

Recommending Burgers based on Pizza Preferences: Addressing Data Sparsity in Recommendation with a Product of Experts

MARTIN MILENKOSKI* and DIEGO ANTOGNINI, École Polytechnique Fédérale de Lausanne, Switzerland
CLAUDIU MUSAT, Swisscom, Switzerland

In this paper, we describe a method to tackle data sparsity and create recommendations in domains with limited knowledge about user preferences. We expand the variational autoencoder collaborative filtering from a single-domain to a multi-domain setting. The intuition is that user-item interactions in a source domain can augment the recommendation quality in a target domain. The intuition can be taken to its extreme, where, in a cross-domain setup, the user history in a source domain is enough to generate high-quality recommendations in a target one. We thus create a Product-of-Experts (POE) architecture for recommendations that jointly models user-item interactions across multiple domains. The method is resilient to missing data for one or more of the domains, which is a situation often found in real life. We present results on two widely-used datasets - Amazon and Yelp, which support the claim that holistic user preference knowledge leads to better recommendations. Surprisingly, we find that in some cases, a POE recommender that does not access the target domain user representation can surpass a strong VAE recommender baseline trained on the target domain.

Additional Key Words and Phrases: Single-Domain Recommendation, Cross-Domain Recommendation, Variational Autoencoder

ACM Reference Format:

Martin Milenkoski, Diego Antognini, and Claudiu Musat. 2021. Recommending Burgers based on Pizza Preferences: Addressing Data Sparsity in Recommendation with a Product of Experts. In *RecSys '21: Workshop of Cross-Market Recommendation*. ACM, New York, NY, USA, 10 pages.

1 INTRODUCTION

Recommender systems are ubiquitous. Due to their prevalence, the user-item interactions are becoming less sparse and collaborative approaches gain the upper hand on content-based methods. This has led to a plethora of neural collaborative filtering methods [3, 12, 16]. However, they suffer from the cold-start problem. Traditionally, it can be alleviated with smarter priors and better defaults. For instance, content-based methods can be integrated as features in complex systems that merge different knowledge sources [31]. Other examples are preference elicitation [7, 8] or critiquing [5, 21], where users express their preferences by interacting with the items' attributes. A major drawback is the assumption of a fixed set of known attributes, especially for domains with expressive features such as hotels or products [1, 19].

There is, however, no escaping the fact that with no knowledge about the users, personalization is impossible. As personalization is what generally makes recommenders useful, we inquire what are the additional information sources that enable it even in the absence of an in-domain interaction history. In this line of thought, known as Cross-domain recommendations [23, 30], we assume consistency between the behaviors of users across different domains. The consistency rests on the belief that observable behaviors in one domain depend partially on some core values and beliefs.

*Work done while at Swisscom.

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.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

The idea of cross-domain recommendations is not new [13, 14, 17, 29]. The main issue with most of these approaches is that they can only be applied in a fully-supervised setting. This means that the model needs to be trained on users with existing interactions on all the considered domains. This is not aligned with one very common use case, when a user has an extensive history in one domain (e.g. buying pizza) but no history in another (e.g. buying burgers).

To fully use existing user-item interactions, we believe a cross-domain recommender system needs to handle cases where user interactions are lacking for one more multiple domains both during training and during inference. To achieve this, we propose a Product of Experts (POE) method, inspired by previous work on multimodal variational autoencoders [28]. By assuming conditional independence between the domains, given the common latent representation, we can ignore the missing domains when calculating the joint latent distribution. We model the joint distribution as a product of experts of the latent distributions of the individual domains. Based on prior work on the usage of variational autoencoders for collaborative filtering [16], we assume that each latent distribution is Gaussian. Also, these assumptions allow us to calculate analytically the joint distribution because the product of Gaussian distributions is itself a Gaussian [6].

We evaluate the performance of the POE model on two different datasets containing Amazon and Yelp reviews. Both datasets contain multiple domains, for instance books and clothing for Amazon and restaurants and shopping for Yelp. We experiment using combinations of two domains at a time, with the combinations ranging from very similar domains - like Books and Kindle in Amazon, to very dissimilar ones, like Books and Clothing. We measure the performance in terms of Recall and NDCG. We do so in two different but related settings. First, we focus on single-domain recommendation accuracy, by training the recommender simultaneously on a pair of domains. The second setting is the pure cross-domain recommendation, where we make recommendations in a target domain for a user without knowing their history in that domain. We train the model with user interactions from both a source and target domains but recommend items for the target domain, only based on the user's interactions in the source one.

