

# Collaborative Synthesis of Patient Records through Multi-Visit Health State Inference

Hongda Sun<sup>1\*</sup>, Hongzhan Lin<sup>1\*</sup>, Rui Yan<sup>12†</sup>

<sup>1</sup> Gaoling School of Artificial Intelligence, Renmin University of China

<sup>2</sup> Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education  
{sunhongda98, linhongzhan, ruiyan}@ruc.edu.cn

## Abstract

Electronic health records (EHRs) have become the foundation of machine learning applications in healthcare, while the utility of real patient records is often limited by privacy and security concerns. Synthetic EHR generation provides an additional perspective to compensate for this limitation. Most existing methods synthesize new records based on real EHR data, without consideration of different types of events in EHR data, which cannot control the event combinations in line with medical common sense. In this paper, we propose **MSIC**, a **M**ulti-visit health **S**tatus **I**nference model for Collaborative EHR synthesis to address these limitations. First, we formulate the synthetic EHR generation process as a probabilistic graphical model and tightly connect different types of events by modeling the latent health states. Then, we derive a health state inference method tailored for the multi-visit scenario to effectively utilize previous records to synthesize current and future records. Furthermore, we propose to generate medical reports to add textual descriptions for each medical event, providing broader applications for synthesized EHR data. For generating different paragraphs in each visit, we incorporate a multi-generator deliberation framework to collaborate the message passing of multiple generators and employ a two-phase decoding strategy to generate high-quality reports. Our extensive experiments on the widely used benchmarks, MIMIC-III and MIMIC-IV, demonstrate that MSIC advances state-of-the-art results on the quality of synthetic data while maintaining low privacy risks.

## Introduction

Electronic health records (EHRs) have received extensive attention in research and applications in the AI-empowered healthcare field. Mainstream healthcare-related tasks such as medical named entity recognition (Gligic et al. 2020; Gorinski et al. 2019; Quimbaya et al. 2016) and drug recommendation (Sun et al. 2022; Yang et al. 2021; Shang et al. 2019) can be carried out based on EHR datasets. However, EHR data sharing is generally limited by privacy/policy regulations, resulting in real records being insufficient for requirements (Neamatullah et al. 2008; El Emam et al. 2011; Benitez and Malin 2010; El Emam, Rodgers, and Malin 2015).

\*These authors contributed equally.

†Corresponding author: Rui Yan (ruiyan@ruc.edu.cn)

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

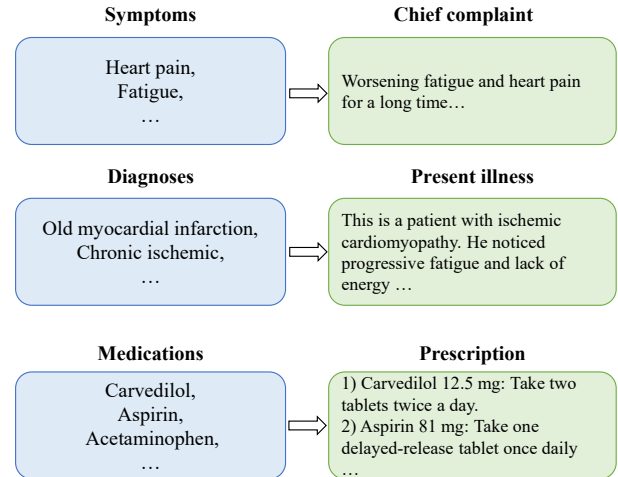


Figure 1: An example of a single-visit EHR data.

Therefore, synthetic EHR generation is a safer alternative for privacy protection (Cui et al. 2020; Zhang et al. 2021). With the promising development of deep generative models (Xu, Li, and Zhou 2015; Salakhutdinov 2015), generative adversarial networks (GANs) and variational autoencoders (VAEs) are widely used for synthetic EHR generation (Li et al. 2021; Rashidian et al. 2020; Yan et al. 2020). These methods learn from the real EHR data to synthesize a series of new records. However, longitudinal EHR data has a patient-visit-event hierarchical structure: (1) each patient has sequential records for multiple visits; (2) each patient visit contains various types of medical events including symptoms, diagnoses, and medications. Previous EHR synthesis methods usually fail to model the logical associations between different types of events. Instead, they generate symptoms, diagnoses, and medications symmetrically, which cannot ensure that their combination conforms to medical common sense, affecting the quality and reliability of synthesized data. For multi-visit synthetic scenarios, how to use previous records to accurately generate current events is also a key challenge for this task. Furthermore, original EHR synthesis methods only focus on generating discrete labels for each medical event, and this limited medical information directly affects their application scope.

To this end, we propose a **Multi-visit health Status Inference model for Collaborative EHR synthesis (MSIC)** to address these limitations. First, we formulate the synthetic EHR generation process as a probabilistic graphical model to capture the correlations between different types of medical events. To better model the generation process, we define a latent health state variable for each patient as a key connection with different events. Specifically, we treat symptoms and diagnoses as the patient’s inherent factors, which determine the distribution of the patient’s health state. Then, by inferring the health state, we can obtain appropriate medications for the patient, and generate symptoms and diagnoses in the next visit. Second, we generalize the graphical model to the multi-visit scenario by linking the connections between events with the same type. Correspondingly, we link the health states for each patient to make it continuously evolve across multiple visits. Thus, we can generate the current synthetic record based on the patient’s previous records and enable sequential synthesis.

Traditional EHR synthesis generates limited medical event labels, restricting task coverage. We then propose to generate medical reports to expand the application scenarios to serve the tasks that labels cannot support, like symptom extraction (Sondhi et al. 2012), diagnosis prediction (Franz, Shrestha, and Paudel 2020), and clinical outcome prediction (Naik et al. 2022). An example of a single-visit EHR data is shown in Figure 1, with three types of medical events on the left and the corresponding medical report paragraphs detailing the events on the right. The patient’s health state acts as also a conditional signal to control the report generation process. We develop a new multi-generator deliberation framework with two-phase decoding to collaborate the paragraph generation among multiple generators. In the draft phase, we sequentially draft three paragraphs through forward message passing. In the polish phase, we use the paragraphs drafted by each generator as deliberation feedback to refine the other two paragraphs.

We conduct extensive experiments on two benchmark datasets, MIMIC-III and MIMIC-IV, to evaluate the quality and privacy of synthetic EHR data. Experimental results demonstrate that our MSIC outperforms state-of-the-art baselines on both medical event synthesis and report generation. Meanwhile, evaluation results of membership and attribute inference attacks show that MSIC can also maintain low privacy risks for safer synthesis.

Our technical contributions are summarized as follows:

- We construct a probabilistic graphical model to formulate the synthetic EHR generation process. We introduce a latent health state variable for each patient to connect the logical associations between various types of medical events.
- We derive a new sequential health state inference method to better leverage a patient’s previous visits to synthesize current and future EHR data more coherently.
- We propose to generate medical reports to detail each medical event using textual descriptions, providing broader applications for synthetic data. We develop a multi-generator deliberation framework and a two-phase decoding strategy to improve the collaboration of multi-paragraph generation.

