

Understanding Triggers for Clarification Requests in Community-Based Software Help Forums

Nathaniel Hudson, Parmit K. Chilana, Xiaoyu Guo, Jason Day, and Edmund Liu
University of Waterloo
Waterloo, ON, Canada

Abstract—Help-seekers on community-based software help forums often face difficulty in composing queries or troubleshooting requests that bring immediate resolution, forcing help givers to request clarification that delays diagnosis. We investigate the characteristics of a forum post that trigger these requests for clarification from help givers (e.g., missing information, unclear goals, non-standard terminology). We created a classification scheme based on such triggers and applied it to 1000 Q&A pairs from four popular consumer software help forums to understand the prevalence of these triggers across different applications. Even though the user interface for posting questions on the four forums that we studied was largely uniform, we found a large difference in the presence of these triggers across the forums. Our findings suggest that instead of trying to create universal automated tools and recommendations for improving question quality on software forums, we should take into account the unique characteristics of the software and its user community.

Keywords—Software Help Forums; Software Troubleshooting; Web-based Q&A

I. INTRODUCTION

As software companies work to decrease the cost of providing one-on-one support to end users, many are turning to software help forums as a cost-saving measure [1]. End users also benefit as community-based forums allow them to post new questions or search through past answers and discussions by other users and support staff. However, often users' questions remain unanswered or forum threads involve lengthy back-and-forth dialogues with no clear resolution [2].

Prior work on software forums suggests that in many cases the help seeker's initial problem description lacks key information needed for diagnosis [1,2,3], triggering requests for clarification or additional details from help givers. For example, help seekers sometimes have difficulty in articulating their goals, actions and other aspects of their situation [3]. Some help seekers fail to provide a description of the steps needed to reproduce the problem [4]. Others fail to explain the context of their problem in words alone and are requested to attach visual information, such as screenshots [5].

Given the importance and growing popularity of forums in facilitating the use of modern software, recent initiatives have sought to improve question quality on these forums. For example, structured problem formulation scripts have been shown to produce more thorough explanations in bug tracking systems [3], and shared visual environments have been shown

to assist remote diagnosis tasks [6]. These works suggest that if we can devise interventions to enhance the quality of help seekers' questions and reports, we can improve the overall diagnosis of software problems. However, to design successful interventions for different kinds of software, we need to better understand the characteristics of help seekers' posts that are likely to stall diagnosis and result in requests for clarification.

In this paper, we investigate the *triggers* in a help seeker's initial forum post that prompt requests for clarification from help givers and the prevalence of these triggers across different software forums. We present a qualitative analysis of 1000 opening posts from help seekers from four different software help forums to characterize the kinds of details and information that users typically omit from initial descriptions. Although the user interfaces for posting questions on the forums that we studied were largely uniform, we found variation in the distribution of these triggers across the forums. Our findings suggest that the characteristics of the software and its user community can affect the style of help seekers' posts on forums and the kinds of information they provide (or do not provide). Instead of trying to create universal recommendations or interventions for improving question quality on software forums, we need to better understand and take into account the diversity of activity and varieties of question-asking behavior in individual question and answer (Q&A) communities.

II. METHOD

To study the triggers in help requests that prompt requests for clarification, we carried out a systematic qualitative analysis of Q&A discussions extracted from forums of four popular software vendors: Microsoft, Google, Adobe, and Autodesk. We selected these particular forums because they represent a range of interactive consumer-facing software used by millions of novice and professional end users and have consistent user interfaces for querying and posting Q&A.

A. Data Set Development

To conduct our analysis, we needed a dataset of Q&A pairs where the help giver had responded to the help seeker by expressing confusion and/or requesting more information. We began by first downloading a random sample of 100 Q&A posts (25 from each forum) and looked for key words or phrases that signaled potential confusion or need for more information (e.g., “*can you provide more details*”, “*what do you mean*”, “*can you post a screenshot*”, “*can you clarify*”, “*I don't understand your question*”). We created a list of 15 such

phrases most commonly seen in our data set and expanded them with synonyms to have a list of 36 signal phrases.

