Using Image Captions and Multitask Learning for Recommending Query Reformulations

Gaurav Verma^{1(✉)}, Vishwa Vinay^{1(✉)}, Sahil Bansal², Shashank Oberoi³,
Makkunda Sharma⁴, and Prakhar Gupta⁵

¹ Adobe Research, Bangalore, India
{gaverma, vinay}@adobe.com

² IBM Research, Bangalore, India

³ Adobe Inc., Noida, India

⁴ Indian Institute of Technology Delhi, Delhi, India

⁵ Carnegie Mellon University, Pittsburgh, USA

Abstract. Interactive search sessions often contain multiple queries, where the user submits a reformulated version of the previous query in response to the original results. We aim to enhance the query recommendation experience for a commercial image search engine. Our proposed methodology incorporates current state-of-the-art practices from relevant literature – the use of generation-based sequence-to-sequence models that capture session context, and a multitask architecture that simultaneously optimizes the ranking of results. We extend this setup by driving the learning of such a model with captions of clicked images as the target, instead of using the subsequent query within the session. Since these captions tend to be linguistically richer, the reformulation mechanism can be seen as assistance to construct more *descriptive* queries. In addition, via the use of a pairwise loss for the secondary ranking task, we show that the generated reformulations are more *diverse*.

Keywords: Query reformulations · Seq-to-seq translation · Captions

1 Introduction

A successful search relies on the engine accurately interpreting the intent behind a user’s query and returning likely relevant results ranked high. There has been much progress allowing search engines to respond effectively even to short keyword queries on rare intents [5, 9, 25]. Despite this, recommendation of queries is an integral part of all search experiences – either in the form of *query autocomplete* (queries that match the prefix the user has currently typed into the search box) or *query suggestions* (reformulation options once an initial query has been provided). In this work, we focus on the query suggestion task.

S. Bansal, S. Oberoi, and M. Sharma contributed equally to this work.

© Springer Nature Switzerland AG 2020

J. M. Jose et al. (Eds.): ECIR 2020, LNCS 12035, pp. 681–696, 2020.

https://doi.org/10.1007/978-3-030-45439-5_45

Fig. 1. The basic idea behind our work. We generate query reformulations using (a) subsequent queries within sessions, and (b) the captions of clicked images, as supervision signals. In both the cases, the task of generating reformulations is done while jointly optimizing the ranking of results.

Original algorithms for this scenario relied on extracting co-occurrence patterns between query pairs, and their constituent terms, within historical logs [3, 12, 16, 18]. Such methods often work well for frequent queries. Recent work utilizing generative approaches common in natural language processing (NLP) scenarios offer generalization in terms of being able to provide suggestions even for rare queries [10, 21]. More specifically, the work by Sordoni et al. [26] focuses on generating query suggestions that are aware of the context of the user’s current session. The current paper is most similar to this work in terms of motivation and the core technical component.

The experiments described here are based on data from a commercial stock image search engine. In this setting, the items in the index are professionally taken high quality images to be used in commercial publishing material. The users of such a system exhibit similar properties to what might be expected on general purpose search engines - i.e., the use of relatively short queries often with multiple reformulations within a session. The logged data therefore contains not only the sequence of within-session queries, but also impression logs listing what images were shown in response to a query and which amongst those were clicked.

The availability of usage data, which provides implicit relevance signals, allows the building of a query reformulation model that includes aspects that have been shown to be useful in related literature: session context capturing information from previous queries in the session, as well as properties of relevant results via a multitask component. Building on state-of-the-art models in this manner, we specialize the solution to our setting by utilizing a novel supervision signal for the reformulation model in the form of linguistically rich captions available for the clicked results (in our case, images) across sessions (Fig. 1).

2 Related Work

A user of a search system provides an input query, typically a short list of keywords, into the search box and expects content relevant to their need ranked high in the result list. There are many reasons why a single iteration of search may not be successful – mis-specified queries (including spelling errors), imperfect ranking, ambiguous intent, and many more. As a result, it is useful to think of a search session as a series of interactions – where the user enters a query, examines and potentially interacts with the returned results, and constructs a refined query that is expected to more accurately represent their intent. Search engines therefore mine historical behavior of users on this query and similar ones in an attempt to optimize the entire search session [24].

Being able to effectively extract these signals from historical logs starts with understanding and interpreting user behavior appropriately. For example, Huang et al. [17] pointed out that successful reformulations, especially those involving changes to words and their order, can be identified as those that retrieve new items which are presented higher in the subsequent results. An automatic reformulation experience involves implementing lessons from such analyses. The first of these is the use of previous queries within the current search sessions to inform the subsequent suggestions – i.e., modeling the *session context*. Earlier papers (e.g. [7]) explicitly captured co-occurrence within sessions which, while being an intuitive and simple strategy, had the disadvantage of not being able to account for rarer queries. Newer efforts (e.g. [21]) therefore utilize distributed representations of terms and queries to help generalize to unseen queries.

