

Cumulative Stay-time Representation for Electronic Health Records in Medical Event Time Prediction*

Takayuki Katsuki¹, Kohei Miyaguchi¹, Akira Koseki¹, Toshiya Iwamori¹,
Ryosuke Yanagiya² and Atsushi Suzuki²

¹IBM Research – Tokyo

²Fujita Health University

kats@jp.ibm.com, miyaguchi@ibm.com, akoseki,iwamori@jp.ibm.com,
yanagiya,aslapin@fujita-hu.ac.jp

Abstract

We address the problem of predicting when a disease will develop, i.e., medical event time (MET), from a patient’s electronic health record (EHR). The MET of non-communicable diseases like diabetes is highly correlated to cumulative health conditions, more specifically, how much time the patient spent with specific health conditions in the past. The common time-series representation is indirect in extracting such information from EHR because it focuses on detailed dependencies between values in successive observations, not cumulative information. We propose a novel data representation for EHR called cumulative stay-time representation (CTR), which directly models such cumulative health conditions. We derive a trainable construction of CTR based on neural networks that has the flexibility to fit the target data and scalability to handle high-dimensional EHR. Numerical experiments using synthetic and real-world datasets demonstrate that CTR alone achieves a high prediction performance, and it enhances the performance of existing models when combined with them.

1 Introduction

Predicting medical events, such as disease progression, from *electronic health records* (EHR) is an important task in medical and healthcare applications [Tan *et al.*, 2020]. The EHR represents a patient’s health history. Such prediction can assist in providing detailed health guidance, e.g., for early disease detection, intervention, and the allocation of limited resources in healthcare organizations [Inaguma *et al.*, 2020].

This paper addresses a scenario in which we predict *when* a patient will develop some disease after an index date, i.e., the *medical event time* (MET), from past observations in EHR, as shown in Fig. 1 [Liu *et al.*, 2018]. This is a common task in survival analysis and time-to-event analysis, and we focus on MET, not just its occurrence. The past observations for each patient come from a window that spans the initial observation time to the index date and contain lab test results at each

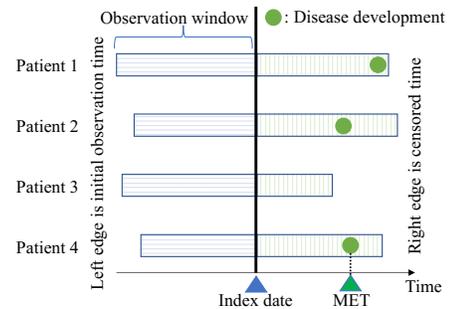


Figure 1: We predict when patient will develop disease after index date from EHR in observation window.

time, as shown in the LHS in Fig. 2. From accumulated EHR datasets, we learn a prediction model for MET.

A patient’s cumulative health conditions appearing in past observations in EHR are of help for MET prediction. They can be interpreted as the *cumulative stay-time* in specific health states—more specifically, how much time a patient has spent with different health conditions. For example, when a patient has high blood pressure, hyperglycemia, or high body fat for a long enough period, diseases can develop [James *et al.*, 2014; Am Diabetes Assoc, 2019]. In particular, for non-communicable diseases, like diabetes, the cumulative stay-time is extremely related to their progress and MET.

To utilize information in EHR, the common approach is to formalize the raw observations in EHR into an ordinary time-series representation [Zhang, 2019; Rubanova *et al.*, 2019; Hussain *et al.*, 2021]. In this approach, at each time, we record the value of each lab test result, as shown in the table in Fig. 2. The focus is on the detailed dependencies between values in successive observations. When we handle the cumulative stay-time with this representation, prediction models, such as recurrent neural networks (RNNs) [Zhang, 2019], need to encode values in an *entire time series* into the cumulative stay-time. This makes modeling the cumulative stay-time indirect.

We therefore propose directly extracting the cumulative stay-time from raw observations in EHR as a novel representation for EHR, that is, the *cumulative stay-time representation* (CTR). In contrast to the time-series representation, we record the cumulative stay-time at each combination of values

* A full version is available at: <https://arxiv.org/abs/2204.13451>

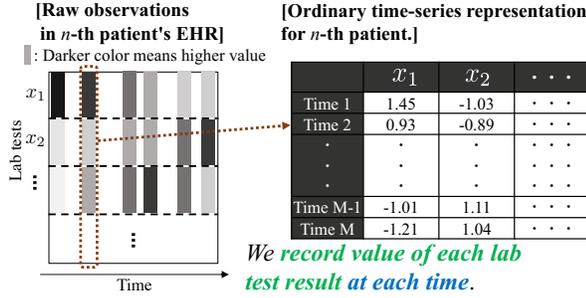


Figure 2: Ordinary time-series representation.

