

Mixture-based Multiple Imputation Model for Clinical Data with a Temporal Dimension

1st Ye Xue
Northwestern University
Evanston, USA
ye.xue@u.northwestern.edu

2nd Diego Klabjan
Northwestern University
Evanston, USA
d-klabjan@northwestern.edu

3rd Yuan Luo
Northwestern University
Chicago, USA
yuan.luo@northwestern.edu

Abstract—The problem of missing values in multivariable time series is a key challenge in many applications such as clinical data mining. Although many imputation methods show their effectiveness in many applications, few of them are designed to accommodate clinical multivariable time series. In this work, we propose a multiple imputation model that capture both cross-sectional information and temporal correlations. We integrate Gaussian processes with mixture models and introduce individualized mixing weights to handle the variance of predictive confidence of Gaussian process models. The proposed model is compared with several state-of-the-art imputation algorithms on both real-world and synthetic datasets. Experiments show that our best model can provide more accurate imputation than the benchmarks on all of our datasets.

Index Terms—Imputation, missing data, EHR, Gaussian process, data mining

I. INTRODUCTION

The computational modeling in clinical applications attracts growing interest with the realization that the quantitative understanding of patient pathophysiological progression is crucial to clinical studies [1]. With a comprehensive and precise modeling, we can have a better understanding of a patient’s state, offer more precise diagnosis and provide better individualized therapies [2]. Researchers are increasingly motivated to build more accurate computational models from multiple types of clinical data. However, missing values in clinical data challenge researchers using analytic techniques for modeling, as many of the techniques are designed for complete data.

Traditional strategies used in clinical studies to handle missing values include deleting records with missing values and imputing missing entries by mean values. However, deleting records with missing values and some other filtering strategies can introduce biases [3] that can impact modeling in many ways, thus limiting its generalizability. Mean imputation is widely used by researchers to handle missing values. However, it is shown to yield less effective estimates than many other modern imputation techniques [4]–[7], such as maximum likelihood approaches and multiple imputation methods (e.g. multiple imputation by chained equations (MICE) [8]). The maximum likelihood and multiple imputation methods are based on solid statistical foundations and has become standard in the last few decades [9], [10].

Multiple imputation is originally proposed by Rubin [11] and with the idea of replacing each missing value with multiple estimates that reflect variation within each imputation model. The classic multiple imputation method treats the complete variables as predictors and incomplete variables as outcomes, and uses multivariate regression models to estimate missing values [12], [13]. Another imputation approach relies on chained equations where a sequence of regression models is fitted by treating one incomplete variable as the outcome and other variables with possible best estimates for missing values as predictors [14]–[16]. Multiple imputation has its success in healthcare applications, where many imputation methods are based on multiple imputation [17]–[23].

In recent years, additional imputation methods are proposed. Although many imputation methods [7], [8], [14], [24]–[28] show their effectiveness in many applications, few of them are designed for time series-based clinical data. These clinical data are usually multivariable time series, where patients have measurements of multiple laboratory tests at different times. Many methods are designed for cross-sectional imputation (measurements taken at the same time point) and do not consider temporal information that is useful in making predictions or imputing missing values. Ignoring informative temporal correlations and only capturing cross-sectional information may yield less effective imputation.

In order to address the limitations mentioned above, we present a mixture-based multiple imputation model (MixMI) for clinical time series ¹. Our model captures both cross-sectional information and temporal correlations to estimate missing values using mixture models. We model the distribution of measurements using a mixture model. The mixture is composed of linear regression to model cross-sectional correlations and Gaussian processes (GPs) to capture temporal correlations. The problem of integrating GP within a standard mixture model is that GP models in all patient cases get the same mixing weights, while the confidence of predictions by GP models can vary largely across different patient cases. We overcome this problem by introducing individualized mixing weights for each patient cases, instead of assigning a fixed weight. We train our model using the Expectation-

This work is supported in part by NIH grant R21LM012618.

¹The implementations of our model is available at: <https://github.com/y-xue/MixMI>.

Maximization (EM) algorithm. We demonstrate the effectiveness of our model by comparing it with several state-of-the-art imputation algorithms on multiple clinical datasets.

Our main contributions are summarized as follows.

1. To the best of our knowledge, we are the first to build imputation model for time series by integrating GP within mixture models. We address the problem that all GP models in all patient cases get a fixed mixing weight by introducing individualized mixing weights.

2. We test the performance of our model on two real-world clinical datasets and several synthetic datasets. Our best model outperforms all comparison models including several state-of-the-art imputation models. Using synthetic datasets, we also explore and discover the properties of the data that benefit our model and/or comparison models. Experiments show that our best model is robust to the variation of these properties and outperforms comparison models on all synthetic datasets.

The remainder of this paper is structured as follows. Section II discusses related work while in Section III, the proposed method is described. The experimental setup, including dataset collection and evaluation procedure, is described in Section IV. Section V discusses the computational results and underlying analyses. The conclusions are drawn in Section VI.

II. RELATED WORK

Research in designing imputation methods for multivariable time series attracts growing interest in recent decades. Previous studies generally fall into two categories. One comes from methods using linear or other simple parametric functions to estimate missing values. The other is the methods treating time series as smooth curves and estimating missing values using GP or other nonparametric methods.

In the first category, multivariable time series are modeled based on either linear models [29], linear mixed models [30], [31] or autoregressive models [32], [33]. However, in these methods, the potential trajectories of variables are only limited to linear or other simple parametric functions. Alternatively, many authors choose GPs or other nonparametric functions to model time series. Compared to linear models, GPs only have locality constraints in which close time points in a time series usually have close values. Therefore, GPs bring in more flexibility in capturing temporal trajectories of variables.

However, directly applying GPs on imputation for multivariable time series usually yields less effective imputation due to the lack of considerations for the similarities and correlations across time series. Futoma et al. [34] extends the GP-based approach to imputation in the multivariable setting by applying the multi-task Gaussian Processes (MTGP) prediction [35]. However, the quality of estimating missing values relies on the estimation of covariance structure among variables when using MTGP or other multi-task functional approaches [36]–[38]. To make a confident estimation of the covariance, a large amount of time points with shared observations of these variables are often required by these multi-task approaches [35], [39]. Due to the fact that many patients only have records with a limited number of time points, time series of inpatient

clinical laboratory tests fall short of such a requirement. Hori et al. [40] extend MTGP to imputation on multi-environmental trial data in a third-order tensor. However, this approach is not applicable to clinical data with a large number of patients, since the covariance matrix of observed values needs to be computed together with its inverse, which is intractable for our datasets.

