



A Worrying Reproducibility Study of Intent-Aware Recommendation Models

Faisal Shehzad
faisal.shehzad@aau.at
University of Klagenfurt
Klagenfurt, Austria

Maurizio Ferrari Dacrema
maurizio.ferrari@polimi.it
Politecnico di Milano
Milano, Italy

Dietmar Jannach
dietmar.jannach@aau.at
University of Klagenfurt
Klagenfurt, Austria

Abstract

Lately, we have observed a growing interest in *intent-aware* recommender systems (IARS). The promise of such systems is that they are capable of generating better recommendations by predicting and considering the underlying motivations and short-term goals of consumers. From a technical perspective, various sophisticated neural models were recently proposed in this emerging and promising area. In the broader context of complex neural recommendation models, a growing number of research works unfortunately indicates that (i) reproducing such works is often difficult and (ii) that the true benefits of such models may be limited in reality, e.g., because the reported improvements were obtained through comparisons with untuned or weak baselines. In this work, we investigate if recent research in IARS is similarly affected by such problems. Specifically, we tried to reproduce five contemporary IARS models that were published in top-level outlets, and we benchmarked them against a number of traditional non-neural recommendation models. In two of the cases, running the provided code with the optimal hyperparameters reported in the paper did not yield the results reported in the paper. Worryingly, we find that all examined IARS approaches are consistently outperformed by at least one traditional model. These findings point to sustained methodological issues and to a pressing need for more rigorous scholarly practices.

CCS Concepts

• Information systems → Recommender systems.

Keywords

Recommender Systems, Intent-Awareness, Reproducibility

ACM Reference Format:

Faisal Shehzad, Maurizio Ferrari Dacrema, and Dietmar Jannach. 2025. A Worrying Reproducibility Study of Intent-Aware Recommendation Models. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3726302.3730307>

1 Introduction

Intent-aware recommender systems (IARS) are designed to capture the underlying motivations and goals of users when they use an online service, with the aim of leveraging this information for

improved item suggestions [14, 34]. In the e-commerce domain, for example, consumers may have different intents when visiting an online shop. A consumer’s goal might be to understand the space of options in a certain category, to compare alternatives, to search for ideas, or just hedonic browsing [42]. Depending on the current intent, different sets of recommendations may be appropriate. Similarly, in the music domain, users might use a streaming app for different reasons, e.g., to discover new artists, to listen to their favorite tracks, or to find some music to play in a certain situation [15, 31]. Generally, adapting the recommendations to the assumed user intent has a strong value potential, as evidenced in the field study in the e-commerce domain reported in [42].

Technically, a number of algorithmic approaches to implement intent-awareness in recommender systems have been proposed over the last decade. Early approaches mainly aimed at the diversification of the recommendations, with the goal of covering multiple possible intents through a single recommendation list [48, 55]. Some other works assume the existence of a fixed set of pre-defined intents in a given domain [31, 42]. Most recently, however, a multitude of *latent intent* modeling approaches were proposed. Such approaches are commonly based on sophisticated deep learning architectures which comprise specific network components designed to model the hidden relationships between observed user interactions and the underlying latent intents. Corresponding proposals have been made both for traditional *top-n* recommendation scenarios, e.g., [29, 53, 54], as well as for sequential recommendation settings, e.g., [8, 25, 47], see [14] for a recent survey on IARS.

Despite the increased sophistication of IARS models and recommendation models in general, and despite the substantial computational complexity and carbon footprint of such models [43, 49], a number of recent studies have indicated that more complex models are not necessarily more effective than longer-existing simpler ones. Ferrari Dacrema et al. [10, 11] for example found that in almost all cases they examined, a recent neural recommendation model was outperformed by traditional algorithms based, e.g., on matrix factorization.¹ Similar observations of the competitive performance of traditional learning techniques were later reported also in [1, 32, 36, 37]. For the case of sequential recommendation settings, related reports on the effectiveness of simple models—also in practical settings—can be found in [16, 19, 21, 28, 41]. Overall, while numerous papers claim to advance the state-of-the-art, the achieved progress in this area may actually be very limited. Similar phenomena where “*improvements don’t add up*” were reported earlier in the field of information retrieval [3, 26] and other application areas of machine learning [30].

¹Notably, the examined cases were all published in top-level conferences or journals.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '25, Padua, Italy*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1592-1/2025/07
<https://doi.org/10.1145/3726302.3730307>

A number of reasons contribute to the problem. First, the level of reproducibility is generally limited, and still today many papers are published without sharing code and data. Second, different methodological issues and certain bad practices such as using sampled metrics may lead to unreliable results [20]. Finally, earlier reproducibility studies mentioned above indicate that researchers in many cases do not properly tune the hyperparameters of the baseline models or only consider baselines that are generally weak.

In this paper, we investigate if similar phenomena of “*phantom progress*” [11] can be observed for the emerging and promising area of Intent-Aware Recommender Systems. To that purpose, we have tried to reproduce five recent neural IARS models—all of them published in highly-ranked outlets—and benchmarked them against traditional *top-n* recommendation models. Specifically, as done in [11], we benchmark the considered IARS models using the code, data, and experimental configuration used in the original papers. Our findings are quite worrying, indicating that the mentioned methodological problems in terms of baseline selection and tuning still persist in the community. In all five cases we found that complex IARS models are *not* outperforming existing non-neural techniques if these are properly tuned. Thus, additional measures seem urgently required to improve the scientific rigor in our community.

The paper is organized as follows. Next, in Section 2, we outline our research methodology. The results of our experiments are then described in Section 3. The paper ends with a discussion of implications of our findings.

2 Methodology

2.1 Identification of Papers

In order to avoid cherry-picking individual models for inclusion in our study, we followed a structured approach similar to [11]. First, we queried Google Scholar and IEEE Xplore for papers containing the terms ‘intent’ or ‘intent awareness’ along with the keywords ‘recommend’ in the title, abstract, or introduction. We reviewed the retrieved papers for relevance and identified more papers by examining each paper’s references. This process left us with 88 articles.² To further focus our analysis, we only retained works that fulfilled the following criteria: (i) the work proposes a model for *top-n* collaborative filtering-based recommendation, (ii) the work was published recently, i.e., during the last four years, (iii) it was published in a CORE A* conference or in a journal with an Impact Factor ≥ 5 . Applying these criteria left us with 13 papers.

In the next phase, our goal was to reproduce the results reported in these papers using the artifacts (code and data) that were used in the original papers. We adopt ACM’s terminology, where reproducibility means that “*that an independent group can obtain the same result using the author’s own artifacts*.”³ If the authors did not share the artifacts needed for reproducibility, we contacted the authors via email. After waiting for 30 days, we considered an article to be a candidate for reproduction if it fulfilled the following criteria:

- A working version of the code is available or requires minimal modification to make it runnable.

- At least one public dataset is used and the authors either share the train-test splits or provide sufficient information in the paper to create the train-test splits from the given dataset.

From the 13 papers, we could get the provided code running five with reasonable effort,⁴ leading to a percentage of executable artifacts of 38.46%. This rate is almost identical to the one observed in a related previous study from 2019 [11], indicating that the level of reproducibility may not have increased much since then. Table 1 shows the works providing runnable code and works without code or non-functioning code.

Table 1: Papers with runnable code vs. papers with missing artifacts or non-runnable code.

