DEPLOYMENT OF A ROBUST AND EXPLAINABLE MORTALITY PREDICTION MODEL: THE COVID-19 PANDEMIC AND BEYOND

Jacob R. Epifano, Stephen Glass, Ravi P. Ramachandran
Department of Electrical and Computer Engineering
Rowan Univsercity
Glassboro NJ, USA
epifanoj0@rowan.edu sglass520@gmail.com, ravi@rowan.edu

Sharad Patel
Department of Critical Care Medicine
Cooper University Hospital
Camden NJ, USA
patel-sharad@cooperhealth.edu

Aaron J. Masino
Department of Biostatistics, Epidemiology, Informatics
University of Pennsylvania Perelman School of Medicine
Philadelphia PA, USA
aaron.masino@pennmedicine.upenn.edu

Ghulam Rasool
Department of Machine Learning
Moffitt Cancer Center
Tampa FL, USA
ghulam.rasool@moffitt.org

ABSTRACT

This study investigated the performance, explainability, and robustness of deployed artificial intelligence (AI) models in predicting mortality during the COVID-19 pandemic and beyond. The first study of its kind, we found that Bayesian Neural Networks (BNNs) and intelligent training techniques allowed our models to maintain performance amidst significant data shifts. Our results emphasize the importance of developing robust AI models capable of matching or surpassing clinician predictions, even under challenging conditions. Our exploration of model explainability revealed that stochastic models generate more diverse and personalized explanations thereby highlighting the need for AI models that provide detailed and individualized insights in real-world clinical settings. Furthermore, we underscored the importance of quantifying uncertainty in AI models which enables clinicians to make better-informed decisions based on reliable predictions. Our study advocates for prioritizing implementation science in AI research for healthcare and ensuring that AI solutions are practical, beneficial, and sustainable in real-world clinical environments. By addressing unique challenges and complexities in healthcare settings, researchers can develop AI models that effectively improve clinical practice and patient outcomes.

*K*eywords Explainable AI · Bayesian Neural Networks · Influence Functions · Dataset-drift · Uncertainty Quantification · COVID-19

1 Introduction

Artificial Intelligence (AI) and Machine Learning (ML) have started to gain traction in medical domains [1]. However, due to the high-stakes nature of medicine, AI/ML models rarely make it to real-world deployments [2]. While a few models have successfully transitioned to deployment [3, 4, 5], user feedback is often negative [6, 7]. Factors such as high false alarm rates and uninterpretable outputs [8, 9] demand several considerations for AI/ML to flourish in high-risk environments like healthcare.

The challenges of AI/ML in healthcare are well-recognized in the field [10]. Consequently, substantial efforts have been made to distill missing elements into a few key requirements. Many emphasize the importance of interpretability and explainability [11, 12, 13]. Interpretability is concerned with understanding the inner workings of a model and is frequently associated with transparent design [10]. Explainability concentrates on the comprehension of how and why a model decision was arrived at. However, with the advent of deep learning, ML practitioners are increasingly trading the transparent design of linear models, whose decisions are easily understood (e.g., logistic regression, decision trees),
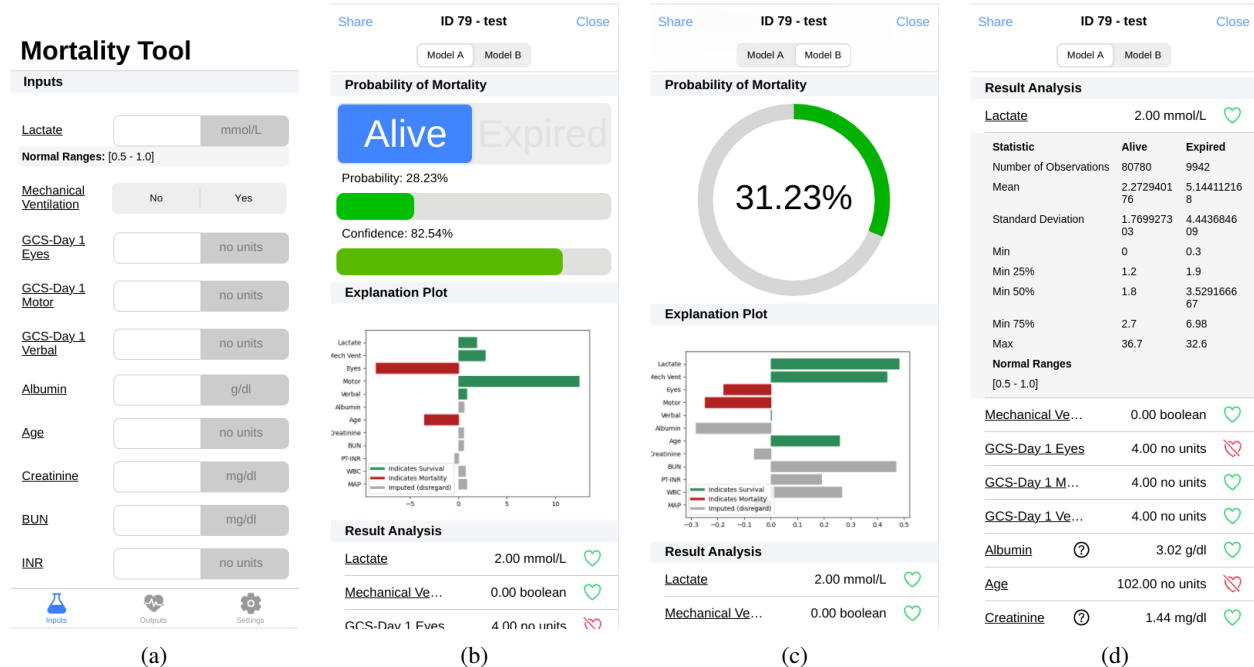
Figure 1: Images of web-app user interface. **1a:** input screen. All variables can be selected and the healthy range for each variable will appear in a drop-down menu. This feature was requested by our users. **1b**: output of model A (VDP). This output pane includes the probability of mortality as well as the confidence score derived from the variance output of the stochastic network. **1c:** output of model B (deterministic/traditional neural net). Below the output probabilities for each model, an explanation plot derived from influence functions is provided. For each input feature, the influence function provides an indication of whether or not the signal contributed towards one class or the other (indicated by the sign) and the magnitude determines how much each signal contributed to the classification. **1d:** dataset statistics. For our users who want a more fine grained understanding of the classification decision, the output screen contains a drop down menu where it shows statistics from the training set. This may provide the user with a better understanding of how the model came to its prediction.

for the better performance of black-box models (e.g., Deep Neural Networks (DNNs)) [14]. In a black-box model, the inputs and outputs are observable. However, the hidden nature of the algorithmic operations gives no idea of how the outputs are arrived at. As a result, a significant body of work has been dedicated to explaining the predictions of black-box DNN models. In particular, several notable methods exist for explaining the predictions of DNNs due to their popularity. Nonetheless, there is a growing concern regarding how accurate an explainability method is (known as faithfulness [15]) as these methods rely on creating local interpretable models for each prediction [16].

Measuring a model's confidence or its uncertainty is equally important in medical AI [11, 10]. DNNs have been shown to be typically overconfident in their predictions [17, 18], i.e., high softmax output does not equate to high prediction confidence [19]. This becomes crucial when clinicians must take action based on a prediction, as their subsequent actions may vary depending on their perception of the situation's severity [20]. Overconfident models lead to high false-alarm rates, which undermine overall trust in the AI model [21].