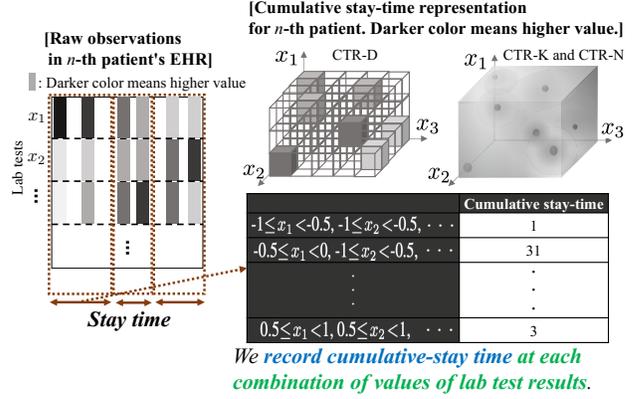


Figure 3: Cumulative stay-time representation.

of lab test results that represents a state, as shown in Fig. 3. This explicitly represents how long a patient stays in a specific health state.

Representations for modeling the cumulative stay-time in specific states and using it in prediction have been proposed in other domains than EHR modeling, such as for the usage history of batteries [Takahashi and Idé, 2012] and GPS trajectories [Liao *et al.*, 2018]. However, they are defined only with discrete state modeling that can be seen as bins of non-overlapping segmented values for lab test results, as shown in the table in Fig. 3. As such, they focus on low-dimensional observations, such as one, two, or three dimensions, and cannot handle more than several dimensions. This is because the number of states increases exponentially against the dimension of observation variables with this state definition. Since observations in EHR have many more dimensions, it is difficult to use these approaches on EHR directly.

This paper addresses the above difficulties by deriving methods for constructing CTR with enough scalability to handle EHR. We first formally derive a general construction of CTR by using the discrete state. This formalization leads to further enhancements of CTR with states defined as continuous measurements, CTR-K and CTR-N, which have states based on kernel functions and neural networks, respectively. They are more practical variants that avoid exponential increases in the number of states and lead to smooth interpolation between states. In addition, CTR-N can be learned from data, which enables flexible state modeling.

Contributions. Our main contributions are the following:

- We propose a novel representation for EHR for MET prediction, CTR, which represents how long a patient stays in a specific health state. This helps to model the cumulative health conditions of patients.
- We derive a trainable construction of CTR based on neural networks that adapts to data flexibly and has scalability for high-dimensional EHR.
- Extensive experiments on multiple MET prediction tasks with synthetic and real-world datasets show the effectiveness of CTR, especially for EHR with relatively longer observation periods, where cumulative health conditions are more crucial for MET. CTR shows modularity high enough to further improve the prediction performance when combined with other models.

2 Preliminary

2.1 Medical Event Time Prediction

Our goal is to construct a model for predicting the medical event time (MET), $y > 0$, after an index date on the basis of pairs of past observations and the corresponding timestamps, $\{\mathbf{X}, \mathbf{t}\}$, which are recorded in the EHR of a patient [Liu *et al.*, 2018], as shown in Fig. 1. The past observations for each patient contain M number of observations in an observation window $\mathbf{X} \equiv \{\mathbf{x}^{(m)}\}_{m=1}^M$, where the m -th observation $\mathbf{x}^{(m)}$ is represented as a D number of lab test results $\mathbf{x}^{(m)} \in \mathbb{R}^D$, and \mathbf{X} thus forms an $M \times D$ matrix. The timestamps are $\mathbf{t} \equiv \{t^{(m)}\}_{m=1}^M$, where the m -th timestamp is $t^{(m)} > 0$. We assume that each patient has an $M \times D$ matrix \mathbf{X} and M vector \mathbf{t} . Note that observation intervals can vary over time, and the length of sequence M can be different over patients.

When we take a machine learning-based approach, the raw observations $\{\mathbf{X}, \mathbf{t}\}$ must be formalized into a tractable representation that contains enough information for the MET prediction. We here denote the representation as a function, $\mathbf{z} \equiv \mathbf{z}(\mathbf{X}, \mathbf{t})$, whose output forms either a vector, matrix, or tensor depending on the formalization. Once $\{\mathbf{X}, \mathbf{t}\}$ is formalized into \mathbf{z} , we use \mathbf{z} as the input of the prediction model, $f(\mathbf{z})$, and learn it with a general scheme for minimizing the expected loss:

$$f^* \equiv \operatorname{argmin}_f E[\mathcal{L}(f(\mathbf{z}), y)], \quad (1)$$

where f^* is the optimal prediction model, \mathcal{L} is the loss function (e.g., the squared error), and E denotes the expectation over $p(y, \mathbf{X}, \mathbf{t})$. In the Experiments section, we will define a specific prediction model and loss function for each specific problem. By using the learned f^* , we can predict y for new data as $\hat{y} = f^*(\mathbf{z})$.

2.2 Property for Cumulative Stay-time Representation

This paper focuses on how to formalize raw observations $\{\mathbf{X}, \mathbf{t}\}$ into a tractable representation, \mathbf{z} . We directly model the cumulative stay-time of a specific patient's states with the construction of \mathbf{z} . We would like the representation to be:

(i) Direct to model stay time

How long a patient stays in a specific health condition can make diseases develop, particularly non-communicable diseases, like diabetes [James *et al.*, 2014; Am Diabetes Assoc, 2019].

(ii) Explicit to handle variable observation intervals

Observation intervals in EHR can vary over time and over patients since the observations are recorded when a patient is treated.

(iii) Scalable to high-dimensional data

Since EHR has high-dimensional observations, the representation should not cause combinatorial explosion against the number of dimensions in observations.

