

# Global-Local Item Embedding for Temporal Set Prediction

SEUNGJAE JUNG, Naver R&D Center NAVER CLOVA, Korea, Republic of

YOUNG-JIN PARK, Naver R&D Center NAVER CLOVA NAVER AI LAB, Korea, Republic of

JISU JEONG and KYUNG-MIN KIM, NAVER CLOVA NAVER AI LAB, Korea, Republic of

HIUN KIM, MINKYU KIM, and HANOCK KWAK, NAVER CLOVA, Korea, Republic of

Temporal set prediction is becoming increasingly important as many companies employ recommender systems in their online businesses, e.g., personalized purchase prediction of shopping baskets. While most previous techniques have focused on leveraging a user’s history, the study of combining it with others’ histories remains untapped potential. This paper proposes *Global-Local Item Embedding* (GLOIE) that learns to utilize the temporal properties of sets across whole users as well as within a user by coining the names as global and local information to distinguish the two temporal patterns. GLOIE uses Variational Autoencoder (VAE) and dynamic graph-based model to capture global and local information and then applies attention to integrate resulting item embeddings. Additionally, we propose to use Tweedie output for the decoder of VAE as it can easily model zero-inflated and long-tailed distribution, which is more suitable for several real-world data distributions than Gaussian or multinomial counterparts. When evaluated on three public benchmarks, our algorithm consistently outperforms previous state-of-the-art methods in most ranking metrics.

CCS Concepts: • **Information systems** → **Information systems applications**; **Recommender systems**.

Additional Key Words and Phrases: Temporal Sets, Set Prediction, Tweedie Distribution, Variational Autoencoder

## ACM Reference Format:

Seungjae Jung, Young-Jin Park, Jisu Jeong, Kyung-Min Kim, Hiun Kim, Minkyu Kim, and Hanock Kwak. 2021. Global-Local Item Embedding for Temporal Set Prediction. In *Fifteenth ACM Conference on Recommender Systems (RecSys '21)*, September 27–October 1, 2021, Amsterdam, Netherlands. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460231.3478844>

## 1 INTRODUCTION

Many recommendation tasks can be viewed as a problem of predicting the next set given the sequence of sets, e.g., predicting the next basket in online markets and the next playlist in streaming services. Previous works mainly focused on the local information (a given user’s history), applying RNNs [2, 4, 20] or self-attention [21] to learn the temporal tendency of a sequence of sets within a user. From the recommender system point of view, temporal set prediction can have a sparsity problem as many users interact with only a small number of items. Throughout the recommender system literature, the sparsity problem has been dealt with low-rank approximation or collaborative filtering [11, 15–17]. However, such attempts have been less explored in temporal set prediction literature.

In this paper, we propose Global-Local Item Embedding (GLOIE) that integrates global and local information for temporal set prediction. To capture the global information, we utilize the Variational Autoencoders (VAEs) [10, 12] which are effective on noisy and sparse data. GLOIE then integrates the local embeddings, that are learned through dynamic graphs [21], with global embeddings by using an attention method. In addition, we enhance the performance

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

Manuscript submitted to ACM

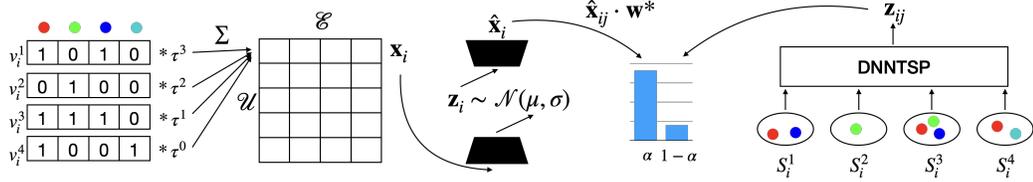


Fig. 1. Overview of GLOIE. We first make sum of time decayed vector  $x_i$  (left part of the figure). VAE maximizes the ELBO of each  $x_i$ . The weighted sum of  $\hat{x}_{ij} \cdot w^*$  and  $z_{ij}$  becomes final embedding for interacted items. The weight  $\alpha$  is determined by the attention mechanism. For the items that the user never interacted, we just use  $\hat{x}_{ij}$ .

of GLOIE by using Tweedie distribution for the likelihood of VAE instead of Gaussian or multinomial distributions. We show that purchase logs follow zero-inflated and long-tail distributions similar to the Tweedie distribution.

