CrowdSTAR: A Social Task Routing Framework for Online Communities

Besmira Nushi[1](✉), Omar Alonso[2], Martin Hentschel[3], and Vasileios Kandylas[2]

[1] Department of Computer Science, ETH Zurich, Zurich, Switzerland
nushib@inf.ethz.ch
[2] Microsoft Corporation, Mountain View, CA, USA
{omalonso,vakandyl}@microsoft.com
[3] SnowFlake Computing, San Mateo, CA, USA
martin.hentschel@snowflakecomputing.com

Abstract. The online communities available on the Web have shown to be significantly interactive and capable of collectively solving difficult tasks. Nevertheless, it is still a challenge to decide how a task should be dispatched through the network due to the high diversity of the communities and the dynamically changing expertise and social availability of their members. We introduce CrowdSTAR, a framework designed to route tasks across and within online crowds. CrowdSTAR indexes the topic-specific expertise and social features of the crowd contributors and then uses a routing algorithm, which suggests the best sources to ask based on the knowledge vs. availability trade-offs. We experimented with the proposed framework for question and answering scenarios by using two popular social networks as crowd candidates: Twitter and Quora.

Keywords: Task routing · Social search · Question answering · Crowdsourcing · Expertise detection

1 Introduction

Social Task Routing is the problem of effectively routing tasks to the right crowds and the right users in online communities and social networks [4]. The need for solving tasks with the help of people is motivated by the fact that human intervention and skills can solve problems that are difficult to tackle by machines only. This motivation is even more crucial in *social search* (*i.e.* forwarding the question to a social network) where the semantic and context awareness of humans can help to increase the quality of Web search results as well as the users' satisfaction [12].

The challenge that we are addressing is how to find the best experts within the best matching crowd for a given task. This challenge is made more difficult because of the dynamically changing characteristics of networks and users. For example, the user base of a crowdsourcing platform or social network might significantly grow or shrink over time. Also, the activity of a single user may vary from being absolutely committed to only being marginally present or not present at all.

The main contribution of this work is a system for social task routing that combines expertise detection with social characteristics of users. We argue that

P. Cimiano et al. (Eds.): ICWE 2015, LNCS 9114, pp. 219–230, 2015.
DOI: 10.1007/978-3-319-19890-3_15

expertise detection is a crucial factor for the accuracy of a social task router, yet it is not sufficient. Equally important factors are the users' interactivity and availability characteristics. Therefore, the second main contribution of this paper is an exploration of the trade-offs between these dimensions. Our system, **CrowdSTAR** (**Crowd**sourced **S**ocial **Ta**sk **R**outing), investigates these aspects in the context of question-answering tasks using two popular social networks, Quora and Twitter.

2 Related Work

Social task routing stands at the boundary of major research fields like collaborative information seeking and crowdsourcing. Dustdar and Gaedke [4] were among the first to envision the general social routing principle supported by Web-scale workflows. Morris, Teevan, and Panovich [12] describe a thorough comparison between Web search and social search (*i.e.* forwarding questions to social networks). Further studies in the context of *community question answering* show that routing questions to Q&A communities increases the users' satisfaction [10,11].

The most relevant work to our problem definition is the one presented by Bozzon et al. [2]. The authors propose an generic resource-to-user graph model to represent any given crowd. Although the method does not take into account the social features of the experts, it provides a solid formal design for expertise matching and detection. Further analyses focus on the Q&A potential of crowds but yet do not make use of the social features [3,17]. Horowitz and Kamvar [6] explore the concept of social availability. Their work characterizes a social search engine (Aardvark) where people ask questions to other users via email and instant messaging. The availability of the members is not part of the user-topic model but works as a general pruning criterion. The expert search is isolated within a single network and within the circle of contacts of the person who is asking the question. Although this can be efficient for personal questions, it might not be as profitable for questions requiring a broader domain of competence. In the "IM-an-Expert" system proposed in [14], availability is not topical but it is defined as the the user status in an instant messaging system. Although the work does not use availability for task routing, it shows that it can impact the answer quality. Sung, Lee, and Lee [15] linearly combine (topical) availability with expertise into a single measure called *question affordance*.

