

Mitigating Matthew Effect: Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning for Conversational Recommendation

Yongsen Zheng^{1,2}, Ruilin Xu³, Guohua Wang^{4†}, Liang Lin^{3,5}, Kwok-Yan Lam^{1,2†}

¹Nanyang Technological University, Singapore ²Digital Trust Centre Singapore

³Sun Yat-sen University ⁴South China Agricultural University ⁵Peng Cheng Laboratory

{yongsen.zheng, kwokyan.lam}@ntu.edu.sg, xurlin5@mail2.sysu.edu.cn

wangguohua@scau.edu.cn, linliang@ieee.org

Abstract

The Matthew effect is a big challenge in Recommender Systems (RSs), where popular items tend to receive increasing attention, while less popular ones are often overlooked, perpetuating existing disparities. Although many existing methods attempt to mitigate Matthew effect in the static or quasi-static recommendation scenarios, such issue will be more pronounced as users engage with the system over time. To this end, we propose a novel framework, Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning for Conversational Recommendation (HiCore), aiming to address Matthew effect in the Conversational Recommender System (CRS) involving the dynamic user-system feedback loop. It devotes to learn multi-level user interests by building a set of hypergraphs (*i.e.*, item-, entity-, word-oriented multiple-channel hypergraphs) to alleviate the Matthew effect. Extensive experiments on four CRS-based datasets showcase that HiCore attains a new state-of-the-art performance, underscoring its superiority in mitigating the Matthew effect effectively. Our code is available at <https://github.com/zysensmile/HiCore>.

1 Introduction

Engaging users in ongoing conversations for personalized recommendations, Conversational Recommendation Systems (CRSs) (Qin et al., 2023; Mishra et al., 2023) have become a prevalent strategy utilized in diverse fields (Liu et al., 2023; Epure and Hennequin, 2023). However, CRSs often face a big challenge known as Matthew effect (Liu and Huang, 2021), captured by the adage "the privileged gain more privilege, while the underprivileged fall further behind." This observation underscores that well-received items or categories in past records garner heightened visibility in future suggestions, whereas less preferred ones frequently face neglect or marginalization.

Lately, a multitude of studies have focused on investigating the Matthew effect in relatively unchanging offline recommendation scenarios (Liu and Huang, 2021; Anderson et al., 2020), identifying two root causes for its occurrence. One cause (Liang et al., 2021; Zheng et al., 2021a; Hansen et al., 2021; Anderson et al., 2020) is the heightened vulnerability of individuals with narrower and uniform preferences or interests to succumb to the pervasive influence of the Matthew effect. This susceptibility often stems from a tendency towards familiarity and comfort, leading to a reinforcement of existing patterns and a limited exploration of diverse alternatives. Another cause (Zheng et al., 2021b) arises from the pervasive favoritism towards mainstream items, resulting in a perpetual reinforcement of their prominence, while lesser-known alternatives linger in the shadows. This bias towards popular choices not only perpetuates existing trends but also limits the discoverability of niche or underappreciated options. Thus, the amplification of visibility for widely favored items can overshadow the potential value and diversity offered by less popular but equally deserving alternatives.

Despite their effectiveness, most existing methods still suffer from two major limitations. 1) *Interactive Strategy*. While many methods have offered valuable insights into the Matthew effect, they often overlook the adverse effects originating from the dynamic user-system feedback loop (Zhang et al., 2021), as they primarily focus on mitigating the Matthew effect in relatively stable offline recommendation settings. In fact, the Matthew effect can intensify as users interact more actively with the system over time, potentially exacerbating concerns such as echo chambers (Ge et al., 2020) and filter bubbles (Steck, 2018). Hence, it is important to address the Matthew effect in the CRS. 2) *Interest Exploration*. Considering that the root cause of the Matthew effect lies in

[†]Corresponding author.

the confinement of user interests (Zheng et al., 2021a; Liang et al., 2021; Hansen et al., 2021; Anderson et al., 2020), most existing methods focus on leveraging hypergraphs to unveil complex high-order user relationship patterns for exploring user interests. However, these hypergraphs often remain single-channel, constraining their capacity to capture diverse user relation patterns since each hypergraph can only represent a specific type of user patterns. Moreover, these single-channel hypergraphs may risk evolving into traditional Knowledge Graphs (KGs) due to the scarcity of user-item interaction data. Thus, the construction of multi-channel hypergraphs is paramount for exploring multi-level user interests.

To address these limitations, we propose the novel framework, **Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning for Conversational Recommendation (HiCore)**, which aims to mitigate the negative impact of Matthew effect when users engage with the system over time in the CRS. It is comprised of Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning and Interest-Boosted CRS. The former devotes to construct multi-hypergraph (*i.e.*, item-oriented, entity-oriented, and word-oriented triple-channel hypergraphs) to learn multi-level user interests (*i.e.*, item-level, entity-level, word-level triple-channel interests), where triple channels contain the group, joint, and purchase channels. The latter aims to utilize the multi-level interests to enhance both conversation and recommendation tasks when users chat with system over time. Concretely, multi-level user interests are used to effectively generate next utterances in the conversational task, and accurately predict users' interested items in the recommendation task. Extensive experimental results on four benchmarks show that HiCore achieves a new state-of-the-art performance compared all the baselines, and the effectiveness of mitigating Matthew effect in the CRS.

Overall, our main contributions are included:

- To the best of our knowledge, this is the first work to build multi-hypergraph from triple-channel settings for learning multi-interest to mitigate Matthew effect in the CRS.
- We proposed a novel end-to-end framework HiCore, aiming to use multi-interest enhance both recommendation and conversation tasks.
- Quantitative and qualitative experimental results show the effectiveness of HiCore and the superiority of HiCore in mitigating Matthew effect.

2 Related Work

2.1 Matthew Effect in Recommendation

The Matthew effect poses a formidable challenge in recommendation systems. To combat this issue, there are two primary research lines. One line of research focuses on understanding a diverse range of user interests to enhance recommendation diversification (Anderson et al., 2020; Hansen et al., 2021; Liang et al., 2021; Zheng et al., 2021a). The other line of research (Zheng et al., 2021b) is dedicated to mitigating popularity bias to ensure a balanced exposure of items across various categories. For example, Wang et al. (Wang et al., 2019) conducted a meticulous quantitative analysis, providing valuable insights into the quantitative characteristics of the Matthew effect in collaborative-based recommender systems. Liu et al. (Liu and Huang, 2021) have confirmed the presence and impact of the Matthew effect within the intricate algorithms of YouTube's recommendation system. However, these methods primarily concentrate on exploring the Matthew effect in static recommendation environments, overlooking the crucial interplay of the user-system feedback loop.

2.2 Conversational Recommender System

Conversational Recommendation System aims to uncover users' genuine intentions and interests through natural language dialogues, thereby offering top-notch recommendations to users. Currently, CRS-based methods can be categorized into two main groups. 1) Attribute-based CRS (Deng et al., 2021a; Lei et al., 2020a,b; Ren et al., 2021; Xu et al., 2021), which seeks to delve into user interests by posing queries about items or their attributes. However, this approach primarily relies on predefined templates for response generation, often falling short in producing fluent, human-like natural language expressions. 2) Generated-based CRS (Chen et al., 2019; Deng et al., 2023; Li et al., 2022; Zhou et al., 2020a, 2022; Shang et al., 2023), which can address the shortcomings of attribute-centric CRS by utilizing the Seq2Seq architecture (Vaswani et al., 2017a) to integrate a conversation component and a recommendation component, resulting in the creation of smooth and coherent human-like responses. Despite their effectiveness, they face challenges in grasping the varied interests of users because of the restricted and scarce character of user-item interaction data.