Using these 36 signal phrases, we queried recent posts on each forum site and extracted Q&A pairs where the help giver's initial response to the help seeker contained one or more of these phrases. We created an initial filtered dataset of over 1200 posts (25% from each forum) in June 2014 and expanded the data set with another 600 posts (25% from each forum) in April 2015 using the web interfaces for all of these forums and standard HTTP queries. Our research team manually inspected each downloaded Q&A pair with signal phrases and removed any false positives in our data set (e.g., forum discussions where the signal phrases did not actually represent a need for clarification from the help giver). About 28% of the posts fit these criteria and were removed from our data set, so we ended up with 1296 unique Q&A pairs.

We next describe the method that we used to classify triggers in a help seeker's posts and how we applied this classification scheme to our data set.

B. Classification of Triggers

Using our filtered data set of Q&A pairs that contained signal phrases, we next wanted to analyze the *triggers* in the initial phrasing of questions by help seekers that led to the confusion or request for more information. From our filtered data set containing 1296 Q&A pairs, we selected a uniform random sample of 200 Q&A pairs (50 from each forum) and did an open coding analysis to look for the triggers in phrasings of the questions by help seekers. Our coding was guided and informed by findings from prior work. For example, we know that users often face difficulty in describing their goals and intentions [3], in giving a visual description of their software issue [5,6], or in providing the needed steps to reproduce the problem and other application/system-related details [4]. We also referenced question-asking guides describing common asking mistakes among users' posts [7]. We employed an inductive analysis approach [8] to classify and reclassify our descriptions of the different question triggers. Three of the authors first independently examined all of the 200 reports, generating descriptions of each trigger. After numerous iterations, our discussions converged on the following coding scheme to describe triggers in help seeker's questions.

Use of Non-Standard Terminology: Help-seeker's description uses unfamiliar terminology in relation to the application. (e.g., "How do I transfer custom snippets to DW2014?")

Lack of Information on Steps Taken So Far: The description lacks information on the steps taken so far. How did they arrive at their current state? (e.g., "What step am I missing to be able to change the stroke on my rectangle? I'm trying to create a border around my document...").

Terse and Ambiguous Description: The description does not contain enough information, is ambiguous or too general. (e.g., HOW CAN I ACCESS MY EMAILS???? Since you sent the email about the new compose I cannot access my emails!!!!).

User's Goal is Unclear: It is uncertain what the user is trying to accomplish. What is their goal or intention? (e.g., "How do I

generate color separations? Using InDesign CS6 ver 8.0, OS X ver 10.8.5, Acrobat X ver 10.1.9.").

Lack of Visual Demonstration: The description is difficult to understand in text—a visual, such as a screenshot, video, or the actual file, is needed to understand or diagnose the underlying issue (e.g., "...I'm trying to convert the nice rounded face into a bunch of flat triangles which will still be angled to form the overall curved surface. Imagine a half circle, but instead of having a nice smooth curve on the top, it's got 5-6 flat faces that still form the alf circlecurve[sic]. If anyone could share tips or methods to make this, it'd be appreciated.").

Lack of Information on Current Application Settings or Software Version: The description is missing parameters or settings of the application needed to accurately assess the problem (e.g., "So I am not sure how to hide this option. This menu comes up every time I move my mouse during the presentation. How do I disable this? I am kind of new to Mac so I did it in windows version, but I [can't] figure out this version.").

Lack of Information on System Settings or Version: The description lacks the settings of the operating system or other system configuration details needed to assess the problem. (e.g., "How to uninstall Microsoft.net framework 4 client profile?").

C. Sampling and Analysis

To test our classification, we next selected a uniform random sample of 1000 Q&A pairs from our data set, excluding those used to develop the classification. Three of the authors applied the coding scheme described above to our sample and consulted with each other regularly to ensure that the codes were being applied consistently. In the event of conflicts between coders, or where multiple codes were possible, we had internal discussions to select a single most appropriate code that best represented the underlying issue. To formally assess the reliability of the coding scheme, we computed the Fleiss Kappa on a subset of 100 Q&A posts (coded by each of the 3 coders) to measure the overall agreement between all coders. We found strong agreement in the coding of the trigger types between the three coders ($\kappa=0.82$).

III. RESULTS

We now present our main findings on the characteristics of Q&A forum posts that trigger requests for clarification from help givers in software help forums.