Lastly, some studies highlight dataset-shift as a crucial factor to consider when designing AI/ML models for healthcare [22]. Dataset-shift occurs most commonly when the model is deployed in a data or environment that is different from the one used during training. A robust model is insensitive to dataset-shift. Dataset-shift is often overlooked during the design and development process or is left for manual review and update. It is naive to assume that the independent and identically distribution (i.i.d.) assumption of many AI/ML algorithms will hold in a field like healthcare, where different patient populations can exhibit vastly different distributions [23].

Algorithmic bias is another related and critical concern in healthcare. Algorithmic bias exists when a repeatable error leads to an unfair outcome. It has been demonstrated that mortality prediction models' performance can vary depending on patient ethnicity [24]. There are common issues in the mitigation of dataset-shift and algorithmic bias. The two collectively address the full spectrum of failure of black-box AI models.
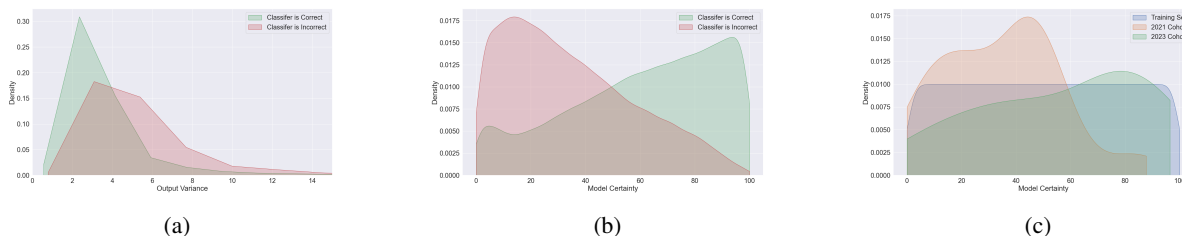
Figure 2: **2a:** Output variance and **2b:** Uncertainty scores computed across training set by estimating the Cummulative Density Function (CDF) of the joint distribution. On average, the variance is lower (model is more confident) when the prediction is correct. This indicates that the Bayesian model is well calibrated to the dataset. **2c:** Uncertainty scores broken down by cohort. COVID-19 afflicted cohort (2021) has lower certainty than unafflicted cohort (2023). The 2023 cohort seemingly follows the training distribution.

Our contributions include:

1. Real world deployment of a first-day mortality prediction model designed with the help of critical care clinicians.

2. Usage of a noise-robust Bayesian Neural Network to quantify prediction uncertainty via covariance propagation.

3. Implementation and deployment of Influence Functions [25] to provide instance-level explanations, a trait highly desired by clinicians [10].

4. A post-hoc analysis of model performance during significant data-shift induced by the COVID-19 pandemic thereby highlighting the robust behavior of our models.

2 Related Work

2.1 Explainable AI

The performance of healthcare AI/ML models has started to become a secondary priority to the interpretability of the model. A tool that only outputs a probability score for the presence of a disease is not one that would likely be endorsed by clinical staff. Several approaches have become mainstream and have made their way into clinical workflows [6, 7]. The most popular class of explainable methods are feature attribution methods or methods which try to attribute the prediction to one or several input features. Influence functions [25] have been used in mortality prediction [26]. Also, SHapley Additive exPlanations (SHAP) values [27] were used in drug discovery [28] and first-day mortality prediction [29, 26]. The power behind feature attribution methods is that they can often be used to estimate the global feature importance as well as give local feature importance for a given instance. These two elements were called out specifically by clinician surveys as being highly desirable for a given clinical model [10].

2.2 Uncertainty Estimation

When discussing uncertainty, it is important to distinguish between the two sub-types. The first is aleatoric uncertainty, which refers to the uncertainty of the data or the environment. The second is epistemic uncertainty, which refers to the uncertainty of the parameters of the model [30]. In a clinical setting, it is becoming highly desirable to include an estimate of the uncertainty of the prediction [11, 10]. This is typically done via ensemble approaches like Deep ensemble averaging [31] or probabilistic approaches like Bayesian Neural Nets (BNNs) [32, 33]. Baking model uncertainty into models comes with benefits. Techniques like stochastic weight averaging [34] and Varitational Density Propagation (VDP) [33] boast better generalization and robustness to noise and artifacts. To our knowledge, models equipped with uncertainty estimation have yet to make it to production settings [35].

2.3 Dataset-shift and algorithmic bias

One obvious example of dataset-shift in healthcare came with the advent of COVID-19. A particular issue is mortality prediction in the Intensive Care Unit (ICU). During several periods of the pandemic, ICU mortality rates spiked [36]. Several biomarkers were found to be better predictors of mortality in COVID-19 patients than standard AI/ML models

[37]. This same model was replicated and tried in the Netherlands and it was found that the mortality prediction accuracy was much lower [38]. This case study shows that even when a tool is tailor made to serve a specific subset of the population, it may only work well in a specific region. When developing these models, it can become easy to fool oneself about the performance of a model. Several strategies should be employed once the model is deployed. These include periodical testing using data from the deployment, model fine-tuning, and full model retraining [39].

3 Methods

3.1 Dataset and Training Scheme

To train our model, we aggregated mortality data from two publicly available datasets: MIMIC-III [40] and the eICU [41] databases. From these databases, we selected data such as demographic information (age, gender), indicator variables (patient on mechanical ventilation, metastatic cancer, etc.) and lab values (blood lactate, glucose, creatinine, etc.) that are given within 24 hours of patient admittance. This results in nearly 200,000 instances and 200 features with a 92%-8% class imbalance where patient mortality is the positive class. In light of limited sample size, to reduce the likelihood of model overfitting, we adopted an aggressive feature selection procedure. First, pair-wise correlations were computed and all features with Pearson correlation coefficient greater than 0.9 were dropped from the dataset. Next, any features with greater than 50% missingness were removed from the dataset. We then computed the mutual information between the remaining features and the mortality outcome and selected the top 20. Out of the top 20, a group of clinicians chose 12 based on clinical opinion. We opted to not take a purely data science feature selection technique since some of the top features chosen did not give clinicians any actionable insights. Additionally by taking user opinion into account during the model selection stage, we hoped to foster stronger understanding between model and user prior to deployment. Our final set of features is as follows: Blood lactate, mechanical ventilation (yes/no), all Glasgow Coma Scale parameters (eyes, motor, verbal), albumin, age, creatinine, Prothrombin Time (PT-INR), White Blood cell Count (WBC), Blood Urea Nitrogen (BUN), and Mean Arterial Pressure (MAP).

Model training was performed in a similar manner to [26] where first-day mortality prediction was investigated. First-day prediction of mortality is a useful indicator for clinicians in order to direct care, although trained models may be used any day after admittance. Based on input of our clinician collaborators, missing data was assumed to be missing at random. We used multiple imputation with chained equations [42] as the imputation method. Prior to training, the data was standardized. The learning algorithm chosen here is a shallow neural network. To optimize parameters and capture performance metrics, we used a 10-fold cross validation scheme to split our data into train and test sets. We used Weights and Biases [43] to optimize the hyper parameters of our neural network. After discussing with our clinician collaborators, we chose to optimize for positive likelihood ratio (LR+ = sensitivity / (100 – specificity)) as it is a metric commonly used for diagnostic tests and therefore readily understood by our users [44]. Given the large class imbalance, various options to rectify the imbalance were given to the hyperparameter optimizer: Synthetic Minority Oversampling (SMOTE) [45, 46], majority undersampling [46], and loss function weighting (penalizing minority class predictions). After the hyperparameters are selected, we trained the models using both splits. Once deployed, we did not perform any retraining or online leaning.

