

ContrastVAE: Contrastive Variational AutoEncoder for Sequential Recommendation

Yu Wang
University of Illinois Chicago
Chicago, United States
ywang617@uic.edu

Hengrui Zhang
University of Illinois Chicago
Chicago, United States
hzhan55@uic.edu

Zhiwei Liu
Salesforce
San Francisco, United States
zhiweiliu@salesforce.edu

Liangwei Yang
University of Illinois Chicago
Chicago, United States
lyang84@uic.edu

Philip S. Yu
University of Illinois Chicago
Chicago, United States
psyu@uic.edu

ABSTRACT

Aiming at exploiting the rich information in user behaviour sequences, sequential recommendation has been widely adopted in real-world recommender systems. However, current methods suffer from the following issues: 1) **sparsity** of user-item interactions, 2) **uncertainty** of sequential records, 3) **long-tail items**. In this paper, we propose to incorporate contrastive learning into the framework of Variational AutoEncoders to address these challenges simultaneously. Firstly, we introduce ContrastELBO, a novel training objective that extends the conventional single-view ELBO to two-view case and theoretically builds a connection between VAE and contrastive learning from a two-view perspective. Then we propose Contrastive Variational AutoEncoder (ContrastVAE in short), a two-branched VAE model with contrastive regularization as an embodiment of ContrastELBO for sequential recommendation. We further introduce two simple yet effective augmentation strategies named *model augmentation* and *variational augmentation* to create a second view of a sequence and thus making contrastive learning possible. Experiments on four benchmark datasets demonstrate the effectiveness of ContrastVAE and the proposed augmentation methods. Codes are available at <https://github.com/YuWang-1024/ContrastVAE>

CCS CONCEPTS

• **Information systems** → **Recommender systems**.

KEYWORDS

Sequential Recommendation, Variational AutoEncoder, Contrastive Learning

ACM Reference Format:

Yu Wang, Hengrui Zhang, Zhiwei Liu, Liangwei Yang, and Philip S. Yu. 2022. ContrastVAE: Contrastive Variational AutoEncoder for Sequential Recommendation. In *Proceedings of the 31st ACM International Conference*

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '22, October 17–21, 2022, Atlanta, GA, USA.

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9236-5/22/10...\$15.00
<https://doi.org/10.1145/3511808.3557268>

on Information and Knowledge Management (CIKM '22), October 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3511808.3557268>

1 INTRODUCTION

Sequential Recommendation (SR) has attracted increasing attention due to its ability to model the temporal dependencies in users' clicking histories, which can help better understand the user behaviors and intentions. Recent research justifies the promising ability of self-attention models [16, 19, 35] in characterizing the temporal dependencies on real-world sequential recommendation tasks. These methods encode a sequence as an embedding via an attention-based weighted sum of items' hidden representations. To name a few, SASRec [16] is a pioneering work adopting the self-attention mechanism to learn transition patterns in item sequences, and Bert4Rec [35] extends it as a bi-directional encoder to predict the next item.

Despite their great representation power, both the *uncertainty* problem and the *sparsity* issue impair their performance.

Firstly, the uncertainty problem is due to the rigorous assumption of sequential dependencies, which may be destroyed by unobserved factors in real-world scenarios. For example, for music recommendation, the genre of music that a user listens may vary according to different circumstances. Nevertheless, those factors are unknown and cannot be fully revealed in sequential patterns.

Secondly, the sparsity issue is a long-existing and not yet a well-solved problem in recommender systems [24, 39, 43]. Supposing that a user only interacts with few items, current methods are unable to learn high-quality representations of the sequences, thus failing to characterize sequential dependencies. Moreover, the sparsity issue increases the deficiency of uncertainty in sequential recommendation. More concretely, if a user has fewer historical interactions, those uncertain factors are of higher dominance over sequential patterns. However, these two issues are seldom studied simultaneously.

Therefore, we investigate the potential of adopting Variational AutoEncoder (VAE) into sequential recommendation. The reasons are threefold. First of all, VAE can estimate the uncertainty of the input data. To be more specific, it characterizes the distributions of those hidden representations via an encoder-decoder learning paradigm, which assumes that those representations follow a Gaussian distribution. Hence, the variances in Gaussian distribution can well

characterize the uncertainty of the input data. Moreover, the decoder maximizes the expected likelihood of input data conditioned on such latent variables, which can thus reduce the deficiency from unexpected uncertainty. Secondly, the posterior distribution estimation in VAE decreases the vulnerability to the sparsity issue. Though a sequence contains few items, we can still characterize its distribution from learned prior knowledge and thus generate the next item. Last but not least, probabilistic modeling of those hidden representations also enhance the robustness of sparse data against uncertainty. Specifically, if we can ensure the estimated posterior of perturbed input still being in-distribution, the decoder in VAE will tolerate such perturbations and yield correct next-item prediction.

However, conventional VAE suffers from posterior collapse issues [21, 46]. Concretely, if the decoder is sufficiently expressive, the estimated posterior distributions of latent factors tend to resemble the standard Gaussian distributions, i.e., these estimations are indistinguishable from each other as they follow the same distribution [2, 21]. Furthermore, VAE might collapse to point estimation for rare classes that simply memorize the locations in latent space. The high skewed distribution of user behaviors will exaggerate these problems. Specifically, the sequential input data consists of *long-tail items* [20], which refer to the infrequent items that rarely appear in the users' historical records. Such items account for a large portion of all items. These limitations prevent VAE from achieving satisfactory performance for SR tasks.

Recent advances in adopting contrastive learning (CL) for alleviating representation degeneration problem [29] motivate us to resort to contrastive learning to mitigate the above issues. Concretely, contrastive learning encourages the uniformity distribution of latent representations of different inputs [38], thus enforcing them distinguishable in latent space. Besides, augmentations in CL encourage perturbed sequences to share similar representations, thus being robust to a large variance of the estimated posterior distribution. To incorporate contrastive learning to the framework of VAE, we first extend the conventional single-variable ELBO to the two-view case and propose ContrastELBO. We theoretically prove that optimizing ContrastELBO induces a mutual information maximization term, which could be effectively optimized with CL [28, 37].

To instantiate ContrastELBO for SR, we propose ContrastVAE, a two-branched VAE model that naturally incorporates CL. ContrastVAE takes two augmented views of a sequence as input and follows the conventional encoder-sampling-decoder architecture to generate the next predicted item. The model is learned through optimizing an additional contrastive loss between the latent representations of two views in addition to the vanilla reconstruction losses and KL-divergence terms. To deal with the potential inconsistency problem led by uninformative data augmentations, we further propose two novel augmentation strategies: *model augmentation* and *variational augmentation*, which introduce perturbations in the latent space instead of the input space. We conduct comprehensive experiments on four benchmark datasets, which verify the effectiveness of the proposed model for sequential recommendation tasks, especially on recommending long-tail items. The contributions of this paper are summarized as follows:

- We derive ContrastELBO, which is an extension of conventional single-view ELBO to two-view case and naturally incorporates contrastive learning to the framework of VAE.
- We propose ContrastVAE, a two-branched VAE framework guided by ContrastELBO for sequential recommendation.
- We introduce model augmentation and variational augmentation to avoid the semantic inconsistency problem led by conventional data augmentation.
- We conduct comprehensive experiments to evaluate our method. The results show that our model achieves state-of-the-art performance on four SR benchmarks. Extensive ablation studies and empirical analysis verify the effectiveness of the proposed components.