Recently, Luo et al. [41] explore the application of GPs in clinical data and propose an algorithm, 3-dimensional multiple imputation with chained equations (3D-MICE), that combines the imputation from GP and traditional MICE based on weighting equations. However, the weighting equations are calculated only based on the standard deviations of GP and MICE predictions for missing values. The weighting strategy is static and not optimized. We postulate that calculating weights through an optimization problem can help to improve the imputation quality. In our work, instead of the predictive mean matching used in [41], we choose linear regression as one component of our model. Our model is also based on a statistical model and thus statistically justified which is not the case for [41].

The mixture model proposed in this work differs from the mixtures of Gaussian processes (MGP) [42] in several ways. The mixing components in our model are not limited to Gaussian process, thus bringing in possibilities of capturing correlations with various choices of parametric functions. More importantly, our model is designed for multivariable time series and captures both temporal and cross-sectional correlations. The latter is ignored in [42] when MGP is applied to multivariable time series, which yields less effective predictions since some clinical tests have strong cross-sectional relations. Without a Gaussian process component, the standard Gaussian mixture model (GMM) is well studied and has been extended to imputation tasks [43]–[46]. These approaches still suffer from losing informative temporal correlations when imputing missing values in longitudinal data.

In order to effectively model the interaction between different aspects, we represent the data as a tensor with each aspect being one mode. Therefore, our method is also considered as a tensor completion approach. Tensor completion problems are extensively studied. However, the classic tensor completion methods [47]–[49] focus on general tensors and usually do not consider temporal aspects. In recent years, many studies explore the application of temporal augmented tensor completion on imputing missing values in time series [50]–[54]. These methods discretize time into evenly sampled intervals. However, due to the fact that inpatient clinical laboratory tests are usually measured at varying intervals, assembling clinical data over regularly sampled time periods might have several drawbacks, such as leading to sparse tensors if discretizing time at fine granularity (e.g. every minute) while some laboratory tests are measured less frequently (e.g. daily). Furthermore, extending these methods to the case, where time is not regularly sampled, is not easy and straightforward, requiring changing design details and the objective functions to be optimized. Recently, Yang et al. [55] propose a tensor com-

pletion method that can deal with irregularly sampled time. However, most components of this approach are tailored to the characteristics of septic patients. We implement this approach with only the component of time-aware mechanism that is general and applicable to our experimental settings. However, this approach is not so effective in our experiments, thus not included as a benchmark in this paper. Lately, Zhe et al. [56] propose a Bayesian nonparametric tensor decomposition model that captures temporal correlations between interacting events. However, this approach is not directly applicable to continuous multivariable time series because it focuses on discrete events and captures temporal correlations between the occurrences of events.

Recurrent Neural Networks (RNNs) are another choice to capture temporal relationships and have been studied in imputation tasks. Early studies [57]–[59] in RNN-based imputation models capture temporal correlations either within each time series or across all time series [60]. Yoon et al. [60] recently propose a novel multi-directional RNN that captures temporal correlations both within and across time series. Compared with traditional statistical models, neural network-based models are usually harder to train and provide limited explainability, while the performance is not always necessarily better. We use [60] as a benchmark since it is the most recent model. Besides imputation, researchers also attempt to build classification models with RNNs that utilize the missing data patterns [61]–[63]. Instead of relying on the estimation of missing values, these models perform predictions by exploring the missingness itself and no separate imputation is required. These models can not be a substitute for imputation methods since they do not provide imputed values.

III. METHODOLOGY

A. Imputation Framework

In many predictive tasks on temporal clinical data, time series are often aligned into the same-length sequences to derive more robust patient phenotypes through matrix decomposition or discover feature groups by applying sequence/graph mining techniques [55], [64]. We model under this assumption. In this work, we use tensor representation, in which patients have the same number of time points. We represent the data collected from P patients with V laboratory tests and B time points as two 3D tensors $\mathcal{X} \in \mathbb{R}^{P \times V \times B}$ and $\mathcal{T} \in \mathbb{R}^{P \times V \times B}$, shown in Fig. 1. Each laboratory test measurement $x_{p,v,b}$ is stored in the measurement tensor \mathcal{X} . Each time $t_{p,v,b}$, when $x_{p,v,b}$ is measured, is stored in the time tensor \mathcal{T} .

Table I lists the main symbols we use throughout the paper. Missing values in the measurement tensor are denoted as $x_{p,v,b}^{mis}$, showing that the value of test v at time index b for patient p is missing. Correspondingly, $x_{p,v,b}^{obs}$ denotes an observed value. The time tensor \mathcal{T} is complete, since we only collect patient records at the time when at least one laboratory test measurement is available. We assume we know the “prescribed” time a missing measurement should have been taken. In the matrix $x_{:,:,b}$ at time index b , the measurement time $t_{p,v,b}$ and $t_{q,v,b}$ can be different when $p \neq q$, whereas for

TABLE I
MAIN SYMBOLS AND DEFINITIONS

Symbol	Definition
\mathcal{X}, \mathcal{T}	Measurement and time tensor
$x_{p,v,-b}$	Measurements in fiber $x_{p,v,:}$, excluding $x_{p,v,b}$
$t_{p,v,-b}$	Times in fiber $t_{p,v,:}$, excluding $t_{p,v,b}$
$Mix_{v,b}$	Mixture model for v and b .
$V_{v,b}$	Concatenation of $x_{:, -v, b}$ and $x_{:, v, -b}$
$\mathcal{N}(\boldsymbol{\mu}_{v,b}^{(k)}, \boldsymbol{\Sigma}_{v,b}^{(k)})$	The k th prior multivariate normal distribution in $Mix_{v,b}$
$m^{(3)}(\cdot), \Sigma^{(3)}(\cdot)$	Predictive mean and variance of a GP model
$\gamma_{v,b}$	The set of all trainable parameters of $Mix_{v,b}$
I	The identity matrix

a given patient p , we have $t_{p,v,b} = t_{p,u,b}$ for $v, u \in [1 : V]$. That is, all tests for a particular patient are taken at the same time.

