

9 July 2011

Research Issues In Community Based Question Answering

Mohan John Blooma

RMIT International University, blooma.john@rmit.edu.vn

Jayan Chirayath Kurian

RMIT International University, jayan.kurian@rmit.edu.vn

ISBN: [978-1-86435-644-1]; Full paper

Recommended Citation

Blooma, Mohan John and Kurian, Jayan Chirayath, "Research Issues In Community Based Question Answering" (2011). *PACIS 2011 Proceedings*. 29.

<http://aisel.aisnet.org/pacis2011/29>

This material is brought to you by the Pacific Asia Conference on Information Systems (PACIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in PACIS 2011 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

RESEARCH ISSUES IN COMMUNITY BASED QUESTION ANSWERING

Mohan John Blooma, Centre of Commerce, RMIT International University, Ho Chi Minh City, Vietnam, blooma.john@rmit.edu.vn

Jayan Chirayath Kurian, Centre of Commerce, RMIT International University, Ho Chi Minh City, Vietnam, jayan.kurian@rmit.edu.vn

Abstract

Community based Question Answering (CQA) services are defined as dedicated platforms for users to respond to other users' questions, resulting in the building of a community where users share and interactively give ratings to questions and answers (Liu et al., 2008). CQA services are emerging as a valuable information resource that is rich not only in the expertise of the user community but also their interactions and insights. However, scholarly inquiries have yet to dovetail into a composite research stream where techniques gleaned from various research domains could be exploited for harnessing the information richness in CQA services. This paper explores the CQA domain by first understanding the service and its modules and then exploring previous studies that was conducted in this domain. This paper then compares a CQA service with traditional question answering (QA) systems to identify possible research challenges that need to be focused. This paper also identifies two nontrivial research issues that are prominent in this domain and proposes various recommendations for addressing them in future.

Keywords: Community based Question Answering, question answering systems, research issues, social computing

1 INTRODUCTION

Over the past decade, the Web has been transformed from a repository of largely static content to an interactive information space (Kirsch et al., 2006) where users are able to participate freely in co-creating and sharing various kinds of content (such as text, image, audio and video). This is facilitated by a set of social computing applications such as blogs, social networking services, wikis, vlogs, and community based question answering (CQA) services (Iskander et al., 2007). Social computing has been defined as a computational integration of social studies and human social dynamics together with the design and use of ICT technologies in a social context (Wang et al., 2007).

In particular, a CQA service is defined as a tool for users to respond to other users' questions (Liu et al., 2008b). In recent years, CQA services like Yahoo! Answers, Naver, and AnswerBag have become very popular, attracting a large number of users who seek and contribute answers to a variety of questions on diverse subjects (Wang et al., 2009). While Yahoo! Answers offered by Yahoo! was launched in 2005, Naver is a South Korean search portal that added its CQA service in 2005 and AnswerBag is a collaborative online database of FAQs which was founded in 2003. Social computing applications such as wikis and blogs provide comments and opinions but may not solicit responses. However, responses from users contributing answers to questions form the backbone of a successful CQA service. These services are dedicated platforms for user-oriented QA and result in building up a community where users share and interactively give ratings and comments to questions and answers. Hence, CQA services are emerging as a valuable information resource that is rich not only in the expertise of a user community but also in the community's interactions and insights. The emergence of CQA services raises new research challenges for information systems researchers.

CQA services are derived from a branch of information retrieval known as question answering (QA). The goal of QA is to build intelligent systems that can provide succinct answers to questions constructed in a natural language. This approach helps in understanding users' information needs in the form of a question and delivers exactly the required information in the form of an answer (Demner-Fushman, 2006). The development of automatic approaches to QA takes place primarily in the framework of large-scale evaluations, such as the QA track at the Text Retrieval Conference (TREC) (Voorhees, 2003). However, QA systems that participate in these evaluations work largely on restricted domains and on closed corpora such as encyclopedia or news articles (Brill et al., 2002). Moreover, these systems focus on fact-based direct questions commonly known as factoid questions, for example: "Who was the US President in 1999?" These questions are involved in finding an exact short string, often representing an entity, such as named entities (person, organization, location), temporal expressions, or numerical expressions. Examples of popular TREC systems are Quanda (Breck et al., 1999), Falcon (Harabagiu et al., 2000), and AskMSR (Brill et al., 2002). Unlike these systems, START (Katz and Levin, 1988), Mulder (Kwok et al., 2001), and AnswerBus (Zheng, 2002) are examples of open-domain QA systems that are scaled to the Web. These systems use the redundancy of information on the Web to answer factoid questions.

