

A Personalized Dense Retrieval Framework for Unified Information Access

Hansi Zeng*
University of Massachusetts Amherst
USA
hzeng@cs.umass.edu

Surya Kallumadi
Lowe's Companies, Inc.
USA
surya@ksu.edu

Zaid Alibadi
Lowe's Companies, Inc.
USA
zalibadi@email.sc.edu

Rodrigo Nogueira
University of Campinas
Brazil
rfn@unicamp.br

Hamed Zamani
University of Massachusetts Amherst
USA
zamani@cs.umass.edu

ABSTRACT

Developing a universal model that can efficiently and effectively respond to a wide range of information access requests—from retrieval to recommendation to question answering—has been a long-lasting goal in the information retrieval community. This paper argues that the flexibility, efficiency, and effectiveness brought by the recent development in dense retrieval and approximate nearest neighbor search have smoothed the path towards achieving this goal. We develop a generic and extensible dense retrieval framework, called UIA, that can handle a wide range of (personalized) information access requests, such as keyword search, query by example, and complementary item recommendation. Our proposed approach extends the capabilities of dense retrieval models for ad-hoc retrieval tasks by incorporating user-specific preferences through the development of a personalized attentive network. This allows for a more tailored and accurate personalized information access experience. Our experiments on real-world e-commerce data suggest the feasibility of developing universal information access models by demonstrating significant improvements even compared to competitive baselines specifically developed for each of these individual information access tasks. This work opens up a number of fundamental research directions for future exploration.

CCS CONCEPTS

• **Information systems** → **Document representation; Learning to rank.**

KEYWORDS

Dense Retrieval; Personalization; Unified Information Access

*A part of this work was done while Hansi Zeng was an intern at Lowe's.

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.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9408-6/23/07...\$15.00
<https://doi.org/10.1145/3539618.3591626>

ACM Reference Format:

Hansi Zeng, Surya Kallumadi, Zaid Alibadi, Rodrigo Nogueira, and Hamed Zamani. 2023. A Personalized Dense Retrieval Framework for Unified Information Access. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539618.3591626>

1 INTRODUCTION

Information access systems, such as search engines and recommender systems, play a key role in the Web ecosystem. Often, a combination of information access systems is required to satisfy different information needs of users. For instance, e-commerce websites provide both search and recommendation functionalities to their users. Decades of research have been dedicated to developing specialized models for each information access scenario. For instance, lexical and semantic matching models developed for document retrieval in response to keyword queries are fundamentally different from the models used in state-of-the-art recommender systems. Inspired by Belkin and Croft's early examination of information retrieval and filtering systems [3], Zamani and Croft [36, 38] hypothesized that joint modeling and optimization of search engines and collaborative filtering can lead to higher generalization and can improve both search and recommendation quality. In this paper, we take this hypothesis even further *towards* developing **universal information access**:

a unified model that can efficiently and effectively perform different information access functionalities

Successful implementation of universal information access could close the gap between the research communities working on various aspects of information access, such as the IR and RecSys communities. Potentially, all information access functionalities can benefit from universal modeling and optimization (i.e., transferring knowledge across information access functionalities). Moreover, such universal models can potentially reduce the engineering efforts required for developing and maintaining multiple information access systems separately. All of these points highlight the importance of research towards universal information access.

Thanks to the flexibility of dense retrieval models and their state-of-the-art performance in various information retrieval scenarios [8, 17, 34], we develop UIA—a novel dense retrieval framework for unified information access. UIA follows a bi-encoder architecture.

The first encoder learns a dense vector for a given information access request (e.g., an item¹ specified by the user or a search query submitted by a user) and an *information access functionality* that the user is interested in (e.g., recommending complementary items or retrieving relevant items). Since a wide range of information access functionalities require personalization, UIA adjusts the obtained dense vectors based on the user’s historical interactions by introducing a novel *Attentive Personalization Network (APN)*. APN performs both content-based and collaborative personalization based on the user’s past interactions with all the information access functionalities (e.g., user’s past interactions with a search engine is also used for personalizing recommendation results). The second encoder learns a dense representation of the each information item in the collection. After training, these representations can be computed offline and an approximate nearest neighbor algorithm can be employed for efficient retrieval. Our proposed approach, the UIA model, utilizes a non-personalized pre-training and personalized fine-tuning strategy to improve representation learning. Additionally, we draw upon recent advancements in dense retrieval by incorporating the strategy of sampling hard negative items from various sources [24, 34, 40]. By combining these training strategies, our approach aims to achieve better performance in various information access scenarios.

To train and evaluate the UIA framework, we focus on three information access functionalities. (1) **Keyword Search**: retrieving relevant items in response to a short textual query (e.g., retrieving ‘iPhone 14 Pro’ for query ‘apple smartphone’), (2) **Query by Example**: retrieving items that are similar to a item specified by the user (e.g., retrieving ‘Google Pixel 7 Pro’ as a similar item to ‘iPhone 14 Pro’), and (3) **Complementary Item Recommendation**: recommending information items that are complementary to a given item specified by the user (e.g., recommending ‘AirPods Pro’ to a user who is interested in ‘iPhone 14 Pro’). UIA uses text description of each information access functionality, thus is *generic* and extensible to a wide range of other information access functionalities.

We evaluate our model on a real-world dataset collected from user interactions with different information access systems on a major e-commerce website.² To improve reproducibility, we also extend our experiments to a dataset constructed from the Amazon ESCI data [25], recently released as part of KDD Cup 2022.³ Extensive experiments demonstrate significant improvements compared to a wide range of competitive baselines for all three information access functionalities. We also demonstrate that up to 45% NDCG@10 improvements can be obtained by jointly modeling all information access functionalities compared to their individual modeling using the same dense retrieval architecture.

To summarize, the main contributions of this work include:

- Developing a generic and extensible dense retrieval framework that performs multiple information access functionalities.

¹In this paper, items refer to unstructured or semi-structured textual documents.

²Even though our both datasets are from the e-commerce domain, the proposed approach is sufficiently generic to be applied to any domain. We are not aware of any other publicly available dataset beyond e-commerce that can be used to evaluate our task.

³KDD Cup 2022: <https://amazonkddcup.github.io/>

- Demonstrating the feasibility of learning a single model that can perform a wide range of information access functionalities.
- Proposing an attentive personalized network with a two-stage training process and negative sampling strategies from various sources to improve personalized dense retrieval performance.
- Evaluating the proposed model on real-world data and demonstrating substantial gain compared to competitive baselines.

We open-source our implementation of the UIA framework to foster research towards developing universal information access.⁴

2 RELATED WORKS

2.1 Personalized Information Access

For one query, different users may have different intents or preferences. Hence, the ability of information access systems to capture the user’s personalized preferences would play an important role in improving user experiences. For personalized search systems, there are several models have been proposed in recent years and have demonstrated search quality improvements compared to non-personalized models. For instance, HEM [2] jointly models the user, query, and item representations using the doc2vec technique [21] and ZAM [1] applies an attention-based method to capture the user dynamic information.