Gathering expertise evidence in social networks is also an active field of research [1,2,5,13]. The generalized approach is to score candidate experts according to the likelihood of a person being an expert on query [1]. Pal and Counts define multiple features around the textual content a person generates [13] which are then also used as a scoring mechanism. The metrics explained in this study are built in the same spirit and elaborated for a better depiction of our vision.

An interesting line of work complementary to our study concerns task routing in networks with local knowledge [7,16]. This approach employs users in further routing tasks among each other to improve task assignment. Even though current

Table 1. Metric definitions of expertise

Metric	Definition
A	answer
CA	correct answer
P	post
OP	original non-conversational post
μ	average value among all users
N	total number of user data points

Table 2. Metric definitions of social availability

Metric	Definition
CP	conversational post
PQ	question presented to the user
AQ	question answered by the user
RT	average response time
LQ	last question presented
LA	last answer provided

results in this field are mainly theoretical and have not been studied in real-world applications, these ideas constitute a promising future work direction for CrowdSTAR and social task routing in general.

3 User Utility Model

CrowdSTAR adapts a multi dimensional user model to catalogue features of users from a utility perspective. The utility of a crowd member (*i.e.* her adequacy to solve a given task) is (1) topic-specific, (2) continuously changing, and (3) strongly affected by the user's social behavior in the network. In contrast to previous work [6], we decide to model more than one feature for each triple $<user, topic, crowd>$ and use them altogether for routing purposes.

First, we identify two main dimensions for a given user part of a certain crowd on a particular topic: *Knowledge* and *Availability*. *Knowledge* is the dimension that captures the passive or active expertise on the topic while *Availability* shows the social involvement in answering questions or conversing on the same topic. Aiming for high knowledge is crucial but not sufficient. Accounts which seem to know a lot on a particular matter can be slow or not helpful in answering questions. In addition, the definition of these two dimensions is improved by decomposing them into two other sub-features. *Knowledge* is further divided into *Qualification* and *Interest* while *Availability* is broken down into *Responsiveness* and *Activity*. Semantically, the meaning of each of the features is as follows:

1. *Qualification*: How much original and qualitative content does the user generate? A user on Quora, for example, may be active on a subject by posting questions but this does not show that he is qualified. This feature also comprises the accuracy of the user since it includes the fraction of correct answers.

$$K_1(c,u,t) = \frac{\mathrm{CA}_{(c,u,t)} + \mu}{\mathrm{A}_{(c,u,t)} + N} + \frac{\mathrm{OP}_{(c,u,t)} + \mu}{\mathrm{P}_{(c,u,t)} + N} \tag{1}$$

2. *Interest*: How active and interested is the user? The aim is to compute the degree of interest on the topic with respect to the overall user content.

$$K_2(c,u,t) = \frac{\mathrm{P}_{(c,u,t)} + \mu}{\mathrm{P}_{(c,u)} + N} \tag{2}$$

3. *Responsiveness*: How responsive is the user to conversations and questions relevant to the topic? This feature can be exploited as a discriminative filter for distinguishing advertisement/company accounts from real human members. The average response time (RT) in the definition is useful to retrieve the answers faster.

$$A_1(c,u,t) = \frac{\mathrm{AQ}_{(c,u,t)} + \mu}{\mathrm{PQ}_{(c,u,t)} + N} + \frac{\mathrm{CP}_{(c,u,t)} + \mu}{\mathrm{P}_{(c,u,t)} + N} + \frac{1}{\mathrm{RT}_{(c,u,t)}} \tag{3}$$

4. *Activity*: How long has it been since the user's last contribution on the topic? Considering that human crowd members cannot be accessed continuously, this metric helps to increase user satisfaction by keeping them engaged without overloading.

$$A_2(c,u,t) = now - \max\{time(LQ_{(c,u,t)}), time(LA_{(c,u,t)})\} \tag{4}$$

The explanation of the acronyms used in the formal definition of features is given in Table 1 and 2. The variables μ and N are used to make the expertise detection less susceptible to low-frequency users (*i.e.* users that post only a few tweets) and spammers. This technique is similar to additive smoothing or Laplace smoothing.

