

# Learning Representations of Missing Data for Predicting Patient Outcomes

Brandon Malone, Alberto García-Durán, and Mathias Niepert

## Abstract

Extracting actionable insight from Electronic Health Records (EHRs) poses several challenges for traditional machine learning approaches. Patients are often missing data relative to each other; the data comes in a variety of modalities, such as multivariate time series, free text, and categorical demographic information; important relationships among patients can be difficult to detect; and many others. In this work, we propose a novel approach to address these first three challenges using a representation learning scheme based on message passing. We show that our proposed approach is competitive with or outperforms the state of the art for predicting in-hospital mortality (binary classification), the length of hospital visits (regression) and the discharge destination (multi-class classification).

## 1 Introduction

Healthcare is an integral service in modern societies. Improving its quality and efficiency has proven to be difficult and costly. Electronic health records (EHRs) provide a wealth of information carrying the potential to improve treatment quality and patient outcomes. Extracting useful and actionable medical insights from EHRs, however, poses several challenges both to traditional statistical and machine learning techniques.

First, not all patients receive the same set of laboratory tests, examinations, consultations, *etc.*, while they are at the hospital. Thus, many patients have *missing and incomplete data*, relative to other patients. Second, the various medical conditions and the corresponding treatment activities yield different kinds of data. For example, a blood oxygen saturation sensor may collect numeric values for a given amount of time at fixed frequency, while a consultation with a physician may produce only free text notes of the physician’s interpretation. That is, there are multiple *modalities* (or *attribute types*) of data, and these modalities have variations due to a number of external factors. Third, patients may share important relationships which are not easily captured in typical data representations. For example, family members often share a similar genetic background. Likewise, patients with similar initial diagnoses may share underlying characteristics which are difficult to capture in traditional

models. Thus, it is important to explicitly model *relationships* among patients that capture some form of disease or treatment affinity.

Recent work on representation learning for graph-structured data has specifically addressed the last two problems albeit in different domains. In that line of work, edges in an *affinity graph* connect instances deemed similar in some way, and a message passing scheme is used to learn representations of the data modalities associated with each instance. While several message-passing approaches to graph representation learning have been proposed (Kipf and Welling 2017; Hamilton, Ying, and Leskovec 2017; Gilmer et al. 2017), we base our work on embedding propagation (EP), a method that is both unsupervised and learns a representation specific to each data modality (García-Durán and Niepert 2017).

EP has several characteristics that make it well-suited for the problem of learning from medical records. First, it provides a method for learning and combining representations, or *embeddings*, of the various data modalities typically found in EHRs. For instance, it is possible to simultaneously learn embeddings of free-text notes and to combine these with embeddings learned for sensor time series data. Second, as we will propose in this work, due to its unsupervised reconstruction loss, EP allows us to learn a vector representation for every data modality and every patient, even if that particular data modality is not observed at all for some of the patients. Third, EP learns representations for each data modality independently; thus, it is possible to distinguish between the influence of these independently learned modality representations on the predictions.

Intuitively, embedding propagation learns data representations such that representations of nodes close in the graph, that is, nodes similar in one or several ways, are more similar to each other than those of nodes far away in the graph. The representations can then be used in traditional downstream machine learning tasks. In the context of EHRs, patients are modeled as nodes in the graph<sup>1</sup>, and similarity relationships between patients are modeled with edges.

In this work, we extend the embedding propagation

<sup>1</sup>Specifically, patient episodes in an intensive care unit (ICU) are modeled as nodes; for ease of exposition, we use “patient”, “episode” and “node” interchangeably in this work. The meaning is clear from the context.

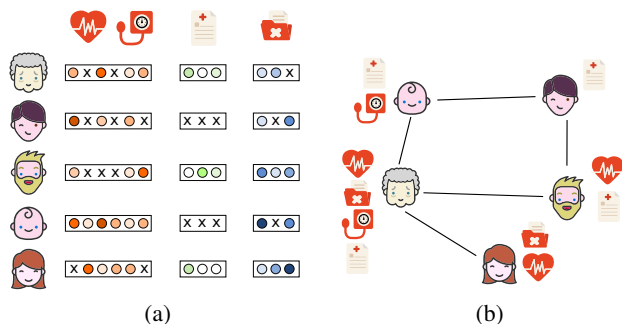


Figure 1: (a) Illustration of patients and the corresponding time-series, text, and demographic data. Missing data is indicated with an x. (b) A patient affinity graph constructed with a subset of the patient attributes. The graph structure is leveraged in the EP framework.

framework to account for missing data. In particular, for each data modality, we learn *two* representations for each patient; the first representation carries information about the observed data, while the second representation is learned to account for missing data. Learning an explicit representation of missing data within the graph-based learning framework has the advantage of propagating representations of missing data throughout the graph such that these representations are also informed by representations of neighboring patients. Combining these learned feature representations gives a complete representation for downstream tasks.

We use a recently-introduced benchmark based on the MIMIC-III dataset (Johnson et al. 2016) to evaluate the proposed approach and show that it is competitive with the state of the art when using a single data modality, numeric time series observations. After augmenting the data with additional data modalities, including free text from physicians’ notes and demographic information, we outperform the state of the art on two of the three tasks we consider in this work, namely, length of hospital stay and discharge destination prediction.

The rest of this paper is structured as follows. In Section 2, we describe the data that we use in this work. Section 3 introduces EP and describes our novel contribution for handling missingness within it. We describe our experimental design and empirical results in Section 4. Section 5 places our contribution in context with related work, while Section 6 concludes the paper.

## 2 Data

In this work, we use the MIMIC-III (Johnson et al. 2016) (Medical Information Mart for Intensive Care III) benchmark datasets created by (Harutyunyan et al. 2017) which contain information about patient intensive care unit (ICU) visits, or *episodes*. However, these benchmarks include only time series data. In order to demonstrate the efficacy of our proposed approach in a multimodal context, we augment the benchmarks to include textual data and (categorical) demographic data. This section describes our dataset as well as the specific tasks we consider.

## Adding Additional Modalities

The original benchmark dataset includes 17 time series variables (see Appendix A in the Supplementary Material), sampled at varying frequencies and many of which are missing for some patients (please see (Harutyunyan et al. 2017) for more details). Further, each admission includes a SUBJECT\_ID and EPISODE which link it back to the complete MIMIC-III dataset. We use these foreign keys to augment the time series variables with text (from the NOTEVENTS table) and demographic (from the ADMISSIONS and PATIENTS tables) data. In both cases, we ensure to only use data available at the time of prediction.<sup>2</sup> For example, we never use “Discharge Summaries” to predict mortality.

