Learning to Predict with Supporting Evidence: Applications to Clinical Risk Prediction

Aniruddh Raghu
araghu@mit.edu
Massachusetts Institute of Technology

John Gutttag
Massachusetts Institute of Technology

Katherine Young
Massachusetts Institute of Technology
Harvard Medical School

Eugene Pomerantsev
Massachusetts General Hospital

Adrian V. Dalca
Massachusetts Institute of Technology
Massachusetts General Hospital
Harvard Medical School

Collin M. Stultz
Massachusetts Institute of Technology
Massachusetts General Hospital
Harvard Medical School

ABSTRACT

The impact of machine learning models on healthcare will depend on the degree of trust that healthcare professionals place in the predictions made by these models. In this paper, we present a method to provide individuals with clinical expertise with domain-relevant evidence about why a prediction should be trusted. We first design a probabilistic model that relates meaningful latent concepts to prediction targets and observed data. Inference of latent variables in this model corresponds to both making a prediction *and* providing supporting evidence for that prediction. We present a two-step process to efficiently approximate inference: (i) estimating model parameters using variational learning, and (ii) approximating *maximum a posteriori* estimation of latent variables in the model using a neural network, trained with an objective derived from the probabilistic model. We demonstrate the method on the task of predicting mortality risk for patients with cardiovascular disease. Specifically, using electrocardiogram and tabular data as input, we show that our approach provides appropriate domain-relevant supporting evidence for accurate predictions.

KEYWORDS

Machine Learning, Interpretability

ACM Reference Format:

Aniruddh Raghu, John Gutttag, Katherine Young, Eugene Pomerantsev, Adrian V. Dalca, and Collin M. Stultz. 2021. Learning to Predict with Supporting Evidence: Applications to Clinical Risk Prediction. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3450439.3451869>

1 INTRODUCTION

There is significant interest in using machine learning (ML) models in high-stakes domains such as medicine, where human experts use model predictions to inform decisions [6, 9, 46]. When used

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8359-2/21/04.

<https://doi.org/10.1145/3450439.3451869>

Figure 1: Learning to Predict with Supporting Evidence. We propose a model that produces both a prediction and accompanying domain-relevant supporting evidence for that prediction: inferred concepts that are clinically meaningful and can assist a clinician in deciding how to act on the prediction. Consistency of the prediction and supporting evidence is ensured using a probabilistic model and domain knowledge.

by a clinician, the utility of an ML model depends not only on the accuracy of its predictions, but also on how much trust the clinician places in the prediction [42].

Building trust in predictions can be tackled in different ways, including (1) providing statistical arguments to reflect the level of certainty in a prediction [31, 35] (2) providing an explanation of how a model reached its prediction, either through understanding important input features that led to the model output, or by using an inherently explainable model [27, 43]; and (3) offering *domain-specific supporting evidence* about the prediction [13].

For models used by domain experts with a limited understanding of machine learning, prior work has emphasised the importance of providing *domain-relevant supporting evidence* for predictions [21, 28]. In medicine, presenting relevant clinical information that supports a model prediction can be crucial, since clinicians draw significantly on medical knowledge when making decisions [42, 45]. For example, a cardiologist might deem a patient to be at high risk of death (prediction) because of their low cardiac output (domain-relevant supporting evidence), which may inform a therapeutic choice [47]. A model that produces a prediction (e.g.,

high risk of death) accompanied by clinically meaningful evidence for this prediction (e.g., poor cardiac health) closely mirrors a clinician's mental model when treating patients, helping to build trust.

In this paper, we present a method to construct predictive models that are accurate *and* provide domain-relevant supporting evidence for predictions (Figure 1). The method is designed to provide users that have a deep understanding of the application domain with an appropriate reason to trust or distrust a prediction made by the model. Our contributions are:

- (1) We design a probabilistic model that relates observed data and prediction targets to domain-relevant latent variables. We embed a rich forward model driven by domain knowledge to ensure that the latent variables correspond to domain-relevant concepts. By construction, *maximum a posteriori* (MAP) inference in this probabilistic model yields both accurate predictions and domain-relevant supporting evidence.
- (2) We demonstrate a two-step learning process that approximates MAP inference: (i) estimating parameters of the probabilistic model using variational learning; (ii) approximating MAP estimation of latent variables in the model, yielding a prediction and supporting evidence, using a neural network trained with an objective derived from the probabilistic model. Importantly, we do not need labels for supporting concepts at training time.
- (3) We demonstrate our method on a real-world clinical dataset for the task of predicting mortality risk for patients with cardiovascular disease using multimodal electrocardiogram and tabular data. We show that our method produces accurate risk predictions jointly with meaningful supporting evidence for predictions. The supporting evidence captures information that is often only obtained from invasive procedures, and could provide important therapeutic insight for clinicians.

2 RELATED WORK

Work on improving trust has often focused on ML model interpretability, including *post hoc* interpretation using input feature attribution [25, 40, 43, 50]; *post hoc* concept-based attribution [10, 16]; and constructing explainable models through regularisation [27, 33] or prototype/input extraction [1, 22, 24]. Our work differs in spirit from these, since we: (1) do not focus on explaining how the predictive model works, and instead on providing explicit supporting evidence that is useful to domain experts with limited ML understanding; (2) embed a mechanism to provide supporting evidence for predictions directly into the model, rather than relying on *post hoc* analysis; and (3) use prior domain knowledge to inform the higher-level abstractions for supporting evidence, rather than learning abstractions that may not resemble inherently meaningful concepts.

One closely related work is on Self-Explaining Neural Networks (SENN) [27], which provide explanations for predictions by forming predictions as a product of input-dependent concepts and weighting terms for these concepts. The concepts for explanations are learned from data (not constrained by domain understanding) and are interpreted by considering input examples that most characterise them. Unlike the supporting evidence we consider, learned concepts in SENN need not resemble meaningful abstractions that a domain expert would find useful in decision making. Finally, SENN

Figure 2: Probabilistic model for LPS. The model relates the class label y_i to (latent) class probability vector π_i , (latent) supporting evidence factors z_i , observed features for each data point x_i , and distributional parameters ϕ and ψ . The circled quantities are random variables, with shading used for quantities observed at training time. At inference time, we only observe the features x_i . The plate signifies variables that are data point-specific, for N data points, and the quantities outside the plate are common across data points.

uses input examples as prototypes to characterise learned concepts. With complex, multimodal data, as in our clinical experiment, such examples can be challenging to visualise and understand.