Runnable	Missing or non-runnable
[33],[54],[35],[59],[53]	[13],[56],[23],[24],[58],[29],[27],[57]

2.2 Evaluation Methodology

Benchmarking Approach. In the literature, two approaches to benchmarking algorithms can be found, each of them having specific advantages. One approach is to integrate all models to be compared in a common evaluation framework, as done, e.g., in [1] or [28]. Such an approach ensures that all algorithms are evaluated under identical conditions, e.g., same data splitting, same metrics. Furthermore, such a setting allows us to assess the effectiveness of different models across a collection of datasets beyond the ones that were used for the original paper. The alternative is to evaluate and benchmark existing models using the exact same configuration as was used in the original paper, as done in [11]. This approach should allow us to exactly reproduce the results that were reported in the original paper. Since the goal of our present paper is to reproduce the original papers, we adopt this latter approach. Technically, this in particular also means that we rely on the code provided by the authors for the evaluation, i.e., no changes to the evaluation methodology are introduced in our experiments. The downside of this approach is that the reproduced models usually cannot be compared to each other, as there are always small differences in the evaluation methodology across papers. Such a comparison of IARS models is however not the goal of this paper.

Instead, the goal of our work is to understand (i) if published results can be reproduced based on the artifacts provided by authors and (ii) to what extent recent IARS models are favorable over traditional models. We therefore selected a number of baseline models of different types for our study. The selection is inspired by previous benchmark studies [1, 11]. The following baselines were considered.

TopPop: A non-personalized algorithm that recommends the most popular items to everyone. The popularity of items is calculated based on the number of interactions in the training data.

²The list of articles can be found online. A similar number of papers on the topic of IARS was identified in a recent survey [14].

³<https://www.acm.org/publications/policies/artifact-review-and-badging-current>.

⁴We refer readers to the discussions in [11], where the authors found that shared code repositories are often incomplete, miss relevant parts, or only contain a skeleton of the model.

ItemKNN: A traditional collaborative filtering algorithm based on item similarities [50]. We use cosine similarity in our study. The hyperparameters considered in our study are the neighborhood size, the item weighting strategy, and a shrinking weight for items with fewer interactions [17, 51].

UserKNN: Like ItemKNN, but considering user similarities instead of item similarities [38].

P³ α : A bipartite graph based algorithm that implements a walk between users and items [5]. Previous work indicates that the approach can be very effective despite its simple nature [11]. The model’s hyperparameter are a similarity scaling coefficient α and the neighborhood size.

RP³ β : An advanced version of P³ α , using an additional hyperparameter β to account for item popularity.

EASE^R: An “embarrassingly shallow” linear model designed for sparse data [44]. The model has only one hyperparameter and has shown to lead to very competitive performance in previous studies [1, 10].

We systematically tuned the hyperparameters for all baselines for all experimental configurations. Technically, we used Bayesian optimization [10, 11] and explored 35 to 50 cases, depending on the complexity of the model. The first five cases were selected randomly to reduce the chances of premature convergence. We share all code and data used in our experiment, as well as hyperparameter ranges and best hyperparameters for each dataset online.⁵

3 Results

In this section, we discuss the results for the five IARS models with runnable code ordered by publication date.

3.1 Disentangled Graph Collaborative Filtering (DGCF)

Summary. The DGCF model [54] was presented at SIGIR ’20. The authors highlight that existing models treat user-item interactions as isolated or uniform, neglecting a user’s diverse intentions on items, such as shopping for others, scrolling through social media to pass the time, etc. Such uniform approach is assumed to lead to a suboptimal representation of user-item interaction data. The proposed DGCF model consists of two key components to obtain disentanglement: (i) a graph disentangling component, which converts user/item embedding into chunks and attaches an intent to each chunk; (ii) an independence component to enhance the independence between intents.

Datasets and Code. Experiments were conducted with datasets called Gowalla, Amazon Book and Yelp2018. For all datasets, only users and items with at least ten interactions were retained. For each dataset, an 80/20 rule was adopted for data splitting. For each user, 80% of the interactions were used for training, the rest was used for testing⁶ and 10% of the interactions of the training data were used for hyperparameter tuning. The authors shared the train-test splits, which were used to obtain the results reported in the published article.

⁵<https://faisalse.github.io/IntentAwareRS/>

⁶We note that user-wise splitting may in general lead to data leakage, as future interactions of other users can inform the predictions of the model of a given user.

Original Evaluation. The authors considered five baselines, including a traditional matrix factorization technique, three graph-based models, and a previous disentanglement method called MacridVAE. Recall@K and NDCG@K are used as performance measures. In terms of hyperparameter tuning, the authors mention that for three models they copied the *optimal* hyperparameters from an own previous paper [52]. Looking up the reference, we find that hyperparameters were apparently not fully tuned in [52], because a central parameter, the embedding size, was not tuned but set to a fixed size. The same fixed size of 64 is also used for two of the graph-based models. Other hyperparameters like the learning rate are tuned. For MacridVAE, embedding sizes were searched between 55 and 75, even though the original paper used an embedding size of 100 (for a different dataset). No code is shared for the baselines, data pre-processing, and hyperparameter tuning.

Our Results. In Table 2 and in subsequent tables we report the results using the metrics and cutoffs used in the original papers. We print numbers in bold when they represent the best results or when a baseline model performs equally or better than the proposed IARS model. R stands for Recall, and N for NDCG.

Table 2: Experimental results for DGCF model. DGCF (Paper) refers to the numbers reported in the original paper, DGCF (Code) are the results obtained by running the provided artifacts, DGCF ([2]) refers to results obtained in a related evaluation by other researchers.

Methods	Gowalla		Yelp 2018		Amazon Book	
	R@20	N@20	R@20	N@20	R@20	N@20
TopPop	0.013	0.031	0.012	0.010	0.002	0.004
ItemkNN	0.156	0.123	0.064	0.054	0.056	0.045
UserkNN	0.155	0.128	0.061	0.050	0.051	0.041
P ³ α	0.175	0.146	0.064	0.052	0.056	0.043
RP ³ β	0.183	0.152	0.067	0.056	0.068	0.056
EASE ^R	0.163	0.136	0.064	0.052	0.071	0.056
DGCF (Paper)	0.179	0.152	0.064	0.052	0.039	0.030
DGCF (Code)	0.133	0.109	0.037	0.028	0.011	0.017
DGCF ([2])	0.173	0.147	0.062	0.050	0.038	0.029

In the table, we report three sets of values for DGCF. *DGCF (Paper)* shows the numbers reported in the original paper. *DGCF (Code)* refers to the numbers when we executed the author-provided code. Only one set of hyperparameters is provided by the authors, which we used for all datasets. The obtained numbers are much lower than the results reported in the original paper.⁷ *DGCF ([2])* refers to the performance results obtained by the authors of [2], using a new PyTorch implementation of DGCF and systematic hyperparameter tuning. We used the best hyperparameters reported in [2] with the original code by the authors, and we then obtained similar results. Comparing the optimal hyperparameters used in [2] and the author-provided ones, we particularly find that using a larger batch size (2000 vs. 1024) may help to improve the performance of the DGCF model. Thus, with proper hyperparameter tuning, the

⁷We contacted the authors on this issue without success.

reported results of DGCF can be roughly reproduced, even though the reproduced results are a bit lower.⁸

However, as also reported in [2], the simple RP_3^β method outperforms DGCF on all metrics and on all datasets. This is true also when comparing the results to the numbers reported in the paper. On the Yelp dataset, all compared methods lead to rather similar results. On the Amazon Book dataset, DGCF is not performing well, and it is actually outperformed by all simple models. We could not run experiments with $EASE^R$ for the Amazon Book dataset due to hardware constraints. We therefore report the results obtained in [60] instead, which suggest that $EASE^R$ is performing best on this dataset. The experiments in [60] are based on the same train-test splits as used in the original DGCF paper. The authors of [60] however do not rely on sampled metrics [20].

