



SE-PEF: a Resource for Personalized Expert Finding

Pranav Kasela
University of Milano-Bicocca
Milan, Italy
ISTI-CNR, Pisa, Italy
Pisa, Italy
pranav.kasela@unimib.it

Gabriella Pasi
University of Milano-Bicocca
Milan, Italy
gabriella.pasi@unimib.it

Raffaele Perego
ISTI-CNR, Pisa, Italy
Pisa, Italy
raffaele.perego@isti.cnr.it

ABSTRACT

The problem of personalization in Information Retrieval has been under study for a long time. A well-known issue related to this task is the lack of publicly available datasets to support a comparative evaluation of personalized search systems. To contribute in this respect, this paper introduces SE-PEF (StackExchange - Personalized Expert Finding), a resource useful for designing and evaluating personalized models related to the Expert Finding (EF) task. The contributed dataset includes more than 250k queries and 565k answers from 3 306 experts, which are annotated with a rich set of features modeling the social interactions among the users of a popular cQA platform. The results of the preliminary experiments conducted show the appropriateness of SE-PEF to evaluate and to train effective EF models.

CCS CONCEPTS

• Information systems → Expert search; Information retrieval.

KEYWORDS

Question Answering, Expert Finding, User Model, Personalization.

ACM Reference Format:

Pranav Kasela, Gabriella Pasi, and Raffaele Perego. 2023. SE-PEF: a Resource for Personalized Expert Finding. In *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region (SIGIR-AP '23)*, November 26–28, 2023, Beijing, China. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3624918.3625335>

1 INTRODUCTION

Expert finding (EF) is a well-studied problem in community question answering (cQA). The aim of EF in a cQA scenario, is to identify users, namely the experts, that might be able to answer correctly a given question on a specific topic. This task is important for many applications, e.g., crowd-sourcing, and for cQA platforms that wish to increase user engagement by precisely identifying the experts to whom to propose the questions about a given topic.

Personalization is gaining traction in many IR [3, 4, 7, 15] and NLP [6] tasks, but it is not largely adopted in EF due to the lack of publicly-available, large-scale datasets containing user-related information. In this research paper, building upon our previous

work [10], in which the authors presented a dataset for personalized community question answering, we introduce SE-PEF (Stack-Exchange - Personalized Expert Finding), a large dataset rich in user-level features that can be leveraged for training, evaluating and comparing both personalized and non-personalized models for the *Expert Finding* task. SE-PEF comprises around 250k questions and 560k associated answers provided by 3,306 experts, and it inherits a rich set of features modeling the social interactions within the user community. To train personalized models, we keep the user-related data as they are provided in the original dataset: users' past questions and answers, their own social autobiography, their reputation score, and the number of profile views that they have received.

In the case of EF, personalization can improve the perceived service quality in different ways. For example, when the requesting user is interested in multiple topics, identifying an expert by considering also the requesting user's interests can improve the trust in the answer received. A similar effect can be obtained by preferring experts that are closer to the requesting users based on past interactions or follower/followee dynamics. In summary, the contribution of this paper is the following:

- We provide and make available the SE-PEF dataset, as a public resource consisting of a comprehensive corpus including around 255k questions and 560k answers provided by 3306 expert users. The richness and variety of features provided with the dataset enable its use for the design and evaluation of personalized EF methods.
- We report a preliminary comparison of the performance of different EF methods applied to the questions, answers, and users in SE-PEF. The results confirm that models based on deep learning outperform in effectiveness traditional retrieval models and that by exploiting personalization features we can obtain a significant performance boost.

The rest of the paper is organized as follows. Section 2 introduces the SE-PEF dataset and reports some statistics about its content. Moreover, the section details the EF task addressed in this paper by using SE-PEF. Section 3 compares SE-PEF with respect to other publicly-available resources in the field. Section 4 presents a preliminary comparison of traditional and personalized models for EF applied to SE-PEF. In Section 5 we discuss the utility and the practical implications of the new resource. Finally, Section 6 concludes the work and draws future lines of investigation.



This work is licensed under a Creative Commons Attribution-Share Alike International 4.0 License.