## Preprocessing

We extracted features from the time series data as previously described (Harutyunyan et al. 2017). Briefly, we create seven sets of observations from each time series (the entire sequence, and the first and last 10%, 25% and 50% segments of the sequence). We then calculate nine features for each segment (count, minimum, maximum, mean, standard deviation, skew, kurtosis, median, and max absolute deviation). Thus, in total, for the time series data we have  $7 \times 9 = 63$  features.<sup>3</sup> Finally, we standardize all observations such that each feature has a mean of 0 and a variance of 1 for the observed values across all patients in the training set.

The notes were first partitioned based on their CATEGORY (see Appendix B) into 6 note types: nursing, radiology, respiratory, ecg, echo and other. As mentioned above, we never use the discharge summaries in this work. We then convert each note into a bag of words representation; we discard all words which appear in less than 0.1% or more than 90% of the notes. The concatenated notes for each note type are used as the text observations for each patient. That is, each patient has 6 bag-of-words text features, one for each type.

Further, we extract admission and demographic information about each episode from the respective tables. In particular, we collect the ADMISSION\_TYPE (“urgent”, “elective”, etc.), ADMISSION\_LOCATION (“emergency room”, etc.), DIAGNOSIS (“a preliminary, free text diagnosis ... [which is] usually assigned by the admitting clinician and does not use a systematic ontology”<sup>4</sup> about the admission. Each of these three fields contains text. We also collect the patient’s ethnicity, gender, age, insurance, and marital status.

Figure 1 gives an example of the dataset after preprocessing and graph construction (see Section 4 for details about graph construction).

<sup>2</sup>We will make the scripts used to create the dataset publicly-available.

<sup>3</sup>This is slightly higher than the prior work because we also include median and max absolute deviation features since they are robust against outliers.

<sup>4</sup><https://mimic.physionet.org/mimictables/admissions/>

## Patient Outcome Labels

In this work, we consider three prediction tasks: in-hospital mortality (MORT), length of stay (LOS), and discharge destination (DD). In all cases, we make predictions 48 hours after admission to the ICU. We adopt the same semantics for the starting time of an episode as (Harutyunyan et al. 2017).

For MORT, this is exactly the same problem in the original benchmarks. For LOS, prediction at 48 has previously been used (Silva et al. 2012).

Several medical studies have considered the DD task (Mauthe et al. 1996; Brauer et al. 2008; Wee, Bagg, and Palepu 1999); however, these approaches have relied on traditional clinical scores such as the FIM or Berg balance score. To the best of our knowledge, the DD task has not been extensively considered in the machine learning literature before. In contrast to the binary MORT classification task, this is a multiclass classification problem. In particular, the set of patients in the MIMIC-III in this study have 6 discharge destinations (after grouping; see Appendix C). Conceptually, we consider DD as a more fine-grained proxy for the eventual patient outcome compared to MORT.

## 3 EP for Learning Patient Representations

The general learning framework of EP (García-Durán and Niepert 2017) proceeds in two separate steps. We use the term *attribute* to refer to a particular node feature. Moreover, an *attribute type* (or *data modality*) is either one node attribute or a set of node attributes grouped together.

In a first step, EP learns a vector representation for every attribute type by passing messages along the edges of the affinity graph. In the context of learning from EHRs, for instance, one attribute type consists of time series data recorded for ICU patients. The attribute types used in this work are: (1) a group of time-series data from ICU measurements; (2) words of free-text notes; and (3) a group of categorical features of demographic data. As described in Section 2, we have several attributes of each type.

In a second step, EP learns patient representations by combining the previous learned attribute representations. For instance, the method combines the representations learned for the physicians’ notes with the representations learned for the time series data. These combined representations are then used in the downstream prediction tasks.

### Attribute Type Representation Learning

Each attribute type  $i$  is associated with a domain  $\mathcal{D}_i$  of possible values. For instance, an attribute type consisting of  $N$  numerical attributes has the domain  $\mathbb{R}^N$ , while an attribute type modeling text data has the domain  $\mathbb{R}^{|T|}$  where  $T$  is the vocabulary. For each attribute type we choose a suitable encoding function  $\mathbf{f}$ . This encoding function  $\mathbf{f}$  is parameterized and maps every  $\mathbf{x} \in \mathcal{D}_i$  to its vector representation  $\mathbf{x}' \in \mathbb{R}^{d_i}$ , that is,  $\mathbf{x}' = \mathbf{f}_i(\mathbf{x})$ . These encoding functions have to be differentiable to update their parameters during learning. Concretely, in that work, a single dense layer is used as the encoding function for each label type. For the time series-derived numeric attribute, this amounts to a standard matrix multiplication, where as for the text and categor-

ical attribute types, this amounts to an embedding lookup table.

The functions  $\mathbf{l}_i : V \rightarrow 2^{\mathcal{D}_i}$  map every vertex  $v$  in the graph to a (possibly empty) set of vectors  $\mathbf{x} \in \mathcal{D}_i$ . We write  $\mathbf{l}(v) = \bigcup_i \mathbf{l}_i(v)$  for the set of all vectors of attribute type  $i$  associated with vertex  $v$ . Moreover, we write  $\mathbf{l}_i(\mathbf{N}(v)) = \{\mathbf{l}_i(u) \mid u \in \mathbf{N}(v)\}$  for the multiset of vectors of attribute type  $i$  associated with the neighbors of vertex  $v$ . EP learns a vector representation for each attribute type of the problem.

- We write  $\mathbf{h}_i(v)$  to denote the current vector representation of attribute type  $i$  for node  $v$ . It is computed as follows:

$$\mathbf{h}_i(v) = \mathbf{g}_i(\{\mathbf{f}_i(\mathbf{x}) \mid \mathbf{x} \in \mathbf{l}_i(v)\}). \quad (1)$$

- We write  $\tilde{\mathbf{h}}_i(v)$  to denote the reconstruction of the representation of attribute type  $i$  for node  $v$ .  $\tilde{\mathbf{h}}_i(v)$  is computed from the attribute type representations of  $v$ ’s neighbors in the graph. It is computed as follows:

$$\tilde{\mathbf{h}}_i(v) = \tilde{\mathbf{g}}_i(\{\mathbf{f}_i(\mathbf{x}) \mid \mathbf{x} \in \mathbf{l}_i(\mathbf{N}(v))\}), \quad (2)$$

where  $\mathbf{g}_i$  and  $\tilde{\mathbf{g}}_i$  are aggregation functions that map a multiset of  $d_i$ -dimensional embeddings to a single  $d_i$ -dimensional embedding. These aggregation functions can be parameterized but are often parameter-free aggregation functions such as the elementwise average or maximum.

The core idea of EP is to make the attribute type representation and its reconstruction similar for each attribute type and each node in the graph. In other words, EP learns attribute type representations such that the distance between  $\mathbf{h}_i(v)$  and  $\tilde{\mathbf{h}}_i(v)$  is small. More formally, for all attribute types EP minimizes the following loss

$$\mathcal{L}_i = \sum_{v \in V} \sum_{u \in V \setminus \{v\}} \left[ \gamma + d_i(\tilde{\mathbf{h}}_i(v), \mathbf{h}_i(v)) - d_i(\tilde{\mathbf{h}}_i(v), \mathbf{h}_i(u)) \right]_+ \quad (3)$$

where  $d_i$  is the Euclidean distance,  $[x]_+$  is the positive part of  $x$ , and  $\gamma > 0$  is a margin hyperparameter.