## Related Work

### Synthetic EHR Generation

Deep generative models have advanced promising performance on synthetic EHR generation. GAN-style models are also widely applied in this field. MedGAN combined an autoencoder and GAN to generate high-dimensional medical codes (Choi et al. 2017). MedGAN was improved into two more effective variants, MedWGAN and MedBGAN, to achieve comparable performance on distribution statistics (Baowaly et al. 2019). LongGAN combined RNN and Wasserstein GAN Gradient Penalty to generate time-series data (Sun et al. 2021). However, GAN-style methods have several drawbacks such as training instability and mode collapse (Durall Lopez et al. 2021; Zhang, Li, and Yu 2018). EVA used a variational autoencoder modeling longitudinal patterns across visits to produce records (Biswal et al. 2021). Language models can also be used in generating synthetic EHR data. PromptEHR uses prompt learning for conditional generation of synthetic data (Wang and Sun 2022). Different from previous methods, we model the logical associations between various types of medical events with multi-visit health state inference.

### Medical Language Modeling and Generation

Language models have been widely used in the clinical and medical fields (Wang et al. 2021). BioBERT pre-trained BERT with biomedical corpora for natural language understanding tasks (Lee et al. 2020). ClinicalBERT (Alsentzer et al. 2019) and BlueBERT (Peng, Chen, and Lu 2020) are also BERT-based models with more corpora. GENRE (De Cao et al. 2021) and BARTNER (Yan et al. 2021) are all auto-regressive models for named entity recognition and entity linking tasks. BioBART is pre-trained on biomedical literature and can be used for downstream generation tasks including medical dialogue and summarization (Yuan et al. 2022). We employ BioBART as the backbone of medical report generation and further collaborate with multiple generators to enhance message passing between them, improving the quality of generated reports.

## Preliminary

### Problem Formulation

An EHR dataset containing  $N$  patients can be denoted as  $\mathbf{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(N)}\}$ . The  $j$ -th patient can be characterized by a sequence of multiple visits:  $\mathbf{x}^{(j)} = [\mathbf{x}_1^{(j)}, \mathbf{x}_2^{(j)}, \dots, \mathbf{x}_{T_j}^{(j)}]$ , where  $T_j$  is the total number of visits of patient  $j$ . Within each visit, the record encompasses various medical event types. Specifically, the  $t$ -th visit of patient  $j$  can be represented as  $\mathbf{x}_t^{(j)} = (\mathbf{s}_t^{(j)}, \mathbf{d}_t^{(j)}, \mathbf{m}_t^{(j)})$ , where  $\mathbf{s}_t^{(j)} \in [0, 1]^{|S|}$ ,  $\mathbf{d}_t^{(j)} \in [0, 1]^{|D|}$ , and  $\mathbf{m}_t^{(j)} \in [0, 1]^{|M|}$  are binary labels of symptoms, diagnoses and medications, respectively. Here,  $S$ ,  $D$ , and  $M$  indicates the overall symptom, diagnosis, and medication sets, and  $|\cdot|$  indicates the cardinality of a given set.

The objective of synthetic EHR generation can be divided into two primary processes. First, we aim to synthesize new

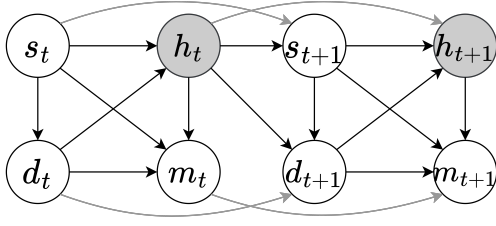


Figure 2: The graphical representation of synthetic EHR generation. Shaded circles represent latent variables.

accurate records from existing partial records. In one case, **cross-type synthesis** needs to generate other types of events based on observed event types (e.g., generate  $\mathbf{m}_t^{(j)}$  given  $\mathbf{s}_t^{(j)}$  and  $\mathbf{d}_t^{(j)}$ ). In another case, **longitudinal synthesis** needs to generate the next visit  $\mathbf{x}_{t+1}^{(j)}$  based on the existing visits  $(\mathbf{x}_1^{(j)}, \dots, \mathbf{x}_t^{(j)})$ . Second, beyond synthesizing discrete event labels, we introduce a supplementary task, namely **medical report generation**, to improve the informativeness of synthetic data. Consistent with the event labels  $\mathbf{x}_t^{(j)}$ , the medical report  $\mathbf{y}_t^{(j)}$  consists of three paragraphs: (1) *chief complaint*  $\mathbf{y}_{s_t}^{(j)}$  describes the patient’s self-reported symptoms, corresponding to  $\mathbf{s}_t^{(j)}$ ; (2) *present illness*  $\mathbf{y}_{d_t}^{(j)}$  specifies the current detailed diagnosis of the patient, corresponding to  $\mathbf{d}_t^{(j)}$ ; and (3) *prescription*  $\mathbf{y}_{m_t}^{(j)}$  details the recommended medications for the patient, corresponding to  $\mathbf{m}_t^{(j)}$ .

### Graphical Model for Synthetic EHR Generation

Previous synthetic EHR generation methods overlook the unique roles of different medical events and treat them as symmetrical modeling positions (Biswal et al. 2021; Choi et al. 2017). To overcome this limitation, we consider the formation logic of real-world EHR data and design a probabilistic graphical model to capture the generative process, displayed in Figure 2. Note that we will omit the superscript  $j$  referring to a specific patient for simplification. The process begins with establishing the diagnoses  $d_t$  from observed symptoms  $s_t$ . Since  $s_t$  and  $d_t$  are inherent factors related to the patient’s health condition, we derive a health state variable  $h_t$  as a crucial conditional signal in generating these events. The medications  $m_t$  can be determined by inferring  $h_t$  in conjunction with  $s_t$  and  $d_t$  to simulate the real-world prescription process. Additionally, the current  $h_t$  also directly influences symptoms  $s_{t+1}$  and diagnoses  $d_{t+1}$  in the next visit. To enable longitudinal synthesis across visits, we expand the graphical model to accommodate multi-visit scenarios. We model the correlations exhibited by these events by adding critical links in the graphical model. Figure 2 illustrates the cross-type relation paths with black arrows, which are mainly centered around the health state to depict the interplay between symptoms, diagnoses, and medications. The gray arrows represent longitudinal relation paths demonstrating the dependencies between consecutive visits, enriching the temporal dynamics modeling in synthesizing EHR data.

