

# RecBaselines2023: a new dataset for choosing baselines for recommender models

Veronika Ivanova<sup>1</sup>, Oleg Lashinin<sup>2</sup>, Marina Ananyeva<sup>1,2</sup> and Sergey Kolesnikov<sup>2</sup>

<sup>1</sup>National Research University Higher School of Economics, Myasnitskaya Ulitsa, 20, Moscow, 101000, Russian Federation

<sup>2</sup>Tinkoff, 2-Ya Khutorskaya Ulitsa, 38A, bld. 26, Moscow, 117198, Russian Federation

## Abstract

The number of proposed recommender algorithms continues to grow. The authors propose new approaches and compare them with existing models, called baselines. Due to the large number of recommender models, it is difficult to estimate which algorithms to choose in the article. To solve this problem, we have collected and published a dataset containing information about the recommender models used in 903 papers, both as baselines and as proposed approaches. This dataset can be seen as a typical dataset with interactions between papers and previously proposed models. In addition, we provide a descriptive analysis of the dataset and highlight possible challenges to be investigated with the data. Furthermore, we have conducted extensive experiments using a well-established methodology to build a good recommender algorithm under the dataset. Our experiments show that the selection of the best baselines for proposing new recommender approaches can be considered and successfully solved by existing state-of-the-art collaborative filtering models. Finally, we discuss limitations and future work.

## Keywords

recommender systems, dataset, baselines

## 1. Introduction

There is an increasing number of publications in the field of recommender systems. Authors need to evaluate the performance of the proposed model against reference models to demonstrate its efficiency. Reference models are usually referred to as baselines. However, there are no rigid guidelines that define a comprehensive list of essential baselines. Inaccurate selection of baselines can lead to incorrect conclusions about the performance of the proposed model. Subsequent papers [1] on the reproducibility and progress of existing work have demonstrated this fact. For example, in two recent papers [2, 3], the authors report that for a particular information retrieval task, some non-neural methods outperform recent neural methods. In 2016, Kharazmi et al. [4] examined previous work on IR and found a tendency to select weak baselines for comparative experiments. In the field of recommender systems, the empirical analyses of session-based recommendation papers showed that sometimes almost trivial methods can outperform the latest neural methods [5, 6]. One of the reasons for such disappointing performance of novel models is the poor choice of baselines, which gives the illusion of better

---

*BIR 2023: 13th International Workshop on Bibliometric-enhanced Information Retrieval at ECIR 2023, April 2, 2023*

✉ veronika.ivanova88@yandex.ru (V. Ivanova); o.a.lashinin@tinkoff.ru (O. Lashinin); m.ananyeva@tinkoff.ru (M. Ananyeva); scitator@gmail.com (S. Kolesnikov)

ORCID 0000-0001-8894-9592 (O. Lashinin); 0000-0002-9885-2230 (M. Ananyeva); 0000-0002-4820-987X (S. Kolesnikov)



© 2023 Copyright © JJJJ for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

results. Another consequence of not choosing appropriate baselines for a new algorithm is that the proposed paper may be rejected [7]. Thus, the choice of baselines is currently one of the major issues in recommender systems research [8, 9].

With accurate baseline selection, the development of recommender systems can progress more quickly. Both researchers and practitioners are faced with an increasing number of models to select for their experiments in order to consider relevant baselines. However, the number of baselines included in the paper is limited for the following reasons. First, more baselines require more computational time. The recent success of deep learning forces the inclusion of complicated algorithms as baselines. Therefore, some researchers cannot afford to choose effective hyperparameters sufficiently [8]. Second, some papers with new recommender algorithms do not have the source code of the implementation [8]. This may lead to poor performance of third party implementations [10, 8]. Finally, due to space limitations, a paper cannot include too many baselines. Thus, it is common practice to study the performance of only 3-7 baselines against the newly proposed method. The problem of selecting a few relevant items from a large set is a well-known task and can be solved by recommender systems [11]. To the best of our knowledge, there is no open source dataset that can be used as a basis for developing the recommender baseline suggestion system.

It is important to note that the baseline recommendations can be applied to other areas of machine learning, such as natural language processing, computer vision, time series prediction, and others. However, in this paper we focus only on recommender systems.

In this paper we describe a process for collecting a novel dataset called RecBaselines2023. It can be considered as a classical dataset with interactions between papers and baselines. In addition, we present the results of experiments that have been performed on RecBaselines2023. Our results show potential advantages of our experiments and open new research directions. Specifically, the main contribution of this paper can be listed as follows:

- We have created a new open source dataset called RecBaseline2023<sup>1</sup> for selecting baselines for experiments on recommender models. We examined 1009 papers for the collection process. After preprocessing, RecBaseline2023 contains information on 363 baselines used in 903 articles published between 2010 and 2022. We also provide a data collection procedure and descriptive statistics.
- We discuss that the problem of baseline selection can be solved by collaborative filtering approaches. We then compare the baseline ranking quality of seven state-of-the-art top-N recommender models on RecBaseline2023. The results show that this problem can be effectively solved by selected algorithms.
- We describe a scenario where a partial list of baselines needs to be completed. The list is given to collaborative filtering approaches that recommend baselines based on the list of methods already used. Some other possible use cases of RecBaseline2023 are also mentioned.

