Zero-Shot Heterogeneous Transfer Learning from Recommender Systems to Cold-Start Search Retrieval

Tao Wu*, Ellie Ka-In Chio*, Heng-Tze Cheng*

Yu Du, Steffen Rendle, Dima Kuzmin, Ritesh Agarwal, Li Zhang, John Anderson, Sarvjeet Singh, Tushar Chandra, Ed H. Chi, Wen Li, Ankit Kumar, Xiang Ma, Alex Soares, Nitin Jindal, Pei Cao
Google Inc.[†]

ABSTRACT

Many recent advances in neural information retrieval models, which predict top- K items given a query, learn directly from a large training set of (query, item) pairs. However, they are often insufficient when there are many previously unseen (query, item) combinations, often referred to as the cold start problem. Furthermore, the search system can be biased towards items that are frequently shown to a query previously, also known as the “rich get richer” (a.k.a. feedback loop) problem. In light of these problems, we observed that most online content platforms have both a search and a recommender system that, while having heterogeneous input spaces, can be connected through their common output item space and a shared semantic representation. In this paper, we propose a new Zero-Shot Heterogeneous Transfer Learning framework that transfers learned knowledge from the recommender system component to improve the search component of a content platform. First, it learns representations of items and their natural-language features by predicting (item, item) correlation graphs derived from the recommender system as an auxiliary task. Then, the learned representations are transferred to solve the target search retrieval task, performing query-to-item prediction without having seen any (query, item) pairs in training. We conduct online and offline experiments on one of the world’s largest search and recommender systems from Google, and present the results and lessons learned. We demonstrate that the proposed approach can achieve high performance on offline search retrieval tasks, and more importantly, achieved significant improvements on relevance and user interactions over the highly-optimized production system in online experiments.

CCS CONCEPTS

• **Computing methodologies** → **Transfer learning**; • **Information systems** → **Information retrieval**.

KEYWORDS

zero-shot learning; transfer learning; search; recommender systems

* Equal contribution.

[†] Corresponding author: Tao Wu: iotao@google.com.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6859-9/20/10.

<https://doi.org/10.1145/3340531.3412752>

ACM Reference Format:

Tao Wu*, Ellie Ka-In Chio*, Heng-Tze Cheng* and Yu Du, Steffen Rendle, Dima Kuzmin, Ritesh Agarwal, Li Zhang, John Anderson, Sarvjeet Singh, Tushar Chandra, Ed H. Chi, Wen Li, Ankit Kumar, Xiang Ma, Alex Soares, Nitin Jindal, Pei Cao. 2020. Zero-Shot Heterogeneous Transfer Learning from Recommender Systems to Cold-Start Search Retrieval. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, October 19–23, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3340531.3412752>

1 INTRODUCTION

Most online content platforms, such as music streaming services or e-commerce websites, have systems that return top- K items either given a natural-language query (i.e., a search retrieval system), or given the user context, which can be the user attributes and user’s interactions with the platform (i.e., a recommender system). These two systems share the same output item space, but have different input feature spaces. In this paper, we study how to improve search retrieval by transferring learnings from recommender systems.

Recently neural information retrieval (neural IR) models have been widely applied in search products across many industries [10, 22, 32]. Such methods can retrieve and score items that do not share keywords with the query. However, they usually require large amount of (query, item) pairs as training data, usually collected from users’ search logs. However, these types of training data may not be available for many (query, item) combinations. This is often referred to as the *cold start* problem [13]. Consider, for example, an online music streaming system, where users mostly listen to music through homepage recommendations or playlists generated by other users. In this case, it is likely that many songs in the system have been listened by users, but not through search requests. This motivated us to ask: Can we utilize the abundant data from the recommender system to cold-start the search retrieval system of the same content platform?

In this paper, we propose a new Zero-Shot Heterogeneous Transfer Learning framework (ZSL), which does not require any input from (query, item) search data, but learns the query and item representations completely via an auxiliary task derived from the recommender system. There are two advantages: first, this method can cold-start the search system as it does not require search data to train. Second, supervised methods trained on (query, item) can suffer from the potential bias and feedback loop [25] introduced by the search system. For large-scale search and recommender systems, there could easily be more than thousands of relevant items matching a query, but users are only able to explore a very limited amount. Therefore, the collected training data are heavily affected by the current search algorithm, exacerbating the “rich get richer”

Figure 1: A comparison on the high-level frameworks between the classic zero-shot learning in image classification [26] (on the left) and the zero-shot heterogeneous transfer learning in this paper (on the right).

effect [3]. Thus, this framework, although motivated by the search retrieval cold-start problem, can also be useful even when adequate search data (query, item) is available.

We assume that (item, item) correlations can be extracted from the recommender system, with text features for each item. The (item, item) correlations commonly exist [12], such as citation network of a bibliography system or music co-listening pairs of an online music platform, where the text features of items are titles, keywords or descriptions. The auxiliary task of the proposed framework is to predict neighbor items given a seed item, where the learned semantic representations will be transferred to the target task which is to predict items given a query. We explore two implementations under this framework. We call the first method Multi-task Item Encoder, where the item and the text features representations are jointly learned by optimizing the two tasks of predicting the item given its text features, and predicting the item given its neighbors. The second approach is Item-to-Text Transformed Encoder, where it only optimizes for a single task of predicting the item given its neighbors, but utilizing the text features to encode the items.

