



Identifying Unclear Questions in Community Question Answering Websites

Jan Trienes¹(✉) and Krisztian Balog²

¹ University of Twente, Enschede, Netherlands
jan.trienes@gmail.com

² University of Stavanger, Stavanger, Norway
krisztian.balog@uis.no

Abstract. Thousands of complex natural language questions are submitted to community question answering websites on a daily basis, rendering them as one of the most important information sources these days. However, oftentimes submitted questions are unclear and cannot be answered without further clarification questions by expert community members. This study is the first to investigate the complex task of classifying a question as clear or unclear, i.e., if it requires further clarification. We construct a novel dataset and propose a classification approach that is based on the notion of similar questions. This approach is compared to state-of-the-art text classification baselines. Our main finding is that the similar questions approach is a viable alternative that can be used as a stepping stone towards the development of supportive user interfaces for question formulation.

1 Introduction

The emergence of community question answering (CQA) forums has transformed the way in which people search for information on the web. As opposed to web search engines that require an information seeker to formulate their information need as a typically short keyword query, CQA systems allow users to ask questions in natural language, with an arbitrary level of detail and complexity. Once a question has been asked, community members set out to provide an answer based on their knowledge and understanding of the question. Stack Overflow, Yahoo! Answers and Quora depict popular examples of such CQA websites.

Despite their growing popularity and well established expert communities, increasing amounts of questions remain ignored and unanswered because they are too short, unclear, too specific, hard to follow or they fail to attract an expert member [2]. To prevent such questions from being asked, the prediction of question quality has been extensively studied in the past [1, 13, 15]. Increasing the question quality has a strong incentive as it directly affects answer quality [9], which ultimately drives the popularity and traffic of a CQA website. However, such attempts ignore the fact that even a high-quality question may

Ask a question		Similar questions talked about....
Title	Command line tool for video editing?	<i>Consider the highlighted aspects to improve the clarity of your question.</i>
Body	I would like to combine two videos into one. The process should be automated (for example a script). Can anyone recommend a video editing tool that allows to do this?	<ul style="list-style-type: none"> * What have you tried with ffmpeg? * Are you using windows linux unix osx BSD? * Merging in the sense of appending them? * Do you have video editing software?
		Ask question

Fig. 1. Envisioned question formulation interface. If a question is found to be unclear, a list of clarification questions (obtained from similar questions) is presented to encourage the user to include information that may be required to provide an answer.

lack an important detail that requires clarification. On that note, previous work attempts to identify what aspects of a question requires editing. While the need for editing can be reliably detected, the prediction of whether or not a question lacks important details has been shown to be difficult [19]. In order to support an information seeker in the formulation of her question and to increase question quality, we envision the following two-step system: (1) determine whether a question requires clarification (i.e., it is unclear), and (2) automatically generate and ask clarifying questions that elicit the missing information. This paper addresses the first step. When successful, the automated identification of unclear questions is believed to have a strong impact on CQA websites, their efforts to increase question quality and the overall user experience; see Fig. 1 for an illustration.

We phrase the unclear question detection as a supervised, binary classification problem and introduce the Similar Questions Model (SQM), which takes characteristics of similar questions into account. This model is compared to state-of-the-art text classification baselines, including a bag-of-words model and a convolutional neural network. Our experimental results show that this is a difficult task that can be solved to a limited extent using traditional text classification models. SQM provides a sound and extendable framework that has both comparable performance and promising options for future extensions. Specifically, the model can be used to find keyphrases for question clarification that may be utilized in a question formulation interface as shown in Fig. 1. Experiments are conducted on a novel dataset including more than 6 million labeled Stack Exchange questions, which we release for future research on this task.¹

2 Related Work

Previous work on question modeling of CQA forums can be roughly grouped into three categories: *question quality prediction*, *answerability prediction* and *question review prediction* [17]. With respect to the prediction of question quality,

¹ The dataset and sources can be found at <https://github.com/jantrienes/ecir2019-qac>.

user reputation has been found to be a good indicator [18]. Also, several machine learning techniques have been applied including topic and language models [1, 15]. However, there is no single objective definition of quality, as such a definition depends on the community standards of the given platform. In this paper, we do not consider question quality itself, since a question may lack an important detail regardless of whether its perceived quality is high or low. Question answerability has been studied by inspecting unanswered questions on Stack Overflow [2]. Lack of clarity and missing information is among the top five reasons for a question to remain unanswered. Here, we do not consider other problems such as question duplication and too specific or off-topic questions [5]. Finally, question review prediction specifically attempts to identify questions that require future editing. Most notably, Yang et al. [19] determine if a question lacks a code example, context information or failed solution attempts based on its contents. However, they disregard the task of predicting whether detail (e.g., a software version identifier) is missing and limit their experiments to the programming domain.

