
Learning Federated Representations and Recommendations with Limited Negatives

Lin Ning
Google Research
linning@google.com

Karan Singhal
Google Research
karansinghal@google.com

Ellie X. Zhou
Google
ezhou@google.com

Sushant Prakash
Google Research
sush@google.com

Abstract

Deep retrieval models are widely used for learning entity representations and recommendations. Federated learning provides a privacy-preserving way to train these models without requiring centralization of user data. However, federated deep retrieval models usually perform much worse than their centralized counterparts due to non-IID (independent and identically distributed) training data on clients, an intrinsic property of federated learning that limits negatives available for training. We demonstrate that this issue is distinct from the commonly studied client drift problem. This work proposes *batch-insensitive losses* as a way to alleviate the non-IID negatives issue for federated movie recommendations. We explore a variety of techniques and identify that batch-insensitive losses can effectively improve the performance of federated deep retrieval models, increasing the relative recall of the federated model by up to 93.15% and reducing the relative gap in recall between it and a centralized model from 27.22% - 43.14% to 0.53% - 2.42%. We also open-source our code framework to accelerate further research and applications of federated deep retrieval models.

1 Introduction

Recent years have witnessed the successes of deep retrieval models in many large-scale recommendation systems [2, 13, 19, 17, 7, 16] and natural language tasks [5, 4, 18, 1]. While these models largely improve user experience by enabling personalization and representation learning, they can also raise privacy concerns as user data typically needs to be sent to a centralized server for training.

Federated learning (FL) is a decentralized training strategy in which clients collaborate with a coordinating server to train a machine learning model [14]. It provides opportunities to leverage distributed client data to learn useful models while preserving user privacy. Bringing deep retrieval models to FL is a promising way to power recommendations and learn representations for users and items while addressing privacy concerns by reducing the centralization of user data.

In this work, we observe a challenge with training federated deep retrieval models: the insufficiency of negative examples available on a client’s device. A deep retrieval model (shown in Figure 1) learns embeddings representing users’ contexts and items like movies, songs, or websites. For the model to learn meaningful embeddings, it generally uses two types of examples: positive examples and negative examples. The positive examples pull embeddings of the training (*context, item*) pairs to be close together in an embedding space, while the negative examples push embeddings of unrelated pairs farther apart. Negative examples are typically required to prevent *embedding space collapse*, where learned representations collapse to a single point and are no longer informative.

Typically, negative examples are produced by sampling items from training data. In centralized training, training data is usually assumed to be independent and identically distributed (IID). Therefore, negatives can be sampled reliably from the overall training distribution. However, the IID data




Figure 1: An illustration of a deep retrieval model.

assumption does not hold for federated learning since clients generate data locally based on their circumstances. That means, on each device, negatives may not be present. Even if they are present, they may be relatively few and relatively similar. In practice, we observe that this leads to significant performance degradation when FL is applied naively, beyond the typical degradation observed for *e.g.*, classification in the FL setting.

This work focuses on understanding and alleviating the non-IID issue for deep retrieval models. We make the following **key contributions**:

- Observe that naive federated training of deep retrieval models causes an unusually steep performance drop. Show that performance degradation is primarily caused by sampling negatives from non-IID data, not the typical client drift issue or other aspects of federated training.
- Introduce *batch-insensitive losses* as a class of objectives to alleviate the issue in this setting.
- Perform empirical evaluation of different training objectives in a movie recommendation setting, showing that batch-insensitive losses enable performant federated deep retrieval.
- Release an open-source framework to accelerate further research and practical applications of federated deep retrieval models.¹

