

PARK: Personalized academic retrieval with knowledge-graphs

Pranav Kasela^{a,b}^{*}, Gabriella Pasi^a, Raffaele Perego^b

^a University of Milan-Bicocca, Milan, Italy

^b ISTI-CNR, Pisa, Italy

ARTICLE INFO

Keywords:

Personalized information retrieval
Knowledge graphs
Neural information retrieval
Dense retrieval

ABSTRACT

Academic Search is a search task aimed to manage and retrieve scientific documents like journal articles and conference papers. Personalization in this context meets individual researchers' needs by leveraging, through user profiles, the user related information (e.g. documents authored by a researcher), to improve search effectiveness and to reduce the information overload. While citation graphs are a valuable means to support the outcome of recommender systems, their use in personalized academic search (with, e.g. nodes as papers and edges as citations) is still under-explored.

Existing personalized models for academic search often struggle to fully capture users' academic interests. To address this, we propose a two-step approach: first, training a neural language model for retrieval, then converting the academic graph into a knowledge graph and embedding it into a shared semantic space with the language model using translational embedding techniques. This allows user models to capture both explicit relationships and hidden structures in citation graphs and paper content. We evaluate our approach in four academic search domains, outperforming traditional graph-based and personalized models in three out of four, with up to a 10% improvement in MAP@100 over the second-best model. This highlights the potential of knowledge graph-based user models to enhance retrieval effectiveness.

1. Introduction

Academic search is a task that focuses on retrieving scientific documents (e.g., research papers, articles, PhD manuscripts) from large repositories, based on queries, generally, formulated by researchers, students, educators, and professionals seeking specific information. Enhancing this process through search personalization offers the potential to tailor results to individual expertise and preferences, going beyond the *one size fits all* approach of traditional Information Retrieval (IR) systems. Personalized Information Retrieval (PIR) systems make use of user models, which encode user-specific information and contextual cues to improve the assessment of relevance of search results.

The focus of this paper is on user modeling for personalized academic search, a task that remains under-explored in the literature. Specifically, our contributions are as follows: We propose a user model for academic search that leverages a knowledge graph constructed from citation data to capture user preferences and document relationships. Furthermore, we develop a complete retrieval pipeline, starting from an academic citation graph to create a personalized information retrieval system for academic search.

Citation graphs, where nodes are academic papers and directed edges denote citations, can help enhance academic search by improving relevance assessment based on citation patterns and relationships

between scientific documents. While citation graphs have been successfully applied in academic recommender systems [1], their potential in PIR remains partially under-explored [2]. To address this issue, in this paper, we propose a novel approach that uses knowledge graph embeddings to generate user models for personalized academic search. These embeddings map entities and relationships into a continuous vector space, capturing both explicit connections (e.g., citations) and latent structures (e.g., topical similarities and social connections) within citation graphs.

Our approach integrates the principles of collaborative filtering by modeling user preferences based on relationships derived from the citation network. By embedding papers and authors into a shared semantic space with a neural language model, our method enables user models to leverage both the structural and textual information of academic content. Specifically, the proposed approach aligns knowledge graph embeddings with neural language models through a two-stage training process. First, we train a bi-encoder neural language model on an academic search dataset. In the second step, we use a translational algorithm for knowledge graph embeddings to embed the knowledge graph, derived from the academic citation network, into the same semantic space as the neural language model. The two-stage training approach

* Corresponding author at: University of Milan-Bicocca, Milan, Italy.

E-mail addresses: pranav.kasela@unimib.it (P. Kasela), gabriella.pasi@unimib.it (G. Pasi), raffaele.perego@isti.cnr.it (R. Perego).

offers several advantages: it allows for seamless integration with any dense encoder with the knowledge graph without requiring retraining of the language model. Additionally, it allows the acceleration of the training process by pre-indexing document embeddings for efficient retrieval.

The results of the performed experiments show that user models based on such knowledge graph significantly enhance retrieval effectiveness compared to other methods, including graph-based methods, achieving state-of-the-art performance in three out of four benchmarks in personalized academic search [3], with improvements of up to 10% over the best-performing baseline. By leveraging the semantic and structural information embedded within citation networks for user modeling, this work represents a significant advancement in personalized academic retrieval, offering a robust framework for future research and applications in this domain.

The work is organized as follows: In Section 2 we describe the training techniques employed for creating text embeddings and knowledge graph embeddings. In Section 3, we propose PARK, an architecture that relies on two-stage training and embeds the knowledge graph in a semantic space shared with neural language models. The experimental settings are reported in Section 4, and the results are presented and discussed in Section 4.3. In Section 5 we present an ablation study on the type of nodes utilized in the knowledge graph. Finally, Section 6 concludes the work by providing new insights and new future research directions.

2. Related work and background

Neural Information Retrieval (NIR) marks a major advancement in IR by using deep learning to enhance retrieval effectiveness [4]. Unlike traditional IR models that depend on manual feature engineering and basic mathematical models, NIR systems automatically learn complex representations and relationships between queries and documents through deep learning. In this work, we focus on the *Siamese Networks* or Bi-encoders class of neural IR approaches [5]. The training of these models is done following the approach presented by Huang et al. [6], which minimizes the triplet margin loss [7] on triplet samples consisting of a query q , a relevant document d^+ , and negative documents D^- randomly sampled from the collection. For efficiency purposes, the negative samples are taken from the samples used in the same batch, a technique known as random in-batch negatives [8,9].

