clinical trial simulations

Clinical trial simulations (CTS) are extensively used in pharmacology, particularly to optimize the design of real-world clinical trials (including their sample size, dosing regimen, choice of control, duration, etc.) Estimating the possible outcomes of the trial before its execution—particularly phases II and III—is crucial, since the wrong selection of a treatment regimen and incorrect responder group identification could result in catastrophic losses in time and money for a drug company. As a result, it is a growing requirement from the regulators (such as FDA) to have a strong justification for the selected dose within the design of the trial (see e.g. [Project Optimus](#) for oncological drug development and [Friends of Cancer Research meeting report](#) [66]). While clinical trial simulations could act as part of this justification, they may rely on simplistic assumptions on how preclinical data from animal and *in vitro* experiments would translate to patients—it is hoped that machine learning could help overcome these assumptions and discover the connections in a data-driven manner. For example, Bedon et al. use machine learning to find the dose-limiting toxicity predictors in phase I of a chemotherapy treatment [67]; similar data-driven predictors could be used to identify relationships between individual features and expected or adverse effects—especially when they are nonlinear and linked to complex, multidimensional patterns. Machine learning, particularly developments in bandit literature, has also been proposed for identifying responder subgroups [51,68], identifying maximum tolerated dose under safety and efficacy constraints [69] and designing better trials [70,71]. Finally, in the post-market phase, clinical trial simulations can be used to propose recommendations in the case of delayed or missed administration of a drug [72,73], which could again be potentially improved with machine learning methods.

5. Optimal dose selection

5.1. Model-informed precision dosing

One of the main purposes of PK/PD modeling is to inform precision dosing for patients on the basis of their individual features, increasing the probability of target attainment, especially when therapeutic windows of the drug are narrow [74,75]. In addition to machine learning methods most widely used in pharmacology, such as linear regression, random forests and gradient boosting—methods used to predict a single exposure marker or probability of target attainment—other approaches can predict the entire time-series of drug concentrations in relation to individual features, the treatment regimen, and initial observations of treatment response. Recurrent neural networks have been recently proposed for estimating treatment outcomes over time [76], and generative adversarial networks have been used to model the effects of continuous-valued treatments [77]. Future work could investigate the potential of using such methods for representing both continuous doses and their administration frequencies together, capturing the full treatment regimen. Another direction would be to look into more powerful representations of popPK models as dynamical systems. Modeling of dynamical systems using machine learning is a field of growing interest, with a range of developments in symbolic regression [78–80], neural ODEs [81–87], or combinations of the two [10] for deriving data-driven and interpretable closed form expressions. In particular, Janssen et al show how combining deep neural networks with ODEs can improve on NLME modeling when proposing the first dose and simulating treatments [88]. More complex methods could be also used to model the changes in rate constants over time. Of course, many of these approaches require large amounts of available data—often a luxury in pharmacology, especially in rare diseases. However, machine learning could also help by integrating additional modalities of pharmacological knowledge beyond popPK models, such as known or latent factors of variability, stages of disease, relations to other drugs and animal studies, causal relationships, and other *inductive biases* that would improve the existing data efficiency for learning individual variability.

5.2. Decision support with therapeutic drug monitoring

In clinical practice, following an observation of an off-target value of an exposure marker (AUC, C₀, C_{max}, among others), the clinician or clinical pharmacologist uses their expertise to decide whether to change the dose by a certain amount, to keep the dose and ask for a control, or to wait until the next planned measurement in

order to improve the outcome. Such “manual” precision dosing decisions can be largely dependent on the past experience and biases of the individual clinician/pharmacologist. In this case the use of the more principled model-informed precision dosing could be helpful [89], although its introduction to routine practice is still challenging [90]. With the same caveats of sufficient data availability for both short- and long-term outcomes, machine learning could be leveraged to develop the next-generation decision support systems for therapeutic drug monitoring. Reinforcement learning is an increasingly popular technique, where therapeutic drug monitoring is reframed as a Markov decision process with an agent taking actions (prescribing treatments) based on the feedback from the environment (patient state) so as to maximize the reward (optimal exposure of the drug and positive treatment outcomes) [91,92]. Such agents are already being developed for sensitive drugs like heparin [93] and warfarin [94] as well as chemotherapy [92,95,96]. As machine learning techniques for pharmacometrics become more advanced and reliable, decision making for therapeutic drug monitoring could evolve into adaptive clinical trial designs, acquiring the most relevant information while minimizing adverse effects, as shown in Figure 5.

6. Systems pharmacology

Unlike pharmacometric disciplines of PK and PD which analyze drug mechanisms in isolation, (quantitative) systems pharmacology (QSP) considers these mechanisms within the broader context of a complex biological system, where a single reaction is only a small part of a large network of interacting metabolic pathways. Indeed, the reasons for why a drug may be effective for an individual patient may have to do with the differences in the downstream cascade of the events initiated by the drug-target binding, or the broader characteristics of individual physiopathology. Jointly modeling the mechanisms of disease through metabolic pathway systems, gene expression data as well as the direction of causality (e.g. considering inflammation and its modification in gene expression as a consequence rather than the cause of autoimmune conditions) may also lead to better understanding of the disease and inform new treatments. One example of such a shift is the *network pharmacology* paradigm, which moves beyond the ‘one gene, one drug, one disease’ approach and instead considers simultaneous treatment of multiple molecular targets within the disease network [97].

Systems pharmacology aims to model and integrate multiple data sources explicitly by drawing on mathematical modeling, experimental drug pharmacology, multi-omic data, and data from previous clinical trials and animal studies (Figure 4). In a white paper report published by the QSP special interest group at the International

Society of Pharmacometrics (ISoP QSP SIG) [98], it is proposed that integration of machine learning and traditional QSP could improve parameter estimation, model structure, and mitigating the problems in dimensionality reduction. Examples of machine learning approaches include integration of multi-omic (e.g., transcriptomic, genomic, epigenetic) data to identify patient subgroups for Sjögren syndrome for potential treatment individualization [99], and latent hybridization models integrating PK/PD data with longitudinal hospital measurements to improve outcomes of a COVID-19 treatment [10]. In particular, by connecting an *in vitro* PB/PK model for dexamethasone and certain inflammatory cytokines [100] to nonspecific patient measurements indicating inflammation (C-reactive protein), the latter approach serves as a compelling demonstration of the potential of machine learning in effectively connecting disparate sources of systems pharmacology and clinical knowledge to improve specific treatments.

