

Think as a Doctor: An Interpretable AI Approach for ICU Mortality Prediction

Qingwen Li^{1,**}, Xiaohang Zhao^{1,**}, Xiao Han^{2,*}, Hailiang Huang^{1,*}, Lanjuan Liu¹

¹ School of Information Management & Engineering,

Shanghai University of Finance and Economics, Shanghai, China

² Key Laboratory of Data Intelligence and Management (Beihang University),

Ministry of Industry and Information Technology, School of Economics and Management,

Beihang University, Beijing, China

* Corresponding Author: Xiao Han, xh_bh@buaa.edu.cn; Hailiang Huang, hlhuang@shufe.edu.cn

** Equal Contribution

Abstract: Intensive Care Unit (ICU) mortality prediction, which estimates a patient’s mortality status at discharge using electronic health records (EHRs) collected early in an ICU admission, is vital in high-stakes critical care. For this task, predictive accuracy alone is insufficient; interpretability is equally essential for building clinical trust and meeting regulatory standards, a topic that has attracted significant attention in information system research. Accordingly, an ideal solution should enable intrinsic interpretability and align its reasoning process with three key elements of the ICU decision-making practices: clinical course identification, demographic heterogeneity, and prognostication awareness. However, conventional interpretable healthcare prediction approaches largely focus on demographic heterogeneity, overlooking clinical course identification and prognostication awareness. Recent prototype learning methods address clinical course identification, yet the integration of demographic heterogeneity and prognostication awareness into such frameworks remains underexplored. To address these gaps, we propose ProtoDoctor, a novel ICU mortality prediction framework that delivers intrinsic interpretability while integrating all three elements of the ICU decision-making practices into its reasoning process. Methodologically, ProtoDoctor features two key innovations: the Prognostic Clinical Course Identification module and the Demographic Heterogeneity Recognition module. The former enables the identification of clinical courses via prototype learning and achieves prognostication awareness using a novel regularization mechanism. The latter models demographic heterogeneity through cohort-specific prototypes and risk adjustments. Extensive empirical evaluations demonstrate that ProtoDoctor outperforms state-of-the-art baselines in predictive accuracy. Human evaluations further confirm that its interpretations are more clinically meaningful, trustworthy, and applicable in ICU practice.

Keywords: Health Information Technology, Healthcare Predictive Analytics, ICU Mortality Prediction, Interpretable AI, Prototype Learning

1. Introduction

Fueled by the widespread adoption of electronic health records (EHRs), Information Systems (IS) researchers have tackled numerous healthcare predictive analytics tasks, ranging from readmission prediction to adverse event prediction using advanced IT solutions, particularly those based on artificial intelligence (AI) techniques (Xie et al. 2021, Lin and Fang 2021, Guo et al. 2024). One important healthcare prediction task is Intensive Care Unit (ICU) mortality prediction, which uses EHR data from the early hours of an ICU admission to predict whether a critically ill patient will survive the hospital stay. This prediction task addresses a significant healthcare need, as in the United States alone, over five million patients are admitted to ICUs each year, and ICU-related expenditures exceed \$100 billion annually (Kannan et al. 2023). Given the life-threatening conditions of ICU patients and the significant consumption of medical resources, accurate ICU mortality prediction is essential for multiple stakeholders. For patients and their families, reliable mortality risk assessments enable timely and informed treatment decisions. For clinicians and hospital administrators, objective mortality predictions support data-driven care planning and efficient ICU resource allocation. Indeed, one study found that deploying an accurate ICU risk model reduced length of stay and yielded approximately \$1091 in savings per patient (Ericson et al. 2022). At the macro level, certain AI-based tools could generate substantial financial benefits for the healthcare system, with estimated annual savings ranging from \$55 million to \$72 million (Na et al. 2023).

Given the importance of the task, numerous methods have been developed for ICU mortality prediction. Traditional approaches rely on scoring systems that quantify abnormalities in patients’ physiological variables recorded in EHRs. Popular examples include the Simplified Acute Physiology Score (SAPS) (Le Gall et al. 1984), the Acute Physiology and Chronic Health Evaluation (APACHE) (Zimmerman et al. 2006), and the Sequential Organ Failure Assessment (SOFA) (De Mendonça et al. 1996). Although easy to interpret, these systems have difficulty in capturing complex and evolving temporal patterns in EHRs, and therefore suffer from limited predictive accuracy and clinical utility (Salluh and Soares 2014). In recent years, deep learning (DL) techniques have substantially improved predictive performance by modeling intricate temporal patterns in EHR data (Morid et al. 2023). These methods use architectures such as recurrent, convolutional, and attention-based neural networks to integrate dynamic physiological measurements with static demographic information (Harutyunyan et al. 2019, Caicedo-Torres

and Gutierrez 2019, Ma et al. 2020). However, these performance gains often come at the cost of interpretability: complex DL models generally function as “black boxes”, which undermines clinician trust and fails to comply with emerging regulatory requirements for algorithmic transparency, e.g., the European Union (EU)’s General Data Protection Regulation (GDPR) (Voigt and Von dem Bussche 2017).

To maintain predictive performance of DL models without compromising interpretability in the high-stakes ICU environment, an ideal solution is to develop an intrinsically interpretable DL model whose reasoning logic closely aligns with the ICU decision-making practices. An illustration of the ICU decision-making practices is provided in Figure 1, where the EHR for an ICU admission consists of the patient’s demographic attributes and physiological sequence. First, faced with the complexity inherent in the multivariate physiological sequence, where multiple interdependent variables evolve over time, a doctor assesses the mortality risk by identifying the existing strengths of various **prototypical clinical courses** within the sequence (Bhattacharyay et al. 2023). Each prototypical clinical course represents a typical temporal pattern of physiological features and offers important insight into the patient’s mortality risk (Wang et al. 2020). For instance, the prototypical clinical course C_1 in Figure 1 is characterized by persistently low Glasgow Coma Scale (GCS) scores (a sign of deep coma) along with extended periods of severely reduced respiratory rates (evidence of respiratory failure). This clinical course is associated with high mortality risk as it exemplifies critical deterioration in neurological and pulmonary systems that strongly predict adverse ICU outcomes (Fathi et al. 2022). A strong presence of the clinical course C_1 in the patient’s physiological sequence in Figure 1 thus indicates the patient is facing a high mortality risk. This clinical course identification procedure offers a structured approach to analyzing and interpreting patients’ physiological data in relation to mortality risk, presenting a promising way of enhancing both the predictive performance and interpretability of the model. Also, the doctor estimates the implied mortality risk of each clinical course while accounting for **demographic heterogeneity**, which means the same clinical course may be associated with different risks across demographic cohorts. Specifically, the doctor first assigns the patient to a cohort based on the patient’s demographic attributes and then assesses how the implied mortality risk of each clinical course should be adjusted according to the assigned cohort. This step is illustrated via the dashed red and blue arrows in Figure 1, which indicate that the clinical course C_1 poses a higher mortality risk for older male patients compared to other cohorts (Gonçalves-Pereira et al. 2023). By integrating the identified existing strengths of prototypical clinical courses with cohort-specific risk implications, the doctor forms a comprehensive and personalized assessment of the patient’s mortality status at discharge. The incorporation of demographic heterogeneity into the model’s reasoning process enables it to




Figure 1 Illustration of the ICU decision-making practices.

account for cohort-specific risk adjustments, leading to more clinically meaningful interpretations and more accurate, personalized mortality predictions. Moreover, a key component of the ICU decision-making practices is **prognostication**, i.e., the forward-looking process of estimating a patient’s near-future health state from historical observations during admission (Rencic et al. 2020). Here, health state refers to the patient’s overall clinical condition at a given time, reflecting the combined severity of illness across multiple physiological systems. Unlike the previous two components, prognostication is future-oriented, focusing on anticipating how the health state will evolve rather than assessing or adjusting current risks. Since the mortality status at discharge, i.e., the prediction target in ICU mortality prediction, is closely related to the patient’s future health state, we embed this forward-looking perspective into model reasoning to enhance predictive performance. In conclusion, grounding the model’s reasoning process in the ICU decision-making practices yields two key advantages. First, incorporating clinical insights from these practices into model design enables state-of-the-art predictive accuracy, consistent with established IS practices of integrating domain knowledge to enhance model performance (Abbasi et al. 2018, He et al. 2019, Lee and Ram 2024, Zhang et al. 2025). Second, the alignment of the model’s intrinsic interpretability with the ICU decision-making practices increases transparency and trustworthiness. This alignment strengthens practitioner acceptance, facilitates clinical adoption, and ultimately amplifies real-world impact in critical care (Bienefeld et al. 2023).

Motivated by the above analysis, our research objective is to build an effective and intrinsically interpretable DL model for ICU mortality prediction, whose reasoning process should incorporate

three key elements of the ICU decision-making practices: clinical course identification, demographic heterogeneity, and prognostication awareness. By capturing these crucial aspects, state-of-the-art predictive performance can be achieved along with clinically meaningful interpretations. However, these three essential elements are seldom considered simultaneously and integrated with intrinsic interpretability in existing interpretable healthcare prediction methods, regardless of whether they are post hoc or intrinsic. On the one hand, post hoc methods, while partially addressing demographic heterogeneity, only explain black-box DL models through local approximations. In this case, they do not reveal the actual reasoning logic of the model as intrinsic interpretations do, which may limit clinical trust and compromise regulatory compliance (Hyland et al. 2020, Kim et al. 2023). On the other hand, attention-based models provide intrinsic interpretations and account for demographic heterogeneity via learned attention weights (Ma et al. 2020, 2023). Yet, neither approach interprets model reasoning in terms of prototypical clinical courses, a cornerstone of the ICU decision-making practices.

One conceptually appealing approach to our objective is Prototype Learning, which offers intrinsic interpretability by interpreting predictions with reference to representative prototypes learned from training data, each reflecting a class-specific typical pattern. In our context of ICU mortality prediction, the key concepts of prototypical clinical courses and demographic cohorts can be recognized as distinct sets of prototypes to support interpretation. Nevertheless, incorporating critical elements of the ICU decision-making practices into prototype learning introduces significant challenges. First, existing approaches generally operate on a single data modality, such as images (Chen et al. 2019), texts (Hong et al. 2023), or univariate time series (Ming et al. 2019). In contrast, our task requires predicting and interpreting ICU mortality status based on both multivariate physiological sequences and static demographic attributes within EHRs. Yet, prototype learning models tailored to handle this multimodal data structure remain underexplored. Second, conventional prototype learning methods typically capture only the influence of a set of unimodal prototypes. When accounting for demographic heterogeneity, however, it becomes essential to model interactions between clinical course prototypes and demographic cohort prototypes, which originate from different modalities. This need reflects a universal principle in clinical practice: the same medical condition may have different implications depending on patient demographics (Gabler et al. 2009, Varadhan et al. 2013). However, the design of principled mechanisms to capture and integrate such cross-modal prototype interactions remains an open problem. Finally, the forward-looking nature of prognostication indicates that historical health states should provide predictive cues about future states. According to the clinical course identification process adopted by ICU doctors to initiate assessments, the health state at each period is represented as the existing strengths of different

prototypical clinical courses reflected in the physiological data. From this perspective, the prognostication process suggests that these clinical-course-based representations should be predictable from preceding observations. Nevertheless, existing prototype learning frameworks offer no explicit means of operationalizing such temporal predictability.

To address the challenges, we propose **ProtoDoctor**, a novel prototype learning model for ICU mortality prediction. Unlike existing interpretable healthcare prediction models and current prototype learning methods, ProtoDoctor provides intrinsic interpretations while integrating all three essential elements of the ICU decision-making practices: clinical course identification, demographic heterogeneity, and prognostication awareness. To achieve this, ProtoDoctor builds on the prototype learning paradigm and introduces two methodological contributions: the Prognostic Clinical Course Identification (PCCI) module and the Demographic Heterogeneity Recognition (DHR) module. Specifically, the PCCI module learns clinical course prototypes for clinical course identification and incorporates a novel regularization strategy to capture prognostication awareness. The DHR module accounts for demographic heterogeneity by introducing demographic cohort prototypes and a novel interaction mechanism. Extensive empirical evaluations demonstrate that ProtoDoctor outperforms both existing interpretable healthcare prediction models and state-of-the-art prototype learning approaches adapted to this task. Human evaluation further validates that the interpretations generated by ProtoDoctor are more clinically meaningful and helpful for ICU practice compared to established interpretable baselines.

2. Literature Review

Two streams of research are closely related to our study: existing methods for interpretable healthcare prediction and state-of-the-art prototype learning models that are adaptable for interpretable ICU mortality prediction. In this section, we first review each stream of related studies and then summarize the key novelties of our work.