A search of the literature suggested that there exist only a limited number of studies that have attempted to consider research gaps in CQA by integrating the QA research (Jeon et al., 2005a; Jeon et al., 2005b; Bian et al., 2008b; Blooma et al., 2009). A research trend analysis of the QA domain revealed that Bian et al. (2008b) was among the first in this field to show that QA is amenable to CQA. They used machine learning methods to automatically answer factoid questions from CQA corpora. This paper compares the research issues in building a QA system by reusing the information resource in collected in CQA services. This paper also points out the research directions in CQA that needs to be addressed in future. Research in CQA needs to evolve to encompass new theories and methodologies that can address gaps and research questions posed in this domain, which is not addressed in traditional QA research. This paper suggests that the information systems community needs to focus on this emerging domain of CQA as a priority topic, and, in the process, evolve the core of research in the discipline itself.

The remainder of the paper is organized as follows. Section 2 presents an overview of CQA services focusing on the processing modules and also gives a review of the previous studies in CQA. Section 3 discusses research issues related to CQA systems by comparing it with QA systems. Section 4 concludes.

2 OVERVIEW OF CQA

CQA services have become a popular medium where users share their expertise to answer other users' questions (Bian et al., 2008b). The success of these services has been largely attributed to the fact that users can obtain quick and precise answers to any natural language question (Liu et al., 2008b). A study conducted by Liu et al. (2008b) showed that users approach CQA services for opinions and to answer complex rather than factoid questions. The growth of CQA services has resulted in expanding repositories for opinions and complex questions. The following sections give a more detailed description of CQA research by first presenting processes in CQA services and then discussing major strands of CQA research. In a typical CQA service, as illustrated in Figure 1, there are different processing modules in a CQA service. Details of each processing module in a CQA service are described below.

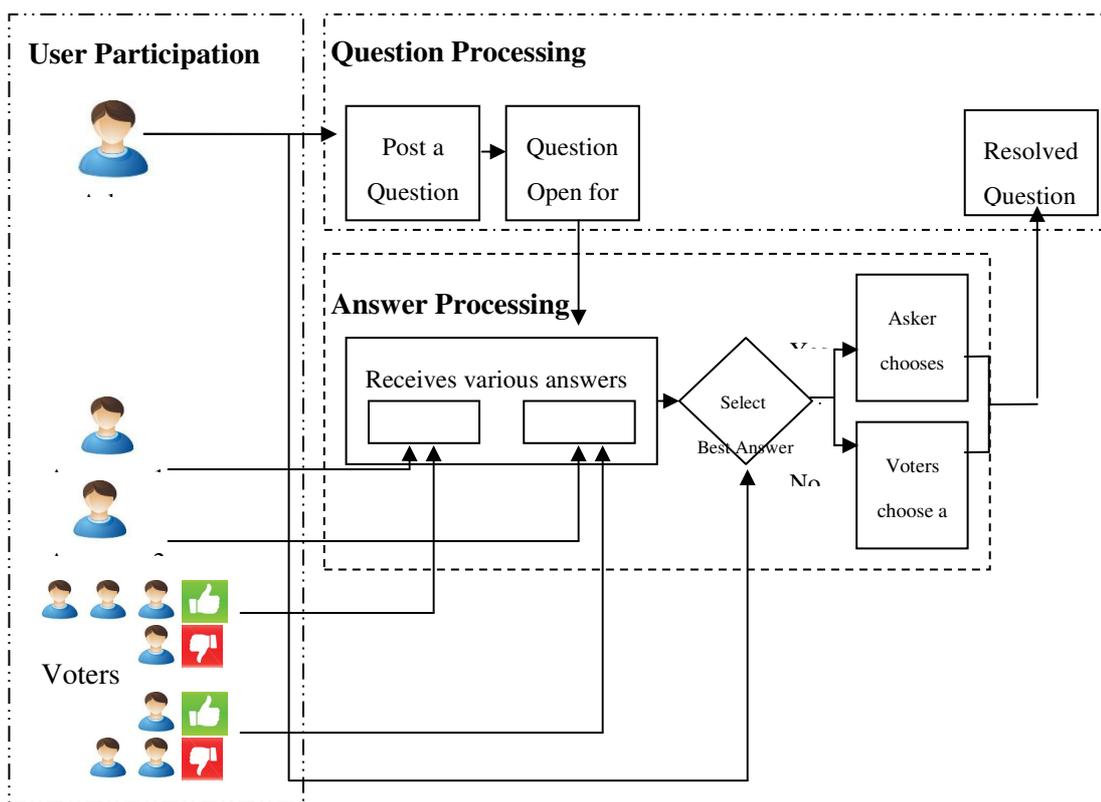


Figure 1. Processing modules in a CQA service

2.1 Processes in CQA Services

As illustrated in Figure 1, there are different processing modules in a CQA service. They are identified as question processing, answer processing, and user participation modules. However, the first two modules are distinct from question processing and answer processing modules of a automatic QA system. In a QA system, these processing modules use machine learning techniques and linguistic processing methods. On the other hand, in a CQA service, question processing and answer processing modules rely solely on the voluntary involvement of users. Details of each processing module in a CQA service are described below.