Due to the diversity and complexity of data, such machine learning methods are rare and highly specialized; however, we believe in their potential and necessity for helping move the field forward in the future.

7. Conclusions and perspectives

In this article, we reviewed the potential synergies between pharmacometrics and machine learning to better model the pharmacokinetics and/or pharmacodynamics of drugs, and to use this knowledge to improve individual patient outcomes through precision medicine. We discussed the methods for each subfield of pharmacometrics—pharmacokinetics, pharmacodynamics—as well as pharmacometric simulations, model-informed precision dosing, therapeutic drug monitoring and quantitative systems pharmacology.

In particular, we expanded the discussion from machine learning for treatment effects [8] and their uncertainty to exposure/effect modeling, the use of surrogate markers as proxies for clinical outcomes of interest, and multimodal data integration. We have left out of scope, however, questions no less important to machine learning for pharmacology, including but not limited to principled data collection, data quality, value of information, integration of data from diverse healthcare systems, and continuous monitoring of patients.

Given the promising machine learning developments so far, one may ask why these approaches are not researched and applied more widely. The most important constraint is the availability of high-quality pharmacometric data (the initial and critical step for any successful machine learning system, as emphasized in Figures 4 and 5). Unlike relatively cheap and ubiquitous vital sign measurements using automatic sensors,

collection of pharmacometric data (such as drug concentrations at fine-grained timescales) is rare and expensive. While there exist some developments in the form of federated learning platforms [101,102] and electronic healthcare record databases, the data is still subject to high regulatory scrutiny and privacy issues, which make it less accessible. Just like for chemists, with no data to feed the data-hungry deep learning systems, the AI revolution for pharmacologists is yet to happen [103].

A potential solution to improving the data accessibility and model reproducibility issues could be the development of large, open datasets [104,105]. Such datasets could be developed by (1) anonymizing and releasing existing data or (2) generating synthetic datasets.

Anonymized datasets have the advantage of representing real patient cases, but, prepared improperly, could reveal sensitive data. They are also expensive to acquire—especially in quantities required for large-scale training—and are subject to regulations, further restricting their accessibility.

Synthetic datasets have been proposed as an alternative source of more accessible, low-cost training data. However, generative modeling—a machine learning strategy for creating synthetic datasets—has its own challenges. The choice of the generative model (such as variational autoencoder, generative adversarial network, diffusion model, etc.) depends on the situation at hand and may not be straightforward. Some models can provide more control of incorporated domain knowledge while being too rigid to fully represent real-world data; others could flexibly adapt to more complex data but uncontrollably overfit (and fail to generalize) due to overparameterization.

Currently rising in their popularity, deep generative models in particular depend on the availability of sufficient and diverse enough data for their training—which is a problem when it is precisely the lack of such data that motivates the use of generative modeling in medicine. For instance, recent successful generative models for image and text data, such as StableDiffusion [106] and ChatGPT [107], have all been trained on much larger datasets than would typically be available in the medical domain, while medical data is arguably even more complex in its structure. Overparameterized deep generative models, having lost the original context and meaning of data, could produce unrealistic or non-physiologically-viable examples, amplify spurious dependencies and biases, and leak private information by reproducing its own training data.

Principled synthetic dataset development and validation in collaboration of both machine learning and pharmacology experts is therefore crucial. A new class of pharmacology-specific generative models will be

needed to account for expert insights and relevant biological and pharmacological nuances, while fitting complex data in an explainable way. Domain expertise could also be used to filter spurious unrealistic data and correct biases. The quality and reliability of the synthetic data should always be validated in an independent setting; possibly by testing how a downstream model trained on generated data performs in unseen scenarios, compared to when trained on an independent real dataset.

The second constraint that prevents adoption of machine learning in pharmacometrics is making sure that the models are principled and trustworthy enough to be used in high-stakes settings for real-world patients, with challenges in relation to explainability of black-box models, quantification of uncertainty guarantees, and the regulatory approval for machine learning models in routine care. The naive ways of applying machine learning to pharmacometric data, as shown in Figure 2, can result in unreliable answers that cannot always be used for real-world patients—but further collaborative research in specialized, pharmacometrics-focused methods is likely to eventually make them more applicable. This is recognized by the regulatory agencies such as the US Food and Drug Administration [108]; however, rapid advancement of machine learning often outpaces the process of establishing a regulatory framework, which in turn requires a lot of careful consideration and time to strike the right balance between innovation and patient safety. Collaboration between researchers, industry stakeholders, and regulatory agencies is essential, and we are hopeful that—as the guidelines become clearer—the demonstration of efficacy and trustworthiness of machine learning methods will become more straightforward, leading to higher adoption.

The final constraint is the lack of appropriate *combined* expertise to build and apply more specialized pharmacometric machine learning models. While standard off-the-shelf machine learning packages (see Keras [109], scikit-learn [110], or tidymodels [111]) are relatively accessible to pharmacologist practitioners, the development and use of more complex machine learning models requires mathematical and statistical skill. The increasing availability of advanced machine learning tools therefore has the danger of increased prevalence of low-quality research due to unprincipled application. For example, a reproducibility study by Kapoor and Narayanan [112] shows that, in more than 300 of the scientific and medical papers surveyed, complex machine learning methods fail to reproduce and do not outperform baseline decades-old models such as logistic regression. Of particular note is the issue of data leakage, which may involve both leakage of test data into the

training procedure, or including information from the future or outside knowledge that would typically be unavailable in real-world settings. Such mistakes might inflate the performance of the model during training and test-time but perform poorly when faced with previously unseen real-world data.

