

DEEP RETRIEVAL: AN END-TO-END LEARNABLE STRUCTURE MODEL FOR LARGE-SCALE RECOMMENDATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

One of the core problems in large-scale recommendations is to retrieve top relevant candidates accurately and efficiently, preferably in sub-linear time. Previous approaches are mostly based on a two-step procedure: first learn an inner-product model and then use maximum inner product search (MIPS) algorithms to search top candidates, leading to potential loss of retrieval accuracy. In this paper, we present Deep Retrieval (DR), an end-to-end learnable structure model for large-scale recommendations. DR encodes all candidates into a discrete latent space. Those latent codes for the candidates are model parameters and to be learnt together with other neural network parameters to maximize the same objective function. With the model learnt, a beam search over the latent codes is performed to retrieve the top candidates. Empirically, we showed that DR, with sub-linear computational complexity, can achieve almost the same accuracy as the brute-force baseline.

1 INTRODUCTION

Recommendation systems have gained great success in various commercial applications for decades. The objective of recommendation systems is to retrieve relevant candidates from an enormous corpus based on user features and historical behaviors. In the era of mobile internet, the amount of candidates from content providers and the amount of active users both grow rapidly to tens of millions to hundred of millions, making it more challenging to design accurate recommendation systems. The scalability and efficiency of the algorithm are the main challenge for modern recommendation systems.

One of the early successful techniques of recommendation systems is the collaborative filtering (CF), which makes predictions based on the simple idea that similar users may prefer similar items. Item-based collaborative filtering (Item-CF) (Sarwar et al., 2001) extends the idea by considering the similarities between items and items, and later lays a foundation for Amazon’s recommendation system (Linden et al., 2003).

Recently, vector-based recommendation algorithms have been widely adopted. The main idea is to embed users and items in a latent vector space, and use the inner product of vectors to represent the preference between users and items. Representative vector embedding methods includes matrix factorization (MF) (Mnih & Salakhutdinov, 2008; Koren et al., 2009), factorization machines (FM) (Rendle, 2010), DeepFM (Guo et al., 2017), Field-aware FM (FFM) (Juan et al., 2016), etc. However, when the number of items is large, the complexity of brute-force computation of the inner product for all items can be prohibitive. Thus, maximum inner product search (MIPS) or approximate nearest neighbors (ANN) algorithms are usually used to retrieve items when the corpus is large. Efficient MIPS or ANN algorithms includes tree-based algorithms (Muja & Lowe, 2014; Houle & Nett, 2014), locality sensitive hashing (LSH) (Shrivastava & Li, 2014; Spring & Shrivastava, 2017), product quantization (PQ) (Jegou et al., 2010; Ge et al., 2013), hierarchical navigable small world graphs (HNSW) (Malkov & Yashunin, 2018), etc.

Despite their success in real world applications, vector-based algorithms has two main deficiencies: (1) The objective of learning vector representation and learning good MIPS structure are not perfectly aligned; (2) The dependency on inner products of user and item embeddings is not sufficient to capture the complicated structure of user-item interactions (He et al., 2017). In order to break these limitations, tree based models (Zhu et al., 2019; 2018) have been proposed. These methods use a tree

as indices and map each item to a leaf node of the tree. Learning objectives for model parameters and tree structures are well aligned to improve the accuracy. However, the number of parameters in these models are proportional to the number of clusters, making the tree structure itself difficult to learn — data available at the leaf level can be scarce and might not provide enough signal to learn a good tree at that finer level.

In this paper, we proposed an end-to-end trainable structure model — Deep Retrieval (DR). Instead of using a tree structure, we propose to use a $D \times K$ matrix as in Figure 1a for indexing, motivated by (Chen et al., 2018). Each item is indexed by one or more “codes” (or equivalently “paths”) with length D and range $\{1, 2, \dots, K\}$. For example, an item related to chocolate may be encoded by $[36, 204, 105]$ and another item related to cake may be encoded by $[36, 227, 120]$. There are K^D possible paths and each path can be interpreted as a cluster of items: each path could contain multiple items and each item could also belong to multiple paths.

