

# Domain Adaptation for Enterprise Email Search

Brandon Tran<sup>1\*</sup>, Maryam Karimzadehgan<sup>2</sup>, Rama Kumar Pasumarthi<sup>2</sup>, Michael Bendersky<sup>2</sup>,  
Donald Metzler<sup>2</sup>

<sup>1</sup>Department of Mathematics, Massachusetts Institute of Technology, MA, USA

<sup>2</sup>Google LLC, Mountain View, CA, USA

<sup>1</sup>btran115@mit.edu    <sup>2</sup>{maryamk, ramakumar, bemike, metzler}@google.com

## ABSTRACT

In the enterprise email search setting, the same search engine often powers multiple enterprises from various industries: technology, education, manufacturing, etc. However, using the same global ranking model across different enterprises may result in suboptimal search quality, due to the corpora differences and distinct information needs. On the other hand, training an individual ranking model for each enterprise may be infeasible, especially for smaller institutions with limited data. To address this data challenge, in this paper we propose a domain adaptation approach that fine-tunes the global model to each individual enterprise. In particular, we propose a novel application of the Maximum Mean Discrepancy (MMD) approach to information retrieval, which attempts to bridge the gap between the global data distribution and the data distribution for a given individual enterprise. We conduct a comprehensive set of experiments on a large-scale email search engine, and demonstrate that the MMD approach consistently improves the search quality for multiple individual domains, both in comparison to the global ranking model, as well as several competitive domain adaptation baselines including adversarial learning methods.

## CCS CONCEPTS

• **Information systems** → **Learning to rank; Enterprise search; Query representation; Novelty in information retrieval;**

## KEYWORDS

Enterprise Search; Domain Adaptation; Learning-to-Rank

## ACM Reference Format:

Brandon Tran, Maryam Karimzadehgan, Rama Kumar Pasumarthi, Michael Bendersky, Donald Metzler. 2019. Domain Adaptation for Enterprise Email Search. In *Proceedings of the 42nd Int'l ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19), July 21–25, 2019, Paris, France*. ACM, NY, NY, USA, 10 pages. <https://doi.org/10.1145/3331184.3331204>

## 1 INTRODUCTION

Traditionally, enterprise email search engines were installed locally on the premises of the organization. In these installations, search

\*Work done while Brandon Tran was at Google.


Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR '19, July 21–25, 2019, Paris, France

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6172-9/19/07.

<https://doi.org/10.1145/3331184.3331204>



**Figure 1: A visualization of our enterprise email search data. The set of all data (the source domain) consists of multiple individual enterprise domains (target domains), with four examples in the figure labeled with A, B, C, and D.**

ranking was generally predetermined by the vendor and kept fixed. In the last several years, cloud-based search engines (e.g., Microsoft Azure, Amazon CloudSearch, or Google Cloud Search) have been gaining traction as an effective tool for search. In these cloud solutions, the corpora, the ranking models, and the search logs are stored in the cloud. This enables the cloud search providers to optimize the quality of their search engines based on user click data, similarly to what was previously done in web search [25].

The transition to the cloud and the abundance of available user interaction data provide a unique opportunity to significantly improve the quality of enterprise email search engines, which traditionally lagged behind web search [23]. In particular, in recent years, deep neural learning-to-rank models were shown to significantly improve the performance of search engines in the presence of large-scale query logs, both in web search [22] and in email search [39, 46] settings.

However, directly applying these advances in neural learning-to-rank models to enterprise email search is not straightforward. An important difference between web and enterprise email search is that in the latter, the model can be applied to multiple, often very different domains, as described in Figure 1. For instance, the same enterprise email search engine can power enterprises from various industries: technology, education, manufacturing, etc. Therefore, using the same global learning-to-rank model across the different enterprises may lead to suboptimal search quality for each individual enterprise, due to the differences between their corpora and information needs.

On the other hand, training an individual ranking model for each industry or enterprise may not be feasible, especially for smaller enterprises with limited search traffic data. This is especially true for deep learning-to-rank models that typically require large amounts of training data. The reason for this requirement is that deep neural

networks are susceptible to overfitting to the training data, and they need enough inputs to actually learn a model rather than simply memorizing the input examples.

As such, a natural question to ask is how to make use of the state-of-the-art deep learning-to-rank models in the enterprise search setting. To this end, we propose the use of domain adaptation techniques [15, 32] to adapt a global model, trained on the entire data, to a particular enterprise domain. In this work, we specifically focus on the enterprise email search setting; however our findings easily generalize to other enterprise search scenarios as well.

Domain adaptation, at a high level, deals with the ability to adapt efficient, high-performing models trained on one domain to perform well on a different domain. Typically, the first domain, called the source domain (in our case the entire dataset – see Figure 1), contains a wealth of data, and so an effective prediction model can be trained. However, due to what is known as dataset bias or domain shift [19], these models do not immediately generalize to new datasets, referred to as the target domains (in our case individual enterprises). The target domains are expected to be significantly smaller than the source domain, so that a model cannot simply be trained with only target training data due to overfitting. As a result, these fully-trained networks are typically fine-tuned to the new dataset. That is, the available labeled data from the target domain is used to slightly alter the parameters of the original model to fit new data. This is, of course, difficult and expensive to carry out.

Work in domain adaptation, then, attempts to reduce the harmful effects of this domain shift. The deep learning model maps both the source and target domain into some latent feature space. Reduction of the domain shift is then accomplished by either minimizing some measure of domain shift, such as maximum mean discrepancy (MMD) [32, 43], or with adversarial adaptation methods [15, 30, 41, 42]. In the latter case, the model is trained to make the two mappings of the source and target domain indistinguishable in feature space to a discriminator model.

While domain adaptation methods have been previously studied, we emphasize that the majority of the research done was in other areas such as image classification. Moreover, the typical focus was on an unsupervised setting with no labeled data from the target domain. We note, though, that there has been one related work by Cohen et al. [9], studying the problem of domain adaptation in learning-to-rank models.

However, there are two major differences between these prior works, including [9], and the problem we study. First, the enterprise email search problem deals with datasets that are *several orders of magnitude* larger than the previous work. The source domain, consisting of the combined inputs from all individual enterprise domains, contains  $O(100M)$  inputs as compared to  $O(10K)$  inputs in the image classification datasets [15, 30, 32, 41–43] and the prior learning-to-rank work [9]. Additionally, we deal with labeled enterprise domains in a weakly supervised setting (using user click data), whereas the aforementioned prior works all assumed unlabeled target domains.