3.2 Knowledge Graph-based Intent Network (KGIN)

Summary. The KGIN approach [53] was presented at TheWebConf '21. KGIN is a recommendation model that uses a knowledge graph which models the user-item interactions from two perspectives: (i) identifying the intent behind each user-item interaction at a higher level of granularity and (ii) proposing a path-aware aggregation mechanism to integrate information from multiple hops and promote the independence of intents for better prediction and interpretability.

Datasets and Code. Three datasets—Amazon Book, Last-FM and Alibaba-iFashion—were used for performance evaluation. Users and items with less than ten interactions were said to be discarded from the data, but our analysis revealed that users and items with fewer interactions exist in the dataset.⁹ An 80/20 data split was applied, and the data splits are shared for all datasets. A leave-one-out splitting strategy was employed to obtain validation data from the shared training data.

Original Evaluation. The authors compare KGIN to six baseline models. One of them is a matrix factorization technique, and the five others are based on Graph Neural Networks and Knowledge Graphs. The embedding size for all models was set to 64 “for a fair comparison”. Again, this is methodologically wrong as the embedding size is a hyperparameter to tune, which can drastically impact model performance [41]. Furthermore, for some baselines, fixed values, e.g., batch size and neighborhood size are used. No code is shared for the baselines, data pre-processing, and hyperparameter tuning.

Our Results. The authors provide scripts to train and evaluate the models for all three datasets, stating that these scripts use the best hyperparameters. The results of running the scripts and the performance results for our baselines are shown in Table 3.

With *KGIN (Paper)*, we again refer to the results in the paper, and *KGIN (Code)* are the numbers obtained by running the scripts by the authors which are said to lead to the best performance. The obtained numbers are often substantially lower, with the exception for Recall for the Last-FM dataset, which is more than twice as high

⁸We print the numbers reported in the paper in gray and base our comparison on the results that could actually be reproduced.

⁹Details can be found in the online material.

Table 3: Experimental results for the KGIN model

Methods	Amazon Book		Alibaba-iFashion	
	R@20	N@20	R@20	N@20
TopPop	0.028	0.012	0.031	0.016
ItemkNN	0.164	0.097	0.122	0.080
UserkNN	0.159	0.091	0.101	0.065
P_3^α	0.166	0.094	0.069	0.046
RP_3^β	0.144	0.084	0.110	0.073
$EASE^R$	0.161	0.092	0.120	0.076
KGIN (Paper)	0.168	0.091	0.114	0.071
KGIN (Code)	0.143	0.064	0.102	0.058
	Last-FM (leaking)		Last-FM*	
TopPop	0.022	0.019	0.015	0.009
ItemkNN	0.182	0.180	0.234	0.202
UserkNN	0.130	0.143	0.200	0.168
P_3^α	0.194	0.196	0.239	0.213
RP_3^β	0.189	0.187	0.253	0.233
$EASE^R$	0.185	0.184	0.227	0.191
KGIN (Paper)	0.097	0.0848	–	–
KGIN (Code)	0.202	0.173	0.163	0.092

as in the paper. We contacted the authors on this issue without success, leading to a situation where the reported results are not reproducible. Given the surprising results for the Last-FM dataset, we investigated the provided train-test splits to find a significant data leakage issue, affecting 17,290 out of 23,566 users.¹⁰ We therefore created the union of the provided train-test splits, i.e., removed the duplicates, and created our own 80/20 split, leading to what we denote the Last-FM* dataset. The results for the different models on this dataset are also shown in Table 3.¹¹

Comparing the performance of our baselines and KGIN, we find that many of our baselines outperform KGIN on all three (correct) datasets. In particular the UserKNN and ItemKNN work very well on all datasets. We note that for all datasets and measurements except for Recall@20 for the Amazon Book dataset, at least one baseline model leads to better results than the non-reproducible higher ones reported in the paper.

3.3 Intent Disentanglement and Feature Self-supervision (IDS4NR)

Summary. The IDS4NR model [33] was published in IEEE TKDE in 2022. The authors present a mechanism that optimizes accuracy and novelty by disentangling the users’ intent for popular items (conformity bias) and their personal interests. To model the relationships between users and items different traditional recommendation models can be used as a *backbone*. Three backbones are considered in the paper, matrix factorization (named LFM), neural collaborative filtering (NCF), and collaborative metric learning (CML).

Datasets and Code. The authors use the MovieLens-100k dataset as well as the Amazon Music and Amazon Beauty datasets for their

¹⁰This problem was also identified recently in [7].

¹¹We used the provided hyperparameters for Last-FM for the Last-FM* dataset.

experiments. The datasets are augmented with side information. Only items and users with at least five interactions are retained. Training and test splits are provided. No code for data preprocessing, baselines, and tuning is shared.

Original Evaluation. The evaluation considers Recall as an accuracy metric, as well as Coverage and Novelty as beyond-accuracy metric. An F-score metric is computed as the harmonic mean of the three metrics. One set of hyperparameters for the proposed model is shared. The optimization goal is not reported. The embedding dimension for all models is set to 50, which, as mentioned, is a major methodological problem. Hyperparameters for the baselines are said to be taken from the original papers or fine-tuned—without further specification—when they were unknown.

Our Results. We could run the code for two of three versions of the proposed model (when using LFM and NCF as backbones) using the shared artifacts and the provided hyperparameters. Running the code for the CML backbone required the specification of a number of parameters that were not shared in the paper or in the provided code. Inquiries to the authors were without success.

The results of executing the provided code and the results for the baselines are shown in Table 4. We observe that the numbers reported in the paper largely deviate from the numbers we obtain when running the provided code. In terms of accuracy measures, the obtained numbers are about 70% higher than what is reported in the paper for the MovieLens dataset. For the other two datasets, the Recall values are in contrast substantially lower. Thus, with the provided artifacts, we could not reproduce the results. In any case, independent of which performance results for IDSNR we consider, i.e., from the paper or ours, all our traditional and simpler baselines lead to higher Recall values and outperform IDSNR. The same holds for the F-measure that combines Recall with Coverage and Novelty.

For Coverage and Novelty, we also observe that our results often deviate largely from the numbers reported in the paper. IDSNR is generally very strong in terms of these metrics, compared to the baselines. However, considering the lower F-measure value, it becomes clear that increasing Coverage and Novelty comes at a too high price in terms of accuracy.

3.4 Disentangled Contrastive Collaborative Filtering (DCCF)

Summary. The DCCF model [35] was presented at SIGIR '23. The model implements intent disentanglement combined with adaptive self-supervised augmentation. The latter component shall help the model to be more robust against noise.

Datasets and Code. The authors evaluated their model on the Gowalla, Amazon Book, and Tmall datasets. Differently from other related works like the DGCF model discussed above, the datasets were apparently not preprocessed, or preprocessed differently than in previous works. This can make a comparison with previous works difficult.¹² The authors however share the training and test data that were used in the experiments. No information about

the splitting procedure are provided. No code for preprocessing, baselines, or hyperparameter tuning is shared.

Original Evaluation. The authors benchmark their method against as many as 12 baseline models from different families. Hyperparameter ranges for the proposed model are described. For the baseline models, a fixed embedding size of 32 and batch size of 1024 are used, again representing a major methodological issue. No information about the selection of these values is provided. Additional hyperparameter values for the considered baselines and the target metric for tuning are not reported either. The authors use Recall and NDCG as evaluation measures. The choice of the cutoff length of 40 is unusual and not explained in the paper.