We conduct experiments on three public benchmarks, i.e. Dunnhumby Carbo, TaoBao, and Tags-Math-Sx. Empirical results demonstrate that the proposed method outperforms state-of-the-art methods on most metrics. The main contributions of the paper are summarized as follows:

- We propose GLOIE which differentiate global and local information and integrates them.
- We claim that using Tweedie output for VAE decoder is beneficial as it naturally models two properties of data distributions from temporal set prediction problems: zero-inflated and long-tailed.
- We achieve state-of-the-art performance on three public benchmarks.

## 2 PRELIMINARIES

### 2.1 Problem Definition

Let  $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$  be set of users and  $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$  be set of items. Given a user  $u_i$ 's sequence of sets  $\mathbb{S}_i = \{\mathbb{S}_i^k \mid 1 \leq k \leq T_i\}$ , our goal is to predict next set  $\mathbb{S}_i^{T_i+1}$ . Each set  $\mathbb{S}_i^k$  can also be represented as a binary vector form  $\mathbf{v}_i^k$ , where each element  $v_{i,j}^k = \mathbb{1}(e_j \in \mathbb{S}_i^k)$ .  $\mathbb{1}(x)$  is an indicator function which returns 1 if  $x$  is true and 0 otherwise. We will use the set notation  $\mathbb{S}_i^k$  and  $\mathbf{v}_i^k$  interchangeably to denote user  $u_i$ 's  $k$ -th set.

### 2.2 Variational Autoencoders

Variational Autoencoders (VAEs) [10, 13] are a class of deep generative models. VAEs provide latent structures that can nicely explain the observed data (e.g., customer's purchase history). Formally speaking, VAEs find the latent variables ( $\mathbf{z}$ ) that maximize the evidence lower bound (ELBO), a surrogate objective function for the maximum likelihood estimation of the given observation  $\mathbf{x}$ :

$$\log p_\theta(\mathbf{x}) \geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\log p_\theta(\mathbf{x}|\mathbf{z})] - D_{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \quad (1)$$

where  $\theta$  and  $\phi$  are model parameters of a decoder and an encoder, respectively.

Encoders are commonly structured by multilayer perceptrons (MLPs) that produce Gaussian distributions over the latent variable  $\mathbf{z}$ . On the other hand, decoders are designed to have different probability distributions of output layer depending on the characteristics of datasets. For example, Gaussian distribution and multinomial distribution are often used to represent the real-valued continuous and binary data, respectively.

### 3 METHOD

#### 3.1 Learning Global-Local Information by VAE

We are interested in modeling other users' history as well as the given user's history. We proceed with this under the VAE framework. However, a difficulty arises: every user has a different length of the sequence of sets. To resolve this, we sum time decayed sequence of sets.

$$\mathbf{x}_i = \sum_{1 \leq k \leq T_i} \mathbf{v}_i^k \cdot \tau^{T_i - k} \quad (2)$$

where  $\mathbf{v}_i^k$  is a vector representation of user  $u_i$ 's  $k$ th set  $S_i^k$ ,  $\tau \in (0, 1)$  is a decay factor and  $T_i$  is the length of the user  $u_i$ 's sequence as explained in Section 2.1. Note that  $k$  in  $\mathbf{v}_i^k$  represents the index and  $T_i - k$  in  $\tau^{T_i - k}$  represents the power. We then update our model to maximize the ELBO of each  $\mathbf{x}_i$ :

$$\mathcal{L}_{vae} = \mathbb{E}_{\mathbf{x}_i \sim \mathcal{D}_u} \left[ -\mathbb{E}_{\mathbf{z}_i \sim q(\mathbf{z} | \mathbf{x}_i)} [\log p(\mathbf{x}_i | \mathbf{z}_i)] + D_{KL}(q(\mathbf{z} | \mathbf{x}_i) || p(\mathbf{z})) \right] \quad (3)$$

where  $\mathcal{D}_u$  is a set of sum of time decayed sequence of sets  $\mathbf{x}_i$ . Equation (3) is merely an expectation of negative of Equation (1) over  $\mathcal{D}_u$ . Equation (3) can be understood as the low-rank approximation with the constraint  $D_{KL}(q(\mathbf{z} | \mathbf{x}_i) || p(\mathbf{z}))$ .

