MACQUARIE
University

Disentangled and Invariant Cross-Domain Recommendation Systems

by

Jiajie Zhu

A thesis submitted in fulfilment of

the requirements for the award of the degree

Doctor of Philosophy

from

School of Computing

Faculty of Science and Engineering

MACQUARIE UNIVERSITY

Supervisor: Prof. Yan Wang

Associate Supervisor: Dr. Zhu Sun

Adjunct Supervisor: Dr. Feng Zhu

2025

Statement of Candidate

I certify that the work in this thesis entitled **"Disentangled and Invariant Cross-Domain Recommendation Systems"** has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree to any other university or institution other than Macquarie University.

I also certify that the thesis is an original piece of research and it has been written by me. Any help and assistance that I have received in my research work and the preparation of the thesis itself have been appropriately acknowledged.

In addition, I certify that all information sources and literature used are indicated in the thesis.

Jiajie Zhu

17 September 2025

To my parents and my wife,

who make me understand the true meaning of love.

Jiajie Zhu.

Abstract

To alleviate the data sparsity problem, cross-domain recommendation (CDR) aims to transfer valuable information from a relatively data-richer source domain to a relatively data-sparser target domain, forming the setting of *single-target CDR*. In contrast, *dual-target CDR* has been proposed in recent years to improve recommendation performance in both domains simultaneously by sharing the common knowledge across domains. However, existing methods still struggle to clearly distinguish different components of user preferences, making it difficult to determine what should or should not be transferred across domains. To effectively advance CDR, it is crucial not only to learn disentangled representations of user preferences and better understand which component of user preferences should be transferred across domains, but also to learn debiased user preferences that are robust to confounders, and generalizable user preferences that remain invariant across distribution shifts.

Firstly, existing dual-target CDR methods fail to recognize and decouple domain-independent user preferences from domain-shared and domain-specific user preferences, leading to incomplete user preference modeling and suboptimal recommendation performance. To address this problem, in this thesis, we propose a novel Disentanglement-based framework with Interpolative Data Augmentation for dual-target CDR, called DIDA-CDR. In DIDA-CDR, we first propose an interpolative data augmentation approach to generating both relevant and diverse augmented user representations to augment sparser domain and explore potential user preferences. We then propose a disentanglement module to effectively decouple domain-specific and domain-independent information to capture comprehensive user preferences, thereby improving the recommendation performance on both domains.

Secondly, existing dual-target CDR methods often overlook that, in addition to

users' true preferences, user-item interactions are also influenced by observed confounders (e.g., free shipping, sales promotion), including *single-domain confounders* and *cross-domain confounders*. However, the existing methods neither explicitly decouple cross-domain confounders, nor preserve the positive impacts of observed confounders on predicted interactions, while eliminating their negative impacts on capturing comprehensive user preferences. To address this problem, in this thesis, we propose a novel Causal Deconfounding framework via Confounder Disentanglement for dual-target CDR, called CD2CDR. In CD2CDR, we first propose a confounder disentanglement module to effectively decouple observed single-domain and cross-domain confounders. We then propose a causal deconfounding module to preserve the positive effects of such observed confounders and eliminate their negative effects via backdoor adjustment, thereby enhancing the recommendation accuracy in each domain.

Thirdly, existing CDR methods typically focus only on *cross-domain distribution shifts*, while assuming that training and testing data share the same distribution within target domain, overlooking potential *single-domain distribution shifts*. In reality, both types of distribution shifts co-exist, creating complex out-of-distribution (OOD) scenarios in CDR. This observation motivates us to propose a new setting of *cross-domain OOD recommendation*, which simultaneously addresses two types of distribution shifts through unified modeling. To this end, in this thesis, we propose a novel Causal-Invariant Cross-Domain Out-of-distribution Recommendation framework, called CICDOR. In CICDOR, we first propose a dual-level causal preference learning module that can effectively infer domain-specific and domain-shared causal-invariant user preferences, respectively. Then, we propose an LLM-guided confounder discovery module that seamlessly integrates LLMs with a conventional causal discovery method to extract observed confounders for effective deconfounding, thereby ensuring accurate causal-invariant user preference inference for OOD recommendation.

All three proposed methods in this thesis have been validated through theoretical analysis and extensive experiments on real-world datasets. The experimental results demonstrate that our methods significantly outperform state-of-the-art CDR methods.

Acknowledgments

Completing this PhD journey has been one of the most challenging and rewarding experiences of my life, and I am fortunate to have received support, encouragement, and guidance from many kind individuals along the way.

First, I am deeply grateful to my supervisor, Prof. Yan Wang, for his unwavering support, guidance, and understanding throughout my PhD journey. He has set an example for me with his rigorous academic attitude and constant pursuit of high-quality research, which has deeply influenced my approach to conducting research and helped shape my future academic path. Under his supervision, I have learned how to stay focused in research, identify novel problems, refine my work through multiple iterations, and communicate ideas in a clear and accessible way. I am especially thankful for his tireless effort in reviewing and revising my manuscripts with great care and timely feedback, which significantly improved the quality of my work. In addition, Prof. Wang has also offered personal support in my daily life whenever possible. He often reminded me to take care of my health and encouraged me to exercise regularly. His thoughtful guidance and care have brought me warmth and strength during my time studying abroad alone.

Second, I would also like to express my sincere gratitude to my Associate Supervisor, Dr. Zhu Sun, and my Adjunct Supervisor, Dr. Feng Zhu, for their academic support throughout my PhD journey. Although based in Singapore, Dr. Zhu Sun generously offered valuable feedback on my paper submissions and revisions. I truly admire her dedication and meticulous approach, which helped improve my work in terms of structure, presentation, and clarity. Furthermore, despite Dr. Feng Zhu's busy schedule in China, he always took the time to review the technical aspects of my research and provide constructive suggestions, especially from an industry perspective.

His insights played a crucial role in the successful completion of my PhD thesis.

Third, I am also sincerely grateful to my colleagues, Dr. Shoujing Wang, Dr. Pengfei Ding, Dr. Yatong Sun, Xuankai Yang, Han Zhang, Hongyang Liu, Rongwei Xu, and Nan Wang, for their support and companionship during my PhD journey. Dr. Shoujing Wang encouraged me to regain confidence and persevere during the early and difficult stages of my PhD, which played an important role in helping me stay motivated and ultimately complete my thesis. Dr. Pengfei Ding generously shared his experience and offered valuable academic suggestions, deepening my understanding of research and serving as a role model throughout my PhD. Dr. Yatong Sun impressed me with his dedication and calm focus on academic work, and his support was crucial to the completion of my thesis. Moreover, I would also like to thank the other colleagues for their support in both my academic journey and daily life. Their encouragement and assistance were essential to the successful completion of my thesis.

Last but not least, I would like to express my heartfelt thanks to my parents. Their unconditional love, selfless support, and unwavering encouragement gave me the strength to face and overcome the many challenges of studying and living abroad. I am also deeply grateful to my wife, whose understanding, patience, and support have been a constant source of motivation during difficult times. Their love and encouragement have been the foundation that enabled me to complete this PhD thesis.

Publications

This thesis is based on the research work I have performed with the help of my supervisors and other colleagues during my PhD program at the School of Computing, Macquarie University between 2021 and 2025. Some parts of my research have been published/accepted/submitted in the following papers:

[1] **Jiajie Zhu**, Yan Wang, Feng Zhu, Pengfei Ding, Hongyang Liu, Zhu Sun: Causal-Invariant Cross-Domain Out-of-Distribution Recommendation, ACM Transactions on Information Systems (TOIS), submitted in May 2025. (**CORE2020**[1] **Rank A**)

[2] **Jiajie Zhu**, Yan Wang, Feng Zhu, Zhu Sun: Causal Deconfounding via Confounder Disentanglement for Dual-Target Cross-Domain Recommendation, ACM Transactions on Information Systems (TOIS), 43.5 (2025): 1-33. (**CORE2020 Rank A**)

[3] **Jiajie Zhu**, Yan Wang, Feng Zhu, Zhu Sun: Domain Disentanglement with Interpolative Data Augmentation for Dual-Target Cross-Domain Recommendation, The 17th ACM Conference on Recommender Systems (RecSys 2023), pages 515 - 527. (**main track, acceptance rate 18.7%, CORE2023 Rank A**)

[1] CORE stands for the Computing Research and Education Association of Australasia (`http://www.core.edu.au`).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and Significance

With the rapid growth of digital content across various online platforms, such as Amazon, YouTube and Facebook, users face considerable difficulty in discovering relevant information that matches their preferences [152]. To address this issue, recommender systems (RSs) have been developed as an effective solution by analyzing user interaction data to predict user preferences and deliver personalized recommendations [186]. The main goal of RSs is to identify items that align with individual user preferences, thereby improving user satisfaction and reducing information search time. Despite their widespread success, conventional single-domain recommendation approaches suffer from a significant limitation, i.e., the data sparsity problem, where insufficient user-item interactions limit RSs' ability to accurately capture user preferences.

To alleviate the data sparsity problem, Cross-Domain Recommendation (CDR) [82] aims to employ abundant information from a relatively richer domain to improve recommendation performance on a sparser domain, forming the so-called *single-target CDR* [183]. Effective CDR requires the source and target domains to share certain relatedness while maintaining distinctions in user intents, user behaviors, or item categories [84]. For instance, on e-commerce platforms like Taobao[1], different purchase scenarios, such as 'what to take when travelling' and 'how to dress up for a party', share overlapping user interests (e.g., finding suitable clothing or accessories), while

[1]Taobao is a Chinese Customer-to-Customer (C2C) platform that facilitates transactions between individual sellers and buyers, similar to eBay.

maintaining distinctions in user intents, such as practicality for travel and aesthetics for parties [22, 170]. Similarly, on Tmall[2], behaviors such as 'add to favorite' and 'purchase' can be regarded as business domains, both reflecting user interest but differing in their focus as exploration- and purchase-oriented actions, respectively [63, 167]. In contrast, on Amazon, domains often refer to different item categories (e.g., 'movie' and 'book') that share user interests in certain features (e.g., genres or styles), while differing in domain-specific item features and user preferences [14, 105]. These relatedness and distinctions together define the principle of domains, thereby ensuring CDR's adaptability to diverse recommendation scenarios.

Moreover, on top of the same above-mentioned principle of domains, Dual-Target CDR [186] has been proposed to capture comprehensive user preferences, and thus enhance the recommendation accuracy in both data-richer and data-sparser domains simultaneously, which are source domains and target domains as well. The existing dual-target CDR methods can be divided into two groups, i.e., (1) conventional methods, (2) disentanglement-based methods. *Conventional methods* mainly utilize various transfer layers [44, 59, 172, 74] to integrate the representations learned by two base encoders in their respective domains. In contrast, *disentanglement-based methods* tend to use the variational autoencoder (VAE) [8] or other disentangling techniques [34, 165, 9] to decouple the domain-shared and domain-specific information, and only transfer the domain-shared information to each domain, which enhances the recommendation accuracy on both domains simultaneously.

The effectiveness of both single-target and dual-target CDR methods fundamentally depends on the quality of user preference modeling. High-quality user preference modeling is crucial not only for making accurate recommendations within individual domains, but also for enabling effective knowledge transfer across domains. Despite growing research efforts in this field, existing CDR methods still suffer from three critical limitations in modeling user preferences. These limitations are discussed in the

[2]Tmall is a Chinese Business-to-Customer (B2C) platform designed for brand merchants and flagship stores, similar to Amazon.

following three subsections.

1.1.1 From Incomplete to Comprehensive User Preference Modeling

Limitation 1: None of the existing CDR methods decouple all three essential components needed to capture comprehensive user preferences. To be specific, the existing CDR methods only decouple two essential components, i.e., domain-shared and domain-specific information, and ignore the existence of domain-independent information. For example, in the movie and music domains, the domain-shared information (e.g., 'Category') extracted from two domains expresses the same meaning in each domain. In addition, the domain-specific information, e.g., 'Frame' in the movie domain, exists only in its own domain. By contrast, although the domain-independent information, e.g., 'Rhythm', exists in each domain, it has different meanings in different domains. Such oversight of domain-independent information results in incomplete preference modeling, leading to suboptimal recommendation results in both domains.

1.1.2 From Biased to Debiased User Preference Modeling

Even with well-disentangled components for capturing comprehensive user preferences, the user-item interactions might also be affected by confounding factors. A confounding factor, termed as *confounder* in causal inference [27, 192], affects both the treatment and the outcome [138, 83], which can be broadly interpreted as user preference and user-item interaction respectively (see Figure 1.1(a)) in the context of RSs [32, 191].

In the dual-target CDR scenario, observed confounders can be divided into two types, i.e., *single-domain confounder* (SDC) and *cross-domain confounder* (CDC). SDC only affects user preference and user-item interaction in one specific domain and has been widely studied in the existing literature [166, 122]. By contrast, CDC affects both domains, which, however, has been overlooked in existing dual-target CDR meth-

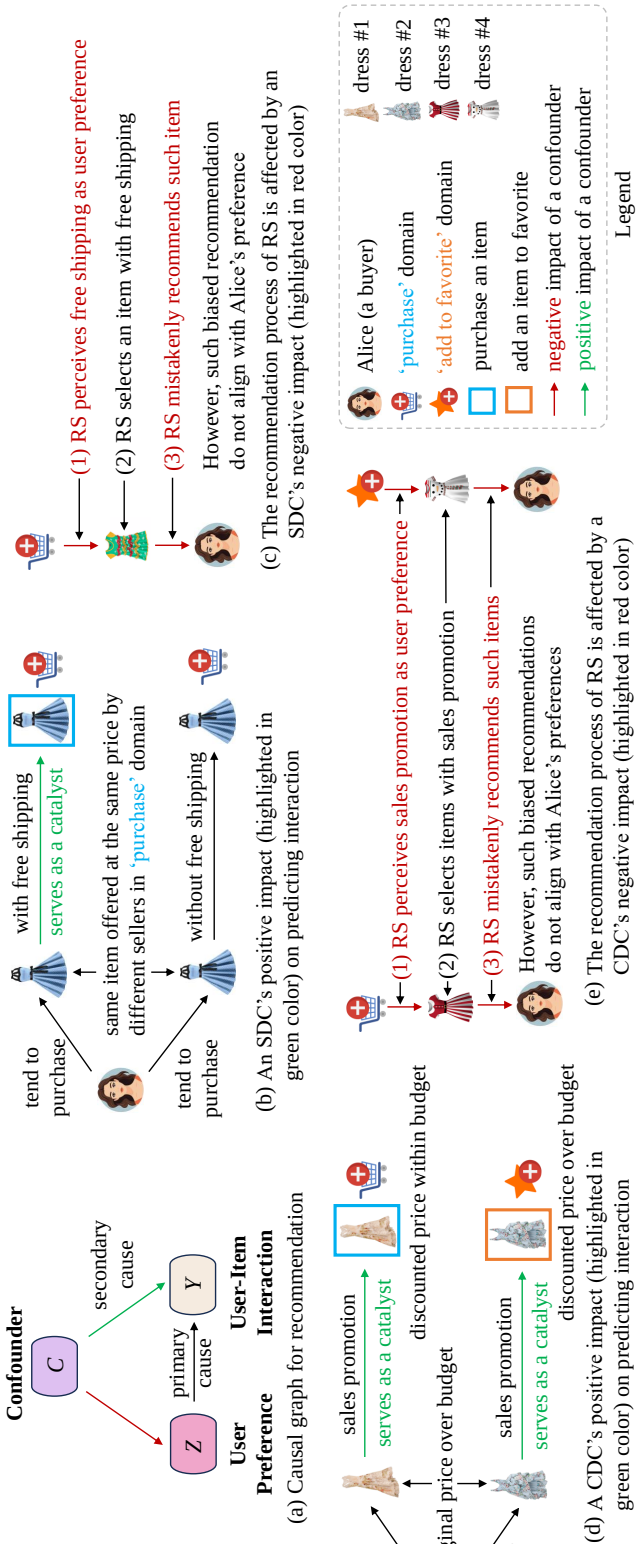

Figure 1.1: Examples to depict single-domain confounder (SDC) and cross-domain confounder (CDC) [20].

ods. Essentially, SDC is a simplified version of CDC. Below we first briefly review SDCs and then provide an in-depth analysis of CDCs, both illustrated with examples from Tmall, where 'purchase' and 'add to favorite' are regarded as two domains, as they align with the principle of domains.

SDCs have both positive and negative impacts on predicting user-item interactions in their corresponding domain. For instance, as shown in Figures 1.1(b)-(c), 'free shipping' is an SDC in the 'purchase' domain. Consider a scenario where the same item is offered at the same price by different sellers. One seller provides free shipping, while the other provides shipping with an additional cost. Thus, the offer with free shipping positively influences Alice's decision to purchase the item with free shipping. As for negative impact, a data-driven RS improperly perceives 'free shipping' (i.e., an SDC shown in Figure 1.1(c)) as Alice's preference in the 'purchase' domain. As a result, the data-driven RS mistakenly recommends an item with free shipping that Alice does not actually like. This misalignment, referred to as *confounding bias* [62], results in biased recommendations

In fact, the confounding biases also exist in cross-domain scenarios. More importantly, CDCs have both positive and negative impacts on predicting user-item interactions in both domains. For example, as illustrated in Figures 1.1(d)-(e), 'sales promotion' is a CDC, because it simultaneously affects 'purchase' and 'add to favorite' domains. On the one hand, this 'sales promotion' CDC has a positive impact. In fact, while Alice's true preference is the primary cause of her behaviors in both domains, 'sales promotion' is a secondary cause that serves as a catalyst. With a new sales promotion on dresses #1 and #2, both of which Alice likes but previously found over her budget, she immediately purchases dress #1 that has become affordable within her budget. By contrast, since the discounted price of dress #2 is still over her budget, she adds it to favorite for future consideration, looking forward to a further price reduction. On the other hand, this 'sales promotion' CDC has a negative impact too. As depicted in Figure 1.1(e), a data-driven RS improperly perceives 'sales promotion' (i.e., a CDC) as Alice's preference in both domains. As a result, the data-driven RS

mistakenly recommends dresses #3 and #4 with sales promotion to Alice, but Alice actually does not like them.

As illustrated by the above examples, both types of observed confounders significantly distort the recommendation process by introducing confounding biases in user preference modeling. These observations reveal a significant limitation as follows.

Limitation 2: Existing dual-target CDR methods fail to account for the influence of observed confounders, especially CDCs that have been entirely overlooked. This oversight results in confounding biases, thereby leading to biased user preference modeling and suboptimal recommendation results in both domains. Addressing these confounding biases is essential for capturing debiased comprehensive user preferences.

1.1.3 From IID to OOD User Preference Modeling

Even if user preferences are modeled comprehensively and without bias, they may still fail to generalize well in dynamic real-world environments where data distributions vary across and within domains. Most existing CDR methods adopt the independent and identically distributed (IID) assumption, positing that training and testing data share the same distribution in the target domain [130]. This assumption simplifies the process of problem formulation and theoretical reasoning, enabling the model to better generalize knowledge learned from training data to testing data within the target domain [65].

However, this assumption often fails to hold in real-world CDR scenarios. In CDR, recommendation performance is affected by two types of distribution shifts in terms of user preferences, interaction patterns, and item features [167]: *cross-domain distribution shift* (CDDS) and *single-domain distribution shift* (SDDS). (i) CDDS results from data distribution differences between source and target domains [63]. (ii) SDDS results from data distribution differences due to temporal and regional variations in target domain [43], leading to mismatches between training and testing data distributions.

Figure 1.2 illustrates these two types of distribution shifts in CDR. Consider an on-

Figure 1.2: An illustrative example of the co-existing distribution shifts (i.e., CDDS and SDDS) in CDR.

line platform that provides reviews for both movie and book domains. To be specific, in this scenario:

• **CDDS:** The inherent differences between domains lead to CDDS. In the movie domain, users often seek immediate gratification through visually engaging Hollywood blockbusters (e.g., *Avengers* and *Avatar*), while the book domain offers a slower, more immersive experience, catering to users who value delayed gratification.

• **SDDS:** In the book domain, regional variations in user preferences result in SDDS. Users in Hong Kong, which is a highly internationalized and modern metropolitan city, tend to reflect openness to international cultural influences and favor internationally renowned literature (e.g., international bestsellers) such as *Harry Potter* and *The Lord of the Rings*. In contrast, users in Beijing, known as China's cultural and political center, tend to prefer Chinese classics like *Journey to the West* and *Three Kingdoms*.

In such a scenario, reliable book recommendations for users in Beijing require: (1) Mitigating CDDS by transferring knowledge from movie domain to book domain, and (2) Mitigating SDDS by generalizing knowledge within book domain from users

in Hong Kong to users in Beijing. While existing approaches focus on handling either CDDS or SDDS in isolation, we observe that beyond learning invariant knowledge under SDDS within book domain, there exists a knowledge transfer path that can capture invariant knowledge under both CDDS and SDDS (illustrated by the red arrow). For example, preferences for heroic adventure stories remain invariant under both CDDS and SDDS. Such invariant preferences learned from movie domain training data (*The Avengers* and *Avatar*) can be transferred to book domain to help make recommendation (e.g., *Journey to the West*) for users in Beijing, as these works share similar narrative patterns of heroic quests and supernatural adventures. This invariant knowledge transfer path helps improve book domain recommendation performance by avoiding the potential information loss and error accumulation that could occur when sequentially applying existing approaches that handle CDDS and SDDS in isolation. The co-existence of these distribution shifts leads to complex OOD environments in CDR [18]. These observations reveal a significant limitation as follows.

Limitation 3: Existing CDR methods lack the capability to simultaneously address these co-existing distribution shifts. This limitation significantly restricts the effectiveness of existing methods in cross-domain OOD recommendation. Learning user preferences that remain invariant across these distribution shifts is essential for improving recommendation performance under OOD environments in CDR.

To address the abovementioned limitations, this thesis focuses on the four significant challenges detailed in the following section.

1.2 Challenges in Cross-Domain Recommendation

1.2.1 Disentangling all Essential Components of User Preferences

To address **Limitation 1** regarding incomplete user preference modeling, the first challenge of this thesis, **CH1**, can be represented by the following question: *How to effectively decouple domain-independent information from domain-specific information, in*

addition to domain-shared information, to capture comprehensive user preferences on each domain, thereby improving the recommendation performance?

The existing dual-target CDR methods either ignore decoupling the domain-specific and domain-shared information [44] (*Group 1*), or directly transfer the domain-shared information to only fuse with the domain-specific information in each domain, overlooking the domain-independent information [8] (*Group 2*). The methods in *Group 1* disregard the discrimination between domain-specific and domain-shared information, which may lead to the negative transfer. The methods in *Group 2* do not differentiate the domain-independent information from domain-specific information and decouple these two types of information, which results in suboptimal recommendation results.

1.2.2 Accurately Extracting Observed Confounders

As the first step in addressing **Limitation 2** regarding the influence of observed confounders, the second challenge of this thesis, **CH2**, can be represented by the following question: *How to effectively extract observed confounders to comprehensively understand user-item interactions?*

Even with well-disentangled components for capturing comprehensive user preferences, user-item interactions might also be affected by observed confounders. However, existing methods face considerable difficulties in effectively extracting these confounders. Existing methods for extracting observed confounders primarily fall into two categories based on their data sources: (i) representation-based extraction from learned embeddings, and (ii) content-based extraction from textual information.

Representation-based extraction from learned embeddings: Existing methods either employ graph clustering strategy [64] and variational information bottleneck [9], or identify unobserved domain-specific confounders first, and then utilize causal techniques, e.g., inverse propensity score (IPS) estimators [63] and invariant learning [167], to obtain debiased representations. However, none of them explicitly decouples observed CDCs, and thus it is hard to obtain a comprehensive understanding of

user-item interactions in each domain.

Content-based extraction from textual information: Recent studies have demonstrated promising results in utilizing large language models (LLMs) to extract causal variables from unstructured text [12, 68]. While these successes suggest the potential of extracting observed confounders from user reviews, LLMs may extract inaccurate confounders without proper theoretical guidance [1]. Thus, effectively extracting accurate observed confounders from user reviews has not been fully explored.

1.2.3 Effectively Deconfounding Observed Confounders

To further address **Limitation 2**, the third challenge of this thesis, **CH3**, can be represented by the following question: *How to preserve the positive impacts of observed confounders on predicted interactions, while eliminating their negative impacts on capturing comprehensive user preferences, thereby enhancing the recommendation accuracy?* Most existing causal methods [136, 146] tend to eliminate the confounders' negative impacts, in order to obtain the debiased comprehensive user preferences for recommendation. However, most of them overlook the confounders' positive impacts, and thus limit their efficacy in enhancing the recommendation accuracy [154].

1.2.4 Effectively Addressing Co-Existing Distribution Shifts

To address **Limitation 3** regarding the inability to handle co-existing distribution shifts, the fourth challenge of this thesis, **CH4**, can be represented by the following question: *How to simultaneously address cross-domain and single-domain distribution shifts to achieve reliable recommendation under OOD environments in CDR?*

For CDDS, existing approaches [33, 108] mainly align shared attributes, user behaviors, or auxiliary data between domains to mitigate distribution shifts [164]. In contrast, to address SDDS, existing approaches either formulate such distribution shift as a debiasing problem, particularly targeting specific biases such as conformity bias [181] or popularity bias [178], or leverage causal discovery to jointly model the causal

structure and invariant user preferences [43]. However, most existing approaches are constrained to handling SDDS or CDDS in isolation, failing to handle two types of distribution shifts simultaneously, thereby leading to degraded recommendation performance in CDR.

Moreover, effectively addressing the co-existing distribution shifts requires leveraging the invariance principle [6, 92], which focuses on identifying invariant relationships across distributions, particularly under OOD environments. Recent studies [43, 65] have proven that such invariant relationships can be effectively represented through a causal structure between user attributes and preferences, as this structure remains invariant across distribution shifts. However, this research direction has not been fully explored.

1.3 Contributions of the Thesis

To address the aforementioned significant challenges in CDR, this thesis proposes a series of interconnected technical frameworks that systematically evolve from incomplete to comprehensive user preference modeling, from biased to debiased user preference modeling, and ultimately from IID to OOD user preference modeling. This thesis has three main contributions: (1) a disentanglement-based CDR framework that, for the first time, recognizes and then decouples domain-independent information in addition to domain-shared and domain-specific information, (2) a causal deconfounding CDR framework that mitigates confounding effects from observed confounders while preserving their beneficial impacts, and (3) a causal-invariant CDR framework that maintains recommendation reliability under complex OOD environments. Together, these contributions establish a robust foundation for reliable CDR systems that can handle complex real-world scenarios involving incomplete user preference modeling, confounding effects, and distribution shifts. To better illustrate how these contributions are structured and interconnected throughout this thesis, Figure 1.3 provides a visual overview of the logical relationships among the three technical chapters, each

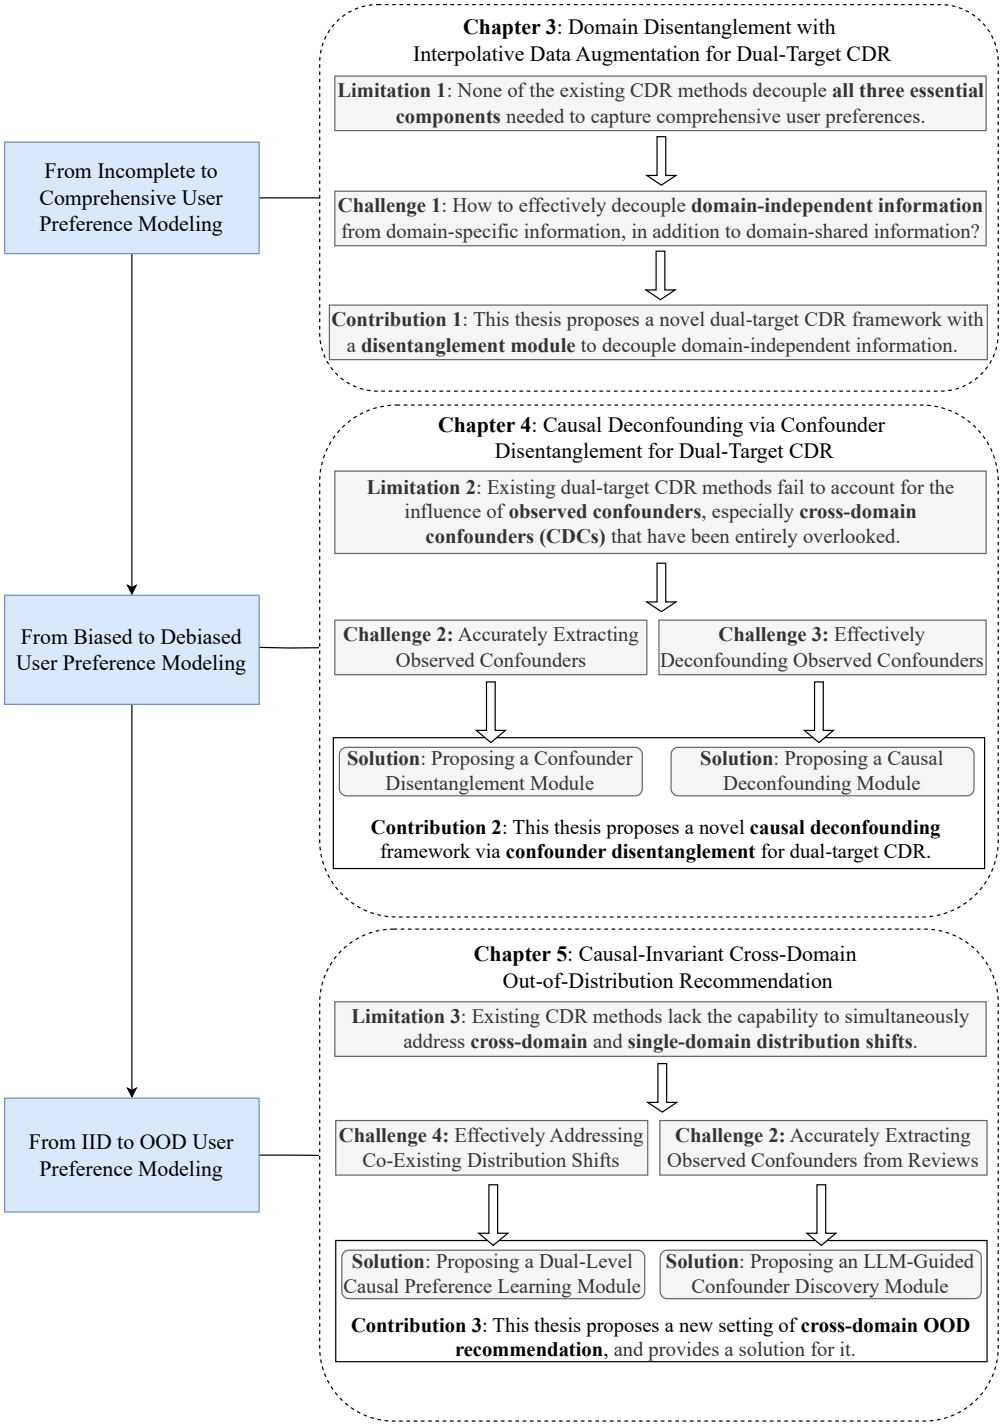

Figure 1.3: The logical relationships among the three technical chapters.

corresponding to one of the main contributions. The three primary contributions are detailed as follows:

1. **Contribution 1:** The first contribution of the thesis is to propose a novel dual-target CDR framework. To the best of our knowledge, this is the first work in the literature that explicitly takes domain-independent information into consideration in addition to domain-shared and domain-specific information, and decouples it to capture more comprehensive user preferences for CDR. The characteristics and contributions of our proposed model are summarized as follows:

   (a) This thesis proposes a **D**isentanglement-based framework with **I**nterpolative **D**ata **A**ugmentation for dual-target **C**ross-**D**omain **R**ecommendation, called DIDA-CDR, which can augment the sparser domain, disentangle three essential components of user preferences and transfer the domain-shared user preferences of common users across domains, thus enhancing the recommendation accuracy on both domains simultaneously.

   (b) This thesis proposes an interpolative data augmentation approach to generating both relevant and diverse augmented user representations, which augments the sparser domain and explores the potential common user preferences and thus improves recommendation performance on both domains.

   (c) Targeting **CH1**, this thesis proposes a disentanglement module to effectively decouple the domain-independent and domain-specific user preferences. The disentanglement module also extracts the domain-shared user preferences from augmented user representations, which can be transferred to both domains to provide the valuable information. The attention mechanism is then applied to combine the above three essential components of user preferences to capture more comprehensive user preferences in each domain, which can improve the recommendation performance on each of both domains.

(d) Extensive experiments conducted on five real-world datasets show that the proposed DIDA-CDR outperforms the best-performing state-of-the-art baselines with an average improvement of 8.54% and 11.10% with respect to HR@10 and NDCG@10, respectively.

2. **Contribution 2:** The second contribution of the thesis is to propose a novel causal deconfounding framework via confounder disentanglement for dual-target CDR. To the best of our knowledge, this is the first work in the literature that explicitly decouples observed CDCs, and incorporates observed confounders' positive impacts into debiased comprehensive user preferences for dual-target CDR. The characteristics and contributions of our framework can be summarized as follows:

   (a) This thesis proposes a **C**ausal **D**econfounding framework via **C**onfounder **D**isentanglement for dual-target **C**ross-**D**omain **R**ecommendation, called CD2CDR, which can disentangle two types of observed confounders (i.e., SDCs and CDCs), eliminate their negative impacts to obtain debiased preferences, and preserve such confounders' positive impacts, thereby enhancing the recommendation accuracy in both domains.

   (b) Targeting **CH2**, this thesis proposes a confounder disentanglement module to effectively disentangle observed SDCs and CDCs. In this module, a dual adversarial structure is devised to disentangle SDCs in each domain and apply half-sibling regression to decouple CDCs, thus obtaining a comprehensive understanding of user-item interactions in both domains.

   (c) Targeting **CH3**, this thesis proposes a causal deconfounding module to deconfound disentangled observed SDCs and CDCs via backdoor adjustment. Specifically, a confounder selection function is designed to mitigate such observed confounders' negative effects, thereby recovering debiased comprehensive user preferences. The observed confounders' positive ef-

fects are incorporated into such debiased user preferences to enhance the recommendation accuracy in both domains.

(d) Extensive experiments conducted on seven real-world datasets demonstrate that the CD2CDR outperforms the best-performing state-of-the-art baseline with an average increase of 6.17% and 8.23% w.r.t. HR@10 and NDCG@10, respectively.

3. **Contribution 3:** The third contribution of the thesis is to propose a new setting of *cross-domain OOD recommendation*, which simultaneously addresses two types of distribution shifts through unified modeling. To provide a solution, this thesis proposes a novel **C**ausal-**I**nvariant **C**ross-**D**omain **O**ut-of-distribution **R**ecommendation framework, called CICDOR. The characteristics and contributions of our proposed model are summarized as follows:

   (a) This thesis proposes CICDOR, a novel framework that discovers invariant causal structures across distributions to tackle both CDDS and SDDS for reliable cross-domain OOD recommendation.

   (b) Targeting **CH4**, this thesis proposes a dual-level causal preference learning module. This module first leverages a user preference disentanglement module to extract domain-specific and domain-shared user preferences. In addition, dual-level causal structures, represented as Directed Acyclic Graphs (DAGs), are learned for both domain-specific and domain-shared levels of modeling. Based on these causal structures, the corresponding causal-invariant user preferences are inferred at each level, thus tackling CDDS and SDDS simultaneously under OOD environments in CDR.

   (c) Targeting **CH2**, this thesis proposes an LLM-guided confounder discovery module. This module first employs an LLM to extract candidate interaction-related causal variables from user reviews and transform them into structured data. The structured data are then fed into a conventional causal dis-

covery method (i.e., the FCI algorithm) to uncover their underlying causal relationships, with conditional independence tests used to eliminate redundant variables. The observed confounders are first identified by leveraging the LLM and then stored in the confounder pool. Next, when such remaining variables cannot fully explain the user-item interactions, feedback is constructed through the LLM to guide the discovery of additional causal variables, forming an iterative process that continuously refines the confounder discovery.

(d) Extensive experiments conducted on two real-world datasets demonstrate that CICDOR outperforms the best-performing state-of-the-art baseline model across various OOD settings with an average increase of 6.28% and 9.42% w.r.t. HR@10 and NDCG@10, respectively.

1.4 Roadmap of the Thesis

The remainder of this thesis is organized as follows.

Chapter 2 presents a comprehensive literature review organized in two parts. The first part focuses on research problems, including CDR methods (both single-target and dual-target), and single-domain OOD recommendation methods, which provide insights for the novel setting of cross-domain OOD recommendation. The second part reviews three key techniques employed in this thesis: disentangled representation learning, causal deconfounding, and causal structure learning.

Chapter 3 presents a novel disentanglement-based framework with interpolative data augmentation for dual-target CDR, called DIDA-CDR, which can augment the sparser domain, disentangle three essential components of user preferences and transfer the domain-shared user preferences of common users across domains, thus enhancing the recommendation accuracy on both domains simultaneously. This chapter includes our paper published at RecSys 2023 [188].

Chapter 4 presents a novel causal deconfounding framework via confounder dis-

entanglement for dual-target CDR, called CD2CDR, which can disentangle two types of observed confounders (i.e., SDCs and CDCs), eliminate their negative impacts to obtain debiased preferences, and preserve such confounders' positive impacts, thereby enhancing the recommendation accuracy in both domains. This chapter includes our paper published by ACM TOIS [189] in 2025.

Chapter 5 presents a novel causal-invariant cross-domain out-of-distribution recommendation framework, called CICDOR, which can discover invariant causal structures across distributions to tackle CDDS and SDDS for reliable cross-domain OOD recommendation. This chapter includes our paper submitted to ACM TOIS in 2025.

Finally, Chapter 6 concludes the thesis and highlights promising directions for future research.

The thesis overview is illustrated in Figure 1.4.

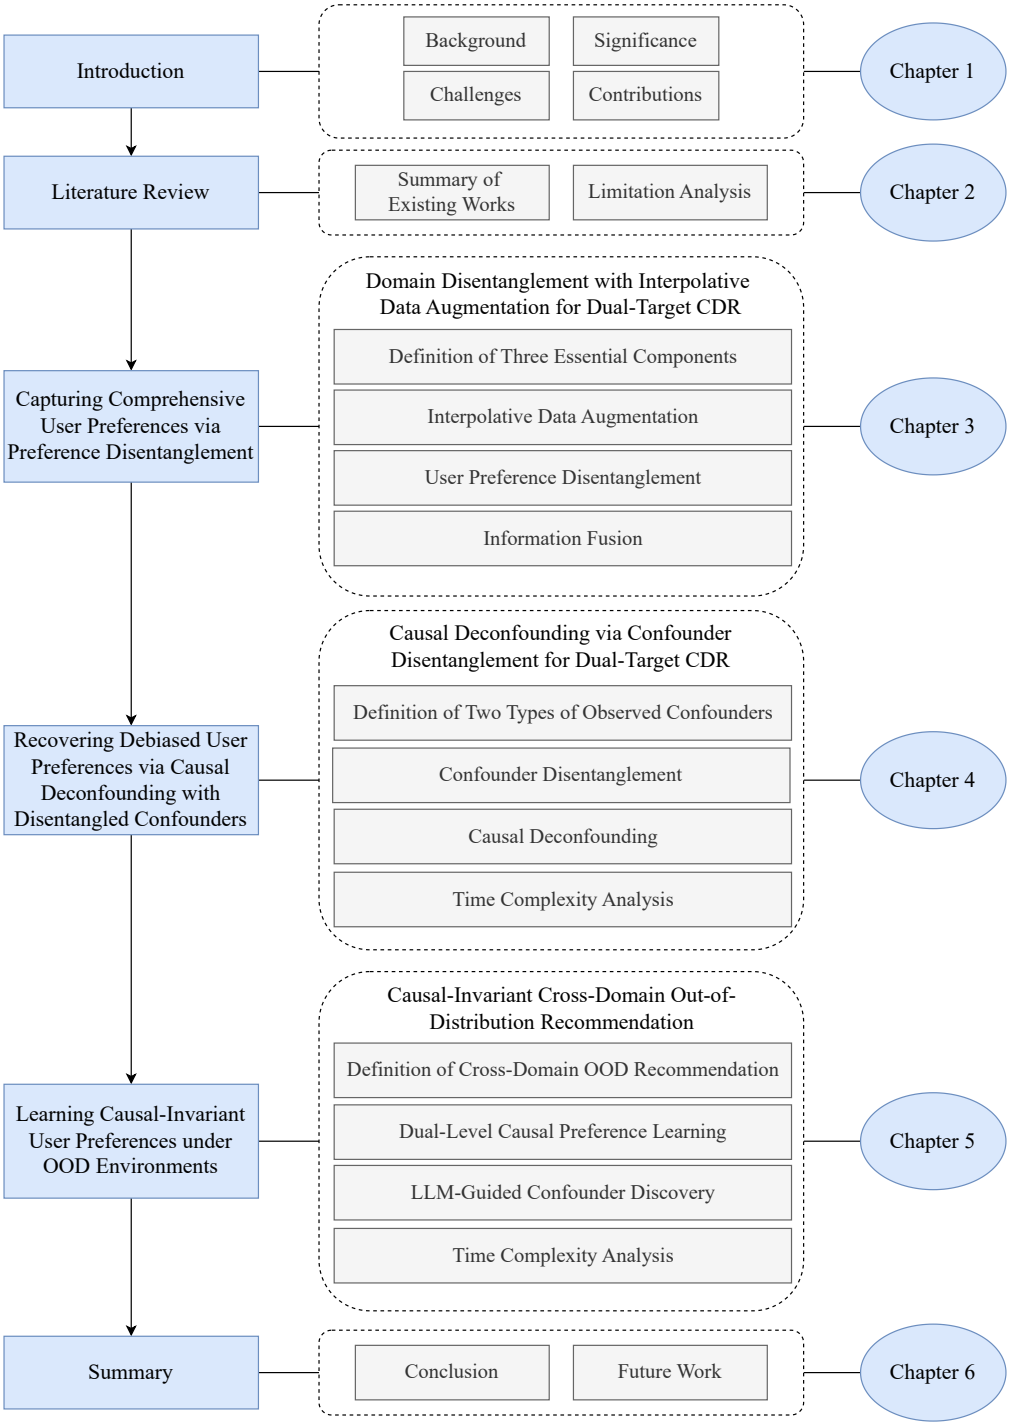

Figure 1.4: The overview of this thesis.