2 RELATED WORKS

2.1 VAE for Sequential Recommendation

Variational AutoEncoder (VAE) approximates the posterior distribution of latent variables given input data through variational inference, and has been introduced to recommender systems recently to model the uncertainty of user preferences. In the context of SR, SVAE [32] and VSAN [45] adopt VAE to learn the dynamic hidden representations of sequences by utilizing Recurrent Neural Network / Self-Attention as the Encoder and Decoder, respectively. Despite their abilities to model the uncertainty of sequential behaviors, empirical studies show that they perform worse than deterministic models on a lot of tasks [16, 35], which is usually attributed to the low quality of approximate posterior distribution [42]. To alleviate such an issue, ACVAE [42] introduced adversarial variational Bayes combined with mutual information maximization between user embeddings and input sequences in order to obtain more salient and personalized representations of different users. However, DIM [14] has revealed that it is insufficient merely optimizing mutual information between latent representations and inputs, so we resort to latent-level contrastive learning through augmentations instead.

2.2 Contrastive Sequential Recommendation

Through maximizing the agreement of the representations of different augmentations of the input data, CL [6, 14, 37] has become a popular method for recommender systems in order to improve the accuracy and robustness of recommendation models [29, 40]. As a pioneering work, S3Rec [47] utilizes the correlations among attributes, items, and sub-sequences through mutual information maximization. CL4Rec [41] introduces sequence masking, cropping, and reordering for sequential data augmentation and applies InfoNCE loss for item-level CL. To generate informative augmentations for user behavior sequences, CCL [4] employs a learnable context-aware generator for data augmentation, and then CL is performed on different augmented samples. To improve the robustness of data augmentations on item sequences, CoSeRec [23] takes the correlations between different items into consideration. DuoRec [29] points out that the representations of items tend to become non-informative for many SR algorithms and highlights the limitations of random data-augmentation based CL for SR. They propose to choose semantically similar sequences (with the same ground truth label) as their positive samples to address the inconsistency problem caused by random data augmentation. Different

from data-augmentation based models, our method works competitively with the simple dropout operation. We further introduce variational dropout [5], an adaptive dropout augmentation method to generate multiple views for a sequence, which naturally fits our VAE framework.

3 PRELIMINARY

3.1 Variational AutoEncoders and ELBO

Given observed data x , latent variable models assume an underlying generative model $p(x, z) = p_\theta(x|z)p(z)$, where z is the latent variable, $p(z)$ is the prior, and $p_\theta(x|z)$ is the likelihood parameterized by θ . Then a natural objective is to maximize the log likelihood of the data distribution:

$$\log p(x) = \log \int_z p(z)p_\theta(x|z)dz \quad (1)$$

However, it is intractable to directly optimize Eq. 1 as it requires integration over all possible z . To mitigate this issue, VAEs [18] adopt variational inference and use an approximate posterior distribution $q_\phi(z|x)$. Then $p_\theta(x|z)$ and $q_\phi(z|x)$ are jointly optimized through maximizing the Evidence Lower Bound (ELBO) to the log likelihood:

$$\begin{aligned} \log p(x) &\geq \mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)] - D_{KL}[q_\phi(z|x)||p(z)]. \\ &= \mathcal{L}_{ELBO}. \end{aligned} \quad (2)$$

In VAE, both $q_\phi(z|x)$ (encoder) and $p_\theta(x|z)$ (decoder) are parameterized by neural networks. In Eq. 2, the first term on the right is the expected conditional joint log-likelihood w.r.t the approximate posterior and is approximately computed through sampling with a reparameterization trick. The second term is the Kullback-Leibler divergence between the approximate posterior $q_\phi(z|x)$ and the prior $p(z)$, whose closed-form solution is easy to compute when assuming both are Gaussian distributions.

3.2 Posterior collapse in VAE

Despite the success of VAEs, a dominant issue named posterior collapse has been observed, which greatly reduces the capacity of the generative model [25, 36, 46]. The posterior collapse is usually formulated as $D_{KL}[q_\phi(z|x)||p(z)] \rightarrow 0$ for every x , and usually occurs when the decoder model is too powerful. This indicates that the learned variational distribution is almost identical to the prior (i.e., standard Gaussian distribution), thus making the latent variables of different inputs indistinguishable in the latent space. Zhao et al. [46] and Takida et al. [36] further point out that posterior collapse makes the mutual information between the input and its latent variable $\mathcal{I}(x, z)$ vanishingly small.

To address this issue, recent methods try to reduce the impact of the KL-divergence term by reducing its weight [21, 27, 33, 45], or introducing an additional regularization term that explicitly maximizes the mutual information between the input and latent [2, 29, 46]. However, this issue is much more serious in SR tasks as the user-item interactions are extremely sparse, and the user's dynamic preferences would be hard to model. Furthermore, we find that these methods are insufficient for better performance on the SR. As a remedy, we address the problem from the two-view CL perspective, where we maximize the mutual information between two views of each sequence in latent space $\mathcal{I}(z, z')$. In

Section 4.2 we show that through extending the vanilla single latent variable generative model to two-view case, the VAE framework could naturally incorporate the mutual information maximization principle, which could be optimized through a CL loss [14, 28, 37].

4 METHODOLOGY

4.1 Problem Definition

In SR tasks, there is a set of users $\mathcal{U} = \{u_1, u_2, \dots, u_M\}$ and a set of items $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$. For each user u , we are given its historical interaction records with the items, which are sorted according to their timestamps. Then in this case user u 's behaviour could be represented as an item sequence $\mathbf{X}^u = x_{1:T}^u = \{x_1^u, x_2^u, \dots, x_T^u\}$, where x_t^u is the index of the item that user u has interacted at time step t , and T is the number of total time steps (i.e., the number of items user u has interacted with). The target of SR is to predict the next item x_{T+1}^u given the previous clicked items $x_{1:T}^u$. For each x_t^u , we further introduce a latent variable z_t^u representing user u 's preference at time t . In Section 4.2, we first derive the general formulation of ContrastELBO in Eq. 2, which extends the traditional single latent-variable ELBO to two-view case. We then discuss how to apply the ContrastELBO in SR scenarios in Section 4.3.

4.2 ContrastELBO

We start by introducing the following theorem about the latent variable model with two observed variables x and x' , which represent two views of input:

THEOREM 4.1. (*ContrastELBO*). *Consider a generative model with two observed variables (x, x') and two latent variables (z, z') with structure $x \leftarrow z - z' \rightarrow x'$, where z is used to generate x , z' is used to generate x' . z and z' are dependent as they are two views of the same input. This generative model indicates that joint distribution can be factorized by $p(z, z', x, x') = p(z, z')p(x|z)p(x'|z')$. Then we have the following lower bound of the log joint probability of the observed variables:*

$$\begin{aligned} \log p(x, x') &\geq \mathbb{E}_{q(z|x)} \log p(x|z) - D_{KL}[q(z|x)||p(z)] \\ &\quad + \mathbb{E}_{q(z'|x')} \log p(x'|z') - D_{KL}[q(z'|x')||p(z')] \\ &\quad + \mathbb{E}_{q(z, z'|x, x')} \log \left[\frac{p(z, z')}{p(z)p(z')} \right] \end{aligned} \quad (3)$$

