

# Multi-Task Learning For Reduced Popularity Bias In Multi-Territory Video Recommendations

PHANIDEEP GAMPA\* and FARNOOSH JAVADI\*, Amazon Prime Video

BELHASSEN BAYAR, Amazon Prime Video

AINUR YESSINALINA, Amazon Prime Video

Various data imbalances that naturally arise in a multi-territory personalized recommender system can lead to a significant item bias for globally prevalent items. A locally popular item can be overshadowed by a globally prevalent item. Moreover, users' viewership patterns/statistics can drastically change from one geographic location to another which may suggest to learn specific user embeddings. In this paper, we propose a multi-task learning (MTL) technique, along with an adaptive upsampling method to reduce popularity bias in multi-territory recommendations. Our proposed framework is designed to enrich training examples with active users representation through upsampling, and capable of learning geographic-based user embeddings by leveraging MTL. Through experiments, we demonstrate the effectiveness of our framework in multiple territories compared to a baseline not incorporating our proposed techniques. Noticeably, we show improved relative gain of up to 65.27% in PR-AUC metric. A case study is presented to demonstrate the advantages of our methods in attenuating the popularity bias of global items.

CCS Concepts: • **Computing methodologies** → **Transfer learning**; *Supervised learning by classification*.

Additional Key Words and Phrases: neural networks, video recommendations, two-tower networks, multi-task learning

## ACM Reference Format:

Phanideep Gampa, Farnoosh Javadi, Belhassen Bayar, and Ainur Yessenalina. 2023. Multi-Task Learning For Reduced Popularity Bias In Multi-Territory Video Recommendations. 1, 1 (September 2023), 9 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

## 1 Introduction

In modern streaming services such as YouTube, Prime Video, Netflix and etc., recommender systems (RS) have become the cornerstone of user engagement by helping the users find relevant contents [4, 8]. ML-based RS are typically trained in a supervised way using the logs of previous user-item interactions as explicit or implicit feedback of user's interest [14, 22]. For RS with wide geographical coverage across multiple countries such as the Amazon Prime Video, theoretically-speaking, one can separately train multiple region-specific ML models to perform recommendations for different regions. Maintaining multiple separate models is challenging from the ML operational perspective. It is therefore, highly preferable for a large-scale recommendation system to create one global model that can learn to make useful recommendations across different geographical regions.

---

\*Both authors contributed equally to this research.  
Copyright held by the author(s).

---

Authors' addresses: Phanideep Gampa, [phanide@amazon.com](mailto:phanide@amazon.com); Farnoosh Javadi, [javadif@amazon.com](mailto:javadif@amazon.com), Amazon Prime Video; Belhassen Bayar, Amazon Prime Video, [bayarb@amazon.com](mailto:bayarb@amazon.com); Ainur Yessenalina, Amazon Prime Video, [yessenal@amazon.com](mailto:yessenal@amazon.com).

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

One common challenge while building ML-based RS is the various imbalances that naturally occur in the user-item interaction training data [12, 21]. In this work, we are particularly concerned with the data imbalances that exist across different geographical partitions of a global RS with worldwide coverage. This data imbalance can occur in any multi-territory setting. For instance, one video title can be the 10<sup>th</sup> popular title in one territory but still can have much more user interaction history than the most popular video title in other territories. Also, user interaction history data of a video title that is only streamed in one specific territory (i.e., a local item) can be significantly smaller than that of a video title that is streamed in multiple territories (i.e., a global item), even though the local item/title is much more popular in that specific territory. If these imbalances are not properly handled, the global RS will inevitably suppress the locally popular items at the expense of globally present items. In other words, a single global model would suffer from popularity biases towards globally present items.