In recommender systems, personalization plays an even bigger role. Early models utilized Collaborative Filtering (CF) [20] to model the user’s personalized preference. Recent studies [15, 18, 31, 32] employ deep neural networks to extract content information from items and seamlessly integrate this information with CF models. However, CF models neglect the order of interactions in user history. To address this issue, models such as GRU4Rec [11], SASRec [16] and BERT4Rec [28] propose deep learning based sequential models to capture the user history preference for next-item prediction.

2.2 Dense Retrieval

Accurate contextual representation of text using large-scale pre-trained language models has led to significant progress in various fields, including information retrieval. Combining these models with efficient approximate nearest neighbor search resulted in the development of dense retrieval models [17, 34]. These models fine-tune the pre-trained language models on the downstream information retrieval task [4] and have shown significant performance improvement over strong lexical matching methods such as BM25 [26]. One line of methods focuses on the optimization of dense retrieval models by producing better negative samples for contrastive loss. For example, DPR [17] uses the BM25 negatives as the source for hard-negative sampling and ANCE [34] applies a self-sampling strategy for negative sampling. Models like RocketQA [24] and Condenser [8] use large batch sizes and have demonstrated that it would be conducive to the stability of training. The other line of methods applies the knowledge distillation techniques that distill the knowledge from re-ranking models to train the dense retrieval models [12, 22, 39]. Adapting dense retrieval models to unseen data is also an active area of research [29].

⁴<https://github.com/HansiZeng/UIA>

Table 1: Properties of the three information access functionalities explored in this paper.

	Information Access Functionality (\mathcal{F})	Information Access Request (\mathcal{R})	User History (\mathcal{H})	Item Information (\mathcal{I})
1	Keyword Search	A short keyword query	User’s past queries and clicks	Content of candidate item
2	Query by Example	Content of an anchor item	User’s past queries and clicks	Content of candidate item
3	Complementary Item Recommendation	Content of an anchor item	User’s past queries and clicks	Content of candidate item

2.3 Joint Search and Recommendation

Recent results show that joint modeling of search and recommendation tasks can lead to better results compared to their individual training. JSR [38] is a framework that uses a task-specific layer on top of a shared representation learning network. It uses multi-task learning for optimization. More recently, SRJGraph [41] extends a similar approach by further applying graph convolution networks to capture higher-order interactions between users, queries, and items. These models are highly specialized for the search and recommendation tasks in hand and they are not simply extendable to other information access functionalities. On the other hand, we propose a dense retrieval model (UIA) that unifies different information access functionalities including search and recommendation. Unlike JSR and SRJGraph, UIA is able to jointly train different tasks without introducing additional parameters. UIA also contains a novel attentive personalization network (APN) that can capture the user’s sequential interaction history and produces significantly better results compared to both JSR and SRJGraph.

3 METHODOLOGY

3.1 Task Formulation

Users can find and access information in a number of different ways. For instance, they can express their needs as a textual query and use a search engine to find relevant information. Alternatively, they can use the outcome of a recommender system without explicitly formulating their needs. This paper studies unified information access: *developing and evaluating a unified model that can efficiently and effectively perform various information access functionalities*. Any information access model can be formulated as a scoring function that takes three input variables:

- (1) **Information Access Request (\mathcal{R}):** it includes information about the current information access request, such as a search query, situational context (e.g., location and time) [37], and short-term context (e.g., session data) [11, 16, 28]. Information Access Request can be empty (i.e., zero-query retrieval) [35].
- (2) **User History (\mathcal{H}):** it includes information about the user who issues \mathcal{R} , such as user’s profile or their long-term interaction history.
- (3) **Candidate Item Information (\mathcal{I}):** it includes the information about the candidate retrieval item, such as the item’s content, author, and source.

In order to develop a unified information access model, we introduce the Information Access Functionality (\mathcal{F}) as the fourth input variable. Without loss of generality, this paper focuses on the following three personalized information access functionalities:⁵

⁵Note that information access functionalities can go beyond these three functionalities, for example, by including multi-modal information (e.g., image queries and video items), contextual requests (e.g., sessions and conversations), and zero query retrieval or recommendation.

- (1) **Keyword Search:** retrieving items in response to a short textual user query.
- (2) **Query by Example:** retrieving items that are similar to a given item specified by a user.
- (3) **Complementary Item Recommendation:** recommending complementary items to a given item specified by a user.

The properties of each of these tasks with respect to the mentioned four input variables are listed in Table 1. We intentionally include two information access functionalities that take identical inputs (i.e., Query by Example and Complementary Item Recommendation) but produce different outputs. They enable us to evaluate the model’s ability in learning different information access functionalities.

Given the input variables introduced above, a unified information access model parameterized by θ for user u at timestamp t can be formally formulated as $f(\mathcal{F}_t^u, \mathcal{R}_t^u, \mathcal{H}_t^u, \mathcal{I}_i; \theta)$, where \mathcal{F}_t^u is a textual description of the information access functionality being applied. For every Information Access Request (\mathcal{R}_t^u) at timestamp t , User History (\mathcal{H}_t^u) denotes a set of interactions that the user u had prior to t . Therefore, $\mathcal{H}_t^u = \{(\mathcal{F}_1^u, \mathcal{R}_1^u, \mathcal{I}_1^u), (\mathcal{F}_2^u, \mathcal{R}_2^u, \mathcal{I}_2^u), \dots, (\mathcal{F}_{t-1}^u, \mathcal{R}_{t-1}^u, \mathcal{I}_{t-1}^u)\}$ is a set of all the past $t - 1$ interactions of the user u , each is represented by a triplet of the past user’s request, the information access functionality that has been used, and the item that the user interacted with. For simplicity, in this paper, Candidate Item Information for the i^{th} item (\mathcal{I}_i) only includes the item’s textual content. As mentioned in Table 1, each \mathcal{R}_t^u either represents a keyword query or the textual content of a given anchor item, depending on \mathcal{F}_t^u .

3.2 Overview of the UIA Framework

This paper proposes UIA, a dense retrieval framework for unified information access. A high-level overview of the UIA framework is depicted in Figure 1. UIA follows a bi-encoder architecture. For each user u at timestamp t , the first encoder takes Information Access Functionality (\mathcal{F}_t^u) and Request (\mathcal{R}_t^u) and produces a dense representation of the request. It then uses a novel *Attentive Personalization Network* to further enhance the request representation through both collaborative and content-based personalization using the user’s historical interaction data. The second encoder takes the content of a candidate item and produces a latent dense representation, which is fed to a feed-forward network in order to adjust the representations based on the personalized request vector. UIA uses the inner product to compute the similarity of encoded request and item. Therefore, both encoders should have the same output dimensionality. In the following subsections, we present the implementation and optimization of UIA.