We conduct experiments on one of the world’s largest recommender systems from Google. Our proposed methods demonstrate promising results on multiple offline retrieval tasks. In a A/B test live experiment study, where the base is the production search system that is already highly fine tuned, with many components of advanced retrieval techniques (e.g., both term-frequency-based scoring and supervised machine learning algorithms), our proposed method improved multiple evaluation metrics. This shows that even for a system with enough search training data available, ensembling our proposed method can successfully introduce new relevant results that are favored by users.

Our contributions in this paper are summarized below:

- (1) To our best knowledge, this is the first work that studies the problem of cold-starting a production-scale search retrieval system from (item, item) correlations in the recommender system from the same online content platform.
- (2) We proposed the Zero-Shot Heterogeneous Transfer Learning framework as a solution to cold-start search retrieval.
- (3) We conduct extensive offline and online experiments on one of the world’s largest recommender systems, and demonstrate the effectiveness of our proposed methods. In addition, our findings regarding the effectiveness of our method on broad queries, are valuable insights for practitioners to apply such techniques to real word search systems.

2 RELATED WORK

Zero-shot learning and transfer learning. For large-scale label domains, it is common to have labels of instances that have never been seen in the training data. The key idea of Zero-shot learning [29] is to utilize some auxiliary information of the unseen labels, and learn a model to connect the auxiliary information to the input space. By mapping the input feature to the auxiliary information, the zero-shot algorithms are then able to find the corresponding unseen labels. Applications of such idea include object detection of unseen classes [26], semantic image retrieval [16], and more recently recommender systems for new users [14]. Transfer learning seeks to improve a learner from one domain by transferring information from a related domain. Our proposed framework in this paper is a case of heterogeneous transfer learning [20], as the input spaces of the auxiliary and target tasks are different (See Figure 1).

Cold-start problem. This mostly refers to modeling new users (no previous interactions with the system) or new items (no records of being consumed by users) in the application of recommender systems. Most methods assume some side feature information of the user or item is available, so that the representation of a new user or new item can be inferred. Such methods include matrix factorization [8, 33], pairwise regression [21], decision trees [27] and recently the zero-shot learning approach that uses linear encode and decode structure [14].

Semantic search. It seeks to retrieve relevant items beyond keyword matching. The early effort of Latent Semantic Analysis (LSA) [6] uses a vector space to represent queries and documents. More recently, neural IR models [5, 18, 19] seek to apply deep learning techniques for building (query, item) scoring models. Supervised models [5, 19] are trained with search logs of (query, item) pairs. In contrast, unsupervised models [9] mostly learn the word and item representations purely based on the item’s text features. Recent work [31] shares the same motivation that search retrieval can learn from recommender system data. We point out several key differences between our paper and theirs. First being the different data assumptions, as their model is built on the (user, item) data with user embedding optimization. Our work does not require any explicit user data, and therefore, can fit to a wider range of applications. Second being the framework differences, as the recommender system data is not necessarily used for prediction target in our framework. Finally, our study focuses on real world search and recommender system with live experiment, which is not covered by their work.

3 PROBLEM STATEMENT

Denote a set of items $\mathcal{Y} = \{y_1, y_2, \dots, y_n\}$ as the corpus of a search and recommender system. Each item y_i has text feature $(x_1^{(i)}, x_2^{(i)}, \dots, x_{k_i}^{(i)})$, which are from a vocabulary of m words: $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$. The size k_i of the text feature can vary for different item. Also the text feature can be either ordered (i.e., sequence) or unordered (i.e., set). Finally, we use a binary matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ to represent (item, item) correlations, where $M_{i,j} = 1$ denotes item y_j is a neighbor of y_i : $j \in \text{Ne}(i)$. In this paper, we do not require $\mathbf{M}$ to be symmetric. For convenience, we also call the neighbor items as context items. When there is no confusion, we use the terms embedding, vector, and representation interchangeably.

Examples of the above problem setting include paper bibliography system, where items are individual papers, with their text features coming from titles, abstract and keywords. The (item, item) correlations can be derived by the citation network, where two papers are correlated if one cites the other. Similarly in an online music streaming system, the correlation matrix $\mathbf{M}$ can come from the music co-listening pairs, and the text features are the title, description or tags of the music.

The task is to retrieve relevant items given a query represented by a sequence of words $x_{e_1}, x_{e_2}, \dots, x_{e_p}$. This would be a traditional supervised learning task if the training data (query, item) were provided. However this data may not be available for newly built systems with limited user search activities. Or the data can be very sparse compared to the total number of queries and items. On the other hand, the (item, item) correlation data commonly exists in most systems [12]. Therefore, we propose to encode the query and item to the same latent vector space by utilizing such (item, item) correlations. Then the search retrieval problem becomes nearest neighbor search [15]. Formally, a model outputs the latent space representation of words $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_m] \in \mathbb{R}^{m \times d}$ and items $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n] \in \mathbb{R}^{n \times d}$. The query embedding is:

$$\mathbf{q} = \text{encoder}(\mathbf{w}_{e_1}, \mathbf{w}_{e_2}, \dots, \mathbf{w}_{e_p}), \quad (1)$$

where the choice encoder can be (but not limited to) Bag-of-Words (BOW) [17], which computes the mean of the word vectors, or (RNN) [4], or self-attention [28], which models the sequential relations of the words. Then the top- K candidate items are the ones with largest scores to the query: $\text{score}(\mathbf{q}, \mathbf{v})$, where the score function can be either vector dot product or cosine similarity.