Clarification questions have been studied in the context of synchronous Q&A dialog systems. Kato et al. [6] analyzed how clarification requests influence overall dialog outcomes. In contrast to them, we consider asynchronous CQA systems. With respect to asynchronous systems, Braslavski et al. [3] categorized clarification questions from two Stack Exchange domains. They point out that the detection of unclear questions is a vital step towards a system that automatically generates clarification questions. To the best of our knowledge, we are the first study to address exactly this novel unclear question detection task. Finally, our study builds on recent work by Rao and Daumé III [14]. We extend their dataset creation heuristic to obtain both clear and unclear questions.

3 Unclear Question Detection

The unclear question detection task can be seen as a binary classification problem. Given a dataset of N questions, $Q = \{q_1, \dots, q_N\}$, where each question belongs to either the *clear* or *unclear* class, predict the class label for a new (unseen) question q . In this section, we propose a model that utilizes the characteristics of similar questions as classification features. This model is compared to state-of-the-art text classification models described in Sect. 4.2.

We define a question to be *unclear* if it received a clarification question, and as *clear* if an answer has been provided without such clarification requests. This information is only utilized to obtain the ground truth labels. Furthermore, it is to be emphasized that it is most useful to detect unclear questions during their creation-time in order to provide user feedback and prevent unclear questions from being asked (see the envisioned use case in Fig. 1). Consequently, the classification method should not utilize input signals available only after question creation, such as upvotes, downvotes or conversations in form of comments. Finally, we do not make any assumptions about the specific representation of a question as it depends on the CQA platform at hand. The representation we employ for our experiments on Stack Exchange is given in Sect. 4.1.

3.1 Similar Questions Model

The Similar Questions Model is centered around the idea that similar existing questions may provide useful indicators about the presence or absence of information. For example, consider the two questions in Table 1. It can be observed that the existing question specifies additional information after a clarification question has been raised. A classification system may extract keyphrases (e.g., *operating system*) from the clarification question and check whether this information is present in the given question (see Fig. 1 and Table 7 for examples). In other words, the system checks if a new question lacks information that was also missing from similar previous questions. It has been shown that this general approach can be successfully employed to find and rank suitable clarification questions [14].

Table 1. Example of a new unclear question (Left) and a similar existing question (Right). The left question fails to specify the operating system. The text in italics has been added in response to the shown comment.

Field	New question	Similar (Unclear) question
Title	Simplest XML editor	XML editing/Viewing software
Body	I need the simplest editor with utf8 support for editing xml files; It's for a non programmer (so no atom or the like), to edit existing files. Any suggestion?	What software is recommended for working with and editing large XML schemas? <i>I'm looking for both Windows and Linux software (doesn't have to be cross platform, just want suggestions for both)</i>
Tags	xml, utf8, editors	Windows, xml, linux
Comments	-	What operating system?

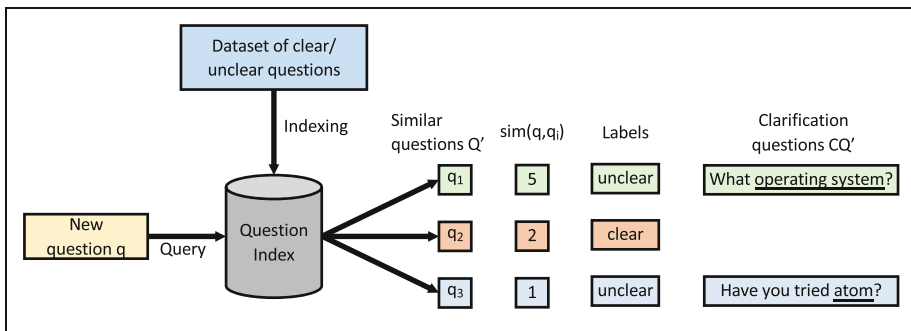


Fig. 2. Illustration of Similar Questions Model. The underlined text of a clarification question indicates a keyphrase.

