

# Chaff from the Wheat : Characterization and Modeling of Deleted Questions on Stack Overflow

Denzil Correa, Ashish Sureka  
Indraprastha Institute of Information Technology  
IIIT-Delhi  
{denzilc, ashish} @iiitd.ac.in

## ABSTRACT

Stack Overflow is the most popular Community based Question Answering (CQA) website for programmers on the web with 2.05M users, 5.1M questions and 9.4M answers. Stack Overflow has explicit, detailed guidelines on *how to post* questions and an ebullient moderation community. Despite these precise communications and safeguards, questions posted on Stack Overflow can be extremely off topic or very poor in quality. Such questions can be *deleted* from Stack Overflow at the discretion of experienced community members and moderators. We present the first study of *deleted questions* on Stack Overflow. We divide our study into two parts – (i) Characterization of *deleted questions* over  $\approx 5$  years (2008-2013) of data, (ii) Prediction of *deletion* at the time of question creation.

Our characterization study reveals multiple insights on question deletion phenomena. We find that it takes substantial time to vote a question to be deleted but once voted, the community takes swift action. We also see that question authors delete their questions to salvage reputation points. We notice some instances of accidental deletion of good quality questions but such questions are voted back to be undeleted quickly. We discover a pyramidal structure of question quality on Stack Overflow and find that *deleted questions* lie at the bottom (lowest quality) of the pyramid. We also build a predictive model to detect the *deletion* of question at the creation time. We experiment with 47 features – based on *User Profile*, *Community Generated*, *Question Content* and *Syntactic* style – and report an accuracy of 66%. Our findings reveal important suggestions for content quality maintenance on community based question answering websites. To the best of our knowledge, this is the first large scale study on poor quality (*deleted*) questions on Stack Overflow.

## Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Information filtering; H.3.5 [Online Information Services]: Web-

based services; K.4.3 [Organizational Impacts]: Computer-supported collaborative work

## Keywords

Question-answering, Question quality, Stack Overflow

## 1. INTRODUCTION

### 1.1 Research Motivation and Aim

There has been an increase in the presence of Community Based Question Answering (CQA) website services like Yahoo! Answers, Ask.com and Quora on the Internet. Stack Exchange is a growing network of thematic question-answering websites with each website dedicated to a specific field of expertise [15]. It consists of 107 CQA websites with 4.1M users, 7.3M questions, 13.2 M answers and 8.5M visits per day.<sup>1</sup> CQA websites on Stack Exchange span across different orthogonal themes like *Technology* (Web Applications, Game Development), *Culture* (Travel, Christianity), *Arts* (Photograph, Scientific Fiction) and *Sciences* (Mathematics, Physics). Stack Overflow is a programming based CQA and the most popular Stack Exchange website consisting of 5.1M questions, 9.4M answers and 2.05 registered users on its website.<sup>2</sup> Stack Overflow has detailed, explicit guidelines on posting questions and it maintains a firm emphasis on following a question-answer format. The community strongly discourages questions which could generate chit-chat, opinions, polls etc. and employs elected moderators to ensure content quality maintenance. Stack Overflow is a free, open (no registration required) website to all users on the Internet and hence, it is a necessity to maintain quality of content on the website [4]. Stack Overflow is driven by the goal to be an exhaustive knowledge base on programming related topics and hence, the community would like to ensure minimal possible *noise* on the website. However, despite of the presence of question posting guidelines and an ebullient moderation community, a significant percentage of questions on Stack Overflow are extremely poor in nature. Questions are a fundamental aspect of any CQA website and the presence of poor quality content may affect user experience. Prior work also shows that poor quality content on a CQA website may drive users away from the platform and adversely affect the traffic [19]. Therefore, it is important to study poor quality questions and develop mechanisms to minimize them on the website. In this work, we focus our attention on

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media.  
WWW'14, April 7–11, 2014, Seoul, Korea.  
ACM 978-1-4503-2744-2/14/04.  
<http://dx.doi.org/10.1145/2566486.2568036>.

<sup>1</sup><https://stackexchange.com/about>

<sup>2</sup><http://stackoverflow.com/>

*deleted* questions (poor quality) on Stack Overflow. Questions which are very poor in quality or extremely off topic in nature are deleted from the website. Table 1 shows examples of deleted questions on Stack Overflow. We see that most of these questions are very poor in quality and of little worth to the community.

**Table 1: shows examples of deleted questions on Stack Overflow**

Id	Title	Score
1644242	Get drive information (free space, etc.) for drives on Windows and populate a memo box	-50
2022741	plzz can u help me	-9
14771464	NDTV iPad app screen design	-8
2351964	I need to see this q fast ..pleash	-8
3077356	PostgreSQL stupidity	-8
2984348	hi question about mathematics	-3

Our research aim is to understand the phenomena of *deleted* questions on Stack Overflow to gain insights about the nature of poor quality questions. In addition, we also develop a predictive framework to detect the probability of a question to be deleted at the time of question creation. Such a predictive framework would help the moderators of the Stack Overflow community to detect a potential *deleted* question on Stack Overflow. However, a deleted question on Stack Overflow is an explicit feedback to the question asker that her question does not conform to the guidelines. A predictive framework can convey immediate feedback to the question asker about her question. This may help her to revise her question and improve it in order to avoid deletion. Therefore, prediction of a deleted question at post creation time has two distinct benefits – (i) An immediate feedback mechanism to the question asker and (ii) A complementary mechanism to aid community moderators in their daily moderation tasks.

## 1.2 Research Contributions

We perform the first large scale study on poor quality or *deleted* questions on Stack Overflow. We make the following research contributions -

- We analyze deleted questions on Stack Overflow posted over  $\approx 5$  years and conduct a characterization study. We perform a longitudinal study of deleted questions, community voting patterns, deletion behavior by question owners and discover question quality structure. We make our data on *deleted questions* publicly available for research purposes.
- We develop a predictive model using a machine learning framework to detect a deleted question at the time of question creation. We experiment with 47 features based on four diverse categories and report 66% accurate predictions. We also perform analysis of our feature space and report features with best discriminatory powers.

