

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/182553>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Explainable Recommender with Geometric Information Bottleneck

Hanqi Yan, Lin Gui, Menghan Wang, Kun Zhang, Yulan He

Abstract—Explainable recommender systems can explain their recommendation decisions, enhancing user trust in the systems. Most explainable recommender systems either rely on human-annotated rationales to train models for explanation generation or leverage the attention mechanism to extract important text spans from reviews as explanations. The extracted rationales are often confined to an individual review and may fail to identify the implicit features beyond the review text. To avoid the expensive human annotation process and to generate explanations beyond individual reviews, we propose to incorporate a geometric prior learnt from user-item interactions into a variational network which infers latent factors from user-item reviews. The latent factors from an individual user-item pair can be used for both recommendation and explanation generation, which naturally inherit the global characteristics encoded in the prior knowledge. Experimental results on three e-commerce datasets show that our model significantly improves the interpretability of a variational recommender using the Wasserstein distance while achieving performance comparable to existing content-based recommender systems in terms of recommendation behaviours.

Index Terms—Recommendation System, Interpretability, Information bottleneck

I. INTRODUCTION

Typically, a recommender system compares users' preferences with item characteristics (e.g., item descriptions or item-associated reviews) or studies user-item historical interactions (e.g., ratings, purchases or clicking behaviours) in order to identify items that are likely of interest to users. In addition to item recommendation, interpretable recommenders also generate the rationale behind the recommendation decision [1], [2]. Most existing interpretable recommenders can either generate rationale or extract text spans from a given user-item review as explanations of model decisions. Rationale generation requires annotated data for training, e.g., short comments provided by users explaining their behaviours of interacting with items. The extractive explainable recommenders either rely on annotated sentiment-bearings aspect spans in reviews for training [3]–[7] or extract sentences from a target item's review as recommendation rationales based on attention weights [8]–[15].

We argue that explanation generation models are often supervised by human-annotated rationales, which are labour-intensive to obtain in practice. The explainable recommenders built on extractive approaches can only extract rationales from an individual user-item review document. They suffer from the following limitations. First, some reviews may be too general to explain the user rating, rendering them useless for

explanation generation. For example, the review '*I really like the smartphone, will recommend it to my friends*' does not provide any clue why the user likes the smartphone. Second, features directly extracted from a review document may fail to reflect some global properties which can only be identified from implicit user-item interactions. Meaningful insights could still be derived from reviews towards items that are not directly purchased/rated by a user but preferred by other like-minded users.

To address the aforementioned limitations, we propose a Geometric InformAtIoN boTtleneck (**GIANT**) framework to incorporate the prior learned from a user-item interaction graph to refine the induced latent factors learnt from reviews. Essentially, we use the prior learned from one modality (i.e., the user-item interaction graph) to constrain the learning of latent factors in another modality (i.e., the review text). In this way, we naturally address the problem of the oversimplified Gaussian prior assumption in approaches built on the Variational Autoencoder (VAE) [16]. More concretely, our prior is derived from the user/item clusters obtained from the user-item interaction graph. Each user or item has its soft cluster assignment represented as a probability distribution over clusters, which is in proportional to its distance to each cluster centroid. For this reason, we call our prior the geometric prior as it translates the geometric properties of users/items into their cluster-associated probability distributions. The prior is subsequently incorporated into the variational network for constraining the learning of latent factors from the review text. For a user-item pair, (u, i) , we feed the user's past reviews and the item's associated reviews (excluding u 's review on i) to the variational network for the inference of user- and item-associated latent factors, which are subsequently used for rating prediction and explanation generation.

Figure 1 illustrates how the proposed geometric prior can be used to develop the explainable recommender. To impose the prior on the latent variables of the variational framework taking the set of user reviews and item reviews as input, we apply the KL divergence to minimise the discrepancy between the posterior distribution of the user latent factor z_u , and the item latent factor z_i , learned from review text, and the geometric prior derived from the user-item interaction graph. The latent factors inferred by the variational network can be treated as topics summarising the key semantics of input reviews. For an input user-item pair, the most prominent dimension of the user (or item) latent factor is considered as the most important topic revealing the user preference (or item characteristic). As will be shown in the experiments section, for each topic (i.e., a dimension of the latent factor), we can

This paper was produced by the IEEE Publication Technology Group. They are in Piscataway, NJ.

Manuscript received April 19, 2021; revised August 16, 2021.

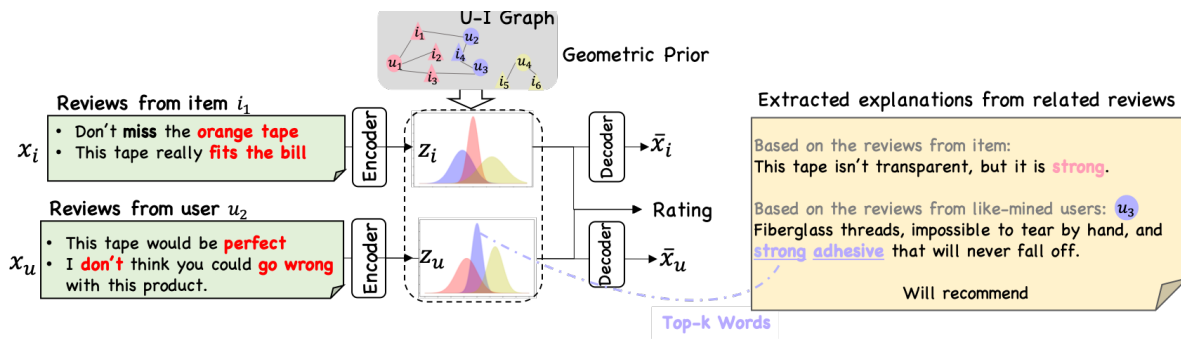


Fig. 1. The proposed geometric prior reflects the U-I distances in their interaction graph, frequently interacted users/items are assigned to the same cluster (illustrated in node colour). Existing extractive methods can only identify indicative words (shown in red) from the given review pair (x_i, x_u) , while our method can extract sentence-level reviews written by other like-minded users or on similar items, which are assigned to the same latent dimension/cluster, i.e., we extract the sentences from the review by user u_3 which is in the same cluster as user u_2 , as the reviews of u_2 are too general to summarise the key features.

retrieve its top- k associated words as its representation. To generate sentence-level explanations, we identify sentences from reviews residing in the same user/item cluster which are most similar to the user- and item-topic representations.¹

Experimental results on the three commonly used benchmarking datasets show that the proposed method achieves performance comparable with several strong baselines in recommendation prediction. Moreover, both quantitative and qualitative analysis shows we can see clear and meaningful clusters distributed in the latent space regularised by the geometric prior, when compared with VAEs using the standard Gaussian prior or employing the Wasserstein distance. The user study demonstrates that our generated explanations are better in *relevance*, *faithfulness* and *informativeness*. Our contributions can be summarized below:

- We proposed a geometric prior derived from user-item interactions in order to incorporate the global features from similar users/items, allowing recommendation explanations to be generated beyond explicit user-item interactions.
- We give a proof that the geometric prior can be effectively incorporated into the variational framework with an upper bound in mutual information between the textual input and the latent representations regularized by the geometric prior.
- Experimental results show the efficiency of our proposed method in both review rating prediction and explanation generation.

II. RELATED WORK

We review recommender systems, with particular attention to those built on VAE or offering explanations of recommendation decisions.

a) *Recommender with VAE*: Text data such as user reviews, item or brand descriptions could be important for developing a high-quality recommender system as they can be exploited to address the sparsity issue in user-item interactions. HFT [17] and CTR [18] adopted Latent Dirichlet Allocation

(LDA) [19] to extract latent topics from review text. VAE can also be used for text modeling due to its ability in extracting latent and interpretable features [20]–[22]. [21] argued that the commonly used isotropic Gaussian in VAE is over-simplified and proposed BiVAE by introducing constrained adaptive prior for learning user- and item-dependent prior distributions. More recently, review information is used to disentangle the latent user intents at the finer granularity, for example, DisenGCN [23] and DisenHan [22] used the graph attention mechanism to differentiate multiple relations and features.

b) *Explainable Recommender System*: One popular interpretable method is to train a language generation model with the ground-truth explanations supplied, which can be the first sentence of a given review or human-annotated text spans in the review text [4], [24]. Some of these approaches firstly extracted user-aspect-opinion and item-aspect-opinion tuples as ground truth, then modelled the interactions between the two tuples to derive the rating and the indicative features [25], [26]. Retrieval-based (aka. extractive) solutions directly select representative text spans (rationales), i.e., words or sentences, from a target item's past reviews to explain the recommender behaviours. Attention-based methods are prevalent in review-based explainable recommenders [14], [15]. Beyond leveraging attention weights to select rationales, [27], [28] proposed different optimization objectives to refine the selection.