Chapter 2

Literature Review

In this chapter, we organize the literature review based on the core research problem of cross-domain recommendation (CDR) and its extension to a novel cross-domain out-of-distribution (OOD) recommendation problem. Section 2.1 reviews both single-target CDR and dual-target CDR methods. To the best of our knowledge, there is no existing method that explicitly addresses the cross-domain OOD recommendation problem. Thus, Section 2.2 reviews single-domain OOD recommendation methods to provide theoretical insights into addressing distribution shifts between training and testing data in recommender systems (RSs).

In addition to the above research problems, we also review three key techniques employed in this thesis. Specifically, we review: (1) disentangled representation learning in Section 2.3, which helps to decouple essential components of user preferences and observed confounders; (2) causal deconfounding in Section 2.4, which mitigates confounding bias from observed confounders in observational data; and (3) causal structure learning in Section 2.5, which enables the discovery of invariant causal relationships among variables from observational data.

This chapter is organized as follows:

- Section 2.1 introduces existing CDR methods, including single-target CDR and dual-target CDR methods.
- Section 2.2 introduces existing single-domain OOD recommendation methods.
- Section 2.3 introduces related disentangled representation learning methods.

- Section 2.4 introduces related causal deconfounding methods.
- Section 2.5 introduces related causal structure learning methods.
- Section 2.6 presents a summary of existing studies.

2.1 Cross-Domain Recommendation (CDR)

As the core research problem of this thesis, CDR has been widely studied under different knowledge transfer paradigms. This section reviews both single-target CDR and dual-target CDR methods.

2.1.1 Single-Target CDR

CDR aims to transfer valuable information from a relatively data-richer source domain to a relatively data-sparser target domain to improve recommendation performance, forming single-target CDR [183]. The existing single-target CDR methods can be divided into three categories, i.e., (1) content-based transfer methods, (2) embedding-based transfer methods and (3) rating pattern-based transfer methods [186].

2.1.1.1 Content-Based Transfer

Content-based transfer methods [49] mainly use the various content information, such as user/item attributes, tags, reviews, etc., to link domains and share their information across domains. TagCDCTR [120] extends collaborative topic regression to the cross-domain setting by leveraging shared tags as semantic bridges between domains, enabling joint topic learning and collective matrix factorization (MF) to enhance recommendation performance under data sparsity. By learning semantic relationships among non-identical tags through word embeddings and clustering, SCT [158] constructs a shared feature space that enables knowledge transfer across domains via semantically aligned user and item representations. The method proposed by Kanagawa

et al. [49] frames CDR for cold-start users as an extreme multi-class classification task and address it through unsupervised domain adaptation using a Domain Separation Network (DSN) combined with a denoising autoencoder for item representation.

2.1.1.2 Embedding-Based Transfer

By contrast, embedding-based transfer methods [45, 23] aim to first get embeddings using various learning techniques, and then transfer them across domains. EMCDR [87] utilizes Bayesian Personalized Ranking (BPR) model [98] as its MF model and maps the latent factors of common users/items across different domains for effective knowledge transfer. DCDCSR [184] integrates the latent factors from both domains by considering the sparsity degrees of individual users/items in each domain, creating more accurate benchmark factors to guide the deep neural network to map the latent factors across domains. PTUPCDR [194] generates personalized bridge functions through a task-optimized meta network, facilitating stable and personalized preference transfer from the source to the target domain. CUT [58] employs a two-phase training strategy that first captures user similarities in the target domain and then transfers source-domain information selectively, leveraging a user transformation module and contrastive learning to avoid relationship distortion.

2.1.1.3 Rating Pattern-Based Transfer

Different from the embedding-based transfer methods, rating pattern-based transfer methods [150] typically focus on capturing generalizable user rating patterns in the source domain, which are then adapted to the target domain to enhance recommendation accuracy. CLFM [28] introduces a cluster-level latent factor model based on joint nonnegative matrix tri-factorization to capture both shared and domain-specific rating patterns, enabling flexible control over cross-domain knowledge transfer and improving recommendation performance through subspace learning. MINDTL [38] addresses data sparsity in CDR by extracting transferable rating patterns from mul-

tiple incomplete source domains using an incomplete orthogonal nonnegative matrix tri-factorization method, and automatically learns source weights to mitigate negative transfer. DARec [150] employs deep domain adaptation to transfer user rating patterns across domains without relying on auxiliary information, using adversarial training and domain classifiers to learn shared and distinct representations from rating matrices alone.

2.1.2 Dual-Target CDR

In contrast to single-target CDR, dual-target CDR aims to achieve better recommendation performance on both domains simultaneously, which can be extended to multiple domains, leading to multi-target CDR [16, 185, 34]. The existing dual-target CDR approaches can be divided into four categories, i.e., (1) conventional approaches, (2) disentanglement-based approaches, (3) causal-based approaches, and (4) alignment-based approaches.

2.1.2.1 Conventional Dual-Target CDR

Conventional approaches primarily utilize two base encoders to transform each domain's interaction data into embeddings, which are then symmetrically incorporated through various transfer layers [59, 74]. CoNet [44] utilizes cross-connection networks to achieve knowledge transfer between two domains. DDTCDR [59] employs a latent orthogonal mapping function to transfer user embeddings across domains. PPGN [172] constructs a cross-domain interaction graph to learn and transfer representations. BiTGCF [74] first leverages two base graph encoders to learn user/item embeddings, and then performs the feature propagation and transfer to fuse user embeddings.

2.1.2.2 Disentanglement-Based Dual-Target CDR

By contrast, disentanglement-based approaches mainly utilize techniques such as variational autoencoders (VAEs) [8, 188] or introduce supervision signals such as adver-

sarial learning [117, 109] and contrastive learning [7] to decouple domain-invariant (shared) user preferences, which remain stable across domains and capture fundamental behavioral patterns. For instance, VAE-based approaches [126] employ the explicit reconstruction loss included in the Evidence Lower Bound (ELBO) and extra regularizers as the disentanglement loss to learn desirable disentangled representations. DisenCDR [8] uses two mutual-information-based regularizers to decouple domain-shared and domain-specific information, transferring only domain-shared information across domains to improve recommendation performance. DIDA-CDR [188] uniquely decouples domain-independent user preferences, as well as domain-shared and domain-specific user preferences, to capture more comprehensive user preferences for recommendation. Other recent works often utilize Graph Convolutional Networks (GCNs) [121], adversarial learning [15, 179], self-supervised learning [56, 34, 160] and fixed or flexible combination strategies [183, 185, 140] to disentangle the latent knowledge, such as domain-shared and domain-specific information. GA-DTCDR [185] generates more representative user/item embeddings by constructing heterogeneous graphs from two domains and applies an element-wise attention mechanism to combine the embeddings of common users to enhance the recommendation accuracy in both domains. CausalCDR [57] incorporates causality into CDR by using causal embeddings to model the joint distribution of interactions and utilizes an adversarial domain classifier to decouple the domain-specific and domain-shared features. GDCCDR [72] leverages two distinct contrastive learning-based constraints for feature disentanglement: one preserves domain-invariant features across domains, and the other disentangles domain-specific features via mutual information, with meta-networks supporting the personalized transfer of domain-invariant features. CrossAug [88] uses intra- and inter-domain data augmentation based on cross-reconstructed representations, while utilizing Householder transformations for domain-shared center alignment to mitigate the domain shift. HJID [19] uses a hierarchical subspace disentanglement method to split user representations into generic shallow and domain-specific deep subspaces, utilizing a causal data generation graph to decouple domain-

shared and domain-specific latent factors, thus enhancing robustness against distribution shifts across domains.

2.1.2.3 Causal-Based Dual-Target CDR

Different from disentanglement-based approaches, causal-based approaches first construct causal graphs to identify confounders [27, 192] that influence both user preferences and user-item interactions [189]. Then, techniques like inverse propensity weighting (IPW) [100] or causal intervention [122] are applied to debias or deconfound such confounders, ensuring that the transferred knowledge accurately reflects true user preferences and strengthens recommendation robustness. SCDGN [64] builds a cross-domain user-cluster graph and employs a debiasing graph convolutional layer to extract and transfer unbiased graph knowledge between domains. CDRIB [9] devises two information bottleneck regularizers to simultaneously model user-item interactions within and across domains, aiming to debias the user and item representations. IPSCDR [63] employs a generalized inverse propensity score (IPS) estimator to mitigate selection bias in cross-domain contexts and devises three types of restrictions to learn propensity scores in the presence of unobserved domain-specific confounders. CD2CDR [189] designs a dual adversarial structure to decouple single-domain confounders (SDCs), leverages half-sibling regression to disentangle cross-domain confounders (CDCs), and then utilize backdoor adjustment to mitigate such confounders' negative effects to obtain debiased comprehensive user preferences.

2.1.2.4 Alignment-Based Dual-Target CDR

In addition, alignment-based approaches mitigate distribution shifts by directly aligning user behaviors and shared user-item interaction patterns [177, 144], or mapping auxiliary features into shared latent spaces [76, 142], thus facilitating effective knowledge transfer. DisAlign [76] combines contrastive learning with two embedding alignment methods, i.e., Stein path alignment and proxy Stein path alignment, which aim

to reduce distribution mismatch between source and target domains. CCDR [142] leverages a diversified preference network and employs both intra-domain and inter-domain contrastive learning strategies, where inter-domain tasks are constructed from user, taxonomy, and neighbor perspectives to facilitate effective transfer. COAST [176] jointly models cross-domain user–item interactions within a unified graph structure and aligns overlapping users' interests through contrastive signals and gradient-based optimization for better representation consistency. HGCCDR [144] leverages a fine-grained heterogeneous graph and domain-aware graph augmentation strategies to model complex user-item interactions, followed by multi-view contrastive learning to align and integrate domain-specific and domain-invariant user embeddings.

2.1.3 CDR: A Summary

Single-Target CDR [186] focuses on addressing the data sparsity problem by utilizing the abundant information available in a data-richer domain to improve the recommendation performance in a data-sparser domain. The existing single-target CDR approaches can be divided into three categories: content-based transfer, embedding-based transfer and rating pattern-based transfer [186]. Content-based transfer [49] leverages user/item attributes and textual information to establish links across domains. By contrast, embedding-based transfer [45, 23] employs machine learning techniques to extract user/item embeddings [150] for cross-domain transfer. In addition, rating pattern-based transfer mainly focus on transferring the learned rating patterns from the source domain to the target domain.

Although single-target CDR has been extensively studied, it suffers from a fundamental limitation: the knowledge transfer is restricted to a unidirectional flow from the source domain to the target domain. While these methods effectively leverage various forms of transferable knowledge such as content, embeddings, and rating patterns, they overlook the fact that even the data-sparser target domain may contain valuable information. This oversight prevents leveraging potentially valuable information from

the target domain that could, in turn, improve the recommendation performance in the source domain.

By contrast, in recent years, dual-target CDR has been proposed to improve the recommendation performance on both domains simultaneously by sharing the common knowledge across domains [186, 152]. The domain-shared information is variously termed as domain-invariant information [81, 109], domain-common information [15], or domain-independent information [73, 99] (note that the domain-independent information in other works does not have the same meaning as in our work, as discussed in Chapter 3). Despite the progress made in dual-target CDR research, existing methods still face three significant limitations. First, they cannot effectively capture comprehensive user preferences because they fail to disentangle all essential components of user preferences, instead entangling domain-independent information with domain-specific information. Second, they primarily focus on modeling user preferences while overlooking that users' final decisions are also influenced by observed confounders, which limits their ability to obtain a comprehensive understanding of user-item interactions in both domains. Third, these approaches either overlook distribution shifts entirely, or solely focus on addressing cross-domain distribution shifts (CDDS) while neglecting single-domain distribution shifts (SDDS), limiting their effectiveness in OOD scenarios. These limitations collectively hinder the performance of existing CDR methods in real-world recommendation tasks.

2.2 Out-of-Distribution (OOD) Recommendation

OOD generalization has emerged as a fundamental research direction in artificial intelligence, with significant developments across various fields such as natural language processing [5], computer vision [4], and others. Among these fields, RSs face particularly challenging OOD scenarios due to the inherently dynamic nature of user behaviors and item distributions. To the best of our knowledge, there is no existing method that explicitly addresses the cross-domain OOD recommendation problem.

Therefore, this section only reviews existing single-domain OOD recommendation methods, which can be broadly categorized into four classes, i.e., (1) disentanglement-based methods, (2) causal-based methods, (3) invariant learning-based methods and (4) adaptation-based methods.

2.2.1 Disentanglement-Based OOD Recommendation

Disentanglement-based methods aim to learn disentangled representations that separate stable user preferences from dynamic factors. For example, DICE [181] disentangles user behaviors into interest and conformity components, training separate embeddings with cause-specific data to capture each factor independently. Moreover, DCCL [178] employs contrastive learning to disentangle user interests from conformity behaviors, integrating item popularity signals to address data sparsity and enhance the robustness of causal representations.

2.2.2 Causal-Based OOD Recommendation

In contrast, causal-based methods focus on modeling the invariant causal mechanisms in recommendation scenarios. For instance, COR [123] views user attribute changes as interventions and formulates OOD recommendation as a post-intervention inference problem, using a variational framework to model the causal relationships between user features and interactions. In addition, CausPref [43] learns causal structures through Directed Acyclic Graphs (DAGs), combining invariant user preference learning with anti-preference sampling to handle implicit feedback. PopGo [155] improves OOD generalization by mitigating interaction-level popularity shortcuts, using a learned shortcut model to adjust predictions and emphasize true user preferences over spurious popularity shortcuts. CausalDiffRec [174] enhances OOD recommendation by eliminating environmental confounders through backdoor adjustment and learning environment-invariant graph representations via a causal diffusion process.

2.2.3 Invariant Learning-Based OOD Recommendation

Different from the above two classes, invariant learning-based methods aim to identify stable patterns across different environments while eliminating spurious correlations that may vary across distributions. For example, InvPref [135] identifies heterogeneous environments from interaction data to separate invariant user preferences from variant ones that are affected by environmental factors, i.e., unobserved confounders. Furthermore, InvCF [156] learns preference representations that remain stable under popularity shifts, utilizing an auxiliary classifier to separate invariant user preferences from dynamic popularity factors.

2.2.4 Adaptation-Based OOD Recommendation

In addition, adaptation-based methods focus on enhancing model adaptability to distribution shifts through various adaptation mechanisms. DR-GNN [116] integrates distributionally robust optimization (DRO) into GNN-based recommendation by treating GNN as a smoothing regularizer and injecting small perturbations into sparse neighbor distributions to enhance robustness against distribution shifts. DT3OR [148] introduces a dual test-time training strategy for OOD recommendation, adapting models to distribution shifts by learning invariant user preferences and variant user/item features through self-distillation and contrastive learning.

2.2.5 OOD Recommendation: A Summary

OOD recommendation addresses the critical challenge of maintaining recommendation quality when facing distribution shifts between training and testing environments. As reviewed above, existing research has developed four main approaches to tackle this challenge. Disentanglement-based methods separate stable user preferences from dynamic contextual factors, enabling models to focus on invariant preference patterns. Causal-based methods identify and model the underlying causal mechanisms governing user-item interactions, viewing distribution shifts as interventions in causal

structures. Invariant learning-based methods aim to discover and leverage preference patterns that remain stable across different environments while eliminating spurious correlations. Adaptation-based methods enhance model flexibility through various mechanisms that allow RSs to adapt to changing distributions during inference time.

Despite these advances, existing OOD recommendation methods suffer from two significant limitations. First, many approaches treat OOD recommendation primarily as a debiasing problem, focusing on specific factors like conformity or popularity as individual confounders while failing to provide a comprehensive framework for modeling multiple observed confounders simultaneously. Second, methods that do consider environmental factors often treat them as unobserved confounders, overlooking the potential impact of observed confounders that could be directly measured and accounted for. This incomplete consideration of the observed confounders limits the ability of existing methods to learn truly robust representations for effective OOD recommendation in complex real-world scenarios.

2.3 Disentangled Representation Learning

Disentangled representation learning has recently attracted increasing attention in RSs due to its ability to capture independent and interpretable preference factors. By separating underlying causes of user behaviors, it enhances both the robustness and generalizability of recommendation models. This section reviews the application of disentangled representation learning in recommendation from three perspectives: First, we provide an overview of the fundamental principles and applications of disentangled representation learning in general recommendation scenarios, including its origins in computer vision and its adaptation to various recommendation tasks. Second, we introduce multi-interest recommendation as a significant application field, focusing on its core methodologies categorized into interaction-based and auxiliary knowledge-based approaches. Finally, we review how disentangled representation learning specifically addresses the unique challenges in multi-interest recommendation, particularly the col-

lapse issue where interest representations lose diversity during training.

2.3.1 Disentangled Recommendation

Disentangled representation learning is originally introduced in the field of computer vision [11, 30] and is mainly used to learn visual features such as shape, color and location features of objects [168]. In addition to computer vision, in recent years, disentangled representation learning has also been used in RSs [85, 129, 86]. The main idea of disentangled representation learning is to focus on decomposing the latent factors behind the observed instances in the low-dimension vector space [3, 125]. For recommendation, MacridVAE [85] performs the macro disentanglement and micro disentanglement based on the user behavior data. DGCF [129] leverages a graph disentangling module to decouple user embeddings learned from user-item interaction data into fine-grained user intents. The method proposed by Ma et al. [86] reconstructs the embedding of future sequence by self-supervised learning to decouple the intentions of users. Moreover, disentangled representation learning has also been applied to causal recommendation. Existing disentanglement-based methods [130] first decouple the semantic-aware intent embeddings, and then employ causal intervention [149, 78] to alleviate the confounding bias.

2.3.2 Multi-Interest Recommendation

Multi-interest recommendation, as a significant branch within the disentangled recommendation paradigm, specifically aims to identify and disentangle users' diverse interest facets. Existing multi-interest recommendation approaches can be categorized into two main types: interaction-based approaches and auxiliary knowledge-based approaches [21]. Interaction-based approaches rely solely on user-item interaction data, typically employing either the capsule network [113, 143] or the attention mechanism [141, 80] to extract diverse interest representations from users' interaction patterns instead of using auxiliary knowledge. For instance, REMI [143] enhances multi-interest

representation learning by introducing an interest-aware hard negative mining strategy alongside routing regularization, effectively preventing the routing collapse in capsule networks. By contrast, auxiliary knowledge-based approaches mainly focus on leveraging external information such as knowledge graphs [71, 134] and multi-type behaviors [90, 61] to guide the learning of multi-interest representations. For example, CKML [90] is a coarse-to-fine framework that combines knowledge-enhanced interest extraction with dynamic routing-based behavior correlation modeling to capture behavior-shared and behavior-specific user interests.

2.3.3 Disentangled Representation Learning for Multi-Interest Recommendation

Disentangled representation learning has also been effectively applied to multi-interest recommendation [13], where the goal is to identify and separate users' diverse preference facets. One key challenge in this area is the collapse issue, where initially distinct interest embeddings become increasingly similar during training, resulting in a loss of diversity and a failure to capture users' multifaceted preferences. To address this challenge, researchers have proposed various disentanglement-based strategies, which can be broadly classified into two groups [21].

The first group of approaches primarily tackles the collapse issue through statistical regularization, aiming to enforce diversity among interest embeddings. Rather than modeling semantic differences directly, these approaches impose various mathematical constraints on the learning process to discourage homogeneity in the representation space. For instance, MDSR [13] frames sequential recommendation as a list generation process, leveraging multi-head attention and dynamic routing to mine multiple user interests, and promotes diversity through a specially designed decoding mechanism and training loss to generate balanced recommendation lists. Furthermore, VALID [114] achieves disentanglement within the regularization framework by iteratively refining personalized item clusters via latent attention.

The second group of approaches mainly handle the collapse issue via representation-guided refinement, focusing on optimizing the bidirectional relationships between interest embeddings and their corresponding items. For example, DisMIR [86] leverages disentanglement techniques to separate user intentions within sequential behavior patterns, applying self-supervised learning in the latent space to maintain interest diversity and prevent representation collapse over time. Moreover, Re4 [163] disentangles user interests by implementing backward flows from interests to items.

2.3.4 Disentanglement Representation Learning: A Summary

As reviewed above, disentangled representation learning in RSs has evolved to address various aspects of user preference modeling, with multi-interest recommendation emerging as a key application field. These methods share the fundamental goal of decomposing the latent factors behind observed user behaviors, thereby enhancing model interpretability, robustness, and effectiveness across recommendation scenarios.

Despite significant advances, several critical limitations persist across this research field. First, while some recent methods have begun to incorporate causal reasoning, many existing methods still predominantly focus on achieving representation disentanglement through architectural designs (such as capsule networks or attention mechanisms) or statistical techniques, without adequately capturing the underlying structural dependencies among different components of user preferences. This limitation constrains their ability to generalize effectively across varying contexts, especially in cross-domain scenarios. Second, in multi-interest recommendation specifically, existing methods face significant challenges in addressing the collapse issue, where initially distinct interest representations become increasingly homogeneous during training, despite various regularization strategies. Even when diversity is maintained, the relationship between the disentangled interests and their corresponding item groups often remains inadequately established. Third, existing methods typically focus on disentangling different components of user preferences, while overlooking the important

distinction between user preferences and various observed confounders. This incomplete consideration limits their effectiveness, particularly when dealing with complex cross-domain scenarios where single-domain confounders (SDCs) and cross-domain confounders (CDCs) coexist.

2.4 Causal Deconfounding

Causal deconfounding has emerged as a powerful technique for mitigating confounding biases in various machine learning domains. This section reviews two major application fields where such techniques have shown significant impact: RSs and domain generalization. In RSs, confounders may obscure true user preferences, resulting in biased recommendations. In domain generalization, they can induce spurious correlations between features and labels, thereby undermining the model's ability to generalize to unseen domains.

2.4.1 Deconfounded Recommendation

In recent years, causal learning [139, 145] has been introduced into RSs due to its ability to tackle confounding biases caused by confounders, which can be classified into two types: observed confounders and unobserved confounders. Moreover, these methods have also been extended to CDR scenarios, where addressing confounding biases becomes even more crucial for effective knowledge transfer across domains.

2.4.1.1 Deconfounded Recommendation for Observed Confounders

For observed confounders, the existing deconfounded RSs adopt inverse propensity weighting (IPW) [100] or backdoor adjustment [122] to address the observed specific confounders, such as item popularity [166] and video duration [154, 42]. DLCE [100] introduces an unbiased learning framework for estimating the causal effect of recommendations, where inverse propensity scoring (IPS) and empirical risk minimization

(ERM) with propensity capping are applied to optimize a debiased ranking objective under finite sample settings. To mitigate the influence of confounders in recommendation, DecRS [122] incorporates a causal graph-based formulation and applies an efficient approximation of backdoor adjustment, enabling dynamic correction of prediction bias based on user-specific conditions. To better utilize popularity bias in recommendation, PDA [166] formulates item popularity as a confounder and proposes a two-stage approach: deconfounding it during training and selectively reintroducing it at inference via causal intervention. To mitigate duration-induced bias in video recommendation, D2Q [154] formulates a causal framework that distinguishes the intrinsic and confounding effects of duration, applying group-wise modeling and label adjustment to ensure both debiasing and scalability. To address the confounding effect of item features, DCR [42] applies do-calculus-based intervention to block the backdoor path, and incorporates a mixture-of-experts (MoE) model to approximate the inference over confounder values with reduced computational cost.

2.4.1.2 Deconfounded Recommendation for Unobserved Confounders

For unobserved confounders, the existing deconfounded RSs either add additional assumptions [69, 70] or infer substitutes for confounders [132, 159] to alleviate the confounding bias. To address confounding bias arising from biased feedback, DIB [69] leverages causal analysis and information theory to separate user embeddings into biased and unbiased parts during training, using only the unbiased part for recommendation at inference time. To overcome the non-identifiability issue in deconfounded recommendation, iDCF [159] introduces a proximal inference framework that uses observable proxies to infer latent confounders and correct feedback prediction under unmeasured confounding. The method proposed by Liang et al. [66] introduce a causal graph incorporating both implicit and explicit feedback to better model user preferences, and employ the front-door adjustment to estimate unbiased user-item interactions in the presence of unobserved confounders. To address the influence of latent confounders on both users and items, MCDCF [46] models the sets of interacted

items and users as multi-cause treatment variables and learns substitute confounders from interaction data.

2.4.1.3 Deconfounded CDR

Moreover, recent research efforts have extended confounder debiasing into CDR scenarios, focusing mainly on unobserved confounders. The unobserved confounders can be further categorized into two classes: domain-specific confounders (e.g., purchase-guided domain setting) and general confounders (e.g., the display position of items) [63]. Most of existing approaches tend to remove the negative influences of domain-specific confounders [63, 167] or general confounders [173, 147]. Grace [167] unifies CDR and domain generalization through a causal invariant framework that disentangles user preferences into domain-invariant and domain-specific components, using adversarial learning and a routing mechanism to enhance transferability and adaptability. ARISEN [173] enhances cross-domain sequential recommendation (CDSR) by aligning user sequences temporally and employing an adaptive representation decomposition framework based on instrumental variables, guided by mutual information for improved causal disentanglement. To enhance CDSR under open-world assumptions, AMID [147] integrates a multi-interest information module that facilitates cross-domain knowledge transfer for both overlapping and non-overlapping users, and a doubly robust estimator for bias reduction in performance estimation.

2.4.2 Deconfounded Domain Generalization

Domain generalization [182, 119] aims to train models on labeled data from source domains to enhance their generalization ability across unseen target domains by learning domain-invariant feature representations. However, confounders influencing both features and labels can undermine such representations, preventing models from capturing the true causal effects. In recent years, causal inference techniques [102, 161] have been employed to address these confounding problems in domain generalization,

thereby enhancing the model's ability to generalize accurately across varied domains. For instance, a line of existing works [79, 169] simply adopts the average value of all domain features in each domain as the confounder, and employs backdoor adjustment to capture the true causality. Another line of works incorporates interventional pseudo-correlation augmentation [93] or adversarial training [118] to remove the confounders to better generalize to the unseen domain. There is also another line of works that exploits the instrumental variables [151] or learns substitutes [47] to eliminate the unobserved confounders and capture the invariant features for domain generalization.

2.4.3 Causal Deconfounding: A Summary

As reviewed above, causal deconfounding has emerged as a powerful method to address the confounding biases in both RSs and domain generalization. In RSs, researchers have developed various methods to mitigate the negative impacts of confounders. For observed confounders such as popularity and item features, approaches like IPW and backdoor adjustment have been applied to remove their confounding effects. For unobserved confounders, researchers have proposed methods that either introduce additional assumptions or infer substitutes to alleviate the confounding bias. These deconfounding techniques have also been extended to CDRs, where they address both domain-specific confounders and general confounders to facilitate effective knowledge transfer. Similarly, in domain generalization, causal deconfounding methods have been employed to mitigate the influence of confounders on feature-label relationships, enhancing models' ability to generalize across unseen domains.

Despite these advances, existing methods exhibit several critical limitations. First, most methods focus primarily on eliminating the negative impacts of confounders, overlooking the potential positive influences that certain confounders might have on predicted interactions. This complete removal strategy may discard valuable information that could enhance recommendation accuracy. Second, while various techniques have been developed for single-domain recommendation and domain generalization

separately, the unique challenge of CDR has not been adequately addressed, particularly the explicit decoupling of observed cross-domain confounders (CDCs). Third, many methods rely on simplified assumptions about confounders, either treating them as given (such as using off-the-shelf features) or neglecting domain-variant features.

2.5 Causal Structure Learning

Causal structure learning seeks to uncover causal relationships among variables from observational data, playing a fundamental role in fields such as healthcare, economics, and RSs [43]. To achieve this goal, causal discovery provides systematic approaches to inferring such relationships through statistical analysis [50]. In this section, we categorize causal structure learning methods into two main classes: conventional causal discovery methods and LLM-guided causal discovery methods.

2.5.1 Conventional Causal Discovery

Conventional causal discovery methods, which are primarily data-driven, can be broadly categorized into three groups, i.e., (1) dependency-based methods, (2) function-based methods, and (3) optimization-based methods. Dependency-based methods rely on statistical strategies to infer causal structures, but are often limited by the Markov equivalence class, making it challenging to determine unique causal directions [36]. In contrast, function-based methods such as LiNGAM [103] utilize asymmetries in the data generation process to identify unique causal structures. In addition, optimization-based methods like NOTEARS [180] reformulate structure learning as a continuous optimization problem, using gradient-based techniques to learn causal relationships.

2.5.2 LLM-Guided Causal Discovery

Furthermore, recent advances in Large Language Models (LLMs) have led to a new paradigm for causal discovery through LLM-guided methods [55, 112]. In contrast

to conventional causal discovery methods, these methods leverage the semantic understanding of variable descriptions to infer causal relationships. Existing studies in this direction primarily follow two streams: one that explores direct causal inference through LLM prompting [51, 153], and another that integrates LLMs into conventional causal discovery methods as prior or posterior knowledge to improve their effectiveness [2, 115]. MAC [55] proposes a multi-agent framework that integrates statistical causal discovery with LLM-based reasoning, where agents collaboratively select causal methods and iteratively refine causal structures using both structured data and metadata. ILS-CSL [2] combines LLM-based inference with data-driven structure learning in an iterative loop, where LLMs supervise and refine edge predictions, resulting in higher-quality and more consistent causal graphs.

2.5.3 Causal Structure Learning: A Summary

As reviewed above, causal structure learning approaches have evolved from conventional data-driven methods to emerging LLM-guided methods, each offering distinct advantages in uncovering causal relationships among variables. Conventional methods, whether dependency-based, function-based, or optimization-based, rely primarily on statistical patterns and mathematical formulations to infer causal structures from observational data. These approaches provide rigorous frameworks for causal discovery but often struggle with challenges such as the Markov equivalence class problem. In contrast, LLM-guided methods leverage the rich semantic understanding and world knowledge embedded in LLMs, either through direct prompting for causal inference or by integrating LLMs with conventional statistical techniques. This integration has shown promising results in improving the quality and consistency of discovered causal structures, particularly in domains with complex variable relationships.

Despite these advances, several limitations persist across the current research field. First, conventional methods typically require pre-defined variables or substantial domain expertise to properly formulate the causal discovery problem, limiting their appli-

cability in domains where causal factors are not well understood or easily identifiable. Second, many LLM-guided approaches rely heavily on the world knowledge embedded in these models, which may not always align with domain-specific causal mechanisms and can potentially introduce biases from the models' training data. Third, while these two types of approaches have shown effectiveness, few attempts have been made to develop frameworks that seamlessly combine the conventional methods with the LLMs, particularly for extracting observed confounders from unstructured data such as user reviews. This gap is especially significant in RSs, where user-generated content contains valuable information about observed confounders that influence both user preferences and user-item interactions.

2.6 Summary

This chapter has presented a comprehensive review focusing on two core research problems: Cross-Domain Recommendation (CDR) and Out-Of-Distribution (OOD) recommendation. In addition, we have reviewed three key techniques essential to our thesis: disentangled representation learning, causal deconfounding, and causal structure learning. Through this systematic analysis, we have identified several critical limitations in the existing literature that lead to four fundamental challenges addressed in this thesis.

Firstly, existing CDR methods either fail to properly disentangle domain-specific and domain-shared preferences or completely overlook domain-independent user preferences, resulting in incomplete user modeling and suboptimal recommendation performance. Addressing this limitation requires advanced disentangled representation learning techniques to effectively decouple all the essential components of user preferences (**CH1**).

Secondly, existing CDR methods fail to explicitly decouple comprehensive observed confounders. These observed confounders, often embedded within various forms of textual information, remain largely underexplored, limiting the comprehen-

sive understanding of user-item interactions. Addressing this limitation requires developing effective techniques to accurately extract and properly represent these observed confounders (**CH2**).

Thirdly, most existing causal deconfounding methods focus solely on eliminating the negative impacts of confounders while neglecting their potentially positive effects on interaction prediction, limiting the recommendation accuracy in real-world scenarios. Addressing this limitation requires developing balanced causal deconfounding methods that can preserve the positive impacts of confounders as well (**CH3**).

Finally, recent OOD recommendation studies typically address cross-domain distribution shifts (CDDS) and single-domain distribution shifts (SDDS) in isolation, failing to handle scenarios where these two types of shifts co-exist, thus leading to degraded recommendation performance in CDR. Addressing this limitation requires leveraging causal structure learning to discover causal relationships that remain invariant across both types of distribution shifts (**CH4**).

Chapter 3

Domain Disentanglement with Interpolative Data Augmentation for Dual-Target Cross-Domain Recommendation

In recent years, dual-target CDR has been proposed to improve the recommendation performance on both domains simultaneously by sharing the common knowledge across domains [186, 152]. However, none of the existing dual-target CDR models decouples all three essential components needed to capture comprehensive user preferences. To be specific, the existing methods only decouple two essential components, i.e., domain-shared and domain-specific information, and ignore the existence of domain-independent information. Since domain-independent information has different meanings from the other two types of information, it cannot be ignored when capturing comprehensive user preferences.

Below we introduce domain-shared, domain-specific and domain-independent information respectively with examples, and further differentiate them.

(i) **Domain-shared Information:** There exists some *domain-shared* information in both movie domain and music domain, such as 'Category'. For example, people who like watching *suspense movies* (i.e., a category in movie domain) tend to like listening to *suspense music* (i.e., a category in music domain), and vice versa. Since

the domain-shared information can provide the valuable information for CDRs, it needs to be first decoupled and then transferred to both domains.

(ii) **Domain-specific Information:** In contrast, there also exists some *domain-specific* information in each domain. For example, the user preference for pictures in a movie (i.e., 'Frame' in movie domain) is not applicable in music domain because 'Frame' is unique to movie domain. Thus, such domain-specific information should be decoupled too, which helps improve the recommendation performance on its own domain, but it should not be transferred to another domain to avoid negative transfer [187].

(iii) **Domain-independent Information:** In addition, some *domain-independent* information also exists in each domain, but should not be transferred to other domains. For instance, 'Rhythm' exists in each of movie, music and book domains. However, in movie domain, it means the use of sound effects, the speed of camera cuts, and the changes in the pace of movie scenes, etc [104]. In music domain, it means the use of beats and the speed of songs, etc [37]. In book domain, it means the fluidity of the writing and the ups and downs of the storyline, etc [29]. In other words, although 'Rhythm' is seemingly common in all three domains, it has different meanings in different domains. Hence, such information is domain-independent and should be extracted from its own domain for capturing comprehensive user preferences, but should not be transferred to other domains.

(iv) **Difference:** Different to the domain-independent information, the domain-shared information (e.g., 'Category') extracted from two domains expresses the same meaning in each domain. In addition, the domain-specific information, e.g., 'Frame' in the movie domain, exists only in its own domain. By contrast, although the domain-independent information, e.g., 'Rhythm', exists in each domain, it has different meanings in different domains.

(v) **Summary:** Therefore, it is vital to recognize the existence of domain-independent information and to clearly differentiate it from domain-shared and domain-specific

information. More importantly, decoupling domain-independent information is crucial for capturing more comprehensive user preferences, otherwise it will cause suboptimal recommendation results. However, existing dual-target CDR methods neglect the above insights. Hence, novel dual-target CDR solutions are needed to incorporate the above insights for capturing more comprehensive user preferences.

Following the above discussions, to target superior dual-target CDR, there is a major challenge that has been introduced in Section 1.2.1 and expressed by **CH1**: *How to effectively decouple domain-independent information from domain-specific information, in addition to domain-shared information, to capture comprehensive user preferences on each domain, thereby improving the recommendation performance?*

Then, this chapter will formulate the target problem of our proposed model. Next, we elaborate on the basic components of DIDA-CDR. Finally, we conduct extensive experiments on five real-world datasets to show the significant superiority of DIDA-CDR over the state-of-the-art methods.

3.1 Problem Statement

This chapter considers the dual-target CDR on two domains D^A and D^B with a shared user set, denoted by $\mathcal{U}$ (of size $m = |\mathcal{U}|$). The sets of items in D^A and D^B are defined as $\mathcal{V}^A$ (of size $n^A = |\mathcal{V}^A|$) and $\mathcal{V}^B$ (of size $n^B = |\mathcal{V}^B|$), respectively. Let $\mathbf{R}^A \in \{0,1\}^{m \times n^A}$ and $\mathbf{R}^B \in \{0,1\}^{m \times n^B}$ denote the binary user-item interaction matrices in D^A and D^B, respectively. By aggregating the interaction data in each domain, we first construct two heterogeneous graphs $G^A = (\mathcal{U}, \mathcal{V}^A, \mathcal{Q}^A)$ and $G^B = (\mathcal{U}, \mathcal{V}^B, \mathcal{Q}^B)$ to learn user embeddings $\mathbf{E}_u^A$, $\mathbf{E}_u^B$ and item embeddings $\mathbf{E}_v^A$, $\mathbf{E}_v^B$ in domains D^A and D^B, respectively, where $\mathcal{Q}^A$ and $\mathcal{Q}^B$ are the edge sets that represent the observed user-item interactions. By linearly interpolating the user embeddings $\mathbf{E}_u^A$ and $\mathbf{E}_u^B$, we then generate augmented user representations $\mathbf{E}_u^{aug}$ to augment the sparser domain D^B. Given two coarse user embeddings $\mathbf{E}_u^A$, $\mathbf{E}_u^B$ and the augmented user representations $\mathbf{E}_u^{aug}$, our goal is to disentangle domain-shared, domain-specific and domain-independent

user preferences, i.e., $\mathbf{Z}_{sha}$, $\mathbf{Z}_{spe}$, and $\mathbf{Z}_{ind}$, and then transfer domain-shared user preferences $\mathbf{Z}_{sha}$ to each domain to capture comprehensive user preferences $\mathbf{E}_u^*$, thereby improving the recommendation performance on both domains.

3.2 The Proposed DIDA-CDR Framework

3.2.1 Framework Overview

To enhance the recommendation accuracy in each of both domains, we propose a novel disentanglement-based framework with interpolative data augmentation for dual-target cross-domain recommendation, called DIDA-CDR. This framework contains five major components, i.e., (1) *Graph Convolution and Propagation Module*, (2) *Interpolative Data Augmentation Module*, (3) *User Preference Disentanglement Module*, (4) *Information Fusion Module*, and (5) *Prediction Module*. The details of each module of our proposed DIDA-CDR are illustrated in Figure 3.1. They are briefly introduced below and described in detail in the following subsections.

(1) Graph Convolution and Propagation. First, we construct two heterogeneous graphs to extract the high-order user-item interaction relationships using the interaction data in domain A and domain B, respectively. Based on the above graphs, we apply the graph convolution and propagation layer in the GCN [53] to generate user and item embeddings.

(2) Interpolative Data Augmentation. Next, we propose an interpolative data augmentation approach to augmenting user embeddings at the representation level. The interpolative data augmentation approach linearly interpolates user embeddings in domains A and B to generate both relevant and diverse augmented user representations.

(3) User Preference Disentanglement. Thereafter, we propose a disentanglement module guided by a domain classifier to decouple more accurate domain-specific and domain-independent user preferences from user embeddings generated by GCNs. This module also disentangles domain-shared user preferences from augmented user repre-

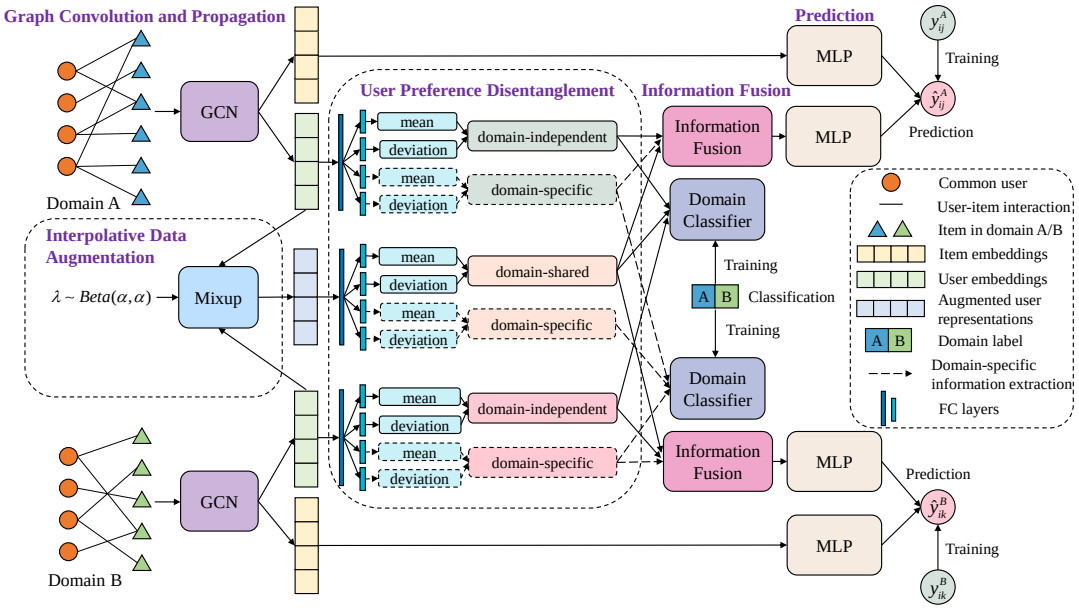

Figure 3.1: The details of each module of our proposed DIDA-CDR.

sentations, which are then transferred to both domains to provide valuable information.

