

THE PENNSYLVANIA STATE UNIVERSITY
SCHREYER HONORS COLLEGE

COLLEGE OF INFORMATION SCIENCES AND TECHNOLOGY

AN EXPLORATION OF COLLABORATIVE ACT DISTRIBUTIONS ON THE STACK
EXCHANGE NETWORK

EVAN FRIEDENBERG
SPRING 2016

A thesis
submitted in partial fulfillment
of the requirements
for a baccalaureate degree
in Information Sciences and Technology
with honors in Information Sciences and Technology

Reviewed and approved* by the following:

Xiaolong Zhang
Professor of Information Sciences and Technology
Thesis Supervisor

Shawn Clark
Professor of Business
Honors Adviser

* Signatures are on file in the Schreyer Honors College.

ABSTRACT

Social question and answering (social Q&A) services have evolved from digital reference services and expert services, which have connected information seekers with experts or librarians for several decades. Social Q&A services provide the ability to combine the knowledge of many and create an indexable knowledge repository that can be accessed by search engines. They allow users to pose questions and then work with the community to find answers to those questions. Use social Q&A sites have grown immensely in recent years as people turn more frequently to the Internet to find information. Popular services in the United States include Yahoo Answers, Quora, and Stack Exchange. Motivating this research, collaboration in the process of information seeking results in answers that are better than the sum of the individual answers. Understanding the nature of collaboration on social Q&A services can inform design changes that improves the quality of the information that the sites contain.

To investigate the types of collaborations that occur on these sites, a comparison of the collaborative acts occurring on Stack Exchange was conducted, extending the work of Tausczik, Kittur, & Kraut (2014). Stack Exchange exists as a network of identically designed sites federated by topic, with each developing its own community. Coding of existing threads on Stack Exchange, using Stack Overflow, DIY, and Parenting as samples, was completed and the distributions were compared in order to understand the differences in collaboration and provide design recommendations. Findings indicate that the distributions of collaborative acts on Stack Overflow differ significantly from those of DIY and Parenting. Contributions of this research 1) a more inclusive schema of collaborative acts on Stack Exchange, 2) a method for collaborative coding and comparing distributions of collaborative acts, 3) evidence that not all sites on the network share the same distribution, and 4) design recommendations based on the findings.

TABLE OF CONTENTS

LIST OF FIGURES	iv
LIST OF TABLES	v
ACKNOWLEDGEMENTS	vi
Chapter 1 Introduction	1
Organization of Thesis	3
Chapter 2 Literature Review	5
2.1 Q&A Services	5
2.1.1 Digital Reference Services	5
2.1.2 Expert Services	6
2.1.3 Social Q&A	7
2.2 Communities of Practice in Social Q&A	11
2.3 CSCW and CIS	12
2.3.1 Computer-Supported Cooperative Work	13
2.3.2 Collaborative Information Seeking	14
Chapter 3 Research Methodology	17
3.1 Stack Exchange Description	17
3.2 Research Approach	20
3.3 Sample Selection	22
3.4 Data Collection	24
Chapter 4 Findings	28
4.1 Analysis of 3 Sites	28
4.2 Pairwise Comparison of 3 Sites	30
4.2.1 Comparison of SO and DIY Frequencies	31
4.2.2 Comparison of SO and P Frequencies	31
4.2.3 Comparison of DIY and P Frequencies	32
Chapter 5 Discussion	33
5.1 Collaborative Acts on Stack Exchange	33
5.2 Comparison of Stack Overflow, DIY, and Parenting	36
5.3 Comparison of Stack Overflow <i>to</i> DIY and Parenting	38
5.4 Design Recommendations	41
5.5 Contributions	44
5.6 Limitations	46
Chapter 6 Conclusion	47

6.1 Research Question and Method47
6.2 Contributions and Limitations.....48
6.3 Future Work.....49
BIBLIOGRAPHY 51

LIST OF FIGURES

Figure 1. Four Layer Model of Information Seeking (Shah, 2009).....	14
Figure 2. A Model for Collaborative Information Seeking (Shah, 2009).....	15
Figure 3. Example of Stack Exchange Question	18
Figure 4. Example of Stack Exchange Answer	19
Figure 5. Histogram of Collaborative Act Frequencies	29

LIST OF TABLES

Table 1. Reputation Points and Privileges	19
Table 2. Questions Sampled	23
Table 3. Coding Schema For Each Post.....	26
Table 4. Distribution of Collaborative Acts.....	28
Table 5. Demographics of Communities Based On Web Traffic	40

ACKNOWLEDGEMENTS

I would like to thank my advisor, Luke Zhang, for his invaluable feedback and support throughout the entire research process and helping me to complete the Master's degree.

Thank you to my thesis committee, Shawn Clark and Michael McNeese, for taking their time to be on the committee and giving their opinions in order to improve the quality of my research.

I would like to thank Madhu Reddy who advised me in the beginning of the Master's program and who first introduced me to research in IST. I would also like to thank Nathan McNeese who taught me, for countless hours, so much of what I know about research.

Thank you to Tamara Peyton who has been an incredible help in completing this research with collaborative coding. I could not have done it without you. Thank you Alison Murphy and Elizabeth Eikey for your feedback on my defense.

Thank you to all of my lab mates throughout the last 5 years, who have given me advice, listened to complaints, entertained me, and made my experience so enjoyable.

Lastly, thank you to my mom and dad, brother and sisters, and all of my friends who I would be nothing without.

Chapter 1

Introduction

While question and answer services have existed before widespread access to the Internet, the web has caused an explosion in the number of users searching for information from the online community. Social question and answering (social Q&A) sites exist as a place for those seeking information to find it using the resources of the online community. Originally, online Q&A was a way to find answers to questions that had a clear solution. Recently, on sites such as Yahoo Answers, Quora, and Stack Exchange, they have evolved into communities of their own, which allow for answering of questions that don't have such clear answers and allow for a significant amount of collaboration. Much previous research in the area has also shown that Q&A sites are online communities which involve groups of people communicating online, with communities being "long or short term, large or small, national or international, and completely or only partially virtual" (Preece & Maloney-Krichmar, 2003). The primary activities on social Q&A sites fit into the definition of collaborative information seeking, which includes collaborative grounding, or when two or more collaborators "construct shared understanding of the information need and seeking process" (Gazan, 2010). This occurs through comments on questions, comments on answers, up voting, down voting, etc. These collaborations are then made permanently available on the web for others to use as a knowledge base. More information on online communities and collaboration will be provided in the following literature review.

Social Q&A sites typically will have either one forum for questions on any subject or a forum for each specific subject. Stack Exchange exists as a network of sites made with the same boilerplate format, each pertaining to its own subject. As of August 2015, there are 152 different

sites in the Stack Exchange network, which are all designed with the exact same format. An advantage of this is that any group that is interested in creating a social Q&A site on a topic can start one.

The way that Stack Exchange is currently designed, there is little consideration for the differences in interactions that might occur between groups that are seeking information or answering questions on different topics. For example, do certain domains require more updating and occur over a longer period of time? Is the new information relevant to the original answerer? Are users more likely to require clarification on a question or an answer within certain domains? As collaboration is a vital aspect of social Q&A and providing high quality answers to complex questions, developing a deeper understanding of it can help lead to tools that better support it.

From this motivation, the main research question will be answered as well as several sub-questions:

Do different Stack Exchange sites lead to different distributions of collaborative acts?

- What are the differences in the distributions of collaborative acts?
- What design changes may be made based on the distribution of collaborative acts that could improve collaboration?

By answering these questions, the researcher seeks to question the efficacy of using the same site design across many, varied domains within Q&A. Designers may use the information found to devise novel interaction mechanisms or better integrate common existing mechanisms for collaboration into each domain's site. This research will extend the work done by Tausczik, Kittur, and Kraut (2014). In their study of collaboration on MathOverflow, they identified five categories of collaborative acts: provided information, clarified the question, critiqued an answer,

revised an answer, and extended an answer. These were used to analyze posts on the website and determine if there was a model that could predict how much the rating of the answer would change followed a collaborative act as well as the frequency of each of the five categories. This study examines several different Stack Exchange sites, including Stack Overflow, the site focusing on programming, Parenting, and Home Improvement. A study of the frequency distributions resulting from this content analysis is then used to call into question qualities of the current Stack Exchange design and provide design recommendations.

Organization of Thesis

This thesis is structured as follows:

- Chapter 2 provides background information on the history of Q&A, describes social Q&A, and discusses related aspects of computer supported cooperative work and collaborative information seeking that provide the theoretical base for this research.
- Chapter 3 describes the Stack Exchange network, research design, the data collection procedure and how rigor was established through multiple initial coders, as well as how the data was analyzed.
- Chapter 4 will report the results of the tests to compare the distributions of collaborative acts across the sites.
- Chapter 5 discusses the findings from this study, including an analysis of the collaborative acts observed, a comparison of the communities, design recommendations, and lastly study contribution and study limitations.

- Chapter 6 summarizes the research questions and method, summarizes contributions and limitations, and provides areas for future work.

Chapter 2

Literature Review

The following section will provide background information on the history of social Q&A, how the field of Computer Supported Cooperative Work (CSCW) and, within that, Collaborative Information Seeking (CIS) applies to this research's investigation into social Q&A, and the current state of social Q&A.

