Received: 14 June 2021 | Revised: 28 September 2021 | Accepted: 28 October 2021

DOI: 10.1002/psp4.12745

ARTICLE

A continued learning approach for model-informed precision dosing: Updating models in clinical practice

Corinna Maier[1,2] | **Jana de Wiljes**[1] | **Niklas Hartung**[1] | **Charlotte Kloft**[3] | **Wilhelm Huisinga**[1]

[1]Institute of Mathematics, University of Potsdam, Potsdam, Germany

[2]Graduate Research Training Program PharMetrX: Pharmacometrics & Computational Disease Modelling, Freie Universität Berlin, Berlin and University of Potsdam, Potsdam, Germany

[3]Department of Clinical Pharmacy and Biochemistry, Institute of Pharmacy, Freie Universität Berlin, Berlin, Germany

Correspondence
Wilhelm Huisinga, Institute of Mathematics, University of Potsdam, Karl-Liebknecht-Str. 24-25, 14476 Potsdam/Golm, Germany.
Email: huisinga@uni-potsdam.de

Funding information
Graduate Research Training Program PharMetrX: Pharmacometrics & Computational Disease Modelling, Berlin/Potsdam, Germany.
Deutsche Forschungsgemeinschaft (DFG) - SFB1294 - 318763901 (associated project).

Abstract

Model-informed precision dosing (MIPD) is a quantitative dosing framework that combines prior knowledge on the drug-disease-patient system with patient data from therapeutic drug/ biomarker monitoring (TDM) to support individualized dosing in ongoing treatment. Structural models and prior parameter distributions used in MIPD approaches typically build on prior clinical trials that involve only a limited number of patients selected according to some exclusion/inclusion criteria. Compared to the prior clinical trial population, the patient population in clinical practice can be expected to also include altered behavior and/or increased interindividual variability, the extent of which, however, is typically unknown. Here, we address the question of how to adapt and refine models on the level of the model parameters to better reflect this real-world diversity. We propose an approach for continued learning across patients during MIPD using a sequential hierarchical Bayesian framework. The approach builds on two stages to separate the update of the individual patient parameters from updating the population parameters. Consequently, it enables continued learning across hospitals or study centers, because only summary patient data (on the level of model parameters) need to be shared, but no individual TDM data. We illustrate this continued learning approach with neutrophil-guided dosing of paclitaxel. The present study constitutes an important step toward building confidence in MIPD and eventually establishing MIPD increasingly in everyday therapeutic use.

Study Highlights

WHAT IS THE CURRENT KNOWLEDGE ON THE TOPIC?

Current strategies to model-informed precision dosing (MIPD) assume a "perfect model scenario" (i.e., the prior model is not adapted or improved as new patient data are observed).

WHAT QUESTION DID THIS STUDY ADDRESS?

How could the underlying model used in MIPD be continuously improved and adapted to the target clinical setting?

WHAT DOES THIS STUDY ADD TO OUR KNOWLEDGE?

We propose a continued learning approach building on a sequential hierarchical Bayesian framework to improve the model while MIPD is applied.

HOW MIGHT THIS CHANGE DRUG DISCOVERY, DEVELOPMENT, AND/OR THERAPEUTICS?

The approach bears the potential to increase the applicability of MIPD by adapting models retrieved from literature to a target clinical setting.

INTRODUCTION

Model-informed precision dosing (MIPD) is a quantitative framework for dose individualization based on modeling and simulation of exposure-response relationships integrating patient-specific data.[1–3] The underlying models are developed based on clinical trial data typically using a nonlinear mixed effects (NLME) framework to describe the pharmacokinetics (PK) or pharmacodynamics (PD) of the drug and the variability between patients.[4] These PK/PD models allow to forecast important aspects of the therapy outcome based on patient characteristics (a priori predictions). Therapeutic drug/biomarker monitoring (TDM) allows to further individualize model predictions (a posteriori predictions) and subsequently to adjust dosing.

When PK/PD models are used for MIPD, a "perfect model scenario" is generally assumed, in which the model represents the drug-patient-disease system sufficiently well, the variability of the outcome is adequately described and the prior study population (used to *develop* the model) is representative of the target individual patient (to which the model will be *applied*). A certain model misspecification or population shift can, however, be expected due to the limited amount of data the models were built on. Specifically, the data from clinical trials only involve a limited number of patients, selected according to strict inclusion/exclusion criteria within a restricted time frame.[1,5,6] Therefore, models underlying MIPD will inevitably be confronted with deviating data in clinical routine, such as differences related to pathophysiology[7] or the patient population (comorbidities, comedications, or special characteristics, e.g., morbidly obese, pregnant, or rare genotypes).[1,5,8–10]

In this "imperfect model scenario," the benefits of MIPD approaches may not reach their full potential. It is therefore prudent to also improve the associated models as clinical routine data on the observed patient population becomes available.

For a given drug-disease-patient-system, there are often numerous models available within the literature, often based on different patient populations (e.g., for warfarin,[11,12] vancomycin[7,13] or cyclosporin).[14] In addition, adjustments to the model used in an MIPD framework were necessitated after treatment of the first patient cohort[15] or in retrospect.[16,17] As an example of high clinical relevance, we focus on paclitaxel causing neutropenia as the most frequent and life-threatening toxicity in oncology. Models describing paclitaxel-induced neutropenia build the basis for neutrophil-guided MIPD to individualize chemotherapy dosing.[18–21] Since the publication of the gold-standard model for neutropenia,[22] many model variants have been developed, which differ not only in parameter estimates,[23–26] but also in their structure.[17,27–29]

The challenge to choose between competing models is often approached via model averaging or model selection.[13] In model averaging, all candidate models are used, weighting the model predictions with the TDM data. In contrast, in model selection, a single model is selected based on a retrospective external evaluation on independent data collected previously in the intended setting (from the same hospital and patient population)[6] and prospective fit-for-purpose verification.[7] None of the approaches, however, integrates the new data collected during routine application of MIPD into the initial models underlying MIPD. Yet, continued learning approaches based on an ever-growing amount of data have enormous potential to improve the predictive capabilities of MIPD in clinical practice.

The problem of transferability is a well-known problem in the machine-learning literature; often called lifelong learning, continual learning,[30,31] transfer learning,[32,33] or domain adaptation.[34] Contrary to typical machine-learning applications with direct access to big data, sensitive patient data may not be accessible and available to this extent across different sources. Consequently, current approaches based on pooling of data are not an option in this case, because they require direct access to TDM data of all patients.[35] Hence, there is a need for approaches for continued model updating that are based on summary information of the data that is extracted locally and can be shared.

In this article, we propose an approach building on a sequential hierarchical Bayesian framework[36,37] for continued model learning. The underlying prior model used within MIPD is improved as new data from the target patient population are assimilated. Importantly, the approach separates inference of the individual model

parameters during a patient's therapy from the update of population parameters across patients. The proposed approach is based on ideas from Bayesian integration of meta-analyses.[36,37] First, we demonstrate how a mismatch between the model and data-generating process could affect MIPD in an in silico trial setting in terms of model parameters and structural misspecifications. We focus in the present article on how to adapt a model on the level of the typical model parameters and the magnitude of the inter-individual variability. The proposed approach aims at bridging the gap between population analyses in model development and application of MIPD in therapeutic use.

METHODS

MIPD framework

In MIPD, a Bayesian framework is used. Here, we very briefly summarize the approach, for details see refs [20 and 21]. Inference for a given (i-th) patient is based on a prior distribution $p(\cdot|\widehat{\theta^{\mathrm{TV}}},\hat{\Omega})$ and the likelihood $L(\cdot|y_i)=p(y_i|\cdot)$ of TDM data $y_i=(y_{i1},\ldots,y_{1n_i})^T$, resulting in the posterior $p(\cdot|y_i,\widehat{\theta^{\mathrm{TV}}},\hat{\Omega})$ of the individual parameters θ_i with

$$p(\theta_i|y_i,\widehat{\theta^{\mathrm{TV}}},\hat{\Omega})\propto p(y_i|\theta_i)p(\theta_i|\widehat{\theta^{\mathrm{TV}}},\hat{\Omega}). \tag{1}$$

The prior distribution in Equation 1 is based on prior population analyses leading to the following statistical model:

$$\Theta_i\sim p(\cdot|\widehat{\theta^{\mathrm{TV}}},\hat{\Omega}) \tag{2}$$

with $\widehat{\theta^{\mathrm{TV}}}$ denoting the estimates of the typical values (TV) and $\hat{\Omega}$ the estimated magnitude of interindividual variability (IIV). We assume in the following a normal distribution for $p(\cdot|\widehat{\theta^{\mathrm{TV}}},\hat{\Omega})$, which can typically be established via transformation (e.g., log-transformation in case of the log-normal distribution). Note that it is possible to consider a covariate dependent model, but we restrict ourselves to a more fixed parametrization of the model parameters $\widehat{\theta^{\mathrm{TV}}}$. The likelihood is based on structural and observational models as follows:

$$\frac{\mathrm{d}x_i}{\mathrm{d}t}(t)=f(x_i(t);\theta_i,d_i),\quad x_i(0)=x_0(\theta_i) \tag{3}$$

$$h_i(t)=h(x_i(t),\theta_i) \tag{4}$$

with state vector $x_i=x_i(t)$ (e.g., drug and neutrophil concentrations) and rates of change $f(x_i;\theta_i,d_i)$ for given doses d_i. The initial conditions $x_0(\theta_i)$ are given by the pretreatment levels (e.g., baseline neutrophil concentration). A statistical model links the observables $h_{ij}(\theta_i)=h_i(t_{ij})$ (i.e., model quantities that are measurable at time points t_{ij}), to observations (t_{ij},y_{ij}) for $j=1,\ldots,n_i$ taking into account measurement errors and potential model misspecifications. This specifies the likelihood in Equation 1; a common likelihood is defined by the following:

$$p(y_{ij}|\theta_i)=h_{ij}(\theta_i)+\varepsilon_{ij};\quad \varepsilon_{ij}\sim_{iid}\mathcal{N}(0,\Sigma) \tag{5}$$

for $j=1,\ldots,n_i$. For ease of notation, we do not explicitly state the dependence on h_{ij} and Σ in $p(y_i|\theta_i)$.