3.3 UIA Architecture

As demonstrated in Figure 1, UIA consists of four major components: (1) Request Encoding, (2) Item Encoding, (3) User History Selection & Encoding, and (4) Attentive Personalization Network.

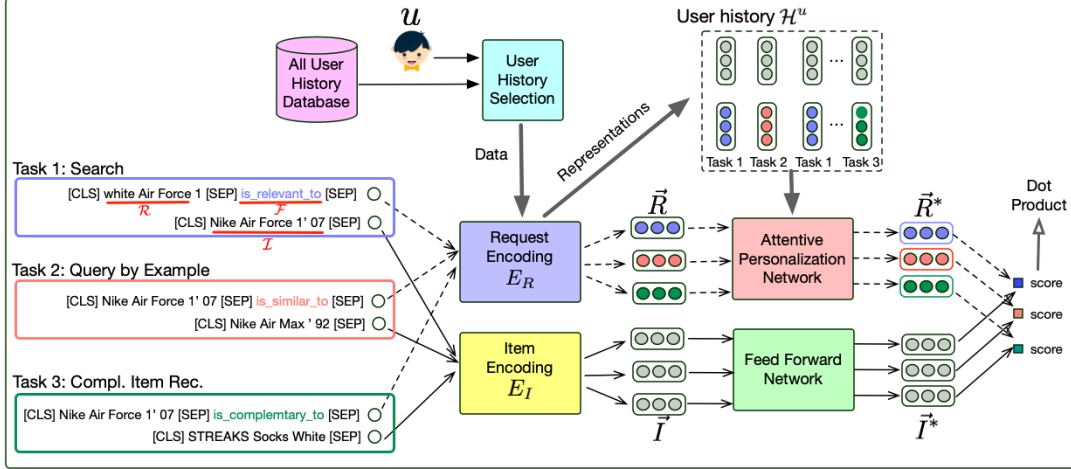


Figure 1: A high-level overview of the UIA framework.

Request Encoding. To implement the request encoder, we use a pre-trained large language model (denoted by E) that often feeds a subword tokenization of the input to an embedding layer followed by a number of Transformer layers and produces a dense vector representation for each input token. In this paper, we use BERT-base [7] to encode each information access functionality (\mathcal{F}_t^u) and request (\mathcal{R}_t^u), as follows:

$$\vec{R}_t^u = E_{\mathcal{R}}([\text{CLS}] \mathcal{R}_t^u [\text{SEP}] \mathcal{F}_t^u [\text{SEP}])$$

An example of the Request Encoding input is shown in Figure 1.

Candidate Item Encoding. Similar to the request encoder, this component also uses a pre-trained BERT-base model for representing each candidate item I_i . It represents the item’s information (i.e., content) by encoding the following input:

$$\vec{I}_i = E_{\mathcal{I}}([\text{CLS}] I_i [\text{SEP}])$$

For both the request and candidate item encoders, we use the [CLS] token representation as the encoder’s output.

User History Selection and Encoding. For content-based personalization of the request encoding \vec{R}_t^u , this paper simply uses the last N interactions of the user, i.e., $\{(\mathcal{F}_{t-N}^u, \mathcal{R}_{t-N}^u, \mathcal{I}_{t-N}^u), \dots, (\mathcal{F}_{t-1}^u, \mathcal{R}_{t-1}^u, \mathcal{I}_{t-1}^u)\}$. We produce two encoding vectors for each of the user’s past interactions: one for the past information access functionality and request ($E_{\mathcal{R}}([\text{CLS}] \mathcal{R}_{t'}^u [\text{SEP}] \mathcal{F}_{t'}^u [\text{SEP}]) : \forall t - N \leq t' \leq t - 1$) and another for the item that the user interacted with ($E_{\mathcal{I}}([\text{CLS}] \mathcal{I}_{t'}^u [\text{SEP}]) : \forall t - N \leq t' \leq t - 1$). UIA uses the same $E_{\mathcal{R}}$ and $E_{\mathcal{I}}$ encoders to ensure that the user history encodings are represented in the same space as the model’s input for timestamp t . All parameters in UIA, including the encoder models, are trained *end-to-end*, and thus they get updated during training.

Therefore, the encoded history would be a set of $2 \times N$ vectors as follows: $\{(\vec{R}_{t-N}^u, \vec{I}_{t-N}^u), (\vec{R}_{t-N+1}^u, \vec{I}_{t-N+1}^u), \dots, (\vec{R}_{t-1}^u, \vec{I}_{t-1}^u)\}$

Attentive Personalized Network. To produce personalized representation of \vec{R}_t^u , we propose a novel Attentive Personalization Network (APN) that performs both content-based and collaborative personalization. For content-based personalization, it learns attention weights from encodings of the user’s past N interactions to

the current request. Note that for each information access functionality, it is able to use the user’s past interactions with all functionalities (e.g., both search and recommendation). For collaborative personalization, it learns a latent representation for each user and information access functionality based on all the past interactions.

A high-level overview of APN architecture is presented in Figure 2. At each timestamp t , APN takes the request encoding \vec{R}_t and the last N user history encodings $\{(\vec{R}_{t-N}^u, \vec{I}_{t-N}^u), \dots, (\vec{R}_{t-1}^u, \vec{I}_{t-1}^u)\}$. It converts the user history encodings to two matrices: (1) $H_t^u \in \mathbb{R}^{N \times d}$ whose rows are equal to the past request encodings, and (2) $C_t^u \in \mathbb{R}^{N \times d}$ whose rows are equal to the past interacted item (clicked) encodings.

APN consists of N_h attention heads, where the j^{th} attention function contains three parameter matrices $\theta_j^Q \in \mathbb{R}^{d \times l}$, $\theta_j^K \in \mathbb{R}^{d \times l}$, and $\theta_j^V \in \mathbb{R}^{d \times l_v}$. It learns attention weights from the user’s interaction history to their current request. Therefore, each APN layer computes the following three vectors:

$$Q_j = \vec{R}_t^u \cdot \theta_j^Q, \quad K_j = H_t^u \cdot \theta_j^K, \quad V_j = C_t^u \cdot \theta_j^V$$

Therefore, $Q_j \in \mathbb{R}^{1 \times l}$, $K_j \in \mathbb{R}^{N \times l}$, and $V_j \in \mathbb{R}^{N \times l_v}$. Based on these three matrices, APN uses the following attention mechanism that is similar to Transformer [30]:

$$\text{Attn}(Q_j, K_j, V_j) = \text{softmax}\left(\frac{Q_j K_j^T}{\sqrt{l}}\right) V_j$$

The output of all attention functions are then concatenated: $\text{concat}(\{\text{Attn}(Q_j, K_j, V_j)\}_{j=1}^{N_h}) \in \mathbb{R}^{1 \times N_h l_v}$. Following Transformer architecture and inspired by residual and layer norm in neural networks, we feed this matrix to an Add & Norm layer.