3.2 Variational Density Propagation (VDP)

In addition to our traditional or deterministic neural net, we also trained a Bayesian Neural Network (BNN) in the same manner. We decided to include this model in order to investigate whether or not using robust models will aid performance in non-stationary environments. We will refer to this as the stochastic neural net. We used Variational Density Propagation (VDP) [33, 47, 48], which is based on Tensor Normal Distributions [49] and variational inference. This method has several advantages over other BNNs such as Bayes By Backprop [32] due to the propagation of the first two moments (mean and covariance) through the network. These advantage are: noise robustness, model averaging, and Uncertainty Quantification (UQ). By including this model we can satisfy the UQ constraint outlined by the clinician survey [10] without relying on an external UQ metrics that may not be faithful to the original model. A full derivation of the VDP framework is provided in A.

To quantify the uncertainty of our prediction, we transform the output of the network. The output of the model contains the mean and covariance. The variance (diagonal elements of the covariance) contain the information about the uncertainty of the prediction. Since this value ranges from $[0, \infty)$, it can be difficult to understand the uncertainty of any one given prediction. To demonstrate how the variance can be used to estimate uncertainty, we have have plotted the distribution of variance outputs (from our training set) by whether or not the classifier correctly classified the instance. We should expect lower variance output when the classifier is correct (less noise in the prediction) than when it is incorrect. Figure 2a shows these distributions. In order to make this metric understandable to our users, we estimate the

Model	ROC AUC	Number of Input Features
[50]	0.83	15
[51]	0.84	37
[26]	0.85	20
[29]	0.86±(0.02)	10
[52]	0.90±(0.00)	20
Determistic (Ours)	0.89±(0.00)	12
Stochastic (Ours)	0.88±(0.00)	12

Table 1: 10-fold cross validated training results for the traditional (deterministic) and Bayesian [33] (stochastic) neural nets compared to several baseline models. All models were designed to predicted first-day mortality and were trained and evaluated on the eICU dataset. The mean score across the 10 runs is shown with the standard error.

CDF of the empirical variance distribution of our training set, $1 - \text{CDF}(\sigma^2)$. This metric will be bounded in the range $[0, 1]$, where 0 indicates low certainty and 1 indicates high certainty (Figure 2b).

3.3 Explanations and Interpretability

To satisfy the interpretability constraints requested by clinicians [10], we provide instance-level explanations via influence functions for each prediction. This method was previously thought to be fragile [53], but has been recently shown to work for small networks [54]. LIME [55] and SHAP [27] are the most commonly used methods to explain black-box models, however they rely on creating a surrogate model and don't provide guarantees that their explanations will be faithful to the original model [56, 57, 58]. Another major weakness of other methods like SHAP, is that they fail when multicollinearity exists in the dataset used to train the model. We solve this problem with our aggressive feature selection scheme which firsts removes highly correlated features. Feature attribution-based explanations are still useful after our feature selection scheme as our final feature set consists of signals that are orthogonal along the treatment space. Determining the most impactful signals can lead to actionable insights for our users.

Consider a standard classification problem where a label y is predicted for each feature vector x. Let $z_i = (x_i, y_i)$, where $i = 1, 2, ..., N$, for N instances in the dataset. It is assumed that we have a trained model where θ represents the trained network parameters. Our loss function can be written as $L(z, \theta) = \sum_{i=1}^{N} L(z_i, \theta)$. Our optimal model parameters are the set of parameters that minimize the loss: $\hat{\theta} = \arg\min_{\theta \in \Theta} \sum_{i=1}^{N} L(z_i, \theta)$ [25]. Recall, from Koh and Liang [25], influence functions estimate the parameter change, $\hat{\theta}_{-z} - \hat{\theta}$ due to removing a training instance z. The estimated parameter change is given in Equation 1, where $H_{\hat{\theta}} = \frac{1}{n} \sum_{i=1}^{n} \nabla_\theta^2 L(z_i, \hat{\theta})$ is the Hessian [25]. Using the chain rule, we can estimate the change in loss on a test instance, z_{test} using Equation 2 [25]. Finally, by perturbing the input features of the training instance to be removed, we can estimate the effect of the perturbation on the loss of the test instance, given by Equation 3 [25].

$$\mathcal{I}_{\text{up,params}}(z) = -H_{\hat{\theta}}^{-1} \nabla_\theta L(z, \hat{\theta}), \tag{1}$$

$$\mathcal{I}_{\text{up,loss}}(z, z_{\text{test}}) = -\nabla_\theta L(z_{\text{test}}, \hat{\theta})^\top H_{\hat{\theta}}^{-1} \nabla_\theta L(z, \hat{\theta}) \tag{2}$$

$$\mathcal{I}_{\text{pert,loss}}(z, z_{\text{test}})^\top = -\nabla_\theta L(z_{\text{test}}, \theta)^\top H_\theta^{-1} \nabla_x \nabla_\theta L(z, \theta). \tag{3}$$

By aggregating Equation 3 over the training set, we can find the average effect that perturbing a single input feature has on loss of the test instance. This local feature importance (FI) score is given in Equation 4 [26]. Explanations generated by influence functions are ideal for this application as the loss difference provides a magnitude and direction indicating feature importance and sentiment, giving insight into which class (in the binary case) the model believes the feature is attributed to.

$$FI_{\text{local}} = \frac{1}{N} \sum_{i=1}^{N} I_{\text{pert,loss}}(z_i, z_{\text{test}}). \tag{4}$$

In addition to the explanations, we have taken steps to make the model more interpretable using the app User Interface (UI). As suggested by clinicians during development, normative ranges for each feature are provided via drop down on

the input screen as shown in Figure 1a. We have also computed some first order statistics for each feature in our dataset in an effort to foster understanding between the model and users. These statistics are provided via drop down on the output screen as shown in Figure 1d.

3.4 App Design for Android/Apple

The design of the app can be split into three components: front-end client, API server, and ML background worker. The parts are deployed in a monolithic architecture on a Google Compute Engine virtual machine (VM) with a graphical processing unit (GPU). Ionic Framework and React library using TypeScript language was used to develop the front-end UI as a Progressive Web Application. The application can be accessed by a publicly accessible domain on a web browser and optionally installed on the device. With communication to the API server, the UI allows users to create accounts, input data using the interface, and view or update previous records with retrospective analysis. Images of the app's UI can be found in Figure 1.

4 Results

4.1 Model Selection

Upon the conclusion of model selection, the optimal parameters for each model were as follows. For the deterministic network, the model had three hidden layers with widths 197, 198, and 112 nodes, respectively. The loss function was weighted to account for the large class imbalance with a positive weight of 14.80. The model was trained for 127 epochs of full-batch gradient descent using stochastic gradient descent (SGD) with parameters: learning rate $\gamma = 0.03104$, weight-decay $\lambda = 0.0104$, and Nesterov momentum $\mu = 0.4204$.