These differences from prior work lead to a more realistic setting for exploring domain adaptation techniques in information retrieval. As we show through extensive experimentation, in this setting MMD outperforms all other domain adaptation techniques, including the state-of-the-art adversarial methods [9], a first such

result in the information retrieval literature. In summary, our key contributions are:

- We propose a general framework for learning-to-rank with domain adaptation, with a particular application to enterprise email search.
- We experimentally demonstrate the shortcomings of simple transfer learning methods, such as re-training or batch-balancing, to individual enterprise domains.
- We propose a novel use of the maximum mean discrepancy (MMD) method for learning-to-rank with domain adaptation, and demonstrate its effectiveness and robustness in the enterprise email search setting.
- We perform a thorough comparative analysis of various domain adaptation methods on realistic, large-scale enterprise email search data.

The rest of our paper is organized as follows. First, we discuss related work in Section 2. Next, in Section 3, we provide motivating evidence that using domain adaptation techniques is feasible for our setting. That is, we show that the distributions of representations of the source and target datasets have nontrivial overlap, and so it is reasonable to try and encourage the model to accurately predict clicks on both sets. In Section 4, we give a detailed explanation of our methodology. Then, we provide our extensive experimental study in Section 5. Finally, we conclude and discuss future work in Section 6.

## 2 RELATED WORK

We split our discussion of related works into four distinct parts. First, we discuss work done in the learning-to-rank literature on which we base our learning-to-rank models. Next, we review research done in enterprise email search setting. Then, we mention work done in developing techniques for domain adaptation in image classification. Finally, we point out other works trying to utilize domain adaptation for information retrieval problems.

### 2.1 Learning-to-Rank

Generally, learning-to-rank refers to the application of machine learning tools and algorithms to rank models for information retrieval. There is a vast literature of learning-to-rank work [5–7, 14, 25, 45], differing in their model and loss function constructions. Recently, with the rise in popularity of deep neural networks, work has been done to use deep neural networks for learning-to-rank [4, 10]. For a complete literature review on neural ranking models for information retrieval, please refer to the survey by Mitra [34].

### 2.2 Enterprise Search

Enterprise search can broadly be viewed as the application of information retrieval techniques towards the specific problem of searching within private organizations. The majority of prior work can be found in the recent surveys by Kruschwitz and Hull [27] and Hawking [23]. To the best of our knowledge, no previous research on enterprise search has studied the problem as an application of domain adaptation.

Enterprise search is also closely related to personal search (e.g., email search), as both deal with searching in private or access

controlled corpora [2, 8, 12, 20, 26, 39, 44]. Even though some success has been found using time-based approaches for personal search [12], relevance-based ranking arising from learning-to-rank deep neural network models has become increasingly popular [39, 46] as the sizes of private corpora increase [20]. However, to the best of our knowledge, our work is the first study on applying deep neural networks specifically in the enterprise search setting.

### 2.3 Domain Adaptation

Extensive prior work on domain adaptation has been done in image classification [19]. These works develop methods to transfer latent representations obtained from deep neural networks from a large, labeled source dataset to a smaller, unlabeled target dataset. The primary strategy focuses on guiding the learning process by encouraging the source and target feature representations to be indistinguishable in some way [15, 17, 30, 32, 40, 41, 43].

One focus of domain adaptation research in image classification aims on minimizing the differences of certain statistics between the source and target distributions. Several methods used the Maximum Mean Discrepancy (MMD) [19] loss, in a variety of ways [32, 40, 43]. The MMD computes the norm of the difference between two domain means. Another direction lies in choosing adversarial losses to minimize domain shift. The goal is to learn a representation from which a classifier can predict source dataset labels while a discriminator cannot distinguish whether a data point is from the source or target datasets. The work in this direction focuses on finding better adversarial losses [15, 17, 41].

### 2.4 Applications of Domain Adaptation to Information Retrieval

As deep learning models became more prevalent in ranking problems, and as more and more transfer learning techniques developed for image classification, work began to study the problem of domain adaptation in the information retrieval domain. The models in Cohen et al. [9] are trained for domain adaptation in ranking with adversarial learning. Specifically, the models are trained using the gradient reversal layers of Ganin et al. [16]. We note that this work focused only on adversarial learning and did not consider maximum mean discrepancy. Moreover, it only compared their adversarial learning technique with very simple baselines: training on all data and training on target data. The work did not consider more interesting baselines such as balancing training batches with a certain number of target data inputs, or fine-tuning a previously trained model on all data with only target data. Additionally, while in the learning-to-rank setting, their datasets are significantly smaller and not as complex as those arising from enterprise search. Similarly, the work of Long et al. [31] also utilizes adversarial learning to solve the problem of domain adaptation to a number of information retrieval tasks, including digit retrieval and image transfers. However, as with Cohen et al. [9], they study only adversarial learning, and their data are also significantly different from enterprise search. Lastly, Mao et al. [33] use a similar statistics-based approach for transferring features from source to target domains. Their method also looks at a certain mean distance computed from the embedded data, but their per-category multi-layer joint kernelized mean distances are quite distinct from an MMD regularization term.

## 3 MOTIVATION

In this section, we provide motivation for using domain adaptation techniques in enterprise email search. First, we take our enterprise email search inputs and map them into a high-dimensional space. As detailed in Section 4.2, we refer to the resulting subset of the high-dimensional space obtained from this mapping as the embedding of our inputs. These embeddings are then passed as inputs into the prediction models, as discussed in Section 4.3. The high-level goal of domain adaptation techniques is to make the embeddings arising from the source and target domains indistinguishable to the prediction models. That way, the model can leverage information from the source domain in order to make predictions on the target domain. However, these techniques crucially rely on the embeddings for the source and target datasets to take a certain form. First, the embeddings cannot be distributionally identical. If this were the case, simply training on all the data or even just the source data would yield a model that generalizes to the target data. Second, the embeddings must have nontrivial overlap. The model can, via gradient descent updates to the embedding weights, push the two distributions closer together. But if they are too far apart to begin with, one cannot hope for this to be successful.