4 ZERO-SHOT HETEROGENEOUS TRANSFER LEARNING FRAMEWORK

Before we introduce our proposed Zero-shot Heterogeneous Transfer Learning methods, we note that it is also possible to learn the semantic vector space of words and items by using only the item text features. The idea is to use the item's text feature as a proxy to the query. We can treat this learning task as a multi-class classification task. For instance, when using softmax to represent the probability of item y_i given the text feature $x_{e_1}, \dots, x_{e_k}$:

$$\Pr(y_i | x_{e_1}, \dots, x_{e_k}) = \frac{\exp(\mathbf{v}_i^T \cdot \mathbf{q})}{\sum_{\ell} \exp(\mathbf{v}_{\ell}^T \cdot \mathbf{q})},$$

where $\mathbf{q} = \text{encode}(\mathbf{w}_{e_1}, \dots, \mathbf{w}_{e_k})$. In this case, the total amount of training data equals the total number of items.

Consider the special case of a BOW encoder. This method generalizes beyond keyword matching by ensuring that the words with similar co-occurrence patterns will be close in the semantic space. For instance, the words "obama" and "president" are likely to co-exist in the same document. As a result, the optimization algorithm will not differentiate these two words much, so that they have similar word vectors. Similarly, items with similar text features will be encoded closely in the semantic space.

However, there are many cases that the words with similar semantic meanings do not co-occur very often. For instance, if the items rarely have both "funny" and "prank" as their text features, then the learned vectors of these two words will be very different, and therefore, by searching the query "funny", it is unlikely to retrieve the "prank" items. Fortunately, it is often common for pairs of items to form a connection in search and recommender system. For the above example, if users find these two kinds of items similar, they may link them as correlated items implicitly via interacting with the search and recommender system.

We propose to improve the target task of search retrieval, by transferring knowledge from (item, item) correlation, so that the semantic links between related words and items can be discovered.

4.1 Zero-shot Learning: Multi-task Item Encoder

Our first proposed way of transfer learning from (item, item) correlation matrix is to jointly optimize for the following two tasks:

- Task 1: Text features of the item $\xrightarrow{\text{predict}}$ item.
- Task 2: Neighbor items of the item $\xrightarrow{\text{predict}}$ item.

The specific model depends on the choice of prediction function and loss function. For instance, cross-entropy loss on a softmax prediction; square loss on a linear dot product prediction; pairwise ranking loss [30] with negative sampling. Here we only present the formulations with cross-entropy and square loss for simplicity.

Formally, the probability of predicting item y_i given text feature encoder $\mathbf{q}$ or given an item y_j is as follows:

$$\Pr(y_i | \mathbf{q}) = \frac{\exp(\mathbf{v}_i^T \cdot \mathbf{q})}{\sum_{\ell} \exp(\mathbf{v}_{\ell}^T \cdot \mathbf{q})}; \quad \Pr(y_i | y_j) = \frac{\exp(\mathbf{v}_i^T \cdot \mathbf{u}_j)}{\sum_{\ell} \exp(\mathbf{v}_{\ell}^T \cdot \mathbf{u}_j)},$$

where $\mathbf{q}_i = \text{encoder}(\mathbf{w}_{e_1}, \mathbf{w}_{e_2}, \dots, \mathbf{w}_{e_p})$ is the encoder of the text feature of y_i and $\mathbf{u}_j$ are the context vectors of y_j . The goal is to jointly optimize the cross-entropy (CE) loss for the above two tasks:

$$\mathcal{L}_{CE} = - \sum_{i=1}^n \left(\underbrace{\log(\Pr(y_i | \mathbf{q}_i))}_{\text{Task 1}} + \sum_{j \in \text{Ne}(i)} \underbrace{\log(\Pr(y_i | y_j))}_{\text{Task 2}} \right),$$

Alternatively when modeling them as regression problem, we can use the following weighted square loss (SL):

$$\begin{aligned} \mathcal{L}_{SL} = & \sum_{i=1}^n \left(\underbrace{(\mathbf{v}_i^T \cdot \mathbf{q}_i - 1)^2 + \omega_0 \sum_{\ell \neq i} (\mathbf{v}_i^T \cdot \mathbf{q}_{\ell})^2}_{\text{Task 1}} \right. \\ & \left. + \underbrace{\sum_{j \in \text{Ne}(i)} (\mathbf{v}_i^T \cdot \mathbf{u}_j - 1)^2 + \omega_0 \sum_{\ell \notin \text{NE}(i)} (\mathbf{v}_i^T \cdot \mathbf{u}_{\ell})^2}_{\text{Task 2}} \right), \end{aligned}$$

Figure 2: Illustration on the differences of the three models: Baseline Single Task Learning (STL) (left) that only trains Task 1 of ZSL_ME. Proposed ZSL_ME (middle) that jointly optimize the two tasks. Proposed ZSL_TE (right) that optimize over (item, item) correlation.

where $\omega_0 < 1.0$ is the weight for implicit observations [12].

The optimization output are word vectors $\mathbf{w}_i$, $1 \leq i \leq m$, item vectors $\mathbf{v}_i$, $q \leq i \leq n$ and context item vectors $\mathbf{u}_i$, $1 \leq i \leq n$. The context vectors are akin to that in the language modeling (see word2vec [17]). The reason we introduce context vectors instead of using the same item vectors is that, we want to encode two items to a close semantic space, if their neighbors are largely overlapped. By contrast, if we eliminate context vectors all together, and use item vectors for both prediction input and target: $\Pr(y_i|y_j) = \exp(\mathbf{v}_i^T \cdot \mathbf{v}_j) / \sum_{\ell} \exp(\mathbf{v}_i^T \cdot \mathbf{v}_j)$, this effectively encodes the item close to its neighbors. From a matrix factorization perspective, eliminating context vectors is similar to symmetric matrix factorization on $\mathbf{M}$ while $\mathbf{M}$ is not necessarily symmetric.