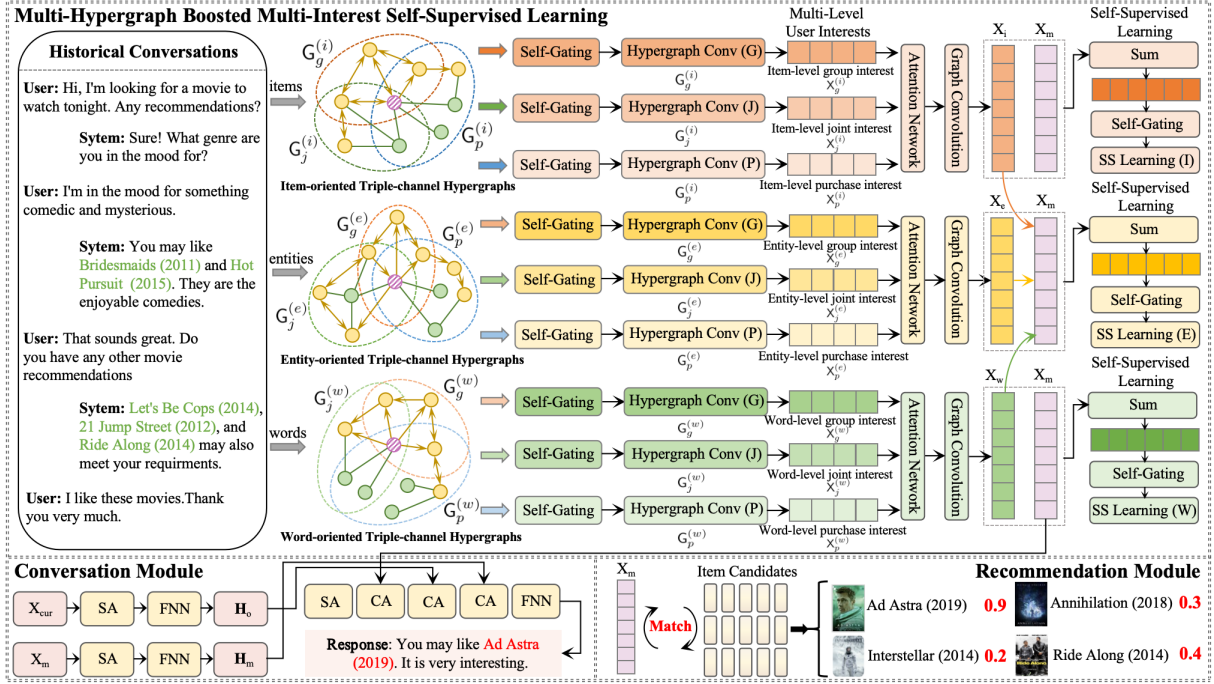


Figure 1: Overview of our HiCore framework. It consists of Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning and Interested-Boosted CRS. The former aims to learn multi-level user interests, while the latter devotes to generate responses in the conversation module and predict items in the recommendation module.

3 HiCore

Most existing methods (Hussein et al., 2020; Liu et al., 2021a; Nguyen et al., 2014) have consistently revealed that individuals with constrained interests are greatly impacted by Matthew effect. Thus, we propose a novel framework, HiCore, which is comprised of Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning and Interest-Boosted CRS. The overall pipeline of the proposed HiCore is illustrated in Fig.1.

3.1 Multi-Hypergraph Boosted Multi-Interest Self-Supervised Learning

In this section, we will establish multi-hypergraph to learn multi-level user interests to mitigate Matthew effect in the CRS.

3.1.1 Multi-Hypergraph Boosts Multi-Interest

Instead of linking only two nodes per edge as in traditional KGs, hypergraphs extend the notion of edges to connect more than two nodes. By utilizing diverse hypergraphs to encode various high-order user relation patterns, we construct multiple knowledge-oriented triple-channel hypergraphs.

Item-oriented triple-channel Hypergraphs. We first build item-oriented hypergraphs from triple channels, *i.e.*, ‘Group Channel (g)’, ‘Joint Channel (j)’, and ‘Purchase Channel (p)’ via the Motif

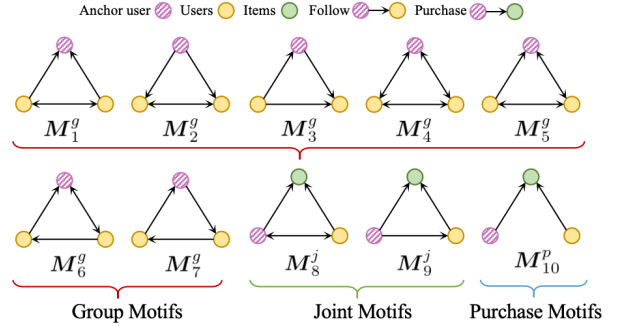


Figure 2: Triangle motifs used in our proposed HiCore.

(Milo et al., 2002; Yu et al., 2021), a commonly utilized tool for capturing complex local structures involving multiple nodes, as illustrated in Fig.2.

Group-channel hypergraph. Group-channel hypergraphs aim to analyze users’ social relations to unveil the dynamics among individuals based on their shared interests, preferences, and characteristics. Understanding group preferences not only consolidates individual tastes but also facilitates collective decisions that benefit the entire group. Formally, we utilize a set of triangular motifs (Milo et al., 2002; Yu et al., 2021) to build the item-oriented group-channel hypergraphs $G_g^{(i)}$ as:

$$G_g^{(i)} = (\mathcal{V}_g^{(i)}, \mathcal{N}_g^{(i)}, \mathbf{A}_{M_k^g}^{(i)}). \quad (1)$$

Here $\mathcal{V}_g^{(i)}$ represents the set of items derived from the historical conversations, while $\mathcal{N}_g^{(i)} = \{M_k^g | 1 \leq k \leq 7\}$ denotes the collection of hyperedges, with each hyperedge representing an occurrence of the specified motif M_k^g in Fig.2. $\mathbf{A}_{M_k^g}^{(i)} \in |\mathcal{V}_g^{(i)}| \times |\mathcal{N}_g^{(i)}|$ is the group-motif-induced adjacency matrices. Firstly, we need to define the matrix computation of each type of motif. Let $\mathbf{H}_k^{(i)}$ be the matrix computation of the motif M_k^g , then we can obtain:

$$\begin{cases} \mathbf{H}_1^{(i)} = (\mathbf{I}^T \mathbf{J}) \otimes \mathbf{I}^T + (\mathbf{J} \mathbf{I}) \otimes \mathbf{I} + (\mathbf{I} \mathbf{I}^T) \otimes \mathbf{J}, \\ \mathbf{H}_2^{(i)} = (\mathbf{I} \mathbf{J}) \otimes \mathbf{I} + (\mathbf{J} \mathbf{I}^T) \otimes \mathbf{I}^T + (\mathbf{I}^T \mathbf{I}) \otimes \mathbf{J}, \\ \mathbf{H}_3^{(i)} = (\mathbf{I} \mathbf{I}) \otimes \mathbf{I} + (\mathbf{I} \mathbf{I}^T) \otimes \mathbf{I} + (\mathbf{I}^T \mathbf{I}) \otimes \mathbf{I}, \\ \mathbf{H}_4^{(i)} = (\mathbf{J} \mathbf{J}) \otimes \mathbf{J}, \\ \mathbf{H}_5^{(i)} = (\mathbf{J} \mathbf{J}) \otimes \mathbf{I} + (\mathbf{J} \mathbf{I}) \otimes \mathbf{J} + (\mathbf{I} \cdot \mathbf{J}) \otimes \mathbf{J}, \\ \mathbf{H}_6^{(i)} = (\mathbf{J} \mathbf{I}) \otimes \mathbf{I}^T + (\mathbf{I} \mathbf{J}) \otimes \mathbf{I}^T + (\mathbf{I} \mathbf{I}) \otimes \mathbf{J}, \\ \mathbf{H}_7^{(i)} = (\mathbf{I} \mathbf{I}) \otimes \mathbf{I}^T, \end{cases} \quad (2)$$

