Connecting Unseen Domains: Cross-Domain Invariant Learning in Recommendation

Yang Zhang
yaoling.zy@antgroup.com
Ant Group
Hangzhou, China

Yue Shen
zhanying@antgroup.com
Ant Group
Hangzhou, China

Dong Wang
yishan.wd@antgroup.com
Ant Group
Hangzhou, China

Jinjie Gu
jinjie.gujj@antgroup.com
Ant Group
Hangzhou, China

Guannan Zhang
zgn138592@antgroup.com
Ant Group
Hangzhou, China

ABSTRACT

As web applications continue to expand and diversify their services, user interactions exist in different scenarios. To leverage this wealth of information, cross-domain recommendation (CDR) has gained significant attention in recent years. However, existing CDR approaches mostly focus on information transfer between observed domains, with little attention paid to generalizing to unseen domains. Although recent research on invariant learning can help for the purpose of generalization, relying only on invariant preference may be overly conservative and result in mediocre performance when the unseen domain shifts slightly. In this paper, we present a novel framework that considers both CDR and domain generalization through a united causal invariant view. We assume that user interactions are determined by domain-invariant preference and domain-specific preference. The proposed approach differentiates the invariant preference and the specific preference from observational behaviors in a way of adversarial learning. Additionally, a novel domain routing module is designed to connect unseen domains to observed domains. Extensive experiments on public and industry datasets have proved the effectiveness of the proposed approach under both CDR and domain generalization settings.

CCS CONCEPTS

• **Information systems** → **Online advertising**.

KEYWORDS

Cross-Domain Recommendation; Domain Generalization; Invariant Learning

ACM Reference Format:

Yang Zhang, Yue Shen, Dong Wang, Jinjie Gu, and Guannan Zhang. 2023. Connecting Unseen Domains: Cross-Domain Invariant Learning in Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23), July 23–27, 2023, Taipei, Taiwan.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3539618.3591965

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

ACM ISBN 978-1-4503-9408-6/23/07…$15.00
https://doi.org/10.1145/3539618.3591965

1 INTRODUCTION

Nowadays, recommender system plays a crucial role in many web applications, such as Amazon and YouTube. The primary objective of a recommender system is to provide users with the most relevant products based on their interests and browsing history. Most existing recommender system research is largely limited to historical user behaviors within a single domain [4, 9, 12]. With the rapid development of the internet industry, web applications provide more in-depth and diversified services. For instance, in an e-commerce platform, there are different business domains serving different purposes. The front page serves as a domain for exploration, while the search results page is a domain focused on purchasing. As users begin to engage with multiple domains, cross-domain recommendation (CDR) has attracted increasing attention in recent years.

Compared with the single-domain recommendation, the key issue of CDR lies in the differences in user behavior patterns and preferences across domains [6, 15]. To tackle this issue, numerous solutions have been put forward in the literature, with a primary focus on learning effective representations that can be transferred across domains. The underlying idea is to bridge the semantic gap between domains through shared parameters [2], feature mapping [6], or semantic space alignment [23]. Despite numerous studies in CDR, they all have a limitation of only improving performance in observed domains and cannot be applied to entirely unseen domains [22]. This issue is particularly important in industries, such as e-commerce platforms, which frequently have new promotional activities. To prepare for such events, a model must be established using historical data from multiple domains (*i.e.,* historical promotional activities) and provide stable performance in unseen domains (*i.e.,* upcoming new promotional activity). However, existing CDR research falls short of meeting this requirement.

To enable generalization in unseen domains, a stream of studies introduces invariant learning. The goal of invariant learning is to capture representations with invariant predictive ability across domains. Along this line, there are mainly four types of methods: kernel-based methods [3, 11], domain adversarial learning [8], explicit feature alignment [17, 18], and invariant risk minimization [1, 13]. However, the latest research suggests that relying solely on invariant parts may be overly conservative [5], leading

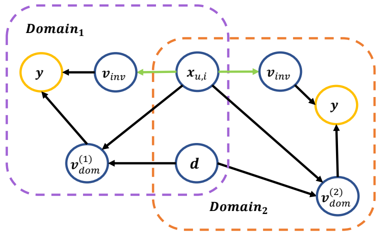

Figure 1: Causal Graph of Cross-Domain Recommendation.

