

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DELTA-XAI: A NOVEL FRAMEWORK FOR EXPLAIN- ING PREDICTION CHANGES IN ONLINE TIME SERIES MONITORING

006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Anonymous authors

Paper under double-blind review

ABSTRACT

Explaining online time series monitoring models is crucial across sensitive domains such as healthcare and finance, where temporal and contextual prediction dynamics underpin critical decisions. While recent XAI methods have improved the explainability of time series models, they mostly analyze each time step independently, overlooking temporal dependencies. This results in further challenges: explaining prediction changes is non-trivial, methods fail to leverage online dynamics, and evaluation remains difficult. To address these challenges, we propose Delta-XAI, which adapts 14 existing XAI methods through a wrapper function and introduces a principled evaluation suite for the online setting, assessing diverse aspects, such as faithfulness, sufficiency, and coherence. Experiments reveal that classical gradient-based methods, such as Integrated Gradients (IG), can outperform recent approaches when adapted for temporal analysis. Building on this, we propose Shifted Window Integrated Gradients (SWING), which incorporates past observations in the integration path to systematically capture temporal dependencies and mitigate out-of-distribution effects. Extensive experiments consistently demonstrate the effectiveness of SWING across diverse settings with respect to diverse metrics. Our code is publicly available at <https://anonymous.4open.science/r/Delta-XAI>.

1 INTRODUCTION

Time series data, inherently gleaned online, span safety- and mission-critical domains such as healthcare (Johnson et al., 2016; Reyna et al., 2020b), transportation (Polson & Sokolov, 2017; Zheng & Huang, 2020), finance (Sezer et al., 2020; Xu et al., 2024), and climate monitoring (Camps-Valls et al., 2021; Reichstein et al., 2019). Despite their success in time series prediction (Mahmoud & Mohammed, 2021; Wang et al., 2024), the black-box nature of deep neural networks (Zhao et al., 2023) and their use in sensitive domains (Wang & Chung, 2022) make explainability essential. When it comes to online time series, practitioners often place greater emphasis on prediction differences between adjacent time steps rather than isolated predictions, as identical predictions can *signify different contextual meanings*. For instance, in clinical settings, a decrease in the probability of sepsis from 90% to 50% indicates patient improvement, while an increase from 10% to 50% signals potentially severe deterioration (Boussina et al., 2024; Shashikumar et al., 2017). This epitomizes the necessity to contextualize prediction changes in online time series and attribute them to the features driving such transitions.

Albeit diverse explainable artificial intelligence (XAI) methods for time series (Ribeiro et al., 2016; Lundberg, 2017; Sundararajan et al., 2017; Shrikumar et al., 2017; Suresh et al., 2017; Tonekaboni et al., 2020; Leung et al., 2021; Crabbé & Van Der Schaar, 2021; Enguehard, 2023; Liu et al., 2024b; Queen et al., 2024; Liu et al., 2024a; Jang et al., 2025) have been suggested, existing approaches mostly explain predictions at isolated time steps, neglecting how predictions have evolved. In this regard, directly applying single-time attribution methods to explain prediction changes poses significant challenges. First, these methods overlook the contextual dynamics of attributions. For instance, this frequently leads to practically irrelevant explanations (see Section 3 for more details). Second, it is in general impossible to explain prediction differences using isolated attributions from individual time steps, which are irrelevant to prediction changes. For example, Table 5 reveals that

054 the existing XAI approach (Crabbé & Van Der Schaar, 2021) produces implausible attributions when
 055 subtracting across two time steps. Last, appropriate evaluation metrics for attribution in explaining
 056 prediction changes remain largely unexplored. Common evaluation metrics, such as performance
 057 degradation after removing salient features, often exhibit limited correlation with attribution quality
 058 in evolving prediction settings.

059 To bridge this gap, we introduce Delta-XAI, a [novel framework](#) for explaining prediction changes in
 060 online time series monitoring. We begin with a problem setup of *explaining prediction changes in*
 061 *online time series*, attributing features to differences between adjacent steps rather than single-step
 062 predictions. Within this framework, we successfully adapt 14 mainstream XAI methods—including
 063 general approaches (Ribeiro et al., 2016; Lundberg, 2017; Sundararajan et al., 2017; Shrikumar et al.,
 064 2017; Suresh et al., 2017) and time series-specific ones (Tonekaboni et al., 2020; Leung et al., 2021;
 065 Crabbé & Van Der Schaar, 2021; Enguehard, 2023; Liu et al., 2024b; Queen et al., 2024; Liu et al.,
 066 2024a; Jang et al., 2025)—via a wrapper function that transforms single-time explanations to directly
 067 quantify prediction changes. We also propose a [novel evaluation](#) suite for online time series XAI,
 068 assessing attributions for faithfulness, sufficiency, completeness, coherence, and efficiency, providing
 069 a holistic standard over fragmented practices. Within this framework, we uncover that conventional
 070 gradient-based methods like Integrated Gradients (IG) (Sundararajan et al., 2017) typically outperform
 071 recent alternatives, consistent with offline setting (Jang et al., 2025).

072 Motivated by these observations, we introduce Shifted Window Integrated Gradients (SWING), a
 073 novel XAI method for explaining prediction changes in online time series. IG typically exploits
 074 integration along a straight path from a zero baseline to the input, failing to capture temporal
 075 dynamics and inducing an out-of-distribution (OOD) problem. SWING addresses these issues by
 076 generalizing the integration path with the retrospective prediction window as baseline and performing
 077 piecewise line integrals through intermediate windows between two time steps. This yields faithful and
 078 interpretable attributions that capture temporal dynamics and mitigate OOD issues, while satisfying
 079 desirable theoretical properties for online time series monitoring—completeness, implementation
 080 invariance, and skew-symmetry. Extensive experiments under a wide range of settings, spanning
 081 diverse datasets and model architectures, verify that SWING outperforms existing baselines.

082 Our contribution can be summarized as follows:

- 083 • We formulate a problem of explaining prediction changes in online time series monitoring
 084 and propose a [novel framework](#) that adapts 14 existing XAI methods via a prediction
 085 wrapper, alongside comprehensive evaluation metrics tailored to this setup.
- 086 • We propose SWING, an XAI method for online time series that extends IG by incorporating
 087 past observations into the integration path, enabling temporal dynamics to be captured while
 088 mitigating OOD effects, while satisfying key theoretical properties.
- 089 • Through extensive experiments across diverse benchmarks and backbone architectures, we
 090 systematically analyze existing time series attribution methods and demonstrate that SWING
 091 surpasses state-of-the-art alternatives under diverse evaluation metrics.

093 2 PROBLEM SETUP

096 Let $f : \mathbb{R}^{W \times D} \rightarrow [0, 1]^C$ be an online time series classifier where W is the fixed lookback window
 097 size, D is the number of features, and C is the number of output classes, which can be either binary or
 098 multiclass. Given an online time series $\mathbf{X} \in \mathbb{R}^{L \times D}$ of length L , the model receives at each time step T
 099 the input window $\mathbf{X}_{T-W+1:T} \in \mathbb{R}^{W \times D}$ and outputs predicted class probabilities $f(\mathbf{X}_{T-W+1:T}) \in$
 100 $\Delta^{C-1} := \{p \in [0, 1]^C \mid \sum_{c=1}^C p_c = 1\}$. In a conventional time series XAI setup, the goal is
 101 to explain the model’s prediction at time T by estimating the contribution of each input feature
 102 within the window: for each t and d , existing approaches calculate the attribution $\varphi(f, \mathbf{X}_{t,d} \mid T)$,
 103 which quantifies the contribution of feature d at time t to the model output $f(\mathbf{X}_{T-W+1:T})_{\hat{c}}$, where
 104 $\hat{c} = \arg \max_c f(\mathbf{X}_{T-W+1:T})_c$.

105 In contrast, our goal is to explain *prediction changes* between two time steps $T_1 < T_2$. We define
 106 the target class as the one with the largest probability increase: $\hat{c} = \arg \max_c f(\mathbf{X}_{T_2-W+1:T_2})_c -$
 107 $f(\mathbf{X}_{T_1-W+1:T_1})_c$, with the corresponding change $\Delta = f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}}$. To
 108 explain this change Δ , we compute attributions over $t \in \{T_1 - W + 1, \dots, T_2\}$, $d \in \{1, \dots, D\}$,

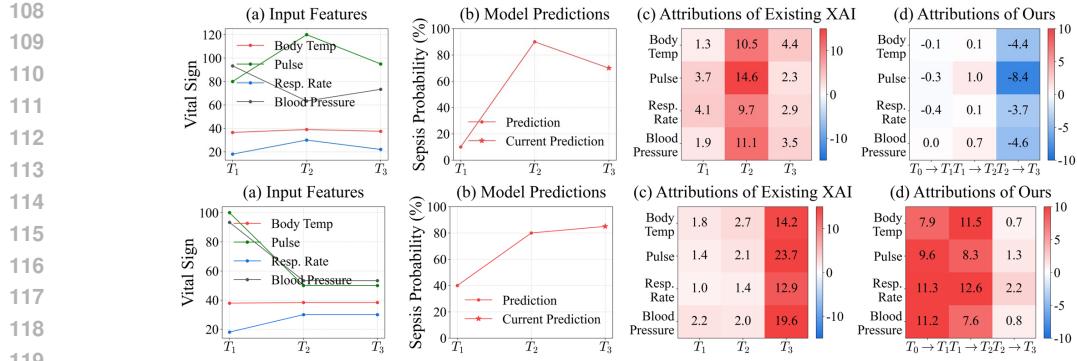


Figure 1: Motivation for explaining prediction changes through illustrative scenarios that are not generated by actual XAI outputs. (top) Vital signs across T_1, T_2, T_3 : risk rises from 10% to 90% then partially recovers (70%). Conventional attribution at T_3 misleads, while our method highlights features driving recovery. (bottom) Risk evolves from 10% at T_1 to 80% at T_2 and slightly increases (85%) at T_3 due to delayed effects. Our method attributes prediction changes between $T_1 \rightarrow T_2$ and $T_2 \rightarrow T_3$, resolving these issues. **Here, T_0 denotes a hypothetical past time step at which the baseline input is placed to compute the initial score.**

where each $\varphi(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2)$ quantifies the contribution of feature d at time t to Δ . By default, we take \hat{c} to be the class with the largest probability increase, motivated by the fact that increasing-probability classes are often the most *practically relevant*. For instance, in disease risk prediction, even if *normal* remains the top class, a sharp rise in *respiratory failure* probability from T_1 to T_2 is clinically significant and merits explanation. However, Delta-XAI also supports flexible setups such as attributing the difference between the predicted classes, *i.e.*, $\Delta = f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}_2} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}_1}$, where $\hat{c}_1 = \arg \max_c f(\mathbf{X}_{T_1-W+1:T_1})_c$ and $\hat{c}_2 = \arg \max_c f(\mathbf{X}_{T_2-W+1:T_2})_c$, though this is not our primary concern. We further restrict to $T_2 - T_1 < W$, as larger gaps yield non-overlapping windows and no common temporal context for attribution.

3 MOTIVATION: CHALLENGES IN ONLINE TIME SERIES EXPLANATION

In time series applications, stakeholders often care less about static predictions and more about *why they change*. Standard XAI methods attribute importance to a single prediction, but clinicians, financiers, or engineers instead want to know which features *caused* risk to rise, fall, or remain stable. For instance, a doctor may ask which signals explain a patient’s sudden recovery, or an analyst may seek factors driving a credit score drop. Explaining such changes requires attribution methods tailored to highlight features responsible for prediction differences over time.

With regard to this, existing XAI methods explain single-time predictions without historical context, making it hard to tell whether a feature caused risk to rise, fall, or stay stable. The top of Figure 1 shows a sepsis model over three steps $T_1 < T_2 < T_3$: risk rises from 10% at T_1 to 90% at T_2 , then drops to 70% at T_3 (b). Clinicians want explanations for the *recovery* between T_2 and T_3 , but standard methods (c) still highlight T_2 ’s features, linking them to high risk. Explaining this 20% reduction instead requires attribution methods explicitly designed for prediction changes (d).

Real-world time series further suffer from irregular sampling and imputation, *e.g.*, by forward filling. As shown in the bottom part of Figure 1, when all values at T_3 are imputed (a), standard methods (c) wrongly assign high attribution to them because models emphasize recent inputs, whereas clinicians require attribution for actual observations. Our setup (d) instead *supports time-wise attribution*, computing T_2 ’s attributions from the change $T_1 \rightarrow T_2$ and T_3 ’s from $T_2 \rightarrow T_3$, thereby highlighting genuine observations while de-emphasizing imputed ones.

4 DELTA-XAI: A FRAMEWORK FOR EXPLAINING PREDICTION CHANGES

In this section, we elaborate on our approach for attributing prediction changes in online time series monitoring. To achieve this, we first introduce a *prediction difference wrapper* that seamlessly adapts existing single-prediction time series explainers to explain prediction changes; we also highlight the

162 special cases where the wrapper function simplifies to differences of attributions due to linearity
 163 properties (Section 4.1). Second, motivated by empirical findings that classical Integrated Gradients
 164 (IG) outperform recent alternatives, we propose Shifted Window Integrated Gradients (SWING),
 165 specifically designed to explain prediction differences, incorporating historical contexts and thereby
 166 addressing the limitations of IG in explaining prediction changes for online time series (Section 4.2).
 167

168 4.1 FROM STATIC TO DYNAMIC: EXTENDING XAI TO PREDICTION CHANGES

170 Given an online time series classifier $f : \mathbb{R}^{W \times D} \rightarrow [0, 1]^C$, our goal is to attribute prediction change
 171 $f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}}$ with $\hat{c} = \arg \max_c f(\mathbf{X}_{T_2-W+1:T_2})_c - f(\mathbf{X}_{T_1-W+1:T_1})_c$
 172 between two time steps $T_1 < T_2$. However, in general, it is impossible to explain prediction differences
 173 using *isolated attributions from individual time steps*, since they are irrelevant to such change. Indeed,
 174 neither computing attributions on differenced inputs (as f is generally nonlinear), nor subtracting
 175 attributions across outputs (as attribution algorithms are also nonlinear), provides valid explanations,
 176 often yielding implausible results; for example, Table 5 shows that Dynamask (Crabbé & Van
 177 Der Schaar, 2021) produces misleading explanations when subtracting attributions across two time
 178 steps. To address this, we devise a prediction difference wrapper g :

$$178 \quad g : \mathbb{R}^{(T_2-T_1+W) \times D} \rightarrow [0, 1]^C, \quad g(\mathbf{X}_{T_1-W+1:T_2}) := f(\mathbf{X}_{T_2-W+1:T_2}) - f(\mathbf{X}_{T_1-W+1:T_1}), \quad (1)$$

180 which allows any single-prediction XAI method φ to be directly applied to explain prediction changes:

$$181 \quad \varphi(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = \varphi(g, \mathbf{X}_{t,d} \mid T_2), \quad \forall t \in \{T_1 - W + 1, \dots, T_2\}, d \in \{1, \dots, D\}. \quad (2)$$

182 This wrapper reformulates attribution of prediction differences as a single-prediction explanation,
 183 making it broadly applicable across existing methods. It further enables single-time XAI approaches
 184 to be adapted with little or no modification, as detailed in Section D.
 185

186 **Special case: linear and complete XAI methods.** While our wrapper function g is broadly
 187 applicable, it requires recomputing attributions for each time pair (T_1, T_2) . When φ satisfies *linearity*
 188 in the attribution space, prediction changes reduce to the difference of single-time attributions,
 189 allowing those from one step to be reused across any pair that includes it, without computing them
 190 through g :

$$191 \quad \varphi(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = \varphi(f, \mathbf{X}_{t,d} \mid T_2) - \varphi(f, \mathbf{X}_{t,d} \mid T_1). \quad (3)$$

192 If φ also satisfies *completeness* in the attribution space, e.g., SHAP variants (Lundberg, 2017),
 193 IG (Sundararajan et al., 2017), DeepLIFT (Shrikumar et al., 2017), this formulation guarantees *online*
 194 *completeness*, ensuring that summed attributions exactly match the prediction change:

195 **Theorem 1** (Attribution Decomposition Theorem for Online Completeness). *Given a linear and*
 196 *complete attribution method φ with a fixed baseline, the following decomposition holds:*

$$197 \quad f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} = \underbrace{\sum_{t=T_1+1}^{T_2} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_2)}_{\text{Addition of newest features}} \\ 198 \quad + \underbrace{\sum_{t=T_2-W+1}^{T_1} \sum_{d=1}^D [\varphi(f, \mathbf{X}_{t,d} \mid T_2) - \varphi(f, \mathbf{X}_{t,d} \mid T_1)]}_{\text{Delayed effect of intermediate features}} - \underbrace{\sum_{t=T_1-W+1}^{T_2-W} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_1)}_{\text{Removal of oldest features}}. \quad (4)$$

206 This theorem indicates that summed attributions across all features and time points exactly equal the
 207 prediction change, providing clear interpretability. The proof is in Section F.
 208

209 4.2 SWING: SHIFTED WINDOW INTEGRATED GRADIENTS

211 Standard Integrated Gradients (IG) remains competitive for explaining prediction changes but suffers
 212 from OOD artifacts and ignores temporal dynamics. We propose SWING, extending IG with i)
 213 retrospective baseline selection (RBS), ii) dual-path integration (DPI) ensuring online completeness,
 214 and iii) piecewise-linear historical integration (PHI), yielding reliable explanations while preserving
 215 key theoretical properties. The overall pipeline is shown in Figure 2, with the detailed procedure
 outlined in Algorithm 1.

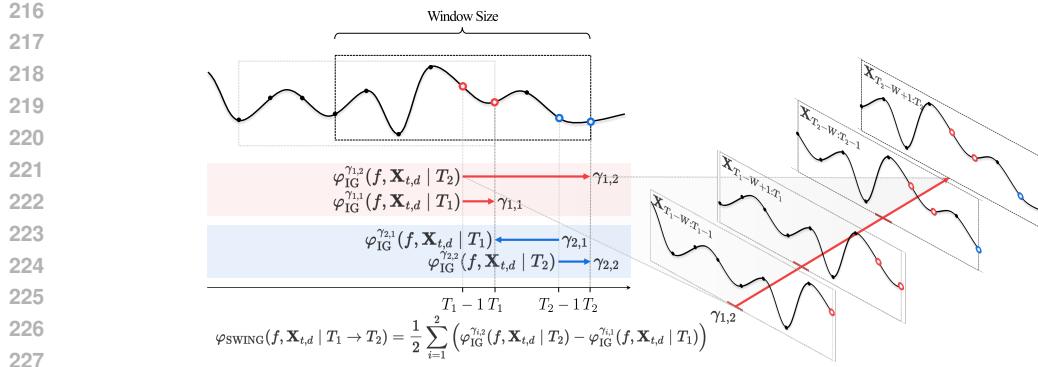


Figure 2: Overview of the proposed SWING framework for explaining prediction changes in online patient monitoring. SWING extends conventional Integrated Gradients (IG) by replacing zero-baseline straight paths with line integrals over shifted historical windows and piecewise-linear paths, capturing temporal dynamics and mitigating out-of-distribution effects.