Popularity bias arising from the data imbalance decreases the fairness of the recommendations [2, 3], specifically for the titles local to a certain country. And the bias would only amplify in a closed feedback setting [12, 44], where the titles present in a lot of territories are always ranked higher and the gap between the global and local titles would either remain the same or increase. To prevent the model from overfitting on the global titles and to improve the fairness of local titles, we propose the use of Hard Parameter Sharing based Multi Task Learning (MTL) [11] approach. MTL has strong regularisation properties [6, 36] that improves the generalisation power of the model by preventing overfitting on global titles. At the same time, MTL learns shared representations that can be used by countries with less data allowing the transfer of knowledge from countries with large data.

Rebalancing and regularisation are among the frequent techniques used for solving the unfairness issue arising from the data imbalance [12] in the field of recommendation. We address the popularity bias issue in our system by means of a data-centric and a model-centric approach, which are our key contributions, as follows: (1) A data upsampling strategy based on active users criteria to enrich the training data with more number of informative examples to learn from. (2) An MTL approach to learn the refined user representations at a territory group level to deal with the distributional differences of viewership patterns across different territories. (3) Through a case study, we show how MTL reduces the popularity bias of global titles by learning group-specific parameters from the conditional distributions at a territory group level.

Our baseline model architecture consists of a two-tower Neural Network (NN) model [24] (one tower for user, and one for title), trained to learn propensity score for a user’s first time content discovery by leveraging representation of user and title embeddings as model input; which are vended by an upstream representation learning feature store. The exact details are described in section 3.1. Through a set of experiments, we compared our baseline recommendation system to a model that incorporates the MTL approach and upsampling strategy. Our experimental results demonstrate the advantage of our proposed model which significantly improved the PR-AUC for the majority of territories, and also helped to reduce the popularity bias towards global titles.

## 2 Related Work

**Popularity Bias.** Popularity bias in ML-based recommender systems occurs when popular items are typically recommended even more frequently than their organic popularity level [2, 3]. Various methods have been proposed to mitigate the effects of popularity bias such as Regularization [1, 13, 28], Adversarial Learning [31] and Causal Graphs [39, 40].

**Unfairness.** A topic that is closely related to the popularity bias is the fairness of ML-based systems. Many ML-based systems have been shown to inadvertently discriminate against certain individuals or groups due to the intrinsic bias that exists in the training data [19, 32]. Various works have proposed mitigating the ML-induced unfairness by means

of techniques such as Rebalancing [5, 9, 20], Regularization [26, 27, 29], Adversarial Learning [7, 10, 18] and Causal Modeling [35, 42, 43].

**Multi-Task Learning.** Multi-task learning (MTL) has been utilized in many Machine Learning applications such as natural language processing and speech recognition [15, 16]. In the related works, MTL has been primarily shown to provide benefits such as improved regularisation capabilities along with efficient transfer learning [37, 41]. There are two main MTL model architectures in Deep neural networks: hard parameter sharing and soft parameter sharing. In hard parameter sharing, a number of hidden layers are shared among all tasks. These common layers are followed by a set of task-specific layers [23, 33]. In soft parameter sharing, on the other hand, each task has a separate model with separate parameters and a regularization term is used in order to encourage the parameters of different models to be close [34]. In this work, we selected the hard parameter sharing architecture because of its regularization capabilities and its advantages of refined task representation learning from a shared part. To the best of our knowledge, our work is the first application of MTL to mitigate the item popularity bias in a global RS.

