Title	Research on Improving Online Recommendations Involving Sparse Interaction Data
Author(s)	李, 智
Citation	大阪大学, 2024, 博士論文
Version Type	VoR
URL	https://doi.org/10.18910/96226
rights	
Note	

様式３

論　文　内　容　の　要　旨

氏　名　（　李　智　Zhi　LI　）	
論文題名	Research on Improving Online Recommendations Involving Sparse Interaction Data （疎な行動データにおけるオンライン推薦の向上に関する研究）

論文内容の要旨

With the explosive growth of the number of online services and the items (i.e., products) they provide, it becomes extremely time-consuming for users to explore their interested products from the countless available ones. Recommender systems (RSs) are developed to generate personal recommendations for users by modeling their preferences towards items, where the recommended items are empirically shown to be superior in satisfying users' needs, resulting in substantial value for service providers as well as users. The tremendous commercial profits of RSs motivate flourishing research that strives to improve the performances of RSs in the past few decades.

Thanks to the rapid development of deep learning models, deep learning-based RSs have demonstrated their effectiveness in accurately modeling users' preferences and also practically boosting commercial profits over the past few years. Such RSs, however, require a large amount of users' item-interacted data to learn their parameters due to the data-hungry nature of deep learning. This requirement significantly limits the practical implementation of deep learning-based RSs, especially in cases where there is only a limited number of interactions for each user or each item, a scenario often referred to as sparse interaction data (i.e., the sparse scenario). As the mainstream solution to tackle the issue of sparse interaction data, transfer learning-based RSs are widely studied. Transfer learning can be applied inner one service or across two services. The former transfers knowledge from users or items that have relatively sufficient interactions to the ones with sparse interactions. The latter transfers knowledge from an external service with sufficient data to facilitate the learning of user preferences in the sparse local service.

However, user preferences vary from service to service and even differ in different periods of the same service. Current transfer learning-based approaches ignore these changing user preferences, leading to a significantly limited performance on recommendations. Since sparse scenarios and changing user preferences are ubiquitous in real-world services, handling changing user preferences in sparse scenarios becomes an essential and critical demand when developing RSs for practical services. By capturing such preferences accurately, RSs are expected to further boost business revenue and better satisfy users' needs by providing high-preference recommendations.

Based on the above observations, our goal in this thesis is to improve recommendation performances by developing RSs that can handle changing user preferences and work well in sparse scenarios. Recall that user preferences change in different periods of the same service and vary from service to service, this thesis studies both types of changes in user preferences. For the first type of changes, we study the RSs for flash sale e-commerce services, which is a typical scenario that user preferences change very frequently (i.e., every few weeks). The challenge of handling changing user preferences in such scenario have not been discussed in the literature yet. If our RSs can effectively handle the frequently changed user preferences in flash sale scenario, it is reasonable to expect that our RSs can work well in other sparse scenarios with changing user preferences. For the scenarios that user preferences are different between two services, we studied RSs that improve

recommendations in the local sparse service by effectively leveraging data from external services to boost the poor recommendation performance in previous RSs. External data include public data and private commercial data, this thesis, hence, studies how to effectively leverage both types of external data. We elaborate on the details of the three practical scenarios we studied and introduce the challenges still required to be tackled in these scenarios below.

i) In a flash sale e-commerce scenario, such as Amazon daily deals, available and discounted items are periodically changing based on the current sale strategies. Users are attracted by the high discounts and show period-specific interactions. The frequent changes of interacted items provoke sparse interactions. Periodically changed interactions also represent users' period-specific preferences. Considering the periodic changes in user preferences, it is reasonable to learn users' preferences in a period by leveraging only the interactions in that period. However, previous works simply combine users' interaction data in all the past periods and learn only a uniform preference, making such works difficult to perform well in different sale periods. Therefore, modeling period-specific user preferences with only limited interactions is still a challenge in flash sale e-commerce.