For **collaborative personalization**, the Attentive Personalization Network also learns a user embedding matrix $\mathbf{E}_{\mathcal{U}} \in \mathbb{R}^{|\mathcal{U}| \times l_u}$, where $|\mathcal{U}|$ denotes the number of users and l_u is the user embedding dimensionality. For each user u , we select the associated user embedding vector \vec{u} from $\mathbf{E}_{\mathcal{U}}$ (i.e., user embedding lookup). To distinguish between the user behavior dealing with different information access functionalities, we also learn an embedding matrix

for information access functionalities $\mathbf{E}_{\mathcal{F}} \in \mathbb{R}^{|\mathcal{F}| \times l_f}$, where $|\mathcal{F}|$ denotes the number of information access functionalities (e.g., three in our case) and l_f is the functionality embedding dimensionality. At timestamp t , we select the associated functionality embedding vector \vec{f} from $\mathbf{E}_{\mathcal{F}}$.

We later concatenate the output of Add & Norm layer with \vec{u} and \vec{f} and feed this vector to a feed-forward layer with non-linear activation (ReLU). This produces a personalized representation of the current request, denoted by \vec{R}_t^{u*} .

Since the personalization component is only applied to the request representation, we use a feed-forward network for the candidate item vectors to adjust their representations with the new personalized semantic space and obtain \vec{I}_i^* (See Figure 1).

3.4 UIA Optimization

We propose a two stage optimization process: (1) non-personalized pre-training, and (2) personalized fine-tuning. The reason is that real-world systems often deals with a large number of new and cold-start users with no or limited historical interactions. Thus, personalization cannot help, yet we can use their data for non-personalized pre-training.

Non-Personalized Pre-Training. We construct a non-personalized training set by aggregating the training data across all users. For each training instance k is in the form of $(\mathcal{F}_k, \mathcal{R}_k, \mathcal{I}_k, \mathcal{Y}_k)$ where \mathcal{Y}_k is the ground truth label. We get the output vector of the Request Encoder and Candidate Item Encoder, i.e., \vec{R}_k and \vec{I}_k , respectively. We compute the non-personalized matching of the request and the candidate item using dot product: $\vec{R}_k \cdot \vec{I}_k$.

The training data only contains positive instances (i.e., user interactions), thus appropriate negative sampling is required. We apply a two-phase negative sampling and training strategy.

Phase 1: for each request in training data, we randomly sample negative items from the top 200 items retrieved by BM25. We set the ratio of negative samples to positive training instances to 1. We then train the model using a cross entropy loss function. Note that in addition to BM25 negative, we also use in-batch negatives.

Phase 2: Once the model is trained in Phase 1, we use the trained Candidate Item Encoder E_I to encode all items in the collection and create an approximate nearest neighbor (ANN) index using the Faiss library [13]. The constructed index is then used to retrieve items for each request in the training data and random negatives are sampled from the top 200 retrieved documents. This self-negative sampling strategy has successfully been used in a number of dense retrieval models, such as ANCE [34] and RANCE [23]. We set the ratio of negative samples to positive training instances to 1. Similar to Phase 1, we re-train the model using a cross-entropy loss function that also uses in-batch negatives.

Personalized Fine-Tuning. In non-personalized pre-training, only the parameters of $E_{\mathcal{R}}$ and E_I are adjusted. We then add the personalization part of the framework and re-create the training data to include the user information and their past interactions. We then use the personalized representation of each request and candidate item in the training data, i.e., \vec{R}_k^* and \vec{I}_k^* respectively (see Figure 1). Then, we use dot product to compute their matching score: $\vec{R}_k^* \cdot \vec{I}_k^*$.

Table 2: Statistics of the datasets constructed for this study.

Property	Lowe's Data	Amazon ESCI
# unique users	893,619	-
# unique queries	953,773	68,139
# items in the collection	2,260,878	1,216,070
# interactions for keyword search	4,075,996	874,087
# interactions for query by example	968,778	2,254,779
# interactions for complementary item rec.	329,992	303,481

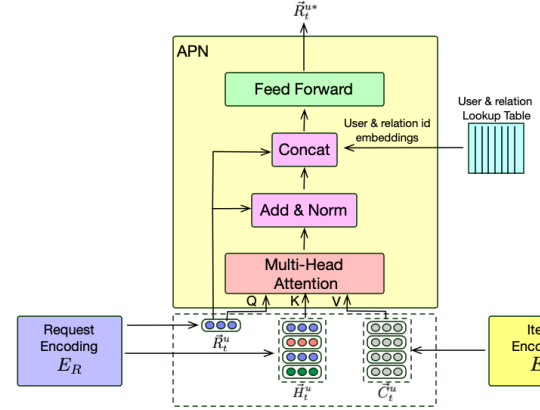


Figure 2: The Attentive Personalization Network (APN).

For negative sampling, we use BM25 results in addition to in-batch negatives (similar to Phase 1 in non-personalized pre-training). We use cross entropy loss function for training.

4 EXPERIMENTS

4.1 Experiment Settings

4.1.1 Dataset. Evaluating the proposed model is challenging, since there is no public dataset that provide large-scale training data for multiple information access functionalities. Therefore, we made significant efforts to create two datasets. The first one is a private real-world data obtained from the Lowe's website—the second-largest hardware chain in the world. The second dataset is built from the public Amazon ESCI data [25] recently released as part of KDD Cup 2022. Table 2 summarizes the statistics of the datasets.

The **Lowe's dataset** is a large-scale e-commerce data containing over 5.3 million user interactions for over obtained from over 890K unique users. The item collection in this data includes over 2.2 million products. The Lowe's dataset contains three types of user-item interaction data: (1) Keyword Search (i.e., query-item click data); (2) Query by Example (i.e., item-item click data for similar items); and (3) Complementary Item Recommendation (i.e., item-item click data for complementary item pairs). This dataset also includes the anonymized user ID and timestamp associated with each interaction.

The **Amazon ESCI dataset** [25] is adopted from KDD Cup 2022 - Task 2 that contains 3 tasks related to product information access. We modify the data to fit our experimental setting. Each data point in KDD Cup 2022 - Task 2 is a (query, item, label) triplet. The label is one of the following four classes: (1) Exact (E) means that the item is relevant to (an exact match for) the query; (2) Substitute (S) means that the item is related to the query but not an exact match (partially

Table 3: Experimental results on the Lowe’s dataset. “-” means the model is not applicable to the search task. Superscripts Δ and \blacktriangle refer to significant improvements compared to all task-specific training baselines and all baselines (including joint training), respectively ($p_value < 0.01$).