In the simulation study, we considered the parameters in log-space and consider neutrophil-guided dosing (i.e., based on the observed neutrophil concentration). From the lowest neutrophil concentration within a treatment cycle (nadir), the neutropenia grade of a chemotherapy cycle is inferred, ranging from no neutropenia (grade 0), mild (grade 1), moderate (grade 2), to severe (grade 3), and life-threatening (grade 4). We used an MIPD approach based on data assimilation (DA) to forecast the neutropenia time course, called DA-guided dosing, as presented in detail in refs. 20, and 21. In short, a particle filter is used to approximate the posterior distribution (Equation 1) at dose selection time points t_c (start of cycle c), integrating data from the i-th patient up to t_c (i.e., (t_{ij},y_{ij}) with $t_{ij}\le t_c$, via a sample approximation):

$$p(\theta_i|y_{i1:t_c})\approx\sum_{m=1}^{M}w_i^{(m)}1_{\left\{\theta_i^{(m)}=\theta_i\right\}}. \tag{6}$$

It is based on an ensemble of weighted particles (samples) $\mathcal{E}_i=\left\{\theta_i^{(m)},w_i^{(m)}:m=1,\ldots,M\right\}$ comprising parameter values $\theta_i^{(m)}$ and importance weights $w_i^{(m)}$. The used particle filter included a resampling and rejuvenation step, see ref. [20]. Solving the structural and observational model (Equation 3 + Equation 4) for each particle (sample) allows to compute the a posteriori probabilities of all neutropenia grades. The optimal dose is determined by minimizing the weighted joint probability of life-threatening grade 4 ($P_{\mathrm{grade\ 4}}$) and subtherapeutic grade 0 neutropenia ($P_{\mathrm{grade\ 0}}$) for the next cycle, i.e.,:

$$d^*=\underset{d\in\mathcal{D}}{\operatorname{argmin}}\left\{\lambda_1P_{\mathrm{grade4}}+\lambda_2P_{\mathrm{grade0}}\right\}. \tag{7}$$

The weighting factors were chosen as $\lambda_1=2/3$ and $\lambda_2=1/3$ to penalize the risk to expose a patient to life-threatening infections (grade 4) more severely than the risk to expose a patient to a subtherapeutic outcome (grade 0). The latter has been associated with reduced median overall survival.[38,39]

Simulation study framework

Paclitaxel-induced neutropenia models

We considered a model for paclitaxel-induced neutropenia which, in ref. [17] had been investigated for model applicability, re-estimated, and structurally modified after new patient data were observed in the clinical trial "Central European Society for Anticancer Research (CESAR) study of Paclitaxel Therapeutic Drug Monitoring (CEPAC-TDM)" ClinicalTrials.gov Identifier: NCT01326767.[40] A schematic representation of the models, corresponding parameter estimates, and typical model predictions are shown in Figure 1. In Supplementary Material Table S2, we list additional models proposed for paclitaxel-induced neutropenia, which illustrates the challenge of choosing a suitable model for MIPD in practice. The initial model (hereafter *gold-standard*) builds on the structure of the gold-standard model for chemotherapy-induced neutropenia[22] with parameter values estimated using a pooled data set of two prior studies,[25,41] including patients with ovarian cancer, non-small cell lung cancer (NSCLC), and various solid tumours.[26] Paclitaxel was given either as monotherapy or in combination with carboplatin. The CEPAC-TDM study included only patients with NSCLC and paclitaxel was given in combination with carboplatin or cisplatin over six treatment cycles. It was observed that the gold-standard model[26] overestimated the neutrophil concentration at

(a)

Parameters	Gold-standard	Gold-standard R	Bone marrow exhaustion
	Joerger et al. (2012)	Henrich (2017) Thesis	Henrich et al. (2017)
TV			
MTT [h]	141	128 (2.03)	145 (2.65)
Slope [L/μmol]	2.6	4.48 (4.55)	13.1 (4.56)
γ	0.2	0.231 (6.79)	0.257 (5.53)
ftr	-	-	0.787 (2.76)
IIV (CV%)			
MTT	27.0	-	-
Slope	44.9	43.8 (8.23)	44.8 (6.54)
ANC_0(=$Circ_0$)	31.6	60.3 (3.27)	51.5 (3.61)
RUV (CV%)			
exp. model	31.6	60.3 (3.27)	51.5 (3.61)

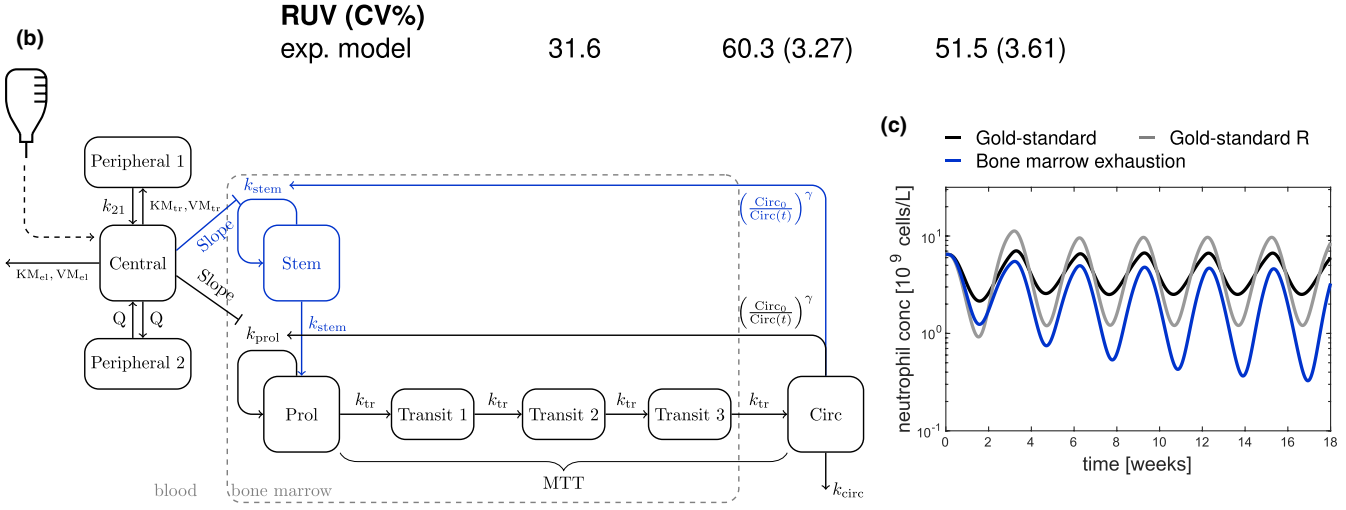

FIGURE 1 Paclitaxel-induced neutropenia models. (a) Parameter estimates of the three considered models with relative standard errors in brackets (if available in literature). (b) Schematic model representation with the PK model for paclitaxel (on the left); gold-standard PK/PD model in black and extension for cumulative neutropenia in black/blue. (c) Time course of neutrophil concentrations for the typical patient in the CEPAC-TDM study during six cycles of chemotherapy (3 weeks each, with drug administration on day 1) for the two model structures (black: gold/standard model trajectory; blue: extended gold-standard) and different parameter values (dark/light black). Abbreviations: $ANC_0 = Circ_0$: absolute neutrophil concentration at baseline or circulating neutrophils (Circ) at baseline; BILI, bilirubin concentration; BSA, body surface area; CEPAC-TDM, Central European Society for Anticancer Research (CESAR) study of Paclitaxel Therapeutic Drug Monitoring; CV, coefficients of variation; ftr, fraction of input in the compartment of proliferating cells (Prol) via replication; IIV, interindividual variability; MTT, mean transit time; PD, pharmacodynamic; PK, pharmacokinetic; RUV, residual variability; Slope, linear drug effect parameter; TV, typical values; γ, feedback exponent. More detailed information on the models is provided in Supplementary Material Section S1.

later cycles because the model does not account for cumulative neutropenia (i.e., an aggravation of neutropenia over multiple cycles,[17] see Figure 1c); a phenomenon that has been reported previously.[42] The parameters were re-estimated (*gold-standard R*) based on the CEPAC-TDM data, and finally the structure was modified to account for bone marrow exhaustion (BME), see Figure 1. Here, we focus our analyses on the more challenging PD models, whereas we considered the PK model to be given, with parameter values inferred previously based on the CEPAC-TDM study data,[17] see Supplementary Material Section S1.

Model adaptation scenarios

A model adaptation may be needed on the level of the structural model (3)+(4), prior parameter distribution (2) and/or likelihood (5). In this article, we consider two types of scenarios where model adaptations may be beneficial:

Structural differences

A divergence in the structural model (e.g., due to the manifestation of phenomena in the target patient population that have not been observed in the prior studies) is considered. To study such a scenario, we used the BME model[17] to generate TDM data, whereas we used the gold-standard model[26] for MIPD. The latter model lacks the structural feature of cumulative neutropenia.