## 2. RELATED WORK

Stack Overflow is a collaborative question answering Stack Exchange website. The underlying theme of Stack Overflow is programming-related topics and the target audience are software developers, maintenance professionals and programmers. Apart from existing as a question-answering

website, the objective of Stack Overflow is to be a comprehensive knowledge base of programming topics. Therefore, Stack Overflow has attracted increasing attention from different research communities like software engineering, human computer interaction, social computing and data mining [6, 9, 10, 21, 22]. Researchers have mined the Stack Overflow knowledge-base to extract interesting and unique insights like API usage obstacles, innovation diffusion via URL link sharing, mobile development issues and programming topic trends [9, 13, 18, 25]. Questions on Stack Overflow have received focused attention from researchers.

Nasehi *et al.* analyze questions on Stack Overflow to understand the quality of a code example [20]. They find nine attributes of good questions like concise code, links to extra resources and inline documentation. Wang and Godfrey analyze iOS and Android developer questions on Stack Overflow to detect API usage obstacles [25]. They used topic models to find a set API classes on iOS and Android documentation which were difficult for developers to understand. Asaduzzaman *et al.* analyze unanswered questions on Stack Overflow and use a machine learning classifier to predict such questions [7]. They observe certain characteristics of unanswered questions which include vagueness, homework questions etc. Allamanis and Sutton perform a topic modeling analysis on Stack Overflow questions to combine topics, types and code [5]. They find that programming languages are a mixture of concepts and questions on Stack Overflow are concerned with the code example rather than the application domain. In contrast to the aforementioned work, our work specifically focuses on quality of content on Stack Overflow.

Measurement of answer quality on CQA has received significant attention using information retrieval models and techniques. Jeon *et al.* use non-textual question features using a maximum entropy classifier to predict answer quality on Naver, a Korean CQA website [16]. Agichtein *et al.* use user relationships, question metadata and content based data to find high quality content on Yahoo! Answers [4]. Shah *et al.* build a classifier to predict answer quality on Yahoo! Answers CQA with the usage of question and answer based textual features [23]. Most of the previous work focuses on answer quality on large scale CQA websites. Prior work reveals that answer quality is a direct function of question quality [4]. Poor quality questions may adversely affect the user experience on the website and therefore, it is important to study such questions [19]. To this end, Li *et al.* study characteristics of question quality in Yahoo! Answers and find that a Mutual Reinforcement Label Propagation approach based on question plus answer features yield good results [17]. Correa and Sureka analyze and predict ‘closed’ questions on Stack Overflow viz. questions which are irrelevant or unfit to the Stack Overflow format [11]. In contrast to these previous studies, our current focus is on analysis and prediction of deleted or poor quality questions on Stack Overflow (a programming related CQA). Deleted questions are extremely poor, off-topic in nature and therefore, cease to exist from the website. The properties of deleted questions are different in both topic and content. To the best of our knowledge, this is the first work which studies poor quality questions on a large-scale CQA website like Stack Overflow.

### 3. DELETED QUESTIONS ON STACK OVERFLOW

In this section, we briefly discuss about deleted questions on Stack Overflow. Figure 1 summarizes various details about deleted questions on Stack Overflow.

#### Why are questions deleted?

The goal of Stack Overflow is to be the most extensive knowledge base of programming related topics. Hence, it would like to keep the *noise* on their website as low as possible. Therefore, questions on Stack Overflow which are extremely off topic or very poor in quality are deleted from the website [1]. In addition, questions which are dormant viz. have no activity over a significant period of time are also deleted. Also, ‘closed’ questions (questions which are deemed unfit for Stack Overflow) which do not serve as useful sign points may also be deleted. The *Why* block of Figure 1 corresponds to this section.

#### Who can delete a question?

Experienced users with 10,000+ reputation points can cast ‘delete votes’ in order to delete a question. Question authors can delete their own questions if – (i) the question has zero answers OR (ii) the question has only one answer but no up votes. However, a community elected moderators (super users of the website) can delete any questions at their discretion [2]. The *Who* block of Figure 1 corresponds to this section.

#### How are questions deleted?

The number of votes required to delete a question scales to the number of answers and the up votes on the answers of those questions. This notwithstanding, a minimum of 3-votes and a maximum of 10-votes are required to delete a question. It must be noted that a deleted question is not physically deleted from the website but *soft deleted*. Moderators, developers and experienced users (10000+ reputation) points are able to view these questions. However, deleted questions do not appear in search results [1]. The *How* block of Figure 1 corresponds to this section.

#### What happens after a question is deleted?

Once a question is deleted there are two possibilities – (i) it remains deleted and (ii) it is *undeleted*. The procedure for *undeleting* questions is similar to that of deleting a question. The *What* block of Figure 1 corresponds to this section.

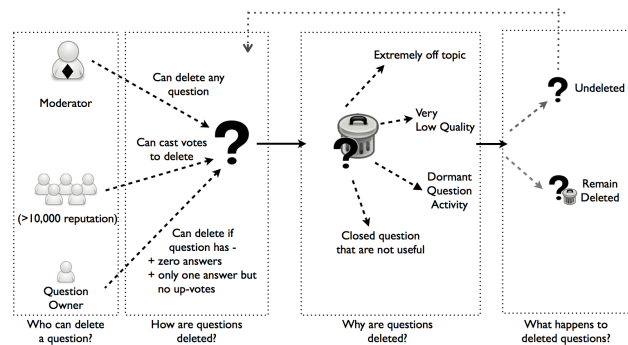


Figure 1: shows the details about procedures and community guidelines to delete a question on Stack Overflow.

### 4. CHARACTERIZATION OF DELETED QUESTIONS

In this section, we present our findings on deleted questions on Stack Overflow.

#### 4.1 Dataset Description

Stack Overflow provides a periodic database dump of all user-generated content under the *Creative Commons Attribute-ShareAlike* [8]. The database dump contains publicly available information of questions, answers, comments, votes and badges from the genesis of Stack Overflow (August 2008) to the release time of the dump. Table 2 shows the months in which Stack Overflow provided database dumps between August 2008 to June 2013.

Table 2: shows the months in which Stack Overflow provided database dumps between August 2008 to June 2013.

Months	2009	2010	2011	2012	2013
January		✓	✓		
February		✓			
March		✓			✓
April		✓	✓	✓	
May		✓			
June			✓		✓
July		✓			
August	✓	✓		✓	
September	✓	✓	✓		–
October	✓	✓			–
November	✓	✓			–
December	✓		✓		–
Total	24 Database Snapshots				
Information	Questions, Answers, Comments, Votes and Badges				

We download all the available 24 database dumps for our study. Table 3 shows the overall statistics of user-generated content on Stack Overflow between August 2008 (inception) to June 2013 (current). The statistics show that Stack Overflow is a very popular programming CQA with 5.1M questions, 9.4M answers and 2.05M registered users.