Extending IG to line integrals over parameterized curves. We generalize IG as a line integral over a parameterized curve $\gamma : [0, 1] \rightarrow \mathbb{R}^{W \times D}$ connecting a baseline \mathbf{X}' to the input $\mathbf{X}_{T-W+1:T}$:

$$\varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d} | T) = \int_0^1 \frac{\partial f(\gamma(\alpha))_{\hat{c}}}{\partial \mathbf{X}_{t,d}} \frac{\partial \gamma_{t,d}(\alpha)}{\partial \alpha} d\alpha. \quad (5)$$

Here, standard IG appears as the special case of a straight-line path $\gamma(\alpha) = (1-\alpha)\mathbf{X}' + \alpha\mathbf{X}_{T-W+1:T}$, where $\partial \gamma_{t,d}(\alpha) / \partial \alpha = \mathbf{X}_{t,d} - \mathbf{X}'_{t,d}$, yielding $\varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d} | T) = (\mathbf{X}_{t,d} - \mathbf{X}'_{t,d}) \int_0^1 \frac{\partial f(\mathbf{X}' + \alpha(\mathbf{X}_{T-W+1:T} - \mathbf{X}'_{t,d}))_{\hat{c}}}{\partial \mathbf{X}_{t,d}} d\alpha$.

Retrospective baseline selection. Motivated by the intuition that the most realistic baseline is the *recent past observation*, we generalize the baseline to the window d steps before the input. For an input window $\mathbf{X}_{T-W+1:T}$, this generalized baseline is $\mathbf{X}_{T-W+1-d:T-d}$. In practice, we set $d = 1$ as the default, since the immediate past provides the most stable and realistic reference. Formally, with $d = 1$, we denote by $\gamma_i : [0, 1] \rightarrow \mathbb{R}^{W \times D}$ the straight-line path from the baseline $\mathbf{X}_{T_i-W:T_i-1}$ to the input $\mathbf{X}_{T_i-W+1:T_i}$, parameterized as: $\gamma_i(\alpha) = (1-\alpha)\mathbf{X}_{T_i-W:T_i-1} + \alpha\mathbf{X}_{T_i-W+1:T_i}$, $\alpha \in [0, 1]$. This keeps $\gamma_i(\alpha)$ near the data manifold and mitigates OOD issues, yielding:

$$\varphi_{\text{RBS}}(f, \mathbf{X}_{t,d} | T_1 \rightarrow T_2) = \varphi_{\text{IG}}^{\gamma_2}(f, \mathbf{X}_{t,d} | T_2) - \varphi_{\text{IG}}^{\gamma_1}(f, \mathbf{X}_{t,d} | T_1). \quad (6)$$

Dual-path integration. Since φ_{RBS} uses distinct baselines at T_1 and T_2 , relying on a single path may lead to incomplete explanations. Thus, we define $\tilde{\gamma}_{i,j} : [0, 1] \rightarrow \mathbb{R}^{W \times D}$ as the straight-line path: $\tilde{\gamma}_{i,j}(\alpha) = (1-\alpha)\mathbf{X}_{T_i-W:T_i-1} + \alpha\mathbf{X}_{T_j-W+1:T_j}$, $\alpha \in [0, 1]$, $i, j \in \{1, 2\}$. DPI integrates along all four baseline–input pairs $\tilde{\gamma}_{i,j}$ and averages the results:

$$\varphi_{\text{DPI}}(f, \mathbf{X}_{t,d} | T_1 \rightarrow T_2) = \frac{1}{2} \sum_{i=1}^2 \left(\varphi_{\text{IG}}^{\tilde{\gamma}_{i,2}}(f, \mathbf{X}_{t,d} | T_2) - \varphi_{\text{IG}}^{\tilde{\gamma}_{i,1}}(f, \mathbf{X}_{t,d} | T_1) \right). \quad (7)$$

This symmetric construction balances both baselines and, more importantly, ensures *online completeness* in Theorem 1, so that summed attributions equal the prediction change.

Piecewise-linear historical integration. When the temporal gap between the baseline at $T_i - 1$ and the target at T_j is large, directly interpolating between $\mathbf{X}_{T_i-W:T_i-1}$ and $\mathbf{X}_{T_j-W+1:T_j}$ may traverse regions far from the data manifold, leading to unstable attributions. To prevent such off-manifold transitions, we introduce a piecewise-linear sliding path $\gamma_{i,j}(\alpha)$ that incrementally shifts the window across intermediate historical time points, ensuring a smooth and temporally consistent progression from baseline to target. We describe the case $i < j$; the case $i > j$ is obtained by reversing the temporal order. Let $M = T_j - T_i + 1$ denote the number of window transitions between $T_i - 1$ and T_j , and for $\alpha \in [0, 1]$ set:

$$K = \underbrace{\lfloor \alpha M \rfloor}_{\text{segment index}}, \quad \tilde{\alpha} = \underbrace{\frac{\alpha M - K}{M}}_{\text{local interpolation ratio}}. \quad (8)$$

270 Then, the integration path is defined as $\gamma_{i,j}(\alpha) = (1 - \tilde{\alpha})\mathbf{X}_{T_i+K-W:T_i+K-1} + \tilde{\alpha}\mathbf{X}_{T_i+K-W+1:T_i+K}$.
 271 As α increases from 0 to 1, the interpolation shifts smoothly from the baseline window at $T_i - 1$ to
 272 the target window at T_j in M small linear steps, ensuring that the trajectory remains close to the
 273 temporal data manifold rather than jumping directly across distant windows. This sliding construction
 274 yields more stable and faithful attributions for long-range temporal changes. SWING then aggregates
 275 contributions over all baseline-target combinations, producing a temporally consistent explanation
 276 while preserving online completeness:

$$\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = \frac{1}{2} \sum_{i=1}^2 \left(\varphi_{\text{IG}}^{\gamma_{i,2}}(f, \mathbf{X}_{t,d} \mid T_2) + \varphi_{\text{IG}}^{\gamma_{i,1}}(f, \mathbf{X}_{t,d} \mid T_1) \right). \quad (9)$$

280 In practice, we uniformly sample α at n_{samples} points within $[0, 1]$ to obtain a tractable approximation
 281 of the path integral.

283 **Theoretical properties of SWING.** SWING extends the axiomatic guarantees of IG—Online Completeness,
 284 Implementation Invariance, and Skew-Symmetry—providing principled interpretations of
 285 online prediction changes. The proofs provided in Section F.

286 **Theorem 2** (Online Completeness). *The sum of SWING attributions equals the prediction difference
 287 between two time steps: $\sum_{t,d} \varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = f(\mathbf{X}_{T_2-W+1:T_2}) - f(\mathbf{X}_{T_1-W+1:T_1})$.*

288 **Theorem 3** (Implementation Invariance). *SWING attributions depend only on the model function,
 289 remaining invariant to equivalent network implementations.*

290 **Theorem 4** (Skew-Symmetry). *For reversed prediction changes, SWING attributions satisfy skew-
 291 symmetry: $\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = -\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_2 \rightarrow T_1)$.*

293 5 EVALUATION METRICS FOR EXPLAINING PREDICTION CHANGES

295 This section introduces new metrics for online time series monitoring. We first examine issues
 296 with zero and average imputation (Section 5.1), then propose metrics for attribution faithfulness
 297 and sufficiency (Section 5.2), and finally extend beyond these to a unified evaluation standard
 298 (Section 5.3).

300 5.1 PROBLEM OF EXISTING EVALUATION METRICS

302 Modern time series XAI methods (Liu et al., 2024b;a) typically 303 assess *faithfulness*—how well attributions reflect model decisions—
 304 and *sufficiency*—how well retained features preserve predictions—
 305 by substituting removed features with zero or average baselines.
 306 These substitutions ignore temporal locality and autocorrelation,
 307 yielding out-of-distribution (OOD) samples. Our analysis on MIMIC-
 308 III (Johnson et al., 2016) (LSTM backbone, SWING attributions) in
 309 Table 1 shows that zero/average substitution produces substantial OOD samples, with high Cumulative
 310 Prediction Difference (CPD) and large OOD scores measured by conditional generative model MSE,
 311 exaggerating prediction differences and distorting correlations. We therefore adopt forward-filling for
 312 faithfulness evaluation, as it reduces OOD effects and improves reliability.

Table 1: Substitution analysis
 for XAI evaluation.

Substitution	CPD	OOD Score
Zero	28.12	0.840
Average	14.85	0.222
Forward-Fill	12.98	0.093

313 5.2 PROPOSED EVALUATION METRICS FOR FAITHFULNESS AND SUFFICIENCY

315 Recently, TIMING (Jang et al., 2025) has identified a key issue in faithfulness and sufficiency
 316 evaluation of time series XAI: removing top/bottom salient points *simultaneously* can inflate scores
 317 by rewarding mere sign alignment. To mitigate this, TIMING introduced Cumulative Prediction
 318 Difference (CPD) and Cumulative Prediction Preservation (CPP), which remove features *sequentially*.
 319 In our wrapper setting g with input \mathbf{X} :

$$\text{CPD}(g, \mathbf{X}, K) = \sum_{k=0}^{K-1} \|g(\mathbf{X}_k^\uparrow) - g(\mathbf{X}_{k+1}^\uparrow)\|_1, \quad \text{CPP}(g, \mathbf{X}, K) = \sum_{k=0}^{K-1} \|g(\mathbf{X}_k^\downarrow) - g(\mathbf{X}_{k+1}^\downarrow)\|_1, \quad (10)$$

323 where \mathbf{X}_k^\uparrow and \mathbf{X}_k^\downarrow denote the inputs obtained by removing the top- k and bottom- k features, respectively,
 from the full set of $(T_2 - T_1 + W) \times D$ features across the entire time window. These metrics

324 showed that gradient-based methods like IG (Sundararajan et al., 2017) often outperform recent
 325 masking-based ones (Liu et al., 2024b; Queen et al., 2024; Liu et al., 2024a), a trend we also find for
 326 online prediction changes.

327 However, these metrics have key limitations: 1) they ignore the relative ranking among top- k features,
 328 2) they overlook evolving attributions across time, and 3) they assess only ranking, not magnitude. To
 329 address 1), we introduce area-based metrics, Area Under Prediction Difference/Preservation (AUPD,
 330 AUPP). AUPD is defined as the average of CPD values over all prefixes of the top- k features:

$$332 \quad \text{AUPD}(g, \mathbf{X}, K) = \frac{1}{2K} \sum_{k=1}^K \left(\text{CPD}(g, \mathbf{X}, k) + \text{CPD}(g, \mathbf{X}, k-1) \right). \quad (11)$$

333 AUPP is defined analogously using CPP instead of CPD. For 2), we aggregate attributions with a
 334 centered sliding window: $\varphi(\mathbf{X}_{t,d} \mid T) = 1/(2W-1) \sum_{T'=t-W+1}^{t+W-1} \varphi(\mathbf{X}_{t,d} \mid T')$. We denote the
 335 resulting macro-level metrics as Macro Prediction Difference/Preservation (MPD, MPP) and their
 336 area-based variants as Area Under Macro Prediction Difference/Preservation (AUMPD, AUMPP).
 337 For 3), we propose Corr., the correlation between ordered attributions and prediction differences:

$$340 \quad \text{Corr.}(\varphi, \mathbf{X}, K) = \text{Corr.} \left(\left[\varphi^{(1)}, \dots, \varphi^{(K)}, \varphi^{(W \times D - K + 1)}, \dots, \varphi^{(W \times D)} \right], \right. \\ 341 \quad \left. \left[|g_1^\uparrow - g_0^\uparrow|, \dots, |g_K^\uparrow - g_{K-1}^\uparrow|, |g_1^\downarrow - g_0^\downarrow|, \dots, |g_K^\downarrow - g_{K-1}^\downarrow| \right] \right), \quad (12)$$

342 Together, these metrics capture ranking consistency, temporal dynamics, and attribution magnitudes,
 343 providing more faithful and interpretable evaluations.

344 5.3 BEYOND FAITHFULNESS AND SUFFICIENCY: BROADER EVALUATION METRICS

345 Existing time series XAI studies (Tonekaboni et al., 2020; Leung et al., 2021; Crabbe & Van
 346 Der Schaar, 2021; Enguehard, 2023; Liu et al., 2024b; Queen et al., 2024; Liu et al., 2024a) mainly
 347 assess faithfulness and sufficiency. While crucial—and expanded here with nine detailed metrics—
 348 these alone are insufficient for practical utility. We therefore incorporate: 1) Coherence, checking
 349 alignment with domain knowledge (case study); and 2) Time/Memory Complexity, measuring real-
 350 time feasibility (empirically). Together, these provide a more comprehensive evaluation of XAI
 351 methods.

352 6 EXPERIMENTS

353 In this section, we comprehensively evaluate SWING against 14 time series XAI baselines within our
 354 Delta-XAI framework. We begin with the experimental setup (Section 6.1), then compare methods
 355 on faithfulness to prediction changes (Section 6.2). Next, we analyze the contribution of each
 356 SWING component through ablation study (Section 6.3), examine qualitative and coherence aspects
 357 (Section 6.4), and finally provide further analyses under diverse settings along with computational
 358 cost analysis (Section 6.5).

359 6.1 EXPERIMENTAL SETUP

360 **Datasets.** Following prior work (Liu et al., 2024a;b), we use two large-scale clinical datasets
 361 commonly adopted for online time series monitoring: MIMIC-III (Johnson et al., 2016) for decom-
 362 pensation prediction and PhysioNet 2019 (Reyna et al., 2020b) for early sepsis detection, where
 363 predictions are updated as new data arrive. To assess generalizability, we also test on Activity, a
 364 human activity recognition dataset (Vidulin et al., 2010), and synthetic benchmarks with controlled
 365 temporal dynamics (Delayed Spike (Leung et al., 2021), Switch-Feature (Tonekaboni et al., 2020)).
 366 Further details are provided in Section G.

367 **Model architectures.** We mainly evaluate XAI methods with LSTM architectures, a fundamental
 368 choice for time series classification (Tonekaboni et al., 2020; Leung et al., 2021). To show our
 369 framework’s versatility, we also implement a CNN with stacked convolutions and a Transformer
 370 encoder for long-range dependencies.

378 **Table 2:** Performance comparison of XAI methods on clinical prediction tasks: MIMIC-III decom-
 379 pensation benchmark using LSTM as backbone architecture. Evaluation is performed by removing
 380 the most or least salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
LIME (Ribeiro et al., 2016)	2.26 \pm 0.04	1.72 \pm 0.03	13.78 \pm 0.08	7.70 \pm 0.08	32.46 \pm 0.16	14.26 \pm 0.10	33.45 \pm 0.14	15.30 \pm 0.10	0.02 \pm 0.00
GradSHAP (Lundberg, 2017)	13.73 \pm 0.06	9.05 \pm 0.04	16.68 \pm 0.08	11.19 \pm 0.05	32.97 \pm 0.13	13.96 \pm 0.06	30.13 \pm 0.17	11.95 \pm 0.09	0.14 \pm 0.00
IG (Sundararajan et al., 2017)	13.42 \pm 0.06	9.10 \pm 0.05	16.14 \pm 0.07	11.31 \pm 0.04	33.55 \pm 0.12	13.97 \pm 0.04	29.46 \pm 0.17	10.85 \pm 0.08	0.17 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	13.58 \pm 0.06	9.42 \pm 0.04	16.03 \pm 0.08	11.25 \pm 0.06	35.96 \pm 0.16	14.61 \pm 0.07	31.53 \pm 0.15	11.41 \pm 0.06	0.19 \pm 0.00
FO (Suresh et al., 2017)	13.14 \pm 0.10	9.92 \pm 0.07	17.79 \pm 0.10	12.92 \pm 0.06	44.02 \pm 0.19	22.32 \pm 0.09	16.92 \pm 0.09	6.24 \pm 0.03	0.26 \pm 0.00
AFO (Tonekaboni et al., 2020)	13.24 \pm 0.07	9.30 \pm 0.05	17.16 \pm 0.09	11.95 \pm 0.03	36.13 \pm 0.22	16.64 \pm 0.09	24.13 \pm 0.16	9.32 \pm 0.07	0.28 \pm 0.00
FIT (Tonekaboni et al., 2020)	3.40 \pm 0.04	2.70 \pm 0.03	7.11 \pm 0.04	6.14 \pm 0.08	35.52 \pm 0.11	17.55 \pm 0.08	12.07\pm0.06	10.19 \pm 0.05	0.06 \pm 0.00
WinIT (Leung et al., 2021)	19.64 \pm 0.07	12.25\pm0.04	24.87\pm0.13	15.45\pm0.08	29.22 \pm 0.07	13.05 \pm 0.05	26.11 \pm 0.12	11.92 \pm 0.06	0.21 \pm 0.00
Dynamask (Crabbé & Van Der Schaar, 2021)	11.72 \pm 0.08	7.56\pm0.04	13.15 \pm 0.08	8.25 \pm 0.04	53.07 \pm 0.24	26.22 \pm 0.08	49.80 \pm 0.16	24.26 \pm 0.06	0.04 \pm 0.00
Extrmask (Enguehard, 2023)	16.66 \pm 0.11	10.47 \pm 0.06	17.51 \pm 0.12	10.63 \pm 0.05	29.91 \pm 0.17	15.13 \pm 0.09	29.64 \pm 0.17	14.84 \pm 0.12	0.08 \pm 0.00
ContraLSP (Liu et al., 2024b)	12.88 \pm 0.36	8.69 \pm 0.26	18.00 \pm 0.16	11.11 \pm 0.18	41.62 \pm 0.30	21.17 \pm 0.10	42.67 \pm 0.28	21.94 \pm 0.10	0.03 \pm 0.00
TimeX (Queen et al., 2024)	16.99 \pm 0.09	11.45 \pm 0.06	19.45 \pm 0.10	12.45 \pm 0.06	50.34 \pm 0.10	24.11 \pm 0.04	51.06 \pm 0.08	24.50 \pm 0.06	0.03 \pm 0.00
TimeX++(Liu et al., 2024a)	11.12 \pm 0.05	7.00 \pm 0.04	13.14 \pm 0.02	7.76 \pm 0.02	34.21 \pm 0.17	13.72 \pm 0.07	32.34 \pm 0.11	13.08 \pm 0.05	0.03 \pm 0.00
TIMING (Jang et al., 2025)	14.99 \pm 0.07	9.71 \pm 0.05	16.50 \pm 0.08	11.53 \pm 0.04	31.22 \pm 0.16	13.36 \pm 0.05	27.19 \pm 0.19	10.24 \pm 0.07	0.19 \pm 0.00
SWING	23.87\pm0.16	16.23\pm0.10	22.27\pm0.19	15.52\pm0.12	17.76\pm0.04	5.85\pm0.04	18.20 \pm 0.06	6.06\pm0.05	0.40\pm0.00

392 **XAI baselines.** We comprehensively implement and evaluate all XAI methods in Delta-XAI, including
 393 our proposed SWING. Existing methods are categorized as: 1) modality-agnostic perturbation-
 394 based (LIME (Ribeiro et al., 2016), FO (Suresh et al., 2017), AFO (Tonekaboni et al., 2020)); 2)
 395 gradient-based (IG (Sundararajan et al., 2017), DeepLIFT (Shrikumar et al., 2017), GradSHAP (Lund-
 396 berg, 2017)); and 3) time series-specific methods, including online explainers (FIT (Tonekaboni et al.,
 397 2020), WinIT (Leung et al., 2021)), masking frameworks (Dynamask (Crabbé & Van Der Schaar,
 398 2021), ExtrMask (Enguehard, 2023), ContraLSP (Liu et al., 2024b), TimeX (Queen et al., 2024),
 399 TimeX++(Liu et al., 2024a)), and TIMING (Jang et al., 2025), augmenting IG with random masking.