However, the aforementioned approaches suffer from two major issues. Some approaches require aspects to be extracted first in order to provide aspect-related predictive features as interpretations, thus heavily dependent on the quality of the aspect extractors employed; while others rely on human-annotated rationales to train supervised language generators for explanation generation. Our proposed approach can instead generate explanations with only supervision signals coming from review ratings.

III. GEOMETRIC INFORMATION BOTTLENECK

To incorporate the global features from actively interacted users/items to explain the recommendation behaviour, we propose to incorporate the geometric prior derived from user-item interaction graphs to learn the latent factors of users and items from review text in a variational network. Although

¹Details about how to extract sentence-level review as explanations can be found in §V: Explanation Generation.

various work has been proposed to adopt richer priors to infer a more complex and realistic posterior in standard VAEs [29], [30], our solution is different from existing approaches as we need to impose the prior knowledge learned from the graph modality to constrain the variables in the text modality. In what follows, we show how this can be done under the theory of information bottleneck.

Based on the Information Bottleneck (IB) theory [31], we can train an encoder-decoder architecture $x \xrightarrow{\text{encode}} z \xrightarrow{\text{decode}} \bar{x}$, where $x \in X$, $z \in Z$, $\bar{x} \in \bar{X}$ are the input, the hidden, and the reconstructed representations, respectively, in order to preserve the meaningful information about \bar{X} in Z while maximally compress the information of X . We achieve this by maximising the following:

$$O_{IB} = I(\bar{X}; Z) - \beta \cdot I(X; Z), \quad (1)$$

where $I(\cdot)$ denotes the mutual information, β is a Lagrange multiplier. The second term is to maximally compress X , equivalent to minimising the mutual information between X and Z , denoted as $I(X; Z) \leq I_c$, I_c is the upper bound.

Intuitively, we assume the input has representations $x^t \in X^t$ and $x^g \in X^g$ under different modalities² satisfying the constraint, $I(X^t; X^g) \geq I_x$, since X^t and X^g are derived from the same input, they should be relevant. Then, two independent well-trained encoder-decoder architectures in the two modalities optimised by Eq. (1) are constrained by $I(X^t; Z^t) \leq I_c$ and $I(X^g; Z^g) \leq I_c$, respectively. We derive a new upper bound for the multi-modality IB as follows:³

$$I(X^t; Z^g) \leq H(X^t) - H(X^g) + H(X^g|X^t) + I_c, \quad (2)$$

where $H(\cdot)$ and $H(\cdot|\cdot)$ are entropy and conditional entropy, respectively. Therefore, we can maximise the objective function using the input from the text modality, with a regularisation from another modality,⁴

$$O_{IB}^t = I(\bar{X}^t; Z^t) - \beta \cdot I(X^t; Z^g). \quad (3)$$

According to [32], the above equation can be approximated by applying the reparameterisation trick [33] with a random Gaussian noise ϵ :

$$O_{IB}^t = \mathbb{E}_{x^t \sim X^t} \mathbb{E}_{\epsilon \sim p(\epsilon)} \left[\underbrace{-\log q(\bar{x}^t | f(x^t, \epsilon))}_{\text{Recons. term: } Q(x^t)} + \beta \cdot \underbrace{\text{KL}(p(Z^t | x^t, \epsilon) | r(Z^g))}_{\text{KL-div. term: } \text{KL}(Z^t | Z^g)} \right], \quad (4)$$

where x^t is one training sample from input X^t , \bar{x}^t is the reconstruction output, $q(\cdot)$ and $p(\cdot)$ are the posterior probabilities estimated by the decoder and the encoder, respectively, $r(Z^g)$ is the estimated distribution from user-item interaction graph. Here, Eq. (4) has a similar form as β -VAE which contains a reconstruction term $Q(x^t)$ defined by the posterior probability and a KL-divergence term $\text{KL}(Z^t | Z^g)$ regularised by the prior from another modality Z^g .

However, two key challenges need to be addressed when incorporating the geometric prior into a variational network:

(1) How to derive prior $r(Z^g)$ from the user-item interaction graph; (2) How to compute $Q(x^t)$, which is defined as the log conditional probability of $\log q(\bar{x}^t | f(x^t, \epsilon))$ in Eq. (4).

a) Geometric prior obtained from clustering with Gaussian kernel.: We now discuss how to derive the geometric prior from the user-item interaction graph, in which a node can be either a user or an item and the frequently interacted nodes are more likely to share similar preferences or characteristics. We can use a graph neural network to learn user/item node representations with review ratings as the supervision. Afterwards, similar users/items are naturally gathered into one of the K clusters, denoted as C_k , $k \in \{1, 2, \dots, K\}$.⁵

We next show how to define the geometric prior based on the clustering results. Here we use the Gaussian kernel to compute the distance ρ_k^g between a user or item and the centroid of its corresponding cluster C_k as in Eq. (5).

$$\rho_k^g = \frac{\exp(-\|e^g - C_k\|^2 / 2\alpha^2)}{\sum_{k'=1}^K \exp(-\|e^g - C_{k'}\|^2 / 2\alpha^2)}, \quad (5)$$

where e^g is graph node embedding for user or item, C_k is the corresponding k -th cluster centroid vector, α is a hyperparameter used to adjust the distribution according to the data density.

It can be easily shown that $\sum_k \rho_k^g = 1$. That is, the distribution of $\{\sqrt{\rho_k^g}\}$ resides in a hyperspherical cap area on a uniform K -ball, weakly approximating to a Gaussian distribution when K is large [34]. We use $r(Z^g)$ to denote the user's or item's probability distribution over clusters obtained from the user-item interaction graph, i.e., $r(Z^g) = \{r(z_u^g), r(z_i^g)\}$, where $r(z_u^g) = [\rho_{u1}^g, \dots, \rho_{uK}^g]$ and $r(z_i^g) = [\rho_{i1}^g, \dots, \rho_{iK}^g]$.

b) Posterior over conditional probability.: The reconstruction term $Q(x^t)$ in Eq. (4) can be derived by the Nadaraya-Watson estimator [35]:

$$Q(x^t) = \log \frac{\hat{p}(\bar{x}_\epsilon^t, x^t)}{\hat{p}(x^t)} = \log \frac{\frac{1}{N} \sum_{j=1}^N \kappa\left(\frac{x^t - x_j^t}{h}\right) \cdot \kappa\left(\frac{\bar{x}_\epsilon^t - x_j^t}{h}\right)}{\frac{1}{N} \sum_{j=1}^N \kappa\left(\frac{x^t - x_j^t}{h}\right)}, \quad (6)$$

where $\kappa(\cdot)$ is a kernel function for density estimation, \bar{x}_ϵ^t is the reconstructed x^t with random noise ϵ , and x_j^t is the j -th training sample. In practice, we choose the RBF kernel $\kappa = \exp(-\|x - x'\|^2)$ and ignore the denominator since the goal is to optimise the decoding of \bar{x}^t . Then we are able to obtain the following bound by the triangle inequality in the learned metric space as $Q(x^t) \propto \log(\exp(-\|\bar{x}_\epsilon^t - x^t\|^2))$. By minimising the loss, we can preserve the key local similarity of nearby representations while ensuring different representations from different feature characteristics [36], [37].

Therefore, we are able to guarantee that the information bottleneck with our proposed geometric prior is:

- **Transferable** – According to Eq. (2), the mutual information between textual input and the compressed latent variable regularized by geometric prior is still upper bounded by the global information bottleneck, i.e., we can

²In the rest of the paper, we use superscript t for text and g for the graph.

³The proof is shown in the Appendix B.

⁴The proof is shown in the Appendix B.

⁵In the following, We omit the subscript u, i for the user and item as they are processed in the same way.

take the bound from one modality to guide the training on the other.

- **Practicable** – With our solution in Eq. (5), we do not have to utilise a pre-defined prior or oversimplified prior, but instead use the Gaussian-kernel based distance distribution to regularize the latent space. The empirical recommendation results in Table III shows the superiority of the geometric prior.
- **Interpretable** – The alignment between representations from different modalities and their geometric space based probabilities makes it possible to use latent factors inferred from review text, i.e., the semantic clusters embedded in various latent dimensions as shown in Table VII, to provide recommendation explanation.