Expertise Detection. There are two main challenges of expertise detection: *candidate selection* and gathering *expertise evidence* [1]. Candidate selection is the problem of finding candidate experts on a particular topic. Gathering expertise evidence is the problem of determining the strength of expertise of a candidate expert given the textual evidence. Candidate selection in our approach is achieved via two steps: (i) finding user-generated documents and (ii) selecting the authors of these documents as candidate experts. In the first step, we find user-generated documents (*e.g.*, tweets, questions, answers, posts) by matching all documents of a social network on a particular topic. Matching in Twitter is performed by checking whether the topic is contained in a tweet. In Quora, the content is tagged by users or editors with the topics it belongs to. In the second step, we choose as candidate experts the set of authors of the matched documents. Gathering expertise evidence is based on the two features of the Knowledge dimension: *Qualification* and *Interest*.

Social Availability. The social dimension of our user model is also topical as for the same level of expertise, people show different response rates on different arguments due to social trends or personal preferences. *Responsiveness* captures the responsiveness of the user to our tasks as well as to posts initiated by other users in the network. At the same time, it also includes the average response time on the topic. *Activity* then keeps track of the last Q&A event with the user on the topic. This means that a user that was recently asked on a topic will not be accessed on the same topic any time soon, yet he might still be a good candidate for other topics on which he is currently idle. The routing strategy described in the next section requires that the underlying features are up to date. From our observations it results that the social *Availability* features tend to change much faster than the *Knowledge* ones and they need to be updated more often.

4 Social Task Routing

According to Law and von Ahn [9] there exist two forms of assigning tasks to crowd members, referred to as push and pull approaches. Pull approaches let the users select the tasks, while the push approaches explicitly match the tasks to users. In our work, the social task routing belongs to the second form of task assignment but CrowdSTAR design is aware of the self-regulating events that happen in dynamic crowds where members can make free choices.

Routing Tasks within a Crowd. In order to consider all the features in the user utility model, the routing algorithm needs to explore the possible trade-offs between the features and access only those users which appear to dominate the rest of the crowd. For this purpose we select as a candidate user set the group of users which is not dominated by others in at least one of the dimensions. We refer to this candidate set as the *crowd skyline* for the topic associated to the task.

Figure 1 illustrates a sample output for two dimensions where the connected points represent the crowd skyline. Depending on the topic, the crowd expertise and how much redundancy one wants from the crowd, it can happen that the number of users in the skyline is not enough. For this purpose, we decide to continue running the skyline algorithm even beyond the first skyline. For example, in Figure 1 the data points connected by the dashed line represent the second skyline. The skyline computation uses the algorithm introduced by Kossmann et al. [8]. It applies a recursive nearest-neighbor search that continuously prunes from the search space regions that are dominated by the actual best data point not yet included in the skyline. The algorithm has a good pruning rate which is a necessary property for our routing algorithm to scale. Furthermore, a good property of the algorithm is the early output of skyline points, which is useful for very large data when it is not possible to wait until the whole computation finishes. We further prune the search space by disregarding users which have very low values in at least one of the axes (the dashed regions in Figure 1) because our experiments showed that these regions contain mostly spammy and non-responsive accounts.

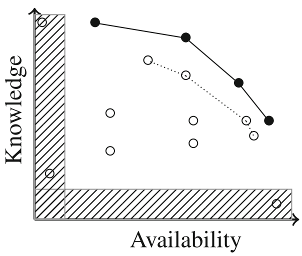

Fig. 1. Crowd skyline example

Whenever the user utility model is updated, the crowd skyline needs to be recomputed since different users may appear in the skyline. For instance, if a user has just answered a question, the respective *activity* is going to be updated with a very low value excluding this way the user from the candidate set to ask. Similarly, if someone gradually changes *interest* from photography to video and starts posting and answering more on the latter topic, the same switch will happen to his or her membership in the topic skyline. In our routing experiments we did not ask all the users in the skyline set since this would be too intrusive. Instead, we start in the middle of the skyline and then incrementally move towards the edges of the skyline in both directions. However, exploring different segments of the skyline is also effective as it gives a chance of participation and improvement to users that do not have the highest scores in all dimensions.