---

<sup>1</sup>We are releasing an online version of the dataset: <https://github.com/fotol1/recbaselines2023>.

**Table 1**

The table shows the corresponding highly cited baselines in each of the three recommender tasks. We use these as a starting point for our data collection. Citation counts are valid as of 27 March 2023.

Recommender task	Mandatory baselines (Citation counts)
conventional top-N	BPR-MF (5297) [18], WMF (3688) [19], MultVAE (844) [20], LightGCN (1245) [21]
next-item	GRU4Rec (2152) [22], SASRec (1146) [23], BERT4Rec (836) [24]
next-basket	TIFUKNN (56) [25], RepeatNet (189) [26]

## 2. Related work

The problem of choosing baselines for research experiments in machine learning is not well studied. A similar problem is citation recommendations. This direction aims at suggesting other papers to cite.

The two main classes of citation recommendations are content-based and collaborative filtering [9]. The content-based methods use textual elements such as abstract and title or metadata elements such as authors. In [12], the authors proposed a content-based approach requiring only textual features and collected the OpenCorpus dataset of 7 million articles. The literature graph was created in [13] using nodes for articles, authors and scientific concepts. Collaborative filtering methods are based on comparing similarities between articles. Liu et al. [14] measured the cosine similarity of article vectors and created article vectors based on co-occurrence in the same citation list. The same concept was used by Haruna et al. [15]. However, they considered the references and citations of the target paper and mined the hidden associations between them using paper-citation relationships. Later, by improving the similarity calculation, this approach was further developed in [16].

Although we know what to cite, it is not clear whether the recommended paper should be used as a baseline. Therefore, researchers are also working on more specialised tasks such as tag or baseline recommendations. The task of tag recommendation has been successfully studied by Wang et al. [17]. They used the collaborative topic regression model. The authors sampled items from the CiteULike dataset, including abstracts, titles and tags for each article. Bedi et al. [7] introduced the task of identifying the papers used as baselines in a given scientific article. The author formulated it as a reference classification problem on a developed ACL anthology corpus dataset, where about 2000 papers were selected and manually annotated.

However, research article datasets are not specifically designed for the task of selecting baselines for recommender system experiments. We hope that our dataset will help to fill this gap and provide researchers with a practical approach to selecting baseline models for their research.

## 3. Dataset

**Collection.** We added several common recommendation tasks to our dataset, including the traditional top-n, next-item and next-basket recommendations. These tasks were used as the basis for the data collection. For each task there are well-established and highly cited baselines, some of which are listed in Table 1. Note that there are no strict guidelines in recommender

**Table 2**

An example of a row in the RecBaselines2023 dataset. This file is called **before\_preprocessing.csv** and is available online.

Column	Description	Example
Paper_id	Unique paper identification	12
URL	WEB link to a paper	https://arxiv.org/pdf/1809.07053.pdf
Title	Paper title	NAIS: Neural Attentive Item Similarity Model for Recommendation
Year	A year of publishing	2018
Baselines	List of used recommender algorithms	MF;MLP;FISM;NAIS

**Table 3**

Examples of different names for the same algorithms

Approach	Occurring names
POP	MostPop, Popular, TopPopular
BPR-MF [18]	BPR, BPR-MF, MF-BPR
NCF [27]	NeuCF, NeuMF, NCF
Mult-VAE [28]	Mult-VAE, Multi-VAE, VAE-CF

**Table 4**

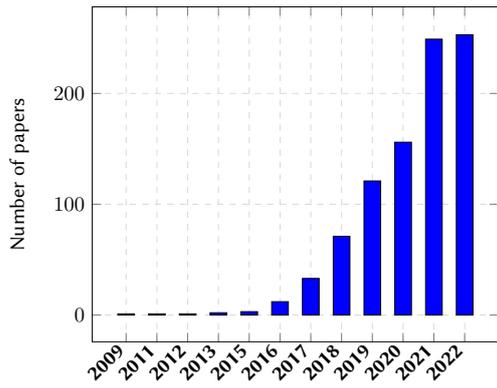
Descriptive statistics of RecBaselines2023 dataset.