2.1.1 Question Processing

Question processing is initiated when a user posts a question by entering a question subject (i.e. title), and may optionally give details (i.e. description). A question title is referred to as a question in this research. After a short delay, which may include checking for abuse, the question appears open for users to contribute answers. Users may rate the question based on how interesting the subject is and may recommend it to other users. As users are not obligated to rate questions, usually only a small proportion of questions attract ratings (Sun et al., 2009).

2.1.2 Answer Processing

Once a question is open, other users can answer the question, vote on other users' answers, comment on the question (e.g., to ask for clarification or provide other, non-answer feedback), or provide various metadata for the question (e.g., give stars for quality). When the asker obtains a satisfactory answer, the question is closed by selecting the satisfactory answer as the best answer. The question might have received numerous other answers depending on its popularity. If the best answer is identified, the question is deemed to have been resolved, and will be available for users for future reference. If the asker fails to select a best answer, the question will remain open for a stipulated number of days depending on the CQA service, and would be closed later. In this case, the answer with the highest number of votes would be marked as the best answer. If answers for a question do not have a best answer marked by an asker or voters, the question would remain as undecided. Thus, it is evident that answer processing in a CQA service is driven solely by users' active involvement (Bian et al., 2009).

2.1.3 User Participation

It is evident from the question processing and answer processing modules that user participation forms the backbone of CQA services. Users participate by asking, answering and rating CQA services, which aids in the success of these services. Users are divided into askers, answerers, and voters. Each user is given points based on roles played, activities in which the user participated, and the quality of questions and answers contributed. The point system and scoring vary depending on CQA services, and user participation details are stored as user profiles. Thus the user participation module records user activities and profiles (Nam et al., 2009).

2.2 Previous Studies in CQA Research

The research focus and central research questions in the relatively young discipline of CQA have evolved over time. This section gives a review of the range of research related to each processing

modules of CQA and also reveals the transformation of CQA as a research domain over the last five years.

2.2.1 *Question Processing*

An emerging area in CQA research lies in processing the questions asked by users. In particular, CQA research on question processing could be discussed in two strands: identification of similar questions (Jeon et al., 2005a; Harper et al., 2008), and classification of questions (Tamura et al., 2005; Li et al., 2008).

On the first strand of studies, two different techniques were tested to identify similar questions, namely the translation-based model (Jeon et al., 2005a; Jeon et al., 2005b; Lee et al., 2008), and the MDL-based (Minimum Description Length) tree cut model (Duan et al., 2008; Cao et al., 2008). First, Jeon et al. (2005a; 2005b) used the translation-based model. In their original study they used similar question pairs to train their model while in the latter study question-answer pairs were used. Lee et al. (2008) addressed the issue of the lexical gap between questions using compact translation models. These studies focused on training the language model for retrieval. Second, Duan et al. (2008) used the MDL-based tree cut model for identifying the topic and focus of a question automatically. The MDL based model requires tree-based syntactic structure, which is adapted from linguistics. Cao et al. (2008) also used the MDL based tree cut model for question recommendation and showed that their model outperformed the vector space model. However, these studies used questions from restricted domains. Another more recent study used quadripartite-clustering approach by considering relationship between questions, answers and users involved (Bloom et al., 2011a). They took a step further by harnessing the collective wisdom garnered in CQA for identifying similar questions.

On the second strand of studies, Tamura et al. (2005) attempted to classify multiple-sentence questions, collected in CQA corpora, by using core sentences. Another study identified subjective orientation of questions using classification techniques (Li et al., 2008). Subjective orientation refers to whether a user is searching for subjective or objective information. Subjective questions seek answers containing private statements, such as personal opinion and experience. In contrast, objective questions request objective, verifiable information, often with support from reliable sources. Liu et al. (2008b) worked on the evolution of CQA services and concluded that CQA corpora were a good source of complex and opinion questions. Another related study attempted to determine if a conversational question posed in a CQA service was intended to start a discussion or was purely informational (Harper et al., 2009). Their study suggested that profiles of users and answer content could be used to classify questions leading to future directions in CQA research.

2.2.2 *Answer Processing*

Studies on CQA services that focus on user-contributed answers can also be divided into two strands of investigation.