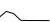


Table 3: overall statistics of user-generated content on Stack Overflow between August 2008 – June 2013

Users	2.05M (936k askers, 630k answerers)
Questions	5.1M (60.22% with accepted answers)
Answers	9.4M (32.67% marked as accepted)
Votes	42.3M (70.5% +ve, 6.7% favorites)
Ratio of Answers to Questions	1.84

However, the database dumps provided by Stack Overflow do not directly contain information about deleted questions. Hence, we analyze the entire 24 database snapshots over  $\approx 5$  years to construct our dataset. Concretely, questions available in the earlier database snapshots (August 2009 – March 2013) but absent in the most recent snapshot (June 2013) are *deleted* from Stack Overflow. We obtain **293,289** ( $\approx 0.29M$ ) deleted questions between September 2009 and June 2013 by following this procedure. However, we notice a small (but sizeable) percentage of questions in this set which are not deleted but wrongly captured. We eliminate such errors by inspecting HTML pages and obtain a final experimental dataset of **270,604** ( $\approx 0.27M$ ) deleted questions. Table 4 shows descriptive statistics of user-generated content of deleted questions on Stack Overflow between August 2008 to June 2013. We would like to point out that our experimental dataset contains a subset of all deleted

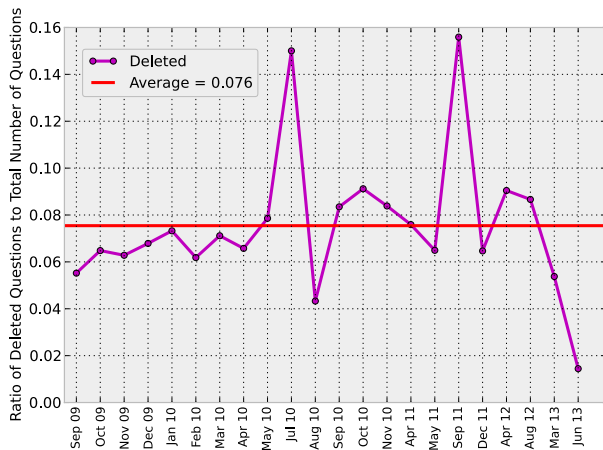
questions on Stack Overflow. This limitation is due to the sporadic sharing of database dumps by Stack Overflow (and not due to the procedure we use to find deleted questions). Questions deleted between two data snapshots would not be captured in our experimental dataset. However, it must be noted that our dataset contains the maximum possible deleted questions which can be obtained given the publicly available Stack Overflow database snapshots. We also make our experimental dataset publicly available for research purposes under the *Creative Commons Attribute-ShareAlike*.<sup>3</sup>

**Table 4:** shows statistics of user-generated content of deleted questions on Stack Overflow between August 2008 to June 2013. The distribution sparklines begin at minimum and end at their corresponding maximum value (log-scale). The distributions are generated using Gaussian kernel density estimates.

	Mean	Median	Min	Max	Distribution
Questions per year	54,120	49,221	17,514	102,623	
Answers	2.96	1.0	0.0	637.0	
Score	0.15	2.22	-56.0	695.0	
Views	221.65	80.0	1.0	296,466	
Favorites	3.59	1.0	0.0	1530.0	
Comments	2.26	1.0	0.0	77.0	

## 4.2 Increase in Deleted Questions Over Time

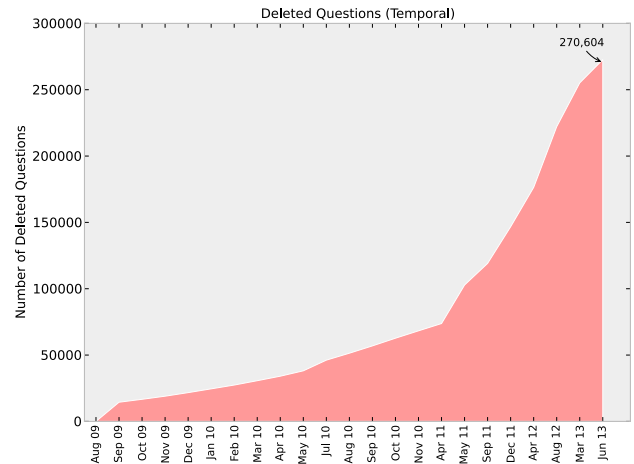
We now perform a temporal trend analysis of deleted questions on Stack Overflow. Figure 2 shows the ratio of deleted questions to total questions in each month over a 49-month period between September 2009 and June 2013. We would like to recollect that the first database snapshot provided by Stack Overflow is on August 2009 and hence, we do not have information of deleted questions prior to this snapshot. We observe that on an average  $\approx 8\%$  of questions are deleted on Stack Overflow. We also notice an anomalous increase in percentage of deleted questions  $\approx 15\%$  after the month of May (2010 2011). We posit that these abrupt rises could be due to periodic *Deletion Question Audits* conducted by the Stack Overflow community around the month of May [3].



**Figure 2:** shows the ratio of deleted questions to total questions in each month over a 49-month period between September 2009 and June 2013.

<sup>3</sup><http://correa.in/datasets.html>

Figure 3 shows the cumulative distribution area chart of deleted questions over a 49-month period between September 2009 and June 2013. The area chart depicts that there is a sharp increase in the total number of deleted questions over time (even with an average of  $\approx 8\%$ ). We notice a particularly steep increase after May 2011. Therefore, despite the presence of comprehensible and explicit question posting guidelines – Stack Overflow receives a high number of extremely poor quality questions which are not fit to exist on its website. Hence, it is important to perform a longitudinal study about deleted questions on Stack Overflow.



**Figure 3:** shows the cumulative distribution of deleted questions over a 49-month period between September 2009 and June 2013. We notice a sharp rise in the number of deleted questions over time.