Stage	number of papers	number of models	number of interactions	density
before preprocessing	1009	2188	7748	0.3%
after preprocessing	903	363	5467	1.6%

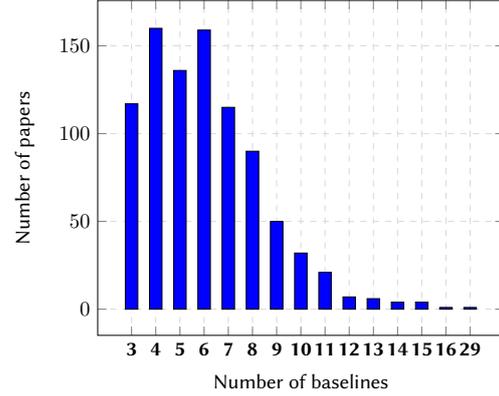
systems research as to which baselines should be used for each of the above tasks. Therefore, we cannot guarantee that other algorithms cannot be used to complement the list of commonly used algorithms. However, the approaches listed in Table 1 have many citations, which is appropriate for the starting point of data collection.

To collect our dataset, we took the following steps:

1. For each model from Table 1, we obtain the list of papers that cited the model in Google Scholar [29]. If a paper included experiments with the model, we included it. We did not include papers with experiments on related problems (such as link prediction or matrix completion, explanation generation). In addition, papers without experiments are not included in the dataset. Note that a paper could cite more than one baseline model of Table 1. Duplicate papers were later filtered out of the dataset during pre-processing. Once we had gone through all the citations of models in Table 1, we continued to process citations of papers that had already been added. This was all done manually by the authors of the paper over the period of one month.
2. Information about each paper collected to build our dataset is presented in Table 2. Each row contains a paper id, URL, paper title, year of publication and a list of recommender models used. The URL and year of publication are taken from the Google Scholar page, while the paper title and list of baselines are taken from the paper itself.

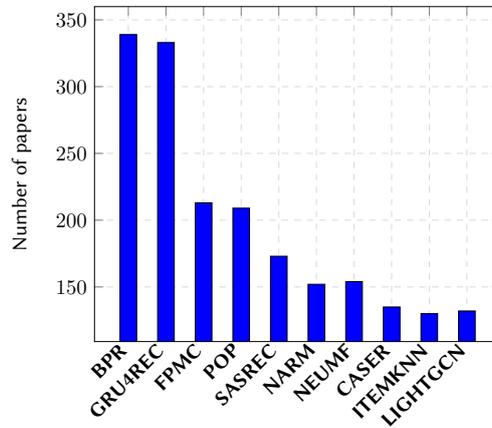


(a) Distribution of the number of papers by the year of publication.



(b) Distribution of the number of papers over the number of algorithms per one paper.