(4) Information Fusion. Next, we use three approaches, i.e., concatenation, element-wise sum, and attention mechanism to incorporate three essential components of user preferences, i.e., domain-shared, domain-specific and domain-independent information, which are decoupled by the disentanglement module, to capture the comprehensive user preferences.

(5) Prediction. Finally, we apply the multi-layer perceptron (MLP) to model the user-item interaction relationships, endowing the non-linearity to our proposed model. Based on the MLP, the predicted user-item interaction matrix can be obtained. The prediction loss between it and the observed user-item interaction matrix, together with two domain classification losses, constitute the final loss for training.

Overall, our model can be easily extended to a multi-target CDR model. Specifically, we can disentangle domain-shared user preferences and transfer them to all domains, decouple the domain-specific and domain-independent user preferences from each domain, and then capture the comprehensive user preferences to improve the recommendation performance on multiple domains simultaneously.

3.2.2 Graph Convolution and Propagation

GCNs excel at capturing the relationships between nodes and learning the representation of graph data, and are well suited for modeling user-item relationships in RSs, because the user-item interaction data can be easily transformed into the graph structure. To distill the high-order user-item interaction relationships, we construct two heterogeneous graphs G^A and G^B for domains A and B, respectively, where nodes refer to entities (i.e., users and items) and edges refer to interactions. In this paper, we apply the graph convolution and propagation layer in the GCN to encode the user and item embeddings according to the user-item interaction matrix $\mathbf{R}^A$ (or $\mathbf{R}^B$). The node embeddings $\mathbf{E}_0^A$ (or $\mathbf{E}_0^B$) are randomly initialized. Given a graph G^A, the propagation rule is represented as:

$$\mathbf{E}_l^A = f(\tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{R}}^A\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{E}_{l-1}^A\mathbf{W}_l + \mathbf{b}_l), \tag{3.1}$$

where $\tilde{\mathbf{R}}^A = \mathbf{R}^A + \mathbf{I}$ is the user-item interaction matrix of graph G^A after adding a self-loop identity matrix $\mathbf{I}$. $\tilde{\mathbf{D}}$ is a degree matrix for normalization. $\mathbf{W}_l$ and $\mathbf{b}_l$ are the trainable weight matrix and bias vector in the l^{th} layer respectively, and $\mathbf{E}_l^A$ is the hidden embedding matrix of graph G^A in the l^{th} layer [75]. $f(\cdot)$ denotes the ReLU activation function.

After l times propagation, we can obtain the global hidden representations $\mathbf{E}^A$ by concatenating multiple embedding matrices from $\mathbf{E}_0^A$ to $\mathbf{E}_l^A$, which can be rearranged into the user embeddings $\mathbf{E}_u^A$ and item embeddings $\mathbf{E}_v^A$ in domain A [172]. Similarly, we can obtain the user embeddings $\mathbf{E}_u^B$ and item embeddings $\mathbf{E}_v^B$ in domain B.

3.2.3 Interpolative Data Augmentation

Although the user embeddings in both domains are obtained by GCNs, the user embeddings learned from the sparser domain are not as accurate as those learned from the richer domain due to the data imbalance between the two domains. To augment

the sparser domain and explore the potential user preferences, inspired by the mixup technique [157], we design an interpolative data augmentation approach, which generates both relevant and diverse augmented user representations. However, the conventional mixup technique cannot be directly utilized for our task, because the interaction data cannot be directly mixed at the pixel level like images. Therefore, we propose to linearly interpolate the embeddings of common users in both domains. The augmentation of common users ensures that augmented user representations maintain the relevance of user preferences in both domains, i.e., relevant augmented user representations, while corresponding to adding more interaction data for users in the sparser domain. In addition, introducing randomness in linear interpolation instead of using a fixed mixing coefficient can generate diverse augmented user representations, which can provide richer information for subsequent disentanglement. The formula for interpolative data augmentation can be expressed as follows:

$$\mathbf{E}_u^{aug} = \lambda \mathbf{E}_u^A + (1-\lambda)\mathbf{E}_u^B, \tag{3.2}$$

where $\mathbf{E}_u^{aug}$ denotes the augmented user representations[1]. Since we aim to generate more diverse augmented user representations through data augmentation, instead of using attention-based methods in this module, we propose the interpolative data augmentation approach. $\lambda \in [0,1]$ is the mixing coefficient sampled from $Beta(\alpha,\alpha)$, $\alpha \in (0,\infty)$. The advantages of adopting a mixing coefficient sampled from $Beta(\alpha,\alpha)$ are as follows. First, since the user embeddings of the common user in each of the two domains should be equivalent, the mixing coefficient should be sampled within the interval $[0,1]$ and be symmetric around 0.5. Beta distribution $Beta(\alpha,\alpha)$ satisfies this characteristic. Second, it has been proven to effectively improve the generalization ability of the model [133]. Meanwhile, λ also introduces randomness into our model, thus weakening the negative transfer that may result from performing a linear interpolation operation with fixed weights.

[1]The two terms, i.e., embedding and representation, are exchangeable in this paper.

3.2.4 User Preference Disentanglement

To capture domain comprehensive preferences of users, we utilize disentangled representation learning to extract the essential components of user preferences from previously obtained user embeddings. Inspired by the method introduced in [24], we design a disentanglement module to decouple the domain-specific and domain-independent user preferences. Specifically, our disentanglement module adopts a similar architecture to the encoder of VAE, but it is quite different from VAE. In particular, VAE only encodes one latent feature, while our disentanglement module aims to better learn to decouple the domain-specific and domain-independent user preferences. Although both domain-specific and domain-independent user preferences cannot be transferred to other domains, they have different meanings, thus their importance in capturing comprehensive user preferences is different. If they are not distinguished, they are equally important in capturing comprehensive user preferences, which is inappropriate. By contrast, if they are decoupled, the subsequent information fusion module can use attention mechanism to learn their weights respectively, thereby capturing more accurate and comprehensive user preferences. This module can also disentangle the domain-shared user preferences from the augmented user representations. In this case, our model can not only extract the domain-shared preferences of common users in both domains, but also explore the domain-specific personalized user preferences and domain-independent user preferences, which enhances the comprehensiveness of capturing user preferences, and thus improves the recommendation performance on both domains. To this end, we feed user embeddings in both domains (see methods introduced in Section 3.2.2) and augmented user representations (see methods introduced in Section 3.2.3) into this module and perform following processing respectively.

Taking domain A as an example, the user embeddings $\mathbf{E}_u^A$ are first entered into the disentanglement module, which consists of several fully connected (FC) layers. More specifically, the first FC layer (see navy blue FC layer in Figure 3.1), followed by the ReLU activation function, is utilized to extract general representations. In addition, the

subsequent FC layers (also see lake blue FC layers in Figure 3.1)) outputs two sets of latent vectors, each representing different mean and deviation information of the input user embeddings. Next, following the method introduced in [24], a reparametrization trick is adopted to generate domain-independent $\mathbf{Z}_{ind}^{A}$ and domain-specific user preferences $\mathbf{Z}_{spe}^{A}$. Similarly, we can obtain the domain-independent $\mathbf{Z}_{ind}^{B}$ and domain-specific user preferences $\mathbf{Z}_{spe}^{B}$ in domain B, as well as the domain-shared $\mathbf{Z}_{sha}^{aug}$ and domain-specific user preferences $\mathbf{Z}_{spe}^{aug}$ decoupled from augmented user representations.

To ensure that the above essential components of user preferences can be accurately disentangled, we introduce a domain classifier $H_{cls}(\cdot)$, which includes a single FC layer to predict the domain probability of user preferences, to guide the disentanglement process. In order to supervise the optimization process of disentanglement module, we further set two domain classification tasks to train our DIDA-CDR. First of all, the disentanglement module is guided to decouple the domain-specific information with stronger domain identification ability, i.e., more accurate domain-specific user preferences, by minimizing the domain classification loss $\mathcal{L}_{cls_1}$. In other words, if the decoupled user preferences can be easily recognized by a domain classifier as belonging to a particular domain, then such user preferences are considered as the domain-specific information. By contrast, the domain-independent and domain-shared information are leveraged to confuse the domain classifier, i.e., to make the domain classifier unable to identify the domain to which they belong, to ensure that they can be distinguished from the domain-specific information. In other words, if the decoupled user preferences are no better than random guesses in identifying the domain to which they belong when fed into a domain classifier, then such user preferences are mutually exclusive with the domain-specific information and should be classified as the domain-independent information. When the input to the disentanglement module is the augmented user representations containing the user preferences from both domains, the above user preferences refer to the domain-shared information. Specifically, the domain classification loss $\mathcal{L}_{cls_1}$ can be defined as follows:

$$\begin{aligned}\mathcal{L}_{cls_1} = \frac{1}{3}\sum [&\ell_1(H_{cls}(\mathbf{Z}^A_{spe}), O^A_{spe}) \\ &+ \ell_1(H_{cls}(\mathbf{Z}^B_{spe}), O^B_{spe}) \\ &+ \lambda \cdot \ell_1(H_{cls}(\mathbf{Z}^{aug}_{spe}), O^{augA}_{spe}) \\ &+ (1-\lambda) \cdot \ell_1(H_{cls}(\mathbf{Z}^{aug}_{spe}), O^{augB}_{spe})],\end{aligned} \tag{3.3}$$

where $\ell_1(P, O)$ denotes the cross-entropy loss function. P is the predicted domain probability of input user preferences and O is the ground truth (GT) domain label. We define the corresponding GT for domain-specific user preferences $\mathbf{Z}^A_{spe}$, $\mathbf{Z}^B_{spe}$ and $\mathbf{Z}^{aug}_{spe}$ as O^A_{spe}, O^B_{spe} and O^{aug}_{spe}, respectively. All the elements in O^A_{spe} and O^{augA}_{spe} are set to 1, while those in O^B_{spe} and O^{augB}_{spe} are set to 0. Here, λ represents the confidence score that an augmented user representation belongs to its initial GT, since augmented user representations are generated by incorporating embeddings of common users in both domains with a mixing coefficient λ [24].

Similarly, we define the corresponding GT for domain-independent user preferences $\mathbf{Z}^A_{ind}$, $\mathbf{Z}^B_{ind}$ and domain-shared user preferences $\mathbf{Z}^{aug}_{sha}$ as O^A_{ind}, O^B_{ind} and O^{aug}_{sha}, respectively. We set all the items in O^A_{ind}, O^B_{ind} and O^{aug}_{sha} as $[0.5, 0.5]$ [24]. This ensures that learned domain-shared and domain-independent user preferences cannot be used to identify the domain to which they belong, and thus be distinguished from domain-specific user preferences. Specifically, the loss function is expressed as follows:

$$\begin{aligned}\mathcal{L}_{cls_2} = \frac{1}{3}\sum [&\ell_2(H_{cls}(\mathbf{Z}^A_{ind}), O^A_{ind}) \\ &+ \ell_2(H_{cls}(\mathbf{Z}^B_{ind}), O^B_{ind}) \\ &+ \ell_2(H_{cls}(\mathbf{Z}^{aug}_{sha}), O^{aug}_{sha})],\end{aligned} \tag{3.4}$$

where $\ell_2(P, O)$ is the Kullback-Leibler divergence loss function.

Table 3.1: Comparison of information fusion approaches [110].

	Formula
Concatenation	$\mathbf{E}_u^* = [\mathbf{Z}_{spe}, \mathbf{Z}_{ind}, \mathbf{Z}_{sha}]$
Element-wise sum	$\mathbf{E}_u^* = \mathbf{Z}_{spe} + \mathbf{Z}_{ind} + \mathbf{Z}_{sha}$
Attention	$\mathbf{C}_u = Softmax(\mathbf{W}_s \cdot \sigma(\mathbf{W}_{spe} \cdot \mathbf{Z}_{spe} + \mathbf{W}_{ind} \cdot \mathbf{Z}_{ind} + \mathbf{W}_{sha} \cdot \mathbf{Z}_{sha}))$ $\mathbf{E}_u^* = [\mathbf{Z}_{spe}, \mathbf{Z}_{ind}, \mathbf{Z}_{sha}] \cdot \mathbf{C}_u$

3.2.5 Information Fusion

The domain-shared, domain-specific and domain-independent information are three essential components of user preferences, which need to be integrated in a reasonable and efficient way to capture comprehensive user preferences. To this end, in this paper, we leverage three approaches, i.e., concatenation, element-wise sum, and attention mechanism, to aggregate individual representations into comprehensive user preferences [110]. The specific operations of these information fusion approaches are expressed in Table 3.1.

3.2.6 Model Prediction and Training

After the information fusion, we obtain the comprehensive user preferences $\mathbf{E}_u^*$ and we also have the corresponding item embeddings $\mathbf{E}_v$ generated by GCNs. To give our model the non-linearity, we adopt a neural network, i.e., MLP, to represent the user-item interactions. Taking the domain A as an example, the input user embeddings and item embeddings in domain A for the MLP are defined as $\mathbf{S}_{in}^A = \mathbf{E}_u^{A*}$ and $\mathbf{T}_{in}^A = \mathbf{E}_v^A$, respectively. Moreover, the output embeddings of user u_i and item v_j of MLP is expressed as:

$$\mathbf{S}_i^A = \mathbf{S}_{out_i}^A = \delta(\dots\delta(\delta(\mathbf{S}_{in_i}^A \cdot W_{\mathbf{S}_1}^A) \cdot W_{\mathbf{S}_2}^A)), \tag{3.5}$$

$$\mathbf{T}_i^A = \mathbf{T}_{out_i}^A = \delta(\dots\delta(\delta(\mathbf{T}_{in_i}^A \cdot W_{\mathbf{T}_1}^A) \cdot W_{\mathbf{T}_2}^A)), \tag{3.6}$$

where $\delta(\cdot)$ is the LeakyReLU activation function. $\mathbf{W}_{\mathbf{S}_1}^A$, $\mathbf{W}_{\mathbf{S}_2}^A \dots$ and $\mathbf{W}_{\mathbf{T}_1}^A$, $\mathbf{W}_{\mathbf{T}_2}^A \dots$ denote the trainable weight matrices of MLP in various layers, respectively.

Next, the predicted user-item interaction $\hat{y}_{ij}^A$ between user u_i and item v_j in domain

A can be formulated as follows:

$$\hat{y}_{ij}^{A} = cosine(\mathbf{S}_i^A, \mathbf{T}_j^A) = \frac{\mathbf{S}_i^A \cdot \mathbf{T}_j^A}{\|\mathbf{S}_i^A\| \|\mathbf{T}_j^A\|}. \tag{3.7}$$

Furthermore, the prediction loss in domain A is defined as follows:

$$\mathcal{L}_{prd}^{A} = \sum_{y \in \mathcal{Y}^{A+} \cup \mathcal{Y}^{A-}} \ell_1(\hat{y}, y) + \gamma(\|\mathbf{S}^A\|_F^2 + \|\mathbf{T}^A\|_F^2), \tag{3.8}$$

where $\ell_1(\hat{y}, y)$ is the cross-entropy loss function. y denotes an observed user-item interaction, and $\hat{y}$ is the corresponding predicted user-item interaction. $\mathcal{Y}^{A+}$ denotes the set of observed interactions and $\mathcal{Y}^{A-}$ is a certain number of negative instances randomly sampled from the set of unseen interactions in domain A to avoid overfitting. $\|\mathbf{S}^A\|_F^2 + \|\mathbf{T}^A\|_F^2$ is a regularizer controlled by γ.

Finally, we utilize a multi-task learning mechanism consisting of a prediction task and two domain classification tasks to optimize our model in domain A. Specifically, the final loss function is formulated as follows:

$$\mathcal{L}^{A} = \mathcal{L}_{prd}^{A} + \mu_1 \cdot \mathcal{L}_{cls_1} + \mu_2 \cdot \mathcal{L}_{cls_2}, \tag{3.9}$$

where μ_1 and μ_2 denote the weights of domain classification losses $\mathcal{L}_{cls_1}$ and $\mathcal{L}_{cls_2}$, respectively. Similarly, we perform the same optimization process for domain B.

3.3 Experiments on DIDA-CDR

In order to demonstrate the superiority of our proposed DIDA-CDR and explore the effectiveness of its various modules, we conduct extensive experiments on five real-world datasets to answer the following five research questions:

- **RQ1.** How does our model perform when compared to representative and state-of-the-art baseline models (see Section 3.3.2.1)?

Table 3.2: Statistics of three dual-target CDR tasks.

Tasks	Datasets	#Users	#Items	#Interactions	Density
Task 1	Douban-Movie	2106	9555	907219	4.508%
	Douban-Book	2106	6777	95974	0.672%
Task 2	Douban-Movie	1666	9555	781288	4.908%
	Douban-Music	1666	5567	69681	0.751%
Task 3	Amazon-Elec	15761	51447	224689	0.027%
	Amazon-Cloth	15761	48781	133609	0.017%

- **RQ2.** How do various modules (i.e., interpolative data augmentation and user preference disentanglement) affect the results of our model (see Section 3.3.2.2)?
- **RQ3.** How do various components of user preferences (i.e., domain-shared, domain-specific and domain-independent information) contribute to the performance improvement of our model (see Section 3.3.2.3)?
- **RQ4.** How do various information fusion approaches influence the performance of our models (see Section 3.3.2.4)?
- **RQ5.** How does the performance of our model change with various values of hyper-parameters (see Section 3.3.2.5)?

3.3.1 Experimental Setting

3.3.1.1 Experimental Datasets and Tasks

In order to verify the recommendation performance of our proposed DIDA-CDR, we conduct extensive experiments on five real-world datasets, i.e., Douban subsets (Douban-Movie, Douban-Book and Douban-Music) released in GA-DTCDR [185] and Amazon subsets (Amazon-Elec and Amazon-Cloth) released in DisenCDR [8]. For these five datasets, we first convert the explicit ratings into implicit feedback, i.e., we binarize the ratings into 0 and 1 to indicate whether the user has interacted with the item or not. Following the methods introduced in [77, 74], we then filter these datasets to remove users and items with less than 5 interactions. Since the two Amazon subsets

are filtered out of the cold-start item entry in the test set, following DisenCDR [8], we perform the same preprocessing operation on the three Douban subsets as well for a fair comparison. Finally, we divide the above five subsets into three pairs of datasets, extract the common users in each pair of datasets, and design three dual-target CDR tasks in a scenario where users completely overlap, which can be listed as follows:

- **Task 1:** Douban-Movie (richer) $\leftrightarrow$ Douban-Book (sparser)
- **Task 2:** Douban-Movie (richer) $\leftrightarrow$ Douban-Music (sparser)
- **Task 3:** Amazon-Elec (richer) $\leftrightarrow$ Amazon-Cloth (sparser)

The details of the above three tasks and corresponding datasets are shown in Table 3.2.

3.3.1.2 Parameter Settings

For the graph convolution and propagation module in Figure 3.1, the layer structure of GCN is '$k \rightarrow k$' and for the disentanglement module, the layer structure is '$2k \rightarrow k$'. In the prediction module, the layer structure of user branch MLP is '$k \rightarrow 2k \rightarrow k$', and the layer structure of item branch MLP is '$2k \rightarrow 2k \rightarrow k$'. k is the embedding dimension. We vary k in the range of $\{64, 128\}$, but in order to balance the trade-off of recommendation accuracy and model training time, we finally set k to 64. The parameters of all these layers are initialized as the Gaussian distribution $X \sim \mathcal{N}(0, 0.01)$. For each observed user-item interaction, following GA-DTCDR [185], we randomly sample 7 unseen interactions as negative instances. For a fair comparison, we leverage the grid search to tune the choice of parameters of all models. For the baseline models, we tune them based on the best parameter settings listed in their original papers. Specifically, we choose the learning rate from $\{0.01, 0.005, 0.001, 0.0005, 0.0001\}$, and search the regularization coefficient in the range of $\{0.001, 0.0001, 0.00001\}$. In addition, we apply the Adam [52] to optimize all the models, and the batch size is 1024. We train our model and other baseline models with 100 epochs in order to guarantee the convergence. Moreover, we investigate the number of GCN layers l in $\{1, 2, 3, 4\}$, α of

$Beta(\alpha, \alpha)$ in $\{0.1, 0.5, 1, 2, 5\}$ and the weights of domain classification losses μ_1, μ_2 in $\{0.1, 0.3, 0.5, 0.7, 1, 3, 5, 10\}$, and analyze their impact on the recommendation performance of our model in Section 3.3.2.5. In the experiments, we set $l = 2$, $\alpha = 1$, and $\mu_1 = \mu_2 = 1$ by default and we resample the mixing coefficient λ once for training each batch of data.

3.3.1.3 Evaluation Metrics

Since the leave-one-out method is ubiquitous in baseline models, such as GA-DTCDR [185] and BiTGCF [74], we also employ it to evaluate the recommendation performance of our proposed DIDA-CDR and baseline models. In other words, we utilize the last interaction record of each test user to form the test set, while all the other interaction records are used as the training set. Following the methods introduced in [54, 8], for each test user-item interaction, we randomly sample 999 items that the test user has not interacted with as negative items, and then predict 1000 candidate scores for ranking. The leave-one-out method contains two main metrics, i.e., Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG) [127], which are widely-used ranking evaluation metrics [187, 77]. In the experiments, we employ them to evaluate the performance of the top-10 ranking results. For a fair evaluation, we perform all experiments 5 times and present the average results.

3.3.1.4 Comparison Methods

We select a total of nine representative and state-of-the-art baseline models to compare with our proposed DIDA-CDR. These nine baseline models can be divided into four categories, i.e., (1) Single-Domain Recommendation (SDR), (2) Single-Target Cross-Domain Recommendation (CDR), (3) Conventional Dual-Target CDR, and (4) Disentanglement-Based Dual-Target or Multi-Target CDR. For a clear comparison, in Table 3.3, we elaborate on the embedding strategies and transfer strategies of the above nine baseline models and our proposed DIDA-CDR.

Table 3.3: The comparison of the baselines and our approach [185].

	Model		Embedding Strategy	Transfer Strategy
Baselines	Single-Domain Recommendation (SDR)	**NGCF** [128]	Graph Embedding	-
		LightGCN [40]	Graph Embedding	-
	Single-Target Cross-Domain Recommendation (CDR)	**BPR_EMCDR** [87]	Linear Matrix Factorization (MF)	MLP
		BPR_DCDCSR [**?**]	Linear MF	Combination & MLP
	Conventional Dual-Target CDR	**PPGN** [172]	Graph Embedding	Combination & MLP
		BiTGCF [74]	Graph Embedding	Transfer Learning
	Disentanglement-Based Dual-Target or Multi-Target CDR	**GA-DTCDR** [185]	Graph Embedding	Combination (Element-wise Attention)
		DR-MTCDR [34]	Graph Embedding & Self-supervised Learning	Domain Adaptation
		DisenCDR [8]	Graph Embedding & VAE	Transfer Learning
Our Approach	Disentanglement-Based Dual-Target CDR	**DIDA-CDR**	Graph Embedding & Disentanglement	Transfer Learning & Combination (Attention)

3.3.2 Experimental Results and Analysis

3.3.2.1 Performance Comparison (for RQ1)

Table 3.4 presents the performance comparison[2] of various approaches for three dual-target CDR tasks according to HR@10 and NDCG@10. Note that since single-target CDR baselines aim to improve the recommendation performance on the sparser domain, we train them on both domains and then only present their experimental results on the sparser domain. From Table 3.4, we have the following observations: (1) Our proposed DIDA-CDR improves other baseline models on sparser domain by a large margin. Specifically, on sparser domain, it outperforms the best-performing baseline model with an average improvement of 9.55% in terms of HR@10 and 11.98% in terms of NDCG@10. This is because we adopt the interpolative data augmentation, which effectively generate both relevant and diverse augmented user representations to augment the sparser domain, and therefore significant improvements in recommendation performance can be obtained on the sparser domain; (2) Disentanglement-based dual-target CDR models improve conventional dual-target CDR models by an average of 14.98% in terms of HR@10 and 20.94% in terms of NDCG@10, which shows that decoupling and then transferring the domain-shared information to both domains is an efficient way for dual-target CDR; (3) Compared with other disentanglement-based dual-target CDR baselines, our model can achieve better recommendation performance over them. Specifically, our proposed DIDA-CDR achieves an average increase of 8.54% in terms of HR@10 and 11.10% in terms of NDCG@10, compared to the best-performing disentanglement-based dual-target CDR baseline. This is because our model particularly takes domain-independent user preferences into consideration, and the disentanglement module can more effectively disentangle all essential components of user preferences, thus capturing more comprehensive user preferences.

[2]Due to space limitation, we only show the results when $k = 64$ in Table 3.4. For the results under other values of k that are omitted, similarly, our model also has a significant improvement over other baseline models.

Table 3.4: Performance comparison (%) of different approaches for three dual-target CDR tasks according to HR@10 and NDCG@10 [185]. While the results of best-performing baselines are underlined, the best results are marked in bold (* indicates $p < 0.05$, paired t-test of our proposed DIDA-CDR vs. the best-performing baselines) [195].

Datasets	SDR Baselines				Single-Target CDR Baselines				Conventional Dual-Target CDR Baselines			
	NGCF		LightGCN		BPR_EMCDR _MLP		BPR_DCDCSR		PPGN		BiTGCF	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Douban-Movie	10.26	5.37	10.53	5.49	-	-	-	-	12.03	6.42	12.11	6.46
Douban-Book	7.31	4.08	7.35	4.15	6.25	3.93	6.74	4.02	10.52	4.78	10.58	4.93
Douban-Movie	9.16	4.23	9.24	4.25	-	-	-	-	10.09	4.35	10.14	4.41
Douban-Music	6.11	3.87	6.36	3.99	5.08	3.45	5.97	3.79	7.24	4.03	7.32	4.10
Amazon-Elec	20.22	11.97	20.03	10.94	-	-	-	-	22.06	12.44	21.79	12.31
Amazon-Cloth	10.95	6.01	11.38	6.10	9.87	5.33	10.90	5.86	13.04	6.91	13.16	6.88

Datasets	**Disentanglement-Based Dual-Target or Multi-Target CDR Baselines**						**Disentanglement-Based Dual-Target CDR (our)**								**Improvement**	
	GA-DTCDR		DR-MTCDR		DisenCDR		**DIDA-CDR _Fixed**		**DIDA-CDR _Base**		**DIDA-CDR _ELBO**		**DIDA-CDR**		**(DIDA-CDR vs. best baselines)**	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Douban-Movie	12.25	6.51	14.74	7.89	15.09	8.02	16.12	8.97	12.98	6.51	15.74	8.51	**16.66***	**9.16***	10.40%	14.21%
Douban-Book	10.71	5.06	12.66	7.45	12.40	7.27	13.08	7.82	10.41	5.52	12.80	7.69	**13.79***	**8.18***	8.93%	9.80%
Douban-Movie	10.35	4.57	11.27	5.74	12.13	5.95	12.76	6.23	10.45	4.99	12.55	6.18	**13.01***	**6.60***	7.25%	10.92%
Douban-Music	7.42	4.19	8.49	4.73	8.92	5.02	9.54	5.49	8.36	4.24	9.36	5.37	**9.97***	**5.79***	11.77%	15.34%
Amazon-Elec	23.87	13.20	22.34	12.98	23.77	13.61	24.68	13.83	22.03	12.97	24.14	13.72	**25.05***	**14.36***	4.94%	5.51%
Amazon-Cloth	13.94	7.09	14.13	7.59	15.46	8.42	16.01	8.94	12.97	6.93	15.76	8.63	**16.69***	**9.33**	7.96%	10.81%

3.3.2.2 Ablation Study (for RQ2)

To show the contribution of each proposed component to the improvement of overall performance, we modify our proposed model to form three variants and conduct an ablation study for three dual-target CDR tasks.

Impact of Interpolative Data Augmentation. We construct a variant of DIDA-CDR, namely **DIDA-CDR_Fixed**, by replacing the interpolative data augmentation module with a fixed mixing strategy. In fact, we conducted experiments to select the best-performing mixing coefficient, i.e., 0.5, from $\{0.1, 0.3, 0.5, 0.7, 0.9\}$ to implement the above variant. From Table 3.4, we can observe that with the interpolative data augmentation, our proposed DIDA-CDR outperforms **DIDA-CDR_Fixed** with an average improvement of 3.7%. This demonstrates that the interpolative data augmentation can not only augment the sparser domain by introducing randomness to increase the diversity of augmentation, but also generate representative augmented user representations by effectively mixing user embeddings for subsequent disentanglement. Meanwhile, the introduction of randomness also weakens the possible negative transfer caused by the linear interpolation operation with fixed weights, which can also be seen from the

experimental results.

Impact of User Preference Disentanglement. Furthermore, another variant, namely **DIDA-CDR_Base**, directly feeds the generated user embeddings to the information fusion module and does not include the disentanglement module, thus the variant only includes the prediction loss $\mathcal{L}_{prd}$, that is $\mu_1 = \mu_2 = 0$. From Table 3.4, we can observe that without the disentanglement module, the recommendation performance of **DIDA-CDR_Base** would degrade to be comparable to that of the conventional dual-target CDR baselines, and weaker than that of disentanglement-based ones. This shows that the disentanglement module can indeed help the model perform more effective cross-domain knowledge transfer without negative transfer by decoupling domain-shared, domain-specific, domain-independent information and transfering only domain-shared information, thus improving the performance of cross-domain recommendations.

Impact of Domain Classifier. In addition, following DisenCDR [8], we modify the disentanglement module in our model, i.e., replace our domain classification losses with the standard ELBO, to form another variant, namely **DIDA-CDR_ELBO**. From Table 3.4, we can observe that our DIDA-CDR outperforms **DIDA-CDR_ELBO** with an average improvement of 5.98%. This demonstrates that the proposed disentanglement module can indeed collaborate well with the domain classifier to decouple more accurate essential components of user preferences, especially the domain-independent information, to capture comprehensive user preferences, thus enabling the model to achieve better recommendation performance through superior disentanglement.

3.3.2.3 Impact of Various Components of User Preferences (for RQ3)

To demonstrate that all three components of user preferences, i.e., domain-share, domain-specific and domain-independent information, are essential and effective for recommendation, and do not require transfer of domain-independent user preferences, we compare DIDA-CDR with its four variants, including DIDA-CDR (w/o sha.), DIDA-CDR (w/o spe.), DIDA-CDR (w/o ind.) and DIDA-CDR (transfer ind.). Figures 3.2(a)-(b) show the performance comparison between our model and the above four

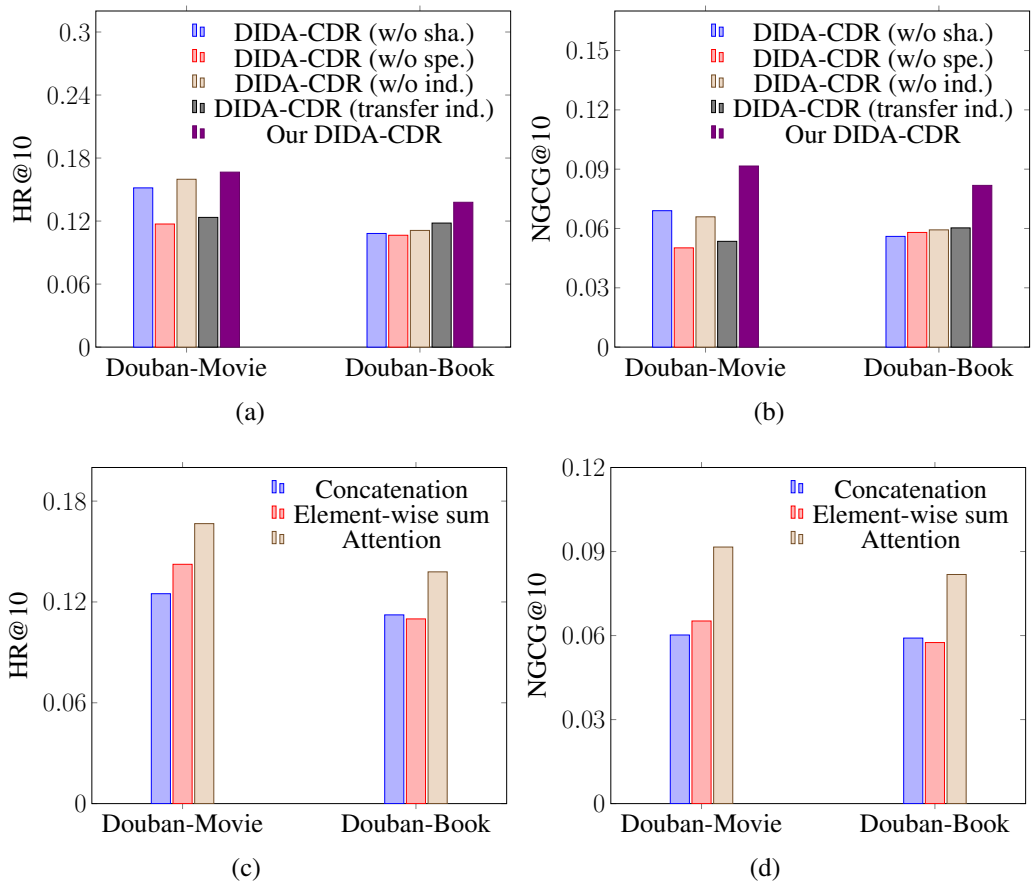

Figure 3.2: (a)-(b): Performance comparison between our model and its four variants. (c)-(d): Performance comparison of adopted information fusion approaches.

variants[3]. The differences between various variants and the impact of each component of user preferences are elaborated in the following subsections.

Impact of Domain-Shared User Preferences. DIDA-CDR (w/o sha.) extracts the domain-shared user preferences from both domains, but does not transfer them to any domain. From Figures 3.2(a)-(b), we can see that our proposed DIDA-CDR outperforms DIDA-CDR (w/o sha.) with an average improvement of 17.21%. This is because the domain-shared user preferences are valuable information, which plays an important role in cross-domain recommendation and can improve the recommenda-

[3]Due to space limitation, we only show the results on Task 1 in Figure 3.2, i.e., only the results on the pair of datasets consisting of Douban-Movie and Douban-Book are presented, and similar trends can be observed for results on the other omitted tasks. Similarly, Figure 3.3 and Figure 3.4 only shows the results on Task 1 for the same reason above.

tion performance on both domains simultaneously.

Impact of Domain-Specific User Preferences. DIDA-CDR (w/o spe.) does not consider the domain-specific user preferences when making recommendations. From Figures 3.2(a)-(b), we can observe that our proposed DIDA-CDR outperforms DIDA-CDR (w/o spe.) with an average improvement of 36.12%. This is because the domain-specific user preferences are inherent personalized preferences of users in each domain. If it is not considered when making recommendations, the recommendation performance of our model will be significantly reduced.

Impact of Domain-Independent User Preferences. DIDA-CDR (w/o ind.) includes the domain-shared and domain-specific user preferences, but does not include the domain-independent user preferences when capturing comprehensive user preferences; **DIDA-CDR (transfer ind.)** disentangles the domain-independent user preferences and then transfer them to another domain. It can be seen from Figures 3.2(a)-(b) that our proposed DIDA-CDR outperforms DIDA-CDR (w/o ind.) and DIDA-CDR (transfer ind.) with an average improvement of 12.40% and 26.04%, respectively. In view of this, we make the following qualitative analysis. The domain-independent user preferences seemingly exist in each domain, but actually have different meanings. If they are not included when capturing user preferences, the captured user preferences are incomplete, resulting in the degraded recommendation performance of the model. Moreover, if the domain-independent information is transferred to other domains, it will provide the useless information and cause the performance degradation.

3.3.2.4 Impact of Various Information Fusion Approaches (for RQ4)

After we obtain three essential user preference components, we compare various information fusion approaches, i.e., concatenation, element-wise sum, and attention, to fuse them into comprehensive user preferences. The performance comparison of three used information fusion approaches is presented in Figures 3.2(c)-(d). We find that when our model utilizes the attention mechanism for information fusion, it improves the variants using concatenation and summation by an average of 28.37% and 20.69%,

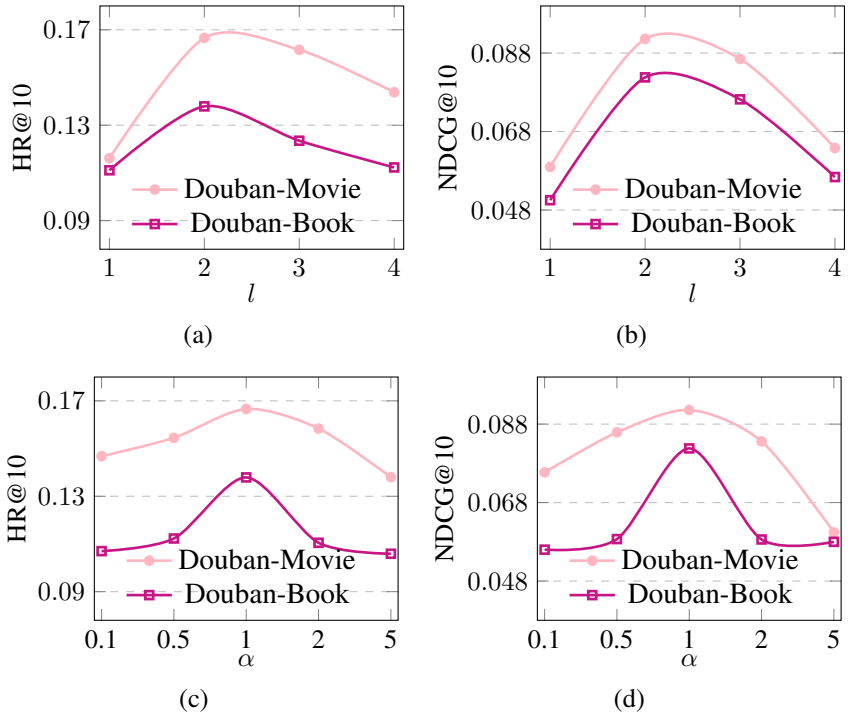

Figure 3.3: (a)-(b): Impact of the number of GCN layers. (c)-(d): Impact of α.

respectively. This is because the attention mechanism can not only capture the relationship between various components, but also selectively highlight the key information and weaken the redundant information by learning weights. In this paper, the attention mechanism measures the importance of domain-shared, domain-specific and domain-independent information to comprehensive user preferences through weights, and weakens redundant information between domain-shared and domain-independent user preferences, thus enabling our model to achieve better recommendation results.

3.3.2.5 Parameter Sensitivity (for RQ5)

Impact of l. Stacking too many layers when training a deep GCN is prone to over-smoothing [60]. In order to explore this effect, we investigate the number of GCN layers l in the range of $\{1, 2, 3, 4\}$ in the experiments and report the experimental results in Figures 3.3(a)-(b). It can be observed that 2-layer GCN is significantly better than 1-layer GCN, which indicates that stacking a moderate number of layers is bene-

ficial for mining higher-order user-item relationships. However, the recommendation performance of our proposed DIDA-CDR drops on some datasets when $l = 3$ and degrades even more when $l = 4$. The reason may be that when the number of layers is greater than 3, the problem of over-smoothing occurs, resulting in the fact that increasing the number of layers at this time will in turn reduce the recommendation performance of our model.

Impact of α. $Beta(\alpha, \alpha)$ is Uniform distribution when $\alpha = 1$, Bimodal distribution when $\alpha < 1$ and Bell-shaped distribution when $\alpha > 1$ [10]. In order to explore from which distribution sampling λ can help the recommendation performance of model the most, we search α in the range of $\{0.1, 0.5, 1, 2, 5\}$. The performance comparison is illustrated in Figures 3.3(c)-(d). We can see that the best performance of our proposed DIDA-CDR is achieved when $\alpha = 1$, which presents that λ sampled from the Uniform distribution can effectively mix the user embeddings in the richer domain and the sparser domain, thereby effectively alleviating the data imbalance and improving the recommendation performance of model. In contrast, when $\alpha < 1$ or $\alpha > 1$, sampling λ from $Beta(\alpha, \alpha)$ will result in the performance degradation.

Impact of μ_1 and μ_2. To explore the effect of weights of domain classification losses on the recommendation performance of our DIDA-CDR, we vary μ_1 and μ_2 in $\{0.1, 0.3, 0.5, 0.7, 1, 3, 5, 10\}$. The results are reported in Figure 3.4. It can be seen from Figures 3.4(a)-(b) that the contribution of domain classification loss $\mathcal{L}_{cls_1}$ to decoupling more accurate domain-specific user preferences is small when $\mu_1 \to 0$. When μ_1 is too large, the domain classification loss $\mathcal{L}_{cls_1}$ receives more attention during the learning process. As a result, the contribution of prediction loss $\mathcal{L}_{prd}$ to the model is weakened, which reduces the recommendation performance of the model. Similarly, from Figures 3.4(c)-(d), a similar trend can be observed for the weights μ_2 of domain classification loss L_{cls_2}. Empirically, we choose $\mu_1 = \mu_2 = 1$.