PROOF. According to the above generative model, we have x and x' that are conditionally independent given z and z' , or formally $p(x, x'|z, z')$, then we can approximate the posterior with a variational distribution $q(z, z'|x, x')$ which could be factorized through:

$$q(z, z'|x, x') = q(z|x)q(z'|x'). \quad (4)$$

Then we have

$$\begin{aligned}
\log p(x, x') &= \log \int p(x, x', z, z') dz dz' \\
&= \log \mathbb{E}_{q(z, z' | x, x')} \left[\frac{p(x, x', z, z')}{q(z, z' | x, x')} \right] \\
&\geq \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(x, x', z, z')}{q(z, z' | x, x')} \right] \\
&= \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(x|z)p(x'|z')p(z, z')}{q(z|x)q(z'|x')} \right] \\
&= \mathbb{E}_{q(z|x)} \log [p(x|z)] + \mathbb{E}_{q(z'|x')} \log [p(x'|z')] \\
&\quad + \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z, z')}{q(z|x)q(z'|x')} \right].
\end{aligned} \tag{5}$$

The last term in the last equation could be further expanded:

$$\begin{aligned}
&\mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z, z')}{q(z|x)q(z'|x')} \right] \\
&= \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z, z')p(z)p(z')}{q(z|x)q(z'|x')p(z)p(z')} \right] \\
&= \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z, z')}{p(z)p(z')} \right] \\
&\quad + \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z)p(z')}{q(z|x)q(z'|x')} \right] \\
&= \mathbb{E}_{q(z, z' | x, x')} \log \left[\frac{p(z, z')}{p(z)p(z')} \right] \\
&\quad - D_{KL}[q(z|x)||p(z)] - D_{KL}[q(z'|x')||p(z')]
\end{aligned} \tag{6}$$

Plugging Eq. 6 into Eq. 5, then we complete the proof. \square

Note that the first four terms on the right of Eq. 3 is identical to that of the vanilla ELBO in Eq. 2 and could be effectively optimized using traditional VAE models. The last term $\mathbb{E}_{q(z, z' | x, x')} \log \frac{p(z, z')}{p(z)p(z')}$, however, is hard to compute. To make this term tractable, we follow the practice in Aitchison [1] that specifies $p(z, z') = q(z, z')$, $p(z) = q(z)$ and $p(z') = q(z')$ through choosing specific prior distributions, and then this term becomes $\mathbb{E}_{q(z, z' | x, x')} \frac{q(z, z')}{q(z)q(z')}$. If taking its expectation under the true data distribution $p_{true}(x, x')$, the last term becomes:

$$\mathbb{E}_{q(z, z')} \frac{q(z, z')}{q(z)q(z')} = D_{KL}[q(z, z')||q(z)q(z')] = I(z, z'). \tag{7}$$

Eq. 7 indicates that we can maximize the mutual information between z and z' from $q(z, z')$. Note that (z, z') are the encoder's output taking a pair of data (x, x') as input, so the mutual information term can be efficiently estimated using its tractable lower bounds through CL [3, 28, 37].

In section 4.3 we present ContrastVAE, a direct instantiation of the contrastELBO, for sequential recommendation.

4.3 ContrastVAE

The framework of ContrastVAE is presented in Fig. 1. ContrastVAE consists of three components: transformer encoder, augmentation strategies for generating the second view, and transformer decoder.

4.3.1 Transformer encoder. Inspired by the recent advances of transformer-based methods for SR [16, 35], we adopt self-attention as the building block of our encoder. Note that our method is agnostic to the structure of the encoder and decoder, so other models such as RNN-based ones would work as well. The objective of the encoder is to infer the posterior distribution $q(z|x)$ (which is assumed to be a multivariate Gaussian distribution $\mathcal{N}(\mu, \sigma^2 \mathbf{I})$), given the input sequence $x_{1:T}$:

$$\begin{aligned}
q(z|x) &= q(z_{1:T}|x_{1:T}) \sim \mathcal{N}(\mu_{1:T}, \sigma_{1:T}^2 \mathbf{I}) \\
&= \prod_t q(z_t | x_{\leq t}) \sim \prod_t \mathcal{N}(\mu_t, \sigma_t^2 \mathbf{I}) \\
\mu_t &= \text{Enc}_\mu(x_{\leq t}) \\
\sigma_t &= \text{Enc}_\sigma(x_{\leq t}),
\end{aligned} \tag{8}$$

where Enc_μ and Enc_σ are two transformer-based encoders to estimate the corresponding mean μ_t and covariance matrix $\sigma_t^2 \mathbf{I}$ of each input sequence. The detailed structure consists of the embedding layer and a stack of multiple self-attention layers.

Embedding layer. To transform the discrete item indices into continuous vectors, we adopt a learnable item embedding table $\mathbf{M} \in \mathbb{R}^{N \times d}$, where N is the total number of items and d is the dimension of item embeddings. For any user's historical item sequence, we follow the common practice [16] that first transforms it into a fixed-length of T , i.e., keeping the most recent T items for sequences of length greater than T , and padding learnable tokens to the left of sequences whose lengths are smaller than T . Through the above operations we obtain the input sequence of a user, which can be represented as a matrix $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_T]^\top \in \mathbb{R}^{T \times d}$, where $\mathbf{e}_t = \mathbf{M}_{x_t}$ is the embedding of item x_t in the look-up table. Since the self-attention mechanism is not aware of the positions of tokens, we explicitly add additional position embeddings to the input embedding, which can be formalized as follows:

$$\hat{\mathbf{E}} = \mathbf{E} + \mathbf{P} = \begin{bmatrix} \mathbf{e}_1 + \mathbf{p}_1 \\ \mathbf{e}_2 + \mathbf{p}_2 \\ \dots \\ \mathbf{e}_T + \mathbf{p}_T \end{bmatrix}, \tag{9}$$

where \mathbf{p}_n is the n -th position embedding, and is a learnable vector.

Self-attention layer. The objective of SR is to recommend users of interest according to their historical interaction records. Previous methods like RNN [13, 44] and Markov Chain [10, 11, 31] have issues of forgetting or short-range attention. The self-attention mechanism learns item representations attending to all items with adaptive weights, thus addressing the above issues. The input of the self-attention layer is $\hat{\mathbf{E}}$, and the operation can be formulated as follows:

$$\text{SA}(\hat{\mathbf{E}}) = \text{softmax} \left(\frac{(\hat{\mathbf{E}}\mathbf{W}^Q)(\hat{\mathbf{E}}\mathbf{W}^K)^\top}{\sqrt{d}} \right) (\hat{\mathbf{E}}\mathbf{W}^V), \tag{10}$$

where \mathbf{W}^Q , \mathbf{W}^K and \mathbf{W}^V are the learnable weight matrices that map the sequence embedding to different space. The self-attention operation compute item hidden representation as a weighted sum of all other item embeddings within this sequence, and the softmax function learns corresponding weights. d is the dimension of latent item embeddings and \sqrt{d} is the normalization factor that avoids