As most people interact with a few items,  $\mathbf{v}_i$  contains many 0s and so does  $\mathbf{x}_i$ . However, the reconstructed vector  $\hat{\mathbf{x}}_i$  generated from the following process:

$$\mathbf{z}_i \sim q(\mathbf{z} | \mathbf{x}_i), \quad \hat{\mathbf{x}}_i \sim p(\mathbf{x}_i | \mathbf{z}_i) \quad (4)$$

is a dense vector and contains the information of expected preference of user  $u_i$  to item  $e_j$  in  $\hat{\mathbf{x}}_{i,j}$  even though the user  $u_i$  never interacted with item  $e_j$ . Consider two users  $u_{i_1}$  and  $u_{i_2}$  with similar histories.  $\mathbf{z}_{i_1}$  and  $\mathbf{z}_{i_2}$  should be close as  $\mathbf{x}_{i_1}$  and  $\mathbf{x}_{i_2}$  are close. In turn, the distance between two reconstructed vectors  $\hat{\mathbf{x}}_{i_1}$  and  $\hat{\mathbf{x}}_{i_2}$  is small. Hence, even if user  $u_{i_1}$  never interacted with item  $e_j$ ,  $\hat{\mathbf{x}}_{i_1,j}$  should be high if  $\mathbf{x}_{i_2,j}$  is high.

Note that Hu et al. [5]'s Personalized Item Frequency (PIF) is similar to our sum of time decayed vectors. However, since their work is based on K-Nearest Neighbor, two shortcomings arise: 1) the inference time grows cubic with the number of users and 2) they cannot use features unlike traditional deep learning approaches.

#### 3.2 Tweedie Output on Decoder

When training VAE, using gaussian output on  $p(\mathbf{x} | \mathbf{z})$  is a straightforward option. However, the distributions of the data generated from temporal set prediction problems are zero-inflated and long-tailed as shown in Figure 2.

Tweedie distribution is a special case of exponential dispersion model (EDM) with a power parameter  $p$  and the variance function  $V(\mu) = \mu^p$  [6]. Tweedie distribution with  $1 < p < 2$  corresponds to a class of compound Poisson distributions [7]. Consider two step sampling process  $N \sim \text{Poisson}(\lambda)$  and  $X_i \sim \text{Gamma}(\alpha, \beta)$  for  $\lambda, \alpha, \beta > 0$  and  $i = 1, \dots, N$ . Now we define a random variable  $Z$  as follows:

$$Z = \begin{cases} 0 & N = 0 \\ X_1 + \dots + X_N & N > 0 \end{cases}. \quad (5)$$

It is straightforward that the distribution defined by Equation (5) is zero-inflated for small enough  $\lambda$  and long-tailed as  $Z$  is addition of  $X_i \sim \text{Gamma}(\alpha, \beta)$  for  $N > 0$ . Hence using Tweedie output on VAE's decoder for temporal set prediction is beneficial as Tweedie distribution easily captures the properties of distributions that are shown in Figure 2.