The Similar Questions Model can be formalized as follows. Given a new question q , we first seek a set of k similar questions $Q' = \{q'_1, q'_2, \dots, q'_k\}$ with their clear and unclear labels. As per the definition of unclear that we employ, the subset of unclear questions $Q'_{unclear}$ has a set of M corresponding clarification questions $CQ' = \{cq'_1, cq'_2, \dots, cq'_M\}$. Within this framework, we design a number of indicative features that are then used to train a classifier to discriminate between the two classes. An illustration of the model can be found in Fig. 2.

Table 2. Features employed by the Similar Questions Model. The example values in the last column are based on the scenario presented in Table 1 and Fig. 2.

(i) Features based on q		Ex
$Len(q)$	Question length in the number of tokens: $ q $	41
$ContainsPre(q)$	Indicator if question contains preformatted elements	0
$ContainsQuote(q)$	Indicator if question contains a quote	0
$ContainsQuest(q)$	Indicator if question contains question mark “?”	1
$Readability(q)$	Coleman-Liau Index (CLI) [4]	16.7
(ii) Features based on Q'		
$SimSum(q, Q')$	Sum of similarity scores: $\sum_{q' \in Q'} sim_{BM25}(q, q')$	8
$SimMax(q, Q')$	Maximum similarity: $\max_{q' \in Q'} sim_{BM25}(q, q')$	5
$SimAvg(q, Q')$	Average similarity: $\frac{1}{ Q' } \sum_{q' \in Q'} sim_{BM25}(q, q')$	2.7
$LenSim(Q')^a$	Number of similar questions retrieved: $ Q' $	3
$LenUnclear(Q')^a$	Number of similar questions that are unclear: $ Q'_{unclear} $	2
$LenClear(Q')^a$	Number of similar questions that are clear: $ Q'_{clear} $	1
$Majority(Q')^a$	Majority vote of labels in Q'	1
$Ratio(Q')^a$	Ratio between clear/unclear questions: $ Q'_{clear} / Q'_{unclear} $	0.5
$Fraction(Q')^a$	Proportion of clear questions among similar: $ Q'_{clear} / Q' $	0.3
(iii) Features based on CQ'^b		
$CQGlobal(q, CQ')$	Cosine similarity between all keyphrases in CQ' and q	0.6
$CQIndividual(q, CQ')$	Sum of cosine similarities between each keyphrase and q	1
$CQWeighted(q, CQ')$	Like above, but weighted by $sim_{BM25}(q, q')$, see Eq. 1	1

^a These features are computed for the top- k similar questions in Q' where $k = \{10, 20, 50\}$.

^b CQ' is obtained from the top $k = 10$ similar questions in $Q'_{unclear}$.

3.2 Features

The features employed by the Similar Questions Model can be grouped into three classes: (i) features based on q only, (ii) features based on the set of similar questions Q' and (iii) features based on the set of clarification questions CQ' . See Table 2 for a summary. We highlight the computation of the scoring features obtained from the set of clarification questions (group (iii) in Table 2). For each clarification question in CQ' , one or more keyphrases are extracted. These keyphrases are the central objects of a clarification question and refer to an aspect of the original question that is unclear (see Table 7 for examples).

Afterwards, we define $f(a) = (p_1, \dots, p_i, \dots, p_L)$ to represent a question or clarification question as a vector, where each element indicates the number of times a keyphrase p_i occurs in q and cq' , respectively. Then, a *question clarity score* is obtained by computing the cosine similarity between these vectors. The scoring features differ in the way the keyphrase vectors are created. The global model constructs a single vector consisting of all keyphrases present in CQ' , whereas the individual model computes the sum of the scores considering each $cq' \in CQ'$ separately. For the individual weighted feature, the final score is given by:

$$CQWeighted(q, CQ') = \sum_{cq' \in CQ'} sim_{cos}(f(q), f(cq'))sim_{BM25}(q, q'), \quad (1)$$

where sim_{cos} is the cosine similarity between the keyphrase vectors and sim_{BM25} is the similarity between q and q' . This gives higher importance to keyphrases belonging to more similar questions.

3.3 Learning Method

We operationalize the Similar Questions Model in a variety of ways:

SimQ Majority. We obtain a simple baseline that classifies q according to the most common label of the similar questions in Q' .