2.3 Representations for EHR

We first discuss the conventional related representations in this subsection considering the above properties; then, we derive the cumulative stay-time representation (CTR) for EHR in Section 3.

Ordinary Time-series Representation

In the ordinary time-series representation, raw observations $\{X, t\}$ are converted into the representation of a matrix form, $z_{ts} \in \mathbb{R}^{M \times D}$, whose two-dimensional index respectively represents timestamps and lab test names, as shown in Fig. 2. This corresponds to directly using the lab test results X as z_{ts} , $z_{ts}(X, t) \equiv X$. In the matrix, successive observations are put in adjacent rows. Thus, such representation helps in modeling the detailed dependencies between values in successive observations by using, for example, RNNs [Xiao *et al.*, 2017; Zhang, 2019], hidden Markov models (HMMs) [Alaa and van der Schaar, 2019; Hussain *et al.*, 2021], convolutional neural networks (CNNs) [Cheng *et al.*, 2016; Makino *et al.*, 2019], ODERNNs [Rubanova *et al.*, 2019], and Transformers [Luo *et al.*, 2020].

However, when we handle the cumulative stay-time, we need to consider how long and what state a patient has been in *over entire observations*, and similar health states are highly related to each other even if they are time-distant. In this case, models need to learn to encode values in an entire time series into the cumulative stay-time completely from data. The learning thus becomes indirect and costly, which leads to degraded performance. These approaches can still handle variable observation intervals by inputting the timestamps or intervals between observations as additional inputs. Also, they usually have good scalability for high-dimensional data. Therefore, we will show the performance difference with the proposed method in our experiments to investigate their indirectness in modeling stay time in a specific health state.

Cumulative Stay-time Representation

As discussed in Introduction, methods for directly modeling the cumulative stay-time in specific states have been proposed in domains other than EHR, e.g., cumulative stay-time for the internal state of batteries [Takahashi and Idé, 2012], location grids against GPS trajectories [Andrienko *et al.*, 2007; Boukhechba *et al.*, 2018], and indoor positioning using RFID [Zuo *et al.*, 2016; Jiang *et al.*, 2018]. They can handle variable observation intervals more naturally than the time-series representation.

The details of this approach and its practical limitations in EHR modeling will be described in Section 3.1. In brief, these methods do not have enough scalability for high-dimensional data because of their discrete state modeling and have not been defined formally with further extendibility. In this paper, we formally derive this approach as the *cumulative stay-time representation with discrete states* (CTR-D) and extend its state definition with a kernel function and trainable neural networks to address a higher dimensional case of EHR.

Other Representations

Representations for temporal observations other than time-series representations and CTR have been studied as reviewed in [Fu, 2011], such as methods based on binarization [Bagnall *et al.*, 2006] and segmented subsequences [Lovrić *et al.*, 2014]. Changing the domain from time into other domains based such as on the Fourier transform [Agrawal *et al.*, 1993] and tensor factorization [Ho *et al.*, 2014; Yu *et al.*, 2016; Yin *et al.*, 2019] is another common way. These methods assume high frequency and regular observation intervals, which is not the case in our scenario.

3 CTR: Cumulative Stay-time Representation

We propose a cumulative stay-time representation, CTR, for directly modeling the cumulative stay-time of a specific patient’s states as a novel formalization of raw observations in EHR.

3.1 CTR-D: CTR with Discrete States

We convert raw observations $\{X, t\}$ into the cumulative stay-time at a finite K number of states as K -dimensional vector z , whose k -th element is $z_k > 0$. Each state represents a combination of observed attribute values and can be seen as a bin segmented by a lattice that defines the value range of each attribute in each state, as shown in Fig. 3. We cumulatively fill each bin with the stay time of which the raw observation falls into the corresponding value ranges.

By using the state function $s(x^{\{m\}}) \in \{0, 1\}^K$, which outputs a one-hot vector representing the current state for input observation $x^{\{m\}}$, CTR z is defined as

$$z(X, t) \equiv \sum_m d^{\{m\}} s(x^{\{m\}}), \tag{2}$$

where $d^{\{m\}} \equiv \lambda^{t^{\{M\}} - t^{\{m\}}} (t^{\{m\}} - t^{\{m-1\}})$

where $d^{\{m\}}$ is the stay time for the m -th observation, which is estimated by calculating the difference between consecutive timestamps $t^{\{m\}}$ and $t^{\{m-1\}}$ with decay for weighting newer observations. λ is the decay rate and is optimized in training. Since the output of the function $s(x^{\{m\}})$ is a one-hot vector, only one element in the vector can become 1, and the others are 0, so the index for the element with value 1 represents the current state of the patient. Thus, for the m -th observation, the element in $d^{\{m\}} s(x^{\{m\}})$ with the current state becomes just $d^{\{m\}}$, and the others are 0. Through the summation of $d^{\{m\}} s(x^{\{m\}})$ over m , each element of z represents the sum of the stay time in each state over the observations. Also, from Eq. (2), this representation can explicitly handle variable