2.1. HIT and ICU Mortality Prediction

Our study falls within the health information technology (HIT) domain and the related area of healthcare predictive analytics (HPA) in IS research (Baird et al. 2020). HIT encompasses “*an array of technologies to store, share, and analyze health information*”. The proliferation of EHRs has yielded vast amounts of clinical data, creating opportunities for IS scholars to develop advanced IT artifacts to address diverse HPA challenges (Morid et al. 2023) such as adverse event prediction (Lin and Fang 2021), length-of-stay prediction (Wang et al. 2022, Guo et al. 2024), and readmission prediction (Bardhan et al. 2015, Xie et al. 2021). In this study, we focus on the critical task of ICU mortality prediction and contribute to the HIT domain with a novel and interpretable solution

for this significant task. Early approaches to ICU mortality prediction utilized prognostic scoring systems to analyze patients’ clinical features in EHRs. Examples included widely used scores like the SAPS (Le Gall et al. 1984), the APACHE (Zimmerman et al. 2006), and the SOFA (Vincent et al. 1996). Although these scoring systems deliver transparent, interpretable frameworks for ICU practitioners, they are fundamentally calibrated for population-level outcome estimation rather than individualized prediction. Their reliance on static measurements and omission of dynamic, context-specific variables prevents them from capturing the distinct temporal patterns present in each patient’s EHR. This limitation directly undermines their predictive precision at the individual level, thereby constraining their applicability to patient-specific decision-making in critical care. To address these limitations, recent studies have leveraged deep learning (DL) models to capture nuanced, patient-specific temporal patterns in longitudinal EHR data, resulting in substantial improvements in prediction accuracy (Morid et al. 2023, Harutyunyan et al. 2019, Bednarski et al. 2022, Yu et al. 2024). These studies formulate ICU mortality prediction as a binary classification problem and train a classifier from historical ICU admission records, each comprising EHR data and the corresponding mortality status at discharge. The trained classifier can then be adopted to predict the mortality status at discharge for new ICU admissions. For instance, Harutyunyan et al. (2019) compared multiple Recurrent Neural Network (RNN)-based approaches to capture temporal patterns in EHRs for ICU mortality prediction and concluded that the multitask channel-wise LSTM method achieved the best performance. Bednarski et al. (2022) proposed a novel patient representation method that captures multi-scale temporal patterns within EHRs using temporal convolutional networks with different kernel sizes. SMART (Yu et al. 2024) addressed the challenge of missing EHR data using an attention mechanism. Collectively, these DL-based approaches have consistently outperformed traditional scoring systems in ICU mortality prediction.

While these DL-based approaches have achieved substantial gains in predictive accuracy, their limited interpretability poses a major barrier to adoption in high-stakes ICU settings. Accordingly, research on interpretable DL models for healthcare prediction has emerged along two main directions: post-hoc interpretable methods and intrinsically interpretable approaches. Post-hoc methods explain a black-box predictor by locally approximating its input–output behavior with a simpler, more interpretable surrogate model (Caicedo-Torres and Gutierrez 2019, Hyland et al. 2020, Thorsen-Meyer et al. 2020, Qiu et al. 2022, Kim et al. 2023). Popular post-hoc techniques, such as DeepLIFT (Caicedo-Torres and Gutierrez 2019) and SHAP (Hyland et al. 2020, Thorsen-Meyer et al. 2020, Qiu et al. 2022), generated time-sensitive importance scores for individual clinical features, which indicate how strongly each feature at a given time point influences the predicted mortality risk and, in some cases, specify whether the influence increases or decreases the risk. More recently, Kim et al. (2023) proposed ROLEX, a post-hoc

framework that provides robust patient-level explanations in formats like feature importance scores and local decision rules. While post-hoc explanations can be helpful, they may not faithfully reflect the actual reasoning process of the underlying model. This concern highlights the need for intrinsically interpretable approaches that make their reasoning logic transparent by design (Rudin 2019).

In response, a growing body of work has focused on intrinsically interpretable methods, which aim to reveal the model’s reasoning logic through specifically designed model components. Many recent approaches achieve this by incorporating attention mechanisms, which highlight important input features by assigning them higher attention weights (Choi et al. 2016, Ma et al. 2017, 2020, 2023, Guo et al. 2024). Accordingly, features with higher attention weights are interpreted as having a greater influence on the predicted outcome. For example, RETAIN (Choi et al. 2016) employed a two-level attention architecture to determine which past hospital visits and which clinical variables within those visits are most relevant for predicting outcomes. Similarly, Dipole (Ma et al. 2017) applied an attention-based bidirectional RNN to quantify the importance of prior visits for diagnosis prediction. Moreover, when incorporating demographic data within EHRs, recent studies, such as ConCare (Ma et al. 2020), AICare (Ma et al. 2023), and the framework proposed by Guo et al. (2024), extended attention mechanisms to model interactions between static demographics and dynamic physiological data. For instance, ConCare quantified the contributions of feature interactions by modeling their interdependencies through a self-attention-based RNN. These approaches not only improve predictive accuracy but also enhance interpretability by revealing how features jointly influence outcomes.

Despite these advantages, attention-based methods face several notable limitations. First, Serrano and Smith (2019) demonstrated that attention weights do not necessarily indicate a model’s true reasoning logic. In other words, higher attention weights assigned to input features do not guarantee stronger correlations with the predicted class, which limits the reliability of attention as an interpretation mechanism. Moreover, prior IS research has emphasized that model interpretations should mirror the way domain experts communicate insights (Kim and Benbasat 2006, Abbasi et al. 2018, Lee and Ram 2024). As outlined in the Introduction, effective ICU mortality prediction requires embedding three essential elements (i.e., clinical course identification, demographic heterogeneity, and prognostication awareness) into the model’s reasoning process. While attention mechanisms can partially address demographic heterogeneity, they fall short of simultaneously integrating all three elements. This gap has fueled growing interest in prototype learning, an emerging class of intrinsically interpretable methods with the potential to incorporate all these elements into model reasoning.

2.2. Prototype Learning

Prototype learning offers a promising path towards building intrinsically interpretable models that can embed all three critical elements of the ICU decision-making practices directly into a model’s reasoning. A pioneering study introduced ProtoPNet, a novel prototype-based framework for interpretable image recognition (Chen et al. 2019). ProtoPNet incorporates a prototype layer after the convolutional layers, where each prototype encodes a representative part of a specific class. By matching image embeddings at different locations to these prototypes, the model provides transparent reasoning for its predictions. For example, users can understand that a bird is classified as a “clay-colored sparrow” because its head shape and wing pattern closely resemble those of the learned prototypes for that species. Since the introduction of ProtoPNet, prototype learning has also been extended to support intrinsic interpretability in various scenarios such as video classification (Trinh et al. 2021), graph classification (Zhang et al. 2022, Seo et al. 2024), federated learning (Tan et al. 2022), and reinforcement learning (Kenny et al. 2023, Ragodos et al. 2022). In parallel, various enhancements have been proposed to further improve interpretability. For example, Nauta et al. (2021) combined decision trees with prototype learning to enable hierarchical reasoning. Rymarczyk et al. (2022) introduced ProtoPShare, which allows some attributes to be shared across prototypes to reduce redundancy. Bontempelli et al. (2023) proposed ProtoPDebug, an effective concept-level debugger that incorporated human feedback into the prototype generation process, offering a promising tool for trustworthy interactive learning.

In the context of ICU mortality prediction, where dynamic physiological sequences and static demographic attributes serve as inputs, prototype learning methods for interpretable sequential modeling are particularly relevant to our study (Ming et al. 2019, Ni et al. 2021, Kuang et al. 2023, Xie et al. 2025, Hong et al. 2023, Xie et al. 2024). For instance, Ming et al. (2019) introduced ProSeNet, the first prototype-based interpretable framework for sequential data. It makes predictions by comparing an input time-series with a set of class-associated prototypical sequences. Similarly, Xie et al. (2024) proposed PahNet, a prototype-based model for ECG classification, where each prototype represents a typical ECG sequence associated with a particular cardiac condition. In both cases, prototypes are interpreted as representative complete sequences observed in the training data. To enable more detailed interpretations, some studies define prototypes at a finer granularity. Specifically, SCNpro (Ni et al. 2021) derived prototypes from local segments, enabling more detailed interpretations in time-series prediction tasks. ProtoryNet (Hong et al. 2023) interpreted document classification through trajectories of sentence-level prototypes, linking each sentence to a representative prototype to improve interpretability. More recently, Xie et al. (2025) introduced a dual-level prototype framework that

models temporal prototype progressions, thereby facilitating interpretable depression detection from sensor data. Building on these advances, prototype learning can be adapted in our setting to represent core concepts such as clinical courses and demographic cohorts, which provides a foundation for embedding the critical elements of the ICU decision-making practices into model reasoning.

2.3. Key Novelties of Our Study

As discussed before, an ideal solution for interpretable ICU mortality prediction should achieve intrinsic interpretability while embedding three critical elements into its reasoning process: clinical course identification, demographic heterogeneity, and prognostication awareness. Our literature review shows that existing interpretable prediction methods all fall short in some of these aspects, as listed in Table 1. Post-hoc methods, while capable of accounting for demographic heterogeneity, provide only local approximations of a DL model’s input–output behavior and therefore lack intrinsic interpretability. Attention-based models can offer intrinsic interpretations and capture demographic heterogeneity through learned attention weights. However, neither post-hoc nor attention-based approaches can analyze and interpret the mortality risks implied by patients’ physiological data in a manner that mimics the clinical course identification process adopted by ICU doctors. Prototype-based methods are conceptually appealing for achieving intrinsic interpretability while potentially incorporating all three elements into model reasoning. However, several challenges remain unresolved. First, although prototypes can be used to represent key concepts, e.g., clinical courses, to embed clinical course identification within model reasoning, it remains unclear how to apply prototype learning to multimodal inputs in our context of ICU mortality prediction. Second, it remains an open problem to capture and integrate cross-modal prototype interactions to effectively address demographic heterogeneity. Moreover, none of the above methods incorporates prognostication awareness into model reasoning. To bridge these gaps, we propose ProtoDoctor, a novel prototype-based framework for ICU mortality prediction. It achieves intrinsic interpretability while integrating all three essential elements into model reasoning, thus achieving state-of-the-art predictive accuracy and delivering clinically meaningful interpretations.

3. Problem Formulation

In this section, we formally define the interpretable ICU mortality prediction problem. Let U denote the set of ICU admissions. The clinical record of each ICU admission u consists of the patient’s EHR and the corresponding in-hospital mortality status during the admission, denoted as $X^{(u)}$ and $y^{(u)}$, respectively. Since a patient may experience multiple ICU admissions, this study focuses on predicting the mortality status for each individual admission. Specifically, the mortality status at

Table 1 Comparison between Our Method and Existing Interpretable Methods.

Category	Representative Methods	Intrinsic Interpretability	ICU Reasoning Alignment		
			Clinical Course Identification	Demographic Heterogeneity	Prognostication Awareness
Post-hoc	SHAP (Hyland et al. 2020)			✓	
	ROLEX (Kim et al. 2023)			✓	
Attention-based	Concare (Ma et al. 2020)	✓		✓	
	Alcare (Ma et al. 2023)	✓		✓	
Prototype-based	ProtoryNet (Hong et al. 2023)	✓	✓		
	PahNet (Xie et al. 2024)	✓	✓		
ProtoDoctor (Ours)		✓	✓	✓	✓

discharge $y^{(u)}$ is a binary variable, where $y^{(u)} = 1$ indicates death and $y^{(u)} = 0$ indicates survival at discharge. The EHR $X^{(u)}$ encompasses the patient’s demographic attributes $\mathcal{D}^{(u)}$ and physiological sequence $\mathcal{P}^{(u)}$, i.e., $X^{(u)} = \langle \mathcal{D}^{(u)}, \mathcal{P}^{(u)} \rangle$. The details of these components are described below.

Demographic attributes. A patient’s demographic attributes capture essential demographic information that is often highly relevant to the patient’s final mortality status during an ICU admission. One established attribute is age, as older patients typically face higher mortality risks. Given an admission u , the demographic attributes is denoted as $\mathcal{D}^{(u)} = \{\mathcal{D}_1^{(u)}, \dots, \mathcal{D}_{n_{\mathcal{D}}}^{(u)}\}$, where $\mathcal{D}_k^{(u)}$ is the k -th demographic attribute and $n_{\mathcal{D}}$ is the number of demographic attributes.

Physiological sequence. The physiological sequence contains continuously monitored data of critical physiological variables such as heart rate and glucose levels during the initial stage of an ICU admission. More specifically, for each admission u , the physiological sequence is recorded as a time series of physiological features over the first T hours of the admission, denoted as $\mathcal{P}^{(u)} = [\mathcal{P}_1^{(u)}, \dots, \mathcal{P}_T^{(u)}]$, where $\mathcal{P}_t^{(u)} = [\mathcal{P}_{t,1}^{(u)}, \dots, \mathcal{P}_{t,n_{\mathcal{P}}}^{(u)}]$ collects the readings of $n_{\mathcal{P}}$ physiological variables at hour t .

Accordingly, we denote the overall dataset as $\mathbf{D} = \{(X^{(u)}, y^{(u)}) \mid u = 1, 2, \dots, |U|\}$, where $|U|$ is the total number of admissions. We now formally define the interpretable ICU mortality prediction problem as follows.

Definition 1 (Interpretable ICU Mortality Prediction). Given a dataset $\mathbf{D}$ containing clinical records of historical ICU admissions, our objective is to develop a model that can accurately predict the mortality status at discharge $\hat{y}^{(v)}$ for a new admission v based on the corresponding EHR $X^{(v)}$, while providing clinically meaningful interpretations as discussed in the Introduction.