SimQ Threshold. We test the scoring features in group (iii) using a threshold classifier where a threshold γ is learned on a held-out dataset. The label is then obtained as follows:

$$\hat{y} = \begin{cases} 0, & \text{if } feat(q, CQ') \geq \gamma \\ 1, & \text{otherwise,} \end{cases}$$

where $feat(q, CQ')$ is the value of the corresponding feature, 0 refers the clear class and 1 refers to the unclear class.

SimQ ML. All features of the Similar Questions Model are combined and provided as input data to a machine learning classifier.

4 Experimental Setup

This section describes our experimental setup including the dataset and methods.

4.1 Dataset Creation

The Stack Exchange CQA platform depicts a suitable data source for our experiments. It is a network of specialized communities with topics varying from programming to Unix administration, mathematics and cooking. A frequent data dump is published consisting of all questions, answers and comments submitted

Table 3. Dataset statistics and class distribution. N is the number of samples, L the median sample length in tokens, and $|V|$ the vocabulary length after tokenization. $|V^*|$ is the vocabulary length with an imposed minimum term-document frequency of 3.

Community	N	L	$ V $	$ V^* $	Clear	Unclear
Stack Overflow	5,859,667	159	8,939,498	1,319,587	35%	65%
Super User	121,998	121	206,249	45,432	33%	67%
Ask Ubuntu	77,712	114	188,476	40,309	27%	73%
Unix & Linux	44,936	133	162,805	31,852	27%	73%
Cross Validated	38,488	157	130,691	24,229	18%	82%

to the site. For any post, a time-stamped revision history is included. We use this dump² to create a labeled dataset consisting of clear and unclear questions.

To obtain unclear questions, we apply a heuristic that has been used in previous research to find clarification questions [3, 14]. A question is considered to be unclear when there is a comment by a different user than the original asker and that comment contains a sentence ending with a question mark. This heuristic is not perfect as it will inevitably miss clarification requests not formulated as a question (e.g., “*Please post your code.*”), while it retains rhetorical questions (e.g., “*Is this a real question?*”). We only keep those questions where the original asker has provided a clarification in form of a comment or question edit.

In order to gather clear questions, we extend the described heuristic as follows. A question is considered to be clear if it has neither edits, nor comments, but it has an accepted answer. An answer can be manually accepted by the question asker if they consider it to adequately answer their question. Again, this heuristic may introduce noise: an answer can make certain assumptions that would have ideally been asked as a clarification question instead of included in the answer itself (e.g., “*Provided you are on system X, the solution is Y*”).

We apply this heuristic to five Stack Exchange communities, each of a different size and with a different domain. The communities considered are Stack Overflow, Ask Ubuntu, Cross Validated, Unix & Linux and Super User, thus covering a broad range of topics. Table 3 summarizes the statistics of each dataset. The text has been preprocessed by combining the question title, body and tags into a single field, replacing URLs with a special token, converting every character to lower-case and removing special characters except for punctuation. Token boundaries are denoted by the remaining punctuation and whitespace. Furthermore, a minimum term-document frequency of 3 is imposed to prevent overfitting.

4.2 System Components

Obtaining Similar Questions. A general purpose search engine, Elasticsearch, is used with the BM25 retrieval model in order to obtain similar questions. The retrieval score is used as $sim_{BM25}(q, q')$ during feature computation.

² Available at <https://archive.org/details/stackexchange>.

We only index the training set of each community but retrieve similar questions for the entire dataset. Queries are constructed by combining the title and tags of a question. These queries are generally short (averaging 13 tokens).³ To ensure efficient querying, we remove stopwords from all queries. Finally, BM25 parameters are set to common defaults ($k_1 = 1.2$, $b = 0.75$) [10].

Extracting Keyphrases. Keyphrases are extracted from clarification questions using the RAKE algorithm [16], which is an efficient way to find noun phrases. This algorithm has been used in a similar setting where CQA comments should be matched to related questions [12]. We tokenize the keyphrases and consider each token individually.

Similar Questions Classifier. Besides applying a threshold-based classifier on a selected set of features presented in Table 2, all features are combined to train a logistic regression classifier with L2 regularization (referred to as SimQ ML). The regularization strength is set to $C = 1$ which has been found to work well for all communities. All features are standardized by removing the mean and scaling to unit variance.