2.1 Q&A Services

Online question and answering services (Q&A) have been around for decades. Online Q&A exists as a way for user to ask questions online and then receive a response online from either his or her peers, or an expert. Three different types of services exist that all approach online Q&A in a different way, leading to different interactions as well as quality. Those include digital reference services, expert services, and social Q&A, the topic of this research (Shah, Oh, & Oh, 2009). The following subsections will describe these services and their essential differences.

2.1.1 Digital Reference Services

Digital references services, also known as virtual reference services, are those that utilize professional reference librarians to provide answers to the service users. Lankes (2004) defined digital reference as the “use of human intermediation to answer questions in a digital environment”. Human intermediation from a reference librarian is the important differentiator

from social Q&A; it means that a professional is helping to find the answer to the question. These are one-on-one interactions and are an extension of the functionality of a traditional reference librarian. Popular digital reference services include Ask an IPL Librarian and the Educators' Reference Desk (Shah, Oh, & Oh, 2009). Also, archives of digital reference service questions and answers are not always made publicly available, making it difficult to use the services to simply search for information (Pomerantz, Nicholson, Belanger, & Lankes, 2004). There is very little social aspect to digital reference services.

The study of digital reference services has been limited; with much of the research conducted analyzing a specific service. To address this, Marilyn White has proposed a framework for evaluating individual digital reference services as well as multiple digital reference services, providing a standard. The framework is based on mission and purpose, structure and responsibilities to client, core functions, and quality control (2001). Digital reference services were also studied in relation to social Q&A services. Hybrid services that utilize the expertise of librarians and the power of the crowd when necessary have been examined due to the different purposes that the two types of Q&A serve (Shah & Kitzie, 2012).

2.1.2 Expert Services

The second type of Q&A service are called Expert services, sometimes referred to as Ask-An-Expert services, Ask-A services, and expert Q&A. Through these services, those seeking information consult experts in the subject area, through both commercial and noncommercial organizations, for their answers. Expertise can be either proven through providing the appropriate degrees and experiences, or sometimes can just be self-declared (Shah,

Oh, & Oh, 2009). Janes, Hill, and Rolfe conducted a study to analyze the effectiveness of expert services (2001). They found that there was a very high response rate to questions on expert services and the answers had high reliability. However, some sites did not have high response rates and the questions that were answered were mostly factual type questions with verifiable answers because experts found these to be the quickest and easiest to answer.

These services may also be free or for pay because users are willing to pay various amounts for the depth or quality of the response. It has also been shown that culture and type of question have an effect on the amount of payment. Culture may dictate when and how much reward should be given for the answer provided. Those who are asking questions that are meant to stimulate conversation or ask for opinions may not give any reward for providing answers. With factual questions, it can be seen as an exchange of goods (information) and thus worthy of reward (Hsieh, Kraut, & Hudson, 2010). Again, because they consist of users consulting an expert, expert services lack an essence of community and gathering information from that community.

2.1.3 Social Q&A

Social Q&A has been studied as a tool for education and learning because it provides answers to questions that are complex and possibly difficult to search for. In order to answer user's questions, social QA repositories can be browsed for paraphrases by preprocessing to turn up existing questions and then present that answer to the user (Bernhard & Gurevych, 2006).

In the past decade, there has been research into better automating question and answering services, referred to generally as QA services, whether it is for automatically generating the

answer from scratch or presenting the best answer to the user (Chen, Subramanian, & Brewer, 2010; Liu, Agichtein, Dror, Gabrilovich, Maarek, Pelleg, & Szpektor, 2011; Shah & Pomerantz, 2010). QA services that create answers are best at responding to single-sentence, fact-based questions that are often answered from a repository of documents, supply no metadata, and provide immediate answers. This is opposed to social Q&A sites that search through existing answers, in which questions and answers can consist of many sentences, are answered from users' knowledge, have varying quality (Bloom & Kurian, 2011). Understanding what the user is asking for is a challenge that may require some input from the user in order to provide clarification (Shah, Radford, Connaway, Choi, & Kitzie, 2012).

For this research, social Q&A services are viewed as communities where there is an exchange of information between askers, answerers, and other community contributors, who all have the ability to comment and up vote or down vote answers based on their quality. Mendes Rodrigues and Milic-Frayling view user intent on social Q&A as fitting into two classes, those with social and non-social intent, for each question. Questions about "social engagement" being termed are classified as social and information, advice, or opinion questions as non-social (2009). They created a typology of questions that described a personal/general perspective, community/individual issue, and social/non-social intent. After analyzing 2000 questions from MSN QnA it was found that 30.5% of the comments of that data set were from social questions, even though those were only 6.5% of the total questions asked. This could indicate that the MSN QnA is supportive of a social atmosphere. This study did not address the intent of the comments themselves. Other researchers have typified questions in a similar manner (Harper, Moy, & Konstan, 2009; Harper, Weinberg, Logie, & Konstan, 2010). Harper et. al. (2010) identified the additional classifications of advice, identification, (dis)approval, quality, prescriptive, and

factual.

Research that discusses community aspects of social Q&A sites has been lacking, even though it seems to be widely accepted that these sites do foster a community (Rosenbaum & Shachaf, 2010; Anderson, Huttenlocher, Kleingberg, & Leskovec, 2012). Rosenbaum and Shachaf assume the definition of online community taken from Preece and Maloney-Krichmar (2010) as “social activity that involves groups of people interacting online” and, from de Moors (2006), online communities “build up a collective history of information created, discussions conducted, tasks performed, and goals accomplished. Especially in virtual communities, traces of these activities can be found in the many technologies used, like mailing list archives, web pages, and document repositories”. Multiple researchers have found higher levels of reputation within a community to be highly correlated with participation and quality of answers (Nam, Ackerman, & Adamic, 2009; Tausczik & Pennebaker, 2011). Anderson et. al. chose to examine all of the answers that are provided to a question on Stack Overflow by temporal aspects as well as reputation in order to predict long-term archival value of an answer (2012).

Several studies have examined usability of online communities for information sharing from the perspective of usability (effectiveness and efficiency) and sociability (how systems support social interactions for the community good). Phang, Kankanhalli, and Sabherwal (2009) used two surveys, one on knowledge seeking and one on knowledge contribution, to compare knowledge seeking and knowledge contribution through usability and sociability. They found a significant difference in how important users view sociability in online communities when contributing knowledge rather than seeking knowledge.

Other research has examined the types of users of social Q&A services (Shah, Oh, & Oh, 2008; Gazan, 2006). Shah, Oh, and Oh characterize users into two groups, consumers and

contributors. Consumers are those that ask questions and use the answers as a source of information. Contributors in addition provide information by answering the questions. Gazan additionally classifies contributors as either synthesizers or specialists. Synthesizers combine information from multiple sources without presenting it as their own and specialists are thought to already have the knowledge and provide the information on their own. In a study of Answerbag, the first official social Q&A service in the United States, he found that synthesizers' answers were more highly rated than those of specialists.

The practices of social Q&A sites also varies from user to user. According to the results of a survey conducted in 2006 of Korean users of Korean Q&A services, "an overwhelming majority of respondents were using question-answer services just for searching and finding knowledge purposes, not making comments or creating knowledge." In fact, 97.7% of respondents said that their main use was searching for knowledge. Only around 20% of users had posted comments at least one time (Lee, 2006). These results are similar to a study in 2007 by Gongyi, Koutrika, Pederson, and Garcia-Molina who examined a different service, Yahoo Answers, and found that of the registered users, over 70% ask questions, only 27% answer questions, and 18% have voted on an answer. This shows that the amount of users that are involved in the community by contributing information is relatively small on these services (2007). These percentages would be much smaller if the unregistered users, who were using the services as a knowledge base, were also included.

Yahoo Answers is split into categories rather than separate communities, but these can still be viewed in a similar manner. Certain categories lead to different types of interactions on the site (Adamic, Zhang, Bakshy, & Ackerman, 2008). Adamic et. al. clustered threads from different categories based off of observations that they had made. They ended with three groups

for the categories based on three characteristics: those that have high proportions of users who pose and answer questions, those who seek and provide advice, and those with factual answers that tend to have shorter thread lengths.

The focus of this research is social Q&A, which takes advantage of the knowledge of the billions of people who are now using the Internet every day. Gazan identified future areas of research for social Q&A. Much research has been done examining the largest of the social Q&A sites, Yahoo Answers. This is only one site and it has a complicated structure of rules and rewards as well as a specific organizational structures. Future research should include an examination of different Q&A sites in order to compare the findings from Yahoo Answers (2011).

The remaining literature review sections will discuss communities of practice as well as the areas of research called CSCW and CIS.

2.2 Communities of Practice in Social Q&A

An important aspect of Stack Exchange is that each of the sites is dedicated to a specific topic. This research views the communities that pertain to a certain knowledge area as communities of practice. Communities of practice (CoP) are defined as “groups of people who share a concern or a passion for something they do and learn how to do it better as they interact regularly” (Wenger, 2006). Wenger, who created the idea of communities of practice, asserts that there are three main characteristics: domain, community, and practice. Each site on Stack Exchange is created revolving around one specific domain. Users who contribute to that site through commenting, asking, or answering questions are members of a community. Lastly, as

users do not simply have a shared interest, they are actively seeking information or contributing information, showing that they are actively involved in the practice of that domain.