Our experiments demonstrate that in a single domain recommendation setup, our method outperforms single-domain VAEs on the majority of cases. In cross-domain recommendations, we show that even completely ignoring the knowledge about a user on the target domain, our model generates recommendations significantly better than the unpersonalized baseline, getting close and even surpassing the recommendation quality obtained with an in-domain VAE.

2 RELATED WORK

There has been a large amount of research dedicated to the usage of Variational Autoencoders in the setting of single-domain recommender systems [9, 11, 15, 16]. Our work is a direct extension of [16] in the setting of cross-domain recommendation. To handle data sparsity, several studies focus on extending the VAE approach in the hybrid setting, incorporating item content as a side information in the recommendation process [2, 9, 11, 15]. Similarly, our work incorporates side information as well. However, we leverage user interactions on other domains instead of item features which might not always be available. Additionally, our model supports weakly-supervised learning and does not require users to have interactions in all domains. Once trained, the model can be used as a single-domain recommender system using only data from the domain of interest, eliminating the need for side information in the recommendation process.

More recent work has been focused on learning the joint distribution $p(x_1, x_2)$ explicitly using a joint inference network $q(z|x_1, x_2)$ with two additional inference networks $q(z|x_1)$ and $q(z|x_2)$ to handle missing data during inference [26, 27]. However, the main drawback of such approaches is that they scale poorly with the number of modalities and are intractable in a general setting. A significant step forward has been made with the multimodal variational autoencoder (MVAE) [28]. MVAE approximates the joint posterior distribution as Product of Experts over the marginal posteriors. This enables cross-modal generation at inference time without the need of additional inference networks. [22]

proposed a mixture of experts model (MMVAE). Its main drawback is that it can only be trained when data is available for all modalities. [25] proposes a novel objective function that utilizes Jensen-Shannon divergence for multiple distributions and a dynamic prior. In this work, we build upon MVAE and adapt it for single and cross-domain recommendation.

The sets of users (or items) in cross-domain recommendation systems might be disjoint, overlap partially, or be the same. The work in this paper focuses on the case of partially overlapping users and items. Early work on this use case has been focused on adopting matrix factorization and transferring shared knowledge based on overlapping users or items [13, 20, 24]. The closest work to ours is [18], where the authors build a cross-domain recommendation framework consisting of domain-specific VAEs for encoding the user interaction vectors and generative adversarial networks for generating user interaction vectors. In our work, we approximate the true joint latent distribution with a product of experts. This allows us to generate recommendations from and to each domain with both partial and full input. Additionally, we propose a cross-domain recommendation model based solely on VAEs without the need for additional networks. Finally, no additional effort in the modeling is needed to extend the solution to more than two domains.

3 METHOD

Before proceeding, we define the following notation used throughout this paper:

- U , I , and D are the number of users, items, and domains. U_d and I_d are the user and item subsets in domain d .
- \mathbf{x}_u^d : A binary vector with length I_d . This is the implicit feedback vector for user u on domain d .
- N_u^d : Total number of interactions of user u with domain d . $N_u^d = \sum_i x_{u,i}^d$

3.1 Variational Autoencoders for Collaborative Filtering

A variational autoencoder (VAE) is a generative model that has been successfully applied in the setting of single-domain collaborative filtering task [9, 15, 16]. Prior work [16] focuses on the application of VAE on single-domain recommendation. Here, we present their work contextualized in the setting of cross-domain recommendation. We explicitly denote the single domain of interest as d and the user feedback vector as \mathbf{x}_u^d . For each user u , the model samples the latent representation \mathbf{z}_u^d from a standard Gaussian prior. Then, a non-linear function $f_{\theta_d}(\cdot)$ is applied on \mathbf{z}_u^d and normalized via softmax to produce a probability distribution $\pi^d(\mathbf{z}_u^d)$. The function $f_{\theta_d}(\cdot)$ is a neural network with parameters θ_d . The feedback vector \mathbf{x}_u^d is assumed to be drawn from a multinomial distribution with probability $\pi^d(\mathbf{z}_u^d)$:

$$\mathbf{z}_u^d \sim \mathcal{N}(0, I_k), \quad \pi^d(\mathbf{z}_u^d) \propto \exp\{f_{\theta_d}(\mathbf{z}_u^d)\}, \quad \mathbf{x}_u^d \sim \text{Mult}(N_u^d, \pi^d(\mathbf{z}_u^d)) \quad \log p_{\theta_d}(\mathbf{x}_u^d | \mathbf{z}_u^d) = \sum_{i=1}^{I_d} x_{u,i}^d \log \pi_i^d(\mathbf{z}_u^d) \quad (1)$$