Figure 2 ProtoDoctor Architecture

4. Method

In this section, we propose **ProtoDoctor** (Prototypical Demographic-and-Clinical-Course-aware Network), a novel framework for interpretable ICU mortality prediction. Figure 2 shows the overall architecture of our method, which is built upon three major components: the Prognostic Clinical Course Identification (PCCI) module, the Demographic Heterogeneity Recognition (DHR) module, and the Mortality Predictor. Given a patient’s EHR, the PCCI module identifies the presence of prototypical clinical courses by first learning a set of clinical course prototypes and then quantifying their existing strengths within the patient’s physiological sequence. To enhance model performance, this module incorporates a novel regularization term that emulates the ICU doctors’ prognostication process. The DHR module accounts for demographic heterogeneity by adjusting the mortality risks associated with different clinical courses according to the patient’s attributed demographic cohort. It learns demographic cohort prototypes, assigns each patient to a specific cohort, and adjusts mortality risks accordingly. This novel cohort-based adjustment mechanism thus enables more personalized predictions and clinically meaningful interpretations. Finally, the Mortality Predictor predicts the patient’s mortality status at discharge by integrating information on the attributed demographic cohort, the existing strengths

of prototypical clinical courses, and the adjusted mortality risks associated with those clinical courses. The PCCI module and the DHR module constitute the major methodological novelties of our method. In the following discussion, we drop the symbol u from most notations by focusing on the prediction of mortality status at discharge for one particular ICU admission. For convenience, we summarize key notations in Table 2.

Table 2 Notation

Notation	Description
$\mathcal{P}$	The physiological sequence recorded for one ICU admission, where $\mathcal{P}_t$ is the physiological features at hour t .
$\mathcal{D}$	The patient’s demographic attributes associated with one ICU admission.
$h_t^{\mathcal{P}}$	The embedding vector of physiological sequence $\mathcal{P}$ up to hour t .
$h^{\mathcal{D}}$	The embedding vector of demographic attributes $\mathcal{D}$.
$p_k^{\mathcal{C}}$	The prototype vector representing clinical course k .
$p_m^{\mathcal{D}}$	The prototype vector representing demographic cohort m .
$\tilde{s}_t^{\mathcal{C}}$	The health state vector at hour t , where $\tilde{s}_{t,k}^{\mathcal{C}}$ is the existing strength of clinical course k up to hour t , defined by Equation 2.
$\hat{s}_t^{\mathcal{C}}$	The predicted health state vector at hour t , defined by Equation 4.
$\tilde{s}^{\mathcal{D}}$	The demographic similarity vector, where $\tilde{s}_m^{\mathcal{D}}$ denotes the patient’s matching degree towards demographic cohort m , defined by Equation 7
$s^{\mathcal{D}}$	The sparse demographic attribution vector, defined by Equation 9.

4.1. Prognostic Clinical Course Identification Module

Given a patient’s multivariate physiological sequence, it is challenging to analyze and interpret it for ICU mortality prediction in a manner consistent with the ICU decision-making practices. We address this challenge by drawing on the well-established concept of the clinical course.

Definition 2 (Clinical course). The clinical course is an abstract concept describing the time-dependent progression of a patient’s medical condition. In the context of ICU mortality prediction, it represents a characteristic temporal pattern formed by the joint evolution of multiple physiological variables, offering critical insights into the patient’s mortality risk (Marino and Sutin 1998).

This definition underscores the importance of modeling temporal patterns across the entire physiological sequence rather than analyzing isolated observations independently. In practice, ICU doctors initiate mortality risk assessment by identifying prototypical clinical courses within

the sequence (Wang et al. 2020, Bhattacharyay et al. 2023). To embed this procedure into model reasoning, we represent the patient’s health state at each period as a vector indicating the existing strengths of different prototypical clinical courses in the physiological sequence. We refer to this as the health state vector. Additionally, a critical step in the ICU decision-making practices is prognostication, where doctors infer patients’ near-future health states based on historical observations. Inspired by the forward-looking nature of this step, we aim to incorporate prognostication awareness into model reasoning. By doing so, we expect the health state vector to better reflect near-future health states, thus facilitating ICU mortality prediction.

To implement these ideas, the PCCI module first learns a set of clinical course prototypes, where each prototype represents a prototypical clinical course associated with either death or survival at discharge. The patient’s health state at each period can then be represented as a vector of existing strengths of these prototypes to facilitate both prediction and interpretation. Finally, we introduce a novel regularization mechanism to incorporate prognostication awareness within model reasoning, which further enhances predictive performance on ICU mortality prediction. The overall architecture of the PCCI module is shown in Figure 3.

4.1.1. Clinical Course Identification. Given the physiological sequence $\mathcal{P} = [\mathcal{P}_1, \dots, \mathcal{P}_T]$ at the first T hours of an ICU admission, we adopt a physiological encoder $R_{\mathcal{P}}$ to generate hourly embeddings $H^{\mathcal{P}} = [h_1^{\mathcal{P}}, \dots, h_T^{\mathcal{P}}]^T \in R^{T \times n_{\mathcal{P}}}$, where $h_t^{\mathcal{P}} \in R^{n_{\mathcal{P}}}$ is the embedding vector of physiological sequence $\mathcal{P}$ up to hour t and $n_{\mathcal{P}}$ is the embedding dimension size. We define $H^{\mathcal{P}}$ as:

$$\begin{aligned} H^{\mathcal{P}} &= [h_1^{\mathcal{P}}, \dots, h_T^{\mathcal{P}}] \\ &= R_{\mathcal{P}}([\mathcal{P}_1, \dots, \mathcal{P}_T]). \end{aligned} \tag{1}$$

The physiological encoder $R_{\mathcal{P}}$ is an RNN layer based on the architecture from Harutyunyan et al. (2019), with specifications detailed in the Benchmark Methods section.

To identify what typical clinical courses are exhibited within the patient’s physiological sequence, we first learn a set of $N_{\mathcal{C}}$ clinical course prototypes, each representing a prototypical clinical course associated with death or survival at discharge. These prototypes, denoted as $P^{\mathcal{C}} = \{p_1^{\mathcal{C}}, \dots, p_{N_{\mathcal{C}}}^{\mathcal{C}}\}$, are learned as trainable parameters and embedded in the same latent space as the physiological embedding vector $h_t^{\mathcal{P}}$. Each prototype vector $p_k^{\mathcal{C}}$ for $k = 1, \dots, N_{\mathcal{C}}$ is interpreted and visualized by the closest physiological sequence in the embedding space, as described in 4.5. Then, we define $\tilde{s}_{t,k}^{\mathcal{C}}$ to represent the existing strength of clinical course k up to hour t . This value is computed as the similarity between the patient’s physiological embedding vector, $h_t^{\mathcal{P}}$, and the corresponding prototype vector $p_k^{\mathcal{C}}$:

$$\tilde{s}_{t,k}^{\mathcal{C}} = \sigma(\varphi \cdot \cos(h_t^{\mathcal{P}}, p_k^{\mathcal{C}})), \tag{2}$$




Figure 3 Architecture of the Prognostic Clinical Course Identification Module

where $\cos(x, y) = x^T y / (\|x\| \cdot \|y\|)$ is the cosine similarity between the input vectors x and y , $\sigma(z) = 1 / (1 + \exp(-z))$ is the sigmoid function, and φ is a positive scaling factor (e.g., $\varphi = 5$). The definition of $\tilde{s}_{t,k}^c$ ensures that its value falls within $(0, 1)$ with a higher value indicating a stronger existence of clinical course k up to hour t . The health state vector $\tilde{s}_t^c$ at hour t is then defined as the collection of similarity scores with respect to all clinical courses. Namely,

$$\tilde{s}_t^c = [\tilde{s}_{t,1}^c, \dots, \tilde{s}_{t,N_c}^c]. \quad (3)$$

By aggregate these vectors $[\tilde{s}_1^c, \dots, \tilde{s}_T^c]$ over time, we can capture the temporal progression of the patient’s health state from the perspective of clinical courses.

4.1.2. Prognostication-aware Regularization. To further improve model performance, we propose a novel regularization mechanism inspired by the prognostication step in the ICU decision-making practices:

Definition 3 (Prognostication). *Prognostication* is a forward-looking step in the ICU decision-making practices, where doctors analyze a patient’s historical records to estimate the near-future health state. Situated cognition theory (Rencic et al. 2020) further requires that doctors should “continually collecting and interpreting patient data, prognosticating, ..., to improve a patient’s well-being” given the emerging clinical data in the ICU environment.

The definition indicates that a patient’s near-future health state can be inferred from historical observations. Also, this forward-looking process should be continuously updated in response to emerging clinical data. We incorporate prognostication awareness into model reasoning by enforcing that the patient’s health state at each period is predictable from preceding states. This constraint guides the model to capture temporal dependencies that are indicative of future conditions, thereby improving the accuracy of ICU mortality prediction.

To operationalize this predictability constraint, we introduce an RNN-based prognostic encoder R_C to predict the health state vector $\hat{s}_t^C$ based on all preceding health state vectors $\tilde{s}_{1:t-1}^C$:

$$\hat{s}_t^C = R_C(\tilde{s}_{1:t-1}^C), t = 2, \dots, T. \quad (4)$$

We then impose a regularization term that minimizes the discrepancy between the predicted and observed health state vectors:

$$L_p^C = \frac{1}{T-1} \sum_{t=2}^T \|\hat{s}_t^C - \tilde{s}_t^C\|_2^2, \quad (5)$$

where $\hat{s}_t^C$ and $\tilde{s}_t^C$ denote the predicted and observed health state vectors at hour t , respectively. Guided by the prognostication process, our method performs step-ahead prediction on the interpretable health state vector derived from clinical course prototypes. Unlike existing healthcare prediction studies that enhance prediction by introducing auxiliary tasks on raw features or discrete clinical events (e.g., decompensation prediction) (Harutyunyan et al. 2019, Lee et al. 2019), we focus on a clinically meaningful representation in the latent space. This representation is less sensitive to noise in raw features and better captures the subtle evolution of patient conditions, thereby providing a more robust and effective regularization mechanism for ICU mortality prediction.

4.2. Demographic Heterogeneity Recognition Module

The DHR module is designed to capture demographic heterogeneity, a well-documented phenomenon in which the same clinical course can imply different mortality risks across patient cohorts (Vallet et al. 2023). To achieve this, the module introduces demographic cohort prototypes and generates an interpretable demographic attribution vector for each patient, mirroring how ICU doctors assign patients to specific cohorts. It then employs a novel interaction component to produce cohort-specific mortality risk adjustments, enabling the same clinical course to carry different implications depending on the patient’s cohort. These adjustments yield personalized mortality risks associated with different clinical courses for each patient.

4.2.1. Demographic Cohort Attribution. Given the patient’s demographic attributes $\mathcal{D}$, we first leverage a demographic encoder $R_{\mathcal{D}}$ to generate a corresponding embedding vector $h^{\mathcal{D}} \in R^{n_{\mathcal{D}}}$, where $n_{\mathcal{D}}$ is the embedding dimension size:

$$h^{\mathcal{D}} = R_{\mathcal{D}}(\mathcal{D}). \quad (6)$$

$R_{\mathcal{D}}$ is implemented as an MLP layer, with details provided in the Benchmark Methods section. To attribute the patient to a specific demographic cohort, we first learn a set of $N_{\mathcal{D}}$ demographic cohort prototypes, denoted as $P^{\mathcal{D}} = \{p_1^{\mathcal{D}}, \dots, p_{N_{\mathcal{D}}}^{\mathcal{D}}\}$. Each prototype vector $p_m^{\mathcal{D}} \in R^{n_{\mathcal{D}}}$ ($m = 1, \dots, N_{\mathcal{D}}$) is

a trainable vector representing a typical demographic cohort. To quantify how closely a patient matches demographic cohort m , we compute the similarity $\tilde{s}_m^{\mathcal{D}}$ between the patient’s demographic embedding vector $h^{\mathcal{D}}$ and the prototype vector $p_m^{\mathcal{D}}$:

$$\tilde{s}_m^{\mathcal{D}} = \sigma(\varphi \cdot \cos(h^{\mathcal{D}}, p_m^{\mathcal{D}})), \quad (7)$$

where $\tilde{s}_m^{\mathcal{D}} \in (0, 1)$, and higher values indicate that the patient exhibits a stronger match towards demographic cohort m . The similarity scores are then assembled into the demographic similarity vector:

$$\tilde{\mathbf{s}}^{\mathcal{D}} = [\tilde{s}_1^{\mathcal{D}}, \dots, \tilde{s}_{N_{\mathcal{D}}}^{\mathcal{D}}]. \quad (8)$$

To attribute the patient to a single demographic cohort, we apply the sparse transformation from ProtoryNet (Hong et al. 2023) to the similarity vector, producing the sparse demographic attribution vector $\mathbf{s}^{\mathcal{D}}$:

$$\mathbf{s}^{\mathcal{D}} = \text{Softmax}(\tilde{\mathbf{s}}^{\mathcal{D}}/\tau, \dots, \tilde{\mathbf{s}}_{N_{\mathcal{D}}}^{\mathcal{D}}/\tau), \quad (9)$$

with τ being a small constant (e.g., 10^{-6}) that ensures $\mathbf{s}^{\mathcal{D}}$ approximates a one-hot vector. The prototype corresponding to the highest similarity score in $\tilde{\mathbf{s}}^{\mathcal{D}}$ is then selected as the activated demographic cohort prototype, representing the patient’s assigned demographic cohort.

