

Personalized Query Auto-Completion Through a Lightweight Representation of the User Context

Manojkumar Rangasamy Kannadasan
eBay Inc.
San Jose, CA
mkannadasan@ebay.com

Grigor Aslanyan
eBay Inc.
San Jose, CA
gaslanyan@ebay.com

ABSTRACT

Query Auto-Completion (QAC) is a widely used feature in many domains, including web and eCommerce search. This feature suggests full queries based on a prefix of a few characters typed by the user. QAC has been extensively studied in the literature in the recent years, and it has been consistently shown that adding personalization features can significantly improve the performance of the QAC model. In this work we propose a novel method for personalized QAC that uses lightweight embeddings learnt through *fastText* [2, 14]. We construct an embedding for the user context queries, which are the last few queries issued by the user. We also use the same model to get the embedding for the candidate queries to be ranked. We introduce ranking features that compute the distance between the candidate queries and the context queries in the embedding space. These features are then combined with other commonly used QAC ranking features to learn a ranking model using the state of the art LambdaMART ranker [3]. We apply our method to a large eCommerce search engine (eBay) and show that the ranker with our proposed feature significantly outperforms the baselines on all of the offline metrics measured, which includes Mean Reciprocal Rank (MRR), Success Rate (SR), Mean Average Precision (MAP), and Normalized Discounted Cumulative Gain (NDCG). Our baselines include the Most Popular Completion (MPC) model which is a commonly used baseline in the QAC literature, as well as a ranking model without our proposed features. The ranking model with the proposed features results in a 20 – 30% improvement over the MPC model on all metrics. We obtain up to a 5% improvement over the baseline ranking model for all the sessions, which goes up to about 10% when we restrict to sessions that contain the user context. Moreover, our proposed features also significantly outperform text based personalization features studied in the literature before, and adding text based features on top of our proposed embedding based features results only in minor improvements.

ACM Reference Format:

Manojkumar Rangasamy Kannadasan and Grigor Aslanyan. 2019. Personalized Query Auto-Completion Through a Lightweight Representation of the User Context. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR 2019 eCom)*. ACM, New York, NY, USA, 8 pages.

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

SIGIR 2019 eCom, July 2019, Paris, France

© 2019 Copyright held by the owner/author(s).

1 INTRODUCTION

Query Auto-Completion (QAC) is a common feature of most modern search engines. It refers to the task of suggesting full queries after the user has typed a prefix of a few characters [6]. QAC can significantly reduce the number of characters typed [26], which is especially helpful to users on mobile devices. QAC can also help reduce the number of spelling errors in queries. In cases when the user is not really sure how to formulate the query, QAC can be of great help. It has been shown that QAC can greatly improve user satisfaction [24]. Moreover, this can reduce the overall search duration, resulting in a lower load on the search engine [1]. Currently QAC has a wide range of applications, including search (such as web, eCommerce, email), databases, operating systems, development environments [6].

Query Auto-Completion has been extensively studied in the literature in the recent years. A detailed survey of the work prior to 2016 can be found in [6], which broadly classifies QAC approaches into two main categories - heuristic models and learning based models. Heuristic models use a few different sources for each possible query completion and compute a final score. These approaches do not use a large variety of features. In contrast, learning based approaches treat the problem as a ranking problem and rely on the extensive research in the literature in the learning-to-rank (LTR) field [15]. Learning based approaches rely on a large number of features and generally outperform heuristic models [6]. The features for both of these approaches can be broadly characterized into three groups - time-sensitive, context-based, and demography based. Time-sensitive features model the query popularity and changes over time, such as weekly patterns. Demographic based features, such as gender and age, are typically limited and may be hard to access. In contrast, context based features rely on the user's previous search activity (short term, as well as long term) to suggest new query completions. This data is essentially free, making context-based features an attractive approach for personalizing QAC. Context-based features for LTR models will be the focus of this work.

In this paper we propose a novel method to learn the query embeddings [2, 14] using a simple and scalable technique and use it to measure similarity between user context queries and candidate queries to personalize QAC. We learn the embeddings in a semi-supervised fashion using *fastText* by taking all the queries in a session as a single document. We design features that measure the similarity between the context and candidate queries, which are then incorporated into a learning-to-rank model. We use the state of the art LambdaMART model [3] for ranking candidate queries for QAC. Even though embedding based features have been studied

for QAC in the literature before, as discussed in Section 2, our work makes the following novel contributions:

- A lightweight and scalable way to represent the user’s context in the embedding space.
- Simple and robust ranking features based on such embeddings for QAC, which can be used in any heuristic or LTR model.
- Training and evaluation of a pairwise LambdaMART ranker for QAC using the proposed features. We show that our proposed features result in significant improvements in offline metrics compared with state-of-the-art baselines.
- We also compare and combine text based features with embedding based features and show that embedding based features result in larger improvements in offline metrics.

