

Towards Generalizable Semantic Product Search by Text Similarity Pre-training on Search Click Logs

Zheng Liu, Wei Zhang, Yan Chen, Weiyi Sun, Michael Du, Benjamin Schroeder

Wayfair LLC

4 Copley Place, Boston, MA USA 02116

{zliu2, wzhang5, ychen4, wsun1, tdu1, beschroeder}@wayfair.com

Abstract

Recently, semantic search has been successfully applied to e-commerce product search and the learned semantic space(s) for query and product encoding are expected to generalize to unseen queries or products. Yet, whether generalization can conveniently emerge has not been thoroughly studied in the domain thus far. In this paper, we examine several general-domain and domain-specific pre-trained Roberta variants and discover that general-domain fine-tuning does not help generalization, which aligns with the discovery of prior art. Proper domain-specific fine-tuning with clickstream data can lead to better model generalization, based on a bucketed analysis of a publicly available manual annotated query-product pair data.

1 Introduction

Semantic Search has been widely studied recently as a promising competitor to token-based inverted index search that had dominated information retrieval (IR) and question answering (QA) for decades. However, in e-commerce product search, the quality of the learned embedding space is not well understood, especially on the quality of generalization to unseen queries, unseen products, or even unseen query-product pairs. The generalization is particularly interesting and important because in Wayfair 80% of queries are new in the monthly basis, and new products are added everyday. In this paper, we will investigate the generalization issue by seeking answers to the two following questions:

1. How well does the learned space generalize to unseen queries/products/query-product combinations?
2. Are we able to improve generalization by pre-training, and how?

Arguably, some pre-trained models, especially the ones pre-trained with domain-specific click-

stream data, may outperform the others on generalization. To validate the hypothesis, we train two-tower semantic search models using three different categories of base models for comparison: not pre-trained (Roberta-base (Liu et al., 2019), Roberta-large), general-domain pre-trained with text similarity (SimCSE (Gao et al., 2021)) and domain-specific pre-trained with text similarity (SimCSE that are continuously pre-trained with clickstream data).

To observe how the models perform, we take a publicly available product search dataset (Chen et al., 2022) and partition the data into buckets according to whether the query/product/query-product pairs have been visited during training. By comparing the performance of the above models, we observe that query representations may have been insufficiently learned during two-tower semantic search fine-tuning as opposed to product representations, and the model variant which is pre-trained on query-query text similarity has a visible impact on models' generalization power.

This paper is organized as follows: 1) we first introduce related work, then 2) explain the proposed pre-training approach, and finally 3) discuss results and analysis.

2 Related Work

Semantic search in product search Recently, semantic search has been explored in e-commerce product search (Nigam et al., 2019; Huang et al., 2020; Zhang et al., 2020; Li et al., 2021), in which queries and products are embedded as vectors through representation learning and search is then conducted in the vector space. These implementations generally follow a two-tower approach: one group of stacked neural network layers, usually referred to as "tower", processes the query, and another tower processes the product. This approach is favored in production due to its ability of decoupling query and product embedding process, so

that products can be embedded offline and search can be accelerated in real-time services. In search systems, the semantic search is usually used as candidate generators focusing on core-relevance, and often followed by dedicated ranking layers that capture other e-commerce attributes such as price, popularity, quality, etc. The two tower model of this work is described in Section 3.5.

Generalization for two-tower neural models

Deep Passage retrieval has become a popular approach recently for IR. In Question Answering, Lewis et al. (2021) observed that when using the two-tower approach for relevant passage retrieval stage, deep learning models tend to memorize training instances and perform substantially worse on unseen questions or unseen answers. They argued the generalization on questions is the major issue, which is similar to our observation in product search. Mao et al. (2021) tried to address the issue by explicitly retrieving question-related texts as query expansions and showed improvements on model generalization to unseen question/answers, yet we take the route of query similarity learning by using automatically mined query pairs from product co-click behaviors.

Pre-training for Neural Retrieval Chang et al. (2020) pre-trained neural retrieval-based QA models with carefully designed synthetic tasks, followed by Guu et al. (2020) with a knowledge-augmented training task and Sachan et al. (2021) with combining ICT and masked salient spans training. Oğuz et al. (2021) instead enhanced training with domain-matching Reddit post comment pairs to pre-train a dense retrieval model and also observed positive influences to open domain QA. Differently, our approach explores the clickstream graph structure instead of mining additional data.

