

# Que2Engage: Embedding-based Retrieval for Relevant and Engaging Products at Facebook Marketplace

Yunzhong He<sup>†1</sup>, Yuxin Tian<sup>†2</sup>, Mengjiao Wang<sup>1</sup>, Feier Chen<sup>1</sup>, Licheng Yu<sup>1</sup>, Maolong Tang<sup>1</sup>,  
Congcong Chen<sup>1</sup>, Ning Zhang<sup>1</sup>, Bin Kuang<sup>1</sup> and Arul Prakash<sup>1</sup>

<sup>1</sup>Meta, <sup>2</sup>University of California, Merced

{yunzhong,mengjiaow,feche,lichengyu,mltang,chenc,ningzhang,binkuang,arulprakash}@meta.com  
ytian8@ucmerced.edu

## ABSTRACT

Embedding-based Retrieval (EBR) is a powerful search retrieval technique in e-commerce to address semantic matches between search queries and products. However, commerce search engines like Facebook Marketplace Search are complex multi-stage systems with each stage optimized for different business objectives. Search retrieval system usually focuses on query-product semantic relevance, while search ranking puts more emphasis on up-ranking products for high quality engagement. As a result, the end-to-end search experience is a combined result of relevance, engagement, and the interaction between different stages of the system. This presents challenges to EBR systems in optimizing overall search experiences. In this paper we present Que2Engage, a search EBR system designed to bridge the gap between retrieval and ranking for better end-to-end optimization. Que2Engage takes a multimodal & multitask approach to infuse contextual information into the retrieval stage and balance different business objectives. We show the effectiveness of our approach via a multitask evaluation framework with thorough baseline comparisons and ablation studies. Que2Engage has been deployed into Facebook Marketplace Search engine and shows significant improvements in user engagement in two weeks of A/B testing.

## CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking**; •  
**Applied computing** → **Online shopping**.

## KEYWORDS

Product search, e-commerce, information retrieval

### ACM Reference Format:

Yunzhong He<sup>†1</sup>, Yuxin Tian<sup>†2</sup>, Mengjiao Wang<sup>1</sup>, Feier Chen<sup>1</sup>, Licheng Yu<sup>1</sup>, Maolong Tang<sup>1</sup>, Congcong Chen<sup>1</sup>, Ning Zhang<sup>1</sup>, Bin Kuang<sup>1</sup> and Arul Prakash<sup>1</sup>. 2023. Que2Engage: Embedding-based Retrieval for Relevant and Engaging Products at Facebook Marketplace. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*, April 30-May 4, 2023, Austin, TX, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3543873.3584633>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*WWW '23 Companion*, April 30-May 4, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 978-1-4503-9419-2/23/04...\$15.00  
<https://doi.org/10.1145/3543873.3584633>

## 1 INTRODUCTION

Embedding-based Retrieval (EBR) has become an important component of e-commerce search engines across Facebook Marketplace, Walmart, Instacart, and more [6, 7, 9, 16]. In general, EBR models focus on learning embedding representations for search queries and documents, so that documents semantically close to a search query can be retrieved via ANN search [5–7, 9]. However, search engines are usually complex multi-stage systems optimized for multiple business objectives, so simply optimizing for semantic relevance may not always lead to the best outcome. For example, [7] points out that integrating EBR systems can lead to NDCG regressions because downstream re-ranking systems may not always be able to rank results retrieved via EBR properly.

In e-commerce platforms like Facebook Marketplace<sup>1</sup>, contextual information such as product price, production condition, seller rating, *etc.* are also important signals to consider to ensure products retrieved are engaging to the searchers. However, we argue that the leverage of contextual information in a search EBR setting towards better searcher engagement is not a trivial problem because (1) traditional EBR modeling techniques based on contrastive learning overly emphasize semantic relevance, so naively apply contextual information in a contrastive learning setting may not work well (2) a product being semantically relevant to a query does not imply that it is engaging to the searcher, and thus simultaneously preserve relevance and engagement is challenging.

In this paper we present Que2Engage, extending Que2Search [7] to address the aforementioned challenges. It takes a multimodal approach to incorporate contextual signals as a unique modality in its transformer fusion backbone. The model is trained with multitask learning that joins contrastive learning with ranker-style training to not only retrieve semantically relevant products, but also up-ranks the more engaging products like a re-ranking model. Similar to [20], we propose a multitask evaluation for EBR models to understand its performances in different domains. We share detailed baseline comparisons and ablation studies using the multitask evaluation framework to illustrate our argument of multi-stage consistency and the effectiveness of our approach in leveraging contextual information.