Personalized Information Retrieval (PIR) is a growing research field, driven by the need to enhance the relevance of search results for individual users [10–13]. Early efforts in personal search focused on leveraging metadata [14], contextual information like query time and location [15,16], and developing training schemes to handle sparse user interactions [17,18]. Central to PIR is the definition of user models that capture various aspects of user behavior, preferences, and context to improve the relevance assessment process for specific user. Initial approaches to building user representations included term-level models [19], topic models [13], and latent representations learned through techniques like matrix factorization [20]. The development of deep learning has significantly impacted also the definition of more accurate user models. Existing techniques rely on representation learning [21] to learn latent semantic representations of queries, documents, and user preferences for enhancing personalized search effectiveness [22–24]. More recent advancements include learning personalized word embeddings [25], and employing Recurrent Neural Networks (RNNs) [22, 26] and Transformers [23,27] to model user representations. A line of research has focused on leveraging hierarchical recurrent neural networks (HRNN) with query-aware attention to capture both short-term and long-term dependencies within user search sessions [28]. By integrating query-aware attention, these models can dynamically adjust the focus based on the specific query, leading to more accurate and more robust user profiles. Additionally, graph neural networks [29] and reinforcement learning techniques [24] have been explored to

model complex user behaviors. Lu et al. [2] propose a method that relies on a knowledge graph for personalized entity linking to improve search results based on user feedback during search sessions. In contrast to purely neural methods, such knowledge-enhanced personalization offers greater ability to capture structural relations across authors and topics. These methods, particularly those with query-aware attention, have shown promising results in capturing the temporal dynamics of user interactions. Despite advancements, these approaches often rely on the availability of session-based data and focus on personalization within search sessions, thus possibly failing to capture users' long-term interests. Furthermore, the reproducibility of results remains a challenge due to the use of proprietary datasets or to a lack of clear instructions for reproducing results [30]. Bassani et al. specifically highlight that many personalized models lack public datasets or open-source code, hindering replicability and comparison. Another research direction explores the use of attention mechanisms to enhance long-term user modeling during query time. Attention mechanisms enable the system to dynamically weigh and aggregate user-related information in relation to the current query, thus refining the personalization process [30]. This method has proven effective in identifying the most relevant aspects of a user's search history related to the given query. Recently, the aspect of controllable personalization has gained attention, addressing concerns about the opacity of personalized systems and the lack of user control. Work in this area has explored providing users with interactive control over personalized search by making user representations, such as term or entity-based profiles, editable through visualization interfaces [31,32]. This allows users to inspect and modify how the system perceives their interests, aiming to improve transparency and user satisfaction. Recent cross-encoder based models like CtrlCE [33] further extend this idea by integrating memory-augmented user profiles and a calibrated control mechanism that decides when personalization is beneficial. This line of work contributes to a growing interest in transparency, fairness, and user agency in IR [34].

2.1. Knowledge graph embeddings

A knowledge graph (KG) is a data structure that encodes the knowledge of the real world in a graph format, whose nodes represent the entities and whose edges are the relationships that connect these entities. Depending on the application, KGs can be modeled as a directed edge-labeled graph, a heterogeneous graph, or a property graph [35]. Deep learning techniques can create dense representations of the graph in a continuous, low-dimensional vector space. One of the most commonly used families of knowledge graph embedding techniques is the translational model, which maps entities and relations in a knowledge graph to a continuous vector space, capturing semantic relationships and preserving their geometry. Among these translational algorithms, *TransE* and *TransH* are particularly popular due to their simplicity and effectiveness.

TransE, Translating Embeddings, was introduced by Bordes et al. [36]. It operates under the principle that relationships between entities can be represented as translations in the embedding space. Given a triplet $\langle h, r, t \rangle$ in the knowledge graph, where h is the head entity, r is the relation, and t is the tail entity, *TransE* embeds these entities and relations into a continuous d -dimensional vector space. The training process involves minimizing a margin-based ranking loss function, which encourages correct triplets to have a smaller distance than incorrect ones. A big drawback is that *TransE* struggles with one-to-many, many-to-one and many-to-many relations. *TransH*, Translating Embeddings on Hyperplanes, was proposed by Wang et al. [37] to address the limitations of *TransE*, particularly in handling complex relations. *TransH* models each relation as a hyperplane, allowing entities to have different embeddings for different relations. *TransH* can overcome the limitations of *TransE* and can handle complex relationships by projecting entities onto relation-specific hyperplanes. This also allows more flexible representations of the entities. Knowledge graphs are

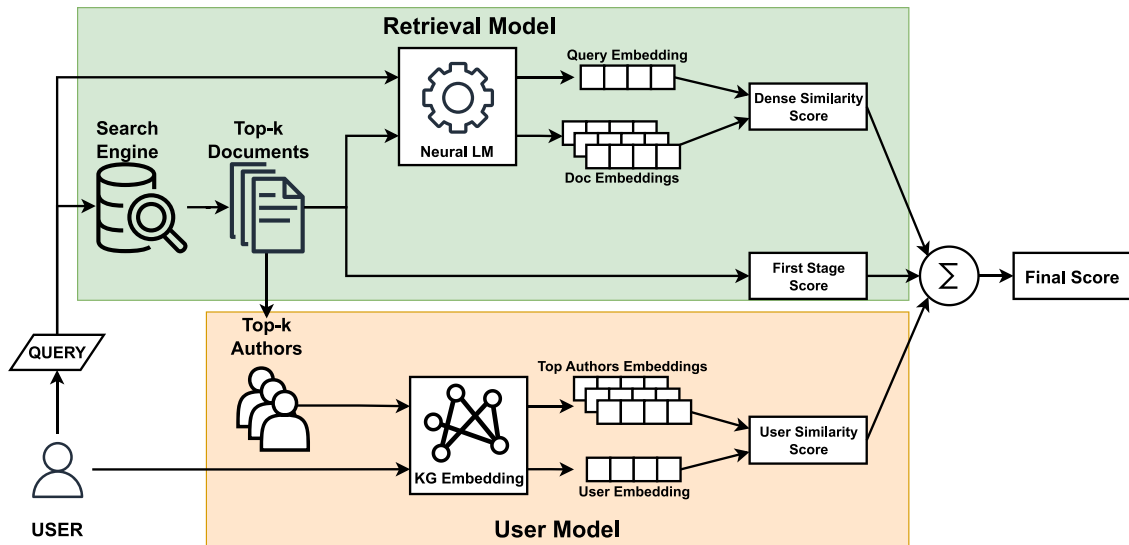


Fig. 1. Overview of the PARK retrieval pipeline.

increasingly used in recommender systems to capture complex entity relationships [38].

The approach we present in this paper aims at developing a personalized retrieval pipeline that leverages a combination of knowledge graph to generate user models and of neural retrieval models. This combination captures both the semantic and structural aspects of academic documents, providing a novel and robust framework for enhancing personalized academic retrieval.

