

MIT Open Access Articles

Interpretable Predictive Models to Understand Risk Factors for Maternal and Fetal Outcomes

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Bosschieter, T.M., Xu, Z., Lan, H. et al. Interpretable Predictive Models to Understand Risk Factors for Maternal and Fetal Outcomes. *J Healthc Inform Res* 8, 65–87 (2024).

As Published: <https://doi.org/10.1007/s41666-023-00151-4>

Publisher: Springer International Publishing

Persistent URL: <https://hdl.handle.net/1721.1/153417>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Interpretable Predictive Models to Understand Risk Factors for Maternal and Fetal Outcomes

This Accepted Manuscript (AM) is a PDF file of the manuscript accepted for publication after peer review, when applicable, but does not reflect post-acceptance improvements, or any corrections. Use of this AM is subject to the publisher's embargo period and AM terms of use. Under no circumstances may this AM be shared or distributed under a Creative Commons or other form of open access license, nor may it be reformatted or enhanced, whether by the Author or third parties. By using this AM (for example, by accessing or downloading) you agree to abide by Springer Nature's terms of use for AM versions of subscription articles: <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

The Version of Record (VOR) of this article, as published and maintained by the publisher, is available online at: <https://doi.org/10.1007/s41666-023-00151-4>. The VOR is the version of the article after copy-editing and typesetting, and connected to open research data, open protocols, and open code where available. Any supplementary information can be found on the journal website, connected to the VOR.

For research integrity purposes it is best practice to cite the published Version of Record (VOR), where available (for example, see ICMJE's guidelines on overlapping publications). Where users do not have access to the VOR, any citation must clearly indicate that the reference is to an Accepted Manuscript (AM) version.

Interpretable Predictive Models to Understand Risk Factors for Maternal and Fetal Outcomes

Tomas M. Bosschier^{1*}, Zifei Xu¹, Hui Lan¹, Benjamin J. Lengerich², Harsha Nori³, Ian Painter⁴, Vivienne Souter MD^{4*} and Rich Caruana^{3*}

¹Stanford University, Stanford, USA.

²Massachusetts Institute of Technology, Cambridge, USA.

³Microsoft Research, Redmond, USA.

⁴Foundation for Healthcare Quality, Seattle, USA.

*Corresponding author(s). E-mail(s): tomasbos@stanford.edu; vsouter@qualityhealth.org; rcaruana@microsoft.com;

Abstract

Although most pregnancies result in a good outcome, complications are not uncommon and can be associated with serious implications for mothers and babies. Predictive modeling has the potential to improve outcomes through better understanding of risk factors, heightened surveillance for high risk patients, and more timely and appropriate interventions, thereby helping obstetricians deliver better care. We identify and study the most important risk factors for four types of pregnancy complications: (i) severe maternal morbidity, (ii) shoulder dystocia, (iii) preterm preeclampsia, and (iv) antepartum stillbirth. We use an Explainable Boosting Machine (EBM), a high-accuracy glass-box learning method, for prediction and identification of important risk factors. We undertake external validation and perform an extensive robustness analysis of the EBM models. EBMs match the accuracy of other black-box ML methods such as deep neural networks and random forests, and outperform logistic regression, while being more interpretable. EBMs prove to be robust. The interpretability of the EBM models reveal surprising insights into the features contributing to risk (e.g. maternal height is the second most important feature for

2 *Interpretable Machine Learning for Pregnancy*

shoulder dystocia) and may have potential for clinical application in the prediction and prevention of serious complications in pregnancy.

Keywords: Interpretability, explainable models, AI for healthcare, pregnancy complications, stillbirth, preeclampsia.

Acknowledgments

We are indebted to Kristin Sitcov, Executive Director Clinical Programs at the Foundation for Health Care Quality (FHCQ), for her support for this project.

1 Introduction

Of the 3.6 million births per year in the U.S. [1], Severe Maternal Morbidity (SMM) happens in as many as 60,000 cases [2], leading to serious short- or long-term consequences for the mother's health.¹ Shoulder dystocia and (preterm) preeclampsia are two other relatively common, serious conditions. Antepartum stillbirth is a relatively rare, but devastating outcome. If accurate models could be trained that were interpretable and provided timely risk prediction, obstetric providers could provide more personalised care, potentially averting some cases of SMM, shoulder dystocia, preterm preeclampsia and antepartum stillbirth. [2, 3]. In this work we use glassbox Explainable Boosting Machines (EBMs) to train interpretable, high accuracy machine learning models for each of these outcomes.

When modeling risk for SMM, preterm preeclampsia and antepartum stillbirth, we only use features known early in pregnancy. For shoulder dystocia we use all features known just prior to delivery. While several analyses and (black-box) models have been developed to predict SMM and its risk factors [4–8], (preterm) preeclampsia [9–12], shoulder dystocia [13, 14], and antepartum stillbirth [15–17], the risk factors for each have remained under-studied using interpretable machine learning. Additionally, few models are externally validated, i.e. validated on a hold-out set containing data from a disjoint set of hospitals [18–20]. Furthermore, models trained on international data (e.g. [21]) might not perform as well in U.S.-based hospitals [22].

Equipped with recent clinical data containing 158,629 de-identified birth events and over 100 features, we use interpretable machine learning models, Explainable Boosting Machines (EBMs), to uncover the most important risk factors for each outcome. We show that EBMs yield an Area Under the Receiver Operating Characteristic curve (AUROC) on par with XGBoost [23], random forests [24], and deep neural networks (DNNs) [25], while outperforming logistic regression [26]. Additionally, we provide empirical evidence that

¹In our work, SMM is a composite term for 6 adverse diagnoses: (i) hysterectomy, (ii) blood transfusion, (iii) disseminated intravascular coagulation, (iv) amniotic fluid embolism, (v) thromboembolism, and (vi) eclampsia.

EBMs are robust by showing they generalize well and yield well-calibrated probabilities, and by evaluating the shape functions of the EBM models as they evolve as a function of the data set size. We compute the discrete Fréchet distances from the shape functions trained on subsets of the data to the shape function trained on the largest data set, which shows the convergence of the sequence of shape functions. Robustness of models is an especially important consideration in healthcare [18, 27].

Our main contributions are:

- Illustrating the value of interpretable models in four healthcare case studies.
- Leveraging a new and robust data set in healthcare to model risk of SMM, shoulder dystocia, preterm preeclampsia, and antepartum stillbirth.
- Demonstrating how intelligible models reveal surprising risk factors not traditionally recognized as important.
- Empirically evaluating the robustness, accuracy and calibration of EBMs.

2 Preliminaries

2.1 Explainable Boosting Machines

For a target variable Y and predictor features x_1, \dots, x_n , Generalized Additive Models (GAMs) [28] generalize linear (regression) models $Y = b_0 + a_1x_1 + \dots + a_nx_n$ to an additive model with univariate shape functions f_i and a situation-dependent link function g (e.g. logistic for classification and identity for regression):

$$g(\mathbb{E}[Y]) = \beta_0 + f_1(x_1) + f_2(x_2) + \dots + f_n(x_n). \quad (1)$$

The Explainable Boosting Machine (EBM) is an algorithm for training the f_i by boosting shallow decision trees in a round-robin fashion [29, 30] (see Figure 1). Additionally, EBMs can be trained to automatically detect important pairwise interaction terms to further boost accuracy while preserving intelligibility [31]. Because EBMs are GAMs, the trees trained on the i -th feature, x_i , combine to provide a component function $f_i(x_i)$ which is piecewise constant. Plotting $f_i(x_i)$ provides the end-user with a visual representation of the function learned by the tree ensemble, e.g. those seen at the bottom of Figure 1. Observe that through this training process and having a main effect $f_i(x_i)$ for each feature x_i , each feature's shape function is learned in parallel with and corrected for all other features in the model. This additivity is an important property for intelligibility and transparency. Observe that while the shape functions shown in Figure 1 look similar to partial dependence plots (PDPs), there is a crucial difference: the PDP of a feature is the feature's marginal effect to the target variable, while a feature's shape function constitutes the feature's contribution to the target variable *while adjusted for all other variables*.