Que2Engage is integrated in Facebook Marketplace Search and powering millions of search queries per day. It has demonstrated significant improvements in searcher engagement via two weeks of online A/B testing.

<sup>1</sup>[www.facebook.com/marketplace](http://www.facebook.com/marketplace)

<sup>†</sup> Both authors contributed equally to this research

## 2 MODELING

Figure 1 presents the overall architecture of our Que2Engage framework, which is a two-tower neural network consisting of a query and a document tower for learning embedding representations of search queries and e-commerce products, respectively. Multimodal and contextual information of products are fed into the document tower using a transformer-fusion approach, and trained with multi-task learning. In the following, we detail our choices in model architecture and our novel multimodal multitask method.

### 2.1 Model Architecture

**2.1.1 Query tower.** Similar to [7], we adopt a multi-granular representation of search queries which consists of both raw query text and character trigrams of the query. Raw query text is encoded using a 2-layer XLM [4] encoder, and character trigrams are encoded using an EmbeddingBag [10] encoder. Different from [7], we combine the two representations using concatenation instead of attention fusion before sending to the final MLP layer as we find the former yields slightly better performance.

**2.1.2 Contextual information as a modality.** For candidate product listings, we use a MLP-based encoder from the document tower to encode their contextual information such as price, category and creation time. Numerical features are represented as single neurons, and categorical features are represented using one-hot encoding. All of the contextual features are concatenated and then fed into a BatchNorm layer followed by a final MLP to ensure a fixed numerical scale and a fixed output length. We call this encoder output a "context token" because it is treated similarly to text and image tokens during the multimodal fusion step covered in section 2.1.3. Essentially, contextual information is treated as a unique modality in our multimodal framework.

**2.1.3 Multimodal fusion.** Besides encoding the contextual information, we use text encoder to convert the textual fields, *i.e.*, product title and description, into a sequence of word tokens and feed them into the transformer to get the textual embedding. A special [CLS] token is used to encode the whole sentence representation. For the variable number of images attached to the document, we take the pre-trained image representations [1] for each of the attached images, apply a shared MLP layer and deep sets [19] fusion to get the image dense representation as image modality token. We borrow the transformer-fusion architecture used in [18], where we feed the concatenation of the text tokens, image token as well as context token to the multimodal fusion encoder. Our text encoder and multimodal fusion encoder are initialized from 6-layer XLM-R [3], a multilingual language model. As in [18], the text encoder inherits its first K layers and the multimodal fusion model inherits its remaining M layers. We extract the hidden output of the [CLS] token at the last layer of multimodal fusion encoder and project it to the desired dimension as the final document embedding.

**2.1.4 Modality dropout.** To ensure that our model does not overly rely on one modality and is robust against missing information during inference time, we introduce a modality dropout mechanism. Specifically, we randomly mask out the output of contextual encoder, image encoder and text encoder with probability of  $\delta_c$ ,  $\delta_i$

and  $\delta_t$  respectively, and replace the masked ones with tensor of zero.

**2.1.5 Learning image representations.** We explore two variants of image encoders to capture visual information from product images. In method one, we directly apply pre-computed image embedding from the GrokNet model [1], with an MLP layer on top to ensure the image embedding size is consistent all the other tokens in transformer fusion. In method two, we include an off-the-shelf ResNet50 encoder from the CommerceMM model [18] into the document tower, and train the entire document tower as a continued CommerceMM fine-tuning process. While the latter is obviously more powerful because the original image encoder is retained and fine-tuned, the former is much simpler to train and costs less memory in both training and serving. We will compare them in more details in section 3.

### 2.2 Multitask Training

**2.2.1 Contrastive learning.** We adopt contrastive learning based on batch negative sampling as part of our training objectives, where positive samples are user engaged <query, product> pairs sampled from de-identified and aggregated search log, and negative samples are generated by randomly combining queries and products within a mini-batch of positive samples. Formally, we introduce the relevance loss  $L_{relevance}$  as follows

$$L_{relevance} = \frac{1}{B} \sum_{i=1}^B -\log \left\{ \frac{\exp \{s \cdot \kappa(q_i, d_i)\}}{\sum_{j=1}^B \exp \{s \cdot \kappa(q_i, d_j)\}} \right\} \quad (1)$$

where  $B$  is the batch size,  $\kappa$  is a similarity kernel that is implemented as the cosine similarity, and  $s$  denotes a scaling factor which is simply a fixed value  $s = 20$  throughout all experiments. The value is the same as the Que2Search [7] model.