Other recent work on using domain-relevant concepts to support predictions is close in spirit to ours [1, 13, 20], but these methods require labelled concepts at training time. In contrast, our work relates the supporting concepts to the predictions through a probabilistic model, such that inference directly yields coupled predictions and supporting evidence for those predictions. We detail further differences from related work (such as Koh et al. [20], Melis and Jaakkola [27]) in the appendix.

3 LEARNING TO PREDICT WITH SUPPORTING EVIDENCE (LPS)

In this section, we present our method, *Learning to Predict with Supporting Evidence (LPS)*, to construct models that produce both predictions and clinically-relevant supporting evidence. We define a probabilistic generative model that relates observed data, domain-relevant concepts, and predictive targets, specifying how to ground the concepts using a forward model. *Maximum a posteriori* (MAP) inference in this model jointly yields predictions and domain-relevant supporting evidence. Since such inference is computationally challenging, we present a two-step learning process to approximate MAP inference.

3.1 Probabilistic Model

We describe the model for a general binary classification problem, where the task is to predict class label y_i and positive class probability π_i , given observed features x_i . The model is summarised in Figure 2.

Formally, for $i = 1, \dots, N$, the class label $y_i \in \{0, 1\}$ is distributed as $y_i \sim \text{Bernoulli}(\pi_i)$ where the positive class probability $\pi_i \in [0, 1]$ is distributed as the flexible prior $\pi_i \sim \text{Beta}(\alpha, \beta)$. Fixed parameters α and β capture the overall balance of the two classes. In clinical risk prediction, π_i could represent an individual patient's 60-day risk of death.

We let $z_i \in \mathbb{R}^m$, $z_i \sim p(z_i|\pi_i, \phi)$ be a latent vector encoding interpretable, domain-specific concepts, where m is the number of concepts, and global parameters ϕ . The variable z_i represents some clinically relevant concept that does not appear directly in the feature space, e.g., estimated glomerular filtration rate (eGFR). The observed features $x_i \in \mathbb{R}^d$, are distributed as $x_i \sim p(x_i|z_i, \psi)$, with global parameters ψ . In risk prediction, x_i could be a patient's creatinine, and low creatinine clearance is associated with low eGFR.

Domain-Relevant Supporting Evidence: We enforce that the variable z encodes domain-relevant concepts by using a well-defined forward model, governed by domain knowledge, that characterises how some subset of the observations x are generated from z . For example, the forward model could be a physiological model that relates observable values to latent cardiac function [7], atlas-based image deformation models [8], physics models [14], or many others. We assume a prior $p(\phi)$ determined by domain understanding.

This modelling directly grounds z to represent domain-relevant concepts by relating z to x via the forward model. Consistency of z and π is enforced through the probabilistic model. We emphasise that the forward model need only characterise a sufficient subset of x to constrain z ; that is, we do not need to define a model to relate z to all of x . This is advantageous and broadens LPS's applicability because x may be very high-dimensional and thus it may be challenging to define a forward model governing x in its entirety.

In this work, we consider only smooth, continuous distributions for the domain knowledge, and a differentiable forward model. In addition to being widely applicable, this enables tractable learning using gradient-based methods.

3.2 Learning

For a given observation x , we use *maximum a posteriori* (MAP) estimates of the class label y , its associated probability π , and the supporting concepts z :

$$\pi^*, z^*, y^* = \arg \max_{\pi, z, y} p(\pi, z, y|x) \quad (1)$$

$$= \arg \max_{\pi, z, y} \int p(\pi, z, y|x, \phi, \psi) p(\phi) p(\psi) d\phi d\psi, \quad (2)$$

Since this is intractable for complex parameter specifications, we approximate it. We use a training set $\mathcal{D} = \{(x_1, y_1), \dots, (x_N, y_N)\}$ to learn informative point estimates for the parameters:

$$\phi^*, \psi^* = \arg \max_{\phi, \psi} \log p(\psi, \phi|\mathcal{D}), \quad (3)$$

and approximate MAP estimates as:

$$\pi^*, z^*, y^* = \arg \max_{\pi, z, y} \log p(\pi, z, y|x, \phi^*, \psi^*). \quad (4)$$

Since the true posterior $p(z, \pi|x, y, \psi, \phi)$ is intractable, we cannot directly use the Expectation Maximization (EM) algorithm to estimate model parameters, because EM requires calculating this

posterior exactly. We therefore resort to variational EM [5, 32] for parameter estimation: we define a variational approximation $q(z, \pi; \theta_q)$ to the true posterior with parameters θ_q , and then construct a lower bound on the log evidence of data and model parameters as follows.

First, we express the joint log likelihood of data and model parameters as:

$$\begin{aligned} \log p(\phi, \psi, \mathcal{D}) &= \log p(\phi, \psi, x_1^N, y_1^N) \\ &= \log p(\phi) + \log p(\psi) + \\ &\quad \sum_{i=1}^N \log p(x_i, y_i|\phi, \psi) \end{aligned} \quad (5)$$

Then, we lower bound the data likelihood terms $\log p(x_i, y_i|\phi, \psi)$ with Jensen's inequality (full derivation in appendix), and obtain the following bound:

$$\begin{aligned} \log p(\psi, \phi, \mathcal{D}) &\geq \log p(\psi) + \log p(\phi) + \\ &\quad \sum_{i=1}^N \left(\mathbb{E}_q \left[\log p(\pi) + \log p(y_i|\pi) + \right. \right. \\ &\quad \left. \left. \log p(z|\pi, \phi) + \log p(x_i|z, \psi) \right] + H(q) \right), \end{aligned} \quad (6)$$

where $H(q)$ is the entropy of q . We obtain MAP estimates by jointly maximising this lower bound w.r.t. ϕ, ψ , and θ_q using stochastic gradient approximations. This results in MAP estimates ϕ^*, ψ^* , and an approximate posterior $q(\cdot; \theta_q^*)$.