To this end, we present some experimental results to show that our data does indeed take the required form. We take a network trained on all data to completion and study the embedding distributions of the source and target datasets. To reiterate what we said in Section 1, the source dataset here refers to the entire search log data, and the target dataset is a specific small enterprise domain. First, we compute the means of the two distributions and compare their norms [35] to the norm of their difference. From Table 1, we can see that the two mean vectors and the difference vector all have a norm that is of the same order. This suggests that the means of the two distributions are indeed quite different, providing evidence that indeed, the source and target domain embeddings are not distributionally identical.

**Table 1: Table of norms of the source and target dataset embedding means as well as their difference.**

Source Mean Norm:	1.0578
Target Mean Norm:	1.3558
Norm of Mean Difference:	0.8558

Additionally, while we cannot visualize the distributions in multidimensional space, we can plot their projections onto a one-dimensional space. To choose the vector on which we project the distributions, we use an idea from robust statistics [11, 28]. We form a matrix consisting of the embedding vector of each example from both the source and target datasets. From this matrix, we compute the vector corresponding to the largest singular value. The intuition is that this vector corresponds to the strongest signal coming from the combined embedding vectors. If the source and target embeddings indeed form two distinct populations, this signal will correlate more strongly with one distribution than the other. From Figure 2, we can see that this is the case. The green values represent the correlations of the target set embeddings onto the top singular

vector, and the blue values are from the source data embeddings. While there is some overlap between the two distributions, they are quite clearly distinct.



**Figure 2: Some statistics of the two embedding distributions from source and target datasets. There are 10,000 examples living in a 508-dimensional space. We show the distributions projected onto the direction of the largest eigenvector of the covariance matrix of embedding vectors. The source dataset is in blue, and the target dataset is in green.**

As we have now established that the embedding distributions are overlapping, but not identical, in the next section we discuss possible domain adaptation techniques to be used.

## 4 METHODOLOGY

In this section, we first formulate our problem and provide definitions of notation we will use. Then, we describe two different solutions to the problem, first using discriminator-based techniques and then statistics-based techniques.

### 4.1 Problem Formulation

The inputs for the enterprise email search problem are queries  $\psi = (q, D, c)$ , where  $q$  represents a user query string,  $D$  represents the query’s list of documents, and  $c$  represents the clickthrough data. For each document  $d \in D$ , we have a feature vector  $x_d$  as its representation, and a boolean  $c_d$  denoting whether the document was clicked. The set of queries from all data is labeled  $\mathcal{S}$ , the source dataset, while the set of queries from a specific domain is labeled  $\mathcal{T}$ , the target dataset, and we are trying to use  $\mathcal{S}$  to help train a model on  $\mathcal{T}$ .

Given the set of queries, our goal is to learn a model  $M$  minimizing a loss defined as:

$$\mathcal{L}(M) = E_{\psi \in \mathcal{T}}(\ell(M, \psi)), \quad (1)$$

where  $\ell(M, \psi)$  denotes the loss of the model  $M$  on query  $\psi$ , and we are taking expectation over queries from the target data. Before we define our loss function, we note that typical neural network models work by approximating the above loss with the training set. We establish notation here by letting the training sets for our deep network be  $S = \{\psi_i^S\} \in \mathcal{S}$  and  $T = \{\psi_j^T\} \in \mathcal{T}$ . Our goal is to learn a model  $M$  that minimizes:

$$\mathcal{L}(M) = \frac{1}{|T|} \sum_{\psi^T \in T} \ell(M, \psi^T).$$

However, in the domain adaptation setting, we assume a scarcity of data in the target distribution. As such, training  $M$  to minimize

$\mathcal{L}$  would result in either an overfitted model or one that cannot generalize to all of  $\mathcal{T}$ . Thus, we instead train  $M$  using training data from both  $\mathcal{S}$  and  $\mathcal{T}$ .

In this paper, our model  $M$  depends on deep neural networks (DNNs) [29]. We choose to use DNNs for a few reasons. First, the number of features from our queries and documents is quite large, and moreover, some features are sparse. While tree-based models [14] can model feature interactions, they are not scalable to a large number of features and cannot handle sparse features, such as those coming from document or query text. On the other hand, DNNs are a natural candidate model for dealing with sparse inputs. Also, DNNs were shown to be quite successful for ranking applications, especially when supplied with large amounts of supervised training data [10].

Over the next few sections, we provide an overview of our model  $M$ . First, we map the query and document features together into a high-dimensional embedding space (specifically, 508-dimensional). Then, a prediction model, which we will call  $P$ , consisting of a DNN is trained on this embedding space. Since the source and target datasets are different, we also expect their embeddings to be different within the embedding space. Thus, we use an additional correction model to make the embeddings indistinguishable to the prediction model, so that source data can be used to predict clicks on the target data.

Figure 3 provides an illustration of our problem. In the next section, we describe how to map enterprise email search inputs into a high-dimensional embedding space.

### 4.2 Input Embeddings

The source and target datasets are mapped to an embedding space via an embedding function (part (1) in Figure 3).

Due to the private nature of our data, our query-document inputs are modeled as bags of frequent n-grams and can be represented in three parts:

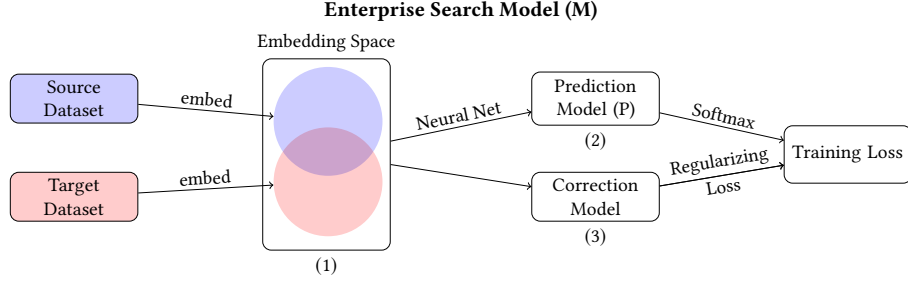
- the features for the query  $q$
- the sparse features for a document  $d$
- the dense features for  $d$  (e.g., document static scores).