## 4.3 Community Takes Long Time to Detect but Swift Action by Moderators

Stack Overflow delineates an elaborate procedure to delete a question. We recall that experienced community members viz. users with 10,000+ reputation points can cast ‘delete votes’ to delete a question. We analyze these ‘delete vote’ patterns to gain insights into the community participation dynamics on poor quality questions on Stack Overflow. Hence in this section, we restrict our analysis to **62,949** deleted questions which have received ‘deleted votes’. Figure 4 shows the distribution of time taken to receive the first ‘delete vote’ on a deleted question on Stack Overflow. We see that  $\approx 80\%$  questions take at least 1 month (or more) to receive its first ‘delete vote’ while  $\approx 50\%$  of the questions take 6 months (or more). Overall, 50 percentile of the questions take  $>164$  days to receive their first ‘delete vote’. Therefore, majority of deleted questions take a significant amount of time to receive its first delete vote.

We now analyze the distribution of ‘deleted votes’ in our dataset. Table 5 shows the ‘deleted vote’ distribution on deleted questions posted and Figure 5 shows the temporal distribution of ‘delete votes’ on deleted questions. We observe that  $\approx 80\%$  questions receive 1 ‘delete vote’ and  $\approx 14\%$  questions receive 3 ‘delete votes’ before they are deleted. This shows that once a question receives the required number of ‘delete vote’(s), moderators move quickly to take appropriate action. The moderators are driven by the Stack Overflow motto to keep a low signal-to-noise ratio in order

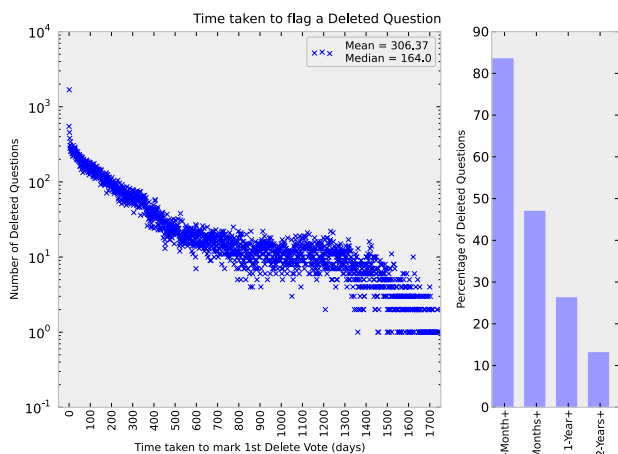


Figure 4: shows the distribution of time taken to receive the first delete vote on a deleted question on Stack Overflow.

to maintain high content quality. We also note that  $\approx 23\%$  of questions in our experimental dataset have received ‘delete votes’ i.e. questions which are voted to be deleted by experienced users. Therefore, 75% of questions in our dataset are never voted for deletion by the community. This reveals that the major duty to detect extremely poor quality questions is on the encumbrance of the moderator. This, coupled with the ‘delete vote’ pattern-action behavior (relatively long time but swift action) motivates the impending need to have a predictive system for the benefit of the moderators on Stack Overflow.

Table 5: shows the ‘deleted vote’ distribution on deleted questions posted between June 2009 and June 2013.

Votes	Deleted Questions
1-vote	50,012 (79.45%)
2-votes	2,736 (4.35 %)
3-votes	9,009 (14.3%)
4-votes	463 (0.74%)
5+-votes	729 (1.16%)
Total	62,949

#### 4.4 Authors Delete Questions to Salvage Reputation

We recall that a question on Stack Overflow can either be deleted by the author of the question or by a moderator. However, this information is not directly available in the publicly available data dumps provide by Stack Overflow. We also recall that questions on Stack Overflow are not digitally deleted i.e. they are hidden from the site and do not appear in search results. We notice that this information about question deletion initiator (question author or moderator) is available on the unique web address of the question. Figure 6 shows the web page screenshots of – (i) question deleted by moderator (left) and (ii) question deleted by author (right).

We download the unique web pages of deleted questions in our experimental dataset and employ a regular expression to extract this information. Figure 7 shows the distribution of question deletion initiator (moderator or author) on Stack Overflow. We notice that  $\approx 87\%$  (237,163) of questions are deleted by a moderator while  $\approx 12\%$  (33,330) or 1 out of 8

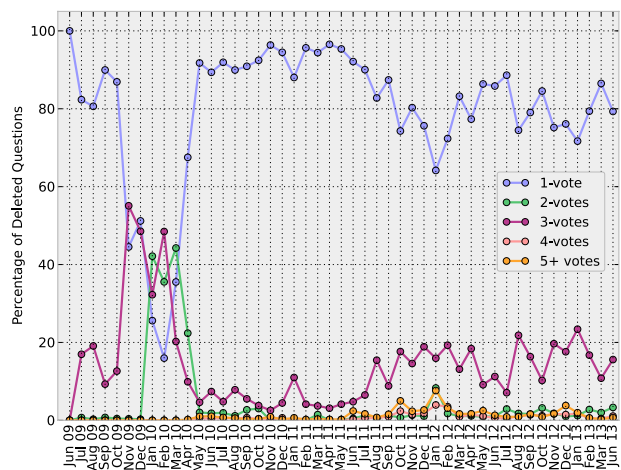


Figure 5: shows the temporal distribution of ‘delete votes’ on deleted questions in Stack Overflow between June 2009 and June 2013.



Figure 6: shows the web page screenshots of two questions deleted by moderator (left) and author (right) on Stack Overflow.

questions are deleted by the question author. A negligible percentage (0.04%) of questions (111) do not belong to either category. This may be a bug in the data dump and hence, we ignore these questions in this section.

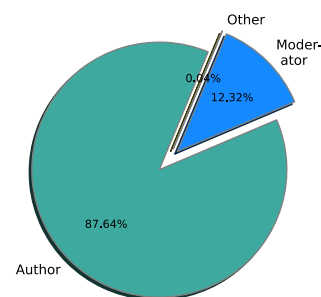
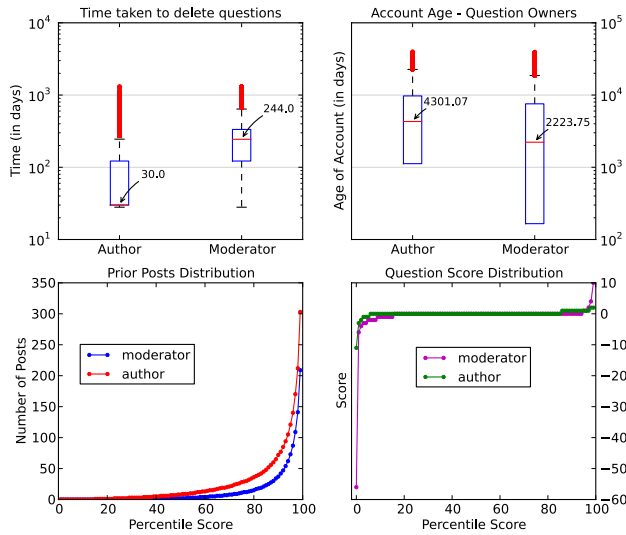