ii) The second scenario is to effectively leverage rich public data to enhance local recommendations. In real-world online services, most users interact with only a small number of items (sparse interactions), occurring not only in new services and small companies but also in established services and large companies. Due to the data-hungry nature of deep learning models, services with sparse interactions have a practical demand for leveraging rich public data to enhance their recommendations, where the public data are released by service providers for the purpose of organizing competitions and promoting researches, such as the Amazon Review data released by Amazon Inc.. Cross-domain recommender systems (CDRSs) are a promising approach that can facilitate the recommendations in the sparse target domain by transferring knowledge from the interaction data in an auxiliary source domain, where the sparse target domain and the source domain are, respectively, the local service and the external service that provides public data.
However, existing CDRSs cannot effectively leverage rich public data due to a so-called negative transfer issue which generally leads to worse recommendation accuracy compared to single-domain methods. This issue is caused by the misleading of the source interactions that present users' unique interests within the source domain, which is also called domain-specific preferences. Existing CDRSs solve this issue by merging individual interactions across domains. With the individually merged interactions, these systems can assign a minimal weight to such source interactions, thereby, alleviating their impact. However, merging individual interactions between domains requires these methods to identify the users having interactions in both domains. These users are commonly referred to as domain-shared users. Identifying such users has to match individuals (i.e. user matching) through shared user information, e.g., user profiles or other characteristics that facilitate individual identification, between domains. In public data, the user characteristics that can identify individuals are definitely unavailable due to the privacy issue. Therefore, it is impractical to identify domain-shared users in the case of leveraging public data. In other words, effectively leveraging public data to improve local recommendations is still a challenge because of negative transfer issue.

iii) The third practical scenario is to effectively leverage commercial data to enhance local recommendations. Public data may suffer from incompleteness or outdated information and may lack quality and accuracy due to limited support. Consequently, it is logical to utilize commercial data to enhance local recommendations by partnering with commercial corporations. However, addressing the negative transfer issue is more challenging in this scenario. While directly merging individual interactions is not feasible in the second scenario, it is possible to merge interactions from the source domain and the target domain at a cluster level by jointly clustering items. Then, the weighting technology also can assign minimal scores for source interactions that present domain-specific preferences at a cluster level to handle the negative transfer issue. However, in the third scenario, merging interactions (i.e. sharing the information on interactions), even cluster-level interactions, is prohibited, as it may disclose personal preferences, behaviors, or patterns of individuals without their explicit consent to other companies. Due to the prohibition of merging interactions, it is impossible to explicitly identify and directly operate on source interactions that present domain-specific preferences when learning user preferences in the target domain. As a result, another type of entity is required to bridge domains and transfer knowledge, e.g., item clusters that contain items having similar textual features from both the source and the target domains. Without sharing interactions, the impact of source interactions has to be implicitly alleviated based on such entities. Different from the explicit and direct methods working on source interactions, it is hard to design implicit methods due to the difficulty in distinguishing how the operation on entities affects the impact of source interactions. This gives the reason why addressing the negative transfer issue is more challenging in this scenario.

In this thesis, we focus on improving recommendations with transfer learning-based methods and tackle the above-mentioned challenges. This thesis consists of five chapters. We introduce the research background and issues for improving recommendations in Chapter 1. In Chapter 2, we address the challenge of modeling period-specific preferences in flash sale e-commerce. In Chapter 3, we address the challenge of leveraging public data to improve target recommendations. In Chapter 4, we address the challenge of boosting target recommendations with external commercial data. Finally, in Chapter 5, we summarize this thesis and discuss our future work.

In Chapter 2, we introduce a novel meta-learning-based RS to model users' period-specific preferences. Moreover, we propose a new hierarchical meta-training algorithm to guide the learning of our recommendation model via user- and period-specific gradients. With the guidance of these gradients, the model can learn user- and period-share prior knowledge, supporting modeling users' period-specific preferences with only limited interactions and several updating steps. By doing so, our RS can quickly adapt to the recommendation tasks for new flash sale periods. Our experimental result on a real-world flash sale e-commerce dataset shows that our proposal remarkably outperforms current state-of-the-art methods. This result also demonstrates the effectiveness of our proposal in modeling period-specific preferences.

In Chapter 3, to effectively leverage public data, we propose a novel CDRS that requires no domain-shared users and can handle the negative transfer issue. To remove the requirement of domain-shared users, we merge cluster-level interactions between domains by jointly clustering items from both domains. Besides, to handle the negative transfer issue, we construct a cross-domain interaction graph to transfer cluster-level interaction between domains, and propose a new debiasing graph convolutional layer to weight source cluster-level interactions. Constructing this graph only requires jointly clustering the items in the source domains and the ones in the target domains based on their semantic embeddings extracted from their text data, i.e., titles of products. This removes the requirement of domain-shared users to bridge the two domains. Our debiasing layer handles the negative transfer issue by adaptively weighting source interactions in the cross-domain graph to alleviate the impact of the source interactions that present users' unique interests within the source domain. Our experimental results on three public datasets and a pair of private datasets verify the advantages of our method over state-of-the-art models in terms of cross-domain recommendations and handling the negative transfer issue.