To learn the generative model we need to approximate the intractable posterior distribution $p(\mathbf{z}_u^d | \mathbf{x}_u^d)$. [16] approximates the posterior distribution using a variational distribution $q(\mathbf{z}_u^d)$ learned with an inference model defined as follows:

$$g_{\phi_d}(\mathbf{x}_u^d) = [\mu_{\phi_d}(\mathbf{x}_u^d), \sigma_{\phi_d}(\mathbf{x}_u^d)] \in \mathbb{R}^{2k}, \quad q_{\phi_d}(\mathbf{z}_u^d | \mathbf{x}_u^d) = \mathcal{N}(\mu_{\phi_d}(\mathbf{x}_u^d), \text{diag}\{\sigma_{\phi_d}^2(\mathbf{x}_u^d)\}) \quad (2)$$

In this way, using the input \mathbf{x}_u^d , the inference model outputs the parameters of the variational distribution $q_{\phi_d}(\mathbf{z}_u^d | \mathbf{x}_u^d)$, which approximates the posterior distribution $p(\mathbf{z}_u^d | \mathbf{x}_u^d)$. Finally, the objective function is the evidence lower bound (ELBO):

$$\mathcal{L}(\mathbf{x}_u^d; \theta_d, \phi_d) = \mathbb{E}_{q_{\phi_d}(\mathbf{z}_u^d | \mathbf{x}_u^d)} [\log p_{\theta_d}(\mathbf{x}_u^d | \mathbf{z}_u^d)] - \beta \cdot \text{KL}(q_{\phi_d}(\mathbf{z}_u^d | \mathbf{x}_u^d) || p(\mathbf{z}_u^d)) \quad (3)$$

The first term can be interpreted as negative reconstruction error and the second term can be interpreted as a regularization term, which is the KL divergence between the variational distribution $q_{\phi_d}(\mathbf{z}_u^d | \mathbf{x}_u^d)$ and the prior $p(\mathbf{z}_u^d)$. The parameter β controls the strength of the regularization and is tuned using a KL annealing strategy [4].

3.2 Variational Autoencoders for Cross-Domain Recommendation

A simple extension of the VAE model in the setting of cross-domain recommendation is treating the set of domains $\{1, \dots, D\}$ as a single domain \tilde{d} . In this case, the input for user u would be a concatenated vector of the individual domain vectors $x_u^{\tilde{d}} = (x_u^1, \dots, x_u^D)$. The dimensionality of the input vector would be $\sum_d I_d$. By setting $d = \tilde{d}$ in Equation 1-3, we can use the same procedure described in Section 3.1 to train a single-domain recommender on the merged domain \tilde{d} .

The main issue with this approach is that it can only be applied in a fully-supervised setting. This means that the model needs to be trained on users with feedback in all domains D . Additionally, since the model is never trained with individual domains, it will not be able to make predictions for users with feedback on only one domain. This fact severely limits the applicability of this approach. We remediate this problem in the next section.

3.3 Product of Experts Variational Autoencoder

In real-life applications, users may have no history on a given domain. For example, on an e-commerce website users might purchase many items from a single category, but they might have never purchased an item from other categories. To make the most effective use of the available data, a cross-domain recommender system needs to handle missing user feedback on some domains both during training and during inference. To tackle this problem, we propose a novel model for cross-domain recommendation inspired by previous work on multimodal variational autoencoders [22, 25, 28].

3.3.1 Model. The proposed model is based on a product of experts variational autoencoder (POE). It is an extension of the collaborative variational autoencoder in the setting of weakly-supervised cross-domain recommendation. In this setting, we work with a set of user feedback vectors $\{x_u^d\}$, where $d \in \{1, \dots, D\}$. As in the multimodal setting [28], we assume that the D domain feedback vectors x_u^d are conditionally independent given the common latent variable z_u . In other words, we assume a generative model of the form $p_\theta(x_u^1, \dots, x_u^D, z_u) = p(z_u)p_{\theta_1}(x_u^1|z_u)\dots p_{\theta_D}(x_u^D|z_u)$ where θ_d are the parameters of the generative model (decoder) associated with domain d . With this factorization, we can ignore missing domains when calculating the marginal likelihood. If a user u has no feedback for domain d , we can omit the term $p_{\theta_d}(x_u^d|z_u)$ from the generative model and still use the model with the known domains for user u .