Text Representations Learning Our approach closely relates to the graph node embeddings within texts (Mikolov et al., 2013) or the linked graphs (Hamilton et al., 2017) rather than node embeddings that involve node relation types (Bordes et al., 2013) or graphical neural networks that aggregate neighbouring nodes in a parametric fashion (Li et al., 2016; Kipf and Welling, 2016). Recently, noisy contrastive learning has enjoyed significant success (Gao et al., 2021) for regularizing text encoders, and shows that such encoders can effectively serve as a starting point for downstream task fine-tuning. In retrieval tasks, Gao and Callan

(2021) pre-trained text encoders with synthetic passage similarity tasks using noisy contrastive loss, assuming that in Wikipedia two adjacent passages in the same article are similar in semantics rather than those from two different articles. In contrast, in product search, we assume two queries leading to the same product browse, or two products purchased under the same query, are similar.

3 Approach

This section describes how models are pre-trained, fine-tuned, and evaluated on clickstream data. Section 3.1 sets up preliminaries and math notations for the product search problem; Section 3.2, 3.3 and 3.4 describe three different pre-training approaches; Section 3.5 discusses how the model is fine-tuned and evaluated in the product search problem.

3.1 Preliminaries

A click graph is denoted as $\mathcal{G} = \{V, E\}$ where V is the set of nodes either being a query q or a product p which is the concatenation of title, brand name, and class name¹. An edge E connecting q and p suggests a customer showed interest to product p when searching with query q .

The core of semantic search is to learn query and product representations. Intuitively, a good representation that generalizes may put similar concepts (p or q) closer in the semantic space, while dissimilar ones further apart.

There are three node type combinations (p, p) , (q, q) , and (p, q) , the last of which is the goal of semantic search. Arguably, learning solely on (p, q) may cause generalization issue when (p, p) and (q, q) relations can not be implicitly learned from (p, q) objective. To verify if this hypothesis holds in our semantic search problem, we design three pre-training tasks to warm up Transformers: the (q, q) pre-training task, the (p, p) pre-training task, and the (q, q) and (p, p) pre-training task. After pre-training, we fine-tune the model on the (p, q) learning objective. In the remaining of this section, we describe the three pre-training and the fine-tuning approaches.

3.2 Query Similarity Learning

qq Model If two queries q_1 and q_2 link to the same product p in graph \mathcal{G} , we regard (q_1, q_2) as

¹A product can have multiple class names with hierarchical structure (e.g. class “sofa” belongs to class “living room furniture”). Here we used the high level product class name (e.g. “living room furniture”).

first query	second query	relation	score
girls study desk	outdoor shed	Entirely different intents	0
bathroom window curtains short	curtains bedroom	Similar (not entail/co-intent)	1
52 inch ceiling fan with light	Ceiling fans with lights	Entailment	2
Fishing storage	Fishing tackle storage	Same intent	3

Table 1: Query similarity manual annotation samples. Note that first and second query is order-sensitive. However, as similarity strength, we disregard directionality.

one co-clicked query pair. After collecting the pairs, we use them to fine-tune a SimCSE-large (Gao et al., 2021) model to obtain a model we call **qq** model.

Pair Sampling We follow the 3 steps below to mine pairs: 1) sample q_1 by $P(q)$, the query distribution obtained from data; 2) for q_1 , we collect top-K popular products p_1, \dots, p_K by frequency, and uniformly sample one product p from them ², and 3) at last, for p , we collect all its related queries and sample one query as q_2 uniformly. We repeat this process N times to obtain N pairs where N is set to 1 million.

Training We train the model with a noisy contrastive learning objective as in (Gao et al., 2021) with default settings. We train the model with 3 epochs. We fine-tune on top of the unsupervised SimCSE-large model and observe the additive effect of open-domain and domain-specific learning.