Algorithm 1 Cumulative stay-time representation

Input: Raw observations $\{\mathbf{X}, \mathbf{t}\}$ and state function $\mathbf{s}(\bullet)$
Output: Cumulative stay-time representation, CTR, \mathbf{z}

- 1: **Initialize:** $\mathbf{z} \leftarrow \mathbf{0}$
 - 2: **for** $m = 1$ **to** M (which can be parallelized over m)
 - 3: $\mathbf{s}^{\{m\}} \leftarrow \mathbf{s}(\mathbf{x}^{\{m\}})$
 - 4: $d^{\{m\}} \leftarrow \lambda^{t^{\{M\}} - t^{\{m\}}} (t^{\{m\}} - t^{\{m-1\}})$
 - 5: $\mathbf{z} \leftarrow \mathbf{z} + d^{\{m\}} \mathbf{s}^{\{m\}}$
-

observation intervals without any additional encoding. The algorithm is described in Algorithm 1.

The state function $\mathbf{s}(\mathbf{x}^{\{m\}})$ is defined by the indication function \mathbf{I} , which always outputs a K -dimensional one-hot vector representing the current state:

$$\mathbf{s}(\mathbf{x}^{\{m\}}) \equiv \mathbf{I}(\mathbf{x}^{\{m\}}, \mathbf{A}), \quad (3)$$

where $\mathbf{A} \in \{\mathbf{a}_k\}_{k=1}^K$ is the K number of non-overlapping collectively exhaustive value segments. The detailed definition of \mathbf{a}_k and the k -th element of the function \mathbf{I} are in the appendix. If $\mathbf{x}^{\{m\}}$ falls into the k -th segment, only the k -th element of $\mathbf{I}(\mathbf{x}^{\{m\}}, \mathbf{A})$ becomes 1, and the others are 0 because of the non-overlapping segmentation. An example segmentation is shown in the table in Fig. 3, which is based on equally spaced boundaries over the value range of $\mathbf{x}^{\{m\}}$, $[-1, -0.5, 0, 0.5, 1]$, where $x_d^{\{m\}}$ is defined in $[-1, 1)$. For example, in a 3-dimensional case, $K = 4^3 = 64$.

We call CTR in Eq. (2) with the state function in Eq. (3) *CTR with discrete states* (CTR-D). The discretely defined state $\mathbf{s}(\mathbf{x}^{\{m\}})$ is easy to understand. When the number of attributes in \mathbf{x} is small enough, we can practically use the function $\mathbf{s}(\mathbf{x}^{\{m\}})$ in Eq. (3) for computing \mathbf{z} .

However, since the number of combinations representing states grows exponentially with the number of attributes D , CTR-D cannot handle more than a few variables. Observations in EHR have many more attributes in general. For example, when we set the number of segments to 100, K becomes 100^D , which quickly causes a combinatorial explosion according to the number of attributes D . Also, the non-continuous boundary prevents generalization between adjacent states, though adjacent states should represent states similar to each other because of the shared boundaries between them in our definition in Eq. (3) (see also Appendix A). We thus extend the function $\mathbf{s}(\mathbf{x}^{\{m\}})$ into a more practical one in the following sections.

3.2 CTR-K: CTR with Kernel-defined States

For mitigating the exponential growth in the number of states, we change the definition of states in Eq. (2) from discrete, i.e., what variable values an observation has, to continuous, i.e., *how close an observation is to some basis vectors*, as shown in Fig. 3. Continuous states are no longer represented as a one-hot vector corresponding to a unique state; they are represented as a weight vector determining at what proportion we assign the current stay time to each state represented by bases. In this case, the number of states is limited to the number of bases and does not grow exponentially. This also leads

to interpolation between states and can smoothly represent intermediate states between the states.

We use a kernel function that represents affinities to bases for observations, where we construct the continuous state vector by assigning different values to multiple elements according to the affinities. The state function $\mathbf{s}_K(\mathbf{x}^{\{m\}}) \in \mathbb{R}^K$ based on the kernel function ϕ is defined as

$$\mathbf{s}_K(\mathbf{x}^{\{m\}}) \equiv \phi(\mathbf{x}^{\{m\}}, \mathbf{B}), \quad (4)$$

where $\mathbf{B} \equiv \{\mathbf{b}^{\{k\}}\}_{k=1}^K$ is the K number of bases, and $\mathbf{b}^{\{k\}} \in \mathbb{R}^D$ is the k -th basis. For example, $\mathbf{s}_K(\mathbf{x}^{\{m\}}) = \{0, 0.3, 0.7, 0, \dots, 0\}$ means that we assign the stay time for the m -th observation with weights of 0.3 and 0.7 to the second and third states, respectively, in the summation in Eq. (2). Bases can be randomly sampled from the training set. When the observation variables are real-valued, as in our scenario, the choice of ϕ is an RBF kernel, whose definition is provided in Appendix B. We can also use other kernels, such as tf-idf vector + cosine similarity [Rajaraman and Ullman, 2011], for binary features.

We call CTR in Eq. (2) with the state function in Eq. (4) *CTR with kernel-defined states* (CTR-K).