In practice, it is not feasible to evaluate the inner summation; thus, it is approximated by sampling a single random node  $u$  which is different than  $v$ .

The margin-based loss defined in Equation 3 updates the parameters (i.e., embedding functions and functions  $\mathbf{g}_i$  and  $\tilde{\mathbf{g}}_i$  in case they are chosen to be parametric) if the distance between  $\mathbf{h}_i(v)$  and  $\tilde{\mathbf{h}}_i(v)$  plus a margin is not smaller than the distance between  $\mathbf{h}_i(u)$  and  $\tilde{\mathbf{h}}_i(v)$ . Intuitively, for each patient node  $v$  and attribute type  $i$ , the vector representation reconstructed from the embeddings of patient nodes neighboring  $v$  are learned to be more similar to the embedding of attribute type  $i$  for  $v$  than to the embedding of attribute type  $i$  of a random patient node in the graph.

The generic working of the attribute type representation learning stage is as follows: in each propagation step, for each node  $v$  in the graph, randomly sample a node  $u$  uniformly from the set of all nodes in the graph; then, for each attribute type  $i$ , the embeddings  $\tilde{\mathbf{h}}_i(v)$ ,  $\mathbf{h}_i(v)$  and  $\mathbf{h}_i(u)$  are computed. Finally, parameters of the model are updated based on the loss defined in Equation (3).

The functions  $g_i(\mathcal{H})$  and  $\tilde{g}_i(\mathcal{H})$  are given as

$$g_i(\mathcal{H}) = \tilde{g}_i(\mathcal{H}) = \frac{1}{|\mathcal{H}|} \sum_{\mathbf{h} \in \mathcal{H}} \mathbf{h}$$

for attribute types  $i$  and sets of embedding vectors  $\mathcal{H}$ .

### Learning Representations for Missing Data

Two characteristics of EP make it suitable for the problem of learning representations for missing data. First, it supports an arbitrary number of attribute types, and one can learn missing data representations tailored to attribute types. Second, EP’s learning principle is based on reconstructing each node’s attribute representation from neighboring nodes’ attribute representations, and this makes it highly suitable for settings wherein a number of nodes have missing data. During training, EP learns how to reconstruct the missing data representation based on a contrastive loss between representations of existing attribute types and, therefore, can learn how to reconstruct a representation when data is missing.

For every attribute type  $i$  and for every node  $v \in V = \{1, \dots, |V|\}$  of the graph we have an input feature vector  $\mathbf{x}$ . Based on this vector, we create two feature vectors. The first feature vector  $\mathbf{x}_1$  is identical to  $\mathbf{x}$  except that all missing values are masked out. The second feature vector  $\mathbf{x}_2 \in \mathbb{R}^{|V|}$  is either (1) all zeros if there are no missing attribute values for attribute type  $i$  or (2) all zeros except for the position  $v$  which is set to 1. The vector  $\mathbf{x}_2$  indicates whether data is missing and is used to learn a latent representation for nodes with missing data. These two vectors are then fed into two encoding functions  $\mathbf{f}_i^1$  and  $\mathbf{f}_i^2$  whose parameters are learned independently from each other. For each input, the output vectors of the encoding functions are added element-wise, and the resulting vector is used in the margin-based contrastive loss function. Therefore, for every attribute type  $i$  we have  $\mathbf{x}' = \mathbf{f}_i(\mathbf{x}) = \mathbf{f}_i^1(\mathbf{x}_1) + \mathbf{f}_i^2(\mathbf{x}_2) = \mathbf{x}_1^T \mathbf{W}_i^1 + \mathbf{x}_2^T \mathbf{W}_i^2$ . The  $\mathbf{f}_i^2$  encoding functions amount to embedding lookup tables. We refer to this instance of EP as EP-MD (for “EP-missing data”).

Figure 2 illustrates our approach to learning separate representations for observed and missing attribute labels. Since the contrastive loss compares representations that are computed both based on missing and observed data, the two representations influence each other. If for some node  $v$  and attribute type  $i$  we have missing data, the representation of that missing data is influenced by the representations of observed and missing data of attribute type  $i$  by neighboring nodes. Thus, the missing data representations are also propagated during the message passing.

As mentioned in Section 2, we standardize features to have a mean of 0;  $\mathbf{x}_2$  allows EP-MD to distinguish between observed 0s and missing values which have been masked.

**Missing data mechanisms** Missing data is often assumed to result from mechanisms such as Missing (Completely) at Random (MAR and MCAR) or Not Missing at Random (NMAR) (Rubin 1976). In contrast to methods like multiple imputation by chained equations (Azur et al. 2011), EP-MD does not estimate (representations of) missing attributes

based on other observed attributes; rather, it leverages the affinity graph to learn representations based on the same attribute observed for similar nodes. Thus, EP-MD does not assume an explicit missingness mechanism. We leave exploration of the statistical missingness properties of EP-MD to future work.

### Generating Patient Representations

Once the learning of attribute type representations has finished, EP computes a vector representation for each patient node  $v$  from the vector representations of  $v$ ’s attribute types. In this work, the patient representation is the concatenation of the attribute type representations:

$$\mathbf{v} = \text{concat} [\mathbf{h}_1(v), \dots, \mathbf{h}_k(v)]. \quad (4)$$

Since we have modeled missing data explicitly, the latent representations  $\mathbf{h}_i(v)$  exist for every node in the graph and every attribute type  $i$ . We also evaluate a patient representation built from *both* the raw features and the learned representations:

$$\mathbf{v} = \text{concat} [\mathbf{h}_1(v), \dots, \mathbf{h}_k(v), \mathbf{x}_1(v), \dots, \mathbf{x}_k(v)], \quad (5)$$

where  $\mathbf{x}_i(v)$  is the raw input for attribute type  $i$  and node  $v$ .

## 4 Experiments

We first describe our experimental design. We then present and discuss our results on the three patient outcome prediction tasks.

**Hyperparameters and computing environment** For EP-MD, we use 200 iterations of training with mini-batches of size 256. We embed all attributes in a 32-dimensional space. We set the margin  $\gamma = 5$ , and we solved the EP-MD optimization problem using ADAM with a learning rate of  $1e-3$ . EP-MD was implemented in KERAS using TENSORFLOW as the backend. All experiments were run on computers with 128GB RAM, four quad-core 2.8 GHz CPU, and a TitanX GPU. The computing cluster was in a shared environment, so we do not report running times in this evaluation. In all cases, though, the running times were modest, ranging from a few minutes up to about an hour for learning embeddings for all episodes for a single attribute type.