The rest of the paper is organized as follows. Section 2 discusses some of the related work in the literature. In Section 3 we describe our methodology. In Section 4 we describe our datasets and experiments. We summarize our work and discuss possible future research in Section 5.

2 RELATED WORK

The user’s previously entered text is used for personalized QAC by Bar-Yossef and Kraus [1]. The method, called NearestCompletion, computes the similarity of query completion candidates to the context queries (user’s previously entered queries), using term-weighted vectors for queries and contexts and applying cosine similarity. This method results in significant improvements in MRR. In addition, the authors proposed the MPC approach, which is based on the overall popularity of the queries matching the given prefix. MPC is a straightforward heuristic approach with good performance and is typically used as a baseline for more complex approaches. We use MPC as one of the baselines in this work as well.

The user’s long term search history is used in [5] to selectively personalize QAC, where a trade-off between query popularity and search context is used to encode the ranking signal. Jiang et. al. [13] study user reformulation behavior using textual features. Shokouhi [22] studies QAC personalization using a combination of context based textual features and demographic features, and shows that the user’s long term search history and location are the most effective for QAC personalization. Su et. al. [25] propose the framework EXOS for personalizing QAC, which also relies on textual features (token level). Jiang et. al. [8] use history-level, session-level, and query-level textual features for personalized QAC. Fei et. al. [4] study features on the observed and predicted search popularity both for longer and shorter time periods for learning personalized QAC. Diversification of personalized query suggestion is studied in [7].

Recurrent Neural Networks (RNN) [11] have also been studied for QAC. Three RNN models - session-based, personalized, and attention based, have been proposed in [12]. Fiorini and Lu [9] use user history based features as well as time features as input to an RNN model. [19] uses RNNs to specifically improve QAC performance on previously unseen queries. An adaptable language model is proposed in [10] for personalized QAC.

Word embeddings, such as *word2vec* [17], *glove* [20], and *fastText* [2, 14], have become increasingly popular in the recent years for a large variety of tasks, including computing similarity between words. Embeddings have also been studied in the context of QAC. Specifically, Mitra [18] studies a Convolutional Latent Semantic Model for distributed representations of queries. Query similarity based on *word2vec* embeddings is studied in [21] where the features are combined with the MPC model. In Section 3, we explain our approach of learning embeddings for the user context in a simple and scalable fashion and the usage of these embeddings and text based features to personalize QAC.

3 PERSONALIZED QUERY AUTO-COMPLETION WITH REFORMULATION

A search session is defined as a sequence of queries $\langle q_1, q_2, \dots, q_T \rangle$ issued by a user within a particular time frame. A query consists of a set of tokens. If the user types a prefix p_T and ends up issuing the query q_T , then the user’s context is the previous queries issued till time step T , $\langle q_1, q_2, \dots, q_{T-1} \rangle$. For example, if the queries issued in a session is $\langle nike, adidas, shoes \rangle$, the prefix used to issue the query *shoes* is *sh*, then $\langle q_1, q_2, \dots, q_{T-1} \rangle = \langle nike, adidas \rangle$, $p_T = sh$, $q_T = shoes$. Given a prefix p_T , the user context $\langle q_1, q_2, \dots, q_{T-1} \rangle$ and candidate queries Q_T matching the prefix, our goal is to provide ranking for the queries $q \in Q_T$ such that we have the best ranking for q_T . The ranking score can be considered as $P(Q_T | \langle q_1, q_2, \dots, q_{T-1} \rangle)$. This can be solved using the learning to rank framework.

The influence of user context features towards the prediction accuracy has already been studied in [13, 18]. In this paper we propose a simple and scalable way to understand the user’s context using query embeddings and use multiple distance related features to compare the user’s context to the candidate queries Q_T .