where \otimes is the element-wise product, \mathbf{S} denote the relation matrix (Yu et al., 2021). $\mathbf{J} = \mathbf{S} \otimes \mathbf{S}$, and $\mathbf{I} = \mathbf{S} - \mathbf{J}$ specify the adjacency matrices of the bidirectional and unidirectional social networks (i.e., group motif), respectively. Then, the group-motif-induced adjacency matrices $\mathbf{A}_{M_k^g}^{(i)}$ is:

$$\mathbf{A}_{M_k^g}^{(i)} = \begin{cases} \mathbf{H}_1^{(i)}, & \text{if } M_1^g, \\ \mathbf{H}_2^{(i)}, & \text{if } M_2^g, \\ \mathbf{H}_3^{(i)} + (\mathbf{H}_5^{(i)})^T, & \text{if } M_3^g, \\ \mathbf{H}_4^{(i)}, & \text{if } M_4^g, \\ \mathbf{H}_5^{(i)} + (\mathbf{H}_3^{(i)})^T, & \text{if } M_5^g, \\ \mathbf{H}_6^{(i)} + (\mathbf{H}_2^{(i)})^T, & \text{if } M_6^g, \\ \mathbf{H}_7^{(i)} + (\mathbf{H}_1^{(i)})^T, & \text{if } M_7^g. \end{cases} \quad (3)$$

If $(\mathbf{A}_{M_k^g}^{(i)})_{n,r} = 1$, it signifies that the node n and the node r co-occur in a single instance of M_k^g . When two nodes appear in multiple instances, it turns to be $(\mathbf{A}_{M_k^g}^{(i)})_{n,r} = \#(n, r \text{ occur in the same instance of } M_k^g)$.

Joint-channel hypergraph. The joint channel reflects the scenario of shared behaviors among friends in a social network. When friends purchase the same items, it not only suggests similarities in tastes and interests but also hints at deeper levels of interaction and trust. This phenomenon of "friends

purchasing the same item" may facilitate information dissemination and interaction within the social network, strengthening social relationships, and to some extent, reflecting influence and collective behavior within the social network. Therefore, by identifying and analyzing the joint motifs, the item-oriented joint-channel hypergraph $\mathbf{G}_j^{(i)}$ is:

$$\begin{cases} \mathbf{G}_j^{(i)} = (\mathcal{V}_j^{(i)}, \mathcal{N}_j^{(i)}, \mathbf{A}_{M_k^j}^{(i)}), \\ \mathbf{H}_8^{(i)} = (\mathbf{R} \mathbf{R}^T) \otimes \mathbf{J}, \\ \mathbf{H}_9^{(i)} = (\mathbf{R} \mathbf{R}^T) \otimes \mathbf{I}, \\ \mathbf{A}_{M_k^j}^{(i)} = \mathbf{H}_8^{(i)}, & \text{if } M_8^j, \\ \mathbf{A}_{M_k^j}^{(i)} = \mathbf{H}_9^{(i)} + [\mathbf{H}_9^{(i)}]^T, & \text{if } M_9^j, \end{cases} \quad (4)$$

where $\mathcal{V}_j^{(i)}$, and $\mathcal{N}_j^{(i)} = \{M_k^j | 8 \leq k \leq 9\}$ denote the item set, and the hyperedge set, respectively. Each hyperedge is induced from each type of joint motif, depicted in Fig.2. \mathbf{R} is a binary matrix that records user-item interactions, and $\mathbf{A}_{M_k^j}^{(i)}$ denotes the joint-motif-induced adjacency matrices.

Purchase-channel hypergraph. Additionally, we should also take into account users who do not have explicit social connections. Therefore, the analysis is non-exclusive and delineates the implicit higher-order social relationships among users who lack direct social ties but still purchase the same items. By considering these users without overt social links, we can uncover hidden patterns of social influence and affiliation that transcend traditional network structures. Thus, the item-oriented purchase-channel hypergraph $\mathbf{G}_p^{(i)}$ can be induced from the purchase motif M_{10}^p as follows:

$$\begin{cases} \mathbf{G}_p^{(i)} = (\mathcal{V}_p^{(i)}, \mathcal{N}_p^{(i)}, \mathbf{A}_{M_k^p}^{(i)}), \\ \mathbf{A}_{M_k^p}^{(i)} = \mathbf{H}_{10}^{(i)} = \mathbf{R} \mathbf{R}^T, & \text{if } M_{10}^p, \end{cases} \quad (5)$$

here $\mathcal{V}_p^{(i)}$ and $\mathcal{N}_p^{(i)} = \{M_k^p | k = 10\}$ are the item set and hyperedge set, respectively. Specifically, the hyperedge set, depicted in Fig.2. $\mathbf{A}_{M_k^p}^{(i)}$ is the purchase-motif-induced adjacency matrices.

Entity-oriented triple-channel hypergraphs. To tackle the sparsity and constraints inherent in historical user-item interaction data, we leverage the rich DBpedia KG (Auer et al., 2007) to build an entity-oriented hypergraph. More precisely, we identify individual items referenced in historical conversations as entities and their k -hop neighbors

to construct each hyperedge. This method enables us to capture shared semantic nuances among the broader network of neighbors. Similar to item-oriented triple-channel hypergraphs, we build the entity-oriented hypergraphs from triple channel setting. Formally, the entity-oriented hypergraphs $\mathbf{G}_c^{(e)}$ from triple channel c can be given as:

$$\mathbf{G}_c^{(e)} = (\mathcal{V}_c^{(e)}, \mathcal{N}_c^{(e)}, \mathbf{A}_{M_k^c}^{(e)}). \quad (6)$$

Here $c \in \{g, j, p\}$ represents triple channels (*i.e.*, group, joint, and purchase channel). $\mathcal{V}_c^{(e)}$ denotes the entities from triple-channel setting. These entities are k -hop neighbors extracted from the historical conversations. $\mathcal{N}_c^{(e)}$ means the hyperedge set induced from different motifs. Each hyperedge is an instance of the Motif. $\mathbf{A}_{M_k^c}^{(e)}$ represents the group-channel, joint-channel, and purchase-channel adjacency matrices, they are defined as Eq.(3), Eq.(4), and Eq.(5), respectively.

Word-oriented triple-channel hypergraphs. The significance of keywords exchanged during conversations is paramount in grasping users' requirements. By scrutinizing notable words, we can pinpoint specific inclinations, a critical aspect in modeling an array of user tastes. To realize this objective, we construct a lexeme-centric hypergraph utilizing the lexicon-focused ConceptNet (Speer et al., 2017) KG to unveil semantic associations such as synonymy, antonyms, and co-occurrence. Based on these analysis, the word-oriented hypergraphs from group-, joint-, and purchase-channel can be expressed as:

$$\mathbf{G}_c^{(w)} = (\mathcal{V}_c^{(w)}, \mathcal{N}_c^{(w)}, \mathbf{A}_{M_k^c}^{(w)}), \quad (7)$$

where $\mathcal{V}_c^{(w)}$ is the words from k -hop neighbors. $\mathcal{N}_c^{(w)}$ denotes the hyperedge set from different motifs, including group, joint, purchase motifs. $\mathbf{A}_{M_k^c}^{(w)}$ are the word-oriented adjacency matrices induced from triple channels, illustrated in Eq.(3) ~ Eq.(5).