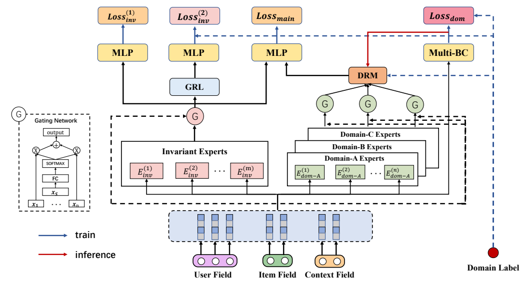

Figure 2: The Overall Architecture of Grace.

to mediocre performance when the distribution shift in unseen domains is relatively small, which greatly limits its applications.

In summary, previous CDR approaches mostly neglected the generalization to unseen domains, while domain generalization methods sacrificed the modeling of specific preferences in observed domains for the sake of cross-domain invariance and generalization. Consequently, these methods are limited in their ability to *simultaneously achieve good performance in both observed and unseen domains*. To achieve this goal, we formulate the above two targets from a united causal invariant view. User behaviors are driven by a combination of domain-invariant preferences (*e.g.,* a preference for cosmetics among women) and domain-specific preferences (*e.g.,* a preference for cheap and interesting items in an exploration-guided domain and expensive items with higher quality in a purchase-guided domain). Especially, the differences in user preferences under different domains are caused by the existence of domain-specific confounders (*e.g.,* domain settings).

In this paper, we propose *Grace* (**G**ene**R**aliz**A**ble **C**ross-domain **E**stimator), a novel framework that elevates CDR performance while exhibiting exceptional domain generalization capabilities. We first disentangle preferences into the domain-invariant preference and domain-specific preferences with respect to each observed domain. Such preference disentanglement enables the transfer of knowledge across domains while also emphasizing the unique preferences of each domain. Furthermore, due to the difficulties in obtaining domain-invariant preference, we employ domain adversarial learning [8] and design two restrictions to ensure both domain invariance and predictive power. Moreover, to overcome the drawback of traditional invariant learning being too conservative, we propose a novel domain routing module that dynamically decides how domain-specific preferences affect predictions. Overall, the main contributions are summarized as that: (1) To the best of our knowledge, this is the first study to unify CDR and domain generalization (DG) through a causal invariant view. (2) We propose Grace which effectively differentiates and utilizes domain-invariant and domain-specific preferences. (3) We conduct extensive experiments on both public and industry datasets. The results demonstrate the effectiveness of Grace over previous methods.

2 CAUSAL INVARIANT VIEW

Assume that the observed user-item interactions are collected from a set of domains, denoted by $\mathcal{U}$, $\mathcal{I}$ and $\mathcal{D}$ respectively, where $u \in \mathcal{U}$ denotes a user, $i \in \mathcal{I}$ denotes an item and $d \in \mathcal{D}$ denotes a domain. It's notable that we do not make specific assumptions about overlapped users or items. Instead, all user-item pairs from different domains share the same attribute set.

We use the causal graph in Figure 1 to define the generation of observational user behaviors. In particular,

- $\boldsymbol{x}_{u,i}$ represents raw features from both user (*e.g.,* profile, historical behaviors) and item (*e.g.,* category, price).
- $\boldsymbol{v}_{inv}$ represents the domain-invariant preference of user for item, which is characterized in the representation space.
- $\boldsymbol{v}_{dom}$ represents the domain-specific preference of user for item affected by domain confounders and the preference is characterized in the representation space.
- $\boldsymbol{y}$ represents the user feedback on the item (*e.g.,* click/buy).
- $\boldsymbol{d}$ represents the domain which is the agents for the set of domain confounders.

The causal graph in Figure 1 indicates that

$$p(y|\boldsymbol{x}_{u,i}, d) = p(y|\boldsymbol{v}_{inv}, \boldsymbol{v}_{dom}) * p(\boldsymbol{v}_{inv}|\boldsymbol{x}_{u,i}) * p(\boldsymbol{v}_{dom}|\boldsymbol{x}_{u,i}, d), \quad (1)$$

where $\boldsymbol{v}_{inv} \in \mathbb{R}^L$ is invariant among different domains while $\boldsymbol{v}_{dom} \in \mathbb{R}^L$ is directly affected by domain confounders.