3. Personalized academic retrieval with knowledge-graphs

In this Section, we first introduce the overall architecture of the proposed academic search system, which we will call Personalized Academic Retrieval with Knowledge-graphs (PARK). Then, we describe the proposed *Retrieval* and *User Models* by detailing how they are designed and trained. Specifically, we focus on how we leverage a citation graph (CG) to create a Knowledge Graph (KG) that is then embedded in a learned latent space, where the KG entities are represented as high-dimensional vectors.

The overall architecture of PARK is illustrated in Fig. 1. Our model has two main components: the first component, in the top part of the Figure, is a classical information retrieval IR pipeline that uses an efficient first retrieval stage based on BM25 and a following re-ranking stage exploiting our dense retrieval model. The second component, depicted in the lower part of the Figure, is the proposed user model based on knowledge graph embeddings, which computes user similarity scores, reflecting the alignment of research profiles and interests of the user issuing the query and the authors of the documents retrieved by the first stage retriever.

PARK integrates user models represented as embeddings derived from a knowledge graph's latent representation, along with document similarity scores generated by a lexical model, in this paper we make use of BM25 [39]. Additionally, it considers the semantic content of the same documents as represented in the dense neural retrieval model. The final documents presented to the user are ranked based on a weighted sum of three scores: BM25, dense similarity, and user similarity scores.

3.1. The retrieval model

We employ a two-stage retrieval pipeline to trade-off retrieval efficiency and effectiveness. In the first stage, we use lexical matching with BM25 to retrieve an initial set of documents for each query. In the second stage, re-ranking is performed using a state-of-the-art

bi-encoder dense retrieval model. The retrieval pipeline outputs for each query-document pair two scores: the first stage score and the semantic similarity score computed by the neural re-ranker (See Fig. 1). The training of the neural retrieval model is performed using the procedure mentioned at the beginning of Section 2. The neural retrieval model is optimized to minimize the distance between the query representation and the associated relevant document representations while increasing the distance between the query representation and the non relevant documents representations using the Triplet Margin Loss [7], $\mathcal{L}(q, d^+, D^-)$, defined as:

$$\sum_{d^- \in D^-} \max(\|q - d^+\|_p - \|q - d^-\|_2 + \gamma, 0)$$

where \mathbf{q} and \mathbf{d} are the vector representation of the query and the document, respectively, $\gamma \in \mathbb{R}$ is the margin and $\|q - d^+\|_2$ is the norm of the difference between the vector representations of the query and of the document.

3.2. The user model: PARK-E and PARK-H

We utilize the *Closed World Assumption* (CWA) [40] when constructing our knowledge graph embeddings with both the TransE and the TransH algorithms. According to the CWA, any statement that is not known to be true is considered false. This assumption enhances the model's ability to learn effectively from the available data. In this context, knowledge graph embeddings map entities and relationships into continuous vector spaces while preserving their structural and semantic characteristics. By applying the CWA, the embedding process allows the model to generate more accurate and robust representations.

Our approach employs knowledge graph embeddings to represent various entities within the academic knowledge graph that we define, including users. The user embeddings are then used to score documents retrieved through the BM25 algorithm, based on the similarity between the user making the query and the authors of the retrieved documents. Below, we outline a two-step procedure, illustrated in Fig. 2, for learning the embeddings from the Citation Graph.

Academic knowledge graph construction. An academic citation graph is a directed graph that describes the citations within a collection of scientific documents. Each vertex (or node) in the graph represents a scientific document, and each edge in the graph is directed from one document to another, indicating that one document cites another. Additionally, each document node contains metadata, such as the author's name, affiliations, publication venue, and timestamp. This information




Fig. 2. Overview of the process to obtain user embeddings from the citation Graph. The first step converts the Citation Graph to an Academic Knowledge Graph, and the second step embeds the Academic Knowledge Graph according to the proposed techniques.

can be used to transform the citation graph into a property graph. A *Property Graph* is a graph model where nodes are connected through edges that can have named properties. These graphs can be effectively used as knowledge graphs for embedding purposes, with each node and edge being considered in the embedding process. In our case, the nodes in the property graph represent papers, authors, institutions, or venues, and edges represent the relationships among these entities. Academic citation graphs, to a reasonable extent, adhere to the CWA because, in these settings, the data is often curated and continuously updated, ensuring a high degree of completeness and accuracy.

To construct an academic knowledge graph from the academic citation graph, we define four types of nodes:

- *Users (or authors)*: These nodes represent authors of the academic documents.
- *Documents*: These nodes represent scholarly works, including papers and articles, identified by their titles and abstracts.
- *Venues*: These nodes denote the publication venues, such as conferences or journals, where scientific documents are published.
- *Affiliations*: These nodes represent academic institutions, such as universities or research organizations, to which authors may belong.

These nodes can be connected by five types of relationships:

- *Wrote* (user $\xrightarrow{\text{wrote}}$ document): Connects a user to all documents they have authored.
- *Cited* (user $\xrightarrow{\text{cited}}$ document): Connects a user to all documents cited by any of their authored works.
- *In Venue* (user $\xrightarrow{\text{in_venue}}$ venue): Connects the user to all venues where their documents were published.
- *Affiliated* (user $\xrightarrow{\text{affiliated}}$ affiliation): Connects the user to their respective affiliation.
- *Co-author* (user $\xleftrightarrow{\text{co_author}}$ user): Connects a user to other users with whom they co-authored a document. This is the only symmetric relationship among the five.

PARK-E and PARK-H. The user models, PARK-E and PARK-H, utilize the TransE and TransH algorithms respectively, to generate user embeddings from knowledge graphs. We selected TransE and TransH for their simplicity and effectiveness in capturing relationships between entities within a continuous vector space, making them well-suited for our user modeling needs. While other KG embedding techniques, such as TransR [41] or RotatE [42], provide more complex representations, TransE and TransH offer a balanced approach between computational efficiency and model performance, an essential consideration given the dimension of academic citation graphs. The complete pipeline, starting from the citation graph, is shown in Fig. 2.