400 **Implementation details.** Our method uses a single hyperparameter ($n_{\text{samples}} = 50$). We set $T_2 - T_1 = 1$ for adjacent-step explanations, remove $K = 50$ points, and absolutize directional attributions
 401 for fairness. Explanations are obtained through the wrapper g , which highlights features driving
 402 the $T_1 \rightarrow T_2$ change. Metrics are scaled by 10^3 except for correlations, memory, and time. Results
 403 are reported as mean \pm standard error over five runs, with best and second-best marked in **bold** and
 404 underline.

407 6.2 RESULTS ON ATTRIBUTION FAITHFULNESS AND SUFFICIENCY

409 **Main results.** Tables 2 and 6 shows that, under the Delta-XAI protocol, SWING achieves the best
 410 performance on most metrics across both clinical datasets (MIMIC-III and PhysioNet 2019). Other
 411 gradient-based explainers such as IG, DeepLIFT, and TIMING also perform consistently well. In
 412 contrast, surrogate-driven methods like LIME, TimeX, and TimeX++ exhibit lower performance
 413 on preservation metrics, likely due to the data- and hyperparameter-sensitivity of their surrogate
 414 models. These findings underscore SWING’s dominance and refine the performance hierarchy of
 415 XAI techniques within Delta-XAI.

417 **Diverse synthetic and real-world benchmarks.** On the real-world Activity dataset (Table 7),
 418 SWING achieves the highest scores across all metrics. On the synthetic State and Switch-Feature
 419 benchmarks, it performs best or second-best on most metrics, with competitive results elsewhere,
 420 underscoring its robustness across both practical and controlled scenarios.

422 **Different backbone architectures.** To assess generalizability across architectures, we re-evaluate
 423 the five strongest clinical baselines and SWING on the MIMIC-III dataset using CNN and Transformer
 424 backbones in Table 8. SWING outperforms competing methods across most of the metrics; these
 425 findings confirm that SWING maintains robust faithfulness across diverse backbones within Delta-
 426 XAI.

428 **Larger time differences.** To assess robustness over longer time gaps, we compare SWING with
 429 five strong baselines on MIMIC-III at $T_2 - T_1 = 6$ and 24 (Table 9), with WinIT reported only for 6
 430 due to generator limits. SWING achieves best or near-best scores on most metrics, with particularly
 431 dominant gains on preservation metrics, though performance gaps narrow for longer intervals as CPD
 432 and AUPD converge across methods.

432 **Table 3:** Ablation study of SWING examining retrospective baseline selection (RBS), dual-path
 433 integration (DPI), and piecewise-linear historical integration (PHI) on the MIMIC-III decompensation
 434 benchmark (Johnson et al., 2016), with LSTM (Hochreiter & Schmidhuber, 1997) backbone and
 435 interval $T_2 - T_1 = 24$. We vary the baseline distance d (default: 1) and remove the most or least
 436 salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
w/o RBS, PHI	45.82 \pm 0.20	26.85 \pm 0.13	40.85 \pm 0.22	26.31 \pm 0.14	104.05 \pm 0.34	48.43 \pm 0.19	77.94 \pm 0.30	29.56 \pm 0.11	0.20 \pm 0.00
w/o RBS	40.55 \pm 0.27	24.54 \pm 0.19	47.00 \pm 0.36	29.35 \pm 0.22	80.62 \pm 0.41	31.59 \pm 0.25	83.91 \pm 0.38	33.72 \pm 0.20	0.16 \pm 0.00
w/o DPI ($\gamma_{1,1}, \gamma_{2,2}$)	55.01\pm0.35	32.39\pm0.21	53.09\pm0.38	33.08\pm0.21	71.01 \pm 0.30	26.12 \pm 0.10	76.26 \pm 0.34	29.29 \pm 0.09	0.19 \pm 0.00
$d = 0$	33.80 \pm 0.23	20.11 \pm 0.13	48.98 \pm 0.44	29.78 \pm 0.24	82.28 \pm 0.29	33.99 \pm 0.10	83.69 \pm 0.35	33.31 \pm 0.12	0.11 \pm 0.00
$d = 3$	42.00 \pm 0.33	26.00 \pm 0.28	47.49 \pm 0.44	29.74 \pm 0.34	68.53 \pm 0.49	25.60 \pm 0.18	71.88 \pm 0.62	27.54 \pm 0.27	0.19 \pm 0.00
$d = 5$	41.98 \pm 0.31	25.65 \pm 0.25	47.18 \pm 0.45	29.44 \pm 0.33	72.31 \pm 0.66	27.88 \pm 0.22	74.62 \pm 0.71	29.11 \pm 0.26	0.18 \pm 0.00
$d = 10$	41.02 \pm 0.28	24.94 \pm 0.18	46.89 \pm 0.38	29.23 \pm 0.25	76.55 \pm 0.48	29.96 \pm 0.16	78.11 \pm 0.44	31.04 \pm 0.16	0.17 \pm 0.00
SWING	41.07 \pm 0.22	26.46 \pm 0.18	50.58 \pm 0.28	32.29 \pm 0.20	60.29\pm0.14	21.60\pm0.08	64.43\pm0.19	23.87\pm0.10	0.21\pm0.00

445 6.3 ABLATION STUDY

446 This subsection examines the contributions of SWING components—retrospective baseline selection
 447 (RBS), piecewise-linear historical integration (PHI), and dual-path integration (DPI). [The meaning of
 448 each ablation configuration and its corresponding mathematical formulation are minutely detailed
 449 in Section H](#). As shown in Table 3, removing both RBS and PHI causes substantial degradation, while
 450 individually removing either module also reduces performance, highlighting their complementary
 451 roles. Without DPI, the model attains the best faithfulness scores (CPD, AUPD, MPD, AUMPD)
 452 but falls behind SWING in preservation metrics, suggesting that DPI primarily stabilizes preserva-
 453 tion. Overall, SWING sacrifices a small margin in prediction-difference metrics to achieve strong
 454 preservation and correlation performance. Varying the baseline offset d further shows that using the
 455 immediate past window ($d = 1$) yields the best trade-off, as both too short ($d = 0$) and longer offsets
 456 ($d \geq 3$) degrade results. Finally, SWING maintains stable performance across varying n_{samples} (10 to
 457 100), demonstrating robustness to hyperparameter changes (Figure 4).

459 6.4 QUALITATIVE ANALYSIS

461 **Case study on feature attributions.** Beyond quantitative evaluation, we qualitatively assess XAI
 462 methods within the Delta-XAI framework. Figures 5 to 9 show raw input trajectories with attribution
 463 heatmaps for fifteen baselines at $T_2 - T_1 = 1$. FO, TimeX, and TimeX++ tend to spread attributions
 464 broadly across the time axis, while Dynamask and TimeX++ align closely with input fluctuations,
 465 reflecting their strong benchmark scores. In contrast, SWING yields sharper, localized attributions
 466 that emphasize recent time steps most responsible for prediction changes.

468 **Coherence analysis.** We further assess whether SWING aligns with clinical knowledge by inspect-
 469 ing attributions on a representative MIMIC-III case in Figure 10 (due to the spatial constraint). [To
 470 maximize visibility, we closely examine how the features at the last time step influence the model’s
 471 predictions.](#) A sharp blood pressure drop at the last time step increases risk (a), consistent with
 472 evidence linking hypotensive episodes to decompensation (Toki et al., 2025). In contrast, an SpO_2
 473 rise lowers risk (b), while an SpO_2 drop increases risk (c), reflecting the destabilizing effect of oxygen
 474 desaturation (Semler et al., 2022). A rise in GCS lowers risk (d), consistent with evidence linking
 475 improved consciousness to better outcomes (Marincowitz et al., 2018). These examples demonstrate
 476 that SWING produces attribution patterns consistent with established clinical findings.

478 6.5 FURTHER ANALYSIS

480 **In-depth analysis under diverse settings.** To further examine SWING’s behavior, we perform
 481 subgroup analyses by splitting cases according to whether predictions or class labels change (Tables 10
 482 and 11), where it consistently preserves explanatory advantages. We also test different temporal
 483 resolutions with 24- and 72-length windows (Table 12), confirming that SWING yields stable
 484 attributions and clear gains over baselines across both horizons. These results demonstrate that
 485 SWING provides reliable explanations under diverse prediction dynamics and temporal contexts,
 486 reinforcing its applicability to real-world clinical monitoring.

486 **Efficiency analysis: runtime and memory.** We assess SWING’s efficiency along two
 487 axes—runtime and memory footprint. As shown in Figure 3, SWING attains the highest AUPD
 488 while requiring only 0.35s per sample, comparable to other gradient-based explainers (DeepLIFT
 489 0.11s, GradSHAP 0.18s). For memory, it consumes 448 MB per sample, identical to IG (448 MB)
 490 and close to GradSHAP and DeepLIFT (438 MB). These results demonstrate that SWING achieves
 491 state-of-the-art explanatory quality without incurring additional computational or memory costs.
 492

493 7 CONCLUSION

494
 495 In this paper, we have introduced the task of explaining prediction changes in online time series
 496 monitoring and proposed Delta-XAI, a [novel](#) framework that integrates 14 XAI methods with a
 497 dedicated evaluation suite for temporal dynamics. Through extensive experiments, we demonstrated
 498 that, when adapted, classical gradient-based methods, such as Integrated Gradients (IG), remain strong
 499 baselines. Motivated by this, we developed SWING, an extension of IG that robustly captures temporal
 500 feature evolution. We believe our contributions significantly advance XAI for online time series by
 501 shifting the paradigm from static interpretations toward dynamic, context-sensitive explanations—an
 502 essential step toward trustworthy AI in time-critical domains.

503 504 ETHICS STATEMENT

505
 506 This work develops explainable AI (XAI) methods for online time series monitoring in domains
 507 such as healthcare and finance. We use only publicly available open-source benchmark datasets (*e.g.*,
 508 MIMIC-III, PhysioNet 2019, Activity), adhering to their usage protocols and ethical standards, and
 509 do not collect new human subject data. Our contributions are methodological, aiming to enhance the
 510 transparency and interpretability of time series models. Potential misuse may arise if explanations are
 511 taken as direct clinical or financial advice; therefore, we emphasize that outputs should be interpreted
 512 by domain experts.

513 514 REPRODUCIBILITY STATEMENT

515
 516 We provide an anonymized implementation of Delta-XAI and SWING at the anonymous repository
 517 link <https://anonymous.4open.science/r/Delta-XAI>. All experimental details—
 518 including dataset preprocessing, model architectures, training protocols, hyperparameter settings, and
 519 evaluation metrics—are provided in Sections 6 and G, while theoretical proofs and pre-processing
 520 steps for MIMIC-III, PhysioNet 2019, Activity, and synthetic benchmarks are elaborated in Sec-
 521 tion G. In addition, several ablation studies and robustness analyses over multiple iterations further
 522 validate the stability of results. All procedures are publicly available and fully reproducible, enabling
 523 independent researchers to reproduce and verify our findings.

524 525 REFERENCES

526 João Bento, Pedro Saleiro, André F Cruz, Mário AT Figueiredo, and Pedro Bizarro. Timeshap:
 527 Explaining recurrent models through sequence perturbations. In *Proceedings of the 27th ACM
 528 SIGKDD conference on knowledge discovery & data mining*, pp. 2565–2573, 2021.

529
 530 Aaron Boussina, Sreeth P Shashikumar, Atul Malhotra, Robert L Owens, Robert El-Kareh, Christo-
 531 pher A Longhurst, Kimberly Quintero, Allison Donahue, Theodore C Chan, Shamim Nemati, et al.
 532 Impact of a deep learning sepsis prediction model on quality of care and survival. *NPJ digital
 533 medicine*, 7(1):14, 2024.

534 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
 535 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
 536 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

537
 538 Gustau Camps-Valls, Devis Tuia, Xiao Xiang Zhu, and Markus Reichstein. *Deep learning for the
 539 Earth Sciences: A comprehensive approach to remote sensing, climate science and geosciences*.
 John Wiley & Sons, 2021.

540 Molnar Christoph. *Interpretable machine learning: A guide for making black box models explainable*.
 541 Leanpub, 2020.
 542

543 Jonathan Crabbé and Mihaela Van Der Schaar. Explaining time series predictions with dynamic
 544 masks. In *International Conference on Machine Learning*, pp. 2166–2177. PMLR, 2021.

545 Arun Das and Paul Rad. Opportunities and challenges in explainable artificial intelligence (xai): A
 546 survey. *arXiv preprint arXiv:2006.11371*, 2020.

547

548 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 549 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
 550 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*
 551 *arXiv:2010.11929*, 2020.

552

553 Joseph Enguehard. Learning perturbations to explain time series predictions. In *International
 554 Conference on Machine Learning*, pp. 9329–9342. PMLR, 2023.

555

556 Andrew Frank. Uci machine learning repository. <http://archive.ics.uci.edu/ml>, 2010.

557

558 John Cristian Borges Gamboa. Deep learning for time-series analysis. *arXiv preprint*
 559 *arXiv:1701.01887*, 2017.

560

561 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 562 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 563 pp. 770–778, 2016.

564

565 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):
 566 1735–1780, 1997.

567

568 Hyeongwon Jang, Changhun Kim, and Eunho Yang. Timing: Temporality-aware integrated gradients
 569 for time series explanation. In *International Conference on Machine Learning (ICML)*, 2025.

570

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

574

575 Changhun Kim, Yechan Mun, Sangchul Hahn, and Eunho Yang. Deltashap: Explaining prediction
 576 evolutions in online patient monitoring with shapley values. *ICML Workshop on Actionable
 577 Interpretability*, 2025.

578

579 Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds,
 580 Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. Captum: A unified and
 581 generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*, 2020.

582

583 Kin Kwan Leung, Clayton Rooke, Jonathan Smith, Saba Zuberi, and Maksims Volkovs. Temporal
 584 dependencies in feature importance for time series predictions. *arXiv preprint arXiv:2107.14317*,
 585 2021.

586

587 Zichuan Liu, Tianchun Wang, Jimeng Shi, Xu Zheng, Zhuomin Chen, Lei Song, Wenqian Dong,
 588 Jayantha Obeysekera, Farhad Shirani, and Dongsheng Luo. Timex++: Learning time-series
 589 explanations with information bottleneck. *arXiv preprint arXiv:2405.09308*, 2024a.

590

591 Zichuan Liu, Yingying Zhang, Tianchun Wang, Zefan Wang, Dongsheng Luo, Mengnan Du, Min
 592 Wu, Yi Wang, Chunlin Chen, Lunting Fan, et al. Explaining time series via contrastive and locally
 593 sparse perturbations. *arXiv preprint arXiv:2401.08552*, 2024b.

594

595 Scott Lundberg. A unified approach to interpreting model predictions. *arXiv preprint*
 596 *arXiv:1705.07874*, 2017.

597

598 Amal Mahmoud and Ammar Mohammed. A survey on deep learning for time-series forecasting.
 599 *Machine learning and big data analytics paradigms: analysis, applications and challenges*, pp.
 600 365–392, 2021.

594 Carl Marincowitz, Fiona Lecky, William Townend, Agostino Fabbri, and Timothy Sheldon. The
 595 risk of deterioration in gcs 13–15 patients with head injury. *Emergency Medicine Journal*, 35(9):
 596 559–565, 2018. doi: 10.1136/emermed-2017-207180.

597

598 Nicholas G Polson and Vadim O Sokolov. Deep learning for short-term traffic flow prediction.
 599 *Transportation Research Part C: Emerging Technologies*, 79:1–17, 2017.

600

601 Owen Queen, Tom Hartvigsen, Teddy Koker, Huan He, Theodoros Tsiligkaridis, and Marinka Zitnik.
 602 Encoding time-series explanations through self-supervised model behavior consistency. *Advances
 603 in Neural Information Processing Systems*, 36, 2024.

604

605 Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno
 606 Carvalhais, and F Prabhat. Deep learning and process understanding for data-driven earth system
 607 science. *Nature*, 566(7743):195–204, 2019.

608

609 Attila Reiss and Didier Stricker. Introducing a new benchmarked dataset for activity monitoring. In
 610 *2012 16th international symposium on wearable computers*, pp. 108–109. IEEE, 2012.

611

612 Matthew A Reyna, Christopher S Josef, Russell Jeter, Supreeth P Shashikumar, M Brandon Westover,
 613 Shamim Nemati, Gari D Clifford, and Ashish Sharma. Early prediction of sepsis from clinical data:
 614 the physionet/computing in cardiology challenge 2019. *Critical care medicine*, 48(2):210–217,
 615 2020a.

616

617 Matthew A Reyna, Christopher S Josef, Russell Jeter, Supreeth P Shashikumar, M Brandon Westover,
 618 Shamim Nemati, Gari D Clifford, and Ashish Sharma. Early prediction of sepsis from clinical data:
 619 the physionet/computing in cardiology challenge 2019. *Critical care medicine*, 48(2):210–217,
 620 2020b.

621

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

625

626 Yulia Rubanova, Ricky TQ Chen, and David K Duvenaud. Latent ordinary differential equations for
 627 irregularly-sampled time series. *Advances in neural information processing systems*, 32, 2019.

628

629 Matthew W Semler, Jonathan D Casey, Bradley D Lloyd, Pamela G Hastings, Margaret A Hays,
 630 Joanna L Stollings, Kevin G Buell, John H Brems, Edward T Qian, Kevin P Seitz, et al. Oxygen-
 631 saturation targets for critically ill adults receiving mechanical ventilation. *New England Journal of
 632 Medicine*, 387(19):1759–1769, 2022.

633

634 Omer Berat Sezer, Mehmet Ugur Gudelek, and Ahmet Murat Ozbayoglu. Financial time series
 635 forecasting with deep learning: A systematic literature review: 2005–2019. *Applied soft computing*,
 636 90:106181, 2020.