3.3 MAP Inference

For a new data point x , we seek the MAP estimates of π, z , and y :

$$\begin{aligned} \pi^*, z^*, y^* &= \arg \max_{\pi, z, y} \log p(\pi, z, y|x, \phi^*, \psi^*) \\ &= \arg \max_{\pi, z, y} \log p(\pi) + \log p(y|\pi) + \\ &\quad \log p(z|\pi, \phi^*) + \log p(x|z, \psi^*). \end{aligned} \quad (7)$$

Depending on the form of each distribution, this maximization might not be solvable in closed form. Computing the MAP estimates on a per-subject basis using EM or gradient ascent is often computationally inefficient with complex $p(x|z, \psi)$. Instead, we use a network $n(x; \theta_n) \rightarrow (\hat{\pi}, \hat{z})$ with parameters θ_n to efficiently approximate the MAP estimates such that:

$$\hat{\pi} \approx \pi^*, \hat{z} \approx z^*, \quad (8)$$

$$\hat{y} = \mathbb{1}[\hat{\pi} \geq \eta] \approx y^*, \quad (9)$$

where $\mathbb{1}[\cdot]$ is an indicator function and η is a threshold that depends on our desired tradeoff between recall and precision. Equation 9 follows from the fact that $p(y|\pi)$ is Bernoulli distributed.

We learn the parameters θ_n by maximising the MAP objective on the set of training data $\mathcal{D} = \{(x_1, y_1), \dots, (x_N, y_N)\}$. Given a network $n(x_i; \theta_n) \rightarrow (\hat{\pi}_i, \hat{z}_i)$, we maximise:

$$\begin{aligned} \mathcal{L}_{\text{MAP}}(\theta_n) &= \sum_{i=1}^N \left(\log p(\hat{\pi}_i) + \log p(y_i|\hat{\pi}_i) + \right. \\ &\quad \left. \log p(\hat{z}_i|\hat{\pi}_i, \phi^*) + \log p(x_i|\hat{z}_i, \psi^*) \right), \end{aligned} \quad (10)$$

which encourages simultaneous accurate risk prediction $\hat{\pi}$ and domain-relevant supporting evidence $\hat{z}$. Consistency between the

supporting evidence $\hat{z}$ and prediction $\hat{\pi}$ is enforced by the MAP objective from the probabilistic model used to train the inference model (via the term $p(\hat{z}|\hat{\pi})$).

To recover π^*, z^*, y^* given a test example x , we directly use $n(x; \theta_n^*)$ to obtain the MAP estimates using (8), and (9).

3.4 Discussion on Modelling

Representation of domain knowledge: In some clinical settings, domain knowledge is typically represented using hard constraints or thresholds. For example, one definition of sepsis uses hard thresholds of a severity score [41]. However, it is often the case that a continuous distribution is often a more realistic representation of the relationship between relevant concepts and the prediction task. For example, the relationship between severity scores and risk of mortality is well represented as a smooth, continuous function, since higher severity scores correlate with higher risk in a smooth fashion.

Generality of the model: This class of models is appropriate for a range of binary classification problems because the Beta distribution is a flexible distribution for π and the distributions $p(z|\pi, \phi)$ and $p(x|z, \psi)$ can be specified as desired. Using a Dirichlet prior for π and a multinomial for y is a natural extension of this model to the multiclass scenario. In this paper, we focus on binary classification for clinical risk stratification.

Supporting Evidence vs. Explanations: Since there is no guarantee that there is a causal link between $\hat{z}$ and the prediction $\hat{\pi}$, we refer to $\hat{z}$ as providing *supporting evidence* for the prediction rather than *explaining* the prediction.

4 RISK PREDICTION IN REAL-WORLD MEDICINE

We instantiate and demonstrate LPS on the task of predicting a cardiovascular patient's risk of mortality within the 6 months following a cardiac event. We show that LPS produces accurate risk predictions and informative supporting evidence. Code implementing LPS is available at <https://github.com/aniruddhraghu/lps>.

Cardiovascular disease affects a large number of people worldwide, and is a major cause of mortality [4]. Predicting mortality risk for patients with cardiovascular disease is an important task that has received significant prior study in the medical literature [2, 11, 26, 29, 30, 34, 37]. In this section, we focus on predicting patient risk and providing supporting evidence for predictions from electrocardiogram (ECG) and tabular data, which are observed for many cardiology patients in hospital settings.

4.1 Data

The dataset we use has 3728 patients from the Massachusetts General Hospital, is de-identified, and was obtained with IRB approval. Each patient has undergone a cardiac catheterisation, which we treat as the index event. Each patient in the dataset has the following measurements and demographic features recorded: 12-lead electrocardiogram (ECG), heart rate (HR), systolic blood pressure, diastolic blood pressure, ethnicity, age, and gender. These values are recorded within 3 days preceding catheterisation event. For each feature, the latest value before the cardiac catheterisation event

is used if multiple are available. Additionally, approximately 80% of the patients have creatinine values recorded within this 3 day period. For those that do not, we use a standard value for a healthy patient (1 mg/dL). From creatinine and the other features, we calculate the estimated glomerular filtration rate (eGFR) for each patient using a standard formula [23]. The data are summarised in Figure 3. Each patient has outcome information: survival ($y = 0$) or mortality ($y = 1$) within 6 months of catheterisation (positive proportion 9.7%).

4.2 Model Instantiation

We let x be the ECG and the tabular features, summarised in Figure 3, and y the mortality outcome. The risk π is the probability of y . We let $z \in \mathbb{R}^5$ be a vector of clinically important quantities: systemic vascular resistance (R), arterial compliance (C), systole time (T_s), diastole time (T_d), and cardiac output (CO). These quantities are **not** observed for many patients; for some, such as R , C , and CO , this is because accurate direct measurement is invasive [3], and estimating them in (non-invasive) physical exams is challenging [12]. When they are available, they play an important role in risk assessment. For example, when known, the values of CO and R are used to guide the choice of therapeutic interventions for patients with heart failure [49].

4.3 Probabilistic Model Specification

We now detail how each distribution in the probabilistic model is specified. We first define $p(\pi) = \text{Beta}(1.0, 1.0)$ and $p(y|\pi) = \text{Bernoulli}(\pi)$. We impose a flat prior on ψ : $p(\psi) \propto 1$.