The embedding dimension is aligned with the semantic space dimension of the neural language model used by the second stage re-ranker, allowing document semantics to be integrated into the knowledge graph embedding space. Specifically, we fix the embeddings of

document nodes and use the embeddings generated by the neural language model to represent the documents. This approach preserves the semantics of the documents while infusing them into the knowledge graph.

The embeddings of all other nodes follow the procedure outlined by knowledge graph embedding algorithms, TransE [36] and TransH [37]. The translational nature of TransE and TransH ensures that entities with similar relationships and contexts are positioned closer together in the embedding space, which is particularly beneficial for user modeling. Although TransE has limitations in handling complex relationships, such as one-to-many or many-to-many, these challenges are mitigated in this context, as clustering similar users together is a desirable property for effective user modeling.

3.3. Integration of retrieval and user models

The final ranking of the top-k retrieved documents by BM25 is based on three distinct scores, as illustrated in Fig. 1. The first score is the dense similarity score, derived from the neural language model used for re-ranking. This score assesses the contextual and semantic relevance of documents to the query. The second score is obtained from the first stage retriever, which extracts the candidate results from the index using BM25. Finally, the user similarity score measures the similarity between the user embeddings of the query writer and the authors of the documents being scored, reflecting the alignment of their research profiles and interests. The documents are ultimately re-ranked using a convex combination of the three normalized scores: BM25 scores, dense similarity scores, and user similarity scores. We describe this in detail in the next section.

4. Experimental evaluation

In the following sections, we will first describe the datasets used in our experiments to evaluate the effectiveness of PARK. Next, we will introduce the competing methods relevant to the academic search task we are addressing. Finally, we will discuss the results of extensive experiments conducted to assess the performance of PARK compared to the competing methods on the specified datasets. The code is made publicly available at the following link: https://github.com/pkasela/PARK-Personalized_Academic_Retrieval_with_Knowledge-graphs.

4.1. Datasets and experimental settings

For our experiments, we utilize four datasets specifically designed for comparing models in the context of personalized academic search, as introduced in [3]. Each dataset focuses on a different domain: Computer Science, Physics, Political Science, and Psychology. Each domain contains approximately 5 million documents, resulting in a total of over 18 million documents and 1.9 million training queries. Queries in this semi-synthetic dataset are generated by processing the titles of academic papers, removing stop words, and applying a Krovetz non-destructive stemmer to ensure they closely resemble real-world search queries. For both the training and validation sets, as in [3],

relevant documents for each query are identified as those cited by the paper and those appearing in the top 100 results from a BM25 search, where the query is the exact paper’s title.

To create a realistic evaluation scenario and prevent data leakage, the dataset is split chronologically. The training sets consist of documents published before specific cutoff dates (2016 for Political Science, 2017 for Computer Science and Physics, and 2019 for Psychology), while the test sets include documents published after these dates.

As the bi-encoder dense retrieval, we rely on *MiniLM* [43], a distilled version of BERT. MiniLM was selected for its computational efficiency and strong performance, making it particularly suitable for large-scale IR tasks. We train MiniLM for 10 epochs, with a learning rate of $5 \cdot 10^{-5}$ and a batch size of 256. The Knowledge Graph embedding model is trained for 100 epoch with a batch size of 16384 and a learning rate of 10^{-3} . In both cases, we use AdamW [44] as the model optimizer.

4.2. Competing methods

The competing baselines include traditional and neural retrieval models, user models, and graph-based approaches. Below, we describe each baseline:

- **BM25** [39]: The probabilistic information retrieval model used as the recall-oriented first-stage ranker in this study. This baseline uses the score returned by our first-stage ranker.
- **MiniLM** [43]: The precision-oriented neural re-ranker based on the pre-trained language model MiniLM and fine-tuned as detailed in Fig. 1. This baseline uses the score returned by our second-stage ranker.
- **Mean** [3]: This user model is defined as the average representation of all documents related to a user.
- **Attention** [45]: A query-aware user modeling technique that uses the Attention mechanism [46] to weigh the contributions of user-related document representations in generating the user model.
- **Self Citation**: This baseline considers the self-citation behavior of authors, where documents cited by their own authors (or co-authors) receive a boost in their relevance score.
- **CrossEnc_{RA}** [47]: A cross-encoder model based on MPNet [48], which inputs the query, the candidate documents and the top-1 document most similar to the query from the user-related information.
- **CTRL_{It}** [49]: This model relies on memory-augmented cross-encoders. It incorporates user profiles as editable memory representations derived from historical interactions, to control and adjust personalization.
- **PageRank** [50]: A graph-based algorithm originally designed to rank web pages, which can also be applied to citation graphs to rank academic papers based on citation links. We use the NetworkX implementation of the PageRank algorithm [51] with α set to 0.85.
- **POP** (Popularity) [3]: This baseline ranks documents based on their popularity, measured by the number of citations they have received. This simple heuristic assumes that more frequently cited documents are more relevant or important.

The code for CTRL_{It} and CrossEnc_{RA} has not been publicly released, so we report the performance metrics as stated in their paper. For all the other baselines (except the first two), the final relevance score for a document is computed by linearly combining the normalized scores from BM25, the neural re-ranker (MiniLM in our case), and the specific baseline method using a convex sum [52]. This approach harnesses the strengths of traditional lexical matching (BM25), semantic understanding (MiniLM), and the unique properties of each baseline to enhance retrieval performance [53,54]. The final score is expressed as follows:

$$S(q, d) = \lambda_1 \cdot \text{BM25}(q, d) + \lambda_2 \cdot \text{MiniLM}(q, d) + \lambda_3 \cdot \text{Baseline}(q, d)$$

where $\lambda_1 + \lambda_2 + \lambda_3 = 1$ and $\text{Baseline}(q, d)$ denotes the score from the respective baseline method (Mean, Attention, Self Citation, PageRank, or POP). The λ_i values are optimized based on performance metrics from the validation set. Additionally, all baseline models are optimized, and those that are trainable are trained end-to-end to ensure a fair comparison.