Similarity Evaluation To evaluate the query representation, we invite 3 domain experts to annotate the degree of query pairs into 4 relation types: not related (0), similar but not entail (1), entailment (2) and identical (3), where the number represents the semantic strengths, the higher the more related they are. A total of 1056 random query pairs are extracted from click sessions and the relation types are evenly distributed. The inter-annotator agreement was 90% after 2 rounds of annotation and curation. A sample of the annotation data is shown in Table 1.

We calculate vector cosine between two queries as similarity score, and compare it to the ground-truth by Spearman Rank Correlation with significant tests. The result is shown in Table 2. As we can see, the model fine-tuned with co-click data achieves the best performance. Surprisingly, “Ours-unsup” learning noisy contrastive loss by masking

²“Uniformly” as a reasonably simple yet arguably effective approach to mitigate exposure bias without complex user behavior modeling

Model	SP (P-value)
Roberta-large	0.287 (1.8e-21)
SimCSE-large-unsup	0.467 (2.9 e-58)
SimCSE-large-sup	0.426 (7.9e-48)
Ours-unsup	0.458 (2.2e-48)
Ours-sup (co-click)	0.546 (5.8e-82)

Table 2: Comparison of different Models’s Spearman Rank Correlation on representation evaluation data. SimCSE-large-unsup is the representation learned *unsupervisedly* with noisy contrastive loss; SimCSE-large-sup is subsequently learned from *supervised* text entailment data. In our settings, the definition of unsup and sup changed a bit: Ours-unsup is the model learned with unsupervised objective (masked language modeling) with noisy contrastive loss, where Ours-sup is learned from query co-click data described in section 3.2.

a single query twice and learning their similarities (similar to SimCSE-unsup with open-domain texts), does not lead to a better performance on query pairs, suggesting that the key of query similarity learning is to capture rich domain-specific text similarity knowledge rather than text robustness and regularity in semantic space. This message encourages us to further extend it to product similarity learning (this time, without manual evaluation).

3.3 Product Similarity Learning

pp Model The product representation learning is performed similarly to the query representation learning. We start from SimCSE-large-unsup model, and fine-tune (p, p) pairs that are obtained from the process similar to the sampling approach in section 3.2 with the role of query and product swapped. The created 1 million pairs are used to fine-tune the SimCSE model and obtain a model dubbed **pp** model. Different from **qq** model, we did not obtain manual annotation data for **pp**, thus we let the model train 2 epochs until convergence.

3.4 A Combined Model

qq+pp Model In this approach, we sequentially pre-trained the model with both (q, q) data and (p, p) data, to examine the additive effect where a model is fine-tuned with both query similarities and product similarities. We take the **qq** model and continue to pretrain it with the above (p, p) data until convergence, and obtain a model called **qq+pp**. Indeed, we may risk catastrophic forgetting since arguably the (p, p) pairs may overwrite the captured (q, q) information. Although we have observed such forgetting to happen, its effect is mild at best.

3.5 Fine-tuning for Semantic Search

pq Learning After pre-training, the semantic search model was fine-tuned on product search tasks: given a search query, the task is to retrieve a list of relevant products from the product catalog. Specifically, both queries and products were represented as vectors in the same latent space, and products were retrieved according to cosine similarities w.r.t. the queries.

As shown in figure 1, we adopted a two tower approach to fine-tune the semantic search model: a user tower embeds search queries into vectors, while a product tower embeds product information (e.g. product name, product brand name, product class name) into vectors. The two towers share the same tokenizer and transformer model. We took the last hidden layer’s CLS token representation and L2-normalized it as the final output embedding vector. Then, the relevance score was calculated as the cosine similarity between search query embedding and product embedding vectors.