To preserve privacy, the inputs are k-anonymized and only query and document n-grams that are frequent in the entire corpus are retained. For more details about the specific features and the anonymization process used in a similar setting see [39, 46]. The query features and sparse document features are passed through a single embedding layer, whereas the dense features are left as they are. We denote the various parts of the input as  $x_q^{\text{embed}}$ ,  $x_d^{\text{embed}}$ , and  $x_d^{\text{dense}}$ , concatenate them together, and then pass them through one additional embedding layer to form the input embedding,  $x_d$ . We denote the function taking an input  $\psi$  and outputting its corresponding embedding  $x_d$  as  $\text{emb}(\psi)$ .

### 4.3 Feed-Forward DNN Model

The discussion of our feed-forward DNN model corresponds to part (2) of Figure 3, the prediction model,  $P$ .

Feed-forward DNNs have become increasingly prominent in research for learning-to-rank problems [13, 24, 46]. Such models have proven to be effective in modeling sparse features naturally arising from inputs such as document text through the use of embeddings.



**Figure 3: Illustration of the Domain Adaptation Problem.** Both the source and target datasets are mapped into a high-dimensional embedding space. Given the embeddings, a prediction model, i.e. a deep neural net, is trained to predict clicks using a softmax cross entropy loss. For domain adaptation, a correction model is used on the embedding space to compute a regularizing loss term that is added to the training loss.

Since they are the basis for our models, we review them in this section.

A deep neural network can be broken down into a series of layers, each applying a matrix multiplication, bias addition, and then activation function to its inputs. We refer to the output of layer  $i$  as  $h_i$ . Letting  $h_0$  be  $x_d$ , the subsequent layers are explicitly obtained as:

$$h_i = \sigma(w_i h_{i-1} + b_i),$$

where  $w_i$  denotes the weight matrix and  $b_i$  the bias vectors at layer  $i$ . We refer to the trainable parameters  $w$  and  $b$  together as the prediction model's  $\theta_p$ . The function  $\sigma$  is what is known as the activation function. We use a hyperbolic function:

$$\sigma(t) = \frac{e^{2t} - 1}{e^{2t} + 1}.$$

If  $h_o$  is the last layer, then the prediction model output is simply:

$$P(\text{emb}(\psi)) = P(x_d) = w_o h_o + b_o.$$

The loss function we aim to minimize utilizes a softmax cross entropy term. Formally, for an input  $\psi$  with  $N$  documents  $d_i$  and click data  $c_i$ , the loss can be written as:

$$\ell_P(\psi, \text{emb}) = - \sum_{i=1}^N c_i \log(p_i), \quad (2)$$

where

$$p_i = \frac{1}{1 + e^{-P(x_{d_i})}}.$$

The overall prediction model loss term is then the average of  $\ell_P$  over all inputs  $\psi$ :

$$\mathcal{L}_P(\{\psi\}, \text{emb}) = \frac{1}{|\{\psi\}|} \sum_{\psi} \ell_P(\psi, \text{emb}) \quad (3)$$

#### 4.4 Domain Adaptation Methods

In this section, we describe our techniques for encouraging the neural network models to make the embeddings of the source and target datasets indistinguishable. This corresponds to part (3) of Figure 3, the correction model. Our first class of methods in Section 4.4.1 is motivated by Generative Adversarial Networks [18], relying on the deep neural networks called discriminators that arise

in Generative Adversarial Networks. Then, in Section 4.4.2, we propose a second class of methods focusing on utilizing various statistics of the embedding distributions for domain adaptation.

**4.4.1 Discriminator-Based Techniques.** Discriminator-based domain adaptation techniques are closely related to adversarial learning methods where a two-player game is created between a discriminator and an adversary. In this setting, each training example is labeled as being from either the source or the target dataset. The discriminator is implemented as a neural network that works to classify an example with its corresponding dataset. At the same time, an adversary updates the embedding weights in such a way as to fool the discriminator.

The goal of discriminator-based techniques is to reach an equilibrium in which the adversary has updated the embedding weights to fool any discriminator. In this case, the two embeddings will be indistinguishable, and a prediction model trained on the source dataset will generalize well to the target dataset.

We then define the total loss function for discriminator-based techniques on the model  $M$  as follows:


$$\mathcal{L} = \mathcal{L}_P + \lambda_D \mathcal{L}_D + \lambda_{adv} \mathcal{L}_{adv}. \quad (4)$$

Here,  $\mathcal{L}_P$  refers to Equation 3, while  $\mathcal{L}_D$  and  $\mathcal{L}_{adv}$  refer to the discriminator and adversarial losses which we will discuss next. The  $\lambda_D$  and  $\lambda_{adv}$  terms are multiplicative factors that control the effects of the discriminator and adversarial losses relative to  $\mathcal{L}_P$ .

The discriminator itself is an additional feed-forward deep neural network, separate from the prediction model, taking the embedding of a query as input, which we will denote as a function  $\mathcal{D}(\text{emb}(\psi))$ . Similar to the parameters for the prediction model, we will denote the trainable parameters for this DNN with  $\theta_D$ . The discriminator loss (i.e.,  $\mathcal{L}_D$ ) is a standard cross-entropy loss, defined as:

$$\begin{aligned} \mathcal{L}_D(\{\psi_i^S\}, \{\psi_j^T\}, \text{emb}) = & - \frac{1}{|S|} \sum_{i=1}^{|S|} \log(D(\text{emb}(\psi_i^S))) \\ & - \frac{1}{|T|} \sum_{j=1}^{|T|} \log(1 - D(\text{emb}(\psi_j^T))). \end{aligned} \quad (5)$$

For adversarial loss (i.e.,  $\mathcal{L}_{adv}$ ), there are different choices that can be made. One standard choice is known as the gradient reversal



**Figure 4: Illustration of the Gradient Reversal Algorithm.** We color code the overall model into three separate parts—green for the embedding, red for the prediction, and blue for the discriminator. An input is mapped by right-facing arrows to the prediction model and discriminator loss terms,  $\mathcal{L}_P$  and  $\mathcal{L}_D$  respectively. Then, during the backpropagation of gradients to train the neural network, each part of the network receives gradients as listed in the diagram. The discriminator part of the network is trained with gradients from  $\mathcal{L}_D$ , and the prediction model is trained with gradients from  $\mathcal{L}_P$ . The embedding part is trained with gradients from both.

loss [16]. The idea with gradient reversal is to directly maximize the discriminator loss, i.e.,  $\mathcal{L}_D$ . The gradient of the discriminator loss is, by definition, the direction of the largest change. While the discriminator will take a gradient step to decrease the loss, the adversary takes a backwards step along this direction. Formally, the adversarial loss is defined as follows:

$$\mathcal{L}_{adv}(\{\psi_i^S\}, \{\psi_j^T\}, \text{emb}) = -\mathcal{L}_D(\{\psi_i^S\}, \{\psi_j^T\}, \text{emb}). \quad (6)$$

For completeness, we mention that there are two other often-used adversarial losses. One uses a cross-entropy term, but with inverted labels [18], labeling each source example as coming from the target dataset, and vice versa. The other computes the cross-entropy of the combined source and target datasets against a uniform distribution drawn from both [41]. In our experiments, we found that these different losses all yielded similar performance and so focus on gradient reversal, which was also used successfully in a related work in domain adaptation for information retrieval [9].