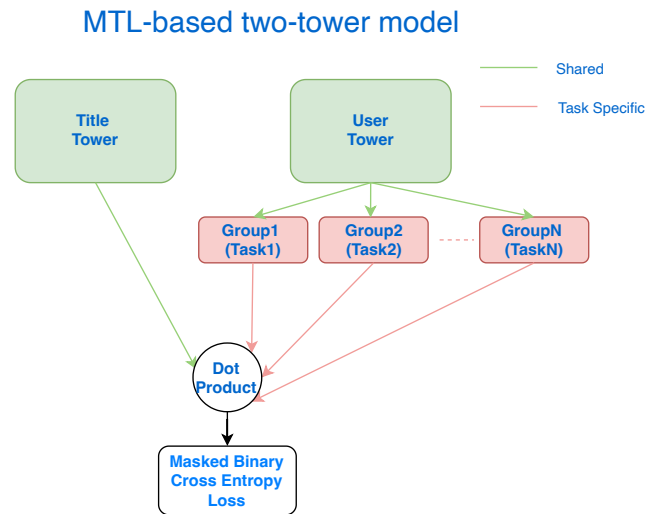


Fig. 1. Block diagram of the proposed MTL-based two-tower model where the user tower is branched out based on the territory grouping. Shared neural network blocks are in green and task-specific blocks are in red. Given a random batch, the loss function is a masked summation of binary cross entropy loss for each task.

### 3 Approach

#### 3.1 Input Embeddings

For all our experiments, we use user and title vector embeddings extracted from other pre-trained models. For the title specific data, we first extract 768 dimensional vectors using a fine-tuned BERT [17] applied on each title’s textual metadata such as name, genres, and synopsis. Next, several content-related categorical vectors are appended to each title’s BERT representation. For users, we use 512-dimensional latent vectors extracted from a ResNet architecture as described in [38]. This model was trained to encode contextual data of Prime Video users, including long-term and near-term streaming history, territory and other features.

### 3.2 Data Upsampling using Active Users

The data creation process generally comprises of selecting the set of users and a set of titles, and then generating the positive and negative user-title pairs. A random sampling of users from the customer universe across all the territories would be biased towards the territories with large number of users. At the same time, the set of active users can have minor representation than the non active streamers. Active users are those who watched at least one title in the recent few months. To reduce these disparities, we add a disjoint set of active users sampled independently to the random user set. This would enrich our training dataset with more number of active user embeddings. Also, this would improve the likelihood of having more number of positive examples in training particularly helping the minor territories.

### 3.3 MTL-based Two-Tower Model

To mitigate the issues arising from data imbalance and to better deal with the distributional differences across territories, we propose the use of MTL-based two-tower model. We use a hard parameter sharing [11] based architecture, where we branch out of the user tower to have task-specific layers, one for each territory group. For this approach, we first group the available 200+ territories into four groups based on the geo-partition they are present in. This criteria groups geographically closer territories together that can potentially have similar viewership patterns. At the same time, we also try to keep a good balance of major and minor territories in each group for efficient transfer learning. We refer a territory as major or minor based on its viewership volume.

We propose a neural network architecture that consists of two parts, shared blocks across all territories and a task-specific block for each territory group as depicted in the figure 1. Specifically, shared blocks are the user tower and title tower as that of the baseline model. We then add separate task/group-specific neural network blocks on top of the user tower. During training time, the shared blocks are always updated using back propagation for all the data points. The group-specific blocks are only updated for the data points coming from the respective group. Such MTL-based techniques are known for having the advantages of transfer learning – minor territories learning from the major territories and vice versa – happening through the shared part. At the same time, we also have independent blocks to take care of the distributional differences across the country groups. These task-specific blocks convert the shared user representations to the group-specific representation. This allows us to acknowledge the differences in the user viewership patterns across various territories and to localize the recommendations at a group level.

Table 1. Statistics of number of positive data points in some major territories signifying the imbalance of positive rewards.

|          |      |     |     |     |      |      |      |      |
|----------|------|-----|-----|-----|------|------|------|------|
|          | T1   | T2  | T3  | T4  | T5   | T6   | T7   | T8   |
| Order of | 1M   | 1M  | 1M  | 1M  | 1M   | 1M   | 100k | 100k |
|          | T9   | T10 | T11 | T12 | T13  | T14  | T15  | T16  |
| Order of | 100k | 10k | 10k | 10k | 100k | 100k | 100k | 100k |

## 4 Experiments

### 4.1 Data

For training and validation, we collected daily first time content discovery datasets for all the 200+ territories we are dealing with. First time content discovery corresponds to generating positive user-title pairs based on the user’s first streaming event of title items within a month. We created two versions of the training dataset where several million users (order of 10M) are sampled randomly and then added the set of active users using the upsampling process