Routing Tasks across Multiple Crowds. The decision of crowd selection is based on an aggregate summary of each crowd. Although we index the features of all users, we do not use all of them to build the crowd summary. The summary includes only those members which will possibly be considered for question asking in the near future, *i.e.* the crowd skyline. The following formulation defines the summary of a crowd c on an arbitrary feature f for a topic t.

$$\text{Summary}(c, t, f) = \frac{\sum_{u \in \text{skyline}(c,t)} f(c, u, t)}{|\text{skyline}(c, t)|} \tag{5}$$

Having the summary on each dimension, the final crowd score of the crowd on the topic can be computed as a weighted linear combination of all the features. Note that *Activity* is excluded from the final score given that it is an individual load-balancing and diversification measure and should not affect the overall accessibility of the crowd.

$$\text{Score}(c, t) = \sum_{f \in \{K_1, K_2, A_1\}} \big(w_f \cdot \text{Summary}(c, t, f)\big) \tag{6}$$

Assigning different weights to the dimensions allows for adapting the routing algorithm to the task requirements. For example, if one is interested in solving a survey task, the highest weights should go to *interest* and *responsiveness* considering that the crowd members will only give their personal opinion and not actually solve a problem. For a fair comparison between crowds the number of users in the skyline of each crowd should be balanced which is very unlikely to happen given the different feature distributions. This problem is solved by choosing for both crowds an equal number of points as skyline representatives and moreover making use of the lower-level skylines.

5 CrowdSTAR System

CrowdSTAR is designed to help end users to solve challenging tasks with the help of human power available on the Web. One possible use case is to employ CrowdSTAR to propose to the user a set of candidate experts given the input query. Afterwards, the user can freely choose how many and which of the presented candidates to ask. In a second use case, the offered service not only finds the possible experts but also contacts them on behalf of the askers and then sends back the answers.

Components. Here we briefly describe each component of the CrowdSTAR system as depicted in Figure 2. All modules are implemented in C# and ran on a large computing cluster. The *Feature Collector* module gathers the textual evidence (*e.g.*, posts, native answers, questions, comments etc.) of users' expertise in Twitter and Quora regardless of their participation in CrowdSTAR.

Feature Monitor monitors in real time the activity of users in answering questions and sends this information to *Feature Index*. The latter uses the incoming data from the previous modules to recompute the changed dimensions of the user utility model. At the moment, the index keeps track of approximately the top 300,000 active users in Twitter and top 45,000 users in Quora. As soon as the *Feature Index* is updated on a certain topic, the *Skyline Builder* gets updated on the same topic by recomputing the skyline which is then used to refresh the crowd summary scores in the *Crowd Summarizer*. Finally, the *Task Router* routes the question according to the summarized crowd scores and the topical skylines. The posting process is done through the Twitter API while Quora does not provide an API yet and needs manual question posting. The budget here refers to the number of users to be asked as a degree of crowdsourcing redundancy that can be specified by the end user.

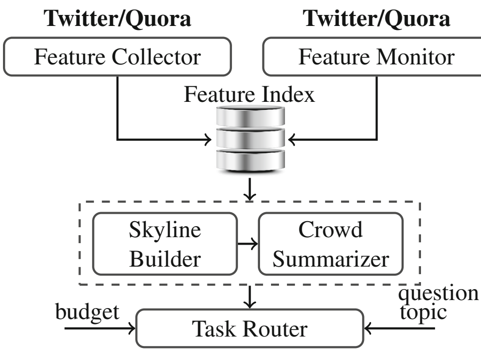

Fig. 2. CrowdSTAR architecture

6 Experimental Evaluation

In this section, we discuss how we collected the data for populating the user utility model and how we performed the routing experiments across and within Twitter and Quora. While we use proprietary data to run CrowdSTAR, we provide as much information as possible to make the experiments reproducible with similar systems.

