
Pretrained deep models outperform GBDTs in Learning-To-Rank under label scarcity

Charlie Hou¹ Kiran K Thekumparampil² Michael Shavlovsky² Giulia Fanti¹ Sujay Sanghavi³

Abstract

On tabular data, a significant body of literature has shown that current deep learning (DL) models perform at best similarly to Gradient Boosted Decision Trees (GBDTs), while significantly underperforming them on outlier data (Gorishniy et al., 2021; Rubachev et al., 2022; McElfresh et al., 2023). However, these works often study idealized problem settings (e.g., fully labeled data). We identify a natural tabular data setting where DL models can outperform GBDTs: tabular Learning-to-Rank (LTR) under label scarcity. Tabular LTR applications, including search and recommendation, often have an abundance of unlabeled data, and *scarce* labeled data. We show that DL rankers can utilize unsupervised pre-training to exploit this unlabeled data. In extensive experiments over both public and proprietary datasets, we show that pretrained DL rankers consistently outperform GBDT rankers on ranking metrics—sometimes by as much as 38%—both overall and on outliers.

1. Introduction

The learning-to-rank (LTR) problem aims to train a model to rank a set of items according to their relevance or user preference (Liu, 2009). An LTR model is typically trained on a dataset of queries and associated *query groups* (i.e., a set of potentially relevant *documents* or *items* per query), as well as an associated (generally incomplete) ground truth ranking of the items in the query group. The model is trained to output the optimal ranking of documents or items in a query group, given a query. LTR is a core ML technology in many real world applications—most notably in search contexts

including Bing web search (Qin & Liu, 2013), Amazon product search (Yang et al., 2022), and Netflix movie recommendations (Lamkhede & Kofler, 2021). More recently, LTR has also found use in Retrieval Augmented Generation (RAG), where a retrieval system is paired with an LLM (large language model) to ground the latter’s responses to truth and reduce hallucination (Lewis et al., 2020; Glass et al., 2022; Pan et al., 2022).

Traditionally, LTR models took *tabular features* (numerical or categorical) of queries and documents as input features (Burges et al., 2005; Cao et al., 2007; Chapelle & Chang, 2011; Qin & Liu, 2013; Lucchese et al., 2016). *Tabular LTR* remains a core problem in machine learning (Bower et al., 2021; Yang et al., 2022; Yu, 2020; Kveton et al., 2022; Pan et al., 2022; Xia et al., 2008; Chen et al., 2009). Today, deep models are largely outperformed by gradient boosted decision trees (GBDTs) (Friedman, 2001) in ML problems with tabular features (Jefferies et al., 2023; Qin et al., 2021), while deep models are state-of-the-art by a significant margin in domains like text (Devlin et al., 2018) and images (He et al., 2016).

Recent breakthroughs in modeling non-tabular data like text and images have been driven by first training a deep neural network to learn from unlabeled data (unsupervised pretraining, or pretraining) (Devlin et al., 2018; Chen et al., 2020), followed by supervised training (finetuning). Models that are pretrained in this way can perform significantly better than models that were only trained on existing labeled data. This remarkable success stems from three factors: (1) limited access to labeled data, (2) large, available sources of unlabeled text and image data, and (3) pretraining methods that can exploit unlabeled data.

A natural question is whether deep models can similarly use unsupervised pretraining to outperform GBDTs on the tabular LTR problem. To this end, many papers have studied unsupervised pretraining techniques—both using deep learning and GBDTs—for tabular data (Yoon et al., 2020; Ucar et al., 2021; Bahri et al., 2021; Verma et al., 2021; Majmundar et al., 2022; Duh & Kirchhoff, 2008; Pseudo-Label, 2013). However, none of these methods have shown convincingly that deep learning (with or without unsupervised pretraining) consistently outperforms GBDTs (semi-

¹Department of Electrical and Computer Engineering, Carnegie Mellon University ²Amazon ³Department of Computer Science, University of Texas at Austin. Correspondence to: Charlie Hou <hou.charlie2@gmail.com>.

Accepted by the Structured Probabilistic Inference & Generative Modeling workshop of ICML 2024, Vienna, Austria. Copyright 2024 by the author(s).

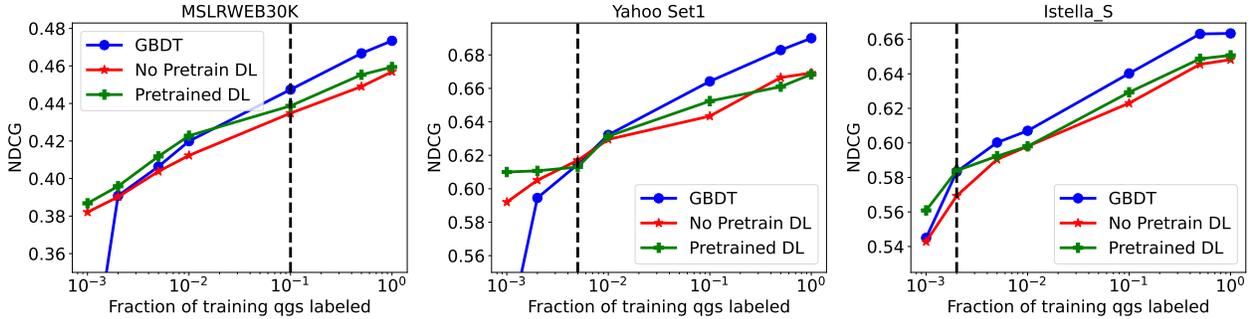


Figure 1: **Main result (Simulated crowdsourcing of labels):** We compare NDCG (\uparrow) as we sweep the percentage of labeled training query groups (QGs); labeled QGs have relevance scores for every element of the QG. For small enough fractions of labeled query groups, pre-trained DL rankers (sometimes significantly) outperform both GBDTs (represents the best of supervised GBDTs and semi-supervised GBDTs) and non-pretrained DL rankers (represents the best of tabular ResNet (Gorishniy et al., 2021), DeepFM (Guo et al., 2017), and DCNv2 (Wang et al., 2021)). To the left of the black dotted line, pretrained rankers perform the best (statistically significant at $p = 0.05$ level). Points are averages over 3 trials.

supervised or supervised); at best, deep methods appear to achieve roughly the same performance as GBDTs (Jeffares et al., 2023; Rubachev et al., 2022; Gorishniy et al., 2021; Schwartz-Ziv & Armon, 2022; Grinsztajn et al., 2022). Furthermore, recent work showed that GBDTs are significantly more robust against outliers in the dataset than deep learning models (McElfresh et al., 2023). Most of these works do not consider the settings of label scarcity and noisy data. We hypothesize that pretrained deep models can beat GBDTs in settings that more accurately reflect real-world challenges.

Through extensive empirical evaluation, we confirm this hypothesis and show that deep models that use unsupervised pretraining consistently outperform GBDTs—sometimes by a large margin—**when labels are scarce**, i.e., only a small number of query groups contain labels. This setting is more standard in most real-world deployments of LTR, rather than the idealized settings studied in prior work. For instance, scarcity can arise when organizations: (a) crowdsource labels for a limited number of query groups due to cost constraints, and/or (b) use rare user feedback (e.g., clicks, purchases) as implicit labels. Our experiments consider both, and are run on the three standard public datasets in the ranking literature: MSLRWEB30K (Qin & Liu, 2013), Yahoo (Chapelle & Chang, 2011), Istella.S (Lucchese et al., 2016), and an industry-scale proprietary ranking dataset.

Contributions: (1) *We demonstrate for the first time that unsupervised pretraining can produce deep models that outperform GBDTs in ranking.* Specifically, pretrained deep rankers can achieve up to 38% higher NDCG (normalized discounted cumulative gain) score (Burgess, 2010) than GBDTs when labels are scarce (Figure 1). For instance, Figure 1 shows on three datasets that as the labels get scarcer, GBDTs start to substantially underperform pretrained deep

rankers (experimental setup in Section 4.1.2).

(2) *We prescribe empirically-justified LTR-specific pretraining strategies*, including a new ranking-specific pretraining loss, SimCLR-Rank, which is a modification of the widely used SimCLR (Chen et al., 2020) loss. In experiments, SimCLR-Rank achieves similar or better performance than most existing pretraining methods at an order-of-magnitude speedup. This allows SimCLR-Rank to scale to datasets such as our proprietary industry-scale ranking dataset.

(3) *We demonstrate that when labels are scarce, pretrained deep rankers can perform significantly better than GBDTs on outlier query groups, boosting NDCG by as much as 80% (Figure 8, Appendix D.2).* This is perhaps surprising, given prior results in the label-rich regime showing that GBDTs are significantly more robust to data with irregularities than deep rankers (McElfresh et al., 2023).

Our main novelty lies in identifying label-scarce LTR as a practical real-world tabular learning problem, and the results, which challenge prior understanding of the modeling choices for tabular learning. We show that pretraining helps deep models consistently outperform GBDTs in the label-scarce setting, even when the labeled data is plentiful in absolute numbers (up to 25% labeled query groups in public datasets, 3 million+ labeled query groups in the industry-scale ranking dataset). Prior to our work, the literature had not identified a widely-applicable setting where (pretrained) deep models consistently outperformed GBDTs (Qin et al., 2021; McElfresh et al., 2023). Due to label scarcity in practical settings, our results suggest a change in how practitioners may approach real-world tabular learning problems in search and recommendation systems.

Related work: We focus on the traditional LTR setting

where the features are all numeric (tabular data). In this setting, gradient boosted decision trees (GBDTs) (Friedman, 2001) are the de-facto models, and deep models have yet to outperform them convincingly (Qin et al., 2021; Joachims, 2006; Ai et al., 2019; Bruch et al., 2019; Ai et al., 2018; Pang et al., 2020; McElfresh et al., 2023). This includes factorization machines (Rendle, 2010), wide & deep architectures (Cheng et al., 2016), and similar architectures designed for recommendation systems (Guo et al., 2017; Wang et al., 2021; Lian et al., 2018; Naumov et al., 2019; Qu et al., 2016). While our paper is architecture agnostic, we perform comparisons of pretrained models against non-pretrained DeepFM and DCNv2 models to show that pre-training is needed to outperform GBDTs consistently. Self-supervised learning (SSL) or unsupervised pretraining has improved performance and robustness to noise (Hendrycks et al., 2019) in settings with significant unlabeled data like text (Devlin et al., 2018; Howard & Ruder, 2018; Kitaev et al., 2018; Liu et al., 2017; Song & Roth, 2015) and images (Chen et al., 2020; Chen & He, 2021; Grill et al., 2020; Ge et al., 2021; Xie et al., 2020).

Inspired by the success of SSL in images and text, many unsupervised pretraining tasks have been proposed for tabular data. We describe some of the most well-known methods. SubTab (Ucar et al., 2021) proposes to train an autoencoder to reconstruct a table using a subset of the columns. VIME-self (Yoon et al., 2020) is a similar method. SCARF and DACL+ (Bahri et al., 2021; Verma et al., 2021) propose augmentations for the SimCLR pretraining task (Chen et al., 2020) that work for the tabular setting. RegCLR (Wang et al., 2022a) is a method for extracting tables from images, which is not directly applicable to our setting. We comprehensively evaluate/compare SubTab, SCARF, DACL+, VIME-self, SimCLR-Rank, SimSiam (Chen & He, 2021), and self-supervised approaches for GBDTs in Section 3.1 and find that SimCLR-Rank along with SimSiam (Chen & He, 2021) are the best for the tabular LTR setting. We discuss further related works in Appendix A.

2. Learning-To-Rank problem and its metrics

In the LTR problem, samples are query groups (QGs) consisting of a query (e.g. shopping query) and L potentially relevant items (e.g. products) for this query. The i -th relevant item here is represented by a feature vector $\mathbf{x}_i \in \mathbb{R}^d$, which captures information about this query-item pair, subsuming any query related features. In labeled data, such as training or testing, we also have an associated scalar relevance label y_i , which could be a binary, ordinal, or real-valued measurement of the relevance of this item for this query (Qin et al., 2021). The objective is to learn a function that ranks these L items such that the items with highest true relevance are ranked at the top.

Most LTR algorithms formulate the problem as learning a scoring function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$ that maps the feature vector associated with each item to a score, and then ranking the items by sorting the scores in descending order. To measure the quality of a ranking induced by our scoring function f_θ , a commonly-used (per-QG) metric is NDCG (\uparrow):

$$\text{NDCG}(\pi_{f_\theta}, \{y_i\}_{i=1}^L) = \frac{\text{DCG}(\pi_{f_\theta}, \{y_i\}_{i=1}^L)}{\text{DCG}(\pi^*, \{y_i\}_{i=1}^L)} \in [0, 1], \quad (1)$$

where $\pi_{f_\theta} : [L] \rightarrow [L]$ ($[L] \triangleq \{1, \dots, L\}$) is a ranking of the L elements induced by the scoring function f_θ on $\{\mathbf{x}_i\}_{i=1}^L$ (i.e. $\pi(i)$ is the rank of the i -th item), while π^* is the ideal ranking induced by the true relevance labels $\{y_i\}_{i=1}^L$, and discounted cumulative gain (DCG) is defined as $\text{DCG}(\pi, \{y_i\}_{i=1}^L) = \sum_{i=1}^L \frac{2^{y_i-1}}{\log_2(1+\pi(i))}$. Typically, a truncated version of NDCG is used that only considers the top- k ranked items, denoted as $\text{NDCG}@k$. In the rest of our paper, we will refer to $\text{NDCG}@5$ as NDCG; this will be our evaluation metric.