4.3 Baseline Models

The Similar Questions Model is compared with a number of baselines and state-of-the-art text classification approaches:

- Random: produce predictions uniformly at random.
- Majority: always predict the majority class (here: *unclear*).
- Bag-of-words logistic regression (BoW LR).
- Convolutional neural network (CNN) [7].

Within the BoW LR model, a question is represented as a vector of TF-IDF weighted n -gram frequencies. Intuitively, this approach captures question clarity on a phrase and topic level. We report model performances for unigrams ($n = 1$) and unigrams combined with phrases of length up to $n = 3$. Using 5-fold cross-validation on the training data, we find that an L2 regularization strength of $C = 1$ works best for all communities. With respect to the CNN model, we use the static architecture variant presented in [7] consisting of a single convolutional layer, followed by a fully connected layer with dropout. Model hyperparameters (number of filters, their size, learning rate and dropout) are optimized per community using a development set.⁴ The network is trained with the Adam optimizer [8], a mini-batch size of 64 and early stopping. We train 300-dimensional word embeddings for each community using word2vec [11] and limit a question to its first 400 tokens (with optional padding). Out-of-vocabulary words are replaced by a special token. There are several other possible neural architectures, but an exploration of those is outside the scope of this paper.

³ We experimented with longer queries that include 100 question body tokens. While computationally more expensive, model performance remained largely unaffected.

⁴ Optimal CNN parameter settings can be found in the online appendix of this paper.

4.4 Evaluation

As the data is imbalanced, we evaluate according to the F1 score of the unclear (positive) class and the ROC AUC score. We argue that it is most important to optimize these metrics based on the envisioned use case. When the classification outcome is used as a quality guard in a user interface, it is less severe to consider a supposedly clear question as unclear as opposed to entirely missing an unclear question. We randomly divide the data for each community into 80% training and 20% testing splits. Of the training set, we use 20% of the instances for hyperparameter tuning and optimize for ROC AUC. We experimented with several class balancing methods, but the classification models were not impacted negatively by the (slight) imbalance. Statistical significance is tested using an approximate randomization test. We mark improvements with Δ ($p < 0.05$) or \blacktriangle ($p < 0.01$), deteriorations with ∇ ($p < 0.05$) or \blacktriangledown ($p < 0.01$), and no significance by \circ .

5 Results and Analysis

This section presents and discusses our experimental results.

5.1 Results

The traditional BoW LR model provides a strong baseline across all communities that outperforms both the random and majority baselines (see Table 5). The generic CNN architecture proposed in [7] does not provide any significant improvements over the BoW LR model. This suggests that a more task-specific architecture may be needed to capture the underlying problem.

Table 4. Results for unclear question detection. The metrics are summarized over the five datasets using both micro-averaging and macro-averaging. F1, precision and recall are reported for the unclear class. Best scores for each metric are in boldface.

Method	Micro-average				Macro-average			
	Acc.	F1	Prec.	Rec.	Acc.	F1	Prec.	Rec.
Random	0.499	0.564	0.649	0.499	0.497	0.586	0.714	0.499
Majority	0.649	0.787	0.649	1.000	0.719	0.835	0.719	1.000
BoW LR ($n = 1$)	0.687	0.786	0.706	0.886	0.736	0.833	0.752	0.933
BoW LR ($n = 3$)	0.699	0.791	0.720	0.877	0.741	0.837	0.752	0.944
CNN	0.699	0.794	0.715	0.893	0.739	0.836	0.749	0.947
SimQ Models								
SimQ Majority	0.566	0.673	0.659	0.688	0.676	0.780	0.739	0.826
CQ Global	0.594	0.727	0.645	0.833	0.626	0.753	0.713	0.803
CQ Individual	0.586	0.721	0.640	0.824	0.632	0.761	0.710	0.824
CQ Weighted	0.604	0.737	0.648	0.855	0.642	0.770	0.716	0.838
SimQ ML	0.673	0.781	0.690	0.902	0.728	0.833	0.736	0.960

Table 5. Model performance for a selected set of communities. F1 scores are reported for the unclear class. Significance for model in line $i > 1$ is tested against line $i - 1$. Additionally, significance of each SimQ model is tested against the BoW LR ($n = 3$) model (second marker).