**Figure 1:** Dataset statistics

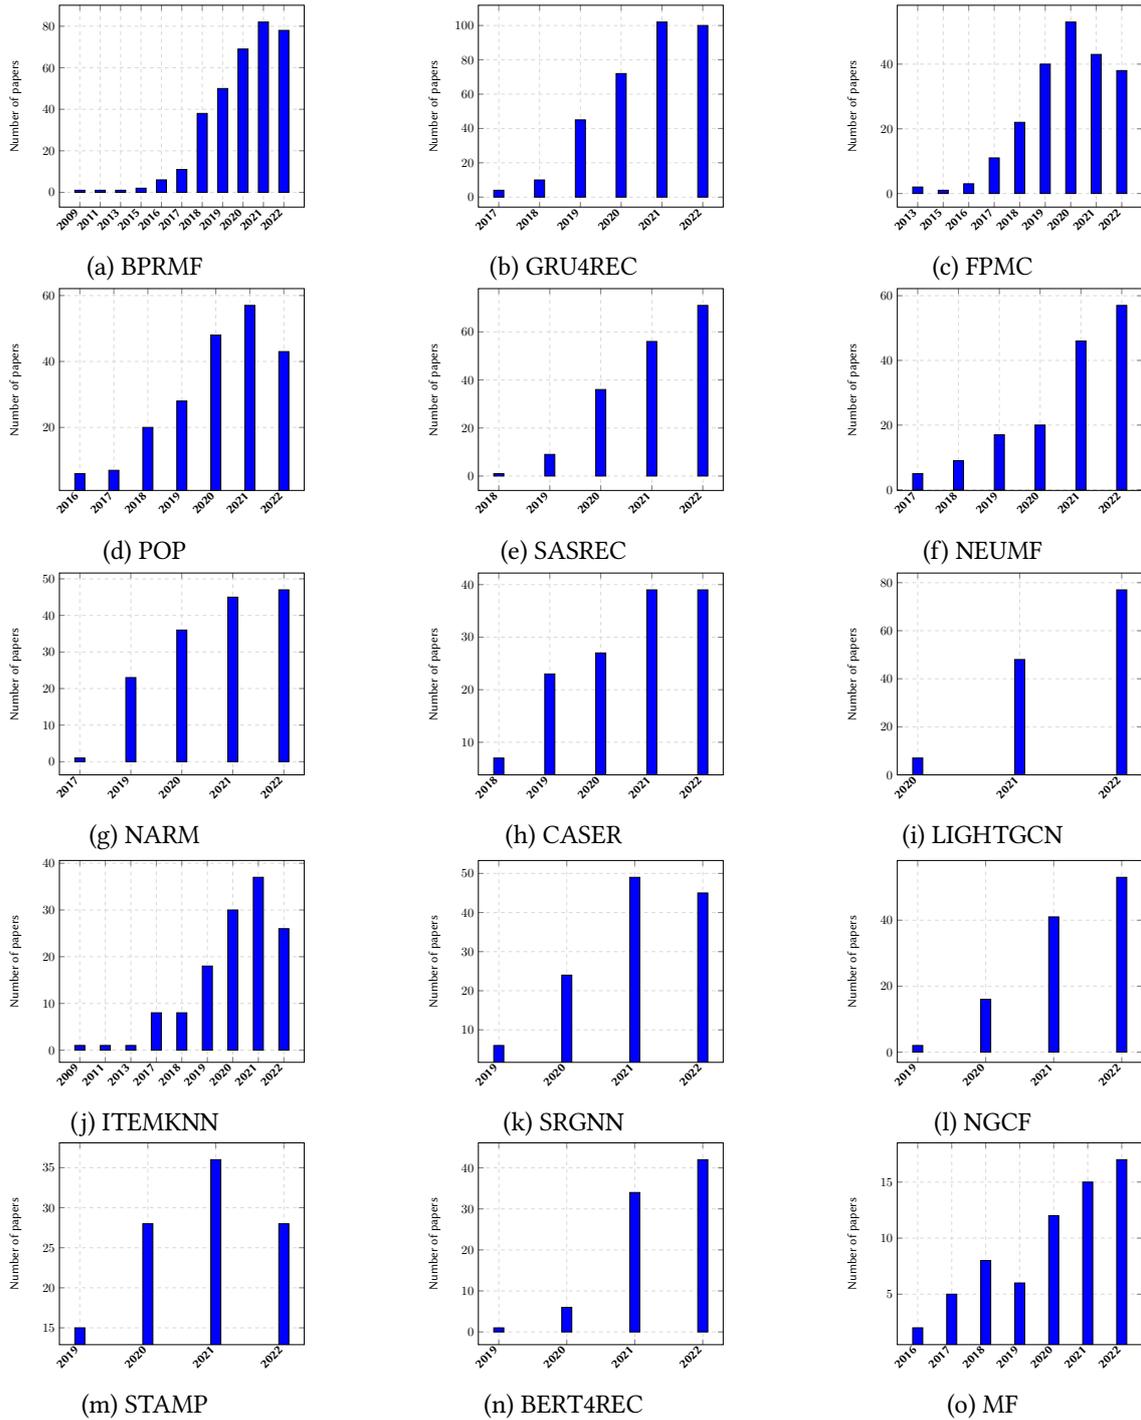


**Figure 2:** Distribution of the number of papers for the top-10 most popular baselines, included in our dataset.

After removing duplicates, we obtained the dataset with 1009 papers and 2187 baselines. A large number of baseline models were only included in one or two papers.

**Preprocessing.** A number of steps were taken to preprocess the data for future research:

1. In some papers, popular models are presented under different names. This is most likely due to space limitations or different names for algorithms that were even proposed in the original paper. For example, the authors of the article Neural Collaborative Filtering (NCF) [27] used a different name, NeuMF, in their experiments. As a result, the cited articles include both NCF and NeuMF. We try to condense common cases and list them in Table 3. To resolve this inconsistency, we have replaced the multiple names of a model with a single option.
2. Some papers modify methods slightly and report different variations of the same methods.



**Figure 3:** The figures show the distribution over the years of the number of papers in which one of the top-15 most popular baselines was included.

For example, the authors introduce three new loss functions in [30] and apply them to different methods such as NeuMF [27], CML [31] and LightGCN [21]. Considering three losses for each of the three models makes our dataset more sparse. To avoid this problem, the preprocessed version of RecBaselines2023 contains only the main algorithms without any specified modifications.

3. To replace rare baselines and papers with extremely few baselines, we then iteratively filtered the dataset until there were only papers with three or more baselines and each baseline was present in three or more papers. The resulting statistics for the filtered dataset can be found in Table 4.

**Statistics.** We briefly present some statistics from the collected dataset. The main characteristics such as number of papers, number of models, number of interactions and density are presented in Table 4.