3 THE PROPOSED APPROACH

3.1 Preference Disentanglement

As illustrated in Figure 2, the proposed Grace first disentangles user preference from two perspectives:

(1) Differentiate between domain-invariant preference $\boldsymbol{v}_{inv}$ and domain-specific preference $\boldsymbol{v}_{dom}$.
(2) Differentiate domain-specific preferences from different domains. Each domain owns a unique preference extraction network, and $\boldsymbol{v}_{dom}^{(d)}$ denotes the domain-specific preference for domain d. Furthermore, $\boldsymbol{V}_{dom} = \{\boldsymbol{v}_{dom}^{(d)} | d \in \mathcal{D}\}$.

Preference Extraction Layer. For each disentangled preference representation $\boldsymbol{v}^{(k)} \in \{\boldsymbol{v}_{inv}\} \cup \boldsymbol{V}_{dom}$, we adopt the same network architecture named as preference extraction layer (PEL). PEL actually takes the embedding vector $\boldsymbol{r}$ as input, which is transformed from raw sparse features $\boldsymbol{x}_{u,i}$.

Inspired by recent research on mixture-of-experts (MoE) [16] and gating networks [20], PEL first utilizes m independent experts to extract knowledge from different representation subspaces. Moreover, the gating network produces a distribution over all experts, and the final output $\boldsymbol{v}^{(k)}$ is formulated as:

$$\boldsymbol{v}^{(k)} = g^{(k)}(\boldsymbol{r})S^{(k)}(\boldsymbol{r}), \quad (2)$$

where $g^{(k)}$ is the weighting function to combine the results from all experts through linear transformation and a SoftMax layer:

$$g^{(k)} = SoftMax(\boldsymbol{W}_g^{(k)} \boldsymbol{r}), \tag{3}$$

where $\boldsymbol{W}_g^{(k)}$ is a parameter matrix. $S^{(k)}(\boldsymbol{r})$ is a selected matrix composed of vectors from all experts:

$$S^{(k)}(\boldsymbol{r}) = [E_{(k,1)}^{\top}, E_{(k,2)}^{\top}, \cdots, E_{(k,m)}^{\top}]^{\top}, \tag{4}$$

where each expert is a single-layer network.

3.2 Invariant Preference Learning

To capture discriminative and domain-invariant preference $\boldsymbol{v}_{inv}$ across multiple domains, it is crucial to balance two objectives: maximizing the predictive power of $\boldsymbol{v}_{inv}$ to y and minimizing the domain information in $\boldsymbol{v}_{inv}$.

3.2.1 Maintaining Predictive Power. To maintain the predictive capability of $\boldsymbol{v}_{inv}$ with respect to user feedback y and prevent it from degrading into mere white noise, we introduce a Multi-Layer Perceptron (MLP), denoted as φ_1 to predict observed feedback y and optimize the loss term for specific recommendation tasks

$$Loss_{inv}^{(1)} = \frac{1}{|O|} \sum_{(u,i,d) \in O} l_{task}(\varphi_1(\boldsymbol{v}_{inv}), y), \tag{5}$$

where O is the observed dataset. Specially, l_{task} denotes binary cross-entropy (BCE) loss for rank-based recommendation tasks (e.g. CTR) and mean square error (MSE) for rating prediction.

3.2.2 Achieving Domain Invariance. To ensure the invariance across domains, $\boldsymbol{v}_{inv}$ should contain no domain-specific information. Specifically, the invariant preference $\boldsymbol{v}_{inv}$ should be constructed to confuse the classifier so that even a well-trained classifier cannot accurately predict the domain label of a given sample. To achieve this, we draw inspiration from Domain Generalization methods [7, 8, 14] and employ domain adversarial learning.