637

638 Lloyd S Shapley. A value for n-person games. In Harold W Kuhn and Albert W Tucker (eds.),
 639 *Contributions to the Theory of Games, Volume II*, pp. 307–317. Princeton University Press, 1953.

640

641 Supreeth P Shashikumar, Matthew D Stanley, Ismail Sadiq, Qiao Li, Andre Holder, Gari D Clifford,
 642 and Shamim Nemati. Early sepsis detection in critical care patients using multiscale blood pressure
 643 and heart rate dynamics. *Journal of electrocardiology*, 50(6):739–743, 2017.

644

645 Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through
 646 propagating activation differences. In *International conference on machine learning*, pp. 3145–
 647 3153. PMIR, 2017.

648

649 Mervyn Singer, Clifford S Deutschman, Christopher Warren Seymour, Manu Shankar-Hari, Djillali
 650 Annane, Michael Bauer, Rinaldo Bellomo, Gordon R Bernard, Jean-Daniel Chiche, Craig M
 651 Coopersmith, et al. The third international consensus definitions for sepsis and septic shock
 652 (sepsis-3). *Jama*, 315(8):801–810, 2016.

653

654 Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In
 655 *International conference on machine learning*, pp. 3319–3328. PMLR, 2017.

648 Harini Suresh, Nathan Hunt, Alistair Johnson, Leo Anthony Celi, Peter Szolovits, and Marzyeh
 649 Ghassemi. Clinical intervention prediction and understanding using deep networks. *arXiv preprint*
 650 *arXiv:1705.08498*, 2017.

651

652 Takayuki Toki, Kazuyuki Mizunoya, Misa Itabashi, Naoki Nishikawa, Koji Hoshino, Hitoshi Saito,
 653 and Yuji Morimoto. Acute decompensated right heart failure potentially triggered by multiple
 654 factors including pulmonary vasodilator removal during plasma exchange: a case report. *JA*
 655 *Clinical Reports*, 11(1):5, 2025.

656 Sana Tonekaboni, Shalmali Joshi, Kieran Campbell, David K Duvenaud, and Anna Goldenberg.
 657 What went wrong and when? instance-wise feature importance for time-series black-box models.
 658 *Advances in Neural Information Processing Systems*, 33:799–809, 2020.

659

660 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 661 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 662 *systems*, 30, 2017.

663 Vedrana Vidulin, Mitja Lustrek, Bostjan Kaluza, Rok Piltaver, and Jana Krivec. Localization Data for
 664 Person Activity. UCI Machine Learning Repository, 2010. DOI: <https://doi.org/10.24432/C57G8X>.

665

666 Yue Wang and Sai Ho Chung. Artificial intelligence in safety-critical systems: a systematic review.
 667 *Industrial Management & Data Systems*, 122(2):442–470, 2022.

668

669 Yuxuan Wang, Haixu Wu, Jiaxiang Dong, Yong Liu, Mingsheng Long, and Jianmin Wang. Deep time
 670 series models: A comprehensive survey and benchmark. *arXiv preprint arXiv:2407.13278*, 2024.

671 Zeda Xu, John Liechty, Sebastian Benthal, Nicholas Skar-Gislinge, and Christopher McComb.
 672 Garch-informed neural networks for volatility prediction in financial markets. In *Proceedings of*
 673 *the 5th ACM International Conference on AI in Finance*, pp. 600–607, 2024.

674 Ziqi Zhao, Yucheng Shi, Shushan Wu, Fan Yang, Wenzhan Song, and Ninghao Liu. Interpretation of
 675 time-series deep models: A survey. *arXiv preprint arXiv:2305.14582*, 2023.

676 Jianhu Zheng and Mingfang Huang. Traffic flow forecast through time series analysis based on deep
 677 learning. *Ieee Access*, 8:82562–82570, 2020.

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A LIMITATIONS AND BROADER IMPACTS**
703704 **Limitations.** Our proposed approach has few limitations. First, while SWING extends IG by
705 incorporating historical points with shifted window paths, it introduces additional computational
706 overhead compared to simpler methods like standard IG. Second, while our prediction wrapper
707 function can seamlessly incorporate most existing single-time XAI algorithms, a few require minor
708 adjustments—particularly those that rely on probability outputs or internal model representations. The
709 specifics of such adaptations are tangential to our main contribution and are left to future algorithm
710 designers.711 **Broader impacts.** Our work significantly advances explainability in time-critical domains by en-
712 abling a nuanced explanation of temporal prediction changes. By providing insights into why and how
713 model predictions evolve, our framework supports decision-making in critical areas such as health-
714 care, finance, and transportation, potentially improving outcomes through enhanced transparency and
715 accountability. While responsible interpretation and privacy considerations remain important, the
716 benefits of improved trustworthiness and actionable insights in high-stakes environments significantly
717 outweigh these concerns.719 **B LLM USAGE DISCLOSURE**
720721 In drafting this manuscript, we made limited use of a Large Language Model (LLM) solely for minor
722 writing improvements, such as grammar polishing and readability enhancement. The LLM was not
723 used for research conception, experimental design, analysis, or generation of substantive content. Its
724 role was strictly restricted to language editing, and all scientific contributions are entirely attributable
725 to the authors.727 **C RELATED WORK**
728729 **Modality-agnostic explainable artificial intelligence.** Although deep neural networks have
730 achieved impressive results across domains like vision (He et al., 2016; Dosovitskiy et al., 2020),
731 language (Vaswani et al., 2017; Brown et al., 2020), and time series (Gamboa, 2017), they often act
732 as black boxes, limiting transparency and accountability—especially in high-stakes areas such as
733 healthcare (Christoph, 2020). To address this, various modality-agnostic XAI methods have been
734 developed. Popular approaches such as LIME (Ribeiro et al., 2016) and SHAP (Shapley, 1953;
735 Lundberg, 2017) attribute predictions to input features by estimating their contribution strength and
736 direction. Variants like KernelSHAP, GradientSHAP, and DeepSHAP (Kokhlikyan et al., 2020) ex-
737 pand their applicability. Gradient-based methods, including Integrated Gradients (IG) (Sundararajan
738 et al., 2017) and DeepLIFT (Shrikumar et al., 2017), compute attributions using model gradients.
739 Perturbation-based methods like Feature Occlusion (FO) (Suresh et al., 2017) and Augmented Feature
740 Occlusion (AFO) (Tonekaboni et al., 2020) measure feature importance by replacing inputs and
741 observing prediction changes. While these methods have enhanced model explainability, most have
742 been evaluated in vision tasks (Das & Rad, 2020). Their application to time series—particularly
743 for explaining prediction changes in online settings—remains limited, despite the importance of
744 capturing temporal dependencies for meaningful explanations.745 **Explainable artificial intelligence for time series.** XAI for time series presents unique challenges
746 due to temporal dependencies, where the order and historical context of observations significantly
747 affect model behavior. Standard modality-agnostic XAI methods, which often assume independently
748 distributed inputs, fail to capture such dynamics. To address this, a number of time series-specific
749 attribution methods have been proposed (Bento et al., 2021; Tonekaboni et al., 2020; Leung et al.,
750 2021; Crabbé & Van Der Schaar, 2021; Enguehard, 2023; Liu et al., 2024b; Queen et al., 2024; Liu
751 et al., 2024a; Kim et al., 2025). More recent methods have improved temporal modeling through
752 dynamic masking (Crabbé & Van Der Schaar, 2021; Enguehard, 2023), contrastive learning (Liu
753 et al., 2024b), and interpretable surrogate modeling (Queen et al., 2024; Liu et al., 2024a), with
754 TimeX++ incorporating an information bottleneck to mitigate trivial explanations. TIMING (Jang
755 et al., 2025) introduces novel evaluation metrics and enhances IG with random masking to improve
sensitivity to temporal variation. Despite recent progress, existing methods fall short in explaining

756 prediction changes in online time series, lacking contextual insight and temporal evaluation. Our
 757 Delta-XAI addresses these gaps by attributing prediction changes directly and introducing metrics
 758 aligned with sequential dynamics.
 759

760 **Explainable artificial intelligence for online time series monitoring.** Among time series XAI
 761 methods, FIT (Tonekaboni et al., 2020) and WinIT (Leung et al., 2021) are particularly relevant to
 762 online prediction tasks. FIT estimates feature importance by comparing predictive distributions under
 763 observed and counterfactual inputs using KL divergence, whereas WinIT models delayed effects
 764 by assessing how past observations influence future predictions. However, our proposed framework
 765 significantly extends beyond these methods by offering a comprehensive and unified approach.
 766 Unlike FIT, which quantifies feature importance solely based on predictive distribution changes
 767 at consecutive time points, and WinIT, which evaluates feature relevance within fixed observation
 768 windows—both producing static attributions across the entire series—our Delta-XAI explicitly
 769 explains prediction changes between distinct time points. Specifically, we attribute changes in model
 770 predictions directly rather than attributing individual predictions independently, thus generating
 771 dynamic and prediction-time-specific attributions that accurately capture temporal evolution in feature
 772 importance. Additionally, we introduce SWING, an advanced attribution method demonstrating
 773 superior empirical performance and fulfilling essential theoretical properties, including linearity,
 774 completeness, and directional attribution. Finally, our framework systematically integrates existing
 775 attribution methods and proposes specialized evaluation metrics tailored explicitly to assessing
 776 prediction change explanations in online time series monitoring.
 777

778 D ADAPTING EXISTING XAI ALGORITHMS

779 In this section, we provide a detailed description of how we adapt existing XAI baselines to our
 780 setting. Specifically, Section D.1 presents the algorithms that can be applied without modification,
 781 *i.e.*, with seamless integration into our prediction-difference framework, while Section D.2 describes
 782 those that require minimal adjustments. In both cases, the adaptation is realized through a wrapper
 783 function g , which standardizes the attribution process for time series inputs. Notably, this wrapper
 784 neither necessitates additional post-training procedures nor alters the underlying models, ensuring
 785 fair and consistent comparison across baselines.
 786

787 D.1 NO MODIFICATION

788 The following attribution algorithms operate without modification: LIME (Ribeiro et al., 2016),
 789 IG (Sundararajan et al., 2017), DeepLIFT (Shrikumar et al., 2017), FO (Suresh et al., 2017),
 790 AFO (Tonekaboni et al., 2020), WinIT (Leung et al., 2021), Dynamask (Crabbé & Van Der Schaar,
 791 2021), Extrmask (Enguehard, 2023), ContraLSP (Liu et al., 2024b), and TIMING (Jang et al., 2025).
 792 In all of these cases, the only adjustment involves the use of the prediction-difference wrapper g
 793 instead of the single-time prediction model f . This wrapper applies uniformly across baselines
 794 without altering their internal mechanisms, so the algorithms remain unmodified in our framework.
 795

796 D.2 MINIMAL MODIFICATION

797 The following attribution algorithms require minimal modification: GradSHAP (Lundberg, 2017),
 798 FIT (Tonekaboni et al., 2020), TimeX (Queen et al., 2024), and TimeX++ (Liu et al., 2024a).
 800

801 **GradSHAP.** Applying GradSHAP (Lundberg, 2017) directly to the wrapper g violates completeness,
 802 since the baseline input for g depends on whether the evaluation is at T_1 or T_2 . For example,
 803 if T_1 corresponds to a window $\mathbf{X}_{T_1-W+1:T_1}$ and T_2 to $\mathbf{X}_{T_2-W+1:T_2}$, then the baseline for g dif-
 804 fers depending on which window is active, leading to inconsistency. To address this, we compute
 805 GradSHAP directly on f , perform two runs with a shared baseline, and subtract the resulting at-
 806 tributions—mirroring the construction of SWING. This procedure preserves completeness while
 807 remaining consistent with the prediction-difference formulation.
 808

809 **FIT.** For FIT (Tonekaboni et al., 2020), we bypass the wrapper g and operate directly on the base
 810 model f . We construct both the current and previous input windows and masks, $\mathbf{X}_{T_2-W+1:T_2}$ and

$\mathbf{X}_{T_1-W+1:T_1}$, and obtain the corresponding predictions $p_{y_{T_2}}$ and $p_{y_{T_1}}$. At each time step and feature dimension, we sample candidate imputations from the generator, evaluate their effect on $p_{y_{T_2}}$, and compute the divergence with respect to the prediction difference. For the divergence metric, we follow the original implementation and support both KL-divergence and mean absolute deviation. This adaptation retains FIT’s perturbation design while aligning it with prediction differences.

TimeX. For TimeX (Queen et al., 2024), the prediction-difference wrapper g is implemented as a Python class, where the `forward()` method returns the prediction difference $f(\mathbf{X}_{T_2-W+1:T_2}) - f(\mathbf{X}_{T_1-W+1:T_1})$. To support the Model Behavior Consistency (MBC) loss, the class defines an auxiliary method that returns the difference between latent embeddings at the two time points. Specifically, while TimeX formulates the MBC loss \mathcal{L}_{MBC} as:

$$\mathcal{L}_{\text{MBC}}(Z, Z^E) = \frac{1}{N^2} \sum_{i,j} [D_Z(z_i, z_j) - D_{Z^E}(z_i^E, z_j^E)]^2, \quad (13)$$

we obtain latent embeddings using the encoder of the model f with $f = \text{dec} \circ \text{enc}$. For an input sequence $\mathbf{X}_{T_1-W+1:T_2}$, we split it into two windows $\mathbf{X}_{T_1-W+1:T_1}$ and $\mathbf{X}_{T_2-W+1:T_2}$, then compute embeddings as:

$$z = \text{enc}(\mathbf{X}_{T_2-W+1:T_2}) - \text{enc}(\mathbf{X}_{T_1-W+1:T_1}). \quad (14)$$

We also experiment with concatenated embeddings $z = [\text{enc}(\mathbf{X}_{T_2-W+1:T_2}), \text{enc}(\mathbf{X}_{T_1-W+1:T_1})]$ and observe negligible differences in performance. In either case, this embedding construction provides a meaningful space over which the surrogate model operates, while remaining consistent with the prediction-difference wrapper g .

TimeX++. For TimeX++ (Liu et al., 2024a), we adopt the same strategy: embeddings are obtained in the same manner as in TimeX, and the surrogate model operates on the difference $z_{T_2} - z_{T_1}$. We further modify the label consistency loss. Instead of the cross-entropy formulation in the original implementation, we employ mean squared error (MSE) loss for stability. Here, \mathbf{X} denotes the original input sequence and $\tilde{\mathbf{X}}$ denotes the perturbed version of the input produced by the explanation method. Concretely, the label consistency objective \mathcal{L}_{LC} is defined as:

$$\mathcal{L}_{\text{LC}}(f(\mathbf{X}), f(\tilde{\mathbf{X}})) = \mathbb{E}[D_{\text{JS}}(f(\mathbf{X}) \parallel f(\tilde{\mathbf{X}}))], \quad (15)$$

following the Jensen–Shannon divergence form in (Queen et al., 2024), but realized with an MSE surrogate. This yields a loss function better aligned with the prediction-difference framework.

E ALGORITHM

We provide the detailed procedure of SWING in Algorithm 1, which computes attribution scores for prediction changes in online time series monitoring. The algorithm explicitly defines integration paths using historically observed shifted windows, computes gradients along piecewise-linear paths via interpolation, and averages attributions from dual integration paths. This approach ensures compliance with the completeness property, provides realistic integration trajectories, and mitigates out-of-distribution (OOD) issues.

F PROOFS

F.1 PROOF OF THEOREM 1

By completeness of φ , for any baseline $\mathbf{X}' \in \mathbb{R}^{W \times D}$, we have:

$$f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}')_{\hat{c}} = \sum_{t=T_2-W+1}^{T_2} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_2), \quad (16)$$

and similarly,

$$f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} - f(\mathbf{X}')_{\hat{c}} = \sum_{t=T_1-W+1}^{T_1} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_1). \quad (17)$$

Algorithm 1 SWING: Shifted Window Integrated Gradients

```

864
865 1: Input: Model  $f$ , inputs  $\mathbf{X}_{T_1-W+1:T_1}, \mathbf{X}_{T_2-W+1:T_2}$ , discretization steps  $n_{\text{samples}}$ .
866 2: Output: Attributions  $\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2)$  for all  $t \in \{T_1 - W + 1, \dots, T_2\}$ ,  $d \in$ 
867  $\{1, \dots, D\}$ .
868 3: for each  $\gamma_{i,j}$  with  $i, j \in \{1, 2\}$  do
869 4:    $M \leftarrow |T_j - (T_i - 1)|$ ,  $\sigma \leftarrow \text{sign}(T_j - (T_i - 1))$ 
870 5:    $\mathbf{X}^{(0)} \leftarrow \mathbf{X}_{T_i-W:T_i-1}$ 
871 6:    $G^{(0)} \leftarrow \partial f(\mathbf{X}^{(0)}) / \partial \mathbf{X}^{(0)}$ 
872 7:    $\varphi_{\text{IG}}^{\gamma_{i,j}} \leftarrow \mathbf{0}_{(T_2-T_1+W) \times D}$ 
873 8:   for  $m = 1$  to  $n_{\text{samples}}$  do ▷ Parallelized in our implementation
874 9:      $\alpha_m \leftarrow m/n_{\text{samples}}$ 
875 10:     $K \leftarrow \underbrace{\min(\lfloor \alpha_m M \rfloor, M - 1)}_{\text{segment index}}$   $\tilde{\alpha} \leftarrow \underbrace{\alpha_m M - K}_{\text{local interpolation ratio}}$ 
876 11:     $s \leftarrow (T_i - 1) + \sigma K$ 
877 12:     $\mathbf{X}^{(m)} \leftarrow (1 - \tilde{\alpha}) \mathbf{X}_{s-W+1:s} + \tilde{\alpha} \mathbf{X}_{(s+\sigma)-W+1:(s+\sigma)}$ 
878 13:     $G^{(m)} \leftarrow \partial f(\mathbf{X}^{(m)}) / \partial \mathbf{X}^{(m)}$ 
879 14:     $\varphi_{\text{IG}}^{\gamma_{i,j}}(t, d) \leftarrow \varphi_{\text{IG}}^{\gamma_{i,j}}(t, d) + (\mathbf{X}_{t,d}^{(m)} - \mathbf{X}_{t,d}^{(m-1)}) \frac{G_{t,d}^{(m)} + G_{t,d}^{(m-1)}}{2}$ ,  $\forall t, d$ 
880 15:  end for
881 16:   $M \leftarrow |T_j - (T_i - 1)|$ ,  $\sigma \leftarrow \text{sign}(T_j - (T_i - 1))$ 
882 17: end for
883 18: 
$$\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) \leftarrow \frac{1}{2} \sum_{i=1}^2 \left( \varphi_{\text{IG}}^{\gamma_{i,2}}(f, \mathbf{X}_{t,d} \mid T_2) - \varphi_{\text{IG}}^{\gamma_{i,1}}(f, \mathbf{X}_{t,d} \mid T_1) \right)$$