Differences in parameters

Differences in the parameter distribution (e.g., the distributional assumption [normal, log-normal, etc.]) as well as the estimated hyperparameter values for a given distribution are potential examples. Here, we only focus on the latter (i.e., the type of distribution is the same, but the hyperparameters differ). To study parameter changes, we used the gold-standard R model[43] to generate TDM data, whereas we used the gold-standard model[26] in MIPD. Both rely on the same structural model, whereas the parameter values of the former were re-estimated to the CEPAC-TDM data.

We compared the performance of MIPD in the presence of structural or parameter bias to (i) MIPD based on an unbiased model (unbiased model scenario), and (ii) standard dosing (i.e., 200 mg/m^2 body surface area), including a dose reduction of 20% if grade 4 was observed based on the neutrophil measurement at day 15.[43]

TDM sampling scenarios

The effect of a potential mismatch between prior model and the (new) data-generating process of MIPD depends on the amount of available TDM data per patient to adapt the model. Therefore, we considered different TDM sampling schemes:

1. *Sparse sampling*: Neutrophil measurements at day 1 and day 15 of each cycle (sampling design of CEPAC-TDM study).
2. *Intermediate sampling*: Weekly neutrophil measurements (as in ref. [25]).
3. *Rich sampling*: Neutrophil measurements taken every third day.

Although the first sampling scheme corresponds to a common outpatient setting, the third is standard in a fully inpatient setting and also mimics the prospective growing availability of point-of-care devices for an outpatient setting (e.g., HemoCue WBC Diff for measuring neutrophil counts[44]), foreseeing richer sampling for monitoring patients.

Hierarchical Bayesian model

To continuously update and learn population parameters, we considered additional hyper priors on the population parameters of the NLME models described in Section **MIPD framework**. The hierarchical structure of fully Bayesian population models thus comprises three stages,[45,46] see Figure 2: (i) the statistical model for the TDM data given by Equation 5 describes the deviations between the individual model predictions and the observational data; (ii) the distributional assumption for inter-individual variability, Equation 2 describes the differences between individuals; (iii) the distributional assumption for (hyper) population parameters, $p\left(\theta^{\mathrm{TV}}\right)$, $p\left(\Omega\right)$ describes the uncertainty in the population parameters.

Population analyses are typically performed in a frequentist NLME setting, reporting maximum likelihood estimates (MLEs) of the population parameters jointly with their relative standard errors (RSEs) or coefficients of variation. This leaves the problem of how to determine suitable hyper prior distributions of the population parameters. We considered a normal distribution for the typical parameters (on log-scale) and an inverse-Wishart for the variances, as suggested in ref. [47]. We chose θ^{TV} to be normally distributed with mean $\overline{\theta^{\mathrm{TV}}}$ identical to the MLE $\widehat{\theta^{\mathrm{TV}}}$ and variance $\mathbf{S}^{\theta^{\mathrm{TV}}}$ identical to the (appropriately transformed) squared standard error $(\mathrm{SE}_{\theta^{\mathrm{TV}}})^2$, i.e. $\theta^{\mathrm{TV}} \sim \mathcal{N}\left(\widehat{\theta^{\mathrm{TV}}}, \left(\mathrm{SE}_{\theta^{\mathrm{TV}}}\right)^2\right)$, see also ref. [48]. The inter-individual variability matrix Ω was assumed to be inverse-Wishart distributed $\mathcal{IW}(\Psi, \nu)$ and diagonal with parameters Ψ, ν such that the population estimate equaled the mean $\Psi/(\nu - n_\Omega - 1)$, i.e., $\Psi = (\nu - n_\Omega - 1)\widehat{\Omega}$, where

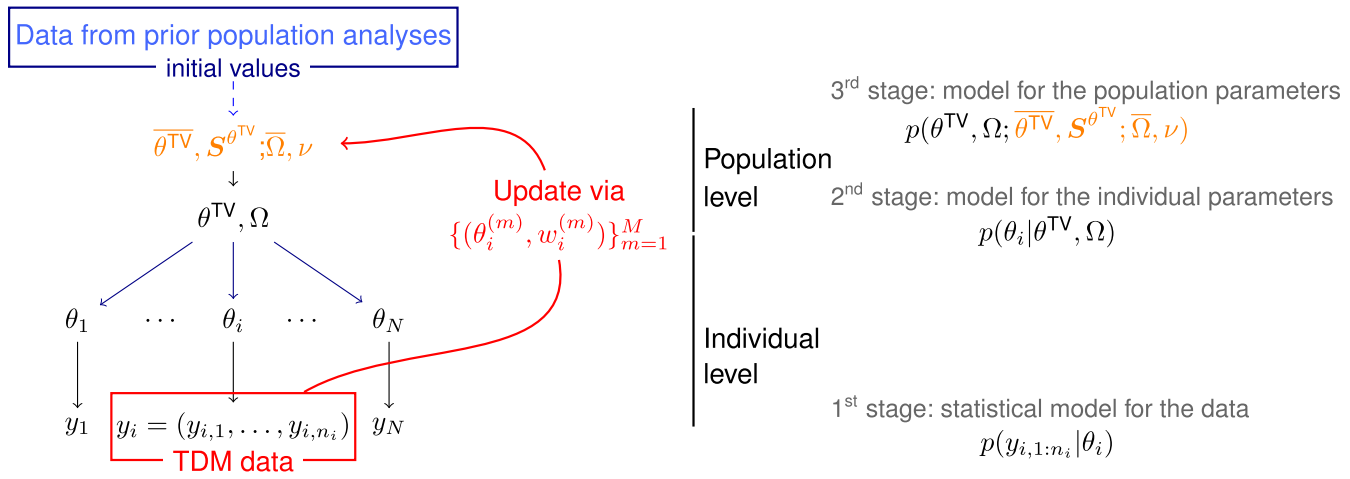

FIGURE 2 Hierarchical Bayesian model framework with separation between the inference on the individual level and the inference on the population level. The standard population analyses typically used to build the prior knowledge based on population estimates $\widehat{\theta^{TV}}, \widehat{\Omega}$ is shown in dark blue. To update the prior population estimates, the population parameters are seen as random variables with parametric probability distributions parametrized with hyperparameters $\overline{\theta^{TV}}, \mathbf{S}^{\theta^{TV}}, \overline{\Omega}$ and ν. A sample representation of the individual posterior, $\left\{\theta_{N+1}^{(m)}\right\}_{m=1}^{M}$ is used to update the hyperparameters of the population parameter distributions (red arrow). On the right the corresponding probability distributions are given for the different levels of the hierarchical model. TDM, therapeutic drug monitoring

n_Ω is the number of random effect parameters and with degrees of freedom ν still to be chosen. The distributions of the typical and variability values were assumed to be independent.

Continued learning across patients

To learn and improve a model across patients, the information provided by the patient-specific TDM data have to be included into the hierarchical model. In mathematical terms, we are interested in the marginal posterior:

$$p(\theta^{TV}, \Omega|y_{1:i}) \propto \int p(\theta_i, \theta^{TV}, \Omega|y_{1:i}) d\theta_i \quad (8)$$

with the joint posterior:

$$p(\theta_i, \theta^{TV}, \Omega|y_{1:i}) \propto p(y_i|\theta_i) p(\theta_i|\theta^{TV}, \Omega) p(\theta^{TV}, \Omega|y_{1:i-1}) \quad (9)$$

determined from a full hierarchical Bayesian procedure. A sample approximation to the joint posterior in Equation 9 allows for a straightforward approximation of the marginal in Equation 8. In the context of particle filter-based inference, this can be realized by augmenting the particle state as well as the parameter space by the population parameters (θ^{TV}, Ω). This approach, however, has two major drawbacks: (i) it is computationally expensive and thus limits real-time inference during the patient's therapy; and (ii) direct access to the individual patient data is required to update the population parameters yet data protection laws and logistical reasons often prohibit this. These are major limitations for a practical application, in particular across different clinics.

Therefore, we propose a two-level sequential Bayesian approach; it is based on previous ideas on Bayesian inference for meta-analyses.[36,37] Importantly, this approach does not change the inference on the individual level (see Algorithm 1 for pseudo-code):

1. *Individual level*: Estimate individual parameters of the ith patient.

$$p(\theta_i|y_i, \theta^{TV}, \Omega) \propto p(y_i|\theta_i) p(\theta_i|\theta^{TV}, \Omega),$$

for example, using a particle filter, sampling importance resampling or a Markov chain Monte Carlo approach.[20] A particle filter ('DA' for data assimilation in the pseudo-code) is used in our analysis, as it was shown to be best suited for the underlying setting.[20] This gives rise to a sample representation $\left\{\theta_i^{(m)}, w_i^{(m)}: m = 1, \ldots, M\right\}$ of the posterior $p(\theta_i|y_i, \theta^{TV}, \Omega)$, summarizing the information provided by the data of the ith patient. This step is identical to the inference step in MIPD without continuous learning, see Equation 1.