Since we will focus on gradient reversal, we provide an illustration of the technique in Figure 4 and will henceforth refer to our proposed training method from the class of discriminator-based techniques as the **gradient reversal method**. The embedding parameters are  $\theta_{\text{emb}}$  (green), the prediction model parameters are  $\theta_P$  (red), and the discriminator parameters are  $\theta_D$  (blue). The gradient updates in each time step are given in the figure as a partial derivative of the loss functions with respect to the parameters.

**4.4.2 Statistics-Based Techniques.** The embeddings from the source and target data points are subsets in our embedding space, which we can think of as distributions. Consequently, we can extract

various statistics from the distributions and encourage the model to match the statistics coming from the two datasets. We focus on using the mean of the distributions. While two distributions can have the same or similar means but still be very different, we found empirically that this statistic actually worked quite well in making the two distributions indistinguishable to the prediction model. Specifically, we add a term to the model loss function consisting of the difference in the means of the source and target embedding distributions. Since neural networks work by minimizing their loss functions, this allows the network to take steps to minimize the difference in the means, drawing the two distributions closer together. We provide a pictorial representation of this technique in Figure 5 which we refer to as **maximum mean discrepancy (MMD)**. The two distributions are given in red and blue, and their means are represented by a bold point. Applying an MMD minimization does not change the shape of either distribution, but brings their means closer together.

The total loss for the model is then defined as:

$$\mathcal{L} = \mathcal{L}_P + \lambda_{MMD} \mathcal{L}_{MMD}. \quad (7)$$

Where  $\mathcal{L}_P$  is the prediction model loss (i.e., Equation 3) and  $\mathcal{L}_{MMD}$  is the maximum mean discrepancy loss. The  $\lambda_{MMD}$  factor also controls the effect of the MMD on the overall loss relative to the prediction loss.

Formally, the MMD loss, i.e.,  $\mathcal{L}_{MMD}$  is given by:

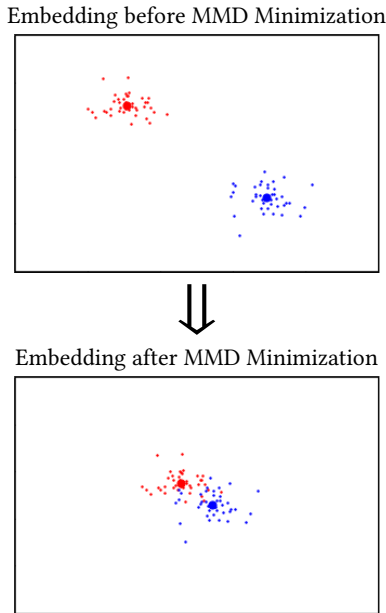
$$\mathcal{L}_{MMD}(\{\psi_i^S\}, \{\psi_j^T\}, \text{emb}) = \left\| \frac{1}{|S|} \sum_{i=1}^{|S|} \text{emb}(\psi_i^S) - \frac{1}{|T|} \sum_{j=1}^{|T|} \text{emb}(\psi_j^T) \right\|_2, \quad (8)$$

Since we are attempting to make two distributions indistinguishable in latent space, it is natural to also include other statistics. Most notably, we could also add a variance term, and aim to minimize not only the mean, but also the variance. We note that since our representations exist in a high dimensional space, we are comparing two covariance matrices in this case. Nonetheless, a discrepancy term can be added to encourage these to be similar. In our experiments, though, we found that adding such a variance term did not alter the results in any significant way.

As a remark, we note that even if two distributions share the same mean and covariance matrix, they can still be quite different. However, if the distributions are close to Gaussian, then they would be characterized by their first two moments. Of course, we cannot prove that the neural network will find representations that map the inputs to a Gaussian subspace of embedding space. However, we did find experimentally that projecting the distributions of representations to random directions yielded one-dimensional distributions that looked Gaussian. Thus, it is reasonable to believe that with a large enough training set, the representations would converge to something close to Gaussian.

## 5 EXPERIMENTS

In this section, we begin with a description of the datasets we use as input in our experiments. Then, we evaluate each of our proposed techniques and baselines using a typical evaluation metric, which we will define. Finally, we discuss the sensitivity of each technique



**Figure 5: Illustration of Maximum Mean Discrepancy Technique.** The two distributions, labeled blue and red, are far apart in embedding space. The Maximum Mean Discrepancy technique adds a regularizing loss that encourages the model to minimize the distance between the means. Note that the distributions themselves do not change in shape. They are only brought closer together in embedding space.

to changes in hyperparameters, as robustness to hyperparameter tuning hastens the training process.

## 5.1 Datasets

Due to the private and sensitive nature of the data, there is no publicly available large-scale enterprise email search dataset. As a result, the data we use comes from the search click logs of Gmail, a commercial email service. For our source dataset, we use click data from the entire Gmail search log. We then study domain adaptation to data arising from logs of four individual enterprise domains. The search logs for the source data consist of hundreds of millions of queries, whereas the target domains are significantly smaller. We chose the four largest domains (selected based on the total number of examples) as our target datasets and label them as A, B, C, and D. Domain A has around 65,000 total examples, B and C have around 50,000 each, and D has 30,000 queries, all significantly smaller than the source dataset. The domains, including the source domain, are split into training and evaluation sets at approximately a 5 : 1 ratio, and in such a way that all queries in the evaluation sets are performed on days after those in the training sets. Each example consists of a query with six candidate documents, one of which is clicked.