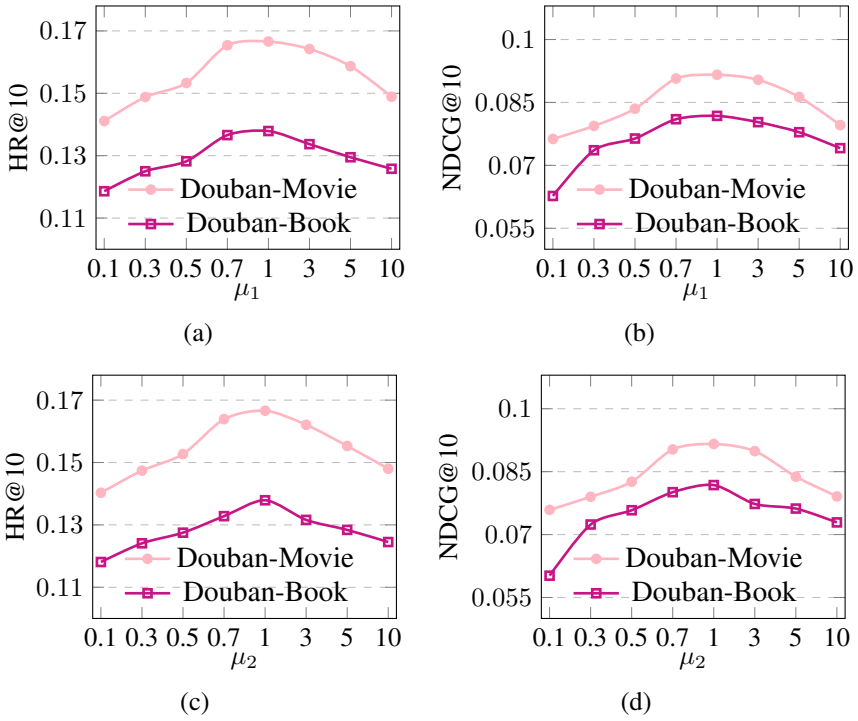

Figure 3.4: (a)-(b): Impact of μ_1. (c)-(d): Impact of μ_2.

3.4 Summary

In this chapter, we have proposed a novel Disentanglement-based framework with Interpolative Data Augmentation for dual-target Cross-Domain Recommendation, called DIDA-CDR. DIDA-CDR consists of an interpolative data augmentation approach to generating both relevant and diverse augmented user representations to augment the sparser domain and explore the potential user preferences, and a user preference disentanglement module to decouple essential components of user preferences to capture comprehensive user preferences, all of which help improve the recommendation performance on both domains simultaneously. Also, we have conducted extensive experiments on five real-world datasets to show the significant superiority of DIDA-CDR over the state-of-the-art methods.

Chapter 4

Causal Deconfounding via Confounder Disentanglement for Dual-Target Cross-Domain Recommendation

Building upon the DIDA-CDR framework introduced in Chapter 3, which focuses on disentangling all essential components of user preferences to capture comprehensive user preferences, this chapter extends the understanding of CDR by addressing a crucial aspect previously underexplored: the role of observed confounders. In this chapter, we further reveal that both single-domain and cross-domain confounders also affect users' final decisions in addition to their true preferences. To this end, an effective dual-target CDR should *deconfound both observed single-domain and cross-domain confounders*, which includes three tasks, namely, (1) identify and decouple such observed confounders, (2) preserve their positive impacts on predicted interactions, and (3) eliminate their negative impacts on user preferences [149]. However, existing dual-target CDR approaches overlook the above observations. Thus, a novel dual-target CDR model should be proposed to incorporate such insights for comprehensively understanding user-item interactions.

To effectively advance dual-target CDR, the following two key challenges need to be addressed.

CH2. *How to effectively extract observed confounders to comprehensively understand user-item interactions?* In the context of dual-target CDR, this challenge can be fur-

ther specified as: *How to effectively decouple observed cross-domain confounders in addition to single-domain confounders to comprehensively understand user-item interactions in dual-target CDR?*

CH3. *How to preserve the positive impacts of observed confounders on predicted interactions, while eliminating their negative impacts on capturing comprehensive user preferences, thereby enhancing the recommendation accuracy?*

This chapter first presents the problem definition and analyzes the problem from a causal perspective. Then, to address the above challenges, CD2CDR is proposed, including three key phases. Finally, we conduct extensive experiments on seven real-world datasets. The experimental results demonstrate that our CD2CDR significantly outperforms the state-of-the-art methods.

4.1 Problem Statement

4.1.1 Problem Definition

The chapter explores a fully overlapping dual-target CDR scenario in the domains D^A and D^B, with a common user set $\mathcal{U}$, the size of which is denoted as $m = |\mathcal{U}|$. Let $\mathcal{V}^A$ (of size $n^A = |\mathcal{V}^A|$) and $\mathcal{V}^B$ (of size $n^B = |\mathcal{V}^B|$) denote the item sets in the domains D^A and D^B, respectively. The raw feature vector of each item in D^A (or D^B) is defined as $\mathbf{E}^A_{vr} \in \mathbb{R}^{d^A}$ (or $\mathbf{E}^B_{vr} \in \mathbb{R}^{d^B}$), where d^A (or d^B) is the dimensionality of features. The interaction matrices are denoted as $\mathbf{R}^A \in \{0,1\}^{m \times n^A}$ and $\mathbf{R}^B \in \{0,1\}^{m \times n^B}$ in D^A and D^B, respectively.

To improve the performance of dual-target CDR, it is crucial to explicitly consider the impacts of observed confounders. These confounders include single-domain confounders $\mathbf{C}_{sd}$ and cross-domain confounders $\mathbf{C}_{cd}$, both of which simultaneously influence user preferences and user-item interactions. Addressing the impacts of these confounders necessitates significant adjustments to existing dual-target CDR models. Hence, it would be beneficial to propose a novel deconfounding framework that is

highly extendable and compatible with most off-the-shelf dual-target CDR models. For this purpose, since DIDA-CDR [188] is a representative and state-of-the-art dual-target CDR model, we choose it as the foundation for our problem definition.

DIDA-CDR has effectively decoupled three essential components of user preferences for modeling comprehensive user preferences $\mathbf{E}_u^*$, thus achieving good recommendation results in each domain. These three components include:

(1) *domain-shared user preferences* $\mathbf{Z}_{sha}$, which have the same meaning in each of both domains. For instance, users might prefer items in the sports 'category', which is the domain-shared preference covering both the 'purchase' and 'add to favorite' domains, reflecting consistent preferences across both domains.

(2) *domain-specific user preferences* $\mathbf{Z}_{spe}$, which are unique to one domain. For example, in the 'add to favorite' domain, users might prefer 'luxurious' items that they cannot afford to purchase but still wish to add them to favorite, while in the 'purchase' domain, users might prefer 'practical' items that offer good value for money.

(3) *domain-independent user preferences* $\mathbf{Z}_{ind}$, which are seemingly common in both domains but have different meanings in each domain [188]. For instance, in the 'purchase' domain, a preference for 'professional features' refers to choosing items that are specialized and match the user's current skill level or needs. Specifically, a beginner photography enthusiast might purchase an entry-level professional digital camera, which can mount different lenses for learning photography, emphasizing practicality and suitability for immediate use. By contrast, in the 'add to favorite' domain, a preference for 'professional features' reflects an aspiration for high-end items with advanced features, such as professional lenses, which are added to favorite for potential future use when the photography enthusiast's skills improve. Unlike domain-specific user preferences that only exist in their corresponding domain, domain-independent user preferences exist in both domains but have different meanings in each domain.

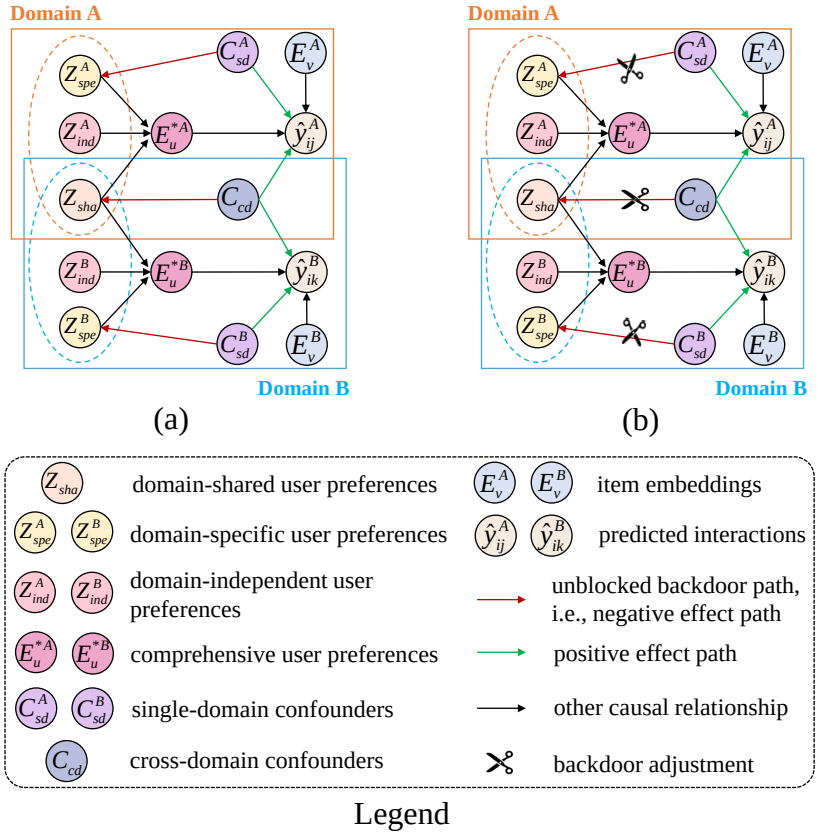

Figure 4.1: Causal graphs of dual-target CDR for deconfounding observed SDCs and CDCs. (a) Original causal graph. (b) Deconfounded causal graph after eliminating such observed confounders' negative effects by blocking backdoor paths via backdoor adjustment, as indicated by scissors [130].

Based on the above notations, the problem of Causal Deconfounding for Dual-Target CDR is defined as follows.

Causal Deconfounding for Dual-Target CDR. *Given the domain-specific and comprehensive user preferences (i.e.,* $\mathbf{Z}_{spe}$ *and* $\mathbf{E}_u^*$*) in each domain, the goal of causal deconfounding for dual-target CDR is to decouple observed single-domain confounders* $\mathbf{C}_{sd}$ *and cross-domain confounders* $\mathbf{C}_{cd}$*, eliminate such observed confounders' negative effects to obtain debiased comprehensive user preferences, and incorporate these confounders' positive effects into such debiased preferences to achieve a comprehensive understanding of user-item interactions, thus enhancing the recommendation accuracy in both domains.*

4.1.2 Causal Graph

A causal graph, i.e., a directed acyclic graph (DAG), where edges represent causal relationships between variables. Taking cross-domain confounders $\mathbf{C}_{cd}$ as an example, as illustrated in Figure 4.1, they affect predicted interactions $\hat{y}$ via two types of paths: $\mathbf{C}_{cd} \rightarrow \hat{y}$ and $\mathbf{C}_{cd} \rightarrow \mathbf{Z}_{sha} \rightarrow \mathbf{E}_u^* \rightarrow \hat{y}$. The first type of path reveals that $\mathbf{C}_{cd}$, even if not the primary cause, i.e., users' true preferences, still have a direct positive impact on predicted interactions. The second type of path indicates that the negative impact of $\mathbf{C}_{cd}$ on domain-shared user preferences $\mathbf{Z}_{sha}$ induces confounding bias in both domains. Such confounding bias, in turn, skews comprehensive user preferences $\mathbf{E}_u^*$, because $\mathbf{Z}_{sha}$ is an essential component for capturing $\mathbf{E}_u^*$ [188]. If the backdoor path $\mathbf{C}_{cd} \rightarrow \mathbf{Z}_{sha}$ is not blocked, $\mathbf{C}_{cd}$ will result in capturing biased comprehensive user preferences, thus yielding inaccurate recommendation results [154]. Similarly, single-domain confounders $\mathbf{C}_{sd}$ also have both positive and negative effects on predicted interactions and user preferences, respectively, thus the backdoor path $\mathbf{C}_{sd} \rightarrow \mathbf{Z}_{spe}$ in each domain should be blocked too.

Overall, the causal graph in Figure 4.1 provides a detailed view of how user preferences, observed confounders, and user-item interactions are causally related in dual-target CDR. In this study, we focus on addressing the confounding bias introduced by observed confounders in cross-domain settings. Firstly, we effectively decouple observed single-domain and cross-domain confounders. Secondly, we perform backdoor adjustment to preserve the positive direct effects of such observed confounders on predicted interactions and eliminate their negative effects on capturing comprehensive user preferences. These steps mitigate confounding biases to a large extent, enable a comprehensive understanding of user-item interactions, and thus improve the recommendation performance in both domains.

4.2 The Proposed CD2CDR Framework

4.2.1 Framework Overview

To enhance the recommendation accuracy in both domains, we propose a novel **C**ausal **D**econfounding framework via **C**onfounder **D**isentanglement for dual-target **C**ross-**D**omain **R**ecommendation, called CD2CDR. As depicted in Figure 4.2, the framework can be divided into three phases, i.e., **Phase 1:** User Preference Disentanglement Pre-Training, **Phase 2:** Confounder Disentanglement, and **Phase 3:** Causal Deconfounding and Cross-Domain Recommendation. In **Phase 1**, we obtain disentangled domain-independent and domain-specific user preferences in each domain and domain-shared user preferences by pre-training the backbone introduced in [188]. In **Phase 2**, we first extract the SDCs in each domain by bidirectionally transforming domain-specific user preferences decoupled in Phase 1. Then, we distill confounding factors that jointly influence comprehensive user preferences in each of both domains as CDCs by adopting half-sibling regression [101]. In **Phase 3**, we utilize the backdoor adjustment to deconfound the observed confounders disentangled in **Phase 2**. Specifically, we design a confounder selection function to mitigate negative effects of such confounders on user preferences, thus recovering debiased comprehensive user preferences. We then incorporate the observed confounders' positive effects into such debiased preferences to predict user-item interactions via a multi-layer perceptron (MLP) in both domains.

4.2.2 Phase 1: User Preference Disentanglement Pre-Training

Accurate disentanglement of user preferences is vital to ensure the robustness of subsequent confounder disentanglement process. Since the method introduced in [188] excels at decoupling three essential components of user preferences for modeling comprehensive user preferences, our CD2CDR adopts it as the backbone for user preference disentanglement. To extract more accurate disentangled user preferences, we consider multi-source content information of users and items, e.g., user reviews and

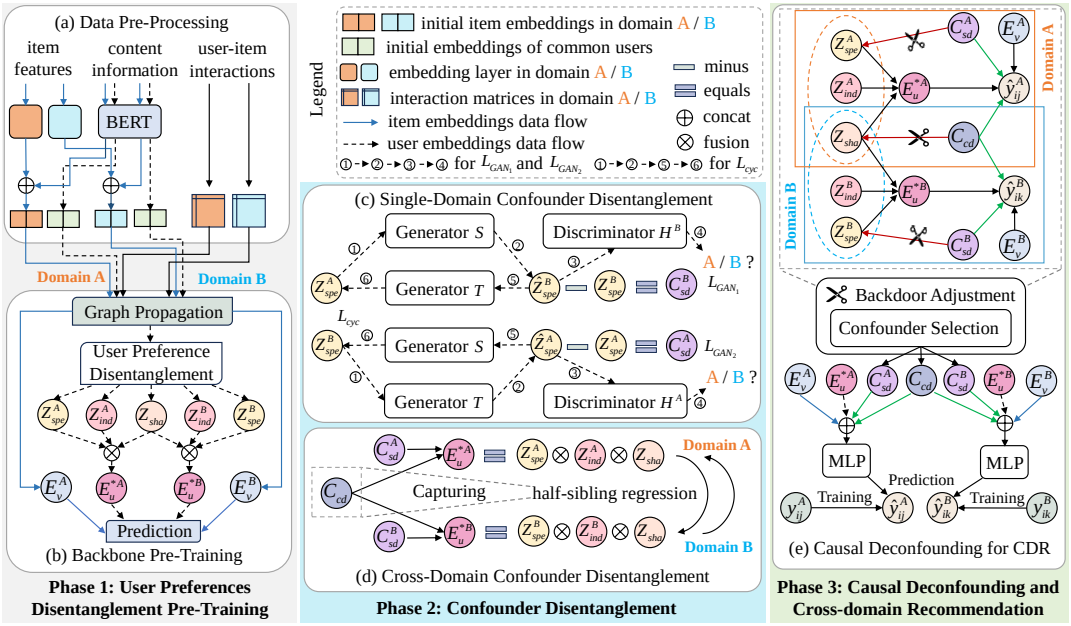

Figure 4.2: The structure of our CD2CDR framework.

item details. Taking domain A as an example, for each categorical feature field of an item (e.g., brand and category), we distill a set of unique features, which are then encoded into vectors using either one-hot or multi-hot encoding. Next, these encoded vectors are concatenated to form the raw feature vector for each item. We then transform the raw feature vectors of items $\mathbf{E}^A_{vr} \in \mathbb{R}^{d^A}$ into the dense embeddings $\mathbf{E}^A_{vd} \in \mathbb{R}^d$ as follows:

$$\mathbf{E}^A_{vd} = \mathbf{W}^A_{rd}\mathbf{E}^A_{vr}, \tag{4.1}$$

where $\mathbf{W}^A_{rd} \in \mathbb{R}^{d^A \times d_d}$ is a trainable mapping matrix. d_d denotes the dimensionality of dense embeddings. Then, for a user u_i, we collect all the user's reviews into a user text document. For an item v_j, we collect its title and all reviews on the item into an item text document. Next, we adopt a pre-trained BERT [17] to map the documents of all users and items in the training set into user text embeddings $\mathbf{E}^A_{ut}$ and item text embeddings $\mathbf{E}^A_{vt}$, respectively. Finally, we concatenate $\mathbf{E}^A_{vd}$ and $\mathbf{E}^A_{vt}$ to form combined item embeddings $\mathbf{E}^A_{vc}$. We then transform $\mathbf{E}^A_{ut}$, $\mathbf{E}^A_{vc}$ into fixed-size initial user embeddings

$\mathbf{E}_{ui}^{A}$ and initial item embedding $\mathbf{E}_{vi}^{A}$ in domain A as follows:

$$\mathbf{E}_{ui}^{A} = \delta_{u}^{A}\mathbf{E}_{ut}^{A}, \quad \mathbf{E}_{vi}^{A} = \delta_{v}^{A}\mathbf{E}_{vc}^{A}, \tag{4.2}$$

where δ_{u}^{A} and δ_{v}^{A} are the mapping functions of MLP layers. Similarly, we can obtain initial user embeddings $\mathbf{E}_{ui}^{B}$ and initial item embeddings $\mathbf{E}_{vi}^{B}$ in domain B. We then leverage such initial embeddings and the interaction matrices as inputs to pretrain the backbone. Specifically, we aggregate interaction data within each domain to build two heterogeneous graphs, which allow us to learn coarse user and item embeddings for each domain. Next, we apply linear interpolation to the user embeddings of both domains to generate augmented user representations, augmenting the sparser domain. With these coarse user embeddings and augmented user representations, we then employ a user preference disentanglement module, guided by a domain classifier, to decouple domain-independent, domain-specific, and domain-shared user preferences, namely, $\mathbf{Z}_{ind}$, $\mathbf{Z}_{spe}$, and $\mathbf{Z}_{sha}$ (for more information, please refer to [188]). By incorporating the above three components of user preferences using attention mechanism in accordance with DIDA-CDR [188], we obtain comprehensive user preferences $\mathbf{E}_{u}^{*}$.

4.2.3 Phase 2: Confounder Disentanglement

Since user-item interactions are also influenced by observed confounders apart from comprehensive user preferences, we propose to decouple SDCs and CDCs, as detailed in the following subsections.

4.2.3.1 Single-Domain Confounder Disentanglement

To explore the SDCs, we utilize bi-directional domain transformation to decouple them from previously obtained domain-specific user preferences. If such SDCs are not identified and explicitly decoupled, their negative effects on domain-specific user preferences can hardly be eliminated. By contrast, if they are well disentangled, the causal deconfounding module can utilize backdoor adjustment to remove the con-

founding bias, thus obtaining debiased domain-specific user preferences. Inspired by CycleGAN [190], we devise a dual adversarial structure, which consists of two domain transformation generators and two discriminators, to disentangle SDCs in each domain. Specifically, we aim to learn two generators, i.e., $S(\cdot) : D^A \rightarrow D^B$ and $T(\cdot) : D^B \rightarrow D^A$, to transform domain-specific user preferences in each domain.

Taking domain B as an example, the generator $S(\cdot)$ takes the domain-specific user preferences $\mathbf{Z}^A_{spe}$ of common users in D^A as inputs to generate $\hat{\mathbf{Z}}^B_{spe} = S(\mathbf{Z}^A_{spe})$ that look similar to domain-specific user preferences in domain B, i.e., $\mathbf{Z}^B_{spe}$. However, if there are still differences between the simulated preferences $\hat{\mathbf{Z}}^B_{spe}$ and the original ones $\mathbf{Z}^B_{spe}$, such differences are not characteristics of domain-specific user preferences in domain B, but should be considered as SDCs [161]. To ensure that the generator $S(\cdot)$ is proficient at domain-specific preference simulation, we introduce a discriminator $H^B(\cdot)$ to recognize which domain the domain-specific user preferences come from. In the adversarial learning paradigm, the discriminator is expected to improve the ability to differentiate domain-specific user preferences in each domain to achieve better discriminative performance, while the generator is supposed to generate indistinguishable simulations of these domain-specific preferences to confuse such discriminator [107]. For training the generator $S(\cdot)$ and the corresponding discriminator $H^B(\cdot)$, we adopt the adversarial loss [31], which can be expressed as follows:

$$\begin{aligned}\mathcal{L}_{GAN_1}(S, H^B, D^A, D^B) &= \mathbb{E}_{\mathbf{Z}^B_{spe}\sim\mathbb{P}^B}[\log H^B(\mathbf{Z}^B_{spe})] \\ &+ \mathbb{E}_{\mathbf{Z}^A_{spe}\sim\mathbb{P}^A}[\log(1 - H^B(S(\mathbf{Z}^A_{spe})))],\end{aligned} \tag{4.3}$$

where $\mathbb{E}$ is the expectation, and $\mathbb{P}^A$, $\mathbb{P}^B$ denote the feature distribution of domain A and domain B, respectively. Similarly, for training the generator $T(\cdot)$ and the corresponding discriminator $H^A(\cdot)$, we adopt the adversarial loss $\mathcal{L}_{GAN_2}(T, H^A, D^B, D^A)$. However, relying solely on adversarial loss is insufficient to ensure that a user's domain-specific preferences remain aligned with the user's preferences after transformation. If transformed domain-specific user preferences no longer reflect the user's prefer-

ences, then such transformation becomes meaningless, serving merely to confuse the discriminator. Hence, the generators should maintain cycle consistency, i.e., $\mathbf{Z}^A_{spe} \rightarrow S(\mathbf{Z}^A_{spe}) \rightarrow T(S(\mathbf{Z}^A_{spe})) \approx \mathbf{Z}^A_{spe}$ and $\mathbf{Z}^B_{spe} \rightarrow T(\mathbf{Z}^B_{spe}) \rightarrow S(T(\mathbf{Z}^B_{spe})) \approx \mathbf{Z}^B_{spe}$ during the training process (see $\mathcal{L}_{cyc}$ in Figure 4.2(c)). To this end, we apply a cycle consistency loss, which is represented as follows:

$$\begin{aligned}\mathcal{L}_{cyc}(S,T) = \mathbb{E}_{\mathbf{Z}^A_{spe}\sim\mathbb{P}^A}[||T(S(\mathbf{Z}^A_{spe})) - Z^A_{spe}||_1] \\ + \mathbb{E}_{\mathbf{Z}^B_{spe}\sim\mathbb{P}^B}[||S(T(\mathbf{Z}^B_{spe})) - \mathbf{Z}^B_{spe}||_1].\end{aligned} \tag{4.4}$$

Moreover, the total objective function for training the generators and discriminators can be formulated as follows:

$$\begin{aligned}\mathcal{L}_{sd}(S,T,H^A,H^B) = \mathcal{L}_{GAN_1}(S,H^B,D^A,D^B) \\ + \mathcal{L}_{GAN_2}(T,H^A,D^B,D^A) + \lambda\mathcal{L}_{cyc}(S,T),\end{aligned} \tag{4.5}$$

where λ controls the importance of cycle consistency loss relative to adversarial losses. Following the method introduced in [161], upon training completion, we calculate the differences between the domain-specific user preferences after transformation and the original ones as candidate SDCs. Even though confounding bias may still exist in the original domain-specific user preferences, the differences calculation helps to decouple candidate SDCs. By performing deconfounding on these decoupled confounders, the impact of such biases can be mitigated to a large extent. The differences are defined as follows:

$$\hat{\mathbf{C}}^A_{sd} = T(\mathbf{Z}^B_{spe}) - \mathbf{Z}^A_{spe}, \quad \hat{\mathbf{C}}^B_{sd} = S(\mathbf{Z}^A_{spe}) - \mathbf{Z}^B_{spe}. \tag{4.6}$$

Recall the causal graph in Figure 4.1(a), the negative effects of single-domain confounders $\mathbf{C}_{sd}$ on domain-specific user preferences $\mathbf{Z}_{spe}$ result in confounding bias, leading to inaccurate estimation of $\mathbf{Z}_{spe}$. For example, as depicted in Figure 1.1(c), in the 'purchase' domain, a data-driven RS improperly perceives the 'free shipping' (i.e., an SDC), as Alice's 'purchase' domain preference, resulting in biased recom-

mendation. Since our well-trained generator excels at simulating Alice's 'purchase' domain preferences based on her 'add to favorite' domain preferences, if there are still differences as per Eq. (4.6), this indicates such differences are not Alice's 'purchase' domain preferences but rather SDCs independent of her preferences. Such SDCs (e.g., 'free shipping'), previously entangled with Alice's 'purchase' domain preferences, are decoupled through our SDC disentanglement process. Note that this process specifically targets biased domain-specific user preferences, because unbiased ones are not entangled with such SDCs. Thus, although this process decouples SDCs from biased domain-specific user preferences, this does not imply a causal relationship $\mathbf{Z}_{spe} \rightarrow \mathbf{C}_{sd}$ in the causal graph, because SDCs are not generated by such biased domain-specific user preferences. To distill representative SDCs and reduce redundancy, we apply K-means clustering on candidate single-domain confounders $\hat{\mathbf{C}}_{sd}^{A}$ (or $\hat{\mathbf{C}}_{sd}^{B}$) and choose J_{sd}^{A} (or J_{sd}^{B}) cluster centroids to form the potential SDC subspace $\mathcal{C}_{sd}^{A}$ (or $\mathcal{C}_{sd}^{B}$).

4.2.3.2 Cross-Domain Confounder Disentanglement

In addition to SDCs, it is more important to identify confounding factors that simultaneously affect user-item interactions in both domains. Inspired by the method introduced in [161], we employ half-sibling regression to disentangle CDCs from the previously obtained comprehensive user preferences in both domains. Half-sibling regression excels at capturing the influence of confounding factors that simultaneously affect multiple observed variables [101], and thus it is well suited for decoupling CDCs in dual-target CDR. As illustrated in Figure 4.1(a), $\mathbf{C}_{cd} \rightarrow \mathbf{E}_{u}^{*A}$ and $\mathbf{C}_{cd} \rightarrow \mathbf{E}_{u}^{*B}$ indicate that CDCs indirectly influence the comprehensive user preferences in each of both domains via $\mathbf{C}_{cd} \rightarrow \mathbf{Z}_{sha} \rightarrow \mathbf{E}_{u}^{*A}$ and $\mathbf{C}_{cd} \rightarrow \mathbf{Z}_{sha} \rightarrow \mathbf{E}_{u}^{*B}$. The core idea of half-sibling regression is: if $\mathbf{E}_{u}^{*A}$ and $\mathcal{C}_{sd}^{B}$ are independent, then predicting $\mathbf{E}_{u}^{*B}$ based on $\mathbf{E}_{u}^{*A}$ becomes a method to selectively capture the influence of $\mathbf{C}_{cd}$ on $\mathbf{E}_{u}^{*B}$ (see Figure 4.2(d)). Similarly, predicting $\mathbf{E}_{u}^{*A}$ based on $\mathbf{E}_{u}^{*B}$ serves to capture the influence of $\mathbf{C}_{cd}$ on $\mathbf{E}_{u}^{*A}$ (for more information, please refer to [101]). Therefore, we can apply half-sibling regression to decouple $\mathbf{C}_{cd}$ from $\mathbf{E}_{u}^{*A}$ and $\mathbf{E}_{u}^{*B}$. Taking the regression from

domain A to domain B as an example, our goal is to estimate a transformation matrix $\mathbf{W}^{A\to B}$ such that:

$$\mathbf{E}_u^{*B} \approx \mathbf{E}_u^{*A}\mathbf{W}^{A\to B}, \tag{4.7}$$

using ridge regression, and regression results are expressed as:

$$\mathbf{W}^{A\to B} = [(\mathbf{E}_u^{*A})^\top \mathbf{E}_u^{*A} + \alpha\mathbf{I}]^{-1}(\mathbf{E}_u^{*A})^\top \mathbf{E}_u^{*B}, \tag{4.8}$$

where α denotes the regularization parameter. We assume that $\mathbf{E}_u^{*A}$ and $\mathbf{C}_{sd}^{B}$ are independent, because $\mathbf{E}_u^{*A}$ are comprehensive user preferences in domain A, while $\mathbf{C}_{sd}^{B}$ are SDCs specific to domain B. When we estimate a transformation matrix $\mathbf{W}^{A\to B}$ to predict $\mathbf{E}_u^{*B}$ using $\mathbf{E}_u^{*A}$, the influence of $\mathbf{C}_{sd}^{B}$ on $\mathbf{E}_u^{*B}$ will not be captured. This is because $\mathbf{E}_u^{*A}$ are independent from $\mathbf{C}_{sd}^{B}$, and as a result, utilizing $\mathbf{E}_u^{*A}$ cannot predict $\mathbf{C}_{sd}^{B}$ and the influence of $\mathbf{C}_{sd}^{B}$ on $\mathbf{E}_u^{*B}$. By contrast, the influence of $\mathbf{C}_{cd}$ on $\mathbf{E}_u^{*B}$ will be captured, because $\mathbf{C}_{cd}$ simultaneously affect $\mathbf{E}_u^{*A}$ and $\mathbf{E}_u^{*B}$, which means the regression results will only capture $\mathbf{C}_{cd}$. Hence, the regression results can be identified as candidate cross-domain confounders:

$$\hat{\mathbf{C}}_{cd}^{A\to B} = \mathbf{E}_u^{*A}\mathbf{W}^{A\to B}. \tag{4.9}$$

Similarly, we can obtain the regression results from domain B to domain A, denoted as $\hat{\mathbf{C}}_{cd}^{B\to A}$. For cross-domain confounders, K-means clustering is also employed on the candidate cross-domain confounders $\hat{\mathbf{C}}_{cd}^{A\to B}$ and $\hat{\mathbf{C}}_{cd}^{B\to A}$, with the J_{cd} cluster centroids forming the potential CDC subspace $\mathcal{C}_{cd}$.

4.2.4 Phase 3: Causal Deconfounding and Cross-Domain Recommendation

After the confounder disentanglement, we obtain the potential SDC subspaces $\mathcal{C}_{sd}^{A}$ and $\mathcal{C}_{sd}^{B}$, and potential CDC subspace $\mathcal{C}_{cd}$. From a causal perspective, if the backdoor paths

(i.e., $\mathbf{C}_{sd} \rightarrow \mathbf{Z}_{spe}$ and $\mathbf{C}_{cd} \rightarrow \mathbf{Z}_{sha}$) are not blocked, the observed confounders C will simultaneously influence user preferences Z and user-item interactions Y (see Figure 1.1(a)), and thus cause biased estimation of comprehensive user preferences. To this end, we perform the *do*-calculus intervention based on backdoor adjustment [193] to block the backdoor paths $C \rightarrow Z$ and enable our model to more accurately estimate the direct effect $Z \rightarrow Y$ (also see Figure 1.1(a)). Formally, the conventional likelihood $P(Y|Z)$ is defined as:

$$P(Y|Z) = \sum\nolimits_c P(Y|Z,c)P(c|Z), \tag{4.10}$$

where c denotes a specific confounder selected from the confounder space $\mathcal{C}$. By applying the *do*-calculus, we exclude all influences directed towards the intervened variable (i.e., Z), and then we have:

$$\begin{aligned} P(Y|do(Z)) &= \sum\nolimits_c P(Y|do(Z),c)P(c|do(Z)) \\ &= \sum\nolimits_c P(Y|Z,c)P(c). \end{aligned} \tag{4.11}$$

For brevity, the detailed proof of the transformations $P(Y|do(Z),c) = P(Y|Z,c)$ and $P(c|do(Z)) = P(c)$ is omitted, which can be found in [95]. In fact, transforming $P(c|do(Z))$ into the prior probability of confounders $P(c)$ blocks backdoor paths $C \rightarrow Z$. As a result, $P(Y|do(Z))$ mainly focus on modeling the direct effect $Z \rightarrow Y$. Specifically, we implement the backdoor adjustment by modeling $P(Y|Z,c)$ with an interaction prediction network, which is expressed as follows:

$$P(Y|do(Z)) = \mathbb{E}_c[P(Y|Z,c)] = \mathbb{E}_c[f(\mathbf{E}_u^*, \mathbf{E}_v, \mathbf{c})], \tag{4.12}$$

where $f(\cdot)$ denotes a neural network, namely, MLP, to predict the probabilities of user-item interactions [41]. $\mathbf{E}_u^*$ and $\mathbf{E}_v$ are comprehensive user preferences and pre-trained item embeddings obtained by the backbone in Phase 1, respectively. In other words, based on two subspaces of disentangled observed confounders in domain A and do-

main B, i.e., $\mathcal{C}^A = \mathcal{C}^A_{sd} \cup \mathcal{C}_{cd}$ and $\mathcal{C}^B = \mathcal{C}^B_{sd} \cup \mathcal{C}_{cd}$, we apply backdoor adjustment to rectify the biased recommendations in each domain using Eq. (4.12). Since the decoupled observed confounders are incorporated as part of the input to MLP, the direct influence of such confounders on user-item interactions $C \rightarrow Y$ is also considered. Moreover, inspired by [161], we devise a confounder selection function to effectively control decoupled confounders for more accurate deconfounding.

Taking domain A as an example, the confounder selection function can be defined as follows:

$$\begin{aligned}\phi(\mathbf{E}^{*A}_u, \mathbf{E}^A_v, \mathbf{c}) = &\frac{\exp(\mathbf{W}^A_u \mathbf{E}^{*A}_u \cdot \mathbf{W}^A_{uc}\mathbf{c})}{2\sum_{\mathbf{c}'}\exp(\mathbf{W}^A_u \mathbf{E}^{*A}_u \cdot \mathbf{W}^A_{uc}\mathbf{c}')} \\ &+ \frac{\exp(\mathbf{W}^A_v \mathbf{E}^A_v \cdot \mathbf{W}^A_{vc}\mathbf{c})}{2\sum_{\mathbf{c}'}\exp(\mathbf{W}^A_v \mathbf{E}^A_v \cdot \mathbf{W}^A_{vc}\mathbf{c}')},\end{aligned} \tag{4.13}$$

where $\mathbf{c}'$ denotes any confounder selected from confounder subspace $\mathcal{C}^A$ and $\cdot$ is the dot product. $\mathbf{W}^A_u$, $\mathbf{W}^A_{uc}$, $\mathbf{W}^A_v$, $\mathbf{W}^A_{vc}$ are trainable matrices for embedding transformation. We then formulate the expectation $\mathbb{E}_c[f(\mathbf{E}^*_u, \mathbf{E}_v, \mathbf{c})]$ as follows:

$$\begin{aligned}&\mathbb{E}_c[f(\mathbf{E}^*_u, \mathbf{E}_v, \mathbf{c})] = f(\mathbf{Q}^A_{in}) \\ &= f[\mathbf{W}_{fc}(\mathbf{E}^{*A}_u || \mathbf{E}^A_v || \sum\nolimits_c p(c)\mathbf{c}\phi(\mathbf{E}^{*A}_u, \mathbf{E}^A_v, \mathbf{c})],\end{aligned} \tag{4.14}$$

where $\mathbf{W}_{fc}$ is the weight matrix of the fully connected (FC) layer and $||$ is the concatenation operator. In practice, we assume a uniform distribution for the prior probability $p(c)$. In addition, $\mathbf{Q}^A_{in} = \mathbf{W}_{fc}(\mathbf{E}^{*A}_u || \mathbf{E}^A_v || \sum_c p(c)\mathbf{c}\phi(\mathbf{E}^{*A}_u, \mathbf{E}^A_v, \mathbf{c})$ denotes the input for MLP in domain A. Moreover, the predicted interaction $\hat{y}^A_{ij}$ between an user u_i and an item v_j in domain A is represented as follows:

$$\hat{y}^A_{ij} = \delta^A_{out}(\delta^A_l(...\delta^A_2(\delta^A_1(\mathbf{Q}^A_{in}))...)), \tag{4.15}$$

where δ^A_l is the mapping function for l-th MLP layer, and there are l MLP layers including δ^A_{out} in domain A. Similarly, the predicted interaction $\hat{y}^B_{ij}$ in domain B can

be obtained.

The essence of our causal deconfounding module lies in blocking the backdoor paths $C \rightarrow Z$, allowing the model to concentrate on the direct effects of users' true preferences on the predicted interactions $Z \rightarrow Y$, and disregard the interference of observed confounders on these preferences. Specifically, the confounder selection function assigns different weights to the potential observed confounders, mitigates the effects of those irrelevant or harmful confounders to the prediction task, and enhances the direct effects of beneficial confounders on predicted interactions. Therefore, this module enables the model to eliminate the negative effects of such observed confounders to learn debiased comprehensive user preferences, and preserve the positive effects of these confounders on predicted interactions, thereby achieving a more comprehensive understanding of user-item interactions in both domains. Finally, we employ cross-entropy loss to fine-tune the user preference disentanglement backbone $g(\cdot)$ and the interaction prediction network $f(\cdot)$. To be specific, the final objective function in domain A can be defined as follows:

$$g^*, f^* = \arg\min_{g,f} \sum_{y \in \mathcal{Y}^{A+} \cup \mathcal{Y}^{A-}} \ell(\hat{y}, y), \tag{4.16}$$

where $\hat{y}$ and y are the predicted interaction and corresponding observed interaction, respectively. $\ell(\hat{y}, y)$ denotes the cross-entropy loss function. $\mathcal{Y}^{A+}$ denotes the observed interaction set, and $\mathcal{Y}^{A-}$ corresponds to a specific quantity of negative samples, which are randomly selected from unseen user-item interaction set in domain A to mitigate the over-fitting. During the fine-tuning process, $g(\cdot)$ serves as the backbone, with the original prediction module being replaced by the interaction prediction network $f(\cdot)$. Likewise, we can obtain the objective function and predicted user-item interaction $\hat{y}_{ik}^{B}$ in domain B.

4.2.5 Time Complexity Analysis

Our CD2CDR mainly focuses on four modules: (1) Graph Propagation, (2) User Preference Disentanglement, (3) Confounder Disentanglement, and (4) Causal Deconfounding and Cross-domain Recommendation. While the first two modules are part of backbone model [188], the latter two constitute our novel framework. For simplicity and consistency, we assume that all embedding dimensions are d and the number of layers in each network structure within each module is L [176]. The time complexity for each module can be analyzed as follows:

(1) Graph Propagation: Assuming the graph has $(m+n)$ nodes, where m and n are the number of users and items, respectively. In addition, assuming the average number of neighboring nodes for each node is $\bar{N}$, the time complexity for the graph propagation process per node is $O(\bar{N}d)$. The total time complexity for graph propagation, using a graph convolutional network (GCN) with L layers, is $O(L(m+n)\bar{N}d)$. Given that $\bar{N} \ll (m+n)$, this simplifies to $O(L(m+n)d)$.

(2) User Preference Disentanglement: Next, we conduct user preference disentanglement using an architecture similar to the VAE encoder. Given that this module is implemented with an MLP consisting of L layers, the time complexity of user preference disentanglement is approximately $O(Lmd^2)$. The time complexity of domain classifier can be ignored as it is relatively simple compared to disentanglement module.

(3) Confounder Disentanglement: Then, we perform the confounder disentanglement module, which involves SDC and CDC disentanglement. For SDC disentanglement, we implement a structure similar to CycleGAN, using an MLP with L layers to decouple candidate SDCs. The time complexity of SDC disentanglement can be roughly estimated as $O(Ld^2m)$. For CDC disentanglement, ridge regression is used to calculate a transformation matrix $\mathbf{W}^{A\to B}$ for obtaining candidate CDCs. The estimated time complexity is $O(md^2+d^3)$. Considering $d \ll m$, it simplifies to $O(md^2)$. To identify representative SDCs and CDCs and eliminate redundancy, we perform K-means clustering on the candidate SDCs and CDCs. Given that the number of clus-

ter centroids is J, the time complexity for the K-means clustering is estimated to be $O(mdJ)$. Thus, the overall time complexity for the confounder disentanglement module is $O(Ld^2m + md^2 + mdJ)$, which simplifies to $O(md(Ld + J))$.

(4) Causal Deconfounding and Cross-domain Recommendation: Finally, we utilize the confounder selection function to effectively control the decoupled observed confounders, achieving more accurate deconfounding with a time complexity of approximately $O(mnJd)$. Subsequently, we concatenate the user embeddings, item embeddings, and selected confounders, feeding them into the MLP for prediction. Given that the prediction module consists of L MLP layers, the time complexity can be estimated as $O(Lmnd^2)$. Thus, the overall time complexity for the causal deconfounding and cross-domain recommendation module is $O(mnJd + Lmnd^2)$, which simplifies to $O(mnd(J + Ld))$.

Overall, the time complexity of our CD2CDR can be approximated as $O(mnd^2(J+L))$, where J is the number of cluster centroids, and m and n are the number of users and items, respectively. This approximation is based on combining the time complexities of all four modules and simplifying by focusing on the dominant terms. The overall time complexity exhibits a non-linear relationship with the number of users, items, observed confounders, and embedding dimensions.

4.3 Experiments on CD2CDR

Extensive experiments are conducted on seven real-world datasets to answer the following four research questions:

- **RQ1.** How does our model perform in comparison with state-of-the-art baseline models (see Section 4.3.2.1)?
- **RQ2.** How do different components, namely, confounder disentanglement, causal deconfounding and cycle consistency loss, influence the recommendation accuracy of our model (see Section 4.3.2.2)?