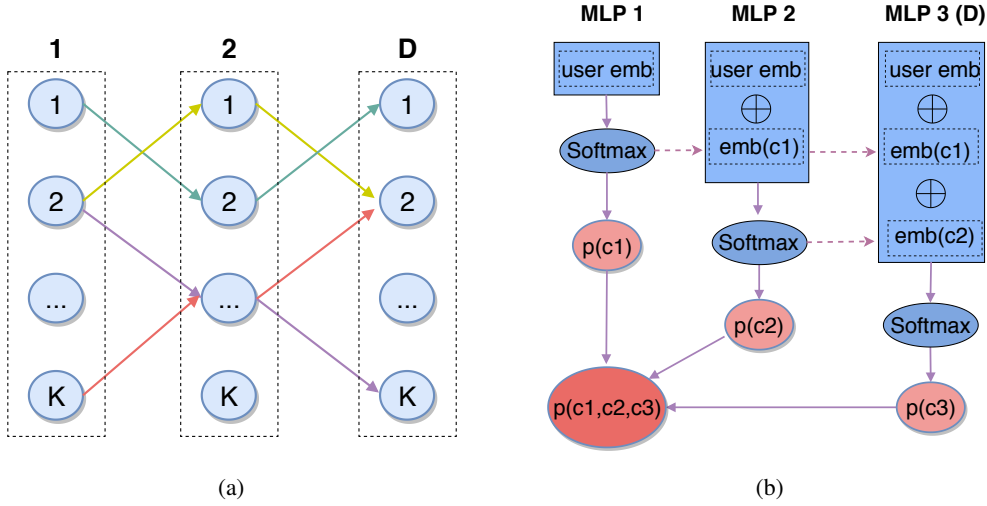


Figure 1: (a) Consider a structure with width $K = 100$ and depth $D = 3$. Assuming an item is encoded by length- D vector $[36, 27, 20]$, which is called a “code” or a “path”. The path denotes that the item is assigned to the $(1, 36)$, $(2, 27)$, $(3, 20)$ indices of the $K \times D$ matrix. In the figure, arrows with the same color form a path. Different paths could intersect with each other by sharing the same index at some layer. (b) Flow chart showing the process for constructing the probability $p(c|x, \theta)$.

The advantage of DR is multi-folded. In training, the item paths can be learnt together with the neural network parameters of the structure model using an expectation-maximization (EM) type algorithm. Therefore the entire training process is end-to-end and can be easily deployed on deep learning platforms. In terms of model capability, the structure model enables DR to have much more clusters than its parameters, which solves the problem of scarce data in TDM and JTM. In addition, the multiple-to-multiple encoding scheme enables DR to learn more complicated relationships among users and items. Experiments show the benefit of being able to learn such complicated relationships.

The rest of the paper is organized as follows. In Section 2, we describe the structure model and the structure objective function used in training in detail. We also introduce a beam search algorithm to find candidate paths in the inference stage. In Section 3, we introduce the EM algorithm for training neural network parameters and paths of items jointly. We adopt a penalization on the sizes of the paths to prevent overfitting. In Section 4, we demonstrate the performance of DR on two public datasets: MovieLens-20M¹ and Amazon books². Experiment results show that DR can almost achieve the brute-force accuracy with sub-linear computational complexity. In Section 5, we conclude the paper and discuss several possible future research directions.

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

²<http://jmcauley.ucsd.edu/data/amazon>

2 STRUCTURE MODEL

In this section, we introduce the structure model in Deep Retrieval in detail. Firstly, we construct the probability for user x_i to select path c_i given the model parameters θ and present the training objective. Then we introduce a multi-path mechanism to enable the model to capture multi-aspect properties of items. In the inference stage, we introduce a beam search algorithm to select candidate paths based on user embeddings.