**Baseline and downstream models** Prior work has shown that linear models perform admirably compared to much more sophisticated methods on these and similar benchmark problems (Lipton et al. 2016; Harutyunyan et al. 2017; Barnes et al. 2016). Thus, we use (multi-class) logistic regression as a base model for MORT and DD and ridge regression as the base model for LOS. Both implementations are from SCIKIT-LEARN (Pedregosa et al. 2011).

We leave a more comprehensive comparison to complex models such as long short-term memory networks (LSTMs) as future work. However, we do include comparison to previously-published performance of an LSTM (Harutyunyan et al. 2017) on MORT. For a baseline, missing values are replaced with the mean of the respective feature. This completed data is then used to train the base model. We compare

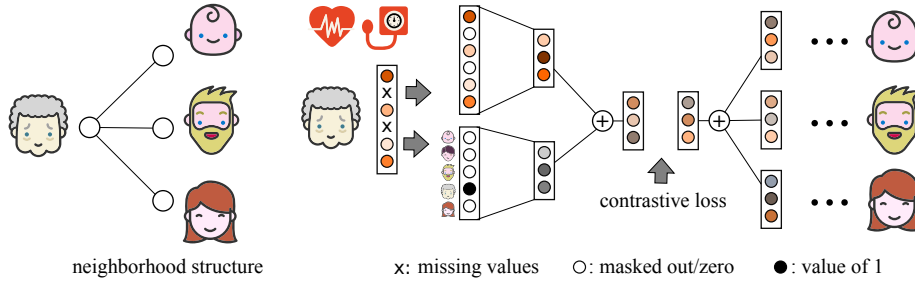


Figure 2: Illustration of the embedding propagation framework for missing data.

the baselines to learning embeddings with EP-MD followed by the same supervised learning approach (either logistic or ridge regression). We select all hyperparameters based on cross-validation (please see Appendix D for more details). EP has been shown to outperform methods like DEEP-WALK (Perozzi, Al-Rfou, and Skiena 2014) and graph convolutional networks (Kipf and Welling 2017); thus, we do not compare to those methods in this work.

**Construction of the affinity graph** An integral component of EP-MD is the affinity graph; it gives a mechanism to capture domain knowledge about the similarity of the nodes to which traditional machine learning approaches do not have access. In this work, since the goal is always to make a prediction about the outcome of an episode, each node in the graph corresponds to an episode.

As described in Section 2, we extract the admission type, location, and initial diagnosis about each episode. We base the construction of our graph on the text from these fields. In particular, we concatenate the three fields to form a “document” for each episode. Next, we use `fastText`<sup>5</sup> to learn custom word embeddings via skipgrams from this corpus. We then use these to calculate a single embedding (a “sentence vector”; please see the software documentation for more details) for each episode. We then define the similarity between two episodes  $i$  and  $j$  as  $s_{i,j} = \exp\{-d_{i,j}\}$ , where  $d_{i,j}$  is the Euclidean distance between the respective embeddings. Finally, we connect all pairs of episodes for which  $s_{i,j} > 0.9$ .

**Data partitioning** (Harutyunyan et al. 2017) created a standard training and testing set split of episodes. The training set includes 17 869 episodes; the testing set includes 3 233 episodes. While modest in comparison to some datasets considered in the machine learning community, this dataset is huge compared to many datasets presented in the medical literature. Thus, we also consider the ability to generalize to the test set when outcome information is available for only smaller numbers of training episodes. In particular, we draw subsets of varying size (see Figure 3) and only observe the class label for those episodes. As described in Section 3, EP-MD can still take advantage of the unlabeled

episodes. On the other hand, completely supervised methods can only use the episodes with class labels. We note that our predictions are always made 48 hours after admission. We ensure that our learning only considers information that would be available in such a realistic scenario.

**Evaluation metrics** We use standard metrics for evaluation. We use the area under the receiver operating characteristic curve (AUROC) to evaluate MORT, mean absolute error (MAE) to evaluate LOS, and a multi-class generalization of AUROC (MC-AUROC) (Hand and Till 2001) for DD.

## Results

We first consider the impact of learning embeddings on only the time series data present in the original benchmarks for use in downstream prediction. Figure 3(top) compares the performance of the RAW baseline (using just the hand-crafted features) compared to first learning representations of the time series attribute type using EP-MD (EMBEDDED). We also include a third strategy (COMBINED) in which we combine the original features with the representations learned by EP-MD as described in Section 3. (Please see Appendix E for detailed tables of all results.)

As the figure shows, the embeddings are particularly helpful in small-data scenarios for MORT and LOS. Prior work in the graph-embedding community has observed similar phenomena, so our findings are in line with that existing work. The embeddings are particularly well-suited for LOS. Indeed, this is the only task in this single data modality setting in which COMBINED meaningfully differs from RAW.

We observed that RAW exhibits particularly high variance in performance for small subsample sizes for LOS, while the variance for EMBEDDED are in line with the other settings. This case particularly highlights the ability of EP-MD to reduce noise and improve performance in missing data scenarios, even when only a single modality is available.

These results also show that both COMBINED and RAW match state-of-the-art performance by LSTMs reported in prior work (Harutyunyan et al. 2017). Further, the LSTMs were trained in a completely supervised manner; thus, their performance is expected to deteriorate when using smaller training sets (Ramachandran, Liu, and Le 2017).