Related Work: Most previous research has focused on mitigating the non-IID data issue [6, 23, 10, 11, 12, 21, 20] for improving general model convergence in federated learning settings. These works study the general problem of heterogeneity in client data causing *client drift* when clients perform multiple local computation steps, as in FEDAVG [9, 8]. However, some of them require sharing some training data across clients [6, 23], and most of them study image classification tasks. The developed techniques are not fully relevant to the deep retrieval model, which sees unusually severe performance degradation due to explicit reliance on sampling negatives for training. Though also caused by non-IID data, this issue is orthogonal to the client drift issue: it occurs even when clients perform one local update step. We can combine the techniques explored in this work with previous techniques for addressing client drift. [21] also studies the effectiveness of hinge loss and spreadout regularizer in federated learning. However, it focuses on classification tasks and only considers extreme cases where each user has no access to negative examples. Our work focuses on a more realistic setting, aiming to understand and increase the performance of representation learning and item recommendation with deep retrieval models when each client has access to some but limited and relatively similar negative examples. More broadly, we propose batch-insensitive losses, which can be generally valuable for alleviating non-IID data issues in other federated learning tasks.

2 Federated Deep Retrieval Model

Deep retrieval models are also referred to as dual encoder, two-tower, or encoder-encoder models depending on the setting [4, 18, 1, 15, 3]. This general framework has been successfully deployed in a variety of real-world applications in embedding and recommendation learning. As illustrated in Figure 1, a deep retrieval model consists of two encoders, each of which can be a fully-connected network, convolutional neural network, Transformer, and so on, depending on the task. The input consists of (*context*, *item*) pairs encoded by the left and right encoders, respectively. For example,

¹https://git.io/federated_dual_encoder

in a movie recommendation use-case, a context might be a sequence of previous movies a user has watched, and the item is the next movie they watch. The goals of applying a deep retrieval model in this setting would be to learn a model that can predict the next movie given an unseen sequence of previous movies and learn encoders that produce embeddings for users and movies.

More formally, we denote the feature vectors representing contexts and items as \mathbf{x} and \mathbf{y} . The two encoders are denoted as two parameterized functions $f(\cdot)$ and $g(\cdot)$, with $f(\mathbf{x})$ and $g(\mathbf{y})$ mapping \mathbf{x} and \mathbf{y} to a shared embedding space. The model outputs the similarity score between the encoded context and item, *e.g.*, the inner product of context and item embeddings, $s(\mathbf{x}, \mathbf{y}) = \langle f(\mathbf{x}), g(\mathbf{y}) \rangle$. A loss function is applied to enforce that positive examples (*i.e.*, similar context and item pairs) have high similarity, and negative examples have low similarity. Once the parameterized functions $f(\cdot)$ and $g(\cdot)$ are learned, the model can predict relevant items given a new context. The representations produced by $f(\cdot)$ and $g(\cdot)$ are general representations of user contexts and items and can also be used for other downstream applications, such as classification [3, 1].

Loss Function: The loss function plays a key role in training a deep retrieval model. The most commonly used one [3, 17, 19] is softmax cross-entropy loss over similarities: $\ell(\mathbf{X}_i, \mathbf{Y}_i) = -\log(e^{s(\mathbf{X}_i, \mathbf{Y}_i)} / \sum_{\mathbf{Y}_j \in \mathcal{N}} e^{s(\mathbf{X}_i, \mathbf{Y}_j)})$. \mathbf{X} and \mathbf{Y} represent all the contexts and items in a batch, and \mathcal{N} is a set of negative labels used to construct negative example pairs $(\mathbf{X}_i, \mathbf{Y}_j)$. The model is incentivized to maximize the similarity between positive example pairs and minimize the similarity between negative example pairs. Note that if it only did the former, then the embedding space produced by $f(\cdot)$ and $g(\cdot)$ would collapse. Therefore, negative examples are important to learn a good model. A standard method for getting negatives is to use in-batch negatives [19, 17], which means given a training batch and any $(\mathbf{X}_i, \mathbf{Y}_i)$ pair in the batch, all other items $\mathbf{Y}_{j, j \neq i}$ in the same batch are treated as negatives for \mathbf{X}_i . We refer to this as *batch softmax* below.

