

Using Clarification Questions to Improve Software Developers' Web Search

Mia Mohammad Imran

Computer Science, Virginia Commonwealth University, Richmond, Virginia, USA

Kostadin Damevski

Computer Science, Virginia Commonwealth University, Richmond, Virginia, USA

Abstract

Context

Recent research indicates that Web queries written by software developers are not very successful in retrieving relevant results, performing measurably worse compared to general purpose Web queries. Most approaches up to this point have addressed this problem with software engineering-specific automated query reformulation techniques, which work without developer involvement but are limited by the content of the original query. In other words, these techniques automatically improve the existing query but can not contribute new, previously unmentioned, concepts.

Objective

In this paper, we propose a technique to guide software developers in manually improving their own Web search queries. We examine a conversational approach that follows unsuccessful queries with a clarification question aimed at eliciting additional query terms, thus providing to the developer a clear dimension along which the query could be improved.

Method

We describe a set of clarification questions derived from a corpus of software developer queries and a neural approach to recommending them for a newly

Email addresses: imranm3@vcu.edu (Mia Mohammad Imran), kdamevski@vcu.edu (Kostadin Damevski)

issued query.

Results

Our evaluation indicates that the recommendation technique is accurate, predicting a valid clarification question 80% of the time and outperforms simple baselines, as well as, state-of-the-art Learning To Rank (LTR) baselines.

Conclusion

As shown in the experimental results, the described approach is capable at recommending appropriate clarification questions to software developers and considered useful by a sample of developers ranging from novices to experienced professionals.

Keywords: software engineering-related search, clarification questions, query refinement

1. Introduction

Instead of printed manuals or books, most of software developers' information needs that support their daily tasks are nowadays found through the Web, using a variety of online sources including Q&A forums, API documentation, tutorials, and chats [1, 2]. Developers regularly search these sites to lookup information, reuse code, and learn new concepts and skills. They usually begin by using Web search engines (e.g., Google, Bing) to locate and access content on popular software engineering sources (e.g., Stack Overflow, W3Schools, Tutorialspoint) [3]. However, numerous developer searches are unsuccessful, consuming valuable developer time and effort. In a recent large-scale study of one million Web search sessions by software developers using Bing, Rao et al. found that software engineering-related queries are less effective than other types of queries [4], resulting in higher rates of query reformulations, fewer clicks, and shorter dwell time compared to non software engineering sessions.

Researchers have studied the Web searches conducted by software developers towards understanding key behaviors [2, 5, 3], observing that the primary problem to failed Web searches are poorly constructed queries, i.e., developers

commonly fail to specify important technical details in their queries (e.g., specific IDE, operating system) [6]. To address the problem of short, incomplete and unspecific queries, researchers have attempted approaches for automated query reformulation that extend a query behind the scenes using synonyms (e.g., using WordNet) or common terms present in highly ranked results [7, 8, 9, 10, 11, 12]. However, these automated approaches are limited by the terms in the original query, as they can only act to extend the meaning already present and cannot contribute additional novel context.

To overcome the shortcomings of automated query reformulation, we can aid searchers in manually extending their Web queries by posing clarification questions that directly address a topic that the original query lacks [13, 14]. Clarification questions pinpoint the exact way a developer should extend their query in order to retrieve improved, more relevant results. For instance, recent work by Zhang et al. demonstrates the promise of prompting developers to manually expand their queries in helping them meet their information needs [15], in their case, for improving question retrieval in Stack Overflow. They proposed conversational query refinement that assists users in recognizing and clarifying the technical details missed in their original queries. Their approach significantly outperforms other approaches to improve question retrieval in Q&A sites and helps users recognize missing technical details in their queries through generating clarification questions based on offline analysis of the top-n similar Stack Overflow questions for a query. However, Zhang et al.’s approach is limited to search within Stack Overflow, which is focused only on retrieving questions which are also tagged with keywords (e.g., java, spring). The goal of this paper is to go beyond Stack Overflow to general Web search, which is where most developers begin their information search [3].

More specifically, we envision a scenario where a software developer, e.g., Alice, writes an initial query, which consists of too few terms and gets a set of inadequate results (see Figure 1). Within the results Web page, our system includes a clarification question that immediately guides Alice in how to expand the query, e.g., *Which operating system are you using?*. The system also provides

a set of common answers to the question, e.g., *Windows*, *Linux*, *Mac OS*, for convenience. Alice selects Mac OS and this adds the term to the query and reissues it to the search engine, retrieving another set of results and perhaps posing another clarification question.

In this paper, we analyze a publicly available dataset of Web queries specific to software engineering to devise a set of clarification questions appropriate for this domain. We then devise a neural network based algorithm for automatic recommendation of clarification questions for a developer query that we train with a large dataset created via data augmentation. We evaluate the clarification question recommendation by comparing to a set of baselines. Finally, we build a prototype tool based on the proposed technique as a browser extension that we use to perform a user study to understand both how our technique performs and is perceived by developers as well as the potential of such approaches for manual query extension.

In summary, the key contributions of this paper are:

- set of clarification questions appropriate for software engineering query expansion;
- algorithm for recommending clarification questions for a Web search query that is able to handle the small number of terms present in most queries (including a data augmentation step to generate appropriate training data);
- quantitative and qualitative evaluation that aids in understanding the value of the technique and the potential of the overall idea.

Significance of contributions.

This paper is one of the first to propose using clarification questions to improve Web searches conducted specifically by software developers with the goal of improving the quality of retrieved results and, through it, the efficiency with which

java ide download Search

Which OS are you using?

Windows Linux Mac OS other

Results:

Eclipse IDE for Java Developers | Eclipse Packages
The essential tools for any Java developer, including a Java IDE, a CVS client, Git client, XML Editor, Mylyn, Maven integration and WindowBuilder.

Download Java IDE - Best Software & Apps - Softonic
Download Java IDE - Best Software & Apps · BlueJ. 5.0.1. 3.8. (2026 votes) · NetBeans IDE. 12.4. 3.6. (795 votes). Free Download · IntelliJ IDEA. 2021.2-build- ...

Download IntelliJ IDEA: The Capable & Ergonomic Java IDE
Download the latest version of IntelliJ IDEA for Windows, macOS or Linux. Developer: JetBrains

Figure 1: An example use of the described system for automatically posing clarification questions to aid developer Web searching.

developers locate relevant information on the Web to aid their daily work. Using clarification questions for general Web search (i.e., so-called conversational search) is recent popular topic in the Information Retrieval community (e.g., see [13, 16]). We contribute a set of clarification questions and a machine learning technique that successfully recommends a relevant question for 4 out of 5 queries. Through a browser extension that implements our technique, we observe that developers value clarification questions as a way to improve their queries and find our technique generally useful.

A replication package containing the dataset and code used for the study is available ¹.

¹<https://anonymous.4open.science/r/Query-Expansion-Questions-5291>

2. Eliciting Clarification Questions

Software developers’ Web queries are often short, consisting of a median of 3 terms [3], and, therefore, it is common for each query to have multiple *facets* (i.e., dimensions of meaning) along which it can be interpreted. Posing clarification questions to developers for an incomplete query aims to identify the specific facet that the developer has in mind, so that it can be specified to the search engine, leading to improved retrieval results.

With the goal of eliciting a set of clarification questions appropriate for software engineering-related searches, we define an approach that, first, identifies the facets of a query and, second, uses the facets to create clarification questions that would unravel developer search intent. In this section, we discuss the process we used for eliciting common clarification questions: 1) selecting a corpus of software developer Web queries; 2) identifying the (usually several) facets for each query; 3) generating facet-based clarification questions; and 4) generalizing the facet-based clarification to a set of common clarification questions.

2.1. Software Development-Related Queries

Logs of actual Web search queries are usually not made publicly available by search engines. Rao et al. note that a significant bottleneck in this area of research is the lack of datasets as search logs can not be made public due to privacy laws (see also Section 5) [4]. In their later work [17], they release a limited dataset of anonymized software development-related search queries mined from Bing logs between September 1, 2019 and August 31, 2020. The dataset contains more than 11,000 real-world search queries related to the C# and Java programming languages, which they group into 7 categories based on their search intent: API, HowTo, Installation, Debug, Learn, Navigational and Miscellaneous. They identify that a query is Java/C# programming language-related when it contains the term ‘*java*’ or ‘*c#*’ respectively, i.e., each query contains either the term ‘*java*’ or ‘*c#*’ and at least one other term.