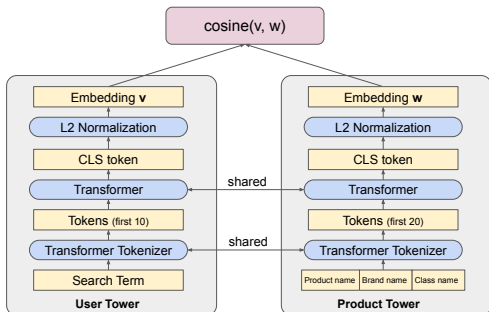


Figure 1: Structure of semantic search models on product search tasks

To train the model, we adopted the contrastive framework and utilized the cross-entropy objec-

tive with in-batch negatives (Chen et al., 2020). In each mini-batch, we sampled n semantically relevant (query, product) pairs $D = \{(q_i, p_i)\}_{i=1}^n$, and calculated the training objective for the i 'th pair (q_i, p_i) as:

$$l_i = -\frac{\cos(f(q_i), g(p_i))}{\tau} + \log \sum_{j=1}^n e^{\frac{\cos(f(q_i), g(p_j))}{\tau}} \quad (1)$$

f, g represent the user and product tower respectively. τ is a temperature hyperparameter. All layers of Transformer models are fine-tuned in the product search task.

Data In the experiment, (q, p) pairs are collected from click graph \mathcal{G} . As show in table 3, each search query and its resulted add-to-cart products (if any) are constructed as semantically relevant (query, product) pairs. The training dataset contains 31M pairs from 2019 and 2020 records of a e-commerce search engine, and the dev dataset contains 100k pairs randomly sampled from 2021. We evaluated model’s performance on the WANDS dataset (Chen et al., 2022), which is a human annotated dataset for evaluating product search relevance. It contains (query, product) pairs with relevance labels annotated by humans in three different levels: relevant, partially relevant, and irrelevant. Table 4 shows details about the training, dev and test dataset.

Training and Evaluation 6 different Roberta variations are fine-tuned in this experiment. They are Roberta base, Roberta large (Liu et al., 2019), unsupervised SimCSE based on Roberta large (Gao et al., 2021), and three in-house pretrained models based on the SimCSE model using query-query, product-product, and query-query plus product-product similarity data respectively.

All of the models have the same training setup. For each model, we ran 4 epochs on the training dataset with mini-batch size of 256. Model checkpoints were saved every 6000 mini-batches, for the purpose of model selection. Adam optimizer with learning rate of $1e^{-5}$ was used to train the model. The temperature hyperparameter was set as 0.07.

In model selection, for each model we select the best performing checkpoint on the dev dataset to be the final trained model. The evaluation metric is the top-k retrieval accuracy, defined as the ratio of (query, product) pairs for which the top-k retrieved products of the query contains the corresponding product. The retrieval space for dev dataset is the unique products in dev dataset.

query	product
rose gold wing chair	alyka adonis wingback chair, red barrel studior, accent chairs
led lights strip lights	doynton strip light, the holiday aisle, under cabinet lighting
sun shelter	messina aluminum patio gazebo, sojag, canopies & gazebos

Table 3: Example of semantically relevant (query, product) pairs

	train	dev	test
total pairs	32M	100K	234K
unique queries	9.5M	76K	483
unique products	2.1M	88K	43K
unique pairs	21M	99K	234K

Table 4: Dataset summary. Train and dev datasets are collected from historical browsing contexts; the test dataset (WANDS) is a human annotated dataset for evaluating product search relevance (Chen et al., 2022).

To evaluate the final trained model, we calculated nDCG@k scores on the WANDS dataset. The ground truth relevance score is derived from human annotation labels: relevant is 1, partial relevant is 0.5, and irrelevant is 0.

For both the retrieval accuracy and nDCG scores, we set $k = 50$, which roughly equals to the number of products shown on the first page of search results.

4 Results

We evaluated the models’ performance on the WANDS dataset. As noted in Table 4, this dataset contains 234k annotated (query, product) pairs. We further expand this dataset exhaustively to 20.6M pairs so that it contain every possible combination of unique query and unique product from the WANDS dataset. If a (query, product) pair is not annotated, we assume its relevance score is 0. Then, we assign these (query, product) pairs into different buckets, according to whether the query, the product, or the (query, product) pair exists in the training data or not. Details about bucket assignments are shown in Table 5. Finally, for each bucket, we calculate the averaged per-query nDCG@50 score.