Federated training of a deep retrieval model involves three main steps in each training round. First, a central server sends the current model to several randomly sampled clients. Second, each sampled client trains on its dataset and updates its model locally. Finally, the local model updates are sent back to the server and aggregated to update the server model. The de-facto standard federated optimization method is FEDAVG [14]. Each client updates its model multiple times before sending the model update back to the server to be averaged. We also later refer to FEDSGD, where each client only runs a single local training step at each round, similar to standard distributed training.

As discussed in Section 4.2, naively applying federated learning to this setting produces a steep performance drop, worse than typical when comparing centralized and federated performance. We will show that training degrades due to non-IID negatives. Note that this is distinct from the *client drift* phenomenon discussed in other works [9], which causes slight performance deterioration when clients take multiple local steps. In contrast, the problem we observe occurs even in the FEDSGD regime (see Figure 3). This work aims to characterize this problem better and propose methods that enable federated deep retrieval models to perform comparably to centralized counterparts.

3 Batch-Insensitive Losses

We propose *batch-insensitive losses* as a potential solution to address the non-IID data issue.

Definition 1 (Batch-Insensitive Loss). *Given a batch of N examples and two parameterized functions $f(\cdot)$ and $g(\cdot)$, with all contexts and items in the batch denoted as \mathbf{X} and \mathbf{Y} , a loss function is batch-insensitive if it satisfies*

$$\ell_{BI}(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{i=0}^N \ell_{BI}(f(\mathbf{X}_i), g(\mathbf{Y}_i)). \quad (1)$$

It follows that applying a batch-insensitive loss over several batches of data in parallel (not in sequence) produces the same average loss and gradient update no matter how the examples are batched. We use this to show that we can produce the same gradient update between federated and centralized learning, providing a natural justification for batch-insensitive losses.

Proposition 1. *Let \mathcal{C} be a collection of clients sampled at round k for federated learning. Denote the aggregated update to the model for that round under FEDSGD as $\Delta_{k, \text{fedsgd}}(\mathcal{C})$. Let \mathcal{E} be the collection of all training examples from clients in \mathcal{C} . Denote the model update of centralized training*

using SGD with all examples in \mathcal{E} in a batch is $\Delta_{k,sgd}(\mathcal{E})$. Using the same batch-insensitive loss ℓ_{BI} as training objective, the same model initialization Θ , and the same learning rate η , we have

$$\Delta_{k,fedsgd}(\mathcal{C}|\ell_{BI}, \Theta) \equiv \Delta_{k,sgd}(\mathcal{E}|\ell_{BI}, \Theta) \quad (2)$$

See Appendix A for proof. Proposition 1 shows that FEDSGD with ℓ_{BI} approximates standard large-batch SGD when the number of clients per round is large (so \mathcal{E} is representative of the centralized data). This motivates using batch-insensitive losses to mitigate the non-IID data issue.

Note that batch softmax loss does not have these properties and is batch-sensitive. We now describe specific batch-insensitive losses for the deep retrieval setting.

3.1 Hinge Loss + Spreadout Regularizer

This loss is a variation of contrastive loss, composed of a positive term pushing positive examples together (hinge loss) and a negative term preventing embedding collapse (spreadout regularization). This combination was first introduced in [21].

Hinge Loss: Given a positive example pair (\mathbf{x}, \mathbf{y}) , where \mathbf{x} and \mathbf{y} are context and item features, the hinge loss is defined as $\ell(\mathbf{x}, \mathbf{y}) = \max(0, \beta - f(\mathbf{x}) \cdot g(\mathbf{y}))^2$, where $f(\mathbf{x})$ is the context embedding, $g(\mathbf{y})$ is the item embedding, and β is a tunable margin set to 0.9 in this work.