The stochastic model also had three hidden layers with widths 31, 93, and 94 respectively. The class imbalance was handled by majority undersampling. The model was trained for 18 epochs with a batch-size of 1000 using mini-batch SGD with parameters: learning rate $\gamma = 0.0022$, weight-decay $\lambda = 0.0064$, and Nesterov momentum $\mu = 0.7589$. We performed a t-test on the test statistics of the two models and no significant difference was found ($p > 0.01$).

Table 1 compares our cross-validated model performance against other first-day mortality prediction models. The metric typically used to compare mortality prediction models is the Receiver Operating Characteristic Area Under the Curve (ROC AUC) [51]. Both of our models achieve a significantly higher ROC AUC than readily available (typically via online calculator) scores (APACHE IV [51], SAPS [50]), as well as a recently published Neural Network model [26] and ensemble model [29] ($p < 0.01$ for all pair-wise t-tests). Our models fall behind a recent Elastic-Net model [52], however more input features were required to achieve this performance. Additional statistics and comparisons of our two models on the union of the eICU and MIMIC-III databases are included in Appendix B.

4.2 Deployment and data collection

Model deployment commenced in January 2021, during which data from two distinct cohorts were gathered. The first cohort consisted of 59 subjects, with data collected between January 2021 and November 2021 (COVID-19 surge). The second cohort included 25 subjects, and their data were obtained from January 2023 to May 2023 (Post COVID-19). In both cohorts combined, a total of 43 subjects had their mortality outcomes documented (21 from 2021 and 22 from 2023). This was accompanied by predictions from a critical care clinician prior to viewing model predictions and true outcomes. We chose to collect data in this way in order to study the effects that dataset drift may have on our models due to the COVID-19 pandemic.

4.3 Dataset drift and model performance

We observed considerable drift in the label distribution between cohorts. To recall, the label distribution of our training set consisted of 92% for negative classes and 8% for positive classes. In contrast, the 2021 and 2023 cohorts exhibited label distributions of 20%-80% and 64%-36%, respectively. We used a one-way Chi-square test to compare the relative frequencies of mortality outcomes in each cohort against the training distribution. The results revealed that for the 2021 cohort, $p < 0.01$, and for the 2023 cohort, $p > 0.01$.

To investigate the potential impact of dataset drift on performance, we utilized a 2-sample Kolmogorov-Smirnov test with Bonferroni correction to compare distributions. For all continuous features in our feature set, we assessed the distribution split by cohort relative to the training set distribution. Blood lactate and blood albumin were identified as significant ($p < 0.01$ for the 2021 cohort and $p > 0.01$ for the 2023 cohort), implying that the distribution of these

Classifier	Clinician	Deterministic Network		Stochastic Network	
Datatset	Accuracy	ROC AUC	Accuracy	ROC AUC	Accuracy
eICU ∪ MIMIC-III	-	0.87±(0.00)	0.77±(0.00)	0.87±(0.00)	0.77±(0.00)
2021 Cohort	0.95	**0.99**	**0.90**	0.99	**0.95**
2023 Cohort	0.83	**0.89**	**0.75**	0.88	**0.83**
Combined	0.90	**0.93**	**0.84**	0.90	**0.90**

Table 2: Performance metrics for each model by cohort. Clinician prediction was collected at assessment prior to model inference. Model performance can be compared directly to clinicians using accuracy. We found that the stochastic model achieved clinician level performance across cohorts while the deterministic model fell behind even though the ROC AUC performance was higher across cohorts
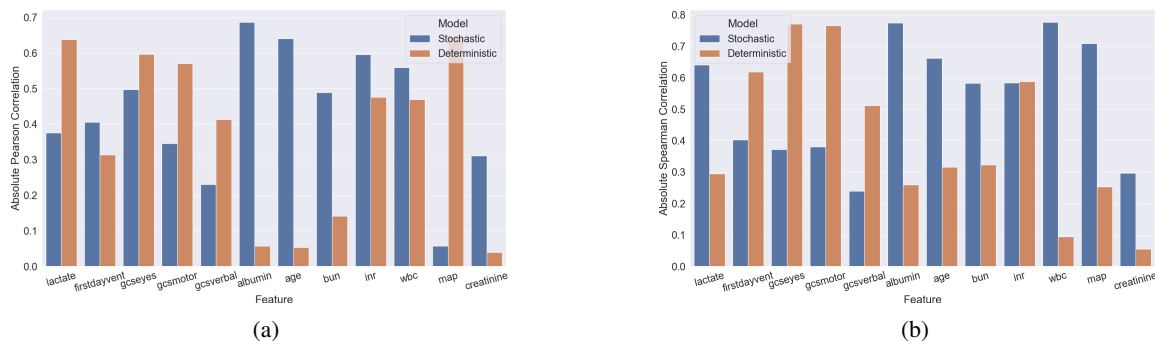

(a)   (b)

Figure 3: To validate influence function performance we adopted a validation scheme used by [25] and computed the **3a:** absolute Pearson correlation and **3b:** absolute Spearman between pair-wise feature differences and pair-wise influence differences. 500 data pairs were randomly sampled from our deployment dataset. The explanations boast moderate to high correlation coefficients ($p < 0.01$) for both models and metrics across most input features demonstrating the effectiveness of the method

variables significantly deviated from the training distribution for the 2021 cohort but not for the 2023 cohort. Kernel density estimates of these two features, separated by cohort are included in the supplementary material.

Turning our attention to model performance, we devoted substantial effort to optimizing the ROC AUC and LR+ during model selection, as these are conventional benchmarks for this problem due to the prevalent class imbalance. Since we also gather clinician predictions during data input, we can directly compare our models to clinicians using accuracy as a metric. Table 2 shows these results. We found that the stochastic model had achieved clinician level performance while the deterministic model fell behind across all cohorts even though the deterministic model had higher ROC AUC during model selection and across cohorts.

4.4 Explainability

To assess the explanations generated by both models, we performed a correlation analysis similar to that which was done in [25]. We computed the correlation between pair-wise feature differences and pair-wise influence differences from our data collected from deployment by randomly sampling subjects and subject pairs without duplicate (N=500). Figure 3 shows the absolute Pearson (Figure 3a) and absolute Spearman (Figure 3b) correlations. All correlations above 0.05 were significant ($p < 0.01$) with bonferonni correction.

To better understand the kinds of explanations generated by each model, we examined the top three features by absolute magnitude for each subject and compared their frequency across both cohorts. The stochastic model displayed greater variability in its explanations, with nine features appearing in the top three positions for both cohorts. Figure 3 shows the distribution of attribution scores for each feature by cohort. In contrast, the deterministic model only had six features in the top three positions. The stochastic model included lactate and albumin in it's top feature sets for both cohorts, which if you recall, were subject to significant distribution shift and therefore should be important.

In addition to importance scores, the influence function output is signed which can be used to evaluate class sentiment. To evaluate the overall sentiment associated with each feature, we applied the signum function to them. A highly

positive/negative score would suggest that the model consistently attributes importance to one class, while a score closer to zero would indicate that the model is equally likely attribute importance to either class. Our findings revealed that the stochastic model had more values near zero compared to the deterministic model. The deterministic model exhibited mostly very positive or very negative values, indicating limited inter-subject variability.

