

XWalk: Random Walk Based Candidate Retrieval for Product Search

Jon Eskreis-Winkler
jeskreiswinkler@etsy.com
Etsy, Inc.

Yubin Kim
ykim@etsy.com
Etsy, Inc.

Andrew Stanton
astanton@etsy.com
Etsy, Inc.

ABSTRACT

In e-commerce, head queries account for the vast majority of gross merchandise sales and improvements to head queries are highly impactful to the business. While most supervised approaches to search perform better in head queries vs. tail queries, we propose a method that further improves head query performance dramatically. We propose XWalk, a random-walk based graph approach to candidate retrieval for product search that borrows from recommendation system techniques. XWalk is highly efficient to train and inference in a large-scale high traffic e-commerce setting, and shows substantial improvements in head query performance over state-of-the-art neural retrievers. Ensembling XWalk with a neural and/or lexical retriever combines the best of both worlds and the resulting retrieval system outperforms all other methods in both offline relevance-based evaluation and in online A/B tests.

KEYWORDS

e-commerce search, product search, graph, random walks, implicit feedback

ACM Reference Format:

Jon Eskreis-Winkler, Yubin Kim, and Andrew Stanton. 2023. XWalk: Random Walk Based Candidate Retrieval for Product Search. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom'23)*. ACM, New York, NY, USA, 5 pages.

1 INTRODUCTION

Modern large-scale search systems are tiered [13] with at least two layers. The *candidate retrieval* layer generates a small subset of potentially relevant documents from a corpus many orders of magnitude larger in size, while emphasizing efficiency and recall. The *re-ranking* layer uses more computationally expensive methods to re-rank the candidates generated by the retrieval stage to produce a high-precision final result list. Better recall in candidate retrieval leads to better overall accuracy. In this paper, we focus on improve search through improving recall in the candidate retrieval layer.

Most evaluations for search systems use an evaluation query set in which every query is assumed to be equally important and has equal impact on the accuracy metric. However, in reality, query frequency distributions are exponential [1]. Consequently, in e-commerce, head queries account for the vast majority of gross merchandise sales and head query performance is far more impactful to business metrics than torso or tail performance. State-of-the-art

supervised neural dense retrievers [2, 6, 7, 14, 16] typically perform better in head queries than tail, due to the higher availability of training data in the head region. However, we show that further substantial improvements to head query performance are possible. We borrow ideas from the recommendation systems community and propose XWalk, a graph-based approach to candidate retrieval.

Historically, graph-based approaches in search were used to create features (e.g. PageRank, click graphs [5, 19]) for the re-ranker layer, but have not been used directly for retrieval. Recently, graph neural networks (GNNs) have achieved state of the art performance in recommendation and are being adapted for search [8, 15, 18, 20]. However, large-scale GNNs are complex and slow to train.

The recommendation systems have long used implicit interaction graphs to directly generate recommendations. Commonly, users and product listings are represented as nodes in a graph and edges represent a logged interaction between a user and product listing (e.g. the user purchasing the listing). Random walks in graphs is a powerful technique used to generate recommendations from interaction graphs [3, 4, 11, 12]. Random walk based approaches are frequently used in large, real-time recommendation systems due to their effectiveness and efficiency [4, 12]. In addition, when using implicit feedback (e.g. logged interaction data such as user clicks) Park et al. [11] showed that random walk based approaches can perform better than matrix factorization approaches.

XWalk uses a random walk based approach to perform candidate retrieval for product search. In XWalk, we cast search as a query-to-listing recommendation problem (as opposed to user-to-listing), that is, we transform our query log into an implicit interaction graph between queries and product listings, and perform candidate retrieval by “recommending” listings to queries. Our approach trains using a fraction of the time and resources used by neural dense retrievers and GNNs, and is highly efficient in inference – XWalk scales to real-time search over graphs of billions of nodes and tens of billions of edges. XWalk also excels in head queries, where implicit feedback signals are plentiful.

While XWalk on its own suffers in tail and novel queries, we show that when results from XWalk are ensembled with a typical retriever that uses text similarity, even one as basic as plain BM25, it substantially improves overall candidate retrieval accuracy compared to strong neural dense retrieval and hybrid retrieval baselines, especially over the head query region, which is responsible for the overwhelming majority of sales in e-commerce. Furthermore, we show that XWalk is complementary to *both* dense retrieval and BM25, and demonstrate the strength of ensembling all three approaches.

