

Learning Hierarchical Representations of Electronic Health Records for Clinical Outcome Prediction

Luchen Liu¹, Haoran Li¹, Zhiting Hu², Haoran Shi¹,
Zichang Wang¹, Jian Tang, PhD^{3,4}, Ming Zhang*, PhD¹

¹ Department of Computer Science, Peking University, Beijing, China

² Carnegie Mellon University, Pittsburgh, PA, US

³ Mila - Québec AI Institute, Montréal, Québec, Canada

⁴ HEC Montréal, Montréal, Québec, Canada

Abstract

Clinical outcome prediction based on Electronic Health Record (EHR) helps enable early interventions for high-risk patients, and is thus a central task for smart healthcare. Conventional deep sequential models fail to capture the rich temporal patterns encoded in the long and irregular clinical event sequences in EHR. We make the observation that clinical events at a long time scale exhibit strong temporal patterns, while events within a short time period tend to be disordered co-occurrence. We thus propose differentiated mechanisms to model clinical events at different time scales. Our model learns hierarchical representations of event sequences, to adaptively distinguish between short-range and long-range events, and accurately capture their core temporal dependencies. Experimental results on real clinical data show that our model greatly improves over previous state-of-the-art models, achieving AUC scores of 0.94 and 0.90 for predicting death and ICU admission, respectively. Our model also successfully identifies important events for different clinical outcome prediction tasks.

1 Introduction

The ever-growing massive Electronic Health Records (EHR) data expose an opportunity for large-scale data-driven health analysis and intelligent medical care. Predicting clinical outcomes, such as death and intensive care unit (ICU) admission, plays an important role in improving the performance of healthcare systems. For instance, accurate clinical prediction based on patients existing medical records can enable advanced and timely medical intervention.

Clinical outcome prediction is challenging because it is hard to utilize rich temporal information encoded in the sequence of clinical events in EHR data.¹ In particular, EHR data usually consist of clinical events with irregular intervals² from heterogeneous sources, including patient health features (vital sign measurements, laboratory test results, etc), medical interventions (procedures, drug inputs, etc), and expert judgments (diagnoses, notes, etc).³ The temporal order of these events is critical for predicting clinical outcome. For example, patient health features can be affected by previous medical interventions, and in turn determine subsequent medical interventions through expert judgments.

Conventional approaches that directly apply classic deep sequential models, such as recurrent neural networks^{4,5} (RNN) and convolutional neural networks⁶ (CNN), usually fail to capture temporal dependencies in such long irregular event sequences,^{7,8} as long-term dependencies can easily exceed the modeling capacity. To handle irregularly timed events, some extensions of the classic models have been developed, such as time-aware RNN⁹ or CNN.¹⁰ However, their performance is still largely unsatisfying due to the limited ability to capture long-term dependencies.

This work aims to address the above challenges. We first make a key observation that, though the clinical events in EHR data can exhibit strong temporal patterns at a *long* time scale, the events occurring within a *short* time period usually do not have a definite order. Specifically, unlike word sequences in natural languages where word tokens are ordered by grammar rules, clinical events recorded in a short period of time are instead a series of events, such as clinical laboratory test results, that only reflect the patients status in different views. Therefore, direct temporal modeling of such short-range events as in previous work can introduce noise and harm the temporal predictive performance. Instead, local dependencies of these events should be modeled as event co-occurrence, and we can further select critical events from each of these short-range event groups as the basis for modeling the real temporal dependencies at

*Corresponding author.

a long time scale. A key difficulty to this end is that the criterion for distinguishing long-term temporal dependencies from the local co-occurrence of critical events in a short range can vary across different diseases and phases, especially in the irregular EHR data.^{11,12}

To address the difficulties mentioned above, we propose a hierarchical neural network for clinical outcome prediction. Specifically, we adaptively segment irregular event sequences into sequential groups to distinguish short-range co-occurrence and long-term temporal dependencies as well as to learn hierarchical representations of the event sequence. At the low level, the model automatically identifies critical clinical events in each group and aggregates the events to form event group representations. At the high level, meaningful long-term dependencies of clinical event groups are captured in a sequence representation vector by a recurrent neural network. Compared to traditional methods, the proposed method has several advantages:

- Our model can deal with the temporal irregularity of clinical event sequences by adaptively segmenting an event sequence into sequential groups.
- Our model learns a hierarchical representation of long and irregular event sequences to capture long-term dependencies of clinical events.
- The model is capable of discovering critical event groups as well as critical events in each group, through a temporal attention and event attention mechanism. This provides useful clinical insights for accurate prediction.