Previous research has also viewed social Q&A services as communities of practice (Rosenbaum & Shachaf, 2010). Rosenbaum & Shachaf viewed CoP's in Q&A through the lens of structuration theory, with the actions of the members of that community being constrained by the structure of the Q&A site itself. In fact, Wenger has said that the idea of CoP's draws on structuration theory. As individuals interact with a structure, they and the structure are continually adapted simultaneously through interaction. The practices of online communities of practice have an effect on the structure of the site, which includes how the norms and functions of the site, as well as the structure, confine the practice of the community. Within the context of Stack Exchange, each community forms a specific structure consisting of social norms and accepted practices which is mutually affected by the behavior of the users.

CoP's exist in different context, whether it is in organizations, in person, or online. Online CoP's have differences from those that exist in the work environment (Baker-Eveleth, Sarker, & Eveleth, 2005). Online CoP's use text as the primary form of communication. Dialogue in these contexts becomes an artifact of the community and enhances articulation, rather than simply existing between two people. Members can make public corrections of previous statements or respond to those statements due to the existence of a historical record.

2.3 CSCW and CIS

The following subsections briefly describe computer-supported cooperative work and collaborative information seeking as it applies to the current research.

2.3.1 Computer-Supported Cooperative Work

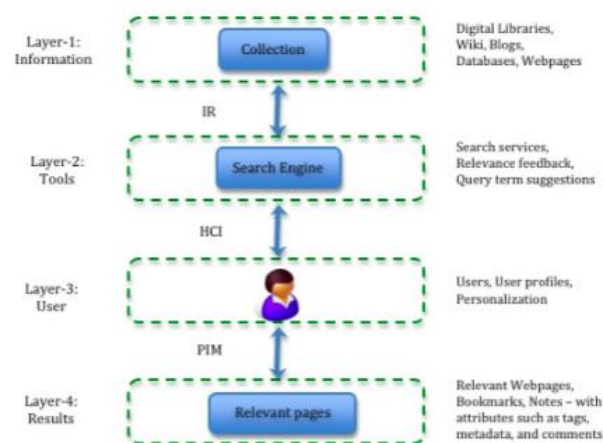
Social Q&A forums can be viewed as a setting for computer-supported cooperative work. First it is important to understand what CSCW is. Schmidt & Bannon define CSCW as “an endeavor to understand the nature and characteristics of cooperative work with the objective of designing adequate computer-based technologies” (1989). This definition makes for a very broad conceptualization that has resulted in debate about what the focus of the research should be (Schmidt & Bannon, 1992). While some researchers prefer to examine the social aspects of cooperative work, others, such as the developers of the tools, prefer to examine the technical aspects. A term that is often used interchangeably with CSCW is Groupware, however they can be considered as two separate things. While CSCW is a field of research, Groupware is defined as “an *application* or set of *tools* that covers needs such as communication, cooperation, and coordination” (Romero, 2011). This research takes the perspective that social Q&A sites can be viewed as Groupware and be studied in the context of CSCW.

A groupware application that shares some characteristics with Stack Exchange is GitHub. Dabbish, Stuart, Tsay, & Herbsleb studied the social aspects of GitHub, a collaborative and open code repository (2012). A grounded approach was used to analyze 24 semi-structured interviews with frequent and infrequent users of GitHub to understand how it was used. Many of the actions were based around the visibility of different aspects of the project status. Users may view projects that feature a lot of comments on code “commits” as an opportunity to make a contribution. “Managing reputation and status” was also a large category that was identified. Actions being made visible causes users to consider that their actions may have lasting results in the project, as well as see it as an opportunity to build their own reputation in the community.

2.3.2 Collaborative Information Seeking

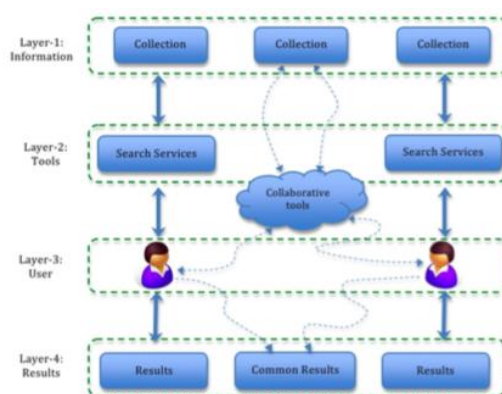
In CSCW research, CIS applies when computer-based, cooperative technologies apply to two or more people seeking information. Shah (2009) presents a model for information seeking that's consists of four layers: information, tools, user, and results. The information layer can come from one or various sources. The user in layer 3 can utilize tools and techniques in order to help him or her find the information that they need. The last layer, results can include knowledge that the user gains. The Four Layer model can be seen below in Figure 1.

Figure 1. Four Layer Model of Information Seeking (Shah, 2009)



The researcher then extends the theory in order to describe collaborative information seeking, as seen below in Figure 2. In this model, multiple users share a collection of results and are also sharing a shared source of information and possibly their own sources of information. They use a collaborative tool to organize, collaborate, and make all of the common results visible. Collaboration according to Shah involves creating a solution that is more than just the sum of the individual efforts.

Figure 2. A Model for Collaborative Information Seeking (Shah, 2009)



Golovchinsky, Qvarfordt, and Pickens use two different personas in order to describe the forms and roles of CIS. The forms of collaboration include explicit and implicit intent, depth of mediation, concurrency, and location. These four characteristics describe the types of computer support for CIS. The researchers also describe roles that people can play during collaboration. These include peer, domain A expert/domain B expert, search expert/search novice or domain expert/domain novice, search expert/domain novice, and prospector/miner (2009). These descriptions can largely be applied to social Q&A, with different users having different knowledge and skills using the same system in order to build a knowledge base. A 2012 study conducted by Moskaliuk, Kimmerle, & Kress examined the reasons for contributing to shared knowledge bases in the area of wikis. An incongruity in the shared knowledge base and the individual's knowledge can lead him or her to making a contribution, especially when there is a medium level of redundancy between what is found on the wiki and his or her own knowledge.

In the area of asynchronous, electronic discussion groups, which could include social Q&A threads, it has been shown that having roles can be important (De Wever, Van Keer, Schellens, & Valcke, 2010; Jonassen & Kwon, 2001). De Wever et. al. identifies that roles, such as moderator, source searcher, and summarizer, can be beneficial in the beginning of the

discussion and then become gradually less important (2010). In the beginning of Q&A threads, there is a question posed, followed by an answer, and then comments are made related to that answer. These can provide some roles such as answerer, critic, or supporter, for example.

Microcollaborations as defined by Gazan (2010) are collaborations that involve a “limited quantity of information exchanged between collaborators, and the casual and informal nature of the interactions, many of which can be maintained simultaneously.” The collaborators must share mutual interest and mutual effort. Gazan conducted a content analysis of 816 “expressions of mutual interest” on Answerbag that can be seen as possible instances of microcollaboration. According to the researcher, to qualify as a microcollaboration there must be a high enough complexity, as well as both facts and opinions involved. He also identified that users do not always follow the sites guidelines on what should be included in comments and answers, and noted a high level of social interactions when microcollaborations occurred. Lastly, he suggests that “a hybrid view of collaboration is warranted, one that takes into account embedded technical infrastructures and affordances, affective needs, the social dynamics of the community, and the information around which collaboration happens.” Collaborative acts are the processes that are involved in the microcollaborations occurring, and are the focus of this research (Tausczik et. al., 2014).

Chapter 3

Research Methodology

The following chapter will describe the research methodology, including a description of Stack Exchange, research approach, sample population, sample selection procedure, data collection procedure, and data analysis methods. Lastly, the research validity will be described.

3.1 Stack Exchange Description

Stack Exchange is a social Q&A site that allows users to search or browse questions, provide answers to questions, up vote or down vote answers, and comment on questions or answers. They can also edit questions or answers. Stack Exchange differs from sites that combine all topics onto one website. Stack Exchange is a federated network of social Q&A sites that are each about a specific topic. Each of those topics form their own community and can operate a meta-level site, where the protocol of managing the sites can be established as well as questions about the site can be posed. There are currently 152 different sites on the network that each cover their own topic, with many more listed in the Area 51 section, which are topics that need to gain more users in order to be moved out of that area. As seen on the All Sites page, the most popular site on the Stack Exchange network is also the original site, Stack Overflow. Stack Overflow has 11 million questions and 5.3 million users. A sample of one of the questions used in the study can be seen below in Figure 3.

Figure 3. Example of Stack Exchange Question

When should I not use WD-40?


▲ Looks like WD-40 can be used for lubricating just about any metal parts.


82 Are there non-obvious cases where using WD-40 will be harmful to the stuff lubricated?

▼ lubrication

★ share improve this question

17

edited Jul 28 '11 at 14:00  Aarthi 842 🗳️ 3 🗳️ 18 🗳️ 28

asked May 10 '11 at 14:04  sharptooth 6,383 🗳️ 7 🗳️ 35 🗳️ 73

We're looking for long answers that provide some explanation and context. Don't just give a one-line answer; explain why your answer is right, ideally with citations. Answers that don't include explanations may be removed.

