

Learning List-Level Domain-Invariant Representations for Ranking

Ruicheng Xian¹ Honglei Zhuang² Zhen Qin²
Hamed Zamani³ Jing Lu² Ji Ma² Kai Hui²
Han Zhao¹ Xuanhui Wang² Michael Bendersky²

¹University of Illinois Urbana-Champaign

²Google Research

³University of Massachusetts Amherst

Learning List-Level Domain-Invariant

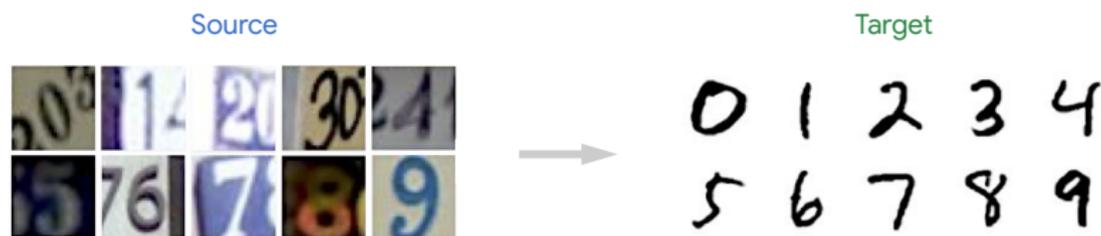
Contributions

Revisit domain adaptation on ranking problems via invariant representation learning.

Whereas prior implementations perform *item-level alignment*, we

- propose *list-level alignment*;
- establish a domain adaptation generalization bound for ranking based on list-level alignment, and
- demonstrate the its empirical benefits.

Domain Adaptation and Invariant Representations



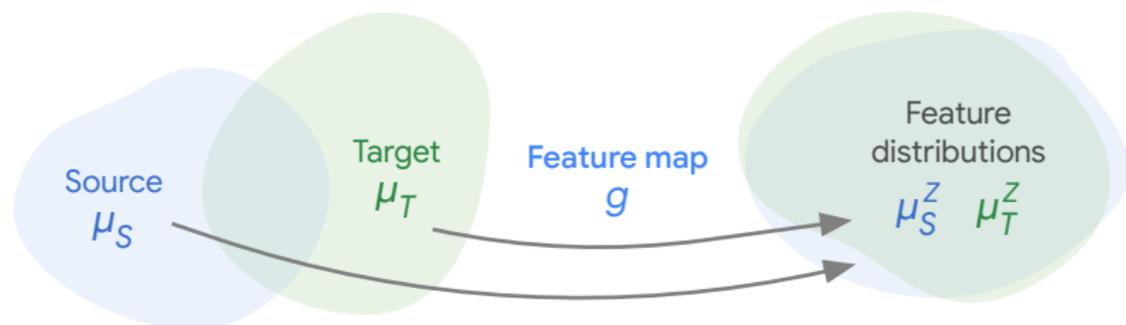
Problem Setup for Domain Adaptation

A (low-resource) target domain μ_T , and a source domain μ_S .

Goal. Train a good model for μ_T using available resources.

Example. In *unsupervised* domain adaptation, have labeled data from source domain, but only unlabeled data from target domain.

Domain Adaptation and Invariant Representations

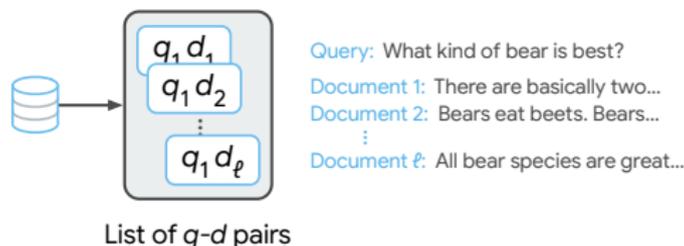


Invariant Representation Learning

1. Learn a mapping $g : \mathcal{X} \rightarrow \mathcal{Z}$ that *matches and aligns* the source and target data distributions on the feature space \mathcal{Z} .
2. Train the model on the learned (and *transferrable*) features.