We consider a natural label-scarce LTR setting with large amounts of unlabeled data and a limited supply of labels for a subset of this data. The labels that do exist can arise from multiple sources, such as crowdsourced workers paid to label QGs, or through implicit user feedback (e.g., clicks or purchases); we model both in our evaluation. Unless specified otherwise, we use all the query groups for unsupervised pretraining, but we use only the labeled query groups for supervised training.

2.1. Outlier-NDCG for outlier performance evaluation

Motivated by a recent observation that GBDTs outperform DL models in datasets with irregularities (McElfresh et al., 2023), we also study their performance on outliers. In interactive ML systems like search, performing well on outlier queries is particularly valuable as it empowers users to search for more outlier queries, which in turn allows the modeler to collect more data and improve the model. To this end, we evaluate the outlier performance by computing the average NDCG on outlier queries groups and refer to it as *Outlier-NDCG* for brevity. In practice, some outlier queries may already be known, and the modeler can define the outlier datasets accordingly, like in our proprietary dataset. For example, since industry data pipelines often have missing data/features one could identify samples with missing features as outliers. When outliers are unknown, detecting them (especially in higher-dimensional data) is challenging. Hence, we create a heuristic algorithm by assuming outliers are rare values that are separated from the bulk of the data. Using it we systematically selected outlier query groups for the public datasets that we use. More details about the outliers are provided in Appendix B.2. Separately

measuring the performance on the outliers is justified by our later observation that performance on the full dataset and the outliers can significantly differ (Tables 2 and 3).

3. Design: Unsupervised pretraining for LTR

In this section, we systematically evaluate different strategies for utilizing unlabeled data in pretraining of tabular LTR models, **including a novel LTR-specific pretraining loss, SimCLR-Rank**. We ultimately find that the top pretraining strategies are SimCLR-Rank or SimSiam, but the best choice among these depends on the dataset.

Semi-supervised learning for GBDTs: While to the best of our knowledge there are no methods to perform unsupervised pretraining in GBDTs, one can use semi-supervised methods like consistency regularization (Jeong & Shin, 2020; Sajjadi et al., 2016; Miyato et al., 2018; Oliver et al., 2018; Berthelot et al., 2019), pseudo-labeling/self-training (Rizve et al., 2021; Pseudo-Label, 2013; Shi et al., 2018; Yarowsky, 1995; McClosky et al., 2006), and PCA (Duh & Kirchoff, 2008) to address limited size of labeled data. In consistency regularization, one increases the size of the training set by using data augmentations. In pseudo-labeling/self-training, a model trained on the available labeled data is used to label unlabeled data. Then, the final model is trained on the resulting “labeled” dataset. In PCA, one projects the features along directions of maximum variability in the union of labeled and unlabeled datasets. Since prior work has found that consistency regularization decreases the performance of GBDTs (Rubachev et al., 2022), we evaluate only pseudo-labeling and PCA.

Pretraining methods for tabular data: There have been many unsupervised pretraining methods proposed for tabular data, which can be applied to the tabular LTR setting. In this paper, we will evaluate (1) SCARF (Bahri et al., 2021) and (2) DACL+ (Verma et al., 2021) which are contrastive learning methods (Chen et al., 2020), (3) VIME-self (Yoon et al., 2020) and (4) SubTab (Ucar et al., 2021) which are autoencoder based methods. These four methods are the most commonly-evaluated baselines in the pretraining literature for tabular data (Ucar et al., 2021; Hajiramezani et al., 2022; Levin et al., 2022; Majmundar et al., 2022).

Domain-agnostic pretraining methods: In the image setting, SimCLR (Chen et al., 2020) and SimSiam (Chen & He, 2021) are two widely-adopted unsupervised pretraining methods. They rely on the idea that two slightly augmented variations of a data point should have similar representations. This simple idea generalizes across domains, so we apply it to tabular LTR. We discuss them below:

SimCLR (Chen et al., 2020): In a batch of B query-groups, the feature vector $\mathbf{x}_{q,i}$ of q -th query group and its i -th item is stochastically augmented twice to obtain $\mathbf{x}_{q,i}^{(0)}$ and $\mathbf{x}_{q,i}^{(1)}$,

which are called a *positive pair*. In this paper, we use combinations of the following two augmentation strategies for all domain-agnostic pretraining methods: (a) *Zeroing*: Zero out features selected randomly and i.i.d. with fixed probability. (b) *Gaussian noise*: Add Gaussian noise of a fixed scale to all features. Further details are given in Appendix B.3.

Next, a base encoder $h(\cdot)$ and projection head $g(\cdot)$ map $\mathbf{x}_{q,i}^{(a)}$ to $\mathbf{z}_{q,i}^{(a)} = g(h(\mathbf{x}_{q,i}^{(a)}))$ for $a = 0, 1$. Then by minimizing the InfoNCE loss (Oord et al., 2018) in this projected representation space, we push the positive pair, $(\mathbf{z}_{q,i}^{(0)}, \mathbf{z}_{q,i}^{(1)})$, closer while pushing them farther from all other augmented data points in the batch, which are called their *negatives*. Formally, for $\mathbf{x}_{q,i}$, we minimize $\ell_{q,i}^{(0)} + \ell_{q,i}^{(1)}$, where:

$$\ell_{q,i}^{(a)} = -\cos(\mathbf{z}_{q,i}^{(a)}, \mathbf{z}_{q,i}^{(\bar{a})})/\tau + \log \sum_{q'=1}^B \sum_{i'=1}^{L_{q'}} \sum_{a'=0}^1 [\mathbf{1}\{(q', i', a') \neq (q, i, a)\} \cdot \exp(\cos(\mathbf{z}_{q,i}^{(a)}, \mathbf{z}_{q',i'}^{(a')})/\tau)] \quad (2)$$

$\bar{a} = 1 - a$, τ is the temperature parameter, $L_{q'}$ is the number of items in q' -th query group, and $\cos(\mathbf{z}, \tilde{\mathbf{z}}) = \langle \mathbf{z}, \tilde{\mathbf{z}} \rangle / \|\mathbf{z}\| \|\tilde{\mathbf{z}}\|$ denotes the cosine similarity. After pretraining, the encoder h is used in downstream applications. SimCLR achieves superior performance in many domains (Chen & He, 2021; Wang et al., 2022b;a). It is known that contrasting a pair of data samples which are hard to distinguish from each other, called *hard negatives* can help an encoder learn good representations (Robinson et al., 2020; Oh Song et al., 2016; Schroff et al., 2015; Harwood et al., 2017; Wu et al., 2017; Ge, 2018; Suh et al., 2019). SimCLR simply contrasts all pairs of data samples in the batch (including from other query groups) with the assumption that a large enough batch is likely to contain a hard negative.

SimSiam (Chen & He, 2021): Similar to SimCLR, for each data point $\mathbf{x}_{q,i}$, this method produces stochastically-augmented *positive pairs* and their *projected* representations. However, we pass these representations further through a predictor $\text{pred}(\cdot)$, to get $\mathbf{p}_{q,i}^{(a)} = \text{pred}(\mathbf{z}_{q,i}^{(a)})$ for $a = 0, 1$. Finally, we maximize the similarity between the projected and predicted representations: $\cos(\mathbf{p}_{q,i}^{(0)}, \text{sg}(\mathbf{z}_{q,i}^{(1)})) + \cos(\mathbf{p}_{q,i}^{(1)}, \text{sg}(\mathbf{z}_{q,i}^{(0)}))$, where sg is a stop-gradient. Unlike SimCLR, there are no *negatives*, i.e., the loss function does not push the representation of an augmented sample away from that of other samples’ augmentations. The asymmetry in the loss due to the stop gradient and the predictor prevents representation collapse to make negatives unnecessary (Tian et al., 2021; Zhang et al., 2021; Wang et al., 2022b). If the number of items in each query group is L , then the time/space complexity for SimSiam is only $O(BL)$ per batch, whereas it is $O(B^2 L^2)$ for SimCLR. Therefore, SimSiam is more efficient and it can scale to larger batchsizes and data.

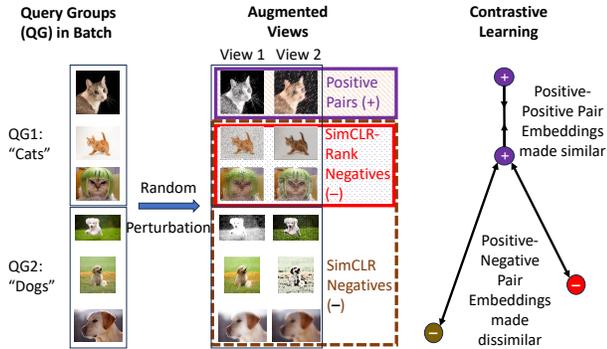


Figure 2: In SimCLR and its variants, *positive pairs*, or augmented versions of the same data sample, are trained to have similar embeddings. Positive-negative pairs, or augmented samples originating from two different data points or classes, are trained to have distant embeddings. In vanilla SimCLR, each positive pair is contrasted with *all other items in the batch*, denoted by the data points contained in the brown dashed line above. In SimCLR-Rank, each positive pair is contrasted with *only the items in the same Query Group (QG)*, denoted by the data points inside the red solid line above. By limiting the negatives to a smaller set of *harder-to-contrast* relevant items, SimCLR-Rank is faster and scales to larger datasets.

Pretraining method for LTR: Motivated by (a) SimCLR’s high complexity and (b) the efficacy of contrasting with hard negatives, we propose SimCLR-Rank, an LTR-specific alternative. Recall that the SimCLR loss for an item in equation 2 uses all the other items in the batch as negatives. Good pretraining strategies usually exploit the structure of the data, e.g. text sequences motivates masked token prediction (Devlin et al., 2018). Studying the query group structure of LTR data, we notice that high-quality hard negatives are freely available as *the other items from the same query group*. These items are retrieved and deemed potentially similarly relevant for this query by an upstream retrieval model, making these items harder negatives than other items in the batch. So, we propose SimCLR-Rank, which modifies SimCLR to contrast only with the other items in the same query group, as illustrated in Figure 2. Formally, we modify the per-augmentation loss $\ell_{q,i}^{(a)}$ as:

$$\ell_{q,i}^{(a)} = -\cos(\mathbf{z}_{q,i}^{(a)}, \mathbf{z}_{q,i}^{(\bar{a})})/\tau + \log \sum_{i'=1}^{L_q} \sum_{a'=0}^1 \mathbf{1}\{(i', a') \neq (i, a)\} \cdot \exp(\cos(\mathbf{z}_{q,i}^{(a)}, \mathbf{z}_{q,i'}^{(a')})/\tau) \quad (3)$$

The SimCLR-Rank loss has a time/space complexity of $O(BL^2)$, making it faster than SimCLR’s $O(B^2L^2)$ complexity, and enjoying SimSiam’s linear complexity in B . Next, we empirically study these methods.

3.1. Evaluation of unsupervised pretraining techniques

In this subsection, we compare the pretraining strategies we discussed above in the following tabular LTR setting.

Dataset. We use MSLRWEB30K (or MSLR for brevity) (Qin & Liu, 2013), Yahoo Set1 (or Set1 for brevity) (Chapelle & Chang, 2011), and Istella_S (Lucchese et al., 2016) (a smaller version of Istella (Dato et al., 2016)). In the tabular LTR literature, these three datasets are generally considered the standard; most recent papers evaluate (only) on these datasets (Qin et al., 2021; Ai et al., 2019; Pang et al., 2020; Yang et al., 2022; Yu, 2020). To simulate scarcity of labels, we let 0.1% of the training query groups in each dataset keep their labels, while the rest of the query groups have no labels. Dataset statistics are in Appendix B.1.

Methodology. For the neural pretraining approaches (which include SimCLR-Rank, SimSiam, SimCLR, SubTab, VIMEself, DACL+, and SCARF), the base of the model is the tabular ResNet¹ (check the footnote for our explanation on the architecture choice). We also evaluate “SimCLR-Sample”, a neural pretraining baseline we create to demonstrate the value of negatives from the same query group. SimCLR-Sample is like SimCLR-Rank, but at the beginning of each epoch the items are first randomly permuted and grouped into fake query groups. Then we apply SimCLR-Rank. Thus SimCLR-Sample has the same time/space complexity as SimCLR-Rank. We pretrain on all the available query groups and finetune/train on only the labeled query groups.

For the GBDT-based approaches, we use the implementation in `lightgbm` (Ke et al., 2017) and adopt a tuning strategy similar to Qin et al. (2021). We report “Supervised GBDT” as GBDT utilizing only the labeled data, and in “Semi-supervised GBDT” we report the best among pseudo-labeled (Pseudo-Label, 2013) and PCA-enriched (Duh & Kirchoff, 2008) GBDTs as well as the combination of the two techniques. We describe our experimental setup in detail in Appendix B.3. This includes hyperparameters, procedures for data augmentation, GBDT pseudo-labeling/PCA details, model training, and finetuning. For comparison to prior baselines on tabular data, our experimental setup is chosen to mimic the original setup in them.

Results. Table 1 charts the test NDCG for each of the pre-

¹We emphasize that we do not prescribe any specific deep learning architecture choice; the choice for the best DL model may depend on the specific LTR task. Our proposed workflow of pretraining deep models in LTR is agnostic to the specific choice of DL architecture. Prior work has proposed DL architectures designed for click-through rate (CTR)/LTR based on factorization machines (Rendle, 2010), such as DeepFM and DCNv2 (Guo et al., 2017; Wang et al., 2021), which we compare to in Section 4.1.2 and Section 4.1.2 against our pretrained models. Note that it is unclear currently how to do unsupervised pretraining using these CTR/LTR specific architectures at the moment.