4.5 Un(Certainty)

To assess the effectiveness of our uncertainty metric, we conducted another 2-sample Kolmogorov-Smirnov test, comparing the "Confidence" metric reported by our stochastic model. This test revealed significant ($p < 0.01$) differences between the training set and the 2021 cohort, while no significant ($p > 0.01$) differences were observed between the training set and the 2023 cohort. Figure 2 presents the kernel density estimates of the training distribution, first broken down by classifier correctness (Figure 2b), and then by cohort (Figure 2c).

5 Discussion

5.1 Robust models and the effect of COVID-19

Given the substantial label shift in the data, we should have expected a sharp decline in performance. We attribute the robust performance of our models, in part, to the class-weighting/resampling approach we took during model selection. By resampling or weighting the loss function (instead of changing the decision threshold), we blind our model to the inherent imbalance of the training dataset. Since, in both cohorts, the label shift actually makes the classes more balanced than the training set, we see an increase in performance.

The explanations provided by influence functions boast moderate to high correlation coefficients for both models across most input features demonstrating the effectiveness of the method. In particular, the explanations generated by the stochastic model displayed increased variability over the deterministic model, in both sentiment and magnitude, that may better capture the unique characteristics of different patients, potentially leading to more accurate and personalized predictions. This difference in behavior is consistent with literature on evaluating explanations produced by BNNs [15]. Future research should continue to explore and develop AI/ML models that can offer more detailed and individualized insights, ultimately enhancing their utility and applicability in real-world clinical scenarios.

Figure 2 offers strong evidence supporting the robustness of BNNs. Our statistical results indicate that the "Confidence" metric may be used to detect dataset shift in small populations. The ability to effectively quantify uncertainty is crucial for AI/ML models in healthcare, as it allows clinicians to better understand the reliability of a given prediction and make more informed decisions based on the model's outputs. The visualizations in Figure 2 demonstrate the utility of our uncertainty metric in identifying differences between various cohorts and the training set, providing valuable insights into the performance of our stochastic model across different conditions. As such, future research should continue to explore and refine methods for estimating uncertainty in AI/ML models, ensuring that these tools can provide healthcare professionals with accurate, reliable, and interpretable information to support their decision-making processes.

These findings underscore the importance of considering the deployment environment when developing AI/ML models. By utilizing stochastic models like BNNs that enhance the robustness of AI models, researchers can develop AI solutions that maintain their performance and reliability even in the face of changing conditions. This ultimately improves their utility and applicability in real-world clinical settings.

5.2 Metric chasing

Even though our models performed well, this was not the main focus of the effort. The prevailing trend in the majority of AI/ML research for healthcare involves developing models to predict various clinical endpoints to demonstrate incremental advancements in key performance metrics. However, only a few extend beyond these improvements to address other essential aspects of AI in healthcare. We firmly believe that future efforts should prioritize implementation science over merely enhancing the performance of predictive models.

Implementation science plays a crucial role in ensuring the successful deployment, adoption, and sustainability of AI systems in real-world healthcare settings. By focusing on this aspect, researchers can better understand the factors that facilitate the integration of evidence-based interventions and strategies into routine clinical practice. This, in turn, can lead to the development of AI solutions that not only excel in their predictive capabilities but also effectively address the unique challenges and complexities of healthcare environments.

Moreover, emphasizing implementation science can promote interdisciplinary collaboration, user-centered design, and stakeholder engagement, all of which contribute to the practical utility and long-term success of AI solutions in

healthcare. By adopting a more comprehensive approach that encompasses these aspects, researchers can help bridge the gap between theoretical advancements and tangible improvements in clinical practice and patient outcomes.

5.3 Limitations

The most significant limitation of our study was the lack of integration into clinicians' workflow. At the outset, we underestimated the time and effort required for ICU clinicians to interrupt their tasks and enter 12 values on their phones. This challenge was further intensified by the COVID-19 pandemic, as clinicians faced longer working hours and managed more critical patients. If Electronic Health Record (EHR) integration had been possible, we estimate that we could have collected a considerably larger volume of data, potentially an order of magnitude greater than our current dataset.

While obtaining more data could have strengthened our results, we believe that it would not have fundamentally altered our findings. Given that our model was trained on 200,000 instances, we anticipate that our models demonstrate satisfactory calibration. Furthermore, the statistical testing we conducted produced positive results that corroborated our hypotheses.

In light of these limitations, future research should prioritize seamless integration into clinical workflows to facilitate data collection and ensure that AI models are practical and beneficial for healthcare professionals. Additionally, researchers should continue to explore methods for improving model robustness and generalizability, ensuring that AI solutions remain effective across diverse patient populations and varying medical conditions.

6 Conclusion

In this study, we examined the performance, explainability, and robustness of AI models for predicting clinical endpoints in the context of the COVID-19 pandemic. To our knowledge, this is the first study of its kind to deploy interpretable and explainable models with uncertainty prediction in a real-world setting. Our findings demonstrate the effectiveness of using BNNs to improve robustness, explanations and to quantify prediction uncertainty. By addressing these challenges, we were able to develop AI/ML models that matched or outperformed clinician predictions even when confronted with substantial dataset drift. Future research should continue to refine these methods to ensure these tools provide accurate, reliable, and interpretable information for healthcare professionals.

A Variational Density Propagation (VDP) Derivation

To describe the VDP method by Dera *et al.* [33], let us use a 2-layer fully connected neural network as an example. In a traditional or deterministic neural network, Equation 5 describes the forward pass of the model.

$$
\begin{aligned}
z &= Wx + b^{[1]}, \\
a &= f(z), \\
\hat{y} &= g(Va + b^{[2]}),
\end{aligned}
\tag{5}
$$

where $W \in \mathbb{R}^{j \times k}$ is the weight matrix of layer 1, $x \in \mathbb{R}^{k \times 1}$ is the input vector, $b^{[1]} \in \mathbb{R}^{j \times 1}$ is the bias vector in layer-1, $z \in \mathbb{R}^{j \times 1}$ is the result of the linear operation in layer-1, f is an arbitrary element-wise non-linear function, g is an arbitrary non-linear activation function that does not operate element-wise, $a \in \mathbb{R}^{j \times 1}$ is the result after applying the non-linear activation function in layer 1, $V \in \mathbb{R}^{l \times j}$ is the weight matrix in layer 2, $b^{[2]} \in \mathbb{R}^{j \times 1}$ is the bias vector in layer 2, $\hat{y} \in \mathbb{R}^{l \times 1}$ is the predicted output and k is the dimensionality of the input vector, j is the number of nodes in layer 1 and l is the number of classes to predict.