The first strand of investigation assesses the quality of user-generated answers in CQA corpora (Suryanto et al., 2009; Agichtein et al., 2008). Among notable research in this strand is the work of Jeon et al. (2006) using Social features, such as the answerer's authority and the asker's answer rating, to predict the quality of answers. Features related to the Textual Content of answers, such as the number of unique words and the length ratio between questions and answers, were ignored. Agichtein et al. (2008) relied on a combination of Social and Textual Content features. In particular, they developed a comprehensive graph-based model of contributor relationships and combined it with Social and Textual Content features to classify high-quality answers. The results suggested that both Social and Textual Content features were found to be equally significant in high-quality answer selection. Kim and Oh (2009) solicited the asker's open-ended comments to uncover the criteria for

selecting high-quality answers. Their study identified the importance of content-oriented relevance judgment criteria. Wang et al. (2009b) attempted to consider questions and their answers as relational data and proposed an analogical reasoning based approach to identify high-quality answers. Their study gave insights into the relation between questions and their answers in quality judgement. Harper et al. (2009) investigated predictors of answer quality through a comparative and controlled field study of responses provided across several online CQA services. They reported qualitative observations to better understand the characteristics of different types of CQA services. Suryanto et al. (2009) used the user expertise of answerers to derive answer quality. They established the fact that good quality answers could be found among answers that are not marked as best by askers. This is in contrast to Jeon et al. (2006), where best answers marked by askers were used for quality judgement. Finally, Blooma et al., (2011b) proposed a comprehensive framework of answer quality that integrated both social and textual features and identified significant features to judge high-quality answers from their low quality counterpart.

The second strand of investigation pertained to answer retrieval (Bian et al., 2008b; Xue et al., 2008). Bian et al. (2008b) described a machine learning based ranking framework for social media that integrated user interactions and content relevance. Their study demonstrated the effectiveness of user interaction and content relevance features for answer retrieval. Bian et al. (2009) proposed a semi-supervised mutual reinforcement framework for simultaneously calculating content quality and user reputation. Xue et al. (2008) proposed a retrieval model with a query likelihood approach for retrieving high-quality answers.

2.2.3 *User Participation*

Studies on CQA services that focus on user participation can be divided into two strands of investigation apart from studies discussed earlier that are related to question processing and answer processing.

The first strand of studies analyses knowledge generation characteristics and users' authority as a result of user participation. Nam et al. (2009) found that altruism, learning, and competency are common motivations enticing top answerers to participate, but such participation is often highly intermittent. A major problem with this approach was determining how many users should be chosen as authoritative, from a ranked list. To address this problem, Bouguessa et al. (2008) proposed a method to identify authoritative participants. This method automatically discriminated between authoritative and non-authoritative users.

The second strand of studies investigated user participation and authority using social network analysis. Rodrigues and Frayling (2009) performed an in-depth content analysis using social network analysis techniques to monitor the dynamics of community ecosystems. Their study concluded that CQA services rely on user participation and the quality of users' contributions. Jurczyk and Agichtein (2007) and Suryanto et al. (2009) used link structure to discover the authority of users. Estimating the authority of users has potential applications for answer ranking, spam detection, and incentive mechanism design.

On reviewing the state-of-the art research based on CQA services it is evident that there is still an immense potential for research in this domain. To explore the possibilities of future research in CQA the following section discusses the research issues in CQA by identifying the research gaps in this domain.

3 RESEARCH ISSUES IN CQA

To better illustrate the research issues in CQA domain, a comparison of QA systems and CQA services is detailed in Table 1. The table presents four differences between a QA system and a CQA service. These differences highlight major research issues in this domain.

The first difference is related to question type. The nature of questions a QA system can handle determines its scope and hence this distinction is important. QA systems handle single-sentence questions that are particularly fact-based. Single-sentence questions are defined as questions composed of one sentence. Since 2006, TREC and other QA systems research groups have started to focus on complex and interactive questions (Dang et al., 2006). However, they are still in the early stages of research. CQA services are rich in multiple-sentence questions, which are defined as questions composed of two or more sentences: For example, "My computer reboots as soon as it gets started. OS is Windows XP. Is there any homepage that tells why it happens?". For conventional QA systems, these questions are not expected and existing techniques are not applicable or barely work in the context of such questions (Tamaru et al., 2005). Therefore, constructing a QA system on CQA corpora that can handle multiple-sentence questions is challenging and desirable to satisfy user needs.

The second difference is related to the source of the answers. This distinction is important as it determines the complexity of processing required to answer questions. Currently, various QA systems have been built on different types of corpora, such as full text news articles or encyclopaedia articles, for closed domain systems (Harabagiu et al., 2000; Brill et al., 2002), and the Web for open domain systems (Katz, 1988; Kwok et al., 2001). However, in a CQA service, questions and answers collected as a CQA corpus are contributed by users, while the producers of content in other types of corpora are professionals, publishers, or journalists (Liu et al., 2008a). Hence they differ in many aspects such as the length of the content, structure and writing style. The noisy nature of the content, probability of spam, and the variance in quality elicits new challenges.

