Self-Supervised Adversarial Distribution Regularization for Medication Recommendation

Yanda Wang[1] , **Weitong Chen**[2] , **Dechang Pi**[1*] , **Lin Yue**[2 3 †] , **Sen Wang**[2] , **Miao Xu**[2]
[1]Nanjing University of Aeronautics and Astronautics, Nanjing, 211106 China
[2]The University of Queensland, Brisbane QLD 4072 Australia
[3]Northeast Normal University, Changchun, 130024 China
{yandawang, dc.pi}@nuaa.edu.cn, {w.chen9, l.yue, sen.wang, miao.xu}@uq.edu.au,

Abstract

Medication recommendation is a significant healthcare application due to its promise in effectively prescribing medications. Avoiding fatal side effects related to Drug-Drug Interaction (DDI) is among the critical challenges. Most existing methods try to mitigate the problem by providing models with extra DDI knowledge, making models complicated. While treating all patients with different DDI properties as a single cohort would put forward strict requirements on models' generalization performance. In pursuit of a valuable model for a safe recommendation, we propose the Self-Supervised Adversarial Regularization Model for Medication Recommendation (SARMR). SARMR obtains the target distribution associated with safe medication combinations from raw patient records for adversarial regularization. In this way, the model can shape distributions of patient representations to achieve DDI reduction. To obtain accurate self-supervision information, SARMR models interactions between physicians and patients by building a key-value memory neural network and carrying out multi-hop reading to obtain contextual information for patient representations. SARMR outperforms all baseline methods in the experiment on a real-world clinical dataset. This model can achieve DDI reduction when considering the different number of DDI types, which demonstrates the robustness of adversarial regularization for safe medication recommendation.

1 Introduction

Technological innovations in deep learning have achieved great success in various clinical applications[Yue *et al.*, 2020], such as disease diagnosis, onset prediction, and physiological condition monitoring [Chen *et al.*, 2018; Li *et al.*, 2019; Malakouti and Hauskrecht, 2019; Wang *et al.*, 2019; Chen *et al.*, 2019], which enable physicians to better diagnose and treat patients since the beginning of the professional practice of medicine. As one of the hottest research topics, medication recommendation aims to assist physicians in making effective prescriptions according to Electronic Healthcare Records (EHRs) that describe patients in terms of vital signs, diagnoses, and procedures. A comprehensive analysis of EHRs based on deep learning provides physicians with insight regarding patients. It helps caregivers cope with peril brought by the fast-paced way of working in ICUs, significantly improving their quality of life.

Deep learning has been adopted for different medication recommendation tasks. Many of these existing methods make great efforts to obtain accurate patient representations to carry out practical medication recommendations, and widely applied approaches include instance-based and longitudinal methods [Lipton *et al.*, 2016; Shang *et al.*, 2019a; Wang *et al.*, 2017]. However, these methods neglect the fatal side effects related to Drug-Drug Interaction (DDI) due to duplication, antagonism, and alternation [Zhang *et al.*, 2017]. Meanwhile, limited researches on this topic apply extra knowledge about DDI to mitigate the problem [Zhang *et al.*, 2017; Ma *et al.*, 2018; Shang *et al.*, 2019b], which requires specific components for knowledge extraction or information fusion, making models complicated and introducing bias that has adverse impacts on the recommendation. Besides, these methods treat patients with different DDI rates, i.e., the fraction of combinations that lead to DDI to the total number of varieties in a set of medications, as a single cohort, which would put strict requirements forwards on the generalization performance of models. However, representations of these patients could negatively affect each other when used to train models due to the different probability distributions they follow.

To address the above problems, we propose a novel **S**elf-Supervised **A**dversarial **R**egularization Model for **M**edication **R**ecommendation (SARMR). Instead of introducing extra knowledge about DDI, SARMR obtains probability distributions of patient representations related to safe medication combinations in the feature space from raw EHRs. Then SARMR applies the knowledge as the true data to adversarially regularize distributions of patient representations with achieving DDI reduction. Patients with different DDI rates are respectively used and regularized as different cohorts, and the adverse impacts on generalization when they are treated as a single cohort could be avoided.

Key challenges for building the model include obtain-

*Contact Author1
†Contact Author2

ing probability distributions that reflect low DDI rates and maintaining recommendation performance while shaping distributions. Discovering meaningful embeddings representing desired probability distributions over data is a significant promise of deep learning. The continuous feature space enables the gradient-based optimization to search for distributions that reflect desired properties [Gómez-Bombarelli *et al.*, 2018; Blaschke *et al.*, 2018; Makhzani *et al.*, 2015; Pan *et al.*, 2018]. However, previous methods require extra tasks for the knowledge-based jump in the latent space, which may introduce bias or result in inaccurate search results. In addition, the DDI rate is not a direct output of patient representations but calculated based on the recommended medications. The lack of direct mapping between latent features and DDI rates makes property-oriented methods unavailable. Also, patients suffer from different diseases, and their representations contain other conditions. Focusing on regularizing distributions for desired DDI rates alone may ignore critical information encoded in patients' representations and lead to poor performance on the recommendation.