3.1 Learning Query Representation for Reformulation

Continuous text representations and embeddings for a text can be learnt through both supervised [18] and semi-supervised approaches [2, 14, 17, 20]. In this paper, we learn the query representations via semi-supervised techniques. We use the publicly available *fastText* library [2, 14] for efficient representation learning to learn the query embeddings. The *fastText* model learns subword representations while taking into account morphology. The model considers subword units, and represents a word by the sum of its character n -grams. The word *iphone* with character n -grams ($n = 3$) is represented as:

“⟨ip⟩”, “iph”, “pho”, “hon”, “one”, “ne”

Some of the previous work learns distinct vector representations for the words thereby ignoring internal structure of the words [17]. If we have a dictionary of n -grams of size G , then the set of n -grams in a word w is denoted as $Gr_w \in \{1, 2, \dots, G\}$. We use the *skipGram* model where the goal is to independently predict the presence or absence of the context words. The problem is framed as a binary classification task. For the word at position t we consider all context words as positive examples and sample negatives at random from the dictionary as described in [2, 14]. For a context word w_c , we

use the negative log likelihood, $l : x \mapsto \log(1 + e^{-x})$, for the binary logistic loss. The objective function is defined as:

$$\sum_{t=1}^T \left[\sum_{c \in \text{Context}} l(s(w_t, w_c)) + \sum_{n \in N_{t,c}} l(-s(w_t, n)) \right] \quad (1)$$

where w_t is the target word, w_c is the context word, $N_{t,c}$ is a set of negative examples sampled from the vocabulary. The scoring function, $s(w, c)$ is defined as

$$s(w, c) = \sum_{g \in G_w} z_g^T v_c \quad (2)$$

where z_g is the vector representation of each n -gram of a word w and v_c is the vector representation of the context. Our goal is to learn scalable and lightweight embeddings for queries based on their reformulation behavior across different users. Here we represent all the queries issued in a session $\langle q_1, q_2, \dots, q_T \rangle$ as one document in the training data. By learning the subword representations using the probability of words in the context of other words present in the queries issued in the same context, we are able to provide a simple and scalable way to encode the query reformulation behavior in the embedding space. We are also able to learn the syntactic and semantic similarity between the vocabulary.

We learn query representations by mining 3 days of eBay’s search logs to get the query reformulations issued by the user. The query log is segmented into sessions with a 30 minute session boundary as followed in [13]. Based on this definition of a session boundary, we collect different queries issued by the users within that session. We remove special characters in the query and convert them to lowercase. We also filter out sessions with only one query in the session. For example, if the user issues q_1, q_2, \dots, q_T in a session, then all of these queries $\langle q_1, q_2, \dots, q_T \rangle$ together, separated by whitespace, are considered as one user context. For example, if a session contains 2 queries in the user context, “*iphone*”, “*iphone xs case*”, then a single document for training will be represented as “*iphone iphone xs case*”.

For training unsupervised character n -Gram representations we consider each user context as one document sample. We tune the model hyperparameters by fixing the dimension of subword representations as 50, minimum occurrence of the words in the vocabulary to be 20 and hierarchical softmax as the loss function. The other hyperparameters of the model are tuned based on the *Embedding Features* model described in Section 4.3. The number of unique words in the vocabulary used to train the model is 189,138. The user context $\langle q_1, q_2, \dots, q_T \rangle$ is then converted to multiple vector representations. Similar vector representations are also created for all candidate target queries in the dataset.

3.2 User Context Features

In this section we propose different user context features based on the queries issued in the session. Vector representations are created for both the individual queries as well as the entire context taking all queries in the session. v_C represents the user context vector and v_{q_T} represents the vector for one query at time step T . We develop four features based on the query representations learned in the previous section. We denote these features as *Embedding Features*. One embedding feature is based on all the queries in the

Table 1: Textual features based on the user context defined across 3 categories. A subset of features are highlighted in the table. Rest of the features are derived from them.

Category	Examples
Token (16 features)	ratio of new terms ratio of used terms average terms in previous queries median terms in previous queries trend of number of terms unique terms added from last query unique terms retained from last query unique terms removed from last query unique terms added from all previous queries occurrence of terms in previous queries
Query (7 features)	frequency in previous queries character n-gram similarity with previous queries token n-gram similarity with previous queries
Session (3 features)	position in session unique terms in session common terms in session

context. Since the median number of searches in a session is approx 3, we considered up to 3 queries previously issued by the user for generating the remaining embedding features. The *Embedding Features* are computed as a distance between 2 vectors using *cosine similarity* [23].

- **user_context_cosine**: Cosine distance between the user context vector v_C and the current target query v_{q_T} .
- **prev_query1_cosine**: Cosine distance between the query vector $v_{q_{T-1}}$ and the current target query v_{q_T} .
- **prev_query2_cosine**: Cosine distance between the query vector $v_{q_{T-2}}$ and the current target query v_{q_T} .
- **prev_query3_cosine**: Cosine distance between the query vector $v_{q_{T-3}}$ and the current target query v_{q_T} .