4.2.2. D-C Interaction. To account for the heterogeneous mortality risks associated with a single clinical course across different demographic cohorts, we introduce a novel *D-C Interaction* component, which explicitly integrates demographic heterogeneity into ProtoDoctor’s reasoning process. Specifically, we define a learnable interaction contribution matrix $\Delta\mathbb{W}^{\mathcal{I}} \in \mathbb{R}^{N_{\mathcal{C}} \times N_{\mathcal{D}}}$, where $N_{\mathcal{C}}$ and $N_{\mathcal{D}}$ denote the number of clinical course prototypes and demographic cohort prototypes, respectively. Each element $\Delta w_{k,m}^{\mathcal{I}}$ represents the mortality risk adjustment for clinical course k when the patient belongs to demographic cohort m :

$$\Delta\mathbb{W}^{\mathcal{I}} = \begin{bmatrix} \Delta w_{1,1}^{\mathcal{I}} & \Delta w_{1,2}^{\mathcal{I}} & \cdots & \Delta w_{1,N_{\mathcal{D}}}^{\mathcal{I}} \\ \Delta w_{2,1}^{\mathcal{I}} & \Delta w_{2,2}^{\mathcal{I}} & \cdots & \Delta w_{2,N_{\mathcal{D}}}^{\mathcal{I}} \\ \vdots & \vdots & \ddots & \vdots \\ \Delta w_{N_{\mathcal{C}},1}^{\mathcal{I}} & \Delta w_{N_{\mathcal{C}},2}^{\mathcal{I}} & \cdots & \Delta w_{N_{\mathcal{C}},N_{\mathcal{D}}}^{\mathcal{I}} \end{bmatrix}. \quad (10)$$

Each row of the $\Delta\mathbb{W}^{\mathcal{I}}$ thus encodes how mortality risk adjustments for a specific clinical course vary across demographic cohorts. Given a patient’s demographic attribution vector $\mathbf{s}^{\mathcal{D}}$, the personalized risk adjustment vector $\Delta\mathbf{W}^{\mathcal{I}} \in \mathbb{R}^{N_{\mathcal{C}}}$ is computed as:

$$\Delta\mathbf{W}^{\mathcal{I}} = \Delta\mathbb{W}^{\mathcal{I}} \mathbf{s}^{\mathcal{D}}. \quad (11)$$

This vector specifies patient-specific mortality risk adjustments for each clinical course, conditioned on the attributed demographic cohort (if $s_m^{\mathcal{D}} \approx 1$ then $\Delta W^{\mathcal{I}}$ is contributed almost entirely by the m th column of $\Delta \mathbb{W}^{\mathcal{I}}$). These adjustments are subsequently combined with the base mortality risks to produce personalized ICU mortality risk predictions. Details of this integration are provided in Section 4.3.

4.3. Mortality Predictor

The Mortality Predictor estimates the patient’s mortality status at discharge by integrating three components: the health state vector at the final hour ($\tilde{s}_T^{\mathcal{C}}$, denoted as $s^{\mathcal{C}}$ for simplicity); the demographic attribution vector ($s^{\mathcal{D}}$); and the risk adjustment vector ($\Delta W^{\mathcal{I}}$). These components jointly produce two complementary mortality risk scores, $z^{\mathcal{C}}$ and $z^{\mathcal{D}}$, as illustrated in Figure 2. The first score $z^{\mathcal{C}}$ captures the patient’s clinical-course-related mortality risks by combining the existing strengths of different clinical courses within the patient’s physiological sequence with the personalized mortality risk associated with each clinical course. Formally:

$$\begin{aligned} z^{\mathcal{C}} &= \tilde{W}^{\mathcal{C}} \cdot s^{\mathcal{C}} \\ &= (\mathbb{W}^{\mathcal{C}} + \Delta W^{\mathcal{I}}) \cdot s^{\mathcal{C}} \\ &= \sum_{k=1}^{N_{\mathcal{C}}} (\mathbb{W}_k^{\mathcal{C}} + \Delta W_k^{\mathcal{I}}) \cdot s_k^{\mathcal{C}}. \end{aligned} \tag{12}$$

where $\mathbb{W}^{\mathcal{C}} \in R^{N_{\mathcal{C}}}$ is a learnable weight vector with each element $\mathbb{W}_k^{\mathcal{C}}$ representing the base mortality risk implied by clinical course k , while $s^{\mathcal{C}}$ reflects the existing strengths of different clinical courses at the final hour T . The additive term $\Delta W^{\mathcal{I}}$ incorporates cohort-specific adjustments, enabling personalized clinical-course-based risk estimation. Notably, this score is computed as the sum of risk contributions from all $N_{\mathcal{C}}$ clinical courses. Each term $(\mathbb{W}_k^{\mathcal{C}} + \Delta W_k^{\mathcal{I}}) \cdot s_k^{\mathcal{C}}$ represents the adjusted mortality risk implied by clinical course k , weighted by its existing strength within the physiological sequence of the given patient. This decomposition clarifies how both the general population-level risks and patient-specific adjustments jointly determine the overall clinical-course-related mortality score. The second score, $z^{\mathcal{D}}$, quantifies the mortality risk from a demographic perspective:

$$\begin{aligned} z^{\mathcal{D}} &= \mathbb{W}^{\mathcal{D}} \cdot s^{\mathcal{D}} \\ &= \sum_{m=1}^{N_{\mathcal{D}}} \mathbb{W}_m^{\mathcal{D}} \cdot s_m^{\mathcal{D}}. \end{aligned} \tag{13}$$

$\mathbb{W}^{\mathcal{D}} \in R^{N_{\mathcal{D}}}$ is a learnable vector where each element represents the mortality risk associated with a demographic cohort. The vector $s^{\mathcal{D}}$ approximates a one-hot vector, with one element typically close to one and the rest near zero. Thus, when $s_m^{\mathcal{D}} \approx 1$, the demographic risk score $z^{\mathcal{D}}$ is determined almost entirely by the corresponding risk weight $\mathbb{W}_m^{\mathcal{D}}$. Finally, the predicted mortality status $\hat{y}$ is

obtained by aggregating the clinical-course-based score z^c and the demographic-based score z^d , followed by a sigmoid activation:

$$\begin{aligned}\hat{y} &= \sigma(z^c + z^d) \\ &= \sigma(\tilde{W}^c \cdot s^c + W^d \cdot s^d)\end{aligned}\tag{14}$$

4.4. ProtoDoctor

Given the clinical record $(X^{(u)}, y^{(u)}) = (\langle \mathcal{D}^{(u)}, \mathcal{P}^{(u)} \rangle, y^{(u)})$ of an admission $u \in U$, ProtoDoctor predicts the patient’s mortality status at discharge according to the following interpretable reasoning process. The model first quantifies the existing strengths of different prototypical clinical courses within the patient’s physiological sequence at the final hour T (Equation 3). Based on the patient’s demographic attributes, it then assigns the patient to a specific demographic cohort (Equation 9). Next, personalized risk adjustments are applied for different clinical courses (Equation 11). Finally, these components are integrated to estimate the mortality status at discharge (Equation 14).

Next, we illustrate how to learn the model parameters based on the training dataset. Let Φ denote the complete set of model parameters in ProtoDoctor, we have:

$$\Phi = \{\Theta, \mathbb{P}, \mathbb{W}\}.\tag{15}$$

Here, $\Theta = \{\theta^p, \theta^c, \theta^d\}$ comprises the parameters of the physiological encoder R_p , prognostication encoder R_c , and demographic encoder R_d , as defined in Equations 1, 4 and 6, respectively. The set $\mathbb{P} = \{P^d, P^c\}$ denotes the trainable prototype vectors of demographic cohorts and clinical courses, respectively. Finally, $\mathbb{W} = \{W^d, W^c, \Delta W^I\}$ contains the weight vectors W^d and W^c , and the interaction contribution matrix ΔW^I . The full parameter set Φ is learned by minimizing the following objective function through mini-batch gradient descent:

$$\Phi^* = \arg \min_{\Phi} \text{CE} + \lambda_d^d L_d^d + \lambda_d^c L_d^c + \lambda_s^c L_s^c + \lambda_p^c L_p^c + \lambda^I |\Delta W^I|,\tag{16}$$

where CE is the standard binary cross-entropy loss computed for the entire dataset $\mathbf{D}$, and the coefficients $(\lambda_d^d, \lambda_d^c, \lambda_s^c, \lambda_p^c, \lambda^I)$ are hyperparameters used to balance the contribution of five regularization terms explained below. Specifically,

$$\text{CE} = \sum_{(X^{(u)}, y^{(u)}) \in \mathbf{D}} y^{(u)} \log \hat{y}^{(u)} + (1 - y^{(u)}) \log(1 - \hat{y}^{(u)})\tag{17}$$

where $y^{(u)}$ is the actual mortality status at discharge for ICU admission u , and $\hat{y}^{(u)}$ is the predicted mortality status computed from EHR $X^{(u)}$ in accordance with Equation 14. In terms of model

regularizations, $|\Delta \mathbb{W}^{\mathcal{I}}|$ represents the L1-norm of the interaction contribution matrix, while $L_d^{\mathcal{D}}$ and $L_d^{\mathcal{C}}$ implement the diversity penalty that promotes the discovery of diverse demographic cohorts and clinical courses by encouraging their respective prototypes to be distinct from other prototypes (Ming et al. 2019). Formally,

$$L_d^{\mathcal{D}} = \sum_{i=1}^{N_{\mathcal{D}}} \sum_{j=i+1}^{N_{\mathcal{D}}} \max(0, d_{\min}^{\mathcal{D}} - \|p_i^{\mathcal{D}} - p_j^{\mathcal{D}}\|_2)^2, \quad (18)$$

$$L_d^{\mathcal{C}} = \sum_{i=1}^{N_{\mathcal{C}}} \sum_{j=i+1}^{N_{\mathcal{C}}} \max(0, d_{\min}^{\mathcal{C}} - \|p_i^{\mathcal{C}} - p_j^{\mathcal{C}}\|_2)^2, \quad (19)$$

where $d_{\min}^{\mathcal{D}}$ and $d_{\min}^{\mathcal{C}}$ define minimum inter-prototype distances and are both set to 3 in our study. The sparsity penalty $L_s^{\mathcal{C}}$, adapted from Crabbé and Van Der Schaar (2021), limits each physiological embedding vector to be close to only a few clinical course prototypes:

$$L_s^{\mathcal{C}} = \frac{1}{|U|} \sum_{i=1}^{|U|} \|\text{sort}(s_i^{\mathcal{C}}) - \mathbf{r}_{\alpha}\|_2^2 \quad (20)$$

where $|U|$ is the number of clinical records, $\|\cdot\|_2$ denotes the L2-norm, and $\text{sort}(\cdot)$ arranges the elements of a health state vector in ascending order. The vector $\mathbf{r}_{\alpha}$ consists of $\text{int}((1-\alpha) \cdot N_{\mathcal{C}})$ zeros followed by $\text{int}(\alpha \cdot N_{\mathcal{C}})$ ones, where $\alpha \in (0, 1)$ specifies the proportion of clinical course prototypes expected to be activated for each physiological embedding. Here, $\text{int}(\cdot)$ denotes the integer part (i.e., floor operation) of its argument. In our study, α is set to 0.1. Finally, we incorporate the prognostication loss $L_p^{\mathcal{C}}$, as defined in Equation 5, to embed prognostication awareness into model reasoning, thereby enhancing ProtoDoctor’s predictive performance.

4.5. Prototype Interpretation

Following Ming et al. (2019) and Chen et al. (2019), to interpret the learned clinical course prototypes and demographic cohort prototypes, we project each clinical course prototype onto the closest physiological embedding from the training dataset:

$$p_k^{\mathcal{C}} \leftarrow h_T^{\mathcal{P}(u^*)} \text{ with } u^* = \arg \max_{\forall u \in U} \cos(h_T^{\mathcal{P}(u)}, p_k^{\mathcal{C}}) \quad (21)$$

where $h_T^{\mathcal{P}(u)}$ is computed using Equation 1 for each admission u . In this case, each clinical course prototype can be interpreted as the complete physiological sequence from a specific ICU admission in the input space.

Similarly, each demographic cohort prototype is visualized as the demographic attributes of an ICU admission in the training dataset, which is closest to the prototype in the embedding space:

$$p_m^{\mathcal{D}} \leftarrow h^{\mathcal{D}(u^*)} \text{ with } u^* = \arg \max_{\forall u \in U} \cos(h^{\mathcal{D}(u)}, p_m^{\mathcal{D}}) \quad (22)$$

where $h^{\mathcal{D}(u)}$ is generated based on Equation 6. These projections enable clinically meaningful interpretation of ProtoDoctor’s reasoning process, which aligns with the ICU decision-making practices. Further details on these interpretations are provided in Section 5.5.