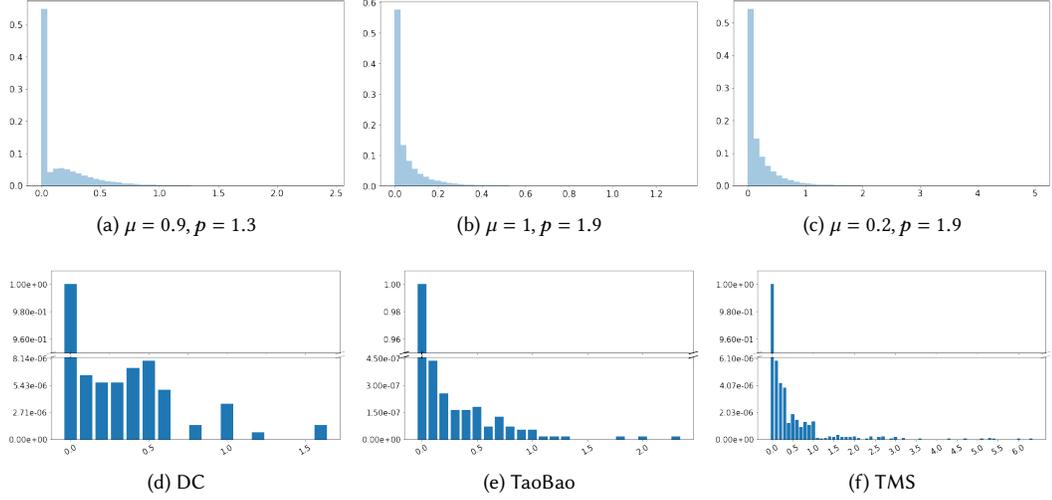


Fig. 2. Upper row: Histogram plot of samples from Tweedie distribution. Lower row: Histogram plot of elements of sum of decayed vector  $\mathbf{x}_j$ . Across all benchmarks, the distributions are zero-inflated and long-tailed.

Learning the mean parameter  $\mu$  and the power parameter  $p$  of Tweedie distribution via maximum likelihood is easy. Minimizing

$$\mathcal{L}_{Tweedie}(z, \mu, p) = -z \cdot \frac{\mu^{1-p}}{1-p} + \frac{\mu^{2-p}}{2-p} \quad (6)$$

maximizes log-likelihood. Here  $z$  is the target. See Yang et al. [19] for details.

A line of works chose distributions other than Gaussian or Bernoulli on matrix factorization and VAE [3, 12, 18]. Poisson distribution or multinomial distribution are usual choices. This paper is the first attempt to apply Tweedie distribution for VAE’s decoder output to the best of our knowledge.

### 3.3 Integrating Global-Local Information

Though VAE with Tweedie output is already competent, it tends to underestimate a user’s preference for the frequently interacted items. This tendency owes to the learning objective of VAE as the model has to maximize the likelihood of 0 for never interacted items. This sometimes sacrifices the ability to maximize the likelihood of values of interacted items. Hence, we integrate item embeddings of frequently interacted items which are learned by state-of-the-art models to our VAE. We use DNNTSP [21] as it is the state-of-the-art method. DNNTSP makes user-dependent embeddings of items that are interacted at least once via dynamic graph neural networks. We denote  $\mathbf{z}_{ij}$  as the embedding of user  $u_i$  for item  $e_j$ .

When it comes to combining VAE with embeddings  $\mathbf{z}_{ij}$ , a problem arises: the reconstructed value  $\hat{x}_{ij}$  of Equation (4) is a scalar while learned embedding  $\mathbf{z}_{ij}$  is a vector. Hence, to match the size between  $\hat{x}_{ij}$  and  $\mathbf{z}_{ij}$ , we simply multiply a vector  $\mathbf{w}^*$ , which is of same size as  $\mathbf{z}_{ij}$ , to  $\hat{x}_{ij}$ . We defer the discussion on the selection of  $\mathbf{w}^*$  to latter part of this section.

Now we combine updated embedding with  $\hat{x}_{ij} \cdot \mathbf{w}^*$ s and  $\mathbf{z}_{ij}$ s. The updated embedding  $\tilde{\mathbf{z}}_{ij}$  is defined as

Table 1. Statistics of public benchmarks

Dataset	# of users	# of sets	# of elements	#E/S	#S/U	#E/U
DC	9,010	42,905	217	1.52	4.76	5.44
TaoBao	113,347	628,618	689	1.10	5.55	4.96
TMS	15,726	243,394	1,565	2.19	15.48	18.05

$$\tilde{z}_{ij} = \begin{cases} Att(\tilde{\mathbf{x}}_{ij} \cdot \mathbf{w}^*, \mathbf{z}_{ij}) & e_j \in \cup_k S_i^k \\ \hat{\mathbf{x}}_{ij} \cdot \mathbf{w}^* & otherwise \end{cases} \quad (7)$$

where  $\tilde{\mathbf{x}}_{ij} = \frac{\hat{\mathbf{x}}_{ij}}{\max_j \hat{\mathbf{x}}_{ij}} - 0.5$  which normalizes all values of  $\hat{\mathbf{x}}_i$  to  $[0.5, -0.5]$  and  $Att(\mathbf{q}, \mathbf{k})$  is defined as follows:

$$Att(\mathbf{q}, \mathbf{k}) = \alpha \cdot \mathbf{q} + (1 - \alpha) \cdot \mathbf{k}, \text{ where } \alpha = \sigma\left(\frac{(\mathbf{W}_q \cdot \mathbf{q})^T (\mathbf{W}_k \cdot \mathbf{k})}{\|\mathbf{W}_q \cdot \mathbf{q}\| \|\mathbf{W}_k \cdot \mathbf{k}\|}\right). \quad (8)$$