Such efforts are part of a wider expansion of techniques originally common within NLP domains to Information Retrieval (IR) scenarios. Conceptually, a generation-based model for query reformulation is obtained by mapping a query to the subsequent one in the same session. Such a model incorporates two signals known to be useful from traditional IR: (1) sequence of terms within a query & (2) sequence of queries within a session. Recent papers have investigated models anchored in the original generic NLP settings but customized to the characteristics of search queries. For example, Dehghani et al. [11] suggest a ‘copy’ mechanism within the sequence-to-sequence (seq-to-seq) models [27] to allow for terms to be carried over across queries in the session. In the current paper, we consider the work of Sordoni et al. [26] as a reference for the core seq-to-seq model. The model, referred to here as *Hierarchical Recurrent Encoder Decoder (HRED)*, is a standard encoder-decoder setup, where word embeddings are aggregated into a query representation, a sequence of which in turn leads to a session representation. A decoder for the hierarchically organized query and session encoders is trained to predict the sequence of query words that compose the subsequent query in the session. Along with being a strong baseline, it serves to illustrate the core components of our work: (a) use of a novel supervision signal in the form of captions of clicked results, and (b) jointly optimizing ranking along with query reformulation. These extensions could similarly be done with other seq-to-seq models used for query suggestion.

Our motivation for using captions of clicked images as supervision signal stems from the fact that captions are often succinct summaries of the content of the actual images as the creators are incentivized to have their images found. In particular, captions indicate which objects are present in the image, their corresponding attributes, as well as relationships with other objects in the same image – for example, “A beautiful girl **wearing** a yellow shirt **standing near** a red car”. These properties make the captions a good target.

Multitask learning [8] has been shown to have success in scenarios where related tasks benefit from common signals. A recent paper [1] shows benefits of such a pairing in a search setting. Specifically, Ahmad et al. show that coupling with a classifier distinguishing clicked results from those skipped helps improve a query suggestion model. We extend this work by utilizing a pairwise loss function commonly used in learning-to-rank [6]. We show that not only does this provide the expected increase in the effectiveness of the ranker component, but also increases the diversity of suggested reformulations. Such diversity has been shown to be important for the query suggestion user experience [20].

We begin by providing details of the mathematical notation in the next section, before describing our models in detail. The subsequent experimental section provides empirical evidence of the benefits that our design choices bring.

3 Notation and Model Architectures

3.1 Notation

We define a session as a sequence of queries, $\mathcal{S} = \{q_1, \dots, q_n\}$. Each query q_i in session $\mathcal{S}$ has a set of displayed images associated with it, $\mathcal{I}_i = \{I_i^1, \dots, I_i^m\}$. A subset of images in $\mathcal{I}_i$ are clicked, we refer to the top-ranked clicked image as I_i^{clicked} . All the images in the set $\mathcal{I}_i$ have a caption describing them, the entire set of which is represented as $\mathcal{C}_i = \{C_i^1, \dots, C_i^m\}$. It follows that every I_i^{clicked} will also have an associated caption with it, given as C_i^{clicked} . Given this, for every successful query q_i in session $\mathcal{S}$, we will have an associated clicked image I_i^{clicked} and a corresponding caption C_i^{clicked} . We consider the size of impression m (number of images) to be fixed for all q_i .

Our models treat each query q_i in any given session, as a sequence of words, $q_i = \{w_1, \dots, w_{l_q}\}$. Captions are represented similarly - as sequences of words, $C_i^j = \{w_1, \dots, w_{l_c}\}$. We use LSTMs [15] to model the sequences, owing to their demonstrated capabilities in modeling various natural language tasks, ranging from machine translation [27] to query suggestion [11].

The input to our models is a query q_i in the session $\mathcal{S}$, and the desired output is a target reformulation q_{reform} . This target reformulation q_{reform} can either be (i) the subsequent query q_{i+1} in the same session $\mathcal{S}$, or (ii) the caption C_i^{clicked} corresponding to the clicked image I_i^{clicked} . Note that obtaining contextual query suggestions via a *translation model* that has learnt a mapping between successive queries within a session (i.e., (i)) has been previously proposed in our reference baseline papers [1, 26]. In the current paper, we utilize a

Fig. 2. An illustration of the (a) query encoder, (b) session encoder, and (c) query decoder

linguistically richer supervision signal, in the form of captions of clicked images (i.e., (ii)), and analyze the behavior of the different models across three high level axes - relevance, descriptiveness and diversity of generated reformulations.

3.2 Model Architectures

In this paper, we evaluate two base models – *HRED* and *HRED* with **Captions** (*HREDCap*), and to study the effect of multitask learning, we add a ranker component to each of these models; giving us two more multitask variants – *HRED + Ranker* and *HREDCap + Ranker*. The underlying architecture of *HRED* and *HREDCap* (and the corresponding variants) is essentially the same, but *HRED* has been trained by using q_{i+1} as target and *HREDCap* has been trained using C_i^{clicked} as target. *HRED* comprises of a query encoder, a session encoder, and a query decoder; all of which are described below.