Table 1: **Pretraining methods:** A comparison across unsupervised pretraining methods on NDCG, averaged over 3 trials. Here only 0.1% of the query groups have labels for finetuning. **Bold** numbers denote the best in a column, and underlined numbers are within the margin of error of the best. Given that SCARF/DACL+/SimCLR are too slow to scale to large LTR datasets (for example on the industrial-scale proprietary online shopping dataset), we find that SimCLR-Rank and SimSiam are the best pretraining methods for tabular LTR.

Method	MSLR (\uparrow)	Yahoo Set1 (\uparrow)	Istella (\uparrow)
Supervised GBDT	0.2801 \pm 0.0002	0.5083 \pm 0.0197	0.5450 \pm 0.0000
Semi-supervised GBDT	0.2839 \pm 0.0004	0.5061 \pm 0.0267	0.4656 \pm 0.0310
SCARF (Bahri et al., 2021)	0.3807 \pm 0.0016	0.5884 \pm 0.0129	0.5542 \pm 0.0024
DACL+ (Verma et al., 2021)	<u>0.3833 \pm 0.0037</u>	0.5887 \pm 0.0012	0.5626 \pm 0.0024
VIME-self (Yoon et al., 2020)	0.3834 \pm 0.0011	0.5839 \pm 0.0058	0.5514 \pm 0.0025
SubTab (Ucar et al., 2021)	0.3748 \pm 0.0025	0.5814 \pm 0.0062	0.5082 \pm 0.0050
SimCLR (Chen et al., 2020)	0.3827 \pm 0.0027	0.5837 \pm 0.0093	<u>0.5602 \pm 0.0055</u>
SimCLR-Sample	0.3498 \pm 0.0541	0.5803 \pm 0.0070	0.5479 \pm 0.0016
SimCLR-Rank	0.3868 \pm 0.0026	0.5843 \pm 0.0062	<u>0.5609 \pm 0.0040</u>
SimSiam (Chen & He, 2021)	0.3790 \pm 0.0028	0.6100 \pm 0.0072	0.5189 \pm 0.0096

training methods we study. SimCLR-Rank and SimSiam perform the best or the second best among these methods for most of the datasets. In particular, SimCLR-Rank outperforms SimCLR-Sample, demonstrating the value and efficiency of negatives from the same query group. As expected, SimCLR-Rank is also 7-14x faster than SimCLR/SCARF/DACL+ (Table 10 in Appendix D.1). We also find that the semi-supervised GBDT baselines do not perform well—surprisingly, they tend to achieve lower NDCG than the supervised GBDT baseline without semi-supervised learning. Given these advantages in terms of test NDCG and speed, we select SimCLR-Rank and SimSiam as the best overall pretraining methods for tabular LTR, and use them for our continued evaluation.

4. Empirical evaluation

Next we compare the best DL pretraining methods for tabular LTR (SimCLR-Rank and SimSiam) against GBDT and non-pretrained DL models on (1) the public datasets under different label scarcity patterns to simulate real-world scenarios², and (2) a private large-scale ranking dataset which is naturally label-scarce.

4.1. Public datasets results

For the public datasets (MSLRWEB30K, Yahoo Set1, Istella_S; details given in Section 3.1), we simulate label scarcity occurring in two scenarios: (1) *crowdsourcing* of labels for a limited number of query groups, and (2) using scarce *implicit binary user feedback* (e.g. clicks) as labels.

²Code is provided at <https://anonymous.4open.science/r/ltr-pretrain-3C01/README.md>.

4.1.1. SIMULATED CROWDSOURCING OF LABELS

Here we simulate a very practical setting where modelers use crowdsourcing to collect labels for only a fraction of the query groups due to limited resources.

Dataset. For each of the three public datasets, we vary the fraction of labeled query groups in the training set in $\{0.001, 0.002, 0.005, 0.1, 0.5, 1.0\}$. Note that within each labeled query group all items are labeled.

Methodology. We repeat the methodology from Section 3.1 for the GBDT and the SimCLR-Rank and SimSiam pretraining methods. For each dataset and labeling fraction, we report the test metric of the best GBDT ranker amongst semi-supervised GBDTs and supervised GBDTs under “GBDT”. Similarly, under “no pretrain DL” we report the best supervised DL model amongst tabular ResNet (Gorishniy et al., 2021), DeepFM (Guo et al., 2017), and DCNv2 (Wang et al., 2021) under the name “no pretrain DL”. Finally, we report the best pretrained ranker amongst SimCLR-Rank and SimSiam under the name “pretrained DL”. Pretrained DL rankers use tabular ResNet (see Appendix B.3 for more details). During hyperparameter tuning, we tune for NDCG and report the resulting test NDCG and test Outlier-NDCG.

Results. We compare pretrained rankers, non-pretrained DL rankers, and GBDTs on public datasets in Figure 1. The results for Outlier-NDCG are provided in Figure 8 in Appendix D.2. We find that pretrained rankers outperform non-pretrained methods (including GBDTs) on NDCG and Outlier-NDCG across all public datasets, up to a dataset-dependent fraction of QGs labeled, as shown in Figure 1 (for NDCG) and Figure 8 (for Outlier-NDCG). We also note that GBDTs have better NDCG than non-pretrained DL models in most of these regimes. This shows that there exist sce-

narios with a limited supply of labeled query groups, where self-supervised pretraining is the factor that allows deep learning models to outperform GBDTs. We also provide a detailed comparison between various semi-supervised and supervised GBDT models in Appendix D.4

4.1.2. SIMULATED IMPLICIT BINARY USER FEEDBACK

Here we simulate a scenario when users of an LTR system provide binary implicit feedback for the items, like the clicking of an advertisement on a webpage. This kind of feedback is usually infrequent in an LTR system, as it is the result of a user choosing to spend more time or resources on an item in a presented list. Therefore, most query groups will not have any labels, and even labeled query groups will only have a few positively labeled items. On the other hand, such labels are often cheaper to collect than crowdsourced labels, so they are used in many industry use cases.

Dataset. To simulate this scenario, we follow the methodology from Yang et al. (2022), to generate stochastic binary labels from the true relevance labels for training and validation sets of each of the public datasets. Note that we still use the true (latent) relevance labels in the test set for evaluation. This models a scenario where we observe binary labels that are noisy observations of a true latent label, but the task is to use these noisy labels to learn to rank according to the true labels. The details of this binarizing transformation are provided in Appendix B.4. In this transformation, a parameter, τ_{target} , implicitly controls how sparse the binary labels are (larger is more sparse). We select $\tau_{\text{target}} = 4.5$, to obtain 8.9%, 5.3%, and 25.2% labeled query groups in MSLR, Yahoo Set1, Istella_S datasets, respectively. We present detailed data statistics in Table 13 of Appendix.

Methodology. We reuse the methodology of Section 4.1.1, except we use a linear scoring head for pretrained models to improve stability and performance (see Appendix B.3). As in Section 4.1.1, we report the best of multiple models for each method and dataset (see Section 4.1.1 for details).

Results. In Table 2, we present the test NDCG and Outlier-NDCG metrics of the models under this setting. We see that pretrained rankers significantly outperform GBDTs and non-pretrained DL rankers in this simulated binary user feedback setting with up to, even with 25.2% labeled query groups (Istella_S). Notice that for MSLRWEB30K and Yahoo Set1 datasets, the second best model is the GBDT, and the non-pretrained DL ranker has a much worse Outlier-NDCG. This aligns with the recent observation that GBDT models perform better than DL models in supervised training with datasets containing irregularities (McElfresh et al., 2023). However, here pretraining helps DL rankers close this gap and at times beat GBDTs in outlier performance. We also present results for $\tau_{\text{target}} \in \{4.25, 5.1\}$ in Appendix D.3. There, we find that making the labels sparser ($\tau_{\text{target}} = 5.1$)

increases the relative improvement of pretrained DL rankers, while in a less sparse setting ($\tau_{\text{target}} = 4.25$) GBDTs dominate. This suggests that pretrained DL models outperform GBDTs when learning from rare, implicit binary user feedback, which is common in search/recommendation systems.

4.2. Large-scale real-world dataset results

In this section we test whether the results from the simulated label scarce settings above, especially Section 4.1.2, translate to a large-scale real-world LTR problem.

Dataset. Our proprietary dataset is derived from the online shopping logs of a large online retailer, numbering in the tens of millions of query groups. In this dataset, query groups consist of items presented to a shopper if they enter a search query. We assign each item a purchase label: 1 if the shopper purchased the item, or 0 if not. Hence, this is a real-world instance of the implicit binary user feedback setting simulated in Section 4.1.2. Further, only about 10% of the query groups have items with non-zero labels. See Appendix B.1 for more dataset details.

Methodology. We compare the performance of pretrained DL models—SimCLR-Rank and SimSiam—with GBDT and non-pretrained DL models. SimSiam and SimCLR-Rank are first pretrained on all the query groups before being finetuned on the labeled query groups with some non-zero purchase labels. The GBDT and the non-pretrained models are only trained on the labeled data. More details about these models and training strategies is given in Appendix B.5.

Results. Our results are summarized in Table 3 and all numbers are given as relative percentage improvements ($\Delta\%$) over the GBDT ranker. First, we note that all the DL models outperform GBDT models in terms of NDCG. Amongst them, the pretrained model using SimSiam performs the best with a gain of over 5.5%. This is substantial, given that in large-scale industry datasets, even a 2% improvement in NDCG is considered significant. In terms of Outlier-NDCG, the story is very different. Here, we see that the non-pretrained DL model and SimCLR-Rank performs up to -7% worse than the GBDT. However, SimSiam pretraining allows the DL ranker to overcome this shortcoming and beat the outlier performance of the GBDT. This verifies our hypothesis from Section 4.1.2 that pretraining can help DL models outperform GBDTs on large-scale real-world LTR problems, especially with implicit user feedback. Lastly, we see a significant performance difference between SimSiam and SimCLR-Rank. We will investigate this using an ablation study with public datasets later in (Section 4.3).

4.3. Ablations: Choices for pretrained models in LTR

Our experiments suggest that the best pretraining and finetuning strategies for the LTR problem are different from the

Table 2: **Main result (Simulated implicit binary user feedback):** When $X\%$ (MSLRWEB30K=8.9%, Yahoo Set1=5.3%, and in Istella_S=25.2%) of query groups are assigned a few binary labels, pretrained DL method outperform both GBDT and non-pretrained DL models. We generate labels with $\tau_{\text{target}} = 4.5$ (Section 4.1.2). \clubsuit means that pretrained models are significantly better in NDCG than non-pretrained models, measured by a t-test with significance $p < 0.05$.

Method	MSLRWEB30K (\uparrow)	Yahoo Set1 (\uparrow)	Istella (\uparrow)
	NDCG (\uparrow)		
GBDT	0.346 \pm 0.001	0.617 \pm 0.000	0.602 \pm 0.000
No Pretrain DL	0.333 \pm 0.001	0.621 \pm 0.000	0.608 \pm 0.002
Pretrained DL	0.356 \pm 0.004\clubsuit	0.622 \pm 0.001	0.613 \pm 0.002\clubsuit
	Outlier-NDCG (\uparrow)		
GBDT	0.288 \pm 0.000	0.540 \pm 0.000	0.678 \pm 0.000
No Pretrain DL	0.238 \pm 0.026	0.536 \pm 0.003	0.687 \pm 0.015
Pretrained DL	0.289 \pm 0.012	0.543 \pm 0.002\clubsuit	0.699 \pm 0.003

Table 3: **Main result (Real-world dataset):** Results on a large-scale proprietary online shopping dataset. We give percentage improvement of rankers pretrained by SimSiam and SimCLR-Rank over GBDTs. Unsupervised pre-training improves (1) overall (full dataset) performance (internally, 2% is a big improvement), and (2) outlier performance.

Method	NDCG (\uparrow)	Outlier-NDCG (\uparrow)
	$\Delta\%$ change w.r.t. GBDT	
GBDT	+0.00%	+0.00%
No Pretrain DL	+3.82% \pm 0.14%	-6.97% \pm 0.43%
SimCLR-Rank	+4.00% \pm 0.14%	-4.85% \pm 0.44%
SimSiam	+5.64% \pm 0.14%	+2.75% \pm 0.43%

known best practices for image or text tasks. In ablation studies, we make a few observations.

(1) SimSiam and SimCLR-Rank perform significantly differently depending on the dataset. This is different from the observations of Chen & He (2021), who proposed SimSiam and found that it performs similarly to SimCLR derivatives. We compare SimCLR-Rank with SimSiam at different levels of label scarcity and qualitatively compare their embeddings in Appendix C.1.

(2) Combining models pretrained by SimCLR-Rank and SimSiam can produce a model that is competitive against SimCLR-Rank and SimSiam across all datasets. While the combined model does not uniformly dominate across all labeled query group percentages, it is a promising first step towards a single unified method for unsupervised pretraining in ranking. Further details can be found in Appendix C.2.

(3) Full finetuning of the model performs better than linear probing. Linear probing is a popular finetuning strategy in text and image domains, which only finetunes a linear

layer over frozen pretrained embeddings (Chen & He, 2021; Chen et al., 2020; Peters et al., 2019; Kumar et al., 2022). We find that in NDCG of both SimCLR-Rank and SimSiam, linear probing performs consistently worse than full finetuning (details in Appendix C.3).