Figure 7: shows the distribution of question deletion initiator (moderator or author) on Stack Overflow

We now analyze multiple question attributes based on the question deletion initiator (author or moderator) and make observations. Figure 8 (Top-Bottom, Left-Right) shows the distributions of – (i) time taken to delete a question (box-and-whisker plot), (ii) account age of question owner (box-and-whisker plot), (iii) posts prior to deleted question cre-

ation (percentile plot) and (iv) deleted question score (percentile plot) for author and moderator deleted questions. The Top-Left box-and-whisker plot reveals that the distribution of time taken to delete a question by the question author is relatively lower than when done by a moderator. The Top-Right box-and-whisker plot shows that the question owner age of account is relatively higher of author deleted questions than moderator deleted questions. This points out that question owners of author-deleted questions are more experienced on the website (in terms of time) than question owners of moderator-deleted questions. The Bottom-Left percentile plot shows that prior posts – posts made by the question owner prior to the time of deleted question creation – of author-deleted questions is higher than those deleted by a moderator. The Bottom-Right percentile plot shows that question scores of author-deleted questions are higher than those of moderator-deleted questions at very low ( $<20$ ) and negative question scores.



**Figure 8:** shows the distributions of – (top-left) time taken to delete a question via box-and-whisker plot, (top-right) account age of question owner via box-and-whisker plot, (bottom-left) posts prior to deleted question creation via percentile plot and (bottom-right) deleted question score via percentile plot for author and moderator deleted questions.

The above observations show that despite being less experienced on the website, question owners of author-deleted questions have more prior posts and have higher question scores on deleted questions than those of moderator-deleted questions. We argue that question owners of author-deleted questions exhibit such a behavior as they want to maintain a healthy reputation earned on Stack Overflow [10]. The owners of author-deleted questions observe that their questions are attracting down votes which affects their *reputation*. In order to stem further decrease reputation points, question owners see deletion of question as a quick fix and therefore, proceed to delete the posted question in an attempt to salvage their reputation.

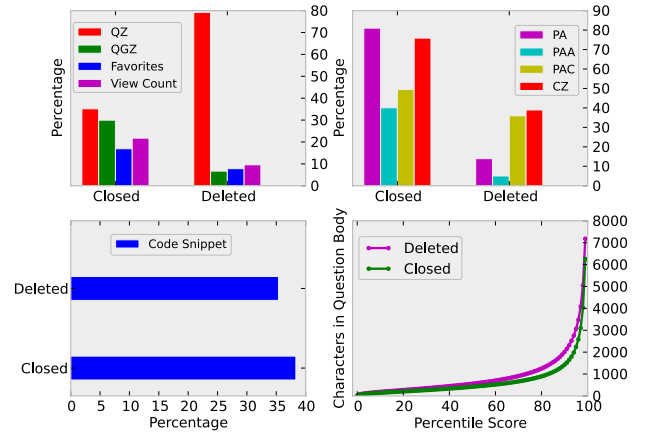
#### 4.5 Question Quality Pyramidal Structure

Questions on Stack Overflow are marked ‘closed’ if they are deemed unfit for the question-answer format on Stack Overflow and indicate low quality. A question can be marked

as ‘closed’ due to five reasons – *duplicate*, *subjective*, *off topic*, *too localized* or *not a real question* [11]. We find that there are **254,446** (0.25M) ‘closed’ questions on Stack Overflow between August 2008 June 2013. In this section, in addition to our experimental dataset we utilize this data to analyze question quality indicators for deleted and ‘closed’ questions. We also make observations about the relative quality of deleted questions in context to ‘closed’ questions on Stack Overflow.

#### Community Value.

Figure 9 shows various quantities of question quality indicators for ‘closed’ and deleted questions on Stack Overflow. We see that deleted questions have higher percentage of questions with zero *score* than ‘closed’ questions. A question *score* is the net worth of the usefulness of a question as determined by the Stack Overflow community.  $\approx 80\%$  of deleted questions have a zero *score* which indicates that most deleted questions are of very little worth to the community. We also observe that in comparison to ‘closed’ questions – deleted questions have a lower percentage of question with greater than zero *score*, *favorite* votes and *view counts*. A *favorite* vote is equivalent to an explicit expression of interest or subscription on a question. Only  $\approx 5\text{--}6\%$  of deleted questions attract a positive question *score* or *favorite* votes. This indicates that deleted questions are generally of very little worth and interest to the Stack Overflow community. In addition,  $\approx 10\%$  of deleted questions have  $>0$  *view counts* which is twice as lower than that of closed questions ( $\approx 23\%$ ). We also find that deleted questions have lesser percentage of code snippets but slightly higher number of characters in their body than ‘closed’ questions. Nonetheless, we observe that deleted questions fare inferior on most indicators in comparison to ‘closed’ questions indicating very poor quality.



**Figure 9:** shows the quantities of question quality indicators for ‘closed’ and deleted questions on Stack Overflow. (QZ=Percentage of Questions with Zero Score, QGZ=Percentage of Questions with Greater than Zero Score, PA=Percentage of Answers, PAA=Percentage of Accepted Answers, PAC=Percentage of Accepted Answers given that a question has an answer, CZ=Percentage of Comments with Zero Scores)

### Question Topics.

Stack Overflow questions contain user supplied *tags* which indicate the topic of the question. We analyze the tag distribution of closed and deleted questions and compare them to the overall tag distribution on Stack Overflow. There are a total of **36,643** tags on all questions in Stack Overflow. Figure 10 shows the venn diagram of tag distributions of questions on Stack Overflow.

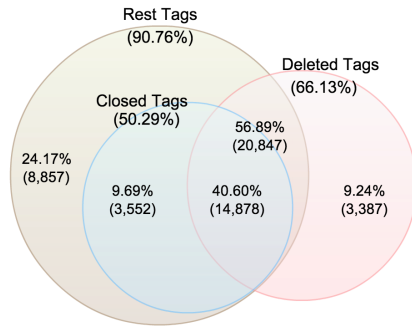


Figure 10: shows the venn diagram of tag distributions of questions on Stack Overflow.