Conversely, machine learning practitioners may not have the necessary pharmacometric expertise available in order to build on the domain knowledge. For example, when applied naively, machine learning models may learn correlations between variables that are not causally related, which would seem to produce good results in the training phase but perform poorly in test-time environments when such correlations do not hold. In this case, the pharmacological expertise may help mitigate such issues in time. Principled and expert-verifiable models (Figures 3-5) are nonstandard and require extensive knowledge of both pharmacometrics and machine learning as well as strong collaborative relationships, and, as always, must be carefully validated in external populations.

With this article, we hope to encourage new partnerships between the experts of both fields, and to inspire new methods and ideas to advance pharmacometric machine learning—the development of new larger datasets that consider the complexity of clinical data, new models that are able to integrate different kinds of pharmacometric domain knowledge, and user-friendly interfaces as well as rigorous evaluation and validation processes to streamline the wider adoption by the practitioners.

Acknowledgements: We thank Zhaozhi Qian, Alexander Terenin, Maciej Wiatrak, and Haoting Zhang for their helpful reviews and suggestions. KS thanks AstraZeneca for research funding.

Funding: No funding was received to assist with the preparation of this manuscript.

Ethics approval: not applicable

Consent to participate: not applicable

Consent for publication: not applicable

Availability of data and material: not applicable

Code availability: not applicable

Competing Interests: JBW, PM and MvdS have no conflict of interest to declare in relation to this work. KS receives research funding from AstraZeneca. RWP is an employee and stockholder of F. Hoffmann-La Roche.

Author Contributions: JBW, KS and RWP conceptualized the ideas, JWB and KS conducted the research and wrote the manuscript, KS made the figures, RWP and PM reviewed and edited the manuscript, MvdS acquired the funding and supervised the project.

References

1. Rawson TM, Gowers SAN, Freeman DME, Wilson RC, Sharma S, Gilchrist M, et al. Microneedle biosensors for real-time, minimally invasive drug monitoring of phenoxymethylpenicillin: a first-in-human evaluation in healthy volunteers. *The Lancet Digital Health*. 2019;1:e335–43.
2. Vanhove T, Annaert P, Kuypers DRJ. Clinical determinants of calcineurin inhibitor disposition: a mechanistic review. *Drug Metab Rev*. 2016;48:88–112.
3. Wicha SG, Mårtson A-G, Nielsen EI, Koch BCP, Friberg LE, Alffenaar J-W, et al. From Therapeutic Drug Monitoring to Model-Informed Precision Dosing for Antibiotics. *Clin Pharmacol Ther*. 2021;109:928–41.
4. Chabot GG. Factors involved in clinical pharmacology variability in oncology (review). *Anticancer Res*. 1994;14:2269–72.
5. Nahata MC. Variability in clinical pharmacology of drugs in children. *J Clin Pharm Ther*. 1992;17:365–8.
6. Wilkinson GR. Drug metabolism and variability among patients in drug response. *N Engl J Med*. 2005;352:2211–21.
7. Johnson M, Patel M, Phipps A, van der Schaar M, Boulton D, Gibbs M. The potential and pitfalls of artificial intelligence in clinical pharmacology. *CPT Pharmacometrics Syst Pharmacol*. 2023;12:279–84.
8. Bica I, Alaa AM, Lambert C, van der Schaar M. From Real-World Patient Data to Individualized Treatment Effects Using Machine Learning: Current and Future Methods to Address Underlying Challenges. *Clin Pharmacol Ther*. 2021;109:87–100.
9. Terranova N, Venkatakrishnan K, Benincosa LJ. Application of Machine Learning in Translational Medicine: Current Status and Future Opportunities. *AAPS J*. 2021;23:74.
10. Qian Z, Zame WR, Fleuren LM, Elbers PWG, Schaar M van der. Integrating Expert ODEs into Neural ODEs: Pharmacology and Disease Progression. *CoRR [Internet]*. 2021;abs/2106.02875. Available from: <https://arxiv.org/abs/2106.02875>
11. Mould DR, Upton RN. Basic concepts in population modeling, simulation, and model-based drug development-part 2: introduction to pharmacokinetic modeling methods. *CPT Pharmacometrics Syst Pharmacol*. 2013;2:e38.
12. Sibieude E, Khandelwal A, Girard P, Hesthaven JS, Terranova N. Population pharmacokinetic model selection assisted by machine learning. *J Pharmacokinet Pharmacodyn*. 2022;49:257–70.
13. Ismail M, Sale M, Yu Y, Pillai N, Liu S, Pflug B, et al. Development of a genetic algorithm and NONMEM workbench for automating and improving population pharmacokinetic/pharmacodynamic model selection. *J Pharmacokinet Pharmacodyn*. 2022;49:243–56.
14. Jaber MM, Yaman B, Sarafoglou K, Brundage RC. Application of Deep Neural Networks as a Prescreening Tool to Assign Individualized Absorption Models in Pharmacokinetic Analysis. *Pharmaceutics*. 2021;13:797.
15. Otalvaro JD, Yamada WM, Hernandez AM, Zuluaga AF, Chen R, Neely MN. A proof of concept reinforcement learning based tool for non parametric population pharmacokinetics workflow optimization. *Journal of Pharmacokinetics and Pharmacodynamics [Internet]*. 2022; Available from: <https://doi.org/10.1007/s10928-022-09829-5>
16. Guidi M, Csajka C, Buclin T. Parametric Approaches in Population Pharmacokinetics. *J Clin Pharmacol*. 2022;62:125–41.
17. Goutelle S, Woillard J-B, Neely M, Yamada W, Bourguignon L. Nonparametric Methods in Population Pharmacokinetics. *J Clin Pharmacol*. 2020;
18. Goutelle S, Woillard J-B, Buclin T, Bourguignon L, Yamada W, Csajka C, et al. Parametric and Nonparametric Methods in Population Pharmacokinetics: Experts' Discussion on Use, Strengths, and Limitations. *J Clin Pharmacol*. 2022;62:158–70.
19. Mould DR, Upton RN. Basic concepts in population modeling, simulation, and model-based drug development. *CPT Pharmacometrics Syst Pharmacol*. 2012;1:e6.
20. Jacqz-Aigrain E, Leroux S, Thomson AH, Allegaert K, Capparelli EV, Biran V, et al. Population pharmacokinetic meta-analysis of individual data to design the first randomized efficacy trial of vancomycin in neonates and young infants. *J Antimicrob Chemother*. 2019;74:2128–38.
21. Sibieude E, Khandelwal A, Hesthaven JS, Girard P, Terranova N. Fast screening of covariates in population models empowered by machine learning. *J Pharmacokinet Pharmacodyn*. 2021;48:597–609.

