Please Help! Patterns of Personalization in an Online Tech Support Board

Sarita Yardi and Erika Shehan Poole
School of Interactive Computing
Georgia Institute of Technology
Atlanta, GA 30308 USA
{syardi3,erika}@gatech.edu

ABSTRACT

We analyze help-seeking strategies in two large tech support boards and observe a number of previously unreported differences between tech support boards and other types of online communities. Tech support boards are organized around technical topics and consumer products, yet the types of help people seek online are often grounded in deeply personal experiences. Family, holidays, school, and other personal contexts influence the types of help people seek online. We examine the nature of these personal contexts and offer ways of inferring need-based communities in tech support boards in order to better support users seeking technical help online.

Categories and Subject Descriptors

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Measurement, Design.

Keywords

Technical support, computer help, consumer electronics, online communities, personalization, help seeking.

1. INTRODUCTION

If you had a question about your health, where might you go to find an answer? Recent studies suggest that roughly 75-80% of American Internet users will go online to look for health or medical information [11]. Indeed, a health board can be a conducive environment for social support, conversation, and sharing of testimonials. But if you instead had a question about your home network, where might you go to find an answer? A related study found that only 2% of people seek technology help online [18]. While broadband ownership at home rose from 33% in 2005 to 55% in 2008 [17], home technologies have become more complex and are increasingly difficult to setup and manage [10]. As people purchase new devices, they are confronted with protocols, tools, and terminology that may be unfamiliar, many of which were designed and architected for skilled or professional users [14]. As more householders confront the difficulties of setting up and maintaining information technologies at home, it becomes more important to understand how to provide resources for technical assistance.

C&T'09, June 25–27, 2009, University Park, Pennsylvania, USA.
ACM 978-1-60558-601-4/09/06.

To better understand current user practices of technical help-seeking online, we examine the ways that people ask for help in two large online tech support discussion boards. What kinds of help are they seeking? What responses do they get? How do patterns of help-seeking relate to special events in people's lives? How can we better support people in seeking and finding technical support online? Prior studies of support boards have used primarily quantitative approaches to model community structure. We know little about the ways that people personalize posts in tech support boards, and the ways that people construct their requests for help in these communities. Tech support boards may be more factual, less discussion-oriented, and more strongly elicit expert and novice roles than other online communities [2]. We argue that tech support boards are different than previously studied online communities and examine the ways that these differences influence help-seeking behaviors.

Using web scraping techniques, we captured over four years of board activity in two large-scale tech support boards focused on consumer support for products such as laptops, mp3 players, and routers. We use a combination of quantitative and qualitative approaches to examine help-seeking strategies and behavior in tech support boards. We examine the ways in which these boards exhibit social and community-oriented characteristics, and the rhetorical strategies used by posters. We analyze strategies used to contextualize and personalize requests for help and the types of response they receive. The questions we address are:

- How do personalized posts differ from other posts?
- What kinds of responses are given?
- How do personalization strategies correlate to real-world events in people's lives?
- What patterns of post strategies emerge, and how can we better support users in seeking help from one another online?

We anticipated that patterns of activity in tech support boards would differ from other types of online communities. Specifically, we hypothesized that most posts would not contain introductions, posts would contain different types of strategies to solicit community attention, and activity would strongly reflect rhythms of users' real lives. Understanding the types of questions posted, when they are posted, and how they relate to the real world can help us to design systems to better support users seeking technical help online.

The paper proceeds as follows. In the Related Work section, we describe examples of online help seeking and expertise systems. We then describe the methods we used to capture datasets from two tech support boards. In the Results section, we describe temporal dynamics of activity on the boards and post strategies used by board participants. In the Discussion section, we describe strategies of post construction, differences between tech support boards and other kinds of online communities, and implications for inferring help-seeking communities. Finally, we conclude with limitations and future directions.

2. RELATED WORK

Tech support online is an understudied genre of online community, but has a long and rich history. People have been sharing technical help on bulletin board systems (BBSes), Internet relay chat rooms (IRC), and multi-user dungeons (MUDs) for decades [6, 32, 37]. However, as broadband access grows, and consumers purchase more and increasingly complex devices for their homes, there is a growing market for tech support beyond these mediums.