RQ1: How well does the learned semantic space generalize? Table 5 shows each model’s nDCG@50 scores on different buckets of the WANDS dataset. The "seen" bucket has much higher nDCG scores than the "unseen" bucket, indicating a large opportunity area of generalization for all models. By looking further into the amortized unseen sub-buckets, it shows that queries and

products follow very different generalization behaviors. For example, the queries seen in the training data (q+, bucket 4, 5) usually performs better than the unseen queries (q-, bucket 6, 7). Yet counter-intuitively, the products seen in training under-perform the unseen ones (bucket 4 vs 5, 6 vs 7), which is against the common understanding where seen products may perform better than unseen ones. Arguably, this difference between query and products can be naively explained as products being relatively well learned compared to queries. We will revisit this point in analysis by graphical evidence.

RQ2: Are we able to improve generalization by pretraining? In Table 4, two domain-pretrained models (the qq model, and the qq+pp model) achieved top-two overall scores, showing the effectiveness of domain-specific pretraining approach. This, again, suggests the importance of query representation learning. Also, the SimCSE-large model performs worse than the Roberta-large model, suggesting that not all general-domain pretrainings are beneficial for domain-specific tasks, which may be due to the fact that the task of semantic search relies on asymmetric query-product relation encoding more than symmetric point-wise similarity. In contrast to SimCSE-large, the domain-specific query-query pretraining may have served the purpose of injecting domain knowledge mined from graph structure into text encoders, rather than a means for representation regularization, and such consistent improvement may have opportunity to be further improved given its moderate improvements over other models. Last, the Roberta-large model outperforms the Roberta-base model as expected due to a larger model capacity, yet the capacity for memorization is a draw between the two, based on the numbers in bucket 2.

5 Analysis

To understand why queries and products follow different generalization pattern in terms of nDCG scores observed in RQ1, we further analyzed the

Bucket index	Unseen sub-buckets						
	1	2	3	4	5	6	7
Bucket name	Overall	Seen	Unseen	q+, p+	q+, p-	q-, p+	q-, p-
Bucket ratio	100%	0.03%	99.9%	29.87%	31.60%	18.72%	19.79%
Roberta base	0.7812	0.9335	0.7722	0.7539	0.7894	0.6962	0.7249
Roberta large	0.7830	0.9337	0.7730	0.7511	0.7893	0.6942*	0.7177
SimCSE large	0.7775*	0.9315	0.7679*	0.7491	0.7875	0.6884*	0.7179*
Query-query (qq) pretrain	0.7866	0.9351	0.7760	0.7477*	0.7918	0.7057	0.7276
Product-product (pp) pretrain	0.7802*	0.9351	0.7700*	0.7460*	0.7875	0.7003	0.7201
qq+pp pretrain	0.7845	0.9319	0.7747	0.7489	0.7886	0.7006	0.7270

Table 5: nDCG@50 of fine-tuned models on product search tasks, evaluated by the WANDS dataset. The model names indicate the base model before fine-tuning. The “overall” bucket contains all (query, product) pairs in the testing dataset, the “seen” bucket contains (query, product) pairs from the testing dataset that also exist in the training dataset, and the “unseen” bucket contains (query, product) pairs from testing dataset that does not exist in the training dataset. The “unseen” bucket is further split into four different sub-buckets, according to whether the specific query or product exists in the training dataset or not. “q” represents query, “p” represents product; “+” represents existing in the training dataset, “-” represents not existing in the training dataset. Bucket ratio represents the size of current bucket with respect to the entire test dataset. In nDCG scores, “*” indicates that number is statistically smaller than that column’s largest (bold-faced) number, tested by one-side t-test with 0.05 p-value. nDCG results at other k’s are presented in the Appendix A.

query-product cosine score distribution for different buckets. To make distribution comparison fair across buckets, in each bucket we only keep a subset of (query, product) pairs that are annotated as “partial relevance”. The model we used to generate the cosine score distributions is the SimCSE-large model that is first **qq**-pretrained and then **pq** fine-tuned on click graph data, the best model according to Table 5.