4 *Interpretable Machine Learning for Pregnancy*

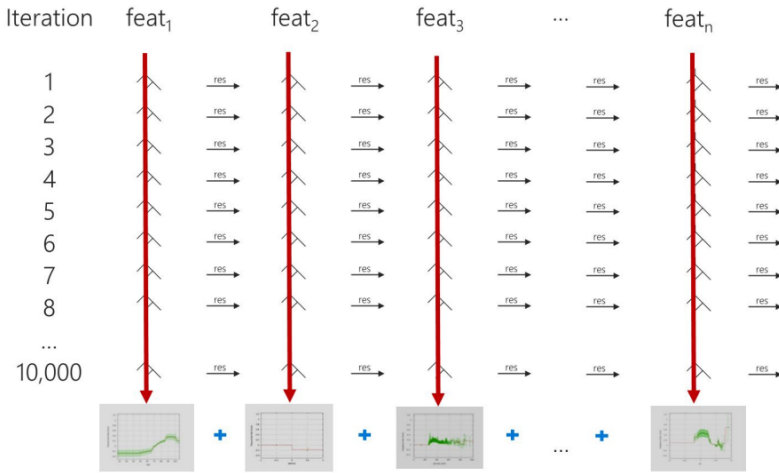


Fig. 1: EBMs are GAMs that perform cyclic gradient boosting with shallow decision trees. As an additive models, the set of trees trained on a feature could be ensembled to form the shape function of that feature. This image illustrates the ensembling (with the red vertical arrows) as well as the training process more generally. Details can be found in [30].

EBM models use a combination of inner- and outer-bagging. At each iteration and for each feature x_i , multiple subsets are taken from the residual, and an ensemble is constructed from all shallow trees trained on the samples. This is called inner-bagging. Outer-bagging refers to bagging the entire process, i.e. retraining the entire model multiple times on random subsamples of the data. This combination of inner- and outer-bagging allows EBM models to improve accuracy, reduce variance, make the learned $f_i(x_i)$ smoother and easier to interpret, and also provide confidence intervals on the $f_i(x_i)$.

EBMs have recently gained popularity because of their high accuracy and local and global interpretability, with applications ranging from healthcare [32] to predicting sports outcomes [33, 34], slope failures [35], and modeling dark matter [36].

2.2 Discrete Fréchet distance

To evaluate the robustness of EBMs we examine how the shape functions of EBMs converge as the number of samples in a dataset increases. Since the models might find different best-fit splits in the feature values when training decision trees for different data sets, the shape functions might learn different bins for each data set size. This makes it hard to compare shape functions using standard measures such as the root-mean-square error (RMSE). Instead, we turn to the discrete Fréchet distance for polygonal chains, a measure of how similar two curves are given that the partitions in each might be different. The Fréchet distance is defined as follows:

Definition 1 (Fréchet distance [37]) Let (S, d) be a metric space and $A, B: [0, 1] \rightarrow S$ curves in S , i.e. continuous mappings. Denoting $\alpha, \beta: [0, 1] \rightarrow [0, 1]$ reparameterizations, the Fréchet distance between A and B is defined as

$$d_F(A, B) = \inf_{\alpha, \beta} \max_{t \in [0, 1]} \{d(A(\alpha(t)), B(\beta(t)))\}. \quad (2)$$

3 Methods

3.1 Data set

The clinical data used in this analysis are from the Foundation for Health Care Quality’s Obstetrical Care Outcomes Assessment Program (OB COAP) [38] and comprised de-identified patient-level clinical data (not administrative or coding data) acquired from medical records from January 2016 till December 2021 at 20 hospitals from Northwestern U.S. collected for quality improvement purposes. The data set contains 158,629 births, but specific filters are applied per outcome, lowering the number of samples in the data set used for each outcome.

For shoulder dystocia we consider only singleton births and no cesarean births, while excluding antepartum stillbirths, so that we consider only the population for whom shoulder dystocia risk prediction is useful. For antepartum stillbirth, we restrict the population to singletons to remove ambiguity about which fetus the outcome is for. The number of samples, positive samples, and features used for each outcome are listed in Table 1. See Section 3.1.1 for details on filtering and preprocessing of the initial data set as a whole.

Table 1: Data set characteristics per outcome.

Outcome	Num. samples	Num. positives	Num. features
SMM	155,935	2262 (1.5%)	49
Shoulder dystocia	82,889	2794 (3.4%)	56
Preterm preeclampsia	153,432	2919 (1.9%)	29
Antepartum stillbirth	153,432	610 (0.4%)	29

3.1.1 Feature engineering and preprocessing assumptions

For each outcome we use clinical expertise to identify and select the appropriate features, i.e. those that can be measured before the outcome occurs (with one exception: the shoulder dystocia model includes birthweight, a feature that can only be estimated before birth). For example, in trying to predict preterm preeclampsia (i.e. preeclampsia before 37 weeks’ gestation)², a feature such as ‘time from hospital admission to delivery’ should not be used by a machine learning model in our setting. This is because labor has not taken place yet and so this feature’s value is not known at the time of prediction, i.e.

²Note that predicting preterm preeclampsia can only be done in practice < 37 weeks into the pregnancy per definition of preterm.

6 *Interpretable Machine Learning for Pregnancy*

< 37 weeks' gestation. However, the feature 'time from hospital admission to delivery' in the data set is selected as a feature to be used in predicting shoulder dystocia because it's known at delivery. Full lists of included features, ranked by importance to each risk prediction outcome, can be found in Appendix A.

As a preprocessing step, we discard birth events with data deemed implausible or impossible by clinical experts (and assumed to be entered into the system erroneously), taking a conservative approach to avoid accidentally discarding valid data. For example, we never include births with a negative length of time from admission till delivery, births where the baby's birthweight is over 8000 grams, and cases where the pregnant patient has a BMI over 120 kg/m². Furthermore, we use the standard procedure of dummy encoding for categorical data, and impute missing values for continuous features with the mean and missing values for categorical features with a unique identifier, typically -1. Data normalization is not required for tree-based methods such as EBMs, random forests and boosted trees because trees are scale-invariant.

3.2 Experimental setup

To obtain training and external validation data sets, we first stratify births by hospital. The set of hospitals \mathcal{H} is then partitioned into 3 sets according to the level of care they provide. There are 7 level-1 hospitals, 6 level-2 hospitals, and 7 level-3 hospitals, each hospital containing a different number of samples. To do external validation, we partition the hospitals \mathcal{H} into 2 sets, where the births in the first set of hospitals constitute the training set, and the second set of hospitals contains the "external validation" set (consisting of births from a disjoint set of hospitals). To ensure (near) equal representation of levels of care in each of the training and external validation sets, we generate many two-way partitions of the hospitals, and select only those which at the aggregate have: 1) a near 75%/25% split between train and test patients, and 2) yield more than 1 hospital of each level in each set so as to minimize the risk of overfitting to any one hospital and level.

The patients in the 25% set are from a disjoint set of hospitals that are used for external validation. External validation is important in healthcare to ensure that the learned models are robust, generalize well to other hospitals, and are well-calibrated for patients at all risk levels. We use a bootstrap analysis to generate confidence intervals for the AUCs reported in Table 2.

3.2.1 Model parameters

The hyperparameters of all models were each determined using 5-fold randomized search Cross-Validation to maximize AUROCs while ensuring good calibration; all parameters not mentioned are common defaults. For SMM, shoulder dystocia and preterm preeclampsia: for EBMs, we use hyperparameters `outer_bags=25`, `inner_bags=25`, `min_samples_leaf=25`, `interactions=10`. For XGBoost we use `eta=0.04`, `subsample=0.7`, and `max_depth=5`. Random forests are trained with `n_estimators=1000`, `min_samples_split=60`, and

`min_samples_leaf=40`. Lastly, the DNN is a multilayer perceptron (MLP) with 7 hidden layers of 100 neurons each. For antepartum stillbirth, the outcome with the smallest number of positives, optimization lowers the number of `min_samples_leaf` to 15 for EBMs, and 25 for random forests.

3.3 Evaluating EBM robustness

While the robustness of many models has been well-studied, it appears this is not fully the case yet for Explainable Boosting Machines. We provide empirical evidence that EBMs are robust in that its constituent shape functions converge rapidly as a function of data sample size and the number of positives therein, and that they generalize well. The latter is shown by performing external validation, see Section 4. The former is done by showing (i) the shape functions' evolutions visually (see Section 4.1), and (ii) the sequence of Fréchet distances of the continuous shape functions trained on subsampled data sets compared to the final shape function trained on all data. For discrete features, we use RMSE instead of Fréchet distance.