Traffic on tech support boards is likely to grow. In the U.S., up to 77% of households had computers in 2008, and 52% had a home broadband connection [21]. Furthermore, almost 60% of nuclear families owned two or more desktop or laptop computers, and just over 60% of these households connected their computers in a home network [21]. Consumer networked products are expected to grow from 492 million units in 2006 to 2.8 billion units in 2010 [9]. Furthermore, as people become more autonomous and as a culture of user produced content pervades the Web, studying the ways that people come together to ask questions and offer answers is valuable.

Despite this optimistic growth and adoption, many people struggle to setup and maintain devices and network connectivity at home. Almost half of users needed help with new devices and a similar number reported that their home Internet access connection failed to work for them sometime in the last 12 months [18]. The remaining untapped consumer market for home networks is low-income individuals [17], who are likely to have less access to help resources. For this demographic, providing quick and free access to online support becomes more critical.

2.1 Helping Others

People come together online to share help in a number of different ways. Torrey, et al. described a form of procedural knowledge sharing called "how-to" sharing [36], Halverson discussed the FAQ [15], and Perkel and Herr-Stephenson described the use of tutorials [31].

In each of these community genres, one or more users describes some set of skills or knowledge they have, and they post it online. Perkel and Herr-Stephenson, for example, explore how people develop media production expertise and create and circulate tutorials based on this expertise on deviantART [31]. Torrey et al.'s "how-to" describes a class of participation by hobbyists engaged in activities such as software use and modification, hardware and electronics, home improvement, knitting, sewing, and woodworking. They describe how some versions of the how-to offer a chronological story of the author's experience, while others are more like recipes, with a list of the necessary tools and step-by-step instructions in order to complete the task [36].

Support boards differ from these classes of sites in three ways: posts are shorter and less personal, posters usually post a question or an answer but do not expect to receive feedback on their own work [36], and topics cover a broad range of interests rather than primarily crafts and hobbies.

2.2 Helping People Help Each Other

Research in expertise matching systems was borne out of offline interactions. Ackerman et al. constructed systems to bring together expertise networks, either by finding those interested in a particular topic or by constructing ad-hoc teams with the required knowledge [1]. They define expertise identification as the problem of knowing what information or special skills other individuals have, and expertise selection as the process of choosing people with the required expertise. An individual may have different levels of expertise about different topics. In other words, expertise is relative, and depends on the contexts in which it is placed [27].

Building off this work, Zhang et al. describe an expertise-finding mechanism that can automatically infer expertise level [41]. They used a set of simulations based on Java Forum data to explore how structural characteristics in the social networks influence the performance of expertise-finding algorithms [40]. In a related study, they explore the effectiveness of a ranking engine that infers expertise by constructing a community asker-helper network based on historical posting-replying data [41]. Similarly, Jurczyk and Agichstein formulate a graph structure and adapt a web link algorithm to estimate topical authority that looks to estimate the authority of people who answered the question, rather than the authority of the answer itself [20].

A complementary approach identifying expert posters is to identify expert posts. In this domain, Adamic et al. examined answer quality and found that they could use replier and answer attributes to predict which answers are more likely to be rated as best [2]. Similarly, Agichtein et al. investigated methods for exploiting community feedback to automatically identify high quality content [3]. Last, Liu and Agichtein developed personalized models of asker satisfaction to predict whether a particular question author will be satisfied with the answers contributed by the community participants [26].

2.3 Yahoo! Answers and Usenet

A number of studies examined quantitative measurements of question and answer boards and other types of discussion boards[1]. Yahoo! Answers (YA) is a strong example of a community-driven question and answer support board. Posters can ask questions on almost any topic and points are awarded for good answers. As of March 2008, YA worldwide had 135 million users and 500 million answers.

Adamic et al. analyzed YA board categories and clustered them according to content characteristics and patterns of interaction among users [2]. They found that some boards resembled technical expertise boards while others were more like support, advice, or discussion-oriented boards. Technical boards with factual answers, like Programming, Chemistry, and Physics, tended to attract few replies but the replies were lengthy (in contrast to Wrestling and Marriage boards, which were more conversational). They also noted that the Programming category had a high out-degree distribution, which reflected the highly active individuals who help others frequently but do not necessarily ask for help themselves.