2. *Population level*: Update population parameters by sampling iteratively from the joint posterior $p(\theta_i, \theta^{TV}, \Omega|y_{1:i})$ via a Metropolis-Hastings-within-Gibbs sampling scheme,[36,48] (i.e., sampling from the full conditionals [see Section S2.1 for a detailed derivation]):

$$p(\theta^{TV}|\theta_i, \Omega, y_{1:i}) \propto p(\theta_i|\theta^{TV}, \Omega) p(\theta^{TV}|\Omega, y_{1:i-1}) \quad (10)$$

$$p(\Omega|\theta_i, \theta^{\mathrm{TV}}, y_{1:i}) \propto p(\theta_i|\theta^{\mathrm{TV}}, \Omega) p(\Omega|\theta^{\mathrm{TV}}, y_{1:i-1}) \quad (11)$$

$$p(\theta_i|\theta^{\mathrm{TV}}, \Omega, y_{1:i}) \propto p(y_i|\theta_i) p(\theta_i|\theta^{\mathrm{TV}}, \Omega) \quad (12)$$

Sampling from Equation 12 is achieved via a Metropolis-Hastings step, using as proposals the posterior samples generated on the individual level $\{\theta_i^{(m)}\}_{m=1}^M$, which are drawn according to weights $w_i^{(m)}$.

To ensure that sampling in Equations 10 and 11 can be performed in closed-form, at the end of assimilating data of the i-th patient, a parametric approximation by a normal-inverse-Wishart distribution is used:

$$p(\theta^{\mathrm{TV}}, \Omega|y_{1:i}) \approx \mathcal{N}\left(\overline{\theta^{\mathrm{TV}}}_i, \mathbf{S}_i^{\theta^{\mathrm{TV}}}\right) \otimes \mathcal{IW}((\nu_i - n_\Omega - 1)\overline{\Omega}_i, \nu_i), \quad (13)$$

with hyperparameters as stated in Algorithm 1. Then, sampling in Equations 10 and 11 corresponds to sampling from a normal and inverse-Wishart distribution, respectively, see Supplementary Material Section S 2.2.

Importantly, through the parametric approximation in Equation 13, $y_{1:i-1}$ are represented implicitly via the updated priors for the i-th step, while y_i enters implicitly through the sample approximation. In no case, the original patient data is needed.

Algorithm 1 Two-level sequential hierarchical Bayesian learning in MIPD

Input: $\widehat{\theta^{\mathrm{TV}}}, \mathrm{SE}_{\widehat{\theta^{\mathrm{TV}}}}, \hat{\Omega}, \nu_0$, ($y_{i,1:n_i}$ only for individual level)
Set hyper prior parameters $\overline{\theta_0^{\mathrm{TV}}} := \widehat{\theta^{\mathrm{TV}}}, \mathbf{S}_0^{\theta^{\mathrm{TV}}} := (\mathrm{SE}_{\widehat{\theta^{\mathrm{TV}}}})^2, \bar{\Omega}_0 := \hat{\Omega}, \nu_0$
for $i = 1 : N_{\mathrm{TDM}}$ **do**
  ▷ Individual level
  Initialize particle ensemble $\{(\theta_{i0}^{(m)}, x_{i0}^{(m)}, w_{i0}^{(m)})\}_{m=1}^M$ based on $p(\theta|\overline{\theta_{i-1}^{\mathrm{TV}}}, \bar{\Omega}_{i-1})$
  for $j = 1 : n_i$ **do**
    $\{(\theta_{ij}^{(m)}, x_{ij}^{(m)}, w_{ij}^{(m)})\}_{m=1}^M \leftarrow \mathrm{DA}\big(y_{ij}, \{(\theta_{ij-1}^{(m)}, x_{ij-1}^{(m)}, w_{ij-1}^{(m)})\}_{m=1}^M\big)$
  end for
  ▷ Population level
  Initialize Markov chain $\theta_i^{\mathrm{TV}(0)} = \overline{\theta_{i-1}^{\mathrm{TV}}}, \Omega_i^{(0)} = \bar{\Omega}_{i-1}$ and $\theta_i^{(0)}$ sampled from $p(\theta|\overline{\theta_{i-1}^{\mathrm{TV}}}, \bar{\Omega}_{i-1})$
  for $l = 1 : L$ **do**
    ▷ Gibbs sampling part
    Draw $\theta_i^{\mathrm{TV}(l)}$ from $p(\theta^{\mathrm{TV}}|\theta_i^{(l-1)}, \Omega_i^{(l-1)})$ ▷ Eq. 10
    Draw $\Omega_i^{(l)}$ from $p(\Omega|\theta_i^{(l-1)}, \theta^{\mathrm{TV}(l)})$ ▷ Eq. 11
    ▷ Metropolis-Hastings part
    Draw proposal $\theta_i^{*(l)}$ from $\{\theta_{in_i}^{(m)}\}_{m=1}^M$ according to $\{w_{in_i}^{(m)}\}_{m=1}^M$ and add rejuvenation
    Accept proposal with probability

$$\alpha = \frac{p(\theta_i^{*(l)}|\theta_i^{\mathrm{TV}(l)}, \Omega_i^{(l)}) / p(\theta_i^{*}|\overline{\theta_{i-1}^{\mathrm{TV}}}, \bar{\Omega}_{i-1})}{p(\theta_i^{(l-1)}|\theta_i^{\mathrm{TV}(l)}, \Omega_i^{(l)}) / p(\theta_i^{(l-1)}|\overline{\theta_{i-1}^{\mathrm{TV}}}, \bar{\Omega}_{i-1})}, \quad (14)$$

  end for
  Parametric approximations of posterior (hyper prior for next individual):
  $p(\theta^{\mathrm{TV}}|y_{1:i}) \approx \mathcal{N}(\overline{\theta_i^{\mathrm{TV}}}, \mathbf{S}_i^{\theta^{\mathrm{TV}}})$ with

$$\overline{\theta_i^{\mathrm{TV}}} = \frac{1}{L}\sum_{l=1}^{L} \theta_i^{\mathrm{TV}(l)}, \quad \mathbf{S}_i^{\theta^{\mathrm{TV}}} = \frac{1}{L-1}\sum_{l=1}^{L} (\theta_i^{\mathrm{TV}(l)} - \overline{\theta_i^{\mathrm{TV}}})(\theta_i^{\mathrm{TV}(l)} - \overline{\theta_i^{\mathrm{TV}}})^T \quad (15)$$

  $p(\Omega|y_{1:i}) \approx \mathcal{IW}((\nu_i - n_\Omega - 1)\bar{\Omega}_i, \nu_i)$ with

$$\bar{\Omega}_i = \frac{1}{L}\sum_{l=1}^{L} \Omega_i^{(l)}, \quad \nu_i = \nu_{i-1} + 1 \quad (16)$$

end for

TABLE 1 Hyper priors for the gold-standard model used in the simulation study

Parameter	Distribution	Hyperparameters
TV parameters		
log(MTT)	$\mathcal{N}$	$\overline{\theta^{TV}}_0 = \log(2.6), \mathbf{S}_0^{\theta^{TV}} = 0.0013$
log(Slope)	$\mathcal{N}$	$\overline{\theta^{TV}}_0 = \log(141), \mathbf{S}_0^{\theta^{TV}} = 0.016$
IIV parameters		
ω^2_{MTT}	$\mathcal{IW}$	$\Psi_{MTT} = 0.6561 = (12-2-1)0.0729,$ $\nu_0 = 12$
ω^2_{Slope}	$\mathcal{IW}$	$\Psi_{Slope} = 1.8144 = (12-2-1)0.2016,$ $\nu_0 = 12$

Abbreviations: IIV, interindividual variability; MTT, mean transit time.

The continued learning approach was sequentially applied to $N_{TDM} = 100$ virtual patients with available TDM data over six treatment cycles depending on the considered sampling scheme. For the analysis, the continuous learning approach was repeated 10 times to account for statistical variability in the individual patient parameters considered for the update. To demonstrate the effect of a mismatch between model and data-generating process, we also applied MIPD alone without continued learning (DA-guided dosing). On the individual level, model parameters MTT, Slope, and ANC_0 were estimated. We restricted the population updates to "MTT" and "Slope" as for "ANC_0" the baseline method B2 described in ref. [49] was used (i.e., no typical parameter was estimated but the baseline value was used to initialize the [empirical Bayes] prior). In addition, we consider a setting which includes in the individual level inference as the value differs across the models. In this study, we neither estimated the residual variability σ on the individual level nor on the population level. The values for σ used to generate the TDM data, however, differed from those that the models in the "imperfect model scenarios" assume. The considered hyper priors (i.e., the distributional assumptions for the population parameters), are summarized in Table 1. Because no RSEs are available for the gold-standard model,[26] the values reported in ref. [25] (one of the two pooled studies) were chosen as conservative choice. The degrees of freedom ν was chosen here to balance confidence in the estimated value while still enabling adaptation. The simulation study was performed in MATLAB 2019b and the code is available under https://doi.org/10.5281/zenodo.5504309.

RESULTS

Current MIPD approaches may not be beneficial in the presence of model bias

For a performance analysis, we generated TDM data (including residual variability) on day 1 and day 15 of each cycle (sparse sampling as in the CEPAC-TDM study). Figure 3 illustrates the performance of MIPD with/without model bias in comparison to standard dosing (median and 90% confidence intervals [CIs]). The aim of individual dose adaptations is to ensure that in each cycle, the nadir concentration stays within the two horizontal lines (i.e., representing the target range of grades 1–3 neutropenia).