Specification for $p(z|\pi, \phi)$: The interpretable latent variable z is 5-dimensional, with components R, C, T_s, T_d , and CO . We denote the components of z as z_m , $m = 1, \dots, 5$. We factor the distribution $p(z|\pi, \phi)$ as follows:

$$p(z|\pi, \phi) = p(R|\pi, \phi_R)p(C|\pi, \phi_C)p(T_s|\pi, \phi_{T_s})p(T_d|\pi, \phi_{T_d})p(CO|\pi, \phi_{CO}).$$

For each z_m , we let $p(z_m|\pi, \phi_{z_m})$ have the following mixture form:

$$p(z_m|\pi, \phi_{z_m}) = \pi \mathcal{LN}(z_m; \mu_{z_m,1}, \sigma_{z_m,1}^2) + (1 - \pi) \mathcal{LN}(z_m; \mu_{z_m,0}, \sigma_{z_m,0}^2),$$

where $\mathcal{LN}$ is a log normal distribution, such that if $X \sim \mathcal{LN}(\mu, \sigma^2)$, then $\log X \sim \mathcal{N}(\mu, \sigma^2)$. This mixture form for each component implies the following definition of the parameters:

$$\phi_{z_m} = \{\mu_{z_m,1}, \sigma_{z_m,1}^2, \mu_{z_m,0}, \sigma_{z_m,0}^2\}.$$

This instantiation for the distribution has an intuitive interpretation of generating each latent variable z_m for each patient as a mixture of two components, weighted by the patient's risk of death. The two components then correspond to the distribution of the latent feature in the high and low risk cases.

Specification for $p(\phi)$: We define $\phi_\mu = \{\mu_{z_m,0}, \mu_{z_m,1}\}_{m=1}^5$, and define $\phi_{\sigma^2} = \{\sigma_{z_m,0}^2, \sigma_{z_m,1}^2\}_{m=1}^5$. For each of these means, we impose an independent, normally distributed prior, i.e.: $p(\mu_{z_m,i}) = \mathcal{N}(\tilde{\mu}_{z_m,i}, 0.01^2)$. For each of the standard deviations, we specify a delta function prior, i.e.: $p(\sigma_{z_m,i}^2) = \delta(\sigma_{z_m,i}^2 - \tilde{\sigma}_{z_m,i}^2)$, which fixes the values of these parameters to be $\tilde{\sigma}_{z_m,i}^2$.

Figure 3: Multimodal data for each patient in our clinical dataset: a 12-lead ECG and 8 tabular features.

We now specify how each of these parameters are determined for each component of z :

- R : Take the roughly 20% of patients in the training dataset that have recorded values for R . Consider the subset of these patients that died. For these patients, fit a lognormal distribution to the resulting R values. The parameters of this fitted distribution become $\tilde{\mu}_{R,1}$ and $\tilde{\sigma}_{R,1}^2$. Repeat this process for the subset of patients that lived; the resulting parameters are $\tilde{\mu}_{R,0}$ and $\tilde{\sigma}_{R,0}^2$.
- C : No patients have measured values for C (vascular compliance is never directly measured). Approximately estimate the time constant $\tau = RC$ for each patient in the training dataset for which we have R recorded. This is done by using the diastole relation from the two element Windkessel model [7] to relate the systolic and diastolic pressures to the time constant. Divide this resulting time constant by R to get approximate values for C . Then, determine parameters by following the same process as with determining distributions for R .
- T_s : No patients have measured values for T_s . Approximately estimate these by using the fact that T_s is approximately 1/3 of the duration of a heart beat (reciprocal of heart rate, recorded for all patients). Having estimated these on the training dataset, follow the same process as with R to determine parameters.
- T_d : No patients have measured values for T_d . Approximately estimate these by using the fact that T_d is approximately 2/3 of the duration of a heart beat (reciprocal of heart rate, recorded for all patients). Having estimated these on the training dataset, follow the same process as with R to determine parameters.
- CO : follow the same process as with determining R , for the $\sim 80\%$ of patients in the training dataset that have recorded values.

Specification for $p(x|z, \psi)$: We partition x into two components: (1) $x^{(g)}$, which is generated from z based on a known forward model $g(z)$; and (2) $x^{(f)}$, which is formed from z based on a forward model $f(z, \psi)$ with parameters ψ that cannot be specified with current domain knowledge, and is thus learned from data.

We let $x^{(g)}$ represent the vital signs (heart rate and blood pressures) and $x^{(f)}$ capture the remaining features. We define: $p(x^{(g)}|z, \psi) = \mathcal{N}(g(z), 0.1^2)$, where the known forward model $g(\cdot)$ has two components: (1) the two element Windkessel model to model the blood pressures using a differential equation [7, 38, 48]; and (2) the definition of the heart rate based on systolic and diastolic times. Concretely, $g(z) \rightarrow (\overline{BP}_{sys}, \overline{BP}_{dias}, \overline{HR})$, corresponding to the estimated means for the end systolic pressure, end diastolic pressure, and heart rate respectively. These are produced as follows:

- **Two element Windkessel model:** this is a differential equation model that models the evolution of the blood pressure waveform as a function of the latent variables R, C, T_s, T_d , and CO . This model separately characterises the systole phase, where blood is ejected out of the heart through the aorta as a result of ventricular contraction, and the diastole phase, where blood flows into the ventricles at the start of the next cardiac cycle. We adopt the formulation from Catanho et al. [7], and the blood pressure in each phase is modelled as follows:

$$C \frac{dP(t)}{dt} + \frac{P(t)}{R} = I(t),$$

where

$$I(t) = \begin{cases} I_0 \sin\left(\frac{\pi t}{T_s}\right) & \text{during systole} \\ 0 & \text{during diastole.} \end{cases}$$

$I(t)$ represents the input blood flow, with a half-sinusoid model used for systole, and zero-input for diastole. The constant I_0 is set based on the fact that over the systole phase, the entire blood flow must be equal to the stroke volume, SV . By definition, stroke volume can be expressed in terms of cardiac output and heart rate $SV = \frac{CO}{HR} = \frac{(T_s + T_d) \times CO}{60}$, and we obtain the following expression for I_0 :

$$\int_{t=0}^{T_s} I_0 \sin\left(\frac{\pi t}{T_s}\right) dt = SV = \frac{(T_s + T_d)CO}{60} \implies I_0 = \frac{\pi CO (T_s + T_d)}{120T_s}.$$

Note that π in these equations refers to the mathematical constant, and not the patient's risk of mortality.