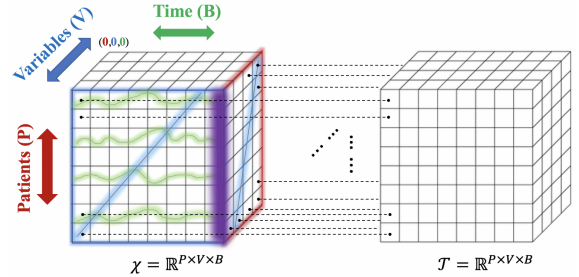


Fig. 1. Measurement and time tensor. An example of the inputs and output of the mixture model $Mix_{v,B}$ is shown as the colorful fiber and matrices. A fiber of the tensor \mathcal{X} is a vector obtained by fixing all but one of \mathcal{X} 's indices. In $Mix_{v,B}$, the target output is $x_{:,v,B}$ shown in purple and the inputs are $x_{:, -v, B}$ and $x_{:, v, -B}$ shown in red and blue matrices, respectively, excluding the purple fiber. We model the output with a mixture model, where we train a linear regression on the red matrix and train GPs or/and another linear regression on the blue matrix.

Disregarding the temporal dimension, the imputation problem is well studied. If one dimension is time, in order to apply imputation methods that are not designed for time series, we need to disregard the temporal aspect or ignore temporal correlations of the data. However, temporal trajectories can reveal patient's underlying pathophysiological evolution, the modeling of which can help to better estimate missing values. For the reason that both cross-sectional information and temporal aspects can impact the estimation of missing measurements, we explore mixture models, which are composed of several base models through either a cross-sectional or temporal view. We introduce these base models in Section III-B.

In our imputation framework, a mixture model is trained for each variable and time index. We use $Mix_{v,b}$ to denote the mixture model to impute missing values of variable v at time index b . The missing values $x_{:,v,b}^{mis}$ in the fiber $x_{:,v,b}$ are imputed by the optimized $Mix_{v,b}$. In each iteration of the algorithm, it is assumed that all other values are known and only $x_{:,v,b}^{mis}$ is imputed. This is wrapped in an outer loop. We call this procedure simple-pass tensor imputation, i.e., one pass through all v, b . Since several simple-pass tensor imputations

are conducted, our approach is also considered as an iterative imputation [65], which can also be regarded as a sampling-based approach where a Gibbs sampler is used to approximate convergence. The convergence of iterative imputation methods can be quite fast with a few iterations [8].

In detail, the iterative imputation approaches start by replacing all missing data with values from simple guesses; we fill in all missing values with initial estimates by taking random draws from observed values. This procedure is called an initial imputation. Then we perform iterative tensor imputation on each copy separately. The training procedure and imputation for fibers are introduced in Section III-C and III-D. We also rely on the concept of multiple imputations, where several iterative imputations are performed and the imputed values are averaged at the end. Each iterative imputation starts with a different iterative imputation tensor and/or uses a different order of v, b .

In summary, the algorithm creates M different copies $\mathcal{X}^1, \dots, \mathcal{X}^M$ of \mathcal{X} , each one filled with different random \mathcal{X}^{mis} . For each $i = 1, \dots, M$, we then perform K simple-pass tensor imputations. Each simple-pass has a loop over all v, b , which uses $Mix_{v,b}$ to adjust $\mathcal{X}^{i,mis}$. At the end of the M imputations, $\mathcal{X}^{i,mis}$ are averaged across all i to yield the final imputed tensor.

The whole imputation process involves $M \times K \times V \times B$ imputation models. We next first focus on the base models behind $Mix_{v,b}$ and then on the actual mixture model.

B. Base Models

Our mixture models are composed of three components that are derived from two base models, linear regression and Gaussian processes. One component consists of GP models and the other two components are linear models through two different views of the measurement tensor. Through a cross-sectional view, the tensor can be considered as a vector of patient-by-variable matrices at different time indices. Through a temporal view, we can view the tensor as a vector of patient-by-time matrices for different variables.

1) *Linear model through cross-sectional view*: We can view the measurement tensor \mathcal{X} as a vector of patient-by-variable matrices. On the slice $x_{:,v,b}$, we use a linear regression model to fit the target variable v as a function of the other variables except v . The target values $x_{:,v,b}$ are modeled as

$$x_{:,v,b} = x_{:,-v,b}\beta_{v,b}^{(1)} + \epsilon_{v,b}^{(1)}, \epsilon_{v,b}^{(1)} \sim \mathcal{N}(0, \sigma_{v,b}^{(1)2} I) \quad (1)$$

where $\beta_{v,b}^{(1)}$ is the column vector of coefficients and $\sigma_{v,b}^{(1)}$ is the standard deviation of the error $\epsilon_{v,b}^{(1)}$, regarding the regression model through cross-sectional view for variable v and time index b .

The likelihood distribution of $x_{:,v,b}$ is then given by

$$x_{:,v,b}|x_{:,-v,b}, \beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2} \sim \mathcal{N}(x_{:,-v,b}\beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2} I). \quad (2)$$

The training data consists of observed target values $(x_{p,v,b})_{p \in P_{v,b}^{tr}}$ and the input data $(x_{p,-v,b})_{p \in P_{v,b}^{tr}}$, where $P_{v,b}^{tr}$ is the training patient set and includes p only if $x_{p,v,b}$ is observed.

2) *Linear model through temporal view*: In addition to the cross-sectional view, we can also view the measurement tensor \mathcal{X} as a vector of patient-by-time matrices. On matrix $x_{:,v,:}$, we use linear regression to model the measurements at time index b against those at other indices.

The target values $x_{:,v,b}$ are modeled as

$$x_{:,v,b} = x_{:,v,-b}\beta_{v,b}^{(2)} + \epsilon_{v,b}^{(2)}, \epsilon_{v,b}^{(2)} \sim \mathcal{N}(0, \sigma_{v,b}^{(2)2} I) \quad (3)$$

where $\beta_{v,b}^{(2)}$ is the column vector of coefficients and $\sigma_{v,b}^{(2)}$ is the standard deviation of the error $\epsilon_{v,b}^{(2)}$, regarding the linear regression model through temporal view for variable v and time index b .

The likelihood distribution of $x_{:,v,b}$ is given by

$$x_{:,v,b}|x_{:,v,-b}, \beta_{v,b}^{(2)}, \sigma_{v,b}^{(2)2} \sim \mathcal{N}(x_{:,v,-b}\beta_{v,b}^{(2)}, \sigma_{v,b}^{(2)2} I). \quad (4)$$

The training data consists of observed target values $(x_{p,v,b})_{p \in P_{v,b}^{tr}}$ and input data $(x_{p,v,-b})_{p \in P_{v,b}^{tr}}$.

3) *Gaussian processes through temporal view*: Gaussian processes are commonly used to capture trajectories of variables, thus used in our mixture model to capture temporal correlations. Through the same temporal view as introduced above, on matrix $x_{:,v,:}$, we fit GPs on time series for each patient.