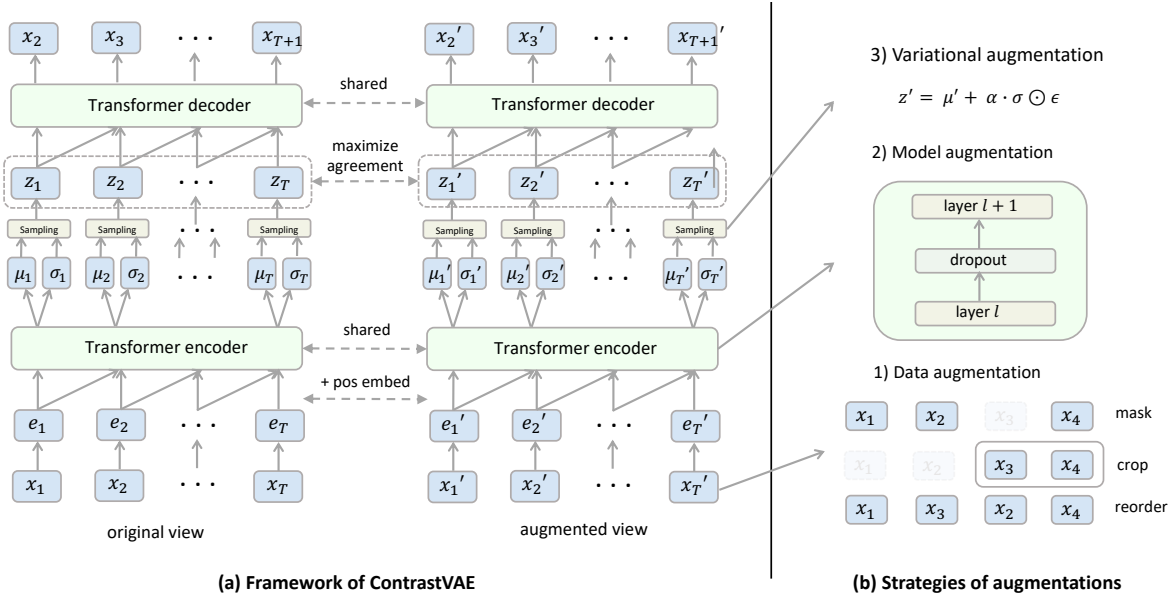


Figure 1: A high-level illustration of ContrastVAE. (a) The model framework: the model takes the original view $x_{1:T}$ and an augmented view $x'_{1:T}$ of the same sequence as input, and estimate the corresponding posterior distributions $q(z|x)$ and $q(z'|x')$ through a Transformer-based encoder. With sampled latent variables $z_{1:T}$ and $z'_{1:T}$, another Transformer-based decoder is used to perform next-item prediction tasks. The agreement of the latent variables of the two views is maximized through contrastive learning. (b) Strategies for generating the augmented view, where we consider three strategies that work at different steps of our model: 1) Data augmentation in input space, such as random masking, cropping or reordering of a sequence; 2) Model augmentation, which performs random dropout in each intermediate layer of the Transformer encoder; 3) Variational augmentation, which introduces a learnable Gaussian dropout rate at sampling step.

large values in softmax function. To avoid information leakage and shortcut-learning, i.e., the model is aware of the items to be clicked in the future of time-step t , we mask all items after time step t when computing the hidden representation of the sequence at t .

Stacking self-attention layers. Since self-attention is the linear operation, we apply multi-layer perceptions over the output of self-attention layers with *relu* as a non-linear activation function. Following Kang and McAuley [16], we also apply batch-normalization [15], residual connection, and dropout as the model goes deeper. Finally, we stack multiple self-attention layers to form our encoder model $q(z|x)$. This encoder takes the sequence embedding matrix as input and outputs the mean $\mu_{1:T}$ and standard deviation $\sigma_{1:T}$ of the amortized posterior Gaussian distribution.

Then we can use the reparameterization trick to sample the latent variable z accordingly:

$$z_{1:T} = \mu_{1:T} + \sigma_{1:T} \odot \epsilon, \quad (11)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is Gaussian noise.

4.3.2 Transformer decoder. The objective of decoder is to learn the underlying generative process of data given the estimated distribution of latent factors, and predict next item of interest, which is

reached by maximizing the reconstruction likelihood term in Eq. 3:

$$\begin{aligned} \mathbb{E}_{q(z|x)} \log[p(x|z)] &= \mathbb{E}_{q(z|x)} \log[p(x_{1:T}|z_{1:T})] \\ &= \sum_t \mathbb{E}_{q(z|x)} [\log p(x_t|z_{<t})]. \end{aligned} \quad (12)$$

The decoder has almost the same structure as the encoder in Sec. 4.3.1, except that it does not require an embedding layer.

4.3.3 Augmentation strategies. To facilitate ContrastVAE with two-view input, we can use some augmentation strategies to create a second view for each input sequence. Although conventional data augmentation methods for sequential data, such as random cropping, shuffling, and masking [41] can be naturally incorporated into our model, Qiu et al. [29] argues that they might lead to inconsistency problems between two augmented views, especially when the sequences are very short. To mitigate this issue, we resort to model-level augmentation methods that directly introduce randomness in the latent space.

Recent studies [9, 29] show that the simple dropout operation is powerful enough to generate informative views for CL. And model augmentation [22] is also studied to improve the sequential recommendation performance. These motivate us to merely use dropout operation at each intermediate layer of the encoder, which could be formulated as:

$$h_{1:T}^{(l+1)} = h_{1:T}^{(l)} + \text{SA}(h_{1:T}^{(l)} \odot \xi) \text{ with } \xi_t \sim \text{Bern}(p), \quad (13)$$

where $\mathbf{h}_{1:T}^{(l)}$ is the embedding at the l -th layer, ξ_t is 0/1 masking sampled from Bernoulli distribution and p is the dropout ratio. Note that we have added the residual connection term in Eq. 13.

Besides conventional Bernoulli dropout, Gaussian dropout [34], which multiplies output with random Gaussian noise, is an effective alternative with a faster convergence rate. Note that in the reparameterization trick in Eq. 11, we also multiply the standard deviation σ with random Gaussian noise to generate samples. This connection motivates us to adopt variational dropout [5], a special Gaussian dropout applied to the reparameterization step with a learnable dropout ratio:

$$\mathbf{z}_{1:T} = \boldsymbol{\mu}_{1:T} + \alpha \cdot \boldsymbol{\sigma}_{1:T} \odot \boldsymbol{\epsilon} \quad (14)$$

where α is the learnable weight parameter to control the Gaussian dropout ratio. Blum et al. [5] further interprets Eq. 14 from the Bayesian regularization perspective, where the KL divergence of the posterior distribution $q_\alpha(\mathbf{z}|\mathbf{x})$ (as a function of α) and the prior $p(\mathbf{z})$ is minimized, if an additional regularization term of α is jointly optimized with the model:

$$\mathcal{L}_\alpha = 0.5 \log \alpha + 1.161\alpha - 1.502\alpha^2 + 0.586\alpha^3. \quad (15)$$

In our implementation, dropout is equipped to both branches to create two views. The first branch always takes the original sequence (without augmentation) as input, and we term the input data and corresponding latent variable by \mathbf{x} and \mathbf{z} . For the second branch, we apply additional strategies like data augmentations or variational augmentations to create a stronger augmented view. The input and corresponding latent variable is termed by \mathbf{x}' and \mathbf{z}' . We provide an empirical study of different augmentation strategies in Sec. 5.3.2.

4.4 Model Optimization

With the model illustrated above, we can optimize the ContrastELBO presented in Sec. 4.2. In this section, we detailedly introduce our objective functions and how they correspond to different terms of the ContrastELBO in Eq. 3. Note that the optimizations of the reconstruction term and KL-divergence term for (\mathbf{x}, \mathbf{z}) and $(\mathbf{x}', \mathbf{z}')$ are identical and we only illustrate that for (\mathbf{x}, \mathbf{z}) .