With these governing differential equations, we use the piecewise solutions for the blood pressures in each phase from Catanho et al. [7]. This solution specifies the blood pressure as a function of time. To obtain the desired quantities, $\overline{BP}_{sys}$ and $\overline{BP}_{dias}$, we step forward the solution for 4 cardiac cycles (each cycle is of duration $T_s + T_d$) so that it reaches steady state. We then evaluate the solution at intervals of T_s and T_d for a further 6 cycles, and average the end systolic and diastolic pressures from these 6 cycles to produce $\overline{BP}_{sys}$ and $\overline{BP}_{dias}$. Note that the computation of these quantities is differentiable.

- **Definition of heart rate:** the heart rate in beats per minute can be defined as $\frac{60}{T_s + T_d}$. Evaluating this yields $\overline{HR}$.

These two models together fully specify $p(x^{(g)}|z, \psi)$. We then let $p(x^{(f)}|z, \psi) = \mathcal{N}(f(z, \psi), \sigma_f^2)$, where $f(\cdot)$ is a neural network with parameters ψ , which are learned as part of the LPS framework. For the ECG, we let $\sigma_f = 5$, and for the remaining features, we

let $\sigma_f = 0.5$. More details on network architecture are provided in Section 4.5 and in the appendix.

4.4 Learning and Inference

As described in Section 3, LPS has two stages: firstly, we obtain approximate MAP estimates $\hat{\phi}^*$ and $\hat{\psi}^*$ for model parameters (learning phase); secondly, we train another predictive model to efficiently output MAP estimates for latent model variables (inference phase), yielding a risk prediction and domain-relevant supporting evidence.

We use variational EM to derive approximate MAP estimates $\hat{\phi}^*$ and $\hat{\psi}^*$. We use a deep neural network to model the variational posterior $q(z, \pi|x; \theta_q)$. Again denoting the components of z as z_m , $m = 1, \dots, 5$, we use a mean field approximation and factor the variational posterior as $\prod_{m=1}^5 q(z_m|x)q(\pi|x)$. Each $q(z_m|x)$ is defined to be a log normal distribution, and $q(\pi|x)$ to be a Beta distribution. The variational posterior network takes the ECG and tabular features as input, and produces mean and variance estimates for the posterior of each z_m , and Beta concentration parameter estimates for the posterior of π . We approximate the expectation in (6) by drawing a single sample from the variational posterior using the reparameterisation trick [18, 36], enabling end-to-end gradient-based training. After running variational EM, we recover approximate MAP estimates $\hat{\phi}^*$ and $\hat{\psi}^*$.

For efficient MAP inference of latent variables, we train a MAP neural network $n(x; \theta_n)$ to take in the ECG and tabular features and directly output MAP estimates of z and π . This is trained using the objective in Equation 10.

Further details on learning and inference are in the appendix.

4.5 Neural Network Architectures

We now outline the architectures for the three neural networks involved in this LPS instantiation: the learned forward model f , the variational posterior q , and the MAP inference network n .

- Learned forward model $f(z, \psi)$: this takes as input z and uses fully connected layers for the tabular features, and a 1D convolutional network with upsampling layers for the ECG.
- Variational posterior $q(z, \pi|x; \theta_q)$: this takes as input the ECG and tabular features, x . The ECG is passed through a 1D convolution residual network, and the tabular features through a two layer fully connected network. These representations are concatenated and passed through additional fully connected layers to produce mean and variance estimates for the posterior of each z_m and Beta concentration parameter estimates for the posterior of π .
- MAP inference network $n(x; \theta_n)$: this is a neural network with the same architecture as the variational posterior network, except that it directly outputs the MAP estimates of z and π , rather than distributional parameters.

Further architectural and training details are in the appendix.

4.6 Experiment Details

Performance Baselines: We train a network with the same architecture as $n(x; \theta_n)$ to predict only the class label without the supporting evidence. We use this to investigate whether simultaneously learning to predict and supporting evidence impacts the quality of the prediction. We use SENN [27] as a second baseline,

which both predicts and provides information designed to supplement the prediction.

Ablation: As an ablation, we consider a variant of LPS that we call LPS- q . Here, instead of training a separate MAP inference network $n(x; \theta_n)$, we take the variational posterior $q(z, \pi|x; \theta_q)$ and use the mode of this posterior to obtain MAP estimates for z and π . The modes of the respective log normal and Beta distributions have simple analytical forms in terms of the distributional parameters that are output by the variational posterior network, so are easy to compute. This simpler variant does not require training a separate MAP inference network.

Evaluation: We use the median/half IQR for ten runs, splitting the dataset into ten train/validation/test sets (60%/20%/20%), and Welch's t -test for statistical significance.

4.7 Results

Table 1 summarises the results on the clinical dataset. While we focus on analysing the supporting evidence, we also find that LPS, SENN, and the baseline perform comparably in terms of AUC, and that LPS- q performs worse than LPS, justifying the use of the MAP inference network. This is an expected result, since the variational posterior must model the entire posterior distribution over latents, and not just the modes. Thus, it may sacrifice accuracy in recovering the modes to better represent the distribution as a whole.

Accuracy of supporting evidence. We first compare the accuracy of the supporting information for the $\sim 80\%$ of patients that have CO measurements. We compare how the estimated CO compares to the measured CO by splitting CO into two groups using a cutoff of 4 L/min, which corresponds to the lower limit of normal range for CO [15] – a clinical standard that informs practice. Computing the resulting F1 score, we observe that LPS estimates do a good job of differentiating between patients who have low and normal cardiac outputs (Table 1). Since CO is usually estimated using invasive procedures and is important in clinical decision making, the fact that LPS can (non-invasively) identify when CO is above/below a meaningful threshold is clinically valuable.