IV. THE GIANT FRAMEWORK

In this section, we provide a practical implementation of our geometric information bottleneck and use it to achieve item recommendation and explanation generation.

a) *Clustering on Graph Node Embedding.*: We construct a user-item interaction graph where nodes are users or items, and edges are created if the corresponding user's rating on the item is above the item's average rating. We then use *LightGCN* [38] to obtain the user/item embeddings. This choice is motivated by its omission of feature transformation and nonlinear activation, making it a popular choice in recent successful recommenders, as demonstrated by studies such as [39]–[41]. Afterwards, we apply K-Means on the learned graph node embeddings to derive the cluster centroid vectors.

b) *Textual Feature Extractor.*: For a given user-item pair, (u, i) , the input to our recommender is a set of reviews written by u and a set of reviews on i (excluding u 's review on i). As each user or item is associated with multiple reviews, and some reviews might be very long, the Convolutional Neural Network (CNN) is used to encode reviews due to its efficiency in encoding long sequences, as has been studied in [8]. We then aggregate multiple users/items reviews by the attention mechanism to derive the final textual vectors, denoted as x_u^t or x_i^t , respectively. Given the textual vectors, we employ two geometric regularised variational networks (i.e., the encoder-decoder structure) with shared weights to infer the latent factors, z_u^t and z_i^t , obtained from the user and item review text, respectively.⁶

c) *Minimising Distribution Discrepancy.*: To minimise the KL-divergence in Eq. (4), we need to project Z^g and Z^t to a metric space with the same size. Specially, we obtain the latent variables $z^t = f(x^t, \epsilon)$ by applying the reparameterisation trick [33]. However, due to the randomness in initialisation, the learned factors Z^t might not find its correspondence in Z^g . Therefore, in our one-layer MLP based encoder, we propose a prior-centralisation trick which builds the connection between the clusters, obtained from the user-item interaction graph, and the encoder by making the encoder

weights close to the cluster centroids $C \in \mathbb{R}^{d \times k}$ obtained from the graph:

$$\mathcal{R}_{centroid} = \sum_{k=1}^K \left\| W_{en}^k - C_k \right\|^2, \quad (7)$$

where W_{en}^k denotes the k -th column of the encoder weight matrix W_{en} and C_k denotes the representation of the k -th cluster centroid vector. Note that $\mathcal{R}_{centroid}$ will force the distribution of Z^t approximate to the prior and ignore the input. Therefore, in practice, this regularisation is only deployed in the initial training epochs.

Once the connection between the encoder weights and cluster centroids is established, the learning of the text latent variable z^t can be regarded as a soft assignment of the text representation x^t to various clusters. Inspired by [42], a Student t -distribution is used to mitigate the crowding problem. Here, we use τ to adjust the distribution tail and map its values to the range of $[0, 1]$. The k -th dimension of a text latent variable z_k^t denotes the probability of the corresponding user or item being assigned to the k -th cluster:

$$\eta_k^t = \frac{1 + e^{\varphi(z_k^t)/\tau}}{\sum_{k'=1}^K e^{\varphi(z_{k'}^t)/\tau}}, \quad (8)$$

where τ is the temperature in the softmax function, $\varphi(z^t)$ is a linear transformation of z^t . The KL divergence of the posterior distribution of Z^t and the prior distribution of Z^g is then defined by:

$$\text{KL}(p(Z^t|x^t, \epsilon)|r(Z^g)) = \sum_{k=1}^K \eta_k^t \log \frac{\eta_k^t}{\rho_k^g} \quad (9)$$

where η is the derived cluster-based distribution defined in Eq. (8), ρ is the prior distribution calculated from the user-item interaction graph, defined in Eq. (5).

d) *Reconstruction Loss.*: As discussed in Section III, the reconstruction term in Eq. (4) has the property, $Q(x^t) \propto \log(\exp(-\|\bar{x}_\epsilon^t - x^t\|^2))$. Therefore, it can be defined as the mean squared error function⁷:

$$Q(x^t) = \mathbb{E}_{x^t \sim X^t} \mathbb{E}_{\epsilon \sim p(\epsilon)} [\|\bar{x}_\epsilon^t - x^t\|^2], \quad (10)$$

where x^t denotes the input user or the item representation, and \bar{x}_ϵ^t denote the reconstructed representation.

e) *Rating Prediction and Final Objective Function.*: After obtaining the regularised latent textual representations for user z_u^t and item z_i^t , we add their corresponding ID features to obtain the final user and item embeddings ζ_u and ζ_i for rating prediction, $\zeta_u = W_u z_u^t + \epsilon_u$, $\zeta_i = W_i z_i^t + \epsilon_i$, where the ID features ϵ_u and ϵ_i are generated by feeding the user ID and item ID to an embedding layer. Inspired by the latent factor model in recommendation systems, we introduce the global biases for users and items in the final prediction layer as $\hat{r}_{ui} = f_{\text{cls}}(\zeta_u, \zeta_i) + b_u + b_i$, where f_{cls} combines the user and item features into a scalar, b_u and b_i are bias derived from ϵ_u and ϵ_i . The regression loss \mathcal{L}_r of the predicted rating is calculated as the MSE on the given user-item pair. Combining all the components above, we derive the training objective as follows:

$$\mathcal{L} = \mathcal{L}_r + O_{IB}^t + \mathcal{R}_{centroid} \quad (11)$$

⁶In what follows, we drop the subscript u and i for simplicity.

⁷The detailed discussion can be found in Appendix C.

f) *Time Complexity and Scalability Analysis*: In this subsection, we conduct a complexity analysis of both **GIANT** and LightGCN, which is our chosen base graph model. Suppose the number of user nodes is m , the number of item nodes is n , the data sparsity is defined as α , and edges in the user-item interaction graph are $E = \alpha \cdot m \cdot n$. The average number of reviews for users and items are $|D_u|$ and $|D_i|$, respectively. Let G denotes the number of GCN layers. d is the graph node embedding size, as well as the dimension of cluster centroid vector \mathcal{C} . K_u is the number of user clusters, K_i is the number of item clusters, T is the maximum number of iterations for K-Means⁸. For simplicity, all the kernels in CNN are assumed having a uniform size of c , and there are a total of f kernels. The padding length for each review text is denoted as L .

We use the user and item embeddings obtained from well-trained LightGCN. Its complexity analysis results were reported in [43]. The additional computations involve several steps, including K-means for user/item clustering, final rating prediction using a VAE framework with rate prediction loss \mathcal{L}_r , information bottleneck regularisation O_{IB}^t , and the centroid regularisation R_{centroid} . For the textual feature extraction, all the reviews associated with a single user or item are aggregated into one, which is then fed into the CNN. This leads to an expected input channel size of $m|D_u|$ for users and $n|D_i|$ for items. Consequently, the expected total computational complexity can be expressed as $\mathcal{O}(m|D_u|Lc^2f) + \mathcal{O}(n|D_i|Lc^2f)$. For rate prediction, we only calculate the dot-product for all user-item interactions, resulting in a complexity of $\mathcal{O}(E)$. For the calculation of the KL divergence in Eq. (9), the complexity is determined by $(K_u + K_i)$ for each user-item pair. Since the dimensions of the user and item latent variables are identical to K_u and K_i , respectively, the reconstruction complexity for each (u, i) pair is also $(K_u + K_i)$. Since the number of

and $86,400 \cdot n$. This allows us to simplify the computational complexity as $\mathcal{O}(|D_u|Lc^2fm + |D_u|Lc^2fn + (dK_u + dK_i)E)$. Taking into account the sparsity of our datasets, as shown in Table II, we can approximate E as $(3.5e - 3)mn$, the overall complexity is primarily determined by $|D_u|Lc^2fm + |D_i|Lc^2fn + 14mn$. *This complexity scales linearly with the number of reviews for each user or item.*

When comparing our cost to that of LightGCN, which is expressed as $2|E|^2G$, where G represents the number of GCN layers, we find that our time complexity becomes comparable to LightGCN only when the result of $|D_u|Lc^2fm + 7mn$ approaches $(\alpha mn)^2$. In practical terms, in our experimental configurations, these two results are within the same level. Therefore, our time complexity is nearly equivalent to that of existing graph-based recommender systems.

V. SENTENCE-LEVEL EXPLANATION GENERATION

Our variational framework can identify the top- k words for each latent dimension. To make it easier for humans to understand the rationales behind model decisions, we extract the most relevant *review sentences* (see the extracted explanations in Figure 1 and examples in Figure 3) from all the candidate reviews as interpretations for a specific user-item pair $(u-i)$.

In particular, our **GIANT** model infers the latent factors z^t . The most prominent dimension in the latent factor represents the assigned cluster for user u or item i . Recall that the number of clusters derived in the graph space equals to the number of latent dimensions. Based on that, we can identify the user candidate reviews as all the past reviews on item i written by the users assigned to the same cluster as u . Similarly, we identify item candidate reviews from the past reviews on the item i which have their most prominent topic the same as the item latent topic. Afterwards, we represent each topic by its top-associated 5 words and derive the topic representation by the aggregated word-level GloVe⁹ word embeddings.¹⁰ These words are selected based on the TFIDF on the reviews in the same topic cluster, with stop words filtered. The review sentence representations are also derived based on the aggregated constituent word GloVe embeddings. The most relevant sentences from user candidate reviews can then be extracted as a summary of the user u 's preferences based on their cosine similarity with the user latent topic representation. Similarly, the most relevant sentences from item candidate reviews are extracted as a summary of item i 's characteristics based on their cosine similarity with the item's latent topic representation.