```

891 By subtracting two equations, we obtain:

$$\begin{aligned}
892 \quad & f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} = \sum_{t=T_1+1}^{T_2} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_2) \\
893 & + \sum_{t=T_2-W+1}^{T_1} \sum_{d=1}^D [\varphi(f, \mathbf{X}_{t,d} \mid T_2) - \varphi(f, \mathbf{X}_{t,d} \mid T_1)] - \sum_{t=T_1-W+1}^{T_2-W} \sum_{d=1}^D \varphi(f, \mathbf{X}_{t,d} \mid T_1).
\end{aligned} \tag{18}$$

894 **F.2 PROOF OF THEOREM 2**

901 **Lemma 5** (Completeness of Integrated Gradients along General Paths). *Let $f : \mathbb{R}^{W \times D} \rightarrow \mathbb{R}^C$ be*
902 *continuously differentiable and let $\gamma : [0, 1] \rightarrow \mathbb{R}^{W \times D}$ be any continuously differentiable curve with*
903 *$\gamma(0) = \mathbf{X}'$ (baseline) and $\gamma(1) = \mathbf{X}_{T-W+1:T}$ (input). Define the generalized Integrated Gradients*
904 *attribution for each coordinate (t, d) by:*

$$\varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d} \mid T) := \int_0^1 \frac{\partial f(\gamma(\alpha))_{\hat{c}}}{\partial \mathbf{X}_{t,d}} \frac{\partial \gamma_{t,d}(\alpha)}{\partial \alpha} d\alpha. \tag{19}$$

905 *Then completeness holds along any such curve:*

$$\sum_{t,d} \varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d} \mid T) = f(\mathbf{X}_{T-W+1:T})_{\hat{c}} - f(\mathbf{X}')_{\hat{c}}. \tag{20}$$

912 *Proof.* Stacking coordinates into a vector gives:

$$\sum_{t,d} \varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d} \mid T) = \int_0^1 \nabla f(\gamma(\alpha))_{\hat{c}}^{\top} \gamma'(\alpha) d\alpha = \int_{\gamma} \nabla f_{\hat{c}} \cdot d\mathbf{s} = f(\gamma(1))_{\hat{c}} - f(\gamma(0))_{\hat{c}} \tag{21}$$

913 by the Fundamental Theorem of Line Integrals, since the integrand is the gradient field of the scalar
914 potential $f_{\hat{c}}$. Substituting $\gamma(1) = \mathbf{X}_{T-W+1:T}$ and $\gamma(0) = \mathbf{X}'$ completes the proof. \square

Proof of Theorem 2. By Lemma 5 we have:

$$\begin{aligned}
& \sum_{i=T_1-W+1}^{T_2} \sum_{d=1}^D \varphi_{\text{IG}}^{\gamma_1, 2}(f, \mathbf{X}_{i,d} \mid T_2) = f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W:T_1-1})_{\hat{c}}, \\
& \sum_{i=T_1-W+1}^{T_2} \sum_{d=1}^D \varphi_{\text{IG}}^{\gamma_2, 2}(f, \mathbf{X}_{i,d} \mid T_2) = f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_2-W:T_2-1})_{\hat{c}}, \\
& \sum_{i=T_1-W+1}^{T_2} \sum_{d=1}^D \varphi_{\text{IG}}^{\gamma_1, 1}(f, \mathbf{X}_{i,d} \mid T_1) = f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} - f(\mathbf{X}_{T_1-W:T_1-1})_{\hat{c}}, \\
& \sum_{i=T_1-W+1}^{T_2} \sum_{d=1}^D \varphi_{\text{IG}}^{\gamma_2, 1}(f, \mathbf{X}_{i,d} \mid T_1) = f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} - f(\mathbf{X}_{T_2-W:T_2-1})_{\hat{c}}.
\end{aligned} \tag{22}$$

SWING averages the attributions of the two paths to $\mathbf{X}_{T_2-W+1:T_2}$ and subtracts the average of the two paths to $\mathbf{X}_{T_1-W+1:T_1}$, yielding:

$$\begin{aligned} \sum_{t,d} \varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) &= \frac{1}{2} \sum_{i=1}^2 \left(\varphi_{\text{IG}}^{\gamma_{i,2}}(f, \mathbf{X}_{t,d} \mid T_2) - \varphi_{\text{IG}}^{\gamma_{i,1}}(f, \mathbf{X}_{t,d} \mid T_1) \right) \\ &= \frac{1}{2} \left[2f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - 2f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}} \right] = f(\mathbf{X}_{T_2-W+1:T_2})_{\hat{c}} - f(\mathbf{X}_{T_1-W+1:T_1})_{\hat{c}}, \end{aligned} \quad (23)$$

which establishes online completeness.

F.3 PROOF OF THEOREM 3

SWING is defined as an average of Integrated Gradients (IG) attributions computed along multiple paths γ . For a continuously differentiable curve $\gamma : [0, 1] \rightarrow \mathbb{R}^{W \times D}$ with $\gamma(0) = \mathbf{X}'$ (baseline) and $\gamma(1) = \mathbf{X}$ (input), the IG attribution is:

$$\varphi_{\text{IG}}^{\gamma}(f, \mathbf{X}_{t,d}) = \int_0^1 \frac{\partial \hat{f}_{\gamma}(\gamma(\alpha))}{\partial \mathbf{X}_{t,d}} \frac{\partial \gamma_{t,d}(\alpha)}{\partial \alpha} d\alpha. \quad (24)$$

Therefore, the resulting SWING attribution $\varphi_{\text{SWING}}(f, \mathbf{X}_{t,d})$, obtained by averaging over its designated paths, depends solely on the function $f(\mathbf{X})$ and not on the particular architecture or parameterization used to realize f . Consequently, any two models implementing the same function yield identical SWING attributions, establishing implementation invariance.

F.4 PROOF OF THEOREM 4

By definition,

$$\sum_{t,d} \varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2) = \frac{1}{2} \sum_{i=1}^2 \left(\sum_{t,d} \varphi_{\text{IG}}^{\gamma_{i,2}}(f, \mathbf{X}_{t,d} \mid T_2) - \sum_{t,d} \varphi_{\text{IG}}^{\gamma_{i,1}}(f, \mathbf{X}_{t,d} \mid T_1) \right). \quad (25)$$

Consider the reversed prediction change $T_2 \rightarrow T_1$. SWING uses the same four designated curves but traversed in reverse, so each IG term flips sign (antisymmetry of line integrals under reversed limits) and the two to T_2 terms and the two to T_1 terms swap roles. Hence,

$$\begin{aligned}
\sum_{t,d} \varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_2 \rightarrow T_1) &= \frac{1}{2} \sum_{i=1}^2 \left(- \sum_{t,d} \varphi_{\text{IG}}^{\gamma_{i,1}}(f, \mathbf{X}_{t,d} \mid T_1) + \sum_{t,d} \varphi_{\text{IG}}^{\gamma_{i,2}}(f, \mathbf{X}_{t,d} \mid T_2) \right) \\
&= - \sum_{t,d} \varphi_{\text{SWING}}(f, \mathbf{X}_{t,d} \mid T_1 \rightarrow T_2).
\end{aligned} \tag{26}$$

972 **Table 4:** We evaluate our method on five datasets: three real-world benchmarks—MIMIC-III (Johnson
 973 et al., 2016), PhysioNet 2019 (Reyna et al., 2020a), and Activity (Reiss & Stricker, 2012)—and two
 974 widely used synthetic datasets—Delayed Spike (Leung et al., 2021) and Switch-Feature (Tonekaboni
 975 et al., 2020; Liu et al., 2024b).

Type	Name	Task	# ID	# Sample / ID	Window Size	Feature	Class
Real-world	MIMIC-III	Decompensation prediction	6,221	5	48	32	2
	PhysioNet 2019	Sepsis prediction	8,066	5	48	40	2
	Activity	Human action recognition	5	200	50	12	7
Synthetic	Delayed Spike	Binary classification	1,000	5	40	3	2
	Switch-Feature	Binary classification	600	5	50	3	2

G DETAILS OF DATASETS

985 **MIMIC-III.** For the MIMIC-III Clinical Database (Johnson et al., 2016), we adopt the decompensa-
 986 tion prediction benchmark defined in Johnson et al. (2016). The dataset consists of over 41,000
 987 ICU stays from 2001–2012, with rich multivariate time series covering vital signs, laboratory values,
 988 interventions, and demographics. Following the benchmark setup, we use sliding windows of length
 989 48 hours with a prediction horizon of 24 hours, labeling an instance positive if the patient dies within
 990 the horizon. This yields roughly 2.5 million prediction windows, of which about 63,000 (2.5%) are
 991 positive. Each window contains irregularly sampled trajectories across up to 32 variables, making it a
 992 challenging setting for both temporal modeling and interpretability.

993 **PhysioNet 2019.** For the PhysioNet 2019 dataset (Reyna et al., 2020b), we adopt the sepsis
 994 prediction task defined under the Sepsis-3 criteria (Singer et al., 2016). The dataset comprises nearly
 995 40,000 ICU stays collected across multiple hospital systems, with each stay providing multivariate
 996 physiological time series such as vitals, labs, and demographics. We define prediction windows
 997 using a 48-hour observation length and a 12-hour prediction horizon, labeling a sample positive
 998 if sepsis onset occurs within the horizon. This yields around 1.1 million prediction instances, of
 999 which approximately 27,000 (2.5%) are positive. Compared to MIMIC-III, this dataset presents
 1000 higher variability in measurement density across hospitals, offering a complementary benchmark for
 1001 evaluating both predictive performance and explanation reliability.

1003 **Activity.** We adopt the Activity dataset (Frank, 2010) and follow the preprocessing protocol from
 1004 Latent ODEs (Rubanova et al., 2019). The dataset comprises 25 sequences from five individuals, each
 1005 having around 6,600 time points. We segment each sequence into overlapping windows of 50 time
 1006 points using a stride of 1 (unlike the stride 25 used in the original Latent ODEs paper). Labels are
 1007 provided at each time point across 11 fine-grained actions; to reduce ambiguity, we merge them into
 1008 seven coarse classes as in (Rubanova et al., 2019): *walking, falling, lying, sitting, standing up, on all*
 1009 *fours, and sitting on the ground*. For evaluation, we split by individual: the first three for training, the
 1010 fourth for validation, and the fifth for testing.

1011 **Delayed Spike.** We adopt the Delayed Spike dataset (Leung et al., 2021), a variant of the Spike
 1012 benchmark originally introduced by Tonekaboni et al. (2020) and later extended by Leung et al.
 1013 (2021). The standard Spike dataset consists of three multivariate NARMA sequences with added
 1014 linear trends and random spikes, where the label flips from 0 to 1 immediately after a spike appears in
 1015 feature 0 and remains positive thereafter. In the Delayed Spike version, however, the label transition is
 1016 shifted by two time steps: it becomes 1 exactly two steps after the spike in feature 0. This modification
 1017 forces explanation methods to correctly identify the causal spike event rather than simply aligning
 1018 with the delayed label change.

1019 **Switch-Feature.** We generate the Switch-Feature dataset following the design in FIT (Tonekaboni
 1020 et al., 2020). Similar to the State dataset, it is constructed based on a three-state hidden Markov model
 1021 with an initial distribution $\pi = (1/3, 1/3, 1/3)$ and the following transition matrix:

$$\begin{pmatrix} 0.95 & 0.02 & 0.03 \\ 0.02 & 0.95 & 0.03 \\ 0.03 & 0.02 & 0.95 \end{pmatrix}.$$

1026 Each hidden state emits a time series from a Gaussian process with an RBF kernel ($\gamma = 0.2$) and
 1027 a fixed marginal variance of 0.1 for all features. The mean vectors of the three states are given
 1028 by $\mu_1 = [0.8, 0.5, 0.2]$, $\mu_2 = [0, 1.0, 0]$, and $\mu_3 = [0.2, 0.2, 0.8]$. Labels $y_i[t]$ are sampled from a
 1029 Bernoulli distribution $\text{Bernoulli}(p_i[t])$, where:

$$1030 \quad 1031 \quad 1032 \quad 1033 \quad p_i[t] = \begin{cases} (1 + \exp(-\mathbf{X}_{t,1}))^{-1} & \text{if } s_t = 0, \\ (1 + \exp(-\mathbf{X}_{t,2}))^{-1} & \text{if } s_t = 1, \\ (1 + \exp(-\mathbf{X}_{t,3}))^{-1} & \text{if } s_t = 2. \end{cases} \quad (27)$$

1034 and s_t denotes the latent state at time t . In our experiments, we generate sequences of length 100 and
 1035 extract multiple online prediction samples from each sequence using a fixed-length sliding window
 1036 of size 50.

1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080 H DETAILS OF ABLATION STUDY

1081
 1082 In this section, we clarify the meaning of each ablation configuration and present its corresponding
 1083 mathematical formulation. We independently verify the contribution of each SWING component
 1084 by removing 1) both RBS and PHI, 2) RBS alone, and 3) DPI ($\gamma_{1,1}, \gamma_{2,2}$). The following variants
 1085 correspond directly to the results in Table 3 and ensure that each ablation is rigorously defined and
 1086 remains faithful to the original algorithm.

1087
 1088 **SWING w/o RBS, PHI.** This variant removes both the retrospective baseline selection (RBS) and
 1089 the piecewise-linear historical integration (PHI), leaving only DPI active. Both integration paths
 1090 start from the zero baseline, and the resulting attribution reduces to subtracting Integrated Gradients
 1091 computed at T_2 and T_1 :

$$\varphi_{\text{SWING (w/o RBS, PHI)}} = \varphi_{\text{IG}}^{\gamma_2}(f, X_{t,d} | T_2) - \varphi_{\text{IG}}^{\gamma_1}(f, X_{t,d} | T_1), \quad (28)$$

1092 where each path is a straight line from zero to the corresponding input window $\gamma_i(\alpha) =$
 1093 $\alpha X_{T_i-W+1:T_i}$. This configuration is mathematically equivalent to applying IG independently at
 1094 two time points and comparing their attributions.

1095
 1096 **SWING w/o RBS.** This configuration removes only the baseline-selection component. The retro-
 1097 spective baselines $X_{T_i-W:T_i-1}$ are replaced with the zero baseline, while all DPI mechanisms and
 1098 the PHI interpolation remain unchanged. As a result, only the initial segment of each integration path
 1099 is modified: the path originates at zero rather than at $X_{T_i-W:T_i-1}$, but all subsequent interpolation
 1100 segments and the dual-path structure are preserved. This ablation isolates the effect of removing
 1101 temporal baseline adaptation while retaining both historical integration and dual-path comparison.

1102
 1103 **SWING w/o DPI ($\gamma_{1,1}, \gamma_{2,2}$).** This variant removes the self-window DPI components $\gamma_{1,1}$ and
 1104 $\gamma_{2,2}$, while keeping the cross-window paths $\gamma_{1,2}$ and $\gamma_{2,1}$ active. Since the cross-window paths are
 1105 essential for attributing prediction changes between T_1 and T_2 , PHI remains applied to them, yielding
 1106 the following attribution:

$$\varphi_{\text{SWING (w/o DPI)}} = \varphi_{\text{IG}}^{\gamma_{2,1}}(f, X_{t,d} | T_2) - \varphi_{\text{IG}}^{\gamma_{1,2}}(f, X_{t,d} | T_1). \quad (29)$$

1107 By eliminating only the self-window paths while retaining the cross-window integration, this ablation
 1108 isolates the contribution of DPI to SWING’s temporal attribution mechanism.

1111 I ADDITIONAL EXPERIMENTS

1112 This section presents additional experiments that complement the main findings. The content is
 1113 organized into extended quantitative evaluations, backbone generalization, longer prediction intervals,
 1114 ablation and robustness analyses, efficiency, qualitative case studies, and clinical coherence.

1115
 1116 **Extended benchmarks.** Tables 6 and 7 provide further results on clinical datasets, synthetic
 1117 benchmarks, and the Activity dataset. Across all settings, SWING achieves the best or second-best
 1118 scores on most metrics, reinforcing its robustness across both controlled and real-world scenarios.

1119
 1120 **Backbone architectures.** Table 8 compares CNN and Transformer backbones on MIMIC-III.
 1121 The results show that SWING retains strong performance across both architectures, demonstrating
 1122 versatility beyond LSTMs.

1123
 1124 **Longer time intervals.** Table 9 evaluates explanations at $T_2 - T_1 = 6$ and 24. Although CPD and
 1125 AUPD converge across methods with longer intervals, SWING consistently achieves best or near-best
 1126 results, and shows particularly dominant gains on preservation metrics.

1127
 1128 **Ablation and robustness.** Table 3 analyzes the contributions of RBS, PHI, and DPI. Removing RBS
 1129 or PHI degrades performance, confirming their complementary roles, while DPI mainly stabilizes
 1130 preservation. Baseline offset analysis further shows that $d = 1$ yields the best trade-off, with $d = 0$
 1131 and larger offsets leading to degradation. Figure 4 confirms robustness across a broad range of
 1132 n_{samples} . Subgroup and temporal resolution analyses (Tables 10 to 12) additionally show that SWING
 1133 preserves its advantages regardless of prediction or label changes and remains stable across both 24-
 and 72-length windows.

1134 **Table 5:** Evaluation metrics for prediction-change explanations on the MIMIC-III decompensation
 1135 benchmark with an LSTM backbone. We follow the same setting as the main experiments, but
 1136 evaluate by removing the most salient 50 feature points per time step with forward-fill substitution.