Query Encoder: The query encoder generates a query level encoding $\mathbf{V}_{q_i}$ for every $q_i \in \mathcal{S}$. This is done by first representing the query q_i using vector embeddings of corresponding words $\{\mathbf{w}_1, \dots, \mathbf{w}_{l_q}\}$, and then sequentially feeding them into a bidirectional LSTM (BiLSTM) [14]. As shown in Fig. 2(a), the query encoder takes each of these word representations as input to the BiLSTM at every encoding step and updates the hidden states based on the forward and backward pass over the input query. The forward and backward hidden states are concatenated, and after applying attention [2] over the concatenated hidden states, we obtain a fixed size vector representation $\mathbf{V}_{q_i}$ for the query $q_i \in \mathcal{S}$.

Session Encoder: The encoded representation $\mathbf{V}_{q_i}$ of query $q_i \in \mathcal{S}$ is used by the session encoder, along with encoded representations $\{\mathbf{V}_{q_1}, \dots, \mathbf{V}_{q_{i-1}}\}$ of previous queries within the same session, to capture the context of the ongoing session thus far. The session encoder, which is modeled by a unidirectional LSTM [15], updates the session context $\mathbf{V}_S^{q_i}$ after each new $\mathbf{V}_{q_i}$ is presented to it. Figure 2(b) illustrates one such update where the session encoding is updated

Fig. 3. The proposed architecture of our multitask model: *HRED + Ranker* (left). For the sake of brevity, we have shown the ranker component separately (right). For *HREDCap + Ranker*, the supervision signals are obtained from captions of clicked images and not subsequent queries.

from $V_S^{q_{i-1}}$ to $V_S^{q_i}$ after V_{q_i} is provided as input to the session encoder by the query encoder. Since it is unreasonable to assume access to future queries in the session while generating a reformulation for the current query, we use a unidirectional LSTM to model the forward sequence of queries within a session. Accordingly, the session encoder updates its hidden state based on the forward pass over the query sequence. As shown in Fig. 2(b), max-pooling is applied over each dimension of the hidden state to obtain the session encoding $V_S^{q_i}$.

Query Decoder: The generated session encoding $V_S^{q_i}$ is used as input by a query decoder to generate a reformulation $\hat{q}_{\text{reform}} = \{\hat{w}_1, \dots, \hat{w}_{l_r}\}$ for the query $q_i \in \mathcal{S}$. As shown in Fig. 2(c), the reformulation is generated word by word using a single layer unidirectional LSTM. With each unfolding of the decoder LSTM at step $t \in \{1, \dots, l_r\}$, a new word $\hat{w}_t$ is generated as per the following probability:¹

$$\begin{aligned} \hat{w}_t &= \arg \max_{w^i \in \mathcal{V}} P(\hat{w}_t = w^i \mid \hat{w}_{1:t-1}, V_S^{q_i}) \\ P(\hat{w}_t = w^i \mid \hat{w}_{1:t-1}, V_S^{q_i}) &= g(\phi(h_d^t)) \end{aligned} \quad (1)$$

Here, h_d^t is the hidden state of the decoder at decoding step t , $\hat{w}_{1:t-1}$ denotes the previous words generated by the decoder, and $\phi(h_d^t)$ is a non-linear operation over h_d^t . The softmax function $g(\cdot)$ provides a probability distribution over the entire vocabulary $\mathcal{V}$. w^i is used to denote the i -th word in $\mathcal{V}$. The joint probability of generating a reformulation $\hat{q}_{\text{reform}} = \{\hat{w}_1, \dots, \hat{w}_{l_r}\}$ can be decomposed into

¹ For $t = 1$, $P(\hat{w}_t = w^i \mid \hat{w}_{1:t-1}, V_S^{q_i})$ reduces to $P(\hat{w}_t = w^i \mid V_S^{q_i})$. However, for the sake of readability, this special consideration for $t = 1$ has been skipped for the following equations.

the ordered conditionals as $P(\hat{q}_{\text{reform}} \mid q_i) = \prod_{t=1}^{l_r} P(\hat{w}_t \mid \hat{w}_{1:t-1}, \mathbf{V}_{\mathcal{S}}^{q_i})$. During training, the decoder compares each word $\hat{w}_t$ in the generated reformulation $\hat{q}_{\text{reform}}$ with the corresponding word w_t in the target reformulation q_{reform} , and aims to minimize the negative log-likelihood. For a given reformulation by the decoder, the loss is

$$\mathcal{L}_{\text{reform}} = - \sum_{t=1}^{l_r} \log P(\hat{w}_t = w_t \mid \hat{w}_{1:t-1}, \mathbf{V}_{\mathcal{S}}^{q_i}) + \mathcal{L}_{\text{reg}} \quad (2)$$

Here, $\mathcal{L}_{\text{reg}} = -\lambda \sum_{w^i \in \mathcal{V}} P(w^i \mid \hat{w}_{1:t-1}, \mathbf{V}_{\mathcal{S}}^{q_i}) \cdot \log P(w^i \mid \hat{w}_{1:t-1}, \mathbf{V}_{\mathcal{S}}^{q_i})$ is a regularization term added to prevent the predicted probability distribution over the words in the vocabulary from being highly skewed. λ is a regularization hyper-parameter. The training loss is the sum of $\mathcal{L}_{\text{reform}}$ over all query reformulations generated by the decoder during training.