SARMR overcomes these obstacles in two steps. In the first step, SARMR uses patient records whose DDI rates are lower than a threshold of D_{rec} to recommend medications, and obtains patient representations related to low DDI rates as the true data for the adversarial regularization in the second step. SARMR firstly encodes EHRs with GRUs, and builds a key-value Memory Neural Network (MemNN) [Miller *et al.*, 2016], whose keys are representations of admissions and values are corresponding medications, to model interactions between physicians and patients. Then SARMR uses the representation of the last admission as query to conduct multi-reading on the MemNN, while Graph Convolutional Network (GCN) [Kipf and Welling, 2016] is presented as the embedding module of the read results. The updated query is used to make the recommendation. Then in the second step, SARMR uses records of all patients regardless of their DDI rates to jointly conducts medication recommendation and adversarial distribution regularization with the Generative Adversarial Network (GAN) [Goodfellow *et al.*, 2014] based on the obtained representations in the first step, so that both effective medication combinations and DDI reduction are achieved.

Our contributions could be summarized as follows:

- We propose a self-supervised strategy to shape distributions of patient representations for safe medication recommendation, which obtains prior distribution reflecting low DDI rates from raw EHRs and conducts adversarial distribution regularization with GAN.
- A key-value MemNN is constructed to learn interactions between doctors and patients by carrying out multi-hop reading on the MemNN, so that contexts in historical EHRs are derived for informative patient features.
- Experimental results demonstrate that SARMR outperforms all the baseline methods, and it can make effective medication recommendation with an F1 at 0.6608, while it also achieves a DDI reduction at -2.72%.

2 Related Work

2.1 Distribution Regularization

Discovering informative embeddings representing meaningful distributions over data is a significant promise of deep learning. The continuous latent space enables models to search for distributions that reflect desired properties. For example, [Lim *et al.*, 2018] incorporate molecular properties into latent features when using conditional variational autoencoder to generate molecules with desired properties. [Gómez-Bombarelli *et al.*, 2018] add a regression task to autoencoder (AE) for property prediction to connect latent features with target properties, searching for desired attribute in the feature space. And [Blaschke *et al.*, 2018] extend the idea by applying Bayesian optimization on the obtained latent space to find new molecular structures with target properties. [De Cao and Kipf, 2018] further utilize reinforcement learning to carry out optimization towards desired chemical properties.

Distribution Regularization for desired properties has also been discussed before. [Makhzani *et al.*, 2015] match the aggregated posterior of hidden vectors of AE with an arbitrary prior distribution to generate meaningful samples. [Kadurin *et al.*, 2017] apply the strategy to drug discovery, and introduce a neuron responsible for reflecting desired properties to develop new molecules. Meanwhile, [Pan *et al.*, 2018] use adversarial regularization to match latent representations of graphs to match a prior distribution, so that meaningful graph embeddings in a continuous vector space are achieved.

SARMR differs from these methods since no extra task is required to guide the search in the latent space. Instead, a prior distribution directly relating to low DDI rates is obtained from raw EHRs for adversarial regularization.

2.2 Medication Recommendation

To model relations covering multiple input views in EHRs, [Le *et al.*, 2018] use memory augmented neural networks to achieve better performance in drug prescription task and disease progression task. But the method does not take relationships between drugs into account, so [Shang *et al.*, 2019b] build a graph based on the co-occurrence of drugs in EHRs, and uses GNN to obtain embeddings for those drugs. Meanwhile, [Zhang *et al.*, 2017] address medication recommendation from the view of sequential decision-making, and use a recurrent decoder to model drug dependency, while drug-to-disease mapping is modeled by content-based attention. But these methods concentrate on visit-level temporal information, so [Choi *et al.*, 2016] present a two-level neural attention model to model both visit-level and variable-level sequential information, which also provides a detailed interpretation of the prediction results. Most of these proposed methods ignored EHRs related to patients with a single visit, so [Shang *et al.*, 2019a] pre-train its transformer-based visit encoder on EHR data from patients with a single visit, then the model is fine-tuned on EHRs of patients with multiple re-admissions.

Compared with previous methods, SARMR applies a self-supervised strategy to obtain target distribution from EHRs for the adversarial distribution regularization. And SARMR conducts multi-hop reading on the key-value MemNN to contextual information for medication recommendations.