To summarize, our novel contributions are: a) showing that XWalk substantially improves performance in the head query region, which accounts for the overwhelming majority of sales in

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

SIGIR eCom'23, July 27, 2023, Taipei, Taiwan
© 2023 Copyright held by the owner/author(s).

e-commerce; b) presenting an efficient random walk inference algorithm that can effectively serve queries at scale; c) showing that XWalk is complementary to other common retrieval methods and showing the strength of a simple ensemble approach that combines XWalk, BM25, and dense retrieval.

2 METHOD

We take inspiration from the recommendation space and recast the search problem as a query-to-product listing recommendation problem using implicit feedback: predict the best k product listings L_{q_i} to “recommend” to a query q_i , by learning from implicit user feedback, i.e. a query log. From the query log, we construct an undirected, weighted bipartite graph $G = (Q, L, E, W)$ where Q are nodes representing queries, L are nodes representing product listings, E are edges $E = \{e_{i,j} = (q_i, l_j) \mid q_i \in Q \wedge l_j \in L\}$, and W are edge weights.

2.1 Graph Construction (Offline Training)

Given a query log which records for each query q_i the set of listings $L_i^{click}, L_i^{cart}, L_i^{purchase}$ that the user clicked on, added to their shopping cart, and purchased, respectively, we construct our graph through the following process:

- (1) For each unique (by text string) query in the query log \hat{q}_i , add \hat{q}_i to Q .
- (2) For each unique (by listing ID) listing in the query log l_j , add l_j to L .
- (3) Collate the query log by query-listing pairs (\hat{q}_i, l_j) , counting the number of occurrences of $click_{i,j}$, $cart_{i,j}$, and $purchase_{i,j}$ interactions for each unique (\hat{q}_i, l_j) pair.
- (4) For each (\hat{q}_i, l_j) , add $e_{i,j}$ to E and its weight $w_{i,j}$ to W , where $w_{i,j}$ is calculated Equation 1.

Intuitively, edge weights represent the popularity or trustworthiness of the edge, i.e. if many different users bought listing l_i from query q_j , $w_{i,j}$ will be higher because we are more confident in the relationship represented by the edge. To weight edges, we use a simple linear combination:

$$w_{i,j} = C_1 \cdot |click_{i,j}| + C_2 \cdot |cart_{i,j}| + C_3 \cdot |purchase_{i,j}| \quad (1)$$

In practice, the best coefficients are $C_1 < C_2 < C_3$, as the goal is to bias walks toward listings which convert well for a given query.

2.1.1 Graph representation for efficient inference. XWalk is designed for sparse graphs scaling up to billions of nodes and tens of billions of edges. The costliest part of random walk graph inference is sampling edges to walk, especially from high degree nodes. For efficient inference, we choose our graph representation carefully.

We store edge weights as cumulative distribution functions in order to use Inverse Transform Sampling, which allows sampling in $O(\log(N))$ time. Note, we choose this approach over the alias method, which allows for constant time sampling, due to the doubling of memory needed for the transform. As XWalk’s space complexity is dominated by edges and corresponding weights, we develop other methods for efficient sampling (Section 2.2).

To transform edge weights in to CDF format, for each node n , we sort its adjacent edges $E_{n,*}$ in decreasing order of their weights

$W_{n,*}$, such that $w_{n,i} > w_{n,i+1}$. We then compute the cumulative distribution of all weights:

$$CDF_{n,i} = \frac{\sum_{j=0}^i w_{n,j}}{\sum_{i=0}^{|E_{n,*}|} w_{n,i}} \quad (2)$$

To sample an edge from $E_{n,*}$, we randomly sample $p \sim Uniform(0, 1)$ and find the corresponding edge through binary search. This formulations provides us a few valuable advantages:

- (1) Weighted sampling is $O(\log(|E_{n,*}|))$. Given some nodes have degrees in the millions, logarithmic growth is critical for performance.
- (2) Normalizing the CDF to 1 allows us to reconstruct the transition probability for outbound edges. This is key for the Metropolis-Hastings sampling strategy (Section 2.2).
- (3) Better cache coherence as the bulk of the weights are located near the front of the distribution.

Finally, we convert the graph into Compressed Sparse Row format, guaranteeing a $O(1)$ lookup cost for edges.