44 Don't combine WD-40 and Duct Tape, you risk creating a black hole. – Adam Prax May 10 '11 at 16:00

1 A lighter seems to be a bad idea to "lubricate". – Spechal May 11 '11 at 4:46

15 Also, it's not for ... *personal* use. – Jay Bazuzi Mar 11 '12 at 1:10

3 By the answers given, it's clear that you asked the exact opposite question of the one you should have asked, to wit: "When should I use WD-40?" So far, we've learned that it's not a lubricant, it will attack plastic, lacquer (insulation on magnet wire in a motor), etc. It's really just for cleaning water off of things, and happens to work as a weak penetrating lubricant. It's kind of like a Swiss Army Knife in that you could use it for all sorts of things, but you really shouldn't. Use the proper lubricant for the job and you'll be happier in the long run. – Andrew Beals Mar 12 '12 at 19:09

9 Perhaps this will help: i.stack.imgur.com/EzXXb.jpg – oscilatingcretin Mar 14 '12 at 18:13 ✎

WD-40 is great for cleaning Stainless Steel as well. Apply lightly and buff off. – Dutchie432 Mar 20 '12 at 15:21

Stack Exchange displays the value of the number of up votes minus the number of down votes, which in this example is show as 82. The comments on the questions are shown in chronological order. The question also displays user-generated tags that describe the content of the questions. These tags help to generate search results, allow users to receive alerts when questions are posted with that tag, and allow them to browse questions with that tag.

Answers on Stack Exchange have a similar appearance. The answer that is selected by the original asker as correct is placed below the question and its selection is represented with a green checkmark. The answer to the example question is shown below in Figure 4.

Figure 4. Example of Stack Exchange Answer



▲ 46 ▼ ✓

1. Firearms - you should never use WD-40 to clean or maintain firearms. It's hygroscopic and will attract moisture to the firearm which will result in rust.
2. Drive chains of any type - because WD-40 isn't a lubricant it really will not work well on Drive chains of any type.
3. Gears of any type - WD-40 doesn't have enough lubrication affect to be useful on any type of gearing. only specified gear lubricants should be used.
4. Locks. Use powdered graphite or silicon spray instead.

share improve this answer

edited Mar 14 '12 at 17:42 answered May 10 '11 at 19:21

 Vebjorn Ljosa 6,339 12 35 76

 nevadaexile 484 4 3

1 WD-40 is a lubricant. – xpdo Mar 17 '12 at 15:15

9 No it's not, really. See other answers. – tomfanning Mar 18 '12 at 21:51

3 @tomfanning "WD-40 cleans/degreases, penetrates to loosen up stuck parts, prevents corrosion and is a light lubricant." wd40.com/faqs There are several more mentions of WD-40 being a lubricant on the site. It contains mineral oil, which is a light lubricant. – xpdo Mar 19 '12 at 4:08

1 I've seen locksmiths use WD-40 in locks. (I can't vouch for their being *good* locksmiths though. The one that changed my ignition didn't rekey it to match my keys.) – aij Sep 13 '14 at 20:57

Stack Exchange implements reputation privileges in order to maintain the quality of site. These are built on reputation points, which are meant to approximate the communities trust in the knowledge of the user. Each user begins with one reputation point. Points can be earned by contributing to the community or lost by negatively contributing. Some examples of ways to earn points are having a question or answer voted up, having an edit accepted, having an answer accepted, and accepting an answer to question posed. Some examples of ways to lose points are voting a question down, having an answer posted be voted down, or receiving 6 “spam” flags. The main privileges that users can earn at varying reputation point levels according to Stack Exchange are represented and described in Table 1.

Table 1. Reputation Points and Privileges

Reputation	Privileges	Description
1	Create posts	Ask a question or contribute an answer

5	Participate in Meta	Discuss the site itself: bugs, feedback, and governance
15	Vote up, flag posts	Indicate when questions and answers are useful Bring content to the attention of the community via flags
20	Talk in chat	Participate in this site's chat rooms
50	Comment everywhere	Leave comments on other people's posts
125	Vote down	Indicate when questions and answers are not useful
1,000	Create gallery chat rooms	Create chat rooms where only specific users may talk
1,500	Create tags	Add new tags to the site
2000	Edit questions and answers	Edits to any question or answer are applied immediately

3.2 Research Approach

The research approach was informed by the methodology used in Tausczik et. al. (2014). This approach involves using existing artifacts on Stack Exchange in order to understand the interactions occurring. Each post was qualitatively coded for each question, correct answer, edit, or comment that was made, based off a schema of the five collaborative acts that were identified in the taxonomy by Tausczik et. al. (2014). These included *provided information*, *clarified the question*, *critiqued an answer*, *revised an answer*, and *extended an answer*. Any interactions that did not fit into the schema or significant aspects that were not previously recorded were added to the schema in a semi-open coding manner, and then iteratively applied to previously coded data.

An initial coding session was conducted with 11 other students and 2 professors, who provided input on possible codes to consider and things to think about. This gave the researcher the ability to interpret the wording and code properly for any nuance that was present. The selection methodology for the questions that were analyzed will be discussed in the following section.

To answer the question of whether there was a significant difference in the distributions of collaborative acts, a quantitative approach was used based on the codes that were established in a schema based off of previous work by Tausczik et. al (2014). A Pearson's chi-squared test was used to test for a difference in the distributions of collaborations between the three different sites. For this test, the hypotheses were as follows:

H_0 : The data from each site follows the same distribution

H_a : The data from at least one site does not follow the same distribution

Pair-wise tests were then conducted to compare differences between the three combinations of the three different sites. Observing the differences in collaborative acts across the sites can then lead to site-specific design recommendations in order to support collaboration in that community.

This approach has also been affected by the level of analysis. Grouping by site and classifying each contribution into collaborative acts places the focus on aspects at the community level. The goal of this research is to compare the differences across sites and not to compare individual actions and motives, so the approach used minimizes the focus on the individual. Collaborative acts that come from the same site are all viewed collectively and only by the classification in the schema.

3.3 Sample Selection

For this research, a subset of the 152 different sites was selected. The sites that have been selected are Parenting, a forum largely based around parenting advice, Home Improvement (DIY), a forum about tools and completing do-it-yourself projects, and Stack Overflow, a forum for specific programming questions. These sites were chosen for two reasons. The first is that they reach across several different demographics and types of questions and interactions. Parenting-type questions come from parents and the site has a casual atmosphere for parents to seek advice about parenting. Home Improvement is an inherently amateur site that often requires technical answers but can be used by professionals. Stack Overflow ranges from amateur to professional computer programmers who require technical, factual responses. These contrast with the site MathOverflow from Tausczik et. al. (2014), which is used by professional mathematicians who are seeking collaborators on scientific problems. By comparing varied sites, the research expected to see differences in the distributions that occur.

The second reason was more pragmatic. Because the threads were being manually coded, it was important that the researcher understands what is being discussed in order to properly interpret the interactions. Parenting, Home Improvement, and Stack Overflow are subjects that are understandable based on the researcher's previous experiences.

From each of the 3 sites that were selected for sampling, 5 questions were selected. These questions should contain significant amounts of collaboration in order to allow for the different types of collaboration to occur and provide a large amount of instances of the codes from each question. The questions on each site were sorted by frequency and the first of the questions to

reach the criteria were chosen. To be selected, the questions must have met the criterion that there must be at least 10 comments on either question or top answer combined.

This criterion was selected in order to provide minimum evidence that there is some sort of collaboration occurring in the thread so that they can be coded meaningfully based on the collaborative acts.

The questions that were selected, as well as the site that they came from and the code that will refer to them, are listed in Table 2 below.

Table 2. Questions Sampled

Social Q&A Site	Question	Code
Stack Overflow	Explain Python's slice notation	SO1
	How to avoid Java code in JSP files?	SO2
	Why is iostream::eof inside a loop condition considered wrong?	SO3
	Should I avoid the use of set(Preferred Maximum Minimum)Size methods in Java Swing?	SO4
	What is the difference between client-side and server-side programming?	SO5
Home Improvement (DIY)	Should I use steel or wood studs for basement exterior walls?	DIY1
	How to prevent heavy rainwater from jumping the gutter?	DIY2

	Is there an easy way to measure the height of a tree?	DIY3
	When should I not use WD-40?	DIY4
	How do I make a height adjustable desk?	DIY5
Parenting	Is bribing children with cash incentives a good idea?	P1
	What's wrong with Dr. Seuss?	P2
	My 4-year-old refers to himself as me	P3
	Is it safer to install a back seat car seat behind the driver or behind the passenger?	P4
	How do you teach a child religious views?	P5

3.4 Data Collection

For each of the 5 questions that were selected on each site, a document was prepared that chronologically listed all questions, answers, comments, and edits, which will be referred to as posts. In order to do this, a screen shot was taken of each post, and then ordered based on the time stamp. Each post was assigned a number, in order to make them easier to reference as well as to compare codes while establishing inter-rater reliability. These documents range from 2-24 pages in length. The median page lengths were 3, 5, and 7 for Parenting, Home Improvement, and Stack Overflow, respectively. The total number of pages was 97. From the 15 questions, there were a total of 238 comments and 77 edits analyzed.

After these documents were created, an initial coding session was done in a large group on one of the questions. After their feedback was considered, the coding schema was adjusted. In order to gain more insight into the collaborations, there were also Stack Exchange specific items that were coded when they occurred (e.g. Picture, Question Protected, Question Unprotected).

Next, to establish inter-rater reliability, collaborative coding was completed with the help of a senior PhD candidate with a lot of previous experience in qualitative research. Using one question from each of the 3 sites, both researchers independently coded each post. They then, together, went through each post and discussed what they coded them as and deliberated until any discrepancies were resolved. This was vital in operationalizing each of the codes in the schema. Collaborative coding for a sample of the data helped to establish the accuracy of the individual coding.