As in the single-domain setting, for each user u , the model samples a latent representation z_u from a standard Gaussian prior. Then, we define D non-linear functions $f_{\theta_d}(\cdot)$ to produce D propability distributions $\pi^d(z_u)$. The function $f_{\theta_d}(\cdot)$ is a domain-specific decoder with parameters θ_d . The feedback vector x_u^d is assumed to be drawn from a multinomial distribution with probability $\pi^d(z_u)$:

$$z_u \sim \mathcal{N}(0, I_k), \quad \pi^d(z_u) \propto \exp\{f_{\theta_d}(z_u)\}, \quad x_u^d \sim \text{Mult}(N_u^d, \pi^d(z_u)) \quad (4)$$

3.3.2 Variational inference. The main issue for training with missing domains is specifying the 2^D inference networks $q(z_u|X_u)$ for each subset of domains $X_u \subseteq \{x_u^1, \dots, x_u^D\}$. Prior work on multimodal learning [28] has shown that under the assumption of conditional independence, the joint posterior distribution can be approximated by a product of experts:

$$q_\phi(z_u|X_u) \propto p(z_u) \prod_{x_u^d \in X_u} q_{\phi_d}(z_u|x_u^d) \quad (5)$$

where $p(z_u)$ is a prior expert, and $q_{\phi_d}(z_u|x_u^d)$ is the domain-specific inference network (encoder) for domain d as defined in Equation 2. The prior expert is a standard Gaussian distribution. The product distribution described above is not solvable in closed form in a general case. However, when both $p(z_u)$ and $q_{\phi_d}(z_u|x_u^d)$ are Gaussian, their product is itself a Gaussian [6] with mean μ and covariance V defined as: $\mu = (\sum_d \mu_d V_d^{-1})(\sum_d V_d^{-1})^{-1}$, $V = (\sum_d V_d^{-1})^{-1}$, where μ_d, V_d are the parameters of the latent distribution for domain d . In our case, both the prior expert and the inference networks

are Gaussian distributions. For this reason, we can compute the mean and covariance of the product distribution as described above. Therefore, we can avoid specifying the 2^D inference networks and train the model efficiently in terms of the D inference networks $q_{\phi_d}(z_u|x_u^d)$.

3.3.3 Training. Let us define the set of user feedback vectors used as input $X_u = \{x_u^d | \text{domain } d \text{ used as input for user } u\}$ and the set of user feedback vectors present in the data $\widetilde{X}_u = \{x_u^d | \text{domain } d \text{ present for user } u\}$. We use the following form of the ELBO objective for a given input X_u :

$$\mathcal{L}(X_u; \theta, \phi) = \mathbb{E}_{q_{\phi}(z_u|X_u)} \left[\sum_{x_u^d \in \widetilde{X}_u} \lambda_d \log p_{\theta_d}(x_u^d|z_u) \right] - \beta * KL(q_{\phi}(z_u|X_u) || p(z_u)) \quad (6)$$

In multimodal learning, prior work consider $X = \widetilde{X}$ [28]. However, treating them as different allows us to teach the model to generate better cross-domain recommendations when using individual domains as input. In this way, when presented with a single domain as input, the model is taught to perform well on all known domains instead of only the domain used as input. λ_d is the weight given to domain d in the loss function.

Optimizing this function on a dataset with fully present domains for all users has an undesirable effect; since the model does not see data points with missing domains during training, it cannot do inference with individual domains. For this reason, we use a modified version of the sub-sampled objective in [28] to train the model for individual inputs:

$$\mathcal{L}(\{x_u^1, \dots, x_u^D\}; \theta, \phi) + \sum_{d=1}^D \mathcal{L}(\{x_u^d\}; \theta, \phi) \quad (7)$$

3.3.4 Inference. We now describe how to make predictions given a trained POE model. We describe the prediction process in the setting of two domains ($D = 2$). Without loss of generality, let us define $s = d$ as the source domain, and $t = d$ as the target domain. In this paper, we present two applications of the model - single-domain and cross-domain recommendation. In the single-domain recommendation setting, we recommend items from the target domain using the user feedback from the target domain. In the cross-domain recommendation setting, we recommend items from the target domain using the user feedback from the source domain.

Let us define the domain used as input as v , with $v = t$ in the single-domain setting, and $v = s$ in the cross-domain setting. In both cases, we pass the input x_u^v through the inference model $g_{\phi_v}(x_u^v)$ to obtain the parameters μ_{ϕ_v} and σ_{ϕ_v} . Then, we calculate the product of experts distribution and take the mean of the distribution as the latent representation z_u . Finally, we rank the items in the target domain based on the un-normalized predicted multinomial probability $f_{\theta_t}(z_u)$.

4 EXPERIMENTS

4.1 Datasets

In order to evaluate the quantitative performance of our proposed POE model, we perform experiments using seven real-world publicly available datasets based on Amazon reviews [19] and Yelp businesses. All datasets contains user and item (partial) overlaps. This means that both users and items can be present in more than one domain.