Note that all of the above graph construction steps are simple ETL (extract, transform, load) operations with no expensive parameter training steps. Compared to neural dense retrievers, “training” an XWalk graph model takes only a fraction of the cost and time.

2.2 Graph Inference (At Query Time)

Inferencing a graph with random walks is challenging to do efficiently. Despite the $O(1)$ edge lookup guarantee of the Compressed Sparse Row format used in graph construction, a naive walk approach that uses depth first search and binary search node lookups create random memory access patterns which result in high rates of costly cache misses [17]. We present an approach for XWalk that scales to graphs of billions of nodes and tens of billions of edges.

At query time, XWalk retrieves relevant listings for a query q_i by sampling nodes in G using k -hop fixed paths [3, 4] with node q_i as the starting point. When k is an odd number, the last node in a k -hop path will always be a listing node (L) due to the bipartite nature of G . XWalk returns listings ranked by the frequency of which they were sampled.

To reduce costly random memory access patterns, we use a breadth first search instead of depth first search for our random walks. We also improve upon the Inverse Transform Sampling strategy by using the Metropolis-Hastings algorithm (a Markov chain Monte Carlo method) in most places. Given the sorted CDF format of edge weights (Eq. 2), we can reconstruct the original edge transition probabilities: $P(n_j|E_{n_i,*}) = w_{n_i,j} - w_{n_i,j-1}$. As Metropolis-Hastings requires a symmetric distribution, we take the absolute value of the proposal index for each edge and sample from the Normal distribution. Ablation testing indicated XWalk is not sensitive to the variance for the proposal distribution, σ^2 . We set $\sigma^2 = 0.2$.

Metropolis-Hastings improves the cost of c edge samples to $O(\log(|E_{n,*}|)) + c$ compared to $c * O(\log(|E_{n,*}|))$ of Inverse Transform Sampling. In cases where c is large (e.g. the initial query node), the computational improvements are substantial. A known limitation of MCMC methods is the auto-correlation of samples, usually requiring a mix time prior to sampling. Therefore, for our first

sample, we use Inverse Transform Sampling to get an unbiased starting point and use Metropolis-Hastings for subsequent samples. In preliminary testing we found no reduction in model accuracy for this implementation compared to using only Inverse Transform Sampling while seeing the expected substantial latency benefits.

Our overall random walk strategy is presented in Algorithm 1.

Algorithm 1: XWalkBFSSampler

```

1 Global variables: Var of Normal distribution  $\sigma^2$ ,
   Dictionary of nodes to counts Counter
2 Input: Starting node n, Number of walks c, Walk-length k,
   Edges E, Weights W, Multiplier m (default 1)
3  $p \sim \text{Uniform}(0, 1)$ 
4  $i = \text{BinarySearch}(E_{n,*}, W_{n,*}, p)$ 
   /* the i'th node of ordered neighbors of n */
5  $\text{Counter}[\text{node}(E_{n,i})] += m$ 
6 for  $\text{step} = \{2, \dots, c\}$  do
7    $j = \text{Metropolis}(i, E_{n_i,*}, W_{n_i,*}, \sigma^2)$ 
   /* the j'th node of ordered neighbors of ni */
8    $\text{Counter}[\text{node}(E_{n_i,j})] += 1$ 
9    $i = j$ 
10 end
11 if  $k > 0$  then
12    $\text{counts} = \emptyset$ 
13   for  $\{n_i, \text{count}\} \in \text{Counter}$  do
14      $\text{counts} =$ 
        $\text{counts} \cup \text{XWalkBFSSampler}(n_i, \text{count}, k - 1, E, W)$ 
15   end
16   return  $\text{counts}$ 
17 else
18    $\text{Nodes} = \text{Sort}(\text{node} \in$ 
      $\text{Counter})$  in non-increasing order  $\{\text{Counter}[n_1] \geq$ 
      $\text{Counter}[n_2] \dots \geq \text{Counter}[n_{|\text{Counter}|}]\}$ 
19   return  $\text{Nodes}$ 
20 end

```

2.3 Extending the Graph

Our e-commerce platform is a two-sided marketplace and our inventory comes from independent sellers. Thus, listings are naturally grouped by shops. In addition, sellers may add tags to their listings to better describe them (e.g. “christmas”, “gift”, etc.).