To propagate the first two moments, several assumptions must be made. First, let us consider $w_m^\top = m^{\text{th}}$ row of W, $m = 1, 2, ..., j$. Then, $z_m = w_m^\top x + b_m^{[1]}$, $m = 1, 2, ..., j$. Next, consider the following assumptions: the input vector x is deterministic, $a \sim \mathcal{N}(\mu_a, \Sigma_a)$, $w_m \sim \mathcal{N}(\mu_{w_m}, \Sigma_{w_m})$, $m = 1, 2, ..., j$, $b_m^{[1]} \sim \mathcal{N}(\mu_{b^{[1]}}, \sigma^2_{b_m^{[1]}})$, $m = 1, 2, ..., j$, $v_n \sim \mathcal{N}(\mu_{v_n}, \Sigma_{v_n})$, $n = 1, 2, ..., l$, $b_n^{[2]} \sim \mathcal{N}(\mu_{b_n}^{[2]}, \sigma^2_{b_n^{[2]}})$, $n = 1, 2, ..., l$ and the weight vectors w_m, a, and bias $b^{[1]}$ and $b^{[2]}$ are uncorrelated to each other for $m = 1, 2, ..., j$. The elements of μ_z and σ^2_z are defined in Equations 6 and 7.

$$\begin{aligned}
\mu_{z_m} &= \mathbb{E}[w_m^\top x + b_m^{[1]}] \\
&= \mathbb{E}[w_m^\top]x + \mathbb{E}[b_m^{[1]}] \\
&= \mu_{w_m}^\top x + \mu_{b_m^{[1]}}
\end{aligned} \tag{6}$$

$$\begin{aligned}
\sigma_{z_m}^2 &= \mathrm{Var}[w_m^\top x + b_m^{[1]}] \\
&= \mathrm{Var}[w_m^\top x_m] + \mathrm{Var}[b_m^{[1]}] \\
&= x_m^\top \mathrm{Var}[w_m]x_m + \mathrm{Var}[b_m^{[1]}] \\
&= x_m^\top \Sigma_{w_m}^2 x_m + \sigma_{b_m^{[1]}}^2
\end{aligned} \tag{7}$$

Since we have assumed the weight vectors and elements of the bias b to be uncorrelated, $\Sigma_{w_p w_q} = 0$ and $\sigma_{b_p b_q} = 0$ for $p \neq q$, where $p, q = 1, 2, ..., j$. Hence, the covariance is zero. To propagate the first two moments through an arbitrary element-wise non-linear function, we utilize the first-order Taylor series approximation shown in Equations 8 and 9.

$$\begin{aligned}
a &= f(z) \\
&= f(\mu_z) + f'(\mu_z)(z - \mu_z) + ... \\
&\approx f(\mu_z) + f'(\mu_z)(z - \mu_z) \\
\mathbb{E}[a] = \mu_a &\approx f(\mu_z) \\
\mu_a &= f(\mu_z)
\end{aligned} \tag{8}$$

$$\Sigma_{a_p a_q} \approx \begin{cases} \sigma_{z_p}^2 f'(\mu_z)^2, & p = q, \\ \sigma_{z_p z_q} f'(\mu_{z_p}) f'(\mu_{z_q}), & p \neq q. \end{cases} \tag{9}$$

For the second layer of the network, we can no longer consider the incoming vector, μ_a, to be deterministic. Additionally we now need to propagate the incoming variance, σ_a^2. The mean and covariance propagated through the second fully-connected layer are given in Equations 10 and 11.

$$\begin{aligned}
\mu_{\tilde{y}} &= \mathbb{E}[v_n^\top a + b_n^{[2]}] \\
&= \mathbb{E}[v_n^\top]\mathbb{E}[a] + \mathbb{E}[b_n^{[2]}] \\
&= \mu_{v_n}^\top \mu_a + \mu_{b_n^{[2]}}
\end{aligned} \tag{10}$$

$$\Sigma_{\tilde{y}_p \tilde{y}_q} = \begin{cases} \mathrm{Tr}(\Sigma_v \Sigma_a) + \mu_v^\top \Sigma_a \mu_v \\ + \mu_a^\top \Sigma_v \mu_a + \Sigma_{b_n^{[2]}}, & p = q, \\ \mu_{v_p}^\top \Sigma_a \mu_{v_q}, & p \neq q. \end{cases} \tag{11}$$

Next, we must use a non-linear activation function such as softmax to obtain the predictions of our model. Since g does not operate element-wise on our mean and variance, we use a slightly different Taylor-series to obtain Equations 12 and 13 [59].

$$\mu_{\hat{y}} \approx g(\mu_{\tilde{y}}) \tag{12}$$

$$\Sigma_{\hat{y}} \approx J_g \Sigma_{\tilde{y}} J_g^\top, \tag{13}$$

where J_g is the Jacobian matrix of the softmax function g with respect to $\tilde{y}$ and calculated at $\mu_{\tilde{y}}$.

Finally, we use the Evidence Lower Bound (ELBO) function, $\mathcal{L}(\phi, D)$, which consists of two parts: the expected log-likelihood of the training data given the weights, and a regularization term, $D = \{x^{(i)}, y^{(i)}\}_{i=1}^N$ given by Equation 14.

$$\mathcal{L}(\phi, D) = \mathbb{E}_{q(\phi)}[\log p(D|\phi)] - \mathrm{KL}[q(\phi)|p(\phi)], \tag{14}$$

where ϕ represents the weights W, V, and biases $b^{[1]}$, $b^{[2]}$. The expected log-likelihood is given in Equation 15 and the KL term is given in Equation 16.

$$
\begin{aligned}
E_{q(\phi)}[\log p(D|\phi)] &\approx \frac{1}{M} \sum_{m=1}^{M} \log p(D|\phi) \\
&\approx -\frac{Nl}{2} \log(2\pi) - \frac{1}{M} \sum_{m=1}^{M} \left[\frac{N}{2} \log(|\Sigma_{\hat{y}}|) \right. \\
&\quad \left. + \frac{1}{2} \sum_{i=1}^{N} (y^{(i)} - \mu_{\hat{y}}^{(m)})^{\top} (\Sigma_{\hat{y}}^{(m)})^{-1} (y^{(i)} - \mu_{\hat{y}}^{(m)}) \right]
\end{aligned}
\tag{15}
$$

where, $y^{(i)}$ is the true label of the i^{th} data point, N is the number of data points and M is the number of Monte Carlo samples needed to approximate the expectation by summation.

$$
\text{KL}[q(\phi)|p(\phi)] = -\frac{1}{2} \sum_{n=1}^{l} (j \log \sigma_{v_n}^2 - ||\mu_{v_n}||_F^2 - j\sigma_{v_n}^2))
\tag{16}
$$

B   Additional Model Selection Metrics

Table 3 shows the results of a 10-fold cross validation using the optimal parameters found for each model. The metric typically used to compare mortality prediction models is the area under the Receiver Operating Characteristic Area Under the Curve (ROC AUC). Both of our models achieve a significantly higher ROC AUC than the currently available and widely used by ICU Clinicians (APACHE [51], SAPS [50]), as well as a recently published Neural Network model [26] ($p < 0.01$ for all pair-wise t-tests). An initially troublesome result was the low precision obtained by both models. Due to the loss function weighting/majority undersampling, both models predict the positive class (right or wrong) about 40% of the time. This results in the model correctly classifying the positive class about 76% of the time. Using Bayes rule, we can estimate a theoretical positive predictive value of 0.15, given the positive class only occurs 8% of the time in the dataset. Therefore, our models are performing as expected given this imbalance.

Table 3: 10-fold cross validated training results for the traditional (deterministic) and Bayesian [33] (stochastic) neural nets. The mean score across the 10 runs is shown with the standard error. Parameter sweeps were run on both the eICU dataset [41] only as well as the union of the eICU [41] and MIMIC-III [40] datasets. The goal of the parameter sweep was to optimize positive likelihood ratio (LR+).