For each of the 4 outcomes we select one important feature to perform this analysis on. We examine 2 continuous features and 2 categorical features, one of which is binary. Specifically, the outcome-feature pairs are:

- **Shoulder dystocia:** *birthweight* (continuous). Birthweight is the most important predictor for shoulder dystocia. See Figure 8.
- **Preterm preeclampsia:** *maternal BMI* (continuous). Maternal BMI is chosen because of its complex shape function: smooth and near-linear in some regions, while also exhibiting significant jumps in other regions, presumably due to clinical interventions. See Figure 10.
- **SMM:** *nulliparity* (binary). Nulliparity is the second most important feature for SMM, and traditionally considered a crucial risk indicator. See Figure 11.
- **Antepartum stillbirth:** *race* (categorical). Race is the second most important feature in predicting antepartum stillbirth, and contains multiple classes. See Figure 12.

4 Results

For each outcome (SMM, shoulder dystocia, preterm preeclampsia, antepartum stillbirth) we compare the EBM's performance to those of XGBoost, random forests, deep neural networks, and logistic regression. The mean Area Under the ROC curve (AUROC) for each model and outcome can be found in Table 2. All results presented are obtained through external validation.

EBMs allow each feature's contribution to the log-odds of the risk to be easily visualized since EBMs are additive models, as discussed in Section 2.1. For each outcome we present 2 shape functions together with a feature importance plot, which shows the top 15 features in terms of their feature importance, i.e. their mean absolute contribution to the log-odds of the outcome prediction. This is more informative than simply a ranking of features, as given in

Table 2: A comparison between the Area Under the ROC curve (AUROC) for each outcome. Higher is better. The EBM has the highest average AUROC.

Outcome	EBM	Logistic Regression	XGBoost	Random Forests	DNN
SMM	0.700 ± 0.013	0.683 ± 0.015	0.679 ± 0.013	0.683 ± 0.015	0.677 ± 0.018
Shoulder dystocia	0.744 ± 0.017	0.742 ± 0.015	0.751 ± 0.018	0.713 ± 0.019	0.752 ± 0.019
Preterm preeclampsia	0.767 ± 0.013	0.735 ± 0.010	0.767 ± 0.010	0.749 ± 0.015	0.754 ± 0.014
Antepartum stillbirth	0.710 ± 0.011	0.691 ± 0.010	0.714 ± 0.012	0.713 ± 0.017	0.706 ± 0.019
Mean AUROC	0.730 ± 0.014	0.713 ± 0.013	0.728 ± 0.013	0.715 ± 0.017	0.722 ± 0.018

Appendix A, as the actual measured feature importance is reported in the figure. Note that for the shape functions, the error bars shown represent the standard deviations yielded by the outer-bagging process employed by EBMs.

4.0.1 SMM

The shape functions of the *maternal age at admission* and *maternal height* features are in Figure 2a, while the feature importance plot is in Figure 2b. The most important features according to the EBM are *preeclampsia/gestational hypertension, labor type*³, *nulliparity, cervical dilation, race*, and the *Distressed Communities Index (DCI) quintile* [39]. The complete, ranked list of features used to train all models for SMM risk prediction can be found in Appendix A.1.

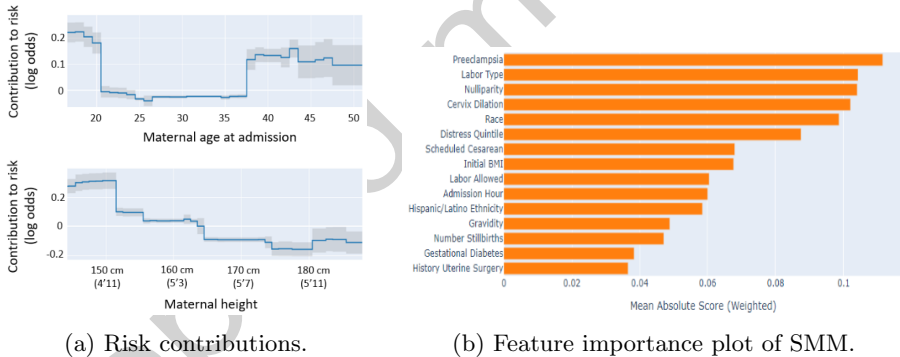



Fig. 2: The shape functions learned by the EBM for SMM of (i) maternal age at admission, and (ii) maternal height, along with SMM’s feature importance plot [40]. The full list of ranked features can be found in Appendix A.1.

4.0.2 Shoulder dystocia

The EBM’s shape functions for *maternal height* and the *baby’s birthweight* can be found in Figure 3a, while Figure 3b shows the feature importance plot. The complete list of importance-ranked features can be found in Appendix A.2.

³Labor type refers to whether labor was spontaneous or induced.



(a) The individual contributions of baby weight and maternal height to shoulder dystocia risk. (b) The feature importance plot for shoulder dystocia. Full list of ranked features can be found in Appendix A.2.

Fig. 3: Two shape functions of the EBM for shoulder dystocia, as well as the feature importance plot. [41]

We observe a clinically surprising near-linear (in log odds) relationship between *birthweight* and risk in Figure 3a for weights between 3250g and 4250g. *Birthweight* is clearly the most important feature in the feature importance ranking in Figure 3b. The second most important feature is *maternal height*, a feature not typically considered a major risk factor for shoulder dystocia. One might speculate that taller women might have larger pelvic openings on average than smaller women, resulting in a lower risk of shoulder dystocia. However, there are other factors that correlate with height, such as socioeconomic background.


4.0.3 Preterm preeclampsia

The EBM's shape functions of *maternal BMI* and *age* are shown in Figure 4a, while the feature importance plot is displayed in Figure 4b. A complete list of features used by the EBM can be found in Appendix A.3, ranked by importance. We find that the most important risk factors for preterm preeclampsia are *maternal BMI*, *nulliparity*, *chronic hypertension*, and *age*.

4.0.4 Antepartum stillbirth

Antepartum stillbirth's most important features are listed in Figure 5b, while the shape functions for the *Distressed Communities Index (DCI) quintile* and *maternal BMI* are shown in Figure 5a. The most important features seem to be *maternal BMI*, while *race*, *DCI quintile*, *maternal age*, *maternal height*, and *Rural-Urban Commuting Area (RUCA) class* [42] also appear to play important roles. A complete list of features used to train the ML models can be found in Appendix A.4, where the features are ranked by feature importance.

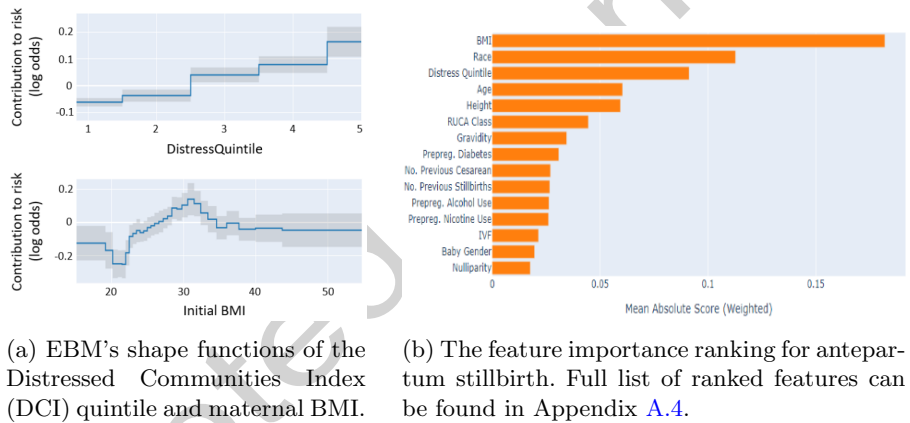
10 *Interpretable Machine Learning for Pregnancy*



(a) The risk contributions of maternal BMI and age to preterm preeclampsia. [11]

(b) Feature importance plot for preterm preeclampsia risk. A full list of features can be found in Appendix A.3.

Fig. 4: Shape functions for preterm preeclampsia along with the most important features.



(a) EBM's shape functions of the Distressed Communities Index (DCI) quintile and maternal BMI.

(b) The feature importance ranking for antepartum stillbirth. Full list of ranked features can be found in Appendix A.4.