Model	Keyword Search			Query by Example			Complementary Item Rec.		
	MRR	NDCG	Recall	MRR	NDCG	Recall	MRR	NDCG	Recall
BM25	0.089	0.095	0.367	0.153	0.167	0.584	0.016	0.014	0.111
Task-Specific Training									
NCF	-	-	-	0.132	0.147	0.351	0.117	0.118	0.236
DPR	0.188	0.192	0.578	0.171	0.180	0.598	0.153	0.156	0.487
ANCE	0.193	0.199	0.582	0.176	0.188	0.601	0.159	0.158	0.494
RocketQA	0.201	0.207	0.595	0.189	0.204	0.613	0.174	0.176	0.507
Context-Aware DPR	0.324	0.377	0.848	0.311	0.356	0.860	0.278	0.283	0.707
Context-Aware ANCE	0.332	0.385	0.856	0.317	0.361	0.866	0.289	0.292	0.714
Context-Aware RocketQA	0.335	0.389	0.861	0.326	0.369	0.874	0.300	0.304	0.723
SASRec++	-	-	-	0.305	0.347	0.836	0.271	0.264	0.695
BERT4Rec++	-	-	-	0.314	0.354	0.851	0.283	0.279	0.703
Joint Training									
JSR	0.324	0.379	0.853	0.349	0.380	0.878	0.325	0.317	0.760
JSR+BERT4Rec++	0.337	0.394	0.871	0.415	0.479	0.919	0.421	0.419	0.820
SRJGraph	0.336	0.392	0.874	0.416	0.478	0.921	0.423	0.420	0.822
UIA	0.340Δ	0.399\blacktriangle	0.880\blacktriangle	0.433\blacktriangle	0.495\blacktriangle	0.945\blacktriangle	0.438\blacktriangle	0.432\blacktriangle	0.836\blacktriangle

relevant); (3) Complement (C) means the item is not relevant to the query but can complement the relevant item with label E; and (4) Irrelevant (I). Let $I_E(q)$, $I_S(q)$, and $I_C(q)$ respectively denote a set of all items with a label E, label S, and label C for query $q \in Q$. We used the following procedure to construct our three datasets: (1) Keyword Search: $\{(q, i) : \forall q \in Q \wedge i \in I_E(q)\}$, (2) Query by Example: $\{(i_1, i_2) : \forall q \in Q \wedge i_1 \in I_E(q) \wedge i_2 \in I_S(q)\}$, and (3) Complementary Item Recommendation: $\{(i_1, i_2) : \forall q \in Q \wedge i_1 \in I_E(q) \wedge i_2 \in I_C(q)\}$.

Note that the Amazon ESCI dataset does not contain user identifiers, hence no personalization is performed for this dataset. To the best of our knowledge, there is no public dataset with user identifiers that includes multiple information access functionalities.

4.1.2 Evaluation Protocols. For the Lowe’s dataset, we use a leave-last-out data splitting strategy, which has been widely used in the literature [5, 9, 14]. For each of the three information access functionalities, we use the user’s most recent interaction for testing, their second most recent interaction for validation, and the rest for training. This results in over 890K interactions in each of the test and validation sets. This realistic data splitting strategy enables us to evaluate the model’s ability based on the future interactions that the user may have with the system. Since there is no user identifiers or timestamp available in the Amazon ESCI dataset, we randomly select 80% of requests (e.g., query text or query item) for training, 10% for validation, and the remaining 10% for testing.

To evaluate the models, we report performance in terms of a wide range of metrics: MRR@10 (MRR for short), NDCG@10 (NDCG for short), and Recall@50 (Recall for short). We use two-tailed paired t-test with Bonferroni correction to identify statistically significant improvements ($p_value < 0.01$).

4.1.3 Implementation Details. We use BERT-base [7] available on HuggingFace [33] as the pre-trained language model in our models. For the Lowe’s dataset, the pre-trained weights is loaded

from the checkpoint ⁶ and for Amazon ESCI dataset (which is smaller), the pre-trained weights are obtained from the checkpoint ⁷. We use the model’s performance on the validation set in terms of NDCG to select the hyper-parameters.

We chose the number of training epoch from [8, 12, 16, 24, 48]. We set the user’s historical interactions (N) to 5. The batch size is empirically set to 384. The learning rate for pre-training and personalized fine-tuning is empirically set to $7e^{-6}$ and $7e^{-5}$, respectively. For personalized fine-tuning, we only keep users with at least 10 interactions in search and query by example tasks, and at least 5 interactions in the complementary item recommendation task. The hidden dimension d is 768, number of heads N_h is 12, hidden dimension of key and value in each head are $l = l_v = 64$. The dimension of user embedding is $l_u = 128$ and of functionality embedding is $l_f = 64$. We use Adam [19] as the optimizer.

4.1.4 Baselines. We use a wide range of baselines in our experiments, from term matching models to dense retrieval models. We also include collaborative and sequential recommendation baselines, when appropriate. The baselines are listed below:

- **BM25** [26]: This is a simple yet effective bag-of-word retrieval model that uses query term frequency, inverse document frequency, and document length to compute relevance scores.
- **NCF** [10]: This is an effective collaborative filtering model, which combines generalized matrix factorization and a multi-layer perceptron approach for recommendation. It only learns from item-item interactions and cannot be applied to keyword search tasks.
- **DPR** [17]: DPR is a dense retrieval model that samples negative documents from the items retrieved by BM25 in addition to in-batch negative sampling. DPR only uses the last request (query text or query item) and does not perform personalization.

⁶<https://huggingface.co/bert-base-uncased>

⁷<https://huggingface.co/sentence-transformers/msmarco-bert-base-dot-v5>

Table 4: Experimental results on the Amazon ESCI dataset. Superscript [▲] refers to significant improvements compared to all baselines (p_value < 0.01).

Model	Keyword Search			Query by Example			Complementary Item Rec.		
	MRR	NDCG	Recall	MRR	NDCG	Recall@50	MRR	NDCG	Recall
BM25	0.513	0.351	0.494	0.017	0.011	0.084	0.030	0.032	0.165
Task-Specific Training									
DPR	0.505	0.347	0.511	0.235	0.174	0.527	0.434	0.450	0.838
ANCE	0.522	0.354	0.519	0.237	0.178	0.531	0.431	0.443	0.825
RocketQA	0.526	0.357	0.525	0.244	0.185	0.538	0.445	0.458	0.847
Joint Training									
JSR	0.528	0.355	0.527	0.243	0.192	0.536	0.477	0.484	0.853
SRJGraph	0.526	0.351	0.522	0.241	0.187	0.540	0.479	0.488	0.855
UIA	0.532[▲]	0.360[▲]	0.533[▲]	0.251[▲]	0.199[▲]	0.543[▲]	0.490[▲]	0.493[▲]	0.868[▲]