Figure 1 and Figure 2 represent three distributions of the dataset: the distribution of the number of papers over the year they were published, the distribution of the number of papers over the number of algorithms included in a paper, and the distribution of the number of papers for the top 10 most popular baselines included in our dataset. The earliest publication date of a paper is 2009, the number of papers remains relatively small and only exceeds 10 in 2017. Then the number of papers increases significantly from year to year. As can be seen in Figure 1, a typical number of baselines included is between 3 and 8. Therefore, the algorithms used to recommend baselines for recommender systems have to work with a small number of available interactions.

Figure 3 describes the distribution of the number of papers over the years for each of the top 15 popular baselines from the RecBaselines2023 dataset. The most popular models are BPR, GRU4Rec, LightGCN, NeuMF and others. These models were used as starting points for the collection of other papers. Therefore, they are represented in the dataset in large numbers.

## 4. Collaborative Filtering for Baseline Selection

Baseline selection can be solved by collaborative filtering (CF) algorithms. For example, the following definition was given in [32].

**Definition 4.1.** *Collaborative filtering is the process of filtering or ranking items using the opinions of other people.*

We can replace the word "items" with "baselines" and the word "people" with "researchers". This definition then provides a justification for the use of the technique. In addition, scientists' "opinions" are often motivated by several reasons. The first is the desire to compare the new algorithm with the best-known or best-performing approaches. The second is to include models based on the same idea. For example, the authors of [33] compare their graph-based model with 4 baselines, 3 of which are also graph-based. These or other reasons explain the choice of models from a large number of options.

Therefore, researchers and practitioners may be interested in baseline recommendations based on a partial list of algorithms already in use. Hopefully, this can be done by applying approaches

**Table 5**Performance comparison on RecBaselines2023. The best value is **in bold**.

Model	R@10	R@20	N@10	N@20	M@10	M@20
Random	0.045	0.08	0.029	0.029	0.008	0.008
BPRMF [18]	0.2281	0.353	0.134	0.134	0.035	0.035
MostPop	0.312	0.339	0.138	0.138	0.035	0.035
MF2020 [41]	0.348	0.446	0.227	0.227	0.067	0.067
EASER [37]	0.397	0.549	0.243	0.243	0.069	0.069
NeuMF [27]	0.42	0.513	0.252	0.252	0.073	0.073
Slim [38]	0.446	0.576	0.264	0.264	0.078	0.078
VAECF [28]	0.455	0.603	0.264	0.264	0.075	0.075
<i>RP3<math>\beta</math></i> [39]	<b>0.473</b>	<b>0.607</b>	<b>0.303</b>	<b>0.303</b>	<b>0.088</b>	<b>0.088</b>

in the inductive scenario [34]. These approaches do not have ID-based user embeddings [33, 20, 35, 36]. They understand user interests based on the set of interactions. Therefore, we can easily adopt such techniques for suggesting baselines based on a partial list of already included methods.

## 5. Experiments

We have experimented with collaborative filtering approaches on the Top-N recommendation task on the RecBaselines2023 dataset. Our experiments aim to answer the following question: "What is the performance of different state-of-the-art collaborative filtering approaches on the RecBaselines2023 dataset?"

**Models.** We included popular approaches of different types: simple random and MostPop; matrix factorisation based BPRMF [18], MF2020; item based EASE [37], SLIM [38]; graph based *RB3* [39], VAE based MultiVAE [20]. According to [40, 8], such models are very strong CF-based baselines.

**Metrics.** Standard quality ranking metrics are chosen, namely Recall@K, NDCG@K and MAP@K.

**Experiment settings.** To provide reproducible experiments we use Elliot [42] similar to [40]. This framework allows experiments to be fully described in a configuration file. This file is available online <sup>2</sup> and hyperparameter ranges can be found there. The total number of hyperparameters set for each model is 20.

**Evaluation protocol.** All interactions are divided into train/valid/test splits. The valid split is used for early stopping and hyperparameter selection. The final quality is estimated on the test split. All papers published before 2021 are used for training. In addition, 80 % of the interactions for papers published in 2021 and 2023 are used for training, and the remaining 20 % are used for validation and test, respectively.

**Results.** To investigate our question, we report quality metrics for different approaches in Table 5. As we can see, the best model *RP3 $\beta$*  is two times better than MostPop’s recommenda-

<sup>2</sup><https://github.com/fotol1/recbaselines2023>



**Figure 4:** The scientist chooses baselines for his experiments. Three baselines have already been chosen. He or she can now pass the three approaches to one of the collaborative filtering algorithms under consideration. As a result, two more baselines are suggested as additional approaches to include in the paper. Based on historical data, these recommended baselines may also be chosen by other researchers.

tions. This shows that there are not many universal baselines in recommendations, and that researchers choose baselines carefully. Surprisingly, the best model,  $RP3_\beta$ , has a Recall@20 of 0.6. This means that we can find more than half of the hidden baselines in lists of length 20.