5. Conclusion

We study the learning-to-rank (LTR) problem with tabular data under an abundance of unlabeled data, but with scarcity of labeled data—a very common scenario in real-world practical LTR systems. Prior works on supervised learning with tabular data have shown that GBDTs outperform deep learning (DL) models, especially on datasets with outliers. In this paper, we find that in the label-scarce setting, one can use DL pretraining methods, which can exploit available unlabeled data, to obtain new state-of-the-art performance. Through experiments with public LTR datasets and a real-world large-scale online shopping dataset, we showed that pretraining allows DL rankers to outperform GBDT rankers, especially on outlier data. Finally, we provided empirically justified guidelines for pretraining on LTR datasets, including a new and faster LTR-specific pretraining method; SimCLR-Rank. There are still several open questions around pretraining for tabular and LTR data. For example, it is still an open question whether we can achieve knowledge transfer across different tabular datasets, as it is commonly done in image and text domains. It will also be useful to investigate if there are other LTR-specific pretraining methods which can further improve the ranking performance.

Acknowledgments

GF and CH gratefully acknowledge the research grant support of Google, Intel, the Sloan Foundation, and Bosch.

References

- Ai, Q., Bi, K., Guo, J., and Croft, W. B. Learning a deep listwise context model for ranking refinement. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pp. 135–144, 2018.
- Ai, Q., Wang, X., Bruch, S., Golbandi, N., Bendersky, M., and Najork, M. Learning groupwise multivariate scoring functions using deep neural networks. In *Proceedings of the 2019 ACM SIGIR international conference on theory of information retrieval*, pp. 85–92, 2019.
- Bahri, D., Jiang, H., Tay, Y., and Metzler, D. Scarf: Self-supervised contrastive learning using random feature corruption. *arXiv preprint arXiv:2106.15147*, 2021.
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., and Raffel, C. A. Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems*, 32, 2019.
- Bower, A., Eftekhari, H., Yurochkin, M., and Sun, Y. Individually fair ranking. *arXiv preprint arXiv:2103.11023*, 2021.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.
- Bruch, S., Zoghi, M., Bendersky, M., and Najork, M. Revisiting approximate metric optimization in the age of deep neural networks. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pp. 1241–1244, 2019.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pp. 89–96, 2005.
- Burges, C. J. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11(23-581):81, 2010.
- Cao, Z., Qin, T., Liu, T.-Y., Tsai, M.-F., and Li, H. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, pp. 129–136, 2007.
- Chapelle, O. and Chang, Y. Yahoo! learning to rank challenge overview. In *Proceedings of the learning to rank challenge*, pp. 1–24. PMLR, 2011.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Chen, W., Liu, T.-Y., Lan, Y., Ma, Z.-M., and Li, H. Ranking measures and loss functions in learning to rank. *Advances in Neural Information Processing Systems*, 22, 2009.
- Chen, X. and He, K. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15750–15758, 2021.
- Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Aradhye, H., Anderson, G., Corrado, G., Chai, W., Ispir, M., et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, pp. 7–10, 2016.
- Darabi, S., Fazeli, S., Pazoki, A., Sankaraman, S., and Sarrafzadeh, M. Contrastive mixup: Self-and semi-supervised learning for tabular domain. *arXiv preprint arXiv:2108.12296*, 2021.
- Dato, D., Lucchese, C., Nardini, F. M., Orlando, S., Perego, R., Tonello, N., and Venturini, R. Fast ranking with additive ensembles of oblivious and non-oblivious regression trees. *ACM Transactions on Information Systems (TOIS)*, 35(2):1–31, 2016.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Duh, K. and Kirchhoff, K. Learning to rank with partially-labeled data. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 251–258, 2008.
- Friedman, J. H. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pp. 1189–1232, 2001.
- Ge, S., Mishra, S., Li, C.-L., Wang, H., and Jacobs, D. Robust contrastive learning using negative samples with diminished semantics. *Advances in Neural Information Processing Systems*, 34:27356–27368, 2021.
- Ge, W. Deep metric learning with hierarchical triplet loss. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 269–285, 2018.
- Glass, M., Rossiello, G., Chowdhury, M. F. M., Naik, A. R., Cai, P., and Gliozzo, A. Re2g: Retrieve, rerank, generate. *arXiv preprint arXiv:2207.06300*, 2022.
- Goren, G., Kurland, O., Tennenholtz, M., and Raiber, F. Ranking robustness under adversarial document manipulations. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 395–404, 2018.

-
- Gorishniy, Y., Rubachev, I., Khruikov, V., and Babenko, A. Revisiting deep learning models for tabular data. *Advances in Neural Information Processing Systems*, 34: 18932–18943, 2021.
- Grill, J.-B., Strub, F., Althché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- Grinsztajn, L., Oyallon, E., and Varoquaux, G. Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems*, 35:507–520, 2022.
- Guo, H., Tang, R., Ye, Y., Li, Z., and He, X. Deepfm: a factorization-machine based neural network for ctr prediction. *arXiv preprint arXiv:1703.04247*, 2017.
- Hager, P., Menten, M. J., and Rueckert, D. Best of both worlds: Multimodal contrastive learning with tabular and imaging data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 23924–23935, June 2023.
- Hajiramezanali, E., Diamant, N. L., Scalia, G., and Shen, M. W. Stab: Self-supervised learning for tabular data. In *NeurIPS 2022 First Table Representation Workshop*, 2022.
- Han, S., Wang, X., Bendersky, M., and Najork, M. Learning-to-rank with bert in tf-ranking. *arXiv preprint arXiv:2004.08476*, 2020.
- Harwood, B., Kumar BG, V., Carneiro, G., Reid, I., and Drummond, T. Smart mining for deep metric learning. In *Proceedings of the IEEE international conference on computer vision*, pp. 2821–2829, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hegselmann, S., Buendia, A., Lang, H., Agrawal, M., Jiang, X., and Sontag, D. Tabllm: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*, pp. 5549–5581. PMLR, 2023.
- Hendrycks, D., Mazeika, M., Kadavath, S., and Song, D. Using self-supervised learning can improve model robustness and uncertainty. *Advances in neural information processing systems*, 32, 2019.
- Howard, J. and Ruder, S. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*, 2018.
- Huang, X., Khetan, A., Cvitkovic, M., and Karnin, Z. Tab-transformer: Tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*, 2020.
- Jeffares, A., Liu, T., Crabbé, J., Imrie, F., and van der Schaar, M. Tangos: Regularizing tabular neural networks through gradient orthogonalization and specialization. *arXiv preprint arXiv:2303.05506*, 2023.
- Jeong, J. and Shin, J. Consistency regularization for certified robustness of smoothed classifiers. *Advances in Neural Information Processing Systems*, 33:10558–10570, 2020.
- Joachims, T. Training linear svms in linear time. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 217–226, 2006.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kitaev, N., Cao, S., and Klein, D. Multilingual constituency parsing with self-attention and pre-training. *arXiv preprint arXiv:1812.11760*, 2018.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., and Iwasawa, Y. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35: 22199–22213, 2022.
- Kumar, A., Raghunathan, A., Jones, R. M., Ma, T., and Liang, P. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022.
- Kveton, B., Meshi, O., Zoghi, M., and Qin, Z. On the value of prior in online learning to rank. In *International Conference on Artificial Intelligence and Statistics*, pp. 6880–6892. PMLR, 2022.
- Lamkhede, S. D. and Kofler, C. Recommendations and results organization in netflix search. In *Proceedings of the 15th ACM Conference on Recommender Systems*, pp. 577–579, 2021.
- Lee, K. and Shin, J. Rényiicl: Contrastive representation learning with skew rényi divergence. *Advances in Neural Information Processing Systems*, 35:6463–6477, 2022.

- Levin, R., Cherepanova, V., Schwarzschild, A., Bansal, A., Bruss, C. B., Goldstein, T., Wilson, A. G., and Goldblum, M. Transfer learning with deep tabular models. *arXiv preprint arXiv:2206.15306*, 2022.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.
- Lian, J., Zhou, X., Zhang, F., Chen, Z., Xie, X., and Sun, G. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 1754–1763, 2018.
- Lin, J., Qu, Y., Guo, W., Dai, X., Tang, R., Yu, Y., and Zhang, W. Map: A model-agnostic pretraining framework for click-through rate prediction. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1384–1395, 2023.
- Liu, G., Yang, J., and Wu, L. Ptab: Using the pre-trained language model for modeling tabular data. *arXiv preprint arXiv:2209.08060*, 2022.
- Liu, H., Chen, Y., Fritz, B., and King, C. Tabcontrast: A local-global level method for tabular contrastive learning. In *NeurIPS 2023 Second Table Representation Learning Workshop*, 2023.
- Liu, P., Qiu, X., and Huang, X. Adversarial multi-task learning for text classification. *arXiv preprint arXiv:1704.05742*, 2017.
- Liu, T.-Y. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval*, 3(3):225–331, 2009.
- Lucchese, C., Nardini, F. M., Orlando, S., Perego, R., Silvestri, F., and Trani, S. Post-learning optimization of tree ensembles for efficient ranking. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 949–952, 2016.
- Majmundar, K., Goyal, S., Netrapalli, P., and Jain, P. Met: Masked encoding for tabular data. *arXiv preprint arXiv:2206.08564*, 2022.
- McClosky, D., Charniak, E., and Johnson, M. Effective self-training for parsing. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pp. 152–159, 2006.
- McElfresh, D., Khandagale, S., Valverde, J., Ramakrishnan, G., Goldblum, M., White, C., et al. When do neural nets outperform boosted trees on tabular data? *arXiv preprint arXiv:2305.02997*, 2023.
- Mitra, B., Craswell, N., et al. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval*, 13(1):1–126, 2018.
- Miyato, T., Maeda, S.-i., Koyama, M., and Ishii, S. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- Nam, J., Song, W., Park, S. H., Tack, J., Yun, S., Kim, J., and Shin, J. Semi-supervised tabular classification via in-context learning of large language models. In *Workshop on Efficient Systems for Foundation Models@ ICML2023*, 2023a.
- Nam, J., Tack, J., Lee, K., Lee, H., and Shin, J. Stunt: Few-shot tabular learning with self-generated tasks from unlabeled tables. *arXiv preprint arXiv:2303.00918*, 2023b.
- Naumov, M., Mudigere, D., Shi, H.-J. M., Huang, J., Sundaraman, N., Park, J., Wang, X., Gupta, U., Wu, C.-J., Azzolini, A. G., et al. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091*, 2019.
- Nogueira, R., Yang, W., Cho, K., and Lin, J. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*, 2019.
- Oh Song, H., Xiang, Y., Jegelka, S., and Savarese, S. Deep metric learning via lifted structured feature embedding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4004–4012, 2016.
- Oliver, A., Odena, A., Raffel, C. A., Cubuk, E. D., and Goodfellow, I. Realistic evaluation of deep semi-supervised learning algorithms. *Advances in neural information processing systems*, 31, 2018.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Pan, F., Caim, M., Glass, M., Gliozzo, A., and Hendler, J. End-to-end table question answering via retrieval-augmented generation. *arXiv preprint arXiv:2203.16714*, 2022.
- Pang, L., Xu, J., Ai, Q., Lan, Y., Cheng, X., and Wen, J. Setrank: Learning a permutation-invariant ranking model for information retrieval. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pp. 499–508, 2020.

-
- Penha, G., Câmara, A., and Hauff, C. Evaluating the robustness of retrieval pipelines with query variation generators. In *European conference on information retrieval*, pp. 397–412. Springer, 2022.
- Peters, M. E., Ruder, S., and Smith, N. A. To tune or not to tune? adapting pretrained representations to diverse tasks. *arXiv preprint arXiv:1903.05987*, 2019.
- Pseudo-Label, D.-H. L. The simple and efficient semi-supervised learning method for deep neural networks. In *ICML 2013 Workshop: Challenges in Representation Learning*, pp. 1–6, 2013.
- Qin, T. and Liu, T. Introducing LETOR 4.0 datasets. *CoRR*, abs/1306.2597, 2013. URL <http://arxiv.org/abs/1306.2597>.
- Qin, Z., Yan, L., Zhuang, H., Tay, Y., Pasumarthi, R. K., Wang, X., Bendersky, M., and Najork, M. Are neural rankers still outperformed by gradient boosted decision trees? 2021.
- Qu, Y., Cai, H., Ren, K., Zhang, W., Yu, Y., Wen, Y., and Wang, J. Product-based neural networks for user response prediction. In *2016 IEEE 16th international conference on data mining (ICDM)*, pp. 1149–1154. IEEE, 2016.
- Rendle, S. Factorization machines. In *2010 IEEE International conference on data mining*, pp. 995–1000. IEEE, 2010.
- Rizve, M. N., Duarte, K., Rawat, Y. S., and Shah, M. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. *arXiv preprint arXiv:2101.06329*, 2021.
- Robinson, J., Chuang, C.-Y., Sra, S., and Jegelka, S. Contrastive learning with hard negative samples. *arXiv preprint arXiv:2010.04592*, 2020.
- Rubachev, I., Alekberov, A., Gorishniy, Y., and Babenko, A. Revisiting pretraining objectives for tabular deep learning. *arXiv preprint arXiv:2207.03208*, 2022.
- Sajjadi, M., Javanmardi, M., and Tasdizen, T. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. *Advances in neural information processing systems*, 29, 2016.
- Schroff, F., Kalenichenko, D., and Philbin, J. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 815–823, 2015.
- Shi, W., Gong, Y., Ding, C., Tao, Z. M., and Zheng, N. Transductive semi-supervised deep learning using min-max features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 299–315, 2018.
- Shwartz-Ziv, R. and Armon, A. Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90, 2022.
- Song, Y. and Roth, D. Unsupervised sparse vector densification for short text similarity. In *Proceedings of the 2015 conference of the North American chapter of the association for computational linguistics: human language technologies*, pp. 1275–1280, 2015.
- Suh, Y., Han, B., Kim, W., and Lee, K. M. Stochastic class-based hard example mining for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7251–7259, 2019.
- Sui, Y., Wu, T., Cresswell, J., Wu, G., Stein, G., Huang, X., Zhang, X., Volkovs, M., et al. Self-supervised representation learning from random data projectors. *arXiv preprint arXiv:2310.07756*, 2023.
- Swezey, R., Grover, A., Charron, B., and Ermon, S. Pi-rank: Scalable learning to rank via differentiable sorting. *Advances in Neural Information Processing Systems*, 34: 21644–21654, 2021.
- Syed, T. and Mirza, B. Self-supervision for tabular data by learning to predict additive homoskedastic gaussian noise as pretext. *ACM Transactions on Knowledge Discovery from Data*, 2023.
- Tian, Y., Chen, X., and Ganguli, S. Understanding self-supervised learning dynamics without contrastive pairs. In *International Conference on Machine Learning*, pp. 10268–10278. PMLR, 2021.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Ucar, T., Hajiramezanali, E., and Edwards, L. Subtab: Subsetting features of tabular data for self-supervised representation learning. *Advances in Neural Information Processing Systems*, 34:18853–18865, 2021.
- Verma, V., Luong, T., Kawaguchi, K., Pham, H., and Le, Q. Towards domain-agnostic contrastive learning. In *International Conference on Machine Learning*, pp. 10530–10541. PMLR, 2021.
- Voorhees, E. M. The trec robust retrieval track. In *ACM SIGIR Forum*, volume 39, pp. 11–20. ACM New York, NY, USA, 2005.
- Wang, R., Shivanna, R., Cheng, D., Jain, S., Lin, D., Hong, L., and Chi, E. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the web conference 2021*, pp. 1785–1797, 2021.