In this work, we select the Java programming language-related queries released by Rao et al. [17], which consists of 6,596 queries. Some of the queries

in the dataset are similar (e.g., *'java api'*, and *'java apis'*, *'java queue'*, and *'java queues'*, *'java for loop'*, and *'for loop java'*), and some are noisy (e.g., *'java chicken'*, *'java apple'*). Therefore, we manually filter out queries that are not software development-related or not unique. From the remaining set, we randomly sample a set of 200 queries, ensuring representation from each of Rao et al.'s intent categories (30 queries each from API, HowTo, Installation, Debug, and Learn). The Miscellaneous category was significantly larger than the rest so we sampled 50 queries from it. The category of Navigational queries was excluded as it includes the specific resource or Web page that the developer is navigating to, which makes it inappropriate to our purpose.

2.2. Facets Identification and Clarification Question Generation

A facet is a topic connected to a specific query that describes a dimension of meaning of the query. Facets are often described by a set of semantically related terms that together define a distinct information need or a topic [18, 19, 20, 21, 22, 23]. A typical query usually has multiple, separate facets, e.g., the query *'java eclipse download'* has facets such as *'Linux'*, *'Windows'*, *'Mac OS'* that describe the operating system used and *'java 8'*, *'java 11'* that describe the version of the software.

In order to devise facets for the selected 200 queries, we used four human annotators (including one of the authors of this paper) that are all experienced software developers (2 graduate students, 1 senior undergraduate and 1 industry). Each annotator processed 40-60 queries, based on a detailed set of instructions for generating facets and clarification questions. In order to make sure the annotators were clearly aware of their task, we provided annotation instructions as a written document that the annotators were asked to read to completion before beginning. The instructions also included detailed examples of what to do as well as common mistakes or incorrect approaches to the annotation. We instructed the annotators to use the following two ways of generating facets for each query:

1. *Search Engine Top 10 Results*: We asked the annotators to re-enter the



Figure 2: (a) Google suggestions for the query - '*java mail api*', (b) top Google results for the query '*java mail api*'. Using the information in (a) and (b) the annotators identified the facets: 1) '*read*', '*compose*', '*write*'; 2) '*example*', '*documentation*', '*tutorial*' etc.

query into a popular search engine and examine the top 10 retrieved results in some detail. After examining the results, the annotators were to select the facets that describe the primary characteristics of relevant groups of retrieved results. Sometimes search engine results for the queries included very specific terms that cannot be grouped into facets. For example, for the query, '*add java to path*', the path - */home/user/path/to/env* - is too specific to belong in a facet, and, for the query, '*java indexOf method*' - the method signature - *indexOf(String str, int start)* - is too specific. In Figure 2, we provide a screenshot of top results provided by Google and a retrieved top page for the query '*java mail api*'. Here, one possible facet included actions such as '*compose*', '*write*', and '*read*'.

2. *Search Engine Auto-complete*: Search engines such as Google, Bing, etc., offer query completion suggestions (see Figure 2 - left) as the user is typing their query in the search box. These suggested words often describe a key characteristic of a query that can sometimes be appropriate as a facet. We asked the annotators to examine the list of suggested query completions to discover common themes or topics, however, we warned the annotators that sometimes the search engines suggest terms that are too specific to form facets, e.g., ('*example in spring boot*' in part (a) of Figure 2).

The annotators discovered a total of 1,577 facets across the 200 input queries (a median of 8 facets per query).

Next, we asked the annotators to focus on identifying clarification questions for each facet they identified. Specifically, the annotators were instructed to: “Put yourself in the position of someone who has searched using the provided query. You are to come up with one clarification question that would produce a specific facet.” It is possible to map several facets to a single clarification question and to have multiple clarification question for a single facet. For example, for the query ‘java mail api’, a possible set of facets is: 1) ‘example’, ‘documentation’, ‘tutorial’, ‘library’; 2) ‘maven’, ‘gradle’; and 3) ‘create’, ‘send’. Based on these facets, the annotators came up with several clarification questions, such as, ‘Are you looking for related API documentation?’, ‘Do you want to add a specific library in your build tool?’, and ‘What type of mail operation do you want to perform?’. Finally, the annotators generated a total of 741 clarification questions across 200 queries (a median of 4 clarification questions per query).

2.3. Grouping into Common Clarification Questions

The clarification questions generated based on the procedure described above (in Section 2.2) are *facet-based clarification questions*. The exact form of these facet-based clarification questions differs from query to query, from facet to facet, and from annotator to annotator. However, it is possible to group the facet-based clarification questions into a set of common, semantically-related clarification questions. For example, the facet-based clarification questions - ‘Are you asking about an example?’, ‘Do you want an example of this API?’, ‘Do you want an example of an abstract class or method?’ - can be expressed by the common clarification question - ‘What type of document are you interested in (example)?’. To group the facet-based clarification questions we: 1) group all the similar query-facet pairs; and 2) examine the corresponding facet-based clarification questions to formulate a more generic way of expressing them.

Two human annotators (from the initial set of four, including one of the authors) generated common clarification questions for all facet-based clarifica-

tion questions. The annotators worked independently and then held a Zoom meeting to discuss disagreements and agree on appropriate clarification question language. There were a number of facet-based clarification questions where the annotators were not able to devise a general-purpose common clarification question as the information need was very specific to the particular query. For example, for the query *'binary search method java'* the clarification question *'Do you want a solution leveraging recursion?'* is very specific to the particular query and cannot be grouped with other questions in our dataset. In total, we used 592 out of the 741 clarification questions to construct 16 common clarification questions, discarding 149 questions that could not be grouped. The common clarification questions, common answers to these questions, and an example query, facet, and facet-based clarification question are listed in Table 1. Common answers are the most common facets regarding the clarification questions. For simplicity, from this point in the paper, we will refer to the common clarification questions as *clarification questions*.

3. Recommending Clarification Questions

In this section, we describe our system that is able to take a user query as input and produce a ranked list of clarification questions according to their suitability to enhance the query. However, while 200 queries are enough to elicit a reasonable set of clarification questions, they are not adequate to formulate and train a general machine learning model. In addition, to train such a system, we require not just queries, but also their corresponding clarification questions. Therefore, we devise an approach to create more training data via a semi-automatic data augmentation process. Our system (called QueryQuestions-SE) for retrieving and recommending clarification questions consists of two parts: 1) data augmentation to generate training data, and 2) neural network architecture to rank clarification questions.

3.1. Data Augmentation

In a recent paper, Chen et al. proposed a technique - Local Additivity based Data Augmentation (LADA) - aimed at utilizing a limited set of labeled textual

	Clarification Question and Common Answers	Example Query, Facet, and Facet-based Clarification Question
CQ1	What type of document are you interested in? A: documentation, example, tutorial, use case, performance, books	<i>Query: 'java reflection api'; Facet: documentation</i> <i>Q: Do you want to read API documentation?</i>
CQ2	What type of source code artifact are you interested in? A: class definition, API, framework, library, tool, plugin	<i>Query: 'java immutablelist api'; Facet: Guava</i> <i>Q: Are you referring to Google library Guava?</i>
CQ3	Which IDE are you using? A: IntelliJ, Eclipse, PyCharm, Jupyter, Visual Studio, Xcode	<i>Query: 'java is not recognized as an internal command'; Facet: intellij</i> <i>Q: Which IDE are you using?</i>
CQ4	What type of operation do you want to perform? A: read, write, print, parse, override, get, find	<i>Query: 'java imageio'; Facet: read, write, resize, show, display</i> <i>Q: What operation on the image do you want to perform?</i>
CQ5	What is your file type? A: text, json, xml, csv, zip, png, jpeg	<i>Query: 'how to open files with java'; Facet: zip, jar, rar</i> <i>Q: What's the file type?</i>
CQ6	Which system development toolkit are you using? A: Java JDK, iOS SDK, .NET SDK, Android SDK	<i>Query: 'java was started return code 1'; Facet: JDK</i> <i>Q: Do you have a JDK installed?</i>
CQ7	If you are using a specific tool, framework, or library, which one? A: -	<i>Query: 'java.lang.classnotfoundexception: com.mysql.jdbc.driver'; Facet: maven, databricks, gradle, pyspark, spark</i> <i>Q: Are you using maven, databricks, gradle, pyspark, or spark?</i>