To summarize, the model encodes the queries, generates session context encodings, and generates the reformulated query using the decoder while updating the model parameters using the gradients of $\mathcal{L}_{\text{reform}}$.

Ranker Component: This additional component is responsible for ranking the m retrieved results for $q_i \in \mathcal{S}$. As shown in Fig. 3 (right), the ranker takes as input the concatenation of query and session encoding $[\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}]$, for every $q_i \in \mathcal{S}$. The concatenated vector representation $[\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}]$ is used to compute the similarity between the query q_i and its candidate results. The concatenation of these encodings is done to ensure that both current query information (as captured in $\mathbf{V}_{q_i}$) and ongoing session context (as captured in $\mathbf{V}_{\mathcal{S}}^{q_i}$) is used by the ranker. To obtain a representation of the images, we use their corresponding captions. Formally, for every query $q_i \in \mathcal{S}$ each image $I_i^j \in \mathcal{I}_i$ is represented using $\mathbf{C}_i^j$. The average of the vector embeddings of words $\{w_1, \dots, w_{l_c}\}$ in $\mathbf{C}_i^j$ is computed for the image I_i^j . The cosine similarities between $[\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}]$ and the image representations $\mathbf{C}_i^j \in \mathcal{C}_i$ are used to rank order the retrieved results. The j -th element of the similarity vector $\mathbf{S}_i$ represents the similarity between $[\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}]$ and $\mathbf{C}_i^j$.

$$S_i^j = \text{sim}([\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}], \mathbf{C}_i^j) \quad (3)$$

During training, the ranker tries to learn model parameters based on one of the following two objectives:

(i) **Cross Entropy Loss:** As described in [1], we utilize the ‘clicked’ versus ‘not-clicked’ boolean event to train a classifier, where the ranker scores the m retrieved results based on the probability of being clicked by the user. In the following equation, $\mathbf{R}_i$ for query q_i is an m -dimensional vector, where each value in the vector indicates whether the corresponding image was clicked or not. I.e., $R_i^j = 0$ if I_i^j was not clicked, and $R_i^j = 1$ if I_i^j was clicked. A sigmoid of the scores from Eq. 3 is taken as the probability of click. Using the $\mathbf{R}_i$ as labels, the ranker can now be trained using a standard cross entropy loss function:

$$\mathcal{L}_{\text{rank}} = \text{BCE}(\sigma(\mathbf{S}_i), \mathbf{R}_i) \quad (4)$$

(ii) **Pairwise Ranking Loss:** As described in [6], the original boolean labels in $\mathbf{R}_i$ can be used to construct an alternate event space where labels $M_{jk} = 1$ when the image at rank j was clicked while the one at k was not. Pairwise ranking loss allows to better model the preferences of certain results over the others.

$$\mathcal{L}_{\text{rank}} = -\frac{1}{m^2} \sum_{j=1}^m \sum_{\substack{k=1 \\ k \neq j}}^m M_{jk} * \log \hat{M}_{jk} + (1 - M_{jk}) * \log(1 - \hat{M}_{jk}) \quad (5)$$

$$\text{where, } \hat{M}_{jk} = P(S_i^j > S_i^k \mid [\mathbf{V}_{q_i} \oplus \mathbf{V}_{\mathcal{S}}^{q_i}]) = \sigma(S_i^j - S_i^k)$$

Since *HRED + Ranker* and *HREDCap + Ranker* are multitask models, their training objective is a weighted combination of $\mathcal{L}_{\text{reform}}$ and $\mathcal{L}_{\text{rank}}$.

$$\mathcal{L}_{\text{multitask}} = \alpha \cdot \mathcal{L}_{\text{reform}} + (1 - \alpha) \cdot \mathcal{L}_{\text{rank}} \quad (6)$$

Here, α is a hyperparameter used for controlling the relative contribution of the two losses. As mentioned earlier, either the regular binary cross-entropy loss or the pairwise-ranking loss can be used for $\mathcal{L}_{\text{rank}}$. We experiment using both and report our results on the effect of using one over the other. The models that are trained using cross entropy loss are appended with (*CE*), and the models that are trained using pairwise ranking objective are denoted as (*RO*).

It is worth noting that since for a given query q_i there can be more than one clicked images, our ranker component allows $\mathbf{R}_i$ to take the value 1 at more than a single place. However, while training the reformulation model, we only consider the caption of the *highest ranked* clicked image.