Our Results. We could run the provided code, leading to the results shown in Table 5. We observe that the obtained numbers are very close to those reported in the paper, sometimes a bit higher, sometimes a bit lower. We therefore consider this work reproducible. We however find that all of the considered baseline models outperform the proposed DCCF model on all datasets.¹³ Among our baselines, no unique winner can be identified, and the relative performance of different algorithms depends on dataset characteristics.

3.5 Bilateral Intent-guided Graph Collaborative Filtering (BIGCF)

Summary. The BIGCF approach [59] was presented at SIGIR '24. It aims at separating a user's individual intent, i.e., personal preferences, from collective intent, which the authors describe as overall awareness. Conceptually, the work is thus related to the IDSNR approach discussed above. Technically, the model is based on Graph Convolution Networks and Graph Contrastive Learning.

Datasets and Code. Datasets from Gowalla, Amazon Book, and Tmall datasets are used for evaluation. The same version of the datasets is used as for the DCCF model discussed above. The train-test splits used for the experiments are shared by the authors. No further information is provided how the splits were exactly created and which parts of the data were used to tune the hyperparameters. The target metric for tuning is not reported either. No code for preprocessing, baselines, or hyperparameter tuning is shared.

Original Evaluation. The authors closely follow the evaluation setup that was used in [35] to evaluate the DCCF approach. Specifically, they also rely on the problematic approach to fix central hyperparameters like the embedding size to a (small) fixed value for all compared models. For the hyperparameters of the proposed model, the search ranges for the hyperparameters are reported. While the evaluation involves three datasets, only one set of optimal values is highlighted in the paper. In the shared code repository, scripts are provided to run the model on three datasets. For these scripts, specific values for the number of epochs and the regularization factor are provided. It remains unclear how the number of epochs was chosen, as no use of a validation set is mentioned.¹⁴

For the compared methods, the authors mention that they used hyperparameters from the original papers or used grid search. No

¹²Often, authors also do not share which *version* of a dataset is used. There are, for example, now three versions of the Amazon Book dataset available by now at https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/.

¹³Again, we could not run the EASE^R model on the largest Amazon Book dataset due to hardware constraints.

¹⁴The study in [46] indicates that a substantial fraction of the literature (37% of the examined papers) were tuned on test data.

Table 4: Experimental results for the IDS4NR model

Methods	MovieLens							
	R@5	R@10	Cov@5	Cov@10	Nov@5	Nov@10	F1@5	F1@10
TopPop	0.052	0.087	0.025	0.040	0.023	0.046	0.049	0.092
ItemKNN	0.090	0.141	0.057	0.080	0.024	0.048	0.085	0.139
UserKNN	0.073	0.110	0.036	0.055	0.023	0.047	0.061	0.110
$P_3\alpha$	0.105	0.162	0.104	0.153	0.024	0.049	0.136	0.224
$RP_3\beta$	0.106	0.174	0.160	0.220	0.025	0.050	0.197	0.303
EASE ^R	0.119	0.185	0.179	0.250	0.025	0.050	0.217	0.338
IDNSR_LFM (Code)	0.080	0.138	0.367	0.486	0.030	0.062	0.431	0.637
IDNSR_LFM (Paper)	0.087	0.143	0.441	0.572	0.3870	0.390	0.049	0.087
IDSNR_NCF (Code)	0.084	0.138	0.390	0.539	0.030	0.060	0.454	0.696
IDSNR_NCF (Paper)	0.085	0.138	0.467	0.615	0.390	0.385	0.049	0.086
Amazon Music								
TopPop	0.020	0.033	0.004	0.007	0.011	0.023	0.015	0.031
ItemKNN	0.153	0.211	0.831	0.918	0.016	0.031	0.885	1.035
UserkNN	0.146	0.201	0.468	0.638	0.014	0.028	0.501	0.721
$P_3\alpha$	0.155	0.212	0.571	0.724	0.014	0.029	0.610	0.816
$RP_3\beta$	0.160	0.219	0.782	0.897	0.015	0.030	0.832	1.007
EASE ^R	0.158	0.214	0.553	0.721	0.014	0.029	0.591	0.812
IDNSR_LFM (Code)	0.061	0.106	0.392	0.524	0.014	0.029	0.423	0.599
IDNSR_LFM (Paper)	0.089	0.136	0.685	0.8324	0.500	0.498	0.071	0.115
IDSNR_NCF (Code)	0.067	0.112	0.800	0.985	0.017	0.033	0.856	1.117
IDSNR_NCF (Paper)	0.090	0.137	0.903	0.974	0.535	0.524	0.085	0.128
Amazon Beauty								
TopPop	0.006	0.010	0.002	0.003	0.007	0.015	0.009	0.018
ItemkNN	0.041	0.058	0.350	0.512	0.008	0.017	0.368	0.555
UserkNN	0.042	0.062	0.447	0.628	0.009	0.018	0.467	0.679
$P_3\alpha$	0.043	0.062	0.481	0.658	0.009	0.018	0.502	0.711
$RP_3\beta$	0.045	0.064	0.656	0.785	0.009	0.018	0.683	0.847
EASE ^R	0.050	0.067	0.456	0.639	0.009	0.018	0.477	0.691
IDNSR_LFM (Code)	0.019	0.032	0.255	0.403	0.011	0.022	0.274	0.453
IDNSR_LFM (Paper)	0.027	0.042	0.602	0.711	0.400	0.412	0.019	0.032
IDSNR_NCF (Code)	0.023	0.037	0.581	0.770	0.010	0.021	0.609	0.840
IDSNR_NCF (Paper)	0.025	0.041	0.902	0.973	0.538	0.536	0.025	0.041

further details are provided. The proposed BIGCF model is then benchmarked against fifteen baseline models from the literature. There is a significant overlap in terms of the baselines with respect to the evaluation of DCCF. Some baselines from the DCCF paper are however not considered, while others are added. For the baselines that are overlapping, the reported numbers are in most cases exactly the same as those reported in [35]. An exception is the NGCF model, where the results in the papers differ.¹⁵

Our Results. Running the scripts provided by the authors, we obtained the numbers reported in Table 6. We observe that the numbers reported in the paper are very close to those obtained when running the code provided to the authors, meaning that the work is reproducible. However, as for the DCCF model, which was evaluated using the same experimental configuration, we find that all simple baselines outperform the BIGCF model on all datasets

and metrics. The $RP_3\beta$ algorithm is consistently among the best-performing models.

Regarding the scripts provided by the authors, we note that the number of epochs is provided as a hyperparameter for model training for each dataset individually. It is however unclear how the optimal number of epochs was determined. In the provided code, we observe that the performance of the model on the test set is observed during training for each epoch, but this information is not used for testing. We re-ran the experiments with an *early-stopping* strategy. These additional experiments showed that early-stopping often leads to substantially lower accuracy values for the Gowalla and Amazon Book datasets. A proper choice of the number of epochs seems therefore essential for the performance of this model. The number of epochs should however be determined based on a validation dataset.

4 Discussion

In this section, we first summarize our observations regarding reproducibility of research published in this area. Then, we will discuss

¹⁵The authors of the BIGFC rely on the evaluation code used for the DCCF model. Similar results are thus generally expected.