22. Karatza E, Papachristos A, Sivolapenko GB, Gonzalez D. Machine learning-guided covariate selection for time-to-event models developed from a small sample of real-world patients receiving bevacizumab treatment. *CPT Pharmacometrics Syst Pharmacol.* 2022;11:1328–40.
23. Asimwe IG, Blockman M, Cohen K, Cupido C, Hutchinson C, Jacobson B, et al. Stable warfarin dose prediction in sub-Saharan African patients: A machine-learning approach and external validation of a clinical dose-initiation algorithm. *CPT Pharmacometrics Syst Pharmacol.* 2022;11:20–9.
24. Ponthier L, Ensuque P, Destere A, Marquet P, Labriffe M, Jacqz-Aigrain E, et al. Optimization of Vancomycin Initial Dose in Term and Preterm Neonates by Machine Learning. *Pharm Res.* 2022;39:2497–506.
25. Franck B, Woillard J-B, Théorêt Y, Bittencourt H, Demers E, Briand A, et al. Population pharmacokinetics of ganciclovir and valganciclovir in paediatric solid organ and stem cell transplant recipients. *Br J Clin Pharmacol.* 2021;87:3105–14.
26. Uster DW, Stocker SL, Carland JE, Brett J, Marriott DJE, Day RO, et al. A Model Averaging/Selection Approach Improves the Predictive Performance of Model-Informed Precision Dosing: Vancomycin as a Case Study. *Clin Pharmacol Ther.* 2021;109:175–83.
27. Leroux S, Jacqz-Aigrain E, Biran V, Lopez E, Madeleneau D, Wallon C, et al. Clinical Utility and Safety of a Model-Based Patient-Tailored Dose of Vancomycin in Neonates. *Antimicrob Agents Chemother.* 2016;60:2039–42.
28. Bououda M, Uster DW, Sidorov E, Labriffe M, Marquet P, Wicha SG, et al. A Machine Learning Approach to Predict Interdose Vancomycin Exposure. *Pharm Res.* 2022;39:721–31.
29. Woillard J-B, Labriffe M, Prémaud A, Marquet P. Estimation of drug exposure by machine learning based on simulations from published pharmacokinetic models: The example of tacrolimus. *Pharmacol Res.* 2021;167:105578.
30. Labriffe M, Woillard J-B, Debord J, Marquet P. Machine learning algorithms to estimate everolimus exposure trained on simulated and patient pharmacokinetic profiles. *CPT Pharmacometrics Syst Pharmacol.* 2022;11:1018–28.
31. Tang B-H, Guan Z, Allegaert K, Wu Y-E, Manolis E, Leroux S, et al. Drug Clearance in Neonates: A Combination of Population Pharmacokinetic Modelling and Machine Learning Approaches to Improve Individual Prediction. *Clin Pharmacokinet.* 2021;60:1435–48.
32. Zwart TC, Moes DJAR, van der Boog PJM, van Erp NP, de Fijter JW, Guchelaar H-J, et al. Model-Informed Precision Dosing of Everolimus: External Validation in Adult Renal Transplant Recipients. *Clin Pharmacokinet.* 2021;60:191–203.
33. Goutelle S, Alloux C, Bourguignon L, Van Guilder M, Neely M, Maire P. To Estimate or to Forecast? Lessons From a Comparative Analysis of Four Bayesian Fitting Methods Based on Nonparametric Models. *Ther Drug Monit.* 2021;43:461–71.
34. 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;10:1150–60.
35. Destere A, Marquet P, Gandonnière CS, 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;61:1157–65.
36. Woillard J-B, Labriffe M, Debord J, Marquet P. Tacrolimus Exposure Prediction Using Machine Learning. *Clin Pharmacol Ther.* 2021;110:361–9.
37. Woillard J-B, Labriffe M, Debord J, Marquet P. Mycophenolic Acid Exposure Prediction Using Machine Learning. *Clin Pharmacol Ther.* 2021;110:370–9.
38. Woillard J-B, Salmon Gandonnière C, Destere A, Ehrmann S, Merdji H, Mathonnet A, et al. A Machine Learning Approach to Estimate the Glomerular Filtration Rate in Intensive Care Unit Patients Based on Plasma Iohexol Concentrations and Covariates. *Clin Pharmacokinet.* 2021;60:223–33.
39. 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;14:1530.
40. Desmée S, Mentré F, Veyrat-Follet C, Sébastien B, Guedj J. Nonlinear joint models for individual dynamic prediction of risk of death using Hamiltonian Monte Carlo: application to metastatic prostate cancer. *BMC Med Res Methodol.* 2017;17:105.
41. Tardivon C, Desmée S, Keroui M, Bruno R, Wu B, Mentré F, et al. Association Between Tumor Size Kinetics and Survival in Patients With Urothelial Carcinoma Treated With Atezolizumab: Implication for Patient Follow-Up. *Clin Pharmacol Ther.* 2019;106:810–20.
42. Rizopoulos D. Dynamic predictions and prospective accuracy in joint models for longitudinal and time-to-event data. *Biometrics.* 2011;67:819–29.