Fig. 5: The relationship between antepartum stillbirth risk and (i) the DCI quintile, and (ii) maternal BMI measured at the initial visit to the mother's healthcare provider. Also shown is the top 15 features ranked according to their feature importance. [15]

4.0.5 Race and ethnicity

Including race in clinical models is controversial [43–45]. However, if race or ethnicity were removed as a feature in the model, this would not necessarily make the model unbiased, because of correlation between race and ethnicity and other features in the model [46]. One way to mitigate bias in an EBM model is to include race and ethnicity as features when the model is trained, but then flatten or zero out the shape functions for race and ethnicity before deployment.

Figure 6 shows the risk contribution of race and ethnicity for shoulder dystocia and SMM. The lowest risk for both shoulder dystocia and SMM seems to be in the group identified as White. The highest risk for shoulder dystocia is associated with pregnant people of missing, Black, and Asian race. Importantly, for SMM, the highest risk is associated with American Indian/Alaskan Native and Multiple races, even higher than for the Black pregnant population. It is noted however that the error bars are large enough for some race categories to overlap. For ethnicity, the lowest risk is associated with having missing data; the reason for this is unclear. Being of Hispanic/Latina ethnicity, when reported, is associated with a higher risk than non-Hispanic/Latina ethnicity.

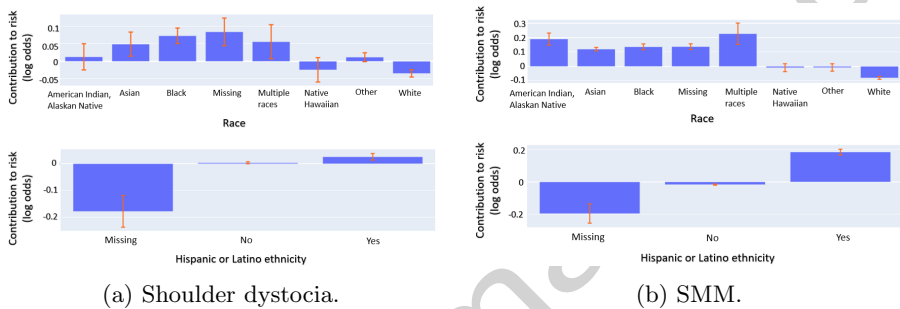


Fig. 6: The risk contribution of race and ethnicity for shoulder dystocia and SMM, after correcting for all other features. Preterm preeclampsia is similar. For more detail of race in antepartum stillbirth, see Section 4.1.4. For SMM, the ‘Missing’ bucket represents 5.6% of the data for the feature ‘Race’, and 3.7% for ‘Hispanic or Latina ethnicity’. For shoulder dystocia, these are 4.8% and 4.0% respectively.

4.0.6 Model calibration

Another important consideration is model calibration. Because SMM, shoulder dystocia, preterm preeclampsia, and antepartum stillbirth are all unbalanced, ML models could be well-calibrated for the majority of patients, yet be uncalibrated for high-risk patients. This is why the within-level splits discussed in Section 3.2 are of such importance. The calibration plots, one for each outcome, are in Figure 7. All calibration plots are generated using 10 bins. Note that only EBM and logistic regression consistently yield good calibration, while XGBoost and Random Forests tend to slightly overfit and are overconfident, e.g. for antepartum stillbirth in Figure 7d, where bad calibration for the (top 10%) highest-risk patients can be seen. Figure 7 shows that EBMs do in fact attain calibration as good as or better than the gold standard, logistic regression.

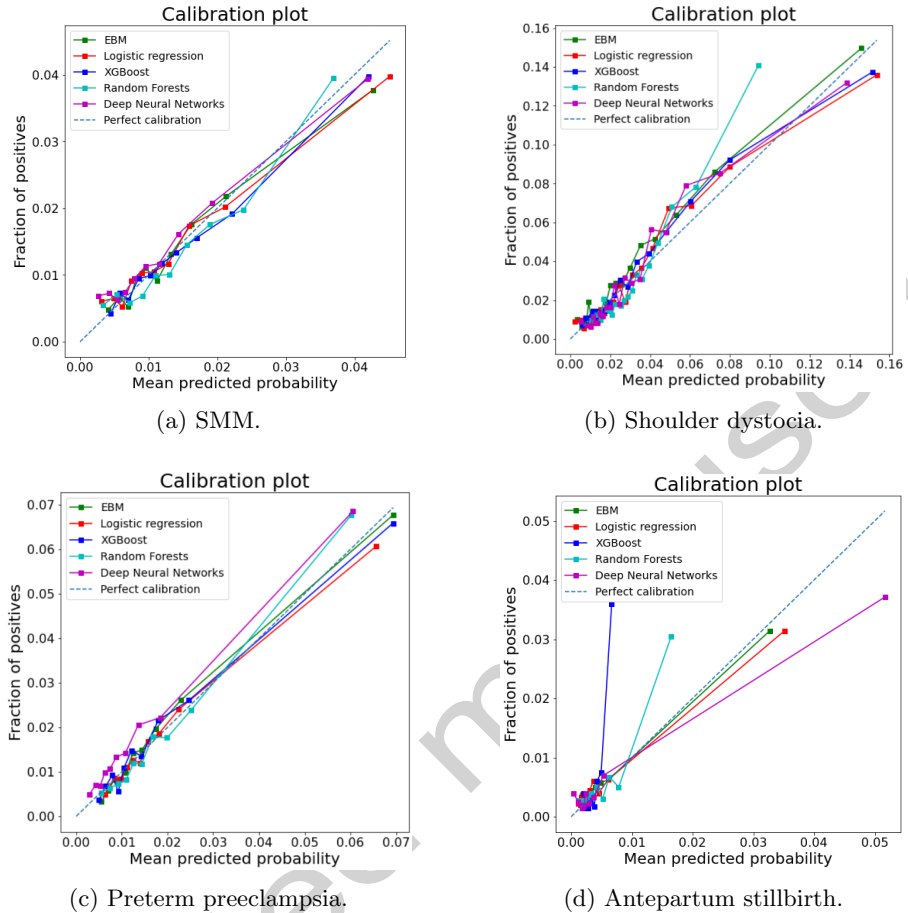


Fig. 7: Calibration plots of the EBM, logistic regression, random forests, XGBoost, and DNN models for SMM, shoulder dystocia, preterm preeclampsia, and antepartum stillbirth. 10 bins are used in the plots.

While the calibration plots appear to show proper calibration of the EBM model, we also provide a quantitative comparison of the calibrations by computing the models' Brier scores, see Table 3.

Table 3: A comparison between the Brier scores for each outcome computed over 5 randomized trials. Lower is better.

Outcome	EBM	Logistic Regression	XGBoost	Random Forests	DNN
SMM	0.013 ± 0.003	0.013 ± 0.003	0.013 ± 0.003	0.014 ± 0.003	0.013 ± 0.003
Shoulder dystocia	0.028 ± 0.004	0.028 ± 0.003	0.031 ± 0.003	0.030 ± 0.004	0.030 ± 0.004
Preterm preeclampsia	0.019 ± 0.002	0.019 ± 0.001	0.020 ± 0.003	0.022 ± 0.002	0.023 ± 0.003
Antepartum stillbirth	0.003 ± 0.001	0.003 ± 0.001	0.005 ± 0.001	0.004 ± 0.001	0.004 ± 0.001

4.1 Results of evaluating EBM robustness

We find that the EBMs are well-calibrated, and learn trends relatively well for individual features even when there are (very) few positives in the data set. Following our experimental setup as described in Section 3.2, we show the evolution of the shape functions for 2 continuous and 2 categorical features: *nulliparity* (categorical) for SMM, *birthweight* (continuous) for shoulder dystocia, *maternal BMI* (continuous) for preterm preeclampsia, and *Race* (categorical) for antepartum stillbirth. We present the evolution of the shape functions for the continuous features first.

4.1.1 Shoulder dystocia

Figure 3b showed that birthweight is the most important feature in shoulder dystocia risk prediction, and we now demonstrate how that feature develops while trained with random subsets of sizes 500, 1000, 2000, 5000, 10000, 20000, 30000, 40000, 50000, 60000, 70000, and the entire data set containing 82889 birth events.