2 Related Works

2.1 Modeling EHR Data

Most existing works based on EHR data have either focused on stationary clinical text^{13,14} and images,^{15,16,17} or ignored irregular time intervals of temporal clinical events.^{18,19,20} For example, previous work trained the semantic embeddings for the categories of clinical events for adverse drug event detection,¹⁹ or proposed a multi-view learning method that generalizes Canonical Correlation Analysis for an arbitrary collection of matrices involving missing data.²⁰ These works make predictions based on the clinical events with regular time intervals, and cannot distinguish short-range order from long-term temporal order of different diseases and patients. Our work addresses the issue by adaptive segmentation of clinical event sequences.

As long-term temporal dependencies are hard to capture, many works use a small subset of the whole EHR information, to avoid dealing with the long clinical event sequences. Some works select a subset of the numerical clinical features (the numerical attributes of clinical events) in the EHR data according to the expertise of clinicians.⁵ For instance, Yoon only uses a set of 21 (temporal) physiological streams comprising a set of 11 vital signs and 10 lab test scores to predict ICU admission.²¹ Some works used graphical models to model patients health status.²² Some techniques transform selected 99 time series features of all the EHR data into a new latent space using the hyper-parameters of multi-task GP (MTGP) models to model patient similarity.¹ Recently, RETAIN used two reversed recursive neural networks (RNN) generating attention variables of sequential international disease classification (ICD) code groups for the prediction.²³ However, the codes are grouped by the fixed-length time slots for distinct patients and diseases, and local dependencies and long-term dependencies may be mixed up. But these works can lose significant information, due to the expert bias when selecting a limit fraction of all clinical features in EHR as the input of the models, and fail to provide new data-driven insights for better healthcare.

2.2 Clinical Outcome Prediction

The clinical outcome prediction problem is studied by many works. However many of them cannot take advantage of the temporal information in EHR data for prediction. Some of these studies used latent variable models to decompose bag-of-words free-text extracted from clinical event descriptions into meaningful features to predict patient mortality.²⁴ “Deep patient”²⁵ arranged all clinical descriptors (features) in a period of time in a sparse vector without temporal information and trained the deep representation of patients with a 3 layer denoising autoencoder for diagnosis. Some work studied how to diagnose and predict Alzheimers disease (AD) with a hybrid manifold learning for non-temporal clinical feature embedding and the bootstrap aggregating (Bagging) algorithm.²⁶ There is also a work model EHR

data by factorizing the medical feature matrix into a latent medical concept mapping matrix and a concept value evolution matrix, and then they averaged all vectors in the evolution matrix to predict heart failure.²⁷ Our model learns the hierarchical representations of clinical event sequences to utilize the temporal information for clinical outcome prediction.

3 Data and Task Descriptions

We give the notations and data descriptions of the predictive tasks in the following.

Clinical Events in EHR

A clinical event is a record in the database of EHR, which describes a clinical activity of a particular patient at a certain time. The events can be measurements of vital signals, injection of drugs, results of laboratory tests, and so on, which are summarized in table 1. Each clinical event has some attributes, including categorical attributes and numerical attributes. For example, the lab test event e_t has 2 categorical features and 1 numerical feature: $e_t = [LabItem : Cholesterol, AbnormalLabel : Abnormal, ResultIndex : 51\mu/L]$. The meaning of this event is the result of the Cholesterol test is $51\mu/L$, which reflects an abnormal health status. An episode of a patient EHR data is a clinical event sequence, which may consist of hundreds of clinical events.

Clinical Outcome Prediction

Clinical outcome prediction is to dynamically predict whether a clinical outcome will happen in 24 hours based on an episode of a patient. We aim to dynamically predict two outcomes in this work. In the first “death prediction task”, the outcome is death in hospital or discharge to home. In the second “ICU admission prediction task”, the outcome is clinical deterioration (need to be immediately transformed to ICU), or stable clinical status.

Patient Cohort Setup

We set up two datasets from one real clinical data source, MIMIC-III³ (Medical Information Mart for Intensive Care III), which is a large, freely-available database comprising de-identified health-related data associated with over forty thousand patients who stayed in intensive care units of the Beth Israel Deaconess Medical Center between 2001 and 2012.