Method	Cross validated			Super user			Stack overflow		
	Acc.	AUC	F1	Acc.	AUC	F1	Acc.	AUC	F1
Random	0.493	0.500	0.618	0.502	0.500	0.575	0.499	0.500	0.563
Majority	0.818 [▲]	0.500 [◦]	0.900[▲]	0.669 [▲]	0.500 [◦]	0.802 [▲]	0.646 [▲]	0.500 [◦]	0.785 [▲]
BoW LR ($n = 1$)	0.819[◦]	0.647 [▲]	0.900[◦]	0.702 [▲]	0.720 [▲]	0.798 [◦]	0.685 [▲]	0.693 [▲]	0.784 [▼]
BoW LR ($n = 3$)	0.818 [◦]	0.659[▲]	0.900[◦]	0.709[▲]	0.731[◦]	0.807[▲]	0.697[▲]	0.718 [▲]	0.788 [▲]
CNN	0.817 [◦]	0.626 [◦]	0.899 [◦]	0.704 [▼]	0.715 [◦]	0.803 [▼]	0.697[◦]	0.720[▲]	0.792[▲]
SimQ Models									
SimQ Majority	0.796 [▼]	0.584 [▼]	0.883 [▼]	0.639 [▼]	0.616 [▼]	0.738 [▼]	0.561 [▼]	0.515 [▼]	0.667 [▼]
CQ Global	0.718 ^{▼▼}	0.515 ^{▼◦}	0.830 ^{▼▼}	0.598 ^{▼▼}	0.536 ^{▼▼}	0.733 ^{◦▼}	0.592 ^{▲▼}	0.520 ^{▲▼}	0.725 ^{▲▼}
CQ Individual	0.713 ^{◦▼}	0.513 ^{◦◦}	0.827 ^{◦▼}	0.591 ^{▼▼}	0.549 ^{▲▼}	0.728 ^{▼▼}	0.584 ^{▼▼}	0.528 ^{▲▼}	0.719 ^{▼▼}
CQ Weighted	0.696 ^{▼▼}	0.496 ^{▼◦}	0.812 ^{▼▼}	0.602 ^{▲▼}	0.534 ^{▼▼}	0.739 ^{▲▼}	0.603 ^{▲▼}	0.503 ^{▼▼}	0.736 ^{▲▼}
SimQ ML	0.819^{▲◦}	0.631 ^{◦◦}	0.900^{▲◦}	0.687 ^{▲▼}	0.671 ^{▲▼}	0.798 ^{▲▼}	0.670 ^{▲▼}	0.666 ^{▲▼}	0.779 ^{▲▼}

We make several observations with respect to the Similar Questions Model. First, a majority vote among the labels of the top $k = 10$ similar questions (SimQ Majority) consistently provides a significant improvement over the random baseline for all datasets (see Table 5). This simplistic model shows that the underlying concept of the Similar Questions Model is promising. Second, the scoring features that take clarification questions into consideration do not work well in isolation (see models prefixed with CQ in Tables 4 and 5). The assumption that one can test for the presence of keyphrases without considering spelling variations or synonyms seems too strong. For example, the phrase “*operating system*” does not match sentences such as “*my OS is X*” and thus results in false positives. Finally, the SimQ ML model outperforms both the random and majority baselines, and has comparable performance with the BoW LR model. It is to be emphasized that the SimQ ML model, in addition to classifying a question as clear or unclear, generates several valuable hints about the aspects of a question that may be unclear or missing (see demonstration in Table 7). This information is essential when realizing the envisioned user interface presented in Fig. 1, and cannot be deduced from the BoW LR or CNN models.

5.2 Feature Analysis

To gain further insights about the performance of the Similar Questions Model, we analyze the features and their predictive power. Features considering the stylistic properties of a question itself such as the length, readability and whether or not the question contains a question mark, are among the top scoring features (see Table 6). Other important features include the distribution of labels among the similar questions and their retrieval scores (*LenUnclear*, *LenClear*, *SimSum*, *Fraction*). With respect to the bag-of-words classifier, we observe that certain question topics have attracted more unclear questions. For example, a question about *windows 10* is more likely to be unclear than a question about *emacs*. Interestingly, also stylistic features are captured (e.g., a “?” token

Table 6. Subsets of learned coefficients for the Similar Questions Model (Left) and BoW LR classifier (Right). Both are trained on the Super User community. Positive numbers are indicative for the unclear class and negative values for the clear class.