Deterministic Network	Precision↑	Sensitivity↑	Specificity↑	ROC AUC↑	PRC AUC↑	Balanced Accuracy↑
eICU Only	0.21±(0.01)	0.79±(0.02)	0.82±(0.02)	0.89±(0.00)	0.43±(0.01)	0.81±(0.00)
eICU + MIMIC-III	0.19±(0.00)	0.80±(0.00)	0.77±(0.00)	0.87±(0.00)	0.39±(0.00)	0.79±(0.00)
Stochastic Network	**Precision↑**	**Sensitivity↑**	**Specificity↑**	**ROC AUC↑**	**PRC AUC↑**	**Balanced Accuracy↑**
eICU Only	0.20±(0.01)	0.79±(0.02)	0.81±(0.02)	0.88±(0.00)	0.37±(0.02)	0.80±(0.00)
eICU + MIMIC-III	0.19±(0.00)	0.80±(0.01)	0.79±(0.00)	0.87±(0.00)	0.35±(0.02)	0.79±(0.00)

Variable Drift by Cohort

Figure 4 shows kernel density estimate by cohort for the features that were determined by a 2-sample Kolmogorov–Smirnov test to be significantly different from the training set by the confidence metric from the stochastic model.
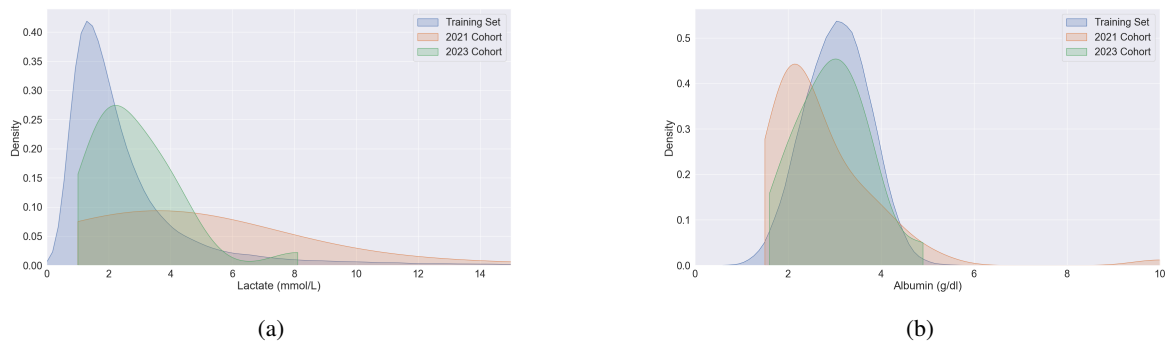
<center>(a)   (b)</center>

Figure 4: Kernel density estimates of two features that showed significant difference in distribution compared to the training set in 2021 but not in 2023 as a result of COVID-19 dataset drift ($p < 0.01$ against the 2021 cohort and $p > 0.01$ against the 2023 cohort using a 2-sample Kolmogorov–Smirnov test). **Left:** Lactate **Right:** Albumin

References

[1] Bertalan Meskó and Marton Görög. A short guide for medical professionals in the era of artificial intelligence. *npj Digital Medicine*, 3(1):1–8, 2020.

[2] Effy Vayena, Alessandro Blasimme, and I Glenn Cohen. Machine learning in medicine: addressing ethical challenges. *PLoS medicine*, 15(11):e1002689, 2018.

[3] Gary B Smith, David R Prytherch, Paul Meredith, Paul E Schmidt, and Peter I Featherstone. The ability of the national early warning score (news) to discriminate patients at risk of early cardiac arrest, unanticipated intensive care unit admission, and death. *Resuscitation*, 84(4):465–470, 2013.

[4] Aaron J Masino, Mary Catherine Harris, Daniel Forsyth, Svetlana Ostapenko, Lakshmi Srinivasan, Christopher P Bonafide, Fran Balamuth, Melissa Schmatz, and Robert W Grundmeier. Machine learning models for early sepsis recognition in the neonatal intensive care unit using readily available electronic health record data. *PloS one*, 14(2):e0212665, 2019.

[5] Madeleine Clare Elish. The stakes of uncertainty: developing and integrating machine learning in clinical care. In *Ethnographic Praxis in Industry Conference Proceedings*, volume 2018, pages 364–380. Wiley Online Library, 2018.

[6] Armando D Bedoya, Meredith E Clement, Matthew Phelan, Rebecca C Steorts, Cara O'Brien, and Benjamin A Goldstein. Minimal impact of implemented early warning score and best practice alert for patient deterioration. *Critical care medicine*, 47(1):49, 2019.

[7] Jessica L Guidi, Katherine Clark, Mark T Upton, Hilary Faust, Craig A Umscheid, Meghan B Lane-Fall, Mark E Mikkelsen, William D Schweickert, Christine A Vanzandbergen, Joel Betesh, et al. Clinician perception of the effectiveness of an automated early warning and response system for sepsis in an academic medical center. *Annals of the American Thoracic Society*, 12(10):1514–1519, 2015.

[8] Maya Dewan, Naveen Muthu, Eric Shelov, Christopher P Bonafide, Patrick Brady, Daniela Davis, Eric S Kirkendall, Dana Niles, Robert M Sutton, Danielle Traynor, et al. Performance of a clinical decision support tool to identify picu patients at high-risk for clinical deterioration. *Pediatric critical care medicine: a journal of the Society of Critical Care Medicine and the World Federation of Pediatric Intensive and Critical Care Societies*, 21(2):129, 2020.

[9] Jennifer C Ginestra, Heather M Giannini, William D Schweickert, Laurie Meadows, Michael J Lynch, Kimberly Pavan, Corey J Chivers, Michael Draugelis, Patrick J Donnelly, Barry D Fuchs, et al. Clinician perception of a machine learning-based early warning system designed to predict severe sepsis and septic shock. *Critical care medicine*, 47(11):1477, 2019.

[10] Sana Tonekaboni, Shalmali Joshi, Melissa D McCradden, and Anna Goldenberg. What clinicians want: contextualizing explainable machine learning for clinical end use. In *Machine learning for healthcare conference*, pages 359–380. PMLR, 2019.

[11] José Jiménez-Luna, Francesca Grisoni, and Gisbert Schneider. Drug discovery with explainable artificial intelligence. *Nature Machine Intelligence*, 2(10):573–584, 2020.

[12] Christine M Cutillo, Karlie R Sharma, Luca Foschini, Shinjini Kundu, Maxine Mackintosh, Kenneth D Mandl, and MI in Healthcare Workshop Working Group Beck Tyler 1 Collier Elaine 1 Colvis Christine 1 Gersing Kenneth 1 Gordon Valery 1 Jensen Roxanne 8 Shabestari Behrouz 9 Southall Noel 1. Machine intelligence in healthcare—perspectives on trustworthiness, explainability, usability, and transparency. *NPJ digital medicine*, 3(1):47, 2020.

[13] Carlo Combi, Beatrice Amico, Riccardo Bellazzi, Andreas Holzinger, Jason H Moore, Marinka Zitnik, and John H Holmes. A manifesto on explainability for artificial intelligence in medicine. *Artificial Intelligence in Medicine*, 133:102423, 2022.

[14] Alex John London. Artificial intelligence and black-box medical decisions: accuracy versus explainability. *Hastings Center Report*, 49(1):15–21, 2019.

[15] Ian E Nielsen, Ravi P Ramachandran, Nidhal Bouaynaya, Hassan M Fathallah-Shaykh, and Ghulam Rasool. Evalattai: A holistic approach to evaluating attribution maps in robust and non-robust models. *IEEE Access*, 2023.