- **Context-Aware DPR:** We extend the DPR model to include user history (personalization). To this aim, we simply concatenate the current user request with their past interactions, separated by a [SEP] token and feed it to the query encoder.
- **ANCE [34]:** ANCE is an effective dense retrieval model that uses the model itself to mine hard negative samples. Similar to DPR, ANCE is not capable of personalization, so we also include a **Context-Aware ANCE** using a similar approach used for Context-Aware DPR.
- **RocketQA [24]:** RocketQA is a state-of-the-art dense retrieval model. It utilizes the large batch size and denoised negative samples for more robust contrastive learning. Similar to DPR and ANCE, RocketQA is not capable of personalization, so we also include a **Context-Aware RocketQA** for user sequential modeling.
- **BERT4Rec++:** BERT4Rec [28] is a sequential recommendation model that represents the user interaction history using BERT for predicting the next item. The original BERT4Rec model takes item IDs and predict the next item ID in the sequence. We improve BERT4Rec by encoding item content too. We use BERT for content embedding. We call this approach BERT4Rec++.
- **SASRec++:** SASRec [16] is a sequential recommendation model uses the self-attention mechanism to identify which items are “relevant” to the user interaction history for the next item prediction. It cannot be used for the keyword search tasks. Similar to the last baseline, we use BERT for content embedding and call it as model SASRec++.
- **JSR [38]:** This is a neural framework that jointly learns the search and recommendation tasks. Each task has a task-specific layer over the base shared network.
- **JSR + BERT4Rec++:** The original JSR uses the user ID to encode user information. We improve the JSR performance by using the representation from BERT4Rec++ to encode the user content.
- **SRJGraph [41]:** This is a recent framework based on neural graph convolution that jointly models the search and recommendation tasks.

We use the provided public code to implement BM25, NCF, DPR, ANCE, RocketQA, SASRec++ and BERT4Rec. We implemented JSR and SRJGraph which don’t have public implementations. Based on how the training data is used, we classify the training of models into two categories: (1) task-specific training; (2) joint training.

Table 5: Ablation study results on the Lowe’s dataset in terms of NDCG for keyword search (Search), query by example (QBE), and complementary item recommendation (CIR). Superscripts [▽] denote significantly lower performance compared to UIA (p_value < 0.01).

	Search	QBE	CIR
- UIA	0.399	0.495	0.432
1 w/o encoding \mathcal{F}	0.371 [▽]	0.347 [▽]	0.284 [▽]
2 w/o joint optimization	0.391 [▽]	0.369 [▽]	0.298 [▽]
3 w/o APN	0.207 [▽]	0.214 [▽]	0.176 [▽]
4 w/o combined content-based personalization	0.397	0.482 [▽]	0.411 [▽]
5 w/o collaborative personalization	0.378 [▽]	0.507	0.419 [▽]

Task-specific training means that the baseline is only trained on the target task (i.e. for evaluation on the search task, they are trained on the search data only). Joint training baselines can access all tasks’ data as UIA does. We use the same BERT backbone with the same initial pre-trained weights to train all dense retrieval baselines (DPR, ANCE, RocketQA, and ours). All context-aware variants consume historical data as their input in reverse chronological order. For SASRec++ and BERT4Rec++, the number of additional transformer layers is chosen from [1,2,4] and each transformer layer contains 12 heads and each head’s hidden dimension is 64. The task-specific layer for JSR and SRJGraph is a single dense layer with hidden dimension 768 and we initiated their share networks by the same BERT model as dense retrieval baselines. The learning rate for all baselines is chosen from [1e-4, 7e-5, 1e-5, 7e-6] based on dev set results.

4.2 Experimental Results

4.2.1 Main Results. The performance of UIA and the baselines on the Lowe’s dataset for all the three information access functionalities is reported in Table 3. The results demonstrate that BM25 and NCF perform poorly compared to the deep learning based models. Note that NCF cannot be used for the keyword search task, as it is a collaborative filtering approach. We also observe that context-aware variation of dense retrieval models substantially outperform the original DPR and ANCE models. This demonstrates the importance of personalization for information access in e-commerce. Besides, BERT4Rec++ has better results than SASRec++, which implies that BERT that is a stack of multiple Transformer layers is

capable of better capturing user history behaviors. Among models that are trained on a corresponding single task, Context-Aware ANCE achieves the best performance.

Table 3 also shows that joint training models perform better than task-specific models, especially for query by example and complementary item recommendation tasks that have substantially smaller training data compared to the keyword search task. All in all, UIA outperforms all the baselines. The improvements are statistically significant in nearly all cases, except for MRR in Keyword Search which is only significant compared to the task-specific models.

We extend our experiments to the Amazon ESCI dataset and report the result in Table 4. Note that ESCI does not contain user identifiers, thus the baselines that cannot perform without personalization are omitted from Table 4. In this dataset, complementary item recommendation benefit the most from joint training. The reason is due to the size of the dataset for this information access functionality. It is smaller than the other two functionalities, thus it benefit the most from knowledge transfer across tasks. These results suggest that UIA again performs better than all the baselines. The improvements are generally smaller than the Lowe’s dataset and it is due to the lack of personalization in UIA for this dataset.

We find a concurrent work [27] which aims to create a universal encoder for all retrieval tasks. We utilized their publicly available pre-trained checkpoint for zero-shot evaluation on the Amazon ESCI dataset, however, its performance was inferior to BM25. As an example, the MRR@10 in the complementary recommendation task was 0.223 compared to 0.230 by BM25. We surmise that this disparity could be due to the method being trained on datasets that prioritize semantic text similarity [4, 6] and overlooks complementary or substitutive item matching signals.

4.2.2 Ablation Study. To demonstrate the impact of each novel component used in UIA, we conduct a thorough ablation study on the Lowe’s dataset that contain temporal information as well as anonymized user identifiers. For the sake of space, we only report NDCG@10 values in Table 5. We make the following observations:

From Row 1, encoding the information access functionality (i.e., \mathcal{F}) is found crucial. The reason is that the model does not know what items are expected to be retrieved given the input request. Thus, it behaves similarly across all information access functionalities. The performance gain in query by example and complementary item recommendation is substantially higher than in keyword search. There are two reasons: (1) the search data in Lowe’s dataset is at least 300% larger than each of the other information access functionalities (see Table 2); and (2) the input to both query by example and complementary item recommendation are identical.

From Row 2, we can conclude that all three information access functionalities benefit from joint optimization. The improvement observed in keyword search task is marginal, but 34% and 45% relative improvements are observed in query by example and complementary item recommendation, respectively. Note that Lowe’s dataset contains fewer data points for these two tasks and thus they can substantially benefit from joint optimization.

Row 3 shows that personalization (both content-based and collaborative) using APN has a significant impact on the models performance. Using APN in UIA leads to 93%, 72%, and 145% NDCG@10 improvements in keyword search, query by example, and complementary item recommendation, respectively. These improvements

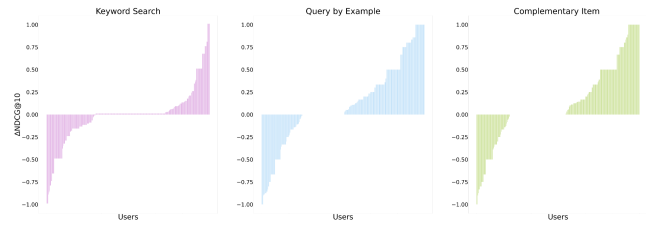


Figure 3: Δ NDCG (sorted) at user level between UIA and TS-UIA (i.e., UIA without joint training) on the Lowe’s dataset.