The goal of the model is to rank the six documents in such a way as to increase the likelihood of a higher ranked document being

clicked. In this way, clicks are regarded as the ground truth from which our model learns.

## 5.2 Model Evaluation

Our neural network models are implemented in TF-Ranking [36], a scalable open-source learning-to-rank Tensorflow [1] library. Our baselines are optimized over a number of hyperparameters, including learning rate, number of hidden layers, dimensions of layers, and batch normalization. Specifically, we use a learning rate of 0.1 and three hidden layers with dimensions 256, 128, and 64. For each training procedure that takes as input both source and target training data, we also tried a number of different ratios of source to target training data. However, we found that none of the procedures were sensitive to this ratio, as long as it was larger than 4 : 1 or 5 : 1. Anything lower would cause overfitting, since the target dataset was so much smaller than the source. Then, for the mean discrepancy (Eq. 7) and gradient reversal (Eq. 4) losses, the  $\lambda$  multipliers are optimized over [0.3, 0.7, 1.0, 3.0, 7.0].

Model performance is evaluated using weighted mean reciprocal rank (WMRR), as proposed in [44]. The weighted MRR (WMRR) is calculated using the one clicked document of any query as:

$$\text{WMRR} = \frac{1}{\sum_1^{|E|} w_i} \sum_{i=1}^{|E|} w_i \frac{1}{\text{rank}_i}, \quad (9)$$

where  $E$  denotes the evaluation set,  $\text{rank}_i$  denotes the position or rank of the clicked document for the  $i$ -th query, and  $w_i$  denotes the bias correction weights. The  $w_i$  are inversely proportional to the probability of observing a click at the clicked position and are set using result randomization, as described in [44]. In addition to reporting the WMRR attained by each model, we also conduct statistical significance tests using the two-tailed paired t-test with 99% confidence.

## 5.3 Results and Discussion

In this section, we provide our main results, as well as accompanying discussion. We first describe the two standard baseline training methods.

### Standard Baselines:

- **Train on all.** This is the simplest baseline. We train a model on the entire source dataset, and then evaluate it on the specific domain test data.
- **Train on domain.** We form a training set consisting of only domain-specific data. As expected, there is not enough data to train a neural network capable of generalizing to test data. Not only does the model have a low WMRR, but we also see severe overfitting.

Since our goal is a thorough analysis of possible ways to train a prediction model for a specific domain, we also tried additional baselines. Both the additional baselines typically outperformed the standard baselines, and we suggest that they should be used for comparison in any domain adaptation study.

### Additional Baselines:

- **Re-train.** This is typically known as vanilla transfer learning. We first load up a **train on all** model. Then, we re-train

this model on only domain data with a lower learning rate. Specifically, we reduce the learning rate by a factor of 10.

- **Batch-balance.** This model is similar to the **train on all**, in that the training set consists of the entire source dataset. The difference is that we enforce a certain proportion of target domain data in each training batch.

The raw WMRR model evaluations from Equation 9 for the baselines are provided in Table 2.

**Table 2: WMRR evaluation results for all four baseline methods trained for each of the four separate domains, labeled A, B, C, and D.**

	A	B	C	D
Train on all	0.659	0.692	0.611	0.598
Train on domain	0.639	0.694	0.573	0.579
Re-train	0.675	0.713	<b>0.608</b>	0.618
Batch-balance	<b>0.682</b>	<b>0.715</b>	<b>0.608</b>	<b>0.621</b>

From Table 2, we can see that generally, **train on domain** performs the worst of all baselines. As noted before, this makes sense due to the fact that the domains do not provide enough training data for neural network models. Then, we have our standard training method **train on all**. Finally, **re-train** and **batch-balance** have roughly the same performance, with the latter performing slightly better for some domains. These are exactly the results we would expect, since these two methods are more involved than only training on domain-specific data or all the source data we have.

Now, we briefly describe our proposed training methods. While they are described in full detail in Section 4.4, we review them here for convenience.

#### Domain Adaptation Methods:

- **Gradient Reversal.** As in the **Batch-balance** baseline, we enforce a certain proportion of domain data in each training batch. A discriminator is then added to try and distinguish, from the embeddings, whether data is from the specific domain. Gradient reversal is used on the embedding weights to fool the discriminator. The resulting loss is described in Equation 4.
- **Mean Discrepancy.** Again, we enforce a certain proportion of domain data in each training batch. A regularization term is added to the standard cross entropy loss consisting of the difference in means of the source and target input embeddings. The resulting loss is described in Equation 7.

Our main table of results is provided in Table 3. Our numbers are recorded as relative improvement of our proposed domain adaptation methods to the corresponding baseline methods.

While the improvement changes from domain to domain, mean discrepancy is consistently the best performing proposed method. In every instance, mean discrepancy achieves a higher WMRR than any of the baselines. To give perspective, we note that improvements of 1% are considered to be highly significant for our

enterprise email search system. Compared to **Re-train**, mean discrepancy achieves at least 1% achievement, although for Domains A and B, it does not quite outperform batch-balance by this much.

While not directly listed as a comparison in the table, we also note that mean discrepancy outperforms gradient reversal in a statistically significant way. In image classification, adversarial methods have shown to be better than those using maximum mean discrepancy [15, 17, 41] for domain adaptation to an unsupervised domain. But in our experiments, the opposite was shown to be true. While we cannot say for certain why this difference exists, we provide a few possible explanations. We first recall that the set of training examples aims to approximate the true distribution of inputs. Since maximum mean discrepancy aims to minimize the difference between source and target distribution means, the method is highly dependent on how well the mean is approximated by training examples. Since the typical size of datasets considered in prior work is quite small, often even smaller than a single target domain, the mean of the training example embeddings may not accurately represent the true mean. With a thousand times more inputs in our source dataset, the mean is more accurately approximated and more useful for maximum mean discrepancy.

But even with good approximation of the true mean, it is fairly common to find situations in which reducing the MMD would not suffice to make two distributions indistinguishable. One can easily imagine two distributions with identical mean that are still drastically different. But as we showed in Table 1, the source and target data embeddings in our problem do have significantly different means. Because of this, using MMD with the way we embed the query-document inputs ends up working very well for our domain adaptation task.