Table 2. Relative gain of PR-AUC score in various territories  $T_1, T_2, \dots, T_{15}, T_{16}$  for the MTL model evaluated on three different test sets. Positive average results are in bold for each territory.

|                                       | Test Set | T1           | T2           | T3           | T4           | T5           | T6           | T7           | T8           |
|---------------------------------------|----------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MTL vs<br>Baseline                    | Test-1   | 9.88         | 34.05        | -7.72        | 11.36        | 15.04        | 27.12        | 10.10        | 21.36        |
|                                       | Test-2   | 35.30        | 65.27        | -21.18       | 24.58        | 24.73        | 23.22        | 0.73         | -0.35        |
|                                       | Test-3   | 31.99        | 16.36        | 11.39        | 18.98        | 9.49         | 20.69        | 33.67        | 5.35         |
|                                       | Average  | <b>25.72</b> | <b>38.56</b> | -5.83        | <b>18.30</b> | <b>16.42</b> | <b>23.67</b> | <b>14.83</b> | <b>8.78</b>  |
| MTL vs<br>Baseline with<br>upsampling | Test-1   | 7.27         | 23.46        | -30.16       | -3.98        | 4.62         | -16.52       | -0.65        | 6.30         |
|                                       | Test-2   | 9.80         | 20.36        | -5.11        | 7.59         | 11.78        | 5.05         | 2.95         | 3.79         |
|                                       | Test-3   | 18.31        | 8.17         | 0.84         | 6.89         | 9.92         | 6.69         | 13.17        | -3.23        |
|                                       | Average  | <b>11.79</b> | <b>17.33</b> | -11.48       | <b>3.49</b>  | <b>8.77</b>  | -1.58        | <b>5.15</b>  | <b>2.28</b>  |
|                                       | Test Set | T9           | T10          | T11          | T12          | T13          | T14          | T15          | T16          |
| MTL vs<br>Baseline                    | Test-1   | 5.21         | 13.43        | 7.68         | 54.75        | 10.45        | 7.87         | 5.51         | 7.36         |
|                                       | Test-2   | -5.85        | 46.97        | 12.88        | 21.64        | 36.42        | 32.31        | 20.20        | 25.68        |
|                                       | Test-3   | 3.75         | -2.92        | 38.39        | 45.69        | 23.06        | 12.61        | 20.40        | 12.70        |
|                                       | Average  | <b>1.03</b>  | <b>19.16</b> | <b>19.65</b> | <b>40.69</b> | <b>23.31</b> | <b>17.59</b> | <b>15.37</b> | <b>15.24</b> |
| MTL vs<br>Baseline with<br>upsampling | Test-1   | 3.91         | 15.52        | 28.45        | 77.35        | 3.38         | 2.39         | 2.41         | 0.48         |
|                                       | Test-2   | 9.47         | 69.10        | 16.07        | 15.89        | 10.28        | 18.59        | 3.86         | 0.13         |
|                                       | Test-3   | -6.86        | 0.25         | 22.26        | 10.85        | 13.68        | 6.79         | 20.35        | 13.43        |
|                                       | Average  | <b>2.17</b>  | <b>28.29</b> | <b>22.26</b> | <b>34.70</b> | <b>9.11</b>  | <b>9.25</b>  | <b>8.88</b>  | <b>4.68</b>  |

Table 3. Title level PR-AUC gain for a set of local and global titles.

| Title          | MTL vs Baseline | MTL vs Baseline Upsampled |
|----------------|-----------------|---------------------------|
| Local Title 1  | 6.50            | 2.63                      |
| Local Title 2  | 3.83            | 6.23                      |
| Global Title 1 | 4.28            | 3.15                      |
| Global Title 2 | 36.22           | 35.27                     |

described in section 3.2. The training data consists of several hundred titles that were promoted within the past several months spread across various territories. Overall, the user-title pairs for training are accumulated over several months time period. The number of negative samples is chosen to be  $\alpha$  times bigger than the number of positives, where  $\alpha$  is approximately equal to the ratio of daily first time streamers among all Prime Video users. As given in table 1, we can observe the data imbalance of positive rewards for all the titles in 16 major territories out of the total 200+. This in turn would result in disparity between titles that are present in multiple territories and titles that are present in few territories. We evaluate our models on three disjoint test data sets for three consecutive days post the end date of the training window. We sample a new set of users disjoint from the train in order of 10 millions for preparing the test data.