## Our Method

### Sequential Health State Inference

Following the formation logic of EHR data, we propose a new variational Bayesian generative model to formulate the generative process. The latent health state  $h_t$  is inferred and evolved over multiple visits and used for a conditional signal for both medical event synthesis and report generation. During the  $t$ -th visit, the observed factors  $s_t$  and  $d_t$  play a crucial role in deriving  $h_t$  since they directly reflect the distribution of the patient’s health condition. Meanwhile, the patient’s health state progresses sequentially across multiple visits, necessitating consideration of the impact of the previous health state  $h_{t-1}$  on the distribution of  $h_t$ . Therefore, how to accurately formulate the distribution of  $h_t$  is crucial for both cross-type and longitudinal EHR synthesis. To infer the health state  $h_t$ , we integrate the effects of both current patient inherent factors ( $s_t, d_t$ ) and previous health state  $h_{t-1}$ , and derive its prior distribution  $P(h_t|h_{t-1}, s_t, d_t)$  parameterized by a prior network  $P_\theta$ . Then  $h_t$  can be drawn based on the prior for generating current medications  $m_t$ . To refine the inference process, we introduce the ground-truth medications  $m_t$  as an additional auxiliary condition and use the posterior distribution  $P(h_t|h_{t-1}, s_t, d_t, m_t)$  to estimate  $h_t$ . We approximate the posterior using an inference network  $P_\phi$  due to complex expectation manipulations of the exact posterior. Overall, we derive the inferred health state  $h_t$  from the inference network  $P_\phi$  in the training stage, or obtain  $h_t$  from the prior network  $P_\theta$  in the testing stage. To optimize both  $P_\theta$  and  $P_\phi$ , we derive the variational evidence lower bound (ELBO), as detailed in the following subsections.

### Synthetic EHR Generation

Based on the inferred health state  $h_t$ , it serves as a conditional signal to generate synthetic EHR data. We divide synthetic EHR generation into the following two steps: (1) multi-visit synthesis for medical events labels and (2) multi-generator deliberation for medical report generation.

**Multi-visit synthesis for medical events.** Given the previous records  $\mathbf{x}_{t-1}$  and the current patient intrinsic factors ( $s_t, d_t$ ), we can perform cross-type synthesis for current medications by estimating  $P(m_t|\mathbf{x}_{t-1}, s_t, d_t)$ , the probability of observing each  $m_t$  in  $\mathbf{x}_t$ . Figure 2 indicates that we need to use the previous health state  $h_{t-1}$  as a key condition of inferring the current  $h_t$ . Then, we can use the derived  $h_{t-1}$  to generate  $m_t$ . From the perspective of ELBO, we can approximate the health state distributions and derive the following probability estimation:

$$\begin{aligned} & \log \hat{P}(m_t|\mathbf{x}_{t-1}, s_t, d_t) \\ &= \mathbb{E}_{P_\phi(h_{t-1})} [\log \mathbb{E}_{P_\theta(h_t)} P(m_t|m_{t-1}, h_t, s_t, d_t)] \\ &\geq \mathbb{E}_{P_\phi(h_{t-1})} [\mathbb{E}_{P_\phi(h_t)} \log g_m(m_t) - \text{KL}(P_\phi(h_t)||P_\theta(h_t))], \quad (1) \end{aligned}$$

where  $\mathbb{E}$  indicates the expectation operation,  $g_m$  is the medication decoder for generating  $m_t$ , and  $\text{KL}$  denotes the Kullback-Leibler divergence.  $P_\phi(h_t)$  and  $P_\theta(h_t)$  are the abbreviations of prior and inference networks, omitting conditionals for simplification. The first term is to reconstruct  $m_t$  from the distribution  $P(m_t|m_{t-1}, h_t, s_t, d_t)$ , and the second

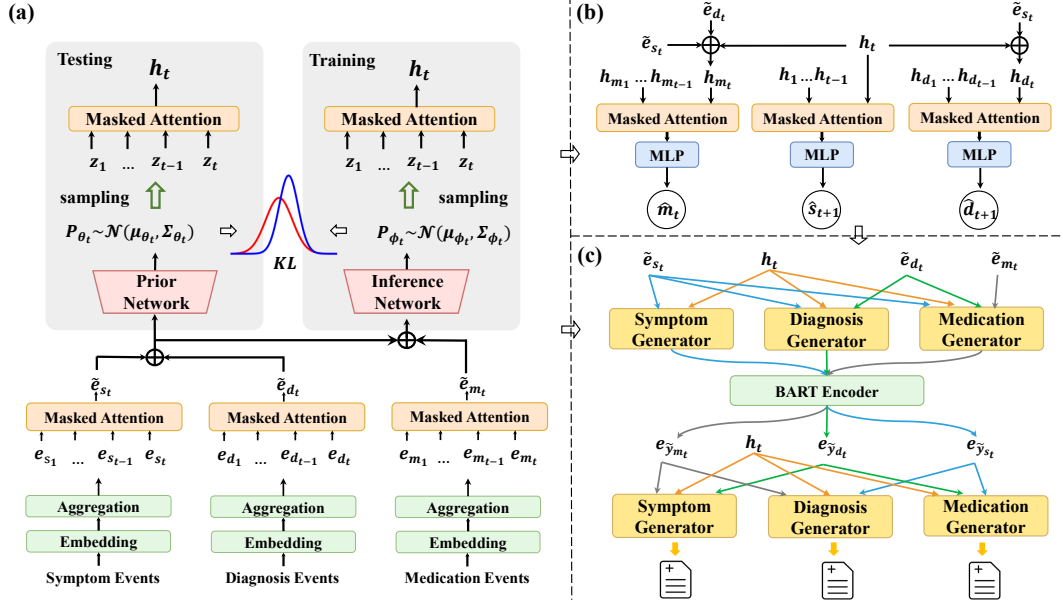


Figure 3: The model overview of MSIC with three modules: (a) health state inference module; (b) medical event synthesis module; (c) medical report generation module.

KL term is to shrink the distance between  $P_{\phi}(h_t)$  to  $P_{\theta}(h_t)$ . Similar to the synthesis of medications  $m_t$ ,  $h_t$  can also be used for cross-type synthesis on the symptoms  $s_{t+1}$  and diagnoses  $d_{t+1}$  in the next visit. Given the current record  $\mathbf{x}_t$  along with inferred  $h_t$ , we can derive the estimated probabilities as:

$$\begin{aligned} & \log \hat{P}(s_{t+1} | \mathbf{x}_t) \\ &= \mathbb{E}_{P_{\phi}(h_{t-1})} [\log \mathbb{E}_{P_{\theta}(h_t)} P(s_{t+1} | s_t, h_t)] \\ &\geq \mathbb{E}_{P_{\phi}(h_{t-1})} [\mathbb{E}_{P_{\phi}(h_t)} \log g_s(s_{t+1}) - \text{KL}(P_{\phi}(h_t) || P_{\theta}(h_t))], \quad (2) \end{aligned}$$

$$\begin{aligned} & \log \hat{P}(d_{t+1} | \mathbf{x}_t, s_{t+1}) \\ &= \mathbb{E}_{P_{\phi}(h_{t-1})} [\log \mathbb{E}_{P_{\theta}(h_t)} P(d_{t+1} | d_t, h_t, s_{t+1})] \\ &\geq \mathbb{E}_{P_{\phi}(h_{t-1})} [\mathbb{E}_{P_{\phi}(h_t)} \log g_d(d_{t+1}) - \text{KL}(P_{\phi}(h_t) || P_{\theta}(h_t))]. \quad (3) \end{aligned}$$