Too few patients have measurements of the other latent variables to enable a meaningful direct comparison. In lieu of doing this, we use the known forward model to reconstruct both the heart rate $\overline{HR}$ (from T_s and T_d) and the blood pressures $\overline{BP}$ from the estimated latent values. We observe high coefficient of determination between the reconstructed quantities and the measured values for HR and BP (Table 1). This suggests effective recovery of the latent parameters.

As further analysis, Figure 4 visualises the true blood pressures and heart rate, and the reconstructed estimates from the forward model using the inferred latent concepts from the model. As can be seen, these show good agreement, indicating successful recovery of the latent factors. The median absolute error in recovering quantities was: BP-systolic: 1.62 mmHg; BP-diastolic: 2.92 mmHg; HR: 4.75 bpm.

Consistency of predictions and supporting evidence. Figure 5 shows histograms for the latent concepts for patients at high and low predicted risk (top and bottom 25% respectively). On average, patients at the highest risk of adverse outcomes have a lower cardiac

Method	Thresholded CO F1 Score	$R^2(HR, \widehat{HR})$	$R^2(BP, \widehat{BP})$	AUC
Baseline	N/A	N/A	N/A	0.75 ± 0.02
SENN	N/A	N/A	N/A	0.73 ± 0.02
LPS	0.78 ± 0.01	0.83 ± 0.01	0.90 ± 0.03	0.74 ± 0.02
LPS- q	0.75 ± 0.04	0.78 ± 0.04	0.71 ± 0.02	0.70 ± 0.01

Table 1: LPS predictions and supporting evidence are accurate. Analysing the supporting evidence factor CO shows that LPS effectively estimates when this value is below/above a meaningful clinical threshold. Using the forward model to reconstruct the heart rate $\widehat{HR}$ and blood pressure $\widehat{BP}$ from the supporting evidence, the reconstructions capture most of the variance of the true values, suggesting that LPS has captured useful information in this evidence space. In addition, LPS predictions are accurate (comparable AUC to the baseline and SENN [27]). Bolded values are significant at $p < 0.05$.

Figure 4: Reconstructing observed tabular features using the forward model. We visualise the true blood pressures and heart rate vs. the reconstructed blood pressures and heart rate calculated from the inferred latent concepts using the forward model. The true and reconstructed quantities show good agreement, indicating successful recovery of the latent factors.

Figure 5: Supporting evidence concepts and risk predictions have distributions that are in alignment with domain understanding. Empirical distributions (histograms) of clinically meaningful factors for patients in the upper and lower quartiles of predicted risk are in accordance with clinical domain knowledge. Such agreement is important for supporting evidence and predictions to be trusted by clinicians [42, 45].

output and higher systemic vascular resistance relative to those who do not have adverse outcomes [44]. This is recovered by LPS, achieving consistency of predictions and supporting evidence with clinical domain knowledge.

Comparing supporting evidence. Figure 6 shows supporting evidence for model decisions from LPS and attributions from Integrated Gradients [43], a commonly used feature attribution method, on the baseline. Normal ranges for the concepts were derived from literature [19, 39, 44]. We focus on showing data for CO and R , since these are most meaningful to clinicians and there are therapeutic

interventions for modifying them [49]. We compare LPS and Integrated Gradients for two patients, one at high risk of death (top), and one at low risk of death (bottom).

LPS produces supporting evidence that is clinically understandable, specifically that the patient shown at the top is at high risk and has elevated R and low CO . These statements give insights beyond the feature space alone. In contrast, the baseline of Integrated Gradients reveals certain features in the tabular data and the ECG that contributed to the decision, but the attributions are not as readily actionable as the supporting evidence from LPS. For example, the patient in Figure 6 had a normal heart rate (HR) of 80 bpm (normal

Figure 6: Supporting evidence from LPS is comprehensible and provides actionable insights. In comparison to attribution methods such as Integrated Gradients, LPS produces supporting evidence that is clinically meaningful and provide insights beyond the input feature space. For a high risk patient (top), LPS produces actionable supporting evidence, namely that that *CO* and *R* (which are hard to observe and important in therapeutic decisions) lie outside their normal ranges. For a low risk patient (bottom), LPS recovers supporting evidence factors within their normal ranges. This supporting evidence for a prediction is actionable and less ambiguous than feature attribution methods such as Integrated Gradients, especially when applied to a high-dimensional input such as the ECG.

Figure 7: LPS supporting evidence could provide actionable insights on challenging, borderline cases. The patient shown had an adverse outcome ($y = 1$), but was predicted by LPS to be at low risk – a misclassification. Analysing the weak supporting evidence for this decision adds more insight.

range: 60-100 bpm), yet the HR feature contributed significantly to the prediction of high risk. It is unlikely that knowing that the HR was a factor in the model's prediction would lead to any action on the part of a clinician.

For the low risk patient (bottom), LPS's supporting evidence for a prediction of low risk is accompanied by the inference that the patient has normal *CO* and *R*. Integrated gradients analysis reveals that most of the tabular features and the QRS complex of the ECG contribute to a low risk prediction, but the HR, recorded as 60 bpm, elevates the patient's risk. For a clinician who is not an ML practitioner, it may be challenging to disentangle these different factors and thus understand/trust the model's predictions using the Integrated Gradients attribution.

LPS supporting evidence could provide therapeutic insight on challenging examples. In a case where LPS makes a misclassification on a hard example, Figure 7, the weakness of the supporting evidence could be clinically informative. LPS predicted that this patient is at low risk, yet they had an adverse outcome (a misclassification). The supporting evidence is that the patient had normal range values for *R* and *CO*, which are only just inside the normal range. This could inform the clinician that although the patient was predicted to be at low risk, the patient might benefit from additional monitoring. Such explanatory insights are challenging to obtain from existing explainability methods.

Patient-specific insights from supporting evidence. Figure 8 compares supporting evidence for a pair of patients predicted to be at high risk. The predictions and supporting evidence are clinically meaningful: on the left, the patient has high *R*, and on the right,

Figure 8: LPS supporting evidence offers patient-specific insights to clinicians. For two patients who were predicted to be at high risk ($\hat{\pi} > 0.8$) LPS supporting evidence captures different factors contributing to risk, potentially informing clinicians of the most suitable patient-specific medical interventions.

the patient has low CO, both of which are typically indicative of poor cardiovascular health [15, 44]. LPS supporting evidence is well-differentiated on a per-patient basis and could inform a clinician that the left patient could benefit from medication to reduce R, and the right patient could benefit from medication to increase CO.