The left column illustrates the scenario of parameter deviation (i.e., the structural model and the class of prior distributions is identical to the data-generating process, but hyperparameter values differ). In this case, MIPD performed comparably to standard dosing (top left, median trajectory, and 90% CI), also in terms of occurrence of grades 4 and 0 neutropenia (bottom left). For reference, in the corresponding model scenario without mismatch, the MIPD approach clearly reduced the occurrence of grades 4 and 0 (bottom and middle panel). It is worth mentioning that the CIs in all panels showed a certain "skewness" toward higher neutrophil concentrations (lower neutropenia grade), because grade 4 is penalized more strongly than grade 0 in Equation 7.

The right column illustrates the more challenging scenario of structural changes (i.e., a model structure differing from the data-generating process). Of note, in this case, both standard dosing and MIPD performed much worse than in the parameter bias scenario (bottom panels). In three out of six cycles (cycles 2–4), MIPD resulted in even larger occurrences of grade 4 compared to standard dosing. The gold-standard model underestimated the drug effect on neutrophil concentrations (see Figure 1) and too high doses were selected, especially in the presence of cumulative neutropenia. Despite relying on an inappropriate structural model, DA was able to correct this initial mismatch on the parameter level over the course of a patient's therapy by integrating TDM data, which lead to a decrease in incidence of grade 4 neutropenia in later cycles. For reference, in the corresponding unbiased model scenario, the MIPD approach clearly and very quickly reduced the occurrence of grades 4 and 0 (bottom and middle panel). In comparison to the parameter adaptation scenario, the occurrence of grades 4 and 0 neutropenia was even further decreased, which might be related to the smaller RUV parameter, see Figure 1a.

In summary, if the underlying model is not consistent with the observational data, MIPD might not be beneficial compared to standard dosing that solely relies on TDM data ("model-free"). As outlined in the introduction, necessary model adaptations can be expected if MIPD is applied in clinical routine, and, therefore, the top panels might better reflect clinical reality than the middle panels. Here, model adaptation during clinical practice is necessary. Most MIPD approaches, however, do not exploit the wealth of TDM data used during MIPD to learn and update their models.

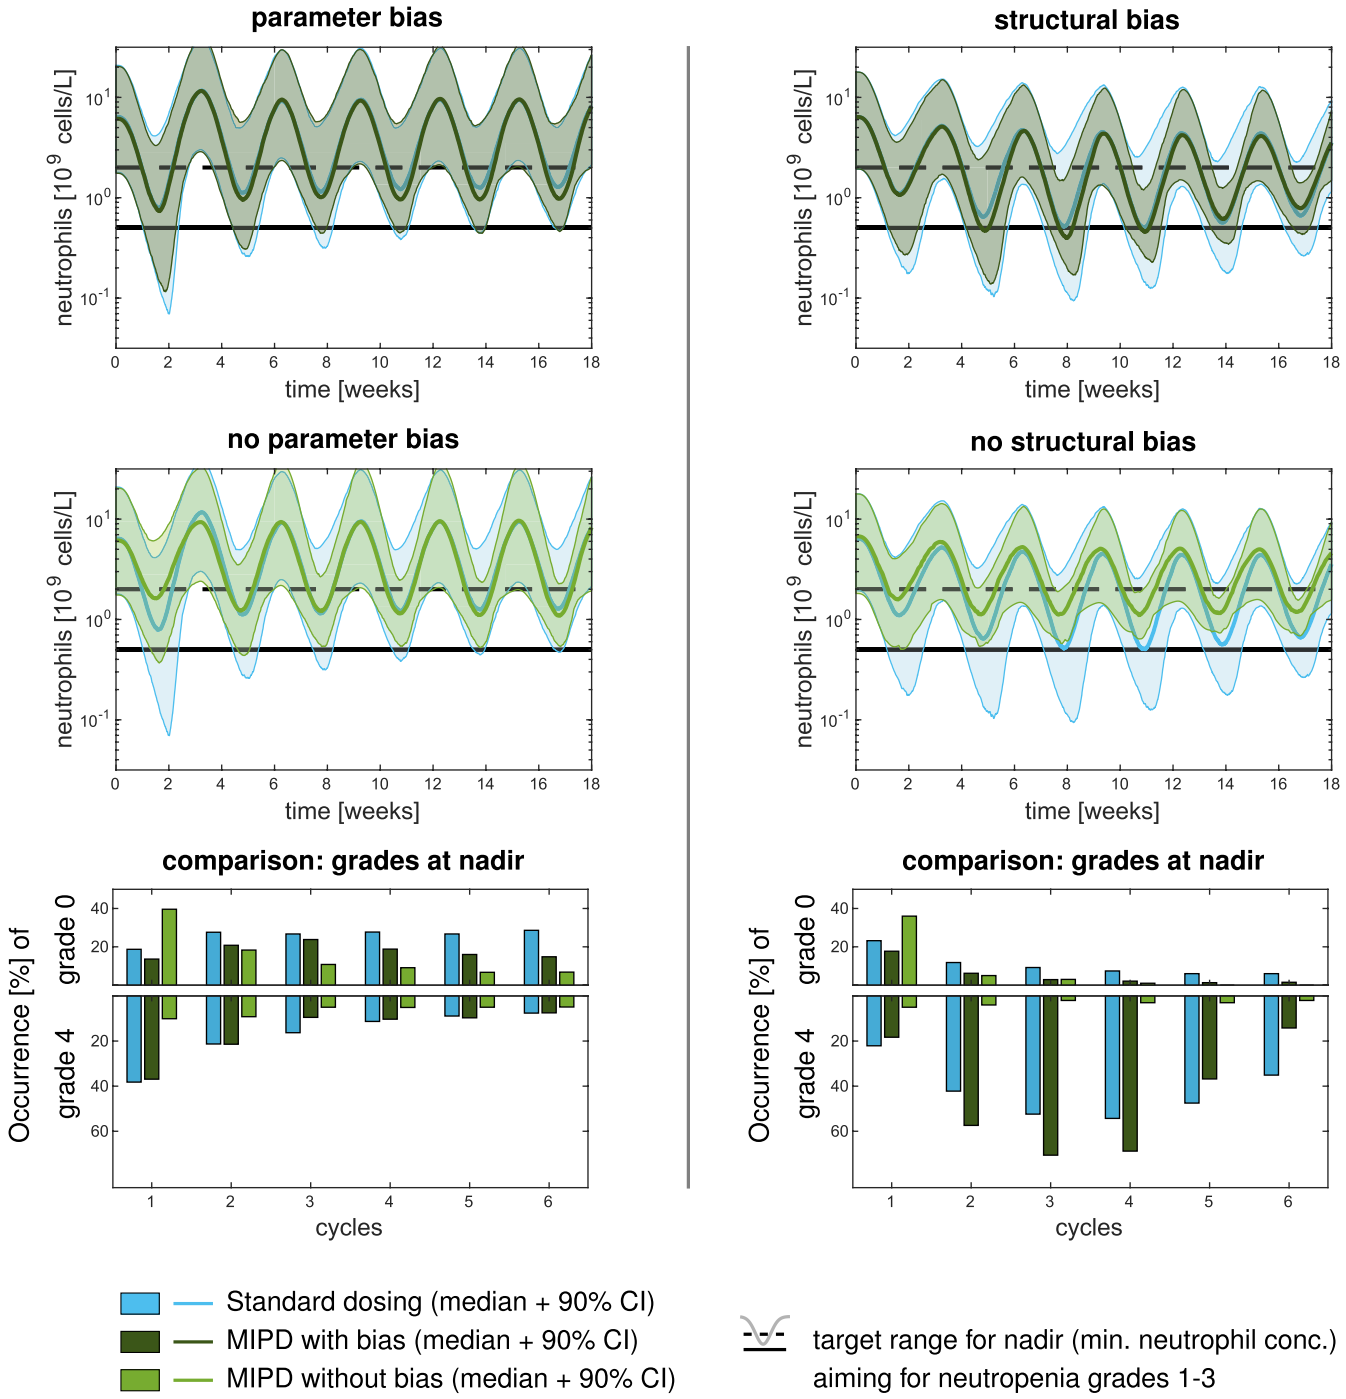

FIGURE 3 MIPD in presence of a mismatch between model and data-generating process. Neutropenia time-courses were simulated for $N_{TDM} = 1000$ virtual patients using MIPD approaches with different underlying models and the standard dosing approach for paclitaxel (200 mg/m^2 with 20% reduction if grade 4 neutropenia is observed in the previous cycle). TDM data were simulated for day 1 and day 15 of each cycle (sparse sampling scenario), including residual variability. Left column: TDM data were generated using the gold-standard R model. Right column: TDM data was generated using the bone marrow exhaustion (BME) model.[17] Top panel: "Imperfect model scenario" (i.e., the MIPD approach uses the gold-standard model), while the data was generated with the gold-standard R model (left) or BME model (right). Middle panel: "Perfect model scenario" (i.e., the MIPD approach uses also the data-generating model). The median time-course is shown along with its 90% confidence interval (CI). Bottom panel: occurrence of life-threatening grade 4 neutropenia and subtherapeutic grade 0 neutropenia for the different scenarios. Note that grade 4 neutropenia is penalized more ($\lambda_1 = 2/3$) compared to grade 0 neutropenia ($\lambda_2 = 1/3$) in Equation (7). This is accounted for in the scale of the bottom panels, which allows to interpret the length of the total bars. MIPD, model-informed precision dosing; TDM, therapeutic drug monitoring

Continued learning MIPD can adapt parameters—depending on the sampling scheme

The proposed continued learning framework was able to adapt the prior parameter distribution across TDM patients. Figure 4 illustrates the sequential updates of the proposed framework for the posterior distributions of the typical parameters of "Slope" and "MTT" across 100 patients for different sampling schemes. For the rich sampling scenario (left), the posterior—95% highest posterior density (HPD) area—evolved over the number of observed patients, moving away from the prior estimate (grey star) toward the value used to generate the data (black star). As more patients were observed, uncertainty about the typical "Slope" and "MTT" parameters decreased, as indicated by the decreasing size of the HPD area. Thus, the proposed framework successfully learned the TVs underlying the TDM data from sample representations of the posterior on the individual level. Note that the parameters γ and σ were not estimated although different values were used to generate the data, which has the effect of introducing an additional bias. The results including γ on the individual level inference are shown in Supplementary Material Figure S1.