The target value $x_{p,v,b}$ is modeled as

$$x_{p,v,b} = \mu_{p,v,b} + f(t_{p,v,b}), \quad (5)$$

$$f(t_{p,v,b}) \sim \mathcal{GP}(0, \mathcal{K}(t_{p,v,b}, t_{p,v,b'}))$$

where $\mu_{p,v,b}$ is the overall mean of the model, $f(\cdot)$ is a Gaussian process with mean of 0 and a covariance matrix $\mathcal{K}(t, t')$ of time pairs (t, t') . Then the likelihood distribution of $x_{p,v,b}$ is written as

$$x_{p,v,b}|x_{p,v,-b} \sim \mathcal{N}(m^{(3)}(\alpha_{p,v,b}), \Sigma^{(3)}(\alpha_{p,v,b})) \quad (6)$$

$$\alpha_{p,v,b} = (\theta_{v,b}, x_{p,v,-b}, t_{p,v,-b})$$

where $\theta_{v,b}$ are the kernel parameters of the GP models, and the predictive mean and variance are given by $m^{(3)}(\cdot)$ and $\Sigma^{(3)}(\cdot)$; see more details in Appendix D. ² For a certain v and b , all GP models share the same kernel parameters $\theta_{v,b}$.

C. The Mixture Model

Given the likelihood distribution of all three components, we model the joint mixture distribution, regarding the variable v and time index b , in the following way

$$p(x_{:,v,b}, V_{v,b})$$

$$= \pi_{v,b}^{(1)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(1)}) \mathcal{N}(x_{:,v,b}|x_{:,-v,b}\beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2} I)$$

$$+ \pi_{v,b}^{(2)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(2)}) \mathcal{N}(x_{:,v,b}|x_{:,v,-b}\beta_{v,b}^{(2)}, \sigma_{v,b}^{(2)2} I)$$

$$+ \pi_{v,b}^{(3)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(3)}) \mathcal{N}(x_{:,v,b}|m^{(3)}(\alpha_{v,b}), \text{diag}(\Sigma^{(3)}(\alpha_{v,b}))) \quad (7)$$

where we define $V_{v,b} = [x_{:,-v,b} \ x_{:,v,-b}]$, $\alpha_{v,b} = (\alpha_{p,v,b})_{p \in P}$, $\delta_{v,b}^{(k)} = (\mu_{v,b}^{(k)}, \Sigma_{v,b}^{(k)})$ and $\pi_{v,b}^{(k)}$ is the mixing weight for the

²The appendix is available in <https://...>

k th component. This model can be interpreted as the joint distribution between observed data $V_{v,b}$ and missing values $x_{:,v,b}$, which consists of a mixture of three distributions. The first one $p_1(x_{:,v,b}, V_{v,b})$ is modeled as

$$\begin{aligned} p_1(x_{:,v,b}, V_{v,b}) &= p(x_{:,v,b}|V_{v,b})p(V_{v,b}) \\ &= \mathcal{N}(x_{:,v,b}|\mu_1(V_{v,b}), \sigma_1(V_{v,b}))\mathcal{N}(V_{v,b}|\delta_{v,b}^{(1)}) \\ &= \mathcal{N}(x_{:,v,b}|x_{:,-v,b}\beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2}I)\mathcal{N}(V_{v,b}|\delta_{v,b}^{(1)}) \end{aligned} \quad (8)$$

the remaining two follow the same logic. By marginalizing over $x_{:,v,b}$, the prior probability distribution $p(V_{v,b})$ is written as $p(V_{v,b}) = \sum_{k=1}^3 \pi_{v,b}^{(k)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(k)})$, which is a mixture of Gaussians. It also follows that

$$\begin{aligned} p(x_{:,v,b}|V_{v,b}) &= \frac{p(x_{:,v,b}, V_{v,b})}{p(V_{v,b})} \\ &= C_{v,b}^{(1)} \mathcal{N}(x_{:,v,b}|x_{:,-v,b}\beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2}I) \\ &\quad + C_{v,b}^{(2)} \mathcal{N}(x_{:,v,b}|x_{:,v,-b}\beta_{v,b}^{(2)}, \sigma_{v,b}^{(2)2}I) \\ &\quad + C_{v,b}^{(3)} \mathcal{N}(x_{:,v,b}|m^{(3)}(\alpha_{v,b}), \text{diag}(\Sigma^{(3)}(\alpha_{v,b}))) \end{aligned} \quad (9)$$

where we define $C_{v,b}^{(k)} = \frac{\pi_{v,b}^{(k)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(k)})}{\sum_{j=1}^3 \pi_{v,b}^{(j)} \mathcal{N}(V_{v,b}|\delta_{v,b}^{(j)})}$.

We train our mixture model on observed target values by maximizing the log likelihood of the joint mixture distribution

$$\hat{\gamma}_{v,b} = \arg \max_{\gamma_{v,b}} \ln p(x_{:,v,b}^{obs}, V_{v,b}^{tr}; \gamma_{v,b})$$

where $V_{v,b}^{tr}$ is the training input data and defined as the concatenation of $(x_{p,-v,b})_{p \in P_{v,b}^{tr}}$ and $(x_{p,v,-b})_{p \in P_{v,b}^{tr}}$, and $\gamma_{v,b}$ is the set of all trainable parameters

$\beta_{v,b}^{(1)}, \sigma_{v,b}^{(1)2}, \beta_{v,b}^{(2)}, \sigma_{v,b}^{(2)2}, \theta_{v,b}$ and $\pi_{v,b}^{(k)}, \mu_{v,b}^{(k)}, \Sigma_{v,b}^{(k)}$ for $k = 1, 2, 3$.

After training, the missing values are imputed using individualized (per-patient) mixing weights that are derived from the conditional distribution. The missing value $x_{p,v,b}^{mis}$ is imputed by $\Pi_{p,v,b}^{(1)} x_{p,-v,b} \hat{\beta}_{v,b}^{(1)} + \Pi_{p,v,b}^{(2)} x_{p,v,-b} \hat{\beta}_{v,b}^{(2)} + \Pi_{p,v,b}^{(3)} m^{(3)}(\hat{\alpha}_{p,v,b})$, where the individualized mixing weight of the k th component for patient p , variable v and time index b is defined as

$$\Pi_{p,v,b}^{(k)} = \frac{\hat{\pi}_{v,b}^{(k)} \mathcal{N}(V_{p,v,b}|\hat{\mu}_{v,b}^{(k)}, \hat{\Sigma}_{v,b}^{(k)})}{\sum_{j=1}^3 \hat{\pi}_{v,b}^{(j)} \mathcal{N}(V_{p,v,b}|\hat{\mu}_{v,b}^{(j)}, \hat{\Sigma}_{v,b}^{(j)})}, k = 1, 2, 3 \quad (10)$$

where $V_{p,v,b}$ is the observed data and defined as the concatenation of $x_{p,-v,b}$ and $x_{p,v,-b}$.