5 CONCLUSION

To assist human experts in decision making, machine learning models should produce both accurate predictions and supporting evidence for these predictions. In healthcare, this consideration is particularly important since clinicians draw significantly on medical principles in their decision making, and therefore can act most effectively on predictions that are accompanied by clinically relevant supporting evidence.

To tackle this problem, we propose a method, *Learning to Predict with Supporting Evidence (LPS)*, to construct models that provide both predictions and supporting evidence using clinically-relevant concepts. We demonstrate that LPS produces accurate predictions and comprehensible supporting evidence for predictions on a real-world medical dataset.

LPS relies on domain knowledge to inform (1) the choice of concepts for supporting evidence, and (2) how these concepts are related to the observed data and the prediction. In medicine, there is a tremendous amount of domain knowledge of this form, obtained from decades of medical practice and modelling. This includes probabilistic models for imaging and disease progression, physiological models for signals, and more [7, 8, 48]. LPS could be applied to other prediction problems by leveraging such medical domain knowledge, helping to further the trustworthiness and actionability of machine learning models for healthcare.

ETHICS STATEMENT

This work represents an initial step towards improving trustworthiness in ML models using supporting evidence. More testing of the method would be necessary prior to any real-world deployment. If the supporting evidence produced by the method ends up misleading care givers it could lead to poor decisions. It is important that prior to any deployment, extensive user studies are performed in order to detect such issues and prevent potential negative impact.

ACKNOWLEDGEMENTS

This work was supported in part by funds from Quanta Computer, Inc. The authors thank the members of the Clinical and Applied Machine Learning group and the Computational Cardiovascular Research group at MIT for all their helpful comments and advice.

REFERENCES

- [1] Maruan Al-Shedivat, Avinava Dubey, and Eric P Xing. 2017. Contextual explanation networks. *arXiv preprint arXiv:1705.10301* (2017).
- [2] Elliott M Antman, Marc Cohen, Peter JLM Bernink, Carolyn H McCabe, Thomas Horacek, Gary Papuchis, Branco Mautner, Ramon Corbalan, David Radley, and Eugene Braunwald. 2000. The TIMI risk score for unstable angina/non-ST elevation MI: a method for prognostication and therapeutic decision making. *Jama* 284, 7 (2000), 835–842.
- [3] J Bajorat, R Hofmoeckel, DA Vagts, M Janda, B Pohl, C Beck, and G Noeldge-Schomburg. 2006. Comparison of invasive and less-invasive techniques of cardiac output measurement under different haemodynamic conditions in a pig model. *European journal of anaesthesiology* 23, 1 (2006), 23–30.
- [4] Emelia J Benjamin, Paul Muntner, Alvaro Alonso, Marcio S Bittencourt, Clifton W Callaway, April P Carson, Alanna M Chamberlain, Alexander R Chang, Susan Cheng, Sandeep R Das, et al. 2019. Heart disease and stroke Statistics-2019 update a report from the American Heart Association. *Circulation* (2019).
- [5] JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, M West, et al. 2003. The variational Bayesian EM algorithm for incomplete data: with application to scoring graphical model structures. (2003).
- [6] Carrie J Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattenberg, Fernanda Viegas, Greg S Corrado, Martin C Stumpe, et al. 2019. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [7] Marianne Catanho, Mridu Sinha, and Varsha Vijayan. 2012. Model of aortic blood flow using the windkessel effect. (2012).
- [8] Adrian V Dalca, Evan Yu, Polina Golland, Bruce Fischl, Mert R Sabuncu, and Juan Eugenio Iglesias. 2019. Unsupervised deep learning for Bayesian brain MRI segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 356–365.
- [9] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 1–12.
- [10] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. 2019. Towards automatic concept-based explanations. In *Advances in Neural Information Processing Systems*. 9273–9282.
- [11] Christopher B Granger, Robert J Goldberg, Omar Dabbous, Karen S Pieper, Kim A Eagle, Christopher P Cannon, Frans Van de Werf, Alvaro Avezum, Shaun G Goodman, Marcus D Flather, et al. 2003. Predictors of hospital mortality in the global registry of acute coronary events. *Archives of internal medicine* 163, 19 (2003), 2345–2353.
- [12] Bart Hiemstra, Geert Koster, Renske Wiersema, Yoran M Hummel, Pim van der Harst, Harold Snieder, Ruben J Eck, Thomas Kaufmann, Thomas WL Scheeren, Anders Perner, et al. 2019. The diagnostic accuracy of clinical examination for estimating cardiac index in critically ill patients: the Simple Intensive Care Studies-I. *Intensive care medicine* 45, 2 (2019), 190–200.
- [13] Michael Hind, Dennis Wei, Murray Campbell, Noel CF Codella, Amit Dhurandhar, Aleksandra Mojsilović, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. 2019. TED: Teaching AI to explain its decisions. In *Proceedings of the 2019*