The extent to which the continuous learning framework could counteract a parameter mismatch depended on the sampling scheme (see Figure 4 middle and right panel). For the intermediate scheme, the posterior distribution moved toward the parameter values used to generate the data. A final parameter bias, however, remained. A potential reason could be parameter identifiability. To assess practical identifiability, we investigated the log-likelihood and log-posterior on the individual patient levels, see Supplementary Material Figure S2. To exclude the possibility of unfavorably chosen sampling time points in the intermediate scheme (weekly), we performed an optimal design analysis, see Supplementary Material Section S3. These analyses showed that the setting is not ideally chosen, but also that optimal design considerations might be patient-specific. For the sparse sampling scheme, the TDM data were not sufficient to adapt the model appropriately. Yet, in the context of the rich sampling, the data was indeed informative enough to move away from the (biased) prior estimate toward the data-generating value, resolving the practical unidentifiability.

Continued learning in MIPD can substantially improve therapy outcome even for structural changes

Finally, we investigated the effects of continued learning of population parameters on MIPD. Here, we show only the more challenging structural bias scenario (a different model used for data generation vs. inference in MIPD); for

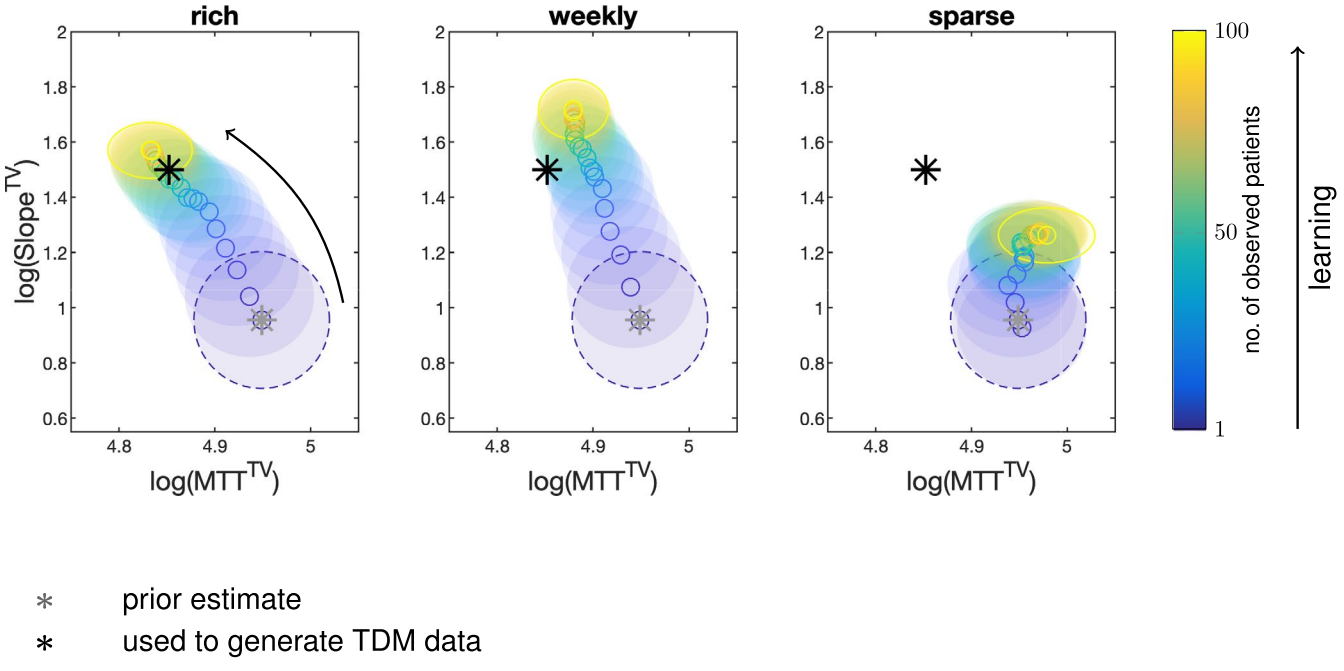

FIGURE 4 Comparison of the sequential updates of the hyper prior for the typical MTT and Slope value for different TDM scenarios. Grey star: prior estimate of the hyperparameters; black star: true hyperparameters (i.e., the values used to generate the TDM data). Sparse sampling consisted of measurements on days 1 and 15 of each cycle, weekly sampling corresponds to an intermediate data situation, and for rich sampling it is assumed that neutrophils are monitored every third day. Mean (circle) and 95% highest posterior density (shaded ellipse) are shown. MTT, mean transit time; TDM, therapeutic drug monitoring

the parameter bias scenario, see Supplementary Material Figure S5. The performance of the different approaches is compared in Figure 5. Standard dosing and DA-guided dosing are set up analogously in Figure 3, except for the fact they are trained with the intermediate sampling design. It becomes clear that continued learning is significantly more effective with more TDM data (see above). We also considered uncertainty with respect to the parameter γ.

Data assimilation-guided dosing was also able to adjust to some extent to cumulative neutropenia over time (see Figure 5 dark green). The "Slope" parameter increased, whereas parameters "$Circ_0$," and γ decreased over the course of the individual therapy, leading to a decrease in occurrence of grade 4 after cycle 3 and a substantial decrease in outcome variability. Effectively, when considering the data points one at a time, the sequential DA framework allowed to account for changes in the parameters over time—a potentially very beneficial property (e.g., in disease progression). Although this might be very desirable for MIPD at the individual patient level, it could be misleading when learning across patients. When the final parameter estimate (after 6 cycles) was used to update the population parameters ($Slope^{TV}$ and MTT^{TV}), this introduced a bias for the first cycle of the next patient, resulting in high occurrence of grade 0 for the first cycle (Figure 5 bottom left). Continued learning was considered across the first 100 patients (blue-green) as well as for some