43. Arisido MW, Antolini L, Bernasconi DP, Valsecchi MG, Rebora P. Joint model robustness compared with the time-varying covariate Cox model to evaluate the association between a longitudinal marker and a time-to-event endpoint. *BMC Med Res Methodol.* 2019;19:222.
44. Lee C, Yoon J, Schaar M van der. Dynamic-DeepHit: A Deep Learning Approach for Dynamic Survival Analysis With Competing Risks Based on Longitudinal Data. *IEEE Trans Biomed Eng.* 2020;67:122–33.
45. Lee H-S, Zhang Y, Zame W, Shen C, Lee J-W, Schaar M van der. Robust Recursive Partitioning for Heterogeneous Treatment Effects with Uncertainty Quantification. 2020.
46. Lee C, Light A, Saveliev ES, van der Schaar M, Gnanapragasam VJ. Developing machine learning algorithms for dynamic estimation of progression during active surveillance for prostate cancer. *NPJ Digit Med.* 2022;5:110.
47. Bulusu KC, Guha R, Mason DJ, Lewis RPI, Muratov E, Kalantar Motamedi Y, et al. Modelling of compound combination effects and applications to efficacy and toxicity: state-of-the-art, challenges and perspectives. *Drug Discovery Today.* 2016;21:225–38.
48. Torng W, Altman RB. Graph Convolutional Neural Networks for Predicting Drug-Target Interactions. *J Chem Inf Model.* 2019;59:4131–49.
49. Al-Rabeah MH, Lakizadeh A. Prediction of drug-drug interaction events using graph neural networks based feature extraction. *Sci Rep.* 2022;12:15590.
50. Feng Y-H, Zhang S-W. Prediction of Drug-Drug Interaction Using an Attention-Based Graph Neural Network on Drug Molecular Graphs. *Molecules.* 2022;27:3004.
51. Lee H-S, Shen C, Zame WR, Lee J-W, Schaar M van der. SDF-Bayes: Cautious Optimism in Safe Dose-Finding Clinical Trials with Drug Combinations and Heterogeneous Patient Groups. *CoRR [Internet].* 2021;abs/2101.10998. Available from: <https://arxiv.org/abs/2101.10998>
52. Crabbe J, Qian Z, Imrie F, van der Schaar M. Explaining Latent Representations with a Corpus of Examples. In: Ranzato M, Beygelzimer A, Dauphin Y, Liang PS, Vaughan JW, editors. *Advances in Neural Information Processing Systems [Internet].* Curran Associates, Inc.; 2021. p. 12154–12166. Available from: <https://proceedings.neurips.cc/paper/2021/file/65658fde58ab3c2b6e5132a39fae7cb9-Paper.pdf>
53. Menden MP, Wang D, Mason MJ, Szalai B, Bulusu KC, Guan Y, et al. Community assessment to advance computational prediction of cancer drug combinations in a pharmacogenomic screen. *Nat Commun.* 2019;10:2674.
54. 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;26:511–24.
55. Barbieri C, Molina M, Ponce P, Tothova M, Cattinelli I, Ion Titapiccolo J, et al. An international observational study suggests that artificial intelligence for clinical decision support optimizes anemia management in hemodialysis patients. *Kidney Int.* 2016;90:422–9.
56. Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X, et al. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *CoRR [Internet].* 2020;abs/2007.15779. Available from: <https://arxiv.org/abs/2007.15779>
57. Wiatrak M, Iso-Sipila J. Simple Hierarchical Multi-Task Neural End-To-End Entity Linking for Biomedical Text. *Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis [Internet].* Online: Association for Computational Linguistics; 2020. p. 12–17. Available from: <https://aclanthology.org/2020.louhi-1.2>
58. Neil D, Briody J, Lacoste A, Sim A, Creed P, Saffari A. Interpretable Graph Convolutional Neural Networks for Inference on Noisy Knowledge Graphs. *CoRR [Internet].* 2018;abs/1812.00279. Available from: <http://arxiv.org/abs/1812.00279>
59. Smith DP, Oechsle O, Rawling MJ, Savory E, Lacoste AMB, Richardson PJ. Expert-Augmented Computational Drug Repurposing Identified Baricitinib as a Treatment for COVID-19. *Front Pharmacol.* 2021;12:709856.
60. Bonner S, Barrett IP, Ye C, Swiers R, Engkvist O, Bender A, et al. A Review of Biomedical Datasets Relating to Drug Discovery: A Knowledge Graph Perspective. *CoRR [Internet].* 2021;abs/2102.10062. Available from: <https://arxiv.org/abs/2102.10062>
61. Paliwal S, de Giorgio A, Neil D, Michel J-B, Lacoste AM. Preclinical validation of therapeutic targets predicted by tensor factorization on heterogeneous graphs. *Sci Rep.* 2020;10:18250.
62. Nguyen THT, Mouksassi M-S, Holford N, Al-Huniti N, Freedman I, Hooker AC, et al. Model Evaluation of Continuous Data Pharmacometric Models: Metrics and Graphics. *CPT Pharmacometrics Syst Pharmacol.* 2017;6:87–109.
63. El-Haffaf I, Guilhaumou R, Velly L, Marsot A. Using a Validated Population Pharmacokinetic Model for Dosing Recommendations of Continuous Infusion Piperacillin for Critically Ill Adult Patients. *Clin Pharmacokinet.* 2022;61:895–906.