In addition to the *Embedding Features*, we also developed various *Textual Features* comparing the user context and the current target query to be ranked as defined in Table 1. We categorize them into three categories, namely *Token*, *Query*, and *Session*. There is a large overlap between the features defined in Table 1 and the features defined in [13, 22]. A query q_T can contain multiple tokens. Users may add or remove tokens between 2 consecutive queries in a session. Based on analyzing the user sessions, between queries q_T and q_{T-1} , tokens can either be added and/or removed. These token reformulation user behavior can be encoded via 16 features, described in Table 1, representing the effectiveness of the tokens in the context C and the target query q_T . Similarly, Query level features represent how users reformulate the queries in a session through repetition and textual similarity between q_T and q_{T-1} . The Session level features represent how users reformulate their queries without taking into account the relationship to the target query q_T .

4 EXPERIMENTS

4.1 Dataset and Experiment Setting

We conduct our ranking experiments on a large scale search query dataset sampled from the logs of eBay Search engine. The query log is segmented into sessions with a 30 minute session boundary




Figure 1: The end to end architecture of the *Textual_Embedding* model. The architecture for the other models is similar, except that some of the features will be excluded.

as described in [13]. For ranking experiments, we do not filter out sessions containing a single query. This is to make sure that we have a single learning to rank model powering sessions with and without user context. The dataset obtained based on the above logic results in about 53% of the sessions with user’s context. This gives us good coverage of user context features to learn a global model.

The labeling strategy used in [13, 22] assume there is at least one query in the context, remove target queries q_T not matching the prefix p_T , setting the first character of the prefix p_T based on q_T . In our method, we use a slightly different labeling strategy for building the personalized QAC. We sample a set of impressions from search logs. For each issued query q_T , we capture the prefix p_T that led to the search page. This is marked as a positive label. For the same prefix p_T we identify rest of the candidate queries $Q_T \setminus q_T$ that were shown to the user and did not lead to the impression. They are marked as negative labels. We also retain sessions without user context.

The above training data now consists of labeled *prefix-query* pairs. To learn the performance of the lightweight query representation of reformulations, we use LambdaMART [3] as the choice of learning to-rank algorithms, a boosted tree version of LambdaRank. LambdaMART is considered as one of the state-of-the-art learning to rank algorithms and has won the Yahoo! Learning to Rank Challenge (2010) [13]. We use a pairwise ranking model and fine tune our parameters based on the *Baseline_Ranker* defined in Section 4.2. We fix these parameters to train and evaluate our models across all of our experiments.

4.2 Baseline System

To evaluate our new personalized QAC ranker we establish two baseline ranking algorithms.

- **MPC:** The *Most Popular Completion* model [1] predicts and provides users with candidate queries which are ranked by the popularity of the query. Popularity of a query is defined as the number of times the query has been issued by all the users in the past.
- **Baseline_Ranker:** The baseline ranker is a Learning to Rank model built using the same methodology for creating and labeling the dataset. The features used in building the model are prefix features, target query features and prefix-query features. We refer to these features as *Baseline_Features*. The

hyperparameters used for the LambdaMART model are exactly the same as in all the experiments for the personalized ranker.

4.3 Personalized Ranking Models

We have developed three personalized ranking models with different combinations of user context features, as described in Section 3.2. These ranking models are compared against the two baseline rankers by experimentally evaluating the improvements on eBay datasets. The results are presented in Section 4.5.

- **Textual:** Ranker with *Baseline_Features* and *Textual_Features* representing the user context.
- **Embedding:** Ranker with *Baseline_Features*, as well as *Embedding_Features* representing the user context.
- **Textual_Embedding:** Ranker with *Baseline_Features*, *Textual_Features*, and *Embedding_Features* representing the user context.

For all of the ranking models we first get the top N candidate queries from the *MPC* model and re-rank them with the ranking model. We show the full end to end architecture for the *Textual_Embedding* model in Figure 1. The architecture for the other models is similar except that they will only include a subset of the features.

4.4 Evaluation Metrics

The quality of our predictions can be measured using the following metrics:

- **Mean Reciprocal Rank (MRR)** - the average of the reciprocal ranks of the target queries in the QAC results. Given a test dataset S , the MRR for algorithm A is computed as

$$MRR(A) = \frac{1}{|S|} \sum_{C_T, q_T \in S} \frac{1}{hitRank(A, C_T, q_T)} \quad (3)$$

where C_T represents the user context at time step T , q_T represents the relevant target query, and the function *hitRank* computes the rank of the relevant query based on the order created by algorithm A . Relevant query refers to the clicked query in QAC.