Zhang et al.'s study of Java Forums found that the majority of users made few posts, and a number of experts mainly answered

[1] In this paper, we view "question and answer" and "support" interchangeably, as well as "forum" and "board".

others' questions without asking many questions themselves. [40]. They note that top repliers answered questions for everyone, whereas less expert users tended to answer questions of others with a lower expertise level [41].

Prior online community research documented a number of post construction strategies for new users based on studies of large discussion boards such as Usenet. Posting on-topic, asking questions, using less complex language, and including introductions have been shown to increase the likelihood of receiving responses, in part by helping create a sense of legitimacy in the group [5, 8, 16, 19]. Introductions, in particular, help signal legitimacy: autobiographical introductions reveal personal connections to the community, topic introductions reveal familiarity with the community, and group introductions reveal commitment to the community.

These approaches have a shared goal of automating components of the question and answer process in order to improve user experience. Other research has looked at roles in technical forums. Combs Turner and Fisher discuss four social types in technical newsgroups on Usenet [39]. They analyze information flow, focusing on how information needs are expressed, sought, and shared. Information seekers users are grouped into four roles: Questioner, Answerer, Community Manager, and Mogul. They show how technical newsgroups exhibit many of the same characteristics of general online communities.

However, none of these studies have focused on examining personalization strategies or the effects of real-world events (such as holidays) on board participation. Tech support boards are a class of discussion board that is characterized by strong roles of novice and expert skills. Many users will have novice skills, and a few users will be experts. Help seeking in tech support boards is less likely to be reciprocal and to rely on decentralized volunteer contributors [2].

3. METHODS

We scraped content from two large online tech support discussion boards. Both are hosted by a large international consumer product technology company that is headquartered in the United States. This company hosts customer-driven online technical support boards for its products. The support forums are split into several sub-boards. The boards are primarily by and for consumers of the brand; while there are a handful of employees who contribute content, the majority of the content comes from consumers. We specifically scraped content from a board intended for new, novice users, which we will call NewbieBoard, as well as a board for more specific network-related problems, which we will call NetworkBoard.

Users come to NetworkBoard seeking help with setting up their home networks. Questions usually relate to network setup, internet connectivity problems, and purchasing advice queries. Users come to NewbieBoard posting a wide range of questions about consumer products, complaints about technical support and service, and requests for purchasing and troubleshooting advice "for a newbie." NewbieBoard is more heavily trafficked because it is a general topic board and includes a range of technologies and digital electronics, as well as a range of user expertise levels. Our data from NetworkBoard and Newbie Board includes content posted from 2003 to 2007. In this time frame, NetworkBoard had 22,000 posts and NewbieBoard had almost 90,000 posts (see Tables 1 and 2). Our methodological approach consisted of five components:

Table 1. Board Descriptions

	Description	First Post	Last Post
NetworkBoard	Networking, Internet	1/31/03	12/03/07
NewbieBoard	New User	9/30/03	12/02/07

Table 2. Board Participation

	NetworkBoard	NewbieBoard
Users	7231(14991[2])	14674 (34441)
Posts	21,890	89,684
Threads	5,321	24,575
Avg posts/thread	4.11	3.65

1. Collecting data from two boards using Perl scripts;
2. Plotting temporal graphs of this data using Spotfire;
3. Isolating categories of posts by performing keyword searches of raw post data in MySQL;
4. Linguistic analyses using LIWC [29];
5. Qualitative ontent analysis with two researchers manually coding posts;

We draw from a variety of prior methodological approaches to studying online communities to inform our study design. Pennebaker, Kramer, and Joyce have studied large quantities of text using linguistic approaches like word count and clustering [4, 19, 30]. Burke et al. and Arguello et al. used machine learning patterns to capture rhetorical patterns in messages [5, 8]. Finally, Joyce et al analyzed a smaller corpus of text using manual coding [19]. Other work has used visualization techniques to portray online communities; Smith and Fiore presented an interface dashboard for navigating and reading discussions in social cyberspaces, and Golder et al. and Yardi et al. visualized temporal trends of social network and blogs, respectively [13, 34].

