

# Temporal Sepsis Modeling: a Fully Interpretable Relational Way

Vincent Lemaire  
Orange Research,  
Lannion, France  
vincent.lemaire@orange.com

Nédra Mellouli  
De Vinci Higher Education,  
De Vinci Research Center,  
LIASD-UP8, France  
Nédra.Mellouli@devinci.fr

Pierre Jaquet  
Recherche Clinique La Fontaine, France  
pierre.jaquet@ch-stdenis.fr

**Abstract**—Sepsis remains one of the most complex and heterogeneous syndromes in intensive care, characterized by diverse physiological trajectories and variable responses to treatment. While deep learning models perform well in the early prediction of sepsis, they often lack interpretability and ignore latent patient sub-phenotypes. In this work, we propose a machine learning framework by opening up a new avenue for addressing this issue: a relational approach. Temporal data from electronic medical records (EMRs) are viewed as multivariate patient logs and represented in a relational data schema. Then, a propositionalisation technique (based on classic aggregation/selection functions from the field of relational data) is applied to construct interpretable features to “flatten” the data. Finally, the flattened data is classified using a selective naive Bayesian classifier. Experimental validation demonstrates the relevance of the suggested approach as well as its extreme interpretability. The interpretation is fourfold: univariate, global, local, and counterfactual.

## I. INTRODUCTION

Sepsis [1] is defined as a life-threatening organ dysfunction caused by an aberrant dysregulated host response to infection, and progression to septic shock markedly increases the risk of multi-organ failure and death. Early identification and risk stratification using validated clinical criteria, such as increases in SOFA score, enable rapid implementation of evidence-based interventions—including prompt broad-spectrum antimicrobial therapy and targeted resuscitation—which are associated with improved outcomes. Despite these advances, sepsis remains a leading cause of morbidity and mortality in intensive care worldwide.

Despite decades of clinical and computational research, early detection remains a major challenge due to the disease’s heterogeneity and temporal complexity. Traditional machine learning models treat sepsis as a binary event, neglecting the underlying patient-specific physiological patterns that drive its onset and progression. However, early detection of sepsis using machine learning algorithms remains an area of research that is still in an ongoing stage of development. Despite recent advances, existing models exhibit limited performance, with some failing to identify up to 67% of sepsis cases [2], which severely restricts their deployment in real-world clinical settings. To ensure that these tools can be used reliably across diverse patient populations, further studies focusing on their clinical implementation and evaluation under real-world conditions are required [3]. Moreover, validating the

reproducibility and generalization capabilities of these models remains a major challenge [4]. This difficulty is attributable to the fact that most existing studies have mainly focused on populations from intensive care units (ICUs) [5], thereby limiting the applicability of these models to other hospital departments. In addition, several retrospective studies have demonstrated that the failure of machine learning models for sepsis detection depends heavily on the integration of clinically relevant and informative variables [3]. Nevertheless, the development and evaluation of predictive models leveraging data from hospitalized patients across all clinical services remain limited. In this context, the MIMIC database provides a broad panel of clinical, biological, and physiological parameters, making it particularly well suited for studying and comparing early sepsis prediction approaches. Accordingly, MIMIC-III [6] has become a widely used benchmark dataset in the literature for the comparative evaluation of sepsis prediction models and constitutes the experimental framework adopted in this study.

This paper is organized as follows. Section II reviews the main families of sepsis modeling. Section III details the relational representation of patient trajectories, the technique of propositionalisation by flattening, and variable selection, emphasizing the ability of this method to produce understandable characteristics. Section IV then presents the performance of the models, the quality of the aggregated variables, and the analysis of their univariate and local importance, while introducing the notion of counterfactual reasoning to explain and influence the model’s decisions. Finally, the last section summarizes the contributions of the proposed approach, highlighting its advantages in terms of explainability and flexibility.

## II. CONTEXT AND RELATED WORKS

Early sepsis prediction from Electronic Health Records (EHRs) has attracted considerable attention in recent years, mainly due to the availability of high-frequency physiological measurements and the clinical importance of early intervention. Existing approaches differ primarily in how they represent temporal information and how they address model interpretability.

### A. Time-series and deep learning approaches

Most recent works formulate early sepsis prediction as a multivariate time-series classification problem. Recurrent neural networks, particularly **Life Saving-Treatment Medicine** (LSTM) architectures, have been widely adopted to capture temporal dependencies in physiological signals [7]. Extensions based on Temporal Convolutional Networks (TCNs) and attention mechanisms further improve performance by modeling long-range dependencies [8]. While these deep learning approaches often achieve state-of-the-art predictive accuracy, they require complex architectures and large amounts of data, and their decision processes remain difficult to interpret in clinical settings.

### B. Feature-based and window-aggregation approaches

An alternative research approach involves transforming time-series data into fixed-length feature vectors by aggregating measurements over predefined temporal windows. This strategy enables the use of classical machine learning models such as logistic regression, random forests, or gradient boosting [9]. Recently, Thiboud et al. [10] proposed a gradient boosting model using aggregated clinical variables extracted from French hospital EHR data, showing competitive performance for early sepsis detection while preserving a degree of interpretability. However, most aggregation schemes are manually designed and do not explicitly control redundancy or model complexity.

### C. MDL-based and Bayesian explainable models

While relational feature extraction provides a systematic way to transform patient trajectories into a flat attribute-value table suitable for classical machine learning, the large number of potential aggregate features can easily lead to overfitting and redundant representations. To address these limitations, Minimum Description Length (MDL)-based methods offer a principled framework for both constructing and selecting informative features. By explicitly balancing model fit and complexity, MDL-based approaches generate compact, high-quality aggregates that capture temporal patterns while remaining suitable for interpretable probabilistic classifiers [11].