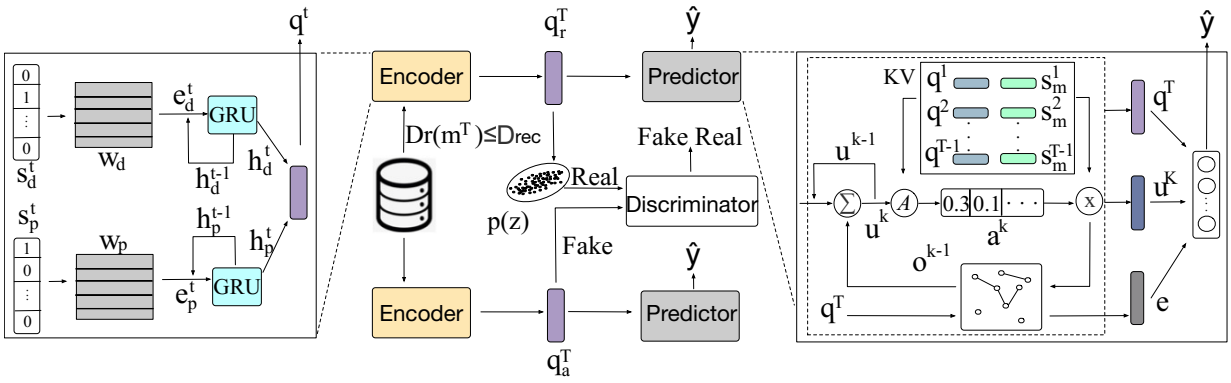

Figure 1: The framework of SARMR. Records of patients whose DDI rates $Dr(m^T) \leq D_{rec}$ are firstly used to train the encoder and predictor for medication recommendation as shown in the upper branch. Then the corresponding patient representations q_r^T are used to fit the Gaussian distribution $p(z)$, which would generate true data for the adversarial regularization. After that, all patient records are used for training, while a discriminator (a multi-layer perceptron) is combined with the encoder to form a GAN model to shape q_a^T for DDI reduction.

3 Proposed Method

Details of SARMR are described in this section. SARMR firstly selects patient records by a threshold D_{rec} for DDI rate, and uses the data to train the Medication Recommendation (MedRec) module that consists of a encoder and a predictor. Then SARMR carries out adversarial regularization on the distributions of all patient representations based on the informative patterns obtained by the encoder to achieve DDI reduction. The overview is demonstrated in Figure 1.

3.1 Notations

EHRs describe patients in terms of different views, and each patient is recorded as a sequence of multivariate observations $X = [x_1, x_2, ..., x_T]$, where T is the number of admissions, and x_t represents records of the t^{th} admission. x_t consists of three sets of medical codes, including diagnoses d^t, procedures p^t, and medications m^t, which belong to the medical code sets S_d, S_p, and S_m respectively, and the corresponding numbers of distinct codes in EHRs are $|S_*|$. These records are transformed into multi-hot vectors $s_d^t \in \mathbb{R}^{|S_d|}$, $s_p^t \in \mathbb{R}^{|S_p|}$ and $s_m^t \in \mathbb{R}^{|S_m|}$ to act as inputs of SARMR.

A drug graph $DG \in \mathbb{R}^{|S_m|*|S_m|}$ is constructed based on all medication combinations m^* to indicate whether two drugs i and j have been prescribed to the same patient, if so, $DG_{i,j} = DG_{j,i} = 1$, otherwise $DG_{i,j} = DG_{j,i} = 0$. Similarly, a DDI graph $IG \in \mathbb{R}^{|S_m|*|S_m|}$ is constructed to indicate polypharmacy interactions with drug pairs or higher-order drug combinations, and $IG_{i,j} = IG_{j,i} = 1$ if medication i and j would cause side effects, otherwise $IG_{i,j} = IG_{j,i} = 0$. Given IG, the DDI rate $Dr(m^t) \in [0, 1]$ for m^t could be calculated as Eq.(1) shows.

$$Dr(m^t) = \frac{\sum_{i,j} |(m_i^t, m_j^t) \in m_t \; and \; IG_{i,j} = 1|}{\sum_{i,j} 1} \tag{1}$$

SARMR uses a threshold D_{rec} for DDI rate to select patients P_{rec} with $Dr(m^T) \leq D_{rec}$, and uses their records to train MedRec to obtain the Gaussian distribution $p(z)$ for adversarial regularization. Given the medical records of a patient $x_t (t < T)$ as well as d^T and p^T, the goal of SARMR is to predict medications m^T.