Table 5: Experimental results for the DCCF model

Methods	Gowalla			
	R@20	R@40	N@20	N@40
TopPop	0.025	0.037	0.013	0.017
UserKNN	0.213	0.295	0.132	0.153
ItemKNN	0.235	0.321	0.146	0.168
$P_3\alpha$	0.231	0.317	0.141	0.164
$RP_3\beta$	0.245	0.337	0.149	0.173
EASE ^R	0.218	0.303	0.132	0.154
DCCF (Code)	0.189	0.266	0.112	0.132
DCCF (Paper)	0.187	0.264	0.112	0.132
Amazon Book				
TopPop	0.009	0.016	0.007	0.009
UserKNN	0.126	0.172	0.104	0.119
ItemKNN	0.152	0.205	0.127	0.144
$P_3\alpha$	0.133	0.188	0.107	0.124
$RP_3\beta$	0.141	0.190	0.117	0.133
EASE ^R	-	-	-	-
DCCF (Code)	0.092	0.138	0.071	0.086
DCCF (Paper)	0.088	0.134	0.068	0.082
Tmall				
TopPop	0.010	0.017	0.007	0.009
UserKNN	0.087	0.131	0.063	0.079
ItemKNN	0.110	0.163	0.080	0.099
$P_3\alpha$	0.109	0.162	0.080	0.099
$RP_3\beta$	0.116	0.173	0.084	0.104
EASE ^R	0.093	0.143	0.067	0.084
DCCF (Code)	0.064	0.101	0.045	0.058
DCCF (Paper)	0.066	0.104	0.046	0.059

methodological issues that we encountered in our reproducibility study, and who these issues may prevent us from making progress.

Reproducibility. Questions of reproducibility in recommender systems and in information retrieval have been reported again and again in the literature and have been around for at least fifteen years [3, 9, 18, 39]. More than ten years ago, related questions were discussed in a dedicated workshop at ACM RecSys. About five years ago, a reproducibility study involving deep learning models [11] revealed that (i) the general level of reproducibility of published research is low and (ii) that then-recent deep learning models were actually not as effective in offline studies as one might assume.

The observations made in our present study in the emerging and highly-promising recent area of intent-aware recommender systems are sobering. At least in the studied area of IARS, we found that only five of the thirteen identified relevant papers shared artifacts that are needed for reproducibility. Moreover, as our study showed, only for a subset of them it was then possible to actually reproduce the results that were reported in the papers based on the provided artifacts and documentation.

Overall, reproducing research in recommender systems therefore remains to be difficult. Besides the issue of lacking reproducibility artifacts and documentation of the proposed models, we found that

Table 6: Experimental results for the BIGCF model

Methods	Gowalla			
	R@20	R@40	N@20	N@40
TopPop	0.025	0.037	0.013	0.017
UserKNN	0.213	0.295	0.132	0.153
ItemKNN	0.235	0.321	0.146	0.168
$P_3\alpha$	0.231	0.317	0.141	0.164
$RP_3\beta$	0.245	0.337	0.149	0.173
EASE ^R	0.218	0.303	0.132	0.154
BIGCF (Code)	0.206	0.289	0.123	0.145
BIGCF (Paper)	0.208	0.288	0.124	0.145
Amazon Book				
TopPop	0.009	0.016	0.007	0.009
UserKNN	0.126	0.172	0.104	0.119
ItemKNN	0.152	0.205	0.127	0.144
$P_3\alpha$	0.133	0.188	0.107	0.124
$RP_3\beta$	0.141	0.190	0.117	0.133
EASE ^R	-	-	-	-
BIGCF (Code)	0.099	0.147	0.076	0.091
BIGCF (Paper)	0.098	0.146	0.076	0.091
Tmall				
TopPop	0.010	0.017	0.007	0.009
UserKNN	0.087	0.131	0.063	0.079
ItemKNN	0.110	0.163	0.080	0.099
$P_3\alpha$	0.109	0.162	0.080	0.099
$RP_3\beta$	0.116	0.173	0.084	0.104
EASE ^R	0.093	0.143	0.067	0.084
BIGCF (Code)	0.075	0.117	0.053	0.068
BIGCF (Paper)	0.075	0.116	0.053	0.068

none of the examined papers shared the codes that were used to obtain the results for the baselines. Codes for hyperparameter tuning and data preprocessing—all required to ensure full reproducibility—were entirely missing as well. Moreover, we found that authors are often not very responsive when approached for guidance in case of reproducibility issues. All in all, this situation is more than worrying, in particular when taking into account that our study focused on papers that were published at high-level venues for which we may assume the highest standards of scientific rigor.

Clearly, the general level of reproducibility in recommender systems research has increased in the last decade. Sharing code was very uncommon fifteen years ago. Nowadays, more and more publication venues encourage authors to share their code and data, and questions related to reproducibility are becoming common in review forms of conferences and journals. However, it seems that the measures that were taken so far are insufficient, and more awareness is needed among various stakeholders to improve the situation [4]. Ultimately, ensuring reproducibility should in most cases be technically feasible, e.g., through containerization or virtualization. Furthermore, various reproducibility guidelines are around for at least ten years [18]. Currently, however, there still seem to be insufficient incentives for researchers to prepare working reproducibility packages.

Methodological Issues and Progress. The existence of fundamental methodological issues in information retrieval and recommender systems are not new either, considering for example Armstrong et al.’s findings from fifteen years ago [3]. Various types of methodological flaws and bad practices were identified over the years, and our study shows that these problems seem to persist in recent research on IARS. The most important issues include the comparison of new models against baselines that were not properly tuned and setups where the potential effectiveness and expressive power of the baselines is artificially restricted by limiting the size of the user and item embeddings. Furthermore, the choice of the baselines also matters. In the examined works, most baselines are rather recent neural models, and only in a few cases more traditional models, e.g., based on matrix factorization, are included. This choice may lead to a continued *propagation of weak baselines* [11].

Furthermore, the establishment of what represents the *state-of-the-art* in terms of a small set of (recent) models that perform well across many use cases and datasets remains difficult and maybe impossible. The relative performance of algorithms may strongly depend on the characteristics of the dataset, the specifics of the evaluation protocol, and how hyperparameter tuning was done [40]. In our study, we found that a range of datasets was used—sometimes even using the same name but applying different preprocessing steps—and that there is also no consistency in terms of cut-off lengths. Specifics of the data splitting approach, which may affect performance outcomes, are also often missing.

Considering the specific papers that were investigated in our study, we may fear that the progress that is made even through the latest models may in fact be quite limited. Despite their conceptual simplicity, we found that longer-existing approaches are highly competitive with the latest complex models, despite the often very high computational costs of such models.

Computational Complexity Analysis. To illustrate the aspect of computational complexity, we consider the running times for the KGIN model discussed above for the largest dataset, Alibaba iFashion, that was used in the original paper for evaluation. The dataset comprises about 114,000 users, 30,000 items, and about 1.7M interactions. We recall that the Netflix Prize dataset that was released fifteen years ago contained 100M interactions. The running times for the KGIN model on our hardware¹⁶ are shown in Table 7. We report the time needed to train the model with the given set of hyperparameters (T-Time), the time needed to make predictions for all users (P-Time), and the time needed to make a prediction for a single user.

In Table 7, we observe that the time needed to train the KGIN model is orders of magnitude higher than our considered baselines. We recall that for this dataset, the decade-old itemKNN approach worked best in terms of accuracy. Running the preprocessing steps¹⁷ for ItemKNN takes less than three minutes, while training KGIN on the GPU may take days to weeks on our hardware. At prediction time, the KGIN model is also 50% slower than ItemKNN. As discussed in Section 3.2, the performance results that we obtained

by running the code provided for the KGIN model were lower than reported in the paper.¹⁸ Still, even if the performance of ItemKNN would be slightly surpassed by KGIN or another complex model, the question arises if the immense computational effort and corresponding carbon footprint would be justified for a modest gain in offline accuracy. In particular, this question has to be addressed bearing in mind that the offline-online gap in recommender systems has been discussed for at least a decade [22], where it is more than unclear if gains in offline accuracy would ever translate into more successful systems in practice.