Spreadout Regularization: Spreadout regularizer [22] maximizes the spread of embeddings in an embedding space. Given an embedding vocabulary V and the corresponding embedding weights W , spreadout regularizer can be formulated as $\ell_{sr}(W) = \sum_{v \in V} \sum_{v' \neq v} (-d^2(w_v, w_{v'}))$, where d is a measure of distance, e.g., Euclidean distance or negative dot product. When combined with L2 normalization, the objective pushes embeddings in W apart on a hypersphere. It can be used with any loss function in the form of $\ell_{(\cdot),sr} = \ell_{(\cdot)} + \alpha \ell_{sr}$, where $\ell_{(\cdot)}$ is the original loss function, ℓ_{sr} is the spreadout regularizer, and α trades off the regularization term and the original loss.

Unlike softmax cross-entropy, hinge loss only considers positive examples; the loss can be trivially minimized by collapsing embeddings into a single point. To avoid collapse, we apply spreadout regularizer to the model’s shared embedding table described in Section 4.1, pushing items in the embedding vocabulary to have orthogonal embeddings. The resulting combined loss pushes positive pairs closer while pushing negatives apart, resulting in a full loss for deep retrieval. Both hinge loss and spreadout regularizer are batch-insensitive based on Definition 1. It is easy to show that a linear combination of them is also batch-insensitive.

3.2 Global Softmax

The *global softmax* loss is defined as $\ell(\mathbf{X}_i, \mathbf{Y}_i) = -\log(e^{f(\mathbf{X}_i) \cdot g(\mathbf{Y}_i)} / \sum_{\mathbf{Y}_j \in V} e^{f(\mathbf{X}_i) \cdot g(\mathbf{Y}_j)})$, where V is the vocabulary of possible items to predict. Global softmax is an extreme case of negative sampling. For any $(\mathbf{X}_i, \mathbf{Y}_i)$ pair, all the items $\mathbf{Y}_{j,j \neq i} \in V$ are treated as negatives for \mathbf{X}_i , and are used to calculate the softmax loss. For a larger scale model, only a subset of items can be sampled randomly instead. In either case, the loss is batch-insensitive on expectation according to Definition 1.

4 Evaluation

We evaluate the efficacy of batch-insensitive losses on the non-IID negatives issue with a movie recommendation task. We train and evaluate a deep retrieval model on the MovieLens 1M dataset². The model takes in a user’s movie-watching history and predicts a relevant next movie for this user.

4.1 Experiment Setting

Below we describe details on the dataset, model, and tasks. Additional details can be found in our open-source code framework for federated deep retrieval (see Section 5).

Dataset: As shown in Table 1, the MovieLens 1M dataset contains approximately 1 million ratings from 6040 users on 3952 movies. Examples are created by taking moving "windows" of the movie

²<https://grouplens.org/datasets/movielens/1m/>

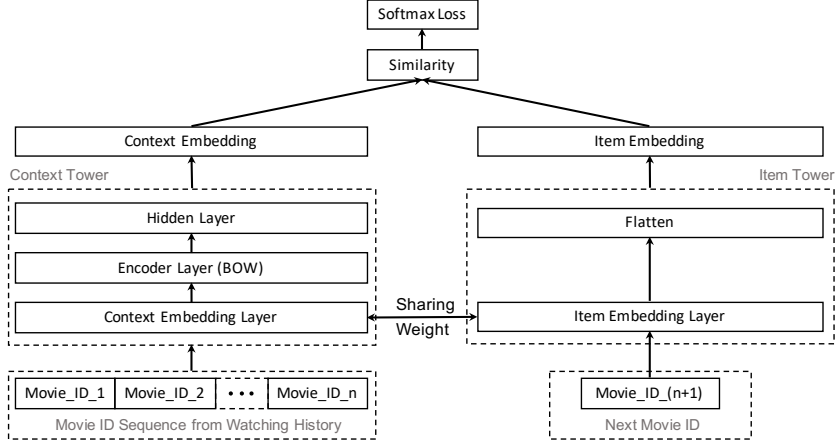


Figure 2: ID-based deep retrieval model for movie recommendation.

sequence (sorted by timestamps) for each user, resulting in context inputs containing ten movie IDs and item inputs representing one next movie ID. For centralized training, examples are randomly shuffled across all users and split to train and test datasets. The train dataset has 894,752 examples, and the test dataset has 99,417 examples. We refer to these as *centralized datasets* later in this section. For federated training, all examples are grouped by user, forming a natural data partitioning across clients. The train and test examples are split by user ids, resulting in 4832 train, 603 validation, and 605 test users. We refer to these as *federated datasets*. We sample 100 clients for each training round.