4.3. Experimental results

Table 1 presents the results of our experiments, showcasing the performance of our proposed PARK-E and PARK-H models in comparison to various baseline methods discussed in the previous section. The evaluation covers four domains: Computer Science, Political Science, Psychology, and Physics. We used three evaluation metrics: Mean Average Precision@100 (MAP@100), Mean Reciprocal Rank@10 (MRR@10), and Normalized Discounted Cumulative Gain@10 (NDCG@10). The findings indicate that the PARK-E and PARK-H models significantly outperform the baseline methods across nearly all domains. The only exception is in Computer Science, where the popularity-based (POP) baseline performs slightly better.

First, we note that all baseline models demonstrate improved retrieval effectiveness over the BM25 baseline. In the Computer Science domain, the PARK-H model shows a significant improvement of 8% in MAP@100 compared to the self-citation user model; however, it falls short by 3% when compared to the popularity-based (POP) baseline. This indicates that popularity, as measured by citation counts, plays a crucial role in determining the relevance of search results in this area. In contrast, the PARK-H model consistently outperforms all other baselines in the domains of Political Science, Psychology, and Physics. Specifically, it improves MAP@100 over the second-best performing model by 9%, 10%, and 9%, respectively. These improvements are statistically significant, as indicated by the asterisk in Table 1, highlighting the robustness of our models.

Overall, these findings suggest that PARK-E and PARK-H are highly effective for personalized academic search, with the exception of the Computer Science domain, where the POP model exhibits slightly better performance. The success of our proposed models in most domains indicates their potential for broader application in academic information retrieval.

5. Ablation study on node types

To better understand the contribution of each node type in our model, we conducted an ablation study across the four benchmarks on PARK-H. The goal was to evaluate the effect of each node type (user, venue, and affiliation) and their associated relations on the model’s retrieval performance. We started by implementing the model with only user nodes and their three associated relations (authored, cited, and co-authored). Next, we added the venue node and venue relation to evaluate any additional impact. Finally, we incorporated the affiliation node, i.e., the complete setup used in PARK (see Table 2).

With **only user nodes** and the associated relations, we observe a notable improvement in retrieval performance across all datasets when compared to the baseline model without any other knowledge graph nodes. This improvement underscores the critical role of user-specific information in personalized academic search tasks. Adding the **venue node** and its associated relation yielded minimal performance gains in terms of NDCG@10. NDCG@10 increased by only 0.3%, 1.5%, 0.5% in the Computer Science, Political Science, and Psychology benchmark. In Physics, NDCG@10 remained constant at 0.383. These limited improvements suggest that publication venue information contributes less significantly to user preference modeling, suggesting that venue context may not heavily influence individual user preferences within these academic domains. The performance increases significantly across all datasets when the **affiliation node** is added, completing the full PARK model. NDCG@10 increases by 2.7%, 2.6%,

Table 1
Effectiveness of PARK-E and PARK-H compared to the competing methods on the four datasets. The best-performing model is highlighted in boldface.

Model	Computer science			Political science			Psychology			Physics		
	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10
BM25	0.123	0.489	0.225	0.133	0.502	0.241	0.126	0.512	0.239	0.128	0.537	0.269
MiniLM	0.193	0.600	0.301	0.186	0.580	0.297	0.218	0.647	0.342	0.183	0.624	0.335
Mean	0.199	0.606	0.308	0.193	0.598	0.306	0.220	0.652	0.347	0.189	0.639	0.345
Attention	0.201	0.612	0.312	0.199	0.612	0.314	0.220	0.656	0.349	0.190	0.648	0.348
Self Citation	0.213	0.624	0.325	0.205	0.613	0.321	0.237	0.689	0.370	0.204	0.671	0.365
CTRL _{IT}	–	0.629	0.322	–	0.648	0.338	–	0.685	0.370	–	0.667	0.366
CrossEnc _{RA}	–	0.635	0.324	–	0.651	0.338	–	0.700	0.380	–	0.673	0.369
PageRank	0.213	0.644	0.331	0.203	0.622	0.324	0.230	0.670	0.360	0.189	0.636	0.346
POP	0.238*	0.684*	0.370*	0.214	0.649	0.345	0.225	0.656	0.356	0.206	0.670	0.370
PARK-E	0.228	0.651	0.344	0.232	0.661	0.356	0.261*	0.716*	0.397*	0.225*	0.695	0.391*
PARK-H	0.230	0.655	0.346	0.233*	0.662*	0.357*	0.255	0.712	0.392	0.225*	0.696*	0.391*

* Indicates a statistically significant difference over the second-best-performing model.

Table 2
Ablation study results for each node type on four datasets.

Node types	Computer science			Political science			Psychology			Physics		
	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10	MAP@100	MRR@10	NDCG@10
Only user	0.223	0.644	0.336	0.220	0.637	0.343	0.244	0.693	0.379	0.221	0.686	0.383
↓ + Venue	0.225	0.642	0.337	0.225	0.646	0.348	0.247	0.694	0.381	0.222	0.686	0.383
↓ + Affiliation	0.230	0.655	0.346	0.233	0.662	0.357	0.255	0.712	0.392	0.225	0.696	0.391

2.9% and 2.1% for Computer Science, Political Science, Psychology, and Physics respectively. These results highlight the significant positive impact of including the affiliation node, which likely captures valuable institutional and collaboration context, further enhancing the retrieval with user research interests.

The ablation study indicates that while user-specific information is the primary driver of performance gains, the affiliation node further enhances retrieval effectiveness by adding valuable institutional context. In contrast, the venue node has limited impact, suggesting it may be less essential for capturing user-specific preferences in academic search tasks.