For the sake of notation simplicity, we described the graph construction and inference above assuming our graph only contains two types of nodes, Q and L . However, in practice, we extend the graph by adding shop nodes (S) and tag nodes (T) to the graph; this allows us to retrieve listings without implicit user feedback (e.g. the cold start problem) and further increase connectivity of the graph. Note that G remains bipartite: $\{Q, S, T\}$ is a separate partition from L and thus the algorithms described in this section can be used unchanged. The weights of edges between shops/tags and listings are set to 1. $w_{q_i, s_j} = w_{q_i, t_j} = 1$.

3 EXPERIMENTS

For our experiments, we sought to closely emulate a real-world e-commerce setting, where the main source of training data is implicit

user feedback from query logs, and models are evaluated under a realistic query popularity distribution. Unfortunately, most public search datasets do not reflect a realistic query distribution and rarely have implicit user feedback as training data. While recommendation system datasets have implicit user feedback, they do not have a text query that is usable by BM25 or dense retrieval, which retrieve based on query to listing text similarity. Therefore, we curated a training and evaluation dataset from our e-commerce platform.

3.1 Dataset Creation

For training data, we collected 365 days of implicit feedback data, comprising of records of queries and the product listings that were clicked, added to cart, or purchased from a given query. Queries are represented by their query text. Listings are represented by their unique ID and the title of the product. In addition, as mentioned in Section 2.3, listings are associated with seller-provided tags, and each listing belongs to exactly one seller’s shop.

Over the time period used for this experiment there were 137,824,871 unique listings, 147,174,817 unique queries, 62,803,463 unique tag, and 3,018,713 unique shops. There were a total of 1,349,734,328 query-listing interactions recorded, where 3.46% were purchases, 6.19% were cart adds and 90.3% were clicks. Altogether, there were 1,395,759,140 edges. Example records are found in Table 1.

Evaluation data was curated to be a representative query distribution, sampled from a single day immediately following the last day of the training data window. We randomly sampled 11,521 queries that resulted in at least one purchase. As the sample is intended to be reflective of the true query popularity distribution, we did not de-duplicate the query set. Figure 1 shows the distribution of the query frequency in the evaluation set.

For each query, the listings that were purchased from that query are considered the relevant document. 82.3% of queries had only one purchase, 12.0% had two purchases and 5.6% had more than two purchases. For each one of these queries, we assigned them to a head/torso/tail frequency bin based on how frequently they occurred in the previous 365 day period. The bins were created such that the total counts of requests are roughly equal among those bins. Of the evaluation queries, 31.0% were in the head bin, 47.9% were in the torso bin, and 43.9% were in the tail bin.




Figure 1: Distribution of evaluation queries’ log frequency, ordered from least frequent to most frequent.

3.2 Experiment set up

We compare XWalk against two other methods of candidate retrieval. First is lexical retrieval using BM25 scoring (BM25). We use

query	listing ID	listing title	interaction type	shop	tags
wedding dress	112	beautiful bridal wedding gown	click	s00	white, gown
wedding gown	112	custom embroidered wedding dress	purchase	s00	white, gown
wedding dress	134	ethereal champagne dress with chiffon skirt	click	s11	fancy, dress, chiffon

Table 1: Hand-created example of implicit feedback training data.

Pyserini [9] to build a Lucene index based on listing titles in our dataset and then retrieve candidates using BM25 rankings using bag-of-words representations. We used the default analyzer and default BM25 parameters ($k1=1.2$, $b=0.75$).

The second baseline is a state of the art neural dense retrieval system [10] trained on search traffic for candidate retrieval (NIR). NIR uses a smaller time window of training data (30 days) due to the time and expense of training on larger data sets. NIR is a Transformer-based, two tower model that uses a multi-part hinge loss to distinguish between interactions that involve a purchase, cart add, favorite, click, or nothing. Model was trained over one epoch for YYY hours. It was designed for better semantic matching between queries and listings by incorporating title, query as well as additional features such as tags, and listing taxonomy.

In addition to the above, we also compare results against hybrid systems of NIR+BM25 [2]. To ensemble the results from each retrieval engine, we use Reciprocal Rank Fusion, a simple but effective fusion technique [2]. Higher recall in candidate retrieval result in higher overall search accuracy [2]. As we are focusing on the first-pass candidate retrieval stage of search, we use recall and mean average precision (MAP) at 100, and 1000 to measure the quality of candidates retrieved.