D. Mixture Parameter Estimation

Let $\ell(\gamma_{v,b})$ be the log likelihood $\ln p(x_{:,v,b}^{obs}, V_{v,b}^{tr}; \gamma_{v,b})$. Explicitly maximizing $\ell(\gamma_{v,b})$ is hard. Instead, we use the EM algorithm to repeatedly construct a lower-bound on $\ell(\gamma_{v,b})$ and then optimize that lower-bound $\mathcal{L}(\gamma_{v,b})$. We first define a latent indicator variable $q_{v,b} \in \{1, 2, 3\}$ that specifies which mixing component that data points come from. Then we use

Jensen's inequality to get the lower-bound $\mathcal{L}(\gamma_{v,b})$, which is given by

$$\begin{aligned} \mathcal{L}(\gamma_{v,b}) &= \sum_{p \in P_{v,b}^{tr}} \sum_{q_{v,b}} Q_p(q_{v,b}) \ln \frac{p(x_{p,v,b}, V_{p,v,b}, q_{v,b}; \gamma_{v,b})}{Q_p(q_{v,b})} \\ &\leq \sum_{p \in P_{v,b}^{tr}} \ln \sum_{k=1}^3 p(q_{v,b} = k) p(x_{p,v,b}, V_{p,v,b}; \gamma_{v,b} | q_{v,b} = k) \end{aligned} \quad (11)$$

where

$$Q_p(q_{v,b} = k) = \frac{p(q_{v,b} = k) p(x_{p,v,b}, V_{p,v,b} | q_{v,b} = k)}{\sum_{j=1}^3 p(q_{v,b} = j) p(x_{p,v,b}, V_{p,v,b} | q_{v,b} = j)}. \quad (12)$$

In (11) and (12), the marginal distribution $p(q_{v,b} = k)$ over $q_{v,b}$ is specified by the mixing coefficients $\pi_{v,b}^{(k)} = p(q_{v,b} = k)$. We can view $Q_p(q_{v,b} = k)$ as the responsibility that component k of the mixture model $Mix_{v,b}$ takes to "explain" $x_{p,v,b}$. We use the standard EM algorithm to maximize the lower-bound $\mathcal{L}(\gamma_{v,b})$; see more details about the estimation of the parameters in Appendix C.

E. The Imputation Model

The proposed mixture model provides flexibility in changing base models and can be solved under various assumptions. For example, a full mixture model consists of all three base models introduced earlier. However, under the assumption of only linear correlations within time series, mixture models composed of only the two linear components (denoted as MixMI-LL) can be used. Although the choice of mixtures can be determined empirically or based on expert knowledge of the data, we propose an imputation model (denoted as MixMI) that chooses the type of mixture automatically.

In MixMI, for each variable and time index, we train both the MixMI-LL and the full mixture model, then select the one with lower training error as the final mixture model that is used to perform imputation in inference. We use the absolute training error, instead of likelihood, as the selection criterion because the likelihood of a mixture model with two linear models is not of the same scale as the full mixture model.

IV. EXPERIMENTAL SETUP

A. Datasets

We collect two real-world datasets from the Medical Information Mart for Intensive Care (MIMIC-III) database [66] and the Northwestern Medicine Enterprise Data Warehouse (NMEDW). Each dataset contains inpatient test results from 13 laboratory tests. These tests are quantitative, frequently measured on hospital inpatients and commonly used in downstream classification tasks [63], [67]. They are the same as those used in [41] in their imputation study. We organize the data by unique admissions. We distinguish multiple admissions of the same patient. Each admission consists of time series of the 13 laboratory tests. Although our model is evaluated only

on continuous variables, they are also applicable on categorical variables with a simple extension (e.g. through predictive mean matching [68]).

In both MIMIC-III and NMEDW datasets, the length of time series varies across admissions. To apply our imputation model on these datasets, we truncate time series so that they have the same length. The length is the average number of time points across all admissions. Before truncating, the average number of time points in MIMIC-III dataset is 11. We first exclude admissions that have less than 11 time points, and then we truncate time series by removing measurements taken after the 11-th time point. We also exclude admissions that contain time series that have no observed values. Our MIMIC-III dataset includes 26,154 unique admissions and the missing rate is about 28.71%. The same data collection procedure is applied on the NMEDW dataset (approved by Northwestern IRB) where we end up with 13,892 unique admissions that have 7 time points. The missing rate of the NMEDW dataset is 24.22%.

We create synthetic datasets to explore and discover the properties of the data that might benefit our model and/or comparison models. In synthetic datasets, we augment the correlation between measurements and times. We do not augment correlations by imposing strong constraints on time series where closer measurements have closer values. Instead, we generate synthetic times by altering real times so that the constraints in synthetic data are “slightly” stronger than real-world data. The synthetic datasets are generated based on the real-world MIMIC-III dataset. We also introduce a scaling factor $d \in (0, 1)$ to control the strength of the constraints in synthetic data. **The synthetic dataset at $d = 0$ is the same as the real-world dataset. As d increases, stronger constraints are introduced to synthetic data (see details in Appendix A).**

B. Evaluation of Imputation Quality

We randomly mask 20% observed measurements in a data set as missing and treat the masked values as the test set. The remaining observed values are used in training. We impute originally missing and masked values together, and compare the imputed values with the ground truth for masked data to evaluate imputation performance. The Mean Absolute Scaled Error (MASE) [69] is used to measure the quality of imputation on the test set. MASE is a scale-free measure of the accuracy of predictions and recommended by Hyndman et al [69], [70] to measure the accuracy of predictions for series. We calculate MASE for all variables and take a weighted average according to the number of masked values of a variable to get an overall MASE per dataset.