6. Conclusion

The experimental results clearly demonstrate that our proposed user models, PARK-E and PARK-H, significantly outperform baseline methods across multiple academic domains, showing particularly strong performance in Political Science, Psychology, and Physics. Although the POP model slightly outperforms PARK models in the Computer Science domain, PARK remains the most effective user-centric model, underscoring the validity and robustness of our approach. These findings highlight the substantial benefits of integrating knowledge graph embeddings into personalized information retrieval. By leveraging user embeddings derived from TransE and TransH algorithms, our models achieve more accurate document re-ranking by effectively capturing user similarities and preferences, thereby significantly enhancing retrieval performance. However, these promising outcomes are accompanied by certain limitations. The effectiveness of knowledge graph embeddings is intrinsically linked to the quality and completeness of citation data. Incomplete or biased citation records can adversely impact embedding quality and, consequently, retrieval performance. Moreover, our model currently operates under a closed-world assumption, which, while suitable for academic search contexts, may restrict its applicability to other domains or tasks. This highlights the need for further exploration and adaptation to enhance generalizability. Additionally, the decision to fix document nodes during training may constrain the model's flexibility. Future research could explore the introduction of soft constraints in the loss function, allowing document nodes to adjust dynamically while still converging toward optimal document embeddings. Other interesting research directions include exploring hybrid models that integrate popularity-based features with

knowledge graph embeddings to improve retrieval effectiveness. Relaxing the fixed document node constraint could also offer a balance between flexibility and accuracy, potentially unlocking even greater performance gains.

In conclusion, the integration of knowledge graph embeddings into user modeling represents a significant advancement in personalized information retrieval. Our results demonstrate substantial improvements in retrieval performance across various academic domains, suggesting that these techniques have the potential to significantly advance the state of the art in academic search. With continued refinement and optimization, these methods are poised to make a lasting impact on the field of personalized information retrieval.

CRedit authorship contribution statement

Pranav Kasela: Writing – original draft, Visualization, Software, Resources, Methodology, Formal analysis, Data curation, Conceptualization. **Gabriella Pasi:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation. **Raffaele Perego:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation.

Funding sources

This work was supported by the European Union – Next Generation EU within the project NRPP M4C2, Investment 1.,3 DD. 341–15 march 2022 – FAIR – Future Artificial Intelligence Research – Spoke 4 - PE00000013 - D53C22002380006.

This work was partially supported by the Spoke “Human-centered AI” of the M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - “FAIR - Future Artificial Intelligence Research”, the “Extreme Food Risk Analytics” (EFRA) project, Grant no. 101093026, funded by European Union – NextGenerationEU.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Pranav Kasela reports financial support was provided by European Union. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data used is already publicly available and the respective resource papers have been cited.