	QA Systems	CQA Services
Question type	Factoid single sentence questions	Multiple sentence questions
Source of answers	Extracted from documents in a corpus	Contributed by users
Quality of answers	High, as answers are extracted from reputed sources	Varying as it depends on answers contributed by users
Availability of metadata	None	Best answer selected by asker and positive and negative ratings given by voters
Time lag	Automatic and immediate	The asker needs to wait for users to post answers

Table 1 Comparison of QA systems with CQA services

The third difference is related to the quality of answers, which eventually determines the quality of a system. It is evident from the discussion on the source of answers that a QA system extracts answers

from a reputed corpus. However, in CQA services, answers are obtained from different kinds of users, with varying reputations (Liu et al., 2008a). Previous studies also showed evidence of variance in the quality of answers in different CQA services. This variance in the quality of answers depends on various factors such as the community that participates, compensation for answers and whether answers are from experts or casual answerers (Harper et al., 2008). In addition, the quality of answers becomes important when there are many responses to a single question (Jeon et al., 2006).

The fourth difference is related to the availability of metadata. CQA services are abundant in metadata such as comments and positive and negative ratings, together with authorship and attribution information created by an explicit support for social interactions between users. These metadata make the content in CQA services rich. Most of the studies on CQA services harnessed these metadata to sieve high-quality content (Liu et al., 2008b; Bian et al., 2008b; Blooma et al., 2011a, 2011b). However, this is not the case in traditional QA systems.

The final difference is related to the time lag in obtaining answers to a newly posed question. QA systems generate answers automatically to questions, and hence eliminate any time lag. However, CQA services involve a time lag and askers have to wait for answers to a question from various users (Jeon et al., 2005a).

Hence, the major research issues in harnessing a CQA corpus could be summarised with respect to the nature of questions posed, source of answers, variance in the quality of answers, richness of metadata available and the time lag in obtaining answers. Considering these issues, two major challenges that need to be addressed in CQA for re-use of CQA corpus is detailed below.

3.1.1 *Identifying similar questions*

The first challenge was to find questions in a corpus that were semantically similar to a users' newly posed question (Jeon et al., 2005b). This facilitates the retrieval of high-quality answers that are associated with similar questions identified in the corpus, reducing the time lag associated with a community-based QA system. Due to the complex nature of questions posed in CQA services as discussed earlier, it is not an easy task to identify similar questions. Measuring semantic similarities between questions is not a trivial task. In addition, two questions having the same meaning may use entirely different wording. For example, "Is downloading movies illegal?" and "Can I share a copy of a DVD online" have almost identical meanings but they are lexically very different. Similarity measures developed for document retrieval work poorly when there is little word overlap. Users tend to express their needs in the form of natural language, with multiple sentences describing a scenario that leads to the question. An example of a question with multiple sentences framed with a scenario is "My computer reboots as soon as it gets started. OS is Windows XP. Is there any homepage that tells why it happens?" This could be formed as a simple direct question given as, "Why does a computer running Windows XP OS reboot as soon as it gets started?" These questions are posed in a free and natural form to a user community when compared to direct factoid questions posed to an automatic QA system (Jeon et al., 2005b).

With respect to previous work in this area that discussed methods for question retrieval, Jeon et al. (2005a) used the similarity between answers in a corpus to estimate probabilities for a translation-based retrieval model. Hence there is a need to work on more comprehensive methods that make use of available metadata to identify similar questions (Bian et al., 2009). More recently, Cao et al., (2011) used question topic and question focus to cluster question search results. They used MDL based Tree Cut Model to identify question topic and clustered questions to improve search results. Blooma et al., (2011) used more complex quadripartite graph-based structure of CQA services considering the relationships between questions, answer, askers and answerers in identifying similar questions. They considered not only the content similarity of questions but also the collective wisdom garnered in CQA services by considering answers, askers and answers. Moreover, as there is an increasing amount of data collected in the CQA corpora, identifying semantically similar questions from the corpora for a

newly posed question not found in it would be more challenging and rewarding for the ongoing research in this domain.