- **Success Rate at Top-K (SR@K)** - the average percentage of relevant queries ranked at or above the position K in the

Table 2: Offline evaluation metrics. We show the ratio of the metrics for four of the ranking models to the MPC model on both test datasets - all data and filtered data to include full coverage for user context.

Dataset	Measure	Baseline_Ranker	Textual_Features	Embedding_Features	Textual_Embedding_Features
Whole	MRR	1.26	1.30	1.31	1.31
	nDCG	1.30	1.32	1.33	1.33
	SR@1	1.24	1.29	1.31	1.31
	SR@3	1.23	1.26	1.27	1.27
	MAP	1.26	1.30	1.31	1.31
	MAP@1	1.24	1.29	1.31	1.31
	MAP@3	1.23	1.27	1.29	1.29
	MRR	1.27	1.34	1.37	1.37
User Context Only	nDCG	1.30	1.35	1.37	1.37
	SR@1	1.26	1.37	1.42	1.41
	SR@3	1.23	1.30	1.33	1.34
	MAP	1.27	1.34	1.37	1.37
	MAP@1	1.26	1.37	1.42	1.41
	MAP@3	1.24	1.33	1.37	1.37

ranked list from QAC. In this paper we will consider only SR@1, SR@2, SR@3.

- **Normalized Discounted Cumulative Gain (nDCG)** - The Discounted Cumulative Gain (DCG) represents the usefulness or gain of the query based on its position in the ranked list from QAC. DCG penalizes the relevance of the query logarithmically based on the position of the query in the ranked list. DCG is defined as

$$DCG_q = \sum_{i=1}^P \frac{2^{rel_i} - 1}{\log_2(i + 1)} \tag{4}$$

where i denotes the rank and rel_i is the relevance of query at rank i . For our purposes rel_i takes values 0 or 1.

$nDCG$ is defined as normalized DCG. Namely, it is the ratio of DCG to $IDCG$ (ideal DCG):

$$nDCG_q = \frac{DCG_q}{IDCG_q} \tag{5}$$

where $IDCG_q$ is the maximum possible value of DCG_q for any ranker.

The overall performance of the ranking algorithm A is measured by the average $nDCG$ across all queries in the dataset:

$$nDCG = \frac{\sum_{q=1}^Q nDCG_q}{Q} \tag{6}$$

- **Mean Average Precision (MAP)** - the mean of the average precision scores for each query across the entire set of queries:

$$MAP = \frac{\sum_{q=1}^Q AvgPrecision(q)}{Q} \tag{7}$$

4.5 Results

We perform our evaluation in two phases. Firstly, we evaluate the quality of our query representations. Secondly, we evaluate the user context embeddings against the user context based textual features using a Learning to Rank framework [3].

To evaluate our query representations we sample a few words across different verticals like fashion, electronics, home and garden,




Figure 2: A three-dimensional t-SNE plot using the vectors learned from user query reformulation, showing how similar intent words are modeled in the embedding space.

to evaluate if the embeddings are representing the syntactic and semantic knowledge of the queries learnt from the query reformulation behavior. We use t-SNE [16] to visualize the embeddings for these sampled queries and show that words like *samsung*, *galaxy*, *tv* are close to each other in the embedding space and far from queries like *adidas* and *iphone*. This verifies that our query embeddings have good subword information to represent the user context in the embedding space. The t-SNE plot for a small sample of queries is shown in Figure 2.

Offline metrics MRR , $SR@k$, $nDCG$, MAP , and $MAP@k$ are shown in Table 2, where we have normalized the metrics with respect to the MPC model. We show results for the whole test dataset, which includes queries with and without user context, as well as the dataset with user context only. To assess statistical significance we have performed 1,000 bootstrap samples over the test queries and computed 95% confidence intervals using those samples. The metrics, together with the 95% confidence intervals, are plotted in Figure 3, where the plots on the left are for the whole dataset and