3.2 Medication Recommendation

The encoder uses two GRUs to obtain temporal information from s_d^t and s_p^t, where $t \leq T$. SARMR firstly embeds s_d^t and s_p^t into continuous representations with embedding matrices $W_d \in \mathbb{R}^{|S_d| \times dim}$ and $W_p \in \mathbb{R}^{|S_p| \times dim}$, where dim is the embedding size that determined with Gaussian Process as a hyper-parameter. Then the model obtains hidden states for each time stamp as Eq.(2) and Eq.(3) show, and combines them with a linear embedding layer to generate the representation q^t for the t^{th} admission.

$$h_d^t = GRU_d(h_d^{t-1}, e_d^t) \tag{2}$$

$$h_p^t = GRU_p(h_p^{t-1}, e_p^t) \tag{3}$$

After encoding patient representations, SARMR carries out multi-hop reading on a key-value MemNN KV to obtain contextual information from interactions between physicians and patients, and uses GCN to transform the results into continuous embeddings.

The MemNN KV uses $q^t (t < T)$ as keys and corresponding ground truth medications $s_m^t (t < T)$ as values. For each hop k, SARMR calculates an attention weight a^k between u^{k-1} and KV_{key}^{T-1} as Eq.(4) shows, where u^{k-1} is the output of last hop, $q^* \in KV_{keys}^{T-1}$, and W_a is a weight matrix. Then SARMR reads weighted memories o^k from KV as Eq.(5) shows, where Z^k are medication embeddings generated by GCN in the k hop, and s_m^* are values in KV_{value}^{T-1}. Given o^k, the query u^{k-1} is updated to u^k according to Eq.(6). The reading process would be repeated for K hops to get u^K.

$$a^k = Softmax(q^* W_a (u^{k-1})^\intercal) \tag{4}$$

$$o^k = a^k s_m^* Z^k \tag{5}$$

$$u^k = u^{k-1} + o^k \tag{6}$$

For the calculation of Z^k, the drug graph DG is processed following GCN procedures as Eq.(7) shows [Kipf and Welling, 2016], where $\hat{D}$ is a diagonal matrix such that $\hat{D}_{ii} = \sum_j DG_{ij}$ and I is a identity matrix. SARMR applies a two-layer GCN on DG to obtain embeddings for drugs as Eq.(8) shows, where W_e is the medication embeddings from DG, and W_1 is the hidden weight parameter matrix. Finally, the correlation between q^T and Z^K is calculated according to Eq.(9), where W_e is a weight matrix for the attention.

$$\hat{A} = \hat{D}^{-1/2}(DG + I)\hat{D}^{-1/2} \tag{7}$$

$$Z = \hat{A}tanh(\hat{A}W_e)W_1 \tag{8}$$

$$e = Softmax(q^T W_e (Z^K)^\intercal) Z^K \tag{9}$$

SARMR predicts the final result as Eq.(10) shows. The equation takes the concatenation of patient representation q^T, multi-hop reading result u^K, and weighted embeddings of medication e as inputs, where $S(.)$ is a sigmoid function and $f(.)$ is a fully connected layer.

$$\hat{y} = S(f([q^T, u^K, e])) \tag{10}$$

3.3 Self-Supervised Adversarial Regularization

The key idea of SARMR is to adversarially match latent features of patients to a prior distribution reflecting low DDI rates, which is obtained in a self-supervised way that requires no extra knowledge, so that patients would be prescribed safe medication combinations. To to so, given the threshold D_{rec}, patients whose DDI rates $Dr(m^T) \leq D_{rec}$ are used to train the MedRec module formed by the encoder and predictor described previously, and SARMR uses the Gaussian distribution $p(z)$ that q_r^T follows [Kadurin *et al.*, 2017; Pan *et al.*, 2018], which represents the distribution over low DDI rates, as the source of real data in the adversarial regularization. Specifically, the mean and covariance matrix of all q_r^T are calculated to get the fitted Gaussian distribution $p(z)$. Then all patients regardless of their DDI rates are used to train the model, and their representations q_a^T contains two parts: q_f^T related to patients whose $Dr(m^T) \in (D_{rec}, 1]$, which follows $p_f(z)$, and ${q'}_r^T$ that follows $p'(z)$, which comes from q_r^T but has been affected by q_f^T.

Patient representations q_a^T obtained by the encoder are adversarially shaped to match $p_f(z)$ to $p(z)$, while $p'(z)$ is also corrected back to $p(z)$, so that they would all follow the same distribution and present the desired low DDI rates. To achieve the goal, a discriminator is attached on top of encoder to form a GAN model, and q_a^T will act as fake data while samples from $p(z)$ would act as real data. Meanwhile, patients suffer from different diseases and their representations contain information of various conditions as well as medications. It is critical to guarantee that the information is mapped to accurate medication combinations when representations are regularized. Thus, SARMR jointly train the GAN model and the MedRec within each mini-batch in two phases: the prediction phase and the regularization phase.