Let $mask_{p,v}$ be the set of cardinality $I_{p,v}$ of all time indices that have been masked for patient p and variable v . Also let $Y_{p,v} = (x_{p,v,j}^{obs})_j$ be the sequence of length $J_{p,v}$ of all observed values for patient p and variable v , and let $\tilde{x}_{p,v,i}$ represent the imputed value. The MASE for variable v is defined as

$$MASE(v) = \frac{1}{\sum_{\bar{p}} I_{\bar{p},v}} \sum_p \frac{\sum_{i \in mask_{p,v}} |\tilde{x}_{p,v,i} - x_{p,v,i}^{obs}|}{\frac{J_{p,v}}{J_{p,v}-1} \sum_{j=2}^{J_{p,v}} |Y_{p,v,j} - Y_{p,v,j-1}|}.$$

To show the effectiveness of our imputation model MixMI and its variant MixMI-LL, we compare MASE scores of MixMI and MixMI-LL with other six imputation methods: (a) MICE [8] with 100 imputations, where the average of all imputations are used for evaluation; (b) the pure Gaussian processes, where a GP model is fitted to the observed data of each time series using GPfit [71] in R and missing values are replaced with the predictions from the fitted GP models; (c) 3D-MICE, a state-of-the-art imputation model [41] for which we obtain their code and adapt it to account for our use of the tensor representation; (d) multi-task Gaussian process prediction (MTGP) [35]; (e) a Gaussian mixture model for imputation (GMM) [43], [72] and (f) M-RNN, a state-of-the-art RNN-based imputation model [60].

To tune hyperparameters if any in these models, we mask out 20% observed measurements in the training set as a validation set and tune hyperparameters on the validation set. We introduce $\pi^{(1)}, \pi^{(2)}$ and $\pi^{(3)}$ as hyperparameters of our imputation model. In our experiments, all mixture models start with the same mixing weights, i.e., $\pi_{v,b}^{(k)} = \pi^{(k)}$ for $k = 1, 2, 3$. Other trainable parameters are initialized directly from the data after initial imputation and do not require a manual specification.

We also evaluate our imputation model under a classification task on our MIMIC dataset and compare it with GRU-D [63], a state-of-the-art RNN model for classifications with multivariable time series, which does not require a separate data imputation stage. We build a RNN model with the standard Gated Recurrent Unit as our classification model (GRU-MixMI) trained on data imputed by MixMI. GRU-D is trained on data without imputation. We predict whether a patient dies within 30 days after admission, instead of 48 hours as in [63], because of severe imbalance in our cohort, where only 0.57% patients have positive labels within 48 hours.

V. RESULTS

Table II and Fig. 2 compare all imputation models on 4 datasets. Table II shows the overall MASE score of each imputation model on all datasets. Fig. 2 provides a comparison for all imputation models in MASE score over 3D-MICE by showing the percentage deviation against 3D-MICE. We select 3D-MICE since it is a state-of-the-art benchmark model. We observe that MixMI outperforms all comparison models on all 4 datasets and is significantly better than the second best model ($p=0.001$, permutation test with 1000 replicates) on all 4 datasets. **We perform analyses on the interpretability of the results, the component weights and individualized weights of our model on the real-world MIMIC dataset (see details in Appendix B).**

Our classification model GRU-MixMI is compared to GRU-D with different random seeds pertaining to algorithms on data with different folds. On average, GRU-MixMI ties with GRU-D in terms of AUC scores (0.7857 vs. 0.7853, $p > 0.5$). However, compared with models that perform imputation and prediction together, our imputation model, as a pure imputation method, offers more opportunities for other supervised and

TABLE II
OVERALL MASE BY DATASET AND IMPUTATION MODEL. THE BOLD NUMBERS ARE THE SIGNIFICANTLY BEST VALUES AMONG ALL IMPUTATION MODELS.

Dataset	GP	MTGP ¹	M-RNN	GMM	MICE	3D-MICE	MixMI-LL	MixMI
Real-world MIMIC ($d=0$)	0.13072	0.12599	0.12512	0.11741	0.11763	0.11186	0.09304	0.09285
Synthetic MIMIC ($d=0.5$)	0.10466	-	0.12320	0.11455	0.11573	0.09087	0.08612	0.08448
Synthetic MIMIC ($d=1$)	0.09220	-	0.12201	0.11436	0.11561	0.07715	0.08427	0.07538
NMEDW	0.18353	0.17370	0.14607	0.13593	0.13718	0.13624	0.11600	0.11589

¹ MTGP cannot have multiple inputs, thus not applicable to the synthetic datasets.




Fig. 2. Percentage deviation of MASE score against 3D-MICE

unsupervised tasks, and modeling choices and effective feature engineering.

VI. CONCLUSIONS

We present and demonstrate a mixture-based imputation model MixMI for multivariable clinical time series. Our model can capture both cross-sectional and temporal correlations in time series. We integrate Gaussian processes with mixture models and introduce individualized mixing weights to further improve imputation accuracy. We show that MixMI can provide more accurate imputation than several state-of-the-art imputation models. Although in this work our model is tested on inpatient clinical data, the proposed idea is general and should work for all multivariable time series where interactions of cross-sectional and temporal correlations exist.