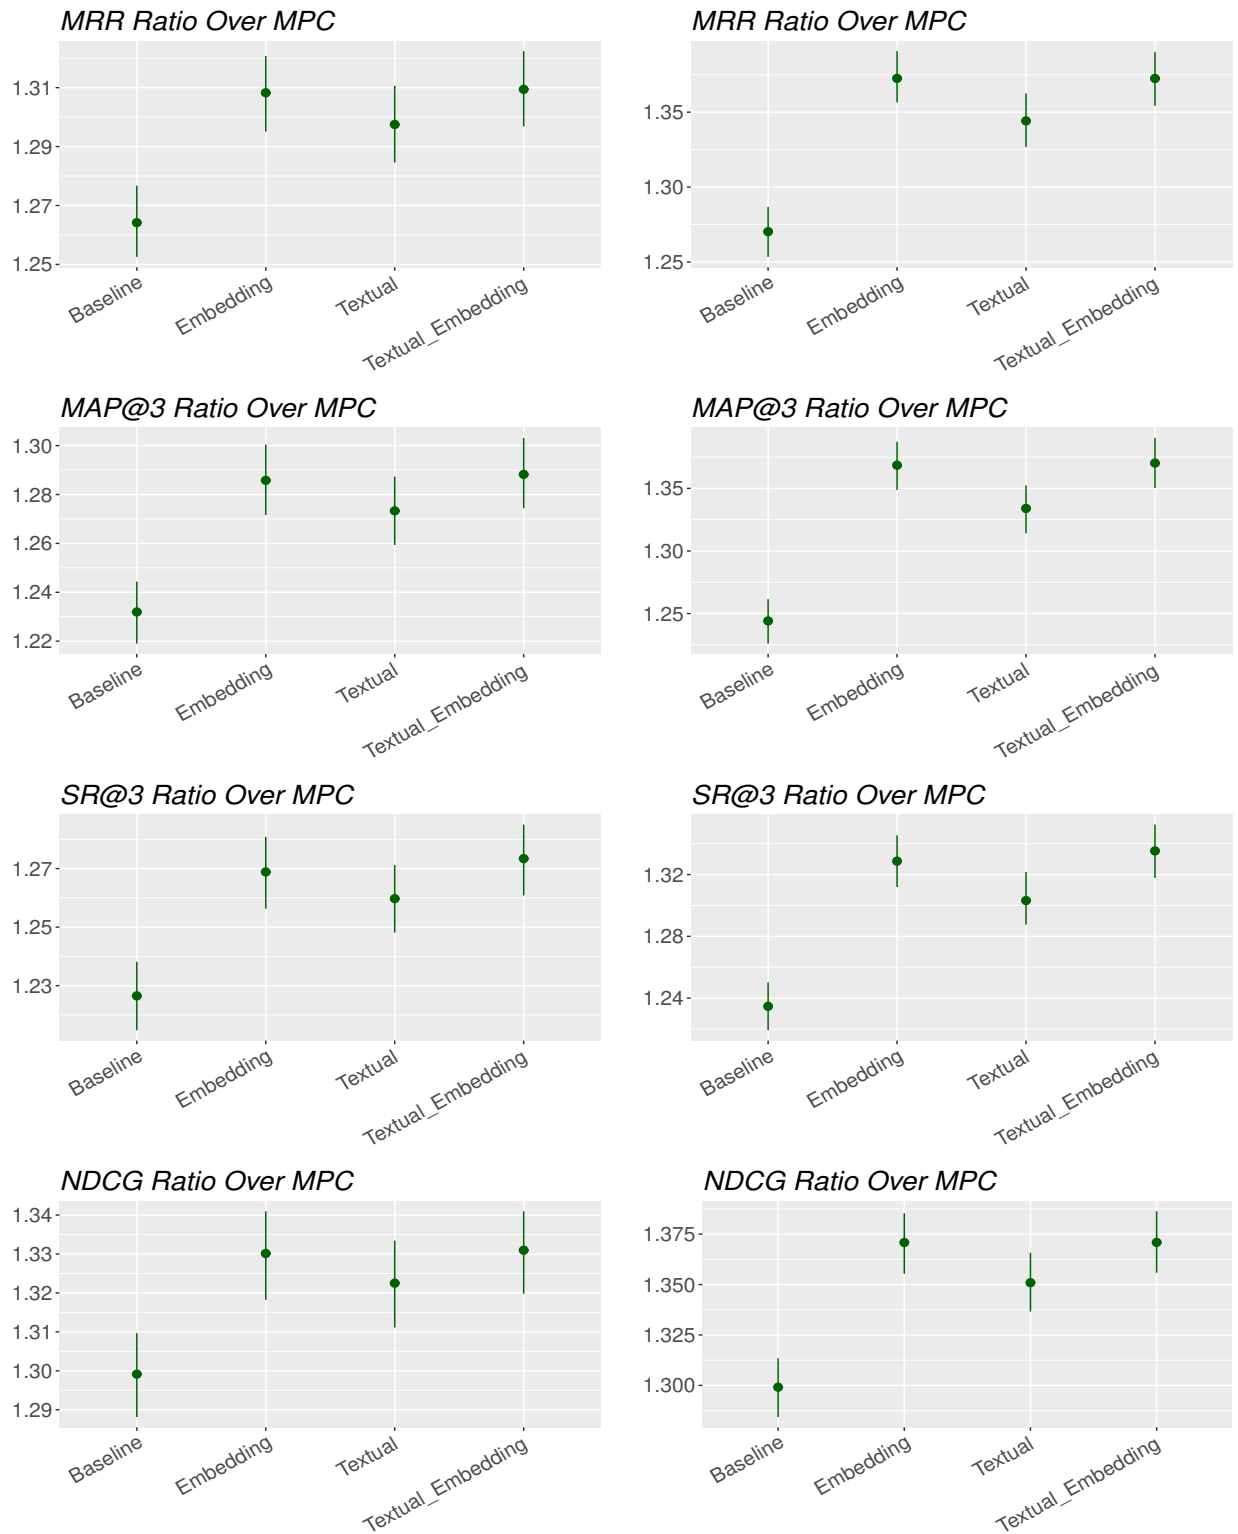


Figure 3: Metrics ratio to MPC for the whole dataset (left) and the user context only dataset (right). Error bars are computed using 1,000 bootstrap samples of the test queries.

the plots on the right are for the context only dataset. We have only plotted one variant of each metric since the others are very similar.

Our results show that all of the LTR models result in 20 – 30% improvements over the MPC model. All three models with contextualization features outperform the **Baseline Ranker** on all the metrics statistically significantly. For example, for **MAP@3 Embedding** outperforms **Baseline Ranker** by 5% on the whole dataset and 10% for the context only dataset. The **Embedding** model also outperforms **Textual** with an improvement of 1.5% for **MAP@3** on the whole dataset and 3% for the context only dataset. The **Textual Embedding** model performs very similarly to **Embedding** which implies that the embedding based features proposed in this work capture all of the information in the textual features (from the perspective of the ranking model), and provide additional significant improvements.

4.6 Feature Analysis

In this section we analyze the user context embedding features through partial dependence plots shown in Figure 4. The partial dependence plot for the *user_context_cosine* feature clearly indicates that the cosine similarity between the user context $(q_1, q_2, \dots, q_{T-1})$ and the target query q_T has a linear relationship with the target. The embedding features based on individual time step (*prev_query1_cosine*, *prev_query2_cosine*, *prev_query3_cosine*) also show a clear monotonic relationship.




Figure 4: Partial dependence plots for user context embedding features learnt from query reformulation.

5 SUMMARY AND FUTURE WORK

In this work we have presented a simple and scalable approach to learn lightweight vector representations (embeddings) for the query reformulations in a user session. These query representations exhibit both syntactic and semantic similarities between different

queries and enable them to model the user context seamlessly. We have leveraged these lightweight embeddings to represent the user context in our personalized ranker for Query Auto-Completion. Different combinations of user context features are created, including textual and embedding features on top of our baseline ranker. We have applied these personalization features to a large scale commercial search engine (eBay) and experimentally verified significant improvements on all the offline ranking metrics. We have evaluated our personalized ranker on the entire dataset and a dataset restricted to sessions containing the user context. We see the biggest improvements on the user context filtered dataset. Furthermore, we show that the ranking model with embedding features outperforms the model with the textual features, whereas the model with combined textual and embedding features results in only minor improvements on top of the model with embedding features alone. The minor improvements from the textual features is likely due to the session level features which are agnostic of the queries in the context. As a future work, we would like to explore different representation learning techniques like *sent2vec*, *doc2vec*, and sequence models, to understand the user context better and incorporate them in the personalized ranker. We also plan to explore the trade offs between short term and long term user contexts in QAC. Lastly, the user context vectors provide a simple and scalable way to understand the user sessions which can be utilized for personalizing different parts of search and recommender systems.