In the study of MathOverflow by Tausczik et. al., each collaborative act was also given a score from 1-5 on the significance of the collaborative act. This was then used to weight the frequency when the collaborative acts on that site were compared. For this research, the quality of the collaborative acts is not being tested. To better represent the actual nature of collaborations that are occurring on the site, all of the collaborative acts were weighted equally.

Collaborative coding was completed iteratively until there was an 80% agreement, based on the percentage of matching codes from 20% of the data set. This process took the examination of 4 of the selected Stack Exchange questions before sufficient agreement was met. A minimum of one question from each site was chosen for this coding, in order to account for differences that may exist in one site and not another. The coding schema that resulted from the collaborative coding can be found below in Table 3.

Table 3. Coding Schema for Each Post

For each post:	Acronym:
Providing information	PI
Clarification	CL
Critique	CR
Revision	RE
Extension	EX
Confirmation feedback	CF
Positive feedback	PF
Metadiscourse	MD

The codes that were used for each of the collaborative acts are defined in the following way:

Providing information. Provided information not directly extending previous contribution

Clarification. Clarified the contribution

Critique. Evaluates or questions the contribution

Revision. Suggests solution or gap for something wrong with the contribution

Extension. Provides new information to extend the contribution

Confirmation feedback. Comments reaffirming the correctness of the contribution

Positive feedback. Comments that compliment a user on contribution or usefulness of the contribution

Metadiscourse. Discussion not directly related to contribution but shows presence in the community

Once that process was complete, each question could be coded individually. As the collaborative coding was an iterative process and codes had changed from the beginning of the process until the schema in Table 2 was established, each of the previously examined questions had to be re-analyzed along with the unexamined questions.

Chapter 4

Findings

The following section will go into detail about the statistical results that came from the data. First the results of analyzing the three sites together will be reported and then pairwise comparisons of those distributions.

4.1 Analysis of 3 Sites


After coding each of the questions on the three sites, the distributions of collaborative act frequencies can be analyzed. The distributions of the acts for each site can be found below in Table 4.

Table 4. Distribution of Collaborative Acts

Count		Type								Total
		CF	CL	CR	EX	MD	PF	PI	RE	
Site	DIY	7	15	10	3	19	16	26	11	107
	P	3	9	15	2	8	20	25	10	92
	SO	4	25	26	10	7	8	23	45	148
Total		14	49	51	15	34	44	74	66	347

There were a total of 347 collaborative acts that occurred on the 15 questions. 148 of those were derived from Stack Overflow, 107 from the DIY site, and 92 from the parenting site. This data is represented visually using a histogram below in Figure 5.

Figure 5. Histogram of Collaborative Act Frequencies



In order to compare the distributions, of the three sites, a Pearson chi-squared test was conducted. The chi-squared test can be used to analyze frequencies of categorical data by comparing values to a theoretical distribution. The null hypothesis is that the given frequency distributions come from the same distribution, and the alternative hypothesis is that they do not come from the same distribution.

H_0 : The data from each site follow the same distribution

H_a : The data from at least one site does not follow the same distribution

This can alternatively be represented by using D to represent the distribution, SO for parenting, DIY for DIY, and P for parenting. The alternate representation of the null and alternative hypotheses are then as follows.

$$H_0: D_{SO} = D_{DIY} = D_P$$

$$H_a: \neg (D_{SO} = D_{DIY} = D_P)$$

In order to use the Pearson chi-squared test, the data must meet certain assumptions. The first assumption is that the data is independent, or not correlated. As the data were all taken from separate sites, this assumption has been met. The second assumption is that there is a large enough sample size. This requires that 80% of the data has to have an expected value of at least 5. For the data of the three sites, there were 4 out of 24 frequencies that were less than 5, meaning that 83.3% of the data were greater than 5, meeting this assumption.

A chi-square test of independence was performed to compare the distributions of collaborative acts between Stack Overflow, DIY, and Parenting, using an alpha level of .05. We observed a strong association between the site and the distribution of collaborative acts, $\chi^2(14) = 56.462, p = .000$. This small p value indicates a rejection of the null hypothesis that the distributions are the same across the sites in favor of the alternative hypothesis that the distributions of collaborative actions are not the same across Stack Overflow, DIY, and Parenting.

4.2 Pairwise Comparison of 3 Sites

Next pairwise comparisons were completed using Pearson's chi-squared test to find out which of pairs of sites also did not come from the same distribution and could be causing the small p-value in the previous hypothesis testing. Three pairwise comparisons were conducted comparing Stack Overflow and DIY, Stack Overflow and Parenting, and Parenting and DIY. Applying the Bonferroni Correction, which adjusts the p value to α/n where n is the number of comparisons to avoid the multiple testing problem, results in a new p value of .017. The pairwise comparison results will be reported in the following subsections.

4.2.1 Comparison of SO and DIY Frequencies

The assumption of independence of the two sites is still sufficient. For these two sites there is one expected value that is less than 5, or 6.3%, meaning that 93.7% of values are greater than 5, satisfying the minimum sample size assumption. There are 255 valid cases. For this test, the null hypothesis is that the distribution of collaborative acts on Stack Overflow is from the same distribution as the distribution of collaborative acts of the DIY site. The alternative hypothesis is that the distributions of collaborative acts on Stack Overflow and DIY do not come from the same distribution.

We observed a strong association between the site and the distribution of collaborative acts, $\chi^2(7) = 37.610, p < .000$. The p value below .017 indicates the rejection of the null hypothesis in favor of the alternative hypothesis, that there is a difference in the collaborative acts across Stack Overflow and DIY.

4.2.2 Comparison of SO and P Frequencies

The assumption of independence of the sites is still sufficient. For these two sites there are three expected values that are less than 5, or 18.8%, meaning that 81.2% of values are greater than 5, satisfying the minimum sample size assumption. There are 240 valid cases. For this test, the null hypothesis is that the distribution of collaborative acts on Stack Overflow is from the same distribution as the distribution of collaborative acts on the Parenting site. The alternative hypothesis is that the distributions of collaborative acts on Stack Overflow and Parenting do not come from the same distribution. We observed a strong association between the site and the

distribution of collaborative acts, $\chi^2(7) = 32.209, p < .000$. This indicates the rejection of the null hypothesis in favor of the alternative hypothesis, that there is a difference in the collaborative acts across Stack Overflow and Parenting.

4.2.3 Comparison of DIY and P Frequencies

The assumption of independence of the sites is still sufficient. For these two sites there are three expected values that are less than 5, or 18.8%, meaning that 81.2% of values are greater than 5, satisfying the minimum sample size assumption. There are 199 valid cases. For this test, the null hypothesis is that the distribution of collaborative acts on the DIY site is from the same distribution as the distribution of collaborative acts on the Parenting site. The alternative hypothesis is that the distributions of collaborative acts on DIY and Parenting do not come from the same distribution.

We did not observe a strong association between the site and the distribution of collaborative acts, $\chi^2(7) = 8.209, p = .315$. This p value greater than .017 indicates that we do not reject the null hypothesis, and maintain that there is not a significant difference between the distributions of collaborative acts between the DIY and Parenting Stack Exchange sites.

Chapter 5

Discussion

The following chapter will provide examples of collaborative acts on the Stack Exchange that were used in the schema, discuss the three communities and the comparison of their collaborative acts, present design recommendations that could be beneficial to the communities, and lastly discuss contributions and limitations of this research.

5.1 Collaborative Acts on Stack Exchange

Each of the comments and edits of the threads on Stack Exchange were able to be coded to the schema shows in Table 3. The following section will provide examples from the communities of each type of collaborative act.

First, “Provides Information” was used in attempt to answer the question but without extending a previous contribution. Providing information was common amongst all three communities, which is fitting of the fact that Stack Exchange is a question & answer service, so new information should be provided in all threads.

[SO1: PI] “I don’t think anyone’s mentioned that None is a valid slice value. I find it’s handy used as a default argument. Eg: crop = lambda l, n=None: l[:n] then crop([1,2,3],) returns [1,2,3] – note that no integer default arg would work here.”

[P2: PI] “It is possible that their objections may be political, rather than about anything to do with the developmental appropriateness. Or maybe they just had to read Oh Say Can You Say? one too many times...”

“Clarification” was used by either the original contributor to clarify an answer they provided or by somebody else to provide, or ask for, clarification for a previous answer.

[SO2: CL] “@ [] Why doesn’t make sense in RESTful web app? Could you explain what you mean? Thanks (I’m a starter as you mentioned)”

[DIY4: CL] “@ [] “WD-40 cleans/degreases, penetrates to loosen up stuck parts, prevents corrosion and is a light lubricant.’ Wd40.com/faqs There are several more mentions of WD-40 being a lubricant on the site. It contains mineral oil, which is a light lubricant.”

“Critique” was used to evaluate or question the contributions, whether that was done in a positive or negative manner.

[DIY1:CR] “There are a few drawbacks to densarmor though: it’s itchy (fiberglass) and you might have to skim-coat it all to get a consistent finish to it (as it has a texture that you won’t see in the taped joints).”

[P2:CR] “@ [] – Curious! But I have a feeling some of that may be the “I am the walrus” effect. “The Cat in the Hat Comes Back, in which the Cat may represent colonial or absolute dictatorial power”- that sounds too much of a stretch even to me, and I’m a life-long expert on progressive-subtext-reading.”