## 5.4 Sensitivity Analysis

In this section, we discuss the sensitivity of the mean discrepancy and gradient reversal methods. In a regime with complicated inputs and long training times, a desired quality of any training method is robustness to hyperparameters. Therefore, we compare the sensitivity of our two proposed training methods, maximum mean discrepancy and gradient reversal, to the changes in their respective parameters. Recall that the set of parameters  $\{\lambda_D, \lambda_{adv}, \lambda_{MMD}\}$  dictates the interactions of the distribution-matching correction terms – part (3) of Figure 3 – with the ranking model’s softmax cross-entropy term  $\mathcal{L}_P$  (see Equations 4 and 7).

Careful parameter tuning is especially important when using methods based on two-player games, such as gradient reversal. In general, two-player games are sensitive to changes in parameters, making a good equilibrium difficult for gradient descent to find [37]. As a result, we hypothesize that maximum mean discrepancy is the better training method for our enterprise email search setting, not only because of better overall WMRR, but also due to its robustness to changes in  $\lambda_{MMD}$ . In what follows, we provide some empirical evidence to back up this hypothesis.

First, in Figure 6, we plot the resulting WMRR from training using a range of  $\lambda$  values. For the mean discrepancy method, this corresponds to  $\lambda_{MMD}$ , and for gradient reversal, we fix two values of  $\lambda_D$  and vary  $\lambda_{adv}$  on the x-axis. The WMRR curves for fixed values of  $\lambda_D$  are shown in order to give a direct comparison between



Table 3: WMRR evaluation results for adapting to four domains. Methods on the x-axis are compared to baselines on the y-axis, and results are given as relative improvement. An asterisk (\*) denotes statistical significance relative to the baseline, according to our two-tailed paired t-test with 99% confidence. Additionally, we underline the numbers where mean discrepancy is statistically significant over gradient reversal.

Domain A	Gradient Reversal	Mean Discrepancy	Domain B	Gradient Reversal	Mean Discrepancy
Train on all	+4.16%*	<u>+4.38%*</u>	Train on all	+2.99%*	<u>+4.00%*</u>
Train on domain	+7.30%*	<u>+7.52%*</u>	Train on domain	+2.77%*	<u>+3.79%*</u>
Re-train	+1.60%*	<u>+1.82%*</u>	Re-train	+0.02%	<u>+1.00%*</u>
Batch-balance	+0.44%*	<u>+0.65%*</u>	Batch-balance	-0.24%*	<u>+0.75%*</u>
Domain C	Gradient Reversal	Mean Discrepancy	Domain D	Gradient Reversal	Mean Discrepancy
Train on all	+1.37%*	<u>+1.49%*</u>	Train on all	+2.45%*	<u>+5.30%*</u>
Train on domain	+8.04%*	<u>+8.17%*</u>	Train on domain	+5.77%*	<u>+8.71%*</u>
Re-train	+1.85%*	<u>+1.97%*</u>	Re-train	-0.92%*	<u>+1.83%*</u>
Batch-balance	+1.78%*	<u>+1.91%*</u>	Batch-balance	-1.44%*	<u>+1.30%*</u>

the mean discrepancy and gradient reversal methods. Additionally, we plot the entire three-dimensional surface of the WMRR metric as a function of both  $\lambda_D$  and  $\lambda_{adv}$  in Figure 7.




Figure 6: WMRR curves on domain A for both the mean discrepancy and the gradient reversal methods. The WMRR is plotted as a function of  $\lambda_{MMD}$  for mean discrepancy and of  $\lambda_{adv}$  for two fixed values of  $\lambda_D$  for gradient reversal.

From Figure 6, we can see that for a large range of  $\lambda_{MMD}$  values, the resulting WMRR is relatively stable. Specifically, in the range of  $\lambda$  values we consider, from 0.3 to 7.0, the WMRR values remain within a  $\pm 0.005$  range. Qualitatively, this robustness to tuning  $\lambda_{MMD}$  holds for all domains.

On the other hand, the plots for gradient reversal support our conjecture that it is not as robust as mean discrepancy. From Figure 6, we see a direct comparison showing the greater sensitivity gradient reversal has relative to mean discrepancy. Moreover, from Figure 7, we observe that the WMRR values are as far apart as 0.05 from each other, nearly ten times as big a difference as we see for the mean discrepancy method. Again, similar trends could be observed for other domains.




Figure 7: WMRR values on domain A when varying both  $\lambda_{adv}$  and  $\lambda_D$  parameters.

## 6 CONCLUSIONS

In this paper, we studied the application of domain adaptation to learning-to-rank, specifically in the setting of enterprise email search. We developed a formal framework for integrating two techniques from the image classification literature for transfer learning, maximum mean discrepancy (MMD) and gradient reversal, in the learning-to-rank models. Our models were implemented in Tensorflow using a deep neural network that efficiently handles various features from query-document inputs, and provides embeddings to which the domain adaptation methods are applied.

The results from our experiments on a large-scale enterprise email search engine indicate that neither a single global model, nor simple transfer learning baselines are sufficient for achieving the optimal performance for individual domains. Overall, we show that maximum mean discrepancy (MMD) is the best technique for adapting the global learning-to-rank model to a small target domain

in enterprise email search. MMD is not only the most effective method on all the tested domains, it also displays robustness to parameter changes.

One possible future direction of study is using regularization terms involving statistics other than the mean. Another is to find more stable equilibrium for discriminator-based methods. While much work has been done [3, 21, 37, 38] to improve the stability of models involving two-player games, these are often very task-specific. Finally, while in this paper we explored domain adaptation specifically in the enterprise email search setting, the proposed methods can easily generalize to other information retrieval tasks.

## REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: a system for large-scale machine learning. In *Operating Systems Design and Implementation*, Vol. 16. 265–283.
- [2] Qingyao Ai, Susan T Dumais, Nick Craswell, and Dan Liebling. 2017. Characterizing email search using large-scale behavioral logs and surveys. In *Proceedings of the 26th International Conference on World Wide Web*. 1511–1520.
- [3] Martin Arjovsky and Léon Bottou. 2017. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862* (2017).
- [4] Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016. A neural click model for web search. In *Proceedings of the 25th International Conference on World Wide Web*. 531–541.
- [5] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning*. 89–96.
- [6] Christopher JC Burges. 2010. From Ranknet to Lambdarank to Lambdamart: An overview. *Learning* 11, 23-581 (2010), 81.
- [7] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th International Conference on Machine Learning*. 129–136.
- [8] David Carmel, Guy Halawi, Liane Lewin-Eytan, Yoelle Maarek, and Ariel Raviv. 2015. Rank by time or by relevance?: Revisiting email search. In *Proc. of the 24th ACM International Conference on Information and Knowledge Management*.
- [9] Daniel Cohen, Bhaskar Mitra, Katja Hofmann, and W Bruce Croft. 2018. Cross domain regularization for neural ranking models using adversarial learning. *arXiv preprint arXiv:1805.03403* (2018).
- [10] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural ranking models with weak supervision. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 65–74.
- [11] Ilias Diakonikolas, Gautam Kamath, Daniel M Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. 2016. Robust estimators in high dimensions without the computational intractability. In *IEEE 57th Annual Symposium on Foundations of Computer Science*. 655–664.
- [12] Susan Dumais, Edward Cutrell, Jonathan J Cadiz, Gavin Jancke, Raman Sarin, and Daniel C Robbins. 2016. Stuff I’ve seen: a system for personal information retrieval and re-use. In *ACM SIGIR Forum*, Vol. 49. 28–35.
- [13] Bora Edizel, Amin Mantrach, and Xiao Bai. 2017. Deep Character-Level Click-Through Rate Prediction for Sponsored Search. *arXiv preprint arXiv:1707.02158* (2017).
- [14] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of Statistics* (2001), 1189–1232.
- [15] Yaroslav Ganin and Victor Lempitsky. 2014. Unsupervised domain adaptation by backpropagation. *arXiv preprint arXiv:1409.7495* (2014).
- [16] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17, 1 (2016), 2096–2030.
- [17] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, David Balduzzi, and Wen Li. 2016. Deep reconstruction-classification networks for unsupervised domain adaptation. In *European Conference on Computer Vision*. 597–613.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. 2672–2680.
- [19] A. Gretton, AJ. Smola, J. Huang, M. Schmittfull, KM. Borgwardt, and B. Schölkopf. 2009. *Covariate shift and local learning by distribution matching*. MIT Press, Cambridge, MA, USA, 131–160.
- [20] Catherine Grevet, David Choi, Debra Kumar, and Eric Gilbert. 2014. Overload is overloaded: Email in the age of Gmail. In *Proceedings of the Sigchi Conference on Human Factors in Computing Systems*. 793–802.
- [21] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of Wasserstein Gans. In *Advances in Neural Information Processing Systems*. 5767–5777.
- [22] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. 55–64.
- [23] David Hawking. 2010. Enterprise Search. In *Modern Information Retrieval, 2nd Edition*, Ricardo Baeza-Yates and Berthier Ribeiro-Neto (Eds.). Addison-Wesley, 645–686. [http://david-hawking.net/pubs/ModernIR2\\_Hawking\\_chapter.pdf](http://david-hawking.net/pubs/ModernIR2_Hawking_chapter.pdf)
- [24] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM International Conference on Information and Knowledge Management*. 2333–2338.
- [25] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 133–142.
- [26] Jin Young Kim, Nick Craswell, Susan Dumais, Filip Radlinski, and Fang Liu. 2017. Understanding and Modeling Success in Email Search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 265–274.
- [27] Udo Kruschwitz and Charlie Hull. 2017. Searching the Enterprise. *Foundations and Trends® in Information Retrieval* 11, 1 (2017), 1–142. <https://doi.org/10.1561/1500000053>
- [28] Kevin A Lai, Anup B Rao, and Santosh Vempala. 2016. Agnostic estimation of mean and covariance. In *IEEE 57th Annual Symposium on Foundations of Computer Science*. 665–674.
- [29] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep Learning. *Nature* 521, 7553 (2015), 436.
- [30] Ming-Yu Liu and Onel Tuzel. 2016. Coupled generative adversarial networks. In *Advances in neural information processing systems*. 469–477.
- [31] Fuchen Long, Ting Yao, Qi Dai, Xinmei Tian, Jiebo Luo, and Tao Mei. 2018. Deep Domain Adaptation Hashing with Adversarial Learning. In *the 41st International ACM SIGIR Conference on Research; Development in Information Retrieval*.
- [32] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I Jordan. 2015. Learning transferable features with deep adaptation networks. *arXiv preprint arXiv:1502.02791* (2015).
- [33] Sitong Mao, Xiao Shen, and Fu-lai Chung. 2018. Deep Domain Adaptation Based on Multi-layer Joint Kernelized Distance. In *the 41st International ACM SIGIR Conference on Research; Development in Information Retrieval*.
- [34] Bhaskar Mitra and Nick Craswell. 2017. Neural Models for Information Retrieval. *arXiv preprint arXiv:1805.03403* (2017).
- [35] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.
- [36] Rama Kumar Pasumarthi, Sebastian Bruch, Xuanhui Wang, Cheng Li, Michael Bendersky, Marc Najork, Jan Pfeifer, Nadav Golbandi, Rohan Anil, and Stephan Wolf. 2019. TF-Ranking: Scalable TensorFlow Library for Learning-to-Rank. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. (to appear).
- [37] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. 2017. Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems*. 2018–2028.
- [38] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*. 2234–2242.
- [39] Jiaming Shen, Maryam Karimzadehgan, Michael Bendersky, Zhen Qin, and Don Metzler. 2018. Multi-Task Learning for Personal Search Ranking with Query Clustering. In *Proceedings of ACM Conference on Information and Knowledge Management*.
- [40] Baochen Sun and Kate Saenko. 2016. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*. 443–450.
- [41] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. 2015. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*. 4068–4076.
- [42] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition*. 4.
- [43] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. 2014. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474* (2014).
- [44] Xuanhui Wang, Michael Bendersky, Donald Metzler, and Marc Najork. 2016. Learning to rank with selection bias in personal search. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 115–124.
- [45] Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. 2008. Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th International Conference on Machine Learning*. 1192–1199.
- [46] Hamed Zamani, Michael Bendersky, Xuanhui Wang, and Mingyang Zhang. 2017. Situational context for ranking in personal search. In *Proceedings of the 26th International Conference on World Wide Web*. 1531–1540.