- **RQ3.** How do different backbone models impact the recommendation accuracy of our model (see Section 4.3.2.3)?
- **RQ4.** How do different hyper-parameter settings affect the recommendation accuracy of our model (see Section 4.3.2.4)?

4.3.1 Experimental Settings

4.3.1.1 Experimental Datasets

Semantic information, such as item titles containing details about free shipping, sales promotion, category, and brand, helps to disentangle user preferences and observed confounders. In e-commerce scenarios, semantic information is easily accessible and crucial for gaining a more comprehensive understanding of user-item interactions. To comprehensively evaluate our CD2CDR model, we conduct experiments in two distinct recommendation scenarios: (1) CDR with fully overlapping user sets and (2) cross-system recommendation (CSR) with only overlapping items and completely non-overlapping users.

For the CDR scenario, we select two real-world e-commerce datasets that provide rich semantic information, ratings, reviews and item metadata, namely, Rec-Tmall[1] dataset [39] and Amazon dataset [8]. For Amazon dataset, we choose two relevant domains, namely, Amazon-Electronics and Amazon-Cloth. Similarly, for Rec-Tmall dataset, we select three relevant behaviors as business domains, namely, Add to Favorite, Purchase, and Add to Cart[2]. In the Tmall-Favorite domain, most users engage in exploration-oriented behaviors, adding items they find appealing to their favorites without an immediate intent to purchase. In contrast, the Tmall-Purchase and Tmall-Cart domains reflect purchase-oriented behaviors, where users are more likely to select items that match their true preferences and may result in actual purchases. By defining

[1] https://tianchi.aliyun.com/dataset/140281

[2] For brevity, we refer to these subsets as Tmall-Favorite, Tmall-Purchase, Tmall-Cart, Amazon-Elec, and Amazon-Cloth, respectively in subsequent discussions.

Table 4.1: The statistics for three dual-target CDR tasks and a dual-target CSR task.

Tasks	Datasets	#Users	#Items	#Interactions	Density
Task #1	Tmall-Favorite	25,434	99,237	500,876	0.020%
	Tmall-Purchase	25,434	28,817	55,057	0.008%
Task #2	Tmall-Favorite	39,657	104,496	807,493	0.019%
	Tmall-Cart	39,657	44,172	354,499	0.020%
Task #3	Amazon-Elec	15,761	51,447	224,689	0.027%
	Amazon-Cloth	15,761	48,781	133,609	0.017%
Task #4	MovieLens	10,000	9,395†	1,462,905	1.56%
	Douban-Movie	2,712	34,893†	1,278,401	1.35%

† There are 4,115 common items between MovieLens and Douban-Movie datasets.

these distinct behaviors as domains, we broaden the concept of domains to encompass varying user intents, thereby enhancing the flexibility of our CDR framework for broader application scenarios [63].

For the CSR scenario, we utilize two widely-used movie recommendation datasets: MovieLens 20M [35] and Douban-Movie [185]. These datasets contain ratings and side information on common movies from different user communities, creating a scenario where user sets are completely non-overlapping while item sets are partially overlapping. This cross-system setting broadens our experimental scope beyond CDR to a more challenging CSR scenario that better testifies the effectiveness of CD2CDR.

4.3.1.2 Experimental Tasks

We construct four experimental tasks: three dual-target CDR tasks using e-commerce datasets and one CSR task using movie datasets. All tasks involve transforming explicit ratings into implicit feedback. For the CDR scenario, we design three tasks with fully overlapping user sets: (1) Tmall-Favorite and Tmall-Purchase, (2) Tmall-Favorite and Tmall-Cart, and (3) Amazon-Elec and Amazon-Cloth. These tasks are chosen to test the model's ability to handle diverse recommendations across different user interactions and preferences in e-commerce settings.

For Task #1, users and items with fewer than 20 interactions are removed from Tmall-Favorite, and those with fewer than 5 interactions are filtered out from Tmall-

Purchase. For the Task #2 and Task #3, users and items with fewer than 20 interactions in Task #2 and those with fewer than 5 interactions in Task #3 are filtered out. In line with the preprocessing operation taken for the two Amazon subsets in DisenCDR [8], we also conduct the same operation on three Rec-Tmall subsets to remove the cold-start item entry for testing.

For the CSR scenario (Task #4: MovieLens and Douban-Movie), we follow the filtering approach in [150, 183], retaining users and items with at least 5 interactions in Douban-Movie and extracting a subset of 10,000 users with at least 5 interactions from MovieLens 20M. We then identify common items across the two datasets, enabling knowledge transfer through overlapping items despite having completely non-overlapping users. The statistics are shown in Table 4.1.

4.3.1.3 Parameter Settings

The settings of our backbone DIDA-CDR are consistent with those listed in its original paper [188], including the number of GCN layers, embedding dimension and information fusion approach, etc. In the interaction prediction network, the structure is $e \rightarrow 32 \rightarrow 16 \rightarrow q$, where e is the combined size after the mapping of FC layer in Eq. (4.14), and q is the output size, i.e., the dimension of latent factors. We vary e in the range of $\{64, 128\}$ and q in the range of $\{8, 16\}$, and finally set $e = 128$ and $q = 8$. The initial parameters for all the above layers are set following a Gaussian distribution $X \sim \mathcal{N}(0, 0.01)$. In line with the approach used in GA-DTCDR [185], for each observed interaction, we randomly select 7 non-interacted items to serve as negative examples. For a fair comparison, we employ grid search to fine-tune the parameters of all models. Specifically, we select the learning rate in $\{0.01, 0.005, 0.001, 0.0005, 0.0001\}$. Moreover, we adopt the Adam optimizer [52] for all models with a batch size of 1024. In addition, we keep the number of cluster centroids $J^A_{sd} = J^B_{sd} = J_{cd}$ and vary them in $\{2, 5, 10, 20, 50\}$. Furthermore, we investigate the weight of cycle consistency loss λ in $\{0.1, 1, 2, 5, 10\}$, and the regularization parameter α in $\{0.1, 1, 10, 20, 50\}$. The influence of the above parameters on

our CD2CDR is particularly discussed in Section 4.3.2.4.

4.3.1.4 Model Training

Since our CD2CDR can be divided into three phases, we first pre-train the backbone of our model with 50 epochs[3] to obtain disentangled user preferences and comprehensive user preferences. Next, we train the generators and discriminators in the dual adversarial structure with 30 epochs to decouple SDCs, apply half-sibling regression, a computational method inherently without a training process [101], to decouple CDCs, and then save cluster centroids of both SDCs and CDCs. Finally, we replace the prediction module in the backbone with the interaction prediction network to fine-tune the overall CD2CDR with 20 epochs.

To enhance the stability of our model training in the dual adversarial structure, inspired by [190], we replace the negative log-likelihood objective with a least-squares loss for the adversarial losses $\mathcal{L}_{GAN_1}(S, H^B, D^A, D^B)$ and $\mathcal{L}_{GAN_2}(T, H^A, D^B, D^A)$. This replacement enables the generator to produce higher-quality outputs and improves training stability. The reasons are as follows. Firstly, the least-squares loss penalizes generated samples far from the decision boundary, guiding the generator to adjust these samples closer to the boundary. This process reduces the discrepancy between generated and real data distributions, improving the quality of generated samples. Secondly, the distance-based penalization produces more gradients to guide the generator's updates, mitigating the gradient vanishing issue and thereby stabilizing the generator's learning process. For further details, please refer to [89]. Taking $\mathcal{L}_{GAN_1}(S, H^B, D^A, D^B)$ as an example, the generator $S(\cdot)$ is trained to minimize $\mathbb{E}_{\mathbf{Z}^A_{spe}\sim\mathbb{P}^A}[(H^B(S(\mathbf{Z}^A_{spe}))-1)^2]$, while the corresponding discriminator $H^B(\cdot)$ is trained to minimize $\mathbb{E}_{\mathbf{Z}^B_{spe}\sim\mathbb{P}^B}[(H^B(\mathbf{Z}^B_{spe})-1)^2]+\mathbb{E}_{\mathbf{Z}^A_{spe}\sim\mathbb{P}^A}[H^B(S(\mathbf{Z}^A_{spe}))^2]$. Likewise, $\mathcal{L}_{GAN_2}(T, H^A, D^B, D^A)$ is optimized in a similar manner. In addition, we adjust the weight of the cycle consistency loss to balance the adversarial process and the cycle consistency constraint, ensuring stable convergence of the dual adversarial training.

[3]The number of training epochs for each phase is chosen in $\{10, 20, 30, 40, 50\}$.

Detailed results of weight adjustment can be found in Sections 4.3.2.2 and 4.3.2.4.

During each epoch, we shuffle and split the training data for both domains into batches. We then iterate through batches, training on domain A and domain B in parallel. This approach allows the model to learn from both domains within the same epoch, ensuring that the model parameters are updated based on information from both domains. This form of joint learning helps improve the generalization performance across domains. Note that Eq. (4.5) and Eq. (4.16) are not optimized jointly. Since observed confounders are no longer entangled with debiased user preferences after deconfounding, the joint optimization of Eq. (4.5) and Eq. (4.16) for decoupling these confounders from such preferences becomes redundant. For a fair comparison, other baselines are trained for 100 epochs to confirm their convergence.

4.3.1.5 Evaluation Metrics

Given the widespread use of leave-one-out approach in baselines, e.g., GA-DTCDR [185] and DisenCDR [8], we adopt it as well to validate the recommendation accuracy of our CD2CDR and baselines. Moreover, the test set is created by the final interaction of each user, while the training set is formed by the remaining interaction records of each user. In line with the methods introduced in [8, 9], for every interaction in the test set, we randomly select 999 non-interacted items as negative samples for the test user, and then predict scores for a total of 1000 items to perform ranking. The leave-one-out approach mainly uses Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG), which are commonly adopted in ranking evaluations [187]. In our experiments, these metrics are applied to validate the recommendation accuracy within top-10 rankings, and all experiments are conducted five times with average results reported in this chapter.

4.3.1.6 Comparison Methods

We choose seventeen state-of-the-art baseline models to conduct a comparison against the proposed CD2CDR. We then categorize the seventeen baseline models into four groups: (I) Single-Domain Recommendation (SDR), (II) Single-Target CDR, (III) Disentanglement-Based Dual-Target CDR, and (IV) Debiasing Dual-Target CDR. To the best of our knowledge, our CD2CDR is the first Deconfounding Dual-Target CDR model in the literature. Thus, we select three representative Debiasing Dual-Target CDR approaches as alternatives for Deconfounding Dual-Target CDR baselines. Moreover, although there are some methods that identify unobserved domain-specific confounders and even unobserved general confounders, or utilize backdoor adjustment in the single-domain manner, they are not selected as baseline models. This is because they focus on different settings, i.e., domain generalization [167, 67], cross-domain sequential recommendation (CDSR) [173, 147, 162], click-through rate (CTR) prediction [131, 91] and different item groups [166, 122], respectively, from our model. Furthermore, in Table 4.2, we present an in-depth analysis of embedding strategies and main ideas of seventeen baselines and our CD2CDR.

Overall, our baselines cover both single-domain and cross-domain recommendation models. In the experiments, we use our CD2CDR framework to extend all the above Disentanglement-Based Dual-Target CDR baselines. Experimental results (see Section 4.3.2.3) demonstrate that CD2CDR is highly extendable and compatible with most off-the-shelf disentanglement-based dual-target CDR backbones, making it suitable for a wide range of recommendation scenarios.

Table 4.2: The comparison of the baselines and our proposed model.

Model			Embedding Strategy	Main Idea
Baselines	Single-Domain Recommendation (SDR)	**NGCF** [128]	Graph Embedding	Devising an embedding propagation layer to encode the collaborative signal
		LightGCN [40]	Graph Embedding	Designing a simplified GCN using linear message propagation for RSs
		DCCF [97]	Graph Contrastive Learning & Disentanglement	Decoupling user intent by adaptively integrating self-supervised augmentation
	Single-Target Cross-Domain Recommendation (CDR)	**BPR_EMCDR** [87]	Linear Matrix Factorization (MF)	Maping the latent factors of common users/items for knowledge transfer
		BPR_DCDCSR [184]	Linear MF	Considering the rating sparsity degrees of individual users and items
		CUT [58]	Contrastive Learning	Filtering users' collaborative information with user similarity constraints
	Disentanglement-Based Dual-Target CDR	**BiTGCF** [74]	Graph Embedding	Integrating in-domain feature propagation and inter-domain feature transfer
		GA-DTCDR [185]	Graph Embedding	Generating and effectively combining more representative embeddings
		DisenCDR [8]	Graph Embedding & VAE	Developing two mutual-information-based disentanglement regularizers
		CausalCDR [57]	Causal Embedding & VAE	Using causal embeddings to model the joint distribution of user interactions
		GDCCDR [72]	Graph Embedding & Contrastive Learning	Using decoulpled graph for feature disentanglement and personalized transfer
		CrossAug [88]	Graph Embedding & Augmentation	Utilizing cross-domain interactions via intra- and inter-domain augmentation
		HJID [19]	Graph Embedding & Disentanglement	Devising a hierarchical subspace disentanglement method for latent factors
		DIDA-CDR [188]	Graph Embedding & Disentanglement	Decoupling three components to capture comprehensive user preferences
	Debiasing Dual-Target CDR	**SCDGN** [64]	Graph Embedding & Semantic Clustering	Distilling unbiased graph knowledge and learning debiasing vectors
		CDRIB [9]	Graph Embedding & Information Bottleneck	Designing two information bottleneck regularizers for debiasing
		IPSCDR [63]	In accordance with the backbone (DIDA-CDR)	Eliminating selection bias and transferring debiased user preferences
Our model	Deconfounding Dual-Target CDR	**CD2CDR**	In accordance with the backbone (DIDA-CDR)	Disentangling fine-grained observed confounders and deconfounding

4.3.2 Experimental Results and Analysis

4.3.2.1 Performance Comparison (for RQ1)

Table 4.3 displays a comparative analysis of the performance[4] of different methods across all four tasks using HR@10 and NDCG@10 as evaluation metrics. It is worth mentioning that Single-Target CDR baseline models are trained in both domains, but only their results in data-sparser domain are reported, because they are designed to enhance recommendation accuracy in data-sparser domain. We find from Table 4.3:

(1) Our CD2CDR improves Disentanglement-Based Dual-Target CDR baselines by an average of 16.73% and 18.40% w.r.t. HR@10 and NDCG@10, respectively. Among this type of baselines, BiTGCF [74] performs well on the Rec-Tmall dataset, outperforming GA-DTCDR [185] and DisenCDR [8], but still falls short compared to CrossAug [88]. CrossAug, which utilizes cross-domain data augmentation and domain-shared center alignment, achieves competitive performance comparable to Debiasing Dual-Target CDR baselines. However, our CD2CDR still outperforms CrossAug by 11.95% and 13.17% w.r.t. HR@10 and NDCG@10, respectively. This is because, in addition to the user preference disentanglement, we adopt the confounder disentanglement, which effectively decouples observed SDCs and CDCs. By decoupling these confounders, we account for the fact that user interactions are not solely driven by their true preferences but also by observed confounders. Such confounders' positive influences can be secondary causes for user-item interactions, while their negative influences will result in capturing biased comprehensive user preferences. Effectively decoupling such observed confounders allows us to consider a more comprehensive range of factors affecting user-item interactions, thereby achieving better recommendation performance in both domains;

(2) Our CD2CDR improves Debiasing Dual-Target CDR baselines by an average of

[4]We only display experimental results when the embedding dimension $d = 64$ in Table 4.3 due to space limitation. For other values of d that are not shown, similarly, our CD2CDR also significantly outperforms other baselines.

13.03% and 14.85% w.r.t. HR@10 and NDCG@10, respectively. This demonstrates that deconfounding the observed confounders in each of both domains effectively benefits the prediction of user-item interactions in dual-target CDR;

(3) Our CD2CDR improves the best-performing baseline, i.e., IPSCDR [63], which is implemented with the same backbone as our model. Specifically, our CD2CDR outperforms IPSCDR with an average increase of 6.17% and 8.23% w.r.t. HR@10 and NDCG@10, respectively. This is because our CD2CDR particularly takes observed CDCs into consideration and our causal deconfounding module can not only eliminate observed confounders' negative effects on user preferences, but also preserve their positive effects on predicted interactions, thus gaining a more comprehensive understanding of user-item interactions;

(4) In the challenging CSR scenario (Task #4) where user sets are completely non-overlapping, our CD2CDR still shows strong performance, outperforming the best-performing baseline model by an average of 4.75% and 6.33% w.r.t. HR@10 and NDCG@10, respectively. This demonstrates that our model adapts effectively to item-wise knowledge transfer through common items across different systems, and can extract item embeddings that are not entangled with observed confounders, enabling more accurate matching with comprehensive user preferences despite the absence of user overlap. This evaluation in the cross-system context extends our experimental scope beyond CDR scenarios, further validating the robustness and effectiveness of our model in more challenging CSR settings.

4.3.2.2 Ablation Study (for RQ2)

To highlight the significance of each component in enhancing the recommendation accuracy of our model, we reconstruct our CD2CDR into four variants and perform an ablation study for all four tasks.

Impact of Confounder Disentanglement. We modify our proposed CD2CDR to form two variants, namely **CD2CDR_Cross** and **CD2CDR_Single**, by removing the

Table 4.3: Comparative performance analysis (%) of different methods in all four tasks using HR@10 and NDCG@10 as evaluation metrics [185]. For experimental results, best results are highlighted in bold and the results of best-performing baseline model are underlined (* denotes $p < 0.05$ in the paired t-test between the best-performing baseline model and CD2CDR) [188].

Datasets	SDR Baselines						Single-Target CDR Baselines					
	NGCF		LightGCN		DCCF		BPR_EMCDR		BPR_DCDCSR		CUT	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Tmall-Favorite	12.39	6.27	12.62	6.35	12.54	6.32	-	-	-	-	-	-
Tmall-Purchase	6.46	4.01	6.74	4.06	6.88	4.21	5.31	3.56	5.88	3.73	7.25	4.34
Tmall-Favorite	12.48	6.29	12.81	6.43	12.65	6.37	-	-	-	-	-	-
Tmall-Cart	10.04	5.23	10.85	5.58	10.98	5.8	8.79	4.87	9.46	5.11	11.48	6.16
Amazon-Elec	21.85	12.36	21.73	11.61	21.57	11.14	-	-	-	-	-	-
Amazon-Cloth	11.62	6.18	12.04	6.22	12.23	6.25	10.69	5.47	11.44	6.15	13.52	7.1
MovieLens	13.17	6.85	13.46	7.09	13.35	7.04	-	-	-	-	-	-
Douban-Movie	10.48	5.38	10.71	5.48	10.83	5.56	9.72	5.19	10.36	5.28	11.25	6.07

Datasets	Disentanglement-Based Dual-Target CDR Baselines							
	BiTGCF		GA-DTCDR		DisenCDR		CausalCDR	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Tmall-Favorite	15.68	8.63	14.87	7.82	15.34	8.25	15.96	8.69
Tmall-Purchase	9.24	5.06	8.44	4.53	9.01	4.94	9.17	5.03
Tmall-Favorite	15.80	8.65	14.91	8.05	15.39	8.53	16.25	9.10
Tmall-Cart	13.45	7.13	12.66	6.38	13.13	6.81	13.56	7.18
Amazon-Elec	23.42	13.65	24.79	13.87	24.53	14.02	25.14	14.47
Amazon-Cloth	14.31	7.59	14.58	7.64	15.81	8.56	15.93	8.65
MovieLens	15.28	8.25	15.53	8.54	19.28	10.53	19.87	10.89
Douban-Movie	11.79	6.19	11.96	6.21	14.72	7.66	15.29	8.27

Datasets	Disentanglement-Based Dual-Target CDR Baselines							
	GDCCDR		CrossAug		HJID		DIDA-CDR	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Tmall-Favorite	16.24	9.05	16.31	9.07	16.45	9.12	16.60	9.14
Tmall-Purchase	9.65	5.19	9.74	5.20	9.88	5.22	10.03	5.28
Tmall-Favorite	16.53	9.13	16.66	9.15	16.79	9.19	17.02	9.26
Tmall-Cart	14.21	7.56	14.38	7.59	14.43	7.60	14.56	7.71
Amazon-Elec	25.69	14.58	25.73	14.62	25.94	14.71	26.12	14.83
Amazon-Cloth	16.72	9.16	16.85	9.24	17.27	9.42	17.75	9.70
MovieLens	20.14	10.92	20.43	10.96	20.64	11.01	21.08	11.09
Douban-Movie	15.58	8.63	15.82	8.67	16.11	8.85	16.34	9.06

Datasets	Debiasing Dual-Target CDR Baselines						Our Model		Improvement (CD2CDR vs. best baselines)	
	SCDGN		CDRIB		IPSCDR		CD2CDR			
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Tmall-Favorite	15.07	8.16	15.9	8.62	17.14	9.38	**18.01***	**9.87***	5.08%	5.22%
Tmall-Purchase	8.73	4.75	9.56	5.17	10.51	5.39	**11.38***	**6.12***	8.28%	13.54%
Tmall-Favorite	15.14	8.21	16.22	9.04	17.43	9.47	**18.35***	**9.98***	5.28%	5.39%
Tmall-Cart	12.82	6.48	13.98	7.51	15.05	8.02	**16.37***	**9.06***	8.77%	12.97%
Amazon-Elec	24.69	13.85	25.06	14.34	26.78	15.19	**28.11***	**16.24***	4.97%	6.91%
Amazon-Cloth	15.19	8.24	16.67	9.18	18.26	9.93	**19.62***	**10.87***	7.45%	9.47%
MovieLens	20.81	11.05	19.36	10.55	21.74	11.67	**22.73***	**12.56***	4.55%	7.63%
Douban-Movie	15.88	8.67	15.13	8.17	16.75	9.18	**17.59***	**9.61***	5.01%	4.68%

SDC disentanglement and CDC disentanglement, respectively. From Table 4.4, we can find that with SDC disentanglement module, CD2CDR outperforms **CD2CDR_Cross** with an average improvement of 5.49 %. This shows that the dual adversarial struc-

Table 4.4: Ablation study of different components in our CD2CDR across three dual-target CDR tasks and a dual-target CSR task. The best results are highlighted in bold.

Datasets	CD2CDR_Cross		CD2CDR_Single		CD2CDR_Coarse		CD2CDR_Cycle		CD2CDR	
	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Tmall-Favorite	17.25	9.41	17.48	9.56	15.03	8.14	17.86	9.70	**18.01**	**9.87**
Tmall-Purchase	10.67	5.40	10.82	5.53	8.68	4.69	11.22	6.04	**11.38**	**6.12**
Tmall-Favorite	17.64	9.59	17.85	9.69	15.11	8.16	18.23	9.89	**18.35**	**9.98**
Tmall-Cart	15.21	8.25	15.56	8.61	12.75	6.43	16.18	9.01	**16.37**	**9.06**
Amazon-Elec	26.83	15.23	26.97	15.32	24.48	13.79	27.94	16.15	**28.11**	**16.24**
Amazon-Cloth	18.32	9.96	18.44	10.01	14.76	7.68	19.37	10.58	**19.62**	**10.87**
MovieLens	21.89	12.38	22.03	12.45	20.23	10.93	22.55	12.49	**22.73**	**12.56**
Douban-Movie	16.91	9.22	17.16	9.40	15.46	8.54	17.38	9.51	**17.59**	**9.61**

ture can effectively disentangle observed SDCs, and SDCs play an important role in predicting user-item interactions in each domain. In addition, our proposed CD2CDR improves **CD2CDR_Single** by an average increase of 4.23%. This indicates that half-sibling regression is well suited for decoupling observed CDCs, which are essential factors for achieving a comprehensive understanding of user-item interactions in both domains. Overall, our confounder disentanglement module can explicitly decouple more accurate observed confounders, especially the CDCs, thus enable CD2CDR to obtain better recommendation performance via accurate causal deconfounding.

Impact of Causal Deconfounding. Moreover, another variant, i.e., **CD2CDR_Coarse**, directly incorporates decoupled observed confounders with biased comprehensive user preferences in each domain and does not include the causal deconfounding module. From Table 4.4, we can observe that without the causal deconfounding module, the recommendation accuracy of **CD2CDR_Coarse** drops by 21.35% on average, making it less effective compared to the Debiasing Dual-Target CDR baselines. This shows that the causal deconfounding module indeed helps the model control the negative effects of SDCs and CDCs on user preferences. By recovering debiased comprehensive user preferences and then incorporating the positive effects of SDCs and CDCs into such preferences, the module enables the model to obtain the better recommendation accuracy in both domains.

Impact of Cycle Consistency Loss. In addition, we construct another variant, namely **CD2CDR_Cycle**, by removing the cycle consistency loss in the SDC disentanglement

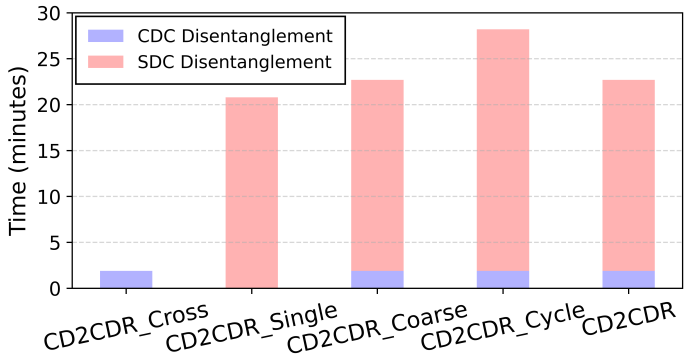

Figure 4.3: Average time comparison of confounder disentanglement phase across four tasks for CD2CDR and its variants.

module. From Table 4.4, we can observe that our CD2CDR improves **CD2CDR_Cycle** by an average of 1.01%. This demonstrates that the cycle consistency loss effectively preserves users' domain-specific preferences during the transformation process, ensuring the transformed preferences accurately reflect the original user preferences rather than merely confusing the discriminator. By incorporating the cycle consistency loss to stabilize the adversarial loss, our model can more accurately disentangle SDCs, providing strong support for explicitly considering the impact of observed confounders on user preferences and user-item interactions.

Overall, our ablation study demonstrates the importance of each component in our CD2CDR model. Similar trends are observed in the CSR scenario (Task #4) where knowledge transfer relies on common items, further confirming the effectiveness of these components across different systems.

Empirical Analysis of Time Complexity. To comprehensively evaluate the trade-off between effectiveness and efficiency of our model and its variants, we further conduct an empirical analysis of their time complexity. Figure 4.3 illustrates the average time consumption of the confounder disentanglement phase across all four tasks. In Figure 4.3, the time of confounder disentanglement phrase is divided into CDC disentanglement time (implemented via half-sibling regression) and SDC disentanglement time (implemented through dual adversarial training).

As shown in Figure 4.3, **CD2CDR_Cross** consumes significantly less time com-

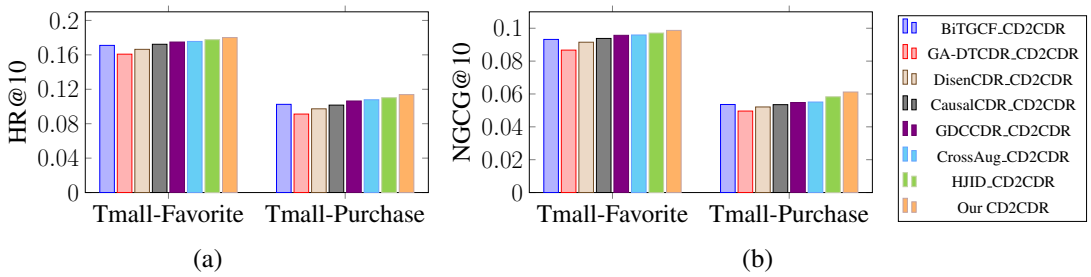

Figure 4.4: (a)-(b): Comparative performance analysis between CD2CDR and its seven variants with different backbones.

pared to other variants since it only employs half-sibling regression, which inherently requires no iterative training as mentioned in Section 4.3.1.4. However, this computational efficiency comes with a performance degradation of 5.49% in recommendation metrics. In contrast, **CD2CDR_Single** requires approximately 10 times more computational time since it relies solely on the dual adversarial structure, which demands multiple training epochs to converge, yet still underperforms our CD2CDR by 4.23% in recommendation metrics. In addition, **CD2CDR_Coarse** shows comparable time consumption to CD2CDR as it uses identical confounder disentanglement processes, despite suffering a substantial 21.35% performance drop. Notably, **CD2CDR_Cycle** consumes more time than our CD2CDR despite removing the cycle consistency loss, while also showing 1.01% lower performance. The increased time consumption occurs because without the stabilizing effect of cycle consistency loss, the adversarial training requires more iterations to reach convergence. Moreover, the absence of cycle consistency loss leads to less accurate SDC disentanglement, which explains the observed performance degradation.

It is worth noting that the confounder disentanglement phase represents only a fraction of the overall computational cost in the entire training pipeline (as discussed in Section 4.2.5). Despite the differences in how CD2CDR and its variants implement confounder disentanglement and causal deconfounding, the major computational cost for both CD2CDR and its variants typically comes from the shared pretraining phase and final recommendation phase. As a result, the time differences observed in confounder disentanglement have a relatively limited impact on total training time. Based

on the above analysis, CD2CDR achieves a good balance between effectiveness and efficiency, providing superior recommendation performance with reasonable computational requirements.

4.3.2.3 Impact of Different Backbones (for RQ3)

Since our CD2CDR can be easily combined with disentanglement-based dual-target CDR backbone models, in addition to DIDA-CDR [188], we select all the other representative and state-of-the-art models from this group as backbones to form seven variants, namely, **BiTGCF_CD2CDR**, **GA-DTCDR_CD2CDR**, **DisenCDR_CD2CDR**, **CausalCDR_CD2CDR**, **GDCCDR_CD2CDR**, and **CrossAug_CD2CDR** as well as **HJID_CD2CDR**. Our aim is to demonstrate the flexibility and effectiveness of CD2CDR by integrating it with various representative and state-of-the-art disentanglement-based dual-target CDR backbone models, thereby highlighting its generalizability and extendability across diverse CDR scenarios. The performance comparison of CD2CDR and its seven variants[5] with corresponding backbones is shown in Figure 4.4. We find that when our model employs DIDA-CDR as the backbone, it improves the above variants, namely, **BiTGCF_CD2CDR**, **GA-DTCDR_CD2CDR**, **DisenCDR_CD2CDR**, **CausalCDR_CD2CDR**, **GDCCDR_CD2CDR**, and **CrossAug_CD2CDR** as well as **HJID_CD2CDR** by an average of 7.97%, 16.87%, 11.42%, 7.74%, 5.09%, 4.47% and 2.34%, respectively. This improvement can be attributed to the ability of DIDA-CDR to effectively decouple three components of user preferences for modeling more accurate comprehensive user preferences. Notably, the ability of DIDA-CDR aligns well with the requirements of our CD2CDR, which relies on this precise disentanglement to accurately decouple observed confounders. In addition, our model, when combined with various backbones, consistently outperforms these backbones in their original form with an average improvement of 9.05% and 8.04% w.r.t. HR@10 and NDCG@10, respectively. This not only shows the superior efficacy of CD2CDR in

[5]Owing to constraints in space, Figure 4.4 and Figure 4.5 solely present the experimental results for Task #1, with similar trends observed in other unshown tasks.

improving recommendation performance in both domains, but also shows its generalizability to various CDR models.

4.3.2.4 Parameter Sensitivity (for RQ4)

Impact of the number of cluster centroids J. To explore the impact of number of cluster centroids J on the efficacy of our proposed CD2CDR, we keep $J_{sd}^A = J_{sd}^B = J_{cd}$ and vary them in $\{2, 5, 10, 20, 50\}$. The corresponding experimental results are depicted in Figures 4.5(a)-(b). We can observe that as J increases, the recommendation performance initially improves but gradually plateaus beyond 10. This suggests that there is a threshold for J, which may vary in different datasets. Beyond this threshold, additional cluster centroids do not significantly improve the recommendation performance. In other words, once J reaches this threshold, the potential confounders represented by these cluster centroids are comprehensive enough for effective deconfounding. With the aim of achieving a balance between model complexity and recommendation accuracy, we finally set $J_{sd}^A = J_{sd}^B = J_{cd} = 10$ in the three dual-target CDR tasks. In particular, the comprehensive confounder disentanglement significantly contributes to more accurate estimation of Eq. (4.12). More importantly, the experimental results show that our confounder disentanglement module can form effective confounder spaces, where even basic clustering techniques can easily identify key confounders, thereby yielding promising deconfounding results.

Impact of the weight of cycle consistency loss λ. To examine the impact of the weight of cycle consistency loss λ on the recommendation performance of our model, we test λ with values from $\{0.1, 1, 2, 5, 10\}$. The experimental results are depicted in Figures 4.5(c)-(d). We can observe that smaller values of λ (e.g., 0.1 or 1) allow the adversarial losses to play a dominant role during training, driving the generator to better align the distributions of domain-specific user preferences across domains. Meanwhile, the cycle consistency loss still enforces moderate consistency, ensuring the transformation remains meaningful. In contrast, when λ is set to larger values (e.g., 5 or 10), the cycle consistency loss becomes overly influential. As a result, the generator focuses primar-

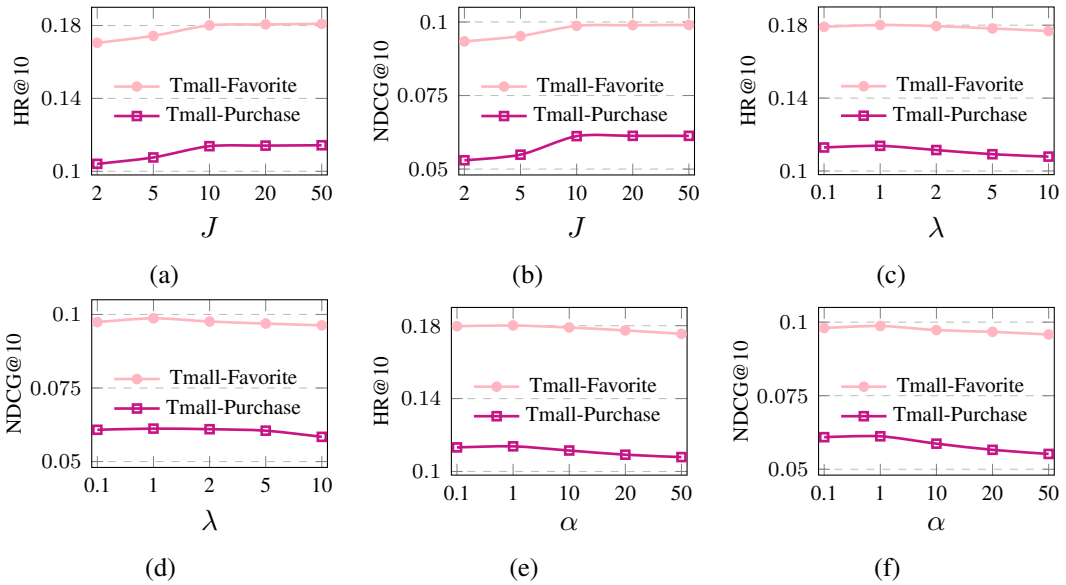

Figure 4.5: (a)-(b): Impact of the number of cluster centroids J. (c)-(d): Impact of the weight of cycle consistency loss λ. (e)-(f): Impact of the regularization parameter α.

ily on minimizing cycle consistency loss, prioritizing outputs that closely resemble the inputs, rather than using feedback from the discriminator to refine cross-domain transformations. This weakens the discriminator's ability to guide the generator towards producing domain-specific user preferences that are indistinguishable from those in the target domain. Therefore, we select $\lambda = 1$ to ensure a proper balance between the optimization of adversarial losses and cycle consistency loss.

Impact of the regularization parameter α. To analyze the impact of the regularization parameter α on decoupling candidate CDCs, we test α with values from $\{0.1, 1, 10, 20, 50\}$. The results are shown in Figures 4.5(e)-(f). We can observe that small values of α (e.g., 0.1) lead to insufficient regularization, making the regression overly sensitive to noise. By contrast, large α values (e.g., 20 or 50) place too much emphasis on the regularization term, leading the model to underfit important features related to the CDCs. This weakens the model's ability to effectively decouple candidate CDCs. As shown in Figures 4.5(e)-(f), when $\alpha = 1$, our CD2CDR achieves the best trade-off between stabilizing numerical computations and maintaining the accuracy of decoupling candidate CDCs, thereby enhancing the effectiveness of subsequent

deconfounding in cross-domain settings.

4.4 Summary

In this chapter, we have proposed a novel Causal Deconfounding framework via Confounder Disentanglement for dual-target CDR, called CD2CDR. CD2CDR not only effectively decouples observed single-domain and cross-domain confounders, but also preserves the positive direct effects of such observed confounders on predicted interactions and eliminates their negative effects on capturing comprehensive user preferences, thereby enhancing the recommendation accuracy in both domains simultaneously. Moreover, we have conducted extensive experiments on seven real-world datasets, which demonstrates that our CD2CDR significantly outperforms the state-of-the-art methods.

Chapter 5

Causal-Invariant Cross-Domain Out-of-Distribution Recommendation

While Chapter 3 proposes DIDA-CDR to disentangle essential components of user preferences, and Chapter 4 proposes CD2CDR to decouple and then deconfound observed confounders, both approaches are developed under the assumption that training and testing data within each domain are independently and identically distributed (IID). However, this IID assumption rarely holds in real-world scenarios. In CDR, the goal is to leverage knowledge from a relatively data-richer source domain to improve recommendation performance in a data-sparser target domain. While traditional CDR methods primarily focus on addressing distribution shifts between domains, i.e., cross-domain distribution shift (CDDS), they often overlook the distribution shifts that occur within a single domain, i.e., single-domain distribution shift (SDDS). The above two co-existing distribution shifts lead to out-of-distribution (OOD) environments that hinder effective knowledge transfer and generalization, ultimately degrading recommendation performance in CDR.

In this chapter, we propose a new setting of cross-domain OOD recommendation, which simultaneously addresses two types of distribution shifts through unified modeling. In addition to previously introduced challenges **CH2** and **CH3** (as discussed in Sections 1.2.2 and 1.2.3), this new setting introduces a critical challenge that has not been fully explored in the literature. As discussed in Section 1.2.4, this challenge **CH4** is formalized as: *How to simultaneously address cross-domain and single-domain dis-*

tribution shifts to achieve reliable recommendation under OOD environments in CDR?

This chapter first formalizes the new setting of cross-domain OOD recommendation. Then, a novel CICDOR framework that discovers invariant causal structures across distributions to tackle both CDDS and SDDS for reliable cross-domain OOD recommendation is proposed. Finally, the detailed framework and experiments are introduced in the following sections.

5.1 Problem Statement

5.1.1 Notations and Background

This chapter considers a source domain D^s with relatively richer data and a target domain D^t with sparser data. Both domains share a common user set $\mathcal{U}$ of size $m = |\mathcal{U}|$. We denote the user reviews in D^s and D^t as R^s and R^t, respectively. The item sets in the source and target domains are $\mathcal{V}^s$ (of size $n^s = |\mathcal{V}^s|$) and $\mathcal{V}^t$ (of size $n^t = |\mathcal{V}^t|$) respectively, with corresponding user-item interaction matrices $\mathbf{Y}^s$ and $\mathbf{Y}^t$. Here, each entry $y_{ij} \in \{0, 1\}$ represents whether user i has interacted with item j.

5.1.2 Problem Formulation

In this chapter, we propose a new setting of cross-domain OOD recommendation.

• **Input:** Data from source domain D^s and training data from target domain D^t, with data distributions $\mathbb{P}^s(\mathcal{U}, \mathcal{V}^s) \neq \mathbb{P}^t_{tr}(\mathcal{U}, \mathcal{V}^t)$.

• **Output:** A cross-domain recommender system that can accurately predict whether any user $i \in \mathcal{U}$ interacts with an item $j \in \mathcal{V}^t$ (i.e., $\hat{y}^t_{ij}$) in the target domain testing data D^t_{te}, where $\mathbb{P}^t_{tr}(\mathcal{U}, \mathcal{V}^t) \neq \mathbb{P}^t_{te}(\mathcal{U}, \mathcal{V}^t)$.

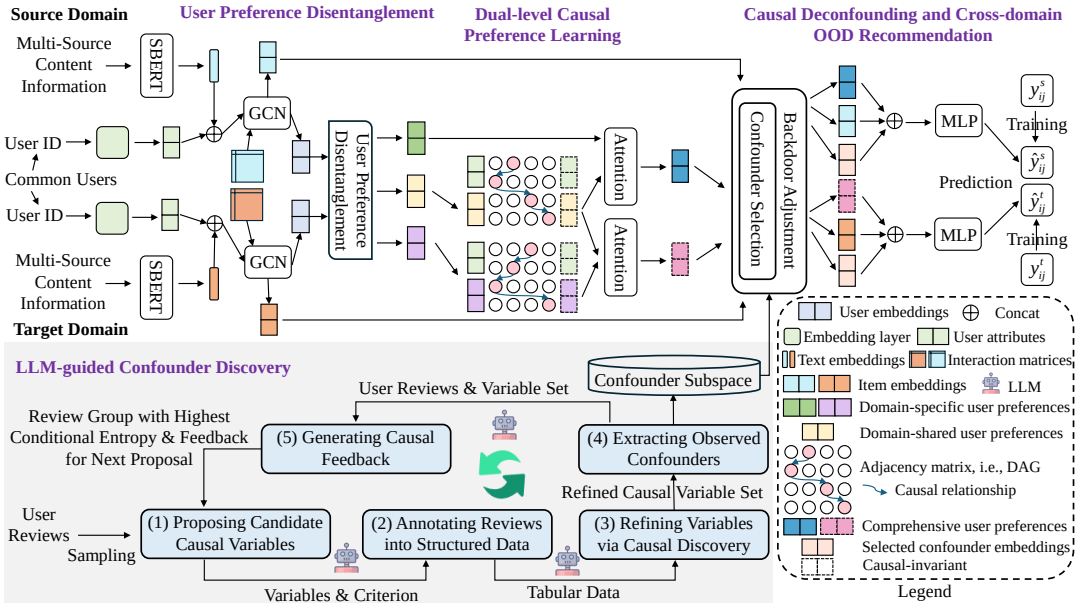

Figure 5.1: The flowchart of the proposed framework.