We extract 18192 kinds of clinical events with their attributes from the database to get event sequences of patients (the events with top frequency are listed in Table 1). The events whose frequency is less than 2500 are dropped out. And we also drop out the admissions, of which the time span from the beginning to the target clinical outcome is less than 36 hours. Each input of a sample is an episode of a patient clinical event sequence 24 hours before the target outcome. The statistics of final clinical event sequences in the two tasks are summarized in Table 2.

4 Methodology

In this section, we introduce the technical details of our proposed model. Our model first segments the whole clinical event sequence into several event groups via the adaptive event sequence segmentation module. Then the model learns hierarchical representations of event sequences with both event attention and temporal attention mechanisms. The architecture of our model is illustrated in Figure 1.

4.1 Adaptive Segmentation

To distinguish long-term temporal dependencies from co-occurrence of important events in short range, we adaptively segment an event sequence for a patient into sequential groups, according to the irregular record time of events. As events in the same group are exchangeable, sequential groups can avoid the influence of the short-range order noisy in clinical events. Moreover, sequential groups reduce the length of the sequences fed to RNN, which makes capturing long-term temporal dependencies easier.

event descriptions	event name (top 2)	frequency (rank-in-all)	coverage (rank-in-all)	frequency per patient
<i>Chart events</i> include routine vital signs, ventilator settings, mental status, and so on.	Heart Rate	5171250 (0.01%)	0.64 (0.16%)	173.4
	SpO2	3410702 (0.01%)	0.479 (0.33%)	153.0
<i>Input events</i> are any fluid which have been administered to the patient: such as oral or tube feedings or intravenous solutions containing medications.	0.9% Normal Saline	2363812 (0.05%)	0.393 (0.66%)	129.2
	Propofol	369103 (0.81%)	0.217 (1.45%)	36.5
<i>Lab events</i> contain all laboratory measurements for a given patient.	Hematocrit	881846 (0.22%)	0.976 (0.01%)	19.4
	Potassium	845825 (0.23%)	0.886 (0.05%)	20.5
<i>Procedure events</i> contain procedures for patients.	Chest X-Ray	32723 (3.2%)	0.204 (1.52%)	3.44
	EKG	13962 (4.35%)	0.167 (1.82%)	1.79
<i>Output events</i> are fluids which have either been excreted by the patient, such as urine output, or extracted from the patient.	Chest Tubes CTICU CT 1	151766 (1.57%)	0.098 (2.67%)	33.2
	Urine	107465 (1.93%)	0.075 (3.05%)	30.8

Table 1: Statistics of High Frequency Examples in Different Types of Clinical Events

Dataset	# of samples	# of events	Avg timespan
death	24301(8%)	20290879	3d 15h 58m
ICU admission	19451(21%)	14515198	4d 18h 31m

Table 2: Statistics of the datasets (the percentage in the second column is the positive sample rate)

We find segmentation points of an event sequence by minimizing the maximum time span of the resulting segments. Formally, given a event sequence $\{e_t\}$, $k-1$ segmentation points can split the sequence into k groups $\{G_i\}_{i=1}^k$, where the event group $G_i = \{e_t | t'_i \leq e_t.time \leq t'_{i+1}\}$ is an episode of clinical events from time t'_i to time t'_{i+1} , where $e_t.time$ is the record time of the event e_t . And the time span of a group is defined as the time difference of the last event and the first in the group, namely $span(G_i) = \max_{e_k, e_j \in G_i} \{e_k.time - e_j.time\}$. So the optimal choice of the segmentation points can be found by minimizing the following:

$$t'_1, \dots, t'_{k+1} = \arg \min_{t'_1, \dots, t'_{k+1}, k \leq M} \max_{1 \leq i \leq k} \{span(G_i)\}$$

where M is the max number of groups and $k \leq M$ is the constraint to avoid the segmentation too fine-grained.

The adaptive segmentation is designed in a way of the combination of greedy method and binary search. We binary search the minimal upper bound of the maximum time span of all groups. And we then verify the searched upper bound of time spans by trying to greedily construct a solution satisfying the constraints of M and the time span upper bound. The time complexity of the algorithm is $O(L \times \log T)$, where L is the length of the event sequence and T is the difference of end time and start time of the sequence. We can regard this algorithm as an algorithm of linear complexity with a big consistent coefficient.