SIGIR-AP '23, November 26–28, 2023, Beijing, China
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0408-6/23/11.
<https://doi.org/10.1145/3624918.3625335>

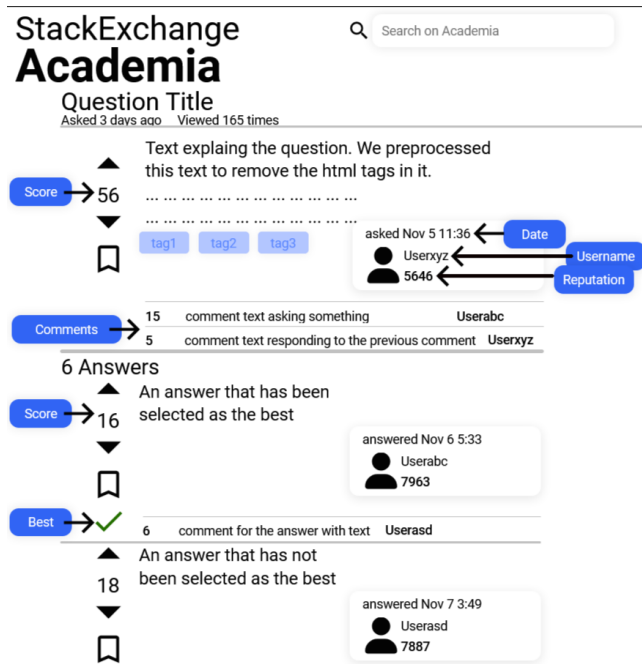


Figure 1: Illustration of StackExchange data.

2 THE SE-PEF DATASET

The dataset proposed in [10] is based on StackExchange¹ and available under a CC BY-SA 4.0 license. It comprises questions and answers from 50 different stackexchange communities, written between 2008-09-10 and 2022-09-25. There are around 1.1 million questions and 2.1 million answers. The training, validation and test splits are based on a temporal condition and are already provided on zenodo[19].

In [10] the authors show that personalization is more useful if multiple communities are used together in this dataset rather than using a single community to create the dataset. Meanwhile, previous works that use StackExchange for EF tasks focus only on a single community or a portion of a community, thus neglecting the domain diversity characterizing the questions and the various experts [8, 13, 17].

2.1 Accessing the SE-PEF dataset

SE-PEF dataset is made publicly available on zenodo² according to the conditions detailed in the included CC BY-SA 4.0. license agreement and the code used for data creation, training, hyper-parameter optimization, and testing are available on github³.

2.2 SE-PEF Definition

In the following, we introduce the specific instance of EF task in which we are interested and illustrate how to address it by using the resources in SE-PEF.

Our EF task shares the same goal as the question-answering task: satisfy users' needs in a cQA forum in the most effective way. In a cQA forum, a user may ask a question that does not have any related answers in the answer collection. Since not receiving any answer can create a sense of frustration in a user posting a question, it is important for the community and the platform to identify and, eventually, notify domain experts who may be able to answer the question correctly. Finding good matches between unanswered questions and expert users can improve remarkably the engagement with the community. In fact, on the one hand, users posting a question can receive correct answers from the alerted experts in a short time; on the other hand, expert users can dedicate their time to answering questions specifically related to their expertise rather than searching for questions that they can respond.

Formally, let \mathcal{E} be a set of expert users $\{e_1, \dots, e_k\}$. Given a question q asked by user u , the EF task consists in retrieving from \mathcal{E} a list of k experts $\{e_{q,1}, \dots, e_{q,k}\}$ ordered by their likelihood of answering correctly to question q .

StackExchange data has been used in several EF papers, e.g., in [8, 13, 17]. These works however mostly focus on solving the expert finding task for a single community. SE-PEF incorporates instead information from multiple communities to provide a dataset that can be used also to investigate models for generalist cQA forums that may not have separate channels for the discussed topics.

To create the dataset, we define as *best answer* for a given question the answer selected as the best one by the user who asked the question, if available; otherwise, we assume the best answer to the one with the highest score, if it has received a score greater than a fixed threshold γ_s ⁴. We note that this assumption, for the best answer being the most voted answer if no answer has been flagged as best by the user asking the question, is used only for the expert detection procedure, which will be explained subsequently and not as relevance judgement for the test data. In the test set we only consider the best answer, the answer explicitly labeled as such by the user asking the question. Exploiting high-scored answers as the best answers allows us to increase the number of questions successfully answered. Indeed this choice is justified by the observation that 87.6% of the answers, which are selected as best ones by the user asking the question, are also the most up-voted ones. On the other hand, we have observed that many users, once they satisfy the information need with a good answer, do not bother to mark the answer as the best.