3.3 CTR-N: CTR with Neural Network-defined States

Additionally, we can consider the requirement for continuous state $\mathbf{s}_K(\mathbf{x}^{\{m\}})$ in Eq. (4) to represent a similar observation with a similar weight vector. Such a vector can also be modeled with neural networks since they are trained to produce similar outputs from similar inputs.

We thus extend $\mathbf{s}_K(\mathbf{x}^{\{m\}})$ to $\mathbf{s}_N(\mathbf{x}^{\{m\}}) \in \mathbb{R}^K$ by replacing kernel function ϕ with a trainable neural network, \mathbf{g} , e.g., multilayer perceptron (MLP), that produces a state-indicating weight vector similar to ϕ , as

$$\mathbf{s}_N(\mathbf{x}^{\{m\}}) \equiv \mathbf{g}(\mathbf{x}^{\{m\}}, \boldsymbol{\theta}_g), \quad (5)$$

where $\boldsymbol{\theta}_g$ are parameters for the neural network. The outputs of the final layer for \mathbf{g} should be normalized, such as by the softmax function, as a weight vector, and the number of the outputs is K . The specific neural network structure for \mathbf{g} is shown in the Experiments section.

We call CTR in Eq. (2) with the state function in Eq. (5) *CTR with neural network-defined states* (CTR-N). This representation can be learned from data and thus provides more flexibility in adjusting the state definition to target data. Also, in contrast to CTR-D and CTR-K, CTR-N does not require having to choose the state boundaries or the bases.

Formally, the following lemma characterizes CTR-K and CTR-N with respect to the three properties: (i) direct to model stay time, (ii) explicit to handle variable observation intervals, and (iii) scalable to high-dimensional data:

Lemma 1. (i) Every element in $\mathbf{z}(\mathbf{X}, \mathbf{t})$ is a linear function of stay time d . Hence, $\mathbf{z}(\mathbf{X}, \mathbf{t})$ is a direct representation of stay time. (ii) $\mathbf{z}(\mathbf{X}, \mathbf{t})$ is a function of an observation interval $t^{\{m\}} - t^{\{m-1\}}$. (iii) The number of dimensions in $\mathbf{s}_K(\mathbf{x})$ and that of the corresponding $\mathbf{z}(\mathbf{X}, \mathbf{t})$ depend on the number of bases, not the number of attributes in \mathbf{x} , D . Also, the

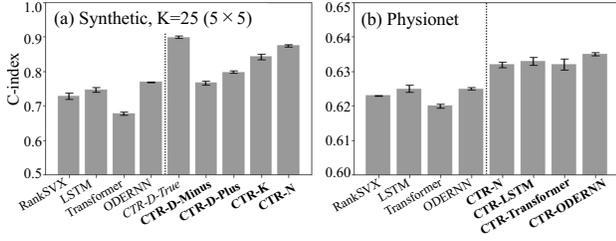


Figure 4: Comparison of C-index on (a) Synthetic and (b) Physionet (higher is better).

number of dimensions in $s_N(\mathbf{x})$ and that of the corresponding $\mathbf{z}(\mathbf{X}, t)$ depend on the number of outputs of g , not D .

The ordinary time-series representation $\mathbf{z}_{ts}(\mathbf{X}, t)$ does not satisfy (i). CTR-D does not satisfy (iii).

Gradients for learning model parameters. We minimize Eq. (1) by using gradient-based optimization methods. We learn f on the basis of Eqs. (1) and (2) by using the gradients for the model parameters for f , θ_f , as

$$\frac{\partial \mathcal{L}}{\partial \theta_f} = E \left[\frac{\partial \mathcal{L}}{\partial f} \frac{\partial f}{\partial \theta_f} \right], \quad (6)$$

where we omit the inputs of the functions for simplicity.

For CTR-N, in addition to learning f , we learn the parameters θ_g of neural network g in (5), which represents the state $s_N(\mathbf{x}^{[m]})$. The gradients for θ_g can be derived as

$$\frac{\partial \mathcal{L}}{\partial \theta_g} = E \left[\frac{\partial \mathcal{L}}{\partial f} \frac{\partial f}{\partial \mathbf{z}} \sum_m d^{[m]} \frac{\partial g(\mathbf{x}^{[m]}, \theta_g)}{\partial \theta_g} \right]. \quad (7)$$

4 Experiments

We assessed the prediction performance of our method CTR in numerical experiments to show the effectiveness of directly modeling the cumulative stay-time.

Evaluation metric. We report the mean and standard error of the concordance index (C-index) [Liu *et al.*, 2018] across 5-fold cross-validation, each with a different randomly sampled training-testing split. The C-index becomes a high value when regressed MET values follow the true ordering, and it can handle censored data, which is essential for our application with EHR. For each fold of the cross-validation, we randomly sampled 20% of the training set for use as a validation set to determine the best hyperparameters for each method, where hyperparameters providing the highest C-index in the validation set were chosen.