We call this framework Zero-shot Learning with Multi-task Item Encoder (ZSL_ME), as items are the shared prediction targets for both tasks. The item vectors $\mathbf{v}_i$ serve the role of bridging the two optimization tasks. Specifically, vectors $\mathbf{v}_i$ and $\mathbf{v}_j$ will be close in the semantic space if their corresponding correlation patterns (i.e., i -th and j -th row of $\mathbf{M}$) are similar. Word vectors on the other hand are indirectly affected to encode such correlation information from $\mathbf{M}$. Consider our previously example of words "funny" and "prank", where they do not co-occur often as the text features from the same item. However if their associated items share similar correlation pattern (i.e., similar neighbors of the items), those items will have similar embedding vectors (due to optimization on Task 2), and therefore the embedding vectors for these two words will also be close (due to optimization on Task 1).

4.2 Zero-shot Learning: Item-to-Text Transformed Encoder

Intuitively, there are two types of data relations, one is between the item and its text features, and other is between the item and its neighbor items (i.e., context items). The above ZSL_ME method treats both relations as prediction targets. In this section, we propose to model the (item, item) correlation as the only prediction task, and utilize the text features to encode context items.

Given the text feature $(x_{e_1}, x_{e_2}, \dots, x_{e_p})$ for an item, the context vector of this item is defined by

$$\mathbf{u} = \text{encoder}(\mathbf{w}_{e_1}, \mathbf{w}_{e_2}, \dots, \mathbf{w}_{e_p}). \quad (2)$$

So the cross-entropy loss can be computed as:

$$\mathcal{L}_{CE} = - \sum_{i=1}^n \sum_{j \in \text{Ne}(i)} \log(\Pr(y_i|y_j)); \quad \Pr(y_i|y_j) = \frac{\exp(\mathbf{v}_i^T \cdot \mathbf{u}_j)}{\sum_{\ell} \exp(\mathbf{v}_i^T \cdot \mathbf{u}_{\ell})}.$$

And similarly the weighted square loss is:

$$\mathcal{L}_{SL} = \sum_{i=1}^n \left(\sum_{j \in \text{Ne}(i)} (\mathbf{v}_i^T \cdot \mathbf{u}_j - 1)^2 + \omega_0 \sum_{\ell \notin \text{NE}(i)} (\mathbf{v}_i^T \cdot \mathbf{u}_{\ell})^2 \right).$$

Since the context vectors are represented by word vectors, the overall output from the above model are just word vectors $\mathbf{w}_i$, $1 \leq i \leq m$ and item vectors $\mathbf{v}_i$, $1 \leq i \leq n$. We call this framework Zero-shot Learning with Item-to-Text Transformed Encoder (ZSL_TE), as the optimization on the correlation matrix trains the item and context item representations, which pass to the text representation due to a transformed encoder on text.

Here we explain our intuition of this proposed method. Text features of a item are usually derived directly from the title or description of the item, which are precise information about the item. In this case, it is well fit to use them as item encoder instead of prediction target. On the other hand, the correlation matrix $\mathbf{M}$ represents how much two items are similar or related content-wise. This relation usually extends beyond the text feature similarities. Therefore they are useful to be the prediction target. To demonstrate how this method generalizes to semantic related words that do not co-occur often, consider the same example of "funny" and "prank". If their associated items (denote y_r and y_s correspondingly) co-occur as the neighbors of a same item y_p , then both $\mathbf{u}_r$ and $\mathbf{u}_s$ will be brought closer from the terms $\mathbf{v}_p^T \cdot \mathbf{u}_r$ and $\mathbf{v}_p^T \cdot \mathbf{u}_s$ in the loss functions. Since $\mathbf{u}_r$ and $\mathbf{u}_s$ are encoded by term "funny" and "prank" correspondingly, these two words end up close in the semantic space. Figure 2 illustrates how each model is different.

4.3 Model Discussion

It is worth mentioning that Neural Collaborative Filtering (NCF) [11] method applies multi-layer perception instead of dot-product as an embedding combiner. We do not compare these two methodologies, as the choice of embedding combiner is not the central importance of this paper. We refer the interested readers to [23] for more details.

In the following, we discuss more details on the model options, training algorithms and practical usage examples.

Regularization. To prevent model overfitting, the regularization term $\lambda \|\theta\|_F^2$ can be added to the model loss functions. The above λ is the co-efficient for the regularization term, and the symbol $\|\cdot\|_F^2$ denotes the matrix Frobenius norm, and θ denotes the matrix of all parameters. $\theta = [\mathbf{w}_1, \dots, \mathbf{w}_m, \mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{v}_1, \dots, \mathbf{v}_n]$ for ZSL_ME and $\theta = [\mathbf{w}_1, \dots, \mathbf{w}_m, \mathbf{v}_1, \dots, \mathbf{v}_n]$ for ZSL_TE.

Weighted Training. Real word correlation data are usually skewed towards popular items (a.k.a. power law) [7]. In other words, the number of non-zeros (i.e., nnz) of the rows (or columns) in $\mathbf{M}$ can be concentrated in only a few rows (or columns). We can reweight the training example to avoid having the loss function dominated by only a few most popular items:

$$\sum_{i=1}^n \sum_{j \in \text{Ne}(i)} \text{loss}(y_i, y_j) \xrightarrow{\text{reweight}} \sum_{i=1}^n \sum_{j \in \text{Ne}(i)} r_i c_j \text{loss}(y_i, y_j),$$

where r_i are row weights and c_j are column weights. In this paper, we set them to be proportional to $1/\sqrt{\text{nnz}}$ of the corresponding rows and columns, and rescale them so that the mean values of both row weights and column weights are 1.0.