3.1.2 *Identifying high-quality answers*

The second challenge was to identify high-quality answers from a set of candidate answers. The quality of user-contributed content in CQA services, in the form of questions, answers and votes, varies drastically from excellent to spam (Agichtein et al., 2008). The open contribution model of the system, the ease of using these services for content creation and sharing, as well as an attraction towards collaboration among like-minded people, have led to significant growth of these services. However, the flipside of growth is that content in CQA services is thematically diverse, lacking a consistent structure together with linguistic style. Moreover, users give poor quality answers for several reasons, including limited knowledge about a question domain, bad intentions (e.g., spam, making fun of others, etc.) or limited time to prepare good answers. In addition, the absence of any editorial control results in the least relevant answers often being classified as spam (Bian et al., 2008b; Heyman et al., 2007), resulting in poor quality of content (Shachaf, 2010). In CQA corpora, each question may have a number of answers from different users. Users may have varying reputations in the community depending on factors such as the quality of their answers or their activity level. There will be a wide variance in the quality, structure and coverage of answers from different users in such communities. Hence there is a need to work on more comprehensive methods that tap into various features associated with answers that in turn help to recommend high-quality answers for newly posed questions.

With respect to previous work in this area, the challenge in identifying a high-quality answer for a new question in a CQA corpus is significant for the following reasons. First, among all answers obtained for a single question, there might be more than one correct answer, although there is a best answer marked by the asker (Agichetien et al., 2008). Second, a large fraction of content contributed as answers in CQA services often reflects the unsubstantiated opinions of users (Su et al., 2007). Finally, the quality of content varies significantly from that of traditional Web content (Bian et al., 2008b). Bloom et al., (2009; 2011) proposed a quality model for that returns high-quality answers to the asker thereby eliminating time the asker waits for the user community to respond, and hopefully paves the way for more robust CQA services. However, there is a need to probe into quality of user-generated answers in CQA services as Shachaf, (2010) compares four CQA services and concludes that although such services provide information that may not resemble the quality of information that professionals provide for example reference librarian.

4 CONCLUSION

CQA offers IS researchers an exciting opportunity for new research, as well as for further evolving the discipline and its practice of research and education. The enabling benefits of addressing the research issues related to questions, answers and users participating in CQA services is not only advantages to CQA services but could be expanded to other social computing tools available. The next step is for IS researchers to take the lead in using them for social change, to deploy new socio-technical systems, and to create new applications. This paper identifies two major research challenges in this domain as identifying similar questions for reusing related answers in CQA corpora as well as identifying high-quality answers from a range of answers obtained for similar questions. Nevertheless, researchers can explore a wide variety of social, economic, and organizational aspects related to questions, answers and users in CQA services that is under the umbrella of social computing. For instance, online communities simultaneously manifest a focus on individualism in the form of personal expression and

altruism in the form of sharing and communal benefits. This interplay can offer rich insights for both social and technical aspects in identifying similar questions as well as high-quality answers.

References

- Agichtein, E., Castillo, C., Donato, D., Gionis, A., & Mishne, G. (2008). Finding high-quality content in social media. In *Proceedings of the International Conference on Web Search and Web Data Mining*, California, United States, (pp. 183-194). NY: ACM.
- Bian, J., Liu, Y., Agichtein, E., & Zha, H. (2008b). Finding the right facts in the crowd: factoid question answering over social media. In *Proceeding of the 17th International Conference on World Wide Web*, Beijing, China (pp. 467-476). NY: ACM.
- Bian, J., Liu, Y., Zhou, D., Agichtein, E., & Zha, H. (2009). Learning to recognize reliable users and content in social media with coupled mutual reinforcement. In *Proceeding of the 18th International Conference on World Wide Web*, Madrid, Spain (pp. 51 – 60). NY: ACM.
- Blooma, M. J., Chua, A. Y. K., & Goh, D. H. (April, 2011a). Quadripartite graph-based clustering of questions? In *Proceedings of the 7th International Conference on Information Technology: New Generations*, Nevada, United States (Accepted for publication). NY: IEEE Computer Society.
- Blooma, M. J., Chua, A. Y. K., & Goh, D. H. (January, 2011b). What makes a high quality user-generated answer? *IEEE Internet Computing*, 15 (1), 66 – 71.
- Blooma, M. J., Chua, A. Y. K., Goh, D. H., & Keong, L. C. (2009). A trend analysis of the question answering domain. In *Proceedings of the 6th International Conference on Information Technology: New Generations 2009*, Nevada, United States (pp. 1522-1527). NY: IEEE Computer Society.
- Bouguessa, M., Dumoulin, B., & Wang, S. (2008). Identifying authoritative actors in question-answering forums - the case of Yahoo! Answers. In *Proceedings of the 14th ACM International Conference on Knowledge Discovery and Data Mining*, Nevada, United States (pp. 866-874). NY: ACM.
- Breck, E., House, D., Light, M., & Mani, I. (1999). Question answering from large document collections. In *Proceedings of AAAI 1999 Fall Symposium on Question Answering Systems*, Florida, United States (pp. 26-31). MA: AAAI Press.
- Brill, E., Dumais, S., & Banko, M. (2002). An analysis of the ASKMSR question-answering system. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing*, Pennsylvania, United States (pp. 257-264). NJ: Association of Computational Linguistics.
- Cao, Y., Duan, H., Lin, C. Y., Yu, Y., & Hon, H. W. (2008). Recommending questions using the MDL-based tree cut model. In Huai, J., Chen, R., Hon, H. W., Liu, Y., Ma, W. Y., Tomkins, A., & Zhang, X. (Eds.), *Proceedings of the 17th International World Wide Web Conference*, Beijing, China (pp. 81-90). NY: ACM.
- Cao, Y. Duan, H., Lin, C. Y., & Yu, Y. (2011) Re-ranking question search results by clustering questions. *Journal of the American Society for Information Science and Technology* (to be published)
- Dang, H. T., Lin, J., & Kelly, D. (2006). Overview of the TREC 2006 question answering track. In *Proceedings of the 15th Text REtrieval Conference (TREC 2006)*, Maryland, United States (pp. 99-116). Gaithersburg: NIST.
- Demner-Fushman, D. & Lin, J. (2006). Answer extraction, semantic clustering, and extractive summarization for clinical question answering. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the ACL*, Sydney, Australia (pp. 841–848). NJ: Association of Computational Linguistics.
- Duan, H., Cao, Y., Lin, C. Y., & Yu, Y. (2008). Searching questions by identifying question topic and question focus. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Ohio, United States (pp. 156–164). NJ: Association of Computational Linguistics.