Implementations for CTRs. We used the loss proposed in [Liu *et al.*, 2018], which can handle censored data well. We used a 3-layer MLP with ReLU [Nair and Hinton, 2010] (more specifically, D -100-1) as the prediction model f for CTR-D, CTR-K, and CTR-N. In CTR-N, we used a 4-layer MLP with ReLU (more specifically, D -100-100-100) as g in Eq. (5), where the final layer is the softmax function for normalization. More details on the implementation, such as the definitions for states in CTR-D and CTR-K, are provided in

the appendix. Note that we could not use CTR-D for real-world EHR experiments since we needed to handle a large number of attributes, which would cause a combinatorial explosion for K . For example, when $D = 37$ and we set the number of segments to 100, K becomes 100^{37} , where one of our real-world datasets contains $D = 37$ number of attributes.

Methods compared. The proposed method was compared with four state-of-the-art methods: **RankSVX** [Liu *et al.*, 2018], **LSTM**, **Transformer** [Luo *et al.*, 2020], and **ODERNN** [Rubanova *et al.*, 2019], where the loss for each method is the same as CTRs, and the number of hidden states in each model is the same as CTR-N. They are based on the time-series representation with the stay time d treated as another column for the input time series. Details on their implementation are provided in the appendix.

Combinations of CTR-N and compared methods. In real-world EHR experiments, we also examined combinations of CTR-N with the compared methods LSTM, Transformer, and ODERNN (**CTR+LSTM**, **CTR+Transformer**, and **CTR+ODERNN**, respectively). In these combinations, the representation just before the final linear layer of each model was extracted, concatenated with CTR \mathbf{z} as a single vector, and fed into the prediction model $f(\mathbf{z})$. They were trained in an end-to-end manner.

Computing infrastructure. All of the experiments were carried out on workstations having 128 GB of memory, a 4.0-GHz CPU, and an Nvidia Tesla V100 GPU. The computational time for each method was a few hours for producing the results for each dataset, *except for ODERNN, which was 10 to 20 times slower than the other methods.*

4.1 Results

We first use a synthetic dataset to investigate 1) *whether the method with CTR can indeed learn to predict what cannot be learned without CTR.* Then, real-world EHR datasets are used to show the 2) *practical effectiveness of CTR.* Finally, we show that 3) *CTR enhances the prediction performance, especially for EHR with relatively longer observation periods, where cumulative health conditions are more crucial for MET.* Details on the datasets are provided in the appendix.

Synthetic. The Synthetic dataset was generated on the basis of our assumed nature, i.e., the cumulative stay-time for each state leads to the development of a disease. The number of records was $N = 1,000$, the observation length for each record was $M = 10$, and the number of attributes was $D = 2$. The observation intervals varied between records. We addressed large ($K = 100$), medium ($K = 49$), and small ($K = 25$) numbers of states settings in data generation. The results are shown in Fig. 4-(a), where the bars represent the means of the C-index across 5-fold cross-validation, and the confidence intervals are standard errors. We show the results with the small ($K = 25$) number of states here, and the others are provided in the appendix. We can see that the overall performance of the proposed method was significantly better than those of the compared methods, which demonstrates that *the proposed method with CTR can learn what cannot be learned without CTR well.* Note that we used multiple settings for CTR-D: the same number of states K for the

data generation (CTR-D-True), $K_d - 1$ (CTR-D-Minus), and $K_d + 1$ (CTR-D-Plus), where $K_d \equiv \sqrt[3]{K}$. We used CTR-D-True as the reference when we knew the true CTR; it thus should achieve the highest score. CTR-K and CTR-N were better than CTR-D with the wrong number of states even if the error was 1, which demonstrates that CTR-K and CTR-N have a better generalization capability than CTR-D against data variation. CTR-N performed the best, which demonstrates that CTR-N learns states from data well.

Physionet. The Physionet dataset is a publicly available real-world EHR dataset (Physionet Challenge 2012 [Silva *et al.*, 2012]). The number of records was $N = 8,000$, and the number of attributes was $D = 37$. The observation intervals varied between records. The results for the MET prediction task for patient death are shown in Fig. 4-(b) with the same configuration as the results of the Synthetic dataset. The performances of the methods with CTR were better than those of the methods without CTR by a sufficient margin in terms of standard error. These results demonstrate that CTR can improve the C-index in the MET prediction problem with real-world EHR. We omitted the results with CTR-K since it was always worse than CTR-N. CTR-N achieved the best performance on average in comparison with the single models. In addition, when looking at results for combinations of CTR-N and other models, CTR+LSTM, CTR+Transformer, and CTR+ODERNN, we can see that adding CTR-N to these models improved their performance further, which shows the high modularity of CTR to work complementarily with other models. *This shows that CTR and the time-series models captured different temporal natures in real-world EHR.* We can automatically determine which type of temporal natures to take into account with the training dataset by training the prediction model f , which is put on top of these models.