Feature Collection. The features of each user were computed from the most recent one month interval. We focus on a broad range of topics (35 in total) from domains like technology, hobbies, news, and entertainment. In Quora we consider that a post falls within a topic if this is claimed from the author or Quora's maintenance staff since the mapping is highly accurate in this network. The posts in Twitter are not as structured. Thus, we categorize a tweet within a topic if the topic word explicitly appears in the tweet text. Involving topic-to-topic relationships would result into misleading outcomes (*e.g.*, a user who talks about *soccer* may not be an expert in *sport* and vice versa).

Table 3 shows an example of retrieving the top five users in Twitter with respect to *qualification* and *responsiveness* for topic *hiking*. Note that the most qualified users are famous accounts on the topic but not necessarily personal accounts, while the most responsive ones match to people who tend to answer and converse more on *hiking*. They are still knowledgeable but their attention is not focused on a single interest only. A similar phenomenon can also be observed in Quora.

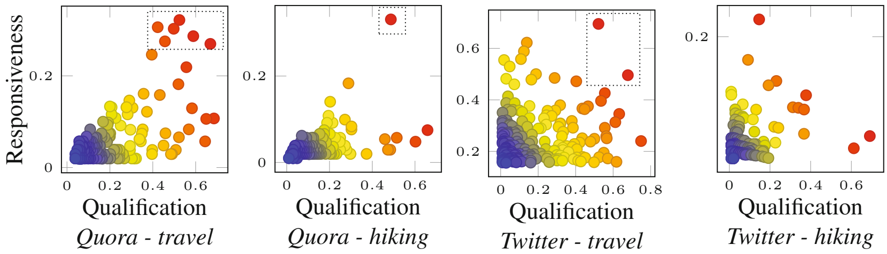

Fig. 3. Qualification and Interest for topics *travel* and *hiking* on Quora and Twitter.

In Figure 3 we show *qualification* and *responsiveness* for 200 most active users of both networks for topics *travel* and *hiking*. Users of the same color gradient would belong to the same skyline level as defined in our method. As expected, there is not necessarily a strong correlation between them (also the case for the other features) which supports once again the fact that using a linear combination or a generalization of all the features (*e.g.*, the total number of posts) is less informative and that the identified dimensions in the user utility model are present in real-world data. User data points of this nature, but of a larger scale, serve as an input for the social task routing algorithm. Ideally, we would like to choose only points that have very good scores on all the features such as those that fall within the dotted rectangles in the figure. In practice, this is not always feasible. For example, comparing the graphs for the two topics we can understand that the skyline region is more dense for popular and general topics like *travel*. For more specific ones like *hiking*, especially on Quora we can notice the existence of very few dominating experts.

Table 3. Example of top 5 Twitter users *w.r.t.* to qualification and responsiveness for topic *hiking*.

TOP 5 WHERE topic='hiking' ORDER BY			
QUALIFICATION		RESPONSIVENESS	
@hiking_camping	0.691	@thatoutdoorguy	0.223
@letsgoforahike	0.612	@nickandriani	0.169
@mightycrack	0.378	@astrogerly	0.141
@etravelhotels	0.367	@melissabravery	0.135
@outdoorgeardotd	0.367	@rsrigda	0.127

Question Posting. We created two different accounts on Twitter and Quora named respectively *@SocialQARouting* (http://twitter.com/SocialQARouting) and *Ada Floyd* (http://www.quora.com/Ada-Floyd), and used them for the purpose of conducting question routing experiments. Both accounts were first bootstrapped by gradually asking questions and posting other non-asking content. In Twitter we alternated two asking strategies: *introductory* and *simple greeting*. Also, we attached the #ask and #<topic> hashtags to the question text

Table 4. Comparison of Twitter and Quora task routing support.