4 Experiments

Dataset: We use logged impression data from Adobe Stock². The query logs contain information about the queries that were issued by users, and the images that were presented in response to those queries. Additionally, they contain information about which of the displayed images were clicked by the user. We consider the top-10 ranked results, i.e., the number of results to be considered for each query is $m = 10$. The queries are segmented into sessions (multiple queries by the same user within a 30 min time window), while maintaining the sequence in which they were executed by a user. We retain both multi-query sessions as well as single-query sessions, leading to a dataset comprising 1,301,888 sessions, 2,122,079 queries, and 10,185,979 unique images. We note that $\sim 24.8\%$ of the sessions are single-query sessions, while rest all are multi-query sessions; each of which, on average, comprise of 2.19 queries. Additionally, we remove all non-alphanumeric characters from the user-entered queries, while keeping spaces, and convert all characters to lowercase.

To obtain the train, test and validation set, we first shuffle the sessions and split them in a 80 : 10 : 10 ratio, respectively. While it is possible for a query to

² <https://stock.adobe.com/>.

be issued by different users in distinct sessions, a given search session occurs in only one of these sets. These sets are kept the same for all experiments, to ensure consistency while comparing the performance of trained models. The validation set is used for hyperparameter tuning.

Experimental Setup: We construct a global vocabulary $\mathcal{V}$ of size 37,648 comprising of words that make up the queries and captions for images. Each word in the vocabulary is represented using a 300-dimensional vector $\mathbf{w}_i$. Each $\mathbf{w}_i \in \mathcal{V}$ is initialized using pre-trained GloVe vectors [23]. Words in our vocabulary $\mathcal{V}$ that do not have a pre-trained embedding available in GloVe (1,941 in number), are initialized using samples from a standard normal distribution. Since the average number of words in a query, average number of words in a caption, and average number of queries within a session are 2.31, 5.22, and 1.63, we limit their maximum sizes to 5, 10, and 5, respectively. For queries and captions that contain less than 5 and 10 words respectively, we pad them using ‘< p >’ tokens. The number of generated words in $\hat{q}_{\text{reform}}$ was limited to 10, i.e., $l_r = 10$.

During training, we use Adam optimizer [19] with a learning rate initialized to 10^{-3} . Across all the models, the regularization coefficient λ is set to be 0.1. For multitask models, the loss trade-off hyperparameter α is set to 0.45. The sizes of the hidden states of query level encoder $\vec{h}_q$ and $\overleftarrow{h}_q$ are set to 256, and that of session level encoder h_S is set to 512. The size of the decoder’s hidden state is kept to be 256. We train all the models for a maximum of 30 epochs, using batches of size 512, with early stopping based on the loss over the validation set. The best trained models are quantitatively and qualitatively evaluated and we discuss the results in the upcoming section.

At test time, we use a beam search-based decoding approach to generate multiple reformulations [2]. For our experiments, we set the beam width $K = 3$. The choice of K was governed by observations that will be discussed later, while analyzing the diversity and relevance of generated reformulations. These three reformulations are rank ordered using their generation probability.

We experiment with a range of hyperparameters and find that the evaluation results are stable with respect to our hyperparameter choices. However, our motivation is less about training the most accurate models, as we wish to measure the effect of the supervision signal and training objective when used alongside the baseline models. While presenting the results in Tables 1 and 2, we report the average of values over 10 different runs, as well the standard deviations.

5 Evaluation and Results

In this section, we evaluate the performance of the aforementioned models using multiple metrics for each of the two tasks: query reformulation and ranking. The metrics used here are largely inspired from [11], and we discuss these below briefly. Towards the end of the section we also provide some qualitative results.

Table 1. Performance of models based on reformulation and ranking metrics

Model	Query reformulation			Ranking
	BLEU (%) (↑)	sim _{emb} (%) (↑)	Diversity <i>Top K</i> = 3 (↑)	MRR <i>Baseline</i> : 0.31 (↑)
HRED	6.92 ± 0.06	40.7 ± 1.3	0.37 ± 0.01	-
HRED + Ranker (CE)	7.63 ± 0.07	43.5 ± 1.2	0.42 ± 0.02	0.35 ± 0.02
HRED + Ranker (RO)	7.51 ± 0.07	40.8 ± 1.4	0.43 ± 0.02	0.39 ± 0.01
HREDCap	7.13 ± 0.09	37.8 ± 1.4	0.39 ± 0.04	-
HREDCap + Ranker (CE)	7.95 ± 0.11	39.4 ± 1.2	0.44 ± 0.06	0.38 ± 0.02
HREDCap + Ranker (RO)	7.68 ± 0.10	37.6 ± 1.4	0.45 ± 0.05	0.41 ± 0.02

5.1 Evaluation Metrics

Evaluation for query reformulation involves comparing the generated reformulation $\hat{q}_{\text{reform}}$ with the target reformulation q_{reform} . For all the models, irrespective of whether they utilize the next query within the session q_{i+1} as the target reformulation, or the caption C_i^{clicked} corresponding to the clicked image, the ground truth reformulation q_{reform} is always taken to be q_{i+1} ³. This consistency has been maintained across all models to ensure that their performance is comparable, no matter what signal was used to train the reformulation model. The metrics used here cover three aspects: ‘Relevance’ (BLEU & sim_{emb}), ‘Ranking’ (MRR) and ‘Diversity’ (analyzed later).