VI. EXPERIMENTS RESULTS ON RECOMMENDATION

A. Experimental Setup

a) *Dataset and Data Processing*: We choose three commonly used e-commerce datasets, including *BeerAdvocate* [44] and two amazon review datasets, *Digital Music* and *Office Products* [45], which not only contain the interaction interactions between user and item but also the review texts.

⁹<https://nlp.stanford.edu/projects/glove/>

¹⁰We also experimented with more representative words but observed less discriminative topic clusters.

LightGCN	Adjacency GCN	$\mathcal{O}(E)$
	BRP Loss	$\mathcal{O}(E ^2Gd)$
		$\mathcal{O}(E d)$
K-Means	clustering	$\mathcal{O}(K_u m T) + \mathcal{O}(K_i n T)$
CNN	Feature extractor	$\mathcal{O}(m D_u Lc^2f) + \mathcal{O}(n D_i Lc^2f)$
VAE	Rate Prediction	$\mathcal{O}(E)$
	O_{IB}	$\mathcal{O}(EK_u + EK_i)$
	R_{centroid}	$\mathcal{O}(dEK_u + dEK_i)$

TABLE I

THE ANALYTICAL TIME COMPLEXITY OF **GIANT**.

clusters is much smaller than the number of items or users, the expected increase in complexity primarily arises from the K-Means clustering, textual feature extraction, and rate prediction steps, roughly speaking:

$$\begin{aligned} &\mathcal{O}((K_u T + |D_u|Lc^2f)m + (K_i T + |D_i|Lc^2f)n) + \\ &\mathcal{O}((dK_u + K_u + dK_i + K_i + 1)E) \end{aligned} \quad (12)$$

In our highest configuration, we set d to 64, K_u and K_i to 64, $|D_u|$ and $|D_i|$ to 100, L to 32, kernel size to 3, and the number of kernels to 3. Consequently, the dominant terms are the second coefficient for m and n , approximately $86,400 \cdot m$

⁸Noted that we use the K-Means package within Scikit-Learn for performing clustering, and the average computational complexity is recorded in the table.

Amazon review data¹¹ is one of the popular datasets, consisting of 24 kinds of products. As our method derive the user/item mainly based on their reviews, we use the dense subset, 5-core, of the amazon review dataset that extracts the user-item pairs, such that each of the remaining users and items have at least 5 reviews each. Table II summarizes the statistics of the three datasets. In order to enhance the performance of our recommender system, we employ a filtering mechanism for entries (user/item) associated with either reviews of significant length or an exceptionally large number of reviews. This is done because both the length and the number of reviews follow a long-tail distribution, as described in [8]. Specifically, we filter out users or items with over 100 reviews/interactions and truncate the reviews to lengths of [200,400,300] for the three datasets, respectively. Additionally, we randomly select 80%, 10%, 10% as the train/validation/test sets.

b) *Evaluation Metrics*: We use Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) to evaluate the rating prediction accuracy. We also report the ranking-based evaluation results, i.e., Precision and Recall to measure the overlapping between the recommended items and the target items.

Datasets	#Train	#users	#items	AveRate	Sparsity
<i>BeerAdvocate</i>	35450	6939	13122	3.74	4.4e-4
<i>Digital Music</i>	51768	5541	3568	4.22	2.9e-3
<i>Office Products</i>	35999	4902	2364	4.33	3.5e-3

TABLE II

ALL THE DATASETS HAVE RATINGS IN THE RANGE OF 1-5. THE BEER DATASET HAS THE SMALLEST SPARSITY.

c) *Baselines*: We compare with several open-source recommenders, including *HFT* [17], *DeepCoNN* [46], and *NARRE* [8]. *HFT* combines reviews with ratings and uses an exponential transformation function to link review text and the ratings. *DeepCoNN* uses a shared layer for interaction modeling the users and items, which is on top of the two encoders for the users and items, respectively. *NARRE* is an interpretable recommender, which introduces review-level attentions to select important reviews and incorporates the user and item IDs as discriminative features for rating prediction and uses the text spans with higher attention as recommender rationales. Our **GIANT** is also built on the *NARRE*, but with a proposed geometric prior in an encoder-decoder structure. More recently, a graph model called *RGCL* [47] was introduced, which incorporates a contrastive loss and takes into account both user-item interactions and review semantics in its node features. Our model similarly emphasises the significance of user-item interactions in review-based recommender systems, but it introduces a novel regularisation technique based on clustering. Furthermore, we compare with several other encoder-decoder structures with different regularizes on top of *NARRE* to learn the latent variable as a baseline to highlight the difference with our proposed regularisation, i.e., *AutoEncoder*, *StandPrior* and *WassersteinVAE* and *IndivPrior*:

- *StandPrior* (aka. *StandVAE*) uses two VAEs for users and items respectively, each with the standard Gaussian distribution $\mathcal{N}(0, 1)$ as the prior.
- *WassersteinVAE* differs from the *StandPrior* in using the Wasserstein metric [48] to calculate the distribution discrepancy between posterior η and prior ρ .
- *IndivPrior* assigns a separate Gaussian prior to each user or item with its mean value calculated from the user's or item's corresponding representation derived from the user-item interaction graph, i.e., $p(z_u^t) = \mathcal{N}(W_u \xi_u, \mathbf{I})$, $p(z_i^t) = \mathcal{N}(W_i \xi_i, \mathbf{I})$, where $\xi_u \in \mathbb{R}^d$, $\xi_i \in \mathbb{R}^d$ denote the user and item graph node, W_u and W_i are learnable parameters in a linear layer of size $(d \times d)$, and d is the dimension of graph node embeddings.

d) *Training Procedure*: We train the model by minimizing the objective function defined in Eq. (11). The β for the KL divergence term is set to 0.01 and the L_2 regularisation is used for all the model parameters with the weight of 0.001. We only train the prior-centralisation term for 0.5 proportion to approximate the cluster centroids for the later training. The KL term in O_{IB} should be introduced later until the encoder centroid is well trained. To do so, we follow a similar cyclical schedule [49] to gradually adjust the anneal factor λ in each epoch. Specifically, we first train the model without the KL term for 0.5 proportion, then anneal it from 0.5 to 1 for 0.25 proportion, and finally fix $\lambda = 1$.¹²

B. Rating Prediction Results

The rating prediction results of our model in comparison with baselines on all datasets are given in Table III. We have the following observations. (1) *DeepCoNN* built on the stack of non-linear neural networks for review semantic modeling outperforms *HFT* which leverages an exponential transformation function to link topic distributions in review text and latent factors derived from ratings. This shows the superiority of deep learning for feature extraction. (2) *AutoEncoder* gives slightly better performance than *NARRE* in *BeerAdvocate*, which shows its effectiveness of extracting key contextual information for rating prediction. (3) *RGCL* emerges as the leading baseline model and even manages to achieve competitive results on both the *Beer* and *Music* datasets when compared to our *GIANT*. However, when considering the *Office* dataset, its overall performance falls short of ours. The average RMSE/MAE values across the three datasets are 84.49/63.49 for *RGCL* and 84.18/62.87 for *GIANT*, indicating that our model outperforms *RGCL*. The results demonstrate the effectiveness and robustness of our proposed information bottleneck regularisation applied on the latent semantic space. By doing this, the users and items can be grouped into different clusters according to the interaction data, which is not the case in other baseline recommenders.

a) *Effect of our proposed information bottleneck regularisation*: We then study the effects of different information bottleneck, i.e., *StandPrior*, *WassersteinVAE* and *IndivPrior*. From the results in the last group of Table III. It can be seen

¹¹<https://jmcauley.ucsd.edu/data/amazon/>

¹²Detailed hyperparameter setting is presented in Appendix A.

that *WassersteinVAE* and *StandPrior* demonstrate overall better results than *IndivPrior* and all the baselines in Table III. The improvement can be explained by the nature of VAE, which is particularly beneficial when dealing with sparse data where few observations are available. However, *IndivPrior*, which imposes a separate Gaussian prior on individual user/item, shows worse performance compared to the other three variational frameworks, and is even worse than *AutoEncoder*. Our model benefits from a rich prior derived from user/item clusters, while avoiding using a global normal prior (*StandPrior* and *WassersteinVAE*) or a separate instance-level prior for each user/item (*IndivPrior*).