Our approach is largely exploratory. We measure large-scale empirical patterns using keyword searches and temporal analyses, and support the results with manual content analyses. We hypothesized that types of questions in a tech support board would differ from other online communities and that existing methods for studying online communities would be insufficient. Because we were interested in examining personalization strategies, we applied artificial filters to select particular subsets of data that were likely to contain personal information. We isolated these subsets using MySQL keyword searches on terms related to family, gifts, holidays, and school (see Section 4.2). We used this subset to run linguistic analyses (see Section 4.4) and qualitatively analyze personalization strategies by coding them for key themes (see Section 4.3).

For the manual coding, two researchers coded 414 threads (with 3-5 posts in each thread for a combined total of 1514 posts)

[2] User Ids were captured for only a subset of the data. The first number is unique ids, the second is anonymous posters (where there may be multiple ids per user).

containing some aspect of personalization. Two coders developed, discussed, and refined the codebook together. We used Cohen's Kappa to compute intercoder reliability. Cohen's Kappa computes the percentage of agreement between two coders, where a Kappa value above 0.75 is be considered excellent agreement [25]. Coders were trained by coding a set of 40 posts and comparing and discussing codes. We threw out the code for whether or not a question posted to the board was answered because determining if questions were answered even if a response was given was too speculative (some posters returned to the board to explicitly report success or failure, but most did not). The Kappa value for this code was K=0.50. Kappa values for all other codes ranged between K=0.72 and K=0.78.

4. RESULTS

We structure our results in three sections. First, we describe activity levels, response rates, and response times in each of the three boards. We compare these baseline metrics to other studies. Second, we visualize temporal patterns in each of the boards, examining variations in use and their correlation to real-world holidays and events. Third, we ran keyword analyses to classify particular rhetorical instances, such as emotion, personal pronouns, and quantifiers. Last, the manual content analysis is performed by coding for instances of personalization strategies.

4.1 Board Participation

We calculated the percentage of posts that received responses in NetworkBoard and NewbieBoard to be 83.3% and 90.2%, respectively (see Table 3). Related studies suggest a range of response rates. Burke et al. showed that 40% of potential thread starting messages in Usenet groups received no response [8]. In a smaller dataset of 6,172 messages from eight Usenet newsgroups, Arguello et al. found that 27.1% of posts did not receive a response [5]. Similarly, Jurczyk found that fewer than 35% of all questions had any user votes cast for any of the answers [20].

Both boards in this study had higher response rates than the Usenet groups. It is possible that the signal-to-noise ratio is perceived to be higher in a tech support board community. Unlike political boards, or pop culture boards (e.g. Angelina Jolie versus Jennifer Aniston comparisons [2]), the majority of people who come to tech support boards are genuinely seeking help. Indeed, Burke et al. found that posts making specific requests (as opposed to general discussion topics) increased the likelihood of getting a reply. However, not all boards exhibit higher response rates when requests are made. Adamic et al. report that on average, only 6% of questions in the Cancer category of Yahoo! Answers go unanswered [2]. Thus, some exceptions may exist where serious topics like terminal illnesses elicit abnormally high response rates.

Furthermore, few posts are *reply sinks* in a tech support board. Smith and Fiore describe reply sinks as those posters who regularly receive large numbers of direct responses [34], thus draining resources and expertise that could otherwise be more evenly spread throughout the community. Unlike in other online communities, however, the nature of tech support boards mitigates the drain of reply sinks. Although both boards exhibited power law behavior (where some posts receive many responses and most receive few), the curve is steep; less than 5% of threads are sinks having threads of 20 posts or more, and a much larger 82% receive one or two replies.