- Harabagiu, S., Moldovan, D., Pasca, M., Mihalcea, R., Surdeanu, M., Bunescu, R., Girju, R., Rus, V., & Morarescu, P. (2000). FALCON: Boosting knowledge for answer engines. In *Proceedings of the 9th Text REtrieval Conference*, Maryland, United States (pp. 479-489). Gaithersburg: NIST.
- Harper, F. M., Moy, D., & Konstan, J. A. (2009). Facts or friends? distinguishing informational and conversational questions in social q&a sites. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, Massachusetts, United States (pp. 759-768). NY: ACM.
- Harper, F. M., Raban, D., Rafaeli, S., & Konstan, J. (2008). Predictors of answer quality in online Q&A sites. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, Florence, Italy (pp. 865-874). NY: ACM.
- Iskandar, D. N. F. A., Pehcevski, J., Thom, J. A., & Tahaghoghi, S. M. M. (2007). Social media retrieval using image features and structured text. In Fuhr, N., Kamps, J., Lalmas, M., & Trotman, A. (Eds.), *Focused Access to XML Documents, INEX 2007*, Dagstuhl, Germany. LNCS 4862 (pp. 358-372). Berlin: Springer.
- Jeon, J., Croft, W. B., & Lee, J. H. (2005a). Finding similar questions in large question and answer archives. In *Proceedings of the 14th ACM Conference on Information and Knowledge Management*, Bremen, Germany (pp. 84-90). NY: ACM.
- Jeon, J., Croft, W. B., & Lee, J. H. (2005b). Finding semantically similar questions based on their answers. In Baeza-Yates, R. A., Ziviani, N., Marchionini, G., Moffat, A., & Tait, J. (Eds.), *Proceedings of the 28th International ACM Conference on Research and Development in Information Retrieval*, Salvador, Brazil (pp. 617-618). NY: ACM.
- Jeon, J., Croft, W. B., Lee, J. H., & Park, S. (2006). A framework to predict the quality of answers with non textual features. In Efthimiadis, E., Dumais, S., Hawking, D., & Järvelin, K. (Eds.), *Proceedings of the 29th International ACM Conference on Research and Development in Information Retrieval*, Seattle, United States (pp. 228 – 235). NY: ACM.
- Jurczyk, P. & Agichtein, E. (2007). Discovering authorities in question answer communities by using link analysis. In *Proceedings of the 16th ACM Conference on Conference on Information and Knowledge Management*, Lisbon, Portugal, (pp. 919-922). NY: ACM.
- Katz, B. & Levin, B. (1988) Exploiting lexical regularities in designing natural language systems. In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary (pp. 316-323).
- Kim, S. & Oh, S. (2009). User's relevance criteria for evaluating answers in a social Q & A site. *Journal of the American Society for Information Science and Technology*, 60(4), 716-727.
- Krisch, S. M., Gnasa, M., Won, M., & Cremers, A. B. (2006). Beyond the web: retrieval in social information spaces. In Lalmas, M., MacFarlane, A., Rüger, S. M. Tombros, A., Tsikrika, T., & Yavlinsky, A. (Eds.), *Advances in Information Retrieval, ECIR 2006, London, United Kingdom. LNCS 3936* (pp. 84-95). Berlin: Springer.
- Kwok, C. C. T., Etzioni, O., & Weld, D. S. (2001). Scaling question answering to the web. *ACM Transactions on Information and Systems*, 19(3), 242-262.
- Lee, J. T., Kim, S. B., Song, Y. I., & Rim, H. C. (2008). Bridging lexical gaps between queries and questions on large online Q&A collections with compact translation models. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Hawaii, United States (pp. 410-418). NJ: Association of Computational Linguistics.
- Li, B., Liu, Y., Ram, A., Garcia, E. V., & Agichtein, E. (2008). Exploring question subjectivity prediction in community QA. In Myaeng, S. H., Oard, D. W., Sebastiani, F., Chua, T. S., & Leong, M. K. (Eds.), *Proceedings of the 31st Annual International ACM Conference on Research and Development in Information Retrieval*, Singapore (pp. 735-736). NY: ACM.
- Liu, Y. & Agichtein, E. (2008a). On the evolution of the Yahoo! Answers QA community. In Myaeng, S. H., Oard, D. W., Sebastiani, F., Chua, T. S., & Leong, M. K. (Eds.), *Proceedings of the 31st Annual International ACM Conference on Research and Development in Information Retrieval*, Singapore, (pp. 737-738). NY: ACM.
- Liu, Y., Bian, J., & Agichtein, E. (2008b). Predicting information seeker satisfaction in community question answering. In Myaeng, S. H., Oard, D. W., Sebastiani, F., Chua, T. S., & Leong, M. K.