Case study. The above experiments on two different datasets have shown that the methods with CTR have superior prediction performance compared with the state-of-the-art methods for MET prediction from EHR. Here, we show a real healthcare use-case, where we predict the onset of complications with diabetes mellitus from a real-world big EHR database. We used datasets provided by one of the largest hospitals in Japan that has maintained a big database of more than 400,000 patients since 2004 [Makino *et al.*, 2019; Inaguma *et al.*, 2020]. We worked with six datasets for six kinds of complications of diabetes mellitus: hyperosmolar (HYP), nephrology (NEP), retinopathy (RET), neuropathy (NEU), vascular disease (VAS), and other complications (OTH), each of which has over $N = 15,000$ records. The number of attributes was $D = 26$, and the observation intervals and lengths varied between records. In this scenario, ODERNN, which is 10 to 20 times slower than the other methods, did not meet the practical needs for this large-scale dataset. Thus, we here show a comparison between the proposed method and the second-best baseline, LSTM, in experiments with the Synthetic and Physionet datasets. The results of the mean and standard error of the C-index across 5-fold cross-validation are listed in Table 1. For most of the six tasks having over 15,000 samples each, the performances of the methods with CTR were better than LSTM by a sufficient

	LSTM	CTR-N	CTR+LSTM
HYP	0.589±0.026	0.612±0.026	0.583±0.036
NEP	0.689±0.012	0.739±0.010	0.708±0.007
RET	0.717±0.013	0.721±0.026	0.745±0.008
NEU	0.569±0.023	0.608±0.020	0.600±0.020
VAS	0.503±0.035	0.481±0.013	0.534±0.022
OTH	0.718±0.015	0.741±0.020	0.734±0.011
Average	0.664	0.687	0.688

Table 1: Comparison of C-index in case study (higher is better). Confidence intervals are standard errors. Best results are in bold.

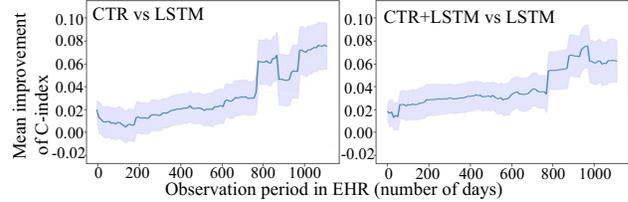


Figure 5: Analysis of performance improvement over different observation periods (higher is better).

margin in terms of standard error. Complications of diabetes mellitus are known to develop due to time-cumulative effects for vessels with an unhealthy status. The results showed that our explicit mechanisms are essential to learning such effects to achieve higher prediction performance.

Performance analysis on different observation periods. We further analyzed the performance improvements of the methods with CTR compared with LSTM over different observation periods by using the Case study dataset containing EHR with more extended periods. We plotted the mean improvements of the C-index between them for data with different observation periods, as shown in Fig. 5, where the confidence intervals are standard errors of the improvements. The right region in the figure show the results for data with longer observation periods. It shows that *CTR improved the performance, especially for data with relatively longer observation periods*, where cumulative health conditions are more crucial for MET prediction.

5 Conclusion

We proposed a cumulative stay-time representation, CTR, for a novel representation of EHR. CTR can efficiently handle the situation in which the development of some disease is related to the cumulative stay-time of a specific patient’s health conditions, e.g., non-communicable diseases. We developed three variations of CTR with discrete states, continuous states with kernel functions, and continuous states with neural networks. In particular, CTR with neural networks, CTR-N, is practical because it has scalability handling high-dimensional data and can be learned from data for flexibility in adjusting to the target data. An experimental evaluation demonstrated that the method with CTR performed better than the methods without CTR. Application to domains other than EHR will be an interesting avenue for future work.