## 6. Selecting baselines for partial lists

This section describes one possible way of using RecBaselines2023. The Figure 4 represents the main idea. A scientist has invented a new recommendation algorithm and wants to compare it with other work. For example, a new approach was inspired by two methods **a** and **b**. So they are automatically included in the experiments of the new paper. In addition, the researchers know that a current state-of-the-art algorithm is a model **c**. So it should also be considered. Given the set of three baselines  $\{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$ , he or she can run the set in one of the adapted collaborative filtering approaches. This will return a list of recommended baselines. The choice of these is consistent with the historical data represented in RecBaselines2023.

In Table 6 we demonstrate recommendations for some sub-lists of baselines. We use SLIM, EASE and  $RP3_\beta$  as recommender models because they are item-based models that can make predictions based on any input list of items. The first three examples emulate iterative updates to the next-item recommender set of baselines. The next examples demonstrate recommendations based on a single baseline using different frameworks. As we can see, SLIM and  $RP3_\beta$  are flexible in changing recommendations as new next-item models are added. When we provide only one element of a particular framework, our models recommend baselines using similar frameworks. For example, RippleNet [43] is a knowledge-based model. If someone includes RippleNet in their experiments, our models will suggest including other knowledge-based

**Table 6**

Examples of recommendations based on partial lists of items.

Input items	SLIM	EASE	$RP3_\beta$
GRU4REC	CASER, SASREC	SASREC, CASER	MIND, DIEN
GRU4REC, SASREC, BERT4REC	CASER, TISASREC	BPR, CASER	TISASREC, JODIE
GRU4REC, SASREC, BERT4REC, TISASREC	CASER, S3REC	BPR, CASER	LESSR, CHORUS
PINSAGE	GCMC, CMN	NGCF, GCMC	CMN, GCMC
VAECF	TRANSCF, LRML	LIGHTGCN, TRANSCF	SGL, TRANSCF
RIPPLENET	PER, CKE	CKE, PER	PER, LIBFM
LIGHTGCN, SGL, VAECF	TRANSCF, LRML	NGCF, BPR	TRANSCF, SBPR

approaches such as PER [44], CKE [45].

## 7. Limitations and future work

Our work has some limitations. In this section we will discuss them and show possible ways to overcome them.

Firstly, the published version of the dataset will become obsolete. We will publish regular updates. In addition, if authors of newly proposed methods want to add their work, we can do this quickly in the repository via a pull request.

Secondly, The dataset may contain misspellings or other errors, with the presence or absence of some baselines in the included works. We have tried to do our best, and have double-checked the interactions several times. If you find any errors, please contact us via issues on Github.

Finally, There are some challenges in publishing baseline recommendations. For example, some of the baselines presented have been used in previous work. However, the latest state-of-the-art approaches replace them. The models considered are not sensitive to this fact. We argue that this problem exists for other datasets as well. It has been shown in [46] that recommending the most recent films can improve quality even for the simple MostPop method. Nevertheless, the practical application can be modified and the most recent baselines with high relevance scores can be treated as more suitable. We leave this as future work.

## 8. Conclusion

This paper investigates the problem of recommending baselines for experiments. We have collected an open source dataset RecBaselines2023, which describes baseline models used for comparative experiments in papers on different types of recommender systems. It consists of 903 papers and 363 baseline models, with 5467 interactions between them. The dataset includes interactions between papers and baseline models, and additional data about each paper, such as web link to a paper, paper title, and year of publication. RecBaselines2023 can be used by researchers to properly compile the baseline list for their experiments. The dataset will be updated as new papers are published. We have used collaborative filtering techniques to identify the best algorithms based on incomplete lists of previously included baselines. Our experiments with hidden predictions of recommender baselines show that state-of-the-art collaborative

filtering techniques can successfully perform this task. We hope that our dataset can open up new lines of research.