Models	<i>BeerAdvocate</i>		<i>Digital Music</i>		<i>Office Products</i>	
	RMSE(↓)	MAE(↓)	RMSE(↓)	MAE(↓)	RMSE(↓)	MAE(↓)
<i>HFT</i>	79.81	62.42	96.42	74.76	89.46	68.57
<i>DeepCoNN</i>	77.28	59.45	94.69	71.03	85.14	64.82
<i>NARRE</i>	76.80	58.94	<u>93.69</u>	69.30	<u>84.40</u>	63.40
<i>RGCL</i>	74.65	57.12	92.02	68.49	86.81	64.87
<i>AutoEncoder</i>	75.94	<u>58.93</u>	93.97	<u>69.13</u>	85.03	<u>64.32</u>
<i>StandPrior</i>	75.80	58.73	<u>93.60</u>	<u>69.21</u>	84.69	63.97
<i>WassVAE</i>	<u>75.62</u>	<u>58.06</u>	93.81	69.32	<u>84.20</u>	<u>63.43</u>
<i>IndivPrior</i>	76.43	59.88	94.57	70.51	85.33	65.05
<i>GIANT</i>	<u>75.36*</u>	<u>57.87*</u>	<u>92.87*</u>	<u>68.68*</u>	84.32	62.05*

TABLE III

PERFORMANCE COMPARISON IN RMSE (%) AND MAE (%) FOR ALL METHODS. * DENOTES THE STATISTICAL SIGNIFICANCE FOR $p < 0.01$, COMPARED TO THE BEST BASELINE. THE BEST RESULTS THROUGHOUT ALL THE ROWS ARE IN BOLD, THE BEST RESULTS IN EACH GROUP ARE UNDERLINED.

b) *Ranking-based Results*: In addition to the accuracy-based metrics, i.e., MSE and MAE, we also evaluate models using the ranking-based metrics, i.e., HitRatio@10 and NDCG@10 for the two best baselines (Table IV). Without evaluation on the whole item list, we followed the common practice in [50], [51] that randomly samples negative 100 items that are not interacted by the user, ranking the test item among the 100 items. Then the ranked list is truncated at 10 for both HitRatio@10 and NDCG@10. The performance trends across different datasets diverge from those observed in accuracy-based metrics. *RGCL* demonstrates comparable performance on the Office dataset, whereas its performance on the *BeerAdvocate* dataset is much worse, achieving a HitRatio of 0.47 compared to 0.62 for *GIANT*. However, our *GIANT* still maintains an overall superior performance in rank-based metrics. Specifically, on the HitRatio@10 metric, our model achieves a score of 0.65, while *RGCL* achieves a score of 0.61. *WassersteinVAE* appears to be the second-best encoder-decoder model (the second group). The relative rank-based performances across the two groups in Table III are different, while the relative performances within encoder-decoder structured models are similar to accuracy-based metrics, i.e., MSE and *GIANT* in general achieves the best results on the two kinds of metrics.

C. Contribution from Various Loss Terms

We study the effects of the three loss terms in rating prediction accuracy and *Diversity* of the latent variables. As variational networks could easily collapse into an unconditional

Models	<i>BeerAdvocate</i>		<i>Digital Music</i>		<i>Office Products</i>	
	HitRatio	NDCG	HitRatio	NDCG	HitRatio	NDCG
<i>PureMF</i>	0.43	0.23	0.78	0.52	0.61	0.33
<i>SVD++</i>	0.41	0.21	0.52	0.28	0.55	0.28
<i>NeuMF</i>	0.41	0.21	0.78	0.52	0.63	0.35
<i>NARRE</i>	0.61	0.52	0.65	0.59	0.59	0.64
<i>RGCL</i>	0.47	0.55	0.72	0.63	0.64	0.72
<i>AutoEnc</i>	0.52	0.49	0.62	0.60	0.61	0.62
<i>StandPrior</i>	0.55	0.54	0.68	0.63	0.61	0.62
<i>WassVAE</i>	0.55	0.54	0.69	0.65	0.61	0.60
<i>GIANT</i>	0.62	0.57	0.69	0.67	0.64	0.64

TABLE IV

THE EVALUATION RESULTS BASED ON RANKING, NAMELY, HITRATIO@10 AND NDCG@10 FOR ALL THE BASELINE AND OUR PROPOSED *GIANT*.

generative model, i.e., in the extreme case, all the input will be mapped into the same latent code [52]. We use the dimension index whose corresponding latent value is the maximum as the cluster assignment ID. We then derive the cluster assignment results $\mathcal{A} \in \mathbb{R}^{N \times K}$, where N is the number of test samples, K is the cluster number. The diversity is calculated based on entropy $H(X) = -\sum_{k \in K} p(k) \log(p(k))$, where $p(k)$ is the fraction of the number of samples falling into the k -th cluster among all the samples. A larger value means better diversity. We record the largest diversity between user and item latent variables. The results are shown in Table V.

Variants	<i>BeerAdvocate</i>			<i>Digital Music</i>			<i>Office Products</i>		
	NDCG	RMSE	Div	NDCG	RMSE	Div	NDCG	RMSE	Div
Full Model	0.57	75.36	1.89	0.67	92.87	3.46	0.64	84.32	1.58
-w/o $\mathcal{R}_{centroid}$	0.55	75.96	0.04	0.65	92.62	0.07	0.62	84.58	0.02
-w/o KL term	0.55	75.64	1.43	0.65	93.83	2.86	0.61	84.23	1.21
-w/o $Q(x^t)$	0.52	76.02	1.04	0.61	93.95	3.10	0.59	84.94	0.93
$Q(x^t)$ Cosine	0.56	75.35	0.11	0.67	92.85	1.76	0.63	84.97	0.07

TABLE V

RATING PREDICTION RESULTS, INCLUDING METRICS SUCH AS NDCG@10, RMSE(%) AND *Diversity* ACROSS MODEL VARIANTS. A LOWER DIVERSITY INDICATES THAT THE LEARNED LATENT VARIABLES FOR DIFFERENT SAMPLES TEND TO COLLAPSE INTO A SIMILAR LATENT CODE.

We observe that the removal of $Q(x^t)$ leads to the largest performance degradation. After removing $\mathcal{R}_{centroid}$, the latent variable diversity shrinks to near zero, indicating nearly all the latent variables fall into the same cluster. This highlights the capability of our prior-centralisation term in enabling the latent variables z^t effectively reflect the soft cluster assignments. We also found that replacing the MSE in $Q(x^t)$ with the cosine similarity $(1 - cosine)$ reduces the latent variable diversity significantly. This can be partly explained by the fact that the cosine similarity focuses on measuring the angle between the input x^t and the reconstructed output \hat{x}^t , ignoring the magnitude of the vectors which is however important in our case.

VII. LATENT VARIABLE AND INTERPRETABILITY EVALUATION

In addition to the rating prediction, we design experiments to evaluate the generated latent variables and explore the improvement source. Firstly, we evaluate the generated latent

clusters in terms of semantic coherence and separability. Then, we examine the effects of the identified important latent dimensions on rating prediction. Finally, we conduct a human evaluation to study the interpretability.

A. Cluster Coherence and Separability

a) *Semantic Coherence*: To verify the capability of creating *semantically coherent clusters*, we propose to measure cluster coherence as the average cosine similarity between every review document pair within a cluster. We first obtain the document-level review representations by feeding the reviews from the test set to our pre-trained CNN encoder. We then calculate the average cosine similarity between the representations of each review pair in a cluster, and are further averaged across all clusters (Table VI). We use the clusters generated via applying KMeans on learnt graph node embedding as a comparison, i.e., ρ in Eq. (5).

Models	BeerAdvocate		Digital Music		Office Products	
	User	Item	User	Item	User	Item
Graph	0.092	0.085	0.155	0.214	0.245	0.200
GIANT	0.490	0.481	0.457	0.471	0.492	0.407

TABLE VI

A LARGER SIMILARITY VALUE MEANS A BETTER COHERENCE. **GIANT** GENERATES SIGNIFICANTLY MORE COHERENT CLUSTERS THAN GRAPH CLUSTERS.

b) *Cluster Separability*: As each dimension of the latent variables z^t corresponds to a cluster, for the k -th dimension, we can search for its relevant reviews which have the highest value in z_k^t and list the most frequent words in the review set as the representative topic words¹³ Representative words in randomly selected three clusters on the three datasets are shown in Table VII. For our model, the most prominent words are different across different clusters (thus being coloured), i.e., ‘brown roasted’, ‘pine, caramel’ and ‘good flavour’. We identify ‘ABBA’ and their popular songs, computer peripherals such as ‘network cable’ and ‘router’ in *Digital Music* and *Office Products* datasets, respectively.

We find it hard to see clear topic separations from the *WassersteinVAE* results. For example, Cluster 2 and 3 share the same 3 words, ‘sweet’, ‘light’ and ‘malt’ out of their top 4 words. In addition, ‘carbonation’ in Cluster 2 and ‘whitehead’ in Cluster 3 are both related to beer foam. The results show that with the incorporation of the priors derived from user/item clusters, our proposed approach is able to learn latent variables in the review semantic space which can produce better separable topic clusters.