Choice of Encoder. The encoder is applied on the text features of an item during training, as well as on the query words during serving. Depending on the text feature characteristics (e.g., ordered vs unordered, long vs short) and the specific application domain, various encoders can be used. Bag-of-Word (BOW) is simplest to use as it doesn't have underlying requirement on the text feature format. RNN can be used to encode a sequence of words. And self-attention mechanism can usually work well for long sentence or document [4]. In this paper, we use BOW in our experiments for simplicity and also for reducing the complexity in order to work with live experiment serving constraints.

Loss Function & Optimization. For large scale search and recommender systems, the number of items can be millions or beyond. This imposes computational challenges when computing the softmax function for $\mathcal{L}_{CE}$ or iterating through all negative examples for $\mathcal{L}_{SL}$, which yields $O(n^2)$ complexity. Approximation solutions include sampled softmax, hierarchical softmax and negative sampling. They share the core idea of avoiding to exhaustively iterate through all items, so that optimization like stochastic gradient descent (SGD) can be applied. In a special case using BOW encoder with square loss $\mathcal{L}_{SL}$, the first-order and second-order derivative of the full loss function can be efficiently computed, without the need of explicitly iterating through $O(n^2)$ negative examples [1]. Therefore the optimization problem can be solved by the coordinate descent algorithm. With this advantage of computing the full negative examples without the need to do negative sampling, we therefore choose to use $\mathcal{L}_{SL}$ instead of $\mathcal{L}_{CE}$ in our experiments.

5 EXPERIMENTS

We first conduct a set of offline evaluations (Sections 5.1, 5.2, 5.3) among our proposed methods ZSL_ME and ZSL_TE and baseline methods. Then we conduct the live experiment (Section 5.4) with the best ZSL method on one of the world's largest search and recommender systems from Google.

Dataset. Each item y_i is a product of the recommender system. We derive the correlation matrix $\mathbf{M}$ from sequential item consumptions. Formally, if y_p is consumed right after y_q by the same user, then y_p is the neighbor of y_q (i.e., $M_{q,p} = 1$). To reduce the noise of the correlation matrix, we rank each seed item's neighbors by the counts of their co-occurrences with the seed item, and only keep the top 250 of them. So each row of $\mathbf{M}$ has at most 250 non-zeros. Titles and descriptions are used as the text features of each item. The word vocabulary contains both unigrams and bigrams from

	STL	SMC	ZSL_ME	ZSL_TE
Recall (%)	17.8	13.8	27.4	35.0

Table 1: Recalls of the correlation matrix reconstruction task. Note a random guess would only yield 0.001% recall in finding candidates from a pool over 17 million items.

those text features. We threshold the minimal occurrences for items and words, and the final vocabulary size is 17,714,821 for items, 2,451,962 for words. The average number of neighbors for an item is 176, and the average number of words of an item is 155.

Experiment Settings. One baseline model is the single task learning version of ZSL_ME, where it only trains on Task 1 that predicts the item given its text features. We call it Single Task Learning (STL). It is important to note that, STL is proven to be a more effective approach than traditional methods like LSI in many of the recommendation tasks [2, 12] as it models implicit feedbacks. In addition we also trained a supervised multiclass classification model directly based on the actual search data. Each record (query, item) corresponds to one item consumption from the search query. Formally, the word and item representations are trained to minimize the cross entropy loss of the sampled softmax. Although this supervised method is outside of the zero-shot framework of this paper, it is interesting to compare them and study how they are different. Here are the detailed setting for each method:

- **STL:** single-task square loss $\mathcal{L}_{SL}$ without combining (i.e., encoder) text feature; $\omega_0 = 0.001$, $\lambda = 4.0$; trained with 10 iterations of coordinate descent.
- **ZSL_ME:** same setting as above multi-task square loss.
- **ZSL_TE:** square loss $\mathcal{L}_{SL}$ with the BOW encoder; same setting as the above two.
- **Supervised Multiclass Classification (SMC):** BOW encoder on words; trained with a total of 1780 millions (query, item) pairs; sampled softmax with 10000 negative samples for each batch with batch size 128; trained with SGD optimizer with learning rate 0.06 for 50 million iterations.

The above hyperparameters are set based on empirical hyperparameter search. For all methods, the embedding dimension is set as 200. The following sections include the semantic retrieval task: given a seed vector $\mathbf{q}$ (e.g., query or item), the task is to retrieve top- K items $\mathbf{v}$ with the highest score($\mathbf{q}, \mathbf{v}$). By default (unless otherwise stated) we use cosine score for STL, ZSL_ME and ZSL_TE, and vectors dot product for SMC¹.

5.1 Correlation Matrix Reconstruction Task

First we evaluate how each method performs in terms of reconstructing the correlation matrix $\mathbf{M}$. We use recall to evaluate percentage of relevant items being retrieved at top positions. This is commonly used for retrieval and top-n item recommendation models. Specifically, for each item $y_i \in \mathcal{Y}$, denote $\mathcal{S}_{true,i}$ as the non-zeros in the i -th row of the correlation matrix $\mathbf{M}$ and k_i as the size of $\mathcal{S}_{true,i}$. Denote $\mathcal{S}_{pred,i}$ as the retrieved top- k_i items $\mathbf{v}_\ell$

¹We found dot product works much better than cosine for SMC for all tasks. More discussions in Section 5.3.