First, the models perform better on seen query-product combinations than unseen combinations (the query and the product has been individually seen in training, but not the combination). The bucket 2 from Figure 2 shows a well learned cosine score distribution for seen query-product pairs, and the scores concentrate around a high similarity range around 0.65. In contrast, for bucket 4, although queries and products are individually seen by the models during training, the combination is novel to them, and it increases the models’ uncertainty, which is demonstrated by a significant similarity score spread-out accompanied by a shift to a lower median around 0.3. In general, the figure shows the difficulty of generalization of deep learning-based semantic search models on even the simpler task of dealing with unseen combinations. This suggests the generalization may be beyond the discussion of learning representations for an individual query or product, and rather about the geometry of the semantic space to properly main-

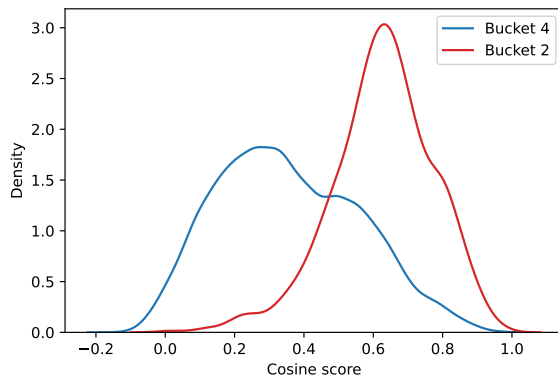


Figure 2: Query-product cosine score distribution for different buckets. The model is firstly pretrained by query-query similarity data, then fine-tuned in the product search task. Bucket 2 represent (query, product) pairs already seen by the model in training. In bucket 4, both query and product are seen, but the combination is not seen by the model in training. Bucket 2’s mean cosine score (0.6253) is significantly higher than bucket 4 (0.3562).

tain the distance between seen and unseen texts after semantic search as asymmetric embedding learning that “disturbs” a well-defined text representation space to encode text rankings.

Second, as shown in Figure 3, the visibility of queries and products in training data poses opposite impacts to a model’s behavior.

1) For seen products, comparing the curves of

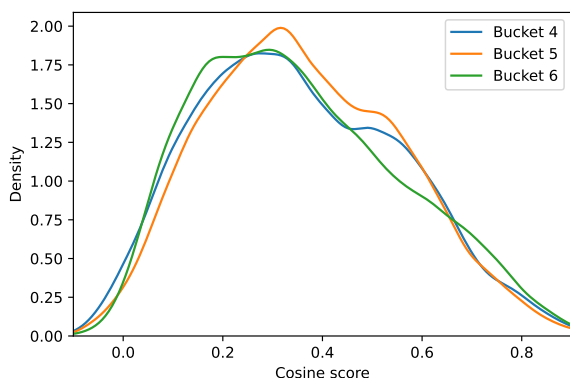


Figure 3: Query-product cosine score distribution for bucket 4, 5 and 6. Their mean cosine scores are 0.3562, 0.3655, and 0.3556 respectively. In bucket 4, both query and product are seen, but the combination is not seen by the model in training. In bucket 5, only the query is seen by the model in training. In bucket 6, only the product is seen by the model in training.

bucket 5 and 4, 1) counter-intuitively, we observe that 4 have a lower peak of around 0.25 as opposed to that of bucket 5 of around 0.3, which suggests that seen products have been “attracted” to seen queries in the training data, thus leading to a biased product representation. Yet, unseen products do not get to go through such an attraction process, and the representation may have maintained a neutral and generalize-able status in semantic space, leading to an overall higher and concentrated peak of 5 as opposed to a flat and lower peak of 4. This phenomenon might have suggested that a large semantic space may be over-parameterized that separately adjusting an individual product’s representation without affecting its geometrically closer neighbours may have become easy. Thus, an un-generalizable learning in generalizable models might have happened.

2) For seen queries, however, by comparing the cures of bucket 6 and 4 (the same as 7 and 5), we found that queries that have been seen in the training data lead to better semantic similarity with products in general than unseen queries, demonstrated by the distribution mass of the blue curve towards a higher similarity range than the green curve. The green line’s peak being extended towards the lower cosine score range demonstrates that the representation of seen queries may have improved instead of overfit. In other words, this may also suggest that query representation may have not been fully learned. From Table 5 we can also observe a simi-

lar hint that query-query pre-trained model is the most effective at improving semantic search performance from product ranking’s perspective. In stark contrast to seen products being overfit, seen queries gets to be learned and improved during semantic search’s **pq** learning process.