5. Empirical Evaluation

5.1. Data Collection and Evaluation Metrics

To evaluate the predictive performance and interpretability of ProtoDoctor, we utilized the Medical Information Mart for Intensive Care version 3 (MIMIC-III) database (Johnson et al. 2016), which contains deidentified electronic health records (EHRs) from 61,532 ICU admissions recorded between 2001 and 2012 at a major academic medical center in Boston. Following the data preprocessing protocol established by Harutyunyan et al. (2019), we excluded all hospitalizations involving multiple ICU admissions to avoid ambiguity in labeling outcomes, since mortality in these cases could be associated with the entire hospitalization rather than a specific ICU admission. We also removed admissions for patients under 18 years old, due to substantial physiological differences between pediatric and adult populations. After applying these criteria, our final dataset for ICU mortality prediction consists of 21,139 unique ICU admissions. For feature construction, we adopted the variable selection strategies of Harutyunyan et al. (2019) and Ma et al. (2020), incorporating 20 clinical variables that include both demographic attributes and physiological features. Demographic attributes are obtained following Ma et al. (2020), while physiological features are selected according to Harutyunyan et al. (2019). By concatenating the sequences of all physiological features, we obtained a 76-dimensional physiological feature vector $\mathcal{P}_t$ at each hour t . The details are provided in Appendix A. For each ICU admission, we extracted sequences of physiological variables from the first 48 hours in addition to demographic attributes. These sequences were preprocessed according to the protocol of Harutyunyan et al. (2019), including imputation, discretization, and standardization to ensure consistency with established clinical time-series modeling benchmarks. The mortality status at discharge is defined as a binary indicator at the admission level: 1 if the patient died during the ICU admission, and 0 otherwise. Among the 21,139 admissions in the final dataset, 2,797 cases (13.23%) result in death, indicating a notable class imbalance. To construct the training and evaluation sets, we partitioned the data into training (14,681), validation (3,222), and test (3,236) sets, reflecting an approximate 70%/15%/15% split. To prevent information leakage and ensure the validity of model evaluation, all admissions from a single patient are assigned exclusively to one partition. We assessed the predictive performance of ProtoDoctor by comparing its predictions with ground-truth labels in the test dataset, using two widely adopted metrics in ICU mortality prediction research: the Area Under the Receiver Operating Characteristic Curve (AUROC) and the Area Under the

Precision-Recall Curve (AUPRC) (Ma et al. 2020, 2023). AUROC measures the model’s capacity to distinguish between mortal and non-mortal ICU admissions, where a score of 1.0 indicates perfect discrimination and 0.5 represents random guessing. AUPRC assesses the model’s ability to accurately identify mortal ICU admissions, capturing the balance between precision (correctly identified mortal admissions among those predicted as mortal) and recall (correctly identified mortal admissions among actual mortal cases). A score of 1.0 reflects perfect performance, while a value near the mortality rate reflects performance close to random guessing. Given the substantial class imbalance in our ICU mortality prediction dataset (mortality rate: 13.23%), we placed greater emphasis on AUPRC as it directly evaluates the model’s effectiveness in detecting actual mortal cases, making it more suitable for the ICU context. For each model, we reported the mean and standard deviation of AUROC and AUPRC over five experimental runs to ensure statistical robustness.

5.2. Benchmark Methods

We evaluated the performance of our proposed model, ProtoDoctor, against two primary categories of benchmark methods: (1) interpretable DL models designed for healthcare prediction tasks, and (2) recent prototype learning approaches adaptable for interpretable ICU mortality prediction. State-of-the-art interpretable healthcare prediction methods commonly employ DL models with attention mechanisms to provide intrinsic interpretability. These models attribute importance scores to input features through attention weights. Accordingly, we first included ConCare (Ma et al. 2020) as a benchmark. ConCare is an attention-based ICU mortality prediction framework that achieves state-of-the-art performance and offers interpretability by modeling cross-feature interdependencies. Then, we also evaluated AICare (Ma et al. 2023), which assigns attention weights to features at each time step. When adapted to ICU data, each attention weight generated by AICare reflects the importance of a clinical feature at a specific hour for ICU mortality prediction. While attention-based models provide interpretability at the feature level, prototype learning methods make it possible to interpret based on representative clinical courses and patient cohorts, mirroring the ICU decision-making practices more closely. Since no prior studies have applied prototype learning methods to ICU mortality prediction, we adapted two representative approaches originally developed for other tasks as benchmark models: ProtoryNet (Hong et al. 2023) and PahNet (Xie et al. 2024). In our context, ProtoryNet models each prototype as a typical static state of physiological features at a specific hour, enabling clinical course-based interpretations through sequences of state-level prototypes. In contrast, PahNet interprets each prototype as a complete physiological sequence, thus directly capturing prototypical clinical courses. To incorporate demographic information without compromising

interpretability, we augmented each prototype-based baseline with a demographic branch, utilizing the same demographic encoder and demographic cohort prototypes as employed in ProtoDoctor. Table 3 provides an overview of all benchmark methods considered in our evaluation.

Table 3 Methods Compared in Our Evaluation

Method	Notes
ProtoDoctor	Our proposed method
ConCare (Ma et al. 2020)	Representative attention-based interpretable ICU mortality prediction method
AICare (Ma et al. 2023)	Representative attention-based interpretable healthcare prediction method
ProtoryNet (Hong et al. 2023)	State-level Prototype learning method adapted for ICU mortality prediction
PahNet (Xie et al. 2024)	Sequence-level Prototype learning method adapted for ICU mortality prediction

Next, we outline the implementation details for all benchmark models. Each model processes both a dynamic physiological sequence and a vector of static demographic attributes as inputs and predicts the mortality status at discharge for each ICU admission. We followed the implementation and configuration of ConCare provided by the authors, since it utilizes the same dataset as ProtoDoctor¹. For all other baseline models, we performed hyperparameter tuning on the validation set and reported predictive performance on the test set. ProtoDoctor encompasses three specialized encoders. The physiological encoder (R_P) employs a channel-wise architecture following Harutyunyan et al. (2019), with the sequence of each physiological feature independently processed by a two-layer bidirectional LSTM (8 hidden units per layer). Feature-specific embeddings are concatenated hourly to form a unified temporal representation, which is then processed by time-wise and feature-wise attention layers. Skip connections and a dropout rate of 0.5 are applied to enhance training stability and prevent overfitting. The demographic encoder (R_D) is a single-layer MLP with 64 hidden units and Tanh activation, generating the embedding vector of demographic attributes. The prognostication encoder (R_C) operates as a single-layer LSTM whose embedding dimension size equals the number of clinical course prototypes (N_C). This encoder predicts the future health state vector based on the sequence of preceding health state vectors to embed prognostication into model reasoning. Other

¹The implementation of ConCare is based on the official code: <https://github.com/Accountable-Machine-Intelligence/ConCare>

key hyperparameters for ProtoDoctor include the number of clinical course prototypes ($N_c = 50$), the number of demographic prototypes ($N_D = 20$), and regularization coefficients ($\lambda_d^D, \lambda_d^C, \lambda_s^C, \lambda_p^C, \lambda^I$), which were set to 1×10^{-3} , 1×10^{-3} , 5×10^{-1} , 5×10^{-2} , and 1×10^{-3} , respectively. For AICare (Ma et al. 2023), the key hyperparameters include the number of hidden units in the channel-wise temporal encoder (8), the embedding size in the multi-head attention layer (32), the feed-forward size (256), and the dropout rate (0.3).² For ProtoryNet (Hong et al. 2023), which was initially developed for text classification, we replaced the original pretrained text embedding module with ProtoDoctor’s physiological encoder R_C , set the LSTM hidden size to 128, and used a dropout rate of 0.5. For PahNet (Xie et al. 2024), originally designed for ECG classification with fewer input channels, we increased the Bi-LSTM hidden units to 64 to accommodate the higher dimensionality of physiological data. For both prototype-based baselines, the numbers of clinical course prototypes and demographic prototypes are set to 50 and 20, respectively. All models were implemented in PyTorch and optimized using the Adam optimizer (Kingma and Ba 2014) with a learning rate of 1×10^{-3} . The batch size was set to 32 for ProtoDoctor and other prototype-based baselines, while it was 256 for attention-based baselines. Training was conducted for up to 100 epochs, with early stopping based on validation loss and a patience of 10 epochs to prevent overfitting.

5.3. Evaluation Results

Table 4 presents the performance metrics (mean $\pm$ standard deviation) for ProtoDoctor and baseline methods across five independent runs, with statistical significance evaluated via t-tests. Specifically, ProtoDoctor achieves the highest AUPRC (0.5417 ± 0.0101), significantly outperforming all baselines ($p < 0.10$), while attaining a competitive AUROC that is statistically indistinguishable from the top-performing ConCare model. As noted by prior studies (Harutyunyan et al. 2019, Ma et al. 2020), the observed AUPRC–AUROC gap reflects the inherent class imbalance in ICU mortality prediction, where mortal cases constitute only a small fraction of admissions. While AUROC measures overall discriminative ability between mortal and non-mortal cases, AUPRC focuses on the mortal class, making it more aligned with the primary clinical goal of accurately identifying high-risk (mortal) ICU patients. Accordingly, ProtoDoctor’s superiority in AUPRC underscores its enhanced clinical utility for ICU mortality prediction compared to existing approaches.³

²The implementation of AICare is based on the official code: <https://github.com/Accountable-Machine-Intelligence/AICare>.

³We also noted that the absolute performance values reported here are comparable to, but slightly lower than those in Ma et al. (2020), likely due to differences in evaluation protocols, such as the use of bootstrapping in prior work.

Table 4 Prediction Performance Comparison between ProtoDoctor and Baselines (T=48)

Model	AUROC	AUPRC
ProtoDoctor(Ours)	0.8668 ± 0.0013	0.5417 ± 0.0101
ConCare (Ma et al. 2020)	0.8673 ± 0.0011	$0.5095 \pm 0.0110^{***}$
AICare (Ma et al. 2023)	$0.8472 \pm 0.0077^{***}$	$0.4800 \pm 0.0119^{***}$
ProtoryNet (Hong et al. 2023)	$0.8579 \pm 0.0057^{**}$	$0.5214 \pm 0.0126^{**}$
PahNet (Xie et al. 2024)	$0.8491 \pm 0.0059^{***}$	$0.4681 \pm 0.0147^{***}$

Note: $*p < 0.10$; $**p < 0.05$; $***p < 0.01$

To ensure the robustness of our evaluation results, we conducted additional experiments by reducing the observation period T of physiological sequences from 48 to 24 hours, which is also a commonly adopted setting in prior ICU mortality prediction research (Iwase et al. 2022). As reported in Table 5, ProtoDoctor consistently achieves the highest AUPRC and the second-highest AUROC across all baseline comparisons. Specifically, ProtoDoctor outperforms the second-best method, ProtoryNet, by approximately 1.0% in AUPRC, and achieves even greater performance margins (up to 10.0%) over other baselines. These statistically significant improvements further validate the robustness of ProtoDoctor’s advantage in predictive performance. We also note that prototype learning-based methods generally achieve higher AUPRC than non-prototype approaches, corroborating their effectiveness in identifying mortal admissions, particularly when the available clinical information is limited.

Table 5 Prediction Performance Comparison between ProtoDoctor and Baselines (T=24)

Model	AUROC	AUPRC
ProtoDoctor(Ours)	0.8146 ± 0.0027	0.4000 ± 0.0062
ConCare (Ma et al. 2020)	0.8161 ± 0.0022	$0.3751 \pm 0.0155^{**}$
AICare (Ma et al. 2023)	$0.7950 \pm 0.0193^*$	$0.3639 \pm 0.0259^{**}$
ProtoryNet (Hong et al. 2023)	$0.8094 \pm 0.0038^{**}$	0.3962 ± 0.0079
PahNet (Xie et al. 2024)	$0.8057 \pm 0.0023^{***}$	$0.3764 \pm 0.0136^{**}$

Note: $*p < 0.10$; $**p < 0.05$; $***p < 0.01$

5.4. Performance Analysis

Methodologically, ProtoDoctor features two key innovations: (i) the prognostication-aware regularization (PAR) component in the PCCI module and (ii) the D-C interaction (DCI) component in the DHR module. To assess the individual contribution of each component, we

conducted ablation studies by sequentially removing these modules and evaluating the impact on predictive performance. First, we evaluated the impact of the DCI component by removing it from ProtoDoctor and referred to the resulting model as ProtoDoctor-D. The performance difference between ProtoDoctor and ProtoDoctor-D reveals the contribution of accounting for the demographic heterogeneity to the performance of ProtoDoctor. Next, we further removed the PAR component from ProtoDoctor-D and referred to the resulting variant as ProtoDoctor-DP. The performance difference between ProtoDoctor-D and ProtoDoctor-DP uncovers the contribution of prognostication awareness to ProtoDoctor’s performance. We compared the performance of ProtoDoctor-DP, ProtoDoctor-D, and ProtoDoctor under the primary evaluation setting ($T = 48$). As reported in Table 6, ProtoDoctor outperforms ProtoDoctor-D by 0.64% in AUROC and 2.01% in AUPRC, whereas ProtoDoctor-D surpasses ProtoDoctor-DP by 0.02% and 1.51% in AUROC and AUPRC, respectively. In summary, both components significantly enhance predictive performance, with particularly notable improvements in AUPRC. These findings confirm that ProtoDoctor’s methodological advances are especially effective in improving the identification of mortal cases, a critical capability given the severe class imbalance in ICU mortality prediction.