	#query	S	average $ S_i $	$\sum_{i=1}^r S_i / S $
Eval Set 1	202	23,377	136.5	1.18
Eval Set 2	125	15,842	174.5	1.38
Eval Set 3	318	40,267	144.7	1.14
Eval Set 4	36	1,797	102.1	2.05

Table 2: Statistics of the evaluation set from human labeled data. Each column represents: the total number of the queries; size of the candidate pool (i.e., the joint of target sets of all queries); average size of target set; average number of appearance for items.

with the highest score (v_i, v_ℓ), the recall for item y_i is defined as:

$$\text{recall}_i = \frac{|S_{\text{true},i} \cap S_{\text{pred},i}|}{|S_{\text{true},i}|},$$

then we report the average recall over all items. We believe using recall for evaluation is more intuitive than directly comparing the square losses of different methods, because the square loss greatly depends on the hyperparameters, and the value itself does not reflect any application meaning.

Results. Table 1 shows the recalls of all methods. We find the two proposed transfer learning methods ZSL_ME and ZSL_TE outperform STL and SMC. This is expected, as STL and SMC are trained without any information of the (item, item) correlation data. This observation shows that the outputs from our proposed framework are indeed influenced by the recommender system data in a positive way. Finally, we notice that between the two proposed methods, ZSL_TE is superior than ZSL_ME. We also observe this pattern for the other tasks as well in this paper. Our hypothesis is that directly optimizing two tasks in ZSL_ME could potentially introduce conflicts [24], while ZSL_TE does not have this issue.

5.2 Offline Retrieval Task on Human Labeled Data

In this section, we evaluate how relevant are the retrieval results when using human labeled (query, item) ground truth pairs. We have a list of queries and their corresponding items that are labeled as relevant by human. We then form a candidate pool as the union of all relevant items from these queries, and a target set for each query as its relevant items.

Formally, consider r different queries, and each query corresponds to a set of relevant items denoted as $S_1 = \{y_1^{(1)}, y_2^{(1)}, \dots, y_{k_1}^{(1)}\}$, $S_2 = \{y_1^{(2)}, y_2^{(2)}, \dots, y_{k_2}^{(2)}\}, \dots, S_r = \{y_1^{(r)}, y_2^{(r)}, \dots, y_{k_r}^{(r)}\}$. We can denote their joint set as $S = S_1 \cup S_2 \cup \dots \cup S_r$. Given a query (for instance the i -th one), we select k_i items from S with the highest score to the query, and denote it as $Z_i = \{\tilde{y}_1^{(i)}, \tilde{y}_2^{(i)}, \dots, \tilde{y}_{k_i}^{(i)}\}$. Then the recall for this query is defined as:

$$\text{recall}_i = \frac{|S_i \cap Z_i|}{|Z_i|}.$$

Intuitively, the above recall denotes the percentage of relevant k items are among the top- k predictions.

We note that it is possible for certain items to exist in multiple the ground truth set $S_i, i = 1, \dots, r$ (i.e., they are not disjoint). This actually makes the task more difficult, because we select the group of queries to share some common attributes, so that for each

Figure 3: Comparison of different methods on the recall for the offline retrieval task on human labeled data. For each evaluation set, the average recall of all queries in this set is reported.

query, the unrelated items (i.e., related to other queries) are not too distinct. We can define the average number of appearance for items as $\sum_{i=1}^r |S_i|/|S|$ in order to show the degree of overlapping on these sets. We have four such datasets, with each representing one category of the queries. For instance, one dataset has all its queries and items being music related. See Table 2 for their statistics.

Results. Figure 3 shows the comparison results for all methods on this task. We can see that our proposed transfer learning approaches ZSL_ME and ZSL_TE outperforms STL. This shows incorporating the (item, item) information from the recommender system can improve the semantic search retrieval task. Consistently, between these two transfer learning approaches, ZSL_TE performs much better. We also notice that ZSL_TE even achieves higher recall than the supervised method SMC. The explanation is that, although the supervised method is directly trained on (query, item) data, but the evaluation sets are extracted from human labeled data, which encodes only the relevance information between query and item. By comparison, the actual user search data (query, item) is affected by popularity of the item or content of the item. Because for a large scale recommender system, hundreds of thousands of items can be relevant to a query, but only those that are appealing enough will be clicked by users. We will later discuss evaluation results in Section 5.3 on the actual search data.

5.3 Offline Retrieval Task on Search Data

In this section, we use the ground-truth pairs (query, item) from the search logs. Note that this is the same data source as the training data of SMC. We hold out 1 million such pairs for evaluation. Different from the human labeled data, the search data reflects not only relevance, but also users' preferences. For large-scale search and recommender system, certain items can be very relevant to a query, but may not necessarily be preferred by users.

We use the metric recall@K for evaluation, which is defined as the ratio of the ground-truth items at the top-K retrieved list of the method.

Results. Figure 4 shows the comparison results of different methods for this task. SMC is not presented in the figure as it is the supervised method that is trained on the exact same data source of the evaluation task, so it is expected to outperform all the unsupervised methods by large margins. In our case, SMC could reach 73.6% for recall@300. We can clearly notice that our proposed ZSL_TE achieves superior recalls compared to ZSL_ME and STL. So far, all

Figure 4: Recall@K for the evaluation task on search data. Dotted lines are two added methods used for demonstration purpose. SMC(cosine) is defined as using the SMC method for retrieval with cosine similarity instead of dot product as score function. ZSL_TE(rescale) is defined as rescaling the item vectors of ZSL_TE method by borrowing the norms of item vectors learned in SMC, and the retrieval score function is dot product.