4 ANALYSIS

As shown in Table 2, when compared independently against other methods, XWalk out-performs other methods in most metrics despite the fact that it is unable to return results for novel queries, due to its strength in the head query bin. When combined with BM25, it outperforms in every metric, both NIR and the hybrid NIR+BM25. Finally, the ensemble of all three methods (XWalk+BM25+NIR) substantially outperforms all other configurations.

We see in Table 2 that BM25 is significantly weaker in performance compared to NIR and XWalk and does not always improve the overall results of NIR and XWalk, especially for MAP. While BM25 can improve recall by adding listings that were not retrieved by NIR and XWalk, its poor ranking drags down MAP in the hybrid systems. For the most popular short queries, BM25 is not able to distinguish between the many listings with titles that token match similarly to the query. Whereas methods like XWalk and NIR are able to provide a more reliable ranking of the highly purchaseable listings based on training data.

However, in Table 3 we see that XWalk is complementary to BM25; XWalk is stronger in the head and torso bins while BM25 outperforms XWalk in the tail bin. This is due to the fact that XWalk suffers from cold start problems: it performs best with many prior examples and is unable to handle novel queries. BM25, as a lexical matching system, is more able to handle novel queries. Furthermore, XWalk+BM25 is still yet complementary with NIR. The semantic matching of dense retrieval excels in the tail, where queries are typically longer. When all three systems are ensemble,

it is the highest performing across all query bins. XWalk’s success in the head query bin is particularly notable – in an e-commerce setting, the head query bin is responsible for a large majority of merchandise sales.

	r@100	r@1000	M@100	M@1000
BM25	0.192	0.394	0.034	0.035
NIR	0.547	0.740	0.107	0.109
XWalk	0.600	0.723	0.153	0.154
NIR+BM25	0.497	0.780	0.097	0.100
XWalk+BM25	0.599	0.829	0.129	0.132
XWalk+BM25+NIR	0.701	0.915	0.194	0.198

Table 2: recall (r@100, r@1000) and MAP (M@100, M@1000) for retrieval models and combinations.

	tail	torso	head
BM25	0.471	0.420	0.299
NIR	0.738	0.728	0.759
XWalk	0.260	0.813	0.899
NIR+BM25	0.804	0.779	0.762
XWalk+BM25	0.595	0.875	0.914
XWalk+BM25+NIR	0.836	0.931	0.942

Table 3: Comparison of retrieval models in terms of recall@1000 stratified by query popularity.

5 ONLINE TESTING

We tested XWalk in a live online A/B experiment on a large e-commerce platform. The experiment ran for 23 and 25 days on our mobile and web version of our platform, respectively. Our search system is a two-stage search system, which uses an ensemble of candidate retrievers in the first pass, followed by a second pass re-ranker. In our A/B experiment, an ensemble of NIR+Solr as the candidate retrieval system was compared against an ensemble of XWalk+NIR+Solr.

We saw a statistically significant and substantial increase in conversion rate for the search system including XWalk in both the web and mobile platforms, +1.2% on web and +1.98% on mobile. In addition, in a production setting, we saw that XWalk was our lowest latency retrieval engine. The 99th percentile latency is only 58% of the NIR engine and 22% that of our Solr inverted index.

6 CONCLUSION

Head queries are responsible for the large majority of purchases in e-commerce. We presented XWalk, a novel candidate retrieval engine, which by frames search as a query-to-product recommendation problem, leverages powerful, highly efficient graph methods to substantially improve head query performance in product search. XWalk is also complementary to other common retrieval engines such as BM25 and dense retrieval, and ensembling produces a powerful retrieval engine.