64. Chan A, Peck R, Gibbs M, van der Schaar M. Synthetic Model Combination: A new machine-learning method for pharmacometric model ensembling. *CPT: Pharmacometrics & Systems Pharmacology* [Internet]. [cited 2023 May 23];n/a. Available from: <https://onlinelibrary.wiley.com/doi/abs/10.1002/psp4.12965>
65. Hughes JH, Tong DMH, Lucas SS, Faldasz JD, Goswami S, Keizer RJ. Continuous Learning in Model-Informed Precision Dosing: A Case Study in Pediatric Dosing of Vancomycin. *Clin Pharmacol Ther.* 2021;109:233–42.
66. Commissioner O of the. Project Optimus. FDA [Internet]. 2022 [cited 2022 Oct 26]; Available from: <https://www.fda.gov/about-fda/oncology-center-excellence/project-optimus>
67. Bedon L, Cecchin E, Fabbiani E, Dal Bo M, Buonadonna A, Polano M, et al. Machine Learning Application in a Phase I Clinical Trial Allows for the Identification of Clinical-Biomolecular Markers Significantly Associated With Toxicity. *Clin Pharmacol Ther.* 2022;111:686–96.
68. Curth A, Hüyük A, Schaar M van der. Adaptively Identifying Patient Populations With Treatment Benefit in Clinical Trials. 2022.
69. Shen C, Wang Z, Villar SS, Schaar M van der. Learning for Dose Allocation in Adaptive Clinical Trials with Safety Constraints. *CoRR* [Internet]. 2020;abs/2006.05026. Available from: <https://arxiv.org/abs/2006.05026>
70. Zame WR, Bica I, Shen C, Curth A, Lee H-S, Bailey S, et al. Machine learning for clinical trials in the era of COVID-19. *Stat Biopharm Res.* 2020;12:506–17.
71. Lee H-S, Shen C, Jordon J, Schaar M van der. Contextual Constrained Learning for Dose-Finding Clinical Trials. 2020.
72. Gu J-Q, Guo Y-P, Jiao Z, Ding J-J, Li G-F. How to Handle Delayed or Missed Doses: A Population Pharmacokinetic Perspective. *Eur J Drug Metab Pharmacokinet.* 2020;45:163–72.
73. Saint-Marcoux F, Woillard J-B, Monchaud C, Friedl J, Bocquentin F, Essig M, et al. How to handle missed or delayed doses of tacrolimus in renal transplant recipients? A pharmacokinetic investigation. *Pharmacol Res.* 2015;100:281–7.
74. Sheiner LB, Beal S, Rosenberg B, Marathe VV. Forecasting individual pharmacokinetics. *Clin Pharmacol Ther.* 1979;26:294–305.
75. Francke MI, Andrews LM, Le HL, van de Wetering J, Clahsen-van Groningen MC, van Gelder T, et al. Avoiding Tacrolimus Underexposure and Overexposure with a Dosing Algorithm for Renal Transplant Recipients: A Single Arm Prospective Intervention Trial. *Clin Pharmacol Ther.* 2021;110:169–78.
76. Bica I, Alaa AM, Jordon J, Schaar M van der. Estimating Counterfactual Treatment Outcomes over Time Through Adversarially Balanced Representations. *CoRR* [Internet]. 2020;abs/2002.04083. Available from: <https://arxiv.org/abs/2002.04083>
77. Bica I, Jordon J, Schaar M van der. Estimating the Effects of Continuous-valued Interventions using Generative Adversarial Networks. *CoRR* [Internet]. 2020;abs/2002.12326. Available from: <https://arxiv.org/abs/2002.12326>
78. Schmidt M, Lipson H. Distilling free-form natural laws from experimental data. *science.* 2009;324:81–85.
79. Brunton SL, Proctor JL, Kutz JN. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences.* 2016;113:3932–3937.
80. Qian Z, Kacprzyk K, van der Schaar M. D-CODE: Discovering Closed-form ODEs from Observed Trajectories. *International Conference on Learning Representations.* 2021.
81. Chen RT, Rubanova Y, Bettencourt J, Duvenaud D. Neural ordinary differential equations. *arXiv preprint arXiv:180607366.* 2018;
82. Yildiz C, Heinonen M, Lähdesmäki H. ODE2VAE: Deep generative second order ODEs with Bayesian neural networks. *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019).* 2019.
83. De Brouwer E, Simm J, Arany A, Moreau Y. Gru-ode-bayes: Continuous modeling of sporadically-observed time series. *33rd Conference on Neural Information Processing Systems (NeurIPS 2019).* 2019.
84. Rubanova Y, Chen RT, Duvenaud D. Latent ordinary differential equations for irregularly-sampled time series. *33th Conference on Neural Information Processing Systems (NeurIPS 2019).* 2019.
85. Kidger P, Morrill J, Foster J, Lyons T. Neural controlled differential equations for irregular time series. *34th Conference on Neural Information Processing Systems (NeurIPS 2020).* 2020.
86. Hochlehnert A, Terenin A, Saemundsson S, Deisenroth M. Learning Contact Dynamics using Physically Structured Neural Networks. In: Banerjee A, Fukumizu K, editors. *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics* [Internet]. PMLR; 2021. p. 2152–2160. Available from: <https://proceedings.mlr.press/v130/hochlehnert21a.html>
87. Saemundsson S, Terenin A, Hofmann K, Deisenroth M. Variational Integrator Networks for Physically Structured Embeddings. In: Chiappa S, Calandra R, editors. *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* [Internet]. PMLR; 2020. p. 3078–3087. Available from: <https://proceedings.mlr.press/v108/saemundsson20a.html>