Different from Chapter 3, utilizing external commercial data requires no user interaction merging between domains due to privacy concerns. We, hence, develop a novel CDRS that does not merge any interactions between domains and can handle the more challenging negative transfer issue in Chapter 4. To remove the mergence of interactions, our CDRS constructs a similarity-based cross-domain graph to bridge domains and transfer knowledge. To handle the more challenging negative transfer issue, our CDRS transfers knowledge based on the cross-domain graph by focusing on the items that are from different domains and have similar textual features, i.e., similar titles. This alleviates the impact of the source items with irrelevant titles which cause the

negative transfer issue. Our experimental results on real-world datasets show that our method significantly outperforms state-of-the-art cross-domain comparisons, which indicates the effectiveness of our proposal in the negative transfer issue.

様式 7

論文審査の結果の要旨及び担当者

氏　名	(Zhi Li)		
論文審査担当者		(職)	氏　名
	主 査	教授	原　隆浩
	副 査	教授	鬼塚　真
	副 査	教授	村田　忠彦
	副 査	教授	松下　康之
	副 査	教授	中野　珠実

論文審査の結果の要旨

オンラインサービスおよび利用（購買）可能なアイテムの増加に伴い、ユーザが自身の興味・趣向に合うアイテムを検索するための時間が非常に長くなりがちである。この問題を解決するシステムとして、推薦システムが多くのサービスで用いられており、推薦システムはユーザそれぞれに個別化した推薦リストを提供する。このシステムにより、ユーザは全てのアイテムを閲覧することなく自身の興味・趣向に合うアイテムを発見できることが多くの実用・研究結果により示されている。この推薦システムで重要な点は、推薦リストの精度であり、例えば精度が低い（興味に全く合わない）リストを提供すると、そのサービスから離脱するユーザは増え、精度が高い場合は利益向上に繋がる。近年の機械学習技術の高度化により、ユーザとアイテムの特徴を捉え、精度の高い推薦リストを提供する深層学習を用いた推薦システムに注目が集まっている。深層学習を用いた推薦システムは、過去のユーザとアイテムのインタラクションおよびアイテムに付随するドキュメント等の周辺の情報を用いて未知のユーザとアイテムのインタラクションを推論する。ここで、過去のインタラクションの数が十分に多い場合は学習精度が高くなることは自明であり、逆に数が少ない場合、学習結果が汎化状態にはならず、精度の高い推薦リストを提供できないという課題がある。本論文では、この疎なインタラクション環境における推薦リストの精度向上に取り組んでおり、以下の主要な結果を示した。

(1) フラッシュセールのように、アイテム（およびセール率）が時々刻々と変化する環境では、ユーザの取る行動（興味）もその時に利用可能なアイテムに大きく依存する。そのため、この環境では疎なインタラクションとなりがちであり、ユーザの興味も時間によって変化する。この環境において、ユーザのある時刻における興味を高精度で表す学習技術を提案し、新たなアイテムが現れた場合でも高速にユーザのモデルを更新できることを示した。

(2) 次に、公開データを用いた転移学習を利用することにより、疎なドメインでの推薦精度向上を考えた。転移学習では一般的にドメイン間で共通のユーザが必要だが、公開データに共通するユーザは存在しない。この研究では、共通ユーザを使うことなく公開データから有用な知識を転移することに成功している。既存技術では別のデータを用いることでノイズを増やし、推薦精度の低下を招いているところ、提案技術は精度を向上できることを実験により示した。

(3) 最後に、外部の商用データを用いた転移学習を考えた。外部の（非公開）データを利用する場合、プライバシーの観点からそのドメインのユーザ情報を流用できないという制約がある。そこで、セマンティックレベルでの知識のみを用いることで、外部データにおけるユーザの情報を使うことなく効果的な転移学習を実現した。また、知識蒸留を用いることにより、ドメイン間で共用できる趣向を抽出し、推薦精度の向上を達成した。

以上のように、本論文は疎なインタラクション環境における推薦システムに関する先駆的な研究として、情報科学に寄与するところが大きい。よって本論文は博士（情報科学）の学位論文として価値のあるものと認める。