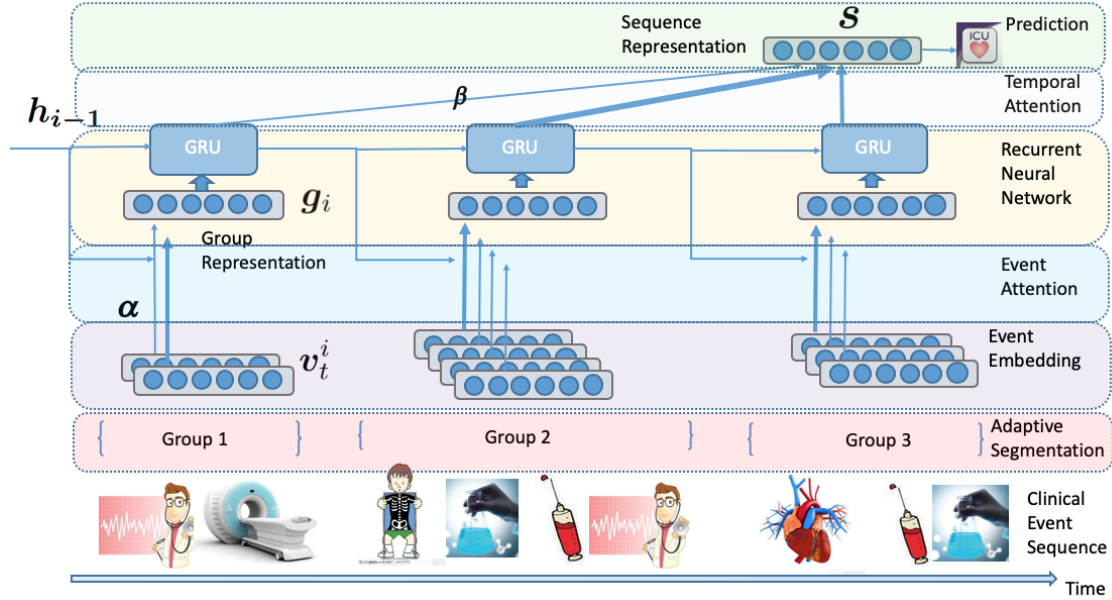


Figure 1: The overall architecture of our model. The original irregular event sequence is segmented into sequential event groups by the adaptive segmentation module. Then our model learns the hierarchical representation of the sequential event groups. In the low-level representations, each event group is represented as a vector g_i by the event attention. In the high-level representations, the embedded sequential groups are modeled by general recurrent units (GRU) with inter-group temporal attention.

Clinical events have many attributes, which are not considered in previous popular methods, such as word2vec, GloVe and so on. To represent clinical events with their attributes, we embed each clinical event e_t into the low dimension space as a vector v_t in the way described in the previous work.²⁸ The representing vector $v_t \in \mathbb{R}^N$ (where N is the event embedding dimension) is the sum of event type vector (as basic event information) and event attribute encoding vector (as the description of event feature).

4.2 Hierarchical Representations with Attention Mechanisms

Based on the sequential event groups, the model can learn hierarchical representations to capture long-term temporal dependencies. In the low-level model, the model automatically identifies critical clinical events in each group via event attention mechanism and aggregates the events to form event group representations. In the high-level model, the meaningful long-term temporal dependencies of clinical event groups are captured by a recurrent neural network with temporal attention mechanism in the sequence representation. The hierarchical representations help to learn long-term temporal dependencies in the original event sequences.

4.2.1 Event Group Representation

To select the significant events in each group and compact events in the same group into one vector as the event group representation, the event attention mechanism are added in the low-level model.

Given sequential groups produced by the adaptive segmentation module $\{G_i\}_{i=1}^T$, where $G_i = \{v_1^i, \dots, v_{n_i}^i\}$, attention score of each event in the group is calculated by the event attention mechanism. The scalars α_t^i are the event attention weights that govern the influence of event embeddings $v_1^i, \dots, v_{n_i}^i$ in the group G_i .

We use a multi-layer perceptron (MLP) with one hidden layer to generate α_t^i based on the event embedding vector v_t^i and the hidden state of the previous time as follows:

$$q_t^i = \mathbf{w}_q \times \tanh(W_e \times \mathbf{v}_t^i + W_h \times \mathbf{h}_{i-1} + \mathbf{b}_h)$$

where $1 \leq i \leq T, 1 \leq t \leq n_i$.

$$\alpha_1^i, \dots, \alpha_{n_i}^i = \text{softmax}(q_1^i, \dots, q_{n_i}^i) \text{ for } 1 \leq i \leq T$$

where $\mathbf{h}_{i-1} \in \mathbb{R}^S$ is the hidden state of the previous gated recurrent units (GRU),²⁹ (which will be described in the following section) q_t^i is the latent layer of the event e_t at group i and $W_e \in \mathbb{R}^{H \times N}, W_h \in \mathbb{R}^{H \times S}, b_h$ and $w_q \in \mathbb{R}^H$, are parameters to learn. Notice that H is the hidden layer dimension and S is the GRU hidden state dimension.