We introduce an MLP denoted as ϕ parametered by θ that identifies domain label using $\boldsymbol{v}_{inv}$. We measure the quality of classification using the cross-entropy (CE) loss

$$Loss_{inv}^{(2)} = \frac{1}{|O|} \sum_{(u,i,d) \in O} \mathrm{CE}(\phi(\boldsymbol{v}_{inv}), d). \tag{6}$$

In summary, the overall objective function is formulated as

$$Loss_{inv} = Loss_{inv}^{(1)} - \alpha Loss_{inv}^{(2)}, \tag{7}$$

where α is a hyperparameter to balance two losses and we aim to find the saddle point $\boldsymbol{v}_{inv}, \theta$ such that

$$\boldsymbol{v}_{inv} = \mathrm{argmin}\{Loss_{inv}^{(1)} - \alpha Loss_{inv}^{(2)}\}, \tag{8}$$

$$\theta = \mathrm{argmin}\{Loss_{inv}^{(2)}\}. \tag{9}$$

Optimizing the above equations directly using gradient descent algorithms has been proven to be infeasible [21]. This may result in the subtraction of gradients instead of their summation.

Inspired by the domain adversarial learning approach [8], we introduce the gradient reversal layer (GRL) between $\boldsymbol{v}_{inv}$ and ϕ to address this issue, which leaves the input unchanged during forward propagation and reverses the gradient by multiplying it by a negative scalar during the backpropagation.

3.3 Domain Routing Module

Utilizing only domain-invariant preference may lead to overly conservative models, especially in real-world scenarios, where unseen domains may share some similarities with observed domains (*e.g.,* upcoming new promotional activity and previously observed ones in the same app).

To address this issue, we propose a novel domain routing module (DRM) that dynamically *selects* the domain-specific preference:

$$\mathrm{DRM}(\boldsymbol{V}_{dom}) = \begin{cases} \boldsymbol{v}_{dom}^{(d)}, & \text{if } d \text{ is observed;} \\ \sum_{t \in \mathcal{D}} \frac{\max(p_t - \lambda, 0)}{1 - \lambda} \boldsymbol{v}_{dom}^{(t)}, & \text{if } d \text{ is unknown.} \end{cases} \tag{10}$$

During training or inferring with observed domain label d (*i.e.,* the CDR setting), DRM actually selects the specific $\boldsymbol{v}_{dom}^{(d)}$.

If the domain is unknown during inference (*i.e.,* the domain generalization setting), the output of DRM can be interpreted as the weighted sum of domain-specific preferences from each observed domain in training data. The weighting coefficient p_t measures the probability of the sample belonging to the t-th domain.

We employ multiple binary classifiers to estimate the coefficients with respect to each observed domain. The loss term is:

$$Loss_{dom} = \frac{1}{|O| \cdot |\mathcal{D}|} \sum_{(u,i,d) \in O} \sum_{t \in \mathcal{D}} \mathrm{BCE}(\Phi_t(r), \mathbb{I}(t = d)), \tag{11}$$

where $\mathbb{I}(\cdot, \cdot)$ is the indicator function, Φ_t is the binary classifier for domain t and $p_t = \Phi_t(r)$. λ is a threshold hyperparameter that filters out noise information.

The proposed Grace predicts user feedback y based on both domain-invariant preference $\boldsymbol{v}_{inv}$ and domain-specific preference selected by DRM. We concatenate the preferences as $\boldsymbol{v}_{main} = \mathrm{Concat}\,(\boldsymbol{v}_{inv}, \mathrm{DRM}(\boldsymbol{V}_{dom}))$ and the loss term is

$$Loss_{main} = \frac{1}{|O|} \sum_{(u,i,d) \in O} l_{task}(\varphi_2(\boldsymbol{v}_{main}), y), \tag{12}$$

where φ_2 is the MLP for the main prediction task.

Finally, we optimize the model parameters by integrating losses for the main prediction task, invariant preference learning, and domain estimation

$$Loss = Loss_{main} + Loss_{inv} + Loss_{dom}. \tag{13}$$

Remark. If the estimated probabilities for each observed domain are below a threshold value λ (*i.e.,* the sample has little similarity with the observed domains), the DRM filters out all domain-specific preferences. The main task tower then downgrades to a conventional invariant learning paradigm, which considers only domain-invariant preferences and achieves stable performance.