5.2 The Proposed CICDOR Framework

5.2.1 Framework

In this section, we propose a novel **C**ausal-**I**nvariant **C**ross-**D**omain **O**ut-of-distribution **R**ecommendation framework, called CICDOR. As illustrated in Figure 5.1, CICDOR is built upon dual-level causal structures and follows the principle of causal invariance to tackle the OOD problem in CDR. Specifically, (1) To simultaneously address CDDS and SDDS, CICDOR proposes a dual-level causal preference learning module that discovers causal structures and infers causal-invariant user preferences at both domain-specific and domain-shared levels. (2) To ensure accurate causal-invariant user preference inference, CICDOR proposes an LLM-guided confounder discovery module. This module identifies interaction-related causal variables from reviews and further extracts observed confounders for effective deconfounding, thus enabling the accurate inference of debiased comprehensive causal-invariant user preferences for reliable cross-domain OOD recommendation.

5.2.2 User Preference Disentanglement

5.2.2.1 Embedding Construction

Taking the target domain D^t as an example, let $\mathbf{u} \in \{0,1\}^m$ and $\mathbf{v}^t \in \{0,1\}^{n^t}$ denote the one-hot vectors of users and items, respectively [179]. We first transform user one-hot vectors into k-dimensional attribute embeddings[1] via $\mathbf{E}^t_{att} = \mathbf{W}^t_{att}\mathbf{u}$, where $\mathbf{W}^t_{att} \in \mathbb{R}^{k\times m}$ is a learnable embedding matrix. To enhance the extraction of disentangled user preferences, we leverage multi-source content information including user reviews and item details [185]. Specifically, we collect all reviews written by user u_i into a user document, while for item v_j, we collect both its details and associated reviews into an item document [189]. Next, we employ a sentence transformer model to encode these documents, converting them into dense representations. This process generates user text embeddings $\mathbf{E}^t_{ut}$ and item text embeddings $\mathbf{E}^t_{vt}$ for all users and items in the training set. We then concatenate $\mathbf{E}^t_{att}$ and $\mathbf{E}^t_{ut}$ to obtrain the combined user embeddings $\mathbf{E}^t_{uc}$. These text embeddings are then projected into fixed-dimensional spaces through multi-layer perceptron (MLP) layers to create initial user embeddings $\mathbf{E}^t_{ui}$ and initial item embedding $\mathbf{E}^t_{vi}$ in the target domain. Likewise, we can obtain $\mathbf{E}^s_{att}$, $\mathbf{E}^s_{ui}$ and $\mathbf{E}^s_{vi}$ in the source domain as well. Using these initial embeddings and the corresponding interaction matrices as inputs, we construct two heterogeneous graphs and then employ the Graph Convolutional Network (GCN) [53] to generate enriched user embeddings $\mathbf{E}_u$ and item embeddings $\mathbf{E}_v$ that capture both content semantics and collaborative patterns within each domain.

5.2.2.2 User Preference Disentanglement

In CDR, a crucial challenge is to determine what knowledge should be transferred across domains [186, 152]. To address this challenge, we disentangle the enriched user embeddings E_u to obtain transferable domain-shared user preferences while preserving domain-specific user preferences that capture unique user behaviors in each

[1]Unless otherwise specified, all embedding dimension mentioned hereafter is k.

domain. Inspired by [15], we implement a user preference disentanglement module with three encoders: one domain-shared encoders and two domain-specific encoders (all using the two-layer MLP). These encoders collaboratively disentangle user preferences into domain-shared and domain-specific components.

For effective disentanglement, we employ a domain discriminator implemented as a fully-connected neural network. This discriminator attempts to identify whether encoded user preferences originate from the source or target domain by minimizing a domain classification loss. Between the domain-shared encoder and discriminator, we utilize a gradient reversal layer (GRL) [25], which multiplies gradients by a negative constant during backpropagation. This causes the domain-shared encoder to maximize the domain classification loss, thereby learning to generate user preferences that confuse the discriminator. Through this domain adversarial training [26], the domain-shared encoder effectively extracts transferable common user preferences. To optimize this disentanglement, we define the domain classification losses using binary cross-entropy as follows:

$$\mathcal{L}^s_{\text{sha}} = -\frac{1}{N^s}\sum_{i=1}^{N^s}\log(1-\hat{d}^s_{\text{sha}}), \quad \mathcal{L}^t_{\text{sha}} = -\frac{1}{N^t}\sum_{i=1}^{N^t}\log(\hat{d}^t_{\text{sha}}), \tag{5.1}$$

where d denotes the binary domain label (0 for source domain and 1 for target domain) and $\hat{d}$ is the discriminator's predicted probability of the input user preferences belonging to the target domain. Similarly, the domain-specific encoders connect directly to the discriminator without GRL and are trained to preserve domain-specific features by minimizing the domain classification loss $\mathcal{L}^s_{\text{spe}}$ and $\mathcal{L}^t_{\text{spe}}$.

Then, the overall domain disentanglement objective can be expressed as:

$$\mathcal{L}^{\text{dom}} = \gamma(\mathcal{L}^s_{\text{sha}} + \mathcal{L}^s_{\text{spe}}) + (1-\gamma)(\mathcal{L}^t_{\text{sha}} + \mathcal{L}^t_{\text{spe}}), \tag{5.2}$$

where γ is a balancing parameter between the source and target domain losses during disentanglement. The details of user preference disentanglement process can be found

in [15]. Finally, we can obtain domain-shared user preferences $\mathbf{E}_{sha}$ and domain-specific user preferences $\mathbf{E}^s_{spe}$ and $\mathbf{E}^t_{spe}$.

5.2.3 Dual-Level Causal Preference Learning

After obtaining domain-specific and domain-shared user preferences, inspired by CDCOR [65], we model the causal relationships between user attributes and user preferences, aiming to identify causal structures that remain invariant across distribution shifts. While CDCOR primarily focuses on learning causal structures at the domain-shared level to assist the learning of causal relationships in the target domain, it overlooks a crucial aspect: the invariant causal structure within the target domain itself, which is fundamental for robust single-domain OOD recommendation. Learning the domain-specific causal structure effectively addresses SDDS by capturing invariant causal relationships within the target domain, while the domain-shared causal structure provides supplementary support for mitigating both SDDS and CDDS. Motivated by the complementary nature of these two causal structures, we propose a dual-level causal preference learning module that learns two Directed Acyclic Graphs (DAGs) to model causal relationships at each level, respectively.

As shown in Figure 5.1, we do not learn the source domain's domain-specific causal structure. This design decision is based on two key considerations: (1) our primary goal is to enhance OOD recommendation performance in target domain, and source domain is only used to provide auxiliary training signals without requiring OOD generalization; (2) source domain's domain-specific user preferences are only used within its own prediction branch and do not participate in target domain inference, making causal invariance unnecessary for these representations. This design reduces model complexity while maintaining essential causal structure learning components.

Taking the domain-specific level as an example, we represent its DAG as a weighted adjacency matrix $\mathbf{A}_{\text{spe}} \in \mathbb{R}^{2k \times 2k}$, where $A_{i,j}$ indicates the strength of the causal influ-

ence from node i to node j. Each node represents one dimension of either user attributes or domain-specific user preferences embedding. The structural causal model (SCM) can be formulated as follows:

$$\mathbf{B}_{\text{spe}} = \mathbf{A}_{\text{spe}}^{\top}\mathbf{B}_{\text{spe}} + \boldsymbol{\epsilon}, \tag{5.3}$$

where $\boldsymbol{\epsilon}$ represents the noise term, and $\mathbf{B}_{\text{spe}} = \mathbf{E}_{\text{att}}^{t} \| \mathbf{E}_{\text{spe}}^{t} \in \mathbb{R}^{2k}$, and $\|$ denotes the operation of concatenation. Eq.(5.3) defines how child nodes are determined by their parent nodes through $\mathbf{A}_{\text{spe}}$. When Eq.(5.3) holds, $\mathbf{A}_{\text{spe}}$ represents the causal structure at the domain-specific level, capturing the invariant relationships between user attributes and domain-specific user preferences. To learn $\mathbf{A}_{\text{spe}}$, we minimize the following reconstruction loss:

$$\mathcal{L}_{\text{spe}}^{\text{rec}} = \frac{1}{N}\sum_{i=1}^{N} \|\mathbf{B}_i - \mathbf{A}_{\text{spe}}^{\top}\mathbf{B}_i\|_2^2. \tag{5.4}$$

To ensure that $\mathbf{A}_{\text{spe}}$ remains acyclic, as required by DAG properties, we optimize the following loss as a constraint [180]:

$$\mathcal{L}_{\text{spe}}^{\text{dag}} = \text{Tr}(\mathbf{e}^{\mathbf{A}\circ\mathbf{A}}) - k, \tag{5.5}$$

where $\text{Tr}(\cdot)$ computes the trace, $\circ$ is the element-wise product, and $\mathbf{e}$ represents the Euler's number. Moreover, to ensure $\mathbf{A}_{\text{spe}}$ aligns with the nature of preference formation, we introduce two structural constraints: (1) causal influences should only flow from user attribute nodes to user preference nodes, and (2) user preference nodes cannot be root nodes [43]. These constraints align with both causal intuition and the characteristics of preference formation. The first constraint reflects that user preferences are shaped by their inherent attributes, while the second constraint ensures that preferences are always derived from underlying personal attributes rather than emerging

spontaneously. To enforce such constraints, we optimize following two losses:

$$\mathcal{L}_{\text{spe}}^{\text{path}} = \|\mathbf{A}_{[k+1:2k,1:k]}\|_1, \tag{5.6}$$

$$\mathcal{L}_{\text{spe}}^{\text{root}} = \sum_{i=k+1}^{2k} -\log \|\mathbf{A}_{[:,i]}\|_1. \tag{5.7}$$

Then, we can obtain the domain-specific causal loss, which is expressed as follows:

$$\mathcal{L}_{\text{spe}}^{\text{cau}} = \mathcal{L}_{\text{spe}}^{\text{rec}} + \alpha_1 \mathcal{L}_{\text{spe}}^{\text{dag}} + \alpha_2 \mathcal{L}_{\text{spe}}^{\text{path}} + \alpha_3 \mathcal{L}_{\text{spe}}^{\text{root}} + \alpha_4 \|\mathbf{A}_{\text{spe}}\|_1. \tag{5.8}$$

Similarly, we can obtain the domain-shared causal loss $\mathcal{L}_{\text{sha}}^{\text{cau}}$. Thus, we formulate the dual-level causal loss as follows:

$$\mathcal{L}^{\text{cau}} = \mathcal{L}_{\text{spe}}^{\text{cau}} + \mathcal{L}_{\text{sha}}^{\text{cau}}. \tag{5.9}$$

When training is completed, we can infer domain-specific causal-invariant user preferences using the learned causal structure $\mathbf{A}_{\text{spe}}$. During inference, we feed $\mathbf{B}_{\text{spe}} = \mathbf{E}_{att}^{t} \| \mathbf{0}$ into Eq.(5.3), where $\mathbf{0}$ is a zero vector replacing domain-specific user preferences. From the output $\hat{\mathbf{H}}_{\text{spe}}$, we extract its posterior k-dimensional vector as the domain-specific causal-invariant user preferences $\mathbf{E}_{spe}^{tinv}$ in target domain. Since $\mathbf{A}_{\text{spe}}$ captures the invariant causal mechanism of how users' attributes generate their preferences, the inferred preferences through $\mathbf{A}_{\text{spe}}$ are inherently causal-invariant and thus more reliable for OOD recommendation. Likewise, we can obtain $\mathbf{E}_{sha}^{sinv}$ and $\mathbf{E}_{sha}^{tinv}$. We then apply an attention mechanism to fuse $\mathbf{E}_{spe}^{tinv}$ and $\mathbf{E}_{sha}^{tinv}$ into comprehensive user preferences $\mathbf{E}_{u}^{*tinv}$, and similarly fuse $\mathbf{E}_{spe}^{s}$ and $\mathbf{E}_{sha}^{sinv}$ into $\mathbf{E}_{u}^{*s}$ [188].

5.2.4 LLM-guided Confounder Discovery

To further enhance the reliability of the inferred causal-invariant preferences, we propose to address the influence of observed confounders during the training phase. User reviews, which contain rich descriptions of user interactions, provide valuable infor-

mation about various interaction-related causal variables. LLMs' strong capabilities in understanding natural language text make them particularly suitable for extracting causal variables from unstructured reviews [1, 12]. However, direct extraction of confounders from reviews could be imprecise due to LLMs' potential hallucinations and the complex nature of causal relationships. To address this challenge, we leverage COAT [68], which provides a robust framework for extracting causal variables using LLMs with a systematic causal discovery approach grounded in causal theory. We propose to extend COAT by extracting observed confounders and working with reviews.

We propose an LLM-guided confounder discovery module that operates in an iterative manner. Specifically, we first use an LLM to propose candidate causal variables related to user interactions, employ a conventional causal discovery method to refine these variables, and then carefully extract observed confounders from the refined variable set. We use superscript τ to denote the τ-th iteration. The module consists of five key steps. Taking the target domain D^t as an example, the complete procedure is outlined in Algorithm 1, and we elaborate on each step in detail in the following subsections. Likewise, the same procedure is also applied to the source domain D^s.

5.2.4.1 Proposing Candidate Causal Variables

Building upon the factor proposal of COAT framework [68], we adapt the prompt design to extract causal variables from user reviews. Specifically, in Figure 5.2, we design a prompt ρ_{pro} for an LLM Φ to induce candidate causal variables that may affect our target variable Y^t (whether a user interacts with an item in D^t). The prompt contains three components: review samples, task instructions, and output format control. (i) *Review samples:* Target domain user reviews R^t are grouped by rating scores at first. Then a few reviews are sampled from each rating group (denoted as $\hat{R} \subset R^t$) due to LLM's context length limitation. The domain name is included in the prompt alongside these samples to establish the appropriate context. These review samples, coming from users who have interacted with items ($Y^t = 1$), tend to exhibit distinct patterns under different ratings, helping capture comprehensive causal variables. The

Algorithm 1: LLM-GUIDED CONFOUNDER DISCOVERY

Input: Target domain user reviews R^t, LLM Φ, maximum number of rounds τ_{max}, causal discovery algorithm $\mathcal{F}$, causal feedback operation ξ.
Output: Target domain confounder subspace $\mathcal{C}^t$.

Initialize $\tau \leftarrow 1$, $\mathcal{Z}^{\leq 0} \leftarrow \emptyset$, $\mathcal{Z}_{pool} \leftarrow \emptyset$, $\mathcal{C}^0_{pool} \leftarrow \emptyset$;
Prepare initial prompt ρ^1_{pro} and randomly select initial review samples $\hat{R}^1$;
while *not converged and* $\tau \leq \tau_{max}$ **do**
    /* Step 1: Proposing Candidate Causal Variables */
    $\mathcal{Z}^\tau \leftarrow \Phi(\rho^\tau_{pro}, \hat{R}^\tau)$;
    $\mathcal{Z}^{\leq \tau} \leftarrow \mathcal{Z}^{\leq \tau - 1} \cup \mathcal{Z}^\tau$;
    /* Step 2: Annotating Reviews into Structured Data */
    for *each review* $r_i \in R^t$ **do**
        for *each variable* $z_j \in \mathcal{Z}^\tau$ **do**
            $q_{i,j} \leftarrow \Phi(r_i, z_j, \rho_{ano})$;
    Form data matrix $\mathcal{Q}^{\leq \tau}$ from all annotation vectors;
    /* Step 3: Refining Variables via Causal Discovery */
    Filter variables from $\mathcal{Z}^{\leq \tau} \cup \mathcal{Z}_{pool}$ using CI test to obtain $\mathcal{Z}_f^{\leq \tau}$;
    $\mathcal{G}^\tau \leftarrow \mathcal{F}(\mathcal{Z}_f^{\leq \tau} \cup \{Y^t\})$;
    Extract Markov Blanket $\mathcal{Z}_{MB}^{\leq \tau} \leftarrow MB(Y^t)$ from $\mathcal{G}^\tau$;
    Update $\mathcal{Z}_{pool}$ with filtered variables and variables not in $\mathcal{Z}_{MB}^{\leq \tau}$;
    /* Step 4: Extracting Observed Confounders */
    if $\mathcal{C}^{\tau-1}_{pool} = \emptyset$ **then**
        Use zero-shot prompting to extract confounders;
    else
        Use few-shot prompting with examples to extract confounders;
    $\mathcal{C}^\tau_{pool} \leftarrow \mathcal{C}^{\tau-1}_{pool} \cup \Phi(\rho_{ext}, \mathcal{Z}_{MB}^{\leq \tau})$;
    /* Step 5: Generating Causal Feedback */
    $(\rho^{\tau+1}_{pro}, \hat{R}^{\tau+1}) \leftarrow \xi(\mathcal{Z}_{MB}^{\leq \tau}, R^t, \rho^\tau_{pro})$;
    $\tau \leftarrow \tau + 1$;
Encode each confounder in $\mathcal{C}^{\tau_{max}}_{pool}$ into embeddings;
Apply K-means clustering to obtain J representative confounders to form $\mathcal{C}^t$;
return $\mathcal{C}^t$;

domain name is provided to help the LLM Φ better interpret the context. (ii) *Task instructions:* To emulate the analytical process of human experts [95], the LLM Φ is guided to propose candidate causal variables through three phases: consideration of potential interaction-related variables, variable filtration to ensure semantic distinct-

(1) Proposing Candidate Causal Variables

You are an expert in the {domain_name} domain and skilled in analyzing user reviews to identify high-level causal variables that influence whether a user interacts with an item.

 # Review Samples
Rating: 5.0
Bag is large and comfortable. Like the color also. I apply a very low price promotion are excited to buy…
Rating: 4.0
The slippers are a perfect fit, easy to walk in (no slipping) and they keep my feet warm. I ususally don't like clogs but these are fine…
Rating: 3.0
These slippers worked great for a month. Then I had pitch them because they had holes in them. Don't buy these poor quality slippers…
Rating: 2.0
Bought for my daughter and we had to return it. Color was very pretty but it ran small. Otherwise, it would have been fine…
Rating: 1.0
This was very low quality and fit like a trash bag. I tried it on when I got it and put it directly in the garbage…

 # Tasks: Variable Proposal.
Your output should follow the three phases below.

Phrase 1: Consideration.
Analyze the review samples and identify high-level causal variables that explain why the user chose to interact with the item.
Hint: Abstract meaningful causal variables from the reviews and design corresponding criteria to assign them values from [-1, 0, 1] based on their expression in the review text.

(1) Proposing Candidate Causal Variables

Phrase 2: Variable Filtration.
Evaluate whether each proposed variable is worth retaining, based on the following criteria:
- Semantic Clarity: The variable captures one concrete, unique aspect and avoids overlap or redundancy with other variables.
- Interpretability: Its presence or absence can be reasonably inferred from the review content, even if not all reviews mention it.
- Causal Relevance: The variable should represent a meaningful factor that may influence the user's decision to interact with the item, rather than a superficial or stylistic feature.

Phrase 3: Final Output.
In this phase, report the causal variables you have selected after filtration. For each variable, follow the instructions below:
- Assign a value from [-1, 0, 1] to indicate how it is expressed in the review.
- For each value, provide a specific criterion based on the review content that determines when it applies.
- The criterion for value 0 must always be: "Otherwise; or not mentioned."

 # Output Format Control
For each proposed variable, use the following format:

{variable_name}
- 1: {positive_criterion}
- 0: Otherwise; or not mentioned.
- -1: {negative_criterion}

Figure 5.2: Illustration of the prompt template for proposing candidate causal variables.

ness and avoid redundancy, and output of filtered variables. Each filtered variable will be presented with its name and criterion that defines positive cases, negative cases, and cases otherwise or not mentioned, corresponding to values 1, -1, and 0, respectively. (iii) *Output format control:* The output format control specifies the required format for presenting the proposed variables and their criteria. The set of variables proposed in the τ-th round is denoted as $\mathcal{Z}^\tau = \Phi(\rho^\tau_{pro}, \hat{R}^\tau)$. We accumulate variables across rounds to form an initial candidate variable set $\mathcal{Z}^{\leq\tau} = \mathcal{Z}^1 \cup \cdots \cup \mathcal{Z}^\tau$.

5.2.4.2 Annotating Reviews into Structured Data

After obtaining candidate causal variables with their criteria, we apply these criteria to annotate all reviews in R^t and transform them into structured data. This two-step approach (first extracting variable criteria from sample reviews, then annotating all reviews based on these criteria) offers significant advantages over directly extracting variables from the entire review set in a single step. It ensures consistent annotation

standards across all reviews and allows us to progressively accumulate variables across rounds. For each review $r_i \in R^t$, we instruct LLM Φ to annotate it against each variable $z_j \in \mathcal{Z}^{\leq\tau}$ based on the variable's criterion, determining its value:

$$q_{i,j} := \Phi(r_i, z_j, \rho_{ano}), \tag{5.10}$$

where $q_{i,j} \in \{-1, 0, 1\}$ is the annotation result, and ρ_{ano} is the annotation prompt. The annotation results for each review r_i form a vector $\mathbf{q}_i = [q_{i,1}, q_{i,2}, ..., q_{i,|\mathcal{Z}^{\leq\tau}|}]$. These annotation vectors collectively form a data matrix $\mathcal{Q}^{\leq\tau} := [\mathbf{q}_1^T; \mathbf{q}_2^T; ...; \mathbf{q}_{|R^t|}^T]$. As we discover new variables in additional rounds, the matrix $\mathcal{Q}^{\leq\tau}$ can be expanded with new columns, providing the structured input required for the subsequent conventional causal discovery method.

5.2.4.3 Refining Variables via Causal Discovery

The candidate causal variables extracted by LLMs may contain inaccuracies, as LLMs might misinterpret review content or propose variables without true causal relationships with the target variable Y^t. To address this challenge, we implement the variable refinement via causal discovery. Inspired by the method introduced in [68], our refinement begins with a filtering process using conditional independence (CI) tests. We evaluate each variable in the candidate variable set $\mathcal{Z}^{\leq\tau}$ through an iterative process. Specifically, for each candidate variable z_j, we test whether it maintains a significant statistical association with Y^t when controlling for the set of variables already identified as relevant in previous iterations of testing. Variables that pass this test are confirmed as relevant and then retained, forming a filtered variable set $\mathcal{Z}_f^{\leq\tau}$, while others are eliminated. The filtering process not only removes irrelevant variables, but also helps address redundancy by eliminating variables that provide similar information as previously confirmed ones, thus providing a more reliable foundation for subsequent causal discovery.

After the initial filtering process, we feed the filtered variable in $\mathcal{Z}_f^{\leq\tau}$ and the tar-

get variable Y^t into the Fast Causal Inference (FCI) algorithm to discover the causal structure $\mathcal{G}^\tau$, which can be expressed as follows.

$$\mathcal{G}^\tau = \mathcal{F}(\mathcal{Z}_f^{\leq\tau} \cup \{Y^t\}), \tag{5.11}$$

where $\mathcal{F}$ denotes the FCI algorithm. FCI algorithm [106] is a conventional causal discovery method that accommodates potential confounders, making it well-suited for our context where causal variables proposed by LLM might be entangled with confounders [48]. In addition, FCI algorithm offers flexibility regarding different functional forms of causal relationships (e.g., linear or non-linear), which corresponds well to the diverse relationship patterns we need to identify in our causal structure learning [68].

From the discovered causal structure $\mathcal{G}^\tau$, we further extract the Markov Blanket of the target variable, denoted as $MB(Y^t)$. This extraction is crucial for our variable refinement process because it allows us to focus only on the most influential variables while discarding those that provide redundant or irrelevant information, thereby achieving the optimal balance between model simplicity and effectiveness for both prediction and causal understanding of user interactions. In the causal structure, $MB(Y^t)$ consists of Y^t's parents, Y^t's children, and the parents of Y^t's children other than Y^t itself [94]. Formally, the Markov Blanket satisfies: $Y^t \perp\!\!\!\perp \mathcal{Z} | MB(Y^t)$ for any set $\mathcal{Z}$ of proposed causal variables disjoint from both $MB(Y^t)$ and Y^t itself. This means that given $MB(Y^t)$, no other proposed variable provides any additional information about Y^t. By focusing on the Markov Blanket, our variable refinement process concludes with the identification of the refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$, which is the most compact set of variables necessary and sufficient for reasoning about the target variable Y^t.

To avoid potentially discarding valuable variables at early stages, we maintain a variable pool $\mathcal{Z}_{pool}$ that stores both variables eliminated during the initial CI test filtering and those from $\mathcal{Z}_f^{\leq\tau}$ that are not included in the Markov Blanket. This mechanism ensures that valuable variables have the opportunity to be reassessed in subsequent analyses, especially when certain variables might reveal their causal relevance only

under specific conditioning with other variables.

Overall, this three-step variable refinement process provides theoretical guarantees that complement the initial LLM-based variable proposal. By integrating the FCI algorithm along with supporting statistical methods, we systematically refine the proposed variables: CI test filtering removes both statistically irrelevant and redundant variables, the FCI algorithm discovers the underlying causal structure among remaining variables, and Markov Blanket extraction determines the minimal sufficient set for causal reasoning. This integration ensures that our refined variable set is theoretically grounded and contains only variables with true causal relationships to the target variable. It provides a reliable foundation for subsequent analyses, unlike approaches that directly use LLM, which may include inaccuracies.

5.2.4.4 Extracting Observed Confounders

After obtaining the refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$ through the variable refinement process, we proceed to extract observed confounders using LLMs. In the context of RSs, confounders are defined as causal variables that simultaneously influence both the treatment (typically manifested as user preferences) and the outcome (represented by user-item interactions) [189]. If not properly addressed, these confounders will introduce confounding biases into user preference modeling, thus resulting in the degraded recommendation performance.

Direct extraction of confounders from raw reviews using LLMs presents several limitations. Specifically, it risks identifying false confounders without theoretical validation, remains vulnerable to the hallucination problem of LLMs, and lacks guarantees of causal relevance to the target variable. To address these limitations, we propose to extract confounders from the refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$, which contains only variables with theoretically validated causal relationships to the target variable. Inspired by the method introduced in [1], we develop a Chain-of-Thought (CoT) prompting strategy instructing the LLM Φ to analyze each variable in $\mathcal{Z}_{MB}^{\leq\tau}$ in a step-by-step manner. This strategy is particularly valuable for confounder extraction as determining causal

(4) Extracting Observed Confounders

You are an expert in causal reasoning and skilled in identifying observed confounders that influence both user preferences and user-item interactions in the context of recommender systems.

 # Input Data and Context
You are given a refined set of causal variables extracted from user reviews in the {domain_name} domain. These variables have been selected based on their potential influence on whether a user interacts with an item.
Refined variable set: [<variable_1>, <variable_2>, ..., <variable_j>]

 # Tasks: Confounder Extraction.
Your output should follow the three phases below to identify observed confounders from the given refined variable set.

Phrase 1: Variable Classification.
For each variable, generate a brief description within the {domain_name} domain and classify it into one of two categories:
(a) item intrinsic attributes or explicit user preferences
(b) marketing, service, or other external factors
Hint: Variables in category (b) are more likely to be observed confounders.

(4) Extracting Observed Confounders

Phrase 2: Confounder Reasoning.
For each variable, analyze whether it qualifies as an observed confounder by evaluating the following two conditions:
- Direct effect: Does the variable directly affect whether users interact with items?
- Indirect effect: Does the variable also influence user preferences, which in turn affect user-item interactions?
Hint: A variable must satisfy both conditions to be considered an observed confounder.
Think step by step and provide your reasoning for each condition.

Phrase 3: Final Output.
In this phrase, report only those variables that satisfy both the direct and indirect effect conditions. If no variables satisfy both conditions, output 'None'.
 # Output Format Control
For each confirmed confounder, use the following format:
{confounder_name}
- Description: {confounder_description}
- Reasoning Process:
  - Direct effect: {direct_effect_reasoning}
  - Indirect effect: {indirect_effect_reasoning}

Figure 5.3: Illustration of the prompt template for extracting observed confounders using zero-shot prompting.

influence patterns requires complex judgment that benefits from explicit intermediate reasoning steps.

As illustrated in Figure 5.3, our prompt ρ_{ext} comprises three essential components to facilitate effective confounder extraction using an LLM Φ: (i) *Input data and context*, which provides the LLM Φ with the refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$ and the target domain name, enabling the accurate identification of observed confounders within the given context. (ii) *Task instructions*, which guide the LLM Φ to first generate a brief description for each variable in $\mathcal{Z}_{MB}^{\leq\tau}$, and then employ CoT reasoning to analyze each variable systematically. To facilitate efficient reasoning, the analysis includes a preliminary classification of variables into two categories: (a) item intrinsic attributes or explicit user preferences, and (b) marketing, service, or other external factors. This classification serves as a heuristic guide, as variables in category (b) are more likely to be confounders, while those in category (a) typically represent direct components of user preferences rather than confounders. Following this preliminary classification, the LLM Φ conducts a thorough evaluation of each variable against the confounder criteria: a variable is identified as a confounder only if it both directly affects user-item interactions and indirectly affects user-item interactions by influencing user prefer-

ences. (iii) *Output format control*, which specifies that the LLM Φ should produce a structured output containing all identified confounders, each with its description and the corresponding reasoning process justifying its identification as a confounder.

In our implementation, we initially employ zero-shot prompting using these three components without examples. After each extraction round, we store newly identified confounders, along with their descriptions and corresponding reasoning processes, into a confounder pool $\mathcal{C}_{pool}$ for subsequent analysis. Formally, at round τ, this process can be represented as:

$$\mathcal{C}_{pool}^{\tau} := \mathcal{C}_{pool}^{\tau-1} \cup \Phi(\rho_{ext}, \mathcal{Z}_{MB}^{\leq \tau}), \tag{5.12}$$

where $\mathcal{C}_{pool}^{0} = \emptyset$ initially, and $\Phi(\rho_{ext}, \mathcal{Z}_{MB}^{\leq \tau})$ represents the newly identified confounders at round τ. Once $\mathcal{C}_{pool}$ contains at least one confounder, we transition to few-shot prompting for subsequent rounds, using both a positive example (identified confounder) and a negative example (non-confounder). This transition involves adding an example demonstration component to our prompt ρ_{ext}, positioned before the input data and context component. For the positive example, we select the first identified confounder from $\mathcal{C}_{pool}$. For the negative example, we randomly select a variable that was explicitly determined not to be a confounder. The pairing of these contrasting examples (few-shot) creates clearer decision boundaries for the LLM, allowing it to better distinguish between confounders and non-confounders based on their causal influence patterns, rather than relying solely on the confounder definition alone (zero-shot) or a positive example alone (one-shot). Each example includes the refined variable set, domain information, variable name, its description, and reasoning process that led to its determination as a confounder or non-confounder. In addition, we instruct the LLM Φ to avoid redundantly identifying confounders that are already present in the confounder pool $\mathcal{C}_{pool}$.

After completing all extraction iterations, we transform each confounder in $\mathcal{C}_{pool}$ into an embedding. First, we concatenate all textual information of the confounder (i.e., name, description, and reasoning process). Next, we encode this concatenated

text using SentenceTransformer (all-MiniLM-L6-v2), chosen for its efficiency, cost-effectiveness, and comparable performance to other embedding models (such as OpenAI's text embeddings) [111]. This encoding produces a 384-dimensional embedding, which we then reduce to k dimensions using Principal Component Analysis (PCA), aligning with our user and item embedding dimensions. To minimize semantic redundancy and extract representative confounders, we apply K-means clustering to all the confounder embeddings and select J^s cluster centroids to form our confounder subspace $\mathcal{C}^t$ [189]. Similarly, we can obtain source domain confounder subspace $\mathcal{C}^s$.

5.2.4.5 Generating Causal Feedback

In the previous step, we have extracted observed confounders from the refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$. However, LLMs typically struggle to propose all relevant variables in a single iteration [137, 96]. When the proposed variable set is incomplete, the subsequent extraction of observed confounders is inevitably limited as well. To address this issue, we develop a causal feedback mechanism to guide the LLM Φ in proposing additional relevant variables in subsequent iterations, thereby progressively improving both variable discovery and confounder extraction throughout our iterative pipeline.

To determine whether our refined variable set $\mathcal{Z}_{MB}^{\leq\tau}$ is sufficient for explaining the target variable Y^t, we employ conditional entropy analysis. Conditional entropy quantifies the reduction in uncertainty about Y^t when conditioning on a given set of variables, providing a statistical measure of information completeness. If our refined variable set is incomplete, there exists a set of additional variables $\hat{\mathcal{Z}}$ that would further reduce the uncertainty about Y^t. We instruct the LLM Φ to propose new variables $\hat{z} \in \hat{\mathcal{Z}}$ that satisfy:

$$H_{\hat{R}^{\tau+1}}(Y^t|\mathcal{Z}_{MB}^{\leq\tau}) - H_{\hat{R}^{\tau+1}}(Y^t|\mathcal{Z}_{MB}^{\leq\tau}, \hat{z}(R^t)) > 0, \tag{5.13}$$

where $H(\cdot \mid \cdot)$ represents conditional entropy measured on $\hat{R}^{\tau+1}$, which is a subset of reviews specifically selected for the $(\tau+1)$-th iteration. When Eq.(5.13) holds,

it indicates that the current variables are insufficient to fully explain Y^t, highlighting the need to discover additional causal variables. Our key insight for identifying these variables is to focus on reviews where the current variable set $\mathcal{Z}_{MB}^{\leq\tau}$ performs poorly in explaining Y^t, as they likely contain information about undiscovered causal variables. We express this selection criterion as $\hat{R}^{\tau+1} = \arg\max_{\hat{R}\subset R^t} H_{\hat{R}}(Y^t|\mathcal{Z}_{MB}^{\leq\tau})$. Following the method introduced in [68], we frame this as a classification problem to identify reviews where the current variables fail to provide adequate predictive power. In our implementation, we use K-means clustering to group reviews based on their representations derived from the current variable set $\mathcal{Z}_{MB}^{\leq\tau}$. After dividing the reviews into multiple clusters, we compute the conditional entropy for each cluster and select the one with the highest conditional entropy as $\hat{R}^{\tau+1}$. This approach effectively identifies the subset $\hat{R}^{\tau+1}$ where significant uncertainty about Y^t remains despite conditioning on current variables, providing informative samples to guide the discovery of additional causal variables in next iteration.

Specifically, we utilize the causal feedback mechanism to generate inputs for the next iteration of variable proposal. This mechanism can be formalized as:

$$(\rho_{pro}^{\tau+1}, \hat{R}^{\tau+1}) = \xi(\mathcal{Z}_{MB}^{\leq\tau}, R^t, \rho_{pro}^{\tau}), \tag{5.14}$$

where ξ represents the causal feedback operation that: (1) selects the samples $\hat{R}^{\tau+1}$ from the highest conditional entropy cluster using the approach described earlier, and (2) generates an enhanced prompt $\rho_{pro}^{\tau+1}$ by modifying ρ_{pro}^{τ} to incorporate such samples and add feedback instructions. These instructions guide the LLM Φ to propose new causal variables not currently included in $\mathcal{Z}_{MB}^{\leq\tau}$, facilitating the construction of a more complete causal variable set in subsequent iterations.

5.2.5 Causal Deconfounding and Cross-domain OOD Recommendation

After obtaining the confounder subspaces $\mathcal{C}^s$ and $\mathcal{C}^t$ through LLM-guided confounder discovery, it becomes crucial to address the confounding biases introduced by the confounders within these subspaces. Without proper handling, such confounding biases can distort the inference of causal-invariant user preferences, thereby compromising the effectiveness of cross-domain OOD recommendation. To effectively mitigate these biases, we adopt the causal deconfounding approach proposed in [189], which leverages backdoor adjustment [193] to estimate the direct causal effect from user preferences E to interactions Y. This approach seamlessly integrates deconfounding into the recommendation training pipeline by modeling $P(Y|E,c)$ with an MLP, allowing us to estimate the causal effect as follows:

$$P(Y|do(E)) = \mathbb{E}_c[P(Y|E,c)] = \mathbb{E}_c[\varphi(\mathbf{E}_u^{*inv}, \mathbf{E}_v, \mathbf{c})], \tag{5.15}$$

where P denotes the probability and $\mathbb{E}$ denotes the expectation. $\varphi(\cdot)$ represents an MLP for predicting interaction probabilities [41] and c is the confounder selected from the corresponding confounder subspace $\mathcal{C}$. $\mathbf{E}_u^{*inv}$ denotes the comprehensive causal-invariant user preferences, and $\mathbf{E}_v$ represents the item embeddings. By implementing backdoor adjustment directly through the prediction process, this approach effectively blocks the backdoor paths from confounders to user-item interactions, enabling the inference of debiased causal-invariant user preferences that are critical for reliable cross-domain OOD recommendation.

Furthermore, inspired by the method introduced in [161], we apply a confounder selection function to control the deconfounding process by appropriately weighting the confounders. Taking the target domain D^t as an example, we formulate this con-

founder selection function as:

$$\psi(\mathbf{E}_u^{*tinv}, \mathbf{E}_v^t, \mathbf{c}) = \frac{\exp(\mathbf{W}_u^t \mathbf{E}_u^{*tinv} \cdot \mathbf{W}_{uc}^t \mathbf{c})}{2\sum_{\mathbf{c}'} \exp(\mathbf{W}_u^t \mathbf{E}_u^{*tinv} \cdot \mathbf{W}_{uc}^t \mathbf{c}')} + \frac{\exp(\mathbf{W}_v^t \mathbf{E}_v^t \cdot \mathbf{W}_{vc}^t \mathbf{c})}{2\sum_{\mathbf{c}'} \exp(\mathbf{W}_v^t \mathbf{E}_v^t \cdot \mathbf{W}_{vc}^t \mathbf{c}')}, \tag{5.16}$$

where $\mathbf{W}_u^t$, $\mathbf{W}_{uc}^t$, $\mathbf{W}_v^t$, and $\mathbf{W}_{vc}^t$ denote learnable transformation matrices. $\cdot$ denotes dot product and $\mathbf{c}'$ represents any confounder from $\mathcal{C}^t$. Next, the expectation $\mathbb{E}_c[\varphi(\mathbf{E}_u^{*tinv}, \mathbf{E}_v^t, \mathbf{c})]$ is expressed as:

$$\mathbb{E}_c[\varphi(\mathbf{E}_u^{*tinv}, \mathbf{E}_v^t, \mathbf{c})] = \varphi[\mathbf{W}_{fc}^t(\mathbf{E}_u^{*tinv} || \mathbf{E}_v^t || \sum\nolimits_c p(c)\mathbf{c}\psi(\mathbf{E}_u^{*tinv}, \mathbf{E}_v^t, \mathbf{c}))], \tag{5.17}$$

where $||$ denotes concatenation operation, and $\mathbf{W}_{fc}^t$ denotes a weight matrix of fully connected (FC) layer in the target domain. For practical purposes, the prior probability $p(c)$ is assumed as a uniform distribution. Moreover, the input of MLP in the target domain is $\mathbf{\Theta}_{in}^t = \mathbf{W}_{fc}^t(\mathbf{E}_u^{*tinv} || \mathbf{E}_v^t || \sum_c p(c)\mathbf{c}\psi(\mathbf{E}_u^{*tinv}, \mathbf{E}_v^t, \mathbf{c})$. Likewise, we can obtain $\mathbf{\Theta}_{in}^s$ in the source domain as well. Furthermore, we obtain the predicted interaction $\hat{y}_{ij}^t$ between user u_i and item v_j within the target domain as follows:

$$\hat{y}_{ij}^t = \delta_{out}(\delta_l(...\delta_2(\delta_1(\mathbf{\Theta}_{in}))...)), \tag{5.18}$$

where δ_l^t represents the transformation operation at the l-th layer of MLP. The MLP comprises l layers, with δ_{out} serving as the final output transformation.

In addition, we employ the cross-entropy loss to formulate the recommendation loss for the target domain as follows:

$$\mathcal{L}_{rec}^t = -\sum_{y \in \mathcal{Y}^{t+} \cup \mathcal{Y}^{t-}} \left[y \log(\hat{y}) + (1-y)\log(1-\hat{y}) \right], \tag{5.19}$$

where $\hat{y}$ represents the predicted interaction probability and y denotes the corresponding ground truth label. The set $\mathcal{Y}^{t+}$ contains all observed positive interactions in the target domain, while $\mathcal{Y}^{t-}$ consists of negative samples randomly selected from unobserved interactions to mitigate overfitting. Likewise, we can obtain $\hat{y}_{ij}^s$ and $\mathcal{L}_{rec}^s$ in

source domain.

Finally, we formulate the total loss function as follows:

$$\mathcal{L}^{total} = \mathcal{L}^{t}_{rec} + \beta_1 \mathcal{L}^{s}_{rec} + \beta_2 \mathcal{L}^{\text{cau}} + \beta_3 \mathcal{L}^{\text{dom}} + \beta_4 \|\mathbf{\Omega}\|_2, \tag{5.20}$$

where $\mathbf{\Omega}$ represents the set of all parameters that are optimized during training. The hyperparameters β_1 through β_4 control the contribution of each term.

5.2.6 Time Complexity Analysis

In this section, we analyze the time complexity of our proposed CICDOR framework. To standardize our analysis, we assume consistent parameters in CICDOR: the number of network layers L and embedding dimensions k. Below, we assess the time complexity of each of four functional modules individually.