By comparing the above two cases, we have seen that queries may have learned slower than products in semantic search. The reason may be due to the fact that queries are usually much richer and more diverse (we have tens of times more unique queries than unique products), requiring a text encoder to observe way more cases for a query to learn representation well; whereas products are relatively fewer, well-formed and limited in diversity in product search. However, for a company that has billions of products in inventory, this observation may not hold true. Yet, the observation indeed resonates with that of question answering and web search, which suggests learning generalizable semantic search through improving query representation learning may still be a promising direction to follow.

6 Conclusions and Future Work

This paper examines the encoder generalization problem in the setting of semantic product search. We observe that query representation learning still is the bottle neck compared to the product representation learning under the semantic search training objective, the problem of which is surfaced by the superior performance of the domain-specific query pair pre-trained encoder model. We also surprisingly found that large, general-domain and unsupervisedly regularized Roberta variants such as SimCSE do not guarantee superior performance compared to Roberta-base or large, which may have suggested that semantic search is an *asymmetric* relationship learning task rather than symmetric, similarity learning task, since the semantic space has to encode much more complex relations (e.g. node priorities in a list) beyond point-wise similarity.

Although being a proof of concept, the proposed pre-training approach could start the conversation about how to properly pre-train encoders for semantic product search, and the results suggest that the knowledge hidden in the click graph can be better explored by defining proper training tasks, either transforming the task into a sequential task learning setting or a multi-task setting. In the fu-

ture, a transfer learning setting that leverage richer domain knowledge and semantic structures from relevant sources such as class and class taxonomy of queries may be promising. Also, data augmentation to enhance query representation learning, or novel learning curriculum/objective/regularization that can help balance query and product learning in the semantic search objective may also be valuable for exploration.

Acknowledgements

The authors would like to acknowledge previous and current members of the Wayfair Search and Recs team for full support on this research: Shujian Liu, Archi Mitra, John Castillo, Pooja Ranawade, John Murray, Avanti Patil, Zhilin Chen, Jeremy Myrland, and Matt Herman. We would also love to thank the suggestions and comments from reviewers to improve this paper.

References

- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26.
- Wei-Cheng Chang, Felix X. Yu, Yin-Wen Chang, Yiming Yang, and Sanjiv Kumar. 2020. [Pre-training tasks for embedding-based large-scale retrieval](#). *CoRR*, abs/2002.03932.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Yan Chen, Shujian Liu, Zheng Liu, Weiyi Sun, Linas Baltrunas, and Benjamin Schroeder. 2022. Wands: Dataset for product search relevance assessment. In *Proceedings of the 44th European Conference on Information Retrieval*.
- Luyu Gao and Jamie Callan. 2021. [Unsupervised corpus aware language model pre-training for dense passage retrieval](#). *CoRR*, abs/2108.05540.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [REALM: retrieval-augmented language model pre-training](#). *CoRR*, abs/2002.08909.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2553–2561.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. Question and answer test-train overlap in open-domain question answering datasets. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008.
- Sen Li, Fuyu Lv, Taiwei Jin, Guli Lin, Keping Yang, Xiaoyi Zeng, Xiao-Ming Wu, and Qianli Ma. 2021. Embedding-based product retrieval in taobao search. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3181–3189.
- Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow. 2016. Gated graph sequence neural networks. In *Proceedings of ICLR’16*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yuning Mao, Pengcheng He, Xiaodong Liu, Yelong Shen, Jianfeng Gao, Jiawei Han, and Weizhu Chen. 2021. [Generation-augmented retrieval for open-domain question answering](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4089–4100, Online. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- Priyanka Nigam, Yiwei Song, Vijai Mohan, Vihan Lakshman, Weitian Ding, Ankit Shingavi, Choon Hui Teo, Hao Gu, and Bing Yin. 2019. Semantic product search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2876–2885.

Barlas Oğuz, Kushal Lakhota, Anchit Gupta, Patrick Lewis, Vladimir Karpukhin, Aleksandra Piktus, Xilun Chen, Sebastian Riedel, Wen-tau Yih, Sonal Gupta, et al. 2021. Domain-matched pre-training tasks for dense retrieval. *arXiv preprint arXiv:2107.13602*.