CQ8	Which version of software are you using? A: -	<i>Query: 'install java on raspberry pi'; Facet:</i> java 8 <i>Q: Are you asking about any specific version?</i>
CQ9	What type of installation-related operation are you interested in? A: update, configure, install, uninstall, download, version check	<i>Query: 'java mongodb'; Facet:</i> driver <i>Q: Are you trying to download mongodb JDBC driver?</i>
CQ10	Which operating system are you using? A: MacOS, Windows, Linux, Android, iOS	<i>Query: 'java ide download'; Facet:</i> Windows 10, Mac, Linux <i>Q: Which operating system?</i>
CQ11	This seems to be a comparison. Is there a topic you want to compare to? A: -	<i>Query: 'kotlin vs java'; Facet:</i> syntax <i>Q: Do you want to know the syntax difference between Java and Kotlin?</i>
CQ12	Which browser are you using? A: Chrome, Firefox, Opera, Safari, Internet Explorer	<i>Query: 'protection menu allow java'; Facet:</i> chrome, firefox, opera <i>Q: Which browser?</i>
CQ13	Which data type are you interested in/referring to? A: integer, string, float, list, map, set, queue	<i>Query: 'java scanner example'; Facet:</i> string, char <i>Q: What data type are you trying to scan?</i>
CQ14	Are you interested in information related to 32-bit or 64-bit architecture? A: 32-bit, 64-bit	<i>Query: 'latest version of java for windows 10'; Facet:</i> 32-bit <i>Q: What's the processor version?</i>
CQ15	What type of an exception-related operation are you interested in? A: handle, catch, throw, avoid, implement	<i>Query: 'missing return statement error java'; Facet:</i> exception <i>Q: Are you throwing an exception?</i>
CQ16	What type of debugging-related artifact are you interested in? A: fix video, fix tutorial, debug, troubleshoot	<i>Query: 'java package does not exist'; Facet:</i> sub-directory <i>Q: Is the package in right directory?</i>

Table 1: List of clarification questions, common answers to each question, and a sample query, corresponding facet, and facet-based clarification questions that were used to generate the specific clarification question.

data to generate a diverse augmented dataset in a semi-supervised way [24]. LADA can create a large amount of realistic labeled data while improving the generalization of learning. Following LADA’s approach, which targets a somewhat different problem - named entity recognition, we design a procedure to generate enough data so that we can effectively train a ML technique for ranking (i.e., recommending) clarification questions based on a software developer query. To augment the data, we leverage RoBERTa model, which is a deep neural model trained on a large-scale corpus of natural language text, to provide prediction of masked words [25]. The masked words are terms used to augment each query.

More specifically, for each of our 200 real-world queries and clarification questions, we perform the following two augmentation procedures to generate additional queries:

- Using RoBERTa’s masked word prediction, we **add** one or two masked terms to the query. That is, we add the {mask} term in all gaps between terms of the query and use RoBERTa’s suggested replacement for the {mask} term. For example, if the original query is ‘*java mail api*’, for one masked term we create new queries as ‘*java {mask} mail api*’, ‘*java mail {mask} api*’, ‘*java mail api {mask}*’, and ‘*{mask} java mail api*’. Therefore, possible generated queries are ‘*java mail api request*’, ‘*java mail api error*’, ‘*java mail api documentation*’, and so on. We repeat the same process with two masked terms. For each instance of a masked query, we consider the top 100 suggestions generated by RoBERTa.
- Using RoBERTa masked word suggestion, we **replace** one or two terms. For example, if the original query is ‘*java mail api*’, the possible structures of new queries are - ‘*java {mask} api*’, ‘*{mask} mail api*’, ‘*java mail {mask}*’, ‘*{mask} {mask} api*’, etc. From these masked queries, possible generated queries are ‘*python mail api*’, ‘*java mail server*’, and so on. As before, we consider the top 100 suggestions for each combination.

As RoBERTa’s suggestions may not be high quality or they may make spe-

cific clarification questions that are mapped to a query redundant, we manually verify each suggested query. That is, we validate augmented queries to be unique, software engineering-related, free of noisy terms, and that the clarification questions from the original query are still applicable. For example, the annotated clarification questions for the query *‘java mail api’* are CQ1 and CQ2. One of RoBERTa suggestions is *‘secure mail api’*, which we consider to be of sufficient quality with the clarification questions still applicable. Alternatively, the suggested query *‘java mail api tutorial’* is removed because both of the clarification questions are not applicable, i.e., it already specifies the answer to CQ1.

Our final training dataset consists of a total of 5,151 unique queries and 11,762 clarification questions, a number that includes the original set of input queries. Each query is mapped to one or more valid clarification questions.

3.2. Neural Network Architecture

Next, we introduce our neural network model for ranking clarification questions. As shown in Figure 3, the model computes the probability of a clarification question’s validity for a user query using a neural network composed of 3 different parts. Each of the parts centers on each of the 3 inputs in our model: user queries, clarification questions, and the set of common answers. We use a Convolutional Neural Network (CNN) architecture for user queries, a Bidirectional Long Short-Term Memory (LSTM) architecture for clarification questions, and a CNN architecture for common answers. A CNN works well for queries and common answers as the order of terms is less important than the different combinations, which a CNN captures well. Clarification questions are natural language sentences so a LSTM would capture this sequence of words well.

For each part of the neural net, as the first layer, in order to introduce additional semantics, we encode the text with GloVe word embeddings that we pre-trained on the entirety of Stack Overflow (using the Stack Overflow data dump as of June, 2020) with default parameters (vector size = 200; window size

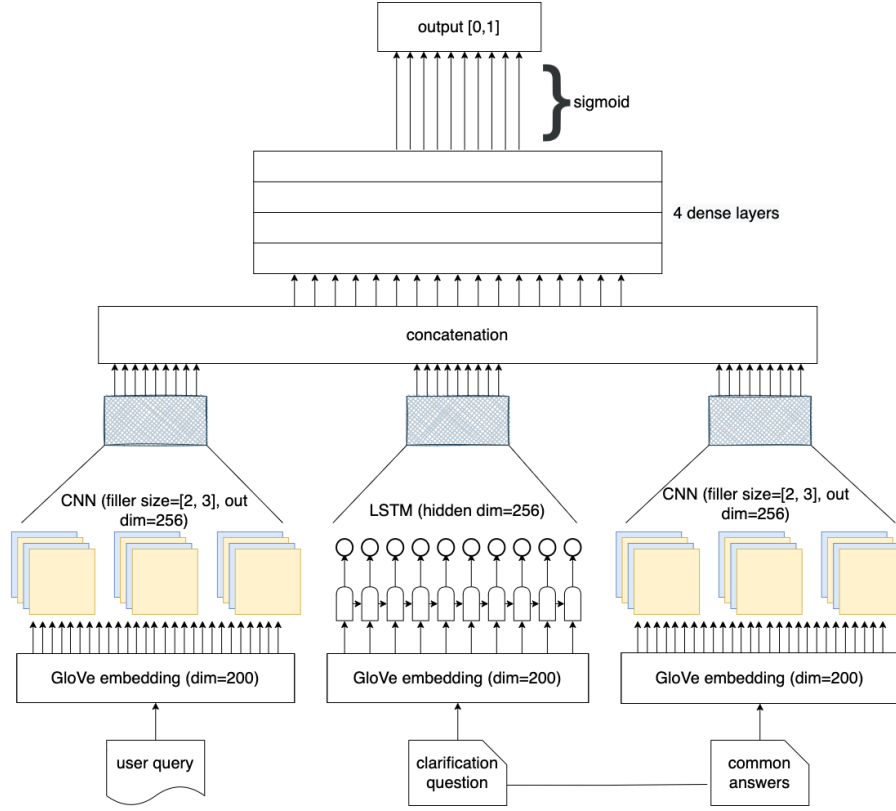


Figure 3: The neural network architecture of the QueryQuestions-SE model aimed at determining valid clarification question for software development queries.

= 15). As the second layer, we train the three neural networks separately and merge their outputs. The output of the two CNN layers are two fully connected layers and the output of LSTM layer is a collection of hidden states. When we concatenate the 3 networks, we directly use the CNN's output, while for the LSTM layer we compute the average of the hidden states. As the third layer, we use a dense neural network. We apply the logistic sigmoid function over this output, producing a single floating-point number which can be interpreted as a probability of how much the current clarification question is applicable for the user query. This probability can be also leveraged as a threshold for when not to recommend any clarification questions, i.e., when the query seems clear

enough as is (see also Section 4.2.1).