At this point, to identify the set of experts \mathcal{E} , we follow the procedure indicated by Dargahi et al. [8] for their StackOverflow dataset:

- For each community C , let \mathcal{U} be the set of users, and \mathcal{B} the set of best answers computed as explained above in the community C . For each user $u \in \mathcal{U}$, let $\mathcal{A}_{u,C} = \{a_{u,1}, a_{u,2}, \dots, a_{u,n}\}$ be the set of answers given by u in C ;
- Remove all users who do not have at least γ_a answers selected as best answers, i.e. define:

$$\mathcal{E}' = \mathcal{U} \setminus \{u \in \mathcal{U}, \text{ s.t. } |\mathcal{A}_{u,C} \cap \mathcal{B}| < \gamma_a\}$$

¹<https://stackoverflow.com>

²<https://doi.org/10.5281/zenodo.8332747>[18]

³<https://github.com/pkasela/SE-PEF>

⁴The γ thresholds used for SE-PEF are reported at the end of Section 3.

Table 1: Comparison between SE-PEF and other cQA datasets for EF. When a specific definition of an expert is provided we distinguish normal users from experts.

Dataset	Questions	Answers	Users	Experts
StackOverflow	123 933	N/A	22 027	1845
Quora	444 138	887 771	95 915	N/A
Wondir D5	752 391 QA pairs	N/A	N/A	17 525
Wondir D20	639 233 QA pairs	N/A	N/A	5 025
Yahoo U10	32 009	97 911	2 515	N/A
Yahoo U15	28 404	89 144	1 339	N/A
Yahoo U20	25 690	80 677	870	N/A
StackExchangeGis	50 718	70 034	N/A	3 168
StackExchangeEnglish	46 692	104 453	N/A	4 781
StackExchangeCodeReview	36 947	57 622	N/A	2 242
SE-PEF	255 352	564 690	81 252	3 306

- Compute the acceptance rate for the users in \mathcal{E}' given by the ratio between the number of accepted answers and the number of total answers of the user in that community. For each user e' we define $ar_{u,C}$:

$$ar_{e',C} = \frac{|\mathcal{A}_{e',C} \cap \mathcal{B}|}{|\mathcal{A}_{e',C}|}$$

- Compute the average acceptance rate \bar{ar}_C for the users in a community and select as experts only those users who have an acceptance rate above the community average one:

$$\mathcal{E}_C = \{e \in \mathcal{E}' \text{ s.t. } ar_{e',C} \geq \bar{ar}_C\}$$

The final set of experts \mathcal{E} is defined as the union of the sets of experts found for each community. The above process ensures that the selected experts have a high level of engagement and write high-quality answers having a high acceptance rate.

In Figure 2 we show the basic structure of the JSON file provided for training, validation, and test. The `user_questions`, `user_answers` contain the identifiers (ids) of the questions and the answers, written before the current question timestamp, of the user asking the question. The `expert_questions`, and `expert_answers` contain the ids of the questions and the answers of the expert that has given the best answer. The data is provided also with a collection of questions and a collection of answers; they are two very simple JSON files, where the keys correspond to the ids of the questions and answers respectively. The values of the keys are constituted by the texts of the questions and answers respectively. The data is provided also with multiple data-frames, curated from the original data found from archive.org, which can be used to add more features. These features are described on the Stack Exchange website.⁵

3 COMPARISON WITH AVAILABLE DATASETS

Concerning the EF task, there are plenty of datasets available [12], and some of them are based on data from cQA websites. For example, StackExchange is used to create a pre-trained BERT model for the EF task in [13]. However, the work focuses only on designing an EF pre-training framework based on a specific augmented masked language model able to learn the question-expert matching task. Other EF

datasets derived from cQA forums come from: StackOverflow [8, 20], Yahoo Answers! [9, 21], Wondir [14] and Quora [22]. Recently, a domain-specific expert finding task was tackled using Avvo [1], a legal cQA website, but in this case, personalization is not possible due to the fact that users are anonymous. In Table 1 we report the basic dataset statistics of some of the commonly used datasets in EF for comparison.