#### 4.2 Experiment Setup

For both the global and the MTL-based model, we use two dense layers along with batch-norm [25] for each tower encoder. For task-specific layers in MTL, we use one dense layer for each task block. We used ReLU activation functions for all layers. We kept the hidden dimension as  $d = 512$  for all of the layers. For training, we use Adam [30] optimiser

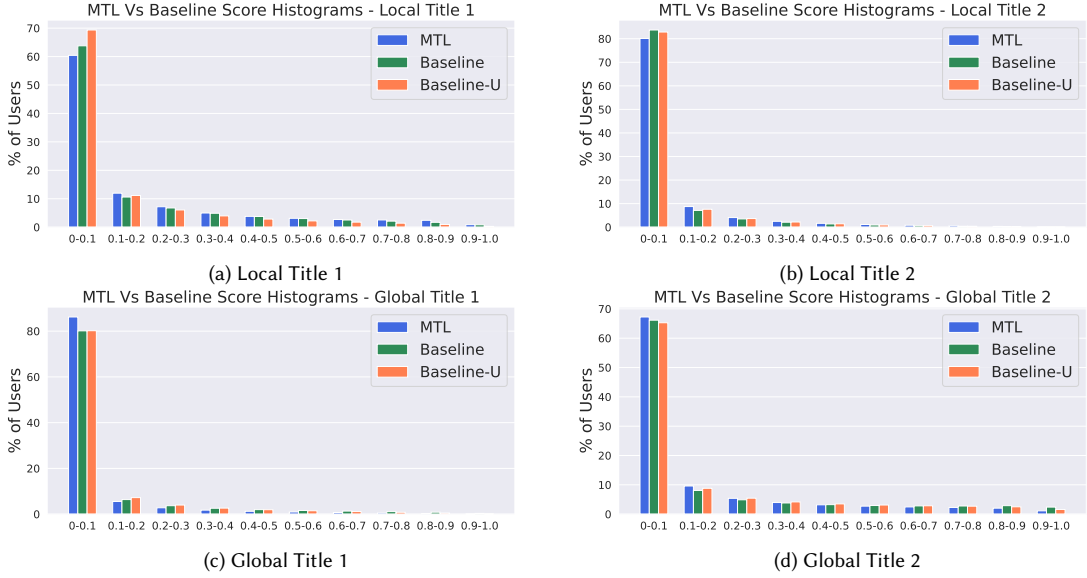


Fig. 2. Each subfigure shows the normalised user histograms of the model propensity scores of all the three models for a few titles. We can observe that the MTL-based model is giving relatively high scores for the local titles, and relatively low scores for the global titles attenuating their popularity bias.

with a learning rate of 0.001 and a batch size of 8192. The models are trained end-to-end using the binary cross entropy loss as we are dealing with binary preferences. For MTL-based experiments, we use random batch strategy where each batch can contain examples from all the tasks. We use one-hot encoded mask while calculating the loss to ensure that only data points from the respective territory group can update the group-specific blocks, blocking the gradients of other blocks. We train the baseline model on both the normal and the upsampled data created using the approach described in section 4.1. For MTL, we report the metrics for the model trained on upsampled data. To adapt the models to the daily patterns, we incrementally train all the models on  $i - 1^{th}$  day train data before evaluating on the  $i^{th}$  day test data.