4 EXPERIMENTS

4.1 Experimental Setup

4.1.1 Datasets. We conduct experiments on both Public and Industry datasets. For each dataset, we collected three observed domains, among which 50% of the samples are allocated as the training set,

Table 1: Results of comparison study on both Public and Industry datasets. "↑" ("↓") indicates larger (smaller) is better.

Category	Method	Amazon (MSE "↓")				Industry (AUC "↑")			
		Book	CD	Movie	Music (OOD)	A_1	A_2	A_3	A_4 (OOD)
Single domain	DeepFM	0.7852	0.8421	0.9498	1.1121	0.8427	0.8165	0.8571	0.6836
CDR	Multi-view DNN	0.7658	0.7891	0.9372	-	0.8839	0.8672	0.9115	-
	Co-Net	0.7642	0.7785	0.9304	-	0.8974	0.8738	0.9183	-
DG	DANN	0.8160	0.8742	0.9757	1.0572	0.8697	0.8409	0.8814	0.7398
	IRM	0.8235	0.8688	0.9801	1.0243	0.8635	0.8393	0.8857	0.7426
Our methods	Grace-Inv	0.7996	0.8259	0.9545	0.9764	0.8754	0.8501	0.8934	0.7415
	Grace	**0.7378**	**0.7439**	**0.9083**	**0.9254**	**0.9117**	**0.8977**	**0.9378**	**0.7483**

30% as the validation set, and 20% as the test set. Since our method can generalize to unseen domains, we also sample data from a fourth domain for out-of-distribution (OOD) evaluation.

Amazon Dataset [19] contains reviews from Amazon, including (user, item, rating) tuples and product metadata (category, price, brand) from different domains. We choose *Book*, *CD* (named "CDs and Vinyl"), *Movie* (named "Movies and TV") as observed domains and evaluate OOD performance on *Music* (named "Digital Music").

Industry Dataset is collected from our online recommender systems, which contains data from four distinct promotional activities, denoted as A_1, A_2, A_3 and A_4, each serving a different business purpose. We choose A_4 for OOD evaluation.

4.1.2 Baselines. We adopt both cross-domain recommendation (CDR) and domain generalization (DG) methods for comparison.

- **DeepFM [9]** is a factorization-machine based single-domain recommendation method for benchmark comparison.
- **Multi-View DNN [6]** is a multi-learning framework for cross-domain user modeling in recommendation, which maps two different views of the data into a shared view.
- **CoNet [10]** is the collaborative cross networks for cross-domain recommendation, which can enable dual knowledge transfer across domains by introducing cross connections.
- **DANN [8]** is the domain-adversarial neural network for domain generalization, which adversarially trains the generator and discriminator to capture invariant representations.
- **IRM [1]** is the invariant risk minimization for domain generalization, which enforces the optimal predictor to be the same across all domains by adding penalty loss terms.
- **Grace-Inv** is a lite version of Grace without DRM, which relies solely on invariant preference when inferring.

4.2 Performance Evaluation

To establish a performance benchmark for each domain, we first implemented DeepFM, a single-domain method. For out-of-distribution (OOD) prediction using DeepFM, the training set consists of all data from observed domains without specifying domain labels.

The comparison study in Table 1 indicates that Grace outperforms all the baseline methods under both CDR and DG settings.

4.2.1 Comparison with CDR. Although CDR baselines have achieved significant improvement by transferring knowledge between multiple domains, our proposed Grace can still outperform such methods by introducing preference disentanglement from a causal perspective. Notably, our approach provides significant improvement in domains with smaller sizes (*i.e.,* CD). Furthermore, the architecture of CDR baselines is specifically designed for observed domains and is limited in their ability to predict samples without domain labels.

4.2.2 Comparison with DG. It's obvious that our proposed Grace achieves better generalization performance on unseen domains by considering the potential similarity between unseen and observed domains. Based on our prior knowledge, the unseen domain in the public dataset is relatively similar to the observed domains (CD and Digital Music), while the unseen domain in the industry dataset is more different from the observed domains (serving different business purposes). Grace outperforms all the DG methods in both cases, which benefits from the dynamic control of the proposed DRM. Unsurprisingly, all DG methods cannot achieve excellent performance on observed domains because they sacrifice modeling domain-specific characteristics for the sake of cross-domain invariance and generalization.