References

- [1] Z. Ali, G. Qi, P. Kefalas, W.A. Abro, B. Ali, A graph-based taxonomy of citation recommendation models, *Artif. Intell. Rev.* 53 (2020) 5217–5260.
- [2] S. Lu, Z. Dou, C. Xiong, X. Wang, J.-R. Wen, Knowledge enhanced personalized search, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 709–718, <http://dx.doi.org/10.1145/3397271.3401089>.
- [3] E. Bassani, P. Kasela, A. Raganato, G. Pasi, A multi-domain benchmark for personalized search evaluation, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 3822–3827, <http://dx.doi.org/10.1145/3511808.3557536>.
- [4] B. Mitra, N. Craswell, An introduction to neural information retrieval, *Found. Trends Inf. Retr.* 13 (1) (2018) 1–126, <http://dx.doi.org/10.1561/1500000061>.
- [5] N. Reimers, I. Gurevych, Sentence-BERT: Sentence embeddings using siamese BERT-networks, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3–7, 2019, Association for Computational Linguistics, 2019, pp. 3980–3990.
- [6] P.-S. Huang, X. He, J. Gao, L. Deng, A. Acero, L. Heck, Learning deep structured semantic models for web search using clickthrough data, in: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13, Association for Computing Machinery, New York, NY, USA, 2013, pp. 2333–2338, <http://dx.doi.org/10.1145/2505515.2505665>, URL <https://doi.org/10.1145/2505515.2505665>.
- [7] V. Balntas, E. Riba, D. Ponsa, K. Mikolajczyk, Learning local feature descriptors with triplets and shallow convolutional neural networks, in: *Bmvc*, Vol. 1, 2016, p. 3.
- [8] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, W.-t. Yih, Dense passage retrieval for open-domain question answering, in: B. Webber, T. Cohn, Y. He, Y. Liu (Eds.), Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP, Association for Computational Linguistics, Online, 2020, pp. 6769–6781, <http://dx.doi.org/10.18653/v1/2020.emnlp-main.550>, URL <https://aclanthology.org/2020.emnlp-main.550>.
- [9] J. Zhan, J. Mao, Y. Liu, M. Zhang, S. Ma, RepBERT: Contextualized text embeddings for first-stage retrieval, 2020, [arXiv:2006.15498](https://arxiv.org/abs/2006.15498). URL <https://arxiv.org/abs/2006.15498>.
- [10] B. Tan, X. Shen, C. Zhai, Mining long-term search history to improve search accuracy, in: T. Eliassi-Rad, L.H. Ungar, M. Craven, D. Gunopulos (Eds.), Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Philadelphia, PA, USA, August 20–23, 2006, ACM, 2006, pp. 718–723, <http://dx.doi.org/10.1145/1150402.1150493>.
- [11] A. Sieg, B. Mobasher, R.D. Burke, Web search personalization with ontological user profiles, in: M.J. Silva, A.H.F. Laender, R.A. Baeza-Yates, D.L. McGuinness, B. Olstad, Ø.H. Olsen, A.O.F. ao (Eds.), Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, CIKM 2007, Lisbon, Portugal, November 6–10, 2007, ACM, 2007, pp. 525–534, <http://dx.doi.org/10.1145/1321440.1321515>.
- [12] M.J. Carman, F. Crestani, M. Harvey, M. Baillie, Towards query log based personalization using topic models, in: J. Huang, N. Koudas, G.J.F. Jones, X. Wu, K. Collins-Thompson, A. An (Eds.), Proceedings of the 19th ACM Conference on Information and Knowledge Management, CIKM 2010, Toronto, Ontario, Canada, October 26–30, 2010, ACM, 2010, pp. 1849–1852, <http://dx.doi.org/10.1145/1871437.1871745>.
- [13] D.A. Sontag, K. Collins-Thompson, P.N. Bennett, R.W. White, S.T. Dumais, B. Billerbeck, Probabilistic models for personalizing web search, in: E. Adar, J. Teevan, E. Agichtein, Y. Maarek (Eds.), Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8–12, 2012, ACM, 2012, pp. 433–442, <http://dx.doi.org/10.1145/2124295.2124348>.
- [14] M. Bendersky, X. Wang, D. Metzler, M. Najork, Learning from user interactions in personal search via attribute parameterization, in: WSDM '17: Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, 2017, pp. 57–66, URL <https://doi.org/10.1145/3018661.3018712>.
- [15] Z. Qin, Z. Li, M. Bendersky, D. Metzler, Matching cross network for learning to rank in personal search, in: Proceedings of the Web Conference 2020, 2020, pp. 1218–1228, URL <https://doi.org/10.1145/3366423.3380046>.
- [16] H. Zamani, M. Bendersky, X. Wang, M. Zhang, Situational context for ranking in personal search, in: WWW '17: Proceedings of the 26th International Conference on World Wide Web, 2017, pp. 1521–1530, URL <https://doi.org/10.1145/3038912.3052648>.
- [17] K. Bi, P. Metrikov, C. Li, B. Byun, Leveraging user behavior history for personalized email search, in: Proceedings of the Web Conference 2021, 2021, pp. 3526–3537, URL <https://doi.org/10.1145/3442381.3450110>.
- [18] X. Wang, M. Bendersky, D. Metzler, M. Najork, Learning to rank with selection bias in personal search, in: SIGIR '16: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2016, pp. 115–124, URL <https://doi.org/10.1145/2911451.2911537>.
- [19] J. Teevan, S.T. Dumais, E. Horvitz, Personalizing search via automated analysis of interests and activities, in: SIGIR '05: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, 2005, pp. 449–456, URL <https://doi.org/10.1145/1076034.1076111>.
- [20] F. Cai, S. Liang, M. de Rijke, Personalized document re-ranking based on Bayesian probabilistic matrix factorization, in: SIGIR '14: Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2014, pp. 915–918, URL <https://doi.org/10.1145/2600428.2609453>.
- [21] Y. Bengio, A.C. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (8) (2013) 1798–1828, <http://dx.doi.org/10.1109/TPAMI.2013.50>, URL <https://doi.org/10.1109/TPAMI.2013.50>.
- [22] Q. Ai, D.N. Hill, S.V.N. Vishwanathan, W.B. Croft, A zero attention model for personalized product search, in: W. Zhu, D. Tao, X. Cheng, P. Cui, E.A. Rundensteiner, D. Carmel, Q. He, J.X. Yu (Eds.), Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3–7, 2019, ACM, 2019, pp. 379–388, <http://dx.doi.org/10.1145/3357384.3357980>.
- [23] J. Zhong, W. Guo, H. Gao, B. Long, Personalized query suggestions, in: J. Huang, Y. Chang, X. Cheng, J. Kamps, V. Murdock, J. Wen, Y. Liu (Eds.), Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2020, Virtual Event, China, July 25–30, 2020, ACM, 2020, pp. 1645–1648, <http://dx.doi.org/10.1145/3397271.3401331>.
- [24] J. Yao, Z. Dou, J. Xu, J. Wen, Rlper: A reinforcement learning model for personalized search, in: Y. Huang, I. King, T. Liu, M. van Steen (Eds.), WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20–24, 2020, ACM / IW3C2, 2020, pp. 2298–2308, <http://dx.doi.org/10.1145/3366423.3380294>.
- [25] J. Yao, Z. Dou, J.-R. Wen, Employing personal word embeddings for personalized search, in: SIGIR '20: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 1513–1516, URL <https://doi.org/10.1145/3397271.3401153>.
- [26] Q. Ai, Y. Zhang, K. Bi, X. Chen, W.B. Croft, Learning a hierarchical embedding model for personalized product search, in: SIGIR '17: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2017, pp. 305–314, URL <https://doi.org/10.1145/3077136.3080813>.
- [27] K. Bi, Q. Ai, W.B. Croft, A transformer-based embedding model for personalized product search, in: SIGIR '20: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 2045–2048, URL <https://doi.org/10.1145/3397271.3401192>.
- [28] S. Ge, Z. Dou, Z. Jiang, J.-Y. Nie, J.-R. Wen, Personalizing search results using hierarchical RNN with query-aware attention, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18, Association for Computing Machinery, New York, NY, USA, 2018, pp. 347–356, <http://dx.doi.org/10.1145/3269206.3271728>.
- [29] Z. Su, Z. Dou, Y. Zhu, X. Qin, J.-R. Wen, Modeling intent graph for search result diversification, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 736–746, <http://dx.doi.org/10.1145/3404835.3462872>.
- [30] E. Bassani, P. Kasela, G. Pasi, Denoising attention for query-aware user modeling, in: K. Duh, H. Gomez, S. Bethard (Eds.), Findings of the Association for Computational Linguistics: NAACL 2024, Association for Computational Linguistics, Mexico City, Mexico, 2024, pp. 2368–2380, URL <https://aclanthology.org/2024.findings-naacl.153>.
- [31] B.A. Zemed, B.J. Gao, Personalized search with editable profiles, in: 2017 IEEE International Conference on Big Data (Big Data), 2017, pp. 4626–4632, URL <https://doi.org/10.1109/BigData.2017.8258572>.
- [32] J. Wook Ahn, P. Brusilovsky, S. Han, Personalized search: Reconsidering the value of open user models, in: IUI '15: Proceedings of the 20th International Conference on Intelligent User Interfaces, 2015, pp. 364–369, URL <https://doi.org/10.1145/2678025.2701410>.
- [33] S. Mysore, G. Dhanania, K. Patil, S. Kallumadi, A. McCallum, H. Zamani, Bridging personalization and control in scientific personalized search, in: Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2025, <http://dx.doi.org/10.1145/3726302.3729913>, URL <https://arxiv.org/abs/2411.02790>.
- [34] J. Chien, D. Danks, Fairness vs. Personalization: Towards equity in epistemic utility, 2023, [arXiv:2309.11503](https://arxiv.org/abs/2309.11503). URL <https://arxiv.org/abs/2309.11503>.
- [35] A. Hogan, E. Blomqvist, M. Cochez, C. D'amato, G.D. Melo, C. Gutierrez, S. Kirrane, J.E.L. Gayo, R. Navigli, S. Neumaier, A.-C.N. Ngomo, A. Polleres, S.M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann, Knowledge graphs, *ACM Comput. Surv.* 54 (4) (2021) <http://dx.doi.org/10.1145/3447772>.