2.1 STRUCTURE OBJECTIVE FUNCTION

The structure model contains D layers³ with K nodes in each layer, where each layer is a multi-layer perceptron (MLP) with skip connections and softmax as output. Each layer takes an input vector and outputs a probability distribution over $\{1, 2, \dots, K\}$ based on parameters θ . Let $\mathcal{V} = \{1, \dots, V\}$ be the labels of all items and $\pi : \mathcal{V} \rightarrow [K]^D$ be the mapping from items to paths. We assume π is fixed in this section and will introduce the algorithm for learning π together with θ in Section 3.

Given a pair of training sample (x, y) , which denotes a positive interaction (click, convert, like, etc.) between user x and item y , as well as the path $c = (c_1, c_2, \dots, c_D)$ associated with the item y , the probability $p(c|x, \theta)$ is constructed layer by layer as follows (see Figure 1b for a flow chart).

- The first layer takes the user embedding $\text{emb}(x)$ as input, and output a probability $p(c_1|x, \theta_1)$ over the K nodes of the first layer, based on parameters θ_1 .
- From the second layer onward, we concatenate the user embedding $\text{emb}(x)$ and the embeddings of all the previous layers $\text{emb}(c_{d-1})$ (called path embeddings) as the input of MLP, and output $p(c_d|x, c_1, \dots, c_{d-1}, \theta_d)$ over the K nodes of layer d , based on parameters θ_d .
- The probability of $p(c|x, \theta)$ is just the product of the probabilities of all the layers.

$$p(c|x, \theta) = \prod_{d=1}^D p(c_d|x, c_1, \dots, c_{d-1}, \theta_d). \quad (1)$$

Given a set of N training samples $\{(x_i, y_i)\}_{i=1}^N$, we maximize the log likelihood function of the structure model.

$$\mathcal{Q}_{\text{str}}(\theta, \pi) = \sum_{i=1}^N \log p(\pi(y_i)|x_i, \theta) = \sum_{i=1}^N \sum_{d=1}^D \log p(\pi(y_i)_d|x_i, \pi(y_i)_1, \dots, \pi(y_i)_{d-1}, \theta_d). \quad (2)$$

The size of the input vector of layer d is the embedding size times d , and the size of the output vector is K . The parameters of layer d have a size of $\Theta(Kd)$. The parameters θ contain the parameters θ_d 's in all layers, as well as the path embeddings. The parameters in the entire model have an order of $\Theta(KD^2)$, which is significantly smaller than the number of possible paths K^D when $D \geq 2$.

2.2 MULTI-PATH STRUCTURE OBJECTIVE

In tree-based deep models (Zhu et al., 2018; 2019) as well as the structure model we introduced before, each item belongs to only one cluster, which limits the capacity of the model to express multi-aspect information in real data. For example, an item related to kebab should belong to a cluster related to food. An item related to flowers should belong to a cluster related to gifts. However, an item related to chocolate or cakes should belong to both clusters in order to be recommended to users interested in either food or gifts. In real world recommendation systems, a cluster might not have an explicit meaning such as food or gifts, but this example motivates us to assign each item to multiple clusters. In DR, we allow each item y_i to be assigned to J different paths $\{c_{i,1}, \dots, c_{i,J}\}$. Let $\pi : \mathcal{V} \rightarrow [K]^{D \times J}$ be the mapping from items to multiple paths and the multi-path structure objective is defined as,

$$\mathcal{Q}_{\text{str}}(\theta, \pi) = \sum_{i=1}^N \log \left(\sum_{j=1}^J p(c_{i,j} = \pi_j(y_i)|x_i, \theta) \right). \quad (3)$$

³Other neural network architectures such as recurrent neural networks (RNN) can also be applied here. Since D is not very large in our settings, for simplicity, we use MLP here.