	SMC	Ensemble	Ensemble (rescale)
Recall@300 (%)	73.6	75.4	76.2

Table 3: Recall@300 for ensembling ZSL_TE or ZSL_LE (rescale) with SMC for the offline retrieval task on search data. Here the recall is improved even with the naive ensemble rule.

the offline evaluations show that ZSL_TE performs consistently better than ZSL_ME, and in this task, ZSL_ME is even slightly worse than STL. Our hypothesis is the same as stated in section 5.1 that multitask optimization could introduce additional conflicts, which is commonly observed in related researches [24]. We also notice that, if we change the SMC retrieval method from vector dot product to cosine (as is used by all other methods), the recalls are beat by our ZSL_TE. This provides us an interesting insight that the vector norms in the SMC method yield useful information about user preferences over items. To verify this hypothesis, we rescaled the item vectors of ZSL_TE based on the norms of SMC, and find out the recall is improved.²

In practice, our proposed methods are not meant to replace the existing supervised models, but to serve as additional sources of candidates generation. Here we show that combining our method to SMC can improve the recall, even with the following trivial ensemble rule. The new item list has the first 150 items the same as SMC, then the next 150 items are from ZSL_TE and the rest of SMC in a interleaved way. Table 3 shows the comparison results. It is also important to note that the benefits of ensembling our ZSL_TE with supervised methods can be more pronounced, because the current evaluation is done on the offline search data, which is already limited in scope as it does not answer the question of what if we recommend the user the other item instead of the current one. This effect can only be assessed from the live experiment (Section 5.4).

²The same observation holds also for ZSL_ME and STL. They are skipped in Figure 4 for the sake of simplicity.

	A/B Test Diff	Notes
Query Coverage	+0.51 (± 0.27)%	
User Interaction	+2.52 (± 1.54)%	
Query Refinement	-0.94 (± 0.57)%	lower is better
Next-page CTR	-1.60 (± 0.75)%	lower is better
Human Rater Score	+14.0 %	

Table 4: Live experiment metrics. The bold numbers are statistically significant, and the numbers after $\pm$ are the standard deviations.

5.4 Live Experiment

Settings. In this section, we evaluate ZSL_TE in an A/B live experiment. The control group is the production search system that is highly-optimized with many components of advanced retrieval techniques (e.g., both term-frequency-based scoring and supervised machine learning algorithms). The experiment group is an ensemble of our ZSL_TE and the production system. Each time when a query comes, the ZSL_TE retrieves top-100 items, and the ensemble algorithm will jointly rank (i.e., based on query features as well as item features) these items with those retrieved by the control production system.

Model Daily Refresh. We do warm-start training with the new data everyday to include new items and words to the model. Specifically, we load the existing model and only need to conduct a few model training iterations to make the embeddings of new items and words converge.

Evaluations. We conduct this live experiment over several millions of queries. We compare the following evaluation metrics:

- Query Coverage: The ratio of queries that receive at least one user interaction.
- User Interaction: The level of interactions between users and the items.
- Query Refinement: The proportion of queries that have a follow up query sharing a keyword. Smaller value is better (meaning users are more satisfied with the results).
- Next-page CTR: The proportion of queries that lead to a next page click by users. Smaller value is better (meaning users are satisfied with the first page of results).
- Human-rated Relevance Score: A numerical score representing the relevance of the retrieved items to a query, as evaluated by trained human raters. We randomly sample 250 instances, each with a query and a list of retrieved items, for both the control group and the experiment group.

We also have several ranking metrics that are similar to the normalized discounted cumulative gain (nDCG), but they are tailored to our specific system, therefore less generic. We do not include them in this paper to avoid confusion. But it is important to mention that our live experiment also shows significant improvements on these metrics as well.

Results. As shown in Table 4, we observe significant improvements including increased query coverage, decreased query refinements (meaning users are more satisfied with the results), increased user interaction, and higher relevance scores from human raters. These results demonstrate that the proposed approach is effective on semantic search task, even without training on any search data.

Figure 5: Live experiment metrics break down by the length of query. The lines represent the mean values, and the colored fills are standard deviations. Two different colors to differentiate higher (blue) or lower (red) is better for the corresponding metrics.

We also notice that the our method has a larger improvement when the query length (number of unigrams) is small (see Figure 5). Since shorter queries often correspond to broader user intent, there are usually more relevant items per query, which means more (query, item) pairs were previously unseen in the supervised training data. This is an evidence that, our zero-shot transfer learning framework can result in more improvement by solving the cold start problem.

6 CONCLUSIONS AND LESSONS LEARNED

In this paper, we explore a Zero-Shot Heterogeneous Transfer Learning framework, that trains a model to learn the natural-language and item representations based on the (item, item) correlations derived from the recommender system, and then uses the learned representations to serve the search retrieval task.

Here are several lessons we learned. 1) Both the proposed ZSL_ME and ZSL_TE show evidences that it is beneficial for search retrieval system to transfer learn from the recommender system. However ZSL_LE works much better with the real data in this paper. This is a valuable finding, as it suggest that treating the text features as rules and predicting the (item, item) correlations works better in practice than predicting both correlations. 2) We find the effectiveness of the supervised method SMC is largely due to "memorization" based on our observation that the norms of embeddings play a crucial role. This is evidence that supervised methods such as SMC are less capable of generalizing beyond seen (query, item) pairs. 3) Most importantly, we find in live experiment, that our proposed ZSL_TE are more effective for broad queries (as inferred by query length). This insight suggests the direction of investigating different query types, so that our model can be deployed in a most effective way.