3.1.2 Multi-Interest Self-Supervised Learning

After constructing a series of hypergraphs from triple-channel setting, we will construct multi-level user interests via the hypergraph convolution network (Yu et al., 2021), which can be written as:

$$\mathbf{P}_c^{(l+1)} = \mathbf{D}_c^{-1} \mathbf{K}_c \mathbf{L}_c^{-1} \mathbf{K}_c^T \mathbf{P}_c^{(l)} = \widehat{\mathbf{D}}_c^{-1} \mathbf{A}_c^{(i)} \mathbf{P}_c^{(l)}, \quad (8)$$

where $\mathbf{P}_c^{(l)}$ and $\mathbf{P}_c^{(l+1)}$ represent the output of the l -th and $(l+1)$ -th layers, respectively. Specifically,

the base user embedding is $\mathbf{P}_c^{(0)} = f_{\text{gate}}^c(\mathbf{P}^{(0)})$, and $f_{\text{gate}}^c(\cdot)$ is self-gating units (SGUs) to control the information flow to different channel from the base user embedding $\mathbf{P}^{(0)}$. \mathbf{D}_c is the degree matrix of \mathbf{A}_c , which is the summation of the motifs without considering self-connections (Yu et al., 2021). In terms of the group motifs, it can be defined as $\mathbf{A}_c^{(i)} = \sum_{k=1}^7 \mathbf{A}_{M_k}^{(i)}$, in terms of joint motifs, $\mathbf{A}_c^{(i)} = \mathbf{A}_{M_8}^{(i)} + \mathbf{A}_{M_9}^{(i)}$, and from the point of the purchase motifs, $\mathbf{A}_c^{(i)} = \mathbf{A}_{M_{10}}^{(i)} - (\mathbf{A}_{M_8}^{(i)} + \mathbf{A}_{M_9}^{(i)})$. Based on these analysis, the item-level interests from triple channel setting (*i.e.*, $\mathbf{X}_g^{(i)}$, $\mathbf{X}_j^{(i)}$, $\mathbf{X}_p^{(i)}$), the entity-level interests from triple channel setting (*i.e.*, $\mathbf{X}_g^{(e)}$, $\mathbf{X}_j^{(e)}$, $\mathbf{X}_p^{(e)}$), the word-level interests from triple channel setting (*i.e.*, $\mathbf{X}_g^{(w)}$, $\mathbf{X}_j^{(w)}$, $\mathbf{X}_p^{(w)}$) can be defined as:

$$\begin{aligned} \mathbf{X}_g^{(h)} &= \widehat{\mathbf{D}}_g^{-1} \left(\sum_{k=1}^7 \mathbf{A}_{M_k^g}^{(h)} \right) \mathbf{P}_g^{(L)}; \\ \mathbf{X}_j^{(h)} &= \widehat{\mathbf{D}}_j^{-1} (\mathbf{A}_{M_8^j}^{(h)} + \mathbf{A}_{M_9^j}^{(h)}) \mathbf{P}_j^{(L)}; \\ \mathbf{X}_p^{(h)} &= \widehat{\mathbf{D}}_p^{-1} (\mathbf{A}_{M_{10}^p}^{(h)} - (\mathbf{A}_{M_8^p}^{(h)} + \mathbf{A}_{M_9^p}^{(h)})) \mathbf{P}_p^{(L)}. \end{aligned} \quad (9)$$

Here $h \in \{i, e, w\}$, and L is the last hypergraph convolution layer. Then, we adopt the attention network $\text{Atta}(\cdot)$ and graph convolution $\text{GConv}(\cdot)$ to learn final multi-interest \mathbf{X}_m as:

$$\begin{aligned} \mathbf{X}_i &= \text{GConv}(\text{Atta}(\mathbf{X}_g^{(i)}; \mathbf{X}_j^{(i)}; \mathbf{X}_p^{(i)})), \\ \mathbf{X}_e &= \text{GConv}(\text{Atta}(\mathbf{X}_g^{(e)}; \mathbf{X}_j^{(e)}; \mathbf{X}_p^{(e)})), \\ \mathbf{X}_w &= \text{GConv}(\text{Atta}(\mathbf{X}_g^{(w)}; \mathbf{X}_j^{(w)}; \mathbf{X}_p^{(w)})), \\ \mathbf{X}_m &= \text{Atta}(\mathbf{X}_i; \mathbf{X}_e; \mathbf{X}_w). \end{aligned} \quad (10)$$

Here $;$ is the concatenation operation. Finally, we use InfoNCE (Yu et al., 2021) as our learning objective to conduct the self-supervised learning as:

$$\begin{aligned} \mathcal{L}_s &= - \sum_h \left\{ \sum_{u \in U} \log \sigma(f(\mathbf{X}_m, z_u^h)) - f(\mathbf{X}_m, \hat{z}_u^h) \right\} \\ &\quad + \sum_{u \in U} \log \sigma(f(z_u^h, \mathbf{k}^h) - f(\hat{z}_u^h, \mathbf{k}^h)). \end{aligned} \quad (11)$$

Here $z_u^h = f_{\text{gate}}^h(f_s(\mathbf{X}_h; \mathbf{p}^h))$, $f_s(\cdot)$ is the sum operation, \hat{z}_u^h is the negative example by shuffling both rows and columns of z_u^h , and h is defined as Eq.(9). $f(\cdot) \in \mathbb{R}^{d \times d}$ serves as the discriminator, evaluating the alignment between two input vectors. Specifically, $\mathbf{k}^h = f_{\text{out}}(Z_h)$, where Z_h and \mathbf{X}_m are ground truths for each other, and $f_{\text{out}}(\cdot)$ aims to perform permutation invariant (Yu et al., 2021).

3.2 Interest-Boosted CRS

To mitigate Matthew effect in the CRS, we employ multi-interest \mathbf{X}_m to enhance both recommendation and conversation tasks.

3.2.1 Recommendation Module

Recommendation module is to precisely forecast items for users via dynamic natural conversations. To improve recommendation diversity, we use \mathbf{X}_m to select target items as $\mathbf{P}_{\text{rec}} = \mathbf{X}_m \times \mathbf{V}_{\text{cand}}$, where \mathbf{V}_{cand} is the embeddings of all candidate items. Finally, we adopt cross-entropy loss (Shang et al., 2023) to learn the recommendation task:

$$\mathcal{L}_{\text{rec}} = - \sum_{b=1}^B \sum_{a=1}^{|\mathcal{I}|} \left\{ - (1 - Y_{ab}) \cdot \log(1 - \mathbf{P}_{\text{rec}}^{(b)}(a)) + Y_{ab} \cdot \log(\mathbf{P}_{\text{rec}}^{(b)}(a)) \right\}, \quad (12)$$

where $Y_{ab} \in \{0, 1\}$, B , and $|\mathcal{I}|$ are the target label, mini-batch size, the size of item set, respectively.