In order to train the model, for each query q_i in the set of all queries Q , we create 16 triplets of query, clarification question CQ_j , and common answers CA_j , one for each of the 16 possible clarification questions. We consider a triplet’s label as positive when a clarification question CQ_j is applicable for the query q_i according to the annotation. Note that there is usually more than one clarification question that is appropriate for a single query. We use Binary Cross Entropy loss and the Adam optimizer [26], our implementation leveraging the popular PYTORCH library.

For a newly submitted user query q , we use the model for each clarification questions CQ_j and common answers CA_j separately, collecting the resulting output probabilities. Based on the output probabilities p_j of each triplet (q, CQ_j, CA_j) , we rank the clarification questions from the highest to the lowest probability and recommend the top one or more.

4. Evaluation

Our evaluation of the proposed approach can be divided into two parts: quantitative evaluation of the QueryQuestions-SE retrieval model and qualitative evaluation via a developer study. The first evaluation strategy aims to find out if clarification questions can be accurately predicted (or recommended) for a newly composed software development-related Web query. The second evaluation strategy focuses on determining if a tool based on the approach described in this paper can be helpful to developers, and what are its strengths and weaknesses. We begin this section by first discussing the quantitative evaluation and then describing the procedure of developer study and collected feedback.

4.1. Retrieval Evaluation

4.1.1. Baselines

As our approach to pose clarification questions is novel to software engineering and there are no similar, comparable techniques that we could compare our model to, we validate the effectiveness of QueryQuestions-SE against: 1)

Simple Baselines – straightforward approaches to ranking that do not require a sophisticated model (e.g., directly computing a similarity based on the GloVe embeddings) and ablation (e.g., using only one part of the model); and 2) *Clarification Question Ranking (CQR) Baselines* – the problem of selecting the highest ranked clarification question from a set of candidate questions is essentially a ranking problem. Therefore, the task of the model is to find the optimal order of candidate questions for a given query. Learning to Rank (LTR) [27] based models rank a set of candidate questions and choose the one with the highest rank, and are a popular approach to compare to when recommending clarification questions [28, 13, 29].

First, we list the Simple Baselines we identified and used:

- *Random*: We generate a random output of 0 (negative) or 1 (positive) for each combination of query and clarification question.
- *Similar Embedding ($\delta \geq 0.5$)*: This baseline generates two vectors: 1) averaged GloVe embedding of the query, and 2) sum of averaged GloVe embedding of the clarification question and averaged GloVe embedding of the common answers. We compute the *cosine similarity* between these two vectors, δ , and set a threshold of $\delta \geq 0.5$ to determine if a query and clarification question are similar.
- *Similar Embedding ($\delta \geq 0.7$)*: This baseline is similar to the previous baseline but with a higher similarity threshold of 0.7.
- *Dissimilar Embedding ($\delta \leq 0.5$)*: In this baseline, we consider if dissimilarity between the GloVe embeddings of a clarification question (and its common answers) and a query are what determines validity. The intuition is that the more dissimilar a clarification question is the more likely the query is in need to the information it offers. We compute the *cosine similarity* δ between these two vectors as before and set a threshold of less than or equal to 0.5.

- *Dissimilar Embedding* ($\delta \leq 0.3$): This baseline is similar to the previous baseline but with a higher dissimilarity threshold of less than or equal to 0.3.
- *Query Only Baseline*: In this baseline, we use only queries to train the model, ignoring the text of the clarification question and the common answers. Instead of clarification question text, we use ids that correspond to each clarification questions (1, 2, 3, ..., 16). In the neural network, we use one-hot encoding of the clarification question id.

We also used the following set of CQR Baselines.

- *LambdaRank* [30]: For this baseline, we adopt the pairwise version of the algorithm, where each query and their relevant clarification question (and answers) pair are considered as positive examples. The query and irrelevant clarification questions (and answers) are considered as negative examples. We use the concatenated Glove word embedding (as described in section 3.2) representation of query, clarification question, and answer used as input features. For LambdaRank implementation, we use Microsoft’s LightGBM² package.
- *XGBoost ranking* [31]: The setup of input documents and input features is as same as LambdaRank. For implementation, we use the open source XGBoost (XGBRanker)³ package.

4.1.2. Metrics

We choose several popular information retrieval evaluation metrics, including *Mean Reciprocal Rank (MRR)*, *Mean Absolute Precision (MAP)*, and *Precision@K (P@K)*.

- *Mean Reciprocal Rank (MRR)*: The goal of MRR is to evaluate how effective is our technique, or a baseline, in locating the *first* valid clarification

²<https://github.com/Microsoft/LightGBM>

³<https://github.com/dmlc/xgboost>

question. It is computed as:

$$MRR = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{1}{rank_i}$$

, where Q is the set of queries in the test set and $rank_i$ refers to the rank position of the first relevant clarification question for the i -th query.

- *Mean Average Precision (MAP)*: For a set of queries, MAP measures how well a model can locate *all* clarification questions relevant to a query. MAP is calculated as the mean of average precision values ($AvgP$) for Q queries:

$$MAP = \frac{1}{|Q|} \sum_{i=1}^{|Q|} AvgP(i)$$

, where Q is the set of queries in the test set and $AvgP(i)$ refers average precision for the query at position i .

- *Precision@K ($P@K$)*: The goal of Precision@K is to measure the number of valid results when considering the top K positions in the ranking. Unlike MRR, it considers all, not only the topmost ranked, results. It is computed as:

$$P@K = \frac{1}{|Q|} \sum_{i=1}^{|Q|} \frac{|v|}{K}$$

, where, as before, Q is the set of queries in the test set and v is the set of valid clarification questions ranked in the top K positions. We use values of 1, 2 and 3 for K as the median of annotated clarification question is 3.

4.1.3. Results and Discussion

We randomly divide our augmented, annotated dataset using a 80-20 split to evaluate QueryQuestions-SE, i.e., 80% train.

We summarize the results of our technique versus the identified baselines in Table 2. The results indicate that our model strongly outperforms all of the baselines, achieving improvements of 0.08 in MRR, 0.2 in MAP, and 0.09 in Precision@1 over the best baseline. The results suggest that the problem

	MRR	MAP	P@1	P@2	P@3
<i>QueryQuestions-SE</i>	0.88	0.77	0.80	0.67	0.57
Simple Baselines:					
Random	0.54	0.43	0.43	0.28	0.22
Similar Emb. ($\delta \geq 0.5$)	0.68	0.55	0.60	0.42	0.32
Similar Emb. ($\delta \geq 0.7$)	0.67	0.54	0.59	0.41	0.31
Dissimilar Emb. ($\delta \leq 0.5$)	0.53	0.45	0.31	0.33	0.31
Dissimilar Emb. ($\delta \leq 0.3$)	0.56	0.44	0.45	0.30	0.23
Query Only	0.56	0.48	0.41	0.31	0.27
CQR Baselines:					
LambdaRank	0.80	0.57	0.70	0.45	0.36
XGBoost	0.79	0.56	0.71	0.45	0.35

Table 2: Evaluation results contrasting our system relative to several baselines.

of finding a valid recommendation question for a query is not easily solved by the simple baselines we devised. Of the simple baselines, we observe that only the similarity baselines noticeably outperform the random choice; both of the similarity baselines, regardless of threshold, perform nearly the same. The two CQR baselines performed similarly to each other and noticeably outperformed the simple baselines, but were in all metrics weaker than QueryQuestions-SE.

From the perspective of whether the QueryQuestions-SE model provides usable results, the key metric is Precision@1 as it indicates how well we recommend a single clarification question, which is the most common use case. This metric for QueryQuestions-SE is 0.8, which indicates 4 out of 5 times we recommend a valid top most result. Precision@2 and Precision@3 decrease significantly from Precision@1 for all techniques including QueryQuestions-SE. The reason for this is that a significant subset of queries contain only one valid clarification question and therefore P@1 has much higher ceiling than P@2 and P@3.

4.2. Software Developer Study

To investigate how our QueryQuestions-SE model would perform in the real world, we conducted a study of software developers using a prototype of our technique as a browser plugin. We conducted a multi-modal study with a total of 6 participants some of which were novices (i.e., new to software development). First, we collected feedback on the clarification questions from the developers leveraging in-use relevance questions posed after each interaction with the tool. Second, the participants completed a survey after interacting with the tool for at least 2 days. Finally, we discussed their experiences and expectations through an in-person interview. In this subsection, we will first explain the functionality of the browser plugin, then describe the study participants and study procedure, and, finally, report and discuss the collected results.