REFERENCES

- [1] R. L. Winslow, N. Trayanova, D. Geman, and M. I. Miller, "Computational medicine: Translating models to clinical care," *Science Translational Medicine*, vol. 4, no. 158, pp. 158rv11–158rv11, 2012.
- [2] I. S. Kohane, "Ten things we have to do to achieve precision medicine," *Science*, vol. 349, no. 6243, pp. 37–38, 2015.
- [3] G. M. Weber, W. G. Adams, E. V. Bernstam, J. P. Bickel, K. P. Fox, K. Marsolo, V. A. Raghavan, A. Turchin, X. Zhou, S. N. Murphy *et al.*, "Biases introduced by filtering electronic health records for patients with 'complete data'," *Journal of the American Medical Informatics Association*, vol. 24, no. 6, pp. 1134–1141, 2017.
- [4] R. L. Brown, "Efficacy of the indirect approach for estimating structural equation models with missing data: A comparison of five methods," *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 1, no. 4, pp. 287–316, 1994.
- [5] W. Wothke, "Longitudinal and multi-group modeling with missing data," in *Modeling longitudinal and multiplegroup data: Practical issues, applied approaches, and specific examples*, T. D. Little, K. U. Schnabel, and J. Baumert, Eds., 2000, pp. 219–240.
- [6] J. L. Arbuckle, "Full information estimation in the presence of incomplete data," in *Advanced Structural Equation Modeling: Issues and Techniques*, G. A. Marcoulides and R. E. Schumacker, Eds., 1996, vol. 243, p. 277.
- [7] A. K. Waljee, A. Mukherjee, A. G. Singal, Y. Zhang, J. Warren, U. Balis, J. Marrero, J. Zhu, and P. D. Higgins, "Comparison of imputation methods for missing laboratory data in medicine," *BMJ Open*, vol. 3, no. 8, p. e002847, 2013.
- [8] S. Buuren and K. Groothuis-Oudshoorn, "mice: Multivariate imputation by chained equations in r," *Journal of Statistical Software*, vol. 45, no. 3, 2011.
- [9] J. W. Graham, "Missing data analysis: Making it work in the real world," *Annual Review of Psychology*, vol. 60, pp. 549–576, 2009.
- [10] J. L. Schafer and J. W. Graham, "Missing data: Our view of the state of the art," *Psychological Methods*, vol. 7, no. 2, p. 147, 2002.
- [11] D. B. Rubin, "Multiple imputations in sample surveys—a phenomenological bayesian approach to nonresponse," in *Proceedings of the survey research methods section of the American Statistical Association*, vol. 1. American Statistical Association, 1978, pp. 20–34.
- [12] —, *Multiple imputation for nonresponse in surveys*. John Wiley & Sons, 1987.
- [13] J. L. Schafer, *Analysis of incomplete multivariate data*. Chapman and Hall/CRC, 1997.
- [14] T. E. Raghunathan, J. M. Lepkowski, J. Van Hoewyk, and P. Solenberger, "A multivariate technique for multiply imputing missing values using a sequence of regression models," *Survey Methodology*, vol. 27, no. 1, pp. 85–96, 2001.
- [15] S. Van Buuren, *Flexible imputation of missing data*. Chapman and Hall/CRC, 2018.
- [16] S. Van Buuren, J. P. Brand, C. G. Groothuis-Oudshoorn, and D. B. Rubin, "Fully conditional specification in multivariate imputation," *Journal of statistical computation and simulation*, vol. 76, no. 12, pp. 1049–1064, 2006.
- [17] O. Harel and X.-H. Zhou, "Multiple imputation for the comparison of two screening tests in two-phase alzheimer studies," *Statistics in Medicine*, vol. 26, no. 11, pp. 2370–2388, 2007.
- [18] L. Qi, Y.-F. Wang, and Y. He, "A comparison of multiple imputation and fully augmented weighted estimators for cox regression with missing covariates," *Statistics in Medicine*, vol. 29, no. 25, pp. 2592–2604, 2010.
- [19] Y.-S. Su, A. Gelman, J. Hill, M. Yajima *et al.*, "Multiple imputation with diagnostics (mi) in r: Opening windows into the black box," *Journal of Statistical Software*, vol. 45, no. 2, pp. 1–31, 2011.
- [20] C.-H. Hsu, J. M. Taylor, S. Murray, and D. Commenges, "Survival analysis using auxiliary variables via non-parametric multiple imputation," *Statistics in Medicine*, vol. 25, no. 20, pp. 3503–3517, 2006.
- [21] Q. Long, C.-H. Hsu, and Y. Li, "Doubly robust nonparametric multiple imputation for ignorable missing data," *Statistica Sinica*, vol. 22, p. 149, 2012.
- [22] S. Van Buuren, H. C. Boshuizen, and D. L. Knook, "Multiple imputation of missing blood pressure covariates in survival analysis," *Statistics in Medicine*, vol. 18, no. 6, pp. 681–694, 1999.
- [23] Y. Deng, C. Chang, M. S. Ido, and Q. Long, "Multiple imputation for general missing data patterns in the presence of high-dimensional data," *Scientific Reports*, vol. 6, p. 21689, 2016.
- [24] D. J. Stekhoven and P. Bühlmann, "Missforest—non-parametric missing value imputation for mixed-type data," *Bioinformatics*, vol. 28, no. 1, pp. 112–118, 2011.
- [25] R. Little and H. An, "Robust likelihood-based analysis of multivariate data with missing values," *Statistica Sinica*, pp. 949–968, 2004.