Model Architecture: Figure 2 illustrates the architecture of the ID-based deep retrieval model used for experiments. It takes a sequence of movie IDs (the movie watching history) as the context and the next movie ID as the item to form the $(context, item)$ pair. The context encoder is a bag-of-words encoder, and the item tower is a simple embedding lookup tower. The two towers share the same bottom embedding layer, which maps from movie ID to dense embedding. The two towers generate context and item embeddings respectively, and similarity is enforced between the positive (context, label) pairs. To make the comparison consistent and fair, we set the batch size to 16 for all experiments. The output dimension of the shared embedding layer is 16. The encoded context embedding and item embedding are also 16-dimensional and L2-normalized.

Federated and Centralized Experiments: One of our goals in this work is to reduce the gap between federated and centralized deep retrieval model performance. We run experiments for both centralized and federated training and compare performance. Interestingly, we observe that batch-insensitive losses can also improve centralized performance, so we also compare against this. We refer to these models as *Improved Centralized* later in this section. In both settings, we measure test recall@k for $k \in [1, 5, 10]$, the fraction of examples for which the correct next movie is within the top k nearest item embeddings for an unseen context.

4.2 Effect of Non-IID Data

To study the effect of non-IID data on federated learning in this setting, we train the ID-based deep retrieval model on federated datasets using FEDAVG [14]. The model is trained with the standard batch softmax cross-entropy loss. We compare recall across items within a batch (batch recall) with centralized training. As shown in Figure 4, the federated model experiences significant performance degradation, especially for recall@1. It is worth noting that performance degradation occurs even when training with FEDSGD, as illustrated in Figure 3. It indicates that client drift, which occurs when clients take multiple local steps with non-IID data, is not the only cause.

To test whether non-IID data causes this performance degradation, we train a *Federated Shuffled* model where data is IID across clients. In this experiment, all examples are shuffled across users while the number of examples per user remains the same. This ensures the same number of local steps taken on each client as before and isolates the effect of non-IID data on federated learning. As shown in Figure 4, the Federated Shuffled result roughly matches the centralized result. We then conclude that non-IIDness causes the performance degradation, not federated training in itself.

Table 1: MovieLens 1M Dataset Statistics. ‘E’ is short for ‘Examples’ and ‘U’ is short for ‘Users’.

	Num	Centralized	Federated
Ratings	1,000,209		
Users	6040		
Movies	3952		
Split Strategy		By Example	By User
Train Data		894,753 E	4832 U
Test Data		99,417 E	605 U




Figure 3: Model performance with batch softmax loss. Federated training uses FEDSGD.




Figure 4: Model performance with batch softmax loss. Federated training uses FEDAVG.

However, although shuffling examples across users could resolve the non-IID data issue, we cannot do it in practice due to privacy and communication limitations.

4.3 Federated and Centralized Results

This section studies the model performance with four different loss functions: batch softmax (BS), batch softmax with spreadout regularizer (BS+S), and two batch-insensitive losses (see Section 3): hinge loss with spreadout (H+S) and global softmax (GS). All the federated models are trained with FEDAVG. We compare recall calculated globally across all items, which has no dependence on examples in a batch, enabling fair comparison.

Batch Softmax (BS): Batch softmax is the standard loss function for training a deep retrieval model. It is calculated with in-batch negatives as described in Section 2. Figure 5(a) shows a large gap between centralized and federated global recalls, similar to the batch recall results in Section 4.2.