- [36] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, O. Yakhnenko, Translating embeddings for modeling multi-relational data, in: C. Burges, L. Bottou, M. Welling, Z. Ghahramani, K. Weinberger (Eds.), *Advances in Neural Information Processing Systems*, Vol. 26, Curran Associates, Inc., 2013, p. 1.
- [37] Z. Wang, J. Zhang, J. Feng, Z. Chen, Knowledge graph embedding by translating on hyperplanes, *Proc. AAAI Conf. Artif. Intell.* 28 (1) (2014) <http://dx.doi.org/10.1609/aaai.v28i1.8870>.
- [38] R. Reinanda, E. Meij, M. de Rijke, et al., Knowledge graphs: An information retrieval perspective, *Found. Trends® Inf. Retr.* 14 (4) (2020) 289–444.
- [39] S.E. Robertson, S. Walker, S. Jones, M. Hancock-Beaulieu, M. Gatford, Okapi at TREC-3, in: D.K. Harman (Ed.), *Proceedings of the Third Text REtrieval Conference, TREC 1994*, Gaithersburg, Maryland, USA, November 2-4, 1994, in: NIST Special Publication, vol. 500–225, National Institute of Standards and Technology (NIST), 1994, pp. 109–126, URL <http://trec.nist.gov/pubs/trec3/papers/city.ps.gz>.
- [40] R. Reiter, On closed world data bases, in: B.L. Webber, N.J. Nilsson (Eds.), *Readings in Artificial Intelligence*, Morgan Kaufmann, 1981, pp. 119–140, <http://dx.doi.org/10.1016/B978-0-934613-03-3.50014-3>, URL <https://www.sciencedirect.com/science/article/pii/B9780934613033500143>.
- [41] Y. Lin, Z. Liu, M. Sun, Y. Liu, X. Zhu, Learning entity and relation embeddings for knowledge graph completion, *Proc. AAAI Conf. Artif. Intell.* 29 (1) (2015) <http://dx.doi.org/10.1609/aaai.v29i1.9491>, URL <https://ojs.aaai.org/index.php/AAAI/article/view/9491>.
- [42] Z. Sun, Z.-H. Deng, J.-Y. Nie, J. Tang, RotatE: Knowledge graph embedding by relational rotation in complex space, in: *International Conference on Learning Representations*, 2019, arXiv:1902.10197. URL <https://arxiv.org/abs/1902.10197>.
- [43] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, M. Zhou, Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers, in: H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems*, Vol. 33, Curran Associates, Inc., 2020, pp. 5776–5788, URL https://proceedings.neurips.cc/paper_files/paper/2020/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- [44] I. Loshchilov, F. Hutter, Fixing weight decay regularization in adam, 2017, CoRR abs1711.05101. arXiv:1711.05101. URL <http://arxiv.org/abs/1711.05101>.
- [45] J. Jiang, T. Wu, G. Roumpos, H. Cheng, X. Yi, E. Chi, H. Ganapathy, N. Jindal, P. Cao, W. Wang, End-to-end deep attentive personalized item retrieval for online content-sharing platforms, in: *WWW '20: The Web Conference 2020*, Taipei, Taiwan, April 20-24, 2020, ACM / IW3C2, 2020, pp. 2870–2877.
- [46] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, *Int. Conf. Learn. Represent.* (2015).
- [47] A. Salemi, S. Mysore, M. Bendersky, H. Zamani, LaMP: When large language models meet personalization, in: L.-W. Ku, A. Martins, V. Srikumar (Eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Bangkok, Thailand, 2024, pp. 7370–7392, <http://dx.doi.org/10.18653/v1/2024.acl-long.399>, URL <https://aclanthology.org/2024.acl-long.399>.
- [48] K. Song, X. Tan, T. Qin, J. Lu, T.-Y. Liu, MPNet: Masked and permuted pre-training for language understanding, 2020, arXiv:2004.09297. URL <https://arxiv.org/abs/2004.09297>.
- [49] S. Mysore, G. Dhanania, K. Patil, S. Kallumadi, A. McCallum, H. Zamani, Memory augmented cross-encoders for controllable personalized search, 2024, arXiv:2411.02790. URL <https://arxiv.org/abs/2411.02790>.
- [50] L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web, Technical Report 1999–66, Stanford InfoLab, 1999, URL <http://ilpubs.stanford.edu:8090/422/>. Previous number=SIDL-WP-1999-0120.
- [51] A.N. Langville, C.D. Meyer, A survey of eigenvector methods for web information retrieval, *SIAM Rev.* 47 (1) (2005) 135–161, <http://dx.doi.org/10.1137/S0036144503424786>, arXiv:<https://doi.org/10.1137/S0036144503424786>.
- [52] B. Mitra, F. Diaz, N. Craswell, Learning to match using local and distributed representations of text for web search, in: *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2017, pp. 1291–1299, <http://dx.doi.org/10.1145/3038912.3052579>, URL <https://doi.org/unimib.idm.oclc.org/10.1145/3038912.3052579>.
- [53] P. Kasela, G. Pasi, R. Perego, Se-pef: a resource for personalized expert finding, in: *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, in: SIGIR-AP '23*, Association for Computing Machinery, New York, NY, USA, 2023, pp. 288–309, <http://dx.doi.org/10.1145/3624918.3625335>, URL <https://doi.org/10.1145/3624918.3625335>.
- [54] P. Kasela, M. Braga, G. Pasi, R. Perego, Se-pqa: personalized community question answering, in: *Companion Proceedings of the ACM Web Conference 2024*, in: *WWW '24*, Association for Computing Machinery, New York, NY, USA, 2024, pp. 1095–1098, <http://dx.doi.org/10.1145/3589335.3651445>, URL <https://doi.org/10.1145/3589335.3651445>.