BLEU Score: This metric [22], commonly used in machine translation scenarios, quantifies the similarity between a predicted sequence of words and the target sequence of words using n-gram precision. A higher BLEU score corresponds to a higher similarity between the predicted and target reformulations.

Embedding Based Query Similarity: This metric takes semantic similarity of words into account, instead of their exact overlap. A phrase-level embedding is calculated using vector extrema [13], for which pretrained GLoVe embeddings were used. The cosine similarity between the phrase-level vectors for the two queries is given by sim_{emb}. A higher value of sim_{emb} is taken to signify a greater semantic similarity between the prediction and the ground truth. Unlike BLEU, we expect sim_{emb} to provide a notion of similarity of the generated query to the target that allows for replacement words that are similar to the observed ones.

Mean Reciprocal Rank (MRR): The ranker’s effectiveness is evaluated using MRR [28], which is given as the reciprocal rank of the first relevant (i.e., clicked) result averaged over all queries, across all sessions. A higher value of MRR will signify a better ranker in the proposed multitask models. To have a standard point of reference to compare against, we computed the observed MRR for the

³ For sessions with less than 5 queries in a session, if q_i is the last query of the session, the model is trained to predict the ‘end of session’ token as the first token of q_{i+1} . The subsequent predicted tokens are encouraged to be the padding token ‘< p >’.

queries in the test set and found it to be 0.31. This means that on average, for queries in our test set, the first image clicked by the users was at rank ~ 3.1 .

5.2 Main Results

Having discussed the metrics, we will now present the performance of our models on the two tasks under consideration, namely query reformulation and ranking. Table 1 provides these results as well as the effect of different ranking losses – denoted by (*RO*) and (*CE*) respectively.

Evaluation Based on Reformulation: For the purpose of this evaluation, we fix the beam width $K = 3$ and report the average of maximum values among all the candidate reformulations, across all queries in our test set.

While comparing *HRED* and *HRED + Ranker* (both *CE* and *RO*), we observe that the multitask version performs better across *all* metrics. A similar trend can be observed when comparing *HREDCap* with its multitask variants. For all the three metrics for query reformulations, the best performing model is a multitask model – this validates the observations from [1] in our context.

When comparing the two core reformulation models – *HRED* & *HREDCap*, we find that the richer captions data that *HREDCap* sees is aiding the model – while *HRED* scores better sim_{emb} , *HREDCap* wins out on BLEU & Diversity. The drop in sim_{emb} values can be explained by noting that on average captions contain more words than queries (5.22 in comparison to 2.31), and hence similarity-based measures, due to additional words in the captions, will not be as high as overlap-based measures (i.e., BLEU). **Evaluation based on Ranking:** To evaluate the performance of the ranker component in our proposed multitask models, we use MRR. We use the observed MRR of clicked results in the test set (0.31) as the baseline. We also analyze the effect of using the pairwise objective as opposed to the binary cross entropy loss.

Looking at the results presented in Table 1, three trends emerge. Firstly, all the proposed multitask models perform better than the baseline. The best performing model, i.e., *HREDCap + Ranker* with pairwise loss (*RO*), outperforms the baseline by about 32%. Secondly, we observe that using pairwise loss leads to an increase in MRR, for both of the cases under consideration, with only marginal drop in reformulation metrics – we revisit this observation in the next section. Lastly, the multitask models that use captions perform better than multitask models that use subsequent queries.

5.3 Analysis

In this section, we concentrate on the following two aspects of the generated query reformulations: (a) diversity, and (b) descriptiveness.

Diverse Query Reformulations due to Multitasking: The importance of suggesting diverse queries to enhance user search experience is well established within the IR community. The mechanism to obtain a diverse set of reformulation alternatives is via the use of beam search based decoding. In scenarios where

a set of top- K candidates are required, we take inspiration from Ma et al. [20] to evaluate the predictions of our models for their diversity. For a beam width of K , a reformulation model will generate $\mathcal{R}_{gen} = \{r_1, r_2, \dots, r_K\}$ candidate reformulations for a given original query. We quantify the diversity in the candidate reformulations by comparing each candidate reformulation r_i with other reformulations $r_j \in \mathcal{R}_{gen} : i \neq j$. The diversity of a set of K queries is evaluated as

$$D(\mathcal{R}_{gen}) = 1 - \frac{1}{K(K-1)} * \left(\sum_{r_i \in \mathcal{R}_{gen}} \sum_{r_j \in \mathcal{R}_{gen} : j \neq i} sim_{emb}(r_i, r_j) \right)$$

In Table 1, it can be observed that multitask models generate more diverse reformulations than models trained just for the task of query reformulation. This is particularly evident when comparing the effect of the ranking loss.