where  $g_s$  and  $g_d$  are symptom and diagnosis decoders for generating  $s_{t+1}$  and  $d_{t+1}$ , respectively. Therefore, we can synthesize all types of medical events given the existing partial records. Building upon the above cross-type synthesis, we can extend our method to make longitudinal synthesis for the next visit  $\mathbf{x}_{t+1}$  based on the current record  $\mathbf{x}_t$ . To achieve this, we intuitively decompose the process of  $\mathbf{x}_{t+1}$  into separate sequential synthesis of  $s_{t+1}$ ,  $d_{t+1}$ , and  $m_{t+1}$ . Based on Equations (1)-(3) The joint probability can be decomposed and rearranged as:

$$\begin{aligned} & \log \hat{P}(\mathbf{x}_{t+1} | \mathbf{x}_t) \\ &= \log [\hat{P}(s_{t+1} | \mathbf{x}_t) \cdot \hat{P}(d_{t+1} | \mathbf{x}_t, s_{t+1}) \cdot \hat{P}(m_{t+1} | \mathbf{x}_t, s_{t+1}, d_{t+1})] \\ &\geq \mathbb{E}_{P_{\phi}(h_{t-1})} [\mathbb{E}_{P_{\phi}(h_t)} \log [g_s(s_{t+1}) g_d(d_{t+1})] - \text{KL}(P_{\phi}(h_t) || P_{\theta}(h_t))] \\ &\quad + \mathbb{E}_{P_{\phi}(h_t)} [\mathbb{E}_{P_{\phi}(h_{t+1})} \log g_m(m_{t+1}) - \text{KL}(P_{\phi}(h_{t+1}) || P_{\theta}(h_{t+1}))] \quad (4) \end{aligned}$$

The sum of two items in the reconstruction indicates that the synthesis of  $\mathbf{x}_{t+1}$  involves two steps: (1) draw  $h_t$  to generate  $s_{t+1}$  and  $d_{t+1}$ ; and (2) estimate  $P_{\theta}(h_{t+1})$  and  $P_{\phi}(h_{t+1})$

based on previously drawn  $h_t$ , which can be used to derive  $h_{t+1}$  and generate  $m_{t+1}$ .

**Multi-generator deliberation for report generation.** We further generate a medical report  $y_t$  in each visit  $t$  based on the medical events  $\mathbf{x}_t$ , so that the paragraphs of  $y_t$  correspond to different types of events to enhance the informativeness of synthetic EHR data. The health state  $h_t$  is also treated as the conditional signal for report generation. In order to ensure flexible collaboration between the generation of chief complaint  $y_{s_t}$ , present illness  $y_{d_t}$  and prescription  $y_{m_t}$ , we employ three generators, each responsible for producing a specific paragraph of  $y_t$ . As these three paragraphs are interdependent, we initially generate  $y_{s_t}$ ,  $y_{d_t}$ , and  $y_{m_t}$  sequentially based on  $h_t$  and  $\mathbf{x}_t$ . The generative process of  $y_t$  can be formulated as:

$$\begin{aligned} & \log P(y_t | \mathbf{x}_t) \\ &= \mathbb{E}_{P_{\phi}(h_{t-1})} [\log \mathbb{E}_{P_{\phi}(h_t)} (g_s^*(h_t, s_t) g_d^*(h_t, s_t, d_t) g_m^*(h_t, s_t, d_t, m_t))] \\ &\geq \mathbb{E}_{P_{\phi}(h_{t-1})} [\mathbb{E}_{P_{\phi}(h_t)} \log [g_s^*(h_t, s_t) g_d^*(h_t, s_t, d_t) g_m^*(h_t, s_t, d_t, m_t)] \\ &\quad - \text{KL}(P_{\phi}(h_t) || P_{\theta}(h_t))], \quad (5) \end{aligned}$$

where  $g_s^*$ ,  $g_d^*$ , and  $g_m^*$  are the symptom, diagnosis, and medication report generators, respectively. Our focus lies in enhancing the collaboration between distinct paragraphs to improve overall generation quality. Consequently, we introduce a multi-generator deliberation framework to facilitate paragraph generation collaboration among multiple generators via message passing and deliberation polish. As shown in Figure 3(c), we assume that report generation consists of a draft phase and a polish phase. In the draft phase, multiple generators share the same patient's health state  $h_t$  as the control signal. Additionally, medical events can serve as constraints for generating corresponding paragraphs, thereby improving the alignment between events and reports. Therefore, we can obtain three draft paragraphs  $\tilde{y}_{s_t}$ ,  $\tilde{y}_{d_t}$ , and

$\tilde{y}_{m_t}$ . To further enhance message passing and collaboration within our multi-generator framework, we draw inspiration from the deliberation networks (Xia et al. 2017; Li et al. 2019) and design a refining mechanism in the polish phase to revise the drafts. Specifically, we encode the draft  $\tilde{y}_{m_t}$  and use the encoded draft  $e_{\tilde{y}_{m_t}}$  as feedback input for the symptom generator and diagnosis generator. Similarly, for the other two drafts, each is passed as feedback to two other generators different from itself. By incorporating draft feedback, the framework promotes iterative refinement and collaboration among the generators, leading to the generation of refined paragraphs  $\hat{y}_{s_t}$ ,  $\hat{y}_{d_t}$ , and  $\hat{y}_{m_t}$ .

## Model Implementation

We will detail the implementation of our method in this subsection. The model overview of MSIC is illustrated in Figure 3, including health state inference, medical event synthesis, and medical report generation.

**Health state inference module.** We employ three separate encoders to encode the input medical events. At the  $t$ -th visit, taking the symptoms  $s_t = (s_{t_1}, \dots, s_{t_{|s_t|}})$  as an example, each symptom item is first encoded through an embedding layer. Then, a transformer encoder is used to further process  $|s_t|$  embeddings and integrate them into a symptom representation  $e_{s_t}$ . To better model the previous events, we use an attention mechanism  $g_{attn}$  to update  $e_{s_t}$  into  $\tilde{e}_{s_t}$ . Similarly, the diagnosis and medication encoder can be used to encode corresponding representations  $e_{d_t}$  and  $e_{m_t}$ . Following the structure of our graphical model, We combine the previously derived  $\tilde{e}_{s_t}$  to get the updated representation  $\tilde{e}_{d_t}$  using two multi-layer perceptions (MLPs)  $g_d$  and  $g_m$ . Similarly,  $\tilde{e}_{m_t}$  can be obtained based on  $\tilde{e}_{s_t}$  and  $\tilde{e}_{d_t}$ . The encoding process can be formulated as:

$$\begin{aligned}\tilde{e}_{s_t} &= g_{attn}(e_{s_1}, \dots, e_{s_t}), \\ \tilde{e}_{d_t} &= g_d(g_{attn}(e_{d_1}, \dots, e_{d_t}) \oplus \tilde{e}_{s_t}), \\ \tilde{e}_{m_t} &= g_m(g_{attn}(e_{m_1}, \dots, e_{m_t}) \oplus \tilde{e}_{s_t} \oplus \tilde{e}_{d_t}),\end{aligned}\quad (6)$$