Explainability is further enhanced by selective Bayesian classifiers such as Fractional Naive Bayes [12], which aim to retain only informative and weakly dependent variables. In contrast to post-hoc explainability techniques applied to black-box models (e.g., SHAP [13]), these approaches provide intrinsic interpretability through probabilistic reasoning and explicit variable selection.

### D. Prediction horizons in early sepsis modeling

The prediction horizon, denoted  $h$ , defines the lead time before sepsis onset at which a model aims to anticipate the condition. In classical machine learning approaches based on aggregated features [9], [10],  $h$  is typically fixed by the aggregation window: only measurements up to  $h$  hours before

onset are included in feature computation. This constrains the models to rely on recent observations and may limit long-range predictive power.

Deep learning models such as LSTM, TCN, and MGP-RNN [7], [8] inherently handle sequences of arbitrary length and can be trained to predict multiple horizons simultaneously. For instance, models can be trained to output risk scores at  $h = 1, 3, 6$  hours before sepsis onset, allowing a flexible trade-off between early detection and predictive accuracy. Multi-horizon prediction is particularly advantageous in ICU settings, where early warnings must be balanced against false alarms.

For relational and MDL-based methods [11], [12], the prediction horizon is explicitly controlled through the number of temporal observations  $p$  included in the secondary table. Each patient trajectory is flattened over  $p$  time steps corresponding to the desired horizon  $h$ , enabling consistent learning of aggregate variables that encode temporal dynamics relevant for early prediction. This approach allows the use of interpretable probabilistic classifiers while maintaining horizon-specific information. Overall, while deep learning methods naturally support multi-horizon prediction, feature-based and relational methods achieve horizon flexibility by careful window selection and aggregation, highlighting a key design choice in early sepsis modeling.

To address the interpretability limitations of existing methods and lack of principled complexity control we represent patient trajectories as a root table linked to a secondary table of temporal observations. MDL-based aggregation, combined with supervised discretization and selective Bayesian classification, yields compact, interpretable features suitable for varying prediction horizons.

Our methodology details (i) relational data representation, (ii) feature extraction and flattening, and (iii) MDL-driven variable selection and Bayesian classification.

## III. MODELING: FROM TEMPORAL DATA TO FLAT DATA

### A. Dataset used

The MIMIC-III database [6] is a publicly available critical care dataset containing de-identified electronic health records of over 40,000 adult patients admitted to intensive care units at the Beth Israel Deaconess Medical Center between 2001 and 2012. Each patient trajectory comprises 36 key physiological parameters, including vital signs (heart rate, blood pressure, respiratory rate, oxygen saturation), laboratory measurements (e.g., white blood cell count, creatinine), and administered interventions relevant to sepsis prediction. As previous studies on this dataset [6], two variables are redrawn: ‘Hour’ and ‘Gender’ (the last one to prevent biased results). Measurements are primarily recorded hourly, although the sampling frequency varies across variables and patients. As is common in real-world electronic health records, missing values occur due to skipped measurements, delayed lab results, or clinical decisions not to collect certain tests. Appropriate preprocessing and imputation strategies are applied to ensure continuous trajectories suitable for temporal modeling and relational feature extraction. As it is frequent to have missing data in

the MIMIC-III database, we removed variables with more than 20% missing observations. Those with less than 20% missing data were imputed using the SMOTE method[14]. A 12-hour observation window is clinically justified, as early hemodynamic and inflammatory dysregulations preceding sepsis onset[15]. When keeping patients of this dataset who have 6 measurements before becoming “positive” plus 6 measurements while being positive (see Figure 1) then the extracted dataset contains 3940 patients that will be used below in the experiments.

The choice of prediction horizons  $h$  equal to 3 and 6 hours before sepsis onset is motivated by both clinical and technical considerations. Clean-ically, early intervention in sepsis is critical, as each hour of delayed treatment significantly increases mortality [16]. A 3-hour horizon provides a realistic window for timely clinical action, while a 6-hour horizon offers earlier warning to plan interventions and allocate resources. From a modeling perspective, MIMIC-III provides hourly measurements that allow the extraction of temporal patterns over 3- to 6-hour windows, which are sufficiently informative for sequential models such as LSTM, TCN, or MDL-based relational feature extraction combined with Bayesian classifiers.

### B. Viewing the data as relational data

The dataset described in the previous section may be viewed as a relational dataset (see Figure 1) with a star schema. In this paradigm, the root table contains two columns: the first one is the ‘Id’ of the patient (a reference to the patient), and the second one contains the class to predict (healthy or sick). In the secondary table, the number of columns corresponds to the  $d$  ( $d = 34$ ) explanatory variables described in the previous section plus the column ‘Id’ to enable table joins. The number of lines of the secondary table is  $pN$ , where  $p$  is the number of (logs or time stamps) of each patient used to do the prediction, which depends on the horizon for which one would like the prediction, as described just below.




Figure 1: Time series relational data flattening

This is represented in the Figure 2. In the initial dataset, each patient is represented by a measure (a vector) of 34 variables, each one measured 12 times ( $t \in \{t_{-6}, \dots, t_{+5}\}$ ) every hour,

so at different time stamps. In the dataset used all the patients become “sick” a  $t = t_0$ . Therefore to predict the nature of the patient one hour before  $t = t_0$  one may use a maximum of 6 measures ( $h = 1$  and  $p = 6$ ) and for a prediction 3 hours before  $t = t_0$  one may use a maximum of 4 measures ( $h = 3$  and  $p = 4$ ) etc. Indeed, the number of lines,  $p$ , to keep in the secondary table depends on the value of  $h$ .