Batch Softmax + Spreadout (BS+S): Figure 5(b) presents results of training with batch softmax combined with spreadout regularization. With spreadout regularizer, the recall values of the federated model almost match those of the batch softmax centralized model, which is trained without spreadout regularizer. However, spreadout regularizer also improves centralized training. The federated model still performs significantly worse compared to the improved centralized model.

The results indicate that spreadout regularization itself is not enough to solve the issue with the non-IID negatives. Although spreadout regularizer pushes embeddings of unrelated pairs farther apart, batch softmax loss still depends on in-batch negatives and can still lead to worse model quality. Therefore, we need a loss function less affected by the training data distribution, motivating batch-insensitive losses.

Batch-Insensitive Losses (H+S and GS): Figure 5(c) and Figure 5(d) show the training results with the two types of batch-insensitive losses. With the combination of hinge loss and spreadout regularizer, both the improved centralized model and the federated model perform much better than the baseline model (Figure 5(d)). Also, the gap between improved centralized and federated models is much smaller than with batch softmax. With global softmax (Figure 5(c)), the federated model performs almost the same as the improved centralized model, and both perform significantly better than the batch softmax centralized model. Both of the results indicate that batch-insensitive loss alleviates the performance degradation caused by non-IID negatives effectively.

We caution that applying global softmax may not be appropriate in all settings. In these experiments, we use all the items in the movie vocabulary as the negatives. In practice, when dealing with items with large or unbounded vocabulary size, we may need other strategies as global softmax is computationally expensive.

Overall Comparison: Table 2 gives an overall performance comparison between centralized, improved centralized, and federated models under different losses.




Figure 5: Comparison of centralized and federated models when training with different loss functions: batch softmax, batch softmax with spreadout regularizer, hinge loss with spreadout regularizer, and global softmax. The last two are batch-insensitive losses.

Table 2: An overall recall comparison between centralized, improved centralized, and federated models. The performance drop is calculated as $(R_c - R_f)/R_c$, where R_c is the centralized global recall, and R_f is the federated global recall.

	Centralized				Federated				Performance Drop			
	BS	BS+S	H+S	GS	BS	BS+S	H+S	GS	BS	BS+S	H+S	GS
R@1	1.02	1.21	1.27	1.24	0.58	0.88	1.17	1.21	43.14%	27.27%	7.87%	2.42%
R@5	4.2	5.15	6.08	5.34	2.88	4.01	5.56	5.25	31.43%	22.14%	8.55%	1.69%
R@10	7.42	8.93	11.15	9.46	5.4	7.29	10.43	9.41	27.22%	18.37%	6.46%	0.53%

Training with batch-insensitive losses (H+S, GS) achieves the highest recall for both centralized and federated models. In particular, hinge loss with spreadout regularizer appears to perform slightly better than global softmax in terms of absolute recall, but both perform significantly better than batch-sensitive losses (BS, BS+S).

We also observe that global softmax incurs the smallest performance gap between centralized and federated training. Hinge loss with spreadout regularizer has the next smallest performance gap, and batch-sensitive techniques have a more significant performance gap as expected. We expect that the remaining performance drop between centralized and federated models results from client drift (clients are still taking multiple local steps). This suggests that combining batch-insensitive losses with approaches to address client drift may be a promising future direction.

5 Open-Source Framework

We are releasing a general code framework for experimenting with federated deep retrieval models built on the TensorFlow Federated library³. The code is released under Apache License 2.0. The framework enables reproduction of our experiments and provides a flexible, well-documented interface for researchers to train federated and centralized deep retrieval models with different models and losses. We provide libraries for training and evaluation for MovieLens next movie prediction, which can be easily extended for new tasks. We hope that this framework spurs further research and lowers the barrier to more practical applications.

6 Conclusion

This work investigates the effect of non-IID negatives on federated training of deep retrieval models and proposes batch-insensitive losses to alleviate the issue. We compare model performance using various loss functions and show that batch-insensitive losses produce better federated deep retrieval models that can approximately match centralized models. We also open-source our code framework to accelerate future research and applications. Note that our proposed techniques do not directly address the separate, well-studied issue of client drift when clients do multiple steps of local training—approaches addressing this issue are complementary and can be combined with our work.