B. Comprehensiveness Evaluation by Perturbing on Latent Variables

To explore the importance of our identified clusters, i.e., the latent dimension with a larger value, we calculate the performance change before and after removing the specific dimension and define *Comprehensiveness* as: $\frac{\sum_i^N (r(z^i) - r(z_{\setminus k}^i))^2}{N}$,

¹³We exclude the stopwords and most frequent words appeared in all clusters.

Where $r(\cdot)$ is the predicted rating, z^i is the latent variable for i -th evaluated user-item pair, $z_{\setminus k}^i$ is to remove the identified top k latent dimensions from z^i . To this end, we replace the values in the top- k dimensions with the average value of the latent variables according to [53]. In the Figure 2 line chart, our model demonstrates a larger performance change as more important dimensions are removed. We further randomly delete k latent dimensions and calculate the relative performance change by subtracting the changes caused by random removal (Table below Figure 2). The relative changes in our model are most predominant, followed by *WassersteinVAE*, while relative changes of *StandPrior* are negative, showing that the random removal of latent dimensions can bring even larger changes to the predominant ones.

C. Human Evaluation of Generated Interpretations

We conduct a human evaluation to validate the interpretability of our proposed method. To make it easier for humans to understand the rationales behind model decisions, we extract the most relevant sentences from user/item reviews as interpretations for a specific user-item pair (user u and item i). Examples of generated interpretations are illustrated in Figure 3, as well as the model’s recommendation suggestion as *will recommend* if the model’s predicted rating score is above the average predicted rating; or *won’t recommend*, otherwise.

We propose three evaluation metrics, *relevance*, *faithfulness* and *informativeness* and ask three English-proficient human evaluators to give 1-5 score to the generated interpretations from *WassersteinVAE*, *NARRE* and *GIANT* on 120 randomly selected user-item pairs from each of the three datasets. The model identifiers are unknown to the evaluators.

- *Relevance*. Do the extracted sentences from user reviews and item reviews pertain to the relevant topic, aspect, or subject? An interpretation with greater overlap in aspects should receive a higher score.
- *Faithfulness*. Do the sentences extracted from user reviews and item reviews influence the recommendation suggestions made by the model? You should consider both the sentiment and the aspects described. For example, if both the selected reviews from the user and the item express positive sentiments regarding the “prices” of the product, the faithfulness score should be rated as 5. Any discrepancies between the sentiment and aspects mentioned can lower this score.
- *Informativeness*. Do the provided interpretations effectively reflect the user preferences and item characteristics? This criterion should remain separate from the previous ones, only focusing on the extent to which various aspects and their associated sentiments are discussed in the extracted reviews. A lower level of coverage of aspects results in a reduced *Informativeness* score, as exemplified by a statement like, “I really like this product!”.

Our human evaluation results presented in Table VIII shows that both *WassersteinVAE* and **GIANT** are better in *relevance*, which can be partly explained that they can extract sentences

	GIANT	WassersteinVAE
Beer	pour , color, brown , feel, roasted , malts, almost, moderate, coffee, dark good , flavour , dark, better , brew, style, bad, Canadian, nose, drinking pine , caramel , lacing, citrus , mouthfeel, hint, ipa, note, body, strong	light, color, malt, glass , poured, drink , sour, pour, mouthfeel, alcohol aroma, light, sweet, malt, white head , mouthfeel, nose , hint, brew , note sweet, coffee , light, malt, aroma, carbonation , pour, mouthfeel, hint, color
Music	dr dre , prince , used known , west coast, mirrors, love, hate, westside story, jadekiss , always	quot, album, Ask Rufus , song, funk , voice , music, band, soul, love
Office	bangles , vicki , vocal , place, live, hero, takes, fall, liverpool, beatles	song, everglow , album, band, beautiful , track, lyrics , suspension, record, Christian,
	dancing queen , take chance , mamma mia , abba, greatest hit, money money, gimme gimme, super trouper, knowing knowing, gold	record, album, song, band, track, time, sound, head heart , folk , live , love
	set , computer , wireless , network, cable, router, edit, pencil, ink	scan , software , pencil, Mac OS , work, printer, small, printed, color, file
	photo , desk , laser , scanner, work, scan, epon, scanning, enough, easy	easy use, dry erase , folder, mouse pad , office, ink cartridge , works, used, post note
	pencil, tape , ink , pen, anyway, pen, rubber , clip, pretty, eraser	recommend , printer, folder, paper , easy use, boxes , feature, stapler, pencil, canon, color

TABLE VII

THE MOST PROMINENT WORDS (SORTED BY OCCURRENCE FREQUENCY) IN THREE RANDOMLY SELECTED CLUSTERS FROM **GIANT** AND **WassersteinVAE** ON *Digital Music* AND *Office Products* DATASETS. FOR EACH DATASET, WE HIGHLIGHT THE TOP 3 WORDS THAT ARE NOT FOUND IN THE OTHER TWO CLUSTERS. **GIANT** GENERATES BETTER SEPARABLE TOPICS WHILE **WassersteinVAE** FAIL TO GENERATE CLEAR TOPIC PATTERN.

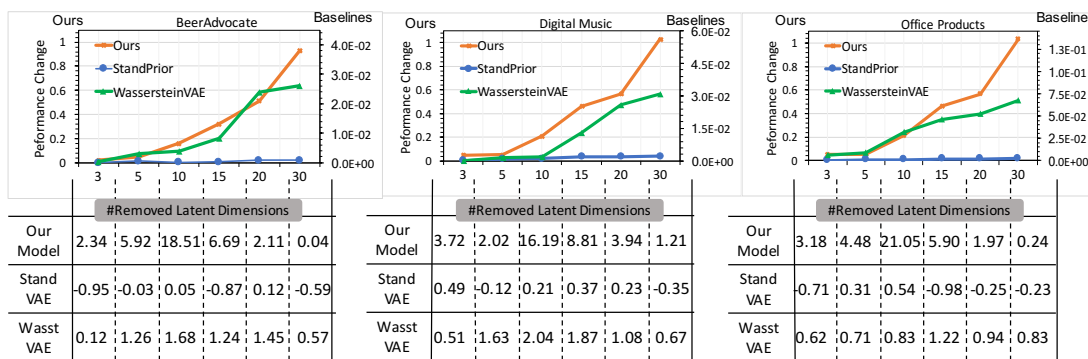


Fig. 2. **Top**: The *Comprehensiveness* values by removing the top k most important latent dimensions identified, $k \in \{3, 5, 10, 15, 20, 30\}$. Ours model uses the left y-axis, the two baselines use the right axis. **Bottom Table**: Relative performance changes after subtracting the changes caused by random removal.

The interpretations between 46 U-I pair. (GIANT)

Considering your preferences and those of users who share your tastes:

- It seems *sturdy enough* for classroom use.
- It is much thinner than duct tape and seems *stronger*.

Furthermore, based on the reviews of this product:

- You *can't just accidentally tear through this tape* you need to *use a sharp blade on purpose* to break the seal.
- I did use some on my grandson's toy lawn mower and *it is holding the broken handle on really well*.

We will recommend this item.

The interpretations between 46 U-I pair. (WassersteinVAE)

Considering your preferences and those of users who share your tastes:

- I don't think you could go wrong* with this product.
- This tape isn't transparent, but it is *strong*.

Furthermore, based on the reviews of this product:

- I used the tape to secure some boxes i m using to store items. *I won't be using for a while*.
- Fiberglass threads, *impossible to tear by hand*, and *strong adhesive* that will never fall off.

We will recommend this item.

Fig. 3. Extracted interpretations from **GIANT** (upper) and **WassersteinVAE** (bottom) for the item, *a tape*. In the example generated by **GIANT**, both the reviews/rationales from the user and item side focus on the 'sturdy and strong', and reveal lots of important aspects for the recommendation. While the user reviews extracted by **WassersteinVAE**, e.g., 'go wrong' fail to capture the key factors, and the extracted item reviews are even contradictory, i.e., 'won't be using' to the recommended behaviour.

from reviews of like-minded users from the same user clustering results when the current reviews don't contain any key information. The superior performance over **WassersteinVAE** further shows the efficiency of our proposed geometric prior.

	Metric	Relevance	Faithfulness	Informativeness
Beer	NARRE	2.7	2.4	3.1
	WasserVAE	3.6	3.5	3.5
	GIANT	3.5	3.7	3.5
Music	NARRE	4.4	3.2	3.8
	WasserVAE	4.0	3.3	3.0
	GIANT	4.7	3.8	4.1
Office	NARRE	2.7	2.8	3.9
	WasserVAE	3.5	2.8	3.8
	GIANT	3.8	3.3	4.2

TABLE VIII

HUMAN EVALUATION RESULTS ON *relevance*, *faithfulness* AND *informativeness* OF THE GENERATED INTERPRETATIONS, WHICH WERE ASSESSED FOR 120 RANDOMLY SELECTED USER-ITEM PAIRS FROM THE THREE DATASETS. THE COHEN'S KAPPA CO-EFFICIENT AMONG EACH PAIR OF ANNOTATORS FOR THE THREE METRICS ARE 0.42, 0.38, AND 0.47, RESPECTIVELY.