<sup>5</sup><https://github.com/facebookresearch/fastText>

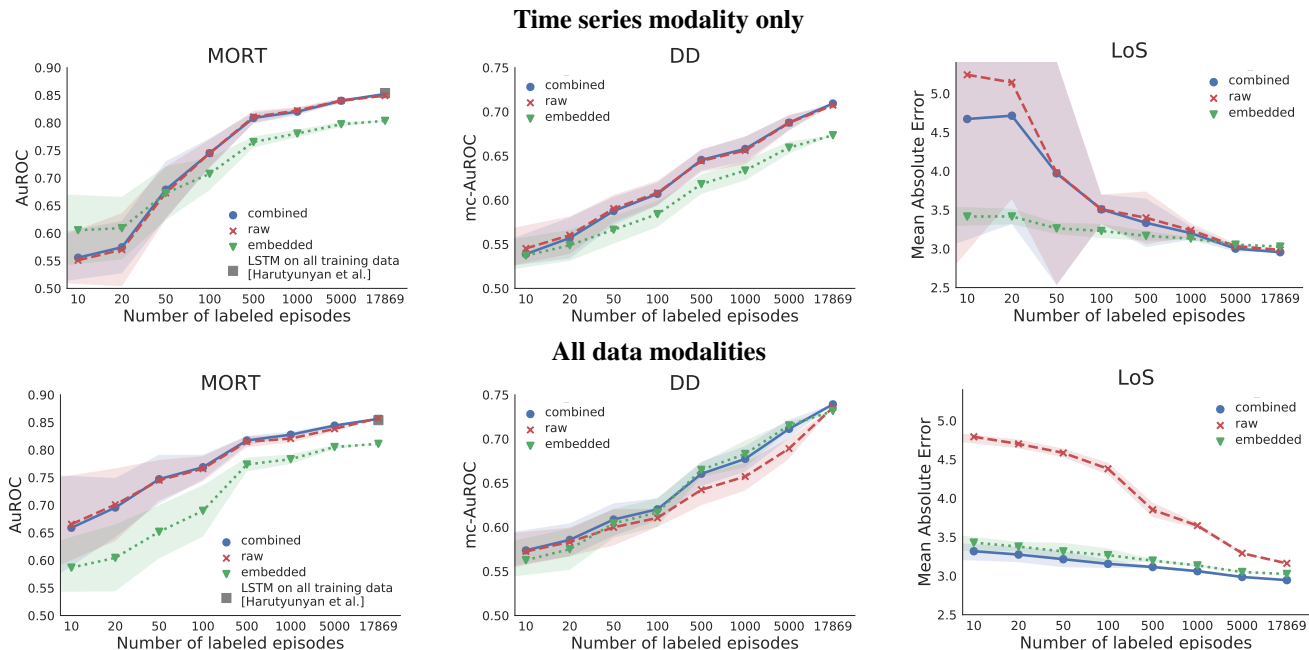


Figure 3: The respective performance of using only the “raw” features, only the “embedded” episode representations, or the “combined” features for training a logistic (for MORT and DD) or ridge (for LoS) regression model using either the single time series modality (top) or all four data modalities (bottom) as input. The x-axis gives the number of labeled training episodes (out of a total of 17 869 episodes in the full training set). 20 random samples of each subset size were evaluated for each representation. The markers give the mean performance on the entire test set, and the shaded regions show one standard deviation above and below the mean. For MORT, we also indicate the performance of the best single-task LSTM used in (Harutyunyan et al. 2017) Please note that *higher AuROC values are better* (for MORT and DD), while *lower MAE values are better* (for LoS).

**Multi-modal data** We then extend our analysis to the multi-attribute case. As described in Section 3, we use EP-MD to learn embeddings for each of the four attribute types (time series features, text notes, demographics, and episode identity within the graph) independently. The final representation for the episode is the concatenation of all attribute type representations. For RAW, we used standard text (term-frequency inverse document-frequency) and categorical (one-hot encoding) variable preprocessing. We construct the graph for EP-MD using additional information about the episode admission (see Section 2). We also encode these data as categorical variables and include them for RAW.

Figure 3(bottom) shows the performance after adding the additional attributes. Relative to their performance in the single-attribute setting, the performance for all three representations improves across nearly all tasks and observed episode subsets. For MORT, as in the single-attribute case, EMBEDDED is clearly worse than using only the hand-crafted features; however, combining the embeddings with the raw features does no harm, and COMBINED and RAW perform virtually identically. Both mildly surpass the state-of-the-art results previously reported (Harutyunyan et al. 2017). The COMBINED representation benefits significantly from the additional data attributes on the LoS task. Indeed, it outperforms both RAW and EMBEDDED across all sub-sample sizes. In this task, COMBINED effectively uses information from *both* RAW and EMBEDDED, rather than simply learning to use the best, as seems the case in other settings.

Finally, all of the representations show significant im-

provement in the DD task when the additional modalities are available. Both EMBEDDED and COMBINED significantly outperform RAW for several of the subset sizes.

**Statistical significance** Across all of the sample sizes, EMBEDDED is significantly outperformed by RAW (paired sample t-test, Benjamini-Hochberg multiple-test corrected  $p$ -value  $< 0.05$ ) and COMBINED on the DD task in the single-modality settings. Both are also significantly better than EMBEDDED on MORT in the multiple modality setting. On the other hand, in the multi-modal setting for LoS, EMBEDDED is significantly better than RAW, and COMBINED is significantly better than both. COMBINED is also significantly better than RAW in the multi-modal case on DD.

**Impact of missing data** We hypothesize that EMBEDDED outperforms RAW for instances with many missing values. Figure 4(left) confirms that, on the LoS task, for patients with few missing values, EMBEDDED and RAW perform similarly. EMBEDDED outperforms RAW the most in cases with 5 to 10 missing values ( $\approx 20 - 40\%$ ); EMBEDDED continues to result in more accurate predictions for larger numbers of missing values, but the difference is not as stark.

We next consider the effect on prediction quality when specific features are missing. As shown in Figure 4(right), predictions based on EMBEDDED for LoS are similar regardless of whether most features are observed or missing; this suggests that EP-MD effectively propagates useful information about missing features throughout the graph.

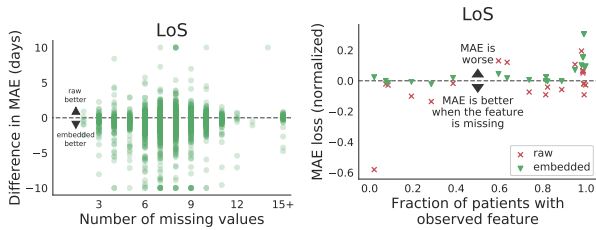


Figure 4: (left) Difference in the mean absolute error on all test instances when using 100 labeled samples for training. Negative values indicate that EMBEDDED produced better predictions than RAW. (right) Comparison of the MAE when a feature is observed for a test instance compared to when it is missing. The x-axis gives the percentage of patients for which a specific feature was observed; the y-axis shows the difference between the MAE for instances when the feature is observed and when it is not. Each marker corresponds to one feature. A negative value indicates that predictions are more accurate for patients where the feature is missing.

On the other hand, RAW features lead to *worse* predictions when most of the less-frequent features are observed. This helps explain the better performance of EMBEDDED compared to RAW on LOS; EMBEDDED-based predictions effectively use all of the available information, while models trained on RAW are actually misled by rare attributes.

For RAW, missing temperature led to the worst change in performance, while features related to the Glasgow coma scale had the largest negative effect for EMBEDDED. RAW results in much better predictions when the capillary refill rate is missing (the bottom left marker in the plot). All of the MAE losses are given in Appendix F.

**Feature importance** As previously described, EP-MD allows us to identify the original features which most influence predictions. For LOS, we consider the average contribution to the final prediction across all test instances (based on the learned model coefficients and embeddings for each instance) as the importance for each feature. The plots in Appendix G show that temperature and pH are the most influential features for EMBEDDED-based predictions; the Glasgow coma scale features are the next most influential. Interestingly, the contributions of temperature and pH are similar regardless of whether they are observed. This again suggests that EP-MD effectively propagates useful information about missing features. On the other hand, the Glasgow coma scale features contribute much less when they are not observed.