3.2.2 Conversation Module

Conversation module centers on crafting appropriate dialogue responses. Next, we use multi-interest \mathbf{X}_m to feed into Transformer $\text{MHA}(\cdot)$ to produce informative responses. Suppose \mathbf{Y}^{n-1} is the output of the last time unit, then the current one \mathbf{Y}^n is:

$$\begin{aligned} \mathbf{A}_0^n &= \text{MHA}(\mathbf{Y}^{n-1}, \mathbf{Y}^{n-1}, \mathbf{Y}^{n-1}), \\ \mathbf{A}_1^n &= \text{MHA}(\mathbf{A}_0^n, \mathbf{X}_m, \mathbf{X}_m), \\ \mathbf{A}_2^n &= \text{MHA}(\mathbf{A}_1^n, \mathbf{X}_{\text{cur}}, \mathbf{X}_{\text{cur}}), \\ \mathbf{A}_3^n &= \text{MHA}(\mathbf{A}_1^n, \mathbf{X}_{\text{his}}, \mathbf{X}_{\text{his}}), \\ \mathbf{A}_4^n &= \beta \cdot \mathbf{A}_2^n + (1 - \beta) \cdot \mathbf{A}_3^n, \\ \mathbf{Y}^n &= \text{FFN}(\mathbf{A}_4^n), \end{aligned} \quad (13)$$

where \mathbf{X}_{cur} and \mathbf{X}_{his} are the current and historical conversations, respectively. β is hyper-parameters to control the information flow. Then, we use cross-entropy loss to learn the conversation task:

$$\mathcal{L}_{\text{conv}} = - \sum_{b=1}^B \sum_{t=1}^T \log(\mathbf{P}_{\text{conv}}(s_t | \{s_{t-1}\})),$$

$$\mathbf{P}_{\text{conv}}(\cdot) = p_1(s_t | Y_i) + p_2(s_t | \mathbf{P}_{\text{rec}}) + p_3(s_t | \mathbf{P}_{\text{rec}}), \quad (14)$$

where $\mathbf{P}_{\text{conv}}(\cdot)$ is the probability of the next token when given a sequence $\{s_{t-1}\} = s_1, s_2, \dots, s_{t-1}$, where s_t signifies the t -th utterance. $p_1(\cdot)$, $p_2(\cdot)$, and $p_3(\cdot)$ denote the vocabulary probability, vocabulary bias, and copy probability, respectively. T is the truncated length of utterances.

3.3 Challenges Discussion

Throughout the developmental journey of hypergraphs, we surmounted several significant challenges, elaborated upon below:

1) **Hypergraph Construction Challenge:** During the project’s initial stages, the real-time construction of hypergraphs presented a bottleneck, resulting in delays. Through the strategic repositioning of this operation to the data preprocessing phase, we adeptly extracted essential subgraphs, leading to a noteworthy reduction in training time. This adjustment enhanced efficiency, streamlined processes, and improved performance.

2) **Graph Storage Challenge:** The transition to sparse graph storage mechanisms is pivotal in enhancing efficiency, streamlining computation time, and optimizing memory utilization. Embracing this shift not only boosts the system’s performance but also establishes a robust foundation for scalable and resource-efficient operations.

3) **Model Training Challenge:** With the emergence of a series of hypergraphs, optimizing the efficiency of model training becomes paramount. Consequently, we redefined our strategy by dispersing hypergraphs across multiple computing cards, enabling parallel computation and achieving a significant boost in the model’s runtime speed.

4 Experiments and Analyses

To fully evaluate the proposed HiCore, we conduct experiments to answer the following questions:

- **RQ1:** How does HiCore perform compared with all baselines in the recommendation task?
- **RQ2:** How does HiCore perform compared with all baselines in the conversation task?
- **RQ3:** How does HiCore mitigate Matthew effect in the CRS?
- **RQ4:** How do parameters affect our HiCore?
- **RQ5:** How do different hypergraphs contribute to the performance?

4.1 Experimental Protocol

Datasets. We assess the effectiveness of our proposed HiCore through comprehensive evaluations on four CRS-based benchmarks: REDIAL (Li et al., 2018b), TG-REDIAL (Zhou et al., 2020b), OpenDialKG (Moon et al., 2019), and DuRecDial (Liu et al., 2021b). The REDIAL dataset comprises 11,348 dialogues involving 956 users and 6,924 items, while the TG-REDIAL dataset encompasses 10,000 dialogues with 1,482 users and

Model	REDIAL						TG-REDIAL					
	R@10	R@50	M@10	M@50	N@10	N@50	R@10	R@50	M@10	M@50	N@10	N@50
TextCNN	0.0644	0.1821	0.0235	0.0285	0.0328	0.0580	0.0097	0.0208	0.0040	0.0045	0.0053	0.0077
SASRec	0.1117	0.2329	0.0540	0.0593	0.0674	0.0936	0.0043	0.0178	0.0011	0.0017	0.0019	0.0047
BERT4Rec	0.1285	0.3032	0.0475	0.0555	0.0663	0.1045	0.0043	0.0226	0.0013	0.0020	0.0020	0.0058
KGSF	0.1785	0.3690	0.0705	0.0796	0.0956	0.1379	0.0215	0.0643	0.0069	0.0087	0.0103	0.0194
TG-ReDial	0.1679	0.3327	0.0694	0.0771	0.0924	0.1286	0.0110	0.0174	0.0048	0.0050	0.0062	0.0076
ReDial	0.1705	0.3077	0.0677	0.0738	0.0925	0.1222	0.0038	0.0165	0.0012	0.0017	0.0018	0.0045
KBRD	0.1796	0.3421	0.0722	0.0800	0.0972	0.1333	0.0201	0.0501	0.0077	0.0090	0.0106	0.0171
BART	0.1693	0.3783	0.0646	0.0744	0.0888	0.1350	0.0047	0.0187	0.0012	0.0017	0.0020	0.0048
BERT	0.1608	0.3525	0.0597	0.0688	0.0831	0.1255	0.0040	0.0194	0.0011	0.0017	0.0018	0.0050
XLNet	0.1569	0.3590	0.0583	0.0677	0.0811	0.1255	0.0040	0.0187	0.0011	0.0017	0.0017	0.0048
KGConvRec	0.1819	0.3587	0.0711	0.0794	0.0969	0.1358	0.0220	0.0524	0.0088	0.0102	0.0119	0.0185
MHIM	0.1966	0.3832	0.0742	0.0830	0.1027	0.1440	0.0300	0.0783	0.0108	0.0129	0.0152	0.0256
HiCore*	0.2192	0.4163	0.0775	0.0874	0.1107	0.1558	0.0270	0.0769	0.0880	0.1074	0.0152	0.0225

Table 1: Recommendation results on REDIAL and TG-REDIAL datasets. * indicates statistically significant improvement ($p < 0.05$) over all baselines.

33,834 items. To provide a holistic evaluation of our proposed methodology, we integrate two cross-domain datasets, OpenDialKG and DuRecDial, which cover a wide array of domains including movies, music, books, sports, restaurants, news, and culinary experiences.

Baselines. We compared our HiCore with the following state-of-the-art methods **TextCNN** (Kim, 2014), **SASRec** (Kang and McAuley, 2018), **BERT4Rec** (Sun et al., 2019), **KBRD** (Chen et al., 2019), **Trans.** (Vaswani et al., 2017b), **ReDial** (Li et al., 2018a), **KGSF** (Zhou et al., 2020a), **KG-ConvRec** (Sarkar et al., 2020), **XLNet** (Yang et al., 2019), **BART** (Lewis et al., 2020), **BERT** (Devlin et al., 2019), **DialoGPT** (Zhang et al., 2020), **UniCRS** (Deng et al., 2021b), **GPT-3** (Brown et al., 2020), **C2-CRS** (Zhou et al., 2022), **LOT-CRS** (Zhao et al., 2023), **MHIM** (Shang et al., 2023), and **HyCoRec** (Zheng et al., 2024).