- [26] Y. Luo, P. Szolovits, A. S. Dighe, and J. M. Baron, "Using machine learning to predict laboratory test results," *American Journal of Clinical Pathology*, vol. 145, no. 6, pp. 778–788, 2016.
- [27] G. Zhang and R. Little, "Extensions of the penalized spline of propensity prediction method of imputation," *Biometrics*, vol. 65, no. 3, pp. 911–918, 2009.
- [28] O. Troyanskaya, M. Cantor, G. Sherlock, P. Brown, T. Hastie, R. Tibshirani, D. Botstein, and R. B. Altman, "Missing value estimation methods for dna microarrays," *Bioinformatics*, vol. 17, no. 6, pp. 520–525, 2001.
- [29] A. Shah, N. Laird, and D. Schoenfeld, "A random-effects model for multiple characteristics with possibly missing data," *Journal of the American Statistical Association*, vol. 92, no. 438, pp. 775–779, 1997.
- [30] M. Liu, J. M. Taylor, and T. R. Belin, "Multiple imputation and posterior simulation for multivariate missing data in longitudinal studies," *Biometrics*, vol. 56, no. 4, pp. 1157–1163, 2000.
- [31] J. L. Schafer and R. M. Yucel, "Computational strategies for multivariate linear mixed-effects models with missing values," *Journal of Computational and Graphical Statistics*, vol. 11, no. 2, pp. 437–457, 2002.
- [32] E. E. Holmes, E. J. Ward, and K. Wills, "Marss: Multivariate autoregressive state-space models for analyzing time-series data," *R Journal*, vol. 4, no. 1, 2012.
- [33] F. Bashir and H.-L. Wei, "Handling missing data in multivariate time series using a vector autoregressive model-imputation (var-im) algorithm," *Neurocomputing*, 2017.
- [34] J. Futoma, S. Hariharan, and K. Heller, "Learning to detect sepsis with a multitask gaussian process rnn classifier," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1174–1182.
- [35] E. V. Bonilla, K. M. Chai, and C. Williams, "Multi-task gaussian process prediction," in *Advances in Neural Information Processing Systems*, 2008, pp. 153–160.
- [36] Y. He, R. Yucel, and T. E. Raghunathan, "A functional multiple imputation approach to incomplete longitudinal data," *Statistics in Medicine*, vol. 30, no. 10, pp. 1137–1156, 2011.
- [37] J.-M. Chiou, Y.-C. Zhang, W.-H. Chen, and C.-W. Chang, "A functional data approach to missing value imputation and outlier detection for traffic flow data," *Transportmetrica B: Transport Dynamics*, vol. 2, no. 2, pp. 106–129, 2014.
- [38] S. Kliethermes and J. Oleson, "A bayesian approach to functional mixed-effects modeling for longitudinal data with binomial outcomes," *Statistics in Medicine*, vol. 33, no. 18, pp. 3130–3146, 2014.
- [39] K. Yu, V. Tresp, and A. Schwaighofer, "Learning gaussian processes from multiple tasks," in *Proceedings of the 22nd International Conference on Machine Learning*. ACM, 2005, pp. 1012–1019.
- [40] T. Hori, D. Montcho, C. Agbangla, K. Ebana, K. Futakuchi, and H. Iwata, "Multi-task gaussian process for imputing missing data in multi-trait and multi-environment trials," *Theoretical and Applied Genetics*, vol. 129, no. 11, pp. 2101–2115, 2016.
- [41] Y. Luo, P. Szolovits, A. S. Dighe, and J. M. Baron, "3d-mice: integration of cross-sectional and longitudinal imputation for multi-analyte longitudinal clinical data," *Journal of the American Medical Informatics Association*, vol. 25, no. 6, pp. 645–653, 2017.
- [42] V. Tresp, "Mixtures of gaussian processes," in *Advances in neural information processing systems*, 2001, pp. 654–660.
- [43] M. Di Zio, U. Guarnera, and O. Luzzi, "Imputation through finite gaussian mixture models," *Computational Statistics & Data Analysis*, vol. 51, no. 11, pp. 5305–5316, 2007.
- [44] O. Delalleau, A. Courville, and Y. Bengio, "Efficient em training of gaussian mixtures with missing data," *arXiv preprint arXiv:1209.0521*, 2012.
- [45] X. Yan, W. Xiong, L. Hu, F. Wang, and K. Zhao, "Missing value imputation based on gaussian mixture model for the internet of things," *Mathematical Problems in Engineering*, vol. 2015, 2015.
- [46] D. S. Silva and C. V. Deutsch, "Multivariate data imputation using gaussian mixture models," *Spatial statistics*, vol. 27, pp. 74–90, 2018.
- [47] M. Filipović and A. Jukić, "Tucker factorization with missing data with application to low- n -rank tensor completion," *Multidimensional Systems and Signal Processing*, vol. 26, no. 3, pp. 677–692, 2015.
- [48] Y. Liu, F. Shang, L. Jiao, J. Cheng, and H. Cheng, "Trace norm regularized candecomp/parafac decomposition with missing data," *IEEE Transactions on Cybernetics*, vol. 45, no. 11, pp. 2437–2448, 2015.
- [49] R. Tomioka, K. Hayashi, and H. Kashima. (2010) Estimation of low-rank tensors via convex optimization. [Online]. Available: <https://arxiv.org/pdf/1010.0789.pdf>
- [50] M. T. Bahadori, Q. R. Yu, and Y. Liu, "Fast multivariate spatio-temporal analysis via low rank tensor learning," in *Advances in Neural Information Processing Systems*, 2014, pp. 3491–3499.
- [51] R. Yu, D. Cheng, and Y. Liu, "Accelerated online low-rank tensor learning for multivariate spatio-temporal streams," in *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, 2015, pp. 238–247.
- [52] H. Ge, J. Caverlee, N. Zhang, and A. Squicciarini, "Uncovering the spatio-temporal dynamics of memes in the presence of incomplete information," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. ACM, 2016, pp. 1493–1502.
- [53] K. Takeuchi, H. Kashima, and N. Ueda, "Autoregressive tensor factorization for spatio-temporal predictions," *IEEE International Conference on Data Mining*, 2017.
- [54] Y. Cai, H. Tong, W. Fan, P. Ji, and Q. He, "Facets: Fast comprehensive mining of coevolving high-order time series," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2015, pp. 79–88.
- [55] X. Yang, Y. Zhang, and M. Chi, "Time-aware subgroup matrix decomposition: Imputing missing data using forecasting events," in *2018 IEEE International Conference on Big Data*. IEEE, 2018, pp. 1524–1533.
- [56] S. Zhe and Y. Du, "Stochastic nonparametric event-tensor decomposition," in *Advances in Neural Information Processing Systems*, 2018, pp. 6855–6865.
- [57] Y. Bengio and F. Gingras, "Recurrent neural networks for missing or asynchronous data," in *Advances in neural information processing systems*, 1996, pp. 395–401.
- [58] V. Tresp and T. Briegel, "A solution for missing data in recurrent neural networks with an application to blood glucose prediction," in *Advances in Neural Information Processing Systems*, 1998, pp. 971–977.
- [59] S. Parveen and P. Green, "Speech recognition with missing data using recurrent neural nets," in *Advances in Neural Information Processing Systems*, 2002, pp. 1189–1195.
- [60] J. Yoon, W. R. Zame, and M. van der Schaar, "Estimating missing data in temporal data streams using multi-directional recurrent neural networks," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 5, pp. 1477–1490, 2018.
- [61] E. Choi, M. T. Bahadori, A. Schuetz, W. F. Stewart, and J. Sun, "Doctor ai: Predicting clinical events via recurrent neural networks," in *Machine Learning for Healthcare Conference*, 2016, pp. 301–318.
- [62] Z. C. Lipton, D. Kale, and R. Wetzel, "Directly modeling missing data in sequences with rnns: Improved classification of clinical time series," in *Machine Learning for Healthcare Conference*, 2016, pp. 253–270.
- [63] Z. Che, S. Purushotham, K. Cho, D. Sontag, and Y. Liu, "Recurrent neural networks for multivariate time series with missing values," *Scientific reports*, vol. 8, no. 1, p. 6085, 2018.
- [64] Y. Xue, D. Klabjan, and Y. Luo, "Predicting icu readmission using grouped physiological and medication trends," *Artificial Intelligence in Medicine*, 2018.
- [65] A. Gelman, "Parameterization and bayesian modeling," *Journal of the American Statistical Association*, vol. 99, no. 466, pp. 537–545, 2004.
- [66] A. E. Johnson, T. J. Pollard, L. Shen, H. L. Li-wei, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "Mimic-iii, a freely accessible critical care database," *Scientific Data*, vol. 3, p. 160035, 2016.
- [67] J.-R. Le Gall, S. Lemeshow, and F. Saulnier, "A new simplified acute physiology score (saps ii) based on a european/north american multi-center study," *Jama*, vol. 270, no. 24, pp. 2957–2963, 1993.
- [68] R. J. Little, "Missing-data adjustments in large surveys," *Journal of Business & Economic Statistics*, vol. 6, no. 3, pp. 287–296, 1988.
- [69] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International Journal of Forecasting*, vol. 22, no. 4, pp. 679–688, 2006.
- [70] P. H. Franses, "A note on the mean absolute scaled error," *International Journal of Forecasting*, vol. 32, no. 1, pp. 20–22, 2016.
- [71] B. MacDonald, P. Ranjan, and H. Chipman, "Gpfit: An r package for fitting a gaussian process model to deterministic simulator outputs," *Journal of Statistical Software*, vol. 64, no. 1, pp. 1–23, 2015.
- [72] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.