In the prediction phase, the MedRec module is updated to minimize the loss between predicted medications and ground truth medications, so that the prediction performance is improved. In the regularization phrase, SARMR updates the discriminator to distinguish real data generated by the target distribution from features generated by the encoder. Once the training procedure is done, patient representations q_a^T are regularized to follow $p(z)$, and these representations could still be mapped to corresponding medications to finish the recommendation by the predictor. Thus, both the expected low DDI rate and prediction performance are guaranteed.

3.4 Training and Inference

Following the joint training strategy, there are two loss functions for SARMR. The GAN regularization process makes use of the typical loss function for GAN model [Goodfellow *et al.*, 2014]. The MedRec module is updated to minimize the weighted loss for medication recommendation shown in Eq.(11) [Shang *et al.*, 2019b], which consists of the binary cross entropy loss $\mathcal{L}_b$ and multi-label margin loss $\mathcal{L}_m$. $\mathcal{L}_m$ is used to make the predicated probability of ground truth medications have at least 1 margin larger than others, which would benefit the predictive performance of SARMR.

Here, $y \in \mathbb{R}^{|S_m|}$ is s_m^T, i.e., the ground truth of medications for the final admission, $\hat{y} \in \mathbb{R}^{|S_m|}$ is the predict result in the form of probability, and $\hat{Y}$ is the predict medication set, so $\hat{y}_i$ and $\hat{y}[\hat{Y}_j]$ are the probabilities that the i^{th} and $\hat{Y}_j^{th}$ medication would be prescribed to the patient. The constraints for the weights α_1 and α_2 are that $\alpha_1 > 0$, $\alpha_2 > 0$ and $\alpha_1 + \alpha_2 = 1$.

$$\mathcal{L} = \alpha_1 \mathcal{L}_b + \alpha_2 \mathcal{L}_m \tag{11}$$

$$\mathcal{L}_b = -\sum_{i=1}^{|S_m|} [y_i log\sigma(\hat{y}_i) + (1 - y_i) log(1 - \sigma(\hat{y}_i))] \tag{12}$$

$$\mathcal{L}_m = \sum_{i=1}^{|S_m|} \sum_{j=1}^{|\hat{Y}|} \frac{max(0, 1 - (\hat{y}[\hat{Y}_j] - \hat{y}[i]))}{L} \tag{13}$$

4 Experiment

In this section, SARMR is compared with different baseline methods on the real-world clinical dataset MIMIC-III v1.4 [Johnson *et al.*, 2016]. The model is implemented with PyTorch and trained on a NVIDIA TITAN Xp GPU, and more information about source code could be found at Github [1].

4.1 Dataset

Patients with at least two admissions to hospitals are selected for experiments, and all their diagnoses and procedures are

[1] https://github.com/yanda-wang/SARMR

mapped to fixed vocabularies to act as inputs, while medications prescribed during the first 24-hour in hospitals are used as ground truth for the recommendation [Shang *et al.*, 2019b]. Totally 6350 patients are selected, and numbers of diagnoses, procedures and medications contained in their records are 1960, 1432, and 153. The top-100 severe DDI types are used to build IG, and Kernel Density Estimate on DDI rates indicates a Gaussian distribution whose mean is 0.4, so D_{rec}=0.4.

4.2 Baselines

The following methods are compared with SARMR.

- **Leap [Zhang *et al.*, 2017]:** Leap addresses medication recommendation as a sequential decision-making task, and uses reinforcement learning to improve accuracy.
- **RETAIN [Choi *et al.*, 2016]:** RETAIN uses a two-level attention model to detects influential past visits for sequential prediction.
- **DMNC [Le *et al.*, 2018]:** DMNC models multi-view interactions and long-term dependencies via memory augmented neural network to recommend medications.
- **GAMNet [Shang *et al.*, 2019b]:** GAMENet uses a dynamic memory network to model historical EHRs, and applies GCN on DDI graph to achieve DDI reduction.
- **MedRec:** MedRec contains the encoder and predictor of SARMR, but without adversarial regularization.

4.3 Metrics

The performance of SARMR is evaluated using Jaccard Similarity (Jaccard), Precision Recall AUC (PRAUC), Average F1 (F1), and changes of average DDI rate ($\triangle DDI$). Among these metrics, $\triangle DDI$, as shown in Eq.(14), indicates the difference between DDI rates of predicted results and ground truth, which shows whether DDI is reduced.

$$\triangle DDI = \frac{1}{N}\sum_{i=1}^{N}\frac{DDI\ rate(\hat{Y}^i) - DDI\ rate(Y^i)}{DDI\ rate(Y^i)} \quad (14)$$

4.4 Evaluation

As shown in Table 1, SARMR achieves the best $\triangle DDI$ at -2.72% and succeeds in attaining DDI reduction. Meanwhile, the figure for MedRec is 3.04%, showing that the GAN model has successfully regularized the distribution of patient representations to reduce the DDI rate. GAMENet uses extra DDI knowledge as a memory component to reduce DDI, and the $\triangle DDI$ is 3.06%. The comparison proves that regularizing distributions of representations could be a more effective strategy for DDI reduction. The remaining methods have much higher DDI rates and fail to predict safe medications.