Algorithm	CPD \uparrow	AUPD \uparrow	Corr. \uparrow
Dynamask w/ Naive Subtraction	0.14 \pm 0.00	0.12 \pm 0.00	-0.19 \pm 0.00
Dynamask in Delta-XAI	11.72\pm0.08	7.56\pm0.04	0.04\pm0.00

1141 **Table 6:** Performance comparison of XAI methods on clinical prediction tasks: PhysioNet 2019
 1142 sepsis benchmark using LSTM as backbone architecture. Evaluation is performed by removing the
 1143 most or least salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
LIME (Ribeiro et al., 2016)	0.29 \pm 0.00	0.21 \pm 0.00	1.83 \pm 0.01	1.02 \pm 0.00	3.66 \pm 0.01	1.82 \pm 0.00	3.96 \pm 0.00	2.09 \pm 0.01	-0.08 \pm 0.00
GradSHAP (Lundberg, 2017)	1.60 \pm 0.00	0.93 \pm 0.00	2.52 \pm 0.01	1.50 \pm 0.00	3.48 \pm 0.00	1.64 \pm 0.00	3.75 \pm 0.01	1.86 \pm 0.00	0.02 \pm 0.00
IG (Sundararajan et al., 2017)	2.68 \pm 0.00	1.54 \pm 0.00	3.00 \pm 0.01	1.93 \pm 0.00	3.19 \pm 0.01	1.43 \pm 0.00	3.43 \pm 0.01	1.61 \pm 0.00	0.11 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	2.72 \pm 0.00	1.60 \pm 0.00	2.87 \pm 0.00	1.86 \pm 0.00	3.06 \pm 0.01	1.35 \pm 0.00	3.38 \pm 0.01	1.56 \pm 0.00	0.16 \pm 0.00
FO (Suresh et al., 2017)	1.17 \pm 0.00	0.78 \pm 0.00	2.75 \pm 0.01	1.80 \pm 0.00	4.63 \pm 0.01	3.13 \pm 0.01	2.44 \pm 0.01	1.13 \pm 0.00	0.02 \pm 0.00
AFO (Tonekaboni et al., 2020)	1.87 \pm 0.00	1.16 \pm 0.00	2.88 \pm 0.01	1.86 \pm 0.00	3.12 \pm 0.00	1.48 \pm 0.00	3.28 \pm 0.01	1.55 \pm 0.00	0.15 \pm 0.00
FIT (Tonekaboni et al., 2020)	0.71 \pm 0.00	0.48 \pm 0.00	1.10 \pm 0.00	0.91 \pm 0.00	3.81 \pm 0.00	1.96 \pm 0.00	1.30\pm0.00	1.02\pm0.00	0.00 \pm 0.00
WinIT (Leung et al., 2021)	1.86 \pm 0.00	1.10 \pm 0.00	2.85 \pm 0.01	1.66 \pm 0.00	3.88 \pm 0.00	2.02 \pm 0.00	3.12 \pm 0.01	1.64 \pm 0.00	0.06 \pm 0.00
Dynamask (Crabbé & Van Der Schaar, 2021)	1.69 \pm 0.00	1.11 \pm 0.00	2.05 \pm 0.00	1.36 \pm 0.00	4.98 \pm 0.01	2.77 \pm 0.00	4.74 \pm 0.01	2.59 \pm 0.00	0.06 \pm 0.00
Extrmask (Enguehard, 2023)	1.16 \pm 0.00	0.72 \pm 0.00	1.91 \pm 0.00	1.11 \pm 0.00	4.04 \pm 0.01	2.41 \pm 0.00	3.98 \pm 0.01	2.36 \pm 0.01	0.05 \pm 0.00
ContralSP (Liu et al., 2024b)	0.73 \pm 0.02	0.34 \pm 0.01	2.47 \pm 0.02	1.32 \pm 0.01	5.32 \pm 0.00	2.93 \pm 0.04	5.35 \pm 0.04	2.97 \pm 0.04	0.04 \pm 0.00
TimeX (Queen et al., 2024)	0.74 \pm 0.00	0.36 \pm 0.00	1.96 \pm 0.01	1.00 \pm 0.00	5.33 \pm 0.01	2.77 \pm 0.00	5.49 \pm 0.01	2.89 \pm 0.01	-0.03 \pm 0.00
TimeX++ (Liu et al., 2024a)	1.76 \pm 0.01	1.06 \pm 0.00	2.07 \pm 0.01	1.19 \pm 0.00	3.99 \pm 0.00	1.81 \pm 0.00	4.02 \pm 0.01	1.82 \pm 0.00	0.05 \pm 0.00
TIMING (Jang et al., 2025)	2.73 \pm 0.00	1.56 \pm 0.00	3.02\pm0.00	1.94\pm0.00	3.12 \pm 0.00	1.41 \pm 0.00	3.38 \pm 0.00	1.60 \pm 0.00	0.13 \pm 0.00
SWING	3.10\pm0.01	1.96\pm0.00	2.81 \pm 0.01	1.78 \pm 0.01	2.27\pm0.01	0.93\pm0.00	2.38\pm0.00	1.01\pm0.00	0.32\pm0.02

1156 **Table 7:** Performance of XAI methods on (top) Activity (Vidulin et al., 2010), (middle) Delayed
 1157 Spike (Leung et al., 2021), and (bottom) Switch-Feature (Tonekaboni et al., 2020) benchmarks with
 1158 an LSTM backbone. Evaluation is performed by removing the most or least salient 50 feature points
 1159 per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	19.80 \pm 0.53	12.53 \pm 0.31	7.81 \pm 0.53	4.76 \pm 0.28	17.51 \pm 0.70	7.33 \pm 0.28	20.23 \pm 1.06	8.77 \pm 0.45	0.09 \pm 0.01
DeepLIFT (Shrikumar et al., 2017)	21.01\pm0.59	13.25 \pm 0.36	8.15 \pm 0.52	5.10\pm0.30	15.26 \pm 0.49	6.22 \pm 0.17	19.07 \pm 0.61	8.20 \pm 0.23	0.17\pm0.02
AFO (Tonekaboni et al., 2020)	15.20 \pm 0.52	10.11 \pm 0.29	7.51 \pm 0.38	4.79 \pm 0.25	13.69 \pm 0.92	5.66\pm0.35	18.23 \pm 0.77	7.99 \pm 0.28	0.17\pm0.01
WinIT (Leung et al., 2021)	6.62 \pm 0.38	3.34 \pm 0.15	6.61 \pm 0.42	3.87 \pm 0.25	24.10 \pm 1.26	12.61 \pm 0.63	20.16 \pm 1.09	9.45 \pm 0.54	0.09 \pm 0.01
TIMING (Jang et al., 2025)	13.75 \pm 0.71	8.08 \pm 0.37	7.18 \pm 0.54	4.35 \pm 0.35	19.39 \pm 1.06	9.28 \pm 0.53	19.44 \pm 0.74	8.69 \pm 0.33	0.08 \pm 0.01
SWING	21.77\pm0.80	14.96\pm0.63	9.20\pm0.83	6.19\pm0.60	7.64\pm0.33	2.70\pm0.13	15.15\pm0.37	6.05\pm0.12	0.46\pm0.03
IG (Sundararajan et al., 2017)	290.74 \pm 1.01	246.74 \pm 1.03	281.63 \pm 1.67	225.76 \pm 1.20	13.14 \pm 0.22	4.30 \pm 0.07	15.78 \pm 0.32	5.59 \pm 0.10	0.46 \pm 0.01
DeepLIFT (Shrikumar et al., 2017)	289.95 \pm 1.37	251.30 \pm 1.05	277.38 \pm 1.08	228.97\pm0.71	11.47 \pm 0.23	3.69 \pm 0.08	14.85 \pm 0.22	5.11 \pm 0.08	0.58\pm0.00
AFO (Tonekaboni et al., 2020)	256.70 \pm 1.90	189.27 \pm 1.42	240.04 \pm 1.70	165.29 \pm 1.22	7.83 \pm 0.19	2.66\pm0.06	12.29 \pm 0.33	4.46 \pm 0.19	0.19 \pm 0.00
WinIT (Leung et al., 2021)	426.69\pm5.33	114.50 \pm 1.37	191.86 \pm 1.38	56.91 \pm 0.31	398.05 \pm 4.40	344.09 \pm 3.76	180.94 \pm 1.96	46.32 \pm 0.48	-0.02 \pm 0.00
TIMING (Jang et al., 2025)	327.29 \pm 3.46	199.72 \pm 2.20	205.82 \pm 2.33	96.03 \pm 0.94	261.30 \pm 2.42	230.69 \pm 2.10	205.20 \pm 2.37	92.31 \pm 0.97	-0.01 \pm 0.00
SWING	317.95 \pm 2.50	252.16\pm1.86	305.09\pm2.06	232.50\pm1.42	3.35\pm0.04	1.13\pm0.02	5.70\pm0.10	1.85\pm0.04	0.44 \pm 0.00
IG (Sundararajan et al., 2017)	464.24 \pm 6.43	422.12 \pm 5.83	474.36 \pm 5.58	395.55 \pm 4.59	10.97 \pm 0.13	4.29 \pm 0.05	17.94 \pm 0.61	6.56 \pm 0.16	0.65 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	525.16 \pm 7.67	470.90 \pm 6.67	523.95 \pm 5.31	430.36 \pm 4.01	8.66 \pm 0.10	3.45\pm0.04	16.35 \pm 0.56	5.97 \pm 0.21	0.63 \pm 0.01
AFO (Tonekaboni et al., 2020)	533.94 \pm 7.78	475.55 \pm 6.78	529.43\pm6.72	431.48 \pm 4.92	9.88 \pm 0.21	3.98 \pm 0.10	13.30\pm0.30	5.37\pm0.14	0.58 \pm 0.00
WinIT (Leung et al., 2021)	507.88 \pm 5.15	447.50 \pm 5.05	496.19 \pm 7.13	391.83 \pm 6.16	12.46 \pm 0.19	7.46 \pm 0.10	311.71 \pm 5.23	113.50 \pm 1.90	0.53 \pm 0.01
TIMING (Jang et al., 2025)	418.65 \pm 3.35	391.51 \pm 3.17	446.03 \pm 4.22	373.96 \pm 3.50	13.97 \pm 0.22	7.54 \pm 0.12	507.71 \pm 6.18	255.11 \pm 3.26	0.64 \pm 0.01
SWING	536.43\pm6.46	481.90\pm5.86	519.55 \pm 5.26	436.18\pm4.42	7.32\pm0.10	2.99\pm0.06	15.84\pm0.53	5.81\pm0.20	0.74\pm0.00

1174 **Efficiency.** Figure 3 demonstrates that SWING achieves state-of-the-art explanatory quality without
 1175 significant computational overhead. Runtime (0.35s/sample) and memory (448 MB) remain
 1176 comparable to gradient-based baselines such as IG and GradSHAP.

1177 **Qualitative and clinical coherence.** Figures 5 to 9 visualize attribution maps, where SWING
 1178 produces sharper, localized explanations than surrogate or masking methods. Finally, Figure 10
 1179 confirms coherence with medical knowledge, correctly highlighting risk increases from SBP drops
 1180 and pH declines, and protective effects from SpO₂ rises.

1181 Overall, these supplementary experiments confirm that SWING consistently delivers reliable, efficient,
 1182 and clinically meaningful explanations across diverse datasets, models, and experimental conditions.

1188

1189

1190

1191

1192

1193

Table 8: Performance comparison of XAI methods on MIMIC-III decompensation prediction task with different backbone architectures: CNN (top) and Transformer (bottom) architectures. Evaluation is performed by removing the most or least salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	28.68 \pm 0.04	16.83 \pm 0.03	32.16 \pm 0.09	20.25 \pm 0.06	41.37 \pm 0.06	19.08 \pm 0.04	37.25 \pm 0.05	15.63 \pm 0.04	0.24 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	29.59 \pm 0.06	17.60 \pm 0.08	32.14 \pm 0.13	20.29 \pm 0.08	41.69 \pm 0.12	19.13 \pm 0.05	37.68 \pm 0.10	15.65 \pm 0.03	0.28 \pm 0.00
AFO (Tonekaboni et al., 2020)	20.64 \pm 0.14	13.42 \pm 0.09	32.35 \pm 0.15	20.16 \pm 0.09	50.97 \pm 0.12	25.25 \pm 0.07	34.30 \pm 0.08	14.49 \pm 0.04	0.30 \pm 0.00
WinIT (Leung et al., 2021)	23.15 \pm 0.07	13.35 \pm 0.04	32.69 \pm 0.09	18.44 \pm 0.06	55.61 \pm 0.16	25.09 \pm 0.08	42.21 \pm 0.07	19.95 \pm 0.04	0.23 \pm 0.00
TIMING (Jang et al., 2025)	29.80 \pm 0.06	17.26 \pm 0.05	32.33 \pm 0.12	20.45 \pm 0.08	40.54 \pm 0.10	18.86 \pm 0.05	36.10 \pm 0.09	15.26 \pm 0.06	0.26 \pm 0.00
SWING	44.02\pm0.17	27.85\pm0.12	40.36\pm0.12	25.86\pm0.10	24.47\pm0.07	8.23\pm0.06	24.97\pm0.03	8.44\pm0.03	0.53\pm0.00

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	16.33 \pm 0.06	9.77 \pm 0.05	15.26 \pm 0.07	10.42 \pm 0.05	33.52 \pm 0.21	14.58 \pm 0.12	27.42 \pm 0.10	10.42 \pm 0.07	0.15 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	15.11 \pm 0.07	9.87 \pm 0.06	15.40 \pm 0.07	10.53 \pm 0.05	33.24 \pm 0.22	14.13 \pm 0.11	28.24 \pm 0.17	10.40 \pm 0.10	0.14 \pm 0.00
AFO (Tonekaboni et al., 2020)	14.61 \pm 0.06	10.22 \pm 0.03	17.51 \pm 0.03	12.06 \pm 0.03	37.26 \pm 0.19	18.10 \pm 0.09	23.78 \pm 0.16	9.23 \pm 0.10	0.23 \pm 0.00
WinIT (Leung et al., 2021)	17.20 \pm 0.05	10.04 \pm 0.04	23.99\pm0.14	14.12 \pm 0.08	34.44 \pm 0.12	16.04 \pm 0.09	25.92 \pm 0.10	12.36 \pm 0.06	0.12 \pm 0.00
TIMING (Jang et al., 2025)	22.53 \pm 0.05	12.59\pm0.03	15.95 \pm 0.07	10.94 \pm 0.05	31.53 \pm 0.13	14.04 \pm 0.08	25.78 \pm 0.16	10.01 \pm 0.08	0.18 \pm 0.00
SWING	25.50\pm0.13	16.38\pm0.09	23.71 \pm 0.15	15.84\pm0.09	17.57\pm0.13	5.98\pm0.06	18.05\pm0.13	6.30\pm0.07	0.26\pm0.00

1205

1206

1207

1208

Table 9: Performance comparison of XAI methods on MIMIC-III decompensation prediction task with various time interval settings between two time steps $T_1 < T_2$: 6 timestamps interval (top) and 24 timestamps interval (bottom) using LSTM as backbone architecture. Evaluation is performed by removing the most or least salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	44.17 \pm 0.15	26.92 \pm 0.10	45.53 \pm 0.23	29.23 \pm 0.13	91.67 \pm 0.46	39.68 \pm 0.17	76.67 \pm 0.28	28.60 \pm 0.14	0.22 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	42.00 \pm 0.18	25.91 \pm 0.10	44.50 \pm 0.22	28.54 \pm 0.12	100.64 \pm 0.54	41.62 \pm 0.21	84.92 \pm 0.43	30.38 \pm 0.18	0.19 \pm 0.00
AFO (Tonekaboni et al., 2020)	35.21 \pm 0.10	22.98 \pm 0.07	47.47\pm0.16	30.01 \pm 0.09	106.36 \pm 0.57	49.81 \pm 0.23	65.85\pm0.47	25.13\pm0.22	0.22 \pm 0.00
WinIT (Leung et al., 2021)	52.98\pm0.02	28.48 \pm 0.01	19.19 \pm 0.01	10.47 \pm 0.00	91.22 \pm 0.09	43.19 \pm 0.04	83.18 \pm 0.51	41.23 \pm 0.25	0.14 \pm 0.00
TIMING (Jang et al., 2025)	52.76 \pm 0.23	30.71 \pm 0.13	47.30 \pm 0.17	30.32 \pm 0.12	87.90 \pm 0.29	38.52 \pm 0.14	72.48 \pm 0.35	27.26 \pm 0.15	0.23 \pm 0.00
SWING	51.46\pm0.24	33.55\pm0.14	55.64\pm0.36	35.81\pm0.24	61.90\pm0.29	20.75\pm0.11	63.95\pm0.19	21.90\pm0.11	0.29\pm0.00

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	45.75 \pm 0.19	26.81 \pm 0.14	40.79 \pm 0.24	26.37 \pm 0.15	103.48 \pm 0.45	48.12 \pm 0.20	77.63 \pm 0.44	29.38 \pm 0.15	0.20 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	40.53 \pm 0.18	24.80 \pm 0.12	40.16 \pm 0.26	25.96 \pm 0.24	109.50 \pm 0.52	49.11 \pm 0.22	83.81 \pm 0.39	30.47 \pm 0.15	0.18 \pm 0.00
AFO (Tonekaboni et al., 2020)	31.04 \pm 0.21	20.46 \pm 0.16	42.44 \pm 0.23	26.92 \pm 0.15	120.00 \pm 0.30	61.17 \pm 0.12	64.36\pm0.35	25.25 \pm 0.08	0.18 \pm 0.00
TIMING (Jang et al., 2025)	58.83\pm0.30	32.38\pm0.22	42.52 \pm 0.26	27.42 \pm 0.17	97.93 \pm 0.41	46.75 \pm 0.18	72.36 \pm 0.35	27.82 \pm 0.15	0.22\pm0.00
SWING	41.07 \pm 0.22	26.46 \pm 0.18	50.58\pm0.28	32.29\pm0.20	60.29\pm0.14	21.60\pm0.08	64.43 \pm 0.19	23.87\pm0.10	0.21 \pm 0.00

1224

1225

1226

Table 10: Subgroup analysis of attribution methods under different score-change conditions (high: $|\Delta| \geq 0.1$ (top), low: $|\Delta| < 0.1$ (bottom), where $\Delta = f(\mathbf{X}_{T_2-W+1:T_2}) - f(\mathbf{X}_{T_1-W+1:T_1})$). Results are reported on the MIMIC-III decompensation benchmark with an LSTM backbone with $T_2 - T_1 = 1$. Evaluation is performed by removing the most or least salient 50 feature points per time step, using forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	281.32 \pm 0.42	222.60 \pm 0.37	319.69 \pm 0.75	251.67 \pm 0.58	381.59 \pm 0.71	186.83 \pm 0.36	349.24 \pm 0.73	147.10 \pm 0.22	0.37 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	296.81 \pm 0.72	236.89 \pm 0.50	326.36 \pm 0.43	257.72 \pm 0.42	371.51 \pm 0.66	181.49 \pm 0.35	334.79 \pm 0.64	140.56 \pm 0.45	0.29 \pm 0.00
AFO (Tonekaboni et al., 2020)	279.71 \pm 0.31	224.13 \pm 0.19	329.23 \pm 0.83	257.53 \pm 0.36	371.05 \pm 0.86	190.70 \pm 0.78	316.06\pm0.52	133.02 \pm 0.38	0.42 \pm 0.00
WinIT (Leung et al., 2021)	341.14 \pm 0.65	244.84 \pm 0.35	408.03\pm0.65	283.17 \pm 0.26	350.94 \pm 1.06	182.03 \pm 0.51	366.12 \pm 1.34	194.42 \pm 0.77	0.23 \pm 0.00
TIMING (Jang et al., 2025)	285.87 \pm 0.53	224.72 \pm 0.42	322.54 \pm 0.81	253.72 \pm 0.76	378.87 \pm 1.29	184.84 \pm 0.49	346.25 \pm 0.39	145.34 \pm 0.21	0.37 \pm 0.00
SWING	397.28\pm0.73	311.94\pm0.52	388.43 \pm 0.93	308.56\pm0.79	261.28\pm0.55	96.11\pm0.24	264.34\pm0.77	98.10\pm0.29	0.57\pm0.00

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	11.44 \pm 0.								