**2.2.2 Learning contextual information.** Although batch negative sampling has proven useful in learning semantic relevance in the search EBR problem [7, 9], we notice that it is not sufficient in learning contextual information towards user engagement. For example, contextual information like product price is a distinguishing factor to identify engaging products among relevant products (*i.e.* a product can be relevant but receives no engagement because the listed price is not reasonable). However, during batch negative training, the model receives little negative supervision from products with very unreasonable prices, since all negative samples are generated from engaged <query, product> pairs. Fundamentally, as [15] points out, this is because batch negative methods implicitly sample from the distribution of engaged products, which may not be the true distribution of the inventory. Methods like mixing random negatives [17] and in-batch hard negative mining [7] are proposed to mitigate the problem. However, we observe that negative samples generated by those approaches are still too easy for the model to pick up the nuances in contextual information. Therefore, we propose an auxiliary training task that optimizes the model directly towards finding engaging products among relevant product. Specifically, we augment the training set in section 2.2.1 by including <query, product> pairs displayed to the searchers but receive no searcher engagements as hard negatives, and computes a BCE loss on those

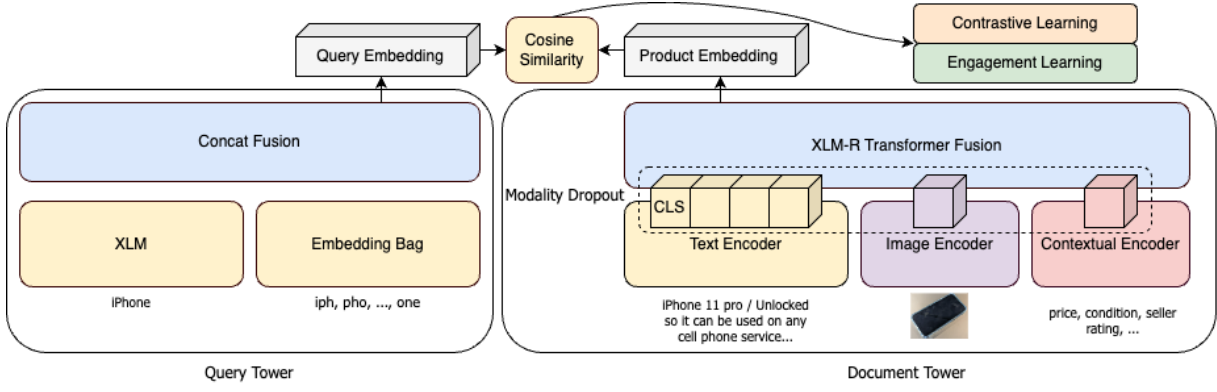


Figure 1: Que2Engage architecture overview

samples. Formally, we define the loss  $L_{engagement}$  as follows

$$L_{engagement} = -(y_i \log(c_i) + (1 - y_i) \log(1 - c_i)) \quad (2)$$

where  $c_i = s \cdot \kappa(q_i, d_i)$ .

To combine  $L_{relevance}$  and  $L_{engagement}$ , we define the final multitask loss as

$$L(\theta) = \lambda_1 \cdot L_{relevance} + \lambda_2 \cdot L_{engagement} \quad (3)$$

where  $\theta$  is the model parameters,  $\lambda_1$  and  $\lambda_2$  are the weighting parameters chosen empirically.

## 3 OFFLINE EXPERIMENTS

### 3.1 Dataset

We collect 150 million <query, product> pairs displayed to searchers from Facebook Marketplace’s search log, with 75 million receiving downstream engagements as positive samples, and the rest of 75 million as negative samples. The data is de-identified and aggregated to remove any personal information before evaluation proceeds. For offline evaluation, we collect 26k human-rated data as the relevance evaluation set. The human-rated dataset is generated by letting raters to decide whether a result is relevant to a query or not. The candidates to be rated are generated by a stratified sampling of the search queries and products, which includes both easy and hard samples. To evaluate user engagement, we reserve one future date among the 150 million de-identified and aggregated search log data.

### 3.2 Baselines and Ablation Studies

We choose Que2Search [7] as our baseline model, which is a two-tower model based on attention fusion of pre-trained XLM encoders [4] and image representations. We further augment the model with encoders based on contextual information, as well as with mixed batch method [17] to incorporate hard negatives from products displayed to searchers. For the treatment group, we use the Que2Engage with pre-computed image embeddings, because in practice its simplicity is preferred during the actual model production, and we share the comparison against an alternative image encoder based on fine-tuning of the CommerceMM encoder separately.