Besides, SARMR outperforms all the baselines with the highest Jaccard, PRAUC, and F1 at 0.5039, 0.7688, and 0.6608 respectively. Among those baselines, DMNC uses an encoder similar to SARMR, and its sub-optimal performance indicates that key-value MemNN plays an important role in modeling EHRs. GAMENet uses a similar key-value MemNN without multi-hop reading, and its lower figures on

Methods	$\triangle DDI$	Jaccard	PRAUC	F1
Leap	7.80%	0.4484	0.6457	0.6109
DMNC	5.93%	0.4933	0.7269	0.6511
RETAIN	10.43%	0.4897	0.7499	0.6494
GAMENet	3.06%	0.4970	0.7589	0.6544
MedRec	3.04%	0.4945	0.7635	0.6519
SARMR	**-2.72%**	**0.5039**	**0.7688**	**0.6608**

Table 1: Performance Comparisons of Different Methods

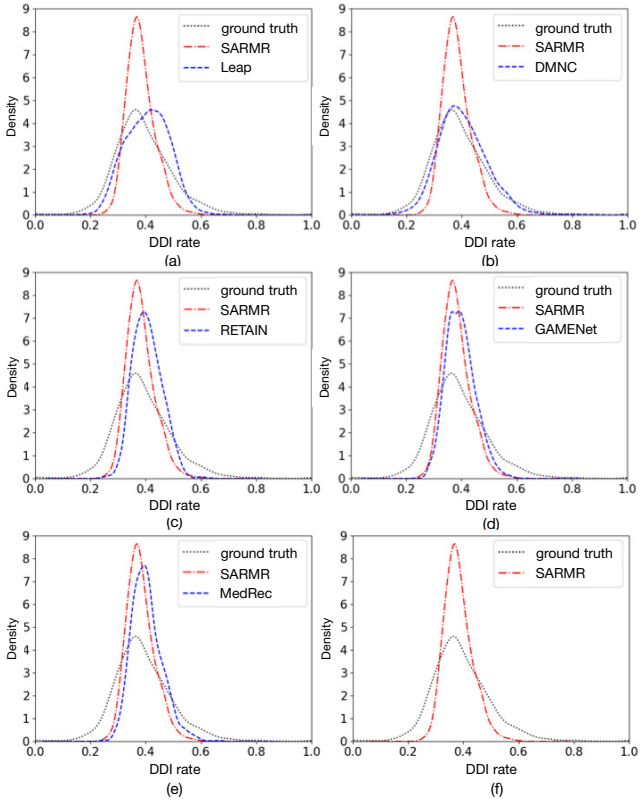

Figure 2: Comparisons between SARMR and Baselines in terms of Kernel Density Estimate on DDI Rates.

metrics show that reasoning over interactions between patients and doctors are essential.

To further illustrate the effectiveness of adversarial regularization, the kernel density estimate is conducted on DDI rates of predict results and ground truth, which are illustrated in Figure 2. As shown in Figure 2.(a) and Figure 2.(b), density curves of DDI rates related to Leap and DMNC shift to the right compared with that of ground truth, which explains why these two methods lead to high DDI rates in the predict results. Similarly, the curve for RETAIN in Figure 2.(c) also shifts to the right, but the density has much higher values at high DDI rates with small variance, and that is why RETAIN has the highest $\triangle DDI$. The results of RETAIN, GAMENet, and MedRec show similar trends, and the densities of DDI rates in their predict results have small variance while the means are more extensive than that of the ground truth. Meanwhile, the density of SARMR has the highest val-

#DDI type	DDI rate	Leap	RETAIN	DMNC	GAMENet	MedRec	SARMR
40	0.0857	-17.80%	8.13%	3.50%	-1.08%	5.43%	-4.15%
60	0.1941	-1.19%	6.91%	11.61%	3.35%	6.12%	-0.60%
80	0.2932	5.69%	7.97%	8.23%	2.62%	3.10%	-3.27%
100	0.3923	7.80%	5.93%	10.43%	3.06%	3.04%	-2.72%

Table 2: Comparisons of $\triangle$ DDI of Different Methods in Terms of Different DDI types

ues in low DDI rates, indicating that the regularization has successfully matched distributions of patient representations to the desired Gaussian distribution to achieve DDI reduction.