## 5 Related Work

We review some of the most related work on missing data and representation learning for clinical datasets.

**Missing data in clinical datasets** A large body of recent research has addressed the problem of predicting patient outcomes from multivariate time series. The availability of large, high-quality datasets like MIMIC-III (Johnson et al. 2016) and those made available through the Computing in Cardiology challenges (Silva et al. 2012) has surely

spurred some of this work. This data entails missingness almost by design since different vital signs are typically observed at different frequencies. Methodologically, much of this work (Rajkomar et al. 2018; Harutyunyan et al. 2017; Lipton et al. 2016; Lipton, Kale, and Wetzel 2016; Che et al. 2015; Suresh, Szolovits, and Ghassemi 2016; Hammerla et al. 2015) borrows from and extends recent developments in the deep learning community, such as long short-term memory (LSTM) networks (Hochreiter and Schmidhuber 1997). Some Bayesian methods (Lasko, Denny, and Levy 2013; Marlin et al. 2012; Lehman et al. 2008) have also been developed. While our approach builds upon this work, such as by extending previously-described hand-crafted features (Harutyunyan et al. 2017; Lipton et al. 2016), the use of multimodal data is clearly novel in our work compared to this body. Further, our novel message passing scheme for missing data distinguishes our contribution from this work.

**Representation learning and clinical data** A second approach to predicting patient outcomes from EHRs has instead focused on learning representations from discrete observations, such as prescribed medications and International Classification of Diseases (ICD) codes assigned during billing. Much of this work (Choi et al. 2017; 2016a; 2016b) builds on the recent successes of deep learning to include domain-specific background knowledge. Some recent work (Choi, Chiu, and Sontag 2016) has also considered text modalities, though they use word co-occurrence matrices.

Our use of a graph to explicitly capture similarity relationships among patients distinguishes our work from these prior approaches. Additionally, these methods typically do not consider the impact of missing data. In contrast, we explicitly learn a representation of the missing observations.

## 6 Discussion

We have presented an approach for learning representations for ICU episodes from multimodal data, which could come from EHRs. Our approach extends an existing representation learning framework, embedding propagation (García-Durán and Niepert 2017), to handle missing data. Empirically, we demonstrate that our approach outperforms a state-of-the-art baseline in predicting hospital visit duration and discharge destination when multiple attribute types are available; when combined with hand-crafted features, our approach is competitive with long short-term memory networks for in-hospital mortality prediction.

Future developments of our approach include extending both the considered tasks and data included in the model. We did not consider the “computational phenotyping” task (Lasko, Denny, and Levy 2013; Che et al. 2015) in this work. This is typically formulated as a multi-label classification problem; the representations learned with EP-MD can be combined with standard models for that setting. Currently, our model does not incorporate medications, procedures and treatments that occur during an episode; clearly, these have bearing on the patient outcomes. Since such data is available in MIMIC-III, we plan to integrate appropriate attribute types for these data.

## References

- Azur, M. J.; Stuart, E. A.; Frangakis, C.; and Leaf, P. J. 2011. Multiple imputation by chained equations: what is it and how does it work? *International Journal of Methods in Psychiatric Research* 20(1):40–49.
- Barnes, S.; Hamrock, E.; Toerper, M.; Siddiqui, S.; and Levin, S. 2016. Real-time prediction of inpatient length of stay for discharge prioritization. *Journal of the American Medical Informatics Association* 23(e1):e2–e10.
- Brauer, S. G.; Bew, P. G.; Kuys, S. S.; Lynch, M. R.; and Morrison, G. 2008. Prediction of discharge destination after stroke using the motor assessment scale on admission: A prospective, multisite study. *Archives of Physical Medicine and Rehabilitation* 89(6):1061–1065.
- Che, Z.; Kale, D.; Li, W.; Bahadori, M. T.; and Liu, Y. 2015. Deep computational phenotyping. In *Proc. KDD*.
- Choi, E.; Bahadori, M. T.; Schuetz, A.; Stewart, W. F.; and Sun, J. 2016a. Doctor AI: Predicting clinical events via recurrent neural networks. In *Proc. MLHC*.
- Choi, E.; Bahadori, M. T.; Sun, J.; Kulas, J.; Schuetz, A.; and Stewart, W. 2016b. RETAIN: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Proc. NIPS*.
- Choi, E.; Schuetz, A.; Stewart, W. F.; and Sun, J. 2017. Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association* 24(2):361–370.
- Choi, Y.; Chiu, C. Y.-I.; and Sontag, D. 2016. Learning low-dimensional representations of medical concepts. In *AMIA Summits on Translational Science Proceedings*, 41–50.
- García-Durán, A., and Niepert, M. 2017. Learning graph representations with embedding propagation. In *Proc. NIPS*.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *Proc. ICML*.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Proc. NIPS*.
- Hammerla, N. Y.; Fisher, J.; Andras, P.; Rochester, L.; Walker, R.; and Ploetz, T. 2015. PD disease state assessment in naturalistic environments using deep learning. In *Proc. AAAI*.
- Hand, D., and Till, R. 2001. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning* 45(2):171–186.
- Harutyunyan, H.; Khachatrian, H.; Kale, D. C.; and Galstyan, A. 2017. Multitask learning and benchmarking with clinical time series data. arXiv:1703.07771 [stat.ML].
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.
- Johnson, A. E.; Pollard, T. J.; Shen, L.; Wei, H.; Lehman, L.; Feng, M.; Ghassemi, M.; Moody, B.; Szolovits, P.; Celi, L. A.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data* 3(160035).
- Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *Proc. ICLR*.
- Lasko, T. A.; Denny, J. C.; and Levy, M. A. 2013. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. *PLoS ONE* 8(6):e66341.
- Lehman, L. H.; Saeed, M.; Moody, G.; and Mark, R. 2008. Similarity-based searching in multi-parameter time series databases. *Computers in Cardiology* 35:653–656.
- Lipton, Z. C.; Kale, D. C.; Elkan, C.; and Wetzel, R. 2016. Learning to diagnose with LSTM recurrent neural networks. In *Proc. ICLR*.
- Lipton, Z. C.; Kale, D. C.; and Wetzel, R. 2016. Modeling missing data in clinical time series with RNNs. In *Proc. MLHC*.
- Marlin, B. M.; Kale, D. C.; Khemani, R. G.; and Wetzel, R. C. 2012. Unsupervised pattern discovery in electronic health care data using probabilistic clustering models. In *Proceedings of the 2<sup>nd</sup> ACM SIGHIT International Health Informatics Symposium*.
- Mauthe, R. W.; Haaf, D. C.; Hayn, P.; and Krall, J. M. 1996. Predictors for patient discharge destination after elective anterior cervical discectomy and fusion. *Archives of Physical Medicine and Rehabilitation* 77(1):10–13.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. DeepWalk: Online learning of social representations. In *Proc. KDD*.
- Rajkomar, A.; Oren, E.; Chen, K.; Dai, A. M.; Hajaj, N.; Liu, P. J.; Liu, X.; Sun, M.; Sundberg, P.; Yee, H.; Zhang, K.; Duggan, G. E.; Flores, G.; Hardt, M.; Irvine, J.; Le, Q.; Litsch, K.; Marcus, J.; Mossin, A.; Tansuwan, J.; Wang, D.; Wexler, J.; Wilson, J.; Ludwig, D.; Volchenboun, S. L.; Chou, K.; Pearson, M.; Madabushi, S.; Shah, N. H.; Butte, A. J.; Howell, M.; Cui, C.; Corrado, G.; and Dean, J. 2018. Scalable and accurate deep learning for electronic health records. *npj Digital Medicine* 1(18).
- Ramachandran, P.; Liu, P. J.; and Le, Q. V. 2017. Unsupervised pretraining for sequence to sequence learning. In *Proc. EMNLP*.
- Rubin, D. B. 1976. Inference and missing data. *Biometrika* 63(3):581–592.
- Silva, I.; Moody, G.; Scott, D. J.; Celi, L. A.; and Mark, R. G. 2012. Predicting in-hospital mortality of patients in ICU: The PhysioNet/computing in cardiology challenge 2012. *Computing in Cardiology* 39:245–248.
- Suresh, H.; Szolovits, P.; and Ghassemi, M. 2016. The use of autoencoders for discovering patient phenotypes. In *NIPS Workshop on Machine Learning for Health*.
- Wee, J. Y.; Bagg, S. D.; and Palepu, A. 1999. The Berg balance scale as a predictor of length of stay and discharge destination in an acute stroke rehabilitation setting. *Archives of Physical Medicine and Rehabilitation* 80:448–452.