(1) User Preference Disentanglement: This module consists of two main computational phases. First, during the embedding construction and graph propagation phase, we transform users and items into enriched embeddings. Since the text embedding is precomputed as a one-time preprocessing step, its time complexity can be excluded from the analysis. Assuming the heterogeneous graph contains $(m+n)$ nodes with an average of $\bar{N}$ neighbors per node, the time complexity of the L-layer GCN operation is $O(L(m+n)\bar{N}k)$, where m and n denote the number of users and items, respectively. Second, in the user preference disentanglement phase, we employ three MLP-based encoders along with a domain discriminator. The time complexity of this phase is dominated by the L-layer MLP operations, approximately $O(Lmk^2)$. Combining both phases and noting that typically $\bar{N} \ll (m+n)$, the overall time complexity of this module simplifies to $O(L(m+n)k + Lmk^2)$.

(2) Dual-level Causal Preference Learning: This module learns casual structure to infer causal-invariant user preferences at both domain-specific and domain-shared levels. The time complexity analysis spans both training and inference phases. In the training phase, learning the causal structure through adjacency matrices involves (i)

matrix multiplication for reconstruction loss computation with complexity $O(mk^2)$, and (ii) acyclicity constraint calculation with complexity $O(k^3)$. Additional regularization terms (e.g., sparsity, path constraints) contribute $O(k^2)$ complexity but are dominated by the higher-order terms. In the inference phase, inferring causal-invariant user preferences requires similar matrix multiplication operations with complexity $O(mk^2)$, while the attention-based fusion of preferences from both levels adds $O(mk)$ complexity. Thus, overall time complexity of this module simplifies to $O(mk^2 + k^3)$.

(3) LLM-guided Confounder Discovery: This module identifies causal variables and confounders through iterative processing across five key steps. (i) For variable proposal, the LLM processes review samples to generate candidate causal variables with complexity $O(\mathcal{T}_\Phi(|\bar{x}|))$, where $\mathcal{T}_\Phi(\cdot)$ represents the LLM inference cost, and $|\bar{x}|$ denotes the average input length. (ii) For review annotation, each review is annotated based on criteria of newly proposed variables, requiring complexity $O(|R||\mathcal{Z}|\mathcal{T}_\Phi(|\bar{x}|))$, where $|R|$ denotes the number of reviews and $|\mathcal{Z}|$ denotes the number of extracted variables. (iii) For variable refinement, CI test requires correlation matrix computation with complexity $O(|R||\mathcal{Z}|^2)$, the FCI algorithm contributes $O(|\mathcal{Z}|^4)$, and Markov Blanket extraction adds $O(|\mathcal{Z}|^2)$, resulting in a combined complexity of $O(|R||\mathcal{Z}|^2 + |\mathcal{Z}|^4)$. (iv) For confounder extraction, the LLM processes the refined variable set in one call with complexity $O(\mathcal{T}_\Phi(|\bar{x}|))$. (v) For causal feedback, K-means clustering and conditional entropy calculation require $O(|R||\mathcal{Z}|JI)$ complexity, where J is the number of cluster centroids and I is the number of clustering iterations. Additionally, after all iterations, we process extracted confounders through embedding generation, dimensionality reduction, and clustering with complexity $O(|\mathcal{C}|(|\bar{x}| + k^2 + kJI))$, where $|\mathcal{C}|$ represents the number of extracted confounders. Considering all steps across τ iterations, and noting that $|\mathcal{Z}| \ll |R|$, $\tau \ll |R||\mathcal{Z}|$, and computation costs are dominated by LLM inference, the overall time complexity of this module simplifies to $O(|R||\mathcal{Z}|\mathcal{T}_\Phi(|\bar{x}|) + \tau|R||\mathcal{Z}|^2)$.

(4) Causal Deconfounding and Cross-domain OOD Recommendation: This module employs backdoor adjustment to mitigate confounding biases and estimate the

direct causal effects of user preferences on interactions. Computing the selection weights for each confounder across all user-item pairs has a complexity of $O(mnJk)$. Subsequently, the process of concatenating user, item, and weighted confounder embeddings and passing them through an L-layer MLP contributes an additional complexity of $O(Lmnk^2)$. Consequently, the overall time complexity of this module is $O(mnk(J + Lk))$.

Overall, the time complexity of our CICDOR framework encompasses two distinct phases: the LLM-guided confounder discovery with complexity $O(|R||\mathcal{Z}|\mathcal{T}_\Phi(|\bar{x}|))$, and the recommendation model training (consisting of the remaining three modules) with complexity $O(mnLk^2)$. These expressions represent the simplified dominant terms from our detailed module-by-module analysis. The LLM-guided confounder discovery phase theoretically scales with the number of processed reviews, extracted variables, and LLM inference cost; however, in common practice, reviews are typically sampled at a computationally tractable scale to balance between representativeness and efficiency. Meanwhile, recommendation training phase scales linearly with user-item interaction space and quadratically with embedding dimensions, which typically becomes the computational bottleneck in large-scale recommendation scenarios.

5.3 Experiments on CICDOR

To validate the efficacy of our proposed CICDOR framework and its different modules, we conduct extensive experiments on two real-world datasets to answer the following research questions:

- **RQ1.** How does our CICDOR perform when evaluated across existing state-of-the-art approaches (see Section 5.3.2.1)?
- **RQ2.** How do various modules within our framework, namely, dual-level causal preference learning and LLM-guided confounder discovery, influence the OOD recommendation performance in the target domain? (see Section 5.3.2.2)?

Table 5.1: Statistics of the datasets.

Datasets	Domains	#Users	#Items	#Interactions	Density
Douban	Movie (source)	2106	9555	907219	4.508%
	Book (target)	2106	6777	95974	0.672%
Douban	Movie (source)	1666	9555	781288	4.908%
	Music (target)	1666	5567	69681	0.751%
Amazon	Electronics (source)	15761	51447	224689	0.027%
	Clothing (target)	15761	48781	133609	0.017%

- **RQ3.** How does our CICDOR's performance vary across different degrees of distribution shift (see Section 5.3.2.3)?

- **RQ4.** How does adjusting different hyperparameters influence the performance of our CICDOR (see Section 5.3.2.4)?

5.3.1 Experimental Settings

5.3.1.1 Experimental Datasets and OOD settings

We conduct experiments on two widely used CDR datasets: Douban [185] and Amazon [8]. For clarity, we refer to the specific domains as Douban-Movie, Douban-Book, Douban-Music, Amazon-Elec, and Amazon-Cloth throughout this paper. Based on these domains, we construct three source-target domain pairs: (1) Douban-Movie $\rightarrow$ Douban-Book, (2) Douban-Movie $\rightarrow$ Douban-Music, and (3) Amazon-Elec $\rightarrow$ Amazon-Cloth, where the first domain in each pair serves as the source domain and the second as the target domain. Table 5.1 presents the statistics of these domain pairs. Both datasets contain user-generated ratings and textual reviews. As the original ratings are explicit feedback, we convert them into implicit feedback by treating interactions with ratings of 4 or higher as positive instances. Moreover, we extract common users in both the source and target domains of each pair, thereby constructing a fully overlapping user scenario.

Following the existing works [43, 65], we conduct experiments under two OOD settings in the target domain:

- **User Degree Shift (OOD #1):** Users with high interaction degrees often exhibit distinct behavior patterns compared to those with lower degrees. In this setting, the training set is randomly sampled from the entire dataset, covering users with varying interaction degrees, while the testing set consists exclusively of high-degree users. This creates a distribution shift in user interaction degrees, challenging the model to generalize from the overall user population to a specific high-degree subset.

- **Region Shift (OOD #2):** User preferences often differ across different regions. In this setting, the training set contains randomly sampled users from all regions, while the testing set is composed of users from Beijing. This simulates a real-world scenario where regional distribution shift occurs between training and testing data.

For all experiments, we adopt an 8:1:1 split ratio for training, validation, and testing data. The user degree shift setting is applied to both Douban and Amazon datasets, while the region shift setting is applied only to the Douban dataset because only this dataset contains region information.

5.3.1.2 Implementation Details

For the hyperparameters in our causal structure learning component, we adopt the weights for different terms in the loss function as used in [65]. For the overall model training, we employ a negative sampling strategy with a 1:3 ratio of positive to negative samples. In addition, we set the trade-off parameters $\beta_1 = 1.0$, $\beta_3 = 1$, $\beta_4 = 0.00001$ and $\gamma = 0.5$ for the corresponding loss terms, while setting the number of cluster centroids $J = J^s = J^t$ to 10 and the dual-level causal loss weight β_2 to 0.5 as our default configuration.

For the prediction network, we implement a three-layer MLP to process the concatenated user, item, and confounder embeddings. Following [189], we set the dimension after the FC layer k_{in} as 128 and the output dimension k_{out} as 8, after exploring

values in ranges {64, 128} and {8, 16}, respectively. For the LLM-guided confounder discovery module, we sample 1000 users from each dataset and collect up to 5 reviews per user as input, utilizing gpt-4o-mini as the LLM with API calls determined at zero temperature to ensure reproducible outputs.]

To find out the optimal hyperparameters, we employ Optuna[2] with 50 trials for each model. The search space includes learning rate in $\{10^{-4}, 10^{-3}, 10^{-2}\}$, batch size in {128, 256, 512}, and embedding dimension k in {32, 64, 128}. Based on these optimization results, we use the Adam optimizer [52] for all models. In Section 5.3.2.4, we present an extensive parameter sensitivity analysis to examine how variations in the number of cluster centroids J in {2, 5, 10, 20, 50} and the dual-level causal loss weight β_2 in {0.1, 0.25, 0.5, 0.75, 1.0} affect the overall model performance.

5.3.1.3 Model Training

Our CICDOR framework adopts a two-phase training strategy designed to ensure stable convergence and accurate causal structure learning. This design is based on the principle that learning reliable causal relationships requires a foundation of high-quality user representations, as attempting to learn causal structures before the representations are well-formed would result in inaccurate or unstable causal structures.

In the first phase, we train the model for 60 epochs[3] to learn user attributes, domain-shared and domain-specific user preferences, while simultaneously performing deconfounding using fixed confounder embeddings extracted via our LLM-guided confounder discovery module. This phase generates debiased domain-shared and domain-specific user preferences, forming debiased comprehensive user preferences.

In the second phase, we continue training for an additional 40 epochs and introduce the dual-level causal loss in Eq. (5.9), which guides the model to learn dual-level causal structures. Specifically, this phase identifies the causal relationships between user attributes and the debiased domain-shared and domain-specific user preferences,

[2]https://optuna.org/

[3]We select the number of training epochs from the range $\{20, 40, 60, 80\}$ in each phase.

respectively. The resulting causal structures capture the invariant causal mechanisms and enable the inference of causal-invariant user preferences at both the domain-shared and domain-specific levels. These causal-invariant user preferences maintain consistency across distribution shifts, significantly improving OOD recommendation performance. For fair comparison, all baseline methods are also trained for 100 epochs to ensure complete convergence.

5.3.1.4 Evaluation Metrics

For our evaluation, we adopt two widely used metrics in the evaluation of RSs: Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG) [65]. During testing, for each positive user-item interaction, we randomly sample 99 items that the user has not interacted with as negative samples. The model then ranks this candidate set of 100 items (1 positive + 99 negative), and we evaluate its performance based on the position of the positive item within this ranked list. Throughout our experimental analysis, we focus specifically on the top-10 positions in the ranking results. To ensure the reliability of our findings, all experiments are conducted five times, and we report the average performance across these runs.

5.3.1.5 Comparison Methods

We compare our proposed CICDOR with twelve representative and state-of-the-art baseline models, which can be categorized into four groups: (I) Disentanglement-based OOD Recommendation, (II) Causality-based OOD Recommendation, (III) Single-target Cross-domain Recommendation (CDR), and (IV) Adaptation-based OOD Recommendation. Given that our work focuses on the novel setting of cross-domain OOD recommendation, which essentially addresses single-target CDR in OOD environments, we select representative and state-of-the-art methods from both single-target CDR and OOD recommendation literature as our baselines. Here, single-target CDR [183] refers to the paradigm of transferring knowledge from data-richer source

domain to improve recommendation performance in the data-sparser target domain. We do not include dual-target CDR methods [186] as baselines, because they aim to enhance the recommendation accuracy in both domains simultaneously, which differs from our setting.

Although several approaches explore invariance and generalization principles, we do not include them in the baseline models as they address different settings from our cross-domain OOD recommendation scenario. These excluded approaches target domain generalization [167, 162], sequential recommendation [171, 147], item cold-start recommendation [124], and recommendation with noisy interaction data [175]. Detailed descriptions of the selected baseline models are provided below.

- **DICE** [181] (I) decouples user interest and conformity by modeling their causal generation process, learning disentangled embeddings through cause-specific sampling guided by the collider effect.
- **DCCL** [178] (I) is a model-agnostic framework that addresses OOD problems and data sparsity by disentangling interest and conformity through contrastive learning with cause-specific sample augmentation.
- **CausPref** [43] (II) learns invariant user preferences from implicit feedback via causal structure learning and enhances robustness to distribution shift through anti-preference negative sampling.
- **COR** [123] (II) formulates user feature shift as an intervention and performs causal inference with a tailored VAE to estimate post-intervention interaction probabilities for robust OOD recommendation.
- **InvCF** [156] (II) identifies causally invariant preference representations by disentangling them from popularity semantics, enabling consistent generalization in real-world scenarios with shifting item popularity.
- **PopGo** [155] (II) improves OOD generalization by mitigating interaction-level popularity shortcuts, using a learned shortcut model to adjust predictions and emphasize

true user preferences over spurious popularity shortcuts.

- **CausalDiffRec** [174] (II) enhances OOD recommendation by eliminating environmental confounders through backdoor adjustment and learning environment-invariant graph representations via a causal diffusion process.
- **PTUPCDR** [194] (III) generates personalized bridge functions via a task-optimized meta network, facilitating stable and personalized preference transfer from the source to the target domain.
- **CUT** [58] (III) employs a two-phase training strategy that first captures user similarities in the target domain and then transfers source-domain information selectively, leveraging a user transformation module and contrastive learning to avoid the relationship distortion.
- **CDCOR** [65] (III) improves OOD recommendation by transferring cross-domain knowledge, leveraging a domain adversarial network to extract the shared user preferences and a causal structure learner to model the invariant relationships under distribution shifts.
- **DR-GNN** [116] (IV) integrates distributionally robust optimization (DRO) into GNN-based recommendation by treating GNN as a smoothing regularizer and injecting small perturbations into sparse neighbor distributions to enhance robustness against distribution shifts.
- **DT3OR** [148] (IV) introduces a dual test-time training strategy for OOD recommendation, adapting models to distribution shifts by learning invariant user preferences and variant user/item features through self-distillation and contrastive learning.

Table 5.2: Quantitative assessment of CICDOR against representative and state-of-the-art baseline models across two datasets under different OOD settings, measured by HR@10 and NDCG@10 metrics. The best results appear in bold text, with the best baseline results identified by underlining. * represents $p < 0.05$ when CICDOR is compared to the best baseline in paired t-test [188].

Dataset		**Douban**								**Amazon**	
Domain: Source → Target		**Movie → Book**				**Movie → Music**				**Elec → Cloth**	
Setting		OOD #1		OOD #2		OOD #1		OOD #2		OOD #1	
Metric		HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
Disentanglement-based OOD Recommendation	DICE	0.3146	0.1735	0.3406	0.1827	0.2825	0.1519	0.3047	0.1729	0.5311	0.2946
	DCCL	0.3524	0.1868	0.3812	0.2165	0.3172	0.1723	0.3438	0.1842	0.5948	0.3385
Causality-based OOD Recommendation	CausPref	0.2911	0.1604	0.3164	0.1781	0.2601	0.1416	0.2840	0.1544	0.4873	0.2682
	COR	0.3115	0.1716	0.3423	0.1839	0.2758	0.1487	0.2996	0.1655	0.5184	0.2778
	InvCF	0.3253	0.1761	0.3610	0.1987	0.2905	0.1589	0.3177	0.1788	0.5441	0.3017
	PopGo	0.2988	0.1652	0.3197	0.1726	0.2682	0.1434	0.2913	0.1606	0.5022	0.2755
	CausalDiffRec	0.3595	0.1981	0.3919	0.2242	0.3246	0.1758	0.3581	0.1968	0.6089	0.3394
Single-target Cross-domain Recommendation	PTUPCDR	0.3108	0.1714	0.3374	0.1791	0.2721	0.1473	0.2892	0.1574	0.5107	0.2763
	CUT	0.2967	0.1642	0.3248	0.1757	0.2653	0.1422	0.2865	0.1557	0.4965	0.2726
	CDCOR	0.3586	0.1974	0.3885	0.2196	0.3217	0.1745	0.3543	0.1956	0.6044	0.3381
Adaptation-based OOD Recommendation	DR-GNN	0.3479	0.1845	0.3803	0.2162	0.3089	0.1704	0.3348	0.1785	0.5792	0.3258
	DT3OR	0.3713	0.2068	0.4071	0.2289	0.3336	0.1778	0.3626	0.1993	0.6239	0.3570
Our Model and its Variants	CICDOR	**0.3951***	**0.2264***	**0.4298***	**0.2493***	**0.3564***	**0.1961***	**0.3882***	**0.2194***	**0.6584***	**0.3867***
	w/o dual-level	0.3368	0.1786	0.3680	0.2051	0.3015	0.1663	0.3280	0.1765	0.5648	0.3185
	w/o specific-level	0.3675	0.2003	0.3994	0.2272	0.3326	0.1774	0.3593	0.1976	0.6140	0.3473
	w/o shared-level	0.3644	0.1997	0.3971	0.2268	0.3309	0.1768	0.3574	0.1962	0.6121	0.3467
	w/o confounder	0.3689	0.2054	0.4032	0.2277	0.3345	0.1781	0.3618	0.1989	0.6166	0.3489
	w/ direct-LLM	0.3816	0.2169	0.4163	0.2381	0.3462	0.1843	0.3745	0.2091	0.6385	0.3624
	w/ gpt2glm	0.3883	0.2194	0.4240	0.2425	0.3521	0.1937	0.3807	0.2163	0.6488	0.3742
Improvement (CICDOR vs. best baselines)		6.41%	9.48%	5.58%	8.91%	6.83%	10.29%	7.06%	10.09%	5.53%	8.32%

5.3.2 Experimental Results and Analysis

5.3.2.1 Performance Comparison (for RQ1)

Table 5.2 presents the performance evaluation of CICDOR[4] against various baseline models across two datasets under different OOD settings, using HR@10 and NDCG@10 as metrics. It should be noted that while Single-target CDR models are trained on both domains, we only report their performance in the target domain, as these models are specifically designed to improve recommendation performance in the target domain. From Table 5.2, we can observe that:

(1) Our CICDOR demonstrates superior performance compared to Disentanglement-based OOD Recommendation baselines across various OOD settings, achieving average improvements of 18.62% and 23.32% w.r.t. HR@10 and NDCG@10, respectively. While DICE [181] and DCCL [178] employ factor disentanglement techniques, they are limited to handling only specific types of distribution shifts. In

[4]Due to page constraints, we only report results with embedding dimension $k = 64$, though CICDOR consistently outperforms baseline models across other embedding dimensions as well.

contrast, our CICDOR can effectively capture the invariant causal structure underlying user preferences. This enables us to infer causal-invariant user preferences for OOD recommendation, resulting in better performance across various types of distribution shifts;

(2) While existing Causality-based OOD Recommendation methods consider knowledge that remains invariant across different distributions, they overlook the potential benefits of leveraging cross-domain invariant knowledge to enhance OOD recommendation in the target domain, thus yielding suboptimal results. Our CICDOR outperforms the best-performing Causality-based OOD Recommendation baseline, CausalDiffRec [174], by an average of 9.18% and 12.49% with respect to HR@10 and NDCG@10, respectively. The improvements stem from CICDOR's ability to utilize source domain data to learn the domain-shared causal structure, enabling the inference of domain-shared causal-invariant user preferences that can be transferred to facilitate OOD recommendation in the target domain;

(3) Compared with the best-performing baseline, DT3OR [148], our CICDOR achieves average improvements of 6.28% and 9.42% w.r.t. HR@10 and NDCG@10, respectively. These improvements stem from two key innovations. First, our dual-level causal preference learning module simultaneously captures both domain-shared and domain-specific causal-invariant user preferences, effectively addressing both CDDS and SDDS. This approach differs from existing single-target CDR baselines that either completely overlook SDDS (e.g., PTUPCDR [194] and CUT [58]) or only partially mitigate it by solely leveraging domain-shared the causal structure while neglecting the domain-specific causal structure in the target domain (e.g., CDCOR [65]). Second, our LLM-guided confounder discovery module accurately identifies and extracts observed confounders from user reviews by leveraging LLM's knowledge and reasoning capabilities. By deconfounding these observed confounders, we ensure that the inferred causal-invariant user preferences can reflect users' true preferences rather than being influenced by confounding bias. The combination of

these two complementary modules enables CICDOR to capture debiased comprehensive causal-invariant user preferences, resulting in superior cross-domain OOD recommendation performance across various OOD settings.

5.3.2.2 Ablation Study (for RQ2)

To evaluate the contribution of each component in improving the OOD recommendation performance of our model, we construct six variants of CICDOR and conduct the ablation study across various OOD settings on two datasets.

Impact of Dual-level Causal Preference Learning. We construct three variants: **w/o dual-level** by removing the entire dual-level causal preference learning module, **w/o specific-level** by removing the domain-specific causal structure learning component, and **w/o shared-level** by removing the domain-shared causal structure learning component. From Table 5.2, we can observe that our CICDOR model outperforms **w/o dual-level** with an average improvement of 19.92%. This significant performance difference demonstrates that the dual-level causal preference learning module is well suited for addressing distribution shifts in CDR. Without this module, the model lacks the ability to infer causal-invariant user preferences that remain invariant across different distributions. By learning the underlying causal structures and inferring causal-invariant user preferences, our CICDOR can achieve superior OOD recommendation performance in scenarios where both CDDS and SDDS co-exist.

Furthermore, the results show that our CICDOR outperforms **w/o specific-level** and **w/o shared-level** with average improvements of 9.32% and 9.81%, respectively. These comparable performance differences indicate that both levels of causal structure learning are essential and contribute almost equally to the model's overall effectiveness. By learning the domain-specific causal structure in the target domain, CICDOR infers domain-specific causal-invariant user preferences that can mitigate SDDS. Meanwhile, by learning the domain-shared causal structure, CICDOR infers domain-shared causal-invariant user preferences that facilitate invariant knowledge transfer from the source domain, thus effectively addressing CDDS and providing valuable

support for OOD recommendation in the target domain. The fusion of both levels of causal-invariant user preferences yields comprehensive causal-invariant user preferences, thereby enabling our CICDOR to obtain better recommendation performance under various OOD settings.

Impact of LLM-guided Confounder Discovery. We modify CICDOR to form a variant, namely **w/o confounder**, by removing the LLM-guided confounder discovery module. From Table 5.2, we can observe that our CICDOR model outperforms **w/o confounder** with an average improvement of 8.53%. This demonstrates that extracting observed confounders for subsequent deconfounding is crucial for capturing accurate comprehensive causal-invariant user preferences, thereby achieving better OOD recommendation performance.

Moreover, we design another variant, namely **w/ direct-LLM**, by replacing our LLM-guided confounder discovery module with a simplified approach that directly uses gpt-4o-mini to extract confounders from user reviews in a single step. We find that the complete CICDOR model outperforms **w/ direct-LLM** by an average of 4.36%. This demonstrates the effectiveness of our module over simply relying on LLM alone. Without the iterative process involving variable proposal, review annotation, variable refinement and causal feedback, simply using LLM to extract observed confounders may generate hallucinations or identify spurious variables as confounders. Our LLM-guided confounder discovery module addresses this limitation by providing theoretical guidance through CI test and FCI algorithm, which help filter and refine the variables before extracting confounders. This process ensures that extracted confounders are causally valid and accurate, leading to more effective deconfounding and ultimately better OOD recommendation performance.

In addition, we construct a third variant, **w/ gpt2glm**, where we replace gpt-4o-mini with glm-4-9b-chat[5] in our LLM-guided confounder discovery module. We can observe that our CICDOR model achieves an average improvement of 1.98% over **w/ gpt2glm**. This shows that the capability of LLM affects the accuracy of extracted

[5]https://huggingface.co/THUDM/glm-4-9b-chat

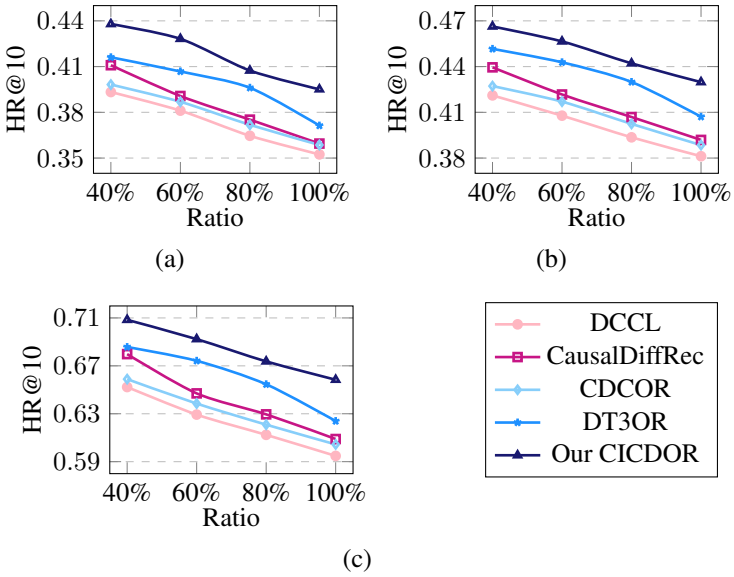

Figure 5.4: Performance comparison (HR@10) between CICDOR and baselines under different degrees of distribution shift: (a) Douban-Movie → Douban-Book OOD #1, (b) Douban-Movie → Douban-Book OOD #2, and (c) Amazon-Elec → Amazon-Cloth OOD #1.

confounders. In our implementation, we select gpt-4o-mini to balance performance benefits with economic considerations, as it enables us to extract more accurate confounders for deconfounding. After deconfounding, we can capture debiased comprehensive causal-invariant user preferences for better OOD recommendation.

5.3.2.3 Stability Analysis (for RQ3)

To investigate our model's recommendation performance under varying degrees of distribution shift, we compared CICDOR with the best-performing baseline from each group across different shift intensities. For user degree shift (OOD #1), we control the proportion of samples from users with high interaction degrees in the test set, ranging from 40% to 100%, with the remaining portions filled by samples from users with low interaction degrees. Similarly, for region shift (OOD #2), we control the proportion of samples from Beijing users in the test set from 40% to 100%, with non-Beijing user samples filling the remainder.

As shown in Figure 5.4, the recommendation performance[6] of all models decreases

[6]Due to space limitations, Figure 5.4 and Figure 5.5 only show results for Douban-Movie →

as distribution shift intensifies. However, CICDOR consistently achieves the best performance across all shift ratios. Moreover, CICDOR exhibits the most gradual performance decline as distribution shift increases, demonstrating superior stability in OOD generalization. Both the superior performance and enhanced stability can be attributed to CICDOR's ability to learn dual-level invariant causal structures from cross-domain knowledge, extract accurate observed confounders, and mitigate these confounders' negative effects on the inferred causal-invariant user preferences. This enables our CICDOR to produce more accurate recommendations across various OOD settings.

5.3.2.4 Parameter Sensitivity (for RQ4)

Impact of the number of cluster centroids J. We examine how varying the number of cluster centroids J affects CICDOR's recommendation performance by testing $J = J^s = J^t$ across $\{2, 5, 10, 20, 50\}$. Figures 5.5(a)-(c) present the results across different OOD settings on both Douban and Amazon datasets. We can observe that the recommendation performance of our CICDOR improves consistently as J increases up to 10, confirming the effectiveness of confounders extracted through our LLM-guided confounder discovery module. At $J = 10$, the observed confounders represented by these cluster centroids provide sufficient information for accurate deconfounding. Beyond this threshold, performance gains become negligible or even slightly decrease in some cases, suggesting that additional centroids may introduce noise rather than capture meaningful confounders. Although our confounder selection function has the ability to select relevant confounders, an excessive number of potential confounders appears counterproductive. Balancing model complexity and performance, we set $J = J^s = J^t = 10$ for all experiments. This threshold indicates that approximately 10 key observed confounders exist in each domain, and our CICDOR can successfully extract them and then mitigate their negative effects on the inferred causal-invariant user preferences to improve the OOD recommendation performance.

Douban-Book and Amazon-Elec $\rightarrow$ Amazon-Cloth. Similar trends are also observed in the Douban-Movie $\rightarrow$ Douban-Music experiments.

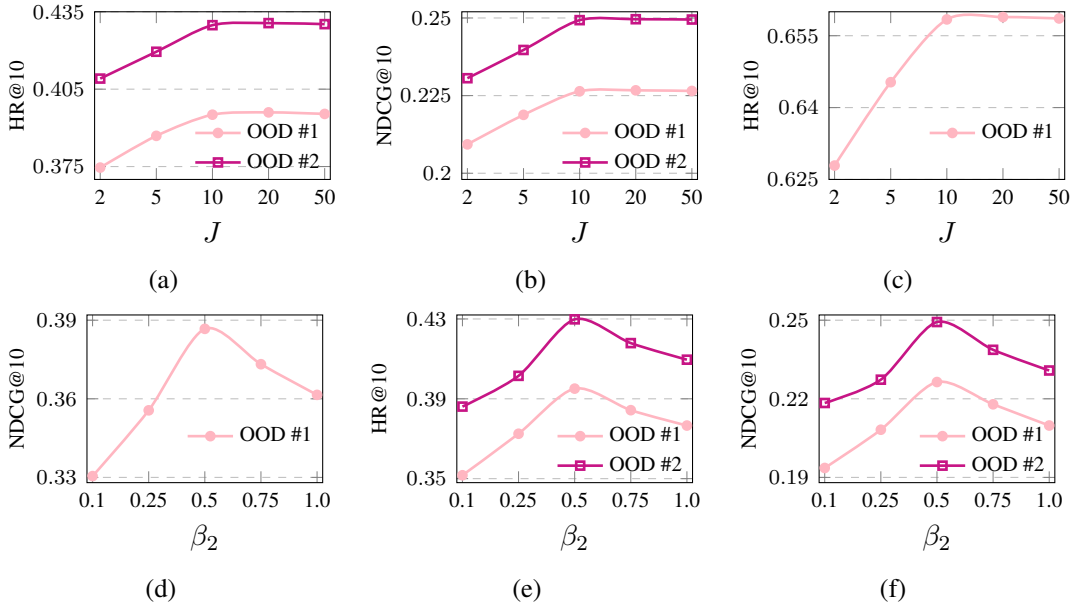

Figure 5.5: (a)-(b): Impact of the number of cluster centroids J on Douban-Movie $\rightarrow$ Douban-Book. (c)-(d): Impact of J and the weight of dual-level causal loss β_2 on Amazon-Elec $\rightarrow$ Amazon-Cloth. (e)-(f): Impact of β_2 on Douban-Movie $\rightarrow$ Douban-Book.

Impact of the weight of dual-level causal loss β_2. We investigate how the weight of dual-level causal loss β_2 affects our model's OOD recommendation performance by varying it across $\{0.1, 0.25, 0.5, 0.75, 1.0\}$. Figures 5.5(d)-(f) illustrate that CICDOR achieves optimal OOD recommendation performance when $\beta_2 = 0.5$. This finding suggests that this specific value creates an ideal balance between the dual-level causal loss and recommendation loss. At this balanced point, our CICDOR can effectively learn domain-shared and domain-specific causal-invariant user preferences while maintaining recommendation accuracy. This balance enables our model to leverage invariant knowledge from the source domain to enhance OOD recommendation performance in the target domain, particularly when both CDDS and SDDS exist. When β_2 is either too large or too small, the model struggles to balance recommendation accuracy with OOD generalization capability, leading to suboptimal results. Based on these observations, we select $\beta_2 = 0.5$ for all our experiments to better optimize the recommendation losses and dual-level causal loss.

5.4 Summary

In this chapter, we have proposed a new setting of cross-domain OOD recommendation, and proposed a novel Causal-Invariant Cross-Domain Out-of-distribution Recommendation framework, called CICDOR. CICDOR consists of a dual-level causal preference learning module to infer domain-specific and domain-shared causal-invariant user preferences, and an LLM-guided confounder discovery module to iteratively identify and refine causal variables from user reviews to extract observed confounders. Through effective deconfounding of the extracted confounders via backdoor adjustment, CICDOR can obtain debiased comprehensive causal-invariant user preferences that significantly improve the OOD recommendation performance in the target domain. Extensive experiments on two real-world datasets validate the superiority of CICDOR over state-of-the-art baselines across various OOD settings.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

Despite the impressive progress achieved by existing CDR methods, significant challenges remain in effectively modeling user preferences in complex real-world scenarios. This thesis identifies and addresses three critical limitations in existing CDR methods. First, existing CDR methods primarily focus on domain-shared and domain-specific user preferences, while overlooking domain-independent user preferences, resulting in incomplete modeling of user preferences. However, even with comprehensive preference modeling, user-item interactions may still be influenced by observed confounders (e.g., free shipping, sales promotion), which simultaneously affect both preferences and interactions, thereby introducing confounding bias into user preference modeling. In addition to these limitations, most CDR methods are built on the IID assumption, failing to model causal-invariant user preferences that remain invariant across distributions, and thus struggle to maintain robust recommendation performance under the co-existing CDDS and SDDS in real-world OOD environments. These limitations collectively impair the effectiveness of existing CDR methods, necessitating the development of more advanced frameworks that can model user preferences in a comprehensive, debiased, and causal-invariant manner.

To advance towards more effective CDR, this thesis tackles four key research challenges derived from the aforementioned limitations. To address the limitation of incomplete user preference modeling, the first challenge is **CH1**: '*How to effectively*

decouple domain-independent information from domain-specific information, in addition to domain-shared information, to capture comprehensive user preferences on each domain, thereby improving the recommendation performance?'. To address the limitation of biased user preference modeling, two additional challenges are identified as follows. **CH2**: '*How to effectively extract observed confounders to comprehensively understand user-item interactions?*' And **CH3**: '*How to preserve the positive impacts of observed confounders on predicted interactions, while eliminating their negative impacts on capturing comprehensive user preferences, thereby enhancing the recommendation accuracy?*' Furthermore, the inability to model causal-invariant user preferences across distributions becomes particularly critical in real-world OOD scenarios where both cross-domain and single-domain distribution shifts co-exist. These co-existing distribution shifts significantly hinder invariant user preference modeling in CDR. Specifically, the fourth challenge is **CH4**: '*How to simultaneously address cross-domain and single-domain distribution shifts to achieve reliable recommendation under OOD environments in CDR?*'

To address the above challenges, we have proposed the following solutions:

- To address **CH1** in Chapter 3, we propose a disentanglement module to effectively decouple the domain-independent and domain-specific user preferences. The disentanglement module also extracts the domain-shared user preferences from augmented user representations, which can be transferred to both domains to provide the valuable information. We then apply the attention mechanism to combine the above three essential components of user preferences to capture more comprehensive user preferences in each domain, which can improve the recommendation performance on each of both domains. Chapter 3 introduces domain-shared, domain-specific and domain-independent information respectively with examples, and further differentiates them. Inspired by these observations, a novel disentanglement-based framework with interpolative data augmentation for dual-target CDR is proposed. To the best of our knowledge, this

is the first work in the literature that explicitly takes domain-independent information into consideration in addition to domain-shared and domain-specific information, and decouples it to capture more comprehensive user preferences for cross-domain recommendation.

- To address **CH2** in Chapter 4, we propose a confounder disentanglement module to effectively disentangle observed SDCs and CDCs. In this module, we devise a dual adversarial structure to disentangle SDCs in each domain and apply half-sibling regression to decouple CDCs, thus obtaining a comprehensive understanding of user-item interactions in each of both domains. Subsequently, to address **CH3**, we propose a causal deconfounding module to deconfound disentangled observed SDCs and CDCs via backdoor adjustment. Specifically, we design a confounder selection function to mitigate such observed confounders' negative effects, thereby recovering debiased comprehensive user preferences. We then incorporate the observed confounders' positive effects into such debiased user preferences to enhance the recommendation accuracy in both domains. Chapter 4 briefly reviews SDCs and then provides an in-depth analysis of CDCs. Inspired by these observations, a novel causal deconfounding framework via confounder disentanglement for dual-target CDR is proposed. To the best of our knowledge, this is the first work in the literature that explicitly decouples observed CDCs, and incorporates observed confounders' positive impacts into debiased comprehensive user preferences for dual-target CDR.

- To address **CH4** in Chapter 5, we propose a dual-level causal preference learning module. This module first leverages a user preference disentanglement module to extract domain-specific and domain-shared user preferences. In addition, dual-level causal structures, represented as DAGs, are learned for both domain-specific and domain-shared levels of modeling. Based on these causal structures, the corresponding causal-invariant user preferences are inferred at each level, thus tackling CDDS and SDDS simultaneously under OOD environ-

ments in CDR. To address **CH2**, Chapter 5 further proposes an LLM-guided confounder discovery module. This module first employs an LLM to extract candidate interaction-related causal variables from user reviews and transform them into structured data. The structured data are then fed into a conventional causal discovery method (i.e., the FCI algorithm) to uncover their underlying causal relationships, with conditional independence tests used to eliminate redundant variables. The observed confounders are first identified by leveraging the LLM and then stored in the confounder pool. Next, when such remaining variables cannot fully explain the user-item interactions, feedback is constructed through the LLM to guide the discovery of additional causal variables, forming an iterative process that continuously refines the confounder discovery. Finally, a confounder selection function is used to control confounders from the confounder pool, enabling effective deconfounding via backdoor adjustment, thereby ensuring accurate causal-invariant user preference inference.

6.2 Future Work

This thesis focuses on dual-target CDR and cross-domain OOD recommendation problems. To address the abovementioned four challenges, we propose three frameworks and validate their superiority through extensive experiments. Nevertheless, several important issues remain unresolved, pointing to promising directions for future research:

- **Privacy-Preserving CDR.** Most existing CDR approaches assume unrestricted access to plaintext data across domains, overlooking the critical data isolation problem in real-world scenarios. This assumption becomes increasingly problematic as CDRs deal with sensitive user information such as browsing histories, purchase records, and demographic profiles held by different commercial entities. Existing privacy-preserving techniques for CDR are limited to simplistic models and struggle to maintain recommendation quality while ensuring data protection. A promising research direction involves developing sophisticated

privacy-preserving frameworks specifically designed for CDR. This could include exploring federated learning approaches for knowledge transfer without sharing raw user data, cryptographic protocols for secure collaborative filtering, and machine unlearning approaches that respect users' right to be forgotten.

- **Explainable CDR.** Most existing CDR approaches operate as black boxes, providing little insight into how preferences from one domain influence recommendations in another. This lack of interpretability significantly undermines user trust, especially when users receive seemingly unrelated recommendations based on their activities in different domains. Unlike single-domain recommendation explainability, CDR faces unique challenges due to the heterogeneous nature of cross-domain data and the complex knowledge transfer mechanisms involved. Existing explanation techniques often treat each domain in isolation, failing to address the conceptual differences between domains and rarely illuminating how user preferences are transferred across domains. A promising research direction involves developing specialized explanation frameworks for CDR that can reveal the underlying connections between user behaviors in different domains while maintaining coherent and intuitive explanations.

- **LLM-Empowered CDR.** Despite growing interest in applying LLMs to CDR, existing approaches often employ simplistic methods, such as using LLMs merely as feature extractors or directly as recommenders. These approaches fail to effectively bridge LLMs' comprehensive world knowledge with the domain-specific expertise embedded in user-item interactions. A promising research direction involves developing vertical domain-specific LLMs tailored for CDR scenarios. This could include exploring fine-tuning strategies on domain-specific data while preserving general reasoning capabilities, enabling models to simultaneously capture unique characteristics of individual domains and leverage common knowledge across them.