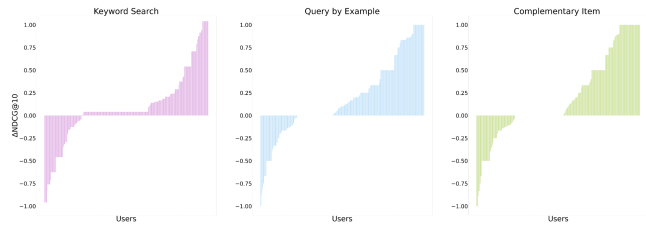


Figure 4: Δ NDCG (sorted) at user level between UIA and UIA without APN on the Lowe’s dataset.

are mostly coming from the content-based personalization. That being said, in the following, we show that APN also benefits from using the combination of all interaction history as well as collaborative personalization.

Row 4 reports the results for UIA where for each information access functionality, only the associated user history is used for personalization. For example, only past search interactions of users are used for the search functionality, as opposed to UIA that uses all the past interactions combined. From this result, we observe that UIA benefits from including historical user interactions from all three information access functionalities. Again, query by example and complementary item recommendation benefit the most, due to their smaller training data size (and thus fewer historical user interactions) compared to the search task.

Finally, Row 5 demonstrates the results for UIA without collaborative personalization, i.e., without user and relation embedding in APN’s final layer (see Figure 2). We show that adding user and relation embedding does not improve the model performance consistently across information access functionalities. Collaborative personalization is found to be helpful in keyword search and complementary item recommendation but not in query by example. The possible explanation might be that the structural information provided by user-relation-item ids is conducive to identify the item-item complementary and query-item relevance relations, but content information which is important in query by example (finding similar items) can be better distilled from content-based personalization without using the user embeddings.

4.2.3 The Impact of Joint Training. Not all users benefit from joint training, since many of them may only use only one of the information access functionalities. To further investigate the effect of joint modeling, we compare the results between UIA and its task-specific (or single task) variation (i.e., the same model only trained on the target information access functionality, denoted by TS-UIA) at user level. We compute the performance of models in

Table 6: The top 1 retrieved item between UIA and TS-UIA. The item is represented as the concatenation of title and its category separated by ";". ✓ means the retrieved item is relevant and ✗ means irrelevant.

Query or query_item	1st Ranked Item by UIA	1st Ranked Item by TS-UIA	Error Reason
Keyword Search: 6x6 privacy fence	Freedom Brighton 6-ft H x 6-ft W White Vinyl Flat-top Fence Panel ; Vinyl Fencing ✓	HOFT Solutions HOFT Kit A6- One 73 in. End Post Black and Hardware ; Outdoor Privacy Screens ✗	Only part of aspect matching
Keyword Search: ev charger	Westinghouse 40-Volt Charger Lithium Ion (li-ion) ; Cordless Power Equipment Batteries & Chargers ✗	LectronTo NEMA 14-50 Plug J1772 Cable EV Charger Level 2 40-Amp Freestanding Single Electric Car Charger ; Electric Car Chargers ✓	Only part of aspect matching
Query by Example: Arrow 10-ft W x 20-ft L x 8-ft H Eggshell Metal Carport; Carports	Arrow 10-ft x 15.27-ft Eggshell Metal Carport ; Carports ✓	Arrow 10-ft W x 20-ft L x 8-ft H Eggshell Metal Carport ; Carports ✗	Exact the same item as query_item.
Compl. Item Rec.: Owens Corning Oakridge 32.8-sq ft Driftwood Laminated Architectural Roof Shingles ; Roof Shingles	Owens Corning VentSure 15-in x 48-in Black Plastic Stick Roof Ridge Vent ; Roof Ridge Vents ✓	Owens Corning DecoRidge 20-lin ft Driftwood Hip and Ridge Roof Shingles ; Roof Shingles ✗	Similar but not complementary item

terms of NDCG for each user in the test set. The performance difference between UIA and TS-UIA for users is sorted and plotted in Figure 3. According to the graph, 60+% of users benefit from joint optimization in query by example and complementary item recommendation tasks. We can observe that Δ NDCG for 5-10% of users in both of these tasks is equal to 1. Meaning that for these users joint optimization produces perfect ranking while task-specific training does not retrieve or recommend any relevant item. This is while we do not observe a significant number of users with Δ NDCG=1. In the keyword search task, the plot is more balanced. The reason is that the training data for this task is substantially larger than the other two tasks and thus it does not benefit a lot from joint optimization. These three plots also explains why the observed improvements for query by example and complementary item recommendation in Table 3 are larger than those for the keyword search task. Given these three plots, we suggest future work to focus on ensemble of task-specific and joint optimization.

4.2.4 The Impact of Personalization. To better understand the impact of personalization in UIA, we compare the results obtained from UIA and its non-personalized variation (i.e., UIA without APN). The performance difference (i.e., Δ NDCG) at the user level (sorted) is plotted in Figure 4 for all three information access functionalities. According to the figure, not all users benefit from personalization. This is expected especially for new and cold-start users and also for situations where user’s information need is different from their past interactions. The plots show that personalization has a negative impact on less than a quarter of users and a larger set of users take advantage of personalization. These plots also show that personalization has a larger impact on query by example and complementary item recommendation compared to keyword search. This is expected, as in keyword search, the user’s information need is often clear from the query, while it is more difficult to infer user’s need in query by example and complementary item recommendation. Predicting the need for personalization in future work can substantially improve UIA.

4.2.5 Case Study. Our case study, which compares the top-1 retrieved item by UIA and its task-specific training counterpart TS-UIA, is presented in Table 6. The results show that both models struggle with partial attribute matching errors in the keyword search task. For example, when searching for “6x6 privacy fence,” the top-1 item returned by TS-UIA is “outdoor privacy screens” which only matches the “privacy” attribute but misses the “fence”

attribute. Similarly, UIA returns an irrelevant item for the query “ev charger.” In the query by example task, the top-1 retrieved item by the task-specific training counterpart (TS-UIA) is often the same as the query item. This may be due to high term overlap in the positive training pairs, resulting in over-fitting of the model. Joint modeling with other tasks can partially mitigate this issue as different tasks serve as regularizers. In the complementary item recommendation task, UIA demonstrates superiority over TS-UIA by accurately identifying complementary items, such as correctly matching “roof ridge vents” with “roof shingle” despite the weak term matching signal.