Amazon Reviews dataset contains information regarding product reviews and metadata obtained from Amazon [19]. We consider each category as separate domain. We chose three of the most popular categories Books, Movies and Clothing to maximize user overlap. Additionally, we chose Kindle Store as a semantically related category with Books in order to study the effect of semantic similarity. We binarize the reviews to create a dataset with implicit feedback. Only reviews ≥ 3.5 are considered as a signal that the user liked the item. All users with less than 5 reviews are filtered out.

Yelp dataset contains information about businesses, reviews and user data on Yelp [10]. We choose some of the

Table 1. Descriptive statistics of the datasets after preprocessing. Interactions are non-zero entries.

#	Dataset	Category	Item threshold	# of users	# of items	# of interactions	Density
1	Amazon	Books	200	63,711	29,124	2,041,610	0.11%
		Kindle Store	30		30,243	1,014,985	0.05%
2	Amazon	Books	200	43,242	29,266	702,081	0.06%
		Movies	20		33,793	671,961	0.05%
3	Amazon	Books	200	42,965	29,354	602,743	0.05%
		Clothing	150		24,244	390,677	0.04%
4	Yelp	Restaurants	100	86,566	7,886	599,587	0.09%
		Hotels	1		5,089	142,516	0.03%
5	Yelp	Restaurants	10	138,801	35,361	1,260,613	0.03%
		Shopping	1		30,998	291,897	0.01%
6	Yelp	Food	1	110,427	29,929	577,291	0.02%
		Shopping	1		30,171	256,196	0.01%
7	Yelp	Burgers	1	67,088	5,340	143,799	0.04%
		Pizza	1		6,364	151,864	0.04%

most popular categories in order to maximize user overlap. The ratings are binarized to create a dataset with implicit feedback. Only ratings ≥ 3.5 were considered as a signal that the user liked the business.

4.2 Metrics

A trained POE model on D domains can be used to produce single-domain recommendation for all domains D . For this reason, a single model can be evaluated on the metric of interest for all domains. To present the performance of a single model on all domains we use the concept of Pareto front. Formally, let M be the metric of interest and let w_1 and w_2 be D -dimensional vectors associated with two models m_1 and m_2 . $w_{i,d}$ is the performance of model m_i on domain d on the metric M . The model m_1 is said to dominate the model m_2 if $\forall d \in \{1, \dots, D\} w_{1,d} \geq w_{2,d}$ and $\exists d \in \{1, \dots, D\} w_{1,d} > w_{2,d}$. A Pareto optimal model is a model m^* which is not dominated by any other model m . The set of all Pareto optimal models is called a Pareto set and its visualization is called a Pareto front. As metrics, we use Recall@K and NDCG@K.

4.3 Experimental Details

We run the experiments on 7 domain combinations. For each combination of categories we choose a category-specific threshold and we filter out all items with number of reviews less than that threshold. The choice of the threshold was dependent on the domain combination and was usually tuned to bring the complexity of the domains close to each other. Table 1 presents the postprocessing attributes of the datasets. We use the same model architecture and hyperparameters for both the VAE and POE approach in order to obtain fair comparison. With the large combination of applications, domains, and combinations, we employed the default parameters used in [16]. We use 95% of users as a train set and 5% as a test set. In the test set, 20% of user interactions are masked and used as ground truth to evaluate the model. The remaining 80% of interactions are used as input.

4.4 Single-Domain Recommendation

In the single-domain setting we compare the POE model to a traditional single-domain VAE model from [16]. On Figure 1, the VAE performance on the individual domains is presented with dashed horizontal and vertical lines. The