Figure 2: Temporal relationship between the fixed observation window ( $p = 4, 6$  hours) and different prediction horizons  $h = 3, 1$ .

Note 1: using the same value for  $p$  and  $h$  (see Figure 1) for all patients does not in any way reduce the scope of the analysis. The proposed approach works in exactly the same way if  $p$  (with  $h$  set to the same value for fairness to patients) were to vary from patient to patient (if the latter had more than 4 measurements in the past before  $t_{-3}$ ).

Note 2: The proposed approach could also be multi-modal. Indeed, it would suffice to have one (or more) other table(s) containing alternative representations of patients and to apply the same process.

### C. Flattening the data

Machine learning models typically require input data in a tabular, flat format where each row represents an instance and each column represents a feature. Relational data stored across multiple tables needs to be combined into a single table to be used effectively for training classifiers. This action is also named propositionalization.

The propositionalization approaches consist of learning a model from relational data by flattening the original relational data, which are stored in several **unlinked** tables (similarly to databases) [17]. These approaches come from the field of Relational Data Mining [18] and are not usually applied to time series. More precisely, relational data contains at least one root table where each row represents a statistical individual (e.g. ‘id’ of the patients) and another secondary table containing detailed records (eg. the measurements re-alized of patients which are represented by each row of the secondary table). The specificity of relational data is to involve one-to-many relationships between the tables (eg. a patient underwent several identical medical examinations over time). The propositionalization problem consists of transforming the relational data into a single attribute-value dataset in order to use regular machine learning methods.

Two kinds of propositionalization approaches can be distinguished (i) *logic-based* methods, such as RSD [19], SINUS [20], tackle the propositionalization task by constructing first-order logic attributes; (ii) the *database-inspired* methods such as RELAGGS [21] apply aggregation functions, such as Min, Max, and Mean in order to generate attributes. The interested reader can find a complete state-of-the-art and a comparative evaluation in [22].

In our study, we used a Minimal Description Length (MDL) based propositionalization approach presented in [11]. This approach exploits a Bayesian formalism to generate informative aggregate variables in a supervised way. To the best of our knowledge, this approach is the only one that avoids overfitting problems by regularizing the complexity of the generated variables.

Our relational data (see Fig. 1) consists of the root table, which contains  $N$  instances, characterized by two variables: the patient identifiers and the target variable (i.e. class values). The secondary table contains  $pN$  detailed medical examination, described by 36 variables already described in Section III-A. The number of aggregate variables to be constructed ( $Q$ ) is the only user parameter. Our relational data is transformed into a regular attribute-value dataset by applying the MDL based propositionalization approach.

#### D. Variable selection and learned classifier

The used MDL approach is able to select the most informative aggregate variables in two ways: i) by filtering uninformative aggregate variables; ii) by finding the most informative and independent subset of variables.

1) *Filtering uninformative variables:* The filtering of the aggregate variables is based on previously developed supervised discretization ([23]) and grouping ([24]) methods. This Bayesian approach turns the learning task into a model selection problem. A prior distribution is defined on the model space that exploits the hierarchy of their parameters. In practice, this approach reaches a good trade-off between robust and accurate models. The prior favors simple models with few intervals, and the likelihood favors models that fit the data regardless of their complexity. The aggregate variables are evaluated, one by one, using a specifically designed MDL optimization criterion ([11]). The complexity of the aggregate variables is taken into account by adding a construction cost in the prior. This criterion can be interpreted as a coding length according to information theory. Compression Gain (CG) compares the coding length of the learned model with the empty model that includes a single interval. CG measures the ability of the learned models to compress the training data, despite the additional construction cost. Only the  $R$  variables with a positive CG are retained ( $R \leq Q$ )

2) *Finding the most informative subset of variables:* All the  $R$  informative variables coming from the previous step (after discretization or grouping values) are gathered together and used to learn a Fractional Naive Bayes classifier

(a Naive Bayes which uses a subset of the  $Q$  variables defined by a selective process). The Fractional Naive Bayes (FNB) aims to select the most informative and independent subset of variables by using a specifically designed MDL optimization criterion ([12]). It selects the most informative and independent subset of variables using a soft selection scheme with variable weights ranging between 0 and 1. The best model is obtained by optimizing a sparse regularization of the model's log-likelihood. The optimization algorithm employed consists of a sequence of forward and backward selection steps, which add or remove variable weights, starting with weight increments of 1. These two selection steps are repeated with decreasing weight increments, each time beginning with a random ordering of the variables. In the end, we retain the most probable subset of weighted variables that comply with the naive Bayes assumption, i.e., that are both informative and independent. At the end, we keep the most probable subset of variables compliant with the naive Bayes assumption, i.e. both informative and independent. This subset contains  $S$  aggregate variables, with  $S \leq R \leq Q$ .