where  $\oplus$  is the concatenation operation. The next critical involves modeling the mapping from these encoded representations to the variational parameters for the latent health state  $h_t$ . To perform sequential health state inference, we estimate the prior and posterior distributions of  $h_t$  using the prior network  $P_\theta$  and the inference network  $P_\phi$ , respectively. For the prior network, we first incorporate the input representations  $\tilde{e}_{s_t}$  and  $\tilde{e}_{d_t}$  to infer the distribution  $P_\theta(h_t) \sim \mathcal{N}(\mu_{\theta_t}, \Sigma_{\theta_t})$ . The prior variables  $\mu_{\theta_t}$  and  $\log \Sigma_{\theta_t}$  are then estimated via two fully connected layers  $g_{\mu_{\theta_t}}$  and  $g_{\Sigma_{\theta_t}}$ . For the posterior network, the input representations include an additional  $\tilde{e}_{m_t}$  to infer  $P_\phi(h_t) \sim \mathcal{N}(\mu_{\phi_t}, \Sigma_{\phi_t})$ , and posterior variables  $\mu_{\phi_t}$  and  $\Sigma_{\phi_t}$  are similarly estimated via  $g_{\mu_{\phi_t}}$  and  $g_{\Sigma_{\phi_t}}$ . Thus, the prior and inference networks can be specified as:

$$\mu_{\theta_t} = g_{\mu_{\theta_t}}(\tilde{e}_{s_t} \oplus \tilde{e}_{d_t}), \quad \mu_{\phi_t} = g_{\mu_{\phi_t}}(\tilde{e}_{s_t} \oplus \tilde{e}_{d_t} \oplus \tilde{e}_{m_t}), \quad (7)$$

$$\Sigma_{\theta_t} = g_{\Sigma_{\theta_t}}(\tilde{e}_{s_t} \oplus \tilde{e}_{d_t}), \quad \Sigma_{\phi_t} = g_{\Sigma_{\phi_t}}(\tilde{e}_{s_t} \oplus \tilde{e}_{d_t} \oplus \tilde{e}_{m_t}). \quad (8)$$

With the re-parameterization trick, we can draw a latent variable  $z_t$  from  $P_\theta(h_t)$  or  $P_\phi(h_t)$ . To model the impact of previous health state  $h_{t-1}$  on  $h_t$ , we further refine  $z_t$  with an attention mechanism  $f_{attn}$  to derive the final health state:

$$h_t = f_{attn}(z_1, \dots, z_{t-1}, z_t) \quad (9)$$

**Medical event synthesis module.** To make cross-type and longitudinal synthesis for different medical events, we use three separate decoders to generate symptoms, diagnoses, and medications, respectively. Based on the inferred health state  $h_t$ , we can generate the current  $m_t$  as well as  $d_{t+1}$  and  $s_{t+1}$  for the next visit. Specifically, we use three MLPs  $f_{\theta_s}$ ,  $f_{\theta_d}$  and  $f_{\theta_m}$  with a sigmoid function  $\sigma(\cdot)$  to predict the probabilities of each item in these events. Finally, the predicted probabilities  $\hat{m}_t$ ,  $\hat{s}_{t+1}$ ,  $\hat{d}_{t+1}$  can be derived as:

$$\hat{m}_t = \sigma(f_{\theta_m}(h_t \oplus \tilde{e}_{s_t} \oplus \tilde{e}_{d_t})), \quad (10)$$

$$\hat{s}_{t+1} = \sigma(f_{\theta_s}(h_t)), \quad (11)$$

$$\hat{d}_{t+1} = \sigma(f_{\theta_d}(h_t \oplus \tilde{e}_{s_{t+1}})). \quad (12)$$

**Medical report generation module.** Following the recent progress on medical text generation (Yuan et al. 2022), we adopt BioBART as the backbone of our generators to generate various paragraphs of the report  $y_t$  based on  $x_t$  and  $h_t$ . We treat the health state  $h_t$  the learned representations of  $x_t$  as continuous prompts and input them into the BART encoder, and then use the BART decoder to generate corresponding report paragraphs. Following the Fixed-Prompt LM-Tuning strategy (Gao, Fisch, and Chen 2021; Logan IV et al. 2022), we only unfreeze the BART decoder parameters to optimize the generation performance. More importantly, we consider enhancing message passing among the paragraphs by two-phase decoding using a multi-generator deliberation framework. In the draft phase, we sequentially generate three paragraphs  $y_{s_t}$ ,  $y_{d_t}$  and  $y_{m_t}$  conditioned on  $h_t$  and the corresponding event representations. Then, we re-feed these drafts into the BART encoder to obtain its paragraph-level representations  $e_{y_{s_t}}$ ,  $e_{y_{d_t}}$  and  $e_{y_{m_t}}$ . Thus, during the polish phase, we can leverage the feedback information from subsequent generators to refine the generation performance of previous generators. For instance,  $e_{y_{d_t}}$  and  $e_{y_{m_t}}$  are used as new continuous prompts to generate a new paragraph  $\tilde{y}_{s_t}$ . Similarly, for the other two drafts, each of them is passed as feedback to refine the other two paragraphs.

## Model Learning

We optimize the model parameters by applying a neural variational inference algorithm. For a specific patient  $j$  at visit  $t$ , we can compute the reconstructed probabilities for each type of events  $\hat{s}_t^{(j)}$ ,  $\hat{d}_t^{(j)}$ , and  $\hat{m}_t^{(j)}$  based on the medical event synthesis module. We denote  $\mathcal{L}_{s_t}^{(j)}$ ,  $\mathcal{L}_{d_t}^{(j)}$ , and  $\mathcal{L}_{m_t}^{(j)}$  as the binary cross-entropy between generated events and ground-truth events. Next, we can reconstruct the report  $\hat{y}_t^{(j)}$  based on the medical report generation module. We denote  $\mathcal{L}_{y_{s_t}}^{(j)}$ ,  $\mathcal{L}_{y_{d_t}}^{(j)}$ , and  $\mathcal{L}_{y_{m_t}}^{(j)}$  as the language modeling loss between generated and ground-truth paragraphs. Then the overall training loss  $\mathcal{L}$  can be given by:

$$\mathcal{L}_{x_t}^{(j)} = \lambda_s \mathcal{L}_{s_t}^{(j)} + \lambda_d \mathcal{L}_{d_t}^{(j)} + \lambda_m \mathcal{L}_{m_t}^{(j)}, \quad (13)$$

$$\mathcal{L}_{y_t}^{(j)} = \lambda_{y_s} \mathcal{L}_{y_{s_t}}^{(j)} + \lambda_{y_d} \mathcal{L}_{y_{d_t}}^{(j)} + \lambda_{y_m} \mathcal{L}_{y_{m_t}}^{(j)}, \quad (14)$$

$$\mathcal{L}_t^{(j)} = \mathcal{L}_{x_t}^{(j)} + \mathcal{L}_{y_t}^{(j)} + \lambda_{kl} \text{KL}(P_\phi(h_t) || P_\theta(h_t)), \quad (15)$$