## References

- [1] M. Ferrari Dacrema, S. Boglio, P. Cremonesi, D. Jannach, A troubling analysis of reproducibility and progress in recommender systems research, *ACM Transactions on Information Systems (TOIS)* 39 (2021) 1–49.
- [2] J. Lin, The neural hype and comparisons against weak baselines, in: *ACM SIGIR Forum*, 2, ACM New York, NY, USA, 2019, pp. 40–51.
- [3] W. Yang, K. Lu, P. Yang, J. Lin, Critically examining the "neural hype" weak baselines and the additivity of effectiveness gains from neural ranking models, in: *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, 2019, pp. 1129–1132.
- [4] S. Kharazmi, F. Scholer, D. Vallet, M. Sanderson, Examining additivity and weak baselines, *ACM Transactions on Information Systems (TOIS)* 34 (2016) 1–18.
- [5] M. Ludewig, D. Jannach, Evaluation of session-based recommendation algorithms, *User Modeling and User-Adapted Interaction* 28 (2018) 331–390.
- [6] M. Ludewig, N. Mauro, S. Latifi, D. Jannach, Performance comparison of neural and non-neural approaches to session-based recommendation, in: *Proceedings of the 13th ACM conference on recommender systems*, 2019, pp. 462–466.
- [7] M. Bedi, T. Pandey, S. Bhatia, T. Chakraborty, Why did you not compare with that? identifying papers for use as baselines, in: *European Conference on Information Retrieval*, Springer, 2022, pp. 51–64.
- [8] M. Ferrari Dacrema, P. Cremonesi, D. Jannach, Are we really making much progress? a worrying analysis of recent neural recommendation approaches, in: *Proceedings of the 13th ACM conference on recommender systems*, 2019, pp. 101–109.
- [9] J. Beel, B. Gipp, S. Langer, C. Breitingner, Paper recommender systems: a literature survey, *International Journal on Digital Libraries* 17 (2016) 305–338.
- [10] A. Petrov, C. Macdonald, A systematic review and replicability study of bert4rec for sequential recommendation, in: *Proceedings of the 16th ACM Conference on Recommender Systems*, 2022, pp. 436–447.
- [11] D. Bollen, B. P. Knijnenburg, M. C. Willemsen, M. Graus, Understanding choice overload in recommender systems, in: *Proceedings of the fourth ACM conference on Recommender systems*, 2010, pp. 63–70.
- [12] C. Bhagavatula, S. Feldman, R. Power, W. Ammar, Content-based citation recommendation, *arXiv preprint arXiv:1802.08301* (2018).
- [13] W. Ammar, D. Groeneveld, C. Bhagavatula, I. Beltagy, M. Crawford, D. Downey, J. Dunkelberger, A. Elgohary, S. Feldman, V. Ha, et al., Construction of the literature graph in semantic scholar, *arXiv preprint arXiv:1805.02262* (2018).
- [14] H. Liu, X. Kong, X. Bai, W. Wang, T. M. Bekele, F. Xia, Context-based collaborative filtering for citation recommendation, *Ieee Access* 3 (2015) 1695–1703.

- [15] K. Haruna, M. Akmar Ismail, D. Damiasih, J. Sutopo, T. Herawan, A collaborative approach for research paper recommender system, *PloS one* 12 (2017) e0184516.
- [16] N. Sakib, R. B. Ahmad, K. Haruna, A collaborative approach toward scientific paper recommendation using citation context, *IEEE Access* 8 (2020) 51246–51255.
- [17] H. Wang, B. Chen, W. Li, Collaborative topic regression with social regularization for tag recommendation, in: F. Rossi (Ed.), *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013, IJCAI/AAAI, 2013*, pp. 2719–2725. URL: <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/7006>.
- [18] S. Rendle, C. Freudenthaler, Z. Gantner, L. Schmidt-Thieme, Bpr: Bayesian personalized ranking from implicit feedback, *arXiv preprint arXiv:1205.2618* (2012).
- [19] Y. Hu, Y. Koren, C. Volinsky, Collaborative filtering for implicit feedback datasets, in: *2008 Eighth IEEE international conference on data mining, Icdm, 2008*, pp. 263–272.
- [20] J. Xu, Y. Ren, H. Tang, X. Pu, X. Zhu, M. Zeng, L. He, Multi-vae: Learning disentangled view-common and view-peculiar visual representations for multi-view clustering, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021*, pp. 9234–9243.
- [21] Y. Shen, Y. Wu, Y. Zhang, C. Shan, J. Zhang, B. K. Letaief, D. Li, How powerful is graph convolution for recommendation?, in: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management, 2021*, pp. 1619–1629.
- [22] B. Hidasi, A. Karatzoglou, L. Baltrunas, D. Tikk, Session-based recommendations with recurrent neural networks, *arXiv preprint arXiv:1511.06939* (2015).
- [23] W.-C. Kang, J. McAuley, Self-attentive sequential recommendation, in: *2018 IEEE International Conference on Data Mining (ICDM), IEEE, 2018*, pp. 197–206.
- [24] F. Sun, J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, P. Jiang, Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer, in: *Proceedings of the 28th ACM international conference on information and knowledge management, 2019*, pp. 1441–1450.
- [25] H. Hu, X. He, J. Gao, Z.-L. Zhang, Modeling personalized item frequency information for next-basket recommendation, in: *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020*, pp. 1071–1080.
- [26] P. Ren, Z. Chen, J. Li, Z. Ren, J. Ma, M. De Rijke, Repeatnet: A repeat aware neural recommendation machine for session-based recommendation, in: *Proceedings of the AAAI Conference on Artificial Intelligence, 01, 2019*, pp. 4806–4813.
- [27] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, T.-S. Chua, Neural collaborative filtering, in: *Proceedings of the 26th international conference on world wide web, 2017*, pp. 173–182.
- [28] D. Liang, R. G. Krishnan, M. D. Hoffman, T. Jebara, Variational autoencoders for collaborative filtering, in: *Proceedings of the 2018 world wide web conference, 2018*, pp. 689–698.
- [29] P. Jacsó, Google scholar: the pros and the cons, *Online information review* 29 (2005) 208–214.
- [30] Z. Gao, Z. Cheng, F. Pérez, J. Sun, M. Volkovs, Mcl: Mixed-centric loss for collaborative filtering, in: *Proceedings of the ACM Web Conference 2022, 2022*, pp. 2339–2347.
- [31] C.-K. Hsieh, L. Yang, Y. Cui, T.-Y. Lin, S. Belongie, D. Estrin, Collaborative metric learning, in: *Proceedings of the 26th international conference on world wide web, 2017*, pp. 193–201.