4.2 Recommendation Performance (RQ1)

In accordance with (Shang et al., 2023), we utilize Recall@ K (R@ K), MRR@ K (M@ K), and NDCG@ K (N@ K) ($K=1, 10, 50$) to assess the recommendation task. Analyzing the results presented in Table 1 and Table 2, it is evident that our proposed method, HiCore, consistently outperforms all the comparison baselines.

There exist multiple crucial facets contributing to the advancement of our proposed HiCore method: (a) Diversification of hypergraphs: we introduced a diverse set of hypergraphs, including item-oriented, entity-oriented, and word-oriented hypergraphs. This expansion aims to go beyond the traditional pairwise interactions, broadening

Model	OpenDialKG		DuRecDial	
	R@1	R@10	R@1	R@10
KBRD	0.1448	0.3162	0.0618	0.3971
KGSF	0.0626	0.1757	0.1395	0.4367
ReDial	0.0008	0.0134	0.0005	0.0336
TGReDial	0.2149	0.4035	0.0956	0.4882
HyCoRec	0.2742	0.4490	0.1279	0.4750
HiCore*	0.2628	0.4526	0.1735	0.5471

Model	OpenDialKG		DuRecDial	
	Dist-2	Dist-3	Dist-2	Dist-3
KBRD	0.3192	1.7660	0.5180	1.5500
KGSF	0.1687	0.5387	0.1389	0.3862
ReDial	0.1579	0.5808	0.1095	0.3981
TGReDial	0.4836	2.1430	0.5453	2.0030
HyCoRec	2.8190	4.7710	1.0820	2.4440
HiCore*	2.8430	4.8120	1.0940	2.4280

Table 2: Results on both recommendation and conversation tasks on OpenDialKG and DuRecDial datasets involving various domains. * indicates statistically significant improvement ($p < 0.05$) over all baselines.

the scope of user interest modeling by incorporating interactions among multiple nodes. (b) Exploration of hypergraph configurations: moving beyond the conventional triple-channel model, we delved into various hypergraph configurations like group-channel, joint-channel, and purchase-channel. These configurations were designed to cater not only to social connections but also individual preferences, enhancing the system’s adaptability. (c) Integration of multi-level user interests: transitioning from the triple-channel structure, we integrated these hypergraphs to capture multi-level user interests. This strategic shift helps alleviate the Matthew effect in the CRS involving the dynamic user-system feedback loop. This comprehensive approach highlights the innovation and adaptabil-

Model	REDIAL			TG-REDIAL		
	Dist-2	Dist-3	Dist-4	Dist-2	Dist-3	Dist-4
ReDial	0.0214	0.0659	0.1333	0.2178	0.5136	0.7960
Trans.	0.0538	0.1574	0.2696	0.2362	0.7063	1.1800
KGSF	0.0572	0.2483	0.4349	0.3891	0.8868	1.3337
KBRD	0.0765	0.3344	0.6100	0.8013	1.7840	2.5977
DialoGPT	0.3542	0.6209	0.9482	1.1881	2.4269	3.9824
GPT-3	0.3604	0.6399	0.9511	1.2255	2.5713	4.0713
UniCRS	0.2464	0.4273	0.5290	0.6252	2.2352	2.5194
C2-CRS	0.2623	0.3891	0.6202	0.5235	1.9961	2.9236
LOT-CRS	0.3312	0.6155	0.9248	0.9287	2.4880	3.4972
MHIM	0.3278	0.6204	0.9629	1.1100	2.3520	3.8200
HiCore*	0.5871	1.1170	1.7500	2.8610	5.7440	8.4160

Table 3: Conversation results on REDIAL and TG-REDIAL datasets. * indicates statistically significant improvement ($p < 0.05$) over all baselines.

ity of HiCore in addressing the intricacies of user interest modeling and enhancing recommendation system performance.

4.3 Conversational Performance (RQ2)

For the conversation task, we use Distinct n -gram (Dist- n) (Shang et al., 2023) ($n=2,3,4$) to estimate the conversation task. Table 2 and Table 3 indicate a significant performance superiority of our HiCore. For example, HiCore gains 123.83%, 138.27%, 77.26%, 65.75%, 62.90%, and 79.10% improvements on Dist-2 against the strong baselines including, C2-CRS, UniCRS, LOT-CRS, DialoGPT, GPT-3, and MHIM on the REDIAL dataset, respectively. It also gains 446.51%, 357.61%, 208.07%, 140.80%, 133.46%, and 157.75% improvements on Dist-2 against the strong baselines including, C2-CRS, UniCRS, LOT-CRS, DialoGPT, GPT-3, and MHIM on the REDIAL dataset, respectively.

The improvement in HiCore can be attributed to the following reasons: (a) Our HiCore focuses on constructing a diverse set of hypergraphs, encompassing item-oriented, entity-oriented, and word-oriented triple-channel hypergraphs. These structures effectively capture intricate local patterns through motif analysis, enabling the exploration of high-order user behaviors. This proves invaluable in generating informative and high-quality response utterances. (b) HiCore is dedicated to mitigating the Matthew effect that may occur as users engage with the system over time. By learning multi-level user interests from the hypergraphs, the system can adapt to users’ evolving preferences. This strategic approach enables the CRS to provide a varied array of responses that align with the diverse interests of the users.




Figure 3: Coverage results of C@k metric.

Model	OpenDialKG			DuRecDial		
	A@5	A@15	A@30	A@5	A@15	A@30
KBRD	0.0025	0.0025	0.0088	0.0318	0.0562	0.0938
KGSF	0.0051	0.0108	0.0182	0.0276	0.0534	0.0952
ReDial	1.0000	0.9375	0.8333	1.0000	0.8824	0.9677
TGReDial	0.0022	0.0043	0.0070	0.0137	0.0399	0.0796
MHIM	0.0022	0.0044	0.0075	0.0228	0.0434	0.0789
HiCore	0.0017	0.0043	0.0065	0.0226	0.0423	0.0751

Model	OpenDialKG			DuRecDial		
	L@5	L@15	L@30	L@5	L@15	L@30
KBRD	0.2921	0.2782	0.2782	0.3758	0.4149	0.3406
KGSF	0.2398	0.2482	0.3343	0.3314	0.4243	0.3302
ReDial	1.0000	0.8750	0.8333	1.0000	0.8235	0.9677
TGReDial	0.2737	0.2482	0.2757	0.3654	0.3803	0.3846
MHIM	0.1919	0.2343	0.2617	0.3315	0.3488	0.2706
HiCore*	0.1906	0.2092	0.2343	0.3122	0.3267	0.2666

Table 4: Results of Average Popularity (A@K) and Long Tail Ratio (L@K).

4.4 Study on Matthew Effect (RQ3)

Given our goal of mitigating the Matthew effect that may arise as users interact with the system over time, we engage in a series of experiments comparing the proposed method with the most robust baselines. This investigation seeks to determine the efficacy of HiCore in effectively alleviating the Matthew effect. Considering the key strategy to mitigate Matthew effect is to improve the recommendation diversification, and thus we use the diversify-based evaluation metrics *Coverage@k* (C@k), *Average Popularity* (A@K) of Recommended Items and *Long Tail Recommendation Ratio* (L@K) to comprehensively evaluate the efficacy of our proposed method in mitigating the Matthew Effect.