We also compared response times across boards. Despite the wide range of activity levels across boards and months, response time remained consistent. There was significant deviation in the boards because some posts remained unanswered for many months so we report the average response time for the third quartile of both boards. The average of the third quartile of NetworkBoard was 220 minutes and of NewbieBoard was 180 minutes. The longer response time for NetworkBoard may be due to its specific networking focus, which draws from a smaller and narrower pool of responses and likely requires more expertise. Zhang et al. show that the average waiting time of a high expertise user asking a more challenging question is about 9 hours, compared with 40 minutes for a low expertise user [41].

Table 3. Response Rate

	Got Replies	No Replies	% Got Replies
NetworkBoard	4,375	946	83.3%
NewbieBoard	22,190	2,273	90.2%

4.2 Rhythms of the Real-World

Figure 1 shows the evolution of views per post on NetworkBoard and NewbieBoard with columns shaded by year from 2003 to 2007. While total traffic on the board has grown over time, views per post average at around 200. Figure 1 shows minor periodic rises during the end of each calendar year and into the beginning of the next. To examine these trends, we plotted number of posts by month of year. Figure 2 shows number of posts to NewbieBoard and NetworkBoard binned by month from 2003-2007. Traffic grows steadily from July to November and drops steeply in December. Given the large size of the dataset, it is unlikely that the monthly trends are by chance.

We hypothesized that the rise and fall of activity might be explained by consumer purchasing trends. To explore this hypothesis, we extracted categories of posts related to real-world events, such as holidays, special events related to family and friends (e.g. birthdays or anniversaries), and purchasing items to prepare for an upcoming school year.

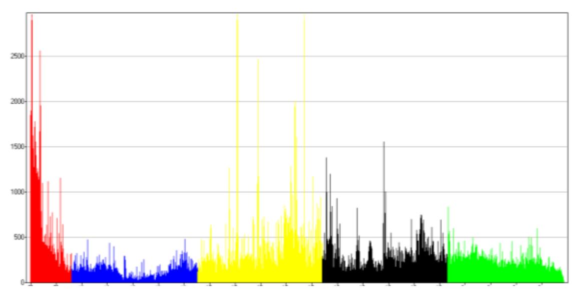

Figure 1. Average viewcount / post.

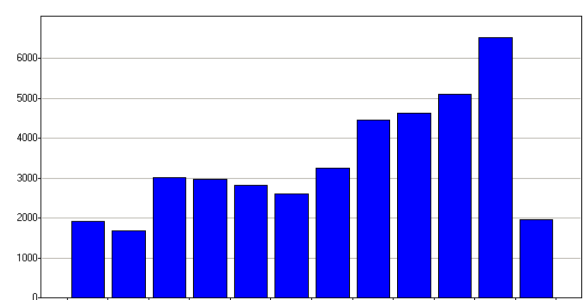

Figure 2. Monthly posts.

We isolated posts containing references to holidays by selecting a range of possible keywords such as “holiday”, “Christmas”, “Kwanzaa”, and “Santa”. Figure 3 shows posts plotted by year, grouped by week for NewbieBoard and NetworkBoard. The sharp spike in traffic in Figure 3 occurs in late December, lagging behind overall traffic in Figure 2. This lag may be caused by a culture of early holiday shopping and the subsequent opening and using of gifts. Indeed, consumer purchasing statistics show that the greatest month-to-month online traffic increases happened between October and November, where consumer electronics sites saw a 30% increase in 2007, a 16% increase in 2006, and a 21% increase in 2005 during these months [9].

We then isolated posts about school to examine correlations to the academic year calendar. The school-related category included four terms: “school”, “university”, “college”, and “semester”. Figure 4 shows a distinct, although more gradual peak. Traffic starts to climb in early July and is highest in August and early September. Traffic slows down again in October and November. It is likely that purchasing of new computers and related consumer products closely align with the beginning of a school year. We observed many posts related to starting a new school year and needing purchasing advice or troubleshooting advice, especially among college students, or parents whose kids were going off to college.

Last, we isolated posts by close relationships. This included references to terms such as “mom”, “dad”, “parent”, “brother”, “son”, “child”, “girlfriend”, and “boyfriend”. We combined family, school, and holiday references to infer some amount of personalization in types of posts. While a number of other keywords could be used to index into personalization (and, certainly not all references to family necessarily are highly personalized posts), we use these three as commonly referenced topics that emerged in our studies of NetworkBoard and NewbieBoard. Subsequent references to personalized posts include the aggregate of these three categories.