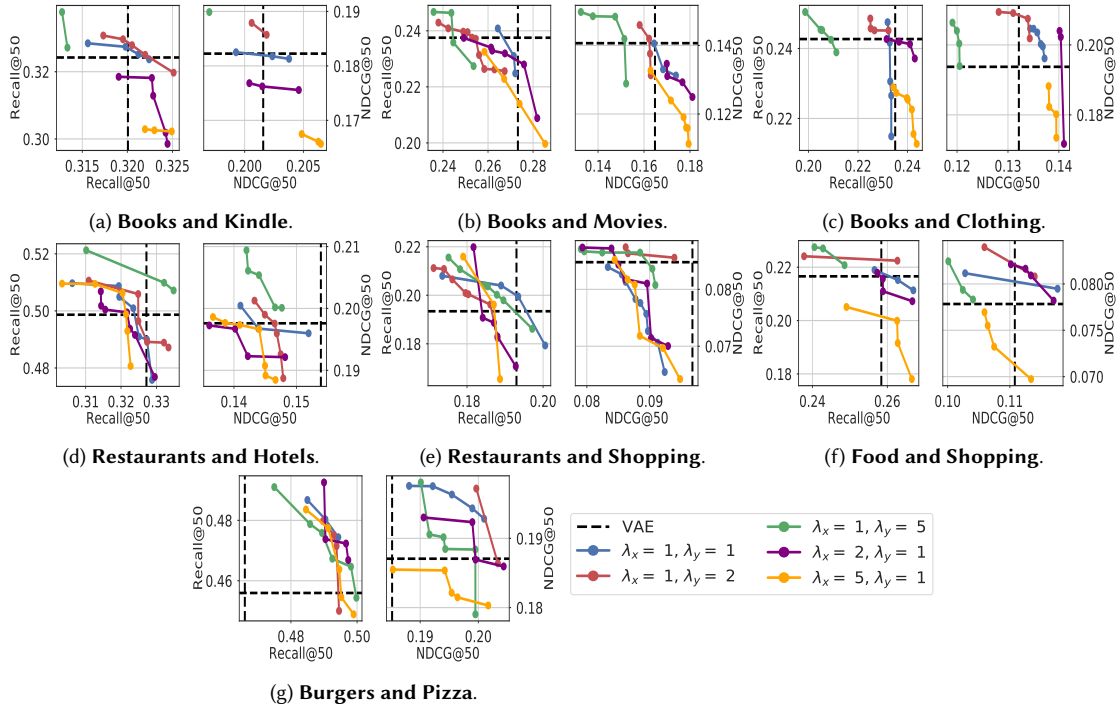


Fig. 1. Pareto fronts for Recall@50 and NDCG@50 on all 7 datasets with different domain weights.

results from the POE model are presented as a Pareto front. For each domain combination we show plots for the performance on the Recall@50 and NDCG@50. In all cases, the best result is in the upper right corner.

First, we train the POE model by giving equal weight to both domains in the loss function ($\forall d \in \{1, \dots, D\} \lambda_d = 1$ in Equation 6). The Pareto fronts are shown in blue on Figure 1. We can observe that in some cases like Pizza and Burgers, the model trained with equal weights is improving upon the VAE performance on both domains simultaneously. However, on other cases like NDCG@50 on Restaurants and Shopping, the POE model is underperforming. Then, we assign different weights to the domains and train the POE model. In some cases like Books and Kindle or Books and Movies our model results in improved single-domain performance compared to the baseline model. However, on other cases increasing the weight of the Restaurant domain leads to a decrease in performance on that domain.

The results show that in most cases, a set of optimal weights exists and we can improve upon the performance of VAE on a single domain of interest. Additionally, for a few combinations, we can improve simultaneously the performance on all domains. Finally, on some pairs, the POE model performs better in one domain only on Recall@50.

4.5 Cross-Domain Recommendation

The cross-domain recommendation task is concerned with generating recommendation on the target domain by exploiting knowledge from the source domain. In our case, we exploit the ability of the POE model to generate recommendations for users with missing domains in order to tackle the cross-domain task. We perform the training and evaluation of POE on the intersection of users present in both domains. During evaluation, the model is shown user history from the source domain but evaluated on the target domain. In this way, we are able to simulate users with