VIII. CONCLUSION

In this paper, we leverage the user/item clusters sharing common interests/characteristics obtained from the user-item interaction graph to refine the review text latent factors via our proposed geometric information bottleneck (**GIANT**). We empirically show that **GIANT** is better in learning a semantically coherent and interpretable latent space and the generated explanations are more faithful to the model decisions, while

achieving comparable rating prediction accuracy on three commonly used datasets.

IX. ACKNOWLEDGEMENT

This work was funded by the UK Engineering and Physical Sciences Research Council (grant no. EP/T017112/1, EP/V048597/1, EP/X019063/1). HY receives the PhD scholarship funded jointly by the University of Warwick and the Chinese Scholarship Council. YH is supported by a Turing AI Fellowship funded by the UK Research and Innovation (grant no. EP/V020579/1).

REFERENCES

- [1] A. Ghazimatin, O. Balalau, R. Saha Roy, and G. Weikum, "Prince: Provider-side interpretability with counterfactual explanations in recommender systems," in *Proceedings of the 13th International Conference on Web Search and Data Mining*, 2020, pp. 196–204.
- [2] Y. Zhang, X. Chen *et al.*, "Explainable recommendation: A survey and new perspectives," *Foundations and Trends® in Information Retrieval*, vol. 14, no. 1, pp. 1–101, 2020.
- [3] Y. Zhang, G. Lai, M. Zhang, Y. Zhang, Y. Liu, and S. Ma, "Explicit factor models for explainable recommendation based on phrase-level sentiment analysis," in *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast, QLD, Australia - July 06 - 11, 2014*, S. Geva, A. Trotman, P. Bruza, C. L. A. Clarke, and K. Järvelin, Eds. ACM, 2014, pp. 83–92. [Online]. Available: <https://doi.org/10.1145/2600428.2609579>
- [4] J. Ni, J. Li, and J. J. McAuley, "Justifying recommendations using distantly-labeled reviews and fine-grained aspects," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Association for Computational Linguistics, 2019, pp. 188–197. [Online]. Available: <https://doi.org/10.18653/v1/D19-1018>
- [5] Z. Chen, X. Wang, X. Xie, T. Wu, G. Bu, Y. Wang, and E. Chen, "Co-attentive multi-task learning for explainable recommendation," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, S. Kraus, Ed. ijcai.org, 2019, pp. 2137–2143. [Online]. Available: <https://doi.org/10.24963/ijcai.2019/296>
- [6] L. Li, Y. Zhang, and L. Chen, "Generate neural template explanations for recommendation," ser. CIKM '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 755–764. [Online]. Available: <https://doi.org/10.1145/3340531.3411992>
- [7] J. Tan, S. Xu, Y. Ge, Y. Li, X. Chen, and Y. Zhang, *Counterfactual Explainable Recommendation*. New York, NY, USA: Association for Computing Machinery, 2021, p. 1784–1793. [Online]. Available: <https://doi.org/10.1145/3459637.3482420>
- [8] C. Chen, M. Zhang, Y. Liu, and S. Ma, "Neural attentional rating regression with review-level explanations," in *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France, April 23-27, 2018*, P. Champin, F. Gandon, M. Lalmas, and P. G. Ipeirotis, Eds. ACM, 2018, pp. 1583–1592. [Online]. Available: <https://doi.org/10.1145/3178876.3186070>
- [9] P. Wang, R. Cai, and H. Wang, "Graph-based extractive explainer for recommendations," in *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, F. Laforest, R. Troncy, E. Simperl, D. Agarwal, A. Gionis, I. Herman, and L. Médini, Eds. ACM, 2022, pp. 2163–2171. [Online]. Available: <https://doi.org/10.1145/3485447.3512168>
- [10] C. Wu, F. Wu, J. Liu, and Y. Huang, "Hierarchical user and item representation with three-tier attention for recommendation," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 1818–1826. [Online]. Available: <https://doi.org/10.18653/v1/n19-1180>
- [11] Y. Tay, A. T. Luu, and S. C. Hui, "Multi-pointer co-attention networks for recommendation," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Y. Guo and F. Farooq, Eds. ACM, 2018, pp. 2309–2318. [Online]. Available: <https://doi.org/10.1145/3219819.3220086>
- [12] C. Li, C. Quan, L. Peng, Y. Qi, Y. Deng, and L. Wu, "A capsule network for recommendation and explaining what you like and dislike," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, B. Piwowarski, M. Chevalier, É. Gaussier, Y. Maarek, J. Nie, and F. Scholer, Eds. ACM, 2019, pp. 275–284. [Online]. Available: <https://doi.org/10.1145/3331184.3331216>
- [13] X. Chen, Y. Zhang, and Z. Qin, "Dynamic explainable recommendation based on neural attentive models," 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, 2019, pp. 53–60. [Online]. Available: <https://doi.org/10.1609/aaai.v33i01.330153>
- [14] H. Liu, W. Wang, Q. Peng, N. Wu, F. Wu, and P. Jiao, "Toward comprehensive user and item representations via three-tier attention network," *ACM Trans. Inf. Syst.*, vol. 39, no. 3, pp. 25:1–25:22, 2021. [Online]. Available: <https://doi.org/10.1145/3446341>
- [15] T. Zhang, C. Sun, Z. Cheng, and X. Dong, "Aenar: An aspect-aware explainable neural attentional recommender model for rating prediction," *Expert Systems with Applications*, vol. 198, p. 116717, 2022.
- [16] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2014. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [17] J. McAuley and J. Leskovec, "Hidden factors and hidden topics: Understanding rating dimensions with review text," in *Proceedings of the 7th ACM Conference on Recommender Systems*, ser. RecSys '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 165–172. [Online]. Available: <https://doi.org/10.1145/2507157.2507163>
- [18] C. Wang and D. M. Blei, "Collaborative topic modeling for recommending scientific articles," in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Diego, CA, USA, August 21-24, 2011*, C. Apté, J. Ghosh, and P. Smyth, Eds. ACM, 2011, pp. 448–456. [Online]. Available: <https://doi.org/10.1145/2020408.2020480>
- [19] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, 2003. [Online]. Available: <http://jmlr.org/papers/v3/blei03a.html>
- [20] H. Fei, Y. Ren, S. Wu, B. Li, and D. Ji, "Latent target-opinion as prior for document-level sentiment classification: A variational approach from fine-grained perspective," in *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, J. Leskovec, M. Grobeldnik, M. Najork, J. Tang, and L. Zia, Eds. ACM / IW3C2, 2021, pp. 553–564. [Online]. Available: <https://doi.org/10.1145/3442381.3449789>
- [21] Q. Truong, A. Salah, and H. W. Lauw, "Bilateral variational autoencoder for collaborative filtering," in *WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021*, L. Lewin-Eytan, D. Carmel, E. Yom-Tov, E. Agichtein, and E. Gabrilovich, Eds. ACM, 2021, pp. 292–300. [Online]. Available: <https://doi.org/10.1145/3437963.3441759>
- [22] X. Wang, H. Jin, A. Zhang, X. He, T. Xu, and T. Chua, "Disentangled graph collaborative filtering," 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*, J. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, and Y. Liu, Eds. ACM, 2020, pp. 1001–1010. [Online]. Available: <https://doi.org/10.1145/3397271.3401137>
- [23] J. Ma, P. Cui, K. Kuang, X. Wang, and W. Zhu, "Disentangled graph convolutional networks," in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 2019, pp. 4212–4221. [Online]. Available: <http://proceedings.mlr.press/v97/ma19a.html>
- [24] L. Li, Y. Zhang, and L. Chen, "Personalized prompt learning for explainable recommendation," *CoRR*, vol. abs/2202.07371, 2022. [Online]. Available: <https://arxiv.org/abs/2202.07371>