“Revision” refers to providing a solution or filling a gap with something that is wrong with a contribution. Within this category falls edits to the question and the answer, which include corrections on spelling, grammar, and tags. Revisions that included additional comment were not as common.

[SO3:RE] “edited tags: c++ iostream c++-faq”

[DIY1:RE] “@ [] DensArmor is paperless wallboard. Should work just fine for fire rating.”

“Extension” involves providing new information to a previous contribution. Extensions did not occur often but do show a high level of collaboration toward an answer because they require involving a minimum of two contributions. The bolded portion of the following example was added by a user to extend the accepted answer to the question.

[SO2:EX] “This was dealing with different result page destinations is easier: redisplaying the form with errors in case of an error (in this particular example you can redisplay it using $\{message\}$ in EL), or just taking the desired target page in case of success.”

[P5:EX] “This is actually a very well considered and thoughtful answer; as an atheist, the only thing I would add is mentioning that “no religion” is also a valid choice – i.e. in addition to exposing the child to people of religions, to introduce them to people who can talk sensibly about the non-religion option. Actively choosing a non-religious philosophy (such as Humanism, as just one example) can also very much be a boon (and liberation) to the individual.”

“Confirmation Feedback” specifically reaffirms the contribution. This type of collaborative act was not as common as “Positive Feedback”.

[SO2:CF] “@ [] I agree with you that this is the best practice approach. Especially with the question being ‘how to avoid’.”

[P3:CF] “I hope it’s normal, because my daughter does it too!”

“Positive Feedback” was more common and was coded when the comments complimented the contributor on their contribution or usefulness of the contribution.

[DIY5:PF] “I would suggest making a new question that links back to this answer (which is very nifty – BTW).”

[P4:PF] “I didn’t even know putting the seat in the center was allowed, let alone recommended! Great question.”

Lastly, “Metadiscourse” was used for comments that did not relate to the answer. These typically involved making a joke or discussion about the community itself, which demonstrates that the member is involved in the community and is a participating member of the thread.

[SO5:MD] “@ [], Haven’t seen a nice canonical one. If you know one with a good answer, please share.”

[P3:MD] “Perhaps he is developing an Oedipus Complex (Please don’t take this comment seriously, this is basically an excuse to prove that my A-level in psychology wasn’t a waste of time).”

5.2 Comparison of Stack Overflow, DIY, and Parenting

This section will compare and contrast the distributions of collaborative acts that were found across the 3 different Stack Exchange Network sites.

The Pearson chi-squared test for independence found that there was a significant difference in the distribution of collaborative acts for at least one site. This does not determine which of the sites is causing the significant result, so post-hoc tests were completed. The

distribution of collaborative acts on Stack Overflow significantly different from that of both DIY and Parenting. DIY and Parenting were not significantly different from each other.

None of the sites that were examined are specifically for professionals. Within Stack Overflow and DIY, there is a wide range of user demographics, from amateurs working on small projects, such as fixing a household item that is broken to master craftsman providing answers. With Stack Overflow, there are people working on their first programs as well as professional software developers who want to leverage the community to help solve a problem they cannot figure out on their own. Stack Overflow and DIY questions also vary from subjective to objective, but with most questions, it can be determined which answer is the best (e.g. the least lines of code or fastest processing). Parenting questions are often questions that are looking for advice on raising a child, and are mostly subjective. Askers may want to hear a variety of answers to inform their own decisions, or there may be multiple, equally correct or best answers.

One difference that should be noted that differs DIY from Stack Overflow and Parenting is the frequency of metadiscourse. Examples of metadiscourse varied from comments about the community, to jokes, to historical facts relevant to the question (e.g. “*Historical fact: that’s how Thales measured the height of the great pyramid in 7 BC ;-)*”). This difference may be explained by the casual nature of DIY, with a large population of hobbyists who enjoy doing do-it-yourself projects and the less time-sensitive nature of the questions.

The following section will describe the difference between Stack Overflow and DIY and Parenting as a pair.

5.3 Comparison of Stack Overflow to DIY and Parenting

The distribution of collaborative acts on Stack Overflow does not come from the same distribution as DIY or Parenting. In this section, differences in the distributions will be discussed and possible reasons for these differences.

The largest difference in the distribution is the large disparity in the number of revisions that occurred (SO=30.4%, DIY=10.3%, P=10.9%). Revisions largely consist of corrections to both spelling and grammar of the question and the answer. In a few cases, they included revisions to the meaning of the question or answer. Also, frequencies of critiquing (SO=17.6%, DIY=9.3%, P=16.3%) were higher in Stack Overflow and Parenting compared to DIY, and clarifications (SO=16.9%, DIY=14.0%, P=9.8%) in Stack Overflow were also higher compared to the other two sites, but especially higher than Parenting. Positive feedback is lower on Stack Overflow than DIY or Parenting (SO=5.4%, DIY=15.0%, P=21.7%).

There are three likely reasons for these differences. The first reason is that Stack Overflow is the original Stack Exchange site, as well as the largest by users by a wide margin. Higher traffic on the site means that there are more views per question and a more competitive environment for earning reputation. Making edits that are accepted is one of the ways to earn reputation, so fixing mistakes more quickly than other site members is one method for improving member status. Some edits that are made are eventually rolled back to the previous state in a future edit, while others may change the purpose of the question from more specific to more generalizable. A higher frequency of critiquing could also originate from the large size of the community for Stack Exchange or for their being more subjective-type responses on Parenting. With many more views per day on the site than the other sites, there is a larger user base that can

detect errors in provided answers. The DIY and Parenting sites have a smaller number of users, so it may be easier to gain reputation through other means than editing. This does not explain the lower frequency of positive feedback.

The second reason for the differences in distributions may be the importance that the site has as a reference or knowledge repository for future users. On Stack Overflow, questions are actively changed to be more searchable as a reference for other site users. Protecting the quality of the answers on the site is important in maintaining its popularity and wide user base. The topics that are higher than DIY and Parenting (revisions, clarifications) are all related to improving answer quality. The higher percentage of critiques can also be explained by this reasoning. Clarifications provide further detail for the users who reference the site in the future. Critiques are meant to directly improve the quality of the answers by providing alternatives. The other two sites also are important as references, but the smaller number of users may be inhibiting its importance as a reference site. The lower frequency of positive feedback may result from the view that if Stack Overflow is a reference site the only positive feedback reported should be in the form of up votes because comments on the user or usefulness of the contribution are not valuable for reference.

The third reason for the differences in collaborative acts may stem from what type of information is being discussed. When looking at Stack Overflow, much of the information that is being discussed is simply knowledge based (e.g. what code could perform this functionality or which method will be faster?) DIY questions will often relate to procedural knowledge such as how to build a certain piece of furniture or how to best divert water. Parenting knowledge will contain both of these types of information but also looks for subjective information. Users of Parenting look to receive feedback on questions that may have more than one correct answer.

Parenting questions may also have the potential to be a more emotional subject, where positive feedback can be more beneficial. The types of knowledge that are being sought across the varying sites would also likely lead to different collaborative activities. This is not limited to the three sites that were examined but could also be expected from any site that uses that type of knowledge.

The fourth reason is more subjective and does not describe all users of any of the sites. There is a difference in demographics and culture for each of the communities that develops from the history of the site as well as the subject area. Quantcast, a secondary source of web traffic data, provides the following relevant demographic information in Table 5.

Table 5. Demographics of Communities Based on Web Traffic

	Stack Overflow	DIY	Parenting
Gender	88% Male 12% Female	74% Male 26% Female	29% Male 71% Female
Largest Age Group	25-34 (27%)	35-44 (24%)	25-34 (34%)
Highest Education Level	College (49%)	College (50%)	College (47%)

As Stack Overflow is on the topic of software development, a large majority of the community would have a background in computer science or some field within information technology. Problems in this field are objective and there are right and wrong answers. Members of this community might be more willing to provide direct critiques of other answers as well as less positive feedback. Stack Exchange also has a large majority of male users. Within the parenting community, naturally most users are parents, with around 71% of the users being

women. Parents are likely more nurturing due to interactions with their children and provide positive feedback frequently to their children. This likely translates to the site where users provide positive feedback to each other. In DIY, projects are not typically time sensitive and users of the site are more casual. Users on DIY are also 74% male. This, combined with the reality that there are often multiple best solutions to the same problem, results in users that are less likely to negatively critique others and interact in a more casual way (e.g. jokes through metadiscourse).

5.4 Design Recommendations

Each site on the Stack Exchange network is created the same. This stems from its origin as a single topic social Q&A site, Stack Overflow. The network was quickly scaled up by applying the layout to other topics or sites. This study has shown that not all sites on the Stack Exchange network come from the same distribution of collaborative acts, based on an examination of Stack Overflow, DIY, and Parenting. Users are appropriating the sites in different ways and collaborating in different ways specific to the communities. Therefore, it is worth exploring potential changes to specific sites that will more closely match the types of collaborations that are occurring on Stack Exchange sites. The following section will provide several design recommendations relevant to one or more of the communities studied that will support their current nature of collaboration. These recommendations are meant for designers of Stack Exchange as well as designers of social Q&A sites that foster different individual communities.