Figure 1 shows the distribution of types of triggers for requests for more information that occurred across all four

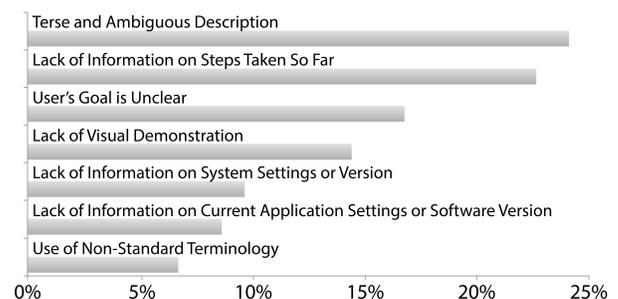


Fig 1. Distribution of Triggers Across All Forums

forums. The three most frequently occurring triggers were: *Terse and Ambiguous Descriptions*, *Lack of Information on Steps Taken So Far*, and *User’s Goal is Unclear*.

This finding is not too surprising given that prior work has shown that users struggle in conveying their software issues and often omit necessary information [3,4,6]. However, analysis of our data shows interesting trends and inter-application differences in users’ posts that have not been investigated in prior work. For example, as seen in Figure 2, we found that while a few of the triggers were prominent across all examined forums, the distribution of the top three triggers was not consistent between forums. The following is an analysis of each trigger type, proceeding from the most common to the least common overall.

A. Terse and Ambiguous Description

Help-seekers’ questions that only provided a *Terse and Ambiguous Description* constituted the largest category of triggers overall, accounting for 23.4% of all triggers. This trigger was most common at Adobe (34.4%) and Google (27.2%), and occurred to a lesser extent at Microsoft (18.4%) and Autodesk (13.6%). These threads were observed to be short and were often abandoned without resolution. One Adobe help seeker’s thread simply read: “*why doesn’t the adobe want to open?*”. Another Adobe thread read, “*Hi, i have question. How to make this effects as in the photo and video. Please good tutorial. [sic].*”, followed by an attached image from a complex professional production. Microsoft and Google posts often reflected help seekers’ frustration at being unable to access personal email, photos and more. One Microsoft user’s post read, “*WHERE HAVE ALL MY HOTMAIL CONTACTS GONE?*”. Similarly, a Google help seeker wrote “*I REALLY NEED TO ACCCESS [sic] MY NEW GMAIL*”.

B. Lack of Information on Steps Taken So Far

The second most common trigger was a *Lack of Information on Steps Taken So Far*, making up 22.0% of all recorded triggers. This trigger was the most common trigger at Google (25.6%) and Microsoft (23.6%), and was the second most common trigger at Adobe (18.0%), and the third most common at Autodesk (20.8%). This trigger was often associated with problems involving troubleshooting, slow trial-and-error resolution processes, and significant back-and-forth dialogue.

C. User’s Goal is Unclear

Threads where help givers requested more information due to the help seeker’s unclear goal occurred in 16.3% of the threads overall. This was primarily an issue at Autodesk with 27.2% of observed threads containing this trigger, while it was less common at Microsoft (16.4%), Adobe (12.0%) and Google (9.6%).

In Autodesk’s help forum we frequently observed very long, detailed and specific questions that needed context before help givers could recommend a more suitable approach. For example, one question post included two paragraphs describing the issues, multiple screenshots and illustrations, as well as a long code sample. Despite this level of detail, the initial replies

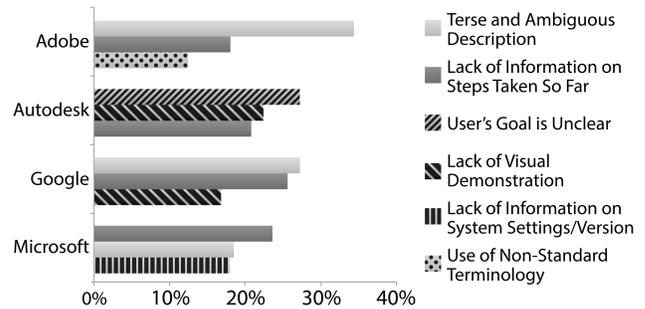


Fig 2. Three Most Frequent Triggers per Help Forum

were confused calls for context. The first reply to the post read: “*Why do you need to access that module?*”.