Table 6 Ablation Analysis of ProtoDoctor

	PAR	DCI	AUROC	AUPRC
ProtoDoctor-DP			0.8611±0.0021	0.5231±0.0085
ProtoDoctor-D	✓		0.8613±0.0020	0.5310±0.0025
ProtoDoctor	✓	✓	0.8668±0.0013	0.5417±0.0101

5.5. Interpretability of ProtoDoctor

Beyond outstanding predictive performance compared with state-of-the-art baselines, ProtoDoctor offers clinically meaningful model interpretability for ICU mortality prediction by identifying the exhibited clinical courses, attributing patients to demographic cohorts, and quantifying personalized mortality risk adjustments. To illustrate the model’s reasoning process, we presented two representative cases, one involving a patient who died and one who survived, as shown in Figures 4 and 5. Each figure consists of three panels: the attributed patient cohort (Panel a), the most aligned clinical course prototype (Panel b), and the mortality risks before and after incorporating the personalized mortality risk adjustment via the DCI component (Panel c). In each figure, Panel (a) summarizes the demographic attributes of the attributed cohort for each patient, including ethnicity, gender, age, height, weight, and body mass index (BMI). Panel (b)

visualizes the clinical course prototype most aligned with the patient’s physiological sequence as sequences of key physiological variables. These variables are organized by the circulatory, respiratory, and neurological systems, respectively (from top to bottom), with their hourly measurements appearing as solid lines. For variables with established normal ranges, the upper and lower bounds are shown using dashed and dotted lines, respectively. These reference lines are labeled with corresponding variable abbreviations followed by (U) for upper bound or (L) for lower bound. For example, “SBP(U)” and “SBP(L)” symbolize the upper and lower bounds of systolic blood pressure. For variables considered abnormal when exceeding a defined threshold, a horizontal dashed line indicates the threshold and is marked with the variable abbreviation followed by (T). For instance, GCS-T(T) denotes the threshold value of the Glasgow Coma Scale (GCS) total score, with values below 15 suggesting neurological impairment according to established medical guidelines (Teasdale and Jennett 1974). Panel (c) presents the predicted mortality probabilities before and after considering personalized mortality risk adjustments, demonstrating how patient-specific cohort information influences the predicted mortality status.

ProtoDoctor predicts the patient in Figure 4 as mortal based on three main reasons. First, as shown in Panel (a), the patient belongs to a demographic cohort characterized by obesity, which has been widely recognized as a risk factor linked to elevated ICU mortality (Hogue et al. 2009). Second, the patient’s physiological sequence closely matches a clinical course prototype that reflects progressive multi-organ failure, as illustrated in Panel (b). A sustained reduction in diastolic and mean blood pressures below established normal thresholds indicates worsening hypotension, a condition strongly associated with elevated mortality in critical care settings (Schuurmans et al. 2024). At the same time, key respiratory indicators such as respiratory rate and oxygen saturation demonstrate deteriorating trends consistent with respiratory failure (Garrido et al. 2018). Neurological evaluation shows persistently low GCS scores, suggesting a prolonged state of deep coma (Teasdale and Jennett 1974). Taken together, these patterns collectively reflect severe multi-organ dysfunction and suggest that the patient faces a very high risk of in-hospital mortality. Finally, cohort-specific mortality risk adjustment elevated the predicted mortality probability from 79.99% to 99.51%, as reported in Panel (c). This adjustment further emphasizes the heightened mortality risk associated with the patient’s attributed cohort.

Similarly, ProtoDoctor predicts survival for the patient in Figure 5 according to three key factors. First, as shown in Panel (a), the patient is assigned to a demographic cohort with a normal body mass index (BMI) of 20.73, a value generally associated with lower ICU mortality risk. Second, the patient’s physiological sequence closely resembles a clinical course prototype that reflects stabilization and recovery, as illustrated in Panel (b): Circulatory variables exhibit some

fluctuation but largely remain within established normal ranges, indicating preserved cardiovascular function. Respiratory measurements also stay near normal levels throughout the observation period. Although the GCS scores initially drop to an abnormal level, they gradually return to normal, signifying neurological recovery. This prototype aligns with a typical clinical course associated with patient survival. Third, the cohort-specific mortality risk adjustment further confirms the survival prediction. The estimated mortality probability remains low before and after personalized adjustment, as reported in Panel (c). This result demonstrates the model’s strong confidence in its survival assessment.




Figure 4 Interpretation of a mortal ICU patient

Furthermore, Figure 6 displays the learned interaction contribution matrix, ΔW^I . In this matrix, prototypes of demographic cohorts and clinical courses are arranged in ascending order according to their associated mortality weights derived from W^D and W^C . The direction of sorting is indicated by the vertical and horizontal arrows along the axes. A positive contribution weight $\Delta w_{k,m}^I$ indicates an elevated mortality risk adjustment of the clinical course k for demographic cohort m . The diversity of the interaction weights between demographic cohorts and clinical courses reveals ProtoDoctor’s capability to capture heterogeneous mortality risk adjustments of each clinical course associated with different demographic cohorts. Some interaction patterns show potential for clinically meaningful interpretations that may assist ICU




Figure 5 Interpretation of a surviving ICU patient




Figure 6 Interaction contribution matrix visualization

doctors in tailoring risk assessment and decision-making. These implications and their alignment with expert judgments are further explored in the Human Evaluation section.

5.6. Human Evaluation

To rigorously evaluate the interpretability of ProtoDoctor, we conducted a structured user study with a panel of 40 ICU experts (8 PhDs and 32 MDs) from a leading Emergency Department in

China. These participants have an average of 15.5 years of ICU experience. Notably, our expert panel is substantially larger than those in previous healthcare IS studies, where the number of experts rarely exceeded eight (Kim et al. 2023, Guo et al. 2024, Xie et al. 2025), and even in clinical diagnostic research, where the largest reported panel was nine (Bertens et al. 2013). This significantly larger panel strengthens the generalizability and robustness of our evaluation. Each expert independently participates in a 20-minute interview comprising two structured surveys designed to evaluate ProtoDoctor’s interpretability in terms of both the clinical reasoning process and the generated clinical statements.

In the first survey, experts are randomly assigned to review model-generated interpretations from either ProtoDoctor (Figure B1) or ProtoryNet (Figure B2), the best-performing prototype-based baseline. Each expert assesses two ICU admission cases, one correctly predicted as death and one as survival, and then rates the model’s interpretability across five dimensions: reasonability, trustworthiness, clinical meaningfulness, usefulness for ICU practice, and consistency with medical guidelines. The five questions, adapted from Xie et al. (2025), are: (1) How reasonable do you think the model’s reasoning process is? (2) How trustworthy do you think the model’s reasoning process is? (3) How clinically meaningful do you think the model’s interpretation is for ICU practices? (4) How useful do you think the model’s summarized typical demographic cohorts and clinical courses are for ICU mortality prediction? (5) How consistent do you think the model’s summarized typical clinical courses are with established medical guidelines? Ratings are collected using a 5-point Likert scale ranging from 0 (“not at all”) to 4 (“extremely”). As shown in Figure 7, ProtoDoctor consistently achieves higher mean scores than the baseline across all criteria, for example, 2.75 versus 1.55 for reasonability, and 2.65 versus 1.30 for trustworthiness. Qualitative feedback from ICU experts further highlights a key limitation of the baseline: relying solely on sequences of typical clinical states is viewed as insufficient for comprehensive mortality assessment. In contrast, ProtoDoctor’s interpretations, incorporating entire clinical courses and personalized mortality risk adjustments, are considered more aligned with the real-world ICU decision-making practices.

The second survey evaluates the clinical plausibility of statements derived from ProtoDoctor’s reasoning, specifically those describing how the adjusted mortality risk of a given clinical course varies across demographic cohorts. Participants are randomly assigned to one of two groups: one reviews two original statements generated by ProtoDoctor, and the other reviews their logical opposites. Following the evaluation format of Guo et al. (2024), experts rate each statement on a 5-point Likert scale (1 = “strongly disagree” to 5 = “strongly agree”). The complete statements and their corresponding mean scores are presented in Table B1. Results show significantly higher agreement with ProtoDoctor’s statements (e.g., 3.65 versus 2.25). Furthermore, qualitative feedback from experts supports the plausibility of the statements. For instance, the first




Figure 7 Evaluation on the Interpretation of Reasoning Process (Mean in Each Group)

statement is consistent with existing clinical guidelines regarding the interaction between obesity and hypotension in critical care, where hypotension signals hemodynamic instability and, when coupled with obesity, is associated with elevated mortality risk (Dickerson et al. 2022). Similarly, the second statement reflects the well-documented observation that elderly patients face higher stroke-related mortality (Soto-Cámara et al. 2020).

6. Conclusion

6.1. Summary and Contributions

Accurate ICU mortality prediction is critical given ICU patients’ severe conditions and the resource-intensive nature of critical care. The high-stakes environment further necessitates model interpretability to support clinical trust and ensure regulatory compliance. To enhance both predictive performance and interpretability, it is suggested that an ICU mortality prediction model should provide intrinsic interpretations and integrate three critical elements of the ICU decision-making practices into its reasoning process: clinical course identification, demographic heterogeneity, and prognostication awareness. However, while some existing interpretable healthcare models provide intrinsic interpretability, they typically only address demographic heterogeneity. Prototype learning approaches show promise for delivering intrinsic interpretations while incorporating clinical course identification; however, how to achieve demographic heterogeneity and prognostication awareness in prototype learning frameworks remains under-investigated. To address these limitations, we propose ProtoDoctor, a novel interpretable framework for ICU mortality prediction. Specifically, it generates intrinsic interpretations via prototype learning and explicitly incorporates all three essential elements using two innovative components: the prototypical clinical course identification module and the demographic heterogeneity recognition module. Extensive empirical evaluations demonstrate that ProtoDoctor not only outperforms leading interpretable healthcare prediction models and state-of-the-art

prototype learning methods in predictive performance but also offers intrinsic interpretations with enhanced interpretability in terms of various dimensions.

Our work belongs to the computational design science (CDS) research genre within the Information Systems (IS) field ([Padmanabhan et al. 2022](#)). In general, CDS research focuses on developing novel computational methods to solve important business and societal problems, thus making methodological contributions to the literature. Specifically, this study addresses a critical healthcare task of ICU mortality prediction by proposing ProtoDoctor, a novel intrinsically interpretable prediction model based on prototype learning. To enable intrinsic interpretability and incorporate critical elements in the ICU decision-making practices, the key novelties of ProtoDoctor lie in the prognostic clinical course identification module and the demographic heterogeneity recognition module, which together form the core methodological contributions of this study. Additionally, our study also contributes to research in the Health Information Technology (HIT) field by introducing a novel interpretative framework that effectively solves a fundamental problem in the high-stakes ICU settings.

6.2. Research and Managerial Implications

Our study offers several important research implications. First, we advance prototype learning by introducing a modeling approach that operates on the multimodal data structure (i.e., physiological sequence and demographic attributes within EHRs) and captures cross-modal prototype interactions. This addresses key limitations of existing prototype learning models, which typically focus on a single data source. Second, we emphasize the critical role of domain knowledge in designing computational artifacts in CDS research. ProtoDoctor’s architecture is explicitly informed by the ICU decision-making practices, which not only improves predictive performance but also enhances interpretability. This supports and extends longstanding CDS literature emphasizing the value of integrating domain expertise into IT artifact development ([Gregor and Hevner 2013](#), [Abbasi et al. 2018](#), [He et al. 2019](#)), especially in high-stakes healthcare contexts ([Wang et al. 2021](#), [Zhang et al. 2025](#)).

ProtoDoctor also brings substantial benefits to various healthcare stakeholders in high-stakes ICU settings ([Imrie et al. 2023](#)). For patients and their families, interpretability grounded in ICU clinical decision-making practices fosters a clearer understanding of mortality risk assessments, enabling more informed treatment decisions. By referring to historical ICU admissions with similar reasoning logics, patients and families can better comprehend potential outcomes and associated burdens, supporting choices that align with their preferences and values. For ICU doctors, ProtoDoctor builds trust in AI-driven predictions by aligning interpretability with clinical practice. Doctors can rapidly assess whether the model’s rationale is consistent with their

expertise. As shown in our human evaluation, ICU doctors found ProtoDoctor’s explanations more clinically meaningful than those of baseline models, supporting safer and more effective care, and promoting seamless AI adoption in ICU workflows. For healthcare regulators and administrators, the transparency inherent in ProtoDoctor provides clear insight into the model’s decision-making logic, thereby fostering clinical trust, enabling rigorous validation, and facilitating its practical deployment in ICU settings. The model interpretability in the form of the desired paradigm is also essential for meeting regulatory standards such as GDPR (Voigt and Von dem Bussche 2017), ensuring accountability in critical care decisions, and advancing stakeholder confidence in health information technology adoption.