From Fig. 4, it can be noted that as more candidate reformulations are taken into consideration, i.e., as the beam width K is increased, the average relevance of the reformulations decreases across all the models. However, the diverseness of $\mathcal{R}_{gen}$ flattens after $K = 3$. This was the reason for setting the beam width to 3 while presenting results in Table 1.

Fig. 4. The trade-off between relevance (as quantified by sim_{emb}) and diversity. As K is increased, the relevance of generated predictions drops across all models.

Descriptive Reformulations using Captions: The motivation for generating more descriptive reformulations is of central importance to our idea of using image captions. To this end, we analyze the generated reformulations to assess if this is indeed the case. We start by noting (see Table 2) that captions corresponding to clicked images for queries in our test set contain, on average, more words than the queries. Following this, we analyze the generated reformulations by two of our multitask models – (i) *HRED + Ranker (RO)*, which guides the process of query reformulation using subsequent queries within a session, and (ii) *HREDCap + Ranker (RO)*, which guides the process of query reformulation using captions corresponding to clicked images. For this entire analysis, we removed stop words [4] from all the queries and captions under consideration.

As can be noted from Table 2, reformulations using captions tend to contain more words than reformulations without them. However, number of words in

a query is only a facile proxy for its descriptiveness. Acknowledging this, we perform a secondary aggregate analysis on the number of novel words inserted into the reformulation and number of words dropped from the original query. We identify novel words as words that were not present in the original query q_i but have been generated in the reformulation $\hat{q}_{\text{reform}}$, and dropped words as the words that were present in the original query but are absent from the generated reformulation. Table 2 indicates that, on average, the model trained using captions tends to insert more novel words while reformulating the query, and at the same time drops fewer words from the query. Interestingly, models trained using subsequent queries inserts almost as many words into the reformulation as it drops from the original query.

To analyze this further, we compute the average similarity between the novel words that were inserted and the words that were dropped, by averaging the GloVe vector based similarity between words, across all queries in our test set. For *HRED + Ranker (RO)* this average similarity is **0.64**, while for *HREDCap + Ranker (RO)* it is **0.41**. A higher similarity value for the former suggests that the model largely *substitutes* the existing words with words having similar semantic meaning. Using captions, on the other hand, is more likely to generate novel words which bring in additional meaning.

Table 2. Analyzing the effect of using captions on length of generated query reformulations, along with influence on generating novel words while dropping the existing ones.

Avg. # of words in queries	2.31 $\pm$ 0.92 word(s)	
Avg. # of words in captions	5.22 $\pm$ 2.37 word(s)	
Models $\rightarrow$	HRED + Ranker (RO)	HREDCap + Ranker (RO)
Avg. # generated words	2.18 $\pm$ 0.61 word(s)	4.91 $\pm$ 1.16 word(s)
Avg. # novel words	1.04 $\pm$ 0.13 word(s)	2.56 $\pm$ 0.47 word(s)
Avg. # dropped words	1.14 $\pm$ 0.15 word(s)	0.89 $\pm$ 0.17 word(s)
Avg. similarity b/w insertions and drops	0.64 $\pm$ 0.03	0.41 $\pm$ 0.04

5.4 Qualitative Results

In Table 3, we present a few examples depicting the descriptive nature of generated reformulations. The generated reformulations by *HRED + Ranker* are compared against those by *HREDCap + Ranker*. We only present the top ranked reformulation among top- K reformulations. We note that using captions as target generates reformulations that are more descriptive and the process of generation results in more insertions of novel words, in comparison to using subsequent queries as targets. These qualitative observations, along with quantitative observations discussed earlier, reinforce the efficacy of using captions of clicked images for the task of query reformulation.

Table 3. Qualitative results comparing the generated reformulation by *HRED + Ranker* and *HREDCap + Ranker*. The words in **bold** are novel insertions.

	Queries	Clicked caption	<i>HRED + Ranker</i> (RO)	<i>HREDCap + Ranker</i> (RO)
Session ₁	q1 traffic	rush hour traffic	traffic jam	traffic jam during rush hour
	q2 traffic jam	traffic jams in the city, road, rush hour	city traffic jam	traffic during rush hour in city
	q3 traffic jam pollution	blurred silhouettes of cars by steam of exhaust	traffic jam cars	dirt and smoke from cars in traffic jam
Session ₂	q1 sleeping baby	sleeping one year old baby girl	cute sleeping baby	little baby sleeping peacefully
	q2 sleeping baby cute	baby boy in white sunny bedroom	sleeping baby	baby sleeping in bed peacefully
	q3 white bed sleeping baby	carefree little baby sleeping with white soft toy	baby sleeping in bed	little baby sleeping in white bed peacefully
Session ₃	q1 chemistry	three dimensional illustration of molecule model	chemical reaction	molecules and structures in chemistry
	q3 molecule reaction	chemical reaction between molecules	reaction molecules	molecules reacting in chemistry
	q3 molecule collision	frozen moment of two particle collision	collision molecules	molecules colliding chemistry reaction

6 Conclusion

In this paper, we build upon recent advances in sequence-to-sequence models based approaches for recommending queries. The core technical component of our paper is the use of a novel supervision signal for training seq-to-seq models for query reformulation – i.e., captions of clicked images instead of subsequent queries within a session, as well as the use of a pairwise preference based objective for the secondary ranking task. The effect of these are evaluated alongside baseline model architectures for this setting. Our extensive analysis evaluated the model and training method combinations towards being able to generate a set of descriptive, relevant and diverse reformulations.