Beam search for inference. In the inference stage, we want to retrieve items from the structure model, given user embeddings as input. To this end, we utilize beam search algorithm (Reddy et al., 1977) to retrieve multiple paths and merge the items in the retrieved paths. In each layer, the algorithm selects top B nodes from the all successors of the selected nodes from the previous layer. Finally it returns B top paths in the final layer. Notice that greedy search is a special case of beam search when $B = 1$. The time complexity of the inference stage is $O(DKB \log B)$, which is independent with the number of items. The detail of the beam search is relegated to the appendix.

3 LEARNING

In the previous section, we introduced the structure model in Deep Retrieval and the structure objective to be optimized. The objective is entirely continuous with respect to the parameters, hence can be optimized by any gradient-based optimizer. However, the objective involves the mapping from items to paths π , which is discrete and can not be optimized by gradient-based optimizer. This mapping acts as the “clustering” of items, which motivates us to use an EM-style algorithm to optimize the mapping and the continuous parameters jointly. In this section, we describe the EM algorithm in detail and introduce a penalization term to prevent overfitting.

3.1 EM ALGORITHM FOR JOINT TRAINING

Given a user-item pair (x_i, y_i) in training data set, let the path associate with the item $(\{\pi_1(y_i), \dots, \pi_J(y_i)\})$ be the latent data in EM algorithm. Along with the continuous parameters θ , the objective function is given by

$$\mathcal{Q}_{\text{str}}(\theta, \pi) = \sum_{i=1}^N \log \left(\sum_{j=1}^J p(\pi_j(y_i) | x_i, \theta) \right) = \sum_{v=1}^V \left(\sum_{i: y_i=v} \log \left(\sum_{j=1}^J p(\pi_j(v) | x_i, \theta) \right) \right). \quad (4)$$

where the outer summation is over all items $v \in \mathcal{V}$ and the inner summation is over all appearances of item v in the training set. We maximize the objective function over all possible mappings π . However, for an item v , there are K^D number of possible paths so we could not maximize it over all $\pi_j(v)$ ’s. Instead, for each appearance, we only record the values for the top paths using beam search and leave the rest paths with zero scores. Unfortunately, the top paths for item v in different appearances may not be the same due to the difference of user and context embeddings, which renders an unstable function due to the presence of $\log(p) + \log(0)$ in Equation (4)⁴. To address this, we approximate it using an upper bound⁵ $\sum_{i=1}^N \log p_i \leq N(\log \sum_{i=1}^N p_i - \log N)$ to obtain

$$\mathcal{Q}_{\text{str}}(\theta, \pi) \leq \bar{\mathcal{Q}}_{\text{str}}(\theta, \pi) = \sum_{v=1}^V \left(N_v \log \left(\sum_{j=1}^J \sum_{i: y_i=v} p(\pi_j(v) | x_i, \theta) \right) - \log N_v \right), \quad (5)$$

where $N_v = \sum_{i=1}^N \mathbb{I}[i : y_i = v]$ denotes the number of occurrences of v in the training set which is independent of the mapping. We denote the score $s[v, c]$ as $\sum_{i: y_i=v} p(c | x_i, \theta)$ for simplicity of notation. In practice, it is impossible to retain all scores as the possible number of paths c is exponentially large, so we only retain a subset of S paths with larger scores through beam search. In M-step, we simply maximize $\sum_{j=1}^J s[v, \pi_j(v)]$ over each $\pi_j(v)$, which is equivalent to pick J highest scores of $s[v, c]$ among the top paths c from beam search. The EM training algorithm is summarized in Algorithm 1.

3.2 PENALIZATION ON SIZE OF PATHS

Overfitting is likely to happen if we do not apply any penalization for Algorithm 1. Imagine that the structure model is overfitted and gives a particular path $(0, 0, 0)$ a very high probability, for any input. Then in M-step, all the items will be assigned to path $(0, 0, 0)$, making the model fail to cluster the

⁴We also tried to use a small value. But it is difficult to figure out what the small value should be.