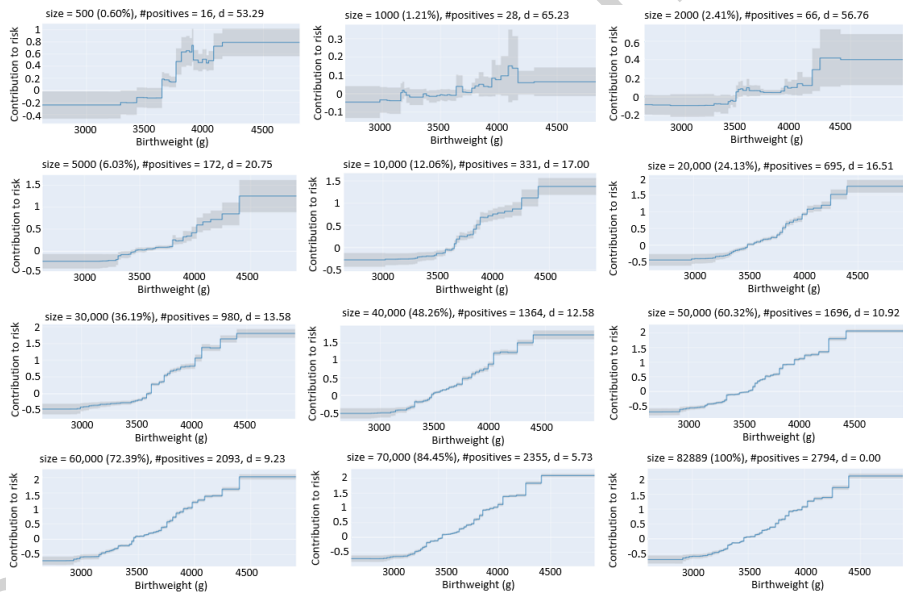


Fig. 8: Evolution of the shape functions of *birthweight* while predicting shoulder dystocia, as the data set size (“size” in the figure) is varied. The parameter “d” represents the discrete Fréchet distance discussed in Section 2.2.

Additionally, Figure 9 shows the performance of the models as a function of data set size while modeling shoulder dystocia. This plot appears to suggest that state-of-the-art performance can be achieved with fewer data points.




Fig. 9: AUC of various models shown as a function of the data set size, where the error bars represent the standard error over 5 randomized trials. This is done for shoulder dystocia only given limited computational power.

4.1.2 Preterm preeclampsia

See Figure 10 for the evolution of the shape function of *initial maternal BMI*, a continuous feature used to train the EBM in predicting preterm preeclampsia risk. We again report the size of the sampled data set and the number of positives therein, where we select data set sizes of 500, 1000, 2000, 5000, 10000, 20000, 50000, 75000, 100000, 125000, 150000, and the entire data set of size 153432. Additionally, “d” represents the discrete Fréchet distance to the final shape function, representing its similarity to the best-fit shape function. We observe that the Fréchet distance decreases rapidly after a data set size of 2000 (1.30%) with only 39 positives, see both Figure 10 for the shape function evolution, as well as the Fréchet distance graph in Figure 13.

4.1.3 SMM

Nulliparity is an important categorical feature that takes on either 0 or 1: 1 represents nulliparity is true, i.e. the mother has never given birth before, and 0 represents false nulliparity. The development of the shape function of nulliparity is visualized in Figure 11, showing the shape functions when trained on random subsets of the data set of sizes 500, 1000, 2000, 5000, 10000, 20000, 50000, 75000, 100000, 125000, 150000, and the entire data set of size 155935. Despite the number of positives in the sampled data set being small, the EBM

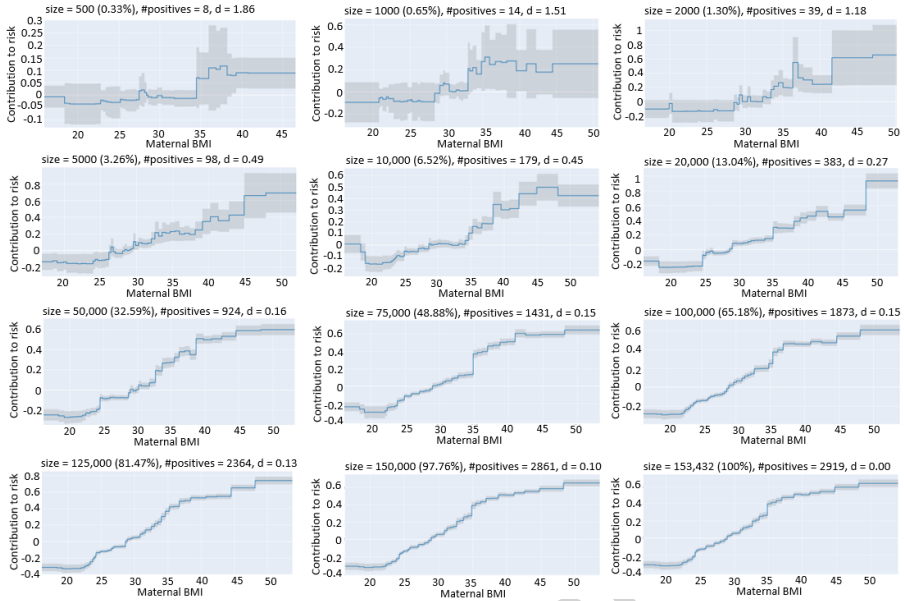


Fig. 10: Evolution of the shape functions of maternal BMI while predicting preterm preeclampsia, as the data set size (“size” in the figure) is varied. The discrete Fréchet distance to the final shape function is reported as “d”.

appears to be able to pick up the signal for nulliparity even with small subsets of the data, presumably because nulliparity is Boolean.

4.1.4 Antepartum stillbirth

For antepartum stillbirth we pick *race* as a feature to explore, which is categorical with multiple classes, ranking second in the feature importance list, see Appendix A.4. Similar to preterm preeclampsia, we select data set sizes of 500, 1000, 2000, 5000, 10000, 20000, 50000, 75000, 100000, 125000, 150000, and 153432 (complete data set). See Figure 12.

4.1.5 Convergence of shape functions

We observe that the EBMs’ shape functions converge quickly as a function of the data set size and number of positives. For continuous features we compute the Fréchet distances, and denote these for each shape function trained on a subset of the full data set, see Figures 8 and 10. Because the scales of risk contributions in Figures 8 and 10 are very different, so are the corresponding Fréchet distances. We normalize the Fréchet distances by dividing the distances by the first in order to more accurately show the relative convergence of the sequence of shape functions. See Figure 13.

For shape function evolutions of categorical features, as in Figures 11 and 12, the main impact of having a larger data set once there are enough samples

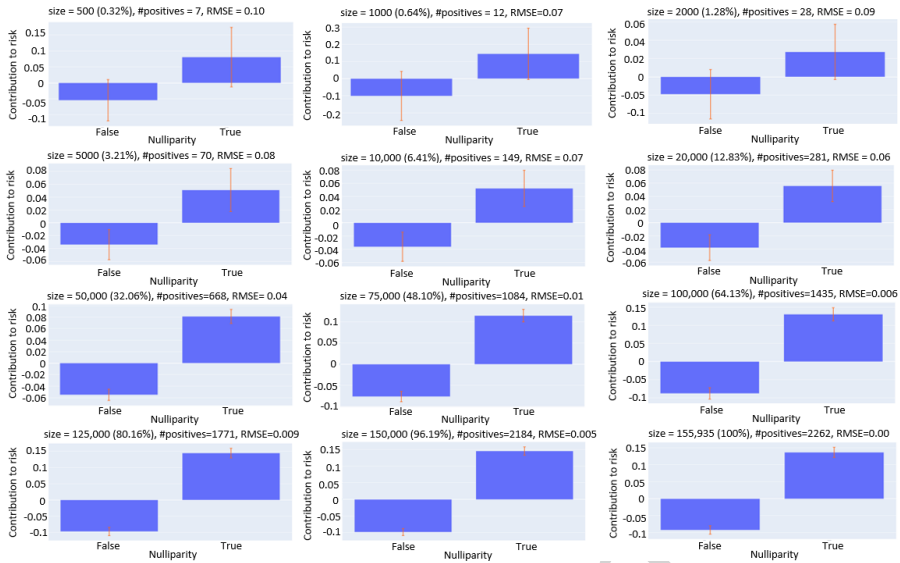


Fig. 11: Evolution of the shape functions of *nulliparity* while predicting SMM, as the data set size (“size” in the figure) is varied. We compute the RMSE by using the final shape function trained on all data as the reference.

to have positive and negative cases for each value seems to be variance reduction (smaller error bars). In each sequence of shape functions of categorical features, we compute the RMSE by using the final shape function trained on all data as the reference, and plot the normalized RMSEs in Figure 13.