Lastly, we calculate the affinity of a user  $u_i$  to item  $e_j$  by

$$\hat{y}_{ij} = \sigma\left(\tilde{z}_{ij}^T \mathbf{w}_0 + b_0\right) \quad (9)$$

where  $\sigma(x) = \frac{1}{1+e^{-x}}$ . We choose multi-label soft margin loss for training:

$$\mathcal{L}_{TSP} = -\frac{1}{N} \sum_i \sum_j^M y_{ij} \log \hat{y}_{ij} + (1 - y_{ij}) \log(1 - \hat{y}_{ij}). \quad (10)$$

We now going back to the selection of  $\mathbf{w}^*$  in Equation (7). We set  $\mathbf{w}^* = \mathbf{w}_0$  in Equation (9) though we can set  $\mathbf{w}^*$  as a learnable parameter. If we set  $\mathbf{w}^* = \mathbf{w}_0$ , Equation (9) becomes  $\hat{y}_{ij} = \sigma\left(\hat{\mathbf{x}}_{ij} \cdot \mathbf{w}_0^T \mathbf{w}_0 + b_0\right)$  for non-interacted item  $e_j \notin \cup_k S_i^k$ , hence it preserves the order of affinity that is learned by VAE. We also run experiments with learnable parameter  $\mathbf{w}^*$  but the performance was on par with  $\mathbf{w}^* = \mathbf{w}_0$ .

## 4 EXPERIMENTS

### 4.1 Benchmarks

We evaluate our method on three public benchmarks: *Dunnhumby Carbo (DC)*, *TaoBao*, and *Tags-Math-Sx (TMS)*. We partitioned each dataset into train, validation and test into 70%, 10% and 20% respectively following Yu et al. [21]. See Table 1 for statistics of benchmarks.

### 4.2 Compared Methods

We compare four methods: *Toppop*, *PersonalToppop*, *Sets2Sets* and *DNNTSP*.

*Toppop* simply serves the items that are interacted the most across all users. *PersonalToppop* serves the items that the given user interacted with the most. *Sets2Sets* uses encoder-decoder framework to predict the next set [4]. Set embeddings are made by pooling operation and set-based attention is used to model temporal correlation relation. This method also models repeated elements. *DNNTSP* is composed of three components: Element Relationship Learning (ERL), Temporal Dependency Learning (TDL) and Gated Information Fusing. ERL is simply a dynamic weighted graph neural networks. TDL captures temporal dependency. By Gated Information Fusing each user shares the embeddings of

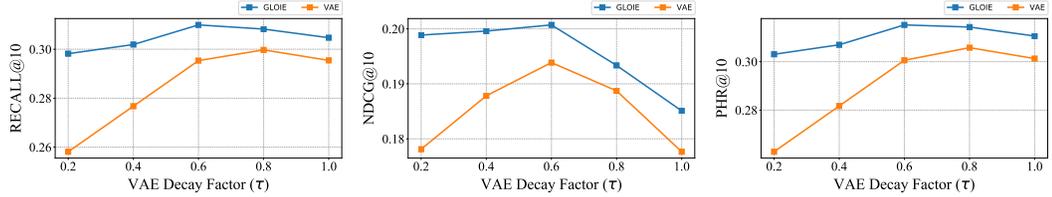


Fig. 3. Metric@10 by varying VAE decay factor ( $\tau$ ) on TaoBao dataset.

uninteracted items. Hence DNNTSP can be seen as an ensemble of model which learns local information and *Toppop* model.

### 4.3 Results & Analyses

The results of our evaluation on three public benchmarks are shown in Table 2. We consider three metrics: *Recall*, *Normalized Discounted Cumulative Gain (NDCG)* and *Personal Hit Ratio (PHR)*.  $PHR@K$  is calculated as  $\sum_{i=1}^{N'} \mathbb{1}(|\hat{S}_i \cap S_i| > 0) / N'$  where  $N'$  is the number of test users,  $\hat{S}_i$  is the predicted top-K elements, and  $S_i$  is the ground truth set.