3) *Learning classifier:* Finally, the used classifier is a naive Bayes classifier, which takes the  $S$  selected aggregate variables as an input. As shown in Eq. 1, the naive Bayes classifier ([25]) estimates the distribution of a particular class value  $C_z$  conditionally to the input variables  $x_k$ .

$$P(C_z|x_k) = \frac{P(C_z) \prod_{j=1}^d P(V_j = x_{jk}|C_z)^{W_j}}{\sum_{t=1}^C [P(C_t) \prod_{j=1}^d P(V_j = x_{jk}|C_t)^{W_j}]} \quad (1)$$

$C$  is the number of class values to be predicted (in this paper  $C = 2$ ),  $V$  the input variables and  $W$  the weights on the variable coming from the process described in [12]. This simple and efficient classifier makes the assumption that the distributions  $P(V_j = x_{jk}|C_z)$  are independent. In practice, these conditional distributions are estimated in a frequentist way, by using the previously learned univariate discretization ([23]) and grouping ([24]) models. The denominator of Eq. 1 normalizes the estimated probability by making a sum of the numerator term over all the class values. At the end, the predicted class value given a particular  $x_k$  is the one that maximizes the conditional probabilities  $P(C_z|x_k)$ .

Note: Other classifiers than the SNB here used could be employed with the  $S$ ,  $R$  or  $Q$  variables found with the process described above. In this paper for the sake of the complete line with a MDL principle (and also since the results are quite good) we did not test this possibility but interested readers may try. All the results are fully reproducible with the Github where the dataset and the used code are available.

## IV. RESULTS

### A. Classifier performances

The average results of the classifier on a 10-fold cross-validation are presented in Table I in terms of both accuracy and AUC, compared to the value of the number of  $Q$  variables to be constructed (as described in the previous section).

Q	R	S	Acc Train	Auc train	Acc Test	Auc Test
10	5	5	0,7941	0,8767	0,7719	0,8526
100	86	44	0,9379	0,9832	0,8963	0,9664
1000	872	98	0,9620	0,9918	0,9009	0,9708
10000	8849	198	0,9713	0,9969	0,9171	0,9751
100000	72331	272	0,9883	0,9992	0,9194	0,9737

Table I: Caption

Name	Level	Importance
Min(EtCO2) where Calcium > 8.075	0.1999	0.02909
Min(HospAdmTime) where BaseExcess > -1.61	0.0976	0.01741
Min(Temp) where Hct ≤ 29.505	0.1275	0.01671
StdDev(EtCO2) where SaO2 ≤ 95.79	0.1490	0.01601
Max(O2Sat) where HR ≤ 86.95	0.1412	0.01572
StdDev(O2Sat)	0.07872	0.01492
Max(EtCO2)	0.1650	0.01467
StdDev(FiO2) where DBP > 59.95	0.1023	0.01458
Max(FiO2) where O2Sat > 99.05	0.3774	0.01451
Min(Temp) where Age > 65.495	0.1214	0.01428
StdDev(BUN) where SaO2 > 95.79	0.01601	0.01331
StdDev(FiO2) where Fibrinogen ≤ 229.15	0.1424	0.01252
Min(pH) where Lactate > 2.496	0.06828	0.01216
StdDev(Alkalinephos) where Glucose > 131.95	0.02323	0.01199
Mean(MAP) where SaO2 > 95.79	0.02699	0.01155
Min(Potassium) where Chloride ≤ 108.55	0.02291	0.01134
Sum(Calcium) where BUN > 19.3	0.04739	0.01127
StdDev(FiO2) where Calcium > 8.075	0.1552	0.01103
Sum(Temp) where AST ≤ 167.5	0.1131	0.01093
Min(pH) where Potassium ≤ 4.162	0.09446	0.01087

Table II: The  $S$  20 most important aggregated variables

Obviously the performances are very good. If a tradeoff between the number of variables incorporated in the model ( $R$ ) and the AUC in test is required, then the value  $Q=1000$ ,  $R=872$  and  $R=98$  could be considered. The 98 aggregate variables to be constructed are sufficient to obtain good performance and a frugal model. But it is up to the final user since the largest model has only 272 input variables. The next section will present some of these aggregate variables.

### B. Example of aggregates elaborated

The Table II, column 1, gives the  $S$  20 most important aggregated variables found by the approach used [11] here. A nice property is their straightforward meaning. On this dataset they have only two operands, even if the method in [11] may produce more complex “rules”. Each of them is easy to read and to understand even for a non-data scientist (or research scientist) and moreover in terms readable by medical practitioners. We also give for the first one a color code where the blue parts are the rules and the orange parts the variables present in the secondary table (to ease the reading). It is worth noting that the clinical interpretation of some of these rules still needs to be validated using longitudinal data collected between 2022 and 2025 at XXXX<sup>1</sup> Hospital.

<sup>1</sup>The name of the hospital is not here given for anonymity reason during submission of this paper.

### C. Univariate variable importance

As described in Section III-D the Khiops classifier is trained in two steps. The first one is to evaluate the univariate predictive information contained in each aggregated variable. For that, a supervised discretization [23] is realized for numerical one and a supervised grouping method [24] is performed on categorical variables. In the dataset used here all the variables are numerical (except the ‘Gender’ and ‘Hour’ variables redrawn at the beginning of the study). One output of this preprocessing is an information of the univariate predictive information, named “level”, of each variable which is in the range  $[0 : 1]$  (0 variable no informative, 1 variable 100 % informative). These values are given in the second column of the Table I.

Since each aggregated variable is discretized, it is also possible to have a look at the information carried by the intervals found. For that the Khiops library provides an interactive results visualization tool, called Khiops Visualization. We present in Figure 3 a copy of a screenshot of this tool for the most important aggregate:  $V_1$  = “Min(EtCO2) where Calcium > 8.075”.

The analysis of the discretization allows us to understand why this aggregate is informative. Small values of  $V_1$  produce intervals relatively pure of Sepsis=1 while the intervals between 26 and 36 are relatively pure of Sepsis=0. The reading of  $P(X|C)$  per interval, and for each variable, is informative to understand the result of the univariate classifications.