Example. In *unsupervised* domain adaptation, the feature map is learned on unlabeled data from source and target, and the model is trained on source labeled data.

Learning to Rank

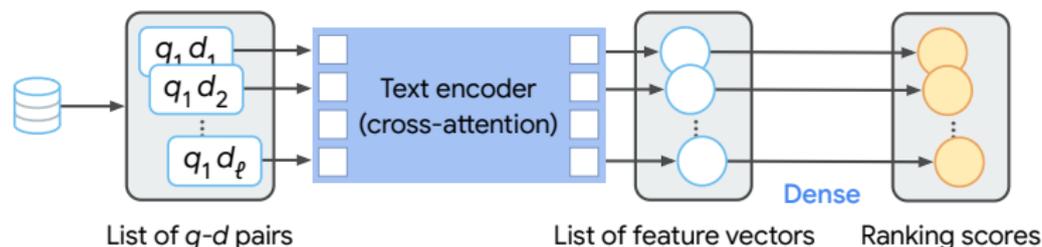


A ranking problem is given by joint distribution μ over length- ℓ lists of items (X_1, \dots, X_ℓ) and scores (Y_1, \dots, Y_ℓ) .

Measure model performance by ranking metrics, e.g., MRR, NDCG.

Goal. Train a ranking model f that ranks the items in agreement with the descending order of the scores.

Learning to Rank



A ranking problem is given by joint distribution μ over length- ℓ lists of items (X_1, \dots, X_ℓ) and scores (Y_1, \dots, Y_ℓ) .

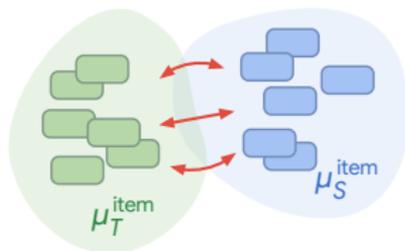
Measure model performance by ranking metrics, e.g., MRR, NDCG.

Goal'. Train a scoring model f that scores the items in agreement with the descending order of the scores.

Feature Space. Model computes a list of feature vectors, $(Z_1, \dots, Z_\ell) \in \mathbb{R}^{\ell \times k}$, where $Z_i \in \mathbb{R}^k$ corresponds to item i .

Item-Level Alignment¹

Feature Space. Model computes a list of feature vectors, $(Z_1, \dots, Z_\ell) \in \mathbb{R}^{\ell \times k}$, where $Z_i \in \mathbb{R}^k$ corresponds to item i .



In item-level alignment, distributions of feature vectors aggregated from all lists are aligned, i.e., $\mu_S^{Z, \text{item}} \approx \mu_T^{Z, \text{item}}$,

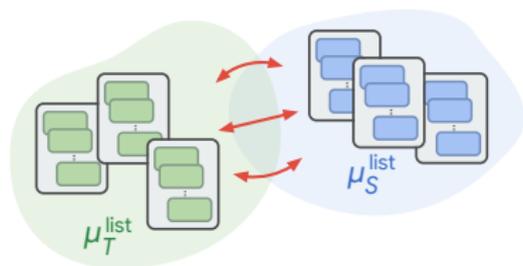
$$\text{supp}(\mu^{Z, \text{item}}) \subseteq \mathbb{R}^k, \quad \mu^{Z, \text{item}}(v) = \mathbb{P}(v \in (Z_1, \dots, Z_\ell)).$$

List structure on data is ignored.

¹Cohen et al., 2018; Tran et al., 2019; Xin et al., 2022.

List-Level Alignment

Feature Space. Model computes a list of feature vectors, $(Z_1, \dots, Z_\ell) \in \mathbb{R}^{\ell \times k}$, where $Z_i \in \mathbb{R}^k$ corresponds to item i .