5 CONCLUSIONS AND FUTURE DIRECTIONS

This paper studied the feasibility of implementing universal information access systems and proposed a generic and extensible dense retrieval framework, called UIA, for this purpose. UIA encodes each information access functionality in addition to the user’s request and takes advantage of approximate nearest neighbor search for efficient retrieval and recommendation of items. In addition, it introduces a novel attentive personalization network to extend the application of UIA to personalized information access tasks. We constructed two datasets to evaluate our approach, one large-scale real-world private data and another data built from the recent Amazon ESCI dataset. We demonstrated that UIA significantly outperforms competitive baselines and conducted extensive empirical exploration and ablation study to evaluate various aspects of the model. Based on the results reported in this paper, we envision a bright future (both short- and long-term) for universal information access. In the future, we will extend this work to other information access functionalities, such as question answering. We are also interested in evaluating the performance of the models beyond e-commerce applications. Therefore, future work will focus on data creation, organizing evaluation campaigns, and exploring ensemble approaches involving task-specific and universal information access models while predicting the need for personalization.

6 ACKNOWLEDGMENTS

This work was supported in part by the Center for Intelligent Information Retrieval, in part by NSF grant number 2143434, in part by Lowes, and in part by the Office of Naval Research contract number N000142212688. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

REFERENCES

- [1] Qingyao Ai, Daniel N. Hill, S. V. N. Vishwanathan, and W. Bruce Croft. 2019. A Zero Attention Model for Personalized Product Search. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (2019).
- [2] Qingyao Ai, Yongfeng Zhang, Keping Bi, Xu Chen, and W. Bruce Croft. 2017. Learning a Hierarchical Embedding Model for Personalized Product Search. *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2017).
- [3] Nicholas J. Belkin and W. Bruce Croft. 1992. Information filtering and information retrieval: two sides of the same coin? *Commun. ACM* 35 (1992), 29–38.
- [4] Daniel Fernando Campos, Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, Li Deng, and Bhaskar Mitra. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. *ArXiv abs/1611.09268* (2016).
- [5] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. 2017. Attentive Collaborative Filtering: Multimedia Recommendation with Item- and Component-Level Attention. *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2017).
- [6] Alexis Conneau and Douwe Kiela. 2018. SentEval: An Evaluation Toolkit for Universal Sentence Representations. *ArXiv abs/1803.05449* (2018).
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- [8] Luyu Gao and Jamie Callan. 2021. Condenser: a Pre-training Architecture for Dense Retrieval. In *EMNLP*.
- [9] Xiangnan He, Zhankui He, Jingkuan Song, Zhenguang Liu, Yu-Gang Jiang, and Tat-Seng Chua. 2018. NAIS: Neural Attentive Item Similarity Model for Recommendation. *IEEE Transactions on Knowledge and Data Engineering* 30 (2018), 2354–2366.
- [10] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. *Proceedings of the 26th International Conference on World Wide Web* (2017).
- [11] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. *CoRR abs/1511.06939* (2016).
- [12] Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy J. Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2021).
- [13] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.
- [14] Santosh Kabbur, Xia Ning, and George Karypis. 2013. FISM: factored item similarity models for top-N recommender systems. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining* (2013).
- [15] Wang-Cheng Kang, Chen Fang, Zhaowen Wang, and Julian McAuley. 2017. Visually-Aware Fashion Recommendation and Design with Generative Image Models. *2017 IEEE International Conference on Data Mining (ICDM)* (2017), 207–216.
- [16] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. *2018 IEEE International Conference on Data Mining (ICDM)* (2018), 197–206.
- [17] Vladimir Karpukhin, Barlas Ögüz, Sewon Min, Patrick Lewis, Ledell Yu Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. *ArXiv abs/2004.04906* (2020).
- [18] Donghyun Kim, Chanyoung Park, Jinoh Oh, Sungyoung Lee, and Hwanjo Yu. 2016. Convolutional Matrix Factorization for Document Context-Aware Recommendation. *Proceedings of the 10th ACM Conference on Recommender Systems* (2016).
- [19] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *CoRR abs/1412.6980* (2015).
- [20] Yehuda Koren and Robert M. Bell. 2011. Advances in Collaborative Filtering. In *Recommender Systems Handbook*.
- [21] Quoc V. Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *ICML*.
- [22] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy J. Lin. 2021. In-Batch Negatives for Knowledge Distillation with Tightly-Coupled Teachers for Dense Retrieval. In *REPLANLP*.
- [23] Prafull Prakash, Julia Killingback, and Hamed Zamani. 2021. Learning Robust Dense Retrieval Models from Incomplete Relevance Labels. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2021).
- [24] Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. RocketQA: An Optimized Training Approach to Dense Passage Retrieval for Open-Domain Question Answering. In *NAACL*.
- [25] Chandan K. Reddy, Lluís Màrquez, Fran Valero, Nikhil Rao, Hugo Zaragoza, Sambaran Bandyopadhyay, Arnab Biswas, Anlu Xing, and Karthik Subbian. 2022. Shopping Queries Dataset: A Large-Scale ESCI Benchmark for Improving Product Search. (2022). *arXiv:2206.06588*
- [26] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.* 3 (2009), 333–389.
- [27] Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2022. One Embedder, Any Task: Instruction-Finetuned Text Embeddings. *ArXiv abs/2212.09741* (2022).
- [28] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (2019).
- [29] Nandan Thakur, Nils Reimers, Andreas Ruckl'e, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. *ArXiv abs/2104.08663* (2021).
- [30] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*.
- [31] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative Deep Learning for Recommender Systems. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2015).
- [32] Suhang Wang, Yilin Wang, Jiliang Tang, Kai Shu, Suhas Ranganath, and Huan Liu. 2017. What Your Images Reveal: Exploiting Visual Contents for Point-of-Interest Recommendation. *Proceedings of the 26th International Conference on World Wide Web* (2017).
- [33] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv abs/1910.03771* (2019).
- [34] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. *ArXiv abs/2007.00808* (2021).
- [35] Liu Yang, Qi Guo, Yang Song, Sha Meng, Milad Shokouhi, Kieran McDonald, and W. Bruce Croft. 2016. Modelling User Interest for Zero-query Ranking. In *European Conference on Information Retrieval (ECIR 2016)*.
- [36] Hamed Zamani. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. *Proceedings of the 13th International Conference on Web Search and Data Mining* (2020).
- [37] Hamed Zamani, Michael Bendersky, Xuanhui Wang, and Mingyang Zhang. 2017. Situational Context for Ranking in Personal Search. *Proceedings of the 26th International Conference on World Wide Web* (2017).
- [38] Hamed Zamani and W. Bruce Croft. 2018. Joint Modeling and Optimization of Search and Recommendation. *ArXiv abs/1807.05631* (2018).
- [39] Hansi Zeng, Hamed Zamani, and Vishwa Vinay. 2022. Curriculum Learning for Dense Retrieval Distillation. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2022).
- [40] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, M. Zhang, and Shaoping Ma. 2021. Optimizing Dense Retrieval Model Training with Hard Negatives. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2021).
- [41] Kai Zhao, Yukun Zheng, Tao Zhuang, Xiang Li, and Xiaoyi Zeng. 2022. Joint Learning of E-commerce Search and Recommendation with a Unified Graph Neural Network. *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining* (2022).