4.2.1. Browser Plugin

To allow study participants to experience receiving clarification questions first hand, we created a Google Chrome browser plugin that monitors user Google queries and reacts with a pop-up if a clarification question is recommended by QueryQuestions-SE. If the prediction score is less than 0.5 for all clarification questions, we do not provide any clarification questions as, in this case, it is likely that the developer has an adequate query that does not need clarifying. The pop-up window appears alongside the Google results page and contains: 1) the highest ranked clarification question; 2) a list of common answers and an input box to provide a typed answer; and 3) two buttons (*‘Update Query’* and *‘Question is Not Relevant’*) at the bottom (see Figure 4). When users press the *‘Update Query’* button, the plugin appends the answer to the original query and performs a Google search, returning a new set of Google search results on the same page. The user can also click on the *‘Question is Not Relevant’* button, which is relevance feedback that we collect for the study. After selecting *‘Update Query’*, alongside the updated Google results page, we ask the user for in-use feedback on the effectiveness of the clarification question via another pop-up that asks *‘Was the prior question useful in suggesting how*

Table 3: Clarification question relevance data collected during use of the plugin.

	Industry Exp.	Active Days	# of Queries	Relevance	Usefulness
P1 (n)	0-3 yrs.	3	12	10/12 (83%)	6/8 (75%), 2 NoAns
P2 (n)	0-3 yrs.	4	23	14/23 (61%)	6/7 (86%), 7 NoAns
P3 (i)	4-6 yrs.	2	21	15/21 (71%)	10/12 (83%), 3 NoAns
P4 (i)	4-6 yrs.	3	24	18/24 (75%)	15/18 (83%)
P5 (i)	4-6 yrs.	4	15	13/15 (87%)	3/4 (75%), 9 NoAns
P6 (i)	6-9 yrs.	6	8	7/8 (87%)	2/6 (33%), 1 NoAns

Table 5: Clarification question relevance data collected during use of the plugin. n = novice, i = industry.

to clarify the query?’ where the user can select either ‘Yes’ or ‘No’. The user can also choose to select no option here at all and close the plugin.

4.2.2. Participants and Procedure

Our goal was to attract two groups of developers for our study: experienced software developers from industry and novice software developers, with the aim of having representatives from both groups. More specifically, through personal contacts, we recruited 4 software developers currently employed in industry (with industry experience ranging from 4 to 9 years) and 2 novice programmers (students that have recently, in the last 1-2 years, learned how to program) that were currently programming on a daily basis.

Figure 4: Browser plugin for user study.

Through a recruitment email, we communicated the basic context of our study and provided a short video on how to install and use the plugin. We asked the participants to watch the video carefully, and then use our plugin during their regular work for software development-related Web searching.

We asked the participants to let us know once they have used the plugin sufficiently to develop an opinion and feel ready to proceed (after at least 2 days), allowing them to take additional time using the tool if they chose. After the participants contacted us that they are ready, we provided them with a link to a short online survey asking several questions regarding their experiences, including:

- How often they used the plugin – *“How often did you use the plugin in your Google searches?”*, with 3 possible answer options – *‘On all searches where I didn’t like the initial Google results’*, *‘On some searches where I didn’t like the initial Google results’*, and *‘On every software development-related Google search’*.
- Clarification question relevance – *“How relevant were the questions provided by the plugin to your searches?”*, answered by a 5 point linear scale (Least Relevant to Most Relevant).

- Clarification question usefulness – *“How useful were the questions provided by the plugin to finding the information you were looking for?”*, answered by a 5 point linear scale (Least Useful to Most Useful).
- Open-ended questions on ideas for improvement – *“How would you improve the clarification questions asked by the plugin?”*, and *“How would you improve the plugin (apart from the questions themselves)?”*.

Following the completion of the study, we organized an online meeting with each participant separately so that they can share their opinions through a semi-structured interview. Our questions in the interview can be organized into 3 topics:

- Questions regarding common search behavior and motivation – *“Do you write queries that are not specific enough?”*, *“When you search, are you usually satisfied with the results you get?”*.
- Questions regarding the clarification questions – *“Do you think asking clarification question on search queries can be useful?”*, *“How did you find the questions the tool posed?”*, *“How do you think the clarification questions can be improved?”*, *“Can you think of any additional clarification questions?”*.
- Questions regarding the plugin – *“Did you find our plugin useful?”*, *“Did the plugin help you get higher quality results from your Google searching?”*, *“What did you like and dislike about the plugin?”*.

4.2.3. Results and Discussion

One of the study participants (P1) was not able to schedule the interview for personal reasons. Otherwise, all other participants completed all aspects of the study, including the in-use (popup-based) feedback and the post-activity survey.

In-Use Feedback. Table 5 lists the number of queries performed by individual software developers, the number of days the plugin was actively used (days with

at least one query performed are counted), and developer feedback on query relevance (as reported by the button on Figure 4) and usefulness captured through an additional popup after the query was reformulated (and the updated Google results shown to the developer). The six developers performed a total of 103 queries, and, of these queries, they deemed 77 (75%) clarification questions as relevant. Out of these 77 queries, 42 clarification questions are selected as useful, 13 are selected as not useful, and no feedback was provided for the remaining 22 queries; the usefulness rate out of the relevant clarification questions was 76% (42/55), and 52% (42/81) overall, not counting the clarification questions without feedback.

Survey. On the question - *“How often did you use the plugin in your Google searches?”* - 3/6 developers selected *‘On all searches where I didn’t like the initial Google results’*, 2/6 developers selected *‘On some searches where I didn’t like the initial Google results’*, and 1/6 developer selected *‘On every software development-related Google search’*.

The relevance and usefulness of the QueryQuestions-SE model as reported by the 6 developers on the survey are shown in Figure 5. The results indicate that 4/6 developers found the clarification questions to be relevant (with a score of least 3 on the linear scale), and 5/6 developers found the clarification questions useful (with a score of least 3 on the linear scale).

The responses to the two open-ended questions - *“How would you improve the clarification questions asked by the plugin?”*, and *“How would you improve the plugin (apart from the questions themselves)?”* - were repeated with more context in the one-on-one interview, therefore, we discuss them there.

Interview. We conducted each interview via Zoom. We started each Zoom meeting with a warm-up question asking about the type of software development-related Web searches performed by each participant. Most of the developers identified more than one search scenario, including both exploratory and lookup [32]. Next, on the question regarding whether they were satisfied with typical search engine results, most participants indicated that they were, however, that they could be improved. When we asked whether clarification ques-




Figure 5: Reported relevance and usefulness of QueryQuestions-SE model.

tions could be useful in improving search, all of them agreed that this is a good mechanism. Specifically, P4 found the experience of our plugin helpful in this regard, stating *“From my experience in using this plugin, I found clarification questions useful.”* P2 and P5 also mentioned that adding terms to queries using clarification questions is a good idea, while P3 pointed out that clarification questions could help to narrow down the search results and provide more direction to generic search queries.

As for the effectiveness of the specific clarification questions provided by our plugin, P2 noted, *“they were good as long as they were relevant.”* P3, P4, and P5 also provided similar opinions, however, P6 disagreed. P6 was concerned that our approach would not help a domain-expert experienced developer and in P6’s opinion, *“I already knew what I was looking for and search it effectively on Google.”* However, others found the clarification questions helpful, e.g., P3 noted, *“It helped me to clarify contexts on some queries when google results were too broad.”* P2 remarked, *“I liked it when I just didn’t know what to add to a search. The tool provided some suggestions.”* P5 noted, *“It did. I was able to get a better query. And because of that better results in the search.”*

The participants offered several ideas to improve the plugin. P4 remarked, *“What I like is that it is easy to use and questions were most of the cases relevant... It may be a good improvement if this can be made more interactive.”* The interactive idea was supported by others as well, e.g., P6 stated *“What kind of queries developers make mostly? - for debugging issues by searching for the*

specific errors, exceptions, etc; tutorials, examples, documentation of specific library, framework, etc. The clarification question can ask for the type of query first, and based on that it can suggest specific terms to improve the query.”

P5 suggested improving the approach by using dynamic options for the common answers suggested by the tool, which would provide common answers that are more tightly integrated to the specific query. P3 and P6 also expressed that dynamically generated answer options could be a way to improve the existing tool. P1 supported this idea as well in her response to the open-ended survey questions. P5 had another suggestion on improving the tool by reformulating the queries multiple ways, e.g., other than adding a term at the end, we could add a term at the beginning, remove unnecessary terms (such as adverbs, pronouns), etc.