88. Janssen A, Leebeek FWG, Cnossen MH, Mathôt RAA, for the OPTI-CLOT study group and SYMPHONY consortium. Deep compartment models: A deep learning approach for the reliable prediction of time-series data in pharmacokinetic modeling. *CPT: Pharmacometrics & Systems Pharmacology*. 2022;11:934–45.
89. Jelliffe R. Goal-oriented, model-based drug regimens: setting individualized goals for each patient. *Ther Drug Monit*. 2000;22:325–9.
90. Woillard J-B, Saint-Marcoux F, Debord J, Åsberg A. Pharmacokinetic models to assist the prescriber in choosing the best tacrolimus dose. *Pharmacol Res*. 2018;130:316–21.
91. Ribba B, Dudal S, Lavé T, Peck RW. Model-Informed Artificial Intelligence: Reinforcement Learning for Precision Dosing. *Clin Pharmacol Ther*. 2020;107:853–7.
92. Maier C, Hartung N, Kloft C, Huisinga W, de Wiljes J. Reinforcement learning and Bayesian data assimilation for model-informed precision dosing in oncology. *CPT Pharmacometrics Syst Pharmacol*. 2021;10:241–54.
93. Nemati S, Ghassemi MM, Clifford GD. Optimal medication dosing from suboptimal clinical examples: a deep reinforcement learning approach. *Annu Int Conf IEEE Eng Med Biol Soc*. 2016;2016:2978–81.
94. Zadeh SA, Street WN, Thomas BW. Optimizing Warfarin Dosing using Deep Reinforcement Learning. *CoRR [Internet]*. 2022;abs/2202.03486. Available from: <https://arxiv.org/abs/2202.03486>
95. Yauney G, Shah P. Reinforcement Learning with Action-Derived Rewards for Chemotherapy and Clinical Trial Dosing Regimen Selection. In: Doshi-Velez F, Fackler J, Jung K, Kale DC, Ranganath R, Wallace BC, et al., editors. *Proceedings of the Machine Learning for Healthcare Conference, MLHC 2018, 17-18 August 2018, Palo Alto, California [Internet]*. PMLR; 2018. p. 161–226. Available from: <http://proceedings.mlr.press/v85/yauney18a.html>
96. Padmanabhan R, Meskin N, Haddad WM. Reinforcement learning-based control of drug dosing for cancer chemotherapy treatment. *Math Biosci*. 2017;293:11–20.
97. Hopkins AL. Network pharmacology: the next paradigm in drug discovery. *Nat Chem Biol*. 2008;4:682–90.
98. Zhang T, Androulakis IP, Bonate P, Cheng L, Helikar T, Parikh J, et al. Two heads are better than one: current landscape of integrating QSP and machine learning. *J Pharmacokinet Pharmacodyn*. 2022;49:5–18.
99. Soret P, Le Dantec C, Desvaux E, Foulquier N, Chassagnol B, Hubert S, et al. A new molecular classification to drive precision treatment strategies in primary Sjögren's syndrome. *Nat Commun*. 2021;12:3523.
100. Dai W, Rao R, Sher A, Tania N, Musante CJ, Allen R. A Prototype QSP Model of the Immune Response to SARS-CoV-2 for Community Development. *CPT Pharmacometrics Syst Pharmacol*. 2021;10:18–29.
101. MELLODDY [Internet]. MELLODDY. [cited 2022 Nov 9]. Available from: <https://www.melloddy.eu>
102. Ogier du Terrail J, Leopold A, Joly C, Béguier C, Andreux M, Maussion C, et al. Federated learning for predicting histological response to neoadjuvant chemotherapy in triple-negative breast cancer. *Nat Med*. 2023;
103. For chemists, the AI revolution has yet to happen. *Nature*. 2023;617:438.
104. Chan AJ, Bica I, Hüyük A, Jarrett D, Schaar M van der. The Medkit-Learn(ing) Environment: Medical Decision Modelling through Simulation. *CoRR [Internet]*. 2021;abs/2106.04240. Available from: <https://arxiv.org/abs/2106.04240>
105. Qian Z, Ceber B-C, Schaar M van der. Synthcity: facilitating innovative use cases of synthetic data in different data modalities. 2023.
106. Stable Diffusion launch announcement [Internet]. Stability AI. [cited 2023 Jun 7]. Available from: <https://stability.ai/blog/stable-diffusion-announcement>
107. Introducing ChatGPT [Internet]. [cited 2023 Jun 7]. Available from: <https://openai.com/blog/chatgpt>
108. Health C for D and R. Artificial Intelligence and Machine Learning in Software as a Medical Device. FDA [Internet]. 2022 [cited 2023 Jun 7]; Available from: <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>
109. Chollet F, others. Keras [Internet]. 2015. Available from: <https://keras.io>
110. Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, et al. API design for machine learning software: experiences from the scikit-learn project. *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 2013. p. 108–122.
111. Kuhn M, Wickham H. Tidymodels: a collection of packages for modeling and machine learning using tidyverse principles. [Internet]. 2020. Available from: <https://www.tidymodels.org>
112. Kapoor S, Narayanan A. Leakage and the Reproducibility Crisis in ML-based Science [Internet]. arXiv; 2022. Available from: <https://arxiv.org/abs/2207.07048>
113. Lee C, van der Schaar M. Temporal Phenotyping using Deep Predictive Clustering of Disease Progression [Internet]. arXiv; 2020. Available from: <https://arxiv.org/abs/2006.08600>
114. Gerlowski LE, Jain RK. Physiologically based pharmacokinetic modeling: principles and applications. *J Pharm Sci*. 1983;72:1103–27.

115. Azer K, Kaddi CD, Barrett JS, Bai JPF, McQuade ST, Merrill NJ, et al. History and Future Perspectives on the Discipline of Quantitative Systems Pharmacology Modeling and Its Applications. *Frontiers in Physiology* [Internet]. 2021 [cited 2022 Dec 9];12. Available from: <https://www.frontiersin.org/articles/10.3389/fphys.2021.637999>
116. Agoram BM, Demin O. Integration not isolation: arguing the case for quantitative and systems pharmacology in drug discovery and development. *Drug Discov Today*. 2011;16:1031–6.

Table 1: Summary of pharmacometric tasks, their goals, methods, and opportunities for machine learning developed in this review.