SimQ ML		BoW LR ($n = 3$)	
Coef	Feature	Coef	Feature
+0.269	$Len(q)$	+4.321	windows 10
+0.166	$CQIndividual(q, CQ')$	+3.956	<URL>
+0.146	$LenUnclear(Q') (k = 50)$	+3.229	problem
+0.120	$LenSim(Q') (k = 20)$	+2.413	nothing
+0.094	$SimSum(q, Q')$	+2.150	help me
+0.081	$Readability(q)$	+1.760	unable to
-0.091	$ContainsPre(q)$	-1.488	documentation
-0.114	$CQIndividual(q, CQ')$	-1.742	difference between
-0.150	$LenClear(Q') (k = 50)$	-1.822	can i
-0.150	$Fraction(Q') (k = 50)$	-1.841	emacs
-0.171	$ContainsQuest(q)$	-3.026	vista
-0.196	$SimMax(q, Q')$	-5.306	?

and the special URL token). Finally, this model reveals characteristics of well-written, clear questions. For example, if a user articulates their problem in the form of “*difference between X and Y,*” such a question is more likely to belong to the clear class. This suggests that it may be beneficial to include phrase-level features in the Similar Questions Model to improve performance.

5.3 Error Analysis and Limitations

The feature analysis above reveals a problem which is common to both the Similar Questions Model and the traditional BoW LR model. Both models suffer from a topic bias. For example, a question about *emacs* is more likely to be classified as clear because the majority of *emacs* questions are clear. Furthermore, stylistic features can be misleading. Consider a post on Stack Overflow that contains an error message. This post does not require an explicit use of a question mark as the implied question most likely is “*How can this error message be fixed?*”. It is conceivable to design features that take such issues into account.

A potential limitation of the Similar Questions Model is its reliance on the existence of similar questions within a CQA website. It is unclear how the model would perform in the absence of such questions. It would make an interesting experiment to process a CQA dataset in chronological order, and measure how the model’s effectiveness changes as more similar questions become available over time. However, we leave the exploration of this idea to future work.

Table 7. Example questions and their clarification questions retrieved by Similar Questions Model. The numbers in parenthesis indicate retrieval score $sim_{BM25}(p, q_i)$. Highlighted text corresponds to the keyphrases extracted by RAKE.

Field	Text
Title	Laptop randomly going in hibernate
Body	I have an Asus ROG G751JT laptop, and a few days ago my battery has died. The problem that I am encountering is that my laptop randomly goes to sleep after a few minutes of use even when plugged in [...].
ClarQ	(20.01) Does this happen if you boot instead from an ubuntu liveusb ? (17.92) Did you enable allow wake timers in power options sleep ? (16.88) Can you pop the battery out of the mouse ? (16.64) Which OS are you using? (16.02) Have you scanned your system for malwares ?
Title	Does ZFS make sense as local storage?
Body	I was reading about ZFS and for a moment thought of using it in my computer, but than reading about its memory requirements I thought twice. Does it make sense to use ZFS as local or only for servers used as storage?
ClarQ	(36.11) What's wrong with more redundancy ? (31.41) What kind of data are you trying to protect ? (30.77) How are you planning on doing backups and or disaster recovery ? (29.70) Is SSD large enough?

6 Conclusion

The paper represents the first study on the challenging task of detecting unclear questions on CQA websites. We have constructed a novel dataset and proposed a classification method that takes the characteristics of similar questions into account. This approach encodes the intuition that question aspects which have been missing or found to be unclear in previous questions, may also be unclear in a given new question. We have performed a comparison against traditional text classification methods. Our main finding is that the Similar Questions Model provides a viable alternative to these models, with the added benefit of generating cues as to why a question may be unclear; information that is hard to extract from traditional methods but that is crucial for supportive question formulation interfaces.

Future work on this task may combine traditional text classification approaches with the Similar Questions Model to unify the benefits of both. Furthermore, one may start integrating the outputs of the Similar Questions Model into a clarification question generation system, which at a later stage is embedded in the user interface of a CQA site. As an intermediate step, it would be important to evaluate the usefulness of the generated cues as to why a question is unclear. Finally, the work by Rao and Daumé III [14] provides a natural extension, by ranking the generated clarification questions in terms of their expected utility.