Besides, how the number of DDI types affects different methods is explained to show the robustness of adversarial regularization on DDI reduction. In addition to the top-100 DDI types, the top-40, top-60, and top-80 types are considered. The results are shown in Table 2. SARMR can achieve DDI reduction regardless of the number of DDI types. Meanwhile, when using the top-40 DDI types, GAMENet succeeds in reducing DDI rates and achieves a $\triangle DDI$ at -1.08%, but the method fails to do so when more types are considered. The figure for Leap drops dramatically from 7.80% to -17.80% as the number of DDI types decreases from 100 to 40, while DMNC and RETAIN cannot achieve DDI reduction in any case. The results demonstrate the robustness of SARMR since it is the only method that leads to a DDI reduction when a different number of DDI types are used.

Ablation Study

To evaluate the effectiveness and necessity of each component of SARMR, the model is deconstructed by replacing or removing these components to build variants as follows:

- **S-GRU:** The variant replaces the two GRUs in SARMR with a Single GRU to model patient representations.
- **S-Encoder:** The variant removes the MemNN and uses the encoder's output to recommend medications directly.
- **No-GCN:** The GCN that embeds medications in SARMR is replaced with an embedding matrix.
- **MedRec:** MedRec removes the discriminator in SARMR and it is not adversarially regularized.

The results in Table 3 and Figure 3 shows that SARMR achieves the best performance. $\triangle DDI$ for MedRec is 3.04%, showing that the GAN model is essential to reduce the DDI rate. Otherwise, the density curve would shift to the right as shown in Figure 3.(d). No-GCN has the second-best performance on Jaccard and F1 as well as the most similar density curve with SARMR, indicating that the clinically meaningful embeddings of medications obtained by GCN could assist SARMR in achieving better performance. Meanwhile, S-Encoder has a dramatically high $\triangle DDI$ at 5.34%, and the density has higher values in high DDI rate, which proves that without an appropriate decoder to interpret information from the encoder, the regularization on patient representations may even lead to an adverse impact on the performance. S-GRU has the lowest Jaccard and F1. The results show that temporal patterns in diagnoses and procedures should be modeled separately and then combined to obtain a comprehensive representation rather than directly treated as a whole.

Methods	$\triangle DDI$	Jaccard	PRAUC	F1
S-GRU	0.75%	0.4923	0.7670	0.6499
S-Encoder	5.34%	0.4961	0.7646	0.6534
No-GCN	-0.15%	0.5026	0.7613	0.6559
MedRec	3.04%	0.4945	0.7635	0.6519
SARMR	**-2.72%**	**0.5039**	**0.7688**	**0.6608**

Table 3: Comparisons of Different Variants of SARMR

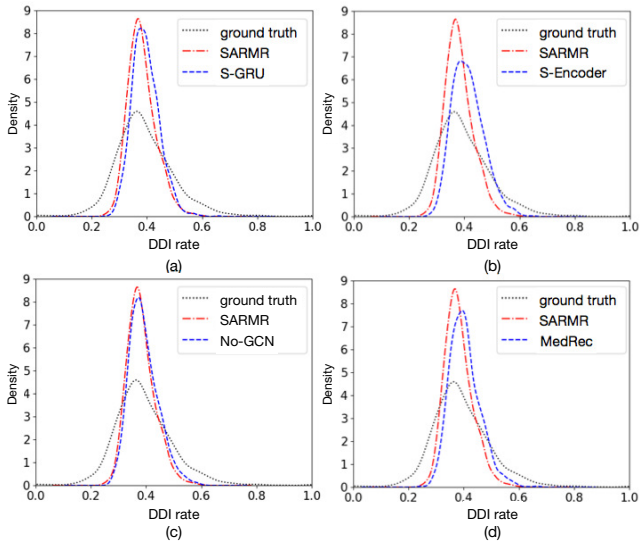

Figure 3: Comparisons of Different Variants of SARMR in terms of Kernel Density Estimate.

5 Conclusion

In this paper, we propose a novel self-supervised adversarial distribution regularization strategy SARMR for safe medication recommendation. Existing methods treat distributions related to different DDI rates as a single cohort, and requires extra tasks to apply knowledge for DDI reduction. SARMR conducts multi-hop reading on MemNN to derive contextual information from EHRs, and adversarially regularizes patient representations based on desired distributions obtained from raw EHRs to achieve DDI reduction. For now, SARMR applies a fixed number of hops for all patients, while their EHRs contain different amounts of information. Given the potential for facilitating accurate information extraction, we expect adaptively determining the number of hops in future work.

Acknowledgments

This research has been supported by the NSFC under Grant No.61972384 and CPSF under Grant No. 2017M621192.

References

[Blaschke *et al.*, 2018] Thomas Blaschke, Marcus Olivecrona, Ola Engkvist, Jürgen Bajorath, and Hongming Chen. Application of generative autoencoder in de novo molecular design. *Molecular informatics*, 37(1-2):1700123, 2018.