1242
1243
1244
1245
1246
12471248 **Table 11:** Subgroup analysis of attribution methods with respect to prediction stability. Cases are
1249 divided into *changed* (top) vs. *unchanged* (bottom), depending on whether the predicted class label
1250 (*i.e.*, $\arg \max f(\mathbf{X}_{T_1-W+1:T_1})$ vs. $\arg \max f(\mathbf{X}_{T_2-W+1:T_2})$) is different or remains the same. Results
1251 are reported on the MIMIC-III decompensation benchmark with an LSTM backbone ($T_2 - T_1 = 1$).
1252 Evaluation is performed by removing the most or least salient 50 feature points per time step, using
1253 forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	227.09 \pm 0.27	177.97 \pm 0.23	256.60 \pm 0.81	202.19 \pm 0.62	277.70 \pm 0.62	138.81 \pm 0.36	252.27 \pm 0.90	107.53 \pm 0.38	0.45 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	230.31 \pm 0.22	181.47 \pm 0.20	256.18 \pm 0.38	202.44 \pm 0.24	272.82 \pm 0.68	136.23 \pm 0.42	246.01 \pm 0.77	104.30 \pm 0.42	0.32 \pm 0.00
AFO (Tonekaboni et al., 2020)	216.81 \pm 0.49	171.28 \pm 0.48	260.44 \pm 0.56	203.22 \pm 0.48	276.04 \pm 1.46	145.05 \pm 1.03	228.50 \pm 1.23	98.48 \pm 0.29	0.45 \pm 0.00
WinIT (Leung et al., 2021)	290.84 \pm 0.41	207.59 \pm 0.35	340.70 \pm 0.77	238.05 \pm 0.40	252.09 \pm 0.64	130.44 \pm 0.50	269.08 \pm 1.12	141.56 \pm 0.70	0.27 \pm 0.00
TIMING (Jang et al., 2025)	230.00 \pm 0.57	179.32 \pm 0.32	258.60 \pm 0.55	203.50 \pm 0.43	276.34 \pm 0.92	137.86 \pm 0.61	249.23 \pm 0.72	106.18 \pm 0.56	0.44 \pm 0.00
SWING	323.55\pm0.25	253.80\pm0.20	311.35\pm0.92	246.80\pm0.74	185.96\pm0.45	69.09\pm0.22	189.32\pm0.58	70.43\pm0.34	0.62\pm0.00
Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	12.39 \pm 0.08	8.26 \pm 0.05	15.01 \pm 0.05	10.36 \pm 0.03	31.93 \pm 0.31	13.25 \pm 0.12	27.87 \pm 0.24	10.14 \pm 0.07	0.13 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	12.61 \pm 0.07	8.60 \pm 0.05	14.88 \pm 0.05	10.32 \pm 0.03	34.14 \pm 0.34	13.79 \pm 0.12	29.75 \pm 0.29	10.70 \pm 0.08	0.16 \pm 0.00
AFO (Tonekaboni et al., 2020)	12.36 \pm 0.08	8.54 \pm 0.05	16.04 \pm 0.09	11.03 \pm 0.06	34.50 \pm 0.30	15.86 \pm 0.11	22.49 \pm 0.26	8.68 \pm 0.07	0.24 \pm 0.00
WinIT (Leung et al., 2021)	18.59 \pm 0.11	11.47 \pm 0.06	23.41 \pm 0.12	14.40 \pm 0.07	27.50 \pm 0.15	12.20 \pm 0.08	24.66 \pm 0.12	11.15 \pm 0.07	0.20 \pm 0.00
TIMING (Jang et al., 2025)	14.07 \pm 0.08	8.92 \pm 0.06	15.37 \pm 0.08	10.58 \pm 0.06	29.36 \pm 0.32	12.58 \pm 0.12	25.55 \pm 0.29	9.57 \pm 0.09	0.15 \pm 0.00
SWING	22.54\pm0.25	15.05\pm0.17	20.94\pm0.15	14.38\pm0.13	16.44\pm0.25	5.39\pm0.07	17.02\pm0.24	5.65\pm0.08	0.36\pm0.00

1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
12751276 **Table 12:** Performance of XAI methods on the MIMIC-III dataset with an LSTM backbone, shown
1277 for window size 24 (top) and window size 72 (bottom). We evaluate by removing the most or least
1278 salient 50 points with forward-fill substitution.

Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	11.98 \pm 2.47	7.84 \pm 1.59	13.91 \pm 3.00	9.40 \pm 2.04	24.08 \pm 6.42	10.69 \pm 2.64	20.67 \pm 6.05	8.35 \pm 2.36	0.21 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	12.40 \pm 2.67	8.31 \pm 1.79	13.93 \pm 3.07	9.45 \pm 2.11	25.33 \pm 7.01	11.05 \pm 2.79	22.36 \pm 6.85	8.90 \pm 2.62	0.24 \pm 0.00
AFO (Tonekaboni et al., 2020)	13.35 \pm 2.51	9.05 \pm 1.71	17.26 \pm 3.32	11.60 \pm 2.24	32.11 \pm 7.22	15.48 \pm 3.30	22.38 \pm 5.35	9.29 \pm 2.17	0.26 \pm 0.00
WinIT (Leung et al., 2021)	16.75 \pm 3.06	10.36 \pm 1.91	22.21 \pm 4.27	13.44 \pm 2.64	27.08 \pm 6.77	12.70 \pm 3.08	25.44 \pm 6.00	12.24 \pm 2.85	0.22 \pm 0.00
TIMING (Jang et al., 2025)	13.35 \pm 3.07	8.43 \pm 1.87	14.22 \pm 3.16	9.59 \pm 2.15	22.42 \pm 5.90	10.21 \pm 2.49	19.27 \pm 5.53	7.91 \pm 2.19	0.23 \pm 0.00
SWING	19.33\pm4.59	12.89\pm2.96	17.96\pm4.25	12.22\pm2.88	13.20\pm4.06	4.37\pm1.35	14.21\pm4.42	4.91\pm1.52	0.49\pm0.03
Algorithm	Removal of Most Salient 50 Points				Removal of Least Salient 50 Points				Corr. \uparrow
	CPD \uparrow	AUPD \uparrow	MPD \uparrow	AUMPD \uparrow	CPP \downarrow	AUPP \downarrow	MPP \downarrow	AUMPP \downarrow	
IG (Sundararajan et al., 2017)	8.71 \pm 0.05	5.93 \pm 0.04	10.23 \pm 0.06	7.16 \pm 0.05	15.64 \pm 0.08	7.27 \pm 0.03	12.40 \pm 0.07	4.76 \pm 0.02	0.23 \pm 0.00
DeepLIFT (Shrikumar et al., 2017)	9.28 \pm 0.02	6.39 \pm 0.02	10.35 \pm 0.05	7.29 \pm 0.04	16.37 \pm 0.07	7.46 \pm 0.03	13.05 \pm 0.07	4.91 \pm 0.01	0.26 \pm 0.00
AFO (Tonekaboni et al., 2020)	8.73 \pm 0.05	6.05 \pm 0.04	10.86 \pm 0.04	7.56 \pm 0.04	16.11 \pm 0.10	8.09 \pm 0.04	10.57 \pm 0.05	4.17 \pm 0.02	0.29 \pm 0.00
WinIT (Leung et al., 2021)	11.44 \pm 0.10	7.22 \pm 0.06	14.88\pm0.08	9.17\pm0.07	12.00 \pm 0.06	5.77 \pm 0.02	11.56 \pm 0.05	5.53 \pm 0.03	0.25 \pm 0.00
TIMING (Jang et al., 2025)	9.14 \pm 0.06	6.08 \pm 0.05	10.36 \pm 0.07	7.24 \pm 0.05	14.73 \pm 0.06	7.00 \pm 0.03	11.55 \pm 0.06	4.51 \pm 0.02	0.24 \pm 0.00
SWING	13.79\pm0.07	9.46\pm0.05	13.05\pm0.07	9.10\pm0.06	7.51\pm0.04	2.45\pm0.02	7.75\pm0.04	2.56\pm0.02	0.48\pm0.00

1291
1292
1293
1294
1295

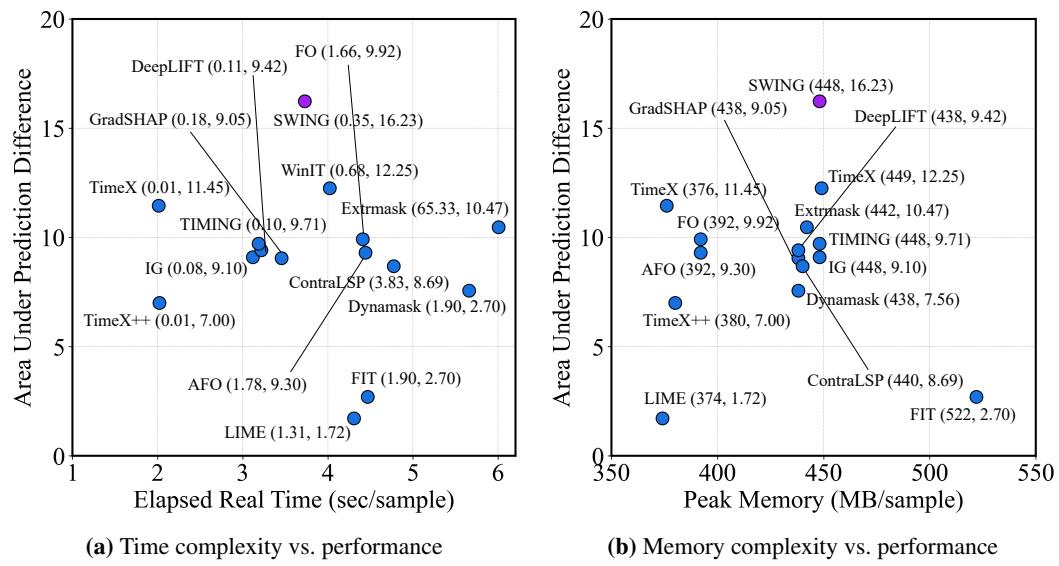


Figure 3: Computational efficiency analysis comparing SWING with baselines on the MIMIC-III benchmark. (a) Elapsed real time per sample (sec/sample, log-scale) versus AUPD ($K = 50$). (b) GPU peak memory consumption per sample (MB/sample) versus AUPD ($K = 50$).

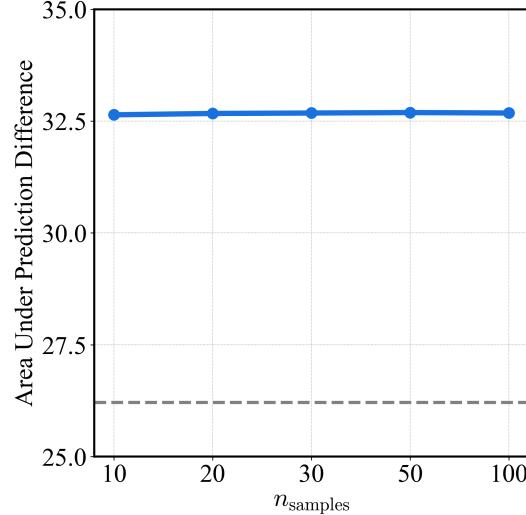


Figure 4: Hyperparameter sensitivity of SWING with respect to $n_{samples}$, compared to IG (dotted gray line).

1350 **Table 13:** List of clinical features used from the MIMIC-III dataset, including their indices and
 1351 descriptive names for model input.

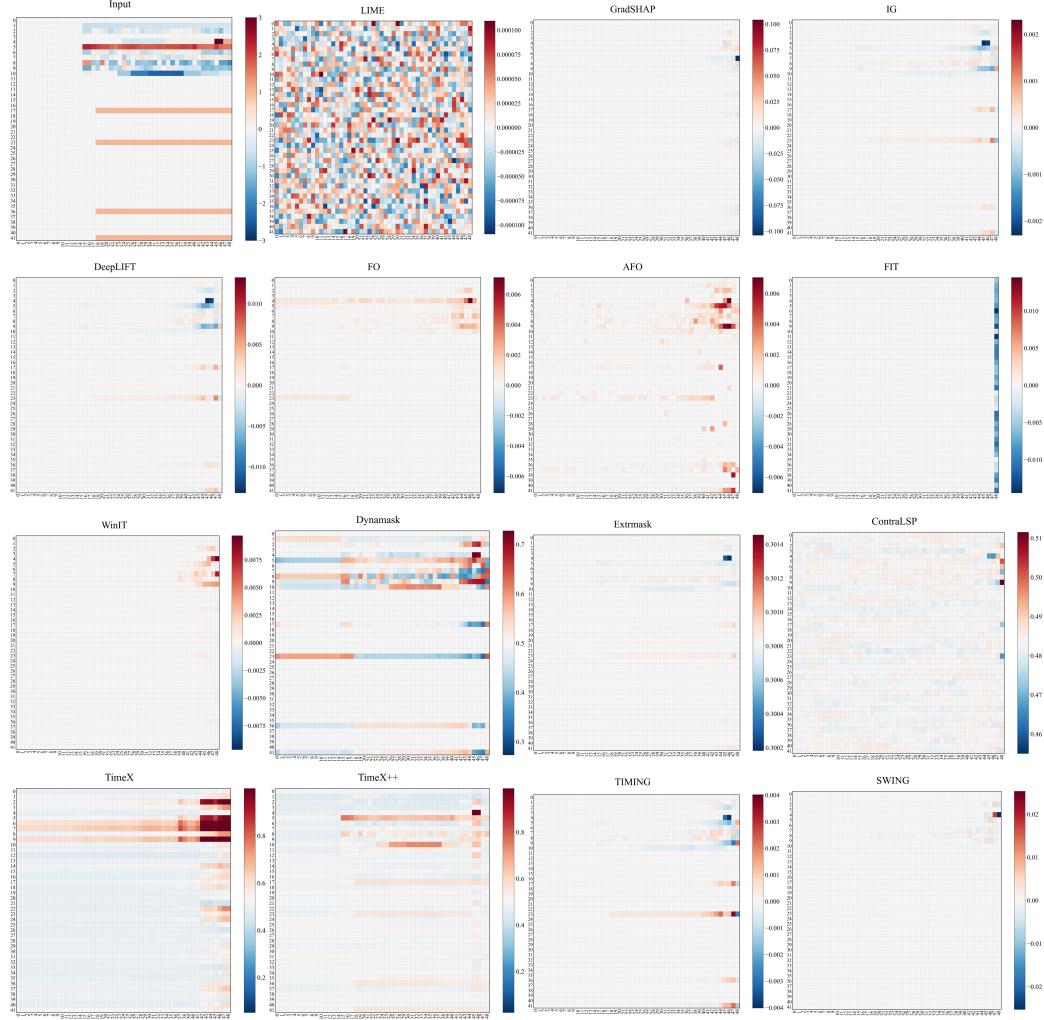
1352

1353	Index	Name	Index	Name	Index	Name	Index	Name
1354	0	Height	11	Weight	22	GCS-Motor4	33	GCS-Total9
1355	1	Hours	12	Blood pH	23	GCS-Motor5	34	GCS-Total10
1356	2	Diastolic BP	13	Cap. Refill	24	GCS-Total0	35	GCS-Total11
1357	3	FiO ₂	14	GCS-Eye0	25	GCS-Total1	36	GCS-Total12
1358	4	Glucose	15	GCS-Eye1	26	GCS-Total2	37	GCS-Verbal0
1359	5	Heart Rate	16	GCS-Eye2	27	GCS-Total3	38	GCS-Verbal1
1360	6	Mean BP	17	GCS-Eye3	28	GCS-Total4	39	GCS-Verbal2
1361	7	SpO ₂	18	GCS-Motor0	29	GCS-Total5	40	GCS-Verbal3
1362	8	Respiratory Rate	19	GCS-Motor1	30	GCS-Total6	41	GCS-Verbal4
1363	9	Systolic BP	20	GCS-Motor2	31	GCS-Total7		
1364	10	Body Temperature	21	GCS-Motor3	32	GCS-Total8		