5 Discussion

We perform external validation to confirm that the models generalize well to hospitals and settings different from the training set, and achieve good performance. In this external validation, we carefully control for hospital level of care to ensure fair representation of different subpopulations, to reduce potential bias in the model as much as possible, and to ensure the models attain good calibration on different demographic groups. Unfortunately, few models are externally validated, and failure to use external validation increases the risk that the model and results will not generalize across different populations. [18–20]. This strengthens our confidence in the models and their findings. However, it is noteworthy that our data set might not have an exhaustive list of features that represent risk factors, i.e. it might be missing known and potentially unknown risk factors as features, such as blood pressure, a baby’s ultrasound images, and so on. Secondly, we would like to reiterate that our data set contains data from 20 hospitals in Washington and that the diversity among our cohort might not be representative of the US population as a whole.

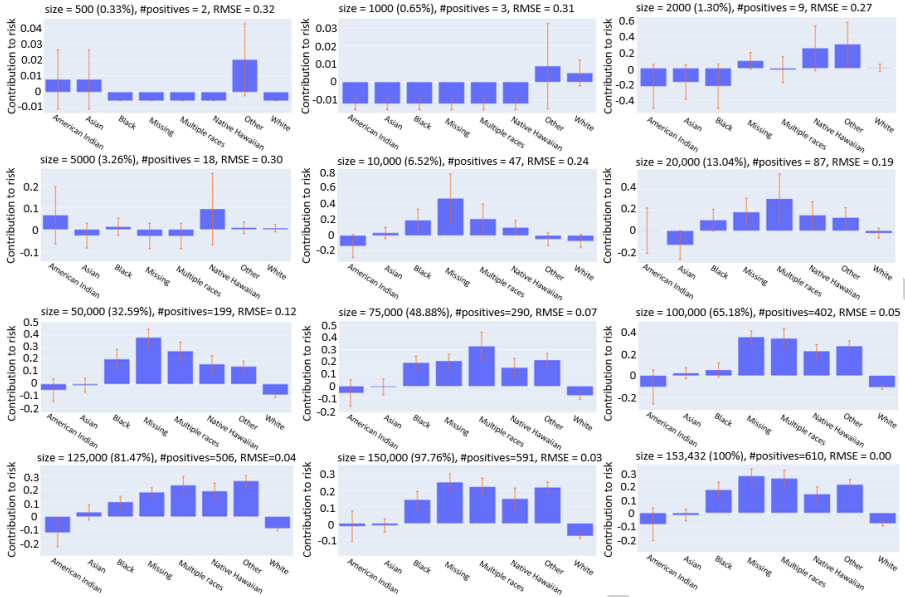


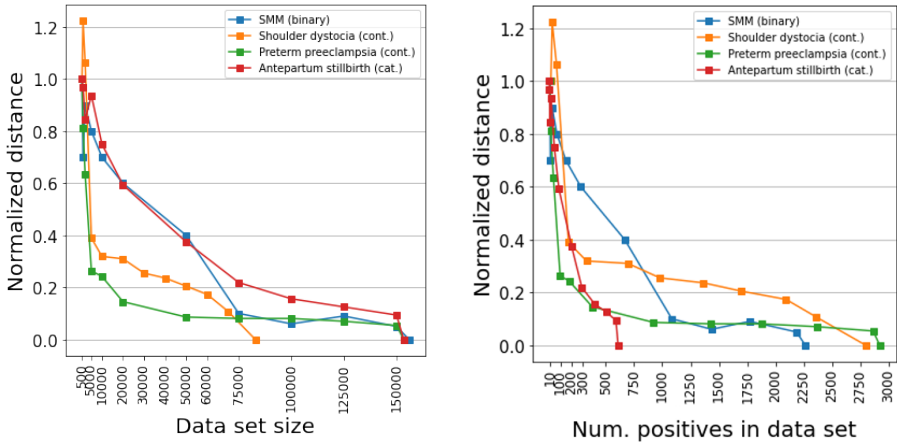
Fig. 12: This shows the development of the EBM’s shape function of *Race* when predicting antepartum stillbirth as the data set size (“size” in the figure) is varied.

As shown in Table 2 and Figure 7, EBMs are well-calibrated models with AUROCs as good as other models. The excellent calibration of the EBM models suggests they perform well across the entire spectrum of patient risk. Additionally, a key distinguishing factor between EBMs and other ML models is the EBM’s interpretability — the shape functions $f_i(x_i)$ in the GAM can easily be visualized, see Section 2.1. Importantly, EBMs find that the largest contributors to risk are not always the traditionally recognized ones. We now examine each of the four outcomes in more detail.

Shoulder dystocia

The main risk factors usually associated with shoulder dystocia are diabetes and baby birthweight. While the EBM reiterates the great influence of birthweight, it also shows that the time from admission to complete cervical dilation and the mother’s height are very important.

Figure 3a uses actual birthweight (as opposed to estimated birthweight) to highlight the importance of fetal weight for shoulder dystocia. This suggests the value of more research and improved prenatal estimation of fetal weight in predicting maternal and fetal outcomes. Additionally, the importance of maternal height to the risk of shoulder dystocia suggests the potential value of more accurate measurement of pelvic dimensions.



(a) Sequence of normalized distances as a function of the data set size.

(b) Sequence of normalized distances as a function of the number of positives in the sampled data set.

Fig. 13: We plot the sequence of normalized distances where we consider the discrete Fréchet distance for continuous features ('cont.') and RMSE for categorical features ('cat.' or 'binary'). By normalization we mean dividing the distances in each sequence by the first (usually largest) distance in the sequence, to better visualize how fast relative convergence occurs. By definition, the distances for the full sample size are zero.

SMM

The main indicators picked up by the EBM for SMM are preeclampsia, labor type, nulliparity, initial cervical dilation, and race. Maternal age appears to be not as important of a predictor, even though it's traditionally recognized as an important factor; see Figure 2b and Appendix A.1.

Preterm preeclampsia

The largest contributions to the risk of preterm preeclampsia are made by the mother's BMI, nulliparity and chronic hypertension. The mother's age also plays a major role, especially for women older than 44, where there is a step increase in risk, see Figure 4a.

Antepartum stillbirth

We find we can model antepartum stillbirth well, even in the face of a very low positive rate. Some of the most important features are BMI, alcohol use, and nicotine use, all of which are potentially modifiable for a subset of the population. This suggests steps could be taken to lower the average antepartum stillbirth risk nationwide. Other important features include race, age, height, and socioeconomic factors. Surprisingly, chronic hypertension was not one of

the more significant factors in predicting antepartum stillbirth, perhaps reflecting greater surveillance and earlier timing of delivery for patients identified as having chronic hypertension.

Race and ethnicity

We recognize that including race and ethnicity in clinical models is a complex issue. We propose one potential approach which is to include race and ethnicity in the model, let the interpretable model tell us what the statistics indicate about race and ethnicity, and then, if needed, mitigate bias by zeroing out the learned effects for these terms.

6 Conclusion

We train a variety of models (logistic regression, EBMs, Random Forests, Boosted Trees and Neural Nets) for each of four important pregnancy complications: SMM, shoulder dystocia, preterm preeclampsia and antepartum stillbirth. Importantly, external validation is used to evaluate all of the models. The results suggest that the interpretable EBM models have accuracy equal to or better than the other models; exhibit better calibration than the other models; and their interpretability allows us to discover surprising clinical effects that in some cases challenge traditional beliefs. An analysis of the convergence of the risk term profiles shows the sample sizes sufficient to train robust models. A key advantage of interpretable learning methods such as EBMs is that they go beyond predicting probabilities to also allow clinicians to understand what the model has learned about the clinical problem. Some of the effects observed in the models may provide opportunities for interventions to improve maternal and fetal outcomes.

Appendix A Lists of features

For each of the four outcomes we provide a full list of features used. The features are ranked according to their feature importance from highest (first feature) to lowest (last feature) as computed by the EBM. Feature importance is measured as the mean absolute contribution to the log odds of the risk prediction over all samples.