REFERENCES

- [1] Ricardo Baeza-Yates and Alessandro Tiberi. Extracting semantic relations from query logs. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '07, page 76–85, New York, NY, USA, 2007. Association for Computing Machinery.
- [2] Tao Chen, Mingyang Zhang, Jing Lu, Michael Bendersky, and Marc Najork. Out-of-Domain Semantics to the Rescue! Zero-Shot Hybrid Retrieval Models. In Matthias Hagen, Suzan Verberne, Craig Macdonald, Christin Seifert, Krisztian Balog, Kjetil Nørkvåg, and Vinay Setty, editors, *Advances in Information Retrieval*, Lecture Notes in Computer Science, pages 95–110, Cham, 2022. Springer International Publishing.
- [3] Fabian Christoffel, Bibek Paudel, Chris Newell, and Abraham Bernstein. Blockbusters and Wallflowers: Accurate, Diverse, and Scalable Recommendations with Random Walks. In *Proceedings of the 9th ACM Conference on Recommender Systems*, RecSys '15, pages 163–170, New York, NY, USA, September 2015. Association for Computing Machinery.
- [4] Chantat Eksombatchai, Pranav Jindal, Jerry Zitao Liu, Yuchen Liu, Rahul Sharma, Charles Sugnet, Mark Ulrich, and Jure Leskovec. Pixie: A System for Recommending 3+ Billion Items to 200+ Million Users in Real-Time. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 1775–1784, Republic and Canton of Geneva, CHE, April 2018. International World Wide Web Conferences Steering Committee.
- [5] Shan Jiang, Yuening Hu, Changsung Kang, Tim Daly, Dawei Yin, Yi Chang, and Chengxiang Zhai. Learning Query and Document Relevance from a Web-scale Click Graph. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, SIGIR '16, pages 185–194, New York, NY, USA, July 2016. Association for Computing Machinery.
- [6] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online, November 2020. Association for Computational Linguistics.
- [7] Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. Latent Retrieval for Weakly Supervised Open Domain Question Answering. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6086–6096, Florence, Italy, July 2019. Association for Computational Linguistics.
- [8] Xiangsheng Li, Maarten de Rijke, Yiqun Liu, Jiaxin Mao, Weizhi Ma, Min Zhang, and Shaoping Ma. Learning Better Representations for Neural Information Retrieval with Graph Information. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 795–804, Virtual Event Ireland, October 2020. ACM.
- [9] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. Pyserini: A Python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, pages 2356–2362, 2021.
- [10] Priyanka Nigam, Yiwei Song, Vijai Mohan, Vihan Lakshman, Weitian Ding, Ankit Shingavi, Choon Hui Teo, Hao Gu, and Bing Yin. Semantic product search. *CoRR*, abs/1907.00937, 2019.
- [11] Haekyu Park, Jinhong Jung, and U. Kang. A comparative study of matrix factorization and random walk with restart in recommender systems. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 756–765, December 2017.
- [12] Bibek Paudel, Fabian Christoffel, Chris Newell, and Abraham Bernstein. Updatable, Accurate, Diverse, and Scalable Recommendations for Interactive Applications. *ACM Trans. Interact. Intell. Syst.*, 7(1):1:1–1:34, December 2016.
- [13] Lidan Wang, Jimmy Lin, and Donald Metzler. A cascade ranking model for efficient ranked retrieval. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '11, page 105–114, New York, NY, USA, 2011. Association for Computing Machinery.
- [14] Shuai Wang, Shengyao Zhuang, and Guido Zuccon. BERT-based Dense Retrievers Require Interpolation with BM25 for Effective Passage Retrieval. In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '21, pages 317–324, New York, NY, USA, August 2021. Association for Computing Machinery.
- [15] Xinlin Xia, Shang Wang, Han Zhang, Songlin Wang, Sulong Xu, Yun Xiao, Bo Long, and Wen-Yun Yang. SearchGCN: Powering Embedding Retrieval by Graph Convolution Networks for E-Commerce Search. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, pages 2633–2634, New York, NY, USA, July 2021. Association for Computing Machinery.
- [16] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *International Conference on Learning Representations*, 2021.
- [17] Ke Yang, Xiaosong Ma, Saravanan Thirumuruganathan, Kang Chen, and Yongwei Wu. Random Walks on Huge Graphs at Cache Efficiency. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles CD-ROM*, pages 311–326, Virtual Event Germany, October 2021. ACM.
- [18] Hamed Zamani and W. Bruce Croft. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 717–725, Houston TX USA, January 2020. ACM.
- [19] Yuan Zhang, Dong Wang, and Yan Zhang. Neural IR Meets Graph Embedding: A Ranking Model for Product Search. In *The World Wide Web Conference*, WWW '19, pages 2390–2400, New York, NY, USA, May 2019. Association for Computing Machinery.
- [20] Kai Zhao, Yukun Zheng, Tao Zhuang, Xiang Li, and Xiaoyi Zeng. Joint Learning of E-commerce Search and Recommendation with a Unified Graph Neural Network. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 1461–1469, Virtual Event AZ USA, February 2022. ACM.