# Supplement: Learning Representations of Missing Data for Predicting Patient Outcomes

Brandon Malone, Alberto García-Durán, and Mathias Niepert

arXiv:1811.04752v1 [cs.LG] 12 Nov 2018

## Appendix A: Time series data

- Capillary refill rate
- Diastolic blood pressure
- Fraction inspired oxygen
- Glasgow coma scale eye opening
- Glasgow coma scale motor response
- Glasgow coma scale total
- Glasgow coma scale verbal response
- Glucose
- Heart Rate
- Height
- Mean blood pressure
- Oxygen saturation
- Respiratory rate
- Systolic blood pressure
- Temperature
- Weight
- pH

## Appendix B: Note category assignments

Table 1 provides the mapping between the CATEGORY field in the NOTEEVENTS table in MIMIC-III and the 6 types used in this work.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Table 1: Note type mapping

CATEGORY	Label
Case Management	NOTE OTHER BOW
Consult	NOTE OTHER BOW
Discharge summary	NOTE DISCHARGE SUMMARY BOW
ECG	NOTE ECG BOW
Echo	NOTE ECHO BOW
General	NOTE NURSING BOW
Nursing	NOTE NURSING BOW
Nursing/other	NOTE NURSING BOW
Nutrition	NOTE OTHER BOW
Pharmacy	NOTE OTHER BOW
Physician	NOTE NURSING BOW
Radiology	NOTE RADIOLOGY BOW
Rehab Services	NOTE OTHER BOW
Respiratory	NOTE RESPIRATORY BOW
Social Work	NOTE OTHER BOW

## Appendix C: Discharge destinations

Table 2 gives the mapping between the DISCHARGE\_LOCATION field in the ADMISSIONS table in MIMIC-III and the 6 types used in this work.

Table 2: Discharge location mapping

DISCHARGE_LOCATION	DD label
DEAD/EXPIRED	MORTALITY_INHOSPITAL
HOME	HOME
HOME HEALTH CARE	HOME
HOME WITH HOME IV PROVIDR	HOME
HOSPICE-HOME	HOME
REHAB/DISTINCT PART HOSP	REHAB
SNF	SNF
SNF-MEDICAID ONLY CERTIF	SNF
LEFT AGAINST MEDICAL ADVI	LEFT
LONG TERM CARE HOSPITAL	TRANSFER
DISCH-TRAN TO PSYCH HOSP	TRANSFER
DISC-TRAN CANCER/CHLDRN H	TRANSFER
OTHER FACILITY	TRANSFER
SHORT TERM HOSPITAL	TRANSFER
HOSPICE-MEDICAL FACILITY	TRANSFER
ICF	TRANSFER
DISC-TRAN TO FEDERAL HC	TRANSFER

## Appendix D: Hyperparameter selection

For the results in Section 4 (Experiments) of the main paper, we use an inner three-fold cross-validation on the training set to select the hyperparameters for all linear models. We never use the testing set during training. In particular, the search grids we use are given in Table 3.

Table 3: Hyperparameter search grid

Logistic regression (MORT and DD)	
Hyperparameter	Values
class weight	balanced, no change
regularization type	l1, l2
regularization parameter	1.0, 0.1, 0.01, 0.001, 0.0001

Ridge regression (LOS)	
Hyperparameter	Values
regularization parameter	$10^{-6}, 10^{-5} \dots 10^{+7}$

Please consult the SCIKIT-LEARN documentation for more details about the hyperparameters.

## Appendix E: Result tables

Tables 4 – 9 give the values shown in Figure 3 of the main text.

## Appendix F: Feature MAE loss

Table 10 gives the normalized MAE loss shown in Figure 4(right) of the main text.

## Appendix G: Feature importance

Figure 1 shows the average contribution of each feature to the final prediction for LOS when training on 100 labeled instances. Specifically, the contribution for each feature on a particular instance is calculated as the sum of the element-wise product of the embedding for the feature with the appropriate coefficients from the learned linear model; we additionally take the absolute value to consider only the magnitude of the contribution.

Figure 2 shows the ratio of the average contributions for each feature when it is observed compared to when it is missing. Larger values indicate that the feature contributes more when it is observed than when it is missing. For example, the embedded Glasgow coma scale motor response feature contributes, on average, about 6 times more to the final prediction when it is observed than when it is missing. In contrast, temperature and pH contribute roughly the same regardless of whether they are observed or not.




Figure 1: The average contribution of each feature to final predictions across all test instances on LOS when trained with 100 labeled instances.




Figure 2: The ratio of the average contribution of each feature to final predictions across all test instances when the feature is observed compared to when it is missing. The dashed line indicates where the ratio is 1, which means the feature contributes equally whether it is observed or not. Higher values indicate that the feature contributes more when it is observed compared to when it is missing.