The first design recommendation is a feature that created issues with the archival value of the information on the sites. Users communicate to each other in the comments by using the “@” symbol followed by their username. However, this does not actually create a link to that users account and it does not get updated when the user changes their name. For this reason, it is sometimes difficult to follow collaborations that are occurring on the threads, as the user to whom the comment is directed is unclear. Directing comments using a mechanism that actually links to a users account and updates automatically will make it easier to follow threads of conversation in the future.

Mechanisms for feedback to the contributors that exist on Stack Exchange currently include up voting and down voting a question or answer. However, this falls short in numerous ways. The assumption inherent in this mechanism is that feedback only occurs on questions or answers to the questions. In reality, all three communities studied have large amounts of feedback to each other in the comments to the questions or answers in the form of critiques and positive feedback statements. This can be expected of all communities on Stack Exchange, so addressing the issue of feedback in comments is important. In several cases across the communities, users comment “+1”, often followed by a reason for provided that positive feedback. Instances of critique or positive feedback occurred both in responses to the question or answer and in response to other comments. Currently there is no official way to provide feedback on comments, even though these may provide correct or incorrect information. Another shortfall is that a simple up or down vote does not allow the contributor to include his or her reasoning, which may provide the information needed to come up with the correct answer. This may be why so much feedback is occurring in the comments section. A mechanism that should be applied across all Stack Exchange sites is the ability to up or down vote both answers and comments, and

then an expandable window where users can provide reasoning for their vote. Providing reasoning in one location will create a collection of feedback that is more easily accessible.

DIY contains a large amount of metadiscourse that occurs on its site compared to the other two. Sites similar to DIY are expected to also contain larger amounts of metadiscourse. Metadiscourse is currently encouraged to take place on the websites “meta” page by site administrators. “Meta” pages exist alongside each Stack Exchange site in order to discuss matters relevant to the site. These “meta” sites include forums to post questions regarding the site itself, and also contains group chat capabilities. However, this is not sufficient because metadiscourse still occurs frequently on the post level. For archival value, comments not pertaining to the question should be moved outside of the thread. However, users still have a desire to make others aware that they are part of the community. If each question were to contain chat window in which metadiscourse could be made and a permanent history of the conversations kept, this information would be easily accessible and also be out of the thread of actually relevant information. Especially for on sites similar to DIY, this feature may help remove metadiscourse from the answer or comments. Live chat windows also give users the ability to directly communicate with each other during the information seeking process and will make it easier to come to a shared understanding of the question or answer.

DIY projects are often described in the question using visuals, such as pictures or gifs, and the answers also often contain visuals. This is because the members of the community are creating something in the real world that may be difficult to describe. Procedural knowledge lends themselves to collaborating using visuals because it is often easier to look at something than it is to describe it, and faster for the other users to understand. Comments do not have the same ability as answers to include visuals. In fact, comments can contain only very limited

formatting, such as italicize, bold, hyperlink, etc. The reason for this is that comment threads can become very long and images would exacerbate the problem. Images that are by default collapsed and show alternate text, but have the ability to expand, would allow users to further discuss their own relevant comments with visuals without having to create a new answer.

The Parenting and DIY sites contained a small number of extensions of previous contributions. Unlike Stack Overflow, where a snippet of code is not specific to one instance, answers in these communities may be relevant to only one person who is providing their experience. Thus, extension of the answer is unlikely to occur unless the answer is more general. Sites similar to DIY and parenting can also expect to have lower rates of the extension collaborative act. To encourage more extending of previous answers, users who post an answer could be granted the ability to import comments that extend their answer into the answer, and share the gain in reputation. This would encourage meaningful contributions toward directly extending previous work and make it easier to quickly view all information that is part of that answer, without having to read through each comments.

These are only a few design recommendations that could facilitate collaboration in the communities, as they exist today. Not all recommendations need to be implemented on all three sites, because they do not affect them all in the same way.

5.5 Contributions

Several contributions have been made through this research and are discussed below.

First, this study began as an extension of the work done using the taxonomy of collaborative acts used in Tausczik et. al. (2014). However, through collaborative, iterative

coding of different Stack Exchange sites other than Math Overflow, it became clear that the coding schema that was being applied is not inclusive of all collaborative acts used on Stack Exchange. Based on the additional three sites examined in this study, this taxonomy should consist of *providing information, clarification, critique, revision, extension, confirmation feedback, positive feedback, and metadiscourse*. Extending this taxonomy to be inclusive of the observed microcollaborations was a major challenge and it is a contribution use to any researchers who are analyzing the interaction that takes place on Stack Exchange or similar community-based Q&A sites.

The second contribution are results that show that different communities on Stack Exchange collaborate in different ways and have different collaborative activities. The findings indicate that distributions of collaborative acts across Stack Exchange sites do not come from the same population distribution. This knowledge is important for comparing different communities and supports that although communities may be given the same set of features, they do not use them in the same way. This finding may generalize to areas outside of social Q&A into other types of online communities.

The last contribution is the design recommendations. To demonstrate how comparing and contrasting the distributions of collaborative acts could lead to design changes, this paper described several ways that the design could be modified for different communities. These recommendations, if applied, should have a positive effect on the amount of collaboration in these communities and help to improve archival values of the threads. These recommendations include hyperlinked user tags, more detailed and pervasive up and down voting, question-level chat windows for metadiscourse, the ability to provide visuals in comments, and the ability to move extension comments into answers and provide reputation incentives.

5.6 Limitations

There are several limitations that should be noted.

- As described previously, each social Q&A site has very specific mechanisms that make it difficult to generalize across all social Q&A sites. As such, the results presented in this paper can generalize to Stack Exchange, but not necessarily any social Q&A site. Each site has its own technical requirements, affordances, cultures, norms, and more that can have an impact on microcollaborations.
- This study examined a sample of 5 questions from each of the three sites examined. While the threads selected were a minimum length that showed popularity and suggested collaboration, it is possible that the sample was not representative of all types of questions that occur on the respective sites.
- This work extended that done by Tausczik, Kittur, and Kraut (2014). The coding schema that was adapted may have formed differently if it were done through open coding, where the schema would have been developed from scratch. The schema also included several new categories that were not previously identified from the taxonomy of collaborative acts, so the schema presented in this paper may not be inclusive of all collaborative acts.
- The initial content analysis that established the coding schema was done in a collaborative manner that resulted in 80% inter-rater reliability. A single researcher then completed all coding. Coding the complete set with multiple researchers would have provided a more rigorous methodology.

Chapter 6

Conclusion

The following chapter will reiterate the research questions and methodology, contributions, and limitations of this work. Lastly, future work in this area will be presented.

6.1 Research Question and Method

Social Q&A services are a specific type of Q&A service, which have existed for decades. However, social Q&A, unlike digital reference services or expert services, leverages the knowledge and experiences of many people all around the globe to help answer questions in virtually any domain. This powerful tool for information seeking, and collaborative information seeking is relatively new and different services are experimenting with different formats with varying success. Stack Exchange is one of the largest and facilitates a massive amount of microcollaborations within each community, separated by topic. Understanding how these communities collaborate and improving these interactions will lead to improved quality of information. This research has been a part of this effort working toward the goal of improving these sites.

This research sought to answer the question of whether different Stack Exchange sites lead to different distributions of collaborative acts. Sub-questions were what are the differences in the distributions of collaborative acts and what changes may be made based on the distribution of collaborative acts that could improve collaboration? The collaborative activities that occur

across the sites are vital for the creators of social Q&A sites to know because what design works for one community that collaborates in one way may not work for other communities.

Answering this question required both qualitative and quantitative methodologies. To understand the collaborative acts, coding of various acts on sample Stack Exchange threads from representative sites, including Stack Overflow, DIY, and Parenting, was conducted.

Collaborative acts can be nuanced and qualitative methods are best at detecting alternate meanings. Quantitative methods utilizing chi-squared tests for independence detected the differences between the distributions of collaborations acts. This answered the question of if there are differences in distributions between communities. Qualitative methods were then used again to examine the differences in frequency between the specific acts that occurred across the three sites, as well as determining design recommendations that arose from the data, answering the two sub-questions.

6.2 Contributions and Limitations

There are three important contributions of this work. The first is the extended schema of collaborative acts that occur on Stack Exchange. The schema from Tausczik et. al. (2014) was insufficient in describing the collaboration acts that occur on other sites on the network other than MathOverflow. This underlines the fact that the collaborations that occur in the different communities are not the same. The second is the finding that, among the three sites that were examined, at least one of the sites had a significant difference in the distribution of collaborative acts. Lastly, this work contributed design recommendations that can be made in response to the

examination of the three Stack Exchange sites that will differ from the current template and should result in improved collaboration.

With all research there are limitations. This study has four notable limitations. The first is with generalizability. Only Stack Exchange was studied so the collaborative acts may be different depending on the affordances and culture on other sites. The second limitation is that with a smaller number of questions studied from each site, there is a chance that they are not representative of the entire population of questions on each site. The third limitation is that extending the work by Tausczik et. al. (2014) may have led to a different schema of collaborative acts than would have arisen from open coding. The representativeness of the final schema based on the questions studied minimizes the effect of this limitation. Lastly, qualitative coding would have been more reliable by collaboratively coding the entire set of questions, although sufficient inter-coder reliability was established before coding as a single researcher.