⁵Since this is an upper bound for the true objective we are maximizing, there is no guarantee as to maximizing a surrogate via a lower bound. However we still find it works well in practice.

Algorithm 1 EM algorithm for structure learning

Input: training set $\{x_i, y_i\}_{i=1}^N$. T is some predefined epoch number.
Initialize $\pi_j^{(0)}$ and θ randomly.
for $t = 1$ **to** T **do**
 Fixed $\pi^{(t-1)}$, optimize parameter θ using a gradient-based optimizer to maximize structure objective $\mathcal{Q}_{\text{str}}(\theta, \pi^{(t-1)})$.
 for $v = 1$ **to** V **do**
 Compute $s[v, c]$ among the top paths c from beam search.
 Let $\{\pi_1^{(t)}(v), \dots, \pi_J^{(t)}(v)\} \leftarrow$ top J scores of $s[v, c]$.
 end for
end for

items. In order to prevent overfitting, we introduce a penalization term on the size of the paths. The penalized \mathcal{Q} function is given by

$$\begin{aligned} \mathcal{Q}_{\text{pen}}(\theta, \pi) &= \overline{\mathcal{Q}}_{\text{str}}(\theta, \pi) - \alpha \cdot \sum_{c \in [K]^D} f(|c|) \\ &= \sum_{v=1}^V \left(N_v \log \left(\sum_{j=1}^J s[v, \pi_j(v)] \right) - \log N_v \right) - \alpha \cdot \sum_{c \in [K]^D} f(|c|). \end{aligned} \quad (6)$$

where α is the penalty factor, $|c|$ denotes the number of items in c and f is an increasing and convex function. A quadratic function $f(|c|) = |c|^2/2$ controls the average size of paths, and higher order polynomials penalize more on larger paths. In our experiments, we use $f(|c|) = |c|^4/4$. It's worth mentioning that the penalty is only applied in M-step, not during training of continuous parameters θ .

Coordinate descent algorithm for path assignment. It is intractable to jointly optimize the penalized objective function (6) over all the path assignments, since the penalization term can not be decomposed into summation of terms of each item. So we use coordinate descent algorithm in M-step. We fix the assignment of all other items while optimizing the path assignments $\pi_j(v)$'s of v . The time complexity for coordinate descent is $O(VJS)$, where S is the pool size for candidate paths. The detail of the coordinate descent algorithm is relegated to the appendix.

3.3 MULTI-TASK LEARNING AND RERANKING WITH SOFTMAX MODELS

In the structure model described above, we only make use of positive instances and maximize the log probability of positive instances. This makes the model able to recognize "good" items, but unable to distinguish "good" items from "bad" ones. So the model is good at retrieval but not ranking. To this end, we introduced a model $f_\phi(x_i, y_i)$ trained with sampled softmax objective described as

$$\mathcal{Q}_{\text{softmax}} = \sum_{i=1}^N \frac{\exp(f_\phi(x_i, y_i))}{\exp(f_\phi(x_i, y_i)) \sum_{j=1}^M \exp(f_\phi(x'_{i,j}, y'_{i,j}))} \quad (7)$$

where $(x'_{i,j}, y'_{i,j})$ are randomly sampled pairs. By sharing the input with the softmax model, we teach DR to distinguish positive instances from random negative instances. The final objective is given by $\mathcal{Q} = \mathcal{Q}_{\text{pen}} + \mathcal{Q}_{\text{softmax}}$. After performing beam search to retrieval a set of candidate items using Algorithm 3, we use the function f_ϕ to rerank those candidates to obtain the final top candidates.

3.4 COMPLEXITY ANALYSIS

We summarize the time complexity of each stage in Deep Retrieval as follows. The time complexity of the inference stage is $O(DKB \log B)$ per sample, which is sub-linear with respect to the number of items. The time complexity of training continuous parameters θ is $O(NJKD^2)$, where N is the number of training samples in one epoch and J be the multiplicity of paths. The time complexity of path assignment by coordinate descent algorithm is $O(VJS)$, where V is corpus size and S is the pool size of candidate paths.