Lastly we asked, *“based on your experience using this plugin [...] would you use it in your daily development work?”* All of the participants except for P6 replied positively. P2 said, *“Yeah, I will. I don’t think I would use it like every search. But I think I would use it a fair amount of time.”* P5 remarked, *“Yes, yes, of course, it was helping me get better queries.”* P3 mentioned, *“Yeah, I will use it as the functionality of the tool narrows down the search domain and saves development time.”* P4 commented, *“Why not if the questions are relevant and the tool can serve my purpose?”* P6, however, didn’t dismiss using improved version of the tool, *“Maybe (I will) when you could improve the tool as the suggestions we have discussed... I like the approach that the tool is trying to help to search and find better answers.”*

Summary. The developer study data collected from the 6 software developers indicated that the idea of clarification questions has potential. The clarification questions were mostly valuable (in-use relevance between 61%-87% and usefulness between 75%-86%, without P6 (33%); survey at or above 3 out of 5 for relevance in 4/6 participants and for usefulness in 5/6 participants). In one-on-one interviews, they noted that clarification questions can provide convenient suggestions, narrow down search space, and save development time. The key suggestion for improvement was to add more dynamicity and interactivity, e.g.,

suggesting more bespoke answers to the questions or engaging the developers in a conversation as in conversational search [16].

Comparison to Google’s Related Searches. In order to understand if the clarification questions can provide the developers with a functionality different from what is currently available by popular search engines, we compared the queries expanded by the developers (using our clarification questions) with the set of query suggestions provided by Google’s Related Searches functionality. Google’s Related Searches appear at the bottom of each results page, and consist of 8 suggested queries that reformulate the current query provided to the search engine. While the algorithm for selecting the 8 suggested queries is not publicly available, Google states that it is based on a corpus of past search queries.

Out of the 42 queries that were successfully expanded by the developers in our study, 9 (21%) had an exact or near exact match among the queries suggested by Google Related Searches (e.g., original query = *‘http vs grpc’*; expanded query = *‘http vs grpc performance’*; Google Related Searches = {*‘grpc vs rest performance’*, *‘grpc vs protobuf’*, ...}). While this indicates that most of the time (79%) the developer’s queries formulated by using our clarifying questions were different, many queries provided by Google Related Searches were topically related to the queries expanded by the clarification questions. However, we observed that in 12/42 (29%) queries, all of the 8 queries in Google Related Searches were completely and significantly different from the expanded queries generated by the developers (e.g., original query = *‘unity mrtk detect pointer in collider’*; expanded query = *‘unity mrtk detect pointer in collider fix tutorial’*; Google Related Searches = {*‘mrtk gestures’*, *‘mrtk disable pointer’*, ...}). Our hypothesis is that this scenario occurs when the original queries are longer and more specific as in that case Google Related Searches suggest more simple and generic alternatives that significantly miss the topical direction that the developer is interested in.

4.3. Threats to Validity

Several limitations may impact the interpretation of our findings on the proposed approach for automatically posing clarification questions to help developers’ Web search. We categorize and list each of them below.

Construct validity. Threats to construct validity include the degree to which our experiments measure what was intended. First, as our original Web query dataset includes only 200 queries from a relatively narrow scope (i.e., Java software development), the model’s training dataset is limited by the topics within those queries as well. More specifically, the validity concern related to this is whether the clarification questions and the model for predicting clarification questions overfit a specific set of topics. One mitigating factor is augmenting our dataset (using RoBERTa) by replacing query terms, which actually generates queries that are no longer Java-specific (e.g., replacing the term ‘*java*’ with the term ‘*python*’). Another means to mitigate this concern is the two ways of evaluation of our model and technique, both indicating that the clarification questions and the model were generally relevant to the developer’s search tasks. For instance, in the developer survey, 77/103 (75%) of clarification questions were deemed relevant, which is very similar to the model’s quantitative evaluation results for Precision@1 of 0.8 (80%).

Another threat is the use of common answers as input to our model. As the common answers are based on facets related to queries, they may be too specific to the particular domain, and, again, lead to overfitting. With different and more diverse common answers, the model may learn different patterns. We mitigate this issue by selecting the broadest common terms, which are likely to be valid across many different software engineering-specific topics.

Internal validity. Threats to internal validity concern study design factors that may influence the results. One threat to our study is the artificially generated dataset. A mitigating factor is the use of a powerful natural language model, i.e., RoBERTa, to suggest terms to augment our queries. However, it is possible that some of the queries may be unrealistic and never occur word-for-

word in real-life. We address this threat by manually checking each augmented query to be semantically correct and to still be in need of a specific clarification question.

External validity. Threats to external validity concern generalizing the experimental findings beyond our settings. One such limitation of our approach is the fixed number of clarification questions, which are also limited by our input set of 200 Java-related Bing queries. It is possible to elicit additional clarification questions with a larger dataset of queries. For example, we observe that the common clarification question ‘*Which programming language are you interested in?*’ is missing in our study. This is because our annotated query dataset contains only Java-related queries and each query explicitly mentions ‘java’. To mitigate this threat, we performed a tool-based study with developers that worked on other development tasks that did not necessarily include Java (e.g., Go, Unity, C#, PHP).

5. Related Work

The related work can be categorized into three parts: 1) empirical studies of software development-related Web search, 2) generating facets from queries, and 3) automatically posing clarification questions.

Empirical Studies of Software Development-Related Web Search.

Developers use general purpose search engines daily to facilitate their work. Therefore, there has been considerable interest among researchers in studying how developers perform these searches [33, 2, 5, 34, 6, 35, 4, 36, 3]. There have been several studies on improving search specifically for source code [37, 34, 38, 39]. For example, Stolee et al.’s survey notes that 85% of the software developers use search engines to search for source codes on the Web at least weekly [40]. They also found that 69% of the respondents used Web search engines (i.e., Google, Bing) instead of bespoke code search tools. Xia et al. noted that developers search for many things other than source code, e.g., debugging,

documentation, best practices, how to use a particular tool, etc. [5] Temla et al. noted that game developers use Web search to get various kinds of help (e.g., instructions, algorithms, and tools) [41]. In addition, there has been research dedicated to developers' Web search for software architecture design patterns [42], algorithms [43], APIs [44], etc.

Sadowski et al. surveyed 27 developers over a two week period to better understand when and why software developers search [2]. Rahman et al. collected search logs from 310 developers, consisting of nearly 150,000 Google search queries and built an automatic classifier that detected software engineering related queries [6]. Hora et al. collected 1.3M developers' queries from a Search Engine Optimization service and concluded that the queries are usually short, tend to omit functional words, and are similar among each other [3]. Xia et al. collected search queries from 60 developers and surveyed 235 software engineers and grouped search tasks into seven groups: General Search, Debugging and Bug Fixing, Programming, Third Party Code Reuse, Tools, Database, and Testing [5]. Rao et al. used search query logs from Bing to analyze Web search behavior for software engineering tasks and classified the search intents into seven categories: API, Debug, HowTo, Installation, Learn, Navigational and Miscellaneous [4].

However, most of the above mentioned datasets are not publicly available to analyze further. In an attempt to encourage contributions to this area, Rao et al. open-sourced a portion of their Bing dataset [17].

Generating Facets from Queries.

Manually and automatically extracting facets from Web query logs has been an area of interest in the research community for a long period of time due to its potential for query reformulation [45, 22, 46, 20, 47, 13, 48, 49]. Li et al. proposed a system that can automatically generate facets for a query using Wikipedia [50]. Dou et al. developed an unsupervised model that automatically extracts facets for a query [22, 46]. Kong et al. developed a supervised graphical

model to extract facets for a query based on a set of Web search results [20, 47]. Jiang et al. proposed a model to generate facets from knowledge bases [23]. Aliannejadi et al. applied a manual facet generation process similar to ours for a corpus of queries [13]. Specifically aimed at software engineering-related search, Zhang et al. used Stack Overflow tags as facets [15].

Automatically Posing Clarification Questions.

Automatically posing clarification questions has recently become an popular area of study in the fields of Natural Language Processing and Information Retrieval [13, 51, 15, 14, 52, 53, 54, 55, 56, 57, 58]. The research on asking clarification questions has been applied within a few different domains and applications such as chatbots [59], open-domain question answering systems [60, 61], domain-specific question answering systems [15, 62, 63], search engines [14], search queries [28], image content [64], and e-commerce [51].