Fig.3 illustrates the experimental outcomes, showcasing the consistent superiority of HiCore in achieving the highest levels of *Coverage* across all datasets in comparison to the most robust baselines. The heightened coverage metric highlights its exceptional ability to encompass a broad spectrum of the recommendation space by incorporating items from diverse categories. Additionally, as

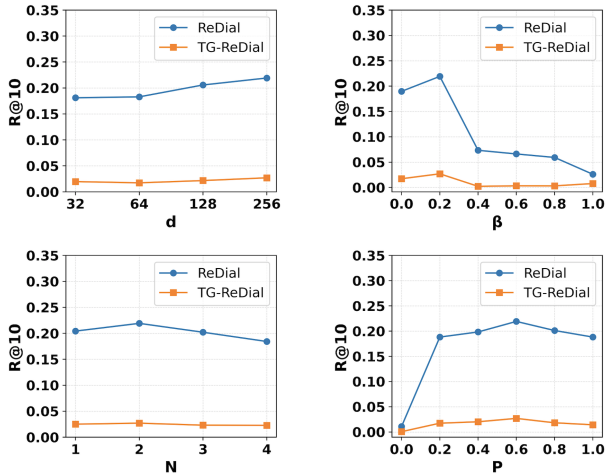


Figure 4: Impact of different hyperparameters.

outlined in Table 4, our proposed method demonstrates the lowest values for *Average Popularity* and *Long Tail Ratio*. This evidence suggests that our method effectively mitigates the adverse effects of item popularity on recommendation outcomes and successfully addresses the long tail distribution of items. These results validate the effectiveness of our proposed approach in combating the Matthew effect in the CRS as users interact with the system over time, attributed to its capability to learn multi-level user interests through a series of hypergraphs from triple-channel setting, including group, joint, and purchase channels.

4.5 Hyperparameters Analysis (RQ4)

Hyperparameters are parameters in a machine learning algorithm that need to be manually set and tuned to optimize model performance, distinct from the parameters that the model learns during training. Next, we will delve into the research on how various hyperparameters influence the performance of recommendations, including the embedding dimension d , comparative learning weight β , hypergraph convolution layers N , and the hyperedge threshold P . From Fig.4, we can obtain: (1) Elevating the feature dimensionality enhances outcomes, as higher dimensions can encapsulate more intricate features effectively; (2) Having too few hyperedges may hinder the capture of intricate local patterns, whereas an excess of hyperedges could impede the model’s convergence; (3) A lower beta value signifies a reduced weight for the comparison term, which show that the recommendation term exerts a more significant influence on the results; (4) A two-layer hyperconv network is sufficient to

Model	REDIAL		TG-REDIAL	
	R@10	R@50	R@10	R@50
HiCore	0.2192	0.4163	0.0270	0.0769
w/o $G_g^{(i)}$	0.2075	0.4160	0.0234	0.0742
w/o $G_j^{(i)}$	0.2012	0.4026	0.0217	0.0706
w/o $G_p^{(i)}$	0.1939	0.4096	0.0220	0.0739
w/o $G_g^{(e)}$	0.2067	0.4044	0.0247	0.0713
w/o $G_j^{(e)}$	0.2142	0.4122	0.0253	0.0756
w/o $G_p^{(e)}$	0.1971	0.4110	0.0243	0.0693
w/o $G_g^{(w)}$	0.2067	0.4142	0.0264	0.0761
w/o $G_j^{(w)}$	0.2151	0.4145	0.0223	0.0733
w/o $G_p^{(w)}$	0.2067	0.3974	0.0263	0.0759

Table 5: Ablation studies on the recommendation task.

encode high-level features for enhancing recommendation performance.

4.6 Ablation Studies (RQ5)

To assess the efficacy of each component within the proposed method, we perform ablation experiments using various iterations of HiCore, including: 1) w/o G_g^i , w/o G_j^i , w/o G_p^i : removing item-oriented group-channel, joint-channel, purchase-channel hypergraph, respectively; 2) w/o G_g^e , w/o G_j^e , w/o G_p^e : removing entity-oriented group-channel, joint-channel, purchase-channel hypergraph, respectively; 3) w/o G_g^w , w/o G_j^w , w/o G_p^w : removing word-oriented group-channel, joint-channel, purchase-channel hypergraph, respectively.

Table 5 outlines the experimental findings, indicating that the removal of any hypergraph type results in a performance decrease. This highlights the effectiveness of each hypergraph type and underscores the superiority of HiCore in learning multi-level user interests through a collection of hypergraphs to mitigate Matthew effect in the CRS.

5 Conclusion

The Matthew effect poses a significant challenge in the CRS due to the dynamic user-system feedback loop, which tends to escalate over time as users engage with the system. In response to these challenges, we proposed a novel framework, HiCore, aimed at mitigating the Matthew effect by capturing multi-level user interests through a variety of hypergraphs, including item-oriented, entity-oriented, and word-oriented triple-channel hypergraphs. Extensive experiments validate that HiCore outperforms all baselines, demonstrating the effectiveness of HiCore in addressing the Matthew effect as users chat with the system over time in the CRS.

6 Limitations

While our HiCore has achieved a remarkable state-of-the-art performance, it does come with certain limitations. Firstly, triple-channel hypergraphs may present challenges due to their computational complexity, interpretational intricacies, and potential issues with sparse data. Secondly, scaling these hypergraphs to larger datasets could introduce scalability hurdles, with a risk of overfitting when the model becomes excessively fine-tuned to the training data. Furthermore, ensuring generalizability and handling resource-intensive computations are crucial factors to consider when leveraging multi-channel hypergraphs.

7 Ethics Statement

The data used in this paper are sourced from open-access repositories, and do not pose any privacy concerns. We are confident that our research adheres to the ethical standards set forth by EMNLP.

8 Acknowledgements

This research / project is supported by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Trust Tech Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore and Infocomm Media Development Authority.

References

- Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic effects on the diversity of consumption on spotify. In *The Web Conference*, pages 2155–2165.
- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. Dbpedia: A nucleus for a web of open data. In *International Semantic Web Conference/Asian Semantic Web Conference*, volume 4825, pages 722–735.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Conference on Neural Information Processing Systems*.
- Qibin Chen, Junyang Lin, Yichang Zhang, Ming Ding, Yukuo Cen, Hongxia Yang, and Jie Tang. 2019. Towards knowledge-based recommender dialog system. *arXiv preprint arXiv:1908.05391*.
- Yang Deng, Yaliang Li, Fei Sun, Bolin Ding, and Wai Lam. 2021a. Unified conversational recommendation policy learning via graph-based reinforcement learning. In *Conference on Research and Development in Information Retrieval*, pages 1431–1441.
- Yang Deng, Yaliang Li, Fei Sun, Bolin Ding, and Wai Lam. 2021b. Unified conversational recommendation policy learning via graph-based reinforcement learning. In *Conference on Research and Development in Information Retrieval*, pages 1431–1441.
- Yang Deng, Wenxuan Zhang, Weiwen Xu, Wenqiang Lei, Tat-Seng Chua, and Wai Lam. 2023. A unified multi-task learning framework for multi-goal conversational recommender systems. *ACM Transactions on Information Systems*, 41(3):1–25.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186. Association for Computational Linguistics.
- Elena V. Epure and Romain Hennequin. 2023. A human subject study of named entity recognition in conversational music recommendation queries. In *European Chapter of the Association for Computational Linguistics*, pages 1273–1288.
- Yingqiang Ge, Shuya Zhao, Honglu Zhou, Changhua Pei, Fei Sun, Wenwu Ou, and Yongfeng Zhang. 2020. Understanding echo chambers in e-commerce recommender systems. In *Conference on Research and Development in Information Retrieval*, pages 2261–2270. ACM.
- Christian Hansen, Rishabh Mehrotra, Casper Hansen, Brian Brost, Lucas Maystre, and Mounia Lalmas. 2021. Shifting consumption towards diverse content on music streaming platforms. In *Conference on Web Search and Data Mining*, pages 238–246. ACM.
- Eslam Hussein, Prerna Juneja, and Tanushree Mitra. 2020. Measuring misinformation in video search platforms: An audit study on youtube. *ACM on Human-Computer Interaction*, 4(CSCW):048:1–048:27.
- Wang-Cheng Kang and Julian J. McAuley. 2018. Self-attentive sequential recommendation. In *IEEE International Conference on Data Mining*, pages 197–206.

- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Empirical Methods in Natural Language Processing (Demonstrations)*, pages 1746–1751.
- Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020a. Estimation-action-reflection: Towards deep interaction between conversational and recommender systems. In *Web Search and Data Mining*, pages 304–312.
- Wenqiang Lei, Gangyi Zhang, Xiangnan He, Yisong Miao, Xiang Wang, Liang Chen, and Tat-Seng Chua. 2020b. Interactive path reasoning on graph for conversational recommendation. In *International Conference on Knowledge Discovery and Data Mining*, pages 2073–2083.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *the Association for Computational Linguistics*, pages 7871–7880. Association for Computational Linguistics.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018a. Towards deep conversational recommendations. *Advances in Neural Information Processing Systems*, 31.
- Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018b. Towards deep conversational recommendations. In *Advances in Neural Information Processing Systems*, pages 9748–9758.
- Shuokai Li, Ruobing Xie, Yongchun Zhu, Xiang Ao, Fuzhen Zhuang, and Qing He. 2022. User-centric conversational recommendation with multi-aspect user modeling. In *Conference on Research and Development in Information Retrieval*, pages 223–233.
- Yile Liang, Tiejun Qian, Qing Li, and Hongzhi Yin. 2021. Enhancing domain-level and user-level adaptivity in diversified recommendation. In *Conference on Research and Development in Information Retrieval*, pages 747–756. ACM.
- Ping Liu, Karthik Shivaram, Aron Culotta, Matthew A. Shapiro, and Mustafa Bilgic. 2021a. The interaction between political typology and filter bubbles in news recommendation algorithms. In *The Web Conference*, pages 3791–3801.
- Ying Chieh Liu and Min Qi Huang. 2021. Examining the matthew effect on youtube recommendation system. In *Conference on Technologies and Applications of Artificial Intelligence*, pages 146–148.
- Yuanxing Liu, Weinan Zhang, Baohua Dong, Yan Fan, Hang Wang, Fan Feng, Yifan Chen, Ziyu Zhuang, Hengbin Cui, Yongbin Li, and Wanxiang Che. 2023. U-NEED: A fine-grained dataset for user needs-centric e-commerce conversational recommendation. In *Conference on Research and Development in Information Retrieval*, pages 2723–2732. ACM.
- Zeming Liu, Haifeng Wang, Zhengyu Niu, Hua Wu, and Wanxiang Che. 2021b. Durecdial 2.0: A bilingual parallel corpus for conversational recommendation. In *Conference on Empirical Methods in Natural Language Processing EMNLP*, pages 4335–4347. Association for Computational Linguistics.
- R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network motifs: Simple building blocks of complex networks. *Science*, 298(5594):824–827.
- Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023. Help me heal: A reinforced polite and empathetic mental health and legal counseling dialogue system for crime victims. In *Association for the Advancement of Artificial Intelligence*, pages 14408–14416.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *Conference of the Association for Computational Linguistics ACL*, pages 845–854. Association for Computational Linguistics.
- Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren G. Terveen, and Joseph A. Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *The Web Conference*, pages 677–686. ACM.
- Libo Qin, Zhouyang Li, Qiying Yu, Lehan Wang, and Wanxiang Che. 2023. Towards complex scenarios: Building end-to-end task-oriented dialogue system across multiple knowledge bases. In *Association for the Advancement of Artificial Intelligence*, pages 13483–13491.
- Xuhui Ren, Hongzhi Yin, Tong Chen, Hao Wang, Zi Huang, and Kai Zheng. 2021. Learning to ask appropriate questions in conversational recommendation. In *Conference on Research and Development in Information Retrieval*, pages 808–817.
- Rajdeep Sarkar, Koustava Goswami, Mihael Arcan, and John Philip McCrae. 2020. Suggest me a movie for tonight: Leveraging knowledge graphs for conversational recommendation. In *Conference on Computational Linguistics*, pages 4179–4189.
- Chenzhan Shang, Yupeng Hou, Wayne Xin Zhao, Yaliang Li, and Jing Zhang. 2023. Multi-grained hypergraph interest modeling for conversational recommendation. *AI Open*, 4:154–164.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Association for the Advancement of Artificial Intelligence*, pages 4444–4451.

- Harald Steck. 2018. Calibrated recommendations. In *Conference on Recommender Systems*, pages 154–162.
- Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In *International Conference on Information and Knowledge Management*, pages 1441–1450. ACM.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017a. Attention is all you need. *Advances in neural information processing systems*, 30.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017b. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Hao Wang, Zonghu Wang, and Weishi Zhang. 2019. Quantitative analysis of matthew effect and sparsity problem of recommender systems. *CoRR*.
- Kerui Xu, Jingxuan Yang, Jun Xu, Sheng Gao, Jun Guo, and Ji-Rong Wen. 2021. Adapting user preference to online feedback in conversational recommendation. In *Web Search and Data Mining*, pages 364–372.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems*, pages 5754–5764.
- Junliang Yu, Hongzhi Yin, Jundong Li, Qinyong Wang, Nguyen Quoc Viet Hung, and Xiangliang Zhang. 2021. Self-supervised multi-channel hypergraph convolutional network for social recommendation. In *World Wide Web WWW*, pages 413–424. ACM / IW3C2.
- Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal intervention for leveraging popularity bias in recommendation. In *Conference on Research and Development in Information Retrieval*, pages 11–20. ACM.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. DIALOGPT : Large-scale generative pre-training for conversational response generation. In *the Association for Computational Linguistics*, pages 270–278. Association for Computational Linguistics.
- Zipeng Zhao, Kun Zhou, Xiaolei Wang, Wayne Xin Zhao, Fan Pan, Zhao Cao, and Ji-Rong Wen. 2023. Alleviating the long-tail problem in conversational recommender systems. In *ACM Conference on Recommender Systems*, pages 374–385. ACM.
- Yongsen Zheng, Ruilin Xu, Ziliang Chen, Guohua Wang, Mingjie Qian, Jinghui Qin, and Liang Lin. 2024. Hycorec: Hypergraph-enhanced multi-preference learning for alleviating matthew effect in conversational recommendation. In *the Association for Computational Linguistics ACL*, pages 2526–2537. Association for Computational Linguistics.
- Yu Zheng, Chen Gao, Liang Chen, Depeng Jin, and Yong Li. 2021a. DGCN: diversified recommendation with graph convolutional networks. In *The Web Conference*, pages 401–412.
- Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Yong Li, and Depeng Jin. 2021b. Disentangling user interest and conformity for recommendation with causal embedding. In *The Web Conference*, pages 2980–2991.
- Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 2020a. Improving conversational recommender systems via knowledge graph based semantic fusion. In *International Conference on Knowledge Discovery and Data Mining*, pages 1006–1014.
- Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020b. Towards topic-guided conversational recommender system. In *International Conference on Computational Linguistics*, pages 4128–4139.
- Yuanhang Zhou, Kun Zhou, Wayne Xin Zhao, Cheng Wang, Peng Jiang, and He Hu. 2022. C²-crs: Coarse-to-fine contrastive learning for conversational recommender system. In *Web Search and Data Mining*, pages 1488–1496. ACM.