Devendra Sachan, Mostofa Patwary, Mohammad Shoeybi, Neel Kant, Wei Ping, William L. Hamilton, and Bryan Catanzaro. 2021. [End-to-end training of neural retrievers for open-domain question answering](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6648–6662, Online. Association for Computational Linguistics.

Han Zhang, Songlin Wang, Kang Zhang, Zhiling Tang, Yunjiang Jiang, Yun Xiao, Weipeng Yan, and Wen-Yun Yang. 2020. Towards personalized and semantic retrieval: An end-to-end solution for e-commerce search via embedding learning. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2407–2416.

A Appendix: Additional nDCG scores of product search performance

Extending the results of Table 5, this section showed three more k values to the nDCG performance of fine-tuned models on product search tasks. Table 6 shows nDCG@1, Table 7 shows nDCG@20, and Table 8 shows nDCG@100.

The results from nDCG@20 and nDCG@100 are consistent with nDCG@50, such that the qq model achieves the best performance on both seen and unseen (query, product) pair segments. However, in nDCG@1 the Roberta-large model performs the best. This reveals that information gained from domain-specific pre-training makes it more difficult for the model to optimize the highest relevance score range; however, it helps the model to produce a generally better ranked list. Since semantic search models are usually used as candidate generators and followed by dedicated ranking models, the nDCG@1 performance is not as important as nDCG@50 or nDCG@100, and thus we still think domain-specific pre-training is a suitable approach to improve encoders’ performance.

Bucket name	Overall	Seen	Unseen	Unseen sub-buckets			
				q+, p+	q+, p-	q-, p+	q-, p-
Roberta base	0.7845	0.8858	0.7793	0.7988	0.8061	0.7204	0.7446
Roberta large	0.7969	0.9032	0.7892	0.7811	0.807	0.7231	0.746
SimCSE large	0.7876	0.8829	0.7793	0.7862	0.8019	0.7124	0.7554
Query-query (qq) pretrain	0.7919	0.9032	0.7816	0.7736	0.8123	0.746	0.7742
Product-product (pp) pretrain	0.7949	0.9017	0.7886	0.787	0.812	0.7272	0.7433
qq+pp pretrain	0.7764	0.8916	0.7712	0.771	0.8123	0.7164	0.7392

Table 6: nDCG@1 of fine-tuned models on product search tasks. Please refer to Table 5 for notation definitions.

Bucket name	Overall	Seen	Unseen	Unseen sub-buckets			
				q+, p+	q+, p-	q-, p+	q-, p-
Roberta base	0.7884	0.921	0.7807	0.7686	0.8053	0.7012	0.7269
Roberta large	0.7888	0.9203	0.7809	0.7641	0.8023	0.7051	0.7216
SimCSE large	0.7818	0.9183	0.7739	0.7631	0.7978	0.6941	0.7223
Query-query (qq) pretrain	0.7914	0.9229	0.7829	0.7607	0.8009	0.7165	0.7316
Product-product (pp) pretrain	0.7863	0.9221	0.7771	0.7598	0.7988	0.7072	0.7288
qq+pp pretrain	0.7877	0.9188	0.78	0.7628	0.8007	0.7151	0.7298

Table 7: nDCG@20 of fine-tuned models on product search tasks. Please refer to Table 5 for notation definitions.

Bucket name	Overall	Seen	Unseen	Unseen sub-buckets			
				q+, p+	q+, p-	q-, p+	q-, p-
Roberta base	0.7681	0.9386	0.7593	0.7603	0.7887	0.7138	0.7321
Roberta large	0.7678	0.9392	0.7583	0.7539	0.7842	0.7102	0.7284
SimCSE large	0.7643	0.9372	0.7545	0.7525	0.7834	0.7059	0.7279
Query-query (qq) pretrain	0.7721	0.9407	0.7624	0.7523	0.7913	0.7229	0.7334
Product-product (pp) pretrain	0.7664	0.9406	0.7571	0.7529	0.7846	0.7157	0.7308
qq+pp pretrain	0.7703	0.9377	0.7609	0.7539	0.789	0.7199	0.7315

Table 8: nDCG@100 of fine-tuned models on product search tasks. Please refer to Table 5 for notation definitions.