	Quora	**Twitter**
Responsiveness	64%	44%
Questions answered	85%	44%
Average response time	∼24 hours (1st response)	12.7 hours
Accuracy	80%-90% (manually evaluated)	
Asking tone	Formal	Informal
Question visibility	Many users	Mainly the assigned user
Human intervention	Thanking (built-in)	Introduction Greeting Thanking
Answer properties	Long and elaborated	140 char max and concise
Quality control	Upvotes and Editors	Candidates: #retweets, #favorites
Types of Q&A	Recommendation multiple items "How to" explanations Only interesting surveys	Recommendation single item Laconic explanations Survey and factual answers

to increase the interest of the user. We noticed that the most famous accounts prefer the introductory strategy while the others prefer a simple greeting. Quora members instead are used to a formal asking tone in contrast to Twitter where people tend to converse in a more relaxed and friendly way. The question promotion process was easier in Quora because it is intentionally designed for Q&A.

The main conclusion of this part of the work is that the networks that are primarily designed for task-solving need less human intervention for both the bootstrapping and the promotion phases because many necessary steps like introduction, thanking, rating, and rewarding are inherently present. Crowds of a more general purpose require additional human steps in the workflow, otherwise people tend to be reluctant to help. Indeed, in earlier stages of this project when we did not include any greeting, introductory or thanking messages the interaction was not satisfactory.

Task Routing. Table 4 shows the main results from routing 100 tasks to the targeted crowds. We received answers to 44% of the questions in Twitter and to 85% of the questions in Quora. Nevertheless, only 64% of the answers in Quora came from the users we pointed. The rest were given by other users interested in the same topic. The answers' accuracy was manually evaluated and varies between 80%–90% which confirms that when people feel confident to answer they are able to provide accurate insights.

Another major difference between the two crowds consists on the type of questions they can accommodate. Due to the message length restrictions in Twitter, it is possible to ask only short questions that can be answered with short replies. The answers in Quora are more elaborated and accordingly argued. We show some examples of questions routed to Quora and Twitter along with the respective retrieved answers in Table 5.

Table 5. Examples of questions routed to Twitter and Quora and the retrieved answers.

Question (Twitter)	**Answer (Twitter)**
@joshuariggins Do you know any good travel coffee mugs preferrably working for both cold and hot weather? #ask #travel	@SocialQARouting try @HydroFlask they are amazing. After 5 hrs with 170° coffee in it, left in 27° snow, it was still 115° hot
@NicoArts Why do you think magic realism is strongly related to the Latin American culture?	@SocialQARouting Sure! Seems to me it's almost entirely because of the works of author Gabriel García Márquez, link: http://en.wikipedia.org/wiki/Gabriel_Garc...
Hi @rickasaurus! Do you know whether there exists an active Machine Learning community within Twitter?	@SocialQARouting you may find this list helpful in your ongoing search https://twitter.com/rickasaurus/...
Question (Quora)	**Answer (Quora)**
What is the best way to rest during rock climbing?	When resting, remember to visualize the moves ahead, focus, breathe and try to release lactic acid from your arms. Read more: http://qr.ae/EsG00
How can a high-school history teacher make the class particularly interesting for the students?	Informative Wall Art. When in the eighth grade I had a history teacher whose room was an engaging learning aid because of the maps and posters on his wall. Read more: http://qr.ae/Estrp
How effective is orthodontics in grown ups?	Orthodontics for adults is very effective and given the recent advancements in orthodontic treatment, adult treatment is quicker and more convenient than ever. Read more: http://qr.ae/Esnpu

In both networks, members preferred to answer questions related to specific topics like *biking*, *hiking*, *poker* rather than general ones like *music*, *sport*, *travel*. A possible reason for this is that people tend to answer more on topics in which they have experience and are particularly enthusiastic of. This phenomenon constitutes an important implicit incentive for most of the Q&A applications and also for our study. According to our profile statistics in Quora, the most difficult questions to answer are those that either (i) belong to a narrow expertise domain (*e.g.*, "How much usability and cognition study is done before starting an architectural project?") or (ii) combine two domains together (*e.g.*, "What is an alternative backup solution for Mac OS X that is similar to Cobian?"). The most popular questions in terms of number of views, followers, and answer quality are queries that contain elements of entertainment, curiosity or professional interest (*e.g.*, "What is the most efficient starting strategy for Settlers of Catan?"). Another successful use-case for Q&A in CrowdSTAR is information gathering for building lists (*e.g.*, "Which are some well-known movie actors who also play on theater stages?'). In these cases, it is difficult to gather the whole answer by asking single individual. Many users instead are able to construct a complete and relevant answer.