4 EXPERIMENTS

In this section, we study the performance of DR on two public recommendation datasets: MovieLens-20M (Harper & Konstan, 2015) and Amazon books (He & McAuley, 2016; McAuley et al., 2015). We compare the performance of DR with brute-force algorithm, as well as several other recommendation baselines including tree-based models TDM (Zhu et al., 2018) and JTM (Zhu et al., 2019). In the end of this section, we investigate the role of important hyperparameters in DR.

4.1 DATASETS AND METRICS

MovieLens-20M. This dataset contains rating and free-text tagging activities from a movie recommendation service called MovieLens. We use the 20M subset which were created by the behaviors of 138,493 users between 1995 and 2015. Each user-movie interaction contains a user-id, a movie-id, a rating between 1.0 to 5.0, as well as a timestamp.

In order to make a fair comparison, we exactly follow the same data pre-processing procedure as TDM. We only keep records with rating higher or equal to 4.0, and only keep users with at least ten reviews. After pre-processing, the dataset contains 129,797 users, 20,709 movies and 9,939,873 interactions. Then we randomly sample 1,000 users and corresponding records to construct the validation set, another 1,000 users to construct the test set, and other users to construct the training set. For each user, the first half of the reviews according to the timestamp are used as historical behavior features and the latter half are used as ground truths to be predicted.

Amazon books. This dataset contains user reviews of books from Amazon, where each user-book interaction contains a user-id, an item-id, and corresponding timestamp. Similar to MovieLens-20M, we follow the same pre-processing procedure as JTM. The dataset contains 294,739 users, 1,477,922 items and 8,654,619 interactions. Notice that Amazon books dataset have much more items but sparser interactions than MovieLens-20M. We randomly sample 5,000 users and corresponding records as the test set, another 5,000 users as the validation set and other users as the training set. The construction procedures of behavior features and ground truths are the same as in MovieLens-20M.

Metrics. We use precision, recall and F-measure as metrics to evaluate the performance for each algorithm. We emphasize that the metrics are computed for each user individually, and averaged without weight across users, following the same setting as both TDM and JTM. We compute the metrics by retrieving top 10 and 200 items for each user in MovieLens-20M and Amazon books respectively.

Model and training. Here we present some details about the model and training procedure in the experiment. Since the dataset is split in a way such that the users in the training set, validation set and test set are disjoint, we drop the user-id and only use the behavior sequence as input for DR. The behavior sequence is truncated to length of 69 if it is longer than 69, and filled with a placeholder symbol if it is shorter than 69. A recurrent neural network with GRU is utilized to project the behavior sequence onto a fixed dimension embedding as the input of DR. We adopt the multi-task learning framework, and rerank the items in the recalled paths by a softmax reranker. We train the embeddings of DR and softmax jointly for the initial two epochs, freeze the embeddings of softmax and train the embeddings of DR for two more epochs. The reason is to prevent overfitting of the softmax model. In the inference stage, the number of items retrieved from beam search is not fixed due to the differences of path sizes, but the variance is not large. Empirically we control the beam size such that the number of items from beam search is 5 to 10 times the number of finally retrieved items.

4.2 EMPIRICAL RESULTS

We compare the performance of DR with the following algorithms: Item-CF (Sarwar et al., 2001), YouTube product DNN (Covington et al., 2016), TDM and JTM. We directly use the numbers of Item-CF, Youtube DNN, TDM and JTM from TDM and JTM papers for fair comparison. Among the different variants of TDM presented, we pick the one with best performance. The result of JTM is only available for Amazon books. We also compare DR with brute-force retrieval algorithm, which directly computes the inner-product of user embedding and all the item embeddings learnt in the softmax model and returns the top K items. The brute-force algorithm is usually computationally

Table 1: Comparison of precision@10, recall@10 and F-measure@10 for DR, brute-force retrieval and other recommendation algorithms on MovieLens-20M.