$$\mathcal{L} = \sum_{j=1}^N \sum_{t=1}^{T_j} \mathcal{L}_t^{(j)}, \quad (16)$$

Dataset	MIMIC-III				MIMIC-IV			
Method	Unigram	Bigram	Seq Bigram	DimWise	Unigram	Bigram	Seq Bigram	DimWise
LSTM+MLP	-0.441	-2.880	-3.222	0.448	0.316	-0.789	-1.061	0.642
MedGAN	0.115	-2.214	-2.211	0.804	0.537	-0.436	-0.178	0.801
MedBGAN	-0.488	-5.909	-6.001	0.859	0.450	-0.374	-0.202	0.838
MedWGAN	0.742	0.386	0.425	0.813	0.771	0.587	0.581	0.857
EVA	0.077	-5.104	-2.644	0.739	0.532	-1.390	-0.502	0.785
MTGAN	0.804	0.608	0.597	0.869	0.790	0.569	0.591	0.875
PromptEHR	0.795	0.616	0.603	0.921	0.775	0.600	0.588	0.855
MSIC w/o $h_t$	0.605	0.169	-0.015	0.872	0.529	0.115	0.001	0.860
MSIC w/o MV	0.820	0.641	0.552	0.897	0.826	0.670	0.643	0.887
<b>MSIC</b>	<b>0.846</b>	<b>0.701</b>	<b>0.645</b>	<b>0.972</b>	<b>0.838</b>	<b>0.684</b>	<b>0.648</b>	<b>0.956</b>

Table 1: Medical event synthesis results of unigram, same record &amp; sequential visit bigram probabilities, and dim-wise fidelity.

where  $\lambda_*$  are hyper-parameters for the weights of each item. The reconstruction term attempts to bring synthetic data close to real data and the KL term aims to close the distance between the posterior and prior of the health state.

## Experiments

### Experimental Setup

**Datasets.** We use real-world EHR data from publicly available datasets MIMIC-III (Johnson et al. 2016) and MIMIC-IV (Johnson et al. 2020). Following the data processing in (Sun et al. 2022), we can directly obtain the diagnosis and medication codes as their labels, and extract symptoms from the clinical notes. We collect the three most important paragraphs of chief complaint, present illness and prescription from the discharge summary, and desensitize the patient information to construct gold medical reports.<sup>1</sup>

**Baselines.** We compare MSIC to mainstream synthetic EHR generation methods as baselines. For medical event synthesis, the baselines include: (1) **LSTM+MLP** is a straightforward sequential prediction model to yield event probabilities across multiple visits; (2) **MedGAN** (Choi et al. 2017) is a pioneer in using GAN for synthetic EHR generation; (3) **MedBGAN** & **MedWGAN** (Baowaly et al. 2019) take advantage of Boundary-seeking GAN (Hjelm et al. 2018) and Wasserstein GAN (Adler and Lunz 2018) to improve the generation performance; (4) **EVA** (Biswal et al. 2021) is a VAE-based model featuring a BiLSTM encoder and CNN decoder; (5) **MTGAN** (Lu et al. 2023) is a time-series generation model via conditional GAN; and (6) **PromptEHR** (Wang and Sun 2022) employs prompt learning for conditional EHR generation.

Medical report generation is our first proposal and we select some representative pre-trained language models as baselines: (1) **Transformer** (Vaswani et al. 2017), the original vanilla architecture for sequence-to-sequence generation; (2) **GPT-2** (Radford et al. 2019) is a decoder-only autoregressive language model pre-trained on open domain corpora; (3) **BioGPT** (Luo et al. 2022) is recently pre-trained

on large-scale biomedical literature; (4) **BART** (Lewis et al. 2020) is an encoder-decoder model pre-trained in the general domain; and (5) **BioBART** (Yuan et al. 2022) adapts BART to the biomedical domain.

**Evaluation metrics.** We leverage the following metrics to measure the reliability of synthetic EHR data. (1) **Statistical similarity**<sup>2</sup>: We count the frequency for each event in real and synthetic data and compute the coefficient of determination for the two frequency lists, yielding the **Unigram** score. Besides, we consider the co-occurrence frequency of each event pair in the same visit to compute the **Bigram** score, or across two consecutive visits for the **Seq Bigram** score. (2) **Fidelity**: We compute the correlation coefficient of the dimension-wise distributions of each event in real and synthetic records as the **Dimwise** score. For medical report generation, we use text generation-based metrics including **BLEU** (Papineni et al. 2002) and **ROUGE** (Lin 2004) to evaluate n-gram overlap between generated and gold reports.

### Quality of Synthetic Data

**Medical event synthesis.** As shown in Table 1, MedGAN and EVA perform poorly on these metrics, indicating that original generative models cannot capture sufficient information from real records. MedBGAN and MedWGAN slightly improve their performance with new GAN-style training. MTGAN and PromptEHR can synthesize better records based on conditional generation patterns. MSIC can outperform all baselines on both statistical similarity and fidelity, indicating that our graphical model can help improve the quality of synthetic EHR data.

**Medical report generation.** Table 2 shows that BLEU-2 (B@2) and ROUGE-L (R@L) of Transformer are lower than pre-trained models. BioGPT and BioBART, pre-trained on the biomedical-domain corpora, can generate better reports than GPT-2 and BART pre-trained on the open-domain datasets. MSIC outperforms the best baseline BioBART through enhanced message passing among multiple generators and two-phase decoding in our deliberation framework.

<sup>2</sup>This metric is adopted by the official function “r2\_score” in “sklearn.metrics”. Lower (even negative) values indicate greater negative correlations.

<sup>1</sup>Our code and data are available at <https://github.com/plnksnow/MSIC>.

Dataset	MIMIC-III		MIMIC-IV	
Method	B@2	R@L	B@2	R@L
Transformer	21.63	9.21	16.66	13.99
GPT-2	22.17	13.95	19.96	9.32
BioGPT	27.26	23.03	26.57	22.42
BART	27.00	20.78	19.13	20.49
BioBART	29.80	24.39	28.57	27.22
MSIC w/o $h_t$	31.17	26.56	28.72	26.97
MSIC w/o MV	32.94	27.55	28.96	26.92
MSIC w/o MG	29.51	25.13	26.93	25.28
<b>MSIC</b>	<b>33.23</b>	<b>27.79</b>	<b>29.64</b>	<b>27.41</b>

Table 2: Report generation comparison on MIMIC-III.

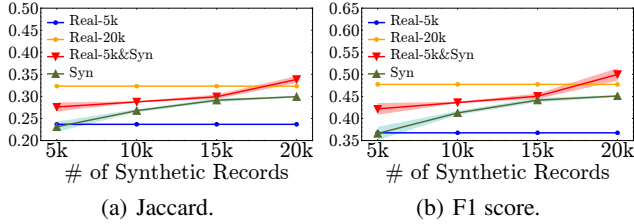
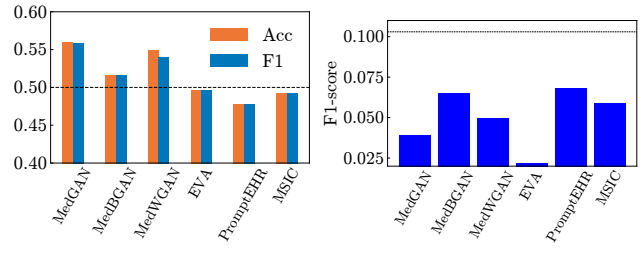


Figure 4: Utility of real, synthetic, and hybrid training data.