- [25] J. Tan, S. Xu, Y. Ge, Y. Li, X. Chen, and Y. Zhang, "Counterfactual explainable recommendation," in *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management, Virtual Event, Queensland, Australia, November 1 - 5, 2021*, G. Demartini, G. Zuccon, J. S. Culpepper, Z. Huang, and H. Tong, Eds. ACM, 2021, pp. 1784–1793. [Online]. Available: <https://doi.org/10.1145/3459637.3482420>
- [26] T. Chen, H. Yin, G. Ye, Z. Huang, Y. Wang, and M. Wang, "Try this instead: Personalized and interpretable substitute 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*, J. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, and Y. Liu, Eds. ACM, 2020, pp. 891–900. [Online]. Available: <https://doi.org/10.1145/3397271.3401042>
- [27] A. Yang, N. Wang, R. Cai, H. Deng, and H. Wang, "Comparative explanations of recommendations," in *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, F. Laforest, R. Troncy, E. Simperl, D. Agarwal, A. Gionis, I. Herman, and L. Médini, Eds. ACM, 2022, pp. 3113–3123. [Online]. Available: <https://doi.org/10.1145/3485447.3512031>
- [28] S. Pan, D. Li, H. Gu, T. Lu, X. Luo, and N. Gu, "Accurate and explainable recommendation via review rationalization," in *Proceedings of the ACM Web Conference 2022, 2022*, pp. 3092–3101.
- [29] J. M. Tomczak and M. Welling, "VAE with a vampprior," *CoRR*, vol. abs/1705.07120, 2017. [Online]. Available: <http://arxiv.org/abs/1705.07120>
- [30] J. J. Zhao, Y. Kim, K. Zhang, A. M. Rush, and Y. LeCun, "Adversarially regularized autoencoders for generating discrete structures," *CoRR*, vol. abs/1706.04223, 2017. [Online]. Available: <http://arxiv.org/abs/1706.04223>
- [31] N. Tishby and N. Zaslavsky, "Deep learning and the information bottleneck principle," 2015. [Online]. Available: <https://arxiv.org/abs/1503.02406>
- [32] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. [Online]. Available: <https://openreview.net/forum?id=HyxQzBceg>
- [33] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," 2013. [Online]. Available: <https://arxiv.org/abs/1312.6114>
- [34] P. Diaconis and D. Freedman, "A dozen de Finetti-style results in search of a theory," *Annales de l'I.H.P. Probabilités et statistiques*, vol. 23, no. S2, pp. 397–423, 1987. [Online]. Available: http://www.numdam.org/item/AIHPB_1987__23_S2_397_0/
- [35] P. Hall, R. C. L. Wolff, and Q. Yao, "Methods for estimating a conditional distribution function," *Journal of the American Statistical Association*, vol. 94, no. 445, pp. 154–163, 1999. [Online]. Available: <http://www.jstor.org/stable/2669691>
- [36] S. Czulbe, O. Krause, I. J. Cox, and C. Igel, "A loss function for generative neural networks based on watsn's perceptual model," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., 2020. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/165a5977cf3b5c4396ba65953d679f17-Abstract.html>
- [37] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004. [Online]. Available: <https://doi.org/10.1109/TIP.2003.819861>
- [38] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "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*, J. X. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, and Y. Liu, Eds. ACM, 2020, pp. 639–648. [Online]. Available: <https://doi.org/10.1145/3397271.3401063>
- [39] K. Mao, J. Zhu, X. Xiao, B. Lu, Z. Wang, and X. He, "Ultragcn: ultra simplification of graph convolutional networks for recommendation," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management, 2021*, pp. 1253–1262.
- [40] S. Peng, K. Sugiyama, and T. Mine, "Svd-gcn: A simplified graph convolution paradigm for recommendation," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022*, pp. 1625–1634.
- [41] F. Wang, Z. Zheng, Y. Zhang, Y. Li, K. Yang, and C. Zhu, "To see further: Knowledge graph-aware deep graph convolutional network for recommender systems," *Information Sciences*, vol. 647, p. 119465, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0020025523010502>
- [42] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [43] J. Wu, X. Wang, F. Feng, X. He, L. Chen, J. Lian, and X. Xie, "Self-supervised graph learning for recommendation," in *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval, 2021*, pp. 726–735.
- [44] J. J. McAuley, J. Leskovec, and D. Jurafsky, "Learning attitudes and attributes from multi-aspect reviews," in *12th IEEE International Conference on Data Mining, ICDM 2012, Brussels, Belgium, December 10-13, 2012*, M. J. Zaki, A. Siebes, J. X. Yu, B. Goethals, G. I. Webb, and X. Wu, Eds. IEEE Computer Society, 2012, pp. 1020–1025. [Online]. Available: <https://doi.org/10.1109/ICDM.2012.110>
- [45] R. He and J. J. McAuley, "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," in *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*, J. Bourdeau, J. Hendler, R. Nkambou, I. Horrocks, and B. Y. Zhao, Eds. ACM, 2016, pp. 507–517. [Online]. Available: <https://doi.org/10.1145/2872427.2883037>
- [46] L. Zheng, V. Noroozi, and P. S. Yu, "Joint deep modeling of users and items using reviews for recommendation," in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM 2017, Cambridge, United Kingdom, February 6-10, 2017*, M. de Rijke, M. Shokouhi, A. Tomkins, and M. Zhang, Eds. ACM, 2017, pp. 425–434. [Online]. Available: <https://doi.org/10.1145/3018661.3018665>
- [47] J. Shuai, K. Zhang, L. Wu, P. Sun, R. Hong, M. Wang, and Y. Li, "A review-aware graph contrastive learning framework for recommendation," in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, ser. SIGIR '22*. New York, NY, USA: Association for Computing Machinery, 2022, p. 1283–1293. [Online]. Available: <https://doi.org/10.1145/3477495.3531927>
- [48] I. O. Tolstikhin, O. Bousquet, S. Gelly, and B. Schölkopf, "Wasserstein auto-encoders," *CoRR*, vol. abs/1711.01558, 2017. [Online]. Available: <http://arxiv.org/abs/1711.01558>
- [49] H. Fu, C. Li, X. Liu, J. Gao, A. Celikyilmaz, and L. Carin, "Cyclical annealing schedule: A simple approach to mitigating KL vanishing," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds. Association for Computational Linguistics, 2019, pp. 240–250. [Online]. Available: <https://doi.org/10.18653/v1/n19-1021>
- [50] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *Proceedings of the 26th International Conference on World Wide Web, ser. WWW '17*. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee, 2017, p. 173–182. [Online]. Available: <https://doi.org/10.1145/3038912.3052569>
- [51] A. M. Elkahky, Y. Song, and X. He, "A multi-view deep learning approach for cross domain user modeling in recommendation systems," in *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015*, A. Gangemi, S. Leonardi, and A. Panconesi, Eds. ACM, 2015, pp. 278–288. [Online]. Available: <https://doi.org/10.1145/2736277.2741667>
- [52] X. Ma, C. Zhou, and E. H. Hovy, "MAE: mutual posterior-divergence regularization for variational autoencoders," in *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [Online]. Available: <https://openreview.net/forum?id=Hke4I2AcKQ>
- [53] R. Fong and A. Vedaldi, "Interpretable explanations of black boxes by meaningful perturbation," *CoRR*, vol. abs/1704.03296, 2017. [Online]. Available: <http://arxiv.org/abs/1704.03296>

X. BIOGRAPHY



Hanqi Yan is a third-year PhD student at the University of Warwick, UK, and a visiting PhD student at King's College London. Her research interest is Natural Language Processing, focusing on interpretable and robust NLP models. She has published several papers in top conferences and high-impact journals. Hanqi obtained her MSc and BEng degrees from Peking University, China and Beihang University, China, respectively.



Lin Gui is a Lecturer at the Department of Informatics, King's College London. His research interests include natural language processing, machine learning, text and data mining, sentiment analysis, and social media analysis. Lin obtained his PhD and MEng degrees in computer science from Harbin Institute of Technology, and his BAsC degree in Mathematics from Nankai University, China.

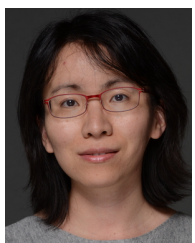


Menghan Wang is an Applied Researcher at eBay inc., working on developing and deploying scalable and extensible recommendation algorithms. Previously, I worked as a senior algorithm engineer in the search and recommendation team at Alibaba inc., where I developed web-scale search & recommendation algorithms and published several papers in top-tier conferences. In 2019, I obtained my PH.D. in Computer Science at Zhejiang University.



Kun Zhang is an Associate Professor with the Department of Philosophy and an Affiliate Faculty Member with the Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA, USA, and an associate professor of Machine Learning at Mohamed bin Zayed University of Artificial Intelligence. His research interests lie in causality, machine learning, and artificial intelligence, especially in causal discovery, hidden causal representation learning, transfer learning, and general-purpose artificial intelligence. Dr. Zhang is a general & program co-

chair for the 2022 Conference on Causal Learning and Reasoning (CLearR), a program co-chair of the UAI 2022, and a general co-chair of the UAI 2023.



Yulan He is a Professor in Natural Language Processing (NLP) at King's College London and an Honorary Professor at the University of Warwick, UK. Her research centers on the integration of machine learning and NLP for text understanding. She currently holds a five-year Turing AI Fellowship, funded by the UK Research and Innovation. She has received several prizes and awards, including a SWSA Ten-Year Award, a CIKM 2020 Test-of-Time Award, and AI 2020 Most Influential Scholar (Honourable Mention) by AMiner. She has served

as a Program Co-Chair for EMNLP 2020 and the General Chair for ACL-IJCNLP 2022. Yulan obtained her PhD degree in Spoken Language Understanding from the University of Cambridge.