4.2.3 Ablation Study. We also test the performance of Grace-Inv for ablation study. Grace-Inv removes DRM and relies solely on invariant preference when inferring. As we can see, Grace-Inv performs worse than Grace on both unseen and observed domains, which is consistent with the conclusion drawn from comparing Grace with DG methods. However, the performance of Grace-Inv is still better than that of other DG methods on observed domains and OOD evaluation, which results from the preference disentanglement from a causal perspective and the stronger modeling capacity provided by the MoE and gating networks.

5 CONCLUSION

In this paper, we proposed Grace, a novel framework that improves CDR performance and demonstrates excellent domain generalization capabilities. The framework disentangles preferences into domain-invariant and domain-specific preferences, allowing for the transfer of knowledge across domains while emphasizing unique domain preferences. Domain adversarial learning is employed to ensure both domain invariance and predictive power, and a novel domain routing module dynamically decides how domain-specific preferences affect predictions. The proposed method is evaluated through extensive experiments, including both observed and unseen domains, demonstrating its effectiveness over baseline approaches. Currently, Grace utilizes the most intuitive method (*i.e.,* multiple binary classifiers) to measure the potential similarity between observed and unseen domains. As future work, advanced open-set recognition techniques may serve as a better way.

REFERENCES

[1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893* (2019).

[2] Shlomo Berkovsky, Tsvi Kuflik, and Francesco Ricci. 2007. Cross-domain mediation in collaborative filtering. In *User Modeling 2007: 11th International Conference, UM 2007, Corfu, Greece, July 25-29, 2007. Proceedings 11*. Springer, 355–359.

[3] Gilles Blanchard, Aniket Anand Deshmukh, Ürun Dogan, Gyemin Lee, and Clayton Scott. 2021. Domain generalization by marginal transfer learning. *The Journal of Machine Learning Research* 22, 1 (2021), 46–100.

[4] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.

[5] Peng Cui and Susan Athey. 2022. Stable learning establishes some common ground between causal inference and machine learning. *Nature Machine Intelligence* 4, 2 (2022), 110–115.

[6] Ali Mamdouh Elkahky, Yang Song, and Xiaodong He. 2015. 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*. 278–288.

[7] Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*. PMLR, 1180–1189.

[8] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The journal of machine learning research* 17, 1 (2016), 2096–2030.

[9] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247* (2017).

[10] 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.

[11] Shoubo Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. 2020. Domain generalization via multidomain discriminant analysis. In *Uncertainty in Artificial Intelligence*. PMLR, 292–302.

[12] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.

[13] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. 2021. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*. PMLR, 5815–5826.

[14] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. 2018. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5400–5409.

[15] Siqing Li, Liuyi Yao, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Tonglei Guo, Bolin Ding, and Ji-Rong Wen. 2021. Debiasing learning based cross-domain recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3190–3199.

[16] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1930–1939.

[17] Divyat Mahajan, Shruti Tople, and Amit Sharma. 2021. Domain generalization using causal matching. In *International Conference on Machine Learning*. PMLR, 7313–7324.

[18] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. 2017. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE international conference on computer vision*. 5715–5725.

[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] Hongyan Tang, Junning Liu, Ming Zhao, and Xudong Gong. 2020. Progressive layered extraction (ple): A novel multi-task learning (mtl) model for personalized recommendations. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 269–278.

[21] Zimu Wang, Yue He, Jiashuo Liu, Wenchao Zou, Philip S Yu, and Peng Cui. 2022. Invariant Preference Learning for General Debiasing in Recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1969–1978.

[22] Tianzi Zang, Yanmin Zhu, Haobing Liu, Ruohan Zhang, and Jiadi Yu. 2022. A survey on cross-domain recommendation: taxonomies, methods, and future directions. *ACM Transactions on Information Systems* 41, 2 (2022), 1–39.

[23] Cheng Zhao, Chenliang Li, Rong Xiao, Hongbo Deng, and Aixin Sun. 2020. CATN: Cross-domain recommendation for cold-start users via aspect transfer network. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 229–238.