Table 4: MORT, time series attribute type only

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$0.556 \pm 0.044$	$0.605 \pm 0.065$	$0.551 \pm 0.041$
20	$0.575 \pm 0.047$	$0.609 \pm 0.057$	$0.570 \pm 0.066$
50	$0.679 \pm 0.051$	$0.673 \pm 0.049$	$0.673 \pm 0.048$
100	$0.745 \pm 0.027$	$0.708 \pm 0.029$	$0.745 \pm 0.025$
500	$0.809 \pm 0.010$	$0.765 \pm 0.010$	$0.811 \pm 0.011$
1 000	$0.820 \pm 0.007$	$0.781 \pm 0.007$	$0.822 \pm 0.006$
5 000	$0.840 \pm 0.003$	$0.797 \pm 0.005$	$0.840 \pm 0.005$
17 869	$0.852 \pm 0.000$	$0.804 \pm 0.000$	$0.849 \pm 0.000$

Table 5: DD, time series attribute type only

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$0.539 \pm 0.014$	$0.537 \pm 0.017$	$0.545 \pm 0.021$
20	$0.557 \pm 0.023$	$0.549 \pm 0.018$	$0.560 \pm 0.021$
50	$0.588 \pm 0.016$	$0.566 \pm 0.016$	$0.590 \pm 0.015$
100	$0.607 \pm 0.013$	$0.584 \pm 0.015$	$0.608 \pm 0.014$
500	$0.645 \pm 0.012$	$0.618 \pm 0.011$	$0.644 \pm 0.013$
1 000	$0.658 \pm 0.014$	$0.634 \pm 0.011$	$0.656 \pm 0.016$
5 000	$0.688 \pm 0.009$	$0.659 \pm 0.007$	$0.688 \pm 0.009$
17 869	$0.709 \pm 0.000$	$0.673 \pm 0.000$	$0.707 \pm 0.000$

Table 6: LOS, time series attribute type only

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$4.672 \pm 1.657$	$3.416 \pm 0.129$	$5.242 \pm 2.612$
20	$4.714 \pm 1.396$	$3.416 \pm 0.100$	$5.143 \pm 1.503$
50	$3.972 \pm 1.443$	$3.263 \pm 0.073$	$3.978 \pm 1.443$
100	$3.508 \pm 0.190$	$3.230 \pm 0.091$	$3.509 \pm 0.190$
500	$3.334 \pm 0.313$	$3.166 \pm 0.064$	$3.399 \pm 0.339$
1 000	$3.202 \pm 0.088$	$3.131 \pm 0.017$	$3.239 \pm 0.089$
5 000	$3.001 \pm 0.038$	$3.053 \pm 0.006$	$3.034 \pm 0.036$
17 869	$2.956 \pm 0.000$	$3.028 \pm 0.000$	$2.985 \pm 0.000$

Table 7: MORT, All attribute types

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$0.659 \pm 0.096$	$0.587 \pm 0.044$	$0.665 \pm 0.082$
20	$0.696 \pm 0.053$	$0.605 \pm 0.061$	$0.701 \pm 0.066$
50	$0.747 \pm 0.044$	$0.652 \pm 0.048$	$0.745 \pm 0.037$
100	$0.769 \pm 0.023$	$0.690 \pm 0.046$	$0.766 \pm 0.023$
500	$0.817 \pm 0.008$	$0.774 \pm 0.012$	$0.815 \pm 0.010$
1 000	$0.828 \pm 0.005$	$0.783 \pm 0.009$	$0.821 \pm 0.007$
5 000	$0.844 \pm 0.003$	$0.805 \pm 0.004$	$0.838 \pm 0.005$
17 869	$0.856 \pm 0.000$	$0.811 \pm 0.000$	$0.858 \pm 0.000$

Table 8: DD, All attribute types

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$0.574 \pm 0.019$	$0.563 \pm 0.020$	$0.573 \pm 0.020$
20	$0.586 \pm 0.019$	$0.575 \pm 0.023$	$0.583 \pm 0.016$
50	$0.609 \pm 0.018$	$0.604 \pm 0.016$	$0.600 \pm 0.020$
100	$0.620 \pm 0.012$	$0.617 \pm 0.016$	$0.611 \pm 0.010$
500	$0.660 \pm 0.014$	$0.665 \pm 0.008$	$0.642 \pm 0.018$
1 000	$0.677 \pm 0.014$	$0.683 \pm 0.015$	$0.657 \pm 0.016$
5 000	$0.711 \pm 0.011$	$0.715 \pm 0.005$	$0.689 \pm 0.012$
17 869	$0.739 \pm 0.000$	$0.732 \pm 0.000$	$0.736 \pm 0.000$

Table 9: LOS, All attribute types

Number of labeled episodes	COMBINED	EMBEDDED	RAW
10	$3.318 \pm 0.109$	$3.429 \pm 0.106$	$4.792 \pm 0.055$
20	$3.275 \pm 0.099$	$3.378 \pm 0.060$	$4.701 \pm 0.057$
50	$3.215 \pm 0.103$	$3.316 \pm 0.111$	$4.584 \pm 0.056$
100	$3.155 \pm 0.053$	$3.267 \pm 0.088$	$4.380 \pm 0.077$
500	$3.115 \pm 0.020$	$3.197 \pm 0.020$	$3.851 \pm 0.093$
1 000	$3.061 \pm 0.019$	$3.135 \pm 0.018$	$3.648 \pm 0.036$
5 000	$2.986 \pm 0.011$	$3.050 \pm 0.009$	$3.292 \pm 0.015$
17 869	$2.945 \pm 0.000$	$3.024 \pm 0.000$	$3.162 \pm 0.000$

Table 10: MAE loss shown in Figure 4(right) of the main text.

Attribute	Embedded MAE loss	Raw MAE loss
Capillary refill rate	-0.023301	0.580816
Diastolic blood pressure	-0.156282	-0.064640
ETHNICITY	0.002419	0.057705
Fraction inspired oxygen	0.022698	0.136004
Glasgow coma scale eye opening	-0.307864	0.019429
Glasgow coma scale motor response	-0.304350	0.020988
Glasgow coma scale total	-0.044169	-0.130732
Glasgow coma scale verbal response	-0.307864	0.019429
Glucose	-0.094687	0.027558
Heart Rate	-0.156282	-0.064640
Height	0.006909	0.100878
MARITAL STATUS	-0.070432	-0.106718
Mean blood pressure	-0.156282	-0.064640
NOTE DISCHARGE SUMMARY BOW	0.233990	0.284654
NOTE ECG BOW	-0.025505	-0.021364
NOTE ECHO BOW	-0.018772	0.017316
NOTE NURSING BOW	-0.016769	-0.120556
NOTE OTHER BOW	0.001341	0.021422
NOTE RADIOLOGY BOW	0.004466	0.041344
NOTE RESPIRATORY BOW	0.014245	0.028457
Oxygen saturation	-0.092130	0.091678
Respiratory rate	-0.150736	-0.050517
Systolic blood pressure	-0.156282	-0.064640
Temperature	-0.101287	-0.194498
Weight	-0.005509	0.073845
pH	-0.001081	0.091238