-
- Wang, W., Kim, B.-H., and Ganapathi, V. Regclr: A self-supervised framework for tabular representation learning in the wild. *arXiv preprint arXiv:2211.01165*, 2022a.
- Wang, X., Fan, H., Tian, Y., Kihara, D., and Chen, X. On the importance of asymmetry for siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16570–16579, 2022b.
- Wang, Z. and Sun, J. Transtab: Learning transferable tabular transformers across tables. *Advances in Neural Information Processing Systems*, 35:2902–2915, 2022.
- Wu, C., Zhang, R., Guo, J., Chen, W., Fan, Y., de Rijke, M., and Cheng, X. Certified robustness to word substitution ranking attack for neural ranking models. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 2128–2137, 2022a.
- Wu, C., Zhang, R., Guo, J., Fan, Y., and Cheng, X. Are neural ranking models robust? *ACM Transactions on Information Systems*, 41(2):1–36, 2022b.
- Wu, C., Zhang, R., Guo, J., De Rijke, M., Fan, Y., and Cheng, X. Prada: practical black-box adversarial attacks against neural ranking models. *ACM Transactions on Information Systems*, 41(4):1–27, 2023.
- Wu, C.-Y., Manmatha, R., Smola, A. J., and Krahenbuhl, P. Sampling matters in deep embedding learning. In *Proceedings of the IEEE international conference on computer vision*, pp. 2840–2848, 2017.
- Xia, F., Liu, T.-Y., Wang, J., Zhang, W., and Li, H. List-wise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th international conference on Machine learning*, pp. 1192–1199, 2008.
- Xie, Q., Luong, M.-T., Hovy, E., and Le, Q. V. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10687–10698, 2020.
- Yang, S., Sanghavi, S., Rahmanian, H., Bakus, J., and SVN, V. Toward understanding privileged features distillation in learning-to-rank. *Advances in Neural Information Processing Systems*, 35:26658–26670, 2022.
- Yarowsky, D. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pp. 189–196, 1995.
- Yates, A., Nogueira, R., and Lin, J. Pretrained transformers for text ranking: Bert and beyond. In *Proceedings of the 14th ACM International Conference on web search and data mining*, pp. 1154–1156, 2021.
- Ye, C., Lu, G., Wang, H., Li, L., Wu, S., Chen, G., and Zhao, J. Ct-bert: learning better tabular representations through cross-table pre-training. *arXiv preprint arXiv:2307.04308*, 2023.
- Yoon, J., Zhang, Y., Jordon, J., and van der Schaar, M. Vime: Extending the success of self-and semi-supervised learning to tabular domain. *Advances in Neural Information Processing Systems*, 33:11033–11043, 2020.
- Yu, H.-T. Pt-ranking: A benchmarking platform for neural learning-to-rank, 2020.
- Zhang, C., Zhang, K., Zhang, C., Pham, T. X., Yoo, C. D., and Kweon, I. S. How does simsiam avoid collapse without negative samples? a unified understanding with self-supervised contrastive learning. In *International Conference on Learning Representations*, 2021.
- Zhang, P., Song, D., Wang, J., and Hou, Y. Bias-variance decomposition of ir evaluation. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pp. 1021–1024, 2013.
- Zhu, B., Shi, X., Erickson, N., Li, M., Karypis, G., and Shoaran, M. Xtab: Cross-table pretraining for tabular transformers. *arXiv preprint arXiv:2305.06090*, 2023.
- Zhuang, H., Qin, Z., Jagerman, R., Hui, K., Ma, J., Lu, J., Ni, J., Wang, X., and Bendersky, M. Rankt5: Fine-tuning t5 for text ranking with ranking losses. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2308–2313, 2023.

A. Detailed discussion of related work

Learning-To-Rank. In our paper, we focus on the traditional LTR setting where the features are all numeric (tabular data). However, there is a line of work in LTR where raw text is also an input. In this case, one can leverage large language models in the ranking setting (Zhuang et al., 2023; Han et al., 2020; Yates et al., 2021; Nogueira et al., 2019; Mitra et al., 2018).

In tabular LTR problems, the dominant models currently used are gradient boosted decision trees (GBDTs) (Friedman, 2001), which are not deep learning models. GBDT models, which perform well on tabular data, are adapted to the LTR setting via losses that are surrogates for ranking metrics like NDCG. Surrogate losses (including LambdaRank/RankNet (Burgess, 2010) and PiRank (Swezey et al., 2021)) are needed because many important ranking metrics (like NDCG) are non-differentiable. The combination of tree-based learners and ranking losses has become the de-facto standard in ranking problems, and deep models have yet to outperform them convincingly (Qin et al., 2021; Joachims, 2006; Ai et al., 2019; Bruch et al., 2019; Ai et al., 2018; Pang et al., 2020).

Deep tabular models. Given the success of neural methods in many other domains, there have been many attempts to adapt deep models to the tabular domain. TANGOS introduced special tabular-specific regularization to try to improve deep models’ performance (Jefferies et al., 2023). FT-transformer and TabTransformer were introduced as transformer-based approaches to tabular data (Gorishniy et al., 2021; Huang et al., 2020). All these models have failed to consistently or convincingly outperform tree-based methods based on their own evaluations.

Self-supervised learning. Self-supervised learning (SSL) or unsupervised pretraining has improved performance in settings where there is a significant source of unlabeled data like text (Devlin et al., 2018) and images (Chen et al., 2020). In SSL, deep models are first pretrained on perturbed unlabeled data using self-supervised tasks to learn useful representations for the data. Then these models are finetuned for a downstream task with labeled data. Finetuning often takes one of two forms: (1) linear probing (a popular finetuning strategy in text and images (Chen & He, 2021; Chen et al., 2020; Peters et al., 2019)), where we freeze the pretrained model and only update the linear head during supervised finetuning, and (2) full finetuning, where we update the whole model during supervised finetuning (Devlin et al., 2018). Sometimes a mix of the two is used (Kumar et al., 2022). The core idea behind prominent SSL approaches like SimSiam and SimCLR is to carefully perturb input training samples, and train a representation that is consistent for perturbations of the same sample. This provides robustness to natural perturbations and noise in data (Hendrycks et al.,

2019).

Inspired by the success of pretraining and self-supervised learning in images and text, several works show how to apply SSL to unlabeled tabular data. One strategy is to corrupt tabular data and train a deep model to reconstruct it (Yoon et al., 2020; Majmundar et al., 2022; Ucar et al., 2021; Nam et al., 2023b; Lin et al., 2023; Syed & Mirza, 2023; Hajiramezanali et al., 2022). Another approach is to use contrastive losses, which have been highly successful in the image domain (Chen et al., 2020). These methods are also applicable to our case because tabular data, like image data, is often composed of fixed-dimensional real vectors (Verma et al., 2021; Bahri et al., 2021; Lee & Shin, 2022; Hager et al., 2023; Liu et al., 2023; Darabi et al., 2021). Rubachev et al. (2022) evaluate a variety of different pretraining methods for tabular learning across many different datasets, finding that there is not a clear state of the art.

Robustness in LTR. There has been prior work on studying worst-case behavior (robustness) of rankers (Voorhees, 2005; Zhang et al., 2013; Goren et al., 2018; Wu et al., 2022b;a; Penha et al., 2022). Some previous metrics measure a model’s robustness against adversarial attack (Goren et al., 2018; Wu et al., 2022a; 2023). Others measure the model’s per-query performance variance on a dataset (Voorhees, 2005; Zhang et al., 2013; Wu et al., 2022b). Our outlier metric, Outlier-NDCG, is a departure from previous work because it is not directly a measure of robustness and it is possible for a model to perform *better* on outlier data.

Transfer learning. We summarize some related work on transfer learning in tabular domain. One direction revolves around pretraining models on common columns across many datasets (Zhu et al., 2023; Wang & Sun, 2022; Sui et al., 2023; Ye et al., 2023). Another direction leverages LLMs (large language models) to do few-shot tabular learning (Hegselmann et al., 2023; Liu et al., 2022; Nam et al., 2023a). Pretrained model features are often useful across datasets. For example, text features are transferable between text datasets (Devlin et al., 2018), which has contributed to the recent success of foundation models (Devlin et al., 2018; Touvron et al., 2023; Kojima et al., 2022; Brown et al., 2020). Unfortunately, it is still unclear how to achieve transferability in tabular LTR, where previous work has leveraged common columns (or at least column documentation) to achieve transfer learning (Levin et al., 2022; Hegselmann et al., 2023). However most LTR datasets do not have common columns or publicly-documented column meanings (Qin & Liu, 2013; Chapelle & Chang, 2011; Lucchese et al., 2016). Developing a method to transfer features in tabular LTR is an important open problem.

B. Experimental Details

B.1. Dataset statistics

In table Table 4, we provide the details of our three public datasets and the proprietary real-world dataset. Note that for the public dataset experiments we sub-sample the pre-training dataset in different ways (as described in Sections 3, 4.1.1 and 4.1.2) to create the training dataset and simulate scarcity of labeled samples. However, we use the same validation dataset without and with the labels for the self-supervised pretraining and supervised learning stages, respectively. Note that in all these datasets, the higher the label value the higher the relevance of the item.

B.2. Outliers query group details

Here we detail how we select the outlier query groups for the Outlier-NDCG metrics (Section 2.1). For our proprietary online shopping dataset, there are known outlier features that cause low quality predictions. These outliers arise from noise introduced in various stages in the data pipeline. We have mitigation strategies in place against these feature outliers, but they are often imperfect. Because we already know some of the outliers features, we deem any query group with items with outlier features as an outlier. These outliers comprise roughly 10% of the test set, which is still a large number of query groups (much larger than the number of query groups in the outlier datasets of the public datasets).

For our public datasets, we systematically select outliers as follows: we generate a histogram with 100 bins for each feature across the *validation dataset*. For example, Istella (Lamkhede & Kofler, 2021) has 220 features, so we have 220 different histograms. For each histogram, we scan from left to right on the bins until we have encountered at least G empty bins in a row, and if there is less than 1% of the validation set above this bin, then all the feature values above this bin are considered outliers. We also repeat this process right to left. Any test query group containing items with outlier feature values is labeled an outlier query group, and placed in the test outlier dataset. In Figure 3, we present an example histogram of a feature in Istella_S dataset.

Because different datasets have differently-sized typical gaps, we tune G for each dataset (MSLR, Yahoo, Istella) such that the resulting percentage of outlier queries is as close to 1% of the test set as we can get. MSLRWEB30K has $G = 5$, with 0.65% (40/6072) outlier queries, Yahoo has $G = 20$, with 1.4% (30/2147) outlier queries, and Istella has $G = 32$, with 0.46% (34/7202) outlier queries. 1% is a hyperparameter that can be tuned according to the user’s goals.

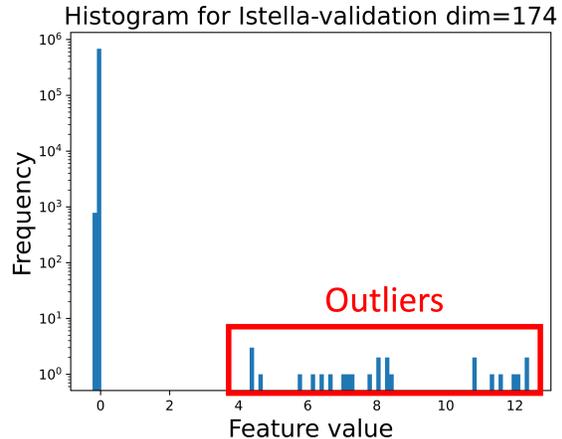


Figure 3: An example of outlier detection in Istella_S (Lucchese et al., 2016) for our Outlier-NDCG metric.