**Ablation study.** To verify the effectiveness of the individual components in MSIC, we compare MSIC with several model variants for ablation analysis: (1) **MSIC w/o  $h_t$**  removes the health state and all paths related to it in Figure 2; (2) **MSIC w/o MV** removes the gray arrows in Figure 2 without sequential impact across visits; (3) **MSIC w/o MG** removes the message passing between multiple generators and generates different paragraphs independently. The ablation results in Tables 1 and 2 show that without modeling the patient’s health state, the performance of MSIC w/o  $h_t$  drops a lot, indicating the health state is indeed a key conditional signal for EHR generation. MSIC w/o MV generating records for a single visit is also inferior to MSIC, suggesting that multi-visit inference is necessary for coherent records. MSIC w/o MG performs worst in report generation, showing the effectiveness of our multi-generator network.

### Utility of Synthetic Data

We also measure the utility of synthetic data by comparing the performance of synthetic, real, and hybrid EHR training data. Specifically, we first use various sizes of fully synthetic records for training and compare them with those trained on 5k and 20k fully real records. Next, we create hybrid training data by combining different proportions of synthetic records with 5k real records. Figure 4 indicates that the performance using fully synthetic data lies between real 5k and real 20k. As the size of the synthetic data increases, it performs better than the low-resourced real data, but not as well as real data of the same size. Furthermore, large-size hybrid data can break the upper limit of real data, indicating that augmenting real data with synthetic data can further improve the utility of patient records.



(a) Membership attack.

(b) Attribute attack.

Figure 5: Privacy risk evaluation results on MIMIC-III.

### Privacy Risk Evaluation

**Membership inference attack.** We follow the evaluation of (Liu et al. 2019) for this analysis, which aims to ascertain whether real records similar to synthetic records belong to the training data. We randomly sample 500 records from the training and test sets respectively, and use the training set to train each model to synthesize records. For each real record, Jaccard is used as a distance measure to identify the closest synthetic record and classify it as either training or test data. The ideal metric value is random guessing (i.e. 0.5), with higher scores indicating a greater degree of privacy leakage. The results in Figure 5(a) show that MSIC enhances the protection of training data inferred from synthetic data.

**Attribute inference attack.** We select the most frequently encountered events as common attributes and the rest as sensitive attributes. We identify the closest synthetic patient to each one in the training set based on the overlap in common attributes. Predicted sensitive attributes are identical to those of the closest synthetic patient, and are evaluated to actual sensitive attributes using the F1 Score. “Real data attack” (RDA) serves as an acceptable attribute inference, using real test records instead of synthetic ones. Figure 5(b) show that MSIC achieves an F1 score of 0.059, lower than RDA (horizontal line at 0.095). This indicates that MSIC can withstand this attack without revealing sensitive attributes of real patient data. Compared to MedGAN and EVA lacking pattern acquisition from real data, MSIC can capture more sufficient information with low privacy risks.

### Conclusion

In this work, we propose MSIC, a **Multi-visit health Status Inference model for Collaborative EHR synthesis**. First, we construct a graphical model to capture logical associations among different types of medical events in synthetic EHR generation. We introduce the latent health state as a key conditional signal for generating different events. Then we extend the health state inference method to the multi-visit scenario to improve the coherence of generated records within multiple visits. Furthermore, we propose a new medical report generation task to provide broader applications for synthetic data. The multi-generator deliberation framework and the two-phase decoding strategy enhance collaboration among different generators, thereby improving the overall performance of EHR synthesis.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC Grant No. 62122089), Beijing Outstanding Young Scientist Program NO. BJWZYJH012019100020098, and Intelligent Social Governance Platform, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China, the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China.

## References

- Adler, J.; and Lunz, S. 2018. Banach wasserstein gan. *Advances in neural information processing systems*, 31.
- Alsentzer, E.; Murphy, J.; Boag, W.; Weng, W.-H.; Jindi, D.; Naumann, T.; and McDermott, M. 2019. Publicly Available Clinical BERT Embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, 72–78.
- Baowaly, M. K.; Lin, C.-C.; Liu, C.-L.; and Chen, K.-T. 2019. Synthesizing electronic health records using improved generative adversarial networks. *Journal of the American Medical Informatics Association*, 26(3): 228–241.
- Benitez, K.; and Malin, B. 2010. Evaluating re-identification risks with respect to the HIPAA privacy rule. *Journal of the American Medical Informatics Association*, 17(2): 169–177.
- Biswal, S.; Ghosh, S.; Duke, J.; Malin, B.; Stewart, W.; Xiao, C.; and Sun, J. 2021. EVA: Generating longitudinal electronic health records using conditional variational autoencoders. In *Machine Learning for Healthcare Conference*, 260–282. PMLR.
- Choi, E.; Biswal, S.; Malin, B.; Duke, J.; Stewart, W. F.; and Sun, J. 2017. Generating multi-label discrete patient records using generative adversarial networks. In *Machine learning for healthcare conference*, 286–305. PMLR.
- Cui, L.; Biswal, S.; Glass, L. M.; Lever, G.; Sun, J.; and Xiao, C. 2020. CONAN: complementary pattern augmentation for rare disease detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 614–621.
- De Cao, N.; Izcard, G.; Riedel, S.; and Petroni, F. 2021. Autoregressive Entity Retrieval. In *9th International Conference on Learning Representations*.
- Durall Lopez, R.; Chatzimichailidis, A.; Labus, P.; and Keuper, J. 2021. Combating Mode Collapse in GAN training: An Empirical Analysis using Hessian Eigenvalues. In *Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, volume 4, 211–218. SciTePress.
- El Emam, K.; Buckeridge, D.; Tamblyn, R.; Neisa, A.; Jonker, E.; and Verma, A. 2011. The re-identification risk of Canadians from longitudinal demographics. *BMC medical informatics and decision making*, 11(1): 1–12.
- El Emam, K.; Rodgers, S.; and Malin, B. 2015. Anonymising and sharing individual patient data. *bmj*, 350.
- Franz, L.; Shrestha, Y. R.; and Paudel, B. 2020. A deep learning pipeline for patient diagnosis prediction using electronic health records. *arXiv preprint arXiv:2006.16926*.
- Gao, T.; Fisch, A.; and Chen, D. 2021. Making pre-trained language models better few-shot learners. In *Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL-IJCNLP 2021*, 3816–3830. Association for Computational Linguistics (ACL).
- Gligic, L.; Kormilitzin, A.; Goldberg, P.; and Nevado-Holgado, A. 2020. Named entity recognition in electronic health records using transfer learning bootstrapped neural networks. *Neural Networks*, 121: 132–139.
- Gorinski, P. J.; Wu, H.; Grover, C.; Tobin, R.; Talbot, C.; Whalley, H.; Sudlow, C.; Whiteley, W.; and Alex, B. 2019. Named entity recognition for electronic health records: a comparison of rule-based and machine learning approaches. *arXiv preprint arXiv:1903.03985*.
- Hjelm, R. D.; Jacob, A. P.; Trischler, A.; Che, G.; Cho, K.; and Bengio, Y. 2018. Boundary Seeking GANs. In *International Conference on Learning Representations*.
- Johnson, A.; Bulgarelli, L.; Pollard, T.; Horng, S.; Celi, L. A.; and Mark, R. 2020. Mimic-iv. *PhysioNet*. Available online at: <https://physionet.org/content/mimiciv/1.0/> (accessed August 23, 2021).
- Johnson, A. E. W.; Pollard, T. J.; Shen, L.; Wei, H.; Lehman, L.; Feng, M.; Ghassemi, M. M.; Moody, B.; Szolovits, P.; Celi, L. A.; and Mark, R. G. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3.
- Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C. H.; and Kang, J. 2020. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4): 1234–1240.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7871–7880.
- Li, J.; Cairns, B. J.; Li, J.; and Zhu, T. 2021. Generating Synthetic Mixed-type Longitudinal Electronic Health Records for Artificial Intelligent Applications. *ArXiv*, abs/2112.12047.
- Li, Z.; Niu, C.; Meng, F.; Feng, Y.; Li, Q.; and Zhou, J. 2019. Incremental Transformer with Deliberation Decoder for Document Grounded Conversations. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 12–21.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, G.; Wang, C.; Peng, K.; Huang, H.; Li, Y.; and Cheng, W. 2019. Socinf: Membership inference attacks on social media health data with machine learning. *IEEE Transactions on Computational Social Systems*, 6(5): 907–921.
- Logan IV, R.; Balažević, I.; Wallace, E.; Petroni, F.; Singh, S.; and Riedel, S. 2022. Cutting Down on Prompts and Parameters: Simple Few-Shot Learning with Language Models. In *Findings of the Association for Computational Linguistics: ACL 2022*, 2824–2835.