Although the experiments were done on data from an image search engine, we believe that similar improvements can be observed if content properties from textual documents can be integrated into the seq-to-seq models. Future work will look into the influence of richer representations on the behavior of the ranker, and in turn on the characteristics of the reformulations.

References

1. Ahmad, W.U., Chang, K.W., Wang, H.: Multi-task learning for document ranking and query suggestion. In: International Conference on Learning Representations (2018)
2. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint [arXiv:1409.0473](https://arxiv.org/abs/1409.0473) (2014)

3. Beeferman, D., Berger, A.: Agglomerative clustering of a search engine query log. In: Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 407–416. ACM (2000)
4. Bird, S., Loper, E.: NLTK: the natural language toolkit. In: Proceedings of the ACL 2004 on Interactive Poster and Demonstration Sessions, p. 31. Association for Computational Linguistics (2004)
5. Broder, A.Z., Fontoura, M., Gabrilovich, E., Joshi, A., Josifovski, V., Zhang, T.: Robust classification of rare queries using web knowledge. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 231–238. ACM (2007)
6. Burges, C., et al.: Learning to rank using gradient descent. In: Proceedings of the 22nd International Conference on Machine Learning, ICML 2005, pp. 89–96. ACM (2005)
7. Cao, H., et al.: Context-aware query suggestion by mining click-through and session data. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 875–883. ACM (2008)
8. Caruana, R.: Multitask learning. *Mach. Learn.* **28**(1), 41–75 (1997)
9. Chirita, P.A., Firan, C.S., Nejdl, W.: Personalized query expansion for the web. In: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 7–14. ACM (2007)
10. Cho, K., et al.: Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint [arXiv:1406.1078](https://arxiv.org/abs/1406.1078) (2014)
11. Dehghani, M., Rothe, S., Alfonseca, E., Fleury, P.: Learning to attend, copy, and generate for session-based query suggestion. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, pp. 1747–1756 (2017)
12. Fonseca, B.M., Golgher, P., Pôssas, B., Ribeiro-Neto, B., Ziviani, N.: Concept-based interactive query expansion. In: Proceedings of the 14th ACM International Conference on Information and Knowledge Management, pp. 696–703. ACM (2005)
13. Forgues, G., Pineau, J., Larchevêque, J.M., Tremblay, R.: Bootstrapping dialog systems with word embeddings. In: Nips, Modern Machine Learning and Natural Language Processing Workshop, vol. 2 (2014)
14. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Netw.* **18**(5–6), 602–610 (2005)
15. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
16. Huang, C.K., Chien, L.F., Oyang, Y.J.: Relevant term suggestion in interactive web search based on contextual information in query session logs. *J. Am. Soc. Inf. Sci. Technol.* **54**(7), 638–649 (2003)
17. Huang, J., Efthimiadis, E.N.: Analyzing and evaluating query reformulation strategies in web search logs. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management, pp. 77–86. ACM (2009)
18. Jones, R., Rey, B., Madani, O., Greiner, W.: Generating query substitutions. In: Proceedings of the 15th International Conference on World Wide Web, pp. 387–396. ACM (2006)
19. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. CoRR [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
20. Ma, H., Lyu, M.R., King, I.: Diversifying query suggestion results. In: AAAI, vol. 10 (2010)

21. Mitra, B.: 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*, pp. 3–12. ACM (2015)
22. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pp. 311–318. Association for Computational Linguistics (2002)
23. Pennington, J., Socher, R., Manning, C.: GloVe: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 (2014)
24. Silvestri, F.: Mining query logs: turning search usage data into knowledge. *Found. Trends® Inf. Retr.* **4**(1), 171–174 (2009)
25. Song, Y., He, L.W.: Optimal rare query suggestion with implicit user feedback. In: *Proceedings of the 19th International Conference on World Wide Web*, pp. 901–910. ACM (2010)
26. Sordoni, A., Bengio, Y., Vahabi, H., Lioma, C., Simonsen, J.G., Nie, J.: A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. *CoRR* [arXiv:1507.02221](https://arxiv.org/abs/1507.02221) (2015)
27. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: *Advances in Neural Information Processing Systems*, pp. 3104–3112 (2014)
28. Voorhees, E.M., Dang, H.T.: Overview of the TREC 2003 question answering track. In: *TREC*, vol. 2003, pp. 54–68 (2003)