Clarification questions are a mechanism to refine user Web search queries that miss key information. Aliannejadi et al. built a model for query refinement to aid open-domain information seeking conversations [13]. Zamani et al. proposed a template-based clarification question generation model based on different aspects of a Web query [58]. Zamani et al. later open-sourced a dataset for open domain clarification question generation [65] and provided an analysis of user interaction with clarification questions posed in a popular Web search engine (Bing) [66]. Sekulić et al. studied how to predict user engagement with posed clarification questions based on the open source Web search dataset released by Zamani et al. [67] In our work, we focus on software development-related clarification questions and release a dataset of annotated queries, facets, and clarification questions.

6. Conclusions and Future Work

This paper illustrates a technique for posing clarification questions for developers’ Web search queries. In this paper, first, we identify common clarification questions based on annotating 200 real-world user queries. Then, we create

an augmented dataset of 5,151 queries that we use to train a machine learning model that can identify the most relevant clarification question for a given user query. Our evaluation shows that the proposed approach outperforms all baselines, while a study of software developers indicates that the clarification question relevance rate aligned with our quantitative (i.e., test set) results. That is, the participants found 75% clarification questions were relevant, while the Precision@1 on the test set was 80%. The feedback from the software developer study, using a prototype we created, suggests that an approach like ours can be useful to software developers.

There are several avenues for future work. First, we intend on making the common answers dynamic, i.e., more adaptable to the specific query. Second, we plan to research an interactive conversation based tool where the users can ask clarification questions successively and engage the search system in a type of a conversation. Third, another line of future work is in identifying different set of clarification questions based on query intents proposed by various literature [5, 4]. Lastly, we can extend the evaluation to a larger set of developers, which would help in further understanding how to improve the current model and clarification questions.

References

- [1] P. Chatterjee, M. A. Nishi, K. Damevski, V. Augustine, L. Pollock, N. A. Kraft, What information about code snippets is available in different software-related documents? an exploratory study, in: 2017 IEEE 24th International Conference on Software Analysis, Evolution and Reengineering (SANER), 2017, pp. 382–386. doi:10.1109/SANER.2017.7884638.
- [2] C. Sadowski, K. T. Stolee, S. Elbaum, How developers search for code: a case study, in: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, 2015, pp. 191–201.
- [3] A. Hora, Googling for software development: What developers search for and what they find, in: 2021 IEEE/ACM 18th Interna-

- tional Conference on Mining Software Repositories (MSR) (MSR), IEEE Computer Society, Los Alamitos, CA, USA, 2021, pp. 317–328. doi:10.1109/MSR52588.2021.00044.
URL <https://doi.ieeecomputersociety.org/10.1109/MSR52588.2021.00044>
- [4] N. Rao, C. Bansal, T. Z. 0001, A. H. Awadallah, N. Nagappan, Analyzing web search behavior for software engineering tasks, in: X. Wu, C. Jermaine, L. X. 0001, X. H. 0001, O. Kotevska, S. Lu, W. Xu, S. Aluru, C. Zhai, E. Al-Masri, Z. C. 0003, J. S. 0001 (Eds.), IEEE International Conference on Big Data, Big Data 2020, Atlanta, GA, USA, December 10-13, 2020, IEEE, 2020, pp. 768–777. doi:10.1109/BigData50022.2020.9378083.
URL <https://doi.org/10.1109/BigData50022.2020.9378083>
- [5] X. Xia, L. Bao, D. Lo, P. S. Kochhar, A. E. Hassan, Z. Xing, What do developers search for on the web?, *Empirical Software Engineering* 22 (6) (2017) 3149–3185.
- [6] M. M. Rahman, J. Barson, S. Paul, J. Kayani, F. A. Lois, S. F. Quezada, C. Parnin, K. T. Stolee, B. Ray, Evaluating how developers use general-purpose web-search for code retrieval, in: *Proceedings of the 15th International Conference on Mining Software Repositories*, 2018, pp. 465–475.
- [7] M. Lu, X. Sun, S. Wang, D. Lo, Y. Duan, Query expansion via wordnet for effective code search, in: *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, IEEE, 2015, pp. 545–549.
- [8] M. M. Rahman, C. Roy, Effective reformulation of query for code search using crowdsourced knowledge and extra-large data analytics, in: *2018 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2018, pp. 473–484. doi:10.1109/ICSME.2018.00057.
- [9] K. Cao, C. Chen, S. Baltes, C. Treude, X. Chen, Automated query reformulation for efficient search based on query logs from stack overflow, in:

- 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE), 2021, pp. 1273–1285. doi:10.1109/ICSE43902.2021.00116.
- [10] S. Haiduc, G. Bavota, A. Marcus, R. Oliveto, A. De Lucia, T. Menzies, Automatic query reformulations for text retrieval in software engineering, in: 2013 35th International Conference on Software Engineering (ICSE), 2013, pp. 842–851. doi:10.1109/ICSE.2013.6606630.
 - [11] F. Zhang, H. Niu, I. Keivanloo, Y. Zou, Expanding queries for code search using semantically related api class-names, IEEE Transactions on Software Engineering 44 (11) (2017) 1070–1082.
 - [12] M. M. Rahman, C. K. Roy, D. Lo, Automatic query reformulation for code search using crowdsourced knowledge, Empirical Software Engineering 24 (4) (2019) 1869–1924.
 - [13] M. Aliannejadi, H. Zamani, F. Crestani, W. B. Croft, Asking clarifying questions in open-domain information-seeking conversations, in: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2019, pp. 475–484.
 - [14] P. Ren, Z. Chen, Z. Ren, E. Kanoulas, C. Monz, M. De Rijke, Conversations with search engines: Serp-based conversational response generation, ACM Trans. Inf. Syst. 39 (4) (Aug. 2021). doi:10.1145/3432726.
URL <https://doi.org/10.1145/3432726>
 - [15] N. Zhang, Q. Huang, X. Xia, Y. Zou, D. Lo, Z. Xing, Chatbot4qr: Interactive query refinement for technical question retrieval, IEEE Transactions on Software Engineering (2020).
 - [16] F. Radlinski, N. Craswell, A theoretical framework for conversational search, CHIIR '17, Association for Computing Machinery, New York, NY, USA, 2017, p. 117–126. doi:10.1145/3020165.3020183.
URL <https://doi.org/10.1145/3020165.3020183>

- [17] N. Rao, C. Bansal, J. Guan, Search4code: Code search intent classification using weak supervision (2021) 575–579doi: 10.1109/MSR52588.2021.00077.
URL <https://doi.ieeecomputersociety.org/10.1109/MSR52588.2021.00077>
- [18] X. Wang, D. Chakrabarti, K. Punera, Mining broad latent query aspects from search sessions, in: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining, 2009, pp. 867–876.
- [19] X. Wang, C. Zhai, Learn from web search logs to organize search results, in: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval, 2007, pp. 87–94.
- [20] W. Kong, J. Allan, Extracting query facets from search results, in: Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval, 2013, pp. 93–102.
- [21] R. Song, M. Zhang, T. Sakai, M. P. Kato, Y. Liu, M. Sugimoto, Q. Wang, N. Orii, Overview of the ntcir-9 intent task., in: NTCIR, Citeseer, 2011.
- [22] Z. Dou, S. Hu, Y. Luo, R. Song, J.-R. Wen, Finding dimensions for queries, in: Proceedings of the 20th ACM international conference on Information and knowledge management, 2011, pp. 1311–1320.
- [23] Z. Jiang, Z. Dou, J.-R. Wen, Generating query facets using knowledge bases, IEEE transactions on knowledge and data engineering 29 (2) (2016) 315–329.
- [24] J. Chen, Z. Wang, R. Tian, Z. Yang, D. Yang, Local additivity based data augmentation for semi-supervised ner, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 1241–1251.