We trained VAE for 30 epochs and then trained DNNTSP for 30 epochs. We used only one layer for the VAEs across all benchmarks. For the decay factor in Equation (2), we set  $\tau = 0.6$ . As illustrated in Figure 3,  $\tau = 0.6$  shows the best performance on TaoBao dataset. We empirically observe that  $\tau = 0.6$  could provide decent results across all metrics on the other datasets as well. The dimension of latent space is 128 for DC and TaoBao, and 512 for TMS. We used Adam optimizer [9] with learning rate 0.001.

Across all benchmarks and metrics, GLOIE with Tweedie output outperforms or is on par with all compared methods. Especially, GLOIE with Tweedie output outperforms the other methods on all metrics on DC and TaoBao datasets. We can see that GLOIE with Tweedie outperforms DNNTSP on every metric when  $K = 10$  which means that the embeddings learned by VAE are richly used as well as the embeddings learned by DNNTSP.

One thing to remark is that VAE with Tweedie output shows comparable performance to all compared methods on most metrics. Given that the number of items a user interacted with is 5.44, 4.96, and 18.05 respectively in DC, TaoBao, and TMS, this shows that VAE with Tweedie output captures the preference of users to non-interacted items.

To investigate the effectiveness of the attention-based integration method illustrated in Equation (8), we compared GLOIE with attention to the one with itemwise learnable weight similar to the method proposed in Yu et al. [21]. For overall datasets and metrics, we could observe performance gains: 3.09%, 0.45%, 0.92% improvement of NDCG@10 on DC, Taobao, and TMS, respectively.

Lastly, we note that the selection of output distribution of decoder on VAE largely affects the performance. We empirically show that both Gaussian and multinomial outputs do not fit temporal set prediction problems even though Gaussian is a popular choice for VAE and multinomial is a common choice in recommender system literature after the advent of VAECF [12].

## 5 CONCLUSION

This paper proposes Global-Local Item Embedding (GLOIE) that learns to utilize the temporal properties of sets across whole users as well as within a user. The proposed model learns global-local information by maximizing ELBO of the sum of time decayed vectors under the VAE framework and integrates local embeddings learned by dynamic graph

Table 2. Comparison between various state-of-the-art methods and ours on three public benchmarks. All highest scores are in **bold** and all second best scores are underlined.