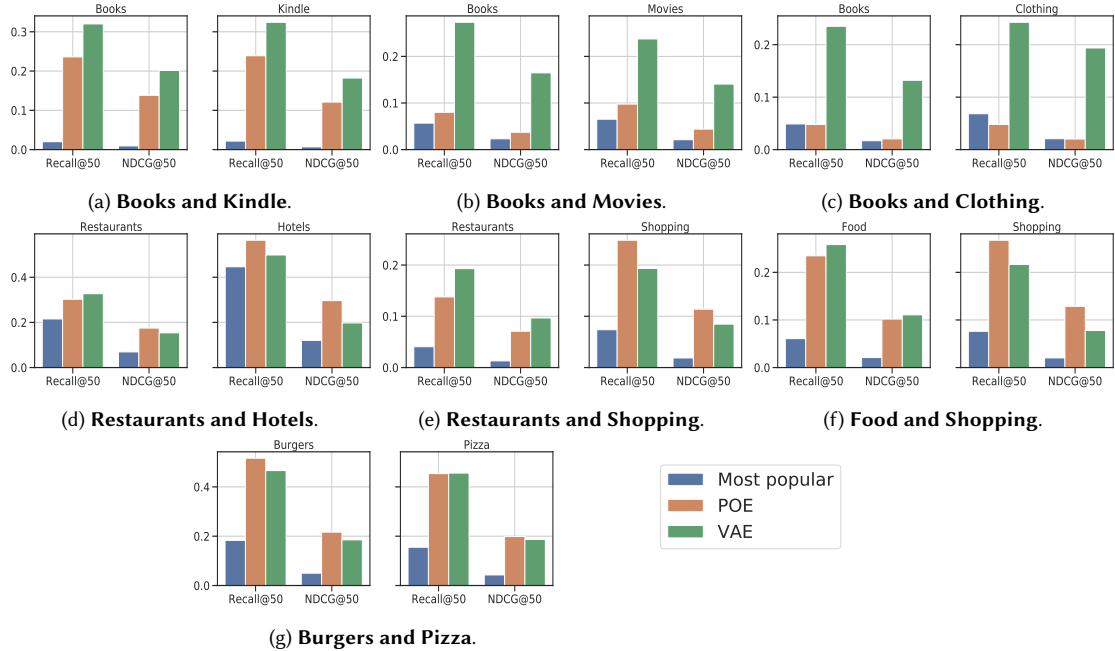


Fig. 2. Performance of the POE model on the domain-specific cold-start problem.

a missing domain and still evaluate the model by using users with known history on both domains.

We compare the POE model with a baseline that recommends the most popular items in the dataset. Even though this is a very simple baseline, it is still widely used in practice in cases when there is no known history for a new user on the domain of interest. We also compare the POE model’s performance with a single-domain VAE trained and evaluated on the target domain. Unlike the POE, the VAE is shown the input from the target domain during evaluation. In this way, we evaluate to which extent can the POE model transfer knowledge from one domain to the other.

In Figure 2, we see that the POE model outperforms the *Popular* baseline in most cases. On 4 out of 7 combinations the POE model is improving upon the VAE baseline on at least one domain. This means that the POE model can improve upon the single-domain recommendation performance of a standard VAE without seeing the user history on the domain of interest. We can conclude that the POE model is a viable solution to the domain-specific cold-start problem.

5 DISCUSSION

Collaborative filtering approaches have shown great potential in the field of recommender systems [3, 12, 16]. A major drawback of traditional collaborative filtering techniques is their inability to handle the cold-start problem. Prior work has proposed different techniques of resolving this issue, such as integration of content-based methods [31], preference elicitation [7, 8] and critiquing [5, 21]. All of these techniques rely on some known attributes about the users or items which are often not available. On the other hand, cross-domain recommendation tries to generate recommendation for a new domain based only on user-item interactions from another domain. However, most of the existing research in this area focuses on the setting where user-item interactions are available for all domain of interest. To resolve this limitation, we propose a new Product of Experts approach that can be both trained and used on users with interactions

in only a subset of the domains. We show that our model can improve recommendations of the target domain by either using the interactions from both domains or the interactions only from the source domain. In this work, we show that the weights of the individual domains in the loss function have a significant impact on the performance. We presented results from five different weight combinations. A potential improvement in this regard is to treat the weights as hyperparameters and tune them with standard hyperparameter tuning methods, or treat them as model parameters and learning them jointly with the model during training. Additionally, we could treat each domain as a separate objective and utilize techniques from multi-objective optimization to obtain the optimal weights. Finally, there is a need to understand why some domain combinations work better than others. In this regard, techniques from multi-task learning can be used to measure the correlation between the different domains. Finally, it might be beneficial to analyze the learned latent distributions by either measuring their similarity or visualizing them using dimensionality reduction techniques.

6 CONCLUSION

We presented a method to alleviate data sparsity in creating recommendations. We base our approach on the single-domain variational autoencoder collaborative filtering and expand it to a multi-domain setup. We created a novel Product-of-Experts approach for joint preference modelling. We show that, using it, a user's preferences in a source domain can be used to improve recommendations in a target domain in two ways: (1) by complementing a shallow representation of the user in the target domain and (2) in a cross-domain setup, where the recommendations are made solely based on the knowledge in the source domain. The results point to an outperformance of POE over single-domain VAE approaches. Moreover, the cross-domain recommendations are competitive, even surpassing the in-domain recommendations in some cases.