We see that tags on ‘closed’ questions are a subset of the overall tags which occur in regular questions. In contrast, an appreciable number of tags on deleted questions ( $\approx 10\%$ ) are not found either in ‘closed’ or regular questions. We extract such tags found in deleted questions for further analysis. The topics of ‘closed’ questions are relevant to the community despite the questions themselves being unfit for the Stack Overflow format. However, some of the topics on deleted questions are extremely off topic to the interests of the community. Figure 11 shows the word cloud of the top-50 tags that occur on deleted questions. We find a very high presence of low quality tags like *homework*, *job-hunting* and *polls* on deleted questions. These topics also show that deleted questions are of extremely poor quality.



Figure 11: shows the word cloud of the top-50 tags that occur on deleted questions. 9.24% of tags are unique to deleted questions on Stack Overflow.

### Effort to Improve Question Quality.

The goal of Stack Overflow is to be a comprehensive knowledge base for programming related topics. Therefore, a question on Stack Overflow may be edited by privileged users

(experienced community members and moderators) to maintain content quality. A question on Stack Overflow has three major sections which can be edited – title, body and tags. In addition, users without sufficient privileges can *suggest edits* to questions. These suggestions are then reviewed by privileged users and are brought into effect at their discretion. In this part, we analyze edit histories of deleted questions on Stack Overflow. Table 6 shows four edit types (edit tags, edit body, edit title and suggested edits) for deleted questions in our experimental dataset. We see that  $\approx 93\%$  of deleted questions receive at least one form of edit. Also, moderator-deleted questions receive more edits than author-deleted questions.

Table 6: shows edit details of deleted questions on Stack Overflow

	Moderator	Author	Total
Deleted Questions	237,163	33,330	270,493
Question Content Edited	226,286 (95.41%)	26,461 (79.39 %)	252,747 (93.44%)
Question 'Closed'	34,209 (15.12%)	2,167 (8.19%)	36,376 (14.38%)

We also notice that 14.38% of deleted questions were marked as ‘closed’ before they were deleted – a higher percentage of moderator-deleted questions were marked as ‘closed’ than author-deleted questions. Figure 12 shows the distribution of questions marked as ‘closed’ due to five reasons for both – author and moderator deleted questions. We see that a higher percentage of author-deleted questions are marked as *too localized*, *duplicate* and *offtopic*. However, a higher percentage of moderator-deleted questions are marked as *subjective* and *not a real question*.

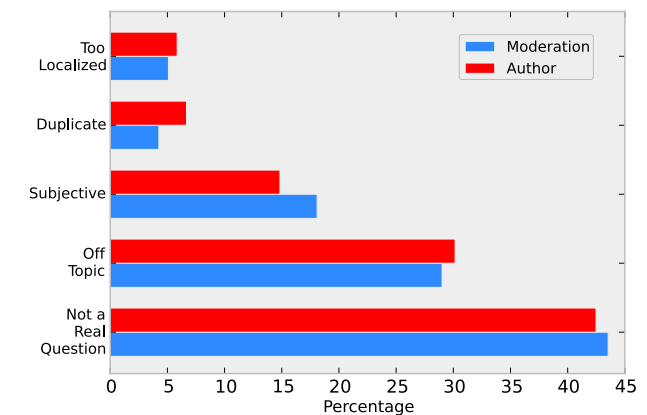


Figure 12: shows the distribution of questions marked as 'closed' on five reasons for author and moderator deleted questions.

Figure 13 shows the distribution of percentages of questions on various edit categories for (left) ‘closed’ versus deleted and (right) moderator versus author deleted questions. We observe that ‘closed’ questions receive a higher percentage of edits than deleted questions. This signifies that the community puts a greater effort to edit and improve ‘closed’ questions than it does for deleted questions. We also see that core content of a question (title and body) for author-deleted questions receive a higher percentage of edits than moderator-deleted questions. This shows that author-deleted

questions are inferior in quality than moderator-deleted questions and require more work to improve their content.

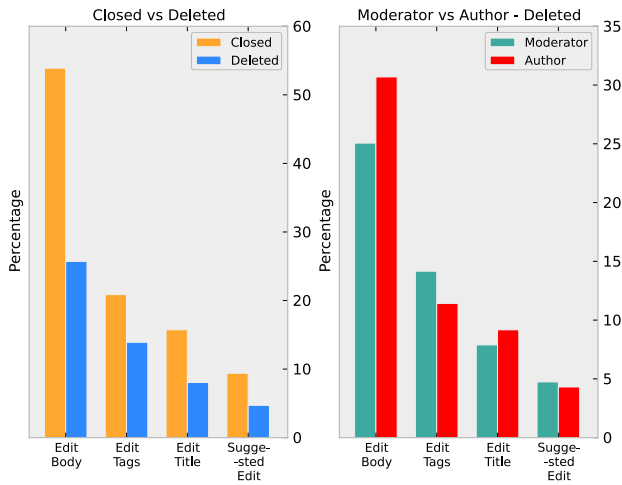


Figure 13: shows the distribution of edit question history on (left) ‘closed’ versus deleted and (right) moderator versus author deleted questions.

### Quality Pyramid.

‘Closed’ questions are questions which are deemed unfit for the Stack Overflow format. A ‘closed’ question is low quality but has the potential to be improved upon. On the other hand, deleted questions are very poor in nature and beyond improvement. Prior analysis in this section revealed that deleted questions fare extremely poor on multiple community value quality indicators. We also find that some topics of deleted questions are entirely irrelevant to the Stack Overflow website. In addition, we see that the community puts in more effort to improve a ‘closed’ question than it does for a deleted question. These findings reveal that question quality on Stack Overflow has a pyramidal structure – regular questions lie at the top of the pyramid (good quality), followed by ‘closed’ questions (bad quality) and deleted questions are placed at the bottom of the pyramid (extremely poor quality). Figure 14 shows this underlying question quality pyramid structure on Stack Overflow.

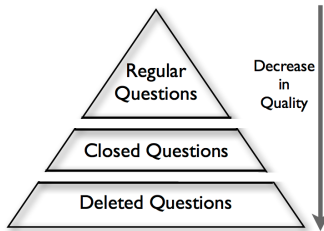


Figure 14: shows the underlying question quality pyramid structure on Stack Overflow.

## 4.6 Accidental Question Deletion

Stack Overflow provides a procedure to *undelete* a deleted question. The similar voting procedure to that of deletion

is followed to *undelete* a question. Table 7 shows some examples of undeleted questions on Stack Overflow.