Dataset	Model	k = 10			k = 20			k = 40		
		Recall	NDCG	PHR	Recall	NDCG	PHR	Recall	NDCG	PHR
DC	Toppop	0.1618	0.0880	0.2274	0.2475	0.1116	0.3289	0.3940	0.1448	0.4997
	PersonalToppop	0.4104	0.3174	0.5031	0.4293	0.3270	0.5258	0.4747	0.3332	0.5785
	Sets2Sets	0.4488	0.3136	0.5458	0.5143	0.3319	0.6162	0.6017	0.3516	0.7005
	DNNTSP	<u>0.4564</u>	<u>0.3165</u>	<u>0.5557</u>	<u>0.5294</u>	<u>0.3369</u>	<u>0.6272</u>	0.6180	<u>0.3568</u>	<u>0.7165</u>
	VAE - <i>Gaussian</i>	0.1618	0.0882	0.2274	0.2507	0.1128	0.3333	0.3847	0.1430	0.4903
	VAE - <i>Multinomial</i>	0.1602	0.0850	0.2230	0.2492	0.1097	0.3311	0.3767	0.1387	0.4786
	VAE - <i>Tweedie</i>	0.4166	0.3000	0.5108	0.5122	0.3267	0.6062	<u>0.6217</u>	0.3517	0.7088
	GLOIE - <i>Gaussian</i>	0.3108	0.2349	0.3971	0.3738	0.2526	0.4664	0.4545	0.2706	0.5563
	GLOIE - <i>Multinomial</i>	0.3265	0.2465	0.4143	0.3870	0.2633	0.4798	0.4615	0.2803	0.5602
	GLOIE - <i>Tweedie</i>	<b>0.4658</b>	<b>0.3264</b>	<b>0.5613</b>	<b>0.5415</b>	<b>0.3477</b>	<b>0.6351</b>	<b>0.6428</b>	<b>0.3708</b>	<b>0.7288</b>
TaoBao	Toppop	0.1567	0.0784	0.1613	0.2494	0.1019	0.2545	0.3679	0.1264	0.3745
	PersonalToppop	0.2190	0.1535	0.2230	0.2260	0.1554	0.2306	0.2433	0.1590	0.2484
	Sets2Sets	0.2811	0.1495	0.2868	0.3649	0.1710	0.3713	0.4672	0.1922	0.4739
	DNNTSP	<u>0.3035</u>	0.1841	<u>0.3095</u>	<u>0.3811</u>	0.2039	<u>0.3873</u>	<u>0.4776</u>	0.2238	<u>0.4843</u>
	VAE - <i>Gaussian</i>	0.1592	0.0750	0.1635	0.2480	0.0974	0.2530	0.3665	0.1219	0.3727
	VAE - <i>Multinomial</i>	0.1588	0.0798	0.1634	0.2494	0.1027	0.2545	0.3660	0.1268	0.3723
	VAE - <i>Tweedie</i>	0.2954	<u>0.1939</u>	0.3006	0.3775	<u>0.2148</u>	0.3827	0.4768	<u>0.2353</u>	0.4822
	GLOIE - <i>Gaussian</i>	0.2982	0.1768	0.3044	0.3790	0.1973	0.3851	0.4769	0.2175	0.4835
	GLOIE - <i>Multinomial</i>	0.2980	0.1791	0.3040	0.3783	0.1995	0.3846	0.4750	0.2195	0.4819
	GLOIE - <i>Tweedie</i>	<b>0.3099</b>	<b>0.2007</b>	<b>0.3152</b>	<b>0.3917</b>	<b>0.2216</b>	<b>0.3972</b>	<b>0.4868</b>	<b>0.2412</b>	<b>0.4924</b>
TMS	Toppop	0.2627	0.1627	0.4619	0.3902	0.2017	0.6243	0.5605	0.2448	0.8007
	PersonalToppop	0.4508	0.3464	0.6440	0.5274	0.3721	0.7146	0.5495	0.3771	0.7374
	Sets2Sets	<u>0.4748</u>	<u>0.3782</u>	<b>0.6933</b>	0.5601	<u>0.4061</u>	0.7594	0.6627	<u>0.4321</u>	<u>0.8570</u>
	DNNTSP	0.4693	0.3473	0.6825	<u>0.5826</u>	0.3839	<b>0.7880</b>	<u>0.6840</u>	0.4097	<b>0.8748</b>
	VAE - <i>Gaussian</i>	0.2731	0.1919	0.4660	0.3913	0.2288	0.6195	0.5496	0.2688	0.7813
	VAE - <i>Multinomial</i>	0.2548	0.1615	0.4431	0.3830	0.2001	0.6020	0.5437	0.2412	0.7740
	VAE - <i>Tweedie</i>	0.4661	0.3744	0.6548	0.5579	0.4040	0.7432	0.6663	0.4316	0.8341
	GLOIE - <i>Gaussian</i>	0.1345	0.0833	0.2486	0.2363	0.1155	0.4018	0.4014	0.1570	0.6033
	GLOIE - <i>Multinomial</i>	0.1479	0.1029	0.2797	0.2192	0.1252	0.3872	0.3259	0.1524	0.5362
	GLOIE - <i>Tweedie</i>	<b>0.4860</b>	<b>0.3823</b>	<u>0.6863</u>	<b>0.5868</b>	<b>0.4144</b>	<u>0.7753</u>	<b>0.6926</b>	<b>0.4418</b>	0.8538

neural networks. As users with similar histories are reconstructed to close vectors, we could model the preference of the given user for a non-interacted item if other users with similar histories frequently interacted with the item. Data analysis and empirical results show that using Tweedie output for VAE’s decoder is effective for modeling temporal set prediction. The proposed method achieves state-of-the-art results by considering global information which is less explored in temporal set prediction literature.

Though the proposed VAE is powerful in itself, there is some room for improvement. Instead of using the sum of time decayed vector, we can model change of sets in continuous time by Neural ODE [1, 14]. Finding an appropriate form of prior that captures long-tailed distribution for the temporal set prediction can also be a future direction. Graph modality can also be used [8]. Last but not least, searching for better ways to integrating embeddings learned by VAE and embeddings learned by other algorithms which focus on local information is also an important future work.