Algorithm	Precision@10	Recall@10	F-measure@10
Item-CF	8.25%	5.66%	5.29%
YouTube DNN	11.87%	8.71%	7.96%
TDM (best)	14.06%	10.55%	9.49%
DR	20.58% \pm 0.47%	10.89% \pm 0.32%	12.32% \pm 0.36%
Brute-force	20.70% \pm 0.16%	10.96% \pm 0.32%	12.38% \pm 0.32%

Table 2: Comparison of precision@200, recall@200 and F-measure@200 for DR, brute-force and other recommendation algorithms on Amazon Books.

Algorithm	Precision@200	Recall@200	F-measure@200
Item-CF	0.52%	8.18%	0.92%
YouTube DNN	0.53%	8.26%	0.93%
TDM (best)	0.56%	8.57%	0.98%
JTM	0.79%	12.45%	1.38%
DR	0.95% \pm 0.01%	13.74% \pm 0.14%	1.63% \pm 0.02%
Brute-force	0.95% \pm 0.01%	13.75% \pm 0.10%	1.63% \pm 0.02%

prohibitive in practical large recommendation systems, but can be used as an upper bound for small dataset for inner-product based models.

Table 1 shows the performance of DR compared to other algorithms and brute-force for MovieLens-20M. Table 2 shows the results for Amazon books. For DR and brute-force, we independently train the same model for 5 times and compute the mean and standard deviation of each metric. We conclude the following results.

- DR performs better than other methods including tree-based retrieval algorithms such as TDM and JTM.
- The performance of DR is very close to or on par with the performance of brute-force method, which is an upper bound of the performance of vector-based methods such as Deep FM and HNSW. However, the inference speed of DR is 4 times faster than brute-force in Amazon books dataset (see Table 3 in Appendix C.1).

4.3 SENSITIVITY OF HYPERPARAMETERS

DR introduces some key hyperparameters which may infect the performance dramatically, including the width of the structure model K , depth of model D , number of multiple paths J , beam size B and penalty factor α . In the MovieLens-20M experiment, we choose $K = 50$, $D = 3$, $B = 25$, $J = 3$ and $\alpha = 3 \times 10^{-5}$. In the Amazon books experiment, we choose $K = 100$, $D = 3$, $B = 50$, $J = 3$ and $\alpha = 3 \times 10^{-7}$. Using the Amazon books dataset, we show the role of these hyperparameters and see how they may affect the performance. We present how the recall@200 change as these hyperparameters change in Figure 2. We fix the value of other hyperparameters unchanged when varying one hyperparameter. Precision@200 and F-measure@200 follow similar trends so we plot them in the appendix.

- **Width of model K** controls the overall capacity of the model. If K is too small, the number of clusters is too small for all the items; if K is too big, the time complexity of training and inference stages grow linearly with K . Moreover, large K may increase the possibility of overfitting. An appropriate K should be chosen depending on the size of the corpus.
- **Depth of model D** . Using $D = 1$ is obviously not a good idea since it fails to capture dependency among layers. In Tabel 4 in Appendix C.2. we present the result for $D = 2, 3$ and 4 for Amazon books dataset, and conclude that $D = 3$ is a good choice as a trade-off between model performance and memory usage in both experiments.




Figure 2: Relationship between recall@200 in Amazon Books experiment and model width K , number of paths J , beam size B and penalty factor α , respectively.

- **Number of paths J** enables the model to express multi-aspect information of items. The performance is the worst when $J = 1$, and keeps increasing as J increases. Large J may not affect the performance, but the time complexity for training grows linearly with J . In practice, choosing J between 3 and 5 is recommended.
- **Beam size B** controls the number of candidate paths to be recalled. Larger B leads a better performance as well as heavier computation in the inference stage. Notice that greedy search is a special case when $B = 1$, whose performance is worse than beam search with larger B .
- **Penalty factor α** controls the number of items in each path. The best performance is achieved when α falls in a certain range. Smaller α leads to a larger path size (see Table 5 in Appendix C.3) hence heavier computation in the reranking stage. Beam size B and penalty factor α should be appropriately chosen as a trade off between model performance and inference speed.