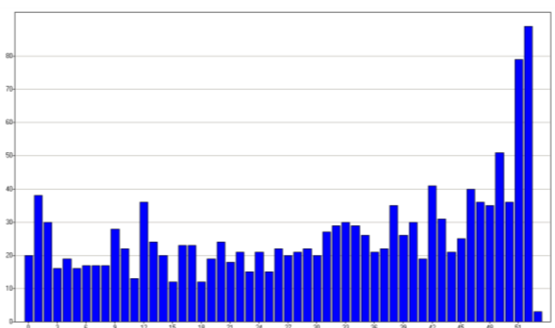

Figure 3: Posts with holiday keywords by week.

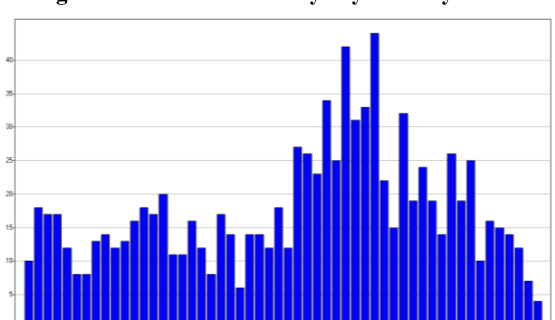

Figure 4: Posts with school keywords by week.

Table 4. Linguistic Patterns in Personalized Posts

	Details	Personalized	Control
Personal Pronouns	I, them, her	8.54	4.47
1st person singular	I, me, mine	6.90	4.00
1st person plural	We, us, our	0.93	0.23
2nd person	You, your	0.09	0.23
3rd person singular	She, her	0.425	0
3rd person plural	They, their	0.29	0
Articles	A, an, the	4.04	9.31
quantifiers	Few, many	1.06	2.93

4.3 Personalization of Posts

Automated text recognition can be a useful index into different styles of help-seeking in a tech support board. We ran a series of linguistic analyses comparing overall posts to the isolated set of personalized posts (discussed in Section 4.2). The question we examine is how explicitly personalized posts (e.g. containing references to family, gifts, holiday, or school) compare to posts that do not contain those keywords.

We first compared response rates and response times of the subset of personalized posts to the values reported in Section 4.1. We found no significant differences in response rate for personalized versus non-personalized posts. This may be because the average response rate was already quite high. However, we found that the average response time increased noticeably, from 220 minutes on NetworkBoard and 180 minutes on NewbieBoard to 112 minutes in the personalized group (consisting of posts from NetworkBoard and NewbieBoard). More interestingly, 69 posts of the 414 received responses in less than 30 minutes, and 24 posts received responses in less than 10 minutes. The higher response time may be a function of multiple variables; posts containing references to personal contexts like family and birthdays may elicit empathy from the community and a quicker response, posts containing personal information may be more novice and thus can be answered by a large body of community members, and references to event-driven deadlines like birthdays and school deadlines may elicit sympathy from the community and a quicker response time. These results suggest personalization strategies can increase response times, but studies with larger sample sizes are needed.

In addition to the common keywords in the personalized posts, we examined if there were particular linguistic patterns. We used LIWC to calculate the degree to which people use different categories of words across texts [29]. The results in Table 4 show consistent differences between “personalized” posts and other posts online; personalized posts contained more personal pronouns, but fewer articles and quantifiers. Personalized posts also contained fewer than half as many positive emotional keywords (1.355 versus 2.82) and over twice as many negative emotional keywords (1.275 versus 0.47). The exception is second person pronouns, where there is higher usage in the control posts than the personalized. The second person pronoun outlier aligns with prior research. Citing Pennebaker et al. [28], Burke et al. suggest that use of first and third person pronouns elicits greater community response, while second person pronouns reduces it [8]. Our results suggest that although response *rates* are not

Table 5. Error Details

	Product Details	Symptoms	Error Message
Yes	465	182	53
No	1049	1332	1461

Table 6. Personal Information

Relevant	Extraneous	Both	Neither
59	225	92	1138

affected, uses of first and third person pronouns versus second person pronouns may vary response time. Different types of posts may reveal important information about help seeking styles with implications for supporting a broader range of help seekers.