D. Lack of Visual Demonstration

The *lack of visual demonstration* trigger occurred in 14% of the threads overall. We observed this trigger type most frequently at Autodesk (22.4%) and Google (16.8%), and less frequently at Adobe (8.4%) and Microsoft (8.4%). This trigger was the second most common trigger in the Autodesk data set, and the third most common in the Google data set. We observed that requests for screenshots or visual examples were used at Autodesk as a way for help givers to quickly understand the help seekers’ design intent, as well as an alternative to requesting download of the help seeker’s file. One thread saw a help giver write: “*Can you share your assembly please ? [...] if it’s a huge assembly, just post a screenshot of what you are seeing.*”.

Google’s forum also contained a number of requests for visual clarification. Unlike Autodesk, we observed that this trigger reflects attempts by help givers to determine precisely what a help seeker was looking at or to troubleshoot cross-browser compatibility issues. One help giver asked, “*Could you please provide a screen capture of the product you are using. There is no such feature either in maps Engine Lite or Pro, maybe you are using Maps Engine, a different product?*”.

E. Less Common Triggers

A smaller number of requests for clarification were triggered by a *Lack of Information on System Settings/Version*. This was a frequent trigger at Microsoft (18.0%), but was less uncommon at Adobe (6.8%), Autodesk (6.0%) and Google (6.8%).

Requests for clarification based on a *Lack of Information on Current Application Settings/Version* accounted for 8.4% of all classified posts. As with requests based on system configuration, these were most frequent at Microsoft (11.6%) with Google (9.6%), Adobe (8.0%), and Autodesk (4.4%) making up the remainder.

While the *Use of Non-Standard Terminology* was the least frequently occurring trigger overall (6.5%), it was the third most frequent in Adobe (12.4%) with only a small number of occurrences at Autodesk (5.6%), Google (4.4%) and Microsoft (3.6%). One Adobe help seeker posted: “*Does anyone know how to go about creating an action for a sparkle using Photoshop CSS?*”. This was quickly met with requests for

clarification on what a sparkle effect was: “By ‘sparkle’ do you mean create an effect like what a photographic star filter would create?”

In summary, our analysis shows that while there are similarities in the prevalence of triggers across different software forums, there are large inter-application differences.

IV. DISCUSSION

Although our analysis is limited to four software forums and we did not look at intra-application variation, it provides some initial evidence of the differences in help seekers’ questions across the forums. We now reflect on our key findings and implications for improving the design of software help forums.

A. Influence of User Community and Software Type

One implication of our findings is that an individual forum’s user community and characteristics of the software being supported influence the types of questions help seekers ask and the kinds of details they provide (or do not provide). For example, we found that terse and ambiguous descriptions frequently occurred in Adobe, Microsoft, and Google’s forums and where novice help seekers often described installation troubles and misunderstood basic software features. However, such posts were a lot less common at Autodesk—rather, we found that help seekers at Autodesk tended to provide verbose descriptions and demonstrated basic knowledge of software features. Given the primarily professional user base of Autodesk products, we found that Autodesk posts tended to be about advanced usage of software features where help seekers often struggled in articulating their actual goals or intentions (by a margin of over 10% compared to the other software).

Our findings also indicate that the nature of the software product being discussed on a forum may affect the types of triggers contained in help seekers’ posts. For example, information on settings and versions was more frequently requested for diagnosing issues at Microsoft and was less common for the other software. We believe that this may be due to the nature of Microsoft’s Windows operating system that works fairly closely with the underlying hardware and where the system version and configuration details can be more critical for diagnosis. Similarly, posts at Autodesk were more frequently followed-up with requests for visuals, confirming findings from prior work that shows that support specialists at Autodesk indeed seek information in multimedia formats due to the graphical nature of the software products [5]. Interestingly, although Adobe also has several graphical software products, requests for visual demonstrations were less common at Adobe, with help givers instead engaging in text-based back-and-forth dialogues. This may again suggest that some sort of community Q&A culture or learned behavior exists in these software forum communities.

B. Redesigning Software Help Forums

Software help and troubleshooting studies in other contexts have proposed and demonstrated a wide range of techniques designed to reduce difficulties responding to help requests [2,5,6,10]. Our study suggests that the triggers in help seekers’

posts were often affected by traits specific to each forum user community or the software being supported, meaning that a technique that is helpful to one forum may not necessarily be helpful to another.