[Chen *et al.*, 2018] Weitong Chen, Sen Wang, Guodong Long, Lina Yao, Quan Z Sheng, and Xue Li. Dynamic illness severity prediction via multi-task rnns for intensive care unit. In *ICDM*, pages 917–922. IEEE, 2018.

[Chen *et al.*, 2019] Weitong Chen, Guodong Long, Lina Yao, and Quan Z Sheng. Amrnn: attended multi-task recurrent neural networks for dynamic illness severity prediction. *World Wide Web*, pages 1–18, 2019.

[Choi *et al.*, 2016] Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter F Stewart. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *NIPS*, pages 3504–3512, 2016.

[De Cao and Kipf, 2018] Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*, 2018.

[Gómez-Bombarelli *et al.*, 2018] Rafael Gómez-Bombarelli, Jennifer N Wei, and David Duvenaud *et al.* Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.

[Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.

[Johnson *et al.*, 2016] Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

[Kadurin *et al.*, 2017] Artur Kadurin, Alexander Aliper, Andrey Kazennov, Polina Mamoshina, Quentin Vanhaelen, Kuzma Khrabrov, and Alex Zhavoronkov. The cornucopia of meaningful leads: Applying deep adversarial autoencoders for new molecule development in oncology. *Oncotarget*, 8(7):10883, 2017.

[Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[Le *et al.*, 2018] Hung Le, Truyen Tran, and Svetha Venkatesh. Dual memory neural computer for asynchronous two-view sequential learning. In *SIGKDD*, pages 1637–1645. ACM, 2018.

[Li *et al.*, 2019] Yuwen Li, Weitong Chen, Deyin Liu, Zhimin Zhang, Shunxiang Wu, and Chengyu Liu. IFFLC: an integrated framework of feature learning and classification for multiple diagnosis codes assignment. *IEEE Access*, 7:36810–36818, 2019.

[Lim *et al.*, 2018] Jaechang Lim, Seongok Ryu, Jin Woo Kim, and Woo Youn Kim. Molecular generative model based on conditional variational autoencoder for de novo molecular design. *Journal of cheminformatics*, 10(1):1–9, 2018.

[Lipton *et al.*, 2016] Zachary C. Lipton, David C. Kale, Charles Elkan, and Randall C. Wetzel. Learning to diagnose with LSTM recurrent neural networks. In *ICLR*, 2016.

[Ma *et al.*, 2018] Tengfei Ma, Cao Xiao, Jiayu Zhou, and Fei Wang. Drug similarity integration through attentive multi-view graph auto-encoders. In *IJCAI*, pages 3477–3483, 2018.

[Makhzani *et al.*, 2015] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.

[Malakouti and Hauskrecht, 2019] Seyedsalim Malakouti and Milos Hauskrecht. Predicting patient's diagnoses and diagnostic categories from clinical-events in EHR data. In *AIME*, page 125–130, 2019.

[Miller *et al.*, 2016] Alexander H. Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. Key-value memory networks for directly reading documents. In *EMNLP*, pages 1400–1409, 2016.

[Pan *et al.*, 2018] Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, Lina Yao, and Chengqi Zhang. Adversarially regularized graph autoencoder for graph embedding. In *IJCAI*, pages 2609–2615, 2018.

[Shang *et al.*, 2019a] Junyuan Shang, Tengfei Ma, Cao Xiao, and Jimeng Sun. Pre-training of graph augmented transformers for medication recommendation. In *IJCAI*, pages 5953–5959, 2019.

[Shang *et al.*, 2019b] Junyuan Shang, Cao Xiao, Tengfei Ma, Hongyan Li, and Jimeng Sun. Gamenet: Graph augmented memory networks for recommending medication combination. In *AAAI*, volume 33, pages 1126–1133, 2019.

[Wang *et al.*, 2017] Meng Wang, Mengyue Liu, Jun Liu, Sen Wang, Guodong Long, and Buyue Qian. Safe medicine recommendation via medical knowledge graph embedding. *arXiv preprint arXiv:1710.05980*, 2017.

[Wang *et al.*, 2019] Yanda Wang, Weitong Chen, Bohan Li, and Robert Boots. Learning fine-grained patient similarity with dynamic bayesian network embedded rnns. In *DASFAA*, page 587–603. Springer, 2019.

[Yue *et al.*, 2020] Lin Yue, Dongyuan Tian, Weitong Chen, Xuming Han, and Minghao Yin. Deep learning for heterogeneous medical data analysis. *World Wide Web*, pages 1–23, 2020.

[Zhang *et al.*, 2017] Yutao Zhang, Robert Chen, Jie Tang, Walter F Stewart, and Jimeng Sun. Leap: learning to prescribe effective and safe treatment combinations for multimorbidity. In *SIGKDD*, pages 1315–1324. ACM, 2017.