### D. Local variable importance

To analyze local behavior (for each example), it is possible to compute the Shapley values of all the input variables of the trained classifier using the method described in [26] and available in the Khiops library V11. This allows to have the variable importance example by example (local importance). We illustrate this information (XAI) in the Table III.

Five examples among the ones with high probability of being positive (Sepsis) (predicted by the classifier) are presented. The first column gives the probability of being positive (Sepsis). Then, for each patient, the two variables that contribute the most to the predicted probability (the number of variables is just defined per user when asking this XAI outcome in the library (see the code in the GitHub provided for reproducibility)). Therefore, here, after the first column, there are 2 triplets of columns. Each triplet gives for each patient the name of the variable, then the value of the variable and finally the Shapley value for this variable. The triplets (so the columns of the file) are sorting according to the Shapley value, allowing a fast understanding of the individual variable importance.

When examining, for example, the second patient (line 1) in this table, we see that the most important variable is “Min(EtCO2) where Calcium > 8,075 ” and the second one is “Min(DBP) where DBP > 59,95 ”. The value of the most important belongs to the value interval  $[30,95,31,15]$  while the value of the second most important value belongs to the value interval  $[66,55,66,95]$  . The associated Shapley values

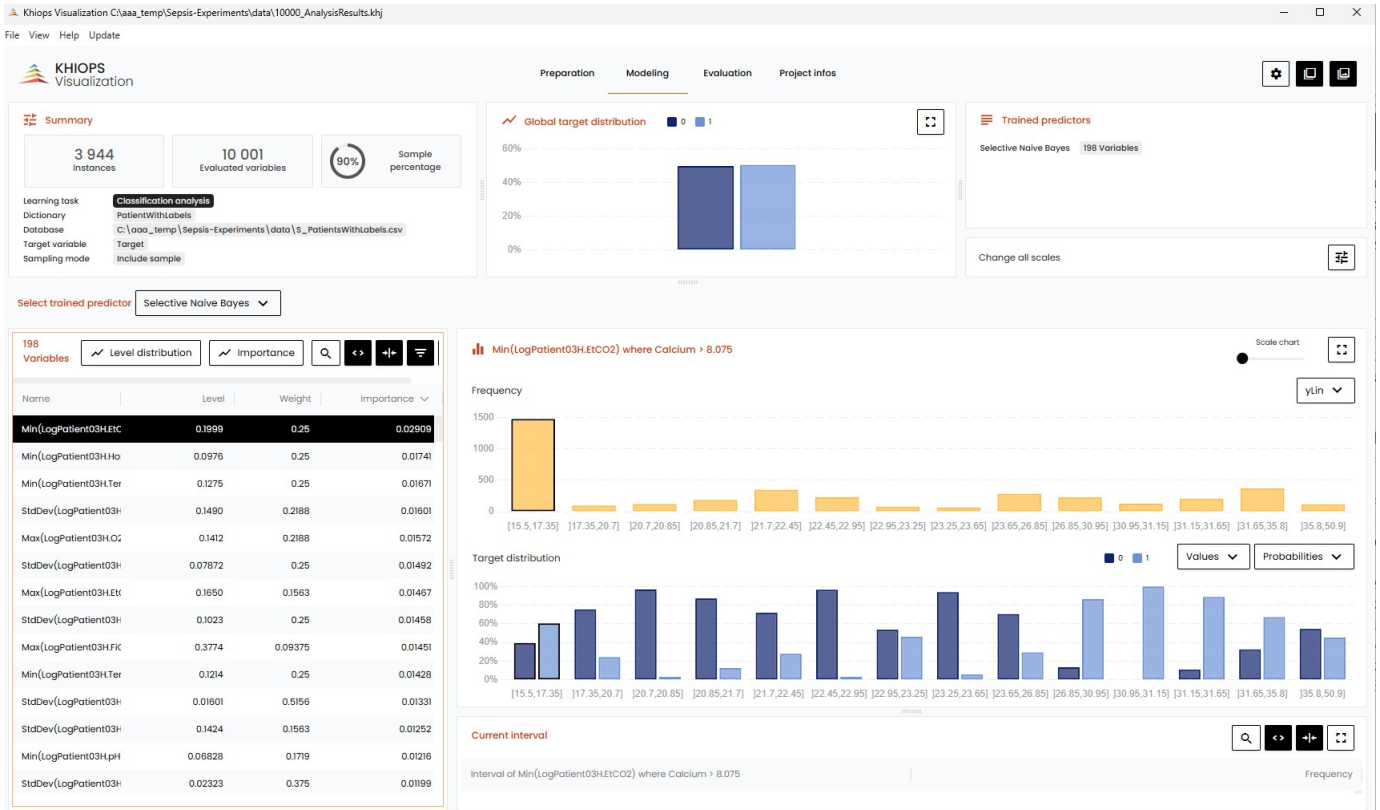


Figure 3: Screenshot of KHIOPS Visualization after analyzing the sepsis database and constructing 10000 aggregates

ProbTarget1	ShapleyVariable_1	ShapleyPart_1	ShapleyValue_1	ShapleyVariable_2	ShapleyPart_2	ShapleyValue_2
0,999936	Min(EtCO2) where Calcium > 8,075	[30,95,31,15]	3,2213	Min(DBP) where DBP > 59,95	[66,55,66,95]	1,9981
0,999936	Max(EtCO2)	[31,35,31,55]	0,5986	Max(EtCO2) where MAP > 77,16	[31,35,31,55]	0,5727
0,999929	Median(Bilirubin_direct) where HospAdmTime > -7,64	[4,27,4,33]	1,4618	Min(HCO3) where Glucose > 131,95	[22,015,22,195]	1,2114
0,999929	Min(EtCO2) where Calcium > 8,075	[30,95,31,15]	3,221	Min(Temp) where Age > 65,495	[37,283,+inf]	0,6948
0,999929	Min(EtCO2) where Calcium > 8,075	[30,95,31,15]	3,221	Max(O2Sat) where HR ≤ 86,95	]-inf,98,1]	0,4083

Table III: Illustration of the knowledge provided by the local importance

are in columns 4 and 7. For this patient, the main causes of a high probability of being positive are therefore very easy to understand. The other lines of this table appear to be equally straightforward to read.