Table 7: shows examples of Undeleted Questions on Stack Overflow.

Id	Title	Score
145	Compressing / Decompressing Folders & Files in C#?	10
249	Accessing a remote form in php	15
5235643	Colon(:) and number in filename in Visio	1
5767118	Facebook text field detection	0

We find a total of **9,350** *undeleted* questions on Stack Overflow. **8,536** of these total questions were originally deleted by the question author while **814** questions were deleted by a moderator. We now analyze the time taken to *undelete* a question from the time of deletion. Figure 15 shows the percentile plot of time taken to *undelete* a question from the time of deletion of author-deleted and moderator-deleted questions. We find that most author-deleted questions are *undeleted* within 2 minutes of its deletion. We attribute this peculiar behavior to accidental deletion. The question author accidentally deletes her question but on realization of her mistake, she *undeletes* the question. We observe that moderator-deleted questions take a relatively longer time to *undelete*. This lag may be due to the time required for the *undelete* voting procedure as defined by the Stack Overflow community guidelines. The guidelines specify that a deleted question must receive a minimum of 3 *undelete* votes to be *undeleted*. This leads to an increase in time taken *undelete* moderator-deleted questions.

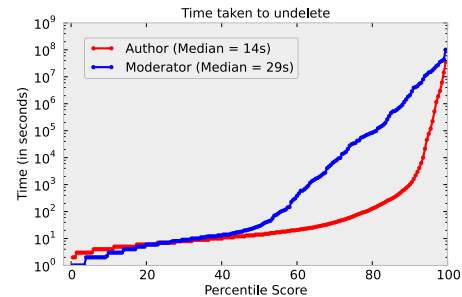


Figure 15: shows the percentile plot of time taken to *undelete* a question from the time of deletion of author-deleted and moderator-deleted questions.

We now analyze the *tags* on *undeleted* questions. Figure 16 shows the word cloud of the top-50 tags that occur in undeleted questions on Stack Overflow. We notice the presence of programming related tags like *objective-c*, *android* and *c#* which points out these *undeleted* questions are relevant to Stack Overflow.

## 5. DELETED QUESTION PREDICTION

In the second phase of our study, we develop a predictive model to detect a deleted question at the time of question creation on Stack Overflow. We frame the problem of deleted question prediction as a binary classification task. In our supervised classification framework, we simulate real-world conditions viz. we only use information available at question creation time. We do not have access to information on answers and other crowdsourced information like



data and make the dataset balanced [14]. Therefore, we randomly select **235,048** questions from the non-deleted class to form the negative class of our dataset. However, such a method may induce a sampling bias. In order to eliminate this bias, we draw 10 random samples of non-deleted questions and perform our prediction experiments on each random sample. We report our results on the average of the experiments drawn from all 10 random samples.

We experiment with multiple classifiers and find that *Adaboost* classifier gives the best performance. *Adaboost* or Adaptive boosting is an ensemble based machine learning framework which combines the performance of a series of weak classifiers to configure a strong classifier [12]. Concretely, *Adaboost* modifies subsequent classifiers in an attempt to correctly classify wrongly classified instances from previous classifiers. Prior work in content quality on community based question-answering websites have also observed best performance with ensemble-based learning methods [4, 17]. We use a 70-30% training-testing split and perform 10-fold cross validation to prevent the problem of over fitting. We use decision tree as the base classifier and SAMME.R for the boosting algorithm [26]. The learning rate and number of estimators parameters are set to their default values. Table 9 shows the experimental setup details for our supervised classification task.

**Table 9: shows the experimental setup for the deleted question prediction task.**

<b>Dataset</b>	470,096 questions
<b>Deleted (+ve class)</b>	235,048 questions
<b>Non-Deleted (-ve class)</b>	235,048 questions (10 times)
<b>Classifier</b>	Adaboost
<b>Learning Rate</b>	1.0
<b>Base Classifier</b>	Decision Tree
<b>Number of Estimators</b>	100
<b>Boosting Algorithm</b>	SAMME.R
<b>Cross Validation</b>	10-fold
<b>Classification Runs</b>	10-times
	(one for each +ve/-ve pair)
<b>Training-Testing split</b>	70-30%
<b>Feature Sets</b>	$\{S_A\}, \{S_A, S_B\}, \{S_A, S_B, S_C\},$ $\{S_A, S_B, S_C, S_D\}$

### 5.3 Evaluation

We now evaluate the performance of our classifier. Table 10 shows the confusion matrix for our supervised classification task. We see that our system is able to accurately classify 65.9% of deleted questions and 66.1% of non-deleted questions with an overall accuracy of 66%.

**Table 10: Confusion Matrix – Prediction Performance**

		Predicted	
		Deleted	Non-Deleted
True	Deleted	<b>65.9%</b>	34.1%
	Non-Deleted	33.9%	<b>66.1%</b>

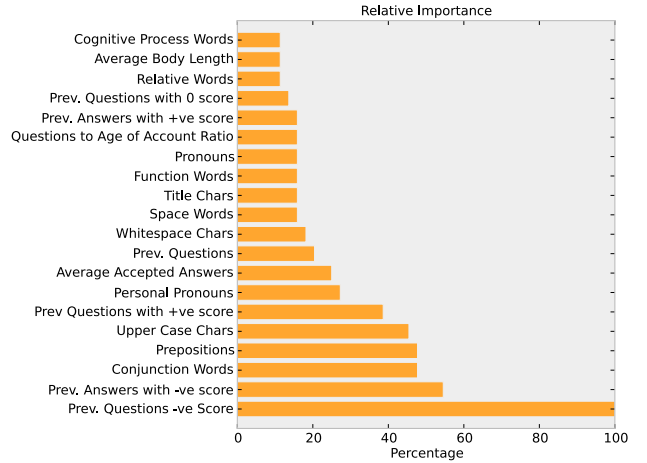
In order to understand the importance of our feature sets, we incrementally add feature sets and evaluate the performance of our classifier. Table 11 shows the classification performance on incremental feature sets based upon multiple evaluation metrics – F1 score, Accuracy and Area-Under-Curve(AUC). We note the improvement in performance of our classifier on each feature set. This shows that our feature sets are important to the classification task.