While previous research has discussed and demonstrated techniques for eliciting different types of software problem descriptions [2,3,5,6,10], our results suggest that universal solutions may be less applicable, given the diversity of triggers in help requests across different forums. For example, the structured problem reporting forms [3] have been successful at eliciting more verbose reports, but may not necessarily be helpful in cases where the user’s goals and intentions are not clear (for example, a major obstacle for questions posted on Autodesk forums).

Despite these caveats, there still are many design opportunities to guide help seekers to post questions that are less likely to cause confusion and more likely to lead to timely resolution. For example, it may be possible to create automated user interface enhancements to guide the posting of questions by help seekers that could be configured by forum administrators to meet their specific needs. For example, recent work [9] has demonstrated an approach to analyzing text using a genetic algorithm that was shown to be useful in classifying the acceptability of questions at Stack Overflow, a large Q&A site related to programming. This type of a tool could be complemented with some sort of a trigger analytics system such that the forum operator could make informed decisions on which interface choices to make to minimize information request triggers. It may also be possible to intelligently recommend improvements based on textual analysis of posts [4] or by crowdsourcing issue reporting in the context of the software application [10]. Such techniques have the potential to be automatically adapted to take into account individual software characteristics and community-specific Q&A behaviors to improve overall resolution time.

V. CONCLUSION

Our main findings demonstrate that how help seekers ask for help and the kinds of information they provide (or do not provide) varies across different software user communities. While our findings point to several opportunities for designing interventions that allow help seekers to avoid potential triggers in their posts, one key implication is that universal one-size-fits-all solutions for improving question quality on forums are less likely to work. Interventions should rather take into account individual characteristics of the software being supported and the forum’s question answering culture to improve help seekers’ posts.

ACKNOWLEDGMENT

We thank Shruti Dembla and Celena Alcock for their assistance. This work was funded in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] K. Xiaole, Y. Bo, Z. Shengdong, C. Jianann, G. Tovi, and F. George, "A Classification of Opening Posts in Commercial Software Help Forums," *APCHI '12*, Aug. 2012.
- [2] V. Singh and M. B. Twidale, "The confusion of crowds: non-dyadic help interactions," in *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, 2008, pp. 699–702.
- [3] M. Nückles and A. Ertelt, "The Problem of Describing a Problem: Supporting Laypersons in Presenting Their Queries to the Internet-based Helpdesk," *Int. J. Hum.-Comput. Stud.*, vol. 64, no. 8, pp. 648–669, Aug. 2006.
- [4] N. Bettenburg, S. Just, A. Schroter, C. Weiss, R. Premraj, and T. Zimmermann, "What Makes a Good Bug Report?," *IEEE Transactions on Software Engineering*, vol. 36, no. 5, pp. 618–643, Sep. 2010.
- [5] P. K. Chilana, T. Grossman, and G. Fitzmaurice, "Modern Software Product Support Processes and the Usage of Multimedia Formats," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2011, pp. 3093–3102.
- [6] L. Karsenty, "Cooperative Work and Shared Visual Context: An Empirical Study of Comprehension Problems in Side-by-side and Remote Help Dialogues," *Hum.-Comput. Interact.*, vol. 14, no. 3, pp. 283–315, Sep. 1999.
- [7] E. S. Raymond and R. Moen. (2014) *How To Ask Questions The Smart Way* [Online]. Available: <http://www.catb.org/esr/faqs/smart-questions.html>
- [8] A. Strauss and J. M. Corbin, *Basics of qualitative research: Grounded theory procedures and techniques*. Thousand Oaks, CA, US: Sage Publications, Inc, 1990.
- [9] L. Ponzanelli, A. Mocci, A. Bacchelli, M. Lanza, and D. Fullerton, "Improving Low Quality Stack Overflow Post Detection," in 2014 IEEE *International Conference on Software Maintenance and Evolution (ICSME)*, 2014, pp. 541–544.
- [10] P. K. Chilana, A. J. Ko, and J. O. Wobbrock, "LemonAid: selection-based crowdsourced contextual help for web applications," in 2012 *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2012, pp. 541–544.