The resulting attention scores α reflect the importance of each event in a group according to the temporal context of the group. Events in the i -th group are weighted averaged with α^i to get the group representation $\mathbf{g}_i \in \mathbb{R}^N$ as the input to the i -th GRU unit.

$$\mathbf{g}_i = \sum_{t=1}^{n_i} \alpha_t^i \times \mathbf{v}_t^i$$

4.2.2 Sequence Representation

To spot the critical phases over the sequence for the final decision and capture long-term temporal dependency of event groups, gated recurrent units (GRU)²⁹ equipped with temporal attention mechanism is employed as the high-level model.

$$\mathbf{h}_i = GRU(\mathbf{h}_{i-1}, \mathbf{g}_i, \theta)$$

where the function $GRU(\cdot)$ represents the recurrent unit, which use the previous hidden state \mathbf{h}_{i-1} and current input vector \mathbf{g}_i to update the hidden state. And θ represents all the parameters of GRU.

The vector $\beta = (\beta_1, \dots, \beta_T)$ contains the temporal attention weights of each group in the sequence. And we use a fully connected feedforward network to generate β from the output of GRU at each time as follow:

$$\beta = \text{softmax}(\mathbf{w}_{temporal} \times O_T)$$

where $O_T = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_T) \in \mathbb{R}^{S \times T}$ is the output matrix and $\mathbf{w}_{temporal} \in \mathbb{R}^S$ is a vector of parameters to learn.

The sequence representation \mathbf{s} is the weighted average of the output matrix O_T . We use \mathbf{s} to predict the true label \hat{y}_T of the sequences.

$$\mathbf{s} = O_T \times \beta$$

$$y_T = \text{sigmoid}(\mathbf{w}_p \times \mathbf{s} + b_p)$$

where $\mathbf{w}_p \in \mathbb{R}^S$ and b_p are parameters to learn.

The cross-entropy loss function is used to calculate the classification loss of each sample as follows:

$$\text{Loss}(\{e_t\}, \hat{y}_T) = \hat{y}_T \times \ln y_T + (1 - \hat{y}_T) \times \ln(1 - y_T)$$

where $\{e_t\}$ is the input event sequence and \hat{y}_T is the label indicating whether the clinical outcome happens. And we can sum up the losses of all the samples in one mini-batch to get the total loss for back propagation.

5 Results and Discussions

5.1 Comparison Methods and Settings

We compare our model with popular models in the literature, which include *bag-of-words vector classifiers* (i.e. support vector machine³⁰ (SVM), logistic regression³¹ (LR), random forest³² (RF)) and *deep sequential models*, such as RETAIN²³ and RNN¹² (implemented with GRU).

Due to the fact that SVM, LR, RF cannot handle a sequence input, the event sequence is compressed into a 0-1 vector in which the i -th element indicates whether the i -th event happens, and then fed into SVM, LR, or RF to make the outcome prediction. We implement these bag-of-words vector classifiers using scikit-learn (<https://scikit-learn.org>).

Models	Death		ICU admission	
	AUC	AUPRC	AUC	AUPRC
SVM	0.7523	0.5154	0.7973	0.7074
LR	0.8843	0.5213	0.8734	0.7266
Random Forest	0.8644	0.5867	0.8389	0.8177
RETAIN	0.8967	0.5808	0.8693	0.8029
RNN	0.9038	0.6234	0.8636	0.8051
Proposed Model	0.9428	0.7476	0.9016	0.8424

Table 3: performance of different models on death and ICU admission prediction tasks

Deep sequential models (i.e. RETAIN²³ and RNN¹²) take the original event sequence as their inputs as described in section 4.1. We implemented our model and neural network based baselines with Keras (<https://keras.io>).

The event embedding size is set to 32 while the hidden layer size is set to 64. The max number of groups M is set to 32. When training the models, we used Adam³³ with the mini-batch of 32 samples and the “early stopping” strategy when the performance of validation set drops down.