Having said that, we in no way claim that research in complex models should be dismissed. In contrast, there are various reports of applying such models successfully in practice for several years, e.g., [6, 45]. Furthermore, there are other reproducibility studies that find that complex models can actually outperform traditional baselines, even if these baselines are extensively tuned, see, e.g., the study by Anelli et al. [2] on models based on Graph Collaborative Filtering. What our study focuses on instead are continuing issues of academic research practices, which seem to continue to prevent us from making faster progress.

Table 7: Comparison of training (T-Time) and prediction times (P-Time) for different models for the Alibaba-iFashion dataset.

Model	T-Time (s)	P-Time (s)	Avg. P-Time/User (ms)
TopPop	0.022	172.280	1.501
ItemkNN	140.771	201.694	1.757
UserkNN	381.969	162.805	1.418
$P_3\alpha$	37.596	170.240	1.483
$RP_3\beta$	38.748	173.301	1.510
EASE ^R	316.519	206.727	1.801
KGIN	117,360.426	362.042	3.150

Limitations. Our research is focused on the recent area of intent-aware recommender systems and to a limited set of five works that were published with artifacts for reproducibility and which we identified through a systematic process. While this process allows us to gauge the level of reproducibility in this particular area, more research is needed in the future to assess if similar methodological issues exist in other works in IARS. This includes approaches that rely on explicitly specified domain-specific intents and works that focus on sequential recommendation problems, see [14]. Furthermore, our study does not allow us to make conclusions about algorithms research in recommender systems in general, or about specific families of models, e.g., graph-based ones. In terms of the baselines considered in our study, we so far only included EASE^R as a non-neural model-based technique. The exploration of traditional models based on matrix factorization, like ALS or BPR, which both were found to be competitive with recent models [32, 37], is part of our ongoing and future work.

Like comparable previous reproducibility studies, we did *not* further tune the hyperparameter of the proposed IARS model. Instead, our assumption is that the hyperparameters that were shared by

¹⁶Hardware specifications: NVIDIA RTX A4000, GPU RAM: 16 GB, and CPU RAM: 125 GB

¹⁷There is no traditional ‘training’ for nearest-neighbor model. Therefore, we report the time that is needed to setup and fill internal data structures for similarity computations.

¹⁸The ItemKNN method is however at least on par with the KGIN model when considering the results in the original paper.

the authors correspond to the optimal ones that they found while developing and tuning the models. Given the computational complexity of some of these models, the authors of the original papers may have however been limited in the number of hyperparameter configurations they could explore. Thus, it may be easily possible that the performance of the proposed IARS models can be further improved. We however recall here that such an additional tuning process is not the focus of our current study and would require substantial computational resources. Instead, the goals of our study were (i) to assess if the results reported in the original papers can be reproduced using only the artifacts (including documentation) shared by the authors, however, a different approach was apparently taken in [2], who partially relied on alternative implementations of the models¹⁹, e.g., in the case of the DGCF model and (ii) how these models compare to well-tuned traditional models.

5 Conclusion & Outlook

We have studied questions of reproducibility and progress in the emerging area of intent-aware recommender systems. Worryingly, our analyses point to continuing methodological issues in this area. Our findings thus call for further action to increase awareness in the community about these issues, intensified education of all involved stakeholders, and better incentives for academic scholars so that we are able to achieve research results that are more reliable and better reproducible in the future. The problems discussed in our work are not new [18], and we for example observe certain improvements over the last decade in particular in terms of sharing of code and data. However, we believe that we as a community still have to strive harder to achieve more rigor in our scientific practices in the future.