REFERENCES

- [1] Diego Antognini and Boi Faltings. 2020. HotelRec: a Novel Very Large-Scale Hotel Recommendation Dataset. In *Proceedings of The 12th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 4917–4923. <https://www.aclweb.org/anthology/2020.lrec-1.605>
- [2] Diego Antognini and Boi Faltings. 2021. Fast Multi-Step Critiquing for VAE-based Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems (Amsterdam, Netherlands) (RecSys '21)*.
- [3] Diego Antognini, Claudiu Musat, and Boi Faltings. 2021. Interacting with Explanations through Critiquing. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, (IJCAI 2021)*.
- [4] Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2015. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349* (2015).
- [5] Robin D. Burke, Kristian J. Hammond, and Benjamin C. Young. 1996. Knowledge-Based Navigation of Complex Information Spaces. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 1 (Portland, Oregon) (AAAI'96)*. AAAI Press, 462–468.
- [6] Yanshuai Cao and David J Fleet. 2014. Generalized product of experts for automatic and principled fusion of Gaussian process predictions. *arXiv preprint arXiv:1410.7827* (2014).
- [7] Urszula Chajewska and Lise Getoor. 1998. Utility Elicitation as a Classification Problem. In *IN PROCEEDINGS OF THE FOURTEENTH CONFERENCE ON UNCERTAINTY IN ARTIFICIAL INTELLIGENCE*. 79–88.
- [8] Urszula Chajewska, Daphne Koller, and Ronald Parr. 2000. Making Rational Decisions using Adaptive Utility Elicitation. In *In Proceedings of the Seventeenth National Conference on Artificial Intelligence*. 363–369.
- [9] Yifan Chen and Maarten de Rijke. 2018. A collective variational autoencoder for top-n recommendation with side information. In *Proceedings of the 3rd Workshop on Deep Learning for Recommender Systems*. 3–9.
- [10] Yelp Dataset. 2016. Yelp Open Dataset. <https://www.yelp.com/dataset>
- [11] Ming He, Qian Meng, and Shaozong Zhang. 2019. Collaborative additional variational autoencoder for top-N recommender systems. *IEEE Access* 7 (2019), 5707–5713.
- [12] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.

- [13] Guangneng Hu, Yu Zhang, and Qiang Yang. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 667–676.
- [14] Liang Hu, Jian Cao, Guandong Xu, Longbing Cao, Zhiping Gu, and Can Zhu. 2013. Personalized recommendation via cross-domain triadic factorization. In *Proceedings of the 22nd international conference on World Wide Web*. 595–606.
- [15] 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*. 305–314.
- [16] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*. 689–698.
- [17] Babak Loni, Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Cross-domain collaborative filtering with factorization machines. In *European conference on information retrieval*. Springer, 656–661.
- [18] Linh Nguyen and Tsukasa Ishigaki. 2018. Domain-to-Domain Translation Model for Recommender System. *arXiv preprint arXiv:1812.06229* (2018).
- [19] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. 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)*. 188–197.
- [20] Weike Pan, Evan Xiang, Nathan Liu, and Qiang Yang. 2010. Transfer learning in collaborative filtering for sparsity reduction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 24.
- [21] James Reilly, Kevin McCarthy, Lorraine McGinty, and Barry Smyth. 2005. Explaining compound critiques. *Artificial Intelligence Review* (2005).
- [22] Yuge Shi, Narayanaswamy Siddharth, Brooks Paige, and Philip HS Torr. 2019. Variational mixture-of-experts autoencoders for multi-modal deep generative models. *arXiv preprint arXiv:1911.03393* (2019).
- [23] Ajit P. Singh and Geoffrey J. Gordon. 2008. Relational Learning via Collective Matrix Factorization. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Las Vegas, Nevada, USA) (KDD '08)*. Association for Computing Machinery, New York, NY, USA, 650–658. <https://doi.org/10.1145/1401890.1401969>
- [24] Ajit P Singh and Geoffrey J Gordon. 2008. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. 650–658.
- [25] Thomas M Sutter, Imant Daunhawer, and Julia E Vogt. 2020. Multimodal Generative Learning Utilizing Jensen-Shannon-Divergence. *arXiv preprint arXiv:2006.08242* (2020).
- [26] Masahiro Suzuki, Kotaro Nakayama, and Yutaka Matsuo. 2016. Joint multimodal learning with deep generative models. *arXiv preprint arXiv:1611.01891* (2016).
- [27] Ramakrishna Vedantam, Ian Fischer, Jonathan Huang, and Kevin Murphy. 2017. Generative models of visually grounded imagination. *arXiv preprint arXiv:1705.10762* (2017).
- [28] Mike Wu and Noah Goodman. 2018. Multimodal generative models for scalable weakly-supervised learning. *arXiv preprint arXiv:1802.05335* (2018).
- [29] Feng Yuan, Lina Yao, and Boualem Benatallah. 2019. DAREC: Deep domain adaptation for cross-domain recommendation via transferring rating patterns. *arXiv preprint arXiv:1905.10760* (2019).
- [30] Xiaoxue Zhao and Weinan Zhang. 2010. Multi-domain collaborative filtering. In *in: Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI)*.
- [31] 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. *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (Oct 2020). <https://doi.org/10.1145/3340531.3411954>