References

- [Agrawal *et al.*, 1993] Rakesh Agrawal, Christos Faloutsos, and Arun Swami. Efficient similarity search in sequence databases. In *FODO*, pages 69–84, 1993.
- [Alaa and van der Schaar, 2019] Ahmed Alaa and Mihaela van der Schaar. Attentive state-space modeling of disease progression. In *NeurIPS*, 2019.
- [Am Diabetes Assoc, 2019] Am Diabetes Assoc. Classification and diagnosis of diabetes: standards of medical care in diabetes. *Diabetes care*, 42(Supplement 1):S13–S28, 2019.
- [Andrienko *et al.*, 2007] Gennady Andrienko, Natalia Andrienko, and Stefan Wrobel. Visual analytics tools for analysis of movement data. *ACM SIGKDD Explorations Newsletter*, 9(2):38–46, 2007.
- [Bagnall *et al.*, 2006] Anthony Bagnall, Eamonn Keogh, Stefano Lonardi, Gareth Janacek, et al. A bit level representation for time series data mining with shape based similarity. *Data Mining and Knowledge Discovery*, 13(1):11–40, 2006.
- [Boukhechba *et al.*, 2018] Mehdi Boukhechba, Philip Chow, Karl Fua, Bethany A Teachman, and Laura E Barnes. Predicting social anxiety from global positioning system traces of college students: feasibility study. *JMIR mental health*, 5(3):e10101, 2018.
- [Cheng *et al.*, 2016] Yu Cheng, Fei Wang, Ping Zhang, and Jianying Hu. Risk prediction with electronic health records: A deep learning approach. In *SDM*, pages 432–440, 2016.
- [Fu, 2011] Tak-chung Fu. A review on time series data mining. *Engineering Applications of Artificial Intelligence*, 24(1):164–181, 2011.
- [Ho *et al.*, 2014] Joyce C Ho, Joydeep Ghosh, and Jimeng Sun. Marble: high-throughput phenotyping from electronic health records via sparse nonnegative tensor factorization. In *KDD*, pages 115–124, 2014.
- [Hussain *et al.*, 2021] Zeshan Hussain, Rahul G Krishnan, and David Sontag. Neural pharmacodynamic state space modeling. In *ICML*, 2021.
- [Inaguma *et al.*, 2020] Daijo Inaguma, Akimitsu Kitagawa, Ryosuke Yanagiya, Akira Koseki, Toshiya Iwamori, Michiharu Kudo, and Yukio Yuzawa. Increasing tendency of urine protein is a risk factor for rapid egfr decline in patients with ckd: A machine learning-based prediction model by using a big database. *PLoS One*, 15(9 September), September 2020.
- [James *et al.*, 2014] Paul A James, Suzanne Oparil, Barry L Carter, William Cushman, Cheryl Dennison-Himmelfarb, Joel Handler, Daniel T Lackland, Michael L LeFevre, Thomas D MacKenzie, Olugbenga Ogedegbe, et al. 2014 evidence-based guideline for the management of high blood pressure in adults: report from the panel members appointed to the eighth joint national committee (jnc 8). *Jama*, 311(5):507–520, 2014.
- [Jiang *et al.*, 2018] Jianguo Jiang, Xiaonan Lv, Yanfang Zhang, and Siye Wang. Research on hotspots finding in indoor space based on regression analysis. In *ICCDE*, pages 98–102, 2018.
- [Liao *et al.*, 2018] Lyuchao Liao, Jianping Wu, Fumin Zou, Jengshyang Pan, and Tingting Li. Trajectory topic modelling to characterize driving behaviors with gps-based trajectory data. *Journal of Internet Technology*, 19(3):815–824, 2018.
- [Liu *et al.*, 2018] Bin Liu, Ying Li, Zhaonan Sun, Soumya Ghosh, and Kenney Ng. Early prediction of diabetes complications from electronic health records: A multi-task survival analysis approach. In *AAAI*, 2018.
- [Lovrić *et al.*, 2014] Miodrag Lovrić, Marina Milanović, and Milan Stamenković. Algorithmic methods for segmentation of time series: An overview. *Journal of Contemporary Economic and Business Issues*, 1(1):31–53, 2014.
- [Luo *et al.*, 2020] Junyu Luo, Muchao Ye, Cao Xiao, and Fenglong Ma. Hitanet: Hierarchical time-aware attention networks for risk prediction on electronic health records. In *KDD*, pages 647–656, 2020.
- [Makino *et al.*, 2019] Masaki Makino, Ryo Yoshimoto, Masaki Ono, Toshinari Itoko, Takayuki Katsuki, Akira Koseki, Michiharu Kudo, Kyoichi Haida, Jun Kuroda, Ryosuke Yanagiya, et al. Artificial intelligence predicts the progression of diabetic kidney disease using big data machine learning. *Scientific reports*, 9(1):1–9, 2019.
- [Nair and Hinton, 2010] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010.
- [Rajaraman and Ullman, 2011] Anand Rajaraman and Jeffrey David Ullman. *Mining of massive datasets*. Cambridge University Press, 2011.
- [Rubanova *et al.*, 2019] Yulia Rubanova, Ricky TQ Chen, and David K Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. In *NeurIPS*, pages 5320–5330, 2019.
- [Silva *et al.*, 2012] Ikaro Silva, George Moody, Daniel J Scott, Leo A Celi, and Roger G Mark. Predicting in-hospital mortality of icu patients: The physionet/computing in cardiology challenge 2012. In *Computing in Cardiology*, pages 245–248. IEEE, 2012.
- [Takahashi and Idé, 2012] Toshihiro Takahashi and Tsuyoshi Idé. Predicting battery life from usage trajectory patterns. In *ICPR*, pages 2946–2949. IEEE, 2012.
- [Tan *et al.*, 2020] Qingxiong Tan, Mang Ye, Baoyao Yang, Siqi Liu, Andy Jinhua Ma, Terry Cheuk-Fung Yip, Grace Lai-Hung Wong, and PongChi Yuen. Data-gru: Dual-attention time-aware gated recurrent unit for irregular multivariate time series. In *AAAI*, pages 930–937, 2020.
- [Xiao *et al.*, 2017] Shuai Xiao, Junchi Yan, Xiaokang Yang, Hongyuan Zha, and Stephen M Chu. Modeling the intensity function of point process via recurrent neural networks. In *AAAI*, pages 1597–1603, 2017.
- [Yin *et al.*, 2019] Kejing Yin, Dong Qian, William K Cheung, Benjamin CM Fung, and Jonathan Poon. Learning phenotypes and dynamic patient representations via rnn regularized collective non-negative tensor factorization. In *AAAI*, pages 1246–1253, 2019.
- [Yu *et al.*, 2016] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. Temporal regularized matrix factorization for high-dimensional time series prediction. In *NIPS*, pages 847–855, 2016.
- [Zhang, 2019] Yuan Zhang. Attain: Attention-based time-aware lstm networks for disease progression modeling. In *IJCAI*, pages 4369–4375, 2019.
- [Zuo *et al.*, 2016] Yi Zuo, Katsutoshi Yada, and ABM Shawkat Ali. Prediction of consumer purchasing in a grocery store using machine learning techniques. In *APWC on CSE*, pages 18–25, 2016.