- [25] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, arXiv preprint arXiv:1907.11692 (2019).
- [26] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, CoRR abs/1412.6980 (2015).
- [27] T.-Y. Liu, et al., Learning to rank for information retrieval, *Foundations and Trends® in Information Retrieval* 3 (3) (2009) 225–331.
- [28] J. Wang, W. Li, Template-guided clarifying question generation for web search clarification, in: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 3468–3472.
- [29] Z. Gao, X. Xia, D. Lo, J. Grundy, Technical q8a site answer recommendation via question boosting, *ACM Transactions on Software Engineering and Methodology (TOSEM)* 30 (1) (2020) 1–34.
- [30] C. Burges, R. Ragno, Q. Le, Learning to rank with nonsmooth cost functions, *Advances in neural information processing systems* 19 (2006).
- [31] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, pp. 785–794.
- [32] G. Marchionini, Exploratory search: from finding to understanding, *Communications of the ACM* 49 (4) (2006) 41–46.
- [33] J. Brandt, P. J. Guo, J. Lewenstein, M. Dontcheva, S. R. Klemmer, Two studies of opportunistic programming: interleaving web foraging, learning, and writing code, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2009, pp. 1589–1598.
- [34] L. Martie, A. v. d. Hoek, T. Kwak, Understanding the impact of support for iteration on code search, in: *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, 2017, pp. 774–785.

- [35] M. Hucka, M. J. Graham, Software search is not a science, even among scientists: A survey of how scientists and engineers find software, *Journal of Systems and Software* 141 (2018) 171–191.
- [36] S. E. Sim, M. Umarji, S. Ratanotayanon, C. V. Lopes, How well do search engines support code retrieval on the web?, *ACM Transactions on Software Engineering and Methodology (TOSEM)* 21 (1) (2011) 1–25.
- [37] S. Bajracharya, T. Ngo, E. Linstead, Y. Dou, P. Rigor, P. Baldi, C. Lopes, Sourcerer: a search engine for open source code supporting structure-based search, in: *Companion to the 21st ACM SIGPLAN symposium on Object-oriented programming systems, languages, and applications*, 2006, pp. 681–682.
- [38] K. Kim, D. Kim, T. F. Bissyandé, E. Choi, L. Li, J. Klein, Y. L. Traon, Facoy: a code-to-code search engine, in: *Proceedings of the 40th International Conference on Software Engineering*, 2018, pp. 946–957.
- [39] R. Holmes, Do developers search for source code examples using multiple facts?, in: *2009 ICSE Workshop on Search-Driven Development-Users, Infrastructure, Tools and Evaluation*, IEEE, 2009, pp. 13–16.
- [40] K. T. Stolee, S. Elbaum, D. Dobos, Solving the search for source code, *ACM Transactions on Software Engineering and Methodology (TOSEM)* 23 (3) (2014) 1–45.
- [41] P. Tamla, T. Böhm, K. Gaisbachgrabner, J. Mertens, M. L. Hemmje, M. Fuchs, Survey: Software search in serious games development., in: *CERC*, 2019, pp. 155–166.
- [42] M. Soliman, M. Wiese, Y. Li, M. Riebisch, P. Avgeriou, Exploring web search engines to find architectural knowledge, in: *2021 IEEE 18th International Conference on Software Architecture (ICSA)*, IEEE, 2021, pp. 162–172.

- [43] S. Bhatia, S. Tuarob, P. Mitra, C. L. Giles, An algorithm search engine for software developers, in: *Proceedings of the 3rd International Workshop on Search-Driven Development: Users, Infrastructure, Tools, and Evaluation*, 2011, pp. 13–16.
- [44] J. Stylos, B. A. Myers, Mica: A web-search tool for finding api components and examples, in: *Visual Languages and Human-Centric Computing (VL/HCC'06)*, IEEE, 2006, pp. 195–202.
- [45] E. Stoica, M. A. Hearst, M. Richardson, Automating creation of hierarchical faceted metadata structures, in: *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, 2007, pp. 244–251.
- [46] Z. Dou, Z. Jiang, S. Hu, J.-R. Wen, R. Song, Automatically mining facets for queries from their search results, *IEEE Transactions on knowledge and data engineering* 28 (2) (2015) 385–397.
- [47] W. Kong, J. Allan, Extending faceted search to the general web, in: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, 2014, pp. 839–848.
- [48] W. Kong, J. Allan, Precision-oriented query facet extraction, in: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 2016, pp. 1433–1442.
- [49] J. Friedrich, C. Lindemann, M. Petrifke, Utilizing query facets for search result navigation, in: *2015 26th International Workshop on Database and Expert Systems Applications (DEXA)*, IEEE, 2015, pp. 271–275.
- [50] C. Li, N. Yan, S. B. Roy, L. Lisham, G. Das, Facetedpedia: dynamic generation of query-dependent faceted interfaces for wikipedia, in: *Proceedings of the 19th international conference on World wide web*, 2010, pp. 651–660.

- [51] Y. Zhang, X. Chen, Q. Ai, L. Yang, W. B. Croft, Towards conversational search and recommendation: System ask, user respond, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, 2018, pp. 177–186.
- [52] C. Rosset, C. Xiong, X. Song, D. Campos, N. Craswell, S. Tiwary, P. Bennett, Leading conversational search by suggesting useful questions, in: Proceedings of The Web Conference 2020, 2020, pp. 1160–1170.
- [53] L. T. Hien, L. T. T. Ly, C. Pham-Nguyen, T. Le Dinh, H. T. Gia, et al., Towards chatbot-based interactive what-and how-question answering systems: the adobot approach, in: 2020 RIVF International Conference on Computing and Communication Technologies (RIVF), IEEE, 2020, pp. 1–3.
- [54] A. Elgohary, D. Peskov, J. Boyd-Graber, Can you unpack that? learning to rewrite questions-in-context, Can You Unpack That? Learning to Rewrite Questions-in-Context (2019).
- [55] S. Vakulenko, S. Longpre, Z. Tu, R. Anantha, Question rewriting for conversational question answering, in: L. Lewin-Eytan, D. Carmel, E. Yom-Tov, E. Agichtein, E. Gabrilovich (Eds.), WSDM '21, The Fourteenth ACM International Conference on Web Search and Data Mining, Virtual Event, Israel, March 8-12, 2021, ACM, 2021, pp. 355–363. doi: 10.1145/3437963.3441748.
URL <https://doi.org/10.1145/3437963.3441748>
- [56] R. Anantha, S. Vakulenko, Z. Tu, S. Longpre, S. Pulman, S. Chappidi, Open-domain question answering goes conversational via question rewriting, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online, 2021, pp. 520–534. doi:10.18653/v1/2021.naacl-main.44.
URL <https://aclanthology.org/2021.naacl-main.44>

- [57] S. Stoyanchev, A. Liu, J. Hirschberg, Towards natural clarification questions in dialogue systems, in: AISB symposium on questions, discourse and dialogue, Vol. 20, 2014.
- [58] H. Zamani, S. Dumais, N. Craswell, P. Bennett, G. Lueck, Generating clarifying questions for information retrieval, in: Proceedings of The Web Conference 2020, 2020, pp. 418–428.
- [59] B. Hancock, A. Bordes, P.-E. Mazaré, J. Weston, Learning from dialogue after deployment: Feed yourself, chatbot!, in: A. Korhonen, D. R. Traum, L. Màrquez (Eds.), Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers, Association for Computational Linguistics, 2019, pp. 3667–3684.
URL <https://www.aclweb.org/anthology/P19-1358/>
- [60] M. De Boni, S. Manandhar, An analysis of clarification dialogue for question answering, in: Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, 2003, pp. 48–55.
- [61] M. De Boni, S. Manandhar, Implementing clarification dialogues in open domain question answering, *Natural Language Engineering* 11 (4) (2005) 343–362.
- [62] J. Trienes, K. Balog, Identifying unclear questions in community question answering websites, in: European Conference on Information Retrieval, Springer, 2019, pp. 276–289.
- [63] P. Braslavski, D. Savenkov, E. Agichtein, A. Dubatovka, What do you mean exactly? analyzing clarification questions in cqa, in: Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, 2017, pp. 345–348.

- [64] N. Mostafazadeh, I. Misra, J. Devlin, M. Mitchell, X. He, L. Vanderwende, Generating natural questions about an image, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers, The Association for Computer Linguistics, 2016.
URL <http://aclweb.org/anthology/P/P16/P16-1170.pdf>
- [65] H. Zamani, G. Lueck, E. Chen, R. Quispe, F. Luu, N. Craswell, Mimics: A large-scale data collection for search clarification, in: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, 2020, pp. 3189–3196.
- [66] H. Zamani, B. Mitra, E. Chen, G. Lueck, F. Diaz, P. N. Bennett, N. Craswell, S. T. Dumais, Analyzing and learning from user interactions for search clarification, in: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2020, pp. 1181–1190.
- [67] I. Sekulić, M. Aliannejadi, F. Crestani, User engagement prediction for clarification in search, in: Proceedings of the Advances in Information Retrieval. ECIR 2021, 2021, pp. 619–633.