In list-level alignment, distributions of lists of feature vectors are aligned, i.e., $\mu_S^{Z, \text{list}} \approx \mu_T^{Z, \text{list}}$,

$$\text{supp}(\mu^{Z, \text{list}}) \subseteq \mathbb{R}^{\ell \times k}, \quad \mu^{Z, \text{list}}(z) = \mathbb{P}(z = (Z_1, \dots, Z_\ell)).$$

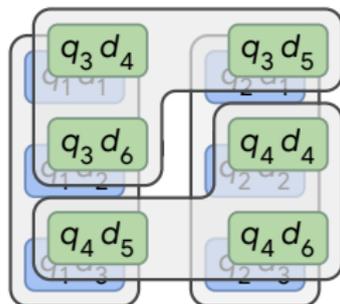
Preserves list structure on data.

List-Level vs. Item-Level Alignment

List-level alignment is stronger than item-level alignment.

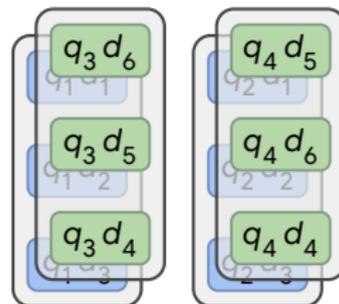
Example. Source and target both have 2 lists of length-3.

Representation 1



Aligned at item-level
but not at list-level

Representation 2



Aligned at item-level
and at list-level

Benefits of List-Level Alignment

1. We establish a domain adaptation generalization bound for ranking based on list-level alignment.

Theorem (Instantiated for mean reciprocal rank)

On ranking problems, under Lipschitz assumptions on model and scores, let $g : \mathcal{X} \rightarrow \mathcal{Z}$, then for all scoring models $h : \mathcal{Z} \rightarrow \mathbb{R}^\ell$,

$$\text{MRR}_T(h \circ g) \geq \text{MRR}_S(h \circ g) - \Theta(\ell) W_1(\mu_S^{Z,\text{list}}, \mu_T^{Z,\text{list}}) - \lambda_g^*,$$

where $\lambda_g^ = \min_{h'} (1 - \text{MRR}_S(h' \circ g) + 1 - \text{MRR}_T(h' \circ g))$. Note that $\text{MRR} \in (0, 1]$.*

Can also be instantiated for other ranking metrics, e.g., NDCG.

Benefits of List-Level Alignment

- List-level alignment achieves better unsupervised domain adaptation performance vs. item-level alignment and zero-shot transfer.

Target domain	Method	MAP	MRR@10	NDCG@10
Robust04	BM25	0.2282	0.6801	0.4088
	Zero-shot	0.2759	0.7977	0.5340
	Item-level alignment	0.2822	0.8037	0.5396
	List-level alignment	0.2901	0.8234	0.5573
TREC-COVID	BM25	0.2485	0.8396	0.6559
	Zero-shot	0.3083	0.9217	0.8200
	Item-level alignment	0.3087	0.9080	0.8142
	List-level alignment	0.3187	0.9335	0.8412
BioASQ	BM25	0.4088	0.5612	0.4653
	Zero-shot	0.5008	0.6465	0.5542
	Item-level alignment	0.4781	0.6383	0.5343
	List-level alignment	0.5191	0.6666	0.5714

Source domain is MS MARCO.

References

- Daniel Cohen, Bhaskar Mitra, Katja Hofmann, and W. Bruce Croft. “Cross Domain Regularization for Neural Ranking Models Using Adversarial Learning”. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 2018, pp. 1025–1028.
- Brandon Tran, Maryam Karimzadehgan, Rama Kumar Pasumarthi, Michael Bendersky, and Donald Metzler. “Domain Adaptation for Enterprise Email Search”. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2019, pp. 25–34.
- Ji Xin, Chenyan Xiong, Ashwin Srinivasan, Ankita Sharma, Damien Jose, and Paul Bennett. “Zero-Shot Dense Retrieval with Momentum Adversarial Domain Invariant Representations”. In *Findings of the Association for Computational Linguistics: ACL 2022*. 2022, pp. 4008–4020.