[16] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

[17] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015.

[18] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.

[19] Sabeen Ahmed, Dimah Dera, Saud Ul Hassan, Nidhal Bouaynaya, and Ghulam Rasool. Failure detection in deep neural networks for medical imaging. *Frontiers in Medical Technology*, 4, 2022.

[20] Craig A Umscheid, Joel Betesh, Christine VanZandbergen, Asaf Hanish, Gordon Tait, Mark E Mikkelsen, Benjamin French, and Barry D Fuchs. Development, implementation, and impact of an automated early warning and response system for sepsis. *Journal of hospital medicine*, 10(1):26–31, 2015.

[21] Peter J Embi and Anthony C Leonard. Evaluating alert fatigue over time to ehr-based clinical trial alerts: findings from a randomized controlled study. *Journal of the American Medical Informatics Association*, 19(e1):e145–e148, 2012.

[22] Christopher J Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. Key challenges for delivering clinical impact with artificial intelligence. *BMC medicine*, 17(1):1–9, 2019.

[23] Bret Nestor, Matthew McDermott, Geeticka Chauhan, Tristan Naumann, Michael C Hughes, Anna Goldenberg, and Marzyeh Ghassemi. Rethinking clinical prediction: why machine learning must consider year of care and feature aggregation. *arXiv preprint arXiv:1811.12583*, 2018.

[24] Irene Chen, Fredrik D Johansson, and David Sontag. Why is my classifier discriminatory? *Advances in neural information processing systems*, 31, 2018.

[25] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR, 2017.

[26] Jacob R Epifano, Ravi P Ramachandran, Sharad Patel, and Ghulam Rasool. Towards an explainable mortality prediction model. In *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2020.

[27] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

[28] Raquel Rodríguez-Pérez and Jürgen Bajorath. Interpretation of compound activity predictions from complex machine learning models using local approximations and shapley values. *Journal of medicinal chemistry*, 63(16):8761–8777, 2019.

[29] Nima Safaei, Babak Safaei, Seyedhouman Seyedekrami, Mojtaba Talafidaryani, Arezoo Masoud, Shaodong Wang, Qing Li, and Mahdi Moqri. E-catboost: An efficient machine learning framework for predicting icu mortality using the eicu collaborative research database. *Plos one*, 17(5):e0262895, 2022.

[30] Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3):457–506, 2021.

[31] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.

[32] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.

[33] Dimah Dera, Ghulam Rasool, and Nidhal Bouaynaya. Extended variational inference for propagating uncertainty in convolutional neural networks. In *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2019.

[34] Pavel Izmailov, Dmitrii Podoprikhin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018.

[35] Benjamin Kompa, Jasper Snoek, and Andrew L Beam. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digital Medicine*, 4(1):4, 2021.

[36] Sara C Auld, Kristin RV Harrington, Max W Adelman, Chad J Robichaux, Elizabeth C Overton, Mark Caridi-Scheible, Craig M Coopersmith, and David J Murphy. Trends in icu mortality from coronavirus disease 2019: a tale of three surges. *Critical Care Medicine*, 50(2):245, 2022.

[37] Li Yan, Hai-Tao Zhang, Jorge Goncalves, Yang Xiao, Maolin Wang, Yuqi Guo, Chuan Sun, Xiuchuan Tang, Liang Jing, Mingyang Zhang, et al. An interpretable mortality prediction model for covid-19 patients. *Nature machine intelligence*, 2(5):283–288, 2020.

[38] Marian JR Quanjel, Thijs C van Holten, Pieternel C Gunst-van der Vliet, Jette Wielaard, Bekir Karakaya, Maaike Söhne, Hazra S Moeniralam, and Jan C Grutters. Replication of a mortality prediction model in dutch patients with covid-19. *Nature Machine Intelligence*, 3(1):23–24, 2021.

[39] Sharon E Davis, Robert A Greevy Jr, Christopher Fonnesbeck, Thomas A Lasko, Colin G Walsh, and Michael E Matheny. A nonparametric updating method to correct clinical prediction model drift. *Journal of the American Medical Informatics Association*, 26(12):1448–1457, 2019.

[40] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

[41] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, 5(1):1–13, 2018.

[42] Stef Van Buuren and Karin Groothuis-Oudshoorn. mice: Multivariate imputation by chained equations in r. *Journal of statistical software*, 45:1–67, 2011.

[43] Lukas Biewald. Experiment tracking with weights and biases, 2020. Software available from wandb.com.

[44] Jonathan J Deeks and Douglas G Altman. Diagnostic tests 4: likelihood ratios. *Bmj*, 329(7458):168–169, 2004.

[45] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.

[46] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research*, 18(17):1–5, 2017.

[47] Dimah Dera, Ghulam Rasool, Nidhal C Bouaynaya, Adam Eichen, Stephen Shanko, Jeff Cammerata, and Sanipa Arnold. Bayes-SAR net: Robust SAR image classification with uncertainty estimation using bayesian convolutional neural network. In *2020 IEEE International Radar Conference (RADAR)*, pages 362–367. IEEE, 2020.

[48] Dimah Dera, Nidhal Bouaynaya, Ghulam Rasool, R Shterenberg, and Hassan Fathallah-Shaykh. Premium-cnn: Propagating uncertainty towards robust convolutional neural networks. *IEEE Transactions on Signal Processing*, 2021.

[49] Ameur M Manceur and Pierre Dutilleul. Maximum likelihood estimation for the tensor normal distribution: Algorithm, minimum sample size, and empirical bias and dispersion. *Journal of Computational and Applied Mathematics*, 239:37–49, 2013.

[50] Jean-Roger Le Gall, Stanley Lemeshow, and Fabienne Saulnier. A new simplified acute physiology score (saps ii) based on a european/north american multicenter study. *Jama*, 270(24):2957–2963, 1993.

[51] Jack E Zimmerman, Andrew A Kramer, Douglas S McNair, and Fern M Malila. Acute physiology and chronic health evaluation (apache) iv: hospital mortality assessment for today's critically ill patients. *Critical care medicine*, 34(5):1297–1310, 2006.

[52] Jacob R. Epifano, Alison Silvestri, Alexander Yu, Ravi P. Ramachandran, Aakash Tripathi, and Ghulam Rasool. A comparison of feature selection techniques for first-day mortality prediction in the icu. In *2023 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2023.

[53] Samyadeep Basu, Philip Pope, and Soheil Feizi. Influence functions in deep learning are fragile. *arXiv preprint arXiv:2006.14651*, 2020.

[54] Jacob R Epifano, Ravi P Ramachandran, Aaron J Masino, and Ghulam Rasool. Revisiting the fragility of influence functions. *Neural Networks*, 162:581–588, 2023.

[55] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.

[56] Dylan Slack, Sophie Hilgard, Emily Jia, Sameer Singh, and Himabindu Lakkaraju. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 180–186, 2020.

[57] Himabindu Lakkaraju and Osbert Bastani. " how do i fool you?" manipulating user trust via misleading black box explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 79–85, 2020.

[58] Carlos Fernández-Loría, Foster Provost, and Xintian Han. Explaining data-driven decisions made by ai systems: the counterfactual approach. *arXiv preprint arXiv:2001.07417*, 2020.

[59] Dan Simon. *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.