- [32] J. B. Schafer, D. Frankowski, J. Herlocker, S. Sen, Collaborative filtering recommender systems, in: *The adaptive web*, Springer, 2007, pp. 291–324.
- [33] Y. Shen, Y. Wu, Y. Zhang, C. Shan, J. Zhang, B. K. Letaief, D. Li, How powerful is graph convolution for recommendation?, in: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 1619–1629.
- [34] T. Schnabel, M. Wan, L. Yang, Situating recommender systems in practice: Towards inductive learning and incremental updates, *arXiv preprint arXiv:2211.06365* (2022).
- [35] Y. Wu, Q. Cao, H. Shen, S. Tao, X. Cheng, Inmo: A model-agnostic and scalable module for inductive collaborative filtering, in: *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, pp. 91–101.
- [36] M. Ananyeva, O. Lashinin, V. Ivanova, S. Kolesnikov, D. I. Ignatov, Towards interaction-based user embeddings in sequential recommender models, in: J. Vinagre, M. Al-Ghossein, A. M. Jorge, A. Bifet, L. Peska (Eds.), *Proceedings of the 5th Workshop on Online Recommender Systems and User Modeling co-located with the 16th ACM Conference on Recommender Systems, ORSUM@RecSys 2022, Seattle, WA, USA, September 23rd, 2022*, volume 3303 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022. URL: <https://ceur-ws.org/Vol-3303/paper10.pdf>.
- [37] H. Steck, Embarrassingly shallow autoencoders for sparse data, in: *The World Wide Web Conference*, 2019, pp. 3251–3257.
- [38] X. Ning, G. Karypis, Slim: Sparse linear methods for top-n recommender systems, in: *2011 IEEE 11th international conference on data mining, IEEE*, 2011, pp. 497–506.
- [39] B. Paudel, F. Christoffel, C. Newell, A. Bernstein, Updatable, accurate, diverse, and scalable recommendations for interactive applications, *ACM Transactions on Interactive Intelligent Systems (TiiS)* 7 (2016) 1–34.
- [40] V. W. Anelli, A. Bellogín, T. Di Noia, D. Jannach, C. Pomo, Top-n recommendation algorithms: A quest for the state-of-the-art, *arXiv preprint arXiv:2203.01155* (2022).
- [41] S. Rendle, W. Krichene, L. Zhang, J. Anderson, Neural collaborative filtering vs. matrix factorization revisited, in: *Fourteenth ACM conference on recommender systems*, 2020, pp. 240–248.
- [42] V. W. Anelli, A. Bellogín, A. Ferrara, D. Malitesta, F. A. Merra, C. Pomo, F. M. Donini, T. D. Noia, Elliot: A comprehensive and rigorous framework for reproducible recommender systems evaluation, in: F. Diaz, C. Shah, T. Suel, P. Castells, R. Jones, T. Sakai (Eds.), *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, ACM, 2021, pp. 2405–2414. URL: <https://doi.org/10.1145/3404835.3463245>. doi:10.1145/3404835.3463245.
- [43] H. Wang, F. Zhang, J. Wang, M. Zhao, W. Li, X. Xie, M. Guo, Ripplenet: Propagating user preferences on the knowledge graph for recommender systems, in: *Proceedings of the 27th ACM international conference on information and knowledge management*, 2018, pp. 417–426.
- [44] X. Yu, X. Ren, Y. Sun, Q. Gu, B. Sturt, U. Khandelwal, B. Norrick, J. Han, Personalized entity recommendation: A heterogeneous information network approach, in: *Proceedings of the 7th ACM international conference on Web search and data mining*, 2014, pp. 283–292.
- [45] H. Wang, M. Zhao, X. Xie, W. Li, M. Guo, Knowledge graph convolutional networks for recommender systems, in: *The world wide web conference*, 2019, pp. 3307–3313.

- [46] N. Neophytou, B. Mitra, C. Stinson, Revisiting popularity and demographic biases in recommender evaluation and effectiveness, in: European Conference on Information Retrieval, Springer, 2022, pp. 641–654.