Note: the KHIOPS library can also output a file with all the Shapley values for all variables and for all the classes, allowing the use of this file with a Python library like Shap [27] to create personalized visualization.

### E. Global variable importance

The global importance of the variables is given in column 3 of the Table II. These importance are defined for each variable as the average on all the train samples of the absolute values of the local Shapley values described in the previous section.

### F. Counterfactuals

Counterfactual reasoning is a concept from psychology and social sciences [28], which involves examining possible alternatives to past events [29]. Humans often use counterfactual

reasoning by imagining what would happen if an event had not occurred, and this is precisely what counterfactual reasoning is. When applied to artificial intelligence, the question is, for example, “Why did the model make this decision instead of another?” (counterfactual explanation) or “How would the decision have differed if a certain condition had been changed?”. This reasoning can take the form of a counterfactual or semi-factual explanation.

A counterfactual explanation might be “If your income had been \$10000 higher, then your credit would have been accepted” [30], [31], [32]. A semi-factual is a special-case of the counterfactual in that it conveys possibilities that “counteract” what actually happened, even if the outcome does not change [33]: “Even if your income had been \$5000 higher your credit would still be denied” (but closer to being accepted).

Within the framework of counterfactual reasoning, we here used the method described in [34] and the notion of trajectory. Indeed, the study of counterfactual trajectories a posteriori is of great interest, as it also makes it possible to identify

when a trajectory is approaching the frontier or pass through the frontier.

In this paper, the trajectory of a counterfactual is the minimum of changes in the input vector to change the predicted class to the opposite<sup>2</sup>.

In the case of Sepsis, one could be interested, for example, when examining Patients :

- predicted as negative (Target=0) but close to the frontier; what are their counterfactual (Target=1) ? The reasons that will cause them to be paid in the opposite case?
- predicted as positive (Target=1) but for which one could be interested in finding reasons to transform them predicted as negative ?

We give below one example of these two possibilities when one considers all the input variables (of the input vector) as “alterable”. But as explained in the literature and in [34] it is up to the user to limit the list of the alterable variables since some of them could not be changed. This must be done in consultation with the expert of the application domain, here a patrician.

In the first case :

- for the patient ‘10001’ (Patient\_Id) : Initial class = ‘0’ - Proba ‘0’ = 0.561962
- the trajectory for the counterfactual as a single stage :
  - Step 1 : if “StdDev(EtCO2)” change of value from 5.795418449 to 0.2537

After the first stage, the initial predicted probability (0.561962) to be negative becomes 0.131336. Therefore, he becomes positive.

In the second case

- for the patient ‘100003’ (Patient\_Id) : Initial class = ‘1’ - Probability ‘1’ = 0.999787
- the trajectory for the counterfactual as 6 stages :
  - Step 1 : ‘Min(Temp)’ from 37.6 to 31.3135 → Proba ‘1’ = 0.999021,
  - Step 2 : ‘Max(HospAdmTime)’ from -253.56 to 8.022499 → Proba ‘1’ = 0.994351,
  - Step 3 : ‘Min(pH)’ from 7.398 to 7.269 → Proba ‘1’ = 0.972920,
  - Step 4 : ‘StdDev(EtCO2)’ from 2.965952629 to ‘8.17991062’ → Proba ‘1’ = 0.912302,
  - Step 5 : ‘Median(FiO2)’ from 0.4 to 1.0095 → Proba ‘1’ = 0.751472,
  - Step 6 : ‘Median(HospAdmTime)’ from 253.56 to 8.022499 → Proba ‘1’ = 0.471857,

After the 6<sup>th</sup> stage, the initial predicted probability (0.561962) to be positive becomes 0.471857. Therefore, the patient becomes negative.

The information provided by the trajectory for the counterfactual is clearly beneficial to understand and to act on the positivity of the Sepsis.

## V. NOTE ON STATE-OF-THE-ART IN MACHINE LEARNING FOR EARLY SEPSIS PREDICTION

Recent advances in machine learning have led to several commercial systems for early prediction of sepsis, evaluated in heterogeneous clinical settings. Among the most representative are *InSight*<sup>®</sup> [35], *Sepsis ImmunoScore*<sup>™</sup> [36], *NAVIOY*<sup>®</sup> Sepsis [37] and *VFusion*<sup>™</sup> Sepsis [38] report high performance, with AUC values generally between 0.83 and 0.91. However, as highlighted by the PREVIA [10] retrospective study and in line with existing meta-analyses [3], direct comparison of these systems remains challenging due to substantial differences in patient populations, prediction horizons, variables used and validation protocols.

### NOTE ON REPRODUCIBILITY AND VALIDITY

The initial dataset (MIMIC-III) which contains 40000 patients as well as the 3940 patients used in the experiments are available on the GitHub (<https://anonymous.4open.science/r/sepsis-project-5DBD/README.md>). The code to train the classifiers (Table I), output the local importance (Table III), the counterfactuals, etc, is also available in the same Github. In this way, all the results are fully reproducible. Moreover one author of this paper is a doctor specialized on Sepsis who validated the found results.