Overall, we can see that DR is fairly stable to hyperparameters since there is a wide range of hyperparameters which leads to near-optimal performances.

5 CONCLUSION AND DISCUSSION

In this paper, we have proposed Deep Retrieval, an end-to-end learnable structure model for large-scale recommender systems. DR uses an EM-style algorithm to learn the model parameters and paths of items jointly. Experiments have shown that DR performs well compared with brute-force baselines in two public recommendation datasets.

There are several future research directions based on the current model design. Firstly, the structure model defines the probability distribution on paths only based on user side information. An useful idea would be how to incorporate item side information more directly into the DR model. Secondly, in the structure model, we only make use of positive interactions such as click, convert or like between user and item. Negative interactions such as non-click, dislike or unfollow should also be considered in future work to improve the model performance. Finally, the current DR still uses a softmax model as a reranker, which might have an upper bound of the performance capped by the softmax model. We are actively working to address this issue as well.

REFERENCES

- Ting Chen, Martin Renqiang Min, and Yizhou Sun. Learning k-way d-dimensional discrete codes for compact embedding representations. *arXiv preprint arXiv:1806.09464*, 2018.
- Paul Covington, Jay Adams, and Emre Sargin. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, pp. 191–198, 2016.
- Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. Optimized product quantization. *IEEE transactions on pattern analysis and machine intelligence*, 36(4):744–755, 2013.
- Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. Deepfm: a factorization-machine based neural network for ctr prediction. *arXiv preprint arXiv:1703.04247*, 2017.
- F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- Ruining He and Julian McAuley. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*, pp. 507–517, 2016.
- Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pp. 173–182, 2017.
- Michael E Houle and Michael Nett. Rank-based similarity search: Reducing the dimensional dependence. *IEEE transactions on pattern analysis and machine intelligence*, 37(1):136–150, 2014.
- Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- Yuchin Juan, Yong Zhuang, Wei-Sheng Chin, and Chih-Jen Lin. Field-aware factorization machines for ctr prediction. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pp. 43–50, 2016.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- Greg Linden, Brent Smith, and Jeremy York. Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing*, 7(1):76–80, 2003.
- Yury A Malkov and Dmitry A Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 43–52, 2015.
- Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. In *Advances in neural information processing systems*, pp. 1257–1264, 2008.
- Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240, 2014.
- D Raj Reddy et al. Speech understanding systems: A summary of results of the five-year research effort. *Department of Computer Science. Camegie-Mell University, Pittsburgh, PA*, 17, 1977.
- Steffen Rendle. Factorization machines. In *2010 IEEE International Conference on Data Mining*, pp. 995–1000. IEEE, 2010.
- Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pp. 285–295, 2001.

Anshumali Shrivastava and Ping Li. Asymmetric lsh (alsh) for sublinear time maximum inner product search (mips). In *Advances in Neural Information Processing Systems*, pp. 2321–2329, 2014.

Ryan Spring and Anshumali Shrivastava. A new unbiased and efficient class of lsh-based samplers and estimators for partition function computation in log-linear models. *arXiv preprint arXiv:1703.05160*, 2017.

Han Zhu, Xiang Li, Pengye Zhang, Guozheng Li, Jie He, Han Li, and Kun Gai. Learning tree-based deep model for recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1079–1088, 2018.

Han Zhu, Daqing Chang, Ziru Xu, Pengye Zhang, Xiang Li, Jie He, Han Li, Jian Xu, and Kun Gai. Joint optimization of tree-based index and deep model for recommender systems. In *Advances in Neural Information Processing Systems*, pp. 3973–3982, 2019.