B.3. Model Settings and Training Details

Pretrained DL model.

Pretraining: (1) the encoder used for pretraining is the tabular ResNet (Gorishniy et al., 2021) with three ResNet blocks, with the final linear layer removed, (2) pretraining is done on the entire dataset with learning rate 0.0005 using Adam (Kingma & Ba, 2014) (3) with SimSiam/SimCLR-Rank/SimCLR-Sample, we tuned among four different augmentations for each pretraining method: randomly zeroing out features (“Zeroing”) with probabilities 0.1 or 0.7, and Gaussian noise with scale 1.0 or 2.0, (4) for SCARF we tune the corruption probability in $\{0.3, 0.6, 0.9\}$, (5) for VIME-self we tune the corruption probability in $\{0.3, 0.5, 0.7\}$, (6) for DACL+ we tune the mixup amount in $\{0.3, 0.6, 0.9\}$, (7) for SubTab we divide input features into 4 subsets with 75% overlap, and tune the masking probability in $\{0.3, 0.6, 0.9\}$, (8) we pretrain for 300 epochs for all methods, and (9) we use a batch size of roughly 200000 items for SimCLR-Rank, SimSiam, VIME-self, and SubTab as they are less GPU memory intensive while using a batch size of roughly 2000 items for SimCLR/DACL+/SCARF which are more GPU memory intensive.

Finetuning: (1) finetuning is done on the labeled train set by adding a three-layer MLP to the top of the pretrained model and training only this head for 100 epochs and then fully finetuning for 100 epochs using Adam with a learning rate of $5e-5$, (2) we use an average batch size of roughly 1000 items (may vary based on query group size), (3) we use the LambdaRank loss (Borges, 2010), (4) we use the validation set to perform early stopping (i.e. using the checkpoint that performed best on the validation set to evaluate on the test set).

Table 4: Dataset statistics on the MSLRWEB30K, Yahoo Set1, Istella_S, and proprietary datasets.

Dataset	Number of query groups					Number of features	List of labels
	Pretraining		Supervised learning		Test		
	Pretrain	SSL Val	Train	Val			
MSLRWEB30K	18151	6072	Subset Pretrain	SSL Val	6072	136	{0,1,2,3,4}
Yahoo Set1	14477	2147	Subset Pretrain	SSL Val	5089	700	{0,1,2,3,4}
Istella_S	19200	7202	Subset Pretrain	SSL Val	6523	220	{0,1,2,3,4}
Proprietary dataset	35.6M	1.62M	3.68M	171k	1.08M	63	{0,1}

Implicit binary user feedback setting. We reuse the methodology of Section 4.1.1, except we use a scoring head of one linear layer (as opposed to a three layer MLP) for pretrained models. We found that this improved stability and performance in the binary label setting.

No Pretrained DL model. There are a few rankers we evaluated for the no-pretrain DL baseline: (1) a 3-layer tabular ResNet from (Gorishniy et al., 2021) with a three-layer MLP on top of it, and was trained for 300 epochs on the labeled training set with learning rate tuned in {0.01, 0.001, 0.0001} using Adam, with early stopping using the validation set, (2) a DeepFM (Guo et al., 2017) model following the model architecture in the paper with learning rate tuned in {0.05, 0.005, 0.0005}, and (3) a DCNv2 (Wang et al., 2021) model with learning rate tuned in {0.05, 0.005, 0.0005} and cross layers tuned in {2,3,4}. DeepFM and DCNv2 were designed to use categorical features, and in MSLRWEB30K/Yahoo Set1/Istella_S do not document their categorical features (if any). We performed an analysis of the datasets and found 5, 48, 6 categorical features in MSLRWEB30K/Yahoo Set1/Istella_S respectively each with cardinality at most 3. We embedded these categorical features into embeddings of dimension 8, following the rule-of-thumb in (Wang et al., 2021).

GBDT model. The GBDT ranker is the one from `lightgbm` (Ke et al., 2017) and we grid search the number of leaves in {31, 96, 200} minimum data in leaf in {20, 60, 200}, and column sample in {0.5, 0.9, 1.0} individually for each data point (for a total of 9 difference choices), while letting the rest of the parameters be the default in `lightgbm` (our tuning strategy is similar to Qin et al. (2021)). We use the LambdaRank loss to train the models.

GBDT + pseudolabeling. We add pseudo-labeled GBDTs as a semi-supervised GBDT baseline. For GBDTs with pseudo-labeling (GBDT-pseudo), we propose the following algorithm as there are no references to our knowledge. First, we train the GBDT on the labeled train set, and then use this GBDT to pseudo-labels for the unlabeled train set. A GBDT ranker outputs real-valued scores. To convert real-valued scores into relevance scores, we take the resulting scores and evenly bin them into 31 buckets. The scores converted

to 0 are the lowest bucket, the scores converted to 1 are the second lowest bucket, and so on. We choose 31 because that is the largest number of labels allowed for GBDT ranking in `lightgbm` (the reason is because the LambdaRank loss requires the calculation of 2^{label} which can get very large).

GBDT + PCA. For GBDT+PCA, we take the entire training set’s features and learn a PCA on it with {5, 10, 15} components, similar to (Duh & Kirchhoff, 2008), to produce {5, 10, 15} new features (i.e. we tune among 5, 10, 15). We then train on the labeled train set, with both the original features and the new PCA features together.

GBDT + PCA + pseudolabeling. We also introduce a baseline where we perform PCA on the full training dataset, and then perform pseudolabeling as above. We tune the PCA similarly to the GBDT + PCA baseline.

Comparing finetuning strategies. When linear probing (LP), we freeze the pretrained model and update only a linear head on top of it for 200 epochs. In multilayer probing (MP), we freeze the pretrained model and update a 3-layer MLP head on top of it for 200 epochs. In full finetuning (FF), we use the finetuning strategy from Section 4.1.1. We use the validation set for early stopping.

All reported values for Outlier-NDCG are those achieved by the corresponding ranker with best overall NDCG. We report all results as averages over 3 trials, and use single V100 GPUs from a shared cluster to run the experiments.

B.4. Details of Binary Label generation used in Section 4.1.2

Following the methodology from Yang et al. (2022), we generate stochastic binary labels y from the relevance labels r as $y = \mathbb{1}\{t \cdot r + G_1 > t \cdot \tau_{\text{target}} + G_0\}$ where t is a temperature parameter, G_1, G_0 are standard Gumbels, and τ_{target} is a parameter controlling how sparse the binary labels are. Higher τ_{target} leads to higher label sparsity. Yang et al. (2022) show that y is 1 with probability $\sigma(t \cdot (r - \tau_{\text{target}}))$ where σ is the sigmoid function. We set $t = 4$ as in Yang et al. (2022). For a given τ_{target} , we produce produce a new dataset for each of MSLRWEB30K, Yahoo Set1, and Istella_S where we convert the relevance labels in the training/validation sets to

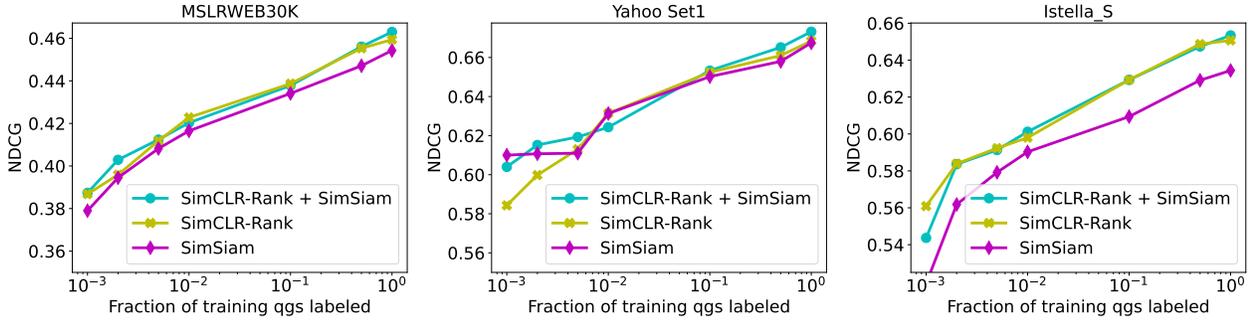


Figure 4: Comparison of SimCLR-Rank and SimSiam pretraining strategies on three public datasets with varying fraction of labeled query groups. Neither method consistently dominate the other in all the datasets. However, by combining SimCLR-Rank and SimSiam encoders, we can produce models that are competitive across all datasets. Data points are averages over 3 trials.

binary labels and keep the true relevance labels in the test set.

B.5. Large industry-scale dataset experimental details

For our DL ranker, first we pass the features (numerical and categorical) through some featurizers. Then we use a simple 4 hidden layer MLP, with and ReLU activation, to predict a single scalar score for each (query, item) pair. We also use dropout and batch normalization in input and hidden layers.

For the SimCLR pretraining stage, we remove the last linear layer of the above ranker to create the main encoder to produce the pretrained embeddings. Projector encoder here is a 2 hidden layer MLP, with ReLU activation and no bias, which outputs an input dimensional embedding. We also do not use any normalization or dropout in this projector.

For SimSiam we use the same embedder and a similar projector. The only difference in the projector being batch normalization in the hidden layers and batch normalization without affine learnable parameters in the output layer. Additionally, SimSiam also employs a predictor MLP, with a single hidden layer and ReLU activation, and with the same input and output sizes. It uses batch normalization and no bias in the hidden layer and no input and hidden dropouts.

For pretraining, we used an augmentation which randomly zeroed out some features. We also tried randomly swapping some features between items in the same batch or query group, but they didn't perform as well as zeros augmentation. During the fine-tuning stage we discard the projectors and predictor and add back the linear layer to complete the non-pretrained ranker structure. DL rankers are trained/fine-tuned using PiRank loss (Swezey et al., 2021), batchsize of 1000, AdamW optimizer with learning rate of 0.00001, weight decay of 0.01 and cosine scheduling with warmup. All training stages and experiments were conducted on ma-

chines with 4 GPUs, where we ran 100k iterations with ddp and picked the best checkpoint based on validation metric.

GBDT rankers were trained with LambdaRank using the `lightgbm` package. They were tuned with Bayesian hyperparameter optimization strategy by tuning number of leaves, learning rate, and minimum data in leaf in a reasonably large range with a maximum of 150 jobs. We used default values of the library for the other hyperparameters.

C. Ablation Studies

C.1. Comparing SimCLR-Rank and SimSiam

SimCLR-Rank and SimSiam perform substantially differently depending on the dataset. This is different from the observations of Chen & He (2021), who proposed SimSiam and found that it performs similarly to SimCLR derivatives. We next compare SimCLR-Rank with SimSiam at different levels of label scarcity.

Methodology. The dataset and methodology follows that of Section 4.1.1, except we let (1) SimCLR-Rank use the Gaussian augmentation with scale 1, and (2) SimSiam use the zeroing augmentation with probability 0.1. These augmentations generally perform well for their respective methods. Experiments illustrating the effect of each augmentation method are included in Appendix C.4.

Results. In Figure 4 we find that SimCLR-Rank performs better on MSLRWEB30K and Istella_S, while SimSiam performs much better on Yahoo Set1 when labels are scarce. To understand these results qualitatively, Figures 5 to 7 show the t-SNE projections of sub-sampled item embeddings learned by SimCLR-Rank, SimSiam, and the pre-final layer of a fully supervised model for Yahoo Set1, MSLRWEB30K, and Istella_S, respectively. We see that SimCLR-Rank and SimSiam also have qualitatively different embed-

dings: SimSiam’s embeddings are more clustered while SimCLR-Rank’s embeddings are less clustered. These observations raise the question of whether we can constructively combine them.

C.2. Unifying SimCLR-Rank and SimSiam

Combining models pretrained by SimCLR-Rank and SimSiam produces a model that has the strengths of both. To make a single recommendation for unsupervised LTR pretraining in LTR, we unify SimCLR-Rank and SimSiam into a single method “SimCLR-Rank + SimSiam”.

Methodology. To produce a “SimCLR-Rank + SimSiam” model, we finetune a linear layer over the embeddings from a finetuned SimCLR-Rank concatenated with the embeddings from a finetuned SimSiam model. In more detail, following the setting of Appendix B.3, we train (1) a SimCLR-Rank model with Gaussian augmentation and noise of scale 1.0, and (2) a SimSiam model with zeroing augmentation and corruption probability 0.1. After pretraining the SimCLR-Rank and SimSiam models, we finetune each of them on the labeled data and keep the checkpoints with the best validation NDCG. Then we combine these two models by concatenating the final embedding layers together—if each model’s penultimate embedding dimension is h ($h = 136$ for us), then the concatenated embedding is of dimension $2 \times h$. Then we pass this concatenated model through a batchnorm (to normalize the scales of each embedding to prevent scale issues), pass it through dropout layer of probability 0.7 (if the dataset is MSLRWEB30K or Istella_S) or 0.0 (if the dataset is Yahoo Set1), and map this vector (of dimension $2 \times h$) to a score of dimension 1 using a linear layer. We finetune only the final linear layer.

Results. In Figure 4 we find that “SimCLR-Rank + SimSiam” performs equal to or better than SimCLR-Rank and SimSiam across many different percentages of labeled query groups. While the combined model does not uniformly dominate across all labeled query group percentages, it is a promising first step towards a single unified method for unsupervised pretraining in ranking.

C.3. Finetuning strategies

Linear probing is a popular finetuning strategy in text and image domains (Chen & He, 2021; Chen et al., 2020; Peters et al., 2019; Kumar et al., 2022). We next compare it with the full finetuning strategy in tabular LTR, and show that full finetuning outperforms linear probing.