A.1 Severe maternal morbidity

(1) Preeclampsia/gestational hypertension, (2) labor type (spontaneous labor or induction of labor), (3) nulliparity, (4) initial cervical dilation, (5) race, (6) Distressed Communities Index quintile (Economic Innovation Group), (7) Scheduled cesarean (planned cesarean birth), (8) initial maternal BMI (either pre-pregnancy or at first prenatal visit), (9) labor allowed (vaginal birth attempted), (10) hour of admission, (11) Hispanic/Latina ethnicity, (12) gravidity (the total number of pregnancies the pregnant person has had including the current pregnancy), (13) number of previous stillbirths, (14) gestational

diabetes, (15) history of uterine surgery, (16) maternal age, (17) month of admission, (18) IVF, (19) final maternal BMI, (20) cervical ripening, (21) maternal height, (22) history of classical incision, (23) insurance type, (24) absence minimal prenatal care, (25) number of previous cesareans, (26) induced labor indication, (27) history of low vertical incision, (28) parity, (29) membrane status (membranes ruptured or intact at the time of admission to labor and delivery), (30) presentation at delivery, (31) placental abruption, (32) history of uterine rupture, (33) day of week at admission, (34) sex of baby, (35) illicit substance use during pregnancy, (36) pre-pregnancy diagnosis of mental illness, (37) Rural-Urban Commuting Area (RUCA) code class, (38) pre-pregnancy diagnosis of diabetes, (39) ASA recommended for use during pregnancy (documentatoin of low dose aspirin recommendation in the medical record), (40) pre-pregnancy use of nicotine, (41) pre-pregnancy use of marijuana, (42) use of nicotine during pregnancy, (43) pre-pregnancy illicit substance use, (44) chronic hypertension, (45) cholestasis of pregnancy, (46) use of marijuana during pregnancy, (47) pre-pregnancy use of alcohol, (48) use of alcohol during pregnancy, (49) number of previous preterm births.

A.2 Shoulder dystocia

(1) Birthweight, (2) maternal height, (3) time from admission to complete cervical dilation, (4) final maternal BMI, (5) oxytocin use, (6) race, (7) Distressed Communities Index quintile, (8) Insurance type, (9) initial cervical dilation, (10) use of regional anesthesia, (11) time from complete dilation till delivery, (12) maternal age, (13) nulliparity, (14) number of weeks on delivery, (15) gravidity, (16) initial maternal BMI, (17) indication for operative vaginal delivery, (18) IVF, (19) cervical effacement, (20) month of admission, (21) membrane status, (22) vacuum use, (23) cervical ripening, (24) hour of admission, (25) pre-pregnancy diagnosis of diabetes, (26) pre-pregnancy diagnosis of mental illness, (27) gestational diabetes, (28) pre-pregnancy illicit substance use, (29) Hispanic/Latina Ethnicity, (30) Forceps, (31) day of week at admission, (32) Rural-Urban Commuting Area (RUCA) code class, (33) indication for induced labor, (34) pre-pregnancy use of alcohol, (35) history of uterine rupture, (36) nicotine use during pregnancy, (37) sex of baby, (38) pre-pregnancy use of marijuana, (39) alcohol use during pregnancy, (40) labor type, (41) parity (number of previous births of ≥ 20 weeks' gestation), (42) cholestasis of pregnancy, (43) preeclampsia/gestational hypertension, (44) use of marijuana during pregnancy, (45) use of illicit substance during pregnancy, (46) absent of minimal prenatal care, (47) ASA recommended for use during pregnancy, (48) pre-pregnancy use of nicotine, (49) number of previous preterm births, (50) number of previous stillbirths, (51) history of uterine surgical history, (52) history of classical incision, (53) pre-pregnancy diagnosis of hypertension, (54) history of low vertical incision, (55) number of previous cesarean births, (56) placental abruption.

A.3 Preterm preeclampsia

(1) BMI (2) nulliparity, (3) chronic hypertension, (4) maternal age, (5) pre-pregnancy diagnosis of mental illness, (6) number of previous stillbirths, (7) gravidity, (8) race, (9) Distressed Communities Index quintile, (10) number of previous preterm births, (11) Rural-Urban Commuting Area (RUCA) code class, (12) pre-pregnancy diagnosis of diabetes, (13) maternal height, (14) Hispanic/Latina ethnicity, (15) illicit substance use during pregnancy, (16) previous low vertical incision, (17) (other) uterine surgical history, (18) number of previous cesarean births, (19) history of uterine rupture, (20) marijuana use during pregnancy (21) pre-pregnancy use of marijuana, (22) pre-pregnancy illicit substance use, (23) IVF, (24) sex of baby, (25) pre-pregnancy use of alcohol, (26) previous classical cesarean, (27) alcohol use during pregnancy, (28) nicotine use during pregnancy, (29) pre-pregnancy use of nicotine.

A.4 Antepartum stillbirth

(1) BMI, (2) race, (3) Distressed Communities Index quintile, (4) maternal age, (5) maternal height, (6) Rural-Urban Commuting Area (RUCA) code class, (7) gravidity, (8) pre-pregnancy diagnosis of diabetes, (9) number of previous cesarean births, (10) number of previous stillbirths, (11) pre-pregnancy use of alcohol, (12) pre-pregnancy use of nicotine, (13) IVF, (14) baby gender, (15) nulliparity, (16) history of classical incision, (17) history of uterine surgical history, (18) previous low vertical incision, (19) pre-pregnancy use of marijuana, (20) chronic hypertension, (21) nicotine use in pregnancy, (22) pre-pregnancy diagnosis of mental illness, (23) marijuana use during pregnancy, (24) illicit substance use during pregnancy, (25) Hispanic/Latina ethnicity, (26) number of previous preterm births, (27) pre-pregnancy illicit substance use, (28) alcohol use during pregnancy, (29) history of uterine rupture.

Declarations

Competing interests Author VS is a medical director at Natera. The other authors declare no competing interests.

Authors' contributions TMB, VS and RC developed and planned the experiments. TMB, ZH and HL carried out the experiments. BJL and HN helped design the framework underlying the experiments. IP and VS prepared and enabled access to the data. VS provided domain expertise. RC developed and supervised the project. All authors took part in writing the manuscript.

Funding Not applicable.

References

- [1] Osterman, M.J., Hamilton, B.E., Martin, J.A., Driscoll, A.K., Valenzuela, C.P.: Births: final data for 2020 (2022)

- [2] Declercq, E., Zephyrin, L.: Maternal mortality in the united states: a primer. Commonwealth Fund (2020)
- [3] Creanga, A.A., Bateman, B.T., Kuklina, E.V., Callaghan, W.M.: Racial and ethnic disparities in severe maternal morbidity: a multistate analysis, 2008-2010. *American journal of obstetrics and gynecology* **210**(5), 435–1 (2014)
- [4] Gao, C., Osmundson, S., Yan, X., Edwards, D.V., Malin, B.A., Chen, Y.: Learning to identify severe maternal morbidity from electronic health records. *Studies in health technology and informatics* **264**, 143 (2019)
- [5] Lengerich, B.J., Caruana, R., Weeks, W.B., Painter, I., Spencer, S., Sitcov, K., Daly, C., Souter, V.: 1017 insights into severe maternal morbidity in the nsv population. *American Journal of Obstetrics & Gynecology* **224**(2), 629–630 (2021)
- [6] Cartus, A.R., Naimi, A.I., Himes, K.P., Jarlenski, M., Parisi, S.M., Bodnar, L.M.: Can ensemble machine learning improve the accuracy of severe maternal morbidity screening in a perinatal database? *Epidemiology* **33**(1), 95–104 (2021)
- [7] Callaghan, W.M., MacKay, A.P., Berg, C.J.: Identification of severe maternal morbidity during delivery hospitalizations, united states, 1991-2003. *American journal of obstetrics and gynecology* **199**(2), 133–1 (2008)
- [8] Lengerich, B.J., Caruana, R., Weeks, W.B., Painter, I., Spencer, S., Sitcov, K., Daly, C., Souter, V.: 46 length of labor and severe maternal morbidity in the nsv population. *American Journal of Obstetrics & Gynecology* **224**(2), 33 (2021)
- [9] Bennett, R., Mulla, Z.D., Parikh, P., Hauspurg, A., Razzaghi, T.: An imbalance-aware deep neural network for early prediction of preeclampsia. *Plos one* **17**(4), 0266042 (2022)
- [10] Moreira, M.W., Rodrigues, J.J., Oliveira, A.M., Saleem, K., Neto, A.J.V.: Predicting hypertensive disorders in high-risk pregnancy using the random forest approach. In: 2017 IEEE International Conference on Communications (ICC), pp. 1–5 (2017). IEEE
- [11] Bosschieter, T.M., Xu, Z., Lan, H., Lengerich, B., Nori, H., Sitcov, K., Painter, I., Caruana, R., Souter, V.: Preterm preeclampsia prediction using intelligible machine learning. *American Journal of Obstetrics & Gynecology* **228**(1), 409 (2023)