4.4 Types of Posts

We manually coded 1514 personalized posts (414 threads, with 3-5 posts per thread) to better understand how people ask for help in tech support boards. We specifically chose posts that received 2-4 responses in order to capture the back and forth nature of the discussions. What kinds of posts are written? What kinds of information is given about the problem? What kinds of responses do they receive? Two coders each coded half of the 1514 posts. Our codebook contained the following themes, which we coded as yes or no (and where not all themes applied to all posts):

- who product was purchased for;
- reason for purchase;
- if product details were included;
- if error messages were included;
- if symptoms were given;
- if steps tried to solve the problem were reported;
- if the post referenced company policy;
- if gave a reason why they should receive help;

The results of our coding are shown in Tables 4-7. Table 4 show that about 30% of users (both question askers and responders) provided details about the product they were discussing. We define details as technical specifications beyond the generic brand name of the product. We found that 44% of users requesting troubleshooting help provided symptoms of the error, and 16% provided the error message in their post. The low number of error messages reported suggests users either did not know what the error was, or they were unable to articulate it in the board post. Related studies have shown that surfacing errors and making them visible to the user is an important goal [33, 35]; being able to recognize and describe the problem in a tech support board can help both users and the tech support community to troubleshoot the problem.

In addition to the above yes/no codes, we also coded personal information in the posts as relevant or extraneous or both. Relevance was coded when the post contained useful information in helping to construct a response to the question, e.g.:

> *"What laptop should I get for my elderly father with poor vision"?*

In contrast, extraneous was personal information that was not relevant to the question being asked:

> *"Can someone help me choose a new laptop? I'm switching jobs and I'm tired of my old one."*

About 30% of the 1514 isolated posts contained personal information, of which over half are extraneous. Although extraneous personal information is off-topic, or peripheral to the main topic, it may contribute to a sense of community and foster social support.

We developed an additional code to categorize post types. The most frequent types of posts from question askers were requests for troubleshooting help, requests for purchasing or warranty advice, and responses reporting back results of trying a step. The most frequent types of posts from community responders were procedural advice containing steps that the original poster might try to solve the problem, and asking for clarification or details from the original poster. Table 7 shows the top categories of post types. Remaining codes (not shown) had less than 100 posts each. These remaining codes were: sympathy, off-topic comments, antagonistic comments, complaints, testimonials, and recommendations to redirect the question to another forum.

The most frequent sequences of back and forth exchanges we observed were initiated with a request for troubleshooting help, followed by procedural advice from a community member, then reporting back the results of the advice, and one more set of recommended procedures. In some cases, a question would be asked and a multiple community members would respond with procedural advice, sometimes also asking for clarifications or details. Requests for help often contained undertones of desperation: college students asking for discounts, people buying gifts for another person, or needing to meet an impending critical deadline and requesting immediate attention.

Table 7. Post Question Types

Type of post	Number
procedural advice	546
request for troubleshooting help*	298
report back of results of trying a step	185
asking for clarification or details	116
request for purchasing or warranty advice*	108

*requests for help

4.5 Detecting Joy and Distress

We define joy and distress indicators in NetworkBoard and NewbieBoard as posts containing instances of multiple punctuation characters (e.g. Help!!! and Why doesn't this work???) as well as heavy use of capitalization. We defined heavy use as three or more words in a row containing three or more capitalized characters in a row (e.g. PLEASE EXPLAIN WHY THIS DOESN'T WORK). Using these two metrics as a rough proxy of joy or distress, we found that 10% of posts (13,594) have at least two exclamations in a row, and slightly over 1% (1,850) use five or more. The text search for strings of capitalized letters revealed 997 posts, e.g.:

> *"I have a dell dinension [gives version number] the fan runs all the time and is loud can anyone help. I JUST GOT IT FOR CHRISTMAS AND DON'T KNOW A LOT ABOUT WHAT I"M DOING YET."*