## VI. CONCLUSION

In this study, we presented an innovative approach for the early detection of sepsis by exploiting electronic medical data in relational form. By combining the relational representation of patient trajectories, the propositionalization technique based on the MDL method, and a selective Bayesian classifier, we succeeded in obtaining a model that is both effective and intrinsically interpretable. This approach not only improves prediction accuracy but also provides a clear understanding of the key factors contributing to sepsis detection, thereby facilitating acceptance by clinicians.

The experimental results highlight the relevance of this approach, particularly its ability to produce understandable features and adapt to different prediction horizons. In addition, the flexibility of relational representation opens up prospects for the integration of multimodal data and the management of patient trajectories of varying sizes.

The proposed approach is based on a three-dimensional interpretation: univariate, global, and local. The first step, univariate, consists of analyzing the importance of each individual variable by evaluating their predictive contribution using measures such as univariate information. The second step, global, provides an overview of the most influential variables across all patients. Finally, local interpretation focuses on each patient individually, using methods such as Shapley value analysis to identify the specific variables that contributed most to the model’s decision for that particular case. This

<sup>2</sup>A notebook that performs this process for the Khiops library is given in the Github for reproducibility

three-step approach provides a comprehensive and hierarchical understanding of the model's mechanisms, facilitating its acceptance by clinicians and its implementation in a medical context.

In the context of sepsis, counterfactuals offer the possibility of identifying the changes needed in clinical variables to shift a prediction from positive to negative (or vice versa). For example, by adjusting certain physiological parameters of a patient, it becomes possible to understand which factors most influence the model's decision. This approach promotes a better understanding of the underlying mechanisms and can guide medical interventions by proposing alternative trajectories to improve care.

This study, which incorporates counterfactual reasoning and local explanation techniques, helps to increase transparency and confidence in predictive models in clinical settings. Continuing this work could contribute to more personalized, accurate, and interpretable medicine, particularly in the fight against complex syndromes such as sepsis.