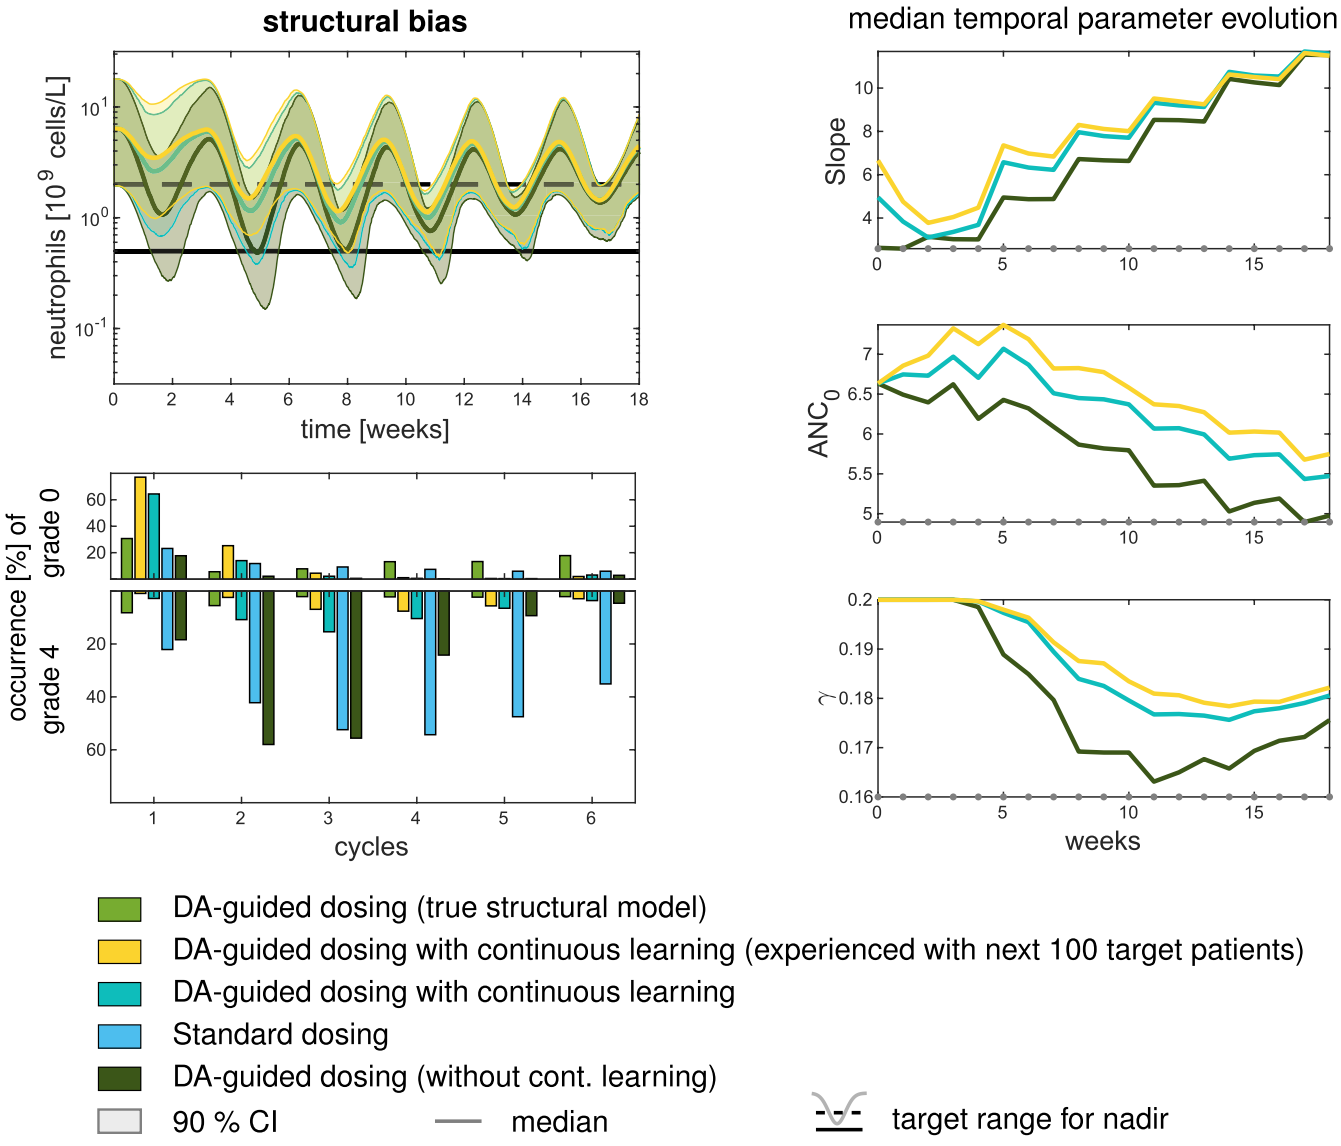

FIGURE 5 Sequential DA allows for temporal parameter changes within the course of a patient's therapy. The TDM data were generated using the BME model and the results are shown for the intermediate sampling scheme (sampling timepoints are indicated as grey dots in the right panels). The median temporal parameter evolution over the course of therapy was computed across all virtual patients ($N_{TDM} = 10 \cdot 100$). BME, bone marrow exhaustion; CI, confidence interval; DA, data assimilation; TDM, therapeutic drug monitoring

second 100 patients (yellow) after learning from the first 100 patients. It can be observed that the typical "Slope" parameter increased (green vs. blue-green vs. yellow) as it was continuously learned across patients (initial Slope value at $t = 0$ in the top right panel).

A major improvement was observed for the continued learning MIPD approach, which reduced the occurrence of grade 4 substantially across all cycles compared to DA-guided dosing or standard dosing. The results for the rich sampling scenario were comparable (Supplementary Material Figure S7); for the sparse sampling scheme, however, the benefits were not so clear (Supplementary Material Figure S6).

DISCUSSION

We proposed a sequential hierarchical Bayesian approach to update the population parameters using posterior samples as a means to exchange information with every treated patient so that the model better reflects the target patient population. We showed that the approach allowed to successfully learn the underlying population parameters of the PK/PD model used to generate the patient data. It is important to note, however, that the results depend on the sampling scheme. In addition, we showed that continued learning has potential to improve MIPD even in the presence of structural changes, again depending on the informativeness of available TDM data.

The proposed approach has two levels and allows to learn sequentially over patients without using patient data on the population level. Thus, the patient data themselves do not need to be stored or shared across centers, which is a crucial advantage compared to pooling approaches.[35] The proposed Bayesian approach is sequential in nature. We propose to update the population hyperparameters with every new patient. The advantage of a continued learning approach is that the prior hyperparameters are always up-to-date, reflecting the current knowledge, and that no decision has to be taken as to when to perform a global (pooled) update. The computational time on the individual level is not impacted by the population level update. The individual level computation depends on the number of particles used in the particle filter, which could also be run in parallel. For more details, see ref. [20]. The population level update is less time critical and can be done in the background.

Thus, the approach builds a basis to develop more informed models, integrating an ever-growing database potentially better reflecting rare covariates. The initial model used to start the continued learning approach could be selected using a retrospective external evaluation based on historical data from the intended clinical setting.[6] Model selection/model averaging approaches do not adapt/improve the underlying model across patients; the a priori forecast remains the same for all patients (based on the covariates). In addition, in their general form, these approaches are implemented in conjunction with MAP estimation, which provides potentially biased predictions in context of nonlinear models.[20] The proposed DA-guided dosing also naturally extents to model averaging and this extension has been considered (on the individual inference level) previously in the context of Bayesian therapy forecasting.[50]

In the context of cytotoxic chemotherapy with neutropenia as dose-limiting toxicity, we show how population-based PK/PD models can be transferred to a different clinical target population, which is often a crucial application hurdle of MIPD in clinical routine. We showed that model misspecification might severely impact MIPD, and, therefore, models should be adapted to the target patient population. The used DA-guided dosing proved to be able to adapt the model to some extent, but only improved MIPD at later cycles, when a certain amount of TDM data was collected. This is a consequence of assimilating data sequentially, thereby allowing to account for temporal changes in the parameters and thus adapting the gold-standard model to some extent to cumulative neutropenia.

The presented analysis revealed that an important aspect for practical implementation is to critically assess the quality of the inference on the individual level.

For rich sampling, the continuous learning approach showed very promising performance. Even in the most challenging case of structural bias, this approach clearly outperformed standard dosing and MIPD without continuous learning. Under sparse sampling with structural bias, inference of hyperparameters proved more difficult; further investigations are required in such a setting, as was demonstrated for time-dependent parameters. With the prospect of novel digital health care devices (e.g., point-of-care devices), more frequent monitoring could become clinical reality. Real-world data is currently underutilized[51] but has great potential to improve MIPD, as shown in this study. The current approach is limited to misspecifications or population shifts on the structural model parameter level. A model change on the structural level (e.g., accounting for cumulative neutropenia), could be corrected for to some extent on the level of parameters. An important extension in the future would be to also estimate the RUV parameter σ, as an increased error in measurement precision or reporting can be expected in clinical routine compared to, for example, controlled clinical study settings. Currently, the IIV parameters ω^2 captured the increased RUV of the data to some extent, which, however, also increased the uncertainty on the individual level, see Supplementary Material Figure S4.

The approach of a learning model (as coined in ref. [1]) for MIPD could be beneficial not only in clinical

practice, but also during drug development, where new (clinical) study data are generated continuously and should be integrated into previously developed models.[52] The study is an important step toward building the underlying models of MIPD on a growing database and thus make MIPD fit-for-purpose in everyday therapeutic use.

ACKNOWLEDGEMENTS

Fruitful discussions with Sven Mensing (AbbVie, Germany), Alexandra Carpentier (Otto-von-Guericke-Universitaet Magdeburg), Sebastian Reich (University of Potsdam, University of Reading), and David Albers (University of Colorado) are kindly acknowledged. Open Access funding enabled and organized by Projekt DEAL.

CONFLICT OF INTEREST

C.K. and W.H. report research grants from an industry consortium (AbbVie Deutschland GmbH & Co. KG, AstraZeneca, Boehringer Ingelheim Pharma GmbH & Co. KG, Grünenthal GmbH, F. Hoffmann-La Roche Ltd., Merck KGaA and Sanofi) for the PharMetrX PhD program. In addition, C.K. reports research grants from the Innovative Medicines Initiative-Joint Undertaking (DDMoRe), the European Commission within the Horizon 2020 framework programme (FAIR), and Diurnal Ltd. All other authors declared no competing interests for this work.

AUTHOR CONTRIBUTIONS

C.M., N.H., J.d.W., C.K., and W.H. wrote the manuscript. C.M., N.H., J.d.W., C.K., and W.H. designed the research. C.M. performed the research. C.M., N.H., J.d.W., and W.H. analyzed the data.