Methodology. We use the methodology from Section 3.1 to compare three finetuning strategies: linear probing (LP), full-finetuning (FF), and multilayer probing (MP) for SimSiam and SimCLR-Rank. MP tunes a 3-layer MLP head on top of the frozen pretrained embeddings. More experimental

Table 5: **Finetuning:** Comparison of fine tuning strategies on MSLRWEB30K dataset (LP = Linear Probing, MP = Multilayer Probing, FF = Full Finetuning). Full finetuning consistently performs best or near-best, while linear probing *and* multilayer probing perform poorly.

Method	SimSiam	SimCLR-Rank
<i>NDCG</i> (\uparrow)		
LP	0.2679 \pm 0.0007	0.3219 \pm 0.0224
MP	0.2764 \pm 0.0001	0.3890 \pm 0.0011
FF	0.3935 \pm 0.0034	0.3959 \pm 0.0022
<i>Outlier-NDCG</i> (\uparrow)		
LP	0.1803 \pm 0.0033	0.2304 \pm 0.0332
MP	0.1749 \pm 0.0023	0.2969 \pm 0.0009
FF	0.3149 \pm 0.0119	0.2892 \pm 0.0025

details are provided in Appendix B.3.

Results. Table 5 provides results on MSLRWEB30K dataset. We find that in terms of NDCG of both SimCLR-Rank and SimSiam, LP is the worst and FF is the best. Interestingly, MP comes as a close second in SimCLR-Rank, but it performs poorly on SimSiam. In Tables 6 and 7 we compare different probing strategies for all datasets when using SimSiam and SimCLR-Rank, respectively, as the pretraining method. The results are consistent across datasets. We thus recommend full fine-tuning as a stable and performant strategy.

To explain the above observations, we we can look again at the t-SNE projections of the embeddings generated by the SimSiam and SimCLR-Rank encoders and the pre-final layer of a fully supervised DL model trained a randomly sampled 1000 samples from the full training set (Figure 5). These plots offer the following explanations for the two phenomena. (1) The embeddings produced by the SimSiam/SimCLR-Rank are not sorted by relevance score in the projection space, unlike the fully supervised encoder. This suggests that a linear ranker cannot use the pretrained embeddings to predict the true relevance labels. (2) MP works better for SimCLR-Rank than SimSiam because SimCLR-Rank’s embeddings are more evenly spread out (less collapsed) than those of SimSiam. This allows an MLP trained on the SimCLR-Rank embeddings to distinguish the items with different relevance labels.

C.4. Data Augmentation Techniques

In Tables 8 and 9, we show the performance of each augmentation choice on SimCLR-Rank and SimSiam, respectively. We use the methodology described in Section 4.1.1, and have 0.2% of training QGs labeled. Overall, we find that no single augmentation works consistently better than others. However, on SimCLR-Rank, Gaussian augmentation

Table 6: SimSiam under different finetuning strategies on NDCG/Outlier-NDCG, averaged over 3 trials (LP = Linear Probing, MP = Multilayer Probing, FF = Full Finetuning). We find that linear probing *and* MLP probing perform extremely poorly (except in Yahoo Set1, where MLP probing performs well).

Method	MSLRWEB30K (\uparrow)	Yahoo Set1 (\uparrow)	Istella (\uparrow)
	NDCG (\uparrow)		
LP	0.2679 ± 0.0007	0.6089 ± 0.0032	0.3805 ± 0.0042
MP	0.2764 ± 0.0001	0.6137 ± 0.0022	0.4484 ± 0.0020
FF	0.3935 ± 0.0034	0.6107 ± 0.0035	0.5618 ± 0.0049
	Outlier-NDCG (\uparrow)		
LP	0.1803 ± 0.0033	0.5088 ± 0.002	0.4407 ± 0.0388
MP	0.1749 ± 0.0023	0.5157 ± 0.0080	0.5324 ± 0.0002
FF	0.3149 ± 0.0119	0.52 ± 0.0133	0.6348 ± 0.0164

Table 7: SimCLR-Rank under different finetuning strategies on NDCG/Outlier-NDCG, averaged over 3 trials (LP = Linear Probing, MP = Multilayer Probing, FF = Full Finetuning). Multilayer probing performs moderately worse than full finetuning, and linear probing performs much worse than all other strategies.

Method	MSLRWEB30K (\uparrow)	Yahoo Set1 (\uparrow)	Istella (\uparrow)
	NDCG (\uparrow)		
LP	0.3219 ± 0.0224	0.5202 ± 0.0093	0.4029 ± 0.0073
MP	0.3890 ± 0.0011	0.5942 ± 0.0016	0.5813 ± 0.0006
FF	0.3959 ± 0.0022	0.6022 ± 0.0013	0.5839 ± 0.0013
	Outlier-NDCG (\uparrow)		
LP	0.2304 ± 0.0332	0.3811 ± 0.0048	0.4717 ± 0.0166
MP	0.2969 ± 0.0009	0.4888 ± 0.0046	0.6369 ± 0.0045
FF	0.2892 ± 0.0025	0.5143 ± 0.0055	0.6352 ± 0.0140

seems to generally work well, and on SimSiam, zeroing with $p = 0.1$ appears to work consistently well across datasets.

D. Additional results

D.1. Runtime comparisons of pre-training strategies

In Table 10, we provide the runtime comparisons between SimSiam, SimCLR, and SimCLR-Rank (new method).

D.2. Simulated crowdsourced label setting

Here we provide the results on label scarcity for Outlier-NDCG in Figure 8. Result interpretation is given in Section 4.1.1.

D.3. Simulated Implicit Binary user feedback setting

In this part of the appendix we provide additional results on binary label generation as detailed in Section 4.1.2. In Tables 11 and 12 we give the dataset statistics when using $\tau_{\text{target}} = 4.25$ and $\tau_{\text{target}} = 5.1$, respectively. For these

two settings, the GBDT baseline does not contain semi-supervised runs as we found that semi-supervision tends to generally degrade GBDT performance; see Appendix D.4. The “No Pretrain” baseline uses the tabular ResNet architecture.

In Tables 14 and 15 we give the comparison between pre-trained and non-pretrained rankers following the methodology in Section 4.1.2 and find that GBDTs perform the best when binary labels are less sparse ($\tau_{\text{target}} = 4.25$) and pre-trained DL rankers perform the best when binary labels are sparser ($\tau_{\text{target}} = 5.1$). This strongly suggests that pretrained DL models outperform GBDTs when the labels arise from binary user feedback that is moderately uncommon or rare (as in online shopping, recommendations, or search).

D.4. Semi-supervised GBDT

Here we provide a more detailed comparison between Supervised GBDTs, GBDT + pseudolabeling, and GBDT + PCA + pseudolabeling in Figure 9 for the simulated crowdsourcing of labels setting and in Table 16 for the simulated

Table 8: We compare different augmentation strategies for SimCLR-Rank. The methodology used is the one in Section 4.1.1, with 0.2% of training QGs labeled.

Augmentation	MSLRWEB30K	Yahoo Set1	Istella
NDCG (\uparrow)			
Zeroing (p=0.1)	0.3830 \pm 0.0007	0.6022 \pm 0.0013	0.5820 \pm 0.0024
Zeroing (p=0.7)	0.3737 \pm 0.0025	0.5998 \pm 0.0053	0.5646 \pm 0.0048
Gaussian noise (scale=1.0)	0.3959 \pm 0.0022	0.5998 \pm 0.0026	0.5839 \pm 0.0013
Gaussian noise (scale=2.0)	0.3907 \pm 0.0025	0.5953 \pm 0.0043	0.5809 \pm 0.0010
Outlier-NDCG (\uparrow)			
Zeroing (p=0.1)	0.2782 \pm 0.0036	0.5143 \pm 0.0055	0.6408 \pm 0.0063
Zeroing (p=0.7)	0.2730 \pm 0.0054	0.5062 \pm 0.0070	0.6141 \pm 0.0064
Gaussian noise (scale=1.0)	0.2892 \pm 0.0025	0.4963 \pm 0.0024	0.6352 \pm 0.0140
Gaussian noise (scale=2.0)	0.2886 \pm 0.0096	0.4875 \pm 0.0054	0.6327 \pm 0.0191

Table 9: We compare different augmentation strategies for SimSiam. The methodology used is the one in Section 4.1.1, with 0.2% of training QGs labeled.

Augmentation	MSLRWEB30K	Yahoo Set1	Istella
NDCG (\uparrow)			
Zeroing (p=0.1)	0.3935 \pm 0.0034	0.6107 \pm 0.0035	0.5618 \pm 0.0049
Zeroing (p=0.7)	0.3911 \pm 0.0003	0.6076 \pm 0.0072	0.5660 \pm 0.0047
Gaussian noise (scale=1.0)	0.3782 \pm 0.0093	0.6010 \pm 0.0026	0.5612 \pm 0.0060
Gaussian noise (scale=2.0)	0.3860 \pm 0.0011	0.6100 \pm 0.0089	0.5587 \pm 0.0047
Outlier-NDCG (\uparrow)			
Zeroing (p=0.1)	0.3149 \pm 0.0119	0.5200 \pm 0.0133	0.6348 \pm 0.0164
Zeroing (p=0.7)	0.3002 \pm 0.0060	0.5081 \pm 0.0097	0.6201 \pm 0.0104
Gaussian noise (scale=1.0)	0.2585 \pm 0.0371	0.4893 \pm 0.0084	0.6145 \pm 0.0160
Gaussian noise (scale=2.0)	0.2929 \pm 0.0021	0.5163 \pm 0.0101	0.5905 \pm 0.0146

implicit binary user feedback setting. We follow the experimental methodology in Appendix B.3. We find that for the most part, semi-supervised learning does not benefit GBDT rankers.

Table 10: Seconds per epoch comparison between pretraining methods. Average over 3 trials. The encoder we use for pretraining is the tabular ResNet (Gorishniy et al., 2021) with the final linear layer taken out.

Method	Complexity	Seconds per epoch		
		MSLRWEB30K	Yahoo Set1	Istella
SimCLR	$O(B^2L^2)$	102.81 ± 1.14	14.45 ± 0.04	69.26 ± 0.16
SimCLR-Rank (new)	$O(BL^2)$	9.4700 ± 0.02	1.970 ± 0.01	5.950 ± 0.04
SimSiam	$O(BL)$	0.8700 ± 0.01	0.130 ± 0.00	0.730 ± 0.00

Table 11: Train and validation dataset statistics on MSLRWEB30K, Yahoo Set1, and Istella_S after binary label generation with $\tau_{\text{target}} = 4.25$. The test set is left untouched. Labeled QGs are those query groups that have at least one item with $y = 1$.

Dataset	# query groups		# labeled QGs		# items per QG		# positives per QG	
	train	val	train	val	train	val	train	val
MSLRWEB30K	18151	6072	15.9%	16.2%	124.35	122.38	1.4%	1.4%
Yahoo Set1	14477	2147	11.1%	10.8%	30.02	30.32	6.2%	6.0%
Istella_S	19200	7202	49.0%	49.5%	106.36	94.98	1.4%	1.6%

Table 12: Train and validation dataset statistics on MSLRWEB30K, Yahoo Set1, and Istella_S after binary label generation with $\tau_{\text{target}} = 5.1$. The test set is left untouched. Labeled QGs are those query groups that have at least one item with $y = 1$.

Dataset	# query groups		# labeled QGs		# items per QG		# positives per QG	
	train	val	train	val	train	val	train	val
MSLRWEB30K	18151	6072	1.1%	1.3%	124.35	122.38	0.8%	0.7%
Yahoo Set1	14477	2147	0.6%	0.6%	30.02	30.32	5.4%	5.1%
Istella_S	19200	7202	3.0%	2.7%	106.36	94.98	1.1%	1.3%

Table 13: Train and validation dataset statistics on MSLRWEB30K, Yahoo Set1, and Istella_S after binary label generation with $\tau_{\text{target}} = 4.5$. The test set is left untouched. Labeled QGs are those query groups that have at least one item with $y = 1$.

Dataset	# query groups		# labeled QGs		# items per QG		# positives per QG	
	train	val	train	val	train	val	train	val
MSLRWEB30K	18151	6072	8.9%	8.9%	124.35	122.38	1.1%	1.1%
Yahoo Set1	14477	2147	5.3%	5.4%	30.02	30.32	5.9%	5.4%
Istella_S	19200	7202	25.2%	25.2%	106.36	94.98	1.3%	1.3%

Table 14: We compare pretrained models to non-pretrained models in the binary label setting with $\tau_{\text{target}} = 4.25$ (Section 4.1.2) on NDCG averaged over three trials. We follow the methodology in Section 4.1.2.

Method	MSLRWEB30K (\uparrow)	Yahoo Set1 (\uparrow)	Istella (\uparrow)
	NDCG (\uparrow)		
GBDT	0.3616 ± 0.0000	0.6233 ± 0.0000	0.6251 ± 0.0000
No Pretrain DL	0.3552 ± 0.0015	0.6297 ± 0.0005	0.6147 ± 0.0006
Pretrained DL	0.3602 ± 0.0007	0.6272 ± 0.0024	0.6243 ± 0.0007
	Outlier-NDCG (\uparrow)		
GBDT	0.2792 ± 0.0000	0.5352 ± 0.0000	0.7393 ± 0.0000
No Pretrain DL	0.3142 ± 0.0063	0.5389 ± 0.0036	0.6631 ± 0.0102
Pretrained DL	0.2923 ± 0.0079	0.5471 ± 0.0050	0.6945 ± 0.0041

Table 15: We compare pretrained models to non-pretrained models in the binary label setting with $\tau_{\text{target}} = 5.1$ (Section 4.1.2) on NDCG averaged over three trials. We follow the methodology in Section 4.1.2. * indicates metrics on which pretrained rankers outperform non-pretrained rankers significantly via a $p < 0.05$ t-test.