## REFERENCES

- [1] J. M. O'Brien, N. A. Ali, S. K. Aberegg, and E. Abraham, "Sepsis," *The American Journal of Medicine*, vol. 120, no. 12, pp. 1012–1022, 2007.
- [2] A. Wong, E. Otlés, J. P. Donnelly, A. Krumm, J. McCullough, O. DeTroyer-Cooley, J. Pestrué, M. Phillips, J. Konye, C. Penzoza *et al.*, "External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients," *JAMA Internal Medicine*, vol. 181, no. 8, pp. 1065–1070, 2021.
- [3] L. M. Fleuren, T. L. T. Klausch, C. L. Zwager, L. J. Schoonmade, T. Guo, L. F. Roggeveen, E. L. Swart, A. R. J. Girbes, P. Thorat, A. Ercole *et al.*, "Machine learning for the prediction of sepsis: A systematic review and meta-analysis of diagnostic test accuracy," *Intensive Care Medicine*, vol. 46, no. 3, pp. 383–400, 2020.
- [4] C. L. Ramspek, K. J. Jager, F. W. Dekker, C. Zoccali, and M. van Diepen, "External validation of prognostic models: What, why, how, when and where?" *Clinical Kidney Journal*, vol. 14, no. 1, pp. 49–58, 2021.
- [5] D. W. Shimabukuro, C. W. Barton, M. D. Feldman, S. J. Mataraso, and R. Das, "Effect of a machine learning-based severe sepsis prediction algorithm on patient survival and hospital length of stay: A randomised clinical trial," *BMJ Open Respiratory Research*, vol. 4, no. 1, p. e000234, 2017.
- [6] A. E. W. Johnson, T. J. Pollard, L. Shen, L.-W. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "Mimic-iii, a freely accessible critical care database," *Scientific data*, vol. 3, no. 1, pp. 1–9, 2016.
- [7] J. Futoma, S. Hariharan, M. Sendak *et al.*, "An improved multi-output gaussian process rnn with real-time validation for early sepsis detection," *arXiv preprint arXiv:1708.05894*, 2017.
- [8] M. Moor, M. Horn, B. Rieck, D. Roqueiro, and K. Borgwardt, "Early recognition of sepsis with gaussian process temporal convolutional networks and dynamic time warping," in *SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020.
- [9] T. Desautels, J. Calvert, J. Hoffman, M. Jay, Y. Kerem, L.-W. Shieh, D. Shimabukuro, U. Chettipally, M. Feldman, and C. Barton, "Prediction of sepsis in the intensive care unit with minimal electronic health record data," *Critical Care Medicine*, vol. 44, no. 2, pp. e94–e101, 2016.
- [10] M. Thiboud, R. Lelong, I. Colombet, T. Durand, and G. J. Escobar, "Development and validation of a machine learning–based early sepsis prediction model using french hospital data," *Scientific Reports*, 2025.
- [11] M. Boullé, C. Charnay, and N. Lachiche, "A scalable robust and automatic propositionalization approach for bayesian classification of large mixed numerical and categorical data," *Mach. Learn.*, 2018.
- [12] C. Hue and M. Boullé, "Fractional naive bayes (fnb): non-convex optimization for a parsimonious weighted selective naive bayes classifier," 2024. [Online]. Available: <https://arxiv.org/abs/2409.11100>
- [13] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [14] N. Bignoumba, S. B. Yahia, and N. Mellouli, "Deep padding and alignment strategies for irregular multivariate clinical time series," in *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 28th International Conference KES-2024, Seville, Spain, 11-13 September 2023*, ser. Procedia Computer Science, C. Toro, S. A. Ríos, R. J. Howlett, and L. C. Jain, Eds., vol. 246. Elsevier, 2024, pp. 3275–3284. [Online]. Available: <https://doi.org/10.1016/j.procs.2024.09.312>
- [15] M. Moor, B. Rieck, M. Horn, C. R. Jutzeler, and K. Borgwardt, "Early prediction of sepsis in the icu using machine learning: A systematic review," *Frontiers in Medicine*, vol. 8, p. 607952, May 2021.
- [16] K. E. Rudd, S. C. Johnson, K. M. Agesa, K. A. Shackelford, D. Tsoi, D. R. Kievlán, D. V. Colombara, K. S. Ikuta, N. Kissoon, S. Finfer *et al.*, "Global, regional, and national sepsis incidence and mortality, 1990–2017: analysis for the global burden of disease study," *The Lancet*, vol. 395, no. 10219, pp. 200–211, 2020.
- [17] N. Lachiche, *Propositionalization*. Boston, MA: Springer US, 2010, pp. 812–817.
- [18] S. Džeroski, "Relational data mining," in *Data Min. Knowl. Disc. Handbook*. Springer, 2009, pp. 887–911.
- [19] N. Lavrač, F. Železný, and P. A. Flach, "RSD: Relational subgroup discovery through first-order feature construction," in *Inductive Logic Programming*. Springer, 2002, pp. 149–165.
- [20] S. Dzeroski and N. Lavrac, *Inductive logic programming: techniques and applications*, E. Horwood, Ed. New York: Prentice Hall, 1994.
- [21] M.-A. Krogel and S. Wrobel, "Transformation-based learning using multirelational aggregation," in *Inductive Logic Programming*. Springer, 2001, pp. 142–155.
- [22] M.-A. Krogel, S. Rawles, F. Železný, P. A. Flach, N. Lavrač, and S. Wrobel, "Comparative evaluation of approaches to propositionalization," in *Inductive Logic Programming*. Springer, 2003, pp. 197–214.
- [23] M. Boullé, "MODL: a Bayes optimal discretization method for continuous attributes," *Mach. Learn.*, vol. 65, no. 1, pp. 131–165, 2006.
- [24] —, "A grouping method for categorical attributes having very large number of values," in *Proceedings of the Fourth International Conference on Machine Learning and Data Mining in Pattern Recognition*, P. Perner and A. Imiya, Eds., vol. 3587, 2005, pp. 228–242.
- [25] P. Langley, W. Iba, and K. Thompson, "An analysis of Bayesian classifiers," in *Proceedings of the Tenth National Conference on Artificial Intelligence*, ser. AAAI'92, 1992, pp. 223–228.
- [26] V. Lemaire, F. Clérot, and M. Boullé, "An efficient shapley value computation for the naive bayes classifier," in *Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, R. Meo and F. Silvestri, Eds. Cham: Springer Nature Switzerland, 2025, pp. 75–90.
- [27] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Neural Information Processing Society (NeurIPS)*, 2017.
- [28] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artificial Intelligence*, vol. 267, pp. 1–38, 2019.
- [29] I. Stepin, J. M. Alonso, A. Catala, and M. Pereira-Fariña, "A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence," *IEEE Access*, vol. 9, pp. 11 781–11 803, 2021.
- [30] S. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual explanations without opening the black box: Automated decisions and the gdpr," *Harvard Journal of Law and Technology*, vol. 31, no. 2, pp. 841–887, 2018.
- [31] V. Lemaire, C. Hue, and O. Bernier, *Data Mining in Public and Private Sectors: Organizational and Government Applications*. IGI Global, 2010, ch. Correlation Analysis in Classifiers, pp. 204–218.
- [32] —, "Correlation explorations in a classification model," in *Workshop Data Mining Case Studies, SIGKDD*, 2009, p. 126.
- [33] S. Aryal and M. T. Keane, "Even if explanations: Prior work, desiderata & benchmarks for semi-factual xai," *arXiv*, 2023. [Online]. Available: <https://arxiv.org/abs/2301.11970>
- [34] V. Lemaire, N. Le Boudec, V. Guyomard, and F. Fessant, "Viewing the process of generating counterfactuals as a source of knowledge: a new approach for explaining classifiers," *2024 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2023.
- [35] T. Desautels, J. Calvert, J. Hoffman, M. Jay, Y. Kerem, L. Shieh, D. Shimabukuro, U. Chettipally, M. D. Feldman, C. Barton *et al.*, "Pre-

diction of sepsis in the intensive care unit with minimal electronic health record data: A machine learning approach,” *JMIR Medical Informatics*, vol. 4, no. 3, p. e5909, 2016.

- [36] I. Taneja, G. L. Damhorst, C. Lopez-Espina, S. D. Zhao, R. Zhu, S. Khan, K. White, J. Kumar, A. Vincent, L. Yeh *et al.*, “Diagnostic and prognostic capabilities of a biomarker and emr-based machine learning algorithm for sepsis,” *Clinical and Translational Science*, vol. 14, no. 4, pp. 1578–1589, 2021.
- [37] I. Persson, A. Macura, D. Becedas, and F. Sjövall, “Early prediction of sepsis in intensive care patients using the machine learning algorithm navoy® sepsis: A prospective randomized clinical validation study,” *Journal of Critical Care*, vol. 80, p. 154400, 2024.
- [38] “VFusion™ sepsis,” <https://www.vivacehealthsolutionsinc.com/vfusion/>, 2024, accessed: November 20, 2024.