### 4.3 Results

We compare the performance of MTL-based model against a baseline without MTL. As described in section 1, our baseline model is a single global two-tower based model that finds the affinity score of a title given user using the dot product similarity. We chose PR-AUC metric for evaluation as it is more sensitive to the positive class performance, and is highly relevant when fraction of positives are much less than the negatives. As shown in the table 2, MTL-based model consistently performs better than the single global model trained on both the normal and upsampled data in most of the territories. If we compare relative gains of MTL against the two baselines, we can observe that the baseline with upsampling improved over the baseline in a good number of territories. MTL performing better than the baseline trained on upsampled data demonstrates the advantages of the refined user representations that the MTL-based model learns.

#### 4.4 Case Study - Attenuating global popularity bias with improved local title affinity scores

To further demonstrate the effectiveness of MTL, we analyze the user-title propensity score distributions for few selected global and local titles with comparable positive examples from two territories. We pick *Local Title 1* and *Global Title 1* from territory  $T_i$  and *Local Title 2* and *Global Title 2* from territory  $T_j$ .

First, we report the overall gain in title level PR-AUC score of the MTL-based model over both the baselines in table 3. Then based on the normalised user histogram of output score bins given in figure 2, we try to answer these two questions: For local titles with comparable positive examples that of a global title, are we moving the score distributions towards the higher side? Yes, MTL-based model relatively moves more users from lower score bins to higher score bins as given in figures 2a and 2b. For global titles with comparable positive examples that of a local title, are we moving the score distributions towards the lower side? Yes, the MTL-based model relatively moves more number of users to the lower score bins from higher score bins attenuating the popularity bias as given in figures 2c and 2d.

## 5 Conclusion

In this paper, we propose an upsampling strategy based on active streamers strategy to deal with data imbalance. We also propose the use of an MTL-based architecture for learning refined user representations to better handle the distributional differences when dealing with multi-territory video recommendations. Through our experiments, we demonstrated the advantage of MTL-based architecture over a single global model trained on data from all the territories. Through our case study, we analysed the changes in the output score distributions of local and global titles thereby showing the advantages of the proposed techniques for attenuating the popularity bias. Overall, the proposed model achieves up to 65.27% improvement in PR-AUC when evaluated on 16 major territories out of 200 overall territories.

## References

- [1] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. Controlling popularity bias in learning-to-rank recommendation. In *Proceedings of the eleventh ACM conference on recommender systems*, pages 42–46, 2017.
- [2] Himan Abdollahpouri and Masoud Mansoury. Multi-sided exposure bias in recommendation. *arXiv preprint arXiv:2006.15772*, 2020.
- [3] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. The unfairness of popularity bias in recommendation. *arXiv preprint arXiv:1907.13286*, 2019.
- [4] Xavier Amatriain and Justin Basilico. Recommender systems in industry: A netflix case study. In *Recommender systems handbook*, pages 385–419. Springer, 2015.
- [5] Abolfazl Asudeh, HV Jagadish, Julia Stoyanovich, and Gautam Das. Designing fair ranking schemes. In *Proceedings of the 2019 international conference on management of data*, pages 1259–1276, 2019.
- [6] Jonathan Baxter. A bayesian/information theoretic model of learning to learn via multiple task sampling. *Machine learning*, 28(1):7–39, 1997.
- [7] Ghazaleh Beigi, Ahmadreza Mosallanezhad, Ruocheng Guo, Hamidreza Alvari, Alexander Nou, and Huan Liu. Privacy-aware recommendation with private-attribute protection using adversarial learning. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 34–42, 2020.
- [8] Shlomo Berkovsky, Ronnie Taib, and Dan Conway. How to recommend? user trust factors in movie recommender systems. In *Proceedings of the 22nd international conference on intelligent user interfaces*, pages 287–300, 2017.
- [9] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*, pages 405–414, 2018.
- [10] Avishek Bose and William Hamilton. Compositional fairness constraints for graph embeddings. In *International Conference on Machine Learning*, pages 715–724. PMLR, 2019.
- [11] R Caruana. Multitask learning: A knowledge-based source of inductive bias1. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 41–48. Citeseer, 1993.
- [12] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. Bias and debias in recommender system: A survey and future directions. *arXiv preprint arXiv:2010.03240*, 2020.
- [13] Zhihong Chen, Rong Xiao, Chenliang Li, Gangfeng Ye, Haochuan Sun, and Hongbo Deng. Esam: Discriminative domain adaptation with non-displayed items to improve long-tail performance. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 579–588, 2020.