To optimize the reconstruction term, we formalize it as a next-item prediction task where the log likelihood could be factorized as follows:

$$\log p(\mathbf{x}|\mathbf{z}) = \sum_{t=1}^T \log p(\mathbf{x}_t|\mathbf{z}_{<t}). \quad (16)$$

Denote the output of the encoder at position t by \mathbf{D}_t , and the embedding of item i by \mathbf{M}_i , then we define the probability that the next item is item i given the hidden state before time t \mathbf{z}_t as:

$$p(\mathbf{x}_{t+1} = i|\mathbf{z}_{\leq t}) \propto \mathbf{D}_t^\top \mathbf{M}_i. \quad (17)$$

For each time step t , we denote the ground truth next item by \mathbf{M}_t . Then following Kang and McAuley [16], we randomly sample another item as the negative example \mathbf{M}_n , thus we can optimize the expected log-likelihood term by minimizing the cross-entropy loss:

$$\mathcal{L}_{CE} = \sum_{t=1}^T \left[\log(\sigma(\mathbf{D}_t^\top \mathbf{M}_t)) + \log(\sigma(1 - \mathbf{D}_t^\top \mathbf{M}_n)) \right], \quad (18)$$

where σ is the sigmoid function.

For the KL-divergence term, we assume the prior $p(\mathbf{z}) = \prod_t p(\mathbf{z}_t)$ and $p(\mathbf{z}_t) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ for all t . Combining with the encoder's design in Sec. 4.3.1 that $q(\mathbf{z}|\mathbf{x})$ is another Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I})$, the closed-form solution of the KL-divergence could be easily computed through:

$$\begin{aligned} \mathcal{L}_{KL} &= D_{KL}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})] \\ &= \sum_{t=1}^T D_{KL}[q(\mathbf{z}_t|\mathbf{x}_t)||p(\mathbf{z}_t)] \\ &= \sum_{t=1}^T \sum_{d=1}^D (\sigma_{t,d}^2 + \mu_{t,d}^2 - 1 - \log \sigma_{t,d}^2) \end{aligned} \quad (19)$$

To maximize the mutual information term between \mathbf{z} and \mathbf{z}' under $q(\mathbf{z}, \mathbf{z}')$, we adopt the InfoNCE loss function [6, 37], which is a multi-sample unnormalized lower bound of mutual information with low variance [28]. Denote u 's hidden representation by \mathbf{z}_u , and its positive sample by \mathbf{z}'_u , then the loss could be formalized as:

$$\mathcal{L}_{InfoNCE} = \frac{1}{M} \sum_{u=1}^M \log \frac{\exp(\mathbf{z}_u^\top \mathbf{z}'_u / \tau)}{\sum_v \exp(\mathbf{z}_u^\top \mathbf{z}'_v / \tau) + \sum_{v \neq u} \exp(\mathbf{z}_u^\top \mathbf{z}_v / \tau)}, \quad (20)$$

where τ is the temperature hyperparameter, $\mathbf{z}_u = \mathcal{R}(\mathbf{z}_{1:T}^u)$ is a representation of the i -th sequence, which summarizes the latent representations of all tokens of the sequence. We use the simple average pooling to implement $\mathcal{R}(\cdot)$ in this paper.

Plugging Eqs. 18, 19, 20 into the ContrastELBO in Eq. 3, we get the final objective function (note that Eq. 15 is required as well when using variational dropout):

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{CE} - \mathcal{L}_{KL} \quad (\text{ELBO of the 1st view}) \\ &+ \mathcal{L}'_{CE} - \mathcal{L}'_{KL} \quad (\text{ELBO of the 2nd view}) \\ &+ \lambda \cdot \mathcal{L}_{InfoNCE} \quad (\text{InfoNCE}) \end{aligned} \quad (21)$$

where λ is the weight of InfoNCE loss as Eq. 20 is an unnormalized estimation of mutual information and does not reflect its magnitude.

5 EXPERIMENT

In this section, we evaluate the proposed ContrastVAE empirically on real-world SR tasks, and we would like to answer the following research questions:

- **RQ1:** How does ContrastVAE perform compared with state-of-the-art SR models?
- **RQ2:** Are the key components in ContrastVAE, such as augmentations and contrastive learning necessary and beneficial for satisfactory improvement?
- **RQ3:** How is the performance of ContrastVAE on items with different frequencies and sequences with different lengths? Does ContrastVAE improve the performance on long-tail items and what's the reasons?
- **RQ4:** How is the robustness of ContrastVAE w.r.t. noisy input sequences, and is ContrastVAE sensitive to some key model hyperparameters?

5.1 Setups

Dataset. We conduct experiments on four datasets collected from Amazon review in various domains [26]: *Beauty*, *Toys and Games*

(Toys), Tools and Home (Tools) and Office Product (Office). We treat all the user-item rating records as implicit feedback and sort them according to the timestamps to form sequences. Following the common settings [16], we filter out users and items with less than five interaction records. For each user, we use the last clicked item for testing, the penultimate one for validation and the remaining clicked items for training. We provide the statistics of the four datasets in Table 1. Furthermore, we split the sequences into five groups according to the latest interacted item’s frequency, and report the population of each group in Table 2, where we can observe that most items lie in the smallest frequency group and exhibits a long-tail distribution.

Table 1: Statistics of datasets, we report the number of users, number of items, number of interactions, number of interactions per item, and the averaged sequence length.

Dataset	#Users	#Items	#Interactions	#Ints / item	Avg. seq. len.
Beauty	22,363	12,101	198,502	16.40	8.3
Toys	19,412	11,924	167,597	14.06	8.6
Tools	16,638	10,217	134,476	13.16	8.1
Office	4,905	2,420	53,258	22.00	10.8

Table 2: Number of sequences end at items of different frequency groups.

Dataset	[≤10]	[10, 20]	[20, 30]	[30, 40]	[≥40]
Beauty	17,353	3,152	1,065	367	426
Toys	16,345	2,320	476	130	141
Tools	13,929	1,769	400	230	310
Office	3,150	1,028	547	97	83

Metrics. For each user, we compute its relevance scores for all items and choose items with top- N scores for the recommendation. Then we adopt two widely used top- N metrics, Recall and NDCG, as our top- N ranking evaluation metrics. We report the experimental results when $N = 20$ and $N = 40$.

Baselines. We compare our methods with four types of representative SR models: 1) VAE-based methods, including SVAE [32] and ACVAE [42]; 2) CL-based methods, including S3Rec [47] and CL4Rec [41]; 3) attention-based methods, including Bert4Rec [35], SASRec [16]; 4) *probabilistic SR*, including STOSA [8], DT4SR [7]. We also compare our methods with collaborative filtering methods that ignore sequential information: LightGCN [12] and BPRMF [30].

Implementation Details. We use PyTorch to implement our model, and all experiments are conducted on a Nvidia V100 GPU with 16G memory. We use Adam [17] to optimize our method. To avoid overfitting, we adopt an early stop strategy to stop experiments when there is no improvement in 100 epochs. We set learning rate as 0.001, hidden dimension as 128, model dropout probability as 0.3 and number of attention heads as 4 for all models.

5.2 Comparative Results

We report the performance of ContrastVAE and other comparative methods in Table 3. Our model ContrastVAE consistently outperforms other methods throughout the four datasets. Concretely,

ContrastVAE achieves 8.86%- 19.34% improvements of Recall@40 and 4.42% - 10.12% improvements of NDCG@40 compared with the best baseline, which shows the effectiveness of our method for the SR. Furthermore, we find that on the Tool dataset, which has the smallest number of interactions per item and per user, our method outperforms the most powerful baselines by over 19%. We attribute this to the better capability of our model to learn from the noise and uncertainty of user behaviors, which is especially beneficial for short sequences (sparsity) and long-tail items.

5.3 Ablation Studies

In order to verify the effectiveness and necessities of the critical components and compare the performance of different designs of our method, we conduct several ablation studies for ContrastVAE.