A common issue with the existing datasets is that the experts are, in many cases, not well-defined, and determining what makes a user an expert is not trivial. Furthermore, most works among those previously cited either rely on a private dataset, or refer to a specific domain and make very strong assumptions simplifying the task addressed. Conversely, SE-PEF will be made publicly available, it has a well-defined definition of an expert, which is inspired by reasonable hypothesis common to other works [11, 13, 17]. Furthermore, it provides a rich set of social features usable for personalization and combines data from multiple communities, which, as we have already stated, increases dataset diversity and opens the possibility of exploiting cross-domain user information for EF.

To build the SE-PEF for EF we followed the procedure detailed in Section 2.2, by setting $\gamma_s = 5$ and $\gamma_a = 10$. Finally, we also remove from the training dataset the questions answered by experts who previously posted less than 5 answers to avoid the cold start problem for expert modeling. Using this procedure, we obtain SE-PEF, from starting from the dataset presented in [10], including 81,252 users, 3,306 experts, 252,501 queries (218,647 for training, 16,710 for validation, and 19,995 for testing), and 564,690 answers.

4 PRELIMINARY EXPERIMENTS WITH SE-PEF

This section provides a concise overview of the experimental setup and introduces the methods employed to showcase the capabilities of SE-PEF in the EF task, defined and discussed in Section 2.2. Finally, we report and discuss the results of the conducted experiments.

4.1 Experimental settings

For our EF task, we use a retrieval-based approach [16], and simply cast the EF task to a cQA task where we use the similarity scores of the retrieved documents as experts' scores. We explain this in detail in the following paragraphs.

We adopt a two-stage ranking architecture that prioritizes efficiency and recall in the first stage. The primary objective of this first stage is to select for each query a set of candidate documents that are eventually re-ranked in a second stage by a precision-oriented ranker. The first stage is based on Elastic Search⁶, and uses BM25 as a fast ranker. We use the same BM25 hyperparameters as indicated in [10]: 1 and 1.75 for `b` and `k1`, respectively. In the second, precision-oriented stage, to re-rank the retrieved documents we utilize a linear combination of the set of available scores that includes the BM25 score, the similarity score computed by a neural re-ranker, and, when used, the score computed by a personalization model exploiting the user history. In all the experiments the second stage re-ranks the top-100 results retrieved with BM25.

⁵<https://meta.stackexchange.com/questions/2677/database-schema-documentation-for-the-public-data-dump-and-sede>

⁶<https://www.elastic.co/>

```

{
  "id": "academia_49906",
  "text": "Including teaching statement in RA position application package [...]",
  "timestamp": 1438693784,
  "user_id": 3422261,
  "user_questions": ["academia_28238", "academia_28240", ...],
  "user_answers": ["genealogy_4058", "academia_37538", ...],
  "tags": ["application"],
  "expert_ids": [339125],
  "expert_questions": [[]],
  "expert_answers": [{"expatriates_2520", "academia_18991", ...}]
}

```

Figure 2: Example of a line of the jsonl file provided.

Non-personalized models. As neural re-ranker in the second stage we use the following two models used also in [10]:

- DistilBERT. This model is obtained by fine-tuning the pre-trained *distilbert-base-uncased* model⁷ for the task of answer retrieval tackled in [10]. We use the same training data and experimental settings used in [10].
- MiniLM, based on MiniLM-L6-H384-uncased⁸. This model is used as it is, without any fine-tuning.

Personalized model for EF. For building the EF personalized model we exploit the folksonomy arising from tags, very similar to the one employed in [10]. This model, which we also call TAG from hereon, aims at capturing the similarities among the topics addressed by the asker in their current and previous questions, and the ones in which a considered expert answered in the past. Given a question \mathbf{q} , asked by user \mathbf{u} at time \mathbf{t} , let $T_{u,t}$ be the set of tags assigned by \mathbf{u} to all their questions posted before \mathbf{t} (including \mathbf{q}). $T_{u,t}$ thus represents the interests of \mathbf{u} as expressed in their previous interactions. The authors of the answers to query \mathbf{q} do not have the possibility of tagging explicitly their answers, so for each answer, we consider the tags associated with the answered questions.