6.3. Limitations and future Work

Our study has limitations and can be extended in multiple ways. Firstly, although ProtoDoctor currently uses only physiological sequences and demographic attributes from a single ICU admission, recent studies show that incorporating additional modalities, such as clinical notes, laboratory results, and medical images, can enhance predictive accuracy in ICU settings (Yang et al. 2021, Chen et al. 2024). Future work could extend ProtoDoctor to integrate these diverse data sources to enhance predictive performance and provide more comprehensive, clinically meaningful interpretations. Secondly, while ProtoDoctor models clinical courses by jointly capturing temporal patterns across all physiological features, ICU doctors often focus on the co-evolution of specific subsets of critical variables. For example, Fraction of Inspired Oxygen (FiO_2), Peripheral Capillary Oxygen Saturation (SpO_2), and respiratory rate (RR) are frequently monitored together, as their combined trends provide key insights into respiratory function and oxygenation (Roca et al. 2019). Future work could enhance interpretability by explicitly identifying and modeling the temporal dynamics of such clinically meaningful feature subsets. Finally, because ProtoDoctor’s interpretation paradigm closely aligns with the ICU decision-making practices, it holds significant potential for broader clinical applications. Future work could adapt this approach to other high-impact prediction tasks in critical care, such as length-of-stay estimation (Guo et al. 2024), ICU readmission prediction (Xie et al. 2021), and decompensation detection (Harutyunyan et al. 2019).

References

- Abbasi A, Zhou Y, Deng S, Zhang P (2018) Text analytics to support sense-making in social media. MIS Quarterly 42(2):427–464.
- Baird A, Angst C, Oborn E (2020) Research curation on health information technology research curation team. MIS Quarterly .

-
- Bardhan I, Oh Jh, Zheng Z, Kirksey K (2015) Predictive analytics for readmission of patients with congestive heart failure. Information Systems Research 26(1):19–39.
- Bednarski BP, Singh AD, Zhang W, Jones WM, Naeim A, Ramezani R (2022) Temporal convolutional networks and data rebalancing for clinical length of stay and mortality prediction. Scientific Reports 12(1):21247.
- Bertens LC, Broekhuizen BD, Naaktgeboren CA, Rutten FH, Hoes AW, van Mourik Y, Moons KG, Reitsma JB (2013) Use of expert panels to define the reference standard in diagnostic research: a systematic review of published methods and reporting. PLoS medicine 10(10):e1001531.
- Bhattacharyay S, Caruso PF, Åkerlund C, Wilson L, Stevens RD, Menon DK, Steyerberg EW, Nelson DW, Ercole A, investigators CT, participants (2023) Mining the contribution of intensive care clinical course to outcome after traumatic brain injury. npj Digital Medicine 6(1):154.
- Bienefeld N, Boss JM, Lüthy R, Brodbeck D, Azzati J, Blaser M, Willms J, Keller E (2023) Solving the explainable AI conundrum by bridging clinicians’ needs and developers’ goals. NPJ Digital Medicine 6(1):94.
- Bontempelli A, Teso S, Tentori K, Giunchiglia F, Passerini A, et al. (2023) Concept-level debugging of part-prototype networks. Proceedings of the The Eleventh International Conference on Learning Representations (ICLR 2023).
- Caicedo-Torres W, Gutierrez J (2019) ISeeU: Visually interpretable deep learning for mortality prediction inside the ICU. Journal of biomedical informatics 98:103269.
- Chen C, Li O, Tao C, Barnett AJ, Su J, Rudin C (2019) This looks like that: deep learning for interpretable image recognition. Proceedings of the 33rd International Conference on Neural Information Processing Systems, 8930–8941.
- Chen J, Wen Y, Pokojovy M, Tseng TLB, McCaffrey P, Vo A, Walser E, Moen S (2024) Multi-modal learning for inpatient length of stay prediction. Computers in Biology and Medicine 171:108121.
- Choi E, Bahadori MT, Kulas JA, Schuetz A, Stewart WF, Sun J (2016) RETAIN: an interpretable predictive model for healthcare using reverse time attention mechanism. Proceedings of the 30th International Conference on Neural Information Processing Systems, 3512–3520.
- Crabbé J, Van Der Schaar M (2021) Explaining time series predictions with dynamic masks. International Conference on Machine Learning, 2166–2177 (PMLR).
- De Mendonça A, Bruining H, Thijs L (1996) The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ. Intensive Care Med 22:707–710.
- Dickerson RN, Andromalos L, Brown JC, Correia MIT, Pritts W, Ridley EJ, Robinson KN, Rosenthal MD, Van Zanten AR (2022) Obesity and critical care nutrition: current practice gaps and directions for future research. Critical Care 26(1):283.

-
- Ericson O, Hjelmgren J, Sjövall F, Söderberg J, Persson I (2022) The potential cost and cost-effectiveness impact of using a machine learning algorithm for early detection of sepsis in intensive care units in Sweden. Journal of health economics and outcomes research 9(1):101.
- Fathi M, Moghaddam NM, Jame SZB, Darvishi M, Mortazavi M (2022) The association of glasgow coma scale score with characteristics of patients admitted to the intensive care unit. Informatics in Medicine Unlocked 29:100904.
- Gabler NB, Duan N, Liao D, Elmore JG, Ganiats TG, Kravitz RL (2009) Dealing with heterogeneity of treatment effects: is the literature up to the challenge? Trials 10(1):43.
- Garrido D, Assioun JJ, Keshishyan A, Sanchez-Gonzalez MA, Goubran B (2018) Respiratory rate variability as a prognostic factor in hospitalized patients transferred to the intensive care unit. Cureus 10(1).
- Gonçalves-Pereira J, Oliveira A, Vieira T, Rodrigues AR, Pinto MJ, Pipa S, Martinho A, Ribeiro S, Paiva JA (2023) Critically ill patient mortality by age: long-term follow-up (CIMbA-LT). Annals of intensive care 13(1):7.
- Gregor S, Hevner AR (2013) Positioning and presenting design science research for maximum impact. MIS Quarterly 337–355.
- Guo T, Bardhan IR, Ding Y, Zhang S (2024) An explainable artificial intelligence approach using graph learning to predict intensive care unit length of stay. Information Systems Research URL <http://dx.doi.org/10.1287/isre.2023.0029>, article in Advance.
- Harutyunyan H, Khachatrian H, Kale DC, Ver Steeg G, Galstyan A (2019) Multitask learning and benchmarking with clinical time series data. Scientific Data 6(1):1–18.
- He J, Fang X, Liu H, Li X (2019) Mobile app recommendation: An involvement-enhanced approach. MIS Quarterly 43(3):827–849.
- Hogue CW, Stearns JD, Colantuoni E, Robinson KA, Stierer T, Mitter N, Pronovost PJ, Needham DM (2009) The impact of obesity on outcomes after critical illness: a meta-analysis. Intensive Care Medicine 35:1152–1170.
- Hong D, Wang T, Baek S (2023) Protorynet-interpretable text classification via prototype trajectories. Journal of Machine Learning Research 24(264):1–39.
- Hyland SL, Faltys M, Hüser M, Lyu X, Gumbsch T, Esteban C, Bock C, Horn M, Moor M, Rieck B, et al. (2020) Early prediction of circulatory failure in the intensive care unit using machine learning. Nature medicine 26(3):364–373.
- Imrie F, Davis R, van der Schaar M (2023) Multiple stakeholders drive diverse interpretability requirements for machine learning in healthcare. Nature Machine Intelligence 5(8):824–829.
- Iwase S, Nakada Ta, Shimada T, Oami T, Shimazui T, Takahashi N, Yamabe J, Yamao Y, Kawakami E (2022) Prediction algorithm for icu mortality and length of stay using machine learning. Scientific reports 12(1):12912.

-
- Johnson AE, Pollard TJ, Shen L, Lehman LwH, Feng M, Ghassemi M, Moody B, Szolovits P, Anthony Celi L, Mark RG (2016) Mimic-iii, a freely accessible critical care database. Scientific data 3(1):1–9.
- Kannan S, Stevens J, Song Z (2023) Growth in patient cost sharing for hospitalizations with and without intensive care among commercially insured patients: study examines cost-sharing for hospitalizations with and without intensive care. Health Affairs 42(9):1221–1229.
- Kenny EM, Tucker M, Shah JA (2023) Towards interpretable deep reinforcement learning with human-friendly prototypes. The Eleventh International Conference on Learning Representations, URL https://openreview.net/forum?id=hWwY_Jq0xsN.
- Kim BR, Srinivasan K, Kong SH, Kim JH, Shin CS, Ram S (2023) Rolex: A novel method for interpretable machine learning using robust local explanations. MIS Quarterly 47(3).
- Kim D, Benbasat I (2006) The effects of trust-assuring arguments on consumer trust in internet stores: Application of toulmin’s model of argumentation. Information Systems Research 17(3):286–300.
- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .
- Kuang J, Xie J, Yan Z (2023) What symptoms and how long? an interpretable ai approach for depression detection in social media. arXiv preprint arXiv:2305.13127 .
- Le Gall JR, Loirat P, Alperovitch A, Glaser P, Granthil C, Mathieu D, Mercier P, Thomas R, Villers D (1984) A simplified acute physiology score for ICU patients. Critical care medicine 12(11):975–977.
- Lee C, Yoon J, Van Der Schaar M (2019) Dynamic-DeepHit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. IEEE Transactions on Biomedical Engineering 67(1):122–133.
- Lee K, Ram S (2024) Explainable deep learning for false information identification: An argumentation theory approach. Information Systems Research 35(2):890–907.
- Lin YK, Fang X (2021) First, do no harm: Predictive analytics to reduce in-hospital adverse events. Journal of Management Information Systems 38(4):1122–1149.
- Ma F, Chitta R, Zhou J, You Q, Sun T, Gao J (2017) Dipole: Diagnosis prediction in healthcare via attention-based bidirectional recurrent neural networks. Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining, 1903–1911.
- Ma L, Zhang C, Gao J, Jiao X, Yu Z, Zhu Y, Wang T, Ma X, Wang Y, Tang W, et al. (2023) Mortality prediction with adaptive feature importance recalibration for peritoneal dialysis patients. Patterns 4(12):100892.
- Ma L, Zhang C, Wang Y, Ruan W, Wang J, Tang W, Ma X, Gao X, Gao J (2020) Concare: Personalized clinical feature embedding via capturing the healthcare context. Proceedings of the AAAI Conference on Artificial Intelligence, volume 34, 833–840.
- Marino PL, Sutin KM (1998) The ICU book (Williams & Wilkins Baltimore).

-
- Ming Y, Xu P, Qu H, Ren L (2019) Interpretable and steerable sequence learning via prototypes. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 903–913.
- Morid MA, Sheng ORL, Dunbar J (2023) Time series prediction using deep learning methods in healthcare. ACM Transactions on Management Information Systems 14(1):1–29.
- Na L, Carballo KV, Pauphilet J, Haddad-Sisakht A, Kombert D, Boisjoli-Langlois M, Castiglione A, Khalifa M, Hebbal P, Stein B, et al. (2023) Patient outcome predictions improve operations at a large hospital network. arXiv preprint arXiv:2305.15629 .
- Nauta M, Van Bree R, Seifert C (2021) Neural prototype trees for interpretable fine-grained image recognition. Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 14933–14943.
- Ni J, Chen Z, Cheng W, Zong B, Song D, Liu Y, Zhang X, Chen H (2021) Interpreting convolutional sequence model by learning local prototypes with adaptation regularization. Proceedings of the 30th ACM International Conference on Information & Knowledge Management, 1366–1375.
- Padmanabhan B, Fang X, Sahoo N, Burton-Jones A (2022) Machine learning in information systems research. MIS Quarterly 46(1).
- Qiu W, Chen H, Dincer AB, Lundberg S, Kaeberlein M, Lee SI (2022) Interpretable machine learning prediction of all-cause mortality. Communications Medicine 2(1):125.
- Ragodos R, Wang T, Lin Q, Zhou X (2022) Protox: Explaining a reinforcement learning agent via prototyping. Advances in Neural Information Processing Systems 35:27239–27252.
- Rencic J, Schuwirth LW, Gruppen LD, Durning SJ (2020) Clinical reasoning performance assessment: using situated cognition theory as a conceptual framework. Diagnosis 7(3):213–221, URL <http://dx.doi.org/10.1515/dx-2019-0051>.
- Roca O, Caralt B, Messika J, Samper M, Sztrymf B, Hernández G, García-de Acilu M, Frat JP, Masclans JR, Ricard JD (2019) An index combining respiratory rate and oxygenation to predict outcome of nasal high-flow therapy. American journal of respiratory and critical care medicine 199(11):1368–1376.
- Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature machine intelligence 1(5):206–215.
- Rymarczyk D, Struski Ł, Górszczak M, Lewandowska K, Tabor J, Zieliński B (2022) Interpretable image classification with differentiable prototypes assignment. European Conference on Computer Vision, 351–368 (Springer).
- Salluh JI, Soares M (2014) ICU severity of illness scores: APACHE, SAPS and MPM. Current opinion in critical care 20(5):557–565.
- Schuurmans J, van Rossem BT, Rellum SR, Tol J, Kurucz VC, van Mourik N, van der Ven WH, Veelo DP, Schenk J, Vlaar AP (2024) Hypotension during intensive care stay and mortality and morbidity: a systematic review and meta-analysis. Intensive Care Medicine 50(4):516–525.