### 3.3 Experimental Setup

**3.3.1 Evaluation metrics.** Similar to [20], we adopt a multitask evaluation framework to measure semantic relevance and searcher engagement separately. Semantic relevance is measured using the 26k human-rated dataset, and searcher engagement is measured using the 220k future engagement dataset. For both datasets, we rank them using the cosine similarity from our model and report ROC\_AUC as the evaluation metric. Note that one significant difference between the two datasets lies in personalization. For the human-rated dataset, raters are asked to make judgement based purely on an objective guideline around textual and visual relevance (e.g. whether there is a catalog or brand mismatch between a query and a product), while more subjective factors outside of the guideline (e.g. whether the listed price or product condition is appealing) play important roles in the engagement dataset.

Therefore, ROC\_AUC on the relevance dataset measures how well the model predicts search relevance (similar to the in relevance degree in [20]), which is the main evaluation metric used in [7]. We have also found it correlates well with search retrieval performance. ROC\_AUC of the engagement dataset essentially measures its performance on the search ranking task, because the candidates are all products with user impressions. In fact, it is also a good indicator of the consistency between search retrieval and search ranking.

**3.3.2 Parameters.** We develop all of the models on Nvidia A100 GPUs using the PyTorch Multimodal framework [11]. Models are trained using batch size of 512, and optimized using Adam optimizer with a learning rate of  $4e - 4$ . We set weighting parameter  $\lambda_1$  and  $\lambda_2$  as 0.8 and 0.2 respectively. For the modality dropout,  $\delta_c$ ,  $\delta_i$  and  $\delta_t$  are 0.5, 0 and 0.5. We directly feed the text tokens and tokens from other modalities into the multimodal transformer as [18], i.e., our multimodal transformer is an early-fusion model with 0-layer text encoder and 6-layer multimodal fusion encoder. One exception to the aforementioned settings is that when comparing pre-computed image embeddings with fine-tuning the CommerceMM encoder, due to the increased GPU memory consumption of the fine-tuning approach, we adjust the batch size to 64 for that particular experiment. Learning rate was also adjusted to  $5e - 5$  for the CommerceMM training to avoid NaN in the loss function computation.

### 3.4 Results

**3.4.1 Analysis of baseline results.** The top part of table 1 outlines the performance of baseline models. We can see that Que2Search does well on the relevance evaluation but performs poorly on the engagement dataset, which is expected because it is optimized for semantic relevance. While adding contextual information to Que2Search itself does not improve the performance on engagement evaluation, changing training objective to include mixed negatives significantly improves the prediction. Mixed negatives also helps Que2Search to better leverage contextual features towards engagement prediction with row 4 in table 1 achieving the highest ROC\_AUC on engagement evaluation. This aligns with our hypothesis in section 2.2.1 - vanilla batch negative is insufficient to learn the nuances in engagement prediction from contextual information due to the sampling bias introduced by generating negatives from the positives, and this can be mitigated by introducing mixed negatives. Note that row 4 and 5 in table 1 suggests a regression in relevance evaluation along with the improvement of engagement evaluation. This is expected because contextual information like product price and condition are irrelevant to the query-product relevance guideline provided to our raters, and thus we do not expect the leverage of contextual signals to improve the relevance evaluation.

**3.4.2 Que2Engage and ablation studies.** The second part of table 1 shows that engagement evaluation can be significantly improved with the Que2Engage approach, with the full Que2Engage using multitask learning and modality dropout achieving the best results across the two evaluation methods. Ablation study on the loss function suggests that multitask training leads to the biggest improvement in engagement evaluation, suggesting that a more focused loss function on hard negatives works better than simply mixing the negatives. Finally, modality dropout further improves both metrics and specially the relevance evaluation, suggesting that forcing missing modalities may prevent the model over-fitting on one task and thus regressing the other.

**3.4.3 Image encoders.** We also compare the two approaches to incorporate image signals into the multimodal fusion framework. We can see that although both methods are based on pre-training tasks, being able to fine-tune the image encoder indeed outperforms the original frozen GrokNet [1] approach used in Que2Search in both evaluations. However, given that fine-tuning the image encoder end-to-end requires significantly more GPU memory and training time, we do not include this technique in the production model for simplicity. And the ablation study is done with a reduced batch size of 64 which significantly regresses the absolute relevance evaluation because larger batch size is proven helpful in contrastive learning [14]. However, we hope to productionalize this technique with memory efficiency optimizations such as [14, 20] and with more powerful hardware.