5.3.1 Effect of CL and variational dropout. First of all, to validate the importance of CL and the regularization loss in variational dropout, we consider two variants of ContrastVAE: 1) removing the mutual information term (i.e., the InfoNCE loss in Eq. 20), 2) removing the regularization term of the variational dropout rate \mathcal{L}_α in Eq. 15. We present the results of Recall@20 and NDCG@20 on Toy and Tool datasets in Table 4. The results show that both the CL and the regularization of variational dropout are important to our model, and the model performance will degrade greatly if they are removed.

Table 4: Performance of ContrastVAE when removing the CL loss (w/o MI) and regularization of variational dropout (w/o \mathcal{L}_α).

Methods	Toy		Tool	
	R@20	N@20	R@20	N@20
w/o MI	0.112 (-3.60%)	0.052 (-5.85%)	0.070 (-4.11%)	0.031 (-5.83%)
w/o \mathcal{L}_α	0.104 (-10.9%)	0.051 (-6.03%)	0.063 (-13.69%)	0.028 (-14.4%)
default	0.116	0.055	0.073	0.033

5.3.2 Comparison of different augmentation strategies. To evaluate the effectiveness of different augmentation methods, we compare the performance when equipping different augmentation strategies to ContrastVAE, including data augmentation (DA), model augmentation (MA), and variational augmentation (VA): 1) For data augmentation, we follow Xie et al. [41] and adopt random cropping, masking and reordering together to generate a perturbed version for each input sequence, and use it as the input of the second branch of our model; 2) For model augmentation, we simply apply the basic dropout to the encoder of the second branch; 3) For variational augmentation, we apply variational dropout at the sampling step with reparameterization as introduced in Sec. 4.3.3. As shown in Table 5 all the three augmentation strategies improve the model’s performance by a large margin, compared with the baseline method AVAE that merely uses single-branch VAE without CL. We also have the following interesting observations: 1) compared with DA, which achieves the best performance on 2 out of 16 metrics, MA and VA are better on more metrics and datasets (6 out of 16 and 8 out of 16 respectively), and this demonstrates the limitations of DA, which

Table 3: Overall Comparison

Dataset	Metric	SVAE	ACVAE	S3Rec	CL4Rec	LightGCN	BPRMF	Bert4Rec	SASRec	STOSA	DT4SR	ContrastVAE	Improv.
Beauty	R@20	0.0268	0.0951	0.0946	0.0398	0.0759	0.0739	0.0890	0.0952	0.0975	<u>0.0982</u>	0.1095	11.51%
	R@40	0.0417	0.1294	0.1348	0.0554	0.1112	0.1089	0.1285	0.1389	0.1337	<u>0.1404</u>	0.1541	9.76%
	N@20	0.0102	0.0467	0.0424	0.0168	0.0306	0.0311	0.0395	0.0420	<u>0.0469</u>	0.0446	0.0496	5.76%
	N@40	0.0132	0.0537	0.0505	0.0200	0.0378	0.0383	0.0476	0.0509	<u>0.0542</u>	0.0533	0.0587	8.30%
Office	R@20	0.0988	0.1327	0.1335	0.0646	0.0532	0.0483	0.1350	0.1478	<u>0.1578</u>	0.1429	0.1708	8.24%
	R@40	0.1647	0.2075	0.2112	0.1025	0.0797	0.0718	0.2230	0.2251	<u>0.2391</u>	0.2186	0.2617	9.45%
	N@20	0.0389	0.0560	0.0571	0.0291	0.0243	0.0218	0.0551	0.0657	<u>0.0694</u>	0.0643	0.0741	6.77%
	N@40	0.0523	0.0713	0.0729	0.0368	0.0297	0.0266	0.0729	0.0815	<u>0.0859</u>	0.0797	0.0925	7.68%
Toy	R@20	0.0178	0.0722	0.0973	0.0392	0.0671	0.0692	0.0699	0.1112	0.1008	<u>0.1130</u>	0.1164	3.01%
	R@40	0.0260	0.1030	0.1307	0.0596	0.0977	0.1007	0.0982	<u>0.1479</u>	0.1357	0.1478	0.1610	8.86%
	N@20	0.0069	0.0359	0.0467	0.0182	0.0287	0.0304	0.0318	<u>0.0539</u>	0.0496	0.0515	0.0547	1.48%
	N@40	0.0086	0.0421	0.0536	0.0224	0.0349	0.0369	0.0376	<u>0.0614</u>	0.0567	0.0560	0.0638	4.42%
Tool	R@20	0.0340	0.0537	0.0632	0.0443	0.0537	0.0505	0.0508	<u>0.0640</u>	0.0615	0.0601	0.0731	14.21%
	R@40	0.0521	0.0759	0.0849	0.0634	0.0751	0.0715	0.0777	<u>0.0879</u>	0.0867	0.0861	0.1049	19.34%
	N@20	0.0149	0.0249	0.0286	0.0194	0.0238	0.0219	0.0213	0.0294	<u>0.0295</u>	0.0289	0.0326	10.51%
	N@40	0.0186	0.0294	0.0330	0.0233	0.0282	0.0262	0.0268	0.0345	<u>0.0346</u>	0.0342	0.0381	10.12%

will introduce inconsistency between augmented views. 2) DA performs best on the Office dataset. This might be because the Office dataset has the largest average sequence length (see Table 1) and thus is less sensitive to the perturbation of data augmentations. 3) MA method performs competitively compared with DA but is much simpler. 4) Our proposed VA achieves comparable or even better results, especially on Toy and Tool datasets with smaller average sequence length and testing item frequency. This shows that the proposed VA can effectively benefit the prediction of short sequences and long-tail items.

Table 5: Comparison of the performance of different augmentation strategies on the four datasets: AVAE is short for the single-branch VAE model that uses attentive encoder and decoder without CL. DA is short for data augmentation, MA is short for model augmentation, and VA is short for variational augmentation.

Dataset	Metric	AVAE	DA	MA	VA
Beauty	R@20	0.0448	0.1059	0.1095	0.1066
	R@40	0.0709	0.1561	0.1541	0.1578
	N@20	0.0180	0.0459	0.0496	0.0464
	N@40	0.0233	0.0562	0.0587	0.0568
Office	R@20	0.1093	0.1745	0.1708	0.1672
	R@40	0.1918	0.2658	0.2617	0.2599
	N@20	0.0419	0.0739	0.0741	0.0722
	N@40	0.0586	0.0924	0.0925	0.0911
Toy	R@20	0.0423	0.1112	0.1130	0.1164
	R@40	0.0700	0.1554	0.1548	0.1610
	N@20	0.0171	0.0503	0.0566	0.0547
	N@40	0.0227	0.0593	0.0652	0.0638
Tool	R@20	0.0380	0.0671	0.0691	0.0731
	R@40	0.0603	0.1004	0.0986	0.1049
	N@20	0.0164	0.0295	0.0310	0.0326
	N@40	0.0209	0.0364	0.0370	0.0381

5.4 Analysis of ContrastVAE

In this section, we study the strengths of the proposed ContrastVAE through analyzing its property from several aspects.

5.4.1 ContrastVAE benefits the prediction of long-tail items and for short sequences. We first investigate ContrastVAE’s performance on long-tail items, which have much fewer interactions in the training set and thus are much harder to predict. To reach this target, we follow the grouping strategies in Table 2 and categorize the user sequences into 5 groups according to the frequencies of their last clicked items. We report Recall@40 of ContrastVAE and representative baseline methods on Toy dataset in Fig. 2a. We can observe that ContrastVAE achieves the highest Recall@40 scores on all groups of sequences. Specifically, on long-tail items (i.e., $[\leq 10]$ and $[10, 20]$), our method outperforms other baseline models by a large margin. The improvement over such user sequences greatly contributes to the overall performance of our method, as most items to be predicted are long-tail items.