7 Conclusions

In this paper, we proposed a general model for task routing in online crowds that combines expertise detection with social availability features. Furthermore, we presented the design and implementation of CrowdSTAR[1], a social task routing system. CrowdSTAR routes questions to responsive experts in an appropriate crowd. Yet, the system makes sure to not overload experts with requests by regulating the number of questions routed to individual users. CrowdSTAR currently supports two popular social networks, Twitter and Quora, but the architecture is extensible to other crowds. Our findings show that the proposed user utility model exists in real social networks and that experts are willing to answer questions which are more specific rather than general.

References

1. Balog, K., Fang, Y., de Rijke, M., Serdyukov, P., Si, L.: Expertise retrieval. Foundations and Trends in Information Retrieval **6**(2–3), 127–256 (2012)
2. Bozzon, A., Brambilla, M., Ceri, S., Silvestri, M., Vesci, G.: Choosing the right crowd: expert finding in social networks. In: Proc. of the 16th EDBT, pp. 637–648. ACM (2013)
3. Difallah, D.E., Demartini, G., Cudré, P.: Pick-a-crowd: tell me what you like, and i'll tell you what to do. In: Proc. of the 22nd WWW, pp. 367–374 (2013)
4. Dustdar, S., Gaedke, M.: The social routing principle. IEEE Internet Computing **15**(4), 80–83 (2011)
5. Ghosh, S., Sharma, N., Benevenuto, F., Ganguly, N., Gummadi, K.: Cognos: Crowdsourcing search for topic experts in microblogs. In: Proc. of the 35th SIGIR, pp. 575–590. ACM (2012)
6. Horowitz, D., Kamvar, S.D.: The anatomy of a large-scale social search engine. In: Proc. of the 19th WWW, pp. 431–440. ACM (2010)
7. Kleinberg, J.: The small-world phenomenon: An algorithmic perspective. In: Proc. of the 32nd Annual STOC, pp. 163–170. ACM (2000)
8. Kossmann, D., Ramsak, F., Rost, S.: Shooting stars in the sky: An online algorithm for skyline queries. In: Proc. of the 28th VLDB, pp. 275–286 (2002)
9. Law, E., von Ahn, L.: Human computation. In: Synthesis Lectures on Artificial Intelligence and Machine Learning, 13 (2011)
10. Li, B., King, I.: Routing questions to appropriate answerers in community question answering services. In: Proc. of the 19th CIKM, pp. 1585–1588. ACM (2010)
11. Liu, Y., Bian, J., Agichtein, E.: Predicting information seeker satisfaction in community question answering. In: Proc. of the 31st SIGIR, pp. 483–490. ACM (2008)
12. Morris, M.R., Teevan, J., Panovich, K.: A comparison of information seeking using search engines and social networks. In: Proc. of the 4th ICWSM, pp. 23–26. AAAI (2010)

[1] The technical report of the project was earlier released at http://arxiv.org/pdf/1407.6714v1.pdf.

13. Pal, A., Counts, S.: Identifying topical authorities in microblogs. In: Proc. of the 4th WSDM, pp. 45–54. ACM (2011)
14. Richardson, M., White, R.W.: Supporting synchronous social q&a throughout the question lifecycle. In: Proc. of the 20th WWW, pp. 755–764. ACM (2011)
15. Sung, J., Lee, J.G., Lee, U.: Booming up the long tails: discovering potentially contributive users in community-based question answering services. In: Proc. