



ARTICLE

Application of SHAP values for inferring the optimal functional form of covariates in pharmacokinetic modeling

Alexander Janssen¹ | Mark Hoogendoorn² | Marjon H. Cnossen³ |
Ron A. A. Mathôt¹ | for the OPTI-CLOT Study Group and SYMPHONY Consortium[†]

¹Department of Clinical Pharmacology, Hospital Pharmacy, Amsterdam University Medical Center, Amsterdam, The Netherlands

²Quantitative Data Analytics Group, Department of Computer Science, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

³Department of Pediatric Hematology, Erasmus MC Sophia Children's Hospital, Erasmus University Medical Center, Rotterdam, The Netherlands

Correspondence

Alexander Janssen and Ron A. A. Mathôt, Amsterdam UMC, locatie AMC, Apotheek, Meibergdreef 9, Postbus 22660, 1100 DD Amsterdam Zuidooost, The Netherlands.
Email: a.janssen@amsterdamumc.nl and r.mathot@amsterdamumc.nl

Funding information

This research received funding from the Dutch Organization for Scientific Research (NWO) in the framework of the NWA-ORC Call grant agreement NWA.1160.18.038

Abstract

In population pharmacokinetic (PK) models, interindividual variability is explained by implementation of covariates in the model. The widely used forward stepwise selection method is sensitive to bias, which may lead to an incorrect inclusion of covariates. Alternatives, such as the full fixed effects model, reduce this bias but are dependent on the chosen implementation of each covariate. As the correct functional forms are unknown, this may still lead to an inaccurate selection of covariates. Machine learning (ML) techniques can potentially be used to learn the optimal functional forms for implementing covariates directly from data. A recent study suggested that using ML resulted in an improved selection of influential covariates. However, how do we select the appropriate functional form for including these covariates? In this work, we use SHapley Additive exPlanations (SHAP) to infer the relationship between covariates and PK parameters from ML models. As a case-study, we use data from 119 patients with hemophilia A receiving clotting factor VIII concentrate peri-operatively. We fit both a random forest and a XGBoost model to predict empirical Bayes estimated clearance and central volume from a base nonlinear mixed effects model. Next, we show that SHAP reveals covariate relationships which match previous findings. In addition, we can reveal subtle effects arising from combinations of covariates difficult to obtain using other methods of covariate analysis. We conclude that the proposed method can be used to extend ML-based covariate selection, and holds potential as a complete full model alternative to classical covariate analyses.

Study Highlights

WHAT IS THE CURRENT KNOWLEDGE ON THE TOPIC?

Covariate selection in pharmacokinetic (PK) modeling is a complex process and is sensitive to bias. Machine-learning (ML) techniques might help to simplify and potentially improve this process, but are difficult to interpret as is.

[†]OPTI-CLOT Study Group and SYMPHONY Consortium author names are included in the Acknowledgment section.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *CPT: Pharmacometrics & Systems Pharmacology* published by Wiley Periodicals LLC on behalf of American Society for Clinical Pharmacology and Therapeutics.

WHAT QUESTION DID THIS STUDY ADDRESS?

Can we use ML models to infer the optimal functional form to represent the relationship between covariates and PK parameters using SHapley Additive exPlanations?

WHAT DOES THIS STUDY ADD TO OUR KNOWLEDGE?

This study presents an extension to covariate selection procedures using ML methods. The resulting framework allows for the detection of intricate effects of covariates, which far exceed the capabilities of classical (linear) covariate analyses. In addition, it is more flexible with respect to covariate importance scores generally used in ML-based covariate selection.

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

By learning the optimal functional form of covariates based on data the complexity of covariate selection and implementation is reduced. This accelerates PK model development and might help improve the accuracy of PK parameter estimates.

INTRODUCTION

In population pharmacokinetic (PK) modeling, identification of the relationship between PK parameters and covariates is important for the explanation of interindividual variation (IIV). The classic stepwise method is among the most popular methods but is not without flaws. In stepwise methods, covariate selection is determined by a significant change in the objective function value following inclusion or exclusion of each covariate. Due to the ordered nature of this process, the method may suffer from bias and multiplicity issues.¹⁻³

The full fixed effects model (FFEM), which is based on a full model fit, was introduced to reduce selection bias.⁴ In this method, all covariates of interest are tested simultaneously and included if they result in a clinically relevant change of the typical PK parameters. Although an improvement over the stepwise method, the FFEM is not able to solve all prior issues. In both methods, an assumption must be made about the functional form describing the relationship between the covariate and the PK parameters. This encourages data dredging because various functional forms can be tested until one satisfies the criteria for inclusion. Furthermore, true covariates may be excluded when suboptimal functional forms are used. In summary, we identify a need for a covariate selection method which performs both a full model fit, while simultaneously estimating the optimal functional form of each covariate.

A recent study describes the use of machine learning (ML) for performing covariate selection for PK models.⁵ Here, the authors discuss how combining ML algorithms with covariate importance scores can be used to obtain a similar or better selection of covariates compared to stepwise methods. Other studies further discuss using such an

approach on real-life data to obtain a set of predictive covariates.^{6,7} ML algorithms might be suitable for this task as they can learn covariate relationships directly from data. These methods might thus reduce the issue of selecting suboptimal functional forms when testing covariates for inclusion. Many ML software packages provide measures of covariate importance. For tree-based methods (e.g., random forests⁸ or gradient boosting trees⁹), examples include counting the number of uses of each covariate, or more sophisticated measures, such as Gini or permutation importance. Although often found to be relatively accurate, there are situations where these measures may be biased.¹⁰ In addition, they only provide a single score of importance without information about the relationship between each covariate and model output. After obtaining a set of important covariates, how do we now select the functional form to implement these covariates without again resorting to stepwise methods?

SHapley Additive exPlanations (SHAP) is a promising model explanation technique due to its strong theoretical base.¹¹ In addition to a more robust benchmark performance compared to other approaches,¹² SHAP allows for identification of the influence of specific covariates and their effect on each individual prediction. The use of SHAP might improve upon importance scores by also allowing for the analysis of the relationship between covariates and model output. Its use for covariate selection has, however, not yet been explored.

In this study, we will focus on tree-based ML algorithms, as there exists an exact method for the computation of SHAP values for these types of models.¹² Specifically, we will use the random forest and XGBoost algorithms.¹³ Both methods create an ensemble model of decision trees. A decision tree is an algorithm that groups observations into bins

(appropriately called leaves), which share a similar value for the response variable. Each tree is composed of multiple layers, where the observation is split into two leaves based on the value of one of the covariates. In a random forest, the model prediction is averaged over multiple independently fit trees. Each tree is fit using a subset of the data adding stochasticity to the learning process aiming to reduce overfitting. In gradient boosting trees (e.g., XGBoost), the trees are built sequentially, so that additional decision trees are added if they improve the prediction of the previous model ensemble. Each tree is thus fit to improve the mistakes of the previous tree. The objective function also contains a regularization term which penalizes the addition of complex models. In contrast to the classic random forest implementation, XGBoost supports missing values.¹³

Our goal is to evaluate the value of combining ML and SHAP for enriching ML-based covariate analysis in the context of PK models. To this end, we will fit a random forest and XGBoost model to predict empirical Bayes estimates of PK parameters and perform a SHAP analysis on the most accurate model. As a case study, we use a retrospective dataset of patients with hemophilia A receiving clotting factor VIII (FVIII) while undergoing surgery.¹⁴ We explore the output of the SHAP analysis and present how it can be used for understanding the relationship between covariates and PK parameters.

METHODS

Dataset

We used retrospective data of 119 individuals with hemophilia A undergoing surgery in five different hemophilia treatment centers in the Netherlands.¹⁴ Patients received clotting factor FVIII concentrate (via bolus or continuous doses) to reach target FVIII levels as set by the Dutch National Hemophilia Consensus. This guideline recommends the following FVIII peak levels during the peri-operative window: 0.80–1.00 IUml⁻¹ at 0–24h, 0.50–0.80 IUml⁻¹ at 24–120h, and 0.30–0.50 IUml⁻¹ beyond 120h post-surgery. A total of 3350 FVIII levels were measured during 197 surgical procedures. All FVIII levels were measured using the one-stage clotting assay. Timing and dosage of measurements was determined at the discretion of the treating physician. For most patients, this resulted in more frequent measurements early in the peri-operative window, and occasional measurements post-surgery to validate if the patient still met target levels. The following 13 covariates were chosen for analysis: treatment center (1–5), pre-assessed surgical risk (low vs. high¹⁵), use of β -domain deleted recombinant FVIII (BDD-FVIII, moroctocog alfa/Refacto AF), hemophilia severity (moderate vs. severe), FVIII baseline levels, blood group (O

vs. non-O), blood loss during surgery, occurrence of a bleeding complication, body weight, body mass index (BMI), age in years, and von Willebrand factor antigen (VWF:Ag) and activity (VWF:act) levels. Five covariates contained missing values. Missing values were either imputed by mean (for continuous variables) or addition of a separate category (for categorical variables).

Prediction of PK parameters using machine learning

Empirical Bayes estimates of the PK parameters were obtained by fitting a base two-compartment model to the data using NONMEM (ICON Development Solutions, Ellicott City, MD¹⁶). Random effects were only estimated for the clearance and central volume parameters in order to improve model stability. A combined additive and proportional error model was used. We fixed the residual error estimates to $\sigma_1 = 0.08$ (additive error) and $\sigma_2 = 0.17$ (proportional error) to improve model stability and shrinkage while matching earlier findings.^{17,18} Random forest (Python scikit-learn package, version 0.23.2) and XGBoost (Python xgboost package, version 1.4.2) models were fit to predict the empirical Bayes estimated clearance and central volume distribution parameters independently. We fit the XGBoost models to both the original (containing missing values) and imputed data set. We performed a 10-fold cross-validation for the estimation of test set error and for SHAP value calculation. Default model hyperparameters were used (see [Material S1](#) for details). Model accuracy was represented as the average mean absolute error (MAE) \pm one SD of PK parameter predictions on the 10 test sets. We also calculated the root mean squared error (RMSE) of predicted FVIII levels by solving a two-compartment model using the test set predicted PK parameters. The empirical Bayes estimated inter-compartmental clearance and peripheral volume parameters were directly used for all patients. FVIII level predictions were performed in the Julia programming language (Julia Computing, Inc., version 1.6.0) using the DifferentialEquations package (version 6.17.1).¹⁹ The RMSE was presented as the mean and SD of the RMSE calculated for each individual patient.

SHAP analysis

A SHAP analysis (Python shap package, version 0.36.0) was performed to explain model output. This method decomposes a model $f(x)$ into a simpler additive model:

$$f(x) = \phi_0 + \sum_{i=1}^M \phi_{x_i}$$

Here, the SHAP value ϕ_{x_i} of covariate $i \in M$ represents its direct effect on the model prediction, whereas ϕ_0 represents the typical prediction. By cumulating the SHAP values for each individual, we can visualize their relationships with each of the covariates. For each of the 10 cross-validations, we calculated SHAP values on the corresponding test set. The SHAP values were pooled and a smoothed representation of the effect was obtained by means of locally estimated scatterplot smoothing (LOESS; Python statsmodels package, version 0.12.2). SHAP values for missing continuous covariates were removed from visualizations.

Model code

All model code, including implementation instructions, will be made available at <https://github.com/Janssen/pkSHAP> at the time of publication.

RESULTS

Patient characteristics and model accuracy

An overview of the patient characteristics, missing data, and the base model parameter estimates is shown in Table 1. RMSE of FVIII level predictions by the base nonlinear mixed effects (NLME) model was $0.23 \text{ IUml}^{-1} \pm 0.27$ (SD). Accuracy of the ML models is depicted in Table 2. The MAE of PK parameter predictions by both ML algorithms fit to the imputed data set was similar. The XGBoost model fit to the original dataset resulted in higher MAE of both clearance (43.8 vs. 40.4 ml/h), as well as central volume predictions (893 vs. 807 ml) compared to the random forest model. In addition, the RMSE of the resulting FVIII level predictions was higher for the XGBoost model (0.36 vs. 0.32 IUml^{-1}). The MAE of PK parameter predictions was indicative of the presence of residual IIV unexplained by the current set of covariates.

SHAP analysis

We present an overview of the SHAP values for the random forest models in Figure 1. This visualization can, for example, be used for the identification of influential covariates, as indicated by the horizontal span of SHAP values. Alternatively, we can use feature importance scores or the mean absolute SHAP value to rank the covariates based on influence. We have provided a comparison of these two scores in Figure S1. Both scores seem to lead to relatively similar results.

TABLE 1 Patient characteristics

	No. of procedures (%) or median [minimum-maximum]	No. of missing data (%)
Covariates		
Weight, kg	75.0 [5–111]	0 (0)
Age, years	39.8 [0.24–77.7]	0 (0)
BMI	24.1 [13.6–32.8]	21 (10.7)
Treatment center		0 (0)
One	40 (20.3)	
Two	45 (22.8)	
Three	76 (38.6)	
Four	16 (8.1)	
Five	20 (10.2)	
Blood group		26 (13.2)
Non-O	82 (41.6)	
O	80 (40.6)	
FVIII concentrate		3 (1.5)
BDD-rFVIII	28 (14.2)	
Non BDD-rFVIII	166 (84.3)	
High pre-assessed surgical risk	97 (49.2)	0 (0)
Has severe hemophilia	147 (74.6)	0 (0)
Blood loss, ml	0 [0–6700]	0 (0)
Had bleeding complication	30 (15.2)	0 (0)
FVIII baseline level, IUml^{-1}	0.0 [0.0–0.05]	0 (0)
VWF:Ag, %	120 [25–250]	79 (40.1)
VWF:Act, %	130 [24–270]	99 (50.3)
NLME model parameters		
CL, ml/h	163 [29.5–387]	
V1, ml	3030 [260–9710]	
Q, ml/h	56.9	
V2, ml	1270	
ω_{CL} (%CV)	65.2	
ω_{V1} (%CV)	83.5	

Abbreviations: %CV, percent coefficient of variation; BDD-rFVIII, β -domain deleted recombinant clotting factor FVIII; BMI, body mass index; CL, clearance; NLME, nonlinear mixed effects; Q, intercompartmental clearance; V1, central volume; V2, peripheral volume; VWF:Act, von Willebrand factor activity; VWF:Ag, von Willebrand factor antigen.

For both PK parameters, patient weight was the most influential covariate. For clearance (Figure 1a), treatment center, blood group, age, and VWF:Ag appeared to be relatively influential. For central volume (Figure 1b), BMI and use of BDD-rFVIII concentrate seem to be the most important covariates aside from patient weight.

TABLE 2 Accuracy of PK parameter and concentration predictions

	Random forest	XGBoost	XGBoost impute
MAE of CL predictions, ml/h	40.4 ± 10.5 SD (R ² = 0.56)	43.8 ± 10.8 SD (R ² = 0.48)	42.4 ± 11.0 SD (R ² = 0.50)
MAE of V1 predictions, ml	807 ± 320 SD (R ² = 0.49)	893 ± 356 SD (R ² = 0.37)	817 ± 308 SD (R ² = 0.47)
RMSE of concentration predictions, IUml ⁻¹	0.32 ± 0.20 SD	0.36 ± 0.26 SD	0.33 ± 0.22 SD

Abbreviations: CL, clearance; MAE, mean absolute error; PK, pharmacokinetic; RMSE, root mean squared error; SD, standard deviation; V1, central volume.




FIGURE 1 Overview of SHAP values for random forest model. SHAP values of the clearance (a) and central volume (b) are shown as calculated for the random forest model. The covariate value is indicated by color. The horizontal span of the SHAP values indicate the change in the parameters value. The larger the span, the larger the changes in PK parameter and thus the more important the covariate. Covariates are ranked from most (top) to least (bottom) influential by means of their mean absolute SHAP value. BDD-rFVIII, β -domain deleted recombinant clotting factor FVIII; BMI, body mass index; PK, pharmacokinetic; SHAP, SHapley Additive exPlanations; VWF, von Willebrand factor.

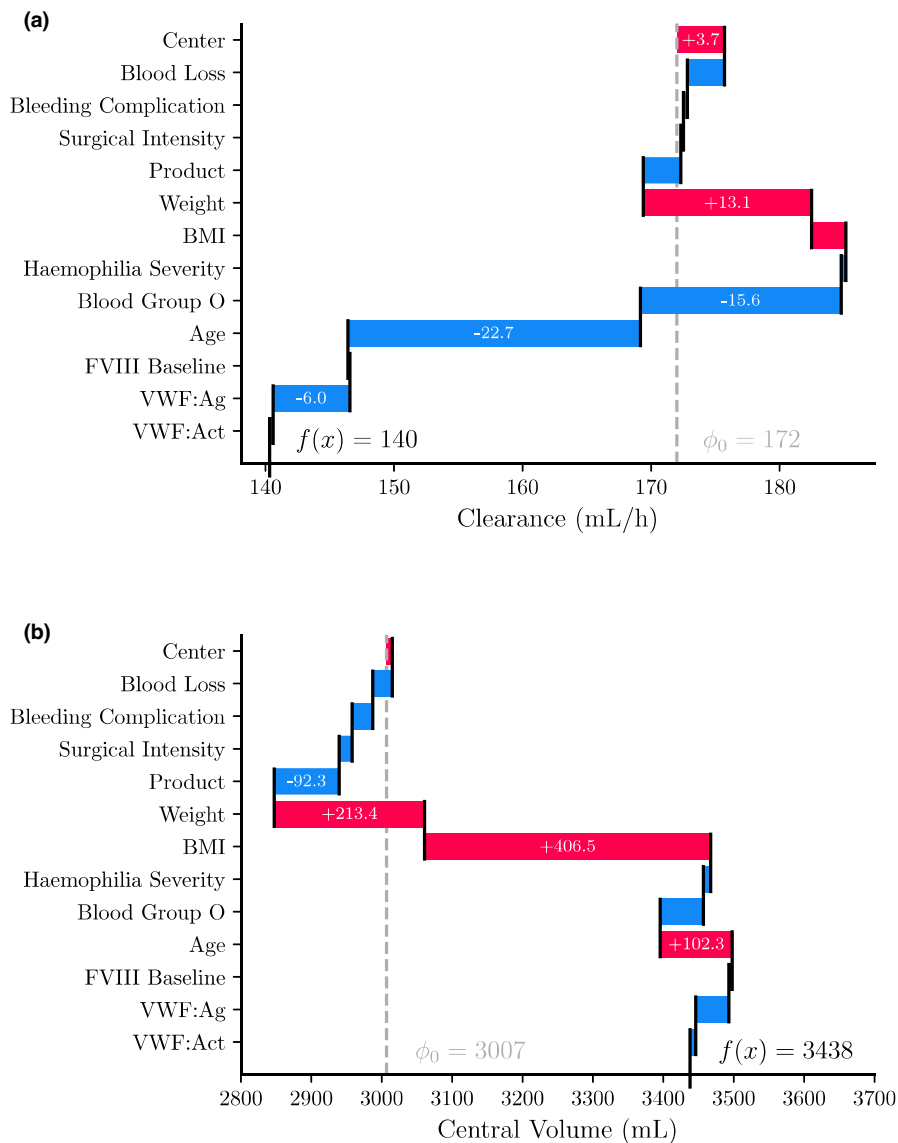
The remaining covariates seem to be less influential for explaining the prediction. We can also take a look at the SHAP values for a single individual (Figure 2). Here, we can see the exact change in clearance and central volume resulting from the inclusion of each covariate.

Our main motivation for performing the SHAP analysis was the ability to visualize the relationship between the calculated SHAP values and each covariate of interest. In Figure 3, we present the resulting relationships for six covariates from the clearance model and three covariates from the central volume model. We observed a positive relationship between body weight and clearance, which flattened for weights above 65 kg (Figure 3a). For age, we saw a negative relationship with clearance, similar to earlier findings.¹⁶ We noticed that individuals with VWF:Ag levels below 100% had higher clearance than those with higher

levels (Figure 3c). In addition, we observed that patients with blood group O displayed an increased clearance compared to non-O individuals (Figure 3d). Both these findings were in line with physiological concepts of hemostasis. Next, we saw that the model predicts a decrease in clearance for individuals in center one, possibly as result of a confounder (Figure 3e). Finally, individuals who received a BDD-rFVIII concentrate displayed slightly increased clearance compared to those who did not (Figure 3f).

For central volume, we also noted a positive relationship with body weight, which flattened slightly with increasing body weight (Figure 3g). We saw a sharp decrease in the SHAP values for central volume for individuals with a BMI 25 (ie, those classified as overweight; Figure 3h). Finally, we saw an increase in the SHAP values for individuals who received BDD-rFVIII concentrate (Figure 3i).

FIGURE 2 SHAP values for a typical patient. SHAP values are shown for the clearance (a) and central volume (b) predictions by the random forest. Data is shown for a 70 kg, 63 year old individual with blood group non-O. SHAP value for each covariate is shown in the corresponding bar. Vertical dashed line indicates expected SHAP value. The SHAP values sum up to the final model prediction. BDD-FVIII, β -domain deleted clotting factor FVIII; BMI, body mass index; SHAP, SHapley Additive exPlanations; VWF, von Willebrand factor.



We could further push the analysis by examining the combined effects of multiple covariates (Figure 4). Because body weight, BMI, and age were correlated, the true effect of either covariate might have been obscured by the others. We combined their respective SHAP values to determine if there was a unique effect of including the separate covariates. After this intervention, there were only small differences between the SHAP values of weight alone versus those of weight and BMI combined for clearance. The same was true for the combined SHAP values of weight and age for central volume. However, combining the SHAP values of weight and age for clearance showed that part of its variance could be well explained by age for individuals with a body weight above 65 kg (Figure 4a). Combining the SHAP values of weight and BMI for central volume resulted in a more pronounced flattening of SHAP values for individuals with a body weight above 65 kg, although considerable variance remained (comparing Figures 3g and

4b). Earlier, we identified a difference in the SHAP values of clearance for patients receiving treatment in center one. The SHAP analysis suggests that individuals without blood group O had SHAP values closer to zero compared to individuals with blood group O (Figure 4c). No such effect is seen for the other centers. For the SHAP values of blood group for clearance, we see a similar result. Here, individuals with lower body weight (65 kg) seem to have SHAP values closer to zero than those with higher body weight (Figure 4d).

A classical approach to obtain intuition on what functional forms to use would be to plot the empirical Bayes estimates of the PK parameters against each of the covariates. This visualization is shown in Figure S2. Here, we see that for highly correlated covariates (ie, weight), it is possible to derive some intuition on the functional form to use, but for most covariates their effect is difficult to discern. This is because we are unable to visualize




FIGURE 3 Relationship between covariates and PK parameters based on SHAP values. Here we visualize the relationship between PK parameter and covariate by plotting SHAP value against covariate value. Points represents the SHAP values, while lines indicate the LOESS fitted smooth representation of the relationship. For the categorical covariates the SHAP value density is also shown by means of a violin plot. We have shown the results for the most important covariates for clearance (a–f) and central volume (g–i). BMI, body mass index; LOESS, locally estimated scatterplot smoothing; PK, pharmacokinetic; SHAP, SHapley Additive exPlanations; VWF, von Willebrand factor.

the contribution of each covariate in isolation. Because unexplained residual variance is also present in the PK parameters, choosing a function to use can be more difficult due to large variation. This can mean that we have to iteratively select functions to implement covariates, reproduce the visualizations, and re-evaluate, thus again resorting to a stepwise approach. With SHAP, we can decide on appropriate functions based on a single full model fit.

Although not shown, the functional forms of the covariates as described by the SHAP values of the two XGBoost models were very similar to those from the random forest. This suggested that the choice between a random forest and XGBoost had only minor effects on the subsequent SHAP analysis.

DISCUSSION

In this study, we aimed to enrich ML-based covariate selection methods using SHAP in order to infer the optimal function form to use when including covariates in PK models. We fit both a random forest and XGBoost model to predict empirical Bayes estimated PK parameters originating from a base NLME model. The random forest resulted in slightly more accurate PK parameter predictions compared to the XGBoost models. Next, influential covariates can, for example, be selected using importance scores.⁵ Finally, after performing a SHAP analysis, we are able to examine the relationship between each covariate and the PK parameters in greater detail. The SHAP analysis also allowed us to explore more complex interaction

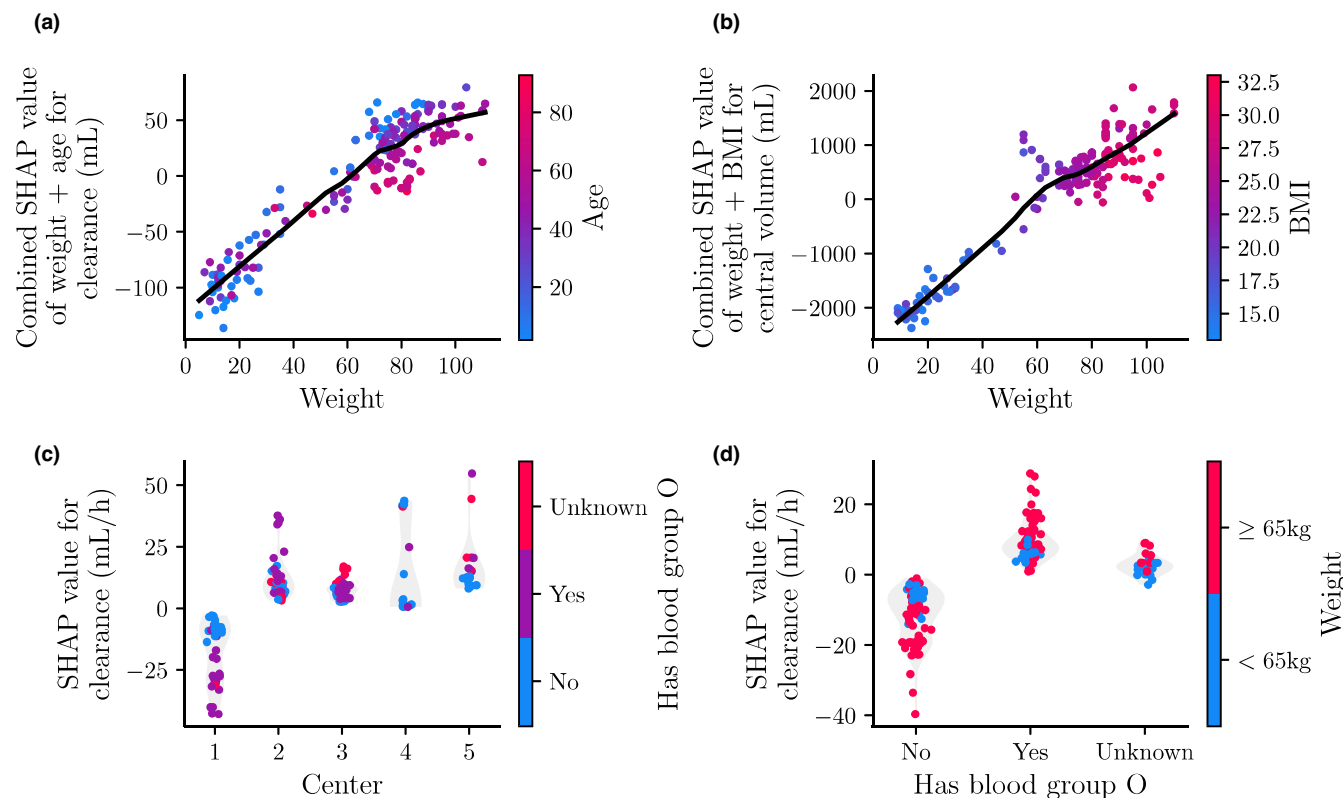


FIGURE 4 Interaction between SHAP values of the covariates. SHAP values of interactions between covariates are shown for the clearance (a, c, and d) and central volume (b) models. Points represents the SHAP values, while lines indicate the LOESS fitted smooth representation of the relationship. The value of the interacting covariate is indicated by color. For the categorical covariates the SHAP value density is also shown by means of a violin plot. BMI, body mass index; LOESS, locally estimated scatterplot smoothing; SHAP, SHapley Additive exPlanations.

effects of covariates resulting from the sequential binning in tree-based methods. Because SHAP values depict the absolute change in output value, the user can intuitively determine clinical relevance. These features display the benefit of SHAP values compared to using importance measures in isolation, where often only a single score of importance is obtained.

The SHAP analysis identified covariates that have previously been associated with the PK of FVIII concentrates. In addition, the suggested relationships of the covariates are similar to their implementation in previous PK models.^{16,20,21} First, we found that patient weight was the most important covariate to explain IIV for both clearance and central volume. The concept of allometric scaling is often applied to the relationship between weight and FVIII clearance. This is mirrored in the flattening of the SHAP values as weight increases (Figure 3a,g). As the central volume compartment represents the blood plasma, a relationship resembling a linear interaction with weight might be expected. An obvious exception exists for overweight individuals, which is represented by the SHAP values in the sharp decline in SHAP values seen for individuals with a BMI greater than 25 (Figure 3h). Measures of fat-free mass have been suggested to better predict

central volume, which could remove the need to model the effect of BMI.²²

Next, we saw a negative interaction between age and clearance. This effect has been demonstrated before,¹⁶ and there might be multiple possible explanations for this effect. One such explanation is the finding that several blood coagulation factors, including VWF, increase with age.^{23,24} It is well known that VWF binds to FVIII to protect it from degradation in the blood circulation. Similar to this effect, SHAP values for patients with blood group O depicted increased FVIII clearance, an effect likely linked to lower VWF:Ag levels seen in patients with blood group O.²⁵ Looking at the interaction between blood group and weight (Figure 4d), we see that individuals below 65 kg (ie, usually younger individuals) with blood group non-O have relatively higher clearance than heavier individuals. This might also be linked to the previously observed increase in VWF:Ag levels with age.^{23,24} It is possible that weight was used by the random forest as a proxy for age. Higher VWF:Ag levels were also directly associated with a decrease in FVIII clearance by the model (Figure 3c). However, considering the large fraction of missing data (40.1%), a low number of patients at the extremes of VWF levels, and the fact that the measurements were outdated

(ie, not measured during the surgical procedure) there remains uncertainty about the observed relationship between VWF:Ag and clearance. Interpreting the effects of covariates with large fractions of missing data should be handled with care.

The SHAP values indicate that individuals from center one had lower clearance compared to other centers. One possible explanation is the use of different assay reagents in this center. The results, however, also indicate that this effect is correlated with the patient blood group (Figure 4c). There could thus be some other factor influencing this effect. Because we worked with retrospective data, it is difficult to underpin the origin of this effect.

Finally, we notice an increase in clearance and central volume associated with patients who received BDD-rFVIII concentrate. It is well known that use of BDD-rFVIII leads to a underestimation of FVIII activity levels when using the one-stage assay versus the chromogenic assay.^{26,27} By changing the phospholipid source in the one-stage assay, similar FVIII activity levels compared to the chromogenic assay are measured. This suggests that this effect is not due to increased clearance or distribution volume of BDD-rFVIII.²⁷ It is possible that this effect leaked into the PK parameter estimates (instead of being part of the estimated error) by the base NLME model. Most of its effect was on increasing the central volume estimate. This can be expected as it would lead to a decrease in predicted FVIII levels.

From the previous discussions, we see the possibility of identifying many subtle effects captured by the random forest model using SHAP. However, the method also has limitations. First, the quality of the empirical Bayes estimated PK parameters is an important factor affecting the accuracy of the ML model and quality of the SHAP analysis. In our case, this required fixing the residual error parameters and only including random effects on clearance and central volume. It might not be clear in advance what measures need to be taken to obtain reliable results. Inspecting the distribution of the resulting PK parameters and comparing these to prior results can be a way to decide on an effective strategy in obtain good quality PK parameter estimates.

Next, we used LOESS to obtain an average representation of the relationship between the covariates and PK parameters. Although this may be helpful for the identification of effects, it might also bias the user to find relationships that do not exist. The method might falsely represent the true effect when SHAP values have high variance or when data are sparse.

Another possible issue lies in the inclusion of covariates that displayed substantial fractions of missing values. For example, roughly 40% of VWF:Ag levels were missing. Although its relationship with clearance suggested by the SHAP values matches previous biological understanding,

we might not want to include the covariate based on the current analysis alone. Previous studies have, however, included this covariate using a function matching the SHAP values.^{20,21}

A more general issue with the application of SHAP values in the context of PK models is that it results in an additive breakdown of the model. Often, covariate effects in PK models are instead implemented as a product of functions. This makes it difficult to compare the outcomes of SHAP analyses with classic methods of covariate analysis, such as forest plots obtained from an FFEM. In addition, by using products, we can prevent the PK parameters from becoming negative. However, because the relationships of the covariates suggested by the SHAP values match those used in previous PK studies, we assume that the functional forms might hold (up to a difference in parameters).^{12,14,15} Such an assumption will have to be validated.