5.2 Evaluation Metrics

Metrics for binary labels such as accuracy are not suitable for measuring the performance on imbalanced datasets. Therefore, similar to the works,^{18,23} we adopt ROC curves (Receiver Operating Characteristic curves) and PRC (Precision-Recall curves) for evaluation metrics. Both of these two curves reflect the overall quality of predicted scores, according to their true labels. To get quantitative measurements, the area under ROC(AUC) and the area under PRC(AUPRC) are utilized.

5.3 Quantitative Results

Table 3 shows the AUC and AUPRC of different models on the death prediction and the ICU admission prediction tasks. From the results shown in Table 3, we can draw the following conclusions:

First, on the whole, deep sequential models(including RETAIN, RNN and the proposed model) outperform non-sequence models(including SVM, LR, and Random Forest) on both tasks, which suggests that temporal information is effective in the outcome prediction tasks.

Second, our model outperforms all the sequential models. For example, on the “ICU admission task”, the proposed model improves AUC by at least 3.4% and AUPRC by at least 3.0% compared other models on all tasks. The improvement verifies our claim that it’s more proper to capture temporal dependencies of clinical event sequences in a hierarchical way.

5.4 Ablation Studies

In this section, we perform ablation studies to examine the effects of our proposed techniques, namely the event attention mechanism, the temporal attention mechanism, and the adaptive segmentation module.

The first study over two attention mechanisms is performed on both tasks. Specifically, we re-train our model by ablating certain components:

- **W/O E-Attn**, where no event attention is performed and a group representation is set as the average of the event embeddings in this group.
- **W/O T-Attn**, where no temporal attention is performed and the sequence representation is set as the final output of the GRU.

Results of attention mechanism ablation studies are represented in Table 4. We can see that both attention mechanisms

Models	Death		ICU admission	
	AUC	AUPRC	AUC	AUPRC
W/O T-Attn	0.9348	0.7181	0.8987	0.8400
W/O E-Attn	0.9170	0.6404	0.8930	0.8376
Full Model	0.9428	0.7476	0.9016	0.8424

Table 4: Ablation study over attention mechanisms

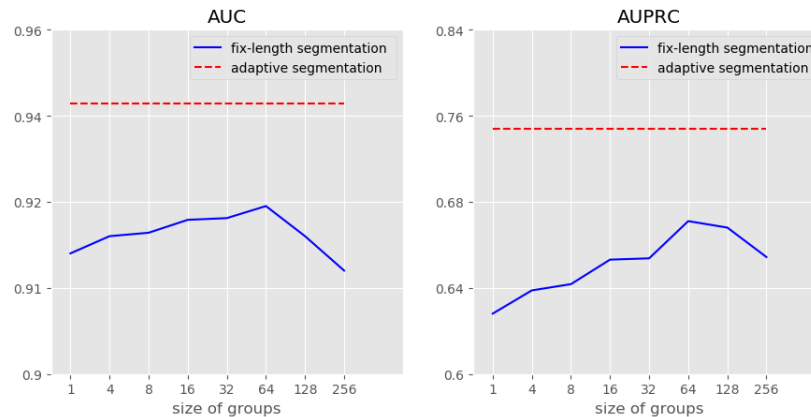


Figure 2: Ablation over the adaptive segmentation on the death prediction task. The fix-length segmentation splits the event sequence into groups of equal size (except the last one).

contribute to the strong empirical results of our model represented previously. It is noteworthy that the event attention, one of the important part of hierarchical representations, plays a more critical role in our model compared to the temporal attention, especially on the death prediction task.

Besides the attention mechanisms, study over the adaptive segmentation is performed on the death prediction task. We re-train our models by replacing the adaptive segmentation module with the fix-length segmentation which splits the original sequence into groups of equal size events (except the last group). Group size of the fix-length segmentation a hyperparameter. Notice that the fix-length segmentation degenerates to no segmentation if group size is set to 1.

Figure 2 shows the AUC and AUPRC of the proposed model where the adaptive segmentation is replaced by the fixed-length segmentation with different group sizes in the death prediction task (the trend in the ICU admission task is similar). We can see that the performance goes down when the group size becomes too small or too large. We infer that if the number is too small, local dependencies of events are modeled as long-term dependencies. And if the number is too large, long-term dependency is lost when the corresponding events are assigned to the same group. Besides, it’s obvious there is a performance gap between the adaptive segmentation and all other segmentation methods, which verifies our claim that the adaptive segmentation can help model long-term dependencies and is suitable for long irregular event sequences.