Task	Goal	Pharmacometric methods	Challenges	Opportunities
Population pharmacokinetic (popPK) modeling	To predict individualized PK parameters and exposure indices based on individual demographic data only (<i>a priori</i>) or additionally incorporating individual drug concentration measurements (<i>a posteriori</i>)	<ul style="list-style-type: none"> • Nonlinear mixed effect (NLME) modeling [11,16,19] • Nonparametric models [17] 	<ul style="list-style-type: none"> • Structural model selection [12–14] • Nonlinearity in covariates/PK parameter associations; collinearity in covariates 	<ul style="list-style-type: none"> • AutoML and neural network-based approaches • Genetic algorithms, neural network and Cox survival based covariate selection methods [20,21] • Xgboost-based individual PK parameter prediction [31,38]
Exposure/effect modeling	To model the relationship between PK exposure indices and microscopic (molecular target) or macroscopic (clinical endpoint) effects on the body	<ul style="list-style-type: none"> • Molecular scale: NLME or nonparametric methods • Population scale: logistic regression, survival analysis, joint modeling of longitudinal outcomes and time-to-event data [42] 	<ul style="list-style-type: none"> • Finding the best exposure indices while taking into account nonlinear relationships, delays in response, changes in dynamics due to cumulative exposure • Leveraging longitudinal exposure marker and outcome observations • Drug interaction modeling 	<ul style="list-style-type: none"> • Knowledge graph models for disease biomarker identification from observational data and for PK/PD relationship modeling [57–59] • Recurrent neural network-based survival analysis [46] • Deep learning methods for temporal subgroup clustering [113] • Graph neural network and deep learning methods [48,50]
Pharmacometric simulations	To evaluate popPK models (via visual predictive checks) and <i>a priori</i> treatment regimens for PK/PD target attainment; to answer counterfactual questions in clinical trials	Monte Carlo simulations [16,17,19]	<ul style="list-style-type: none"> • Unknown underlying distributions for accurate Monte Carlo simulation • Simplistic assumptions on preclinical, animal, in vitro data for clinical trial simulations • Maximum tolerated dose identification under safety and efficacy constraints • popPK model ensembling at population or 	<ul style="list-style-type: none"> • Generative model-based simulations • Data-driven utilization of prior information [67] • ML-based (adaptive) clinical trial design [70,71] • Multi-armed bandit approaches [69] • ML-based or population-level Bayesian model averaging [26,28] • Representation learning-based individualized model

			individualized levels	averaging [64] <ul style="list-style-type: none"> Continual learning for dynamic refinement of popPK priors [65]
Optimal dose selection	To propose an optimal treatment regimen based on popPK and exposure/effect data, and to dynamically adjust the regimen as more data is collected	Joint modeling [42]	<ul style="list-style-type: none"> Refining dose proposal on the basis of individual features (individual concentration measurements and covariates) to increase the probability of target attainment Development of more powerful representations of popPK models as dynamical systems: Symbolic regression [78–80], neural ODEs [81–87], and combinations [88] 	<ul style="list-style-type: none"> Recurrent neural networks for estimating treatment outcomes over time [76] Use of generative models to model the effects of continuous-valued treatments [77] Reinforcement learning to improve dose selection [91,92] ML integration of pharmacological knowledge beyond popPK models (e.g. known or latent factors of variability, stages of disease, relations to other drugs and animal studies, omics data, causal relationships, and other inductive biases) [10,99]
Systems pharmacology	To integrate physiological and other modalities of biological knowledge into pharmacometrics and drug development	<ul style="list-style-type: none"> Physiologically based PK [114] Gene ontology and pathway analysis [115] Compound target and target pathway networks for different compounds and targets [116] 	<ul style="list-style-type: none"> Integration of multi-source and multi-scale data with linear or nonlinear relationships Improvement in parameter estimation, model structure and dimensionality reduction 	

List of Figures

Figure 1: Stages of pharmacometric modeling. (a) Data. A population-level dataset is collected, including attributes such as age, weight, sex, genetic information, as well as concentration measurements of the drug. (b) Compartmental model. An expert picks a model (such as the two-compartment model) that fits the drug dynamics. The compartmental model is associated with a set of differential equations parameterized by the rates of drug absorption, distribution, metabolism and elimination (ADME). (c) Pharmacokinetic (PK) modeling. The compartmental model with parameters fit to population data define an individualized concentration-time curve with associated exposure indices, such as peak and trough concentrations (C_{max}, C_0), and area under curve (AUC). (d) Pharmacodynamic (PD) modeling. Therapeutic and toxic effects of the drug depend on drug exposure. The drug dose is adjusted to fall within the therapeutic range, maximizing the efficacy of the drug while minimizing its toxicity.

Figure 2: Direct machine learning approaches to precision dosing. Machine learning models are often used simply as an off-the-shelf tool, involving none to minimal pharmacometric domain knowledge. (a) Supervised machine learning. A machine learning model, such as a neural network (alternatively a support vector machine, random forest, etc.) is trained to take in the patient inputs to return the correct output (such as a personalized dose or an exposure marker). These large nonlinear models can fit the datasets very well but may not provide much ground behind their predictions and are difficult to trust in high-stakes situations such as in healthcare. (b) Pharmacometric curves as the prediction target. Alternatively, instead of returning the final result immediately, the machine learning model could be trained to predict earlier parts of the pharmacometric pipeline—the use of neural ordinary differential equations (ODEs) to find the pharmacometric curve being one example. An expert can perform a visual predictive check to verify the plausibility of the curve, but in most cases there are no reliability guarantees.

Figure 3: Incorporation of pharmacometric domain practices. A more complicated pipeline showing how machine learning could be incorporated as part of the standard pharmacometric practice (see Figure 1). At the first stage of model selection, machine learning could be used to select the most appropriate compartmental structure, with techniques like symbolic machine learning used to derive the precise differential equations driving the dynamics. These explicit outputs can be observed and verified by the expert. After the parameter fitting stage, machine learning could be further used for individualized error correction and adjustment.

Figure 4: Enriching the models with broad pharmacometric knowledge through multi-source, multi-scale data. A variety of big data sources—including scientific studies, chemical properties for drugs and proteins, metabolic networks, past clinical trials—can be processed and consolidated using advanced machine learning techniques (such as natural language processing and graph neural networks) into large knowledge bases. Combined with richer individual data (including full medication and treatment history, full genome and proteome), joint pharmacometric and machine learning techniques could account for more complex factors (e.g. drug-drug interactions and systems biology), producing an individualized model of drug dynamics.

Figure 5: Use of machine learning models for decision making and further data acquisition. As joint machine learning (ML) and pharmacometric (PMx) developments progress further to be suitable for real-world decision making, techniques such as reinforcement learning could be used to design new adaptive clinical trials, collecting relevant information while improving individual patient outcomes.