References

- [1] Vito Walter Anelli, Alejandro Bellogín, Tommaso Di Noia, Dietmar Jannach, and Claudio Pomo. 2022. Top-N Recommendation Algorithms: A Quest for the State-of-the-Art. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization* (Barcelona, Spain) (UMAP '22). Association for Computing Machinery, New York, NY, USA, 121–131. <https://doi.org/10.1145/3503252.3531292>
- [2] Vito Walter Anelli, Daniele Malitesta, Claudio Pomo, Alejandro Bellogín, Eugenio Di Sciascio, and Tommaso Di Noia. 2023. Challenging the Myth of Graph Collaborative Filtering: A Reasoned and Reproducibility-Driven Analysis. In *Proceedings of the 17th ACM Conference on Recommender Systems* (Singapore, Singapore) (RecSys '23). 350–361. <https://doi.org/10.1145/3604915.3609489>
- [3] Timothy G. Armstrong, Alistair Moffat, William Webber, and Justin Zobel. 2009. Improvements that Don't Add Up: Ad-hoc Retrieval Results Since 1998. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management* (Hong Kong, China) (CIKM '09). Association for Computing Machinery, New York, NY, USA, 601–610. <https://doi.org/10.1145/1645953.1646031>
- [4] Christine Bauer, Maik Fröbe, Dietmar Jannach, Udo Kruschwitz, Paolo Rosso, Damiano Spina, and Nava Tintarev. 2023. Overcoming Methodological Challenges in Information Retrieval and Recommender Systems through Awareness and Education. In *Dagstuhl Seminar 23031: Frontiers of Information Access Experimentation for Research and Education*.
- [5] Colin Cooper, Sang Hyuk Lee, Tomasz Radzik, and Yiannis Siantos. 2014. Random Walks in Recommender Systems: Exact Computation and Simulations (WWW '14 Companion). Association for Computing Machinery, New York, NY, USA, 811–816. <https://doi.org/10.1145/2567948.2579244>
- [6] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (Boston, Massachusetts, USA) (RecSys '16). 191–198. <https://doi.org/10.1145/2959100.2959190>
- [7] Maurizio Ferrari Dacrema, Michael Benigni, and Nicola Ferro. 2025. Reproducibility and Artifact Consistency of the SIGIR 2022 Recommender Systems Papers Based on Message Passing. arXiv:2503.07823 [cs.LG]. <https://arxiv.org/abs/2503.07823>
- [8] Weiqiang Di. 2022. A Multi-Intent based Multi-Policy Relay Contrastive Learning for Sequential Recommendation. *PeerJ Computer Science* 8 (Aug. 2022), e1088. <https://doi.org/10.7717/peerj-cs.1088>
- [9] Michael D. Ekstrand, Michael Ludwig, Joseph A. Konstan, and John T. Riedl. 2011. Rethinking the Recommender Research Ecosystem: Reproducibility, Openness, and LensKit. In *Proceedings of the Fifth ACM Conference on Recommender Systems* (Chicago, Illinois, USA) (RecSys '11). 133–140. <https://doi.org/10.1145/2043932.2043958>
- [10] Maurizio Ferrari Dacrema, Simone Boglio, Paolo Cremonesi, and Dietmar Jannach. 2021. A Troubling Analysis of Reproducibility and Progress in Recommender Systems Research. *ACM Trans. Inf. Syst.* 39, 2, Article 20 (Jan. 2021), 49 pages. <https://doi.org/10.1145/3434185>
- [11] Maurizio Ferrari Dacrema, Paolo Cremonesi, and Dietmar Jannach. 2019. Are we Really Making much Progress? A Worrying Analysis of Recent Neural Recommendation Approaches. In *Proceedings of the 13th ACM Conference on Recommender Systems* (Copenhagen, Denmark) (RecSys '19). Association for Computing Machinery, New York, NY, USA, 101–109. <https://doi.org/10.1145/3298689.3347058>
- [12] Balázs Hidasi and Ádám Tibor Czapp. 2023. The Effect of Third Party Implementations on Reproducibility. In *Proceedings of the 17th ACM Conference on Recommender Systems*, RecSys 2023. 272–282. <https://doi.org/10.1145/3604915.3609487>
- [13] Linmei Hu, Siyong Xu, Chen Li, Cheng Yang, Chuan Shi, Nan Duan, Xing Xie, and Ming Zhou. 2020. Graph Neural News Recommendation with Unsupervised Preference Disentanglement. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 4255–4264. <https://doi.org/10.18653/v1/2020.acl-main.392>
- [14] Dietmar Jannach and Markus Zanker. 2024. A Survey on Intent-Aware Recommender Systems. *ACM Trans. Recomm. Syst.* 3, 2, Article 23 (Dec. 2024), 32 pages. <https://doi.org/10.1145/3700890>
- [15] Komal Kapoor, Vikas Kumar, Loren Terveen, Joseph A. Konstan, and Paul Schrater. 2015. "I Like to Explore Sometimes": Adapting to Dynamic User Novelty Preferences. In *Proc. RecSys '15* (Vienna, Austria), 19–26. <https://doi.org/10.1145/2792838.2800172>
- [16] Barrie Kersbergen, Olivier Sprangers, and Sebastian Schelter. 2022. Serenade - Low-Latency Session-Based Recommendation in e-Commerce at Scale. In *Proceedings of the 2022 International Conference on Management of Data* (SIGMOD '22). 150–159. <https://doi.org/10.1145/3514221.3517901>
- [17] Peter Knees, Dominik Schnitzer, and Arthur Flexer. 2014. Improving Neighborhood-based Collaborative Filtering by Reducing Hubness. In *Proceedings of International Conference on Multimedia Retrieval* (Glasgow, United Kingdom) (ICMR '14). Association for Computing Machinery, New York, NY, USA, 161–168. <https://doi.org/10.1145/2578726.2578747>
- [18] Joseph A. Konstan and Gediminas Adomavicius. 2013. Toward Identification and Adoption of Best Practices in Algorithmic Recommender Systems Research. In *Proceedings of the International Workshop on Reproducibility and Replication in Recommender Systems Evaluation* (Hong Kong, China) (RepSys '13). 23–28. <https://doi.org/10.1145/2532508.2532513>
- [19] Pigi Kouki, Ilias Fountalis, Nikolaos Vasiloglou, Xiquan Cai, Edo Liberty, and Khalifeh Al Jadda. 2020. From the Lab to Production: A Case Study of Session-based Recommendations in the Home-Improvement Domain. In *Proceedings of the 14th ACM Conference on Recommender Systems* (Virtual Event, Brazil) (RecSys '20). Association for Computing Machinery, New York, NY, USA, 140–149. <https://doi.org/10.1145/3383313.3412235>
- [20] Walid Krichene and Steffen Rendle. 2022. On Sampled Metrics for Item Recommendation. *Commun. ACM* 65, 7 (2022), 75–83. <https://doi.org/10.1145/3535335>
- [21] Sara Latifi, Noemi Mauro, and Dietmar Jannach. 2021. Session-Aware Recommendation: A Surprising Quest for the State-of-the-Art. *Information Sciences* 573 (2021), 291–315. <https://doi.org/10.1016/j.ins.2021.05.048>
- [22] Mark Levy. 2013. Offline Evaluation of Recommender Systems: All Pain and No Gain?. In *Proceedings of the International Workshop on Reproducibility and Replication in Recommender Systems Evaluation* (Hong Kong, China) (RepSys '13). 1. <https://doi.org/10.1145/2532508.2532509>
- [23] Qiming Li, Zhao Zhang, Fuzhen Zhuang, Yongjun Xu, and Chao Li. 2023. Topic-Aware Attention Network for Explainable Recommendation with Knowledge Enhancement. *ACM Trans. Inf. Syst.* 41, 4, Article 92 (April 2023), 23 pages. <https://doi.org/10.1145/3579993>
- [24] Shaosong Li, Bo Yang, and Dongsheng Li. 2023. Entity-driven User Intent Inference for Knowledge Graph-based Recommendation. *Applied Intelligence* 53, 9 (2023), 10734–10750. <https://doi.org/10.1007/s10489-022-04048-4>
- [25] Xuewei Li, Aitong Sun, Mankun Zhao, Jian Yu, Kun Zhu, Di Jin, Mei Yu, and Ruiguo Yu. 2023. Multi-Intention Oriented Contrastive Learning for Sequential Recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 411–419. <https://doi.org/10.1145/3539597.3570411>
- [26] Jimmy Lin. 2019. The Neural Hype and Comparisons Against Weak Baselines. *SIGIR Forum* 52, 2 (Jan. 2019), 40–51. <https://doi.org/10.1145/3308774.3308781>

¹⁹Moreover, reimplementing a model may however come with its own risks [12, 39].