1364

1365 Previous Score: 0.04 → Current Score: 0.04 (Score Diff: -0.01)

1366



1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

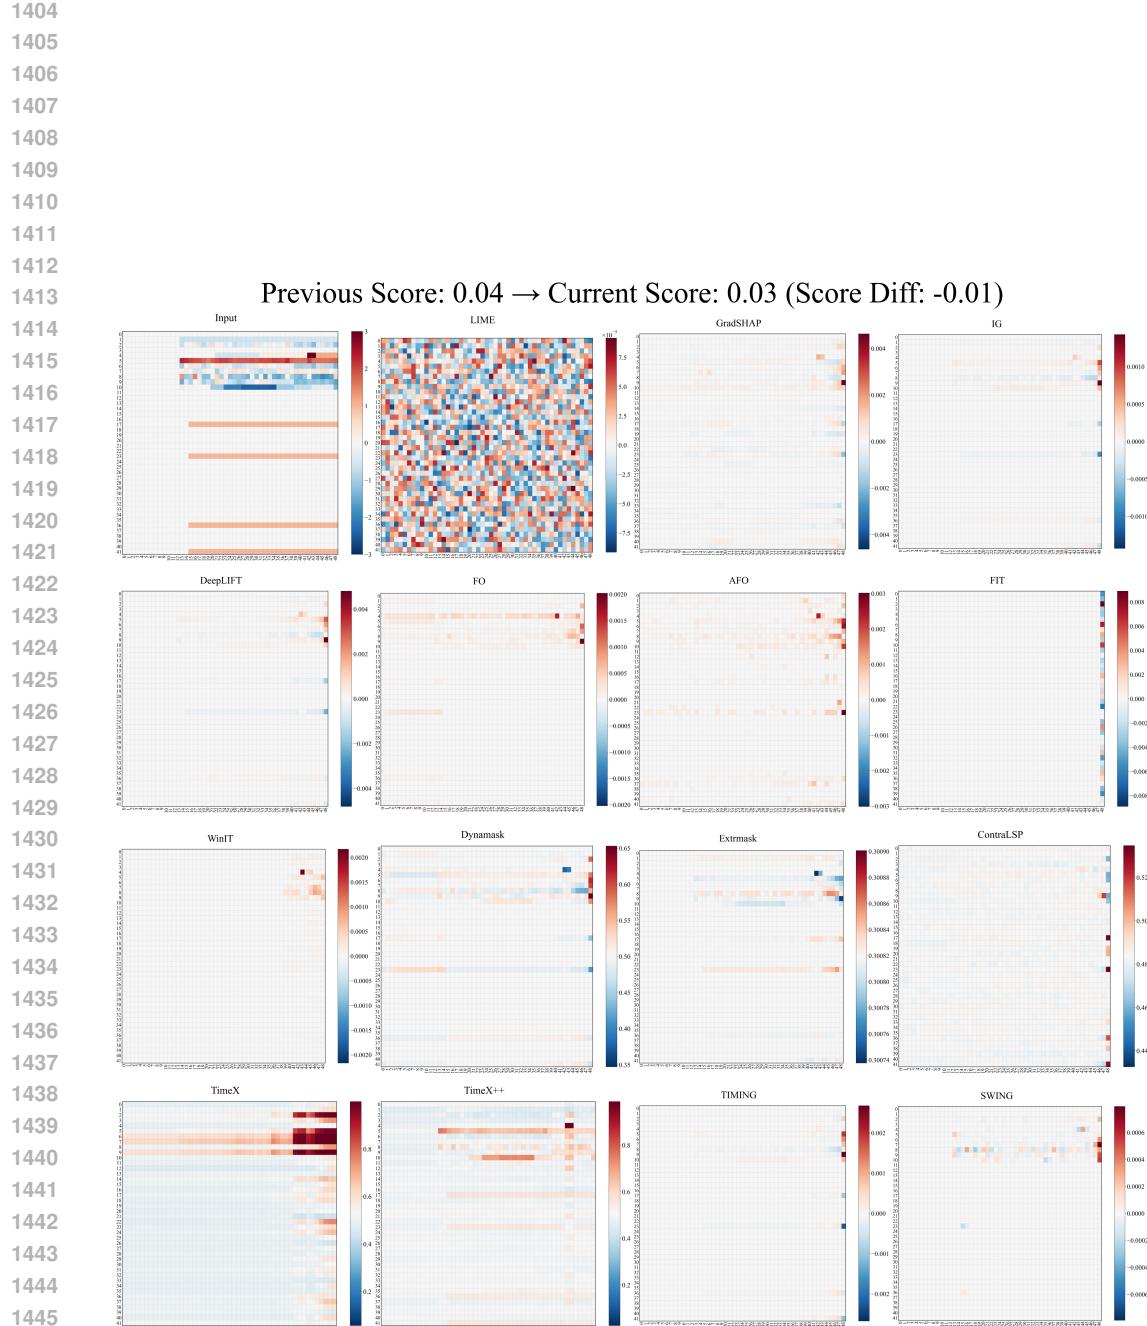
1400

1401

1402

1403

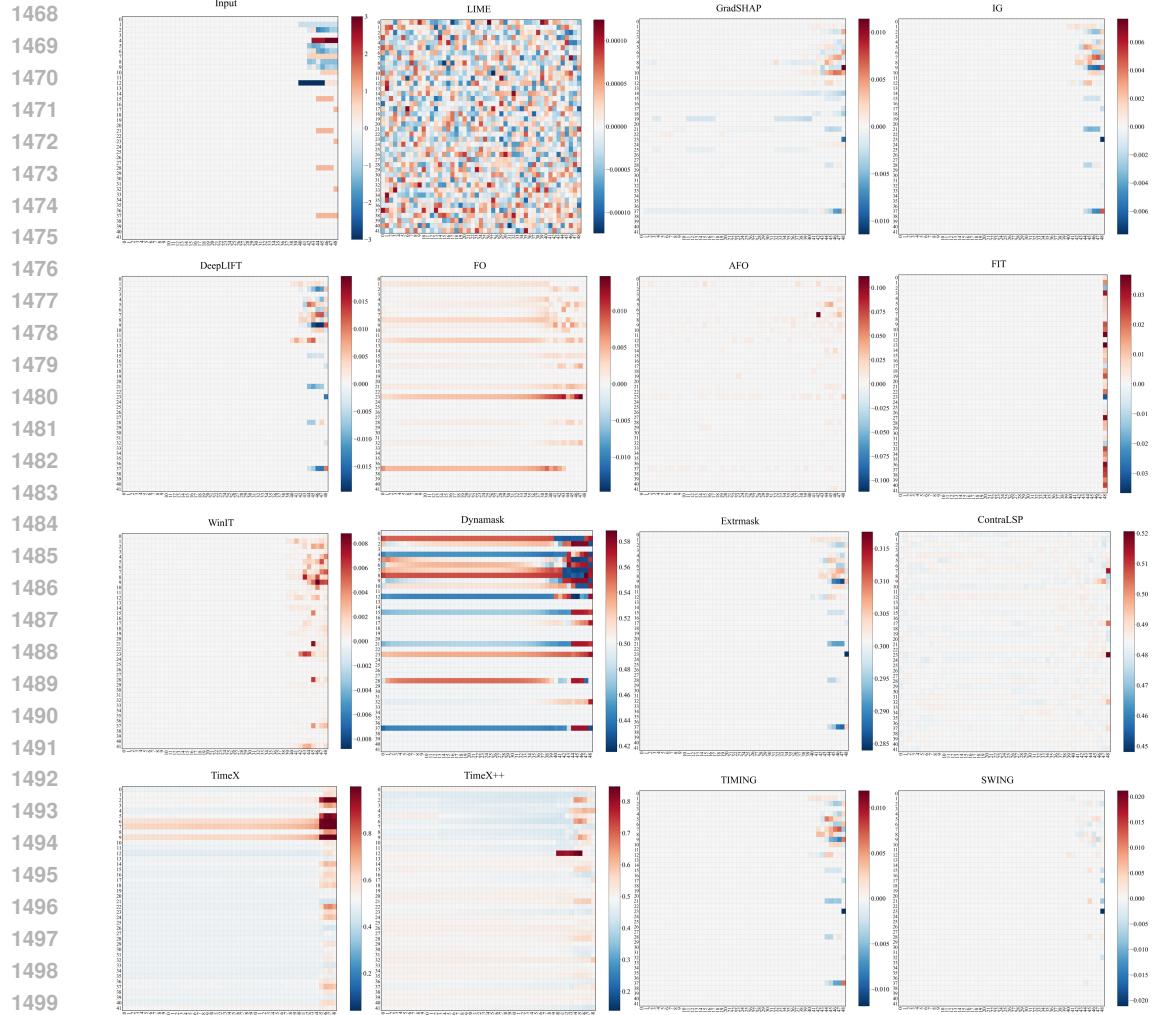
1400 **Figure 5:** Qualitative case study showing attributions extracted with XAI methods on MIMIC-
 1401 III (Johnson et al., 2016) using LSTM (Hochreiter & Schmidhuber, 1997) backbone with $T_1 = 47$
 1402 and $T_2 = 48$, i.e., $T_2 - T_1 = 1$. The uppermost-left heatmap displays the normalized input features,
 1403 while the remaining fifteen panels illustrate the attribution heatmaps generated by each XAI method
 1404 under the Delta-XAI framework, reflecting their respective explanations of the score changes.



1446 **Figure 6:** Qualitative case study showing attributions extracted with XAI methods on MIMIC-
1447 III (Johnson et al., 2016) using LSTM (Hochreiter & Schmidhuber, 1997) backbone with $T_1 = 47$
1448 and $T_2 = 48$, i.e., $T_2 - T_1 = 1$. The uppermost-left heatmap displays the normalized input features,
1449 while the remaining fifteen panels illustrate the attribution heatmaps generated by each XAI method
1450 under the Delta-XAI framework, reflecting their respective explanations of the score changes.
1451
1452
1453
1454
1455
1456
1457

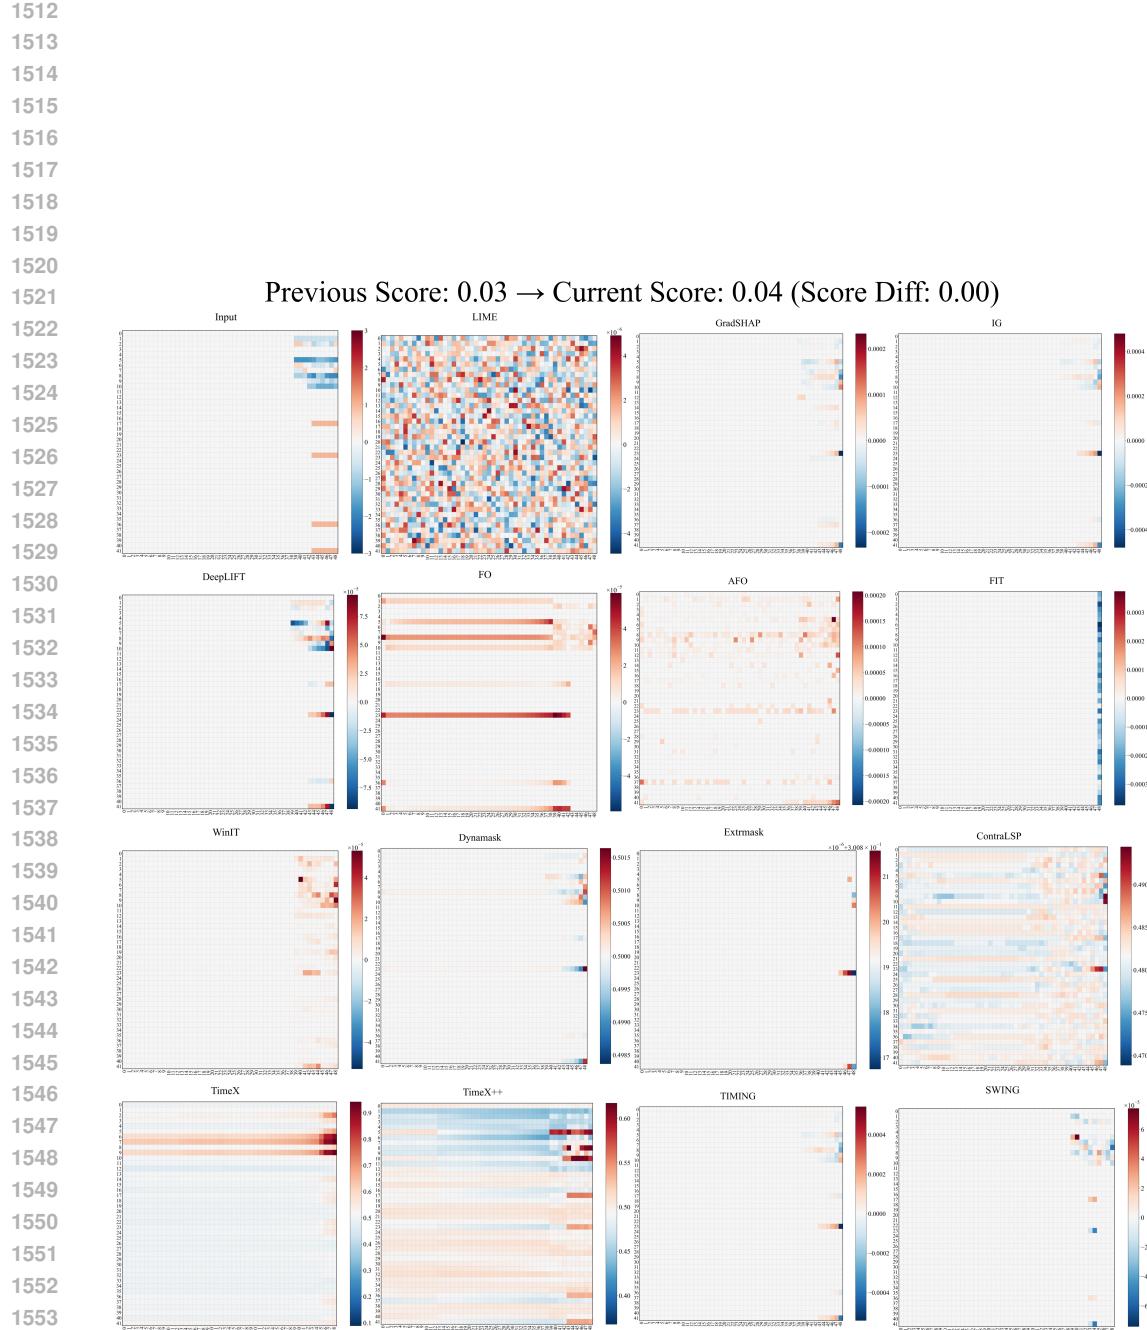
1458
1459
1460
1461
1462
1463
1464
1465
1466

1467 Previous Score: 0.03 → Current Score: 0.03 (Score Diff: 0.00)

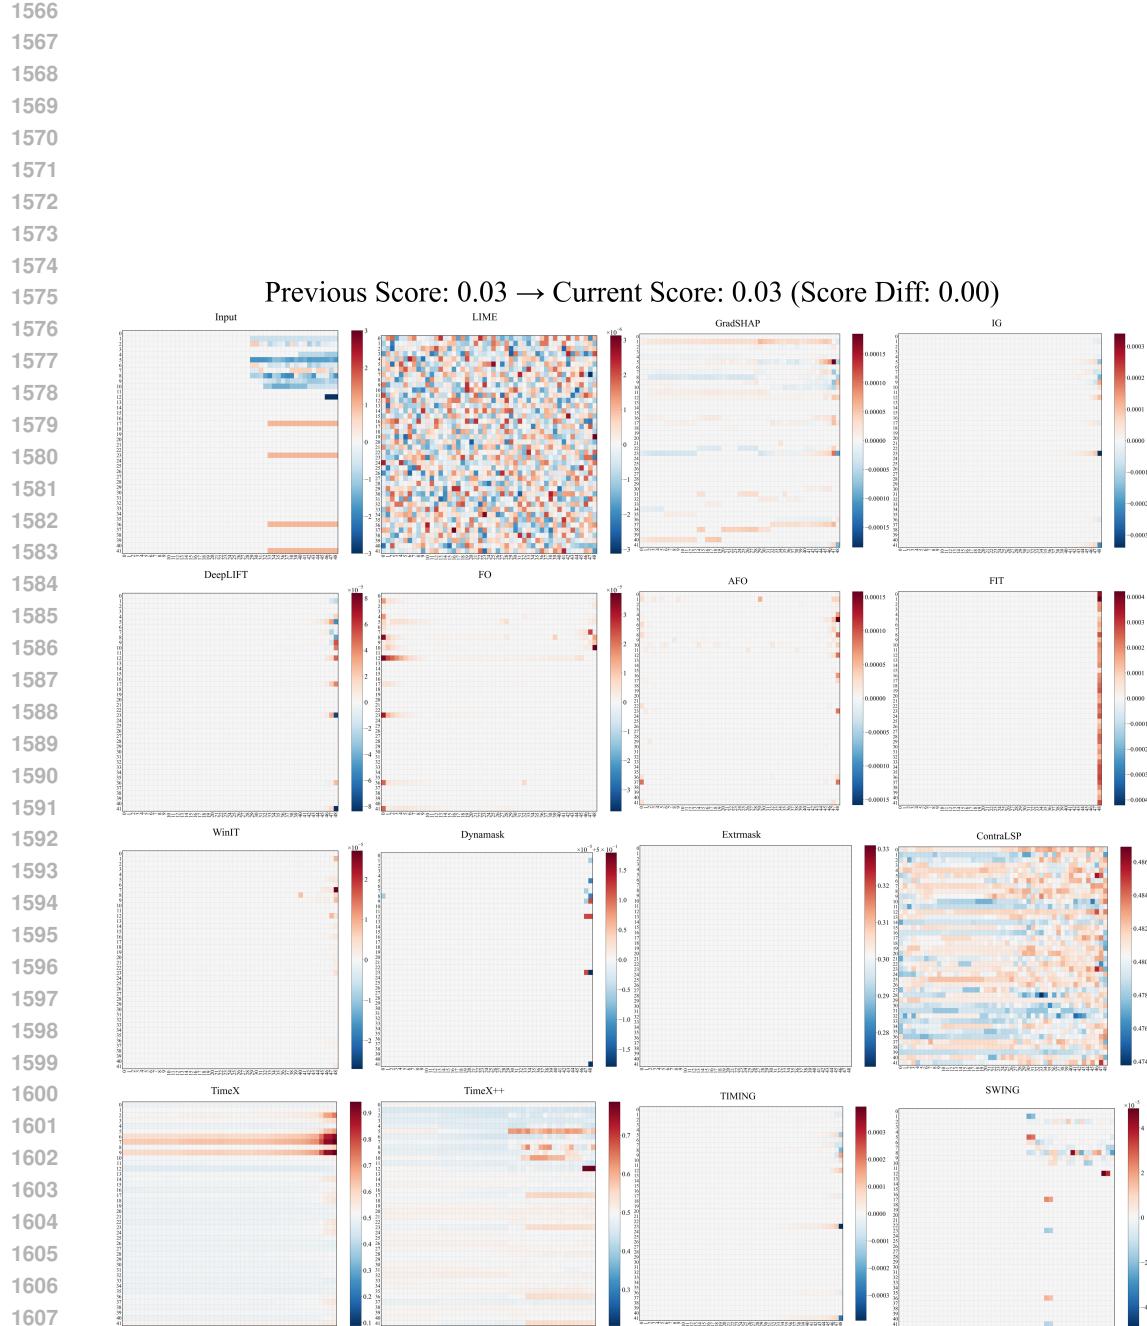


1500 **Figure 7:** Qualitative case study showing attributions extracted with XAI methods on MIMIC-
1501 III (Johnson et al., 2016) using LSTM (Hochreiter & Schmidhuber, 1997) backbone with $T_1 = 47$
1502 and $T_2 = 48$, i.e., $T_2 - T_1 = 1$. The uppermost-left heatmap displays the normalized input features,
1503 while the remaining fifteen panels illustrate the attribution heatmaps generated by each XAI method
1504 under the Delta-XAI framework, reflecting their respective explanations of the score changes.

1505
1506
1507
1508
1509
1510
1511



1554 **Figure 8:** Qualitative case study showing attributions extracted with XAI methods on MIMIC-
1555 III (Johnson et al., 2016) using LSTM (Hochreiter & Schmidhuber, 1997) backbone with $T_1 = 47$
1556 and $T_2 = 48$, i.e., $T_2 - T_1 = 1$. The uppermost-left heatmap displays the normalized input features,
1557 while the remaining fifteen panels illustrate the attribution heatmaps generated by each XAI method
1558 under the Delta-XAI framework, reflecting their respective explanations of the score changes.
1559
1560
1561
1562
1563
1564
1565



1608 **Figure 9:** Qualitative case study showing attributions extracted with XAI methods on MIMIC-
 1609 III (Johnson et al., 2016) using LSTM (Hochreiter & Schmidhuber, 1997) backbone with $T_1 = 47$
 1610 and $T_2 = 48$, i.e., $T_2 - T_1 = 1$. The uppermost-left heatmap displays the normalized input features,
 1611 while the remaining fifteen panels illustrate the attribution heatmaps generated by each XAI method
 1612 under the Delta-XAI framework, reflecting their respective explanations of the score changes.
 1613
 1614
 1615
 1616
 1617
 1618
 1619



Figure 10: Qualitative case study for coherence analysis on decompensation risk. Each subplot indicates two heatmaps for a MIMIC-III (Johnson et al., 2016) sample processed with a LSTM (Hochreiter & Schmidhuber, 1997) backbone. Bars for individual features represent the attribution contributed by newly observed inputs at the current time step, while the 'Delayed Effect of Intermediate Features' and 'Removal of Oldest Features' summarize the aggregated attributions of their respective feature groups. These panels provide an intuitive understanding of how each variable contributes to the observed change in decompensation risk score and illustrate the clinical plausibility of SWING's explanations.