Finally, although SHAP might be able to explain the covariate relationships in the ML model, this does not mean that the results are biologically interpretable. ML algorithms remain black box models, simply deconstructing the model in components does not guarantee that the results are humanly interpretable. For example, we found an effect of center one on clearance, which was correlated with patient blood group. With the current data we are unable to provide an explanation of this effect. Consequently, not every effect found by the SHAP analysis should necessarily be included in PK models.

In summary, we show that combining ML and SHAP allows for an in-depth review of the relationships between covariates and PK parameters. We have mainly focused on using SHAP values for visualizing covariate relationships in ML models. SHAP values can also be used to perform covariate selection. Its benefit over importance scores will have to be evaluated. Covariate selection is a difficult issue, and our method is one of the first to allow one to infer the optimal function form to include covariates based on ML algorithms. The method can prove useful for covariate analysis and hypothesis generation.

AUTHOR CONTRIBUTIONS

A.J., M.H., M.H.C., and R.A.A.M. wrote the manuscript. A.J., M.H., M.H.C., and R.A.A.M. designed the research. A.J. performed the research. A.J. analyzed the data.

ACKNOWLEDGEMENTS

The SYMPHONY consortium, which aims to orchestrate personalized treatment in patients with bleeding disorders, is a unique collaboration among patients, healthcare professionals, and translational and fundamental researchers specializing in inherited bleeding disorders, as well as experts from multiple disciplines. It aims to identify best treatment choice for each individual based on bleeding

phenotype. To achieve this goal, work packages (WP) have been organized according to three themes (eg, Diagnostics [WPs 3 and 4], Treatment [WPs 5–9], and Fundamental Research [WPs 10–12]). This research received funding from the Netherlands Organization for Scientific Research (NWO) in the framework of the NWA-ORC Call grant agreement NWA.1160.18.038. Principal investigator: M.H. Cnossen; project manager: S.H. Reitsma. Beneficiaries of the SYMPHONY consortium: Erasmus MC and Erasmus MC Sophia Children's Hospital, University Medical Center Rotterdam, project leadership and coordination, Sanquin Diagnostics, Sanquin Research, Amsterdam University Medical Centers, University Medical Center Groningen, University Medical Center Utrecht, Leiden University Medical Center, Radboud University Medical Center, Netherlands Society of Hemophilia Patients, Netherlands Society for Thrombosis and Hemostasis, Bayer B.V., CSL Behring B.V., and Swedish Orphan Biovitrum (Belgium) BVBA/SPRL. This study was also performed as part of the OPTI-CLOT international multicenter research consortium, "Patient Tailored Pharmacokinetic (PK) Guided Dosing of Clotting Factor Concentrates and Desmopressin in Bleeding Disorders," which is currently WP 6 within the SYMPHONY consortium. This paper is written on behalf of the international multicenter OPTI-CLOT and To WiN studies that aim to implement a PK-guided approach for the treatment of bleeding disorders using population PK models for desmopressin, factor concentrates, and other alternative drugs. OPTI-CLOT and To WiN study group members are: Steering committee: M.H. Cnossen (principal investigator and chair), F.W.G. Leebeek, Erasmus MC Sophia Children's Hospital and Erasmus MC, University Medical Center Rotterdam, Rotterdam; R.A.A. Mathôt (co-leading investigator), K. Fijnvandraat, M. Coppens, Amsterdam University Medical Center, Amsterdam, University Medical Center, Amsterdam; K. Meijer, University Medical Center Groningen, Groningen; S.E.M. Schols, Radboud University Medical Centre, Nijmegen; H.C.J. Eikenboom, Leiden University Medical Centre, Leiden; R.E.G. Schutgens, University Medical Center Utrecht, Utrecht; E.A.M. Beckers, Maastricht University Medical Center, Maastricht; and P. Ypma, Haga Hospital, The Hague. Principal investigators and local collaborators in the Netherlands: M.J.H.A. Kruip, S. Polinder, Erasmus MC, University Medical Center Rotterdam, Rotterdam; R.Y.J. Tammenga, University Medical Centre Groningen, Groningen; P. Brons, Radboud University Medical Centre, Nijmegen; K. Fischer, K.P.M. van Galen, University Medical Centre Utrecht, Utrecht; F.C.J.I. Heubel-Moenen, Maastricht University Medical Centre, Maastricht; L. Nieuwenhuizen, Maxima Medical Centre, Eindhoven; M.H.E. Driessens, The Netherlands Hemophilia Patient Society; I. van Vliet, Erasmus MC, University Medical

Centre Rotterdam, Rotterdam. OPTI-CLOT/To WiNs: J. Lock, H.C.A.M. Hazendonk, I. van Moort, J.M. Heijdra, M.H.J. Goedhart, W. Al Arashi, Erasmus MC, University Medical Center Rotterdam, Rotterdam; T. Preijers, N.C.B. de Jager, L.H. Bukkems, M.E. Cloesmeijer, A. Janssen, Amsterdam University Medical Centers, Amsterdam. Principal investigators and local collaborators in the United Kingdom—P.W. Collins, Arthur Bloom Haemophilia Centre, Institute of Infection and Immunity, School of Medicine, Cardiff University, Cardiff; R. Liesner, Great Ormond Street Haemophilia Centre, Great Ormond Street Hospital for Children NHS Trust, London; P. Chowdary, Katharine Dormandy Hemophilia Centre and Thrombosis Unit, Royal Free London NHS Foundation Trust, London; C.M. Millar, Hammersmith Hospital-Imperial College Healthcare NHS Trust, London; D. Hart, Department of Haematology, The Royal London Hospital Barts Health NHS Trust, London; and D. Keeling, Oxford Haemophilia and Thrombosis Centre, Oxford University Hospitals, Churchill Hospital, Oxford.