³<https://www.tensorflow.org/federated>

Acknowledgements

We thank Warren Morningstar, Chung-Ching Chang, and Zachary Garrett and for their helpful comments and discussions. We also thank Warren Morningstar for his contribution to the federated training pipeline.

References

- [1] Muthuraman Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. Learning cross-lingual sentence representations via a multi-task dual-encoder model. *arXiv preprint arXiv:1810.12836*, 2018.
- [2] Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems*, New York, NY, USA, 2016.
- [3] Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldrige, Eugene Ie, and Diego Garcia-Olano. Learning dense representations for entity retrieval. *arXiv preprint arXiv:1909.10506*, 2019.
- [4] Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. End-to-end retrieval in continuous space. *arXiv preprint arXiv:1811.08008*, 2018.
- [5] Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. Efficient natural language response suggestion for smart reply. *ArXiv e-prints*, 2017.
- [6] Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip Gibbons. The non-iid data quagmire of decentralized machine learning. In *International Conference on Machine Learning*, pages 4387–4398. PMLR, 2020.
- [7] Jyun-Yu Jiang, Tao Wu, Georgios Roumpos, Heng-Tze Cheng, Xinyang Yi, Ed Chi, Harish Ganapathy, Nitin Jindal, Pei Cao, and Wei Wang. End-to-end deep attentive personalized item retrieval for online content-sharing platforms. In *Proceedings of The Web Conference 2020*, pages 2870–2877, 2020.
- [8] Sai Praneeth Karimireddy, Martin Jaggi, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Mime: Mimicking centralized stochastic algorithms in federated learning. *arXiv preprint arXiv:2008.03606*, 2020.
- [9] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.
- [10] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018.
- [11] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. In *International Conference on Learning Representations*, 2020.
- [12] Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*, 2021.
- [13] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1930–1939, 2018.
- [14] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.

- [15] Pavel Sountsov and Sunita Sarawagi. Length bias in encoder decoder models and a case for global conditioning. *arXiv preprint arXiv:1606.03402*, 2016.
- [16] Maksims Volkovs, Guang Wei Yu, and Tomi Poutanen. Dropoutnet: Addressing cold start in recommender systems. In *NIPS*, pages 4957–4966, 2017.
- [17] Ji Yang, Xinyang Yi, Derek Zhiyuan Cheng, Lichan Hong, Yang Li, Simon Xiaoming Wang, Taibai Xu, and Ed H Chi. Mixed negative sampling for learning two-tower neural networks in recommendations. In *Companion Proceedings of the Web Conference 2020*, pages 441–447, 2020.
- [18] Yinfei Yang, Steve Yuanc, Daniel Cera, Sheng-yi Konga, Noah Constanta, Petr Pilarc, Heming Gea, Yun-Hsuan Sunga, Brian Stropea, and Ray Kurzweila. Learning semantic textual similarity from conversations. *ACL 2018*, page 164, 2018.
- [19] Xinyang Yi, Ji Yang, Lichan Hong, Derek Zhiyuan Cheng, Lukasz Heldt, Aditee Ajit Kumthekar, Zhe Zhao, Li Wei, and Ed Chi, editors. *Sampling-Bias-Corrected Neural Modeling for Large Corpus Item Recommendations*, 2019.
- [20] Tehrim Yoon, Sumin Shin, Sung Ju Hwang, and Eunho Yang. Fedmix: Approximation of mixup under mean augmented federated learning. In *International Conference on Learning Representations*, 2021.
- [21] Felix X. Yu, Ankit Singh Rawat, Aditya Krishna Menon, and Sanjiv Kumar. Federated learning with only positive labels, 2020.
- [22] Xu Zhang, Felix X. Yu, Sanjiv Kumar, and Shih-Fu Chang. Learning spread-out local feature descriptors, 2017.
- [23] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.