Appendix A

The Notations in the Thesis

Table A.1: The Notations in Chapter 3

Notations	Explanations
k	the dimension of embedding matrix
m	the number of users
n	the number of items
R	the rating matrix
$\mathcal{U}$	the set of users
$\mathcal{V}$	the set of items
P	the predicted domain probability
O	the ground truth domain label
$\mathbf{E}_u$	the graph embedding matrix of users
$\mathbf{E}_v$	the graph embedding matrix of items
$\mathbf{Z}_{sha}$	the domain-shared user preferences
$\mathbf{Z}_{spe}$	the domain-specific user preferences
$\mathbf{Z}_{ind}$	the domain-independent user preferences
$\mathbf{E}_u^*$	the comprehensive user preferences
y	the user-item interaction
$*^{aug}$	the notation for data augmentation
$\mathbf{E}_u^{aug}$	the augmented user representations of common users
$*^A$, $*^B$	the notations for domains A and B, e.g., n^A represents the number of items on domain A
$\hat{*}$	the predicted notations, e.g., $\hat{y}_{ij}$ represents the predicted interaction of user u_i on item v_j
G	the heterogeneous graph, where Q is the set of user-item relationships

Table A.2: The Notations in Chapter 4

Notations	Explanations
d	the embedding dimension
m	the number of users
n	the number of items
$\mathcal{U}$	the set of users
$\mathcal{V}$	the set of items
$\mathbf{R} \in \{0,1\}^{m \times n}$	the interaction matrix
y_{ij}	the interaction of user u_i on item v_j
$\hat{y}_{ij}, \hat{y}_{ik}$	the predicted user-item interactions
$*^A$, $*^B$	the notation for domain A and B, respectively
$\mathbf{Z}_{sha}$	the domain-shared user preferences
$\mathbf{Z}_{spe}$	the domain-specific user preferences
$\mathbf{Z}_{ind}$	the domain-independent user preferences
$\mathbf{E}_u^*$	the comprehensive user preferences
$\mathbf{E}_v$	the item embeddings
$\mathbf{C}_{sd}$	the single-domain confounders
$\mathbf{C}_{cd}$	the cross-domain confounders
$S(\cdot), T(\cdot)$	the generator in domain A and B, respectively
$H(\cdot)$	the discriminator
J	the number of cluster centroids
$p(c)$	the uniform distribution for prior probability
λ	the weight of cycle consistency loss
α	the regularization parameter
$\mathbf{W}$	the weight matrix

Table A.3: The Notations in Chapter 5

Notations	Explanations
$\mathcal{U}$	the set of users
$\mathcal{V}$	the set of items
m	the number of users
n	the number of items
R	the user reviews
$\mathbf{Y}$	the interaction matrix
y_{ij}	the interaction of user u_i on item v_j
$\hat{y}_{ij}$	the predicted user-item interaction
k	the embedding dimension
$*^s$, $*^t$	the notation for source and target domain, respectively
$\mathbf{E}^{inv}$	the causal-invariant user preferences
$\mathbf{E}_{sha}$	the domain-shared user preferences
$\mathbf{E}_{spe}$	the domain-specific user preferences
$\mathbf{E}_{att}$	the user attributes
$\mathbf{E}_u^*$	the comprehensive user preferences
$\mathbf{E}_v$	the item embeddings
$\mathbf{A}$	the weighted adjacency matrix, i.e., DAG
Φ	the LLM
ρ^τ	the prompt used at the τ-th round
$\mathcal{F}$	the causal discovery algorithm
ξ	the causal feedback operation
$\hat{R}^\tau$	the subset of review samples used at the τ-th round
$\mathcal{Z}^{\leq\tau}$	the set of variables proposed at the first τ-th round
$\mathcal{Z}_f^{\leq\tau}$	the filtered variable set at the first τ-th round
$\mathcal{Z}_{MB}^{\leq\tau}$	the refined variable set at the first τ-th round
$\mathcal{Z}_{pool}$	the variable pool
$MB(Y)$	the Markov Blanket of the target variable Y
$\mathcal{C}_{pool}$	the confounder pool
$\mathcal{C}$	the confounder subspace
$\mathcal{G}^\tau$	the causal structure obtained at τ-th round over $\mathcal{Z}_f^{\leq\tau} \cup \{Y\}$
$q \in \{-1, 0, 1\}$	the annotation result
$\mathcal{Q}^{\leq\tau}$	the data matrix of annotated values of variables proposed at the first τ-th round
J	the number of cluster centroids
$p(c)$	the uniform distribution for prior probability
α	the weight of causal structure learning loss
β	the weight of different terms of total loss
$\mathbf{W}$	the weight matrix

Appendix B

The Acronyms in the Thesis

Table B.1: The Acronyms in All the Chapters

Sections	Explanations	Acronyms
Chapter 1&2&3&4&5&6	Cross-Domain Recommendation	CDR
Chapter 2&4	Cross-Domain Sequential Recommendation	CDSR
Chapter 1&5&6	Independent and Identically Distributed	IID
Chapter 1&2&5&6	Out-of-Distribution	OOD
Chapter 1&2&3&4&5	Recommender System	RS
Chapter 1&2&5&6	Large Language Model	LLM
Chapter 3&4&5	Hit Ratio	HR
Chapter 3&4&5	Normalized Discounted Cumulative Gain	NDCG
Chapter 2&3&4&5	Graph Convolutional Network	GCN
Chapter 1&2&3&4&5	Variational Autoencoder	VAE
Chapter 1&2&4	Inverse Propensity Score	IPS
Chapter 3&4&5	Multi-Layer Perceptron	MLP
Chapter 1&2&4&6	Single-Domain Confounder	SDC
Chapter 1&2&4&6	Cross-Domain Confounder	CDC
Chapter 1&2&5&6	Single-Domain Distribution Shift	SDDS
Chapter 1&2&5&6	Cross-Domain Distribution Shift	CDDS
Chapter 1&2&4&5&6	Directed Acyclic Graph	DAG

Bibliography

[1] S. Abdali, A. Parikh, S. Lim, and E. Kiciman. Extracting self-consistent causal insights from users feedback with LLMs and in-context learning. *arXiv preprint arXiv:2312.06820*, 2023.

[2] T. Ban, L. Chen, D. Lyu, X. Wang, and H. Chen. Causal structure learning supervised by large language model. *arXiv preprint arXiv:2311.11689*, 2023.

[3] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE TPAMI*, 35(8):1798–1828, 2013.

[4] G. Blanchard, A. A. Deshmukh, U. Dogan, G. Lee, and C. Scott. Domain generalization by marginal transfer learning. *JMLR*, 22(2):1–55, 2021.

[5] D. Borkan, L. Dixon, J. Sorensen, N. Thain, and L. Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *TheWebConf*, pages 491–500, 2019.

[6] P. Bühlmann. Invariance, causality and robustness. *Statistical Science*, 35(3):404–426, 2020.

[7] J. Cao, S. Li, B. Yu, X. Guo, T. Liu, and B. Wang. Towards universal cross-domain recommendation. In *WSDM*, pages 78–86, 2023.

[8] J. Cao, X. Lin, X. Cong, J. Ya, T. Liu, and B. Wang. DisenCDR: Learning disentangled representations for cross-domain recommendation. In *SIGIR*, pages 267–277, 2022.

[9] J. Cao, J. Sheng, X. Cong, T. Liu, and B. Wang. Cross-domain recommendation

to cold-start users via variational information bottleneck. In *ICDE*, pages 2209–2223, 2022.

[10] H. Chen, C.-C. M. Yeh, F. Wang, and H. Yang. Graph neural transport networks with non-local attentions for recommender systems. In *TheWebConf*, pages 1955–1964, 2022.

[11] R. T. Q. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud. Isolating sources of disentanglement in VAEs. In *NeurIPS*, page 2615–2625, 2018.

[12] T. Chen, S. Zuo, C. Li, M. Zhang, Q. Mei, and M. Bendersky. Unlocking the 'Why' of buying: Introducing a new dataset and benchmark for purchase reason and post-purchase experience. *arXiv preprint arXiv:2402.13417*, 2024.

[13] W. Chen, P. Ren, F. Cai, F. Sun, and M. De Rijke. Multi-interest diversification for end-to-end sequential recommendation. *ACM TOIS*, 40(1):1–30, 2021.

[14] X. Chen, Y. Zhang, I. W. Tsang, Y. Pan, and J. Su. Toward equivalent transformation of user preferences in cross domain recommendation. *ACM TOIS*, 41(1):1–31, 2023.

[15] Y. Choi, J. Choi, T. Ko, H. Byun, and C.-K. Kim. Review-based domain disentanglement without duplicate users or contexts for cross-domain recommendation. In *CIKM*, pages 293–303, 2022.

[16] Q. Cui, T. Wei, Y. Zhang, and Q. Zhang. HeroGRAPH: A heterogeneous graph framework for multi-target cross-domain recommendation. In *RecSys*, 2020.

[17] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019.

[18] P. Ding, Y. Wang, G. Liu, N. Wang, and X. Zhou. Few-shot causal representation learning for out-of-distribution generalization on heterogeneous graphs. *IEEE TKDE*, 37(4):1804–1818, 2025.

[19] J. Du, Z. Ye, B. Guo, Z. Yu, and L. Yao. Identifiability of cross-domain recommendation via causal subspace disentanglement. In *SIGIR*, page 2091–2101, 2024.

[20] X. Du, Z. Wu, F. Feng, X. He, and J. Tang. Invariant representation learning for multimedia recommendation. In *MM*, pages 619–628, 2022.

[21] Y. Du, Z. Wang, Z. Sun, Y. Ma, H. Liu, and J. Zhang. Disentangled multi-interest representation learning for sequential recommendation. In *KDD*, pages 677–688, 2024.

[22] Z. Du, X. Wang, H. Yang, J. Zhou, and J. Tang. Sequential scenario-specific meta learner for online recommendation. In *KDD*, pages 2895–2904, 2019.

[23] W. Fu, Z. Peng, S. Wang, Y. Xu, and J. Li. Deeply fusing reviews and contents for cold start users in cross-domain recommendation systems. In *AAAI*, pages 94–101, 2019.

[24] Y. Fu, Y. Fu, and Y.-G. Jiang. Meta-FDMixup: Cross-domain few-shot learning guided by labeled target data. In *MM*, pages 5326–5334, 2021.

[25] Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015.

[26] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. March, and V. Lempitsky. Domain-adversarial training of neural networks. *JMLR*, 17(59):1–35, 2016.

[27] C. Gao, Y. Zheng, W. Wang, F. Feng, X. He, and Y. Li. Causal inference in

recommender systems: A survey and future directions. *ACM TOIS*, 42(4):1–32, 2022.

[28] S. Gao, H. Luo, D. Chen, S. Li, P. Gallinari, and J. Guo. Cross-domain recommendation via cluster-level latent factor model. In *ECML-PKDD*, pages 161–176, 2013.

[29] A. Gibbs. Writing as method: Attunement, resonance and rhythm. *Affective methodologies: Developing cultural research strategies for the study of affect*, pages 222–236, 2015.

[30] A. Gonzalez-Garcia, J. van de Weijer, and Y. Bengio. Image-to-image translation for cross-domain disentanglement. In *NeurIPS*, page 1294–1305, 2018.

[31] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.

[32] S. C. Gowda, S. Joshi, H. Zhang, and M. Ghassemi. Pulling up by the causal bootstraps: Causal data augmentation for pre-training debiasing. In *CIKM*, pages 606–616, 2021.

[33] L. Guo, H. Liu, L. Zhu, W. Guan, and Z. Cheng. DA-DAN: A dual adversarial domain adaption network for unsupervised non-overlapping cross-domain recommendation. *ACM TOIS*, 42(2):1–27, 2023.

[34] X. Guo, S. Li, N. Guo, J. Cao, X. Liu, Q. Ma, R. Gan, and Y. Zhao. Disentangled representations learning for multi-target cross-domain recommendation. *ACM TOIS*, 41(4):1–27, 2023.

[35] F. M. Harper and J. A. Konstan. The movielens datasets: History and context. *ACM TIIS*, 5(4):1–19, 2015.

[36] U. Hasan, E. Hossain, and M. O. Gani. A survey on causal discovery methods for I.I.D. and time series data. *TMLR*, 2023.

[37] M. D. Hauser and J. McDermott. The evolution of the music faculty: A comparative perspective. *Nature neuroscience*, pages 663–668, 2003.

[38] M. He, J. Zhang, P. Yang, and K. Yao. Robust transfer learning for cross-domain collaborative filtering using multiple rating patterns approximation. In *WSDM*, pages 225–233, 2018.

[39] T. He, K. Li, S. Chen, H. Wang, Q. Liu, X. Wang, and D. Wang. DMBIN: A dual multi-behavior interest network for click-through rate prediction via contrastive learning. In *SIGIR*, pages 1366–1375, 2023.

[40] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang. LightGCN: Simplifying and powering graph convolution network for recommendation. In *SIGIR*, pages 639–648, 2020.

[41] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *WWW*, pages 173–182, 2017.

[42] X. He, Y. Zhang, F. Feng, C. Song, L. Yi, G. Ling, and Y. Zhang. Addressing confounding feature issue for causal recommendation. *ACM TOIS*, 41(3):1–23, 2023.

[43] Y. He, Z. Wang, P. Cui, H. Zou, Y. Zhang, Q. Cui, and Y. Jiang. CausPref: Causal preference learning for out-of-distribution recommendation. In *TheWebConf*, pages 410–421, 2022.

[44] G. Hu, Y. Zhang, and Q. Yang. CoNet: Collaborative cross networks for cross-domain recommendation. In *CIKM*, pages 667–676, 2018.

[45] G. Hu, Y. Zhang, and Q. Yang. Transfer meets hybrid: A synthetic approach

for cross-domain collaborative filtering with text. In *TheWebConf*, pages 2822–2829, 2019.

[46] Z. Huang, S. Zhang, D. Cheng, J. Li, L. Liu, and G. Zhang. Multi-cause deconfounding for recommender systems with latent confounders. *Knowledge-Based Systems*, 329:114345, 2024.

[47] X. Jin, N. Li, W. Kong, J. Tang, and B. Yang. Unbiased semantic representation learning based on causal disentanglement for domain generalization. *ACM TOMM*, 20(8):1551–6857, 2024.

[48] Z. Jin, J. Liu, Z. LYU, S. Poff, M. Sachan, R. Mihalcea, M. T. Diab, and B. Schölkopf. Can large language models infer causation from correlation? In *ICLR*, 2024.

[49] H. Kanagawa, H. Kobayashi, N. Shimizu, Y. Tagami, and T. Suzuki. Cross-domain recommendation via deep domain adaptation. In *ECIR*, pages 20–29, 2019.

[50] E. Khatibi, M. Abbasian, Z. Yang, I. Azimi, and A. M. Rahmani. ALCM: Autonomous LLM-augmented causal discovery framework. *arXiv preprint arXiv:2405.01744*, 2024.

[51] E. Kiciman, R. Ness, A. Sharma, and C. Tan. Causal reasoning and large language models: Opening a new frontier for causality. *TMLR*, 2023.

[52] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[53] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

[54] W. Krichene and S. Rendle. On sampled metrics for item recommendation. *CACM*, pages 75–83, 2022.

[55] H. D. Le, X. Xia, and Z. Chen. Multi-agent causal discovery using large language models. *arXiv preprint arXiv:2407.15073*, 2024.

[56] C. Li, Y. Xie, C. Yu, B. Hu, Z. Li, G. Shu, X. Qie, and D. Niu. One for all, all for one: Learning and transferring user embeddings for cross-domain recommendation. In *WSDM*, pages 366–374, 2023.

[57] F. Li, H. Liu, J. He, and X. Du. CausalCDR: Causal embedding learning for cross-domain recommendation. In *SDM*, pages 553–561, 2024.

[58] H. Li, W. Ma, P. Sun, J. Li, C. Yin, Y. He, G. Xu, M. Zhang, and S. Ma. Aiming at the target: Filter collaborative information for cross-domain recommendation. In *SIGIR*, pages 2081–2090, 2024.

[59] P. Li and A. Tuzhilin. DDTCDR: Deep dual transfer cross domain recommendation. In *WSDM*, pages 331–339, 2020.

[60] Q. Li, Z. Han, and X.-M. Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*, pages 3538–3545, 2018.

[61] Q. Li, H. Ma, W. Jin, Y. Ji, and Z. Li. Multi-interest network with simple diffusion for multi-behavior sequential recommendation. In *SDM*, pages 734–742, 2024.

[62] Q. Li, X. Wang, Z. Wang, and G. Xu. Be causal: De-biasing social network confounding in recommendation. *ACM TKDD*, 17(1):1–23, 2023.

[63] S. Li, L. Yao, S. Mu, W. X. Zhao, Y. Li, T. Guo, B. Ding, and J.-R. Wen. Debiasing learning based cross-domain recommendation. In *KDD*, pages 3190–3199, 2021.

[64] Z. Li, D. Amagata, Y. Zhang, T. Hara, S. Haruta, K. Yonekawa, and M. Kurokawa. Debiasing graph transfer learning via item semantic clustering for cross-domain recommendations. In *Big Data*, pages 762–769, 2022.

[65] Z. Li and N. Yang. Cross-domain causal preference learning for out-of-distribution recommendation. In *DASFAA*, pages 245–260, 2024.

[66] Y. Liang, E. Yang, G. Guo, W. Cai, L. Jiang, and X. Wang. Deconfounding user preference in recommendation systems through implicit and explicit feedback. *ACM TKDD*, 18(8):1–18, 2024.

[67] Z. Lin, H. Ding, N. T. Hoang, B. Kveton, A. Deoras, and H. Wang. Pre-trained recommender systems: A causal debiasing perspective. In *WSDM*, pages 424–433, 2024.

[68] C. Liu, Y. Chen, T. Liu, M. Gong, J. Cheng, B. Han, and K. Zhang. Discovery of the hidden world with large language models. In *NeurIPS*, volume 37, pages 102307–102365, 2024.

[69] D. Liu, P. Cheng, H. Zhu, Z. Dong, X. He, W. Pan, and Z. Ming. Mitigating confounding bias in recommendation via information bottleneck. In *RecSys*, pages 351–360, 2021.

[70] D. Liu, P. Cheng, H. Zhu, Z. Dong, X. He, W. Pan, and Z. Ming. Debiased representation learning in recommendation via information bottleneck. *ACM TORS*, 1(1):1–27, 2023.

[71] D. Liu, Y. Yang, M. Zhang, W. Wu, X. Xie, and G. Sun. Knowledge enhanced multi-interest network for the generation of recommendation candidates. In *CIKM*, pages 3322–3331, 2022.

[72] J. Liu, L. Sun, W. Nie, P. Jing, and Y. Su. Graph disentangled contrastive learning with personalized transfer for cross-domain recommendation. In *AAAI*, pages 8769–8777, 2024.

[73] J. Liu, P. Zhao, F. Zhuang, Y. Liu, V. S. Sheng, J. Xu, X. Zhou, and H. Xiong. Exploiting aesthetic preference in deep cross networks for cross-domain recommendation. In *TheWebConf*, pages 2768–2774, 2020.

[74] M. Liu, J. Li, G. Li, and P. Pan. Cross domain recommendation via bi-directional transfer graph collaborative filtering networks. In *CIKM*, pages 885–894, 2020.

[75] S. Liu, Z. Meng, C. Macdonald, and I. Ounis. Graph neural pre-training for recommendation with side information. *ACM TOIS*, 41(3):1–28, 2023.

[76] W. Liu, J. Su, C. Chen, and X. Zheng. Leveraging distribution alignment via stein path for cross-domain cold-start recommendation. *NeurIPS*, 34:19223–19234, 2021.

[77] W. Liu, X. Zheng, J. Su, M. Hu, Y. Tan, and C. Chen. Exploiting variational domain-invariant user embedding for partially overlapped cross domain recommendation. In *SIGIR*, pages 312–321, 2022.

[78] X. Liu, T. Yu, K. Xie, J. Wu, and S. Li. Interact with the explanations: Causal debiased explainable recommendation system. In *WSDM*, pages 472–481, 2024.

[79] Y. Liu, Y. Chen, W. Dai, C. Li, J. Zou, and H. Xiong. Causal intervention for generalizable face anti-spoofing. In *ICME*, pages 01–06, 2022.

[80] Y. Liu, X. Zhang, M. Zou, and Z. Feng. Attribute simulation for item embedding enhancement in multi-interest recommendation. In *WSDM*, pages 482–491, 2024.

[81] Z. Liu, L. Zheng, J. Zhang, J. Han, and S. Y. Philip. JSCN: Joint spectral convolutional network for cross domain recommendation. In *Big Data*, pages 850–859, 2019.

[82] B. Loni, Y. Shi, M. Larson, and A. Hanjalic. Cross-domain collaborative filtering with factorization machines. In *ECIR*, pages 656–661, 2014.

[83] H. Luo, F. Zhuang, R. Xie, H. Zhu, D. Wang, Z. An, and Y. Xu. A survey on causal inference for recommendation. *The Innovation*, 5(2):100590, 2024.

[84] L. Luo, Y. Li, B. Gao, S. Tang, S. Wang, J. Li, T. Zhu, J. Liu, Z. Li, and S. Pan. MAMDR: A model agnostic learning framework for multi-domain recommendation. In *ICDE*, pages 3079–3092, 2023.

[85] J. Ma, C. Zhou, P. Cui, H. Yang, and W. Zhu. Learning disentangled representations for recommendation. In *NeurIPS*, pages 5711–5722, 2019.

[86] J. Ma, C. Zhou, H. Yang, P. Cui, X. Wang, and W. Zhu. Disentangled self-supervision in sequential recommenders. In *KDD*, pages 483–491, 2020.

[87] T. Man, H. Shen, X. Jin, and X. Cheng. Cross-domain recommendation: An embedding and mapping approach. In *IJCAI*, pages 2464–2470, 2017.

[88] Q. Mao, Q. Liu, Z. Li, L. Wu, B. Lv, and Z. Zhang. Cross-reconstructed augmentation for dual-target cross-domain recommendation. In *SIGIR*, pages 2352–2356, 2024.

[89] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley. Least squares generative adversarial networks. In *ICCV*, pages 2794–2802, 2017.

[90] C. Meng, Z. Zhao, W. Guo, Y. Zhang, H. Wu, C. Gao, D. Li, X. Li, and R. Tang. Coarse-to-fine knowledge-enhanced multi-interest learning framework for multi-behavior recommendation. *ACM TOIS*, 42(1):1–27, 2023.

[91] K. Menglin, J. Wang, Y. Pan, H. Zhang, and M. Hou. C^2DR: Robust cross-domain recommendation based on causal disentanglement. In *WSDM*, pages 341–349, 2024.

[92] M. Oberst, N. Thams, J. Peters, and D. Sontag. Regularizing towards causal invariance: Linear models with proxies. In *ICML*, pages 8260–8270, 2021.

[93] C. Ouyang, C. Chen, S. Li, Z. Li, C. Qin, W. Bai, and D. Rueckert. Causality-inspired single-source domain generalization for medical image segmentation. *IEEE TMI*, 42(4):1095–1106, 2022.

[94] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Elsevier, 2014.

[95] J. Pearl and D. Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic books, 2018.

[96] S. Qiao, Y. Ou, N. Zhang, X. Chen, Y. Yao, S. Deng, C. Tan, F. Huang, and H. Chen. Reasoning with language model prompting: A survey. In *ACL*, pages 5368–5393, 2023.

[97] X. Ren, L. Xia, J. Zhao, D. Yin, and C. Huang. Disentangled contrastive collaborative filtering. In *SIGIR*, pages 1137–1146, 2023.

[98] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*, page 452–461, 2009.

[99] A. K. Sahu and P. Dwivedi. Knowledge transfer by domain-independent user latent factor for cross-domain recommender systems. *FGCS*, pages 320–333, 2020.

[100] M. Sato, S. Takemori, J. Singh, and T. Ohkuma. Unbiased learning for the causal effect of recommendation. In *RecSys*, pages 378–387, 2020.

[101] B. Schölkopf, D. W. Hogg, D. Wang, D. Foreman-Mackey, D. Janzing, C.-J. Simon-Gabriel, and J. Peters. Modeling confounding by half-sibling regression. *PNAS*, 113(27):7391–7398, 2016.

[102] P. Sheth, R. Moraffah, K. S. Candan, A. Raglin, and H. Liu. Domain generalization–a causal perspective. *arXiv preprint arXiv:2209.15177*, 2022.

[103] S. Shimizu, P. O. Hoyer, A. Hyvärinen, A. Kerminen, and M. Jordan. A linear non-gaussian acyclic model for causal discovery. *JMLR*, 7(10), 2006.

[104] K. Shirahama, K. Iwamoto, and K. Uehera. Video data mining: Rhythms in a movie. In *ICME*, pages 1463–1466, 2004.

[105] Z. Song, W. Zhang, L. Deng, J. Zhang, Z. Wu, K. Bian, and B. Cui. Mitigating negative transfer in cross-domain recommendation via knowledge transferability enhancement. In *KDD*, pages 2745–2754, 2024.

[106] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2001.

[107] H. Su, J. Li, Z. Du, L. Zhu, K. Lu, and H. T. Shen. Cross-domain recommendation via dual adversarial adaptation. *ACM TOIS*, 42(3):1–26, 2023.

[108] H. Su, J. Li, Z. Du, L. Zhu, K. Lu, and H. T. Shen. Cross-domain recommendation via dual adversarial adaptation. *ACM TOIS*, 42(3):1–26, 2024.

[109] H. Su, Y. Zhang, X. Yang, H. Hua, S. Wang, and J. Li. Cross-domain recommendation via adversarial adaptation. In *CIKM*, pages 1808–1817, 2022.

[110] J. Sun, Y. Zhang, C. Ma, M. Coates, H. Guo, R. Tang, and X. He. Multi-graph convolution collaborative filtering. In *ICDM*, pages 1306–1311, 2019.

[111] Z. Sun, K. Feng, J. Yang, X. Qu, H. Fang, Y.-S. Ong, and W. Liu. Adaptive in-context learning with large language models for bundle generation. In *SIGIR*, pages 966–976, 2024.

[112] M. Takayama, T. Okuda, T. Pham, T. Ikenoue, S. Fukuma, S. Shimizu, and A. Sannai. Integrating large language models in causal discovery: A statistical causal approach. *arXiv preprint arXiv:2402.01454*, 2024.

[113] Y. Tian, J. Chang, Y. Niu, Y. Song, and C. Li. When multi-level meets multi-interest: A multi-grained neural model for sequential recommendation. In *SIGIR*, pages 1632–1641, 2022.

[114] N.-T. Tran and H. W. Lauw. Multi-representation variational autoencoder via iterative latent attention and implicit differentiation. In *CIKM*, pages 2462–2471, 2023.

[115] A. Vashishtha, A. G. Reddy, A. Kumar, S. Bachu, V. N. Balasubramanian, and A. Sharma. Causal inference using LLM-guided discovery. In *AAAI 2024 Workshop on "Are Large Language Models Simply Causal Parrots?"*, 2023.

[116] B. Wang, J. Chen, C. Li, S. Zhou, Q. Shi, Y. Gao, Y. Feng, C. Chen, and C. Wang. Distributionally robust graph-based recommendation system. In *TheWebConf*, pages 3777–3788, 2024.

[117] C. Wang, M. Niepert, and H. Li. RecSys-DAN: Discriminative adversarial networks for cross-domain recommender systems. *IEEE TNNLS*, 31(8):2731–2740, 2019.

[118] J. Wang, Y. Jiang, Y. Long, X. Sun, M. Pagnucco, and Y. Song. Deconfounding causal inference for zero-shot action recognition. *IEEE TMM*, 26:3976–3986, 2023.

[119] J. Wang, C. Lan, C. Liu, Y. Ouyang, T. Qin, W. Lu, Y. Chen, W. Zeng, and S. Y. Philip. Generalizing to unseen domains: A survey on domain generalization. *IEEE TKDE*, 35(8):8052–8072, 2022.

[120] J. Wang and J. Lv. Tag-informed collaborative topic modeling for cross domain recommendations. *Knowledge-Based Systems*, 203:106119, 2020.

[121] K. Wang, Y. Zhu, H. Liu, T. Zang, C. Wang, and K. Liu. Inter-and intra-domain relation-aware heterogeneous graph convolutional networks for cross-domain recommendation. In *DASFAA*, pages 53–68, 2022.

[122] W. Wang, F. Feng, X. He, X. Wang, and T.-S. Chua. Deconfounded recommendation for alleviating bias amplification. In *KDD*, pages 1717–1725, 2021.

[123] W. Wang, X. Lin, F. Feng, X. He, M. Lin, and T.-S. Chua. Causal representation learning for out-of-distribution recommendation. In *TheWebConf*, pages 3562–3571, 2022.

[124] W. Wang, X. Lin, L. Wang, F. Feng, Y. Wei, and T.-S. Chua. Equivariant learning for out-of-distribution cold-start recommendation. In *MM*, pages 903–914, 2023.

[125] X. Wang, H. Chen, S. Tang, Z. Wu, and W. Zhu. Disentangled representation learning. *IEEE TPAMI*, 46(12):9677–9696, 2024.

[126] X. Wang, H. Chen, Y. Zhou, J. Ma, and W. Zhu. Disentangled representation learning for recommendation. *IEEE TPAMI*, 45(1):408–424, 2022.

[127] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua. KGAT: Knowledge graph attention network for recommendation. In *KDD*, pages 950–958, 2019.

[128] X. Wang, X. He, M. Wang, F. Feng, and T.-S. Chua. Neural graph collaborative filtering. In *SIGIR*, pages 165–174, 2019.

[129] X. Wang, H. Jin, A. Zhang, X. He, T. Xu, and T.-S. Chua. Disentangled graph collaborative filtering. In *SIGIR*, pages 1001–1010, 2020.

[130] X. Wang, Q. Li, D. Yu, P. Cui, Z. Wang, and G. Xu. Causal disentanglement for semantics-aware intent learning in recommendation. *IEEE TKDE*, 35(10):9836–9849, 2022.

[131] Y. Wang, H. Guo, B. Chen, W. Liu, Z. Liu, Q. Zhang, Z. He, H. Zheng, W. Yao, M. Zhang, et al. CausalInt: Causal inspired intervention for multi-scenario recommendation. In *KDD*, pages 4090–4099, 2022.

[132] Y. Wang, D. Liang, L. Charlin, and D. M. Blei. Causal inference for recommender systems. In *RecSys*, pages 426–431, 2020.

[133] Y. Wang, W. Wang, Y. Liang, Y. Cai, and B. Hooi. Mixup for node and graph classification. In *TheWebConf*, pages 3663–3674, 2021.

[134] Y. Wang, X. Wang, X. Huang, Y. Yu, H. Li, M. Zhang, Z. Guo, and W. Wu. Intent-aware recommendation via disentangled graph contrastive learning. In *IJCAI*, pages 2343–2351, 2023.

[135] Z. Wang, Y. He, J. Liu, W. Zou, P. S. Yu, and P. Cui. Invariant preference learning for general debiasing in recommendation. In *KDD*, pages 1969–1978, 2022.

[136] Z. Wang, S. Shen, Z. Wang, B. Chen, X. Chen, and J.-R. Wen. Unbiased sequential recommendation with latent confounders. In *TheWebConf*, pages 2195–2204, 2022.

[137] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, pages 24824–24837, 2022.

[138] A. Wu, K. Kuang, R. Xiong, and F. Wu. Instrumental variables in causal inference and machine learning: A survey. *ACM Computing Surveys*, 57(11):1–36, 2025.

[139] P. Wu, H. Li, Y. Deng, W. Hu, Q. Dai, Z. Dong, J. Sun, R. Zhang, and X.-H. Zhou. On the opportunity of causal learning in recommendation systems: Foundation, estimation, prediction and challenges. In *IJCAI*, pages 5646–5653, 2022.

[140] S. Xiao, D. Zhu, C. Tang, and Z. Huang. CATCL: Joint cross-attention transfer and contrastive learning for cross-domain recommendation. In *DASFAA*, pages 446–461, 2023.

[141] Z. Xiao, L. Yang, W. Jiang, Y. Wei, Y. Hu, and H. Wang. Deep multi-interest network for click-through rate prediction. In *CIKM*, pages 2265–2268, 2020.

[142] R. Xie, Q. Liu, L. Wang, S. Liu, B. Zhang, and L. Lin. Contrastive cross-domain recommendation in matching. In *KDD*, pages 4226–4236, 2022.

[143] Y. Xie, J. Gao, P. Zhou, Q. Ye, Y. Hua, J. B. Kim, F. Wu, and S. Kim. Rethinking multi-interest learning for candidate matching in recommender systems. In *RecSys*, pages 283–293, 2023.

[144] Y. Xie, C. Yu, X. Jin, L. Cheng, B. Hu, and Z. Li. Heterogeneous graph contrastive learning for cold start cross-domain recommendation. *Knowledge-Based Systems*, 299:112054, 2024.

[145] S. Xu, J. Ji, Y. Li, Y. Ge, J. Tan, and Y. Zhang. Causal inference for recommendation: Foundations, methods and applications. *ACM TIST*, 16(3):1–51, 2025.

[146] S. Xu, J. Tan, S. Heinecke, V. J. Li, and Y. Zhang. Deconfounded causal collaborative filtering. *ACM TORS*, 1(4):1–25, 2023.

[147] W. Xu, Q. Wu, R. Wang, M. Ha, Q. Ma, L. Chen, B. Han, and J. Yan. Rethinking cross-domain sequential recommendation under open-world assumptions. In *TheWebConf*, pages 3173–3184, 2024.

[148] X. Yang, Y. Wang, J. Chen, W. Fan, X. Zhao, E. Zhu, X. Liu, and D. Lian. Dual test-time training for out-of-distribution recommender system. *IEEE TKDE*, 37(6):3312–3326, 2025.

[149] D. Yu, Q. Li, X. Wang, and G. Xu. Deconfounded recommendation via causal intervention. *Neurocomputing*, 529:128–139, 2023.

[150] F. Yuan, L. Yao, and B. Benatallah. DARec: Deep domain adaptation for cross-domain recommendation via transferring rating patterns. In *IJCAI*, pages 4227–4233, 2019.

[151] J. Yuan, X. Ma, R. Xiong, M. Gong, X. Liu, F. Wu, L. Lin, and K. Kuang. Instrumental variable-driven domain generalization with unobserved confounders. *ACM TKDD*, 17(8):1–21, 2023.

[152] T. Zang, Y. Zhu, H. Liu, R. Zhang, and J. Yu. A survey on cross-domain recommendation: Taxonomies, methods, and future directions. *ACM TOIS*, 41(2):1–39, 2022.

[153] M. Zečević, M. Willig, D. S. Dhami, and K. Kersting. Causal parrots: Large language models may talk causality but are not causal. *arXiv preprint arXiv:2308.13067*, 2023.

[154] R. Zhan, C. Pei, Q. Su, J. Wen, X. Wang, G. Mu, D. Zheng, P. Jiang, and K. Gai. Deconfounding duration bias in watch-time prediction for video recommendation. In *KDD*, pages 4472–4481, 2022.

[155] A. Zhang, W. Ma, J. Zheng, X. Wang, and T.-S. Chua. Robust collaborative filtering to popularity distribution shift. *ACM TOIS*, 42(3):1–25, 2024.

[156] A. Zhang, J. Zheng, X. Wang, Y. Yuan, and T.-S. Chua. Invariant collaborative filtering to popularity distribution shift. In *TheWebConf*, pages 1240–1251, 2023.

[157] H. Zhang, M. Cissé, Y. N. Dauphin, and D. Lopez-Paz. *mixup*: Beyond empirical risk minimization. In *ICLR*, 2018.

[158] Q. Zhang, P. Hao, J. Lu, and G. Zhang. Cross-domain recommendation with semantic correlation in tagging systems. In *IJCNN*, pages 1–8, 2019.

[159] Q. Zhang, X. Zhang, Y. Liu, H. Wang, M. Gao, J. Zhang, and R. Guo. Debiasing recommendation by learning identifiable latent confounders. In *KDD*, page 3353–3363, 2023.

[160] R. Zhang, T. Zang, Y. Zhu, C. Wang, K. Wang, and J. Yu. Disentangled contrastive learning for cross-domain recommendation. In *DASFAA*, pages 163–178, 2023.

[161] S. Zhang, X. Feng, W. Fan, W. Fang, F. Feng, W. Ji, S. Li, L. Wang, S. Zhao, Z. Zhao, et al. Video-audio domain generalization via confounder disentanglement. In *AAAI*, pages 15322–15330, 2023.

[162] S. Zhang, Q. Miao, P. Nie, M. Li, Z. Chen, F. Feng, K. Kuang, and F. Wu. Transferring causal mechanism over meta-representations for target-unknown cross-domain recommendation. *ACM TOIS*, 42(4):1–27, 2024.

[163] S. Zhang, L. Yang, D. Yao, Y. Lu, F. Feng, Z. Zhao, T.-S. Chua, and F. Wu. Re4: Learning to re-contrast, re-attend, re-construct for multi-interest recommendation. In *TheWebConf*, pages 2216–2226, 2022.

[164] W. Zhang, Y. Bei, L. Yang, H. P. Zou, P. Zhou, A. Liu, Y. Li, H. Chen, J. Wang, Y. Wang, et al. Cold-start recommendation towards the era of large language models (LLMs): A comprehensive survey and roadmap. *arXiv preprint arXiv:2501.01945*, 2025.

[165] X. Zhang, J. Li, H. Su, L. Zhu, and H. T. Shen. Multi-level attention-based domain disentanglement for BCDR. *ACM TOIS*, 41(4):1–24, 2023.

[166] Y. Zhang, F. Feng, X. He, T. Wei, C. Song, G. Ling, and Y. Zhang. Causal intervention for leveraging popularity bias in recommendation. In *SIGIR*, pages 11–20, 2021.

[167] Y. Zhang, Y. Shen, D. Wang, J. Gu, and G. Zhang. Connecting unseen domains: Cross-domain invariant learning in recommendation. In *SIGIR*, pages 1894–1898, 2023.

[168] Y. Zhang, Z. Zhu, Y. He, and J. Caverlee. Content-collaborative disentanglement representation learning for enhanced recommendation. In *RecSys*, pages 43–52, 2020.

[169] Y.-F. Zhang, Z. Zhang, D. Li, Z. Jia, L. Wang, and T. Tan. Learning domain invariant representations for generalizable person re-identification. *IEEE TIP*, 32:509–523, 2022.

[170] Z. Zhang, H. Gao, H. Yang, and X. Chen. Hierarchical invariant learning for domain generalization recommendation. In *KDD*, pages 3470–3479, 2023.

[171] Z. Zhang, X. Wang, H. Chen, H. Li, and W. Zhu. Disentangled dynamic graph attention network for out-of-distribution sequential recommendation. *ACM TOIS*, 43(1):1–42, 2024.

[172] C. Zhao, C. Li, and C. Fu. Cross-domain recommendation via preference propagation graphnet. In *CIKM*, pages 2165–2168, 2019.

[173] C. Zhao, X. Li, M. He, H. Zhao, and J. Fan. Sequential recommendation via an adaptive cross-domain knowledge decomposition. In *CIKM*, pages 3453–3463, 2023.

[174] C. Zhao, E. Yang, Y. Liang, P. Lan, Y. Liu, J. Zhao, G. Guo, and X. Wang. Graph representation learning via causal diffusion for out-of-distribution recommendation. In *TheWebConf*, pages 334–346, 2025.

[175] C. Zhao, E. Yang, Y. Liang, J. Zhao, G. Guo, and X. Wang. Distributionally robust graph out-of-distribution recommendation via diffusion model. In *TheWebConf*, pages 2018–2031, 2025.

[176] C. Zhao, H. Zhao, M. He, J. Zhang, and J. Fan. Cross-domain recommendation via user interest alignment. In *TheWebConf*, pages 887–896, 2023.

[177] C. Zhao, H. Zhao, X. Li, M. He, J. Wang, and J. Fan. Cross-domain recommendation via progressive structural alignment. *IEEE TKDE*, 36(6):2401–2415, 2023.

[178] W. Zhao, D. Tang, X. Chen, D. Lv, D. Ou, B. Li, P. Jiang, and K. Gai. Disentangled causal embedding with contrastive learning for recommender system. In *TheWebConf*, pages 406–410, 2023.

[179] X. Zhao, N. Yang, and P. S. Yu. Multi-sparse-domain collaborative recommendation via enhanced comprehensive aspect preference learning. In *WSDM*, pages 1452–1460, 2022.

[180] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing. DAGs with NO TEARS: Continuous optimization for structure learning. *NeurIPS*, page 9492–9503, 2018.

[181] Y. Zheng, C. Gao, X. Li, X. He, Y. Li, and D. Jin. Disentangling user interest and conformity for recommendation with causal embedding. In *TheWebConf*, pages 2980–2991, 2021.

[182] K. Zhou, Z. Liu, Y. Qiao, T. Xiang, and C. C. Loy. Domain generalization: A survey. *IEEE TPAMI*, 45(4):4396–4415, 2022.

[183] F. Zhu, C. Chen, Y. Wang, G. Liu, and X. Zheng. DTCDR: A framework for dual-target cross-domain recommendation. In *CIKM*, pages 1533–1542, 2019.

[184] F. Zhu, Y. Wang, C. Chen, G. Liu, M. Orgun, and J. Wu. A deep framework for cross-domain and cross-system recommendations. In *IJCAI*, pages 3711–3717, 2018.

[185] F. Zhu, Y. Wang, C. Chen, G. Liu, and X. Zheng. A graphical and attentional framework for dual-target cross-domain recommendation. In *IJCAI*, pages 3001–3008, 2020.

[186] F. Zhu, Y. Wang, C. Chen, J. Zhou, L. Li, and G. Liu. Cross-domain recommendation: Challenges, progress, and prospects. In *IJCAI*, pages 4721–4728, 2021.

[187] F. Zhu, Y. Wang, J. Zhou, C. Chen, L. Li, and G. Liu. A unified framework for cross-domain and cross-system recommendations. *IEEE TKDE*, 35(2):1171–1184, 2021.

[188] J. Zhu, Y. Wang, F. Zhu, and Z. Sun. Domain disentanglement with interpolative data augmentation for dual-target cross-domain recommendation. In *RecSys*, pages 515–527, 2023.

[189] J. Zhu, Y. Wang, F. Zhu, and Z. Sun. Causal deconfounding via confounder disentanglement for dual-target cross-domain recommendation. *ACM TOIS*, 43(5):1–33, 2025.

[190] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2223–2232, 2017.

[191] X. Zhu, Y. Zhang, X. Yang, D. Wang, and F. Feng. Mitigating hidden confounding effects for causal recommendation. *IEEE TKDE*, 36(9):4794–4805, 2024.

[192] Y. Zhu, J. Ma, and J. Li. Causal inference in recommender systems: A survey of strategies for bias mitigation, explanation, and generalization. *arXiv preprint arXiv:2301.00910*, 2023.

[193] Y. Zhu, X. Ren, J. Yi, and Z. Chen. Deep deconfounded content-based tag recommendation for UGC with causal intervention. *arXiv preprint arXiv:2205.14380*, 2022.

[194] Y. Zhu, Z. Tang, Y. Liu, F. Zhuang, R. Xie, X. Zhang, L. Lin, and Q. He. Personalized transfer of user preferences for cross-domain recommendation. In *WSDM*, pages 1507–1515, 2022.

[195] Y. Zhu, R. Xie, F. Zhuang, K. Ge, Y. Sun, X. Zhang, L. Lin, and J. Cao. Learning to warm up cold item embeddings for cold-start recommendation with meta scaling and shifting networks. In *SIGIR*, pages 1167–1176, 2021.