- Lu, C.; Reddy, C. K.; Wang, P.; Nie, D.; and Ning, Y. 2023. Multi-label clinical time-series generation via conditional gan. *IEEE Transactions on Knowledge and Data Engineering*.
- Luo, R.; Sun, L.; Xia, Y.; Qin, T.; Zhang, S.; Poon, H.; and Liu, T.-Y. 2022. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6).
- Naik, A.; Parasa, S.; Feldman, S.; Wang, L.; and Hope, T. 2022. Literature-Augmented Clinical Outcome Prediction. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 438–453.
- Neamatullah, I.; Douglass, M. M.; Lehman, L.-W. H.; Reiser, A.; Villarroel, M.; Long, W. J.; Szolovits, P.; Moody, G. B.; Mark, R. G.; and Clifford, G. D. 2008. Automated de-identification of free-text medical records. *BMC medical informatics and decision making*, 8(1): 1–17.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Peng, Y.; Chen, Q.; and Lu, Z. 2020. An Empirical Study of Multi-Task Learning on BERT for Biomedical Text Mining. *BioNLP 2020*, 205.
- Quimbaya, A. P.; Múnera, A. S.; Rivera, R. A. G.; Rodríguez, J. C. D.; Velandia, O. M. M.; Peña, A. A. G.; and Labbé, C. 2016. Named entity recognition over electronic health records through a combined dictionary-based approach. *Procedia Computer Science*, 100: 55–61.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rashidian, S.; Wang, F.; Moffitt, R.; Garcia, V.; Dutt, A.; Chang, W.; Pandya, V.; Hajagos, J.; Saltz, M.; and Saltz, J. 2020. SMOOTH-GAN: towards sharp and smooth synthetic EHR data generation. In *Artificial Intelligence in Medicine: 18th International Conference on Artificial Intelligence in Medicine, AIME 2020, Minneapolis, MN, USA, August 25–28, 2020, Proceedings 18*, 37–48. Springer.
- Salakhutdinov, R. 2015. Learning deep generative models. *Annual Review of Statistics and Its Application*, 2: 361–385.
- Shang, J.; Xiao, C.; Ma, T.; Li, H.; and Sun, J. 2019. Gamenet: Graph augmented memory networks for recommending medication combination. In *proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1126–1133.
- Sondhi, P.; Sun, J.; Tong, H.; and Zhai, C. 2012. Symp-Graph: a framework for mining clinical notes through symptom relation graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1167–1175.
- Sun, H.; Xie, S.; Li, S.; Chen, Y.; Wen, J.-R.; and Yan, R. 2022. Debaised, Longitudinal and Coordinated Drug Recommendation through Multi-Visit Clinic Records. *Advances in Neural Information Processing Systems*, 35: 27837–27849.
- Sun, S.; Wang, F.; Rashidian, S.; Kurc, T.; Abell-Hart, K.; Hajagos, J.; Zhu, W.; Saltz, M.; and Saltz, J. 2021. Generating Longitudinal Synthetic EHR Data with Recurrent Autoencoders and Generative Adversarial Networks. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2021 and DMAH 2021, Virtual Event, August 20, 2021, Revised Selected Papers 7*, 153–165. Springer.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, B.; Xie, Q.; Pei, J.; Tiwari, P.; Li, Z.; et al. 2021. Pre-trained language models in biomedical domain: A systematic survey. *arXiv preprint arXiv:2110.05006*.
- Wang, Z.; and Sun, J. 2022. PromptEHR: Conditional Electronic Healthcare Records Generation with Prompt Learning. In *Conference on Empirical Methods in Natural Language Processing*.
- Xia, Y.; Tian, F.; Wu, L.; Lin, J.; Qin, T.; Yu, N.; and Liu, T.-Y. 2017. Deliberation networks: Sequence generation beyond one-pass decoding. *Advances in neural information processing systems*, 30.
- Xu, J.; Li, H.; and Zhou, S. 2015. An overview of deep generative models. *IETE Technical Review*, 32(2): 131–139.
- Yan, C.; Zhang, Z.; Nyemba, S.; and Malin, B. A. 2020. Generating electronic health records with multiple data types and constraints. In *AMIA annual symposium proceedings*, volume 2020, 1335. American Medical Informatics Association.
- Yan, H.; Gui, T.; Dai, J.; Guo, Q.; Zhang, Z.; and Qiu, X. 2021. A Unified Generative Framework for Various NER Subtasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 5808–5822.
- Yang, C.; Xiao, C.; Ma, F.; Glass, L.; and Sun, J. 2021. SafeDrug: Dual Molecular Graph Encoders for Safe Drug Recommendations. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*.
- Yuan, H.; Yuan, Z.; Gan, R.; Zhang, J.; Xie, Y.; and Yu, S. 2022. BioBART: Pretraining and Evaluation of A Biomedical Generative Language Model. In *Proceedings of the 21st Workshop on Biomedical Language Processing*, 97–109.
- Zhang, Z.; Li, M.; and Yu, J. 2018. On the convergence and mode collapse of GAN. *SIGGRAPH Asia 2018 Technical Briefs*.
- Zhang, Z.; Yan, C.; Lasko, T. A.; Sun, J.; and Malin, B. A. 2021. SynTEG: a framework for temporal structured electronic health data simulation. *Journal of the American Medical Informatics Association*, 28(3): 596–604.