The way we represent the expert user is slightly different: the expertise, in this case, is based on a pre-computed, static representation T_e of each expert e in SE-PEF. This representation considers the tags T'_e of all the questions answered by e included only in the training set. To build T_e from T'_e we perform an additional step consisting in discarding the tags with a frequency lower than the median frequency of tags in T'_e . This tag pruning step reduces the noise coming from the possible presence of non-relevant tags that might have appeared as additional tags in a few questions answered but the expert might not be an expert on those topics. As for the previous task, the EF TAG score $s_{e,q}$ for expert e is finally computed as

$$s_{e,q} = \frac{|T_e \cap T_{u,q}|}{|T_{u,q}| + 1}$$

Score computation and combination. Given the list of answers A retrieved and re-ranked with the above models, we observe that

⁷<https://huggingface.co/distilbert-base-uncased>

⁸<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

Table 2: Results for the SE-PEF EF task.

Model	P@1	R@3	R@5	MRR@5	λ
BM25	0.134	0.255	0.314	0.200	-
BM25 + TAG	0.150*	0.286*	0.361*	0.226*	(.4,.6)
MiniLM _{SBERT}	0.126	0.238	0.296	0.188	-
MiniLM _{SBERT} + TAG	0.143*	0.276*	0.348*	0.217*	(.4,.6)
DistilBERT	0.147	0.274	0.334	0.216	-
DistilBERT + TAG	0.163*	0.304*	0.375*	0.240*	(.5,.5)

some users could have authored multiple answers included in A . This is potentially an important feature for characterizing their expertise. Therefore, to obtain the expert-level score, we sum up the scores assigned to all the answers in A coming from the same expert. Moreover, since the TAG model returns a score for all experts in the dataset, even those not having an answer in A , we assume that these experts receive a score contribution equal to 0 from BM25 and the non-personalized models. Finally, the scores from BM25, personalized and non-personalized models are combined by computing the weighted sum of the normalized scores from the models, using weights λ_{BM25} , $\lambda_{DistilBERT/MiniLM}$, and λ_{TAG} , with $\sum_i \lambda_i = 1$. The λ values are optimized on the validation set by performing a grid search in the interval $[0, 1]$ with step 0.1.

Evaluation Metrics. For the task of expert finding we utilize the following evaluation metrics: Precision at 1 (P@1), Recall at 3 (R@3), Recall at 5 (R@5), and Mean Reciprocal Rank at 5 (MRR@5) as our evaluation metrics. The cutoffs are set low as we prioritize identifying experts at the top of the ranked lists. All the metrics are computed by using the *ranx* library [2, 5].

4.2 Experimental Results

The results are reported instead in Table 2. The symbol * indicates a statistically significant improvement over the respective non-personalized method not using any contribution from the TAG model. Statistical significance is assessed with a Bonferroni-corrected two-sided paired student's t-test with 99% confidence. The column labeled λ reports the optimized weights, found using the

validation set, used for combining the scores computed by BM25, DistilBERT / MiniLM, and TAG models. In the cases in which the optimal weight for the BM25 score is equal to 0 – i.e., BM25 does not contribute to re-ranking – we omit BM25 from the name of the model and $\lambda_1 = 0$ from the weights column.

Differently from the cQA task tackled by the authors of [10], we observe that on EF the performance gap of DistilBERT vs. MiniLM_{SBERT} is sensibly reduced. The best-performing model among the ones tested is in fact DistilBERT + TAG which significantly outperforms both DistilBERT and MiniLM_{SBERT}. Analogously to the cQA task, personalization is very effective for EF. The contribution of the TAG model allows for significantly improving all the non-personalized methods, with a performance boost exceeding three points in MRR@5 for the DistilBERT model. By looking at the optimized λ weights reported in all three tables, we see that the TAG model contribution is much higher for the EF task ($\lambda_{TAG} \geq .5$) than for the one obtained by the authors of [10] ($\lambda_{TAG} \leq .3$).

5 UTILITY AND PREDICTED IMPACT

The SE-PEF resource we make available to the research community is a step ahead toward a fair and robust evaluation of personalization approaches in Expert Finding. The features inherited from [10] include explicit signals to create relevance judgments and a large amount of historical user-level information to design and test classical and novel personalization methods.