## 4 ONLINE EXPERIMENTS

We deploy Que2Engage on Facebook Marketplace Search as a parallel retrieval source to the traditional lexical-based search retrieval. For online A/B testing, we compare Que2Engage with Que2Search [7], which is our previous production model solely optimized for

Model	Engagement	Relevance
Que2Search[7]	55.88	67.14
Que2Search w/ contextual encoder	55.85	66.74
Que2Search w/ mixed batch loss	63.63	63.79
Que2Search w/ contextual encoder + mixed batch loss	64.45	61.17
Que2Engage w/ mixed batch loss	64.70	60.36
Que2Engage w/ multitask training	76.13	65.63
Que2Engage w/ multitask training + modality dropout	<b>76.90</b>	<b>67.21</b>

**Table 1: Results for baseline comparison and ablation studies**

Method	Engagement	Relevance
pre-computed image embedding	74.35	60.19
fine-tuned CommerceMM encoder	<b>74.67</b>	<b>60.55</b>

**Table 2: Results for image encoder comparison**

semantic relevance. We measure both NDCG and searcher engagement for this A/B testing. Note that NDCG is calculated from human-rated labels using simulated search results similar to how the human-rated evaluation set is generated. Two weeks of online A/B testing shows that Que2Engage improves online searcher engagement by 4.5% while keeping NDCG neutral, which aligns with our offline multitask evaluation results.

## 5 RELATED WORK

In recent years, Embedding-based Retrieval (EBR) has been adopted in e-commerce search to retrieve semantically relevant products as a complement of lexical retrieval [6, 7, 9, 20]. Siamese neural networks [2, 12] trained with contrastive learning loss [13, 15] are among the popular modeling choices for EBR in both search and recommendation systems [8, 14, 17]. However, naive contrastive learning using in-batch negatives can suffer from missing interesting negative samples [17, 20] and being memory-hungry [14]. Variants of contrastive learnings are proposed to address them by incorporating smarter negative sampling [7, 17, 20] and optimizing memory usage [14, 20]. Sometimes, teacher-student learning is also used as an auxiliary task to improve relevance [5, 20]. In search retrieval, Pre-trained Language Models (PLM) are widely adopted because the main focus of retrieval is often textual relevance [8, 9]. Yet, [6, 7, 20] point out that contextual information beyond text relevance and consistency with re-ranking stage are also important factors to EBR systems' end-to-end performance. Recently, [20] developed a unified training scheme to balance multiple optimization objectives, yet the role of contextual information in the multi-objective setting is rarely discussed.

## 6 CONCLUSION

We present the need of EBR modeling to balance relevance and engagement in the real-world applications like Facebook Marketplace, and introduce Que2Engage, our latest search EBR system, to address the challenges. Through baseline comparisons and ablations studies, we show the effectiveness of our innovations in incorporating contextual signals, multimodal techniques, representation learning and multitask learning. We have deployed Que2Engage on Facebook Marketplace Search. Through two weeks of A/B testing, we show that it outperforms our existing state-of-the-art search EBR system [7] and significantly improved searcher engagement on product listed at Facebook Marketplace.