-
- Seo S, Kim S, Park C (2024) Interpretable prototype-based graph information bottleneck. Advances in Neural Information Processing Systems 36.
- Serrano S, Smith NA (2019) Is attention interpretable? Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 2931–2951.
- Soto-Cámara R, González-Bernal JJ, González-Santos J, Aguilar-Parra JM, Trigueros R, López-Liria R (2020) Age-related risk factors at the first stroke event. Journal of Clinical Medicine 9(7):2233.
- Tan Y, Long G, Liu L, Zhou T, Lu Q, Jiang J, Zhang C (2022) Fedproto: Federated prototype learning across heterogeneous clients. Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 8432–8440.
- Teasdale G, Jennett B (1974) Assessment of coma and impaired consciousness: a practical scale. The Lancet 304(7872):81–84.
- Thorsen-Meyer HC, Nielsen AB, Nielsen AP, Kaas-Hansen BS, Toft P, Schierbeck J, Strøm T, Chmura PJ, Heimann M, Dybdahl L, et al. (2020) Dynamic and explainable machine learning prediction of mortality in patients in the intensive care unit: a retrospective study of high-frequency data in electronic patient records. The Lancet Digital Health 2(4):e179–e191.
- Trinh L, Tsang M, Rambhatla S, Liu Y (2021) Interpretable and trustworthy deepfake detection via dynamic prototypes. 2021 IEEE Winter Conference on Applications of Computer Vision (WACV), 1972–1982 (IEEE Computer Society).
- Vallet H, Guidet B, Boumendil A, De Lange DW, Leaver S, Szczeklik W, Jung C, Sviri S, Beil M, Flaatten H (2023) The impact of age-related syndromes on ICU process and outcomes in very old patients. Annals of Intensive Care 13(1):68.
- Varadhan R, Segal JB, Boyd CM, Wu AW, Weiss CO (2013) A framework for the analysis of heterogeneity of treatment effect in patient-centered outcomes research. Journal of clinical epidemiology 66(8):818–825.
- Vincent JL, Moreno R, Takala J, Willatts S, De Mendonça A, Bruining H, Reinhart K, Suter P, Thijs L (1996) The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. Intensive care medicine 22(7):707–710.
- Voigt P, Von dem Bussche A (2017) The EU General Data Protection Regulation (GDPR). A Practical Guide, 1st Ed., Cham: Springer International Publishing 10(3152676):10–5555.
- Wang K, Hussain W, Birge JR, Schreiber MD, Adelman D (2022) A high-fidelity model to predict length of stay in the neonatal intensive care unit. INFORMS Journal on Computing 34(1):183–195.
- Wang X, High A, Wang X, Zhao K (2021) Predicting users’ continued engagement in online health communities from the quantity and quality of received support. Journal of the Association for Information Science and Technology 72(6):710–722.
- Wang Y, Lu X, Li Y, Chen H, Chen T, Su N, Huang F, Zhou J, Zhang B, Yan F, et al. (2020) Clinical course and outcomes of 344 intensive care patients with COVID-19. American journal of respiratory and critical care medicine 201(11):1430–1434.

-
- Xie J, Wang Z, Yu Z, Ding Y, Guo B (2024) Prototype learning for medical time series classification via human-machine collaboration. Sensors 24(8):2655.
- Xie J, Zhang B, Ma J, Zeng D, Lo-Ciganic J (2021) Readmission prediction for patients with heterogeneous medical history: A trajectory-based deep learning approach. ACM Transactions on Management Information Systems (TMIS) 13(2):1–27.
- Xie J, Zhao X, Liu X, Fang X (2025) Care for the mind amid chronic diseases: An interpretable AI approach using IoT. Management Science Fothcoming.
- Yang H, Kuang L, Xia F (2021) Multimodal temporal-clinical note network for mortality prediction. Journal of Biomedical Semantics 12(1):1–14.
- Yu Z, Xu C, Jin Y, Wang Y, Zhao J (2024) Smart: Towards pre-trained missing-aware model for patient health status prediction. Advances in Neural Information Processing Systems 37:63986–64009.
- Zhang D, Zhou L, Tao J, Zhu T, Gao G (2025) Ketch: A knowledge-enhanced transformer-based approach to suicidal ideation detection from social media content. Information Systems Research 36(1):572–599.
- Zhang Z, Liu Q, Wang H, Lu C, Lee C (2022) Protgnn: Towards self-explaining graph neural networks. Proceedings of the AAAI Conference on Artificial Intelligence, volume 36, 9127–9135.
- Zimmerman JE, Kramer AA, McNair DS, Malila FM (2006) Acute Physiology and Chronic Health Evaluation (APACHE) IV: hospital mortality assessment for today’s critically ill patients. Critical care medicine 34(5):1297–1310.

A. Description of Clinical Variables

Table A1 provides descriptions of input clinical variables and the output label. Following the preprocessing steps of Harutyunyan et al. (2019), we first impute missing values in each physiological time series using the most recent observation if available, or a pre-specified “normal” value otherwise. Normal values are provided in Harutyunyan et al. (2019). A binary mask is then added for each variable at every hour to indicate whether the value is observed or imputed. Categorical and binary variables are one-hot encoded, while numeric inputs are standardized by subtracting the mean and dividing by the standard deviation, computed per variable after imputation. After preprocessing, each physiological variable i is represented by two sequences: $\{\mu_t^{(i)}\}_{t=1}^T$, where $\mu_t^{(i)}$ is a binary mask indicating observation status at hour t , and $\{c_t^{(i)}\}_{t=1}^T$, the sequence of observed or imputed values. By concatenating these pairs across all physiological variables, we finally yield a 76-dimensional physiological feature vector per hour.

Table A1: Description of Clinical variables and Output Variable

Type	Variable name	Description and unit of measure
Output label	<i>Mortality status</i>	Binary (1=the patient is mortal at discharge)
Demographic attributes	<i>Ethnicity</i>	Categorical (0-4, de-identified ethnicity ID)
	<i>Gender^F</i>	Binary (1=patient is female)
	<i>Age</i>	Patient age in years
	<i>Height</i>	Patient height in centimeters (CM)
	<i>Weight</i>	Patient weight in kilograms (KG)
Physiological variables	<i>CRR</i>	Binary (1=Capillary refill rate of the patient is abnormal)
	<i>DBP</i>	Diastolic blood arterial pressure in millimeters of mercury (mmHg)
	<i>MBP</i>	Mean blood arterial pressure in millimeters of mercury (mmHg)
	<i>SBP</i>	Systolic blood arterial pressure in millimeters of mercury (mmHg)
	<i>FiO₂</i>	Fraction inspired oxygen (float)
	<i>GCS-E</i>	Eye response score in Glasgow Coma Scale (GCS) (1–4)
	<i>GCS-M</i>	Motor response score in GCS (1–6)
	<i>GCS-V</i>	Verbal response score in GCS (1–5)
	<i>GCS-T</i>	Total GCS score (3–15)
	<i>Glucose</i>	Blood glucose in milligrams per deciliter (mg/dL)
	<i>HR</i>	Heart rate in beats per minute
	<i>SPO₂</i>	Oxygen saturation in percentage
	<i>RR</i>	Respiration rate in breaths per minute

Type	Variable name	Description and unit of measure
	<i>Temperature</i>	Body temperature in Celsius
	<i>pH</i>	Body pH (float)

B. Results of Human Evaluation

The survey presented for human evaluation on interpretations generated by ProtoDoctor for a mortal ICU patient and a survival ICU patient is shown in Figure B1. The survey presented for human evaluation on interpretations generated by ProtoryNet for a mortal ICU patient and a survival ICU patient is shown in Figure B2. The results of the survey about medical implications generated from ProtoDoctor are shown in Table B1.

Input: Demographic attributes and hourly physiological features from the first 48 hours of an ICU admission
Output: Final predicted mortality rate (99.51%)

Model explanation includes three components:

1. As shown in Table 1A, based on the patient's demographic profile, the AI model assigns the patient to a representative cohort with similar characteristics.
2. As shown in Plot 1, the model predicts a high mortality risk because the patient's physiological features during the first 48 hours in the ICU closely resemble those of a typical clinical course associated with final mortality. This similarity is based on pattern matching of physiological changes over time. We illustrate (a) circulatory indicators, (b) respiratory indicators, and (c) neurological indicators. In each plot, the lower panel shows the patient's raw data, while the upper panel displays the most similar prototype. Different colors represent different indicators, and dashed/dotted lines indicate the upper/lower bounds of normal ranges.
3. As shown in Table 1B, the model further refines its prediction by adjusting for the patient's demographic-specific sensitivity to different clinical courses (based on Table 1A), increasing the predicted mortality probability from 79.99% to 99.51%.

Table 1A: Attributed Patient Cohort

Attribute	Value
Ethnicity	4
Gender	0
Age	66.2
Height (CM)	160
Weight (KG)	81.8
BMI	31.95

Table 1B: Mortality risk adjustment

	Predicted Mortality Probability
Before Adjustment	79.99%
After Adjustment	99.51%

Plot 1: Exhibited Clinical Course



(a) Mortal ICU patient case

Input: Demographic attributes and hourly physiological features from the first 48 hours of an ICU admission
Output: Final predicted mortality rate (0.08%)

Model explanation includes three components:

1. As shown in Table 1A, based on the patient's demographic profile, the AI model assigns the patient to a representative cohort with similar characteristics.
2. As shown in Plot 1, the model predicts a high mortality risk because the patient's physiological features during the first 48 hours in the ICU closely resemble those of a typical clinical course associated with final survival. This similarity is based on pattern matching of physiological changes over time. We illustrate (a) circulatory indicators, (b) respiratory indicators, and (c) neurological indicators. In each plot, the lower panel shows the patient's raw data, while the upper panel displays the most similar prototype. Different colors represent different indicators, and dashed/dotted lines indicate the upper/lower bounds of normal ranges.
3. As shown in Table 1B, the model further refines its prediction by adjusting for the patient's demographic-specific sensitivity to different clinical courses (based on Table 1A), decreasing the predicted mortality probability from 0.09% to 0.08%.


Table 1A: Attributed Patient Cohort

Attribute	Value
Ethnicity	4
Gender	1
Age	77.5
Height (CM)	170
Weight (KG)	59.9
BMI	20.73

Table 1B: Mortality risk adjustment

	Predicted Mortality Probability
Before Adjustment	0.09%
After Adjustment	0.08%

Plot 1: Exhibited Clinical Course




(b) Survival ICU patient case

Figure B1 Human evaluation on interpretations of a mortal and a survival ICU patient generated by ProtoDoctor

Input: Demographic attributes and hourly physiological features from the first 48 hours of an ICU admission
Output: Final predicted mortality rate (88.60%).

Model explanation includes three components:

1. As shown in Table 1, based on the patient's demographic profile, the AI model assigns the patient to a representative cohort with similar characteristics.
2. As shown in Plot 1, the high predicted mortality risk is due to strong point-by-point similarity between the patient's hourly vital signs during the first 48 hours in the ICU and the model's learned typical states. The model matches each hour's combination of indicators to its most similar prototype state. We present (a) circulatory indicators, (b) respiratory indicators, and (c) neurological indicators. In each plot, the lower panel shows the patient's raw data, while the upper panel shows the most similar prototype state identified for each hour. Different colors represent different indicators; dashed and dotted lines mark the upper and lower bounds of normal ranges.



Input: Demographic attributes and hourly physiological features from the first 48 hours of an ICU admission
Output: Final predicted mortality rate (0.01%).

Model explanation includes three components:

1. As shown in Table 1, based on the patient's demographic profile, the AI model assigns the patient to a representative cohort with similar characteristics.
2. As shown in Plot 1, the high predicted mortality risk is due to strong point-by-point similarity between the patient's hourly vital signs during the first 48 hours in the ICU and the model's learned typical states. The model matches each hour's combination of indicators to its most similar prototype state. We present (a) circulatory indicators, (b) respiratory indicators, and (c) neurological indicators. In each plot, the lower panel shows the patient's raw data, while the upper panel shows the most similar prototype state identified for each hour. Different colors represent different indicators; dashed and dotted lines mark the upper and lower bounds of normal ranges.




Figure B2 Human evaluation on interpretations of a mortal and a survival ICU patient generated by ProtoryNet

Table B1 Mean responses to different statements regarding ICU mortality predictions

Statements	Mean Score (Likert Scale)
Statements derived from the clinical implications generated by ProtoDoctor.	
1. Among patients who consistently present with preserved consciousness, low blood pressure, rapid respiration, and high oxygen supplementation requirements, those with a high body mass index (BMI) are predicted to have a higher mortality risk compared to those with a normal BMI.	3.65
2. Among patients who consistently present with deep coma, elevated systolic blood pressure, and reduced diastolic blood pressure, those of older age are predicted to have a higher mortality risk compared to younger patients.	4.15
Statements that contrast with the implications generated by ProtoDoctor.	
3. Among patients who consistently present with preserved consciousness, low blood pressure, rapid respiration, and high oxygen supplementation requirements, those with a high body mass index (BMI) are predicted to have a lower mortality risk compared to those with a normal BMI.	2.25
4. Among patients who consistently present with deep coma, elevated systolic blood pressure, and reduced diastolic blood pressure, those of older age are predicted to have a lower mortality risk compared to younger patients.	1.90