6.3 Future Work

This research is only the beginning of studying the different communities that exist on social Q&A sites. While much work has been done looking at the automation of information retrieval, work studying the collaboration on these sites is not common. Future work can go several directions. Further studies should be done comparing other sites on Stack Exchange. Other social Q&A sites could be studied in a similar way by sorting by topic. The distributions of the Stack Exchange communities studied in this work could be compared to other sites to see if the behavior reaches across the community through all social Q&A services. Researchers can also conduct qualitative interviews with community users in order to further support the findings

and elicit more design recommendations. Lastly, researchers could implement the design recommendations given in a controlled study and test them with users to see if collaboration improves.

BIBLIOGRAPHY

- Adamic, L. A., Zhang, J., Bakshy, E., & Ackerman, M. S. (2008). Knowledge sharing and yahoo answers: Everyone knows something. In *Proceedings of the 17th International Conference On World Wide Web* (pp. 665-674).
- Anderson, A., Huttenlocher, D., Kleinberg, J., & Leskovec, J. (2012). Discovering value from community activity on focused question answering sites: A case study of stack overflow. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 850-858).
- Baker-Eveleth, L., Sarker, S., & Eveleth, D. M. (2005). Formation of an online community of practice: An inductive study unearthing key elements. In *System Sciences, 2005. HICSS'05. Proceedings of the 38th Annual Hawaii International Conference on System Sciences* (pp. 254b-254b).
- Bernhard, D., & Gurevych, I. (2008). Answering learners' questions by retrieving question paraphrases from social Q&A sites. In *Proceedings of the Third Workshop on Innovative Use of NLP for Building Educational Applications* (pp. 44-52).
- Blooma, M. J., & Kurian, J. C. (2011). Research issues in community based question answering. In *PACIS* (p. 29).
- Blooma, M. J., Kurian, J. C., Chua, A. Y. K., Goh, D. H. L., & Lien, N. H. (2013). Social question answering: Analyzing knowledge, cognitive processes and social dimensions of micro-collaborations. *Computers & Education*, 69, 109-120.

- Chen, J., Subramanian, L., & Brewer, E. (2010). SMS-based web search for low-end mobile devices. In *MobiCom '10: Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking* (pp. 125-136).
- Dabbish, L., Stuart, C., Tsay, J., & Herbsleb, J. (2012). Social coding in GitHub: Transparency and collaboration in an open software repository. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work* (pp. 1277-1286).
- de Moor, A. (2006). Community memory activation with collaboration patterns. In L. Stillman & G. Johanson (Eds.), *Proceedings of the Third Prato Community Informatics Research Network Conference (CIRN 2006)*.
- De Wever, B., Van Keer, H., Schellens, T., & Valcke, M. (2010). Roles as a structuring tool in online discussion groups: The differential impact of different roles on social knowledge construction. *Computers in Human Behavior*, 26(4), 516-523.
- Gazan, R. (2006). Specialists and synthesists in a question answering community. *Proceedings of the American Society for Information Science and Technology*, 43(1), 1-10.
- Gazan, R. (2010). Microcollaborations in a social Q&A community. *Information Processing & Management*, 46(6), 693-702.
- Gazan, R. (2011). Social q&a. *Journal of the American Society for Information Science and Technology*, 62(12), 2301-2312.
- Golovchinsky, G., Qvarfordt, P., & Pickens, J. (2009). Collaborative information seeking. *Computer*, (3), 47-51.
- Gyongyi, Z., Koutrika, G., Pedersen, J., & Garcia-Molina, H. (2007). Questioning Yahoo! Answers.

- Harper F., Moy D., and Konstan J., 2009. "Facts or friends? Distinguishing informational and conversational questions in social Q&A sites," *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, pp. 759–768.
- Harper, F. M., Weinberg, J., Logie, J., & Konstan, J. A. (2010). Question types in social Q&A sites. *First Monday*, 15(7).
- Hsieh, G., Kraut, R. E., & Hudson, S. E. (2010). Why pay?: Exploring how financial incentives are used for question & answer. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 305-314).
- Janes, J., Hill, C., & Rolfe, A. (2001). Ask-an-expert services analysis. *Journal of the American Society for Information Science and Technology*, 52(13), 1106-1121.
- Jonassen, D. H., & Kwon II, H. (2001). Communication patterns in computer mediated versus face-to-face group problem solving. *Educational Technology Research and Development*, 49(1), 35-51.
- Lankes, R. D. (2004). The digital reference research agenda. *Journal of the American Society for Information Science and Technology*, 55(4), 301-311.
- Lee, Y. S. (2006). *Toward a new knowledge sharing community: Collective intelligence and learning through web-portal-based question-answer services*. Thesis.
- Liu, Q., Agichtein, E., Dror, G., Gabrilovich, E., Maarek, Y., Pelleg, D., & Szpektor, I. (2011). Predicting web searcher satisfaction with existing community-based answers. In *Proceedings of the 34th international ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 415-424).

- Mendes Rodrigues, E., & Milic-Frayling, N. (2009). Socializing or knowledge sharing?: Characterizing social intent in community question answering. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management* (pp. 1127-1136).
- Moskaliuk, J., Kimmerle, J., & Cress, U. (2012). Collaborative knowledge building with wikis: The impact of redundancy and polarity. *Computers & Education*, 58(4), 1049-1057.
- Nam, K. K., Ackerman, M. S., & Adamic, L. A. (2009). Questions in, knowledge in?: A study of naver's question answering community. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 779-788).
- Phang, C. W., Kankanhalli, A., & Sabherwal, R. (2009). Usability and sociability in online communities: A comparative study of knowledge seeking and contribution*. *Journal of the Association for Information Systems*, 10(10), 721.
- Pomerantz, J., Nicholson, S., Belanger, Y., & Lankes, R. D. (2004). The current state of digital reference: Validation of a general digital reference model through a survey of digital reference services. *Information Processing & Management*, 40(2), 347-363.
- Preece, J., & Maloney-Krichmar, D. (2003). Online communities: Focusing on sociability and usability. *Handbook of Human-Computer Interaction*, 596-620.
- Quantcast. (2016). *Stackoverflow.com*.
- Quantcast. (2016). *Diy.stackexchange.com*.
- Quantcast. (2016). *Parenting.stackexchange.com*.
- Romero, N. L. (2011). Review and brief history of collaborative systems: Taxonomy, services and classification. *Systems Science and Collaborative Information Systems: Theories, Practices and New Research: Theories, Practices and New Research*, 139.

- Rosenbaum, H., & Shachaf, P. (2010). A structuration approach to online communities of practice: the case of Q&A communities. *Journal of the American Society for Information Science and Technology*, 61(9), 1933–1944.
- Schmidt, K., & Bannon, L. (1992). Taking CSCW seriously. *Computer Supported Cooperative Work (CSCW)*, 1(1-2), 7-40.
- Shah, C. (2009). Toward collaborative information seeking (CIS). *ArXiv Preprint ArXiv:0908.0709*.
- Shah, C., & Kitzie, V. (2012). Social Q&A and virtual reference-comparing apples and oranges with the help of experts and users. *Journal of the American Society for Information Science and Technology*, 63(10), 2020-2036.
- Shah, C., Oh, J. S., & Oh, S. (2008). Exploring characteristics and effects of user participation in online social Q&A sites. *First Monday*, 13(9).
- Shah, C., Oh, S., & Oh, J. S. (2009). Research agenda for social Q&A. *Library & Information Science Research*, 31(4), 205-209.
- Shah, C., & Pomerantz, J. (2010). Evaluating and predicting answer quality in community QA. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 411-418).
- Shah, C., Radford, M. L., Connaway, L. S., Choi, E., & Kitzie, V. (2012). How much change do you get from 40? Analyzing and addressing failed questions on social Q&A. *Proceedings of the American Society for Information Science and Technology*, 49(1), 1-10.

- Tausczik, Y. R., Kittur, A., & Kraut, R. E. (2014). Collaborative problem solving: A study of mathoverflow. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 355-367).
- Tausczik, Y. R., & Pennebaker, J. W. (2011). Predicting the perceived quality of online mathematics contributions from users' reputations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1885-1888).
- Wenger, E. (2011). Communities of practice: A brief introduction.
- White, M. D. (2001). Digital reference services. *Library & Information Science Research*, 23(3), 211-231.

ACADEMIC VITA

Academic Vita of Evan Friedenberg
evan.friedenberg@gmail.com

Education:

M.S. in Information Sciences and Technology

B.S. in Information Sciences and Technology

Minor in Supply Chain Management

Honors: Information Sciences and Technology

Thesis Title: An Exploration of Collaborative Type Distribution on the Stack Exchange Network

Thesis Supervisor: Xiaolong Zhang

Work Experience:

Fall 2015-Spring 2016

Teaching Assistant for IST 440W

The Pennsylvania State University, College of IST

Dr. Shawn Clark, Professor John Hill

Summer 2014

IT Infrastructure Intern

The Boeing Company, Seattle, Washington

Stacy Brening

Summer 2013

Project Employee

Wipro Technologies, Bangalore, India

Alok Jain

Spring 2012-Fall 2014

Undergraduate Research Assistant

The Pennsylvania State University, College of IST

Dr. Madhu Reddy

Publications:

McNeese, N. & Reddy, M., & Friedenberg, E. (2014). Towards a Team Mental Model of Collaborative Information Seeking. *Proceedings of the 2014 Annual Meeting of the Human Factors and Ergonomic Society*.

Activities:

IST Diplomats

IST Club Co-President

Startup Week Student Committee