- [12] Jhee, J.H., Lee, S., Park, Y., Lee, S.E., Kim, Y.A., Kang, S.-W., Kwon, J.-Y., Park, J.T.: Prediction model development of late-onset preeclampsia using machine learning-based methods. *PLoS One* **14**(8), 0221202 (2019)
- [13] Tsur, A., Batsry, L., Toussia-Cohen, S., Rosenstein, M., Barak, O., Brezinov, Y., Yoeli-Ullman, R., Sivan, E., Sirota, M., Druzin, M., *et al.*: Development and validation of a machine-learning model for prediction of shoulder dystocia. *Ultrasound in Obstetrics & Gynecology* **56**(4), 588–596 (2020)
- [14] Bartal, M.F., Tsur, A., Sibai, B.M., Aran, D., Bicocca, M.J., Chauhan, S.P.: 651 clinical efficacy of a machine learning model for prediction of shoulder dystocia. *American Journal of Obstetrics & Gynecology* **224**(2), 409 (2021)
- [15] Bosschieter, T.M., Lan, H., Xu, Z., Lengerich, B., Nori, H., Sitcov, K., Painter, I., Caruana, R., Souter, V.: Unique insights into risk factors for antepartum stillbirth using explainable ai. *American Journal of Obstetrics & Gynecology* **228**(1), 403–404 (2023)
- [16] Allotey, J., Whittle, R., Snell, K., Smuk, M., Townsend, R., von Dadelszen, P., Heazell, A., Magee, L., Smith, G., Sandall, J., *et al.*: External validation of prognostic models to predict stillbirth using international prediction of pregnancy complications (ippic) network database: individual participant data meta-analysis. *Ultrasound in Obstetrics & Gynecology* **59**(2), 209–219 (2022)
- [17] Malacova, E., Tippaya, S., Bailey, H.D., Chai, K., Farrant, B.M., Gebremedhin, A.T., Leonard, H., Marinovich, M.L., Nassar, N., Phatak, A., *et al.*: Stillbirth risk prediction using machine learning for a large cohort of births from western australia, 1980–2015. *Scientific reports* **10**(1), 1–8 (2020)
- [18] Steyerberg, E.W., Harrell, F.E.: Prediction models need appropriate internal, internal-external, and external validation. *Journal of clinical epidemiology* **69**, 245–247 (2016)
- [19] Wynants, L., Collins, G.S., Van Calster, B.: Key steps and common pitfalls in developing and validating risk models. *BJOG: An International Journal of Obstetrics & Gynaecology* **124**(3), 423–432 (2017)
- [20] Kleinrouweler, C.E., Cheong-See, F.M., Collins, G.S., Kwee, A., Thangaratinam, S., Khan, K.S., Mol, B.W.J., Pajkrt, E., Moons, K.G., Schuit, E.: Prognostic models in obstetrics: available, but far from applicable. *American journal of obstetrics and gynecology* **214**(1), 79–90 (2016)

24 *Interpretable Machine Learning for Pregnancy*

- [21] Squires, D.A.: The us health system in perspective: a comparison of twelve industrialized nations. Issue Brief (Commonwealth Fund) **16**, 1–14 (2011)
- [22] Yamada, G., Hayakawa, K., Asai, Y., Matsunaga, N., Ohtsu, H., Hojo, M., Hashimoto, M., Kobayashi, K., Sasaki, R., Okamoto, T., *et al.*: External validation and update of prediction models for unfavorable outcomes in hospitalized patients with covid-19 in japan. *Journal of Infection and Chemotherapy* **27**(7), 1043–1050 (2021)
- [23] Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining, pp. 785–794 (2016)
- [24] Breiman, L.: Random forests. *Machine learning* **45**(1), 5–32 (2001)
- [25] Gardner, M.W., Dorling, S.: Artificial neural networks (the multilayer perceptron)—a review of applications in the atmospheric sciences. *Atmospheric environment* **32**(14–15), 2627–2636 (1998)
- [26] Kleinbaum, David G and Dietz, K and Gail, M and Klein, Mitchel and Klein, Mitchell: *Logistic Regression*, p. 536. Springer, New York, NY (2002)
- [27] Qayyum, A., Qadir, J., Bilal, M., Al-Fuqaha, A.: Secure and robust machine learning for healthcare: A survey. *IEEE Reviews in Biomedical Engineering* **14**, 156–180 (2020)
- [28] Hastie, T., Tibshirani, R.: Generalized Additive Models. *Statistical Science* **1**(3), 297–310 (1986). <https://doi.org/10.1214/ss/1177013604>
- [29] Lou, Y., Caruana, R., Gehrke, J.: Intelligible models for classification and regression. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 150–158 (2012)
- [30] Nori, H., Jenkins, S., Koch, P., Caruana, R.: Interpretml: A unified framework for machine learning interpretability. arXiv preprint arXiv:1909.09223 (2019)
- [31] Lou, Y., Caruana, R., Gehrke, J., Hooker, G.: Accurate intelligible models with pairwise interactions. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 623–631 (2013)
- [32] Zhou, J., Tse, G., Lee, S., Liu, T., Wu, W.K., Cao, Z., Zeng, D.D., Wong, I.C.K., Zhang, Q., Cheung, B.M.Y.: Identifying main and interaction

- effects of risk factors to predict intensive care admission in patients hospitalized with covid-19: a retrospective cohort study in hong kong. medRxiv (2020)
- [33] Decroos, T., Davis, J.: Interpretable prediction of goals in soccer. In: Proceedings of the AAAI-20 Workshop on Artificial Intelligence in Team Sports (2019)
- [34] Xenopoulos, P., Freeman, W.R., Silva, C.: Analyzing the differences between professional and amateur esports through win probability. In: Proceedings of the ACM Web Conference 2022, pp. 3418–3427 (2022)
- [35] Maxwell, A.E., Sharma, M., Donaldson, K.A.: Explainable boosting machines for slope failure spatial predictive modeling. *Remote Sensing* **13**(24), 4991 (2021)
- [36] Cannarozzo, C.: The merger-driven evolution of early-type galaxies and the connection with their dark matter halos (2021)
- [37] Alt, H., Godau, M.: Computing the fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications* **5**(01n02), 75–91 (1995)
- [38] Foundation for Health Care Quality: Obstetrical Care Outcome Assessment Program (OBCOAP). <https://www.qualityhealth.org/obcoap/>
- [39] Group, E.I.: Introduction to the Distressed Communities Index (DCI). <https://eig.org/distressed-communities/> Accessed Jan. 28, 2023
- [40] Xu, Z., Bosschieter, T.M., Lan, H., Lengerich, B., Nori, H., Sitcov, K., Painter, I., Souter, V., Caruana, R.: Predicting severe maternal morbidity at admission for delivery using intelligible machine learning. *American Journal of Obstetrics & Gynecology* **228**(1), 404–405 (2023)
- [41] Lan, H., Bosschieter, T.M., Xu, Z., Lengerich, B., Nori, H., Sitcov, K., Painter, I., Souter, V., Caruana, R.: Understanding risk factors for shoulder dystocia using interpretable machine learning. *American Journal of Obstetrics & Gynecology* **228**(1), 753 (2023)
- [42] Economic Research Service, U.S.D.O.A.: Rural-Urban Commuting Area Codes. <https://www.ers.usda.gov/data-products/rural-urban-commuting-area-codes/> Accessed Jan. 28, 2023
- [43] Jones, C.P.: Invited commentary: “race,” racism, and the practice of epidemiology. *American journal of epidemiology* **154**(4), 299–304 (2001)
- [44] Bedoya, A.D., Economou-Zavlanos, N.J., Goldstein, B.A., Young, A.,

Jelovsek, J.E., O'Brien, C., Parrish, A.B., Elengold, S., Lytle, K., Balu, S., *et al.*: A framework for the oversight and local deployment of safe and high-quality prediction models. *Journal of the American Medical Informatics Association* **29**(9), 1631–1636 (2022)

- [45] Paulus, J.K., Kent, D.M.: Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *NPJ digital medicine* **3**(1), 99 (2020)
- [46] Chang, C.-H., Tan, S., Lengerich, B., Goldenberg, A., Caruana, R.: How interpretable and trustworthy are gams? In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 95–105 (2021)