**Table 11: shows the classification performance on incremental feature sets based upon multiple evaluation metrics – F1 score, Accuracy and Area-Under-Curve(AUC)**

Feature Set	F1	Accuracy	AUC
$\{S_A\}$	58.91	56.81	57.19
$\{S_A, S_B\}$	61.83	59.59	61.11
$\{S_A, S_B, S_C\}$	63.38	62.17	65.24
$\{S_A, S_B, S_C, S_D\}$	65.8	66.0	70.01

### 5.4 Feature Analysis

We now analyze our feature set in order to understand important features for deleted question prediction. The *Adaboost* classifier informs the discriminator power of each feature for classification. Figure 17 shows the relative importance of top 20 features for deleted question prediction as given by our classifier. We see that the top-20 features are a mixture of features from all four categories of feature sets. It points out the importance of the use of different categories of feature sets for prediction.



**Figure 17: shows the relative importance of the top 20 features for deleted question prediction.**

## 6. CONCLUSIONS

We conduct the first large scale study of deleted questions on Stack Overflow. We observe an increasing trend in the number of deleted questions on Stack Overflow over the last 2 years. The community takes significant time to detect a potential deleted question but moderators take swift appropriate action. However, we notice that moderators bear most of the encumbrance of detecting a deleted question. We also find that authors delete their own questions to salvage reputation points on the website. In general, deleted questions are extremely poor in worth to the Stack Overflow community. Also, deleted questions are significantly low in quality than ‘closed’ questions. We discover the existence of a pyramidal structure of question quality in which deleted questions lie at the bottom of the pyramid (extremely poor quality). We also notice accidental question deletion by authors. We build a predictive model to detect deleted questions on Stack Overflow and report 66% accurate predictions. We employ four categories of feature sets – *user profile*, *community based*, *content based* and *stylistic* features – and report most discriminatory features.

## 7. REFERENCES

- [1] Why and how are some questions deleted? <http://stackoverflow.com/help/deleted-questions>.
- [2] How does deleting work? what can cause a post to be deleted, and what does that actually mean? what are the criteria for deletion? <http://meta.stackoverflow.com/q/5221/214223>, September 2008.
- [3] The great question deletion audit of 2010. <http://meta.stackoverflow.com/questions/51097/the-great-question-deletion-audit-of-2010>, May 2010.
- [4] E. Agichtein, C. Castillo, D. Donato, A. Gionis, and G. Mishne. Finding high-quality content in social media. In *Proceedings of the international conference on Web search and web data mining*, pages 183–194. ACM, 2008.
- [5] M. Allamanis and C. Sutton. Why, when, and what: analyzing stack overflow questions by topic, type, and code. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 53–56. IEEE Press, 2013.
- [6] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec. Steering user behavior with badges. 2013.
- [7] M. Asaduzzaman, A. S. Mashiyat, C. K. Roy, and K. A. Schneider. Answering questions about unanswered questions of stack overflow. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 97–100. IEEE Press, 2013.
- [8] J. Atwood. Stack overflow creative commons data dump. <http://blog.stackoverflow.com/2009/06/stack-overflow-creative-commons-data-dump/>, June 2009.
- [9] A. Barua, S. W. Thomas, and A. E. Hassan. What are developers talking about? an analysis of topics and trends in stack overflow. *Empirical Software Engineering*, pages 1–36, 2012.
- [10] A. Bosu, C. S. Corley, D. Heaton, D. Chatterji, J. C. Carver, and N. A. Kraft. Building reputation in stackoverflow: an empirical investigation. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 89–92. IEEE Press, 2013.
- [11] D. Correa and A. Sureka. Fit or unfit: analysis and prediction of ‘closed questions’ on stack overflow. In *Proceedings of the first ACM conference on Online social networks*, COSN ’13, pages 201–212, New York, NY, USA, 2013. ACM.
- [12] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Computational learning theory*, pages 23–37. Springer, 1995.
- [13] C. Gómez, B. Cleary, and L. Singer. A study of innovation diffusion through link sharing on stack overflow. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 81–84. IEEE Press, 2013.
- [14] H. He and E. A. Garcia. Learning from imbalanced data. *Knowledge and Data Engineering, IEEE Transactions on*, 21(9):1263–1284, 2009.
- [15] J. S. Jeff Atwood. Stack exchange platform. <http://stackexchange.com>, September 2009.
- [16] J. Jeon, W. B. Croft, J. H. Lee, and S. Park. A framework to predict the quality of answers with non-textual features. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, SIGIR ’06, pages 228–235, New York, NY, USA, 2006. ACM.
- [17] B. Li, T. Jin, M. R. Lyu, I. King, and B. Mak. Analyzing and predicting question quality in community question answering services. In *Proceedings of the 21st international conference companion on World Wide Web*, WWW ’12 Companion, pages 775–782, New York, NY, USA, 2012. ACM.
- [18] M. Linares-Vásquez, B. Dit, and D. Poshyvanik. An exploratory analysis of mobile development issues using stack overflow. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 93–96. IEEE Press, 2013.
- [19] L. Mamykina, B. Manoim, M. Mittal, G. Hripcsak, and B. Hartmann. Design lessons from the fastest q&a site in the west. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 2857–2866. ACM, 2011.
- [20] S. M. Nasehi, J. Sillito, F. Maurer, and C. Burns. What makes a good code example?: A study of programming q&a in stackoverflow. In *Software Maintenance (ICSM), 2012 28th IEEE International Conference on*, pages 25–34. IEEE, 2012.
- [21] A. Pal, F. M. Harper, and J. A. Konstan. Exploring question selection bias to identify experts and potential experts in community question answering. *ACM Transactions on Information Systems (TOIS)*, 30(2):10, 2012.
- [22] L. Ponzanelli, A. Bacchelli, and M. Lanza. Seahawk: stack overflow in the ide. In *Proceedings of the 2013 International Conference on Software Engineering*, pages 1295–1298. IEEE Press, 2013.
- [23] C. Shah and J. Pomerantz. Evaluating and predicting answer quality in community qa. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 411–418. ACM, 2010.
- [24] Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1):24–54, 2010.
- [25] W. Wang and M. W. Godfrey. Detecting api usage obstacles: a study of ios and android developer questions. In *Proceedings of the Tenth International Workshop on Mining Software Repositories*, pages 61–64. IEEE Press, 2013.
- [26] J. Zhu, S. Rosset, H. Zou, and T. Hastie. Multi-class adaboost. *Ann Arbor*, 1001(48109):1612, 2006.