- AAAI/ACM Conference on AI, Ethics, and Society. 123–129.
- [14] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frdo Durand. 2019. DiffTaichi: Differentiable Programming for Physical Simulation. *arXiv preprint arXiv:1910.00935* (2019).
 - [15] J Hurst, C Rackley, E Sonnenblick, and N Wenger. 1990. *The heart, arteries and veins*. Vol. 1. McGraw-Hill.
 - [16] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *International Conference on Machine Learning*. 2668–2677.
 - [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
 - [18] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
 - [19] Mary E Klingensmith et al. 2008. *The Washington manual of surgery*. Lippincott Williams & Wilkins.
 - [20] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept Bottleneck Models. *arXiv preprint arXiv:2007.04612* (2020).
 - [21] Todd Kulesza, Simone Stumpf, Margaret Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too much, too little, or just right? Ways explanations impact end users' mental models. In *2013 IEEE Symposium on Visual Languages and Human Centric Computing*. IEEE, 3–10.
 - [22] Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. *arXiv preprint arXiv:1606.04155* (2016).
 - [23] Andrew S Levey, Lesley A Stevens, Christopher H Schmid, Yaping Zhang, Alejandro F Castro III, Harold I Feldman, John W Kusek, Paul Eggers, Frederick Van Lente, Tom Greene, et al. 2009. A new equation to estimate glomerular filtration rate. *Annals of internal medicine* 150, 9 (2009), 604–612.
 - [24] Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. 2018. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
 - [25] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*. 4765–4774.
 - [26] Robert L McNamara, Kevin F Kennedy, David J Cohen, Deborah B Diercks, Mauro Moscucci, Stephen Ramee, Tracy Y Wang, Traci Connolly, and John A Spertus. 2016. Predicting in-hospital mortality in patients with acute myocardial infarction. *Journal of the American College of Cardiology* 68, 6 (2016), 626–635.
 - [27] David Alvarez Melis and Tommi Jaakkola. 2018. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*. 7775–7784.
 - [28] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1–38.
 - [29] David A Morrow, Elliott M Antman, Lori Parsons, James A de Lemos, Christopher P Cannon, Robert P Giugliano, Carolyn H McCabe, Hal V Barron, and Eugene Braunwald. 2001. Application of the TIMI risk score for ST-elevation MI in the National Registry of Myocardial Infarction 3. *Jama* 286, 11 (2001), 1356–1359.
 - [30] Paul D. Myers, Wei Huang, Fred Anderson, and Collin M. Stultz. 2019. Choosing Clinical Variables for Risk Stratification Post-Acute Coronary Syndrome. *Scientific Reports* 9, 1 (2019), 14631. <https://doi.org/10.1038/s41598-019-50933-3>
 - [31] Paul D Myers, Kenney Ng, Kristen Severson, Uri Kartoun, Wangzhi Dai, Wei Huang, Frederick A Anderson, and Collin M Stultz. 2020. Identifying unreliable predictions in clinical risk models. *NPJ digital medicine* 3, 1 (2020), 1–8.
 - [32] Radford M Neal and Geoffrey E Hinton. 1998. A view of the EM algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*. Springer, 355–368.
 - [33] Gregory Plumb, Maruan Al-Shedivat, Eric Xing, and Ameet Talwalkar. 2019. Regularizing Black-box Models for Improved Interpretability. *arXiv preprint arXiv:1902.06787* (2019).
 - [34] Sushravya Raghunath, Alvaro E Ulloa Cerna, Linyuan Jing, Joshua Stough, Dustin N Hartzel, Joseph B Leader, H Lester Kirchner, Martin C Stumpe, Ashraf Hafez, Arun Nemani, et al. 2020. Prediction of mortality from 12-lead electrocardiogram voltage data using a deep neural network. *Nature Medicine* (2020), 1–6.
 - [35] Carl Edward Rasmussen. 2003. Gaussian processes in machine learning. In *Summer School on Machine Learning*. Springer, 63–71.
 - [36] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on International Conference on Machine Learning-Volume 32*. II–1278.
 - [37] Matthew T Roe, Anita Y Chen, Laine Thomas, Tracy Y Wang, Karen P Alexander, Bradley G Hammill, W Brian Gibler, E Magnus Ohman, and Eric D Peterson. 2011. Predicting long-term mortality in older patients after non-ST-segment elevation myocardial infarction: the CRUSADE long-term mortality model and risk score. *American heart journal* 162, 5 (2011), 875–883.
 - [38] Kiichi Sagawa, Reidar K Lie, and Jochen Schaefer. 1990. Translation of Otto Frank's paper "Die Grundform des arteriellen Pulses" zeitschrift fr biologische 483–526 (1899). *Journal of molecular and cellular cardiology* 22, 3 (1990), 253–254.
 - [39] N Saouti, N Westerhof, PE Postmus, and A Vonk-Noordegraaf. 2010. The arterial load in pulmonary hypertension. *European Respiratory Review* 19, 117 (2010), 197–203.
 - [40] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034* (2013).
 - [41] Mervyn Singer, Clifford S Deutschman, Christopher Warren Seymour, Manu Shankar-Hari, Djillali Annane, Michael Bauer, Rinaldo Bellomo, Gordon R Bernard, Jean-Daniel Chiche, Craig M Coopersmith, et al. 2016. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA* 315, 8 (2016), 801–810.
 - [42] Collin M Stultz. 2019. The Advent of Clinically Useful Deep Learning. *JACC. Clinical electrophysiology* 5, 5 (2019), 587.
 - [43] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 3319–3328.
 - [44] Thenappan Thenappan, Kurt W Prins, Marc R Pritzker, John Scandurra, Karl Volmers, and E Kenneth Weir. 2016. The critical role of pulmonary arterial compliance in pulmonary hypertension. *Annals of the American Thoracic Society* 13, 2 (2016), 276–284.
 - [45] Sana Tonekaboni, Shalmali Joshi, Melissa D McCradden, and Anna Goldenberg. 2019. What Clinicians Want: Contextualizing Explainable Machine Learning for Clinical End Use. In *Machine Learning for Healthcare Conference*. 359–380.
 - [46] Eric J Topol. 2019. High-performance medicine: the convergence of human and artificial intelligence. *Nature medicine* 25, 1 (2019), 44–56.
 - [47] Jean-Louis Vincent and Daniel De Backer. 2013. Circulatory shock. *New England Journal of Medicine* 369, 18 (2013), 1726–1734.
 - [48] Nico Westerhof, Jan-Willem Lankhaar, and Berend E Westerhof. 2009. The arterial windkessel. *Medical & biological engineering & computing* 47, 2 (2009), 131–141.
 - [49] Clyde W Yancy, Mariell Jessup, Biykem Bozkurt, Javed Butler, Donald E Casey, Mark H Drazner, Gregg C Fonarow, Stephen A Geraci, Tamara Horwich, James L Januzzi, et al. 2013. 2013 ACCF/AHA guideline for the management of heart failure: executive summary: a report of the American College of Cardiology Foundation/American Heart Association Task Force on practice guidelines. *Journal of the American College of Cardiology* 62, 16 (2013), 1495–1539.
 - [50] Matthew D Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European conference on computer vision*. Springer, 818–833.