- [14] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, pages 7–10, 2016.
- [15] Ronan Collobert and Jason Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167, 2008.
- [16] Li Deng, Geoffrey Hinton, and Brian Kingsbury. New types of deep neural network learning for speech recognition and related applications: An overview. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 8599–8603. IEEE, 2013.
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [18] Harrison Edwards and Amos Storkey. Censoring representations with an adversary. *arXiv preprint arXiv:1511.05897*, 2015.
- [19] Batya Friedman and Helen Nissenbaum. Bias in computer systems. *ACM Transactions on information systems (TOIS)*, 14(3):330–347, 1996.
- [20] Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th acm sigkdd international conference on knowledge discovery & data mining*, pages 2221–2231, 2019.
- [21] Elizabeth Gómez. Characterizing and mitigating the impact of data imbalance for stakeholders in recommender systems. In *Fourteenth ACM Conference on Recommender Systems*, pages 756–757, 2020.
- [22] Carlos A Gomez-Urbe and Neil Hunt. The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4):1–19, 2015.
- [23] Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. A joint many-task model: Growing a neural network for multiple nlp tasks. *arXiv preprint arXiv:1611.01587*, 2016.
- [24] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 2333–2338, 2013.
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [26] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Enhancement of the neutrality in recommendation. In *Decisions@ RecSys*, pages 8–14. Citeseer, 2012.
- [27] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Efficiency improvement of neutrality-enhanced recommendation. In *Decisions@ RecSys*, pages 1–8. Citeseer, 2013.
- [28] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Correcting popularity bias by enhancing recommendation neutrality. In *RecSys Posters*, 2014.
- [29] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Issei Sato. Model-based approaches for independence-enhanced recommendation. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, pages 860–867. IEEE, 2016.
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Adit Krishnan, Ashish Sharma, Aravind Sankar, and Hari Sundaram. An adversarial approach to improve long-tail performance in neural collaborative filtering. In *Proceedings of the 27th ACM International Conference on information and knowledge management*, pages 1491–1494, 2018.
- [32] Kun Lin, Nasim Sonboli, Bamshad Mobasher, and Robin Burke. Crank up the volume: preference bias amplification in collaborative recommendation. *arXiv preprint arXiv:1909.06362*, 2019.
- [33] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Philip S Yu. Learning multiple tasks with multilinear relationship networks. *Advances in neural information processing systems*, 30, 2017.
- [34] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3994–4003, 2016.
- [35] Razieh Nabi and Ilya Shpitser. Fair inference on outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [36] Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.
- [37] Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI global, 2010.
- [38] Jianling Wang, Ainur Yessenalina, and Alireza Roshan-Ghias. Exploring heterogeneous metadata for video recommendation with two-tower model. *arXiv preprint arXiv:2109.11059*, 2021.
- [39] Wenjie Wang, Fuli Feng, Xiangnan He, Xiang Wang, and Tat-Seng Chua. Deconfounded recommendation for alleviating bias amplification. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1717–1725, 2021.
- [40] Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1791–1800, 2021.
- [41] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):1–40, 2016.
- [42] Yongkai Wu, Lu Zhang, and Xintao Wu. On discrimination discovery and removal in ranked data using causal graph. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2536–2544, 2018.
- [43] Yongkai Wu, Lu Zhang, and Xintao Wu. Counterfactual fairness: Unidentification, bound and algorithm. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019.

- [44] Ziwei Zhu, Yun He, Xing Zhao, and James Caverlee. Popularity bias in dynamic recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2439–2449, 2021.