REFERENCES

- [1] Ziv Bar-Yossef and Naama Kraus. 2011. Context-sensitive Query Auto-completion. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. ACM, New York, NY, USA, 107–116. <https://doi.org/10.1145/1963405.1963424>
- [2] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [3] Christopher JC Burges. 2010. From ranknet to lambdarank to lambdamart: An overview. *Learning* 11, 23-581 (2010), 81.
- [4] Fei Cai, Wanyu Chen, and Xinliang Ou. 2017. Learning search popularity for personalized query completion in information retrieval. *Journal of Intelligent & Fuzzy Systems* 33, 4 (2017), 2427–2435.
- [5] Fei Cai and Maarten de Rijke. 2016. Selectively Personalizing Query Auto-Completion. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16)*. ACM, New York, NY, USA, 993–996. <https://doi.org/10.1145/2911451.2914686>
- [6] Fei Cai and Maarten de Rijke. 2016. A Survey of Query Auto Completion in Information Retrieval. *Foundations and Trends® in Information Retrieval* 10, 4 (2016), 273–363. <https://doi.org/10.1561/15000000055>
- [7] Wanyu Chen, Fei Cai, Honghui Chen, and Maarten de Rijke. 2017. Personalized query suggestion diversification. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 817–820.
- [8] Jiang Danyang, Fei Cai, and Honghui Chen. 2018. Personalizing Query Auto-Completion for Multi-Session Tasks. 203–207. <https://doi.org/10.1109/CCET.2018.8542201>
- [9] Nicolas Fiorini and Zhiyong Lu. 2018. Personalized neural language models for real-world query auto completion. *arXiv preprint arXiv:1804.06439* (2018).
- [10] Aaron Jaech and Mari Ostendorf. 2018. Personalized language model for query auto-completion. *arXiv preprint arXiv:1804.09661* (2018).
- [11] L. C. Jain and L. R. Medsker. 1999. *Recurrent Neural Networks: Design and Applications* (1st ed.). CRC Press, Inc., Boca Raton, FL, USA.
- [12] D. Jiang, W. Chen, F. Cai, and H. Chen. 2018. Neural Attentive Personalization Model for Query Auto-Completion. In *2018 IEEE 3rd Advanced Information Technology, Electronic and Automation Control Conference (IAEAC)*. 725–730. <https://doi.org/10.1109/IAEAC.2018.8577694>
- [13] Jyun-Yu Jiang, Yen-Yu Ke, Pao-Yu Chien, and Pu-Jen Cheng. 2014. Learning User Reformulation Behavior for Query Auto-completion. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '14)*. ACM, New York, NY, USA, 445–454. <https://doi.org/10.1145/2600428.2609614>

- [14] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2016. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759* (2016).
- [15] Tie-Yan Liu. 2009. Learning to Rank for Information Retrieval. *Foundations and Trends® in Information Retrieval* 3, 3 (2009), 225–331. <https://doi.org/10.1561/15000000016>
- [16] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [17] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'13)*. Curran Associates Inc., USA, 3111–3119. <http://dl.acm.org/citation.cfm?id=2999792.2999959>
- [18] Bhaskar Mitra. 2015. Exploring session context using distributed representations of queries and reformulations. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*. ACM, 3–12.
- [19] Dae Hoon Park and Rikio Chiba. 2017. A neural language model for query auto-completion. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 1189–1192.
- [20] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543. <http://www.aclweb.org/anthology/D14-1162>
- [21] Taihua Shao, Honghui Chen, and Wanyu Chen. 2018. Query Auto-Completion Based on Word2vec Semantic Similarity. In *Journal of Physics: Conference Series*, Vol. 1004. IOP Publishing, 012018.
- [22] Milad Shokouhi. 2013. Learning to Personalize Query Auto-completion. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '13)*. ACM, New York, NY, USA, 103–112. <https://doi.org/10.1145/2484028.2484076>
- [23] Amit Singhal et al. 2001. Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.* 24, 4 (2001), 35–43.
- [24] Yang Song, Dengyong Zhou, and Li-wei He. 2011. Post-ranking Query Suggestion by Diversifying Search Results. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 815–824. <https://doi.org/10.1145/2009916.2010025>
- [25] F. Su, M. Somaiya, S. Mishra, and R. Mukherjee. 2015. EXOS: Expansion on session for enhancing effectiveness of query auto-completion. In *2015 IEEE International Conference on Big Data (Big Data)*. 1154–1163. <https://doi.org/10.1109/BigData.2015.7363869>
- [26] Aston Zhang, Amit Goyal, Weize Kong, Hongbo Deng, Anlei Dong, Yi Chang, Carl A. Gunter, and Jiawei Han. 2015. adaQAC: Adaptive Query Auto-Completion via Implicit Negative Feedback. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*. ACM, New York, NY, USA, 143–152. <https://doi.org/10.1145/2766462.2767697>