CONFLICT OF INTEREST

The authors declared no competing interests for this work.

ORCID

Alexander Janssen  <https://orcid.org/0000-0003-0719-0831>

REFERENCES

1. Derksen S, Keselman HJ. Backward, forward and stepwise automated subset selection algorithms: Frequency of obtaining authentic and noise variables. *Brit J Math Stat Psychol*. 1992;45:265-282.
2. Ribbing J, Jonsson EN. Power, selection bias and predictive performance of the population pharmacokinetic covariate model. *J Pharmacokinet Pharmacodyn*. 2004;31:109-134.
3. Flack VF, Chang PC. Frequency of selecting noise variables in subset regression analysis: A simulation study. *Am Stat*. 1987;41:84-86.
4. Gastonguay M (2011) Full covariate models as an alternative to methods relying on statistical significance for inferences about covariate effects: A review of methodology and 42 case studies. In: Annual meeting of the population approach group in Europe. Athens, Greece.
5. 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.
6. Wang R, Shao X, Zheng J, et al. A machine-learning approach to identify a prognostic cytokine signature that is associated with nivolumab clearance in patients with advanced melanoma. *Clin Pharmacol Therap*. 2020;107(4):978-987.
7. Ciuculete DM, Bandstein M, Benedict C, et al. A genetic risk score is significantly associated with statin therapy response in the elderly population. *Clin Genet*. 2017;91(3):379-385.
8. Breiman L. Random forests. *Mach Learn*. 2001;45(1):5-32.

9. Friedman JH. Greedy function approximation: a gradient boosting machine. *Annal Stat.* 2001;29:1189-1232.
10. Strobl C, Boulesteix A-L, Zeileis A, Hothorn T. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinform.* 2007;8:1-21.
11. Lundberg S, Lee S-I (2017) A unified approach to interpreting model predictions. arXiv preprint arXiv:170507874
12. Lundberg SM, Erion G, Chen H, et al. From local explanations to global understanding with explainable AI for trees. *Nat Machine Intellig.* 2020;2:56-67.
13. Chen T, Guestrin C (2016) Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining. pp. 785-794
14. Hazendonk H, Fijnvandraat K, Lock J, et al. A population pharmacokinetic model for perioperative dosing of factor VIII in hemophilia a patients. *Haematologica.* 2016;101:1159.
15. Koshy M, Weiner SJ, Miller ST, et al. (1995) Surgery and anesthesia in sickle cell disease. Cooperative Study of Sickle Cell Diseases
16. Beal SL, Sheiner LB, Boeckmann AJ, Bauer RJ (Eds). NONMEM 7.4 users guides; 1989-2018. <https://nonmem.iconplc.com/nonmem743/guides>
17. Björkman S, Oh M, Spotts G, et al. Population pharmacokinetics of recombinant factor VIII: The relationships of pharmacokinetics to age and body weight. *Blood.* 2012;119:612-618.
18. McEneny-King A, Chelle P, Foster G, Keepanasseril A, Iorio A, Edginton AN. Development and evaluation of a generic population pharmacokinetic model for standard half-life factor VIII for use in dose individualization. *J Pharmacokinetic Pharmacodyn.* 2019;46:411-426.
19. Rackauckas C, Nie Q. Differentialequations. jl—a performant and feature-rich ecosystem for solving differential equations in julia. *J Open Res Softw.* 2017;5(1). doi:10.5334/jors.151
20. Nestorov I, Neelakantan S, Ludden TM, Li S, Jiang H, Rogge M. Population pharmacokinetics of recombinant factor VIII fusion protein. *Clin Pharmacol Drug Dev.* 2015;4:163-174.
21. Zhang Y, Roberts J, Tortorici M, et al. Population pharmacokinetics of recombinant coagulation factor VIII-SingleChain in patients with severe hemophilia a. *J Thromb Haemost.* 2017;15:1106-1114.
22. Moort I van, Preijers T, Hazendonk HC, et al. (2021) Dosing of factor VIII concentrate by ideal body weight is more accurate in overweight and obese haemophilia a patients. *Br J Clin Pharmacol* 87:2602-2613
23. Gill JC, Endres-Brooks J, Bauer PJ, et al. The effect of ABO blood group on the diagnosis of von willebrand disease. *Blood.* 1987;69:1691-1695.
24. Franchini M. Hemostasis and aging. *Crit Rev Oncol Hematol.* 2006;60:144-151.
25. Klarmann D, Eggert C, Geisen C, et al. Association of ABO (h) and i blood group system development with von willebrand factor and factor VIII plasma levels in children and adolescents. *Transfusion.* 2010;50:1571-1580.
26. Mikaelsson M, Oswaldsson U, Jankowski MA. Measurement of factor VIII activity of B-domain deleted recombinant factor VIII. *Semin Hematol.* 2001;38:13-23.
27. Lee C, Kessler C, Varon D, et al. Influence of phospholipids on the assessment of factor VIII activity. *Haemophilia: State of the Art.* 1998;4:646-650.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Janssen A, Hoogendoorn M, Cnossen MH, Mathôt RAA, for the OPTI-CLOT Study Group and SYMPHONY Consortium. Application of SHAP values for inferring the optimal functional form of covariates in pharmacokinetic modeling. *CPT Pharmacometrics Syst Pharmacol.* 2022;11:1100-1110. doi:[10.1002/psp4.12828](https://doi.org/10.1002/psp4.12828)