We expect the SE-PEF dataset being useful for many researchers and practitioners working in personalized IR and the application of machine/deep learning techniques for personalization. In recent years, significant efforts have been dedicated to the study of personalization techniques. However, there is still a lack of a comprehensive dataset for evaluating and comparing different approaches, which makes the comparison between different methods less reliable or, worse, not possible at all.

For this reason, we expect that the proposed dataset will impact the research community working on personalized EF as it provides a common ground of evaluation built on questions, answers, and experts from real users socially interacting via a community-oriented web platform.

In this proposal, the expert can have different domain backgrounds and share interests and knowledge in various communities. We also expect that training training on such rich and diverse data, like SE-PEF, should produce a more robust and generalizable model.

6 CONCLUSION AND FUTURE WORK

SE-PEF (StackExchange - Personalized Expert Finding) is an extension of a previous work [10], which presents a large real-world dataset for personalized cQA. The data inherits a rich set of user-level features modeling the interactions among the members of the online communities.

Our study provided a detailed description of the data creation and training process. Furthermore, we illustrated the methodologies adopted, explicitly focusing on IR techniques. We discussed how the similarity scores computed can be aggregated and combined to target the EF task adopted. For the retrieval, we adopted

a two-stage architecture, where the second stage utilizes for re-ranking an optimized combination of the scores generated by BM25, DistilBERT/MiniLM_{SBERT}, and TAG models.

The preliminary experiments conducted proved the effectiveness of personalization on this dataset, surpassing methods that rely on pre-trained and fine-tuned large language models by a statistically significant margin. We expect other researchers to develop more complex strategies to improve results on the SE-PEF resource. We leave such research as future work for us and the IR community working on personalized IR.

Acknowledgements. Funding for this research has been provided by: PNRR - M4C2 - Investimento 1.3, Partenariato Esteso PE00000013 - "FAIR - Future Artificial Intelligence Research" - Spoke 1 "Human-centered AI" funded by the European Union (EU) under the NextGeneration EU programme; the EU's Horizon Europe research and innovation programme EFRA (Grant Agreement Number 101093026). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the EU or European Commission-EU. Neither the EU nor the granting authority can be held responsible for them.

REFERENCES

- [1] Arian Askari, Suzan Verberne, and Gabriella Pasi. 2022. Expert Finding in Legal Community Question Answering. In *European Conference on Information Retrieval*. Springer, Springer International Publishing, Cham, 22–30.
- [2] Elias Bassani. 2022. ranx: A Blazing-Fast Python Library for Ranking Evaluation and Comparison. In *ECIR (2) (Lecture Notes in Computer Science, Vol. 13186)*. Springer, Cham, 259–264.
- [3] Elias Bassani, Pranav Kasela, and Gabriella Pasi. 2023. Denoising Attention for Query-aware User Modeling in Personalized Search. arXiv:2308.15968 [cs.LG].
- [4] Elias Bassani, Pranav Kasela, Alessandro Raganato, and Gabriella Pasi. 2022. A Multi-Domain Benchmark for Personalized Search Evaluation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 3822–3827. <https://doi.org/10.1145/3511808.3557536>
- [5] Elias Bassani and Luca Romelli. 2022. Ranx.Fuse: A Python Library for Metasearch. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 4808–4812. <https://doi.org/10.1145/3511808.3557207>
- [6] Marco Braga, Alessandro Raganato, and Gabriella Pasi. 2023. Personalization in BERT with Adapter Modules and Topic Modelling. In *Italian Information Retrieval Workshop*. <https://api.semanticscholar.org/CorpusID:260500844>
- [7] Silvia Calegari and Gabriella Pasi. 2013. Personal ontologies: Generation of user profiles based on the YAGO ontology. *Information Processing & Management* 49, 3 (2013), 640–658. <https://doi.org/10.1016/j.ipm.2012.07.010> Personalization and Recommendation in Information Access.
- [8] Arash Dargahi Nobari, Sajad Sotudeh Gharebagh, and Mahmood Neshati. 2017. Skill Translation Models in Expert Finding. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (Shinjuku, Tokyo, Japan) (SIGIR '17)*. Association for Computing Machinery, New York, NY, USA, 1057–1060. <https://doi.org/10.1145/3077136.3080719>
- [9] Jinwen Guo, Shengliang Xu, Shenghua Bao, and Yong Yu. 2008. Tapping on the Potential of Q&a Community by Recommending Answer Providers. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (Napa Valley, California, USA) (CIKM '08)*. Association for Computing Machinery, New York, NY, USA, 921–930. <https://doi.org/10.1145/1458082.1458204>
- [10] Pranav Kasela, Gabriella Pasi, and Raffaele Perego. 2023. SE-PQA: Personalized Community Question Answering. arXiv:2306.16261 [cs.LG]
- [11] Zeyu Li, Jyun-Yu Jiang, Yizhou Sun, and Wei Wang. 2019. Personalized Question Routing via Heterogeneous Network Embedding. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (Jul. 2019), 192–199. <https://doi.org/10.1609/aaai.v33i01.3301192>
- [12] Shuyi Lin, Wenxing Hong, Dingding Wang, and Tao Li. 2017. A survey on expert finding techniques. *Journal of Intelligent Information Systems* 49, 2 (2017), 255–279.
- [13] Hongtao Liu, Zhepeng Lv, Qing Yang, Dongliang Xu, and Qiyao Peng. 2022. ExpertBERT: Pretraining Expert Finding. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM*