## REFERENCES

- [1] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural Ordinary Differential Equations. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf>
- [2] Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F Stewart, and Jimeng Sun. 2016. Doctor ai: Predicting clinical events via recurrent neural networks. In *Machine learning for healthcare conference*. PMLR, 301–318.
- [3] Prem Gopalan, Jake M. Hofman, and David M. Blei. 2015. Scalable Recommendation with Hierarchical Poisson Factorization. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (Amsterdam, Netherlands) (UAI'15)*. AUAI Press, Arlington, Virginia, USA, 326–335.
- [4] Haoji Hu and Xiangnan He. 2019. Sets2sets: Learning from sequential sets with neural networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1491–1499.
- [5] Haoji Hu, Xiangnan He, Jinyang Gao, and Zhi-Li Zhang. 2020. Modeling personalized item frequency information for next-basket recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1071–1080.
- [6] Bent Jørgensen. 1987. Exponential Dispersion Models. *Journal of the Royal Statistical Society. Series B (Methodological)* 49, 2 (1987), 127–162. <http://www.jstor.org/stable/2345415>
- [7] Bent Jørgensen. 1997. *The Theory of Dispersion Models*. Chapman and Hall/CRC.
- [8] Kyung-Min Kim, Dong-Hyun Kwak, Hanock Kwak, Young-Jin Park, Sangkwon Sim, Jae-Han Cho, Minkyu Kim, Jihun Kwon, Nako Sung, and Jung-Woo Ha. 2019. Tripartite Heterogeneous Graph Propagation for Large-scale Social Recommendation. In *Proceedings of ACM RecSys 2019 Late-Breaking Results*, Marko Tkalcić and Sole Pera (Eds.). <http://ceur-ws.org/Vol-2431/paper12.pdf>
- [9] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [10] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, Yoshua Bengio and Yann LeCun (Eds.). <https://openreview.net/forum?id=33X9fd2-9FyZd>
- [11] Yehuda Koren. 2008. Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Las Vegas, Nevada, USA) (KDD '08)*. Association for Computing Machinery, New York, NY, USA, 426–434. <https://doi.org/10.1145/1401890.1401944>
- [12] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*. 689–698.
- [13] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*. PMLR, 1278–1286.
- [14] Yulia Rubanova, Ricky T. Q. Chen, and David K Duvenaud. 2019. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/file/42a6845a557bef704ad8ac9cb4461d43-Paper.pdf>
- [15] Ruslan Salakhutdinov and Andriy Mnih. 2007. Probabilistic Matrix Factorization. In *Proceedings of the 20th International Conference on Neural Information Processing Systems (Vancouver, British Columbia, Canada) (NIPS'07)*. Curran Associates Inc., Red Hook, NY, USA, 1257–1264.
- [16] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-Based Collaborative Filtering Recommendation Algorithms. In *Proceedings of the 10th International Conference on World Wide Web (Hong Kong, Hong Kong) (WWW '01)*. Association for Computing Machinery, New York, NY, USA, 285–295. <https://doi.org/10.1145/371920.372071>
- [17] Nathan Srebro, Jason Rennie, and Tommi Jaakkola. 2005. Maximum-Margin Matrix Factorization. In *Advances in Neural Information Processing Systems*, L. Saul, Y. Weiss, and L. Bottou (Eds.), Vol. 17. MIT Press. <https://proceedings.neurips.cc/paper/2004/file/e0688d13958a19e087e12314855e4b4-Paper.pdf>
- [18] Quoc-Tuan Truong, Aghiles Salah, and Hady W. Lauw. 2021. Bilateral Variational Autoencoder for Collaborative Filtering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (Virtual Event, Israel) (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 292–300. <https://doi.org/10.1145/3437963.3441759>
- [19] Yi Yang, Wei Qian, and Hui Zou. 2016. Insurance Premium Prediction via Gradient Tree-Boosted Tweedie Compound Poisson Models. [arXiv:1508.06378 \[stat.ME\]](https://arxiv.org/abs/1508.06378)
- [20] Feng Yu, Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. A dynamic recurrent model for next basket recommendation. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 729–732.
- [21] Le Yu, Leilei Sun, Bowen Du, Chuanren Liu, Hui Xiong, and Weifeng Lv. 2020. Predicting Temporal Sets with Deep Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1083–1091.