REFERENCES

1. Keizer RJ, ter Heine R, Frymoyer A, Lesko LJ, Mangat R, Goswami S. Model-informed precision dosing at the bedside: scientific challenges and opportunities. *CPT Pharmacometrics Syst Pharmacol*. 2018;7:785-787.
2. Peck RW. Precision dosing: an industry perspective. *Clin Pharmacol Ther*. 2021;109:47-50.
3. Kluwe F, Michelet R, Mueller-Schoell A, et al. Perspectives on model-informed precision dosing in the digital health era: challenges, opportunities, and recommendations. *Clin Pharmacol Ther*. 2021;109:29-36.
4. Lavielle M. *Mixed effects models for the population approach*. Chapman and Hall/CRC; 2014.
5. Polasek TM, Shakib S, Rostami-Hodjegan A. Precision dosing in clinical medicine: present and future. *Expert Rev Clin Pharmacol*. 2018;11:743-746.
6. Zhao W, Kaguelidou F, Biran V, et al. External evaluation of population pharmacokinetic models of vancomycin in neonates: the transferability of published models to different clinical settings. *Br J Clin Pharmacol*. 2013;75:1068-1080.
7. Heine R, Keizer RJ, Steeg K, et al. Prospective validation of a model? Informed precision dosing tool for vancomycin in intensive care patients. *Br J Clin Pharmacol*. 2020;86:2497-2506.
8. Sheiner LB, Ludden TM. Population pharmacokinetics/dynamics. *Annu Rev Pharmacol Toxicol*. 1992;32:185-209.
9. Deitchman AN. The risk of treating populations instead of patients. *CPT Pharmacometrics Syst Pharmacol*. 2019;8:256-258.
10. Powell JR, Cook J, Wang Y, Peck R, Weiner D. Drug dosing recommendations for all patients: a roadmap for change. *Clin Pharmacol Ther*. 2021;109(1):65-72.
11. Hamberg AK, Dahl M-L, Barban M, et al. A PK-PD model for predicting the impact of age, CYP2C9, and VKORC1 genotype on individualization of warfarin therapy. *Clin Pharmacol Ther*. 2007;81:529-538.
12. Ohara M, Takahashi H, Lee MTM, et al. Determinants of the over-anticoagulation response during warfarin initiation therapy in Asian patients based on population pharmacokinetic-pharmacodynamic analyses. *PLoS One*. 2014;9:e105891.
13. Uster DW, Stocker SL, Carland JE, 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(1):175-183.
14. Mao JJ, Jiao Z, Yun H-Y, et al. External evaluation of population pharmacokinetic models for ciclosporin in adult renal transplant recipients. *Br J Clin Pharmacol*. 2018;84:153-171.
15. Jodrell DI, Reyno LM, Sridhara R, et al. Suramin: development of a population pharmacokinetic model and its use with intermittent short infusions to control plasma drug concentration in patients with prostate cancer. *J Clin Oncol*. 1994;12:166-175.
16. Conley BA, Forrest A, Egorin MJ, Zuhowski EG, Sinibaldi V, Van Echo DA. Phase I trial using adaptive control dosing of hexamethylene Bisacetamide (NSC 95580). *Cancer Res*. 1989;49:3436-3440.
17. Henrich A, Joerger Markus, Kraff Stefanie, et al. Semimechanistic bone marrow exhaustion pharmacokinetic/pharmacodynamic model for chemotherapy-induced cumulative neutropenia. *J Pharmacol Exp Ther*. 2017;362:347-358.
18. Wallin JE, Friberg LE, Karlsson MO. A tool for neutrophil guided dose adaptation in chemotherapy. *Comput Methods Programs Biomed*. 2009;93:283-291.
19. Netterberg I, Nielsen EI, Friberg LE, Karlsson MO. Model-based prediction of myelosuppression and recovery based on frequent neutrophil monitoring. *Cancer Chemother Pharmacol*. 2017;80:343-353.
20. Maier C, Hartung N, Wiljes J, Kloft C, Huisinga W. Bayesian data assimilation to support informed decision making in individualized chemotherapy. *CPT Pharmacometrics Syst Pharmacol*. 2020;9:153-164.
21. 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;10:241-254.
22. Friberg LE, Henningsson A, Maas H, Nguyen L, Karlsson MO. Model of chemotherapy-induced myelosuppression with parameter consistency across drugs. *J Clin Oncol*. 2002;20:4713-4721.
23. Kloft C, Wallin J, Henningsson A, Chatelut E, Karlsson MO. Population pharmacokinetic-pharmacodynamic model for neutropenia with patient subgroup identification: comparison across anticancer drugs. *Clin Cancer Res*. 2006;12:5481-5490.

24. Hansson EK, Wallin JE, Lindman H, Sandström M, Karlsson MO, Friberg LE. Limited inter-occasion variability in relation to inter-individual variability in chemotherapy-induced myelosuppression. *Cancer Chemother Pharmacol*. 2010;65:839-848.
25. Joerger M, Huitema ADR, Richel DJ, et al. Population pharmacokinetics and pharmacodynamics of paclitaxel and carboplatin in ovarian cancer patients: a study by the European Organization for research and treatment of cancer-pharmacology and molecular mechanisms group and new drug development group. *Clin Cancer Res*. 2007;13:6410-6418.
26. Joerger M, Kraff S, Huitema ADR, et al. Evaluation of a pharmacology-driven dosing algorithm of 3-weekly paclitaxel using therapeutic drug monitoring. *Clin Pharmacokinet*. 2012;51:607-617.
27. Kaefer A, Yang J, Noertersheuser P, et al. Mechanism-based pharmacokinetic/pharmacodynamic meta-analysis of navitoclax (ABT-263) induced thrombocytopenia. *Cancer Chemother Pharmacol*. 2014;74:593-602.
28. Pujo-Menjouet L. Blood cell dynamics: half of a century of modelling. *Math Model Nat Phenom*. 2016;11:92-115.
29. Craig M. Towards quantitative systems pharmacology models of chemotherapy-induced neutropenia. *CPT Pharmacometrics Syst Pharmacol*. 2017;6:293-304.
30. Chen Z, Ma N, Liu B Lifelong learning for sentiment classification. In Proc. 53rd Annual Meeting of the Association for Computional Linguistics. 7th International Joint Conference of the National Language Processing (Volume 2 Short Pap., 750–756, Association for Computational Linguistics; 2015. doi:10.3115/v1/P15-2123.
31. Silver DL, Yang Q, Li L. Lifelong machine learning systems: beyond learning algorithms. In AAAI Spring Symposium - Tech. Rep., vol. SS-13-05, 49–55; 2013.
32. Pan SJ, Yang Q. A survey on transfer learning. *IEEE Trans Knowl Data Eng*. 2010;22:1345-1359.
33. Torrey L, Shavlik J. Transfer learning. In: Olivas ES, Guerrero JDM, Martinez-Sober M, Magdalena-Benedito JR, Serrano López AJ, editors, *Handbook Research Machine Learning Application*, chap. 11. IGI Global; 2010:242–264.
34. Jiang J. A Literature Survey on Domain Adaptation of Statistical Classifiers (2008). https://www.mysmu.edu/faculty/jingjiang/papers/da_survey.pdf.
35. 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(1):233-242.
36. Lunn D, Barrett J, Sweeting M, Thompson S. Fully Bayesian hierarchical modelling in two stages, with application to meta-analysis. *J R Stat Soc Ser C (Appl Stat)*. 2013;62(4):551-572.
37. Hooten MB, Johnson DS, Brost BM. Making recursive Bayesian inference accessible. *Am Stat*. 2021;75(2):185-194.
38. Di Maio M, Gridelli C, Gallo C, et al. Chemotherapy-induced neutropenia and treatment efficacy in advanced non-small-cell lung cancer: a pooled analysis of three randomised trials. *Lancet Oncol*. 2005;6:669-677.
39. Di Maio M, Gridelli C, Gallo C, Perrone F. Chemotherapy-induced neutropenia: a useful predictor of treatment efficacy? *Nat Clin Pract Oncol*. 2006;3:114-115.
40. Joerger M, von Pawel J, Kraff S, et al. Open-label, randomized study of individualized, pharmacokinetically (PK)-guided dosing of paclitaxel combined with carboplatin or cisplatin in patients with advanced non-small-cell lung cancer (NSCLC). *Ann Oncol*. 2016;27:1895-1902.
41. Joerger M. Quantitative effect of gender, age, liver function, and body size on the population pharmacokinetics of paclitaxel in patients with solid tumors. *Clin Cancer Res*. 2006;12:2150-2157.
42. Huizing MT, Giaccone G, van Warmerdam LJ, et al. Pharmacokinetics of paclitaxel and carboplatin in a dose-escalating and dose-sequencing study in patients with non-small-cell lung cancer. The European Cancer Centre. *J Clin Oncol*. 1997;15:317-329.
43. Henrich A. *Pharmacometric modelling and simulation to optimise paclitaxel combination therapy based on pharmacokinetics, cumulative neutropenia and efficacy*. Ph.D. thesis, Freie Universität Berlin; 2017.
44. Dunwoodie EH. *Home testing of blood counts in patients with cancer*. PhD, The University of Leeds; 2018.
45. Duffull SB, Friberg LE, Dansirikul C. Bayesian hierarchical modeling with Markov Chain Monte Carlo methods. In: Ette EI, Williams PJ, eds. *Pharmacometrics*. John Wiley & Sons, Inc.; 2007:137-164.
46. Wakefield J. Bayesian individualization via sampling-based methods. *J Pharmacokinet Biopharm*. 1996;24:103-131.
47. Gisleskog PO, Karlsson MO, Beal SL. Use of prior information to stabilize a population data analysis. *J Pharmacokinet Pharmacodyn*. 2002;29:473-505.
48. Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A, Rubin DB. *Bayesian data analysis*, 3rd ed. Chapman and Hall/CRC; 2014.
49. Dansirikul C, Silber HE, Karlsson MO. Approaches to handling pharmacodynamic baseline responses. *J Pharmacokinet Pharmacodyn*. 2008;35:269-283.
50. Albers DJ, Levine M, Gluckman B, Ginsberg H, Hripcsak G, Mamykina L. Personalized glucose forecasting for type 2 diabetes using data assimilation. *PLoS Comput Biol*. 2017;13:e1005232.
51. Tyson RJ, Park CC, Powell JR, et al. Precision dosing priority criteria: drug, disease, and patient population variables. *Front Pharmacol*. 2020;11:1-18.
52. Lalonde RL, Kowalski KG, Hutmacher MM, et al. Model-based drug development. *Clin Pharmacol Ther*. 2007;82:21-32.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Maier C, de Wiljes J, Hartung N, Kloft C, Huisinga W. A continued learning approach for model-informed precision dosing: Updating models in clinical practice. *CPT Pharmacometrics Syst Pharmacol*. 2022;11:185-198. doi:10.1002/psp4.12745