## REFERENCES

- [1] Sean Bell, Yiqun Liu, Sami Alsheikh, Yina Tang, Edward Pizzi, M. Henning, Karun Singh, Omkar Parkhi, and Fedor Borisjuk. 2020. GrokNet: Unified Computer Vision Model Trunk and Embeddings For Commerce. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Virtual Event, CA, USA) (KDD '20). Association for Computing Machinery, New York, NY, USA, 2608–2616. <https://doi.org/10.1145/3394486.3403311>
- [2] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature Verification Using a "Siamese" Time Delay Neural Network. In *Proceedings of the 6th International Conference on Neural Information Processing Systems* (Denver, Colorado) (NIPS'93). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 737–744.
- [3] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8440–8451. <https://doi.org/10.18653/v1/2020.acl-main.747>
- [4] Alexis Conneau and Guillaume Lample. 2019. *Cross-Lingual Language Model Pretraining*. Curran Associates Inc., Red Hook, NY, USA.
- [5] Miao Fan, Jiacheng Guo, Shuai Zhu, Shuo Miao, Mingming Sun, and Ping Li. 2019. MOBIOUS: Towards the Next Generation of Query-Ad Matching in Baidu's Sponsored Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Anchorage, AK, USA) (KDD '19). Association for Computing Machinery, New York, NY, USA, 2509–2517. <https://doi.org/10.1145/3292500.3330651>
- [6] 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 and Data Mining* (Virtual Event, Singapore) (KDD '21). Association for Computing Machinery, New York, NY, USA, 3181–3189. <https://doi.org/10.1145/3447548.3467101>
- [7] Yiqun Liu, Kaushik Rangadurai, Yunzhong He, Siddarth Malreddy, Xunlong Gui, Xiaoyi Liu, and Fedor Borisjuk. 2021. Que2Search: Fast and Accurate Query and Document Understanding for Search at Facebook. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Virtual Event, Singapore) (KDD '21). Association for Computing Machinery, New York, NY, USA, 3376–3384. <https://doi.org/10.1145/3447548.3467127>
- [8] Wenhao Lu, Jian Jiao, and Ruofei Zhang. 2020. TwinBERT: Distilling Knowledge to Twin-Structured Compressed BERT Models for Large-Scale Retrieval. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (Virtual Event, Ireland) (CIKM '20). Association for Computing Machinery, New York, NY, USA, 2645–2652. <https://doi.org/10.1145/3340531.3412747>
- [9] Alessandro Magnani, Feng Liu, Suthee Chaidaroon, Sachin Yadav, Praveen Reddy Suram, Ajit Puthenpuhussery, Sijie Chen, Min Xie, Anirudh Kashi, Tony Lee, and Ciya Liao. 2022. Semantic Retrieval at Walmart. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 3495–3503. <https://doi.org/10.1145/3534678.3539164>
- [10] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- [11] Amanpreet Singh, Vedanuj Goswami, Vivek Natarajan, Yu Jiang, Xinlei Chen, Meet Shah, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. 2020. MMF: A multimodal framework for vision and language research. <https://github.com/facebookresearch/mmf>.
- [12] Yuxin Tian, Xueqing Deng, Yi Zhu, and Shawn Newsam. 2020. Cross-Time and Orientation-Invariant Overhead Image Geolocalization Using Deep Local Features. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (WACV).
- [13] Yuxin Tian, Shawn Newsam, and Kofi Boakye. 2022. Image Search with Text Feedback by Additive Attention Compositional Learning.
- [14] Jinpeng Wang, Jieming Zhu, and Xiuqiang He. 2021. Cross-Batch Negative Sampling for Training Two-Tower Recommenders. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 1632–1636. <https://doi.org/10.1145/3404835.3463032>
- [15] Jiancan Wu, Xiang Wang, Xingyu Gao, Jiawei Chen, Hongcheng Fu, Tianyu Qiu, and Xiangnan He. 2022. On the Effectiveness of Sampled Softmax Loss for Item Recommendation. [arXiv:2201.02327](https://arxiv.org/abs/2201.02327) [cs.LG]
- [16] Yuqing Xie, Taesik Na, Xiao Xiao, Saurav Manchanda, Young Rao, Zhihong Xu, Guanghua Shu, Esther Vasiete, Tejaswi Tenneti, and Haixun Wang. 2022. An Embedding-Based Grocery Search Model at Instacart. [arXiv preprint arXiv:2209.05555](https://arxiv.org/abs/2209.05555) (2022).
- [17] Ji Yang, Xinyang Yi, Derek Zhiyuan Cheng, Lichan Hong, Yang Li, Simon Wang, Taibai Xu, and Ed H. Chi. 2020. Mixed Negative Sampling for Learning Two-tower Neural Networks in Recommendations.
- [18] Licheng Yu, Jun Chen, Animesh Sinha, Mengjiao Wang, Yu Chen, Tamara L. Berg, and Ning Zhang. 2022. CommerceMM: Large-Scale Commerce MultiModal Representation Learning with Omni Retrieval. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 4433–4442. <https://doi.org/10.1145/3534678.3539151>
- [19] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R. Salakhutdinov, and Alexander J. Smola. 2017. Deep Sets. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf>
- [20] Jianjin Zhang, Zheng Liu, Weihao Han, Shitao Xiao, Ruicheng Zheng, Yingxia Shao, Hao Sun, Hanqing Zhu, Premkumar Srinivasan, Denvy Deng, Qi Zhang, and Xing Xie. 2022. Uni-Retriever: Towards Learning The Unified Embedding Based Retriever in Bing Sponsored Search. [arXiv:2202.06212](https://arxiv.org/abs/2202.06212) [cs.LG]