Method	MSLRWEB30K	Yahoo Set1	Istella
	NDCG (\uparrow)		
GBDT	0.2844 ± 0.0000	0.5782 ± 0.0000	0.5638 ± 0.0000
No Pretrain DL	0.3432 ± 0.0065	0.5703 ± 0.0246	0.5722 ± 0.0080
Pretrained DL	$0.3564 \pm 0.0054^*$	$0.6031 \pm 0.0064^*$	$0.599 \pm 0.0014^*$
	Outlier-NDCG (\uparrow)		
GBDT	0.2416 ± 0.0000	0.5011 ± 0.0000	0.6640 ± 0.0000
No Pretrain DL	0.2441 ± 0.0089	0.4593 ± 0.0317	0.6334 ± 0.0054
Pretrained DL	0.2553 ± 0.0161	$0.54 \pm 0.0055^*$	$0.6730 \pm 0.0090^*$

Table 16: We compare GBDT models in the binary label setting with $\tau_{\text{target}} = 4.5$ (Section 4.1.2) on NDCG averaged over three trials. We follow the methodology in Section 4.1.2.

Method	MSLRWEB30K	Yahoo Set1	Istella
	NDCG (\uparrow)		
Supervised GBDT	0.3335 ± 0.0000	0.6168 ± 0.0000	0.6024 ± 0.0000
GBDT + PCA	0.3105 ± 0.0236	0.6141 ± 0.0022	0.5508 ± 0.0410
GBDT + pseudolabeling	0.3467 ± 0.0011	0.5868 ± 0.0012	0.5178 ± 0.0068
GBDT + PCA + pseudolabeling	0.3241 ± 0.0270	0.5944 ± 0.0012	0.5315 ± 0.0120

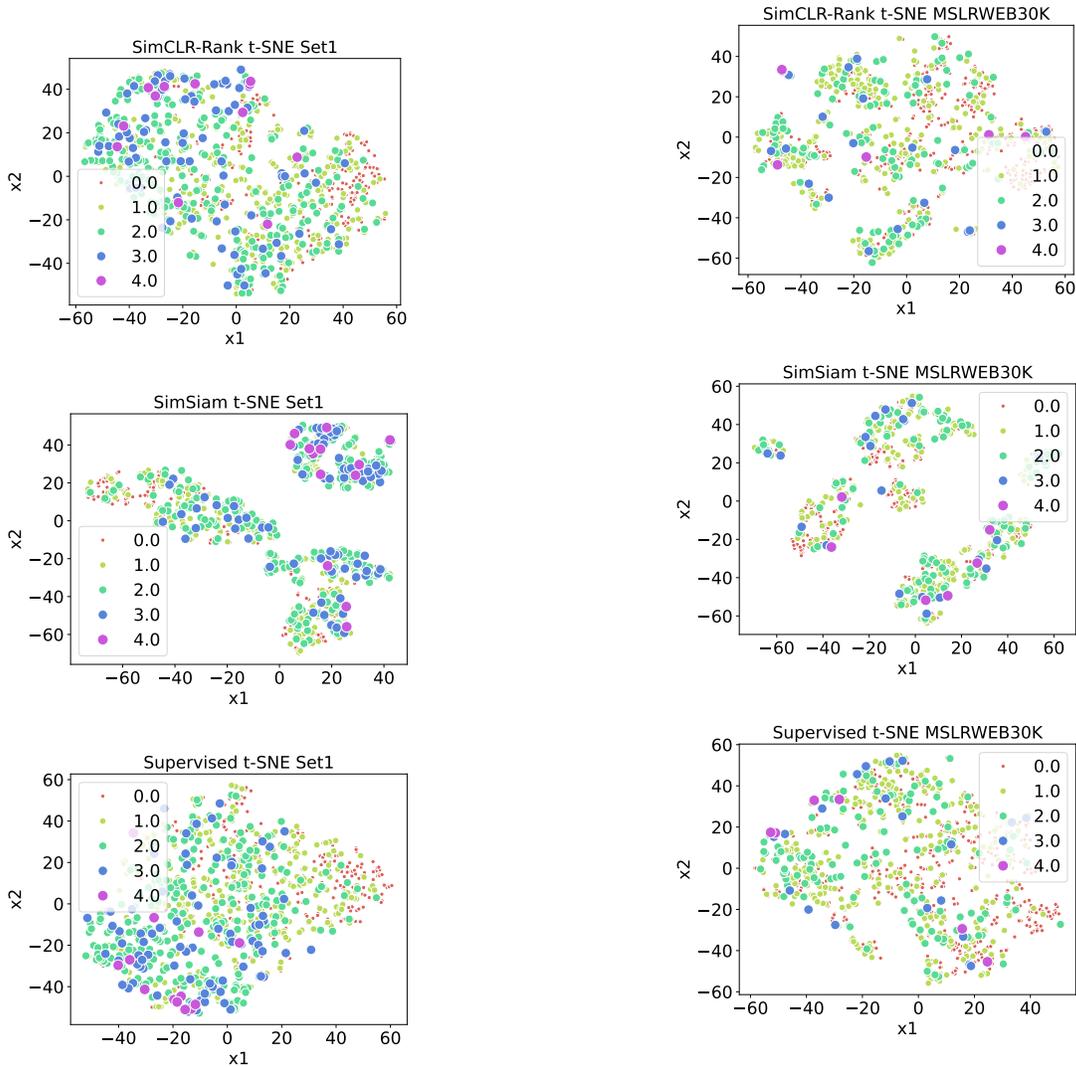


Figure 5: t-SNE plots of SimCLR-Rank, SimSiam, and fully supervised embeddings on Yahoo Set1 datasets. Relevance label is indicated by marker size/color. We see that SimCLR-Rank and SimSiam embeddings are qualitatively very different. Both SimCLR-Rank and SimSiam embeddings do not sort the data as well as fully supervised embeddings, which most effectively groups data points with relevance label 4 together.

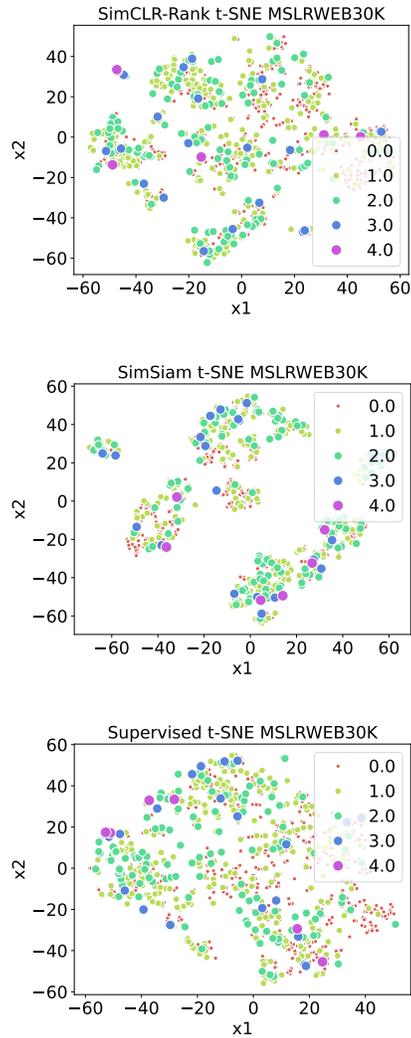


Figure 6: We plot the t-SNE plots of embeddings produced by three different encoders for the MSLRWEB30K dataset: (1) pretrained by SimCLR-Rank, (2) pretrained by SimSiam, (3) trained via supervised training on the entire training set on roughly 1000 samples from MSLR. Marker size and color indicates relevance. We find (1) SimCLR-Rank/SimSiam cluster different relevances effectively but do not order them as well as the supervised encoder. (2) SimCLR-Rank produces more spread-out embeddings with less defined clusters than SimSiam.

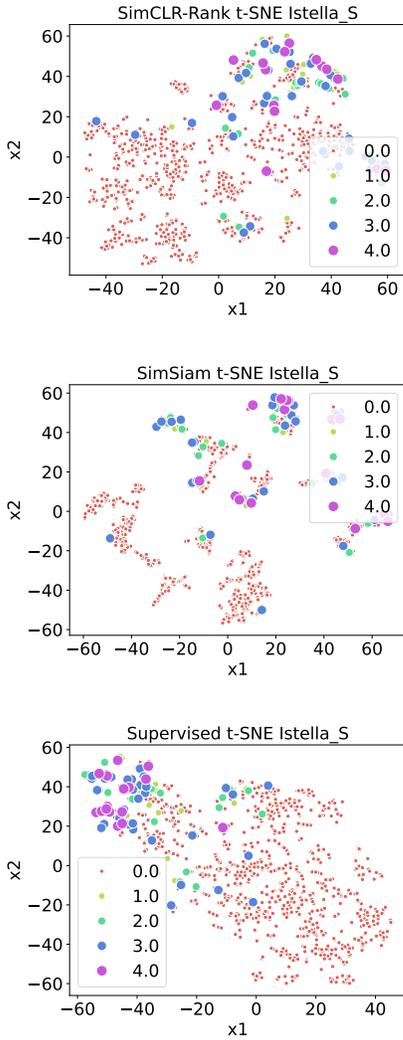


Figure 7: We plot the t-SNE plots of embeddings produced by three encoders on the Istella_S dataset: (1) pretrained by SimCLR-Rank, (2) pretrained by SimSiam, (3) trained via supervised training on the entire training set on roughly 1000 samples from Istella_S. Marker size and color indicates relevance. We find (1) SimCLR-Rank/SimSiam cluster different relevances effectively but do not order them as well as the supervised encoder. (2) SimCLR-Rank produces more spread-out embeddings with less defined clusters than SimSiam.

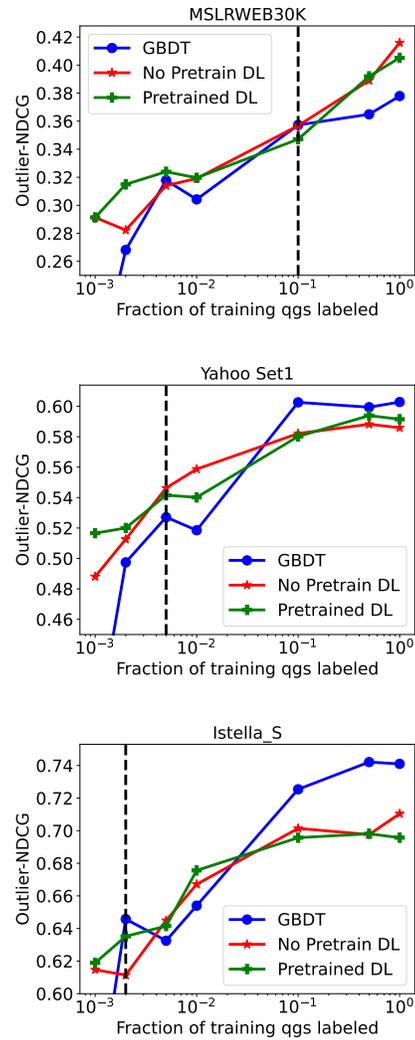


Figure 8: **Simulated crowdsourcing of labels:** We compare Outlier-NDCG (\uparrow) pretrained rankers, non-pretrained DL rankers, and GBDT rankers as we change the percentage of training query groups that are labeled. To the left of the black dotted line, pretrained rankers perform the best. Points are averages over three trials. To the left of the black dotted vertical line, pretrained rankers are (1) significantly better on outliers than GBDTs at the $p = 0.05$ level using a one-sided t -test, and (2) on average better on outliers than all other non-pretrained methods.

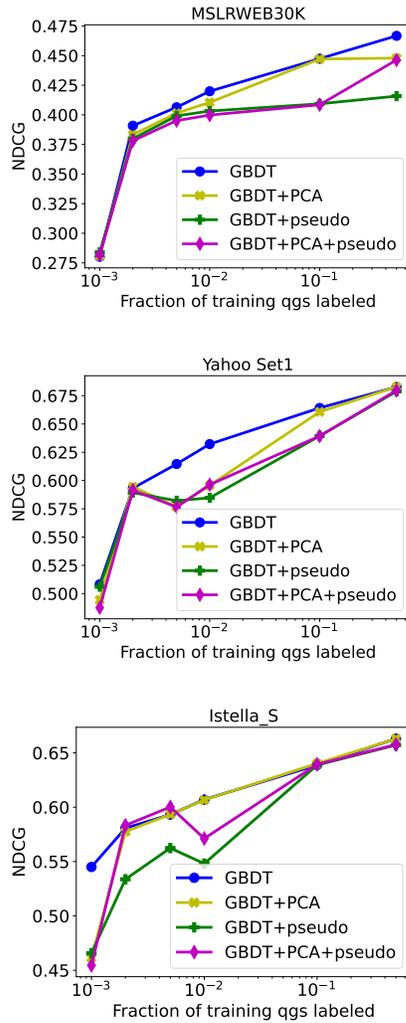


Figure 9: A comparison of Supervised GBDT, GBDT + PCA, GBDT + pseudolabeling, and GBDT + PCA + pseudolabeling. We find that semi-supervised learning may sometimes benefit GBDT rankers, though it is not very consistently beneficial.