We further study how the sequence length affects the model’s performance. Similar to the item frequency, we split user sequences into 5 groups according to their lengths, and we report the performance of ContrastVAE and other models on Toy dataset in Fig. 2b. It demonstrates that ContrastVAE consistently exhibits good performance on sequences with various lengths, which highlights the effectiveness of our model. Similar to the results on long-tail items, ContrastVAE greatly improves the performance of short sequences (i.e., user sequences which have less than 20 interactions). However, on longer sequences (e.g., $[\geq 40]$), ContrastVAE does not show superior performance compared with other models, and we guess that for long sequences, the users’ preferences tend to become certain and easy to predict, in which case the uncertainty and randomness introduced by our method would not help the prediction results.

5.4.2 ContrastVAE alleviates posterior collapse and point-estimation in latent space. We then study why ContrastVAE can improve VAEs for sequential representation through analyzing the posterior distributions learned from ContrastVAE and the vanilla VAE model AVAE. We adopt two metrics to evaluate the quality of learned latent variables z : 1) the average KL divergence between the posterior distribution $p(z|x)$ of sequences and the standard Gaussian distribution $\mathcal{N}(0, I)$, which reflects the extent of posterior collapse problem (posterior collapse induces small KL-divergence); 2) the

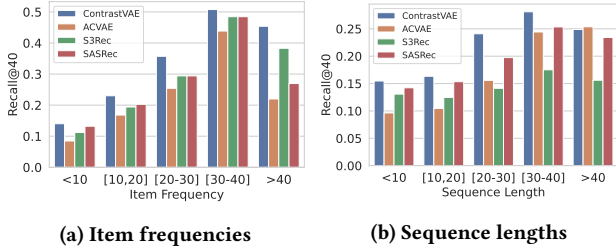


Figure 2: Subgroup sequences analysis on Toy dataset

average variance of latent variables, which reflect the extent of variance vanishing.

We report the Recall@40 scores and these two metrics of sequences targeted items with difference frequencies in Fig. 3. For AVAE model, the KL-divergences between latent variable estimation and standard Gaussian distribution are very low, especially for infrequent items, indicating that they suffer from posterior collapse problems. Then in this case the representations of sequences that target different items would be hard to discriminate in latent space, leading to poor performance of AVAE model on infrequent items. By contrast, ContrastVAE alleviates the posterior degeneration through CL, which implicitly encourages different sequences to have diverse latent representations. An interesting observation is that the KL-divergence of both ContrastVAE and AVAE decrease when item’s frequency gets extremely large (e.g., > 30), and we guess it is because popular items are much easier to predict and thus the model does not require a fine-grained posterior estimation to model it. From Fig. 3b we notice that AVAE model have vanishing variance over infrequent items, which indicates that AVAE collapses to point estimation for such infrequent items. On the contrary, ContrastVAE managed to increase the average variance over such sequences, thus increasing the robustness of noise.

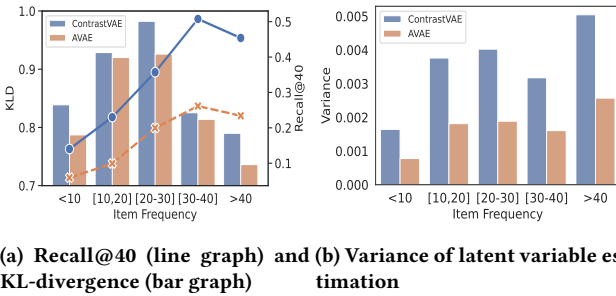


Figure 3: Latent space analysis on Toy dataset.

5.4.3 *Robustness analysis.* We further study the impact of corrupted input sequences for ContrastVAE to analyze its robustness w.r.t. to noisy data. We consider two corrupting strategies: 1) randomly deleting a proportion of items in each sequence (random deletion); 2) randomly replacing a proportion items by other items in each sequence (random replacement). Consistent with Sec. 5.4.1, we compare our method with typical baseline models ACVAE, S3Rec

and SASRec. As shown in Fig. 4, the performance of all models exhibits a drop as we increase the corruption ratio. However, ContrastVAE always outperforms other baseline models by a large margin whatever the corruption method and the corruption ratio, which indicates that ContrastVAE can still exhibit good performance for noisy input data.

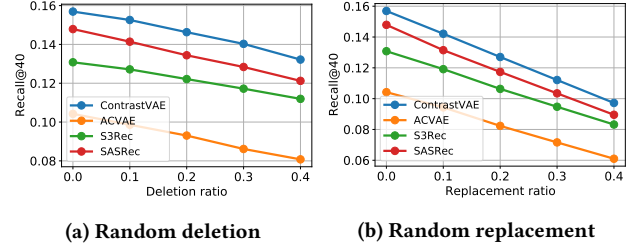


Figure 4: Robustness analysis on Toy dataset.

5.4.4 *Hyper-parameter sensitivity analysis.* We finally study the performance variation of our model w.r.t. the intensity of the CL by tuning the weight of MI loss λ , and we present the Recall@40 scores on Beauty and Office dataset in Fig. 5. We find that compared with not using mutual information (i.e., the weight λ is set as 0), a proper weight can lead to great improvements (up to 7.6%). Also, the weight cannot be too large otherwise it would constrain the model’s learning from the next-item prediction tasks.

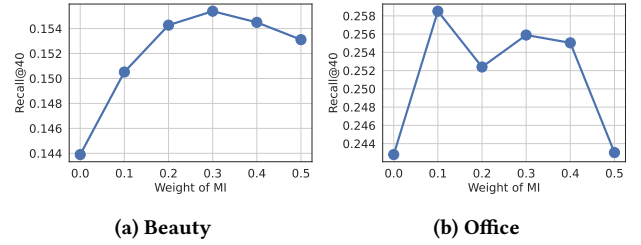


Figure 5: Sensitivity analysis of the weight of contrastive loss.

6 CONCLUSION

In this paper, we have proposed ContrastVAE, a novel method for SR. We start from extending single-view evidence lower-bound to two-view cases, and derive ContrastELBO. Then to optimize ContrastELBO for SR tasks, we propose ContrastVAE, which takes two views of sequential data as input and optimizes an additional mutual information maximization term besides conventional reconstruction loss and kl-divergence loss of two views. We further propose model augmentation and variational augmentation for generating another view for an input sequence to solve the inconsistency problem led by conventional data augmentation methods. Experiment results and analysis show that our architecture combined with augmentations outperforms other competitive baselines.

7 ACKNOWLEDGEMENTS

This work is supported in part by NSF under grants III-1763325, III-1909323, III-2106758, and SaTC-1930941.