REFERENCES

- [1] Immanuel Bayer, Xiangnan He, Bhargav Kanagal, and Steffen Rendle. 2017. A generic coordinate descent framework for learning from implicit feedback. In *WWW*. 1341–1350.
- [2] Jesús Bobadilla, Fernando Ortega, Antonio Hernando, and Abraham Gutiérrez. 2013. Recommender systems survey. *Knowledge-based systems* 46 (2013), 109–132.
- [3] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H Chi. 2019. Top-K Off-Policy Correction for a REINFORCE Recommender System. In *WSDM*. 456–464.
- [4] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* (2014).
- [5] Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. Convolutional neural networks for soft-matching n-grams in ad-hoc search. In *WSDM*. 126–134.
- [6] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science* 41, 6 (1990), 391–407.
- [7] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. 1999. On power-law relationships of the internet topology. In *ACM SIGCOMM computer communication review*, Vol. 29. ACM, 251–262.
- [8] Zeno Gantner, Lucas Drumond, Christoph Freudenthaler, Steffen Rendle, and Lars Schmidt-Thieme. 2010. Learning Attribute-to-Feature Mappings for Cold-Start Recommendations. In *ICDM*, Vol. 10. Citeseer, 176–185.
- [9] Christophe Van Gysel, Maarten De Rijke, and Evangelos Kanoulas. 2018. Neural vector spaces for unsupervised information retrieval. *TOIS* 36, 4 (2018), 1–25.
- [10] Malay Haldar, Mustafa Abdool, Prashant Ramanathan, Tao Xu, Shulin Yang, Huizhong Duan, Qing Zhang, Nick Barrow-Williams, Bradley C Turnbull, Brendan M Collins, et al. 2019. Applying deep learning to Airbnb search. In *KDD*. 1927–1935.
- [11] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW. International World Wide Web Conferences Steering Committee*, 173–182.
- [12] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *ICDM*, Vol. 8. 263–272.
- [13] Xuan Nhat Lam, Thuc Vu, Trong Duc Le, and Anh Duc Duong. 2008. Addressing cold-start problem in recommendation systems. In *IMCOM*. ACM, 208–211.
- [14] Jingjing Li, Mengmeng Jing, Ke Lu, Lei Zhu, Yang Yang, and Zi Huang. 2019. From Zero-Shot Learning to Cold-Start Recommendation. *Age* 25 (2019), 30.
- [15] Ting Liu, Andrew W Moore, Ke Yang, and Alexander G Gray. 2005. An investigation of practical approximate nearest neighbor algorithms. In *NeurIPS*. 825–832.
- [16] Yang Long, Li Liu, Yuming Shen, and Ling Shao. 2018. Towards affordable semantic searching: Zero-shot retrieval via dominant attributes. In *AAAI*.
- [17] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*. 3111–3119.
- [18] Bhaskar Mitra, Nick Craswell, et al. 2018. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* 13, 1 (2018), 1–126.
- [19] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to match using local and distributed representations of text for web search. In *WWW*. 1291–1299.
- [20] Seungwhan Moon and Jaime G Carbonell. 2017. Completely Heterogeneous Transfer Learning with Attention-What And What Not To Transfer. In *IJCAI*, Vol. 1. 1–2.
- [21] Seung-Taek Park and Wei Chu. 2009. Pairwise preference regression for cold-start recommendation. In *RecSys*. ACM, 21–28.
- [22] Rohan Ramanath, Hakan Inan, Gungor Polatkan, Bo Hu, Qi Guo, Cagri Ozcaglar, Xianren Wu, Krishnaram Kenthapadi, and Sahin Cem Geyik. 2018. Towards Deep and Representation Learning for Talent Search at LinkedIn. In *CIKM*. 2253–2261.
- [23] Steffen Rendle, Walid Krichene, Li Zhang, and John Anderson. 2020. Neural Collaborative Filtering vs. Matrix Factorization Revisited. *arXiv preprint arXiv:2005.09683* (2020).
- [24] Ozan Sener and Vladlen Koltun. 2018. Multi-task learning as multi-objective optimization. In *NeurIPS*. 527–538.
- [25] Ayan Sinha, David F Gleich, and Karthik Ramani. 2016. Deconvolving feedback loops in recommender systems. In *NeurIPS*. 3243–3251.
- [26] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-shot learning through cross-modal transfer. In *NeurIPS*. 935–943.
- [27] Mingxuan Sun, Fuxin Li, Joonseok Lee, Ke Zhou, Guy Lebanon, and Hongyuan Zha. 2013. Learning multiple-question decision trees for cold-start recommendation. In *WSDM*. ACM, 445–454.
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*. 5998–6008.
- [29] Wei Wang, Vincent W Zheng, Han Yu, and Chunyan Miao. 2019. A Survey of Zero-Shot Learning: Settings, Methods, and Applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 2 (2019), 13.
- [30] Fabian Wauthier, Michael Jordan, and Nebojsa Jojic. 2013. Efficient ranking from pairwise comparisons. In *ICML*. 109–117.
- [31] Hamed Zamani and W Bruce Croft. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *WSDM*. 717–725.
- [32] Andrew Zhai, Hao-Yu Wu, Eric Tzeng, Dong Huk Park, and Charles Rosenberg. 2019. Learning a Unified Embedding for Visual Search at Pinterest. In *KDD*. 2412–2420.
- [33] Ke Zhou, Shuang-Hong Yang, and Hongyuan Zha. 2011. Functional matrix factorizations for cold-start recommendation. In *SIGIR*. ACM, 315–324.