5.5 Important Events

We analyze the events to which our proposed model pays most attention in prediction. In particular, we use the median of all event attention scores of an event type on a specific task as the importance of the event type on the task.

Top important events on two tasks are listed in Table 5. We can see that even though our model mainly focuses on laboratory tests (such as “Heart Rhythm” and “Blood PH”) on both tasks, the specific events attracting the model on two tasks are different due to their different prediction targets. It is also perhaps surprising that owing to our data-driven approach, our model can select “Family Communication” as an important event type on the death prediction task, which may be ignored by doctors.

Top Events (Sorted by Median of Event Attention Scores)			
Death		ICU admission	
event	attn score	event	atten score
Blood Products	0.9965	Blood PH	0.9998
Radiologic Study: thoracic lumbar sac	0.9896	Vancomycin	0.9995
NV#2 Waveform Appear: overshoot	0.9713	Hematocrit (35-51)	0.9967
Heart Rhythm	0.9702	Edimentation rate	0.9885
Pain Location: periumbilical	0.9668	Daily Weight	0.9850
Family Communication	0.9523	Bilirubin Total	0.9834

Table 5: Top important events on the death prediction task and the ICU-admission prediction task.

6 Conclusion

In this paper, we proposed a model to learn hierarchical representations of long and irregular clinical event sequences of EHR data for clinical outcome prediction. We validate the performance of our model on real clinical datasets for death and ICU admission prediction tasks. The significant improvements indicated that our model is suitable for irregular timed EHR data and can capture long-term temporal dependencies of clinical event sequences for precise clinical outcome predictions.

Acknowledgements

This paper is partially supported by Beijing Municipal Commission of Science and Technology under Grant No. Z181100008918005, National Key Research and Development Program of China with Grant No. SQ2018AAA010010, and the National Natural Science Foundation of China (NSFC Grant No. 61772039 and No. 91646202).

References

- [1] Edward Choi, Andy Schuetz, Walter F Stewart, and Jimeng Sun. “Using recurrent neural network models for early detection of heart failure onset”. In: *Journal of the American Medical Informatics Association* (2016), ocw112.
- [2] Edward Choi, Mohammad Taha Bahadori, Elizabeth Searles, et al. “Multi-layer Representation Learning for Medical Concepts”. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016, pp. 1495–1504.
- [3] Alistair EW Johnson, Tom J Pollard, Lu Shen, et al. “MIMIC-III, a freely accessible critical care database”. In: *Scientific data* 3 (2016).
- [4] Trang Pham, Truyen Tran, Dinh Phung, and Svetha Venkatesh. “Deepcare: A deep dynamic memory model for predictive medicine”. In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer. 2016, pp. 30–41.
- [5] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. “Recurrent neural networks for multivariate time series with missing values”. In: *arXiv preprint arXiv:1606.01865* (2016).
- [6] Phuoc Nguyen, Truyen Tran, Nilmini Wickramasinghe, and Svetha Venkatesh. “Deep: A convolutional net for medical records”. In: *IEEE Journal of Biomedical and Health Informatics* (2016).
- [7] Isak Karlsson and Henrik Bostrom. “Predicting Adverse Drug Events Using Heterogeneous Event Sequences”. In: *IEEE International Conference on Healthcare Informatics*. 2016, pp. 356–362.
- [8] Zachary C Lipton, David C Kale, and Randall Wetzal. “Directly Modeling Missing Data in Sequences with RNNs: Improved Classification of Clinical Time Series”. In: *arXiv preprint arXiv:1606.04130* (2016).
- [9] Inci M Baytas, Cao Xiao, Xi Zhang, Fei Wang, Anil K Jain, and Jiayu Zhou. “Patient subtyping via time-aware LSTM networks”. In: *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. ACM. 2017, pp. 65–74.
- [10] Phuoc Nguyen, Truyen Tran, Nilmini Wickramasinghe, and Svetha Venkatesh. “Deep: a convolutional net for medical records”. In: *IEEE journal of biomedical and health informatics* 21.1 (2016), pp. 22–30.
- [11] Marzyeh Ghassemi, Marco AF Pimentel, Tristan Naumann, et al. “A multivariate timeseries modeling approach to severity of illness assessment and forecasting in icu with sparse, heterogeneous clinical data”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Conference on Artificial Intelligence. Vol. 2015. NIH Public Access. 2015, p. 446.