- '22). Association for Computing Machinery, New York, NY, USA, 4244–4248. <https://doi.org/10.1145/3511808.3557597>
- [14] Xiaoyong Liu, W Bruce Croft, and Matthew Koll. 2005. Finding experts in community-based question-answering services. In *Proceedings of the 14th ACM international conference on Information and knowledge management*. Association for Computing Machinery, New York, NY, USA, 315–316.
- [15] Zhengyi Ma, Zhicheng Dou, Guanyue Bian, and Ji-Rong Wen. 2020. PSTIE: Time Information Enhanced Personalized Search. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (Virtual Event, Ireland) (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 1075–1084. <https://doi.org/10.1145/3340531.3411877>
- [16] Craig Macdonald and Iadh Ounis. 2006. Voting for Candidates: Adapting Data Fusion Techniques for an Expert Search Task. In *Proceedings of the 15th ACM International Conference on Information and Knowledge Management (Arlington, Virginia, USA) (CIKM '06)*. Association for Computing Machinery, New York, NY, USA, 387–396. <https://doi.org/10.1145/1183614.1183671>
- [17] Qiyao Peng, Hongtao Liu, Yinghui Wang, Hongyan Xu, Pengfei Jiao, Minglai Shao, and Wenjun Wang. 2022. Towards a Multi-View Attentive Matching for Personalized Expert Finding. In *Proceedings of the ACM Web Conference 2022 (Virtual Event, Lyon, France) (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 2131–2140. <https://doi.org/10.1145/3485447.3512086>
- [18] Kasela Pranav, Pasi Gabriella, and Perego Raffaele. 2023. *SE-PEF: a Resource for Personalized Expert Finding*. <https://doi.org/10.5281/zenodo.8332748>
- [19] Kasela Pranav, Pasi Gabriella, and Perego Raffaele. 2023. *SE-PQA: a Resource for Personalized Community Question Answering*. <https://doi.org/10.5281/zenodo.7940964>
- [20] Fatemeh Riahi, Zainab Zolaktaf, Mahdi Shafiei, and Evangelos Milios. 2012. Finding Expert Users in Community Question Answering. In *Proceedings of the 21st International Conference on World Wide Web (Lyon, France) (WWW '12 Companion)*. Association for Computing Machinery, New York, NY, USA, 791–798. <https://doi.org/10.1145/2187980.2188202>
- [21] Fei Xu, Zongcheng Ji, and Bin Wang. 2012. Dual Role Model for Question Recommendation in Community Question Answering. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval (Portland, Oregon, USA) (SIGIR '12)*. Association for Computing Machinery, New York, NY, USA, 771–780. <https://doi.org/10.1145/2348283.2348387>
- [22] Zhou Zhao, Lijun Zhang, Xiaofei He, and Wilfred Ng. 2015. Expert Finding for Question Answering via Graph Regularized Matrix Completion. *IEEE Transactions on Knowledge and Data Engineering* 27, 4 (2015), 993–1004. <https://doi.org/10.1109/TKDE.2014.2356461>