- (Eds.), *Proceedings of the 31st Annual International ACM Conference on Research and Development in Information Retrieval*, Singapore (pp. 483-490). NY: ACM.
- Nam, K. K., Ackerman, M. S., & Adamic L. A. (2009). Questions in, knowledge in? a study of naver's question answering community. In Olsen, D. R., Arthur, R. B., Hinckley, K., Morris, M. R., Hudson, S. E., & Greenberg, S. (Eds.), *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, Massachusetts, United States (pp. 779-788). NY: ACM.
- Rodrigues, E. M. & Frayling, N. M. (2009). Socializing or Knowledge Sharing? Characterizing social intent in community question answering. In Cheung, D. W. K., Song, I., Chu, W. W., Hu, X., & Lin, J. (Eds.), *Proceeding of the 18th ACM Conference on Information and Knowledge Management*, Hong Kong, China (pp. 1127-1136). NY: ACM.
- Sun, K., Cao, Y., Sog, X., Song, Y. I., Wang, X., & Lin, C. Y. (2009). Learning to recommend questions based on user ratings. In Cheung, D. W. K., Song, I., Chu, W. W., Hu, X., & Lin, J. (Eds.), *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, Hong Kong, China (pp. 751-758).
- Suryanto, M. A., Lim, E. P., Sun, A., & Chiang, R. H. L. (2009). Quality-Aware collaborative question answering: methods and evaluation. In *Proceedings of the WSDM '09 Workshop on Exploiting Semantic Annotations in Information Retrieval*, Barcelona, Spain, 142-151. NY: ACM.
- Tamura, A., Takamura, H., & Okumura, M. (2005). Classification of Multiple-Sentence Questions. In Kaelbling, L. P., & Saffiotti, A. (Eds.), *Proceedings of the 2nd International Joint Conference on Natural Language Processing*, Scotland, United Kingdom. LNAI 3651 (pp. 426-437). Berlin: Springer.
- Voorhees, E. M. (2003). Overview of the TREC 2003 question answering track. In *Proceedings of 12th Text REtrieval Conference*, Maryland, United States (pp. 54). Gaithersburg: NIST.
- Wang, F. Y., Carley, K., M., Zeng, D., & Mao, W. (2007). Social computing: From social informatics to social intelligence, *IEEE Intelligent Systems*, 22(2), 79-83.
- Wang, K., Ming, Z., & Chua, T. S. (2009). A syntactic tree matching approach to finding similar questions in community-based qa services. In Allan, J., Aslam, J. A., Sanderson, M., Zhai, C. X., & Zobel, J. (Eds.), *Proceedings of the 32nd International ACM Conference on Research and Development in Information Retrieval*, Massachusetts, United States (pp. 187-194). NY: ACM.
- Xue, X., Jeon, J., Croft. (2008). Retrieval models for question and answer archives. In Myaeng, S. H., Oard, D. W., Sebastiani, F., Chua, T. S., & Leong, M. K. (Eds.), *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, Singapore (pp. 475-482). NY: ACM.
- Zheng, Z. (2002). AnswerBus question answering system. In *Proceedings of the 2nd International Conference on Human Language Technology Research*, California, United States (pp. 399-404). NY: ACM.