Acknowledgments. We would like to thank Dolf Trieschnigg and Djoerd Hiemstra for their insightful comments on this paper. This work was partially funded by the University of Twente Tech4People Datagrants project.

References

1. Arora, P., Ganguly, D., Jones, G.J.F.: The good, the bad and their kins: identifying questions with negative scores in stackoverflow. In: Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015, ASONAM 2015, pp. 1232–1239. ACM, New York (2015)
2. Asaduzzaman, M., Mashiyat, A.S., Roy, C.K., Schneider, K.A.: Answering questions about unanswered questions of stack overflow. In: Proceedings of the 10th Working Conference on Mining Software Repositories, MSR 2013, pp. 97–100. IEEE Press, Piscataway (2013)
3. Braslavski, P., Savenkov, D., Agichtein, E., Dubatovka, A.: What do you mean exactly?: analyzing clarification questions in cqa. In: Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval, CHIIR 2017, pp. 345–348. ACM, New York (2017)
4. Coleman, M., Liau, T.L.: A computer readability formula designed for machine scoring. *J. Appl. Psychol.* **60**(2), 283 (1975)
5. Correa, D., Sureka, A.: Chaff from the wheat: characterization and modeling of deleted questions on stack overflow. In: Proceedings of the 23rd International Conference on World Wide Web, WWW 2014, pp. 631–642. ACM, New York (2014)
6. Kato, M.P., White, R.W., Teevan, J., Dumais, S.T.: Clarifications and question specificity in synchronous social q&a. In: CHI 2013 Extended Abstracts on Human Factors in Computing Systems, CHI EA 2013, pp. 913–918. ACM, New York (2013)
7. Kim, Y.: Convolutional neural networks for sentence classification. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, 25–29 October 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL, pp. 1746–1751 (2014)
8. Kingma, D.P., Ba, J.L.: Adam: a method for stochastic optimization. In: Proceedings of the 3rd International Conference on Learning Representations (ICLR) (2015)
9. Li, B., Jin, T., Lyu, M.R., King, I., Mak, B.: Analyzing and predicting question quality in community question answering services. In: Proceedings of the 21st International Conference on World Wide Web, WWW 2012 Companion, pp. 775–782. ACM, New York (2012)
10. Manning, C.D., Raghavan, P., Schütze, H.: Introduction to Information Retrieval. Cambridge University Press, New York (2008)
11. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Proceedings of the 26th International Conference on Neural Information Processing Systems, NIPS 2013, vol. 2, pp. 3111–3119. Curran Associates Inc., USA (2013)
12. Nandi, T., et al.: IIT-UHH at SemEval-2017 Task 3: exploring multiple features for community question answering and implicit dialogue identification. In: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017), pp. 90–97. Association for Computational Linguistics (2017)

13. Ponzanelli, L., Mocci, A., Bacchelli, A., Lanza, M.: Understanding and classifying the quality of technical forum questions. In: Proceedings of the 2014 14th International Conference on Quality Software, QSIC 2014, pp. 343–352. IEEE Computer Society, Washington (2014)
14. Rao, S., Daumé III, H.: Learning to ask good questions: ranking clarification questions using neural expected value of perfect information. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 2737–2746. Association for Computational Linguistics (2018)
15. Ravi, S., Pang, B., Rastagori, V., Kumar, R.: Great question! question quality in community q&a. *Int. AAAI Conf. Weblogs Social Media* **1**, 426–435 (2014)
16. Rose, S., Engel, D., Cramer, N., Cowley, W.: Automatic keyword extraction from individual documents. *Text Min. Appl. Theory*, 1–20 (2010)
17. Srba, I., Bielikova, M.: A comprehensive survey and classification of approaches for community question answering. *ACM Trans. Web* **10**(3), 18:1–18:63 (2016). ISSN 1559–1131
18. Tausczik, Y.R., Pennebaker, J.W.: Predicting the perceived quality of online mathematics contributions from users’ reputations. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2011, pp. 1885–1888. ACM, New York (2011)
19. Yang, J., Hauff, C., Bozzon, A., Houben, G.-J.: Asking the right question in collaborative q&a systems. In: Proceedings of the 25th ACM Conference on Hypertext and Social Media, HT 2014, pp. 179–189. ACM, New York (2014)