- [27] Zheng Liu, Jianxun Lian, Junhan Yang, Defu Lian, and Xing Xie. 2020. Octopus: Comprehensive and Elastic User Representation for the Generation of Recommendation Candidates. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 289–298. <https://doi.org/10.1145/3397271.3401088>
- [28] Malte Ludewig, Sara Latifi, Noemi Mauro, and Dietmar Jannach. 2021. Empirical Analysis of Session-Based Recommendation Algorithms. *User Modeling and User-Adapted Interaction* 31, 1 (2021), 149–181. <https://doi.org/10.1007/s11257-020-09277-1>
- [29] Jianxin Ma, Chang Zhou, Peng Cui, Hongxia Yang, and Wenwu Zhu. 2019. Learning Disentangled Representations for Recommendation. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc.
- [30] Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. 2018. Statistical and Machine Learning Forecasting Methods: Concerns and Ways Forward. *PLoS one* 13, 3 (2018). <https://doi.org/10.1371/journal.pone.0194889>
- [31] Rishabh Mehrotra, Mounia Lalmas, Doug Kenney, Thomas Lim-Meng, and Golli Hashemian. 2019. Jointly Leveraging Intent and Interaction Signals to Predict User Satisfaction with Slate Recommendations. In *The World Wide Web Conference*. 1256–1267. <https://doi.org/10.1145/3308558.3313613>
- [32] Aleksandr Milogradskii, Oleg Lashinin, Alexander P. Marina Ananyeva, and Sergey Kolesnikov. 2024. Revisiting BPR: A Replicability Study of a Common Recommender System Baseline. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (RecSys '24). Association for Computing Machinery, New York, NY, USA, 267–277. <https://doi.org/10.1145/3640457.3688073>
- [33] Tiejun Qian, Yile Liang, Qing Li, Xuan Ma, Ke Sun, and Zhiyong Peng. 2023. Intent Disentanglement and Feature Self-Supervision for Novel Recommendation. *IEEE Transactions on Knowledge and Data Engineering* 35, 10 (2023), 9864–9877. <https://doi.org/10.1109/TKDE.2022.3175536>
- [34] Wajid Rafique, Abdelhakim Senhaji Hafid, and Junaid Qadir. 2023. Developing Smart City Services using Intent-Aware Recommendation Systems: A Survey. *Transactions on Emerging Telecommunications Technologies* 34, 4 (April 2023), e4728. <https://doi.org/10.1002/ett.4728>
- [35] Xubin Ren, Lianghao Xia, Jiashu Zhao, Dawei Yin, and Chao Huang. 2023. Disentangled Contrastive Collaborative Filtering. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Taipei, Taiwan) (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 1137–1146. <https://doi.org/10.1145/3539618.3591665>
- [36] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural Collaborative Filtering vs. Matrix Factorization Revisited. In *Proceedings of the 14th ACM Conference on Recommender Systems* (Virtual Event, Brazil) (RecSys '20). Association for Computing Machinery, New York, NY, USA, 240–248. <https://doi.org/10.1145/3383313.3412488>
- [37] Steffen Rendle, Walid Krichene, Li Zhang, and Yehuda Koren. 2022. Revisiting the Performance of iALS on Item Recommendation Benchmarks. In *Proceedings of the 16th ACM Conference on Recommender Systems* (Seattle, WA, USA) (RecSys '22). Association for Computing Machinery, New York, NY, USA, 427–435. <https://doi.org/10.1145/3523227.3548486>
- [38] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. 1994. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (Chapel Hill, North Carolina, USA) (CSCW '94). Association for Computing Machinery, New York, NY, USA, 175–186. <https://doi.org/10.1145/192844.192905>
- [39] A. Said and A. Bellogin. 2014. Comparative Recommender System Evaluation: Benchmarking Recommendation Frameworks. In *8th ACM Conference on Recommender Systems* (RecSys). 129–136. <https://doi.org/10.1145/2645710.2645746>
- [40] Faisal Shehzad and Dietmar Jannach. 2023. Everyone's a Winner! On Hyperparameter Tuning of Recommendation Models (RecSys '23). Association for Computing Machinery, New York, NY, USA, 652–657. <https://doi.org/10.1145/3604915.3609488>
- [41] Faisal Shehzad and Dietmar Jannach. 2024. Performance Comparison of Session-based Recommendation Algorithms based on GNNs. In *Advances in Information Retrieval*, Nazli Goharian, Nicola Tonello, Yulan He, Aldo Lipani, Graham McDonald, Craig Macdonald, and Iadh Ounis (Eds.). Springer Nature Switzerland, Cham, 115–131. https://doi.org/10.1007/978-3-031-56066-8_12
- [42] Fanjuan Shi and Chirine Ghedira. 2017. Improving Recommender Systems with an Intent-based Algorithm Switching Strategy. In *Proceedings of the Symposium on Applied Computing*. 1668–1673. <https://doi.org/10.1145/3019612.3019761>
- [43] Giuseppe Spillo, Allegra De Filippo, Cataldo Musto, Michela Milano, and Giovanni Semeraro. 2024. Towards Green Recommender Systems: Investigating the Impact of Data Reduction on Carbon Footprint and Algorithm Performances. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (RecSys '24). Association for Computing Machinery, New York, NY, USA, 866–871. <https://doi.org/10.1145/3640457.3688160>
- [44] Harald Steck. 2019. Embarrassingly Shallow Autoencoders for Sparse Data. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 3251–3257. <https://doi.org/10.1145/3308558.3313710>
- [45] Harald Steck, Linas Baltrunas, Ehtsham Elahi, Dawen Liang, Yves Raimond, and Justin Basilico. 2021. Deep Learning for Recommender Systems: A Netflix Case Study. *AI Mag.* 42, 3 (2021), 7–18. <https://doi.org/10.1609/AIMAG.V42I3.18140>
- [46] Zhu Sun, Di Yu, Hui Fang, Jie Yang, Xinghua Qu, Jie Zhang, and Cong Geng. 2020. Are We Evaluating Rigorously? Benchmarking Recommendation for Reproducible Evaluation and Fair Comparison. In *Proceedings of the 14th ACM Conference on Recommender Systems* (RecSys '20). 23–32. <https://doi.org/10.1145/3383313.3412489>
- [47] Md Mehrab Tanjim, Congzhe Su, Ethan Benjamin, Diane Hu, Liangjie Hong, and Julian McAuley. 2020. Attentive Sequential Models of Latent Intent for Next Item Recommendation. In *Proceedings of The Web Conference 2020*. 2528–2534. <https://doi.org/10.1145/3366423.3380002>
- [48] Saul Vargas, Pablo Castells, and David Vallet. 2011. Intent-Oriented Diversity in Recommender Systems. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1211–1212. <https://doi.org/10.1145/2009916.2010124>
- [49] Tobias Vente, Lukas Wegmeth, Alan Said, and Joeran Beel. 2024. From Clicks to Carbon: The Environmental Toll of Recommender Systems. In *Proceedings of the 18th ACM Conference on Recommender Systems* (Bari, Italy) (RecSys '24). Association for Computing Machinery, New York, NY, USA, 580–590. <https://doi.org/10.1145/3640457.3688074>
- [50] Jun Wang, Arjen P. de Vries, and Marcel J. T. Reinders. 2006. Unifying User-based and Item-based Collaborative Filtering Approaches by Similarity Fusion (SIGIR '06). Association for Computing Machinery, New York, NY, USA, 501–508. <https://doi.org/10.1145/1148170.1148257>
- [51] Jun Wang, Stephen Robertson, Arjen P de Vries, and Marcel JT Reinders. 2008. Probabilistic Relevance Ranking for Collaborative Filtering. *Information Retrieval* 11 (2008), 477–497. <https://doi.org/10.1007/s10791-008-9060-1>
- [52] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (SIGIR '19). 165–174. <https://doi.org/10.1145/3331184.3331267>
- [53] Xiang Wang, Tinglin Huang, Dingxian Wang, Yancheng Yuan, Zhenguang Liu, Xiangnan He, and Tat-Seng Chua. 2021. Learning Intents behind Interactions with Knowledge Graph for Recommendation. In *Proceedings of the Web Conference 2021*. ACM, 878–887. <https://doi.org/10.1145/3442381.3450133>
- [54] Xiang Wang, Hongye Jin, An Zhang, Xiangnan He, Tong Xu, and Tat-Seng Chua. 2020. Disentangled Graph Collaborative Filtering. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1001–1010. <https://doi.org/10.1145/3397271.3401137>
- [55] Jacek Wasilewski and Neil Hurley. 2016. Intent-Aware Diversification using a Constrained PLSA. In *Proceedings of the 10th ACM Conference on Recommender Systems*. 39–42. <https://doi.org/10.1145/2959100.2959177>
- [56] Haolun Wu, Yingxue Zhang, Chen Ma, Wei Guo, Ruiming Tang, Xue Liu, and Mark Coates. 2023. Intent-Aware Multi-Source Contrastive Alignment for Tag-enhanced Recommendation. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 1112–1125. <https://doi.org/10.1109/ICDE55515.2023.00090>
- [57] Longqi Yang, Michael Sobolev, Yu Wang, Jenny Chen, Drew Dunne, Christina Tsangouri, Nicola Dell, Mor Naaman, and Deborah Estrin. 2019. How Intention Informed Recommendations Modulate Choices: A Field Study of Spoken Word Content. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 2169–2180. <https://doi.org/10.1145/3308558.3313540>
- [58] Ruoyi Zhang, Huifang Ma, Qingfeng Li, Yike Wang, and Zhixin Li. 2023. Fire: Knowledge-Enhanced Recommendation with Feature Interaction and Intent-Aware Attention Networks. *Applied Intelligence* 53, 13 (2023), 16424–16444. <https://doi.org/10.1007/s10489-022-04300-x>
- [59] Yi Zhang, Lei Sang, and Yiwen Zhang. 2024. Exploring the Individuality and Collectivity of Intents behind Interactions for Graph Collaborative Filtering. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 1253–1262. <https://doi.org/10.1145/3626772.3657738>
- [60] Jieming Zhu, Quanyu Dai, Liangcai Su, Rong Ma, Jinyang Liu, Guohao Cai, Xi Xiao, and Rui Zhang. 2022. BARS: Towards Open Benchmarking for Recommender Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (SIGIR '22). 2912–2923. <https://doi.org/10.1145/3477495.3531723>