REFERENCES

- [1] Laurence Aitchison. 2021. InfoNCE is a variational autoencoder. *arXiv preprint arXiv:2107.02495* (2021).
- [2] Junwen Bai, Weiran Wang, and Carla P. Gomes. 2021. Contrastively Disentangled Sequential Variational Autoencoder. *CoRR abs/2110.12091* (2021).
- [3] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, R. Devon Hjelm, and Aaron C. Courville. 2018. Mutual Information Neural Estimation. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10–15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*. PMLR, 530–539.
- [4] Shuqing Bian, Wayne Xin Zhao, Kun Zhou, Jing Cai, Yancheng He, Cunxiang Yin, and Ji-Rong Wen. 2021. Contrastive Curriculum Learning for Sequential User Behavior Modeling via Data Augmentation. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*. ACM, 3737–3746.
- [5] Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. 2015. Variational Dropout and the Local Reparameterization Trick. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7–12, 2015, Montreal, Quebec, Canada*, Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (Eds.). 2575–2583.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 1597–1607.
- [7] Zhiwei Fan, Zhiwei Liu, Shen Wang, Lei Zheng, and Philip S. Yu. 2021. Modeling Sequences as Distributions with Uncertainty for Sequential Recommendation. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*. ACM, 3019–3023.
- [8] Zhiwei Fan, Zhiwei Liu, Yu Wang, Alice Wang, Zahra Nazari, Lei Zheng, Hao Peng, and Philip S. Yu. 2022. Sequential Recommendation via Stochastic Self-Attention. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*. ACM, 2036–2047.
- [9] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7–11 November, 2021*. Association for Computational Linguistics, 6894–6910.
- [10] Ruining He and Julian J. McAuley. 2016. Fusing Similarity Models with Markov Chains for Sparse Sequential Recommendation. In *IEEE 16th International Conference on Data Mining, ICDM 2016, December 12–15, 2016, Barcelona, Spain*. IEEE Computer Society, 191–200.
- [11] Ruining He and Julian J. McAuley. 2016. Fusing Similarity Models with Markov Chains for Sparse Sequential Recommendation. In *IEEE 16th International Conference on Data Mining, ICDM 2016, December 12–15, 2016, Barcelona, Spain*. IEEE Computer Society, 191–200.
- [12] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25–30, 2020*. ACM, 639–648.
- [13] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*.
- [14] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net.
- [15] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6–11 July 2015 (JMLR Workshop and Conference Proceedings, Vol. 37)*. JMLR.org, 448–456.
- [16] Wang-Cheng Kang and Julian J. McAuley. 2018. Self-Attentive Sequential Recommendation. In *IEEE International Conference on Data Mining, ICDM 2018, Singapore, November 17–20, 2018*. IEEE Computer Society, 197–206.
- [17] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*.
- [18] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*.
- [19] Jiacheng Li, Yujie Wang, and Julian J. McAuley. 2020. Time Interval Aware Self-Attention for Sequential Recommendation. In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3–7, 2020*. ACM, 322–330.
- [20] Xiaopeng Li and James She. 2017. Collaborative Variational Autoencoder for Recommender Systems. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*. ACM, 305–314.
- [21] 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 on World Wide Web, WWW 2018, Lyon, France, April 23–27, 2018*. ACM, 689–698.
- [22] Zhiwei Liu, Yongjun Chen, Jia Li, Man Luo, Philip S Yu, and Caiming Xiong. 2022. Improving Contrastive Learning with Model Augmentation. *arXiv preprint arXiv:2203.15508* (2022).
- [23] Zhiwei Liu, Yongjun Chen, Jia Li, Philip S. Yu, Julian J. McAuley, and Caiming Xiong. 2021. Contrastive Self-supervised Sequential Recommendation with Robust Augmentation. *CoRR abs/2108.06479* (2021).
- [24] Zhiwei Liu, Ziwei Fan, Yu Wang, and Philip S Yu. 2021. Augmenting sequential recommendation with pseudo-prior items via reversely pre-training transformer. In *Proceedings of the 44th international ACM SIGIR conference on Research and development in information retrieval*. 1608–1612.
- [25] James Lucas, George Tucker, Roger B. Grosse, and Mohammad Norouzi. 2019. Don't Blame the ELBO! A Linear VAE Perspective on Posterior Collapse. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*. 9403–9413.
- [26] Julian J. McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-Based Recommendations on Styles and Substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9–13, 2015*. ACM, 43–52.
- [27] Preksha Nema, Alexandros Karatzoglou, and Filip Radlinski. 2021. Disentangling Representations for Recommendation Critiquing with β -VAE. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*. ACM, 1356–1365.
- [28] Ben Poole, Sherjil Ozair, Aäron van den Oord, Alexander A. Alemi, and George Tucker. 2019. On Variational Bounds of Mutual Information. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 5171–5180.
- [29] Ruihong Qiu, Zi Huang, Hongzhi Yin, and Zijian Wang. 2022. Contrastive Learning for Representation Degeneration Problem in Sequential Recommendation. In *WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022*. ACM, 813–823.
- [30] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18–21, 2009*. AUAI Press, 452–461.
- [31] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26–30, 2010*. ACM, 811–820.
- [32] Noven Sachdeva, Giuseppe Manco, Ettore Ritacco, and Vikram Pudi. 2019. Sequential Variational Autoencoders for Collaborative Filtering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM 2019, Melbourne, VIC, Australia, February 11–15, 2019*. ACM, 600–608.
- [33] Ilya Shenbin, Anton Alekseev, Elena Tutubalina, Valentin Malykh, and Sergey I. Nikolenko. 2020. RecVAE: A New Variational Autoencoder for Top-N Recommendations with Implicit Feedback. In *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3–7, 2020*. ACM, 528–536.
- [34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [35] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3–7, 2019*. ACM, 1441–1450.
- [36] Yuhta Takida, Wei-Hsiang Liao, Toshimitsu Uesaka, Shusuke Takahashi, and Yuki Mitsufuji. 2021. Preventing Posterior Collapse Induced by Oversmoothing in Gaussian VAE. *CoRR abs/2102.08663* (2021). <https://arxiv.org/abs/2102.08663>
- [37] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748* (2018).
- [38] Tongzhou Wang and Phillip Isola. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020,*

- 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119). PMLR, 9929–9939.
- [39] Yu Wang, Zhiwei Liu, Ziwei Fan, Lichao Sun, and Philip S Yu. 2021. Dskreg: Differentiable sampling on knowledge graph for recommendation with relational gnn. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 3513–3517.
- [40] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised Graph Learning for Recommendation. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11–15, 2021*. ACM, 726–735.
- [41] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Bolin Ding, and Bin Cui. 2020. Contrastive learning for sequential recommendation. *arXiv preprint arXiv:2010.14395* (2020).
- [42] Zhe Xie, Chengxuan Liu, Yichi Zhang, Hongtao Lu, Dong Wang, and Yue Ding. 2021. Adversarial and Contrastive Variational Autoencoder for Sequential Recommendation. In *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19–23, 2021*. ACM / IW3C2, 449–459.
- [43] Liangwei Yang, Zhiwei Liu, Yu Wang, Chen Wang, Ziwei Fan, and Philip S Yu. 2022. Large-scale Personalized Video Game Recommendation via Social-aware Contextualized Graph Neural Network. In *Proceedings of the ACM Web Conference 2022*. 3376–3386.
- [44] 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, SIGIR 2016, Pisa, Italy, July 17–21, 2016*. ACM, 729–732.
- [45] Jing Zhao, Pengpeng Zhao, Lei Zhao, Yanchi Liu, Victor S. Sheng, and Xiaofang Zhou. 2021. Variational Self-attention Network for Sequential Recommendation. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19–22, 2021*. IEEE, 1559–1570.
- [46] Shengjia Zhao, Jiaming Song, and Stefano Ermon. 2019. InfoVAE: Balancing Learning and Inference in Variational Autoencoders. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. AAAI Press, 5885–5892.
- [47] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-Rec: Self-Supervised Learning for Sequential Recommendation with Mutual Information Maximization. In *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19–23, 2020*. ACM, 1893–1902.