- [12] Chonho Lee, Zhaojing Luo, Kaiping Zheng, Gang Chen, and Beng Chin Ooi. “Big Healthcare Data Analytics: Challenges and Applications”. In: ().
- [13] Shang Gao, Michael T Young, John X Qiu, et al. “Hierarchical attention networks for information extraction from cancer pathology reports”. In: *Journal of the American Medical Informatics Association* 25.3 (2017), pp. 321–330.
- [14] Haoran Shi, Pengtao Xie, Zhiting Hu, Ming Zhang, and Eric P Xing. “Towards automated ICD coding using deep learning”. In: *arXiv preprint arXiv:1711.04075* (2017).
- [15] Hayit Greenspan, Bram Van Ginneken, and Ronald M Summers. “Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique”. In: *IEEE Transactions on Medical Imaging* 35.5 (2016), pp. 1153–1159.
- [16] Yuan Li, Xiaodan Liang, Zhiting Hu, and Eric P Xing. “Hybrid retrieval-generation reinforced agent for medical image report generation”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 1530–1540.
- [17] Christy Y Li, Xiaodan Liang, Zhiting Hu, and Eric P Xing. “Knowledge-driven encode, retrieve, paraphrase for medical image report generation”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 2019.
- [18] Chuanren Liu, Fei Wang, Jianying Hu, and Hui Xiong. “Temporal Phenotyping from Longitudinal Electronic Health Records: A Graph Based Framework”. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2015, pp. 705–714.
- [19] A Henriksson, Jing Zhao, H Bostrom, and H Dalianis. “Modeling electronic health records in ensembles of semantic spaces for adverse drug event detection”. In: *IEEE International Conference on Bioinformatics and Biomedicine*. 2015, pp. 343–350.
- [20] Vijay Huddar, Bapu Koundinya Desiraju, Vaibhav Rajan, Sakyajit Bhattacharya, Shourya Roy, and Chandan K Reddy. “Predicting Complications in Critical Care Using Heterogeneous Clinical Data”. In: *IEEE Access* 4 (2016), pp. 7988–8001.
- [21] Jinsung Yoon, EDU Ahmed M Alaa, Scott Hu, and Mihaela van der Schaar. “ForecastICU: A Prognostic Decision Support System for Timely Prediction of Intensive Care Unit Admission”. In: *Proceedings of The 33rd International Conference on Machine Learning*. 2016, pp. 1680–1689.
- [22] Karla L Caballero Barajas and Ram Akella. “Dynamically modeling patient’s health state from electronic medical records: A time series approach”. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 2015, pp. 69–78.
- [23] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. “RE-TAIN: An Interpretable Predictive Model for Healthcare using Reverse Time Attention Mechanism”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 3504–3512.
- [24] Marzyeh Ghassemi, Tristan Naumann, Finale Doshi-Velez, et al. “Unfolding physiological state: Mortality modelling in intensive care units”. In: *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM. 2014, pp. 75–84.
- [25] Riccardo Miotto, Li Li, Brian A Kidd, and Joel T Dudley. “Deep patient: an unsupervised representation to predict the future of patients from the electronic health records”. In: *Scientific reports* 6 (2016), p. 26094.
- [26] Peng Dai, Femida Gwadry-Sridhar, Michael Bauer, and Michael Borrie. “Bagging ensembles for the diagnosis and prognostication of alzheimer’s disease”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. AAAI Press. 2016, pp. 3944–3950.
- [27] Jiayu Zhou, Fei Wang, Jianying Hu, and Jieping Ye. “From micro to macro: data driven phenotyping by densification of longitudinal electronic medical records”. In: *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2014, pp. 135–144.
- [28] Luchen Liu, Jianhao Shen, Ming Zhang, Zichang Wang, and Jian Tang. “Learning the Joint Representation of Heterogeneous Temporal Events for Clinical Endpoint Prediction”. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. 2018.
- [29] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. “Empirical evaluation of gated recurrent neural networks on sequence modeling”. In: *arXiv preprint arXiv:1412.3555* (2014).
- [30] Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [31] A Yalcin, S Reis, AC Aydinoglu, and T Yomralioglu. “A GIS-based comparative study of frequency ratio, analytical hierarchy process, bivariate statistics and logistics regression methods for landslide susceptibility mapping in Trabzon, NE Turkey”. In: *Catena* 85.3 (2011), pp. 274–287.
- [32] Andy Liaw, Matthew Wiener, et al. “Classification and regression by randomForest”. In: *R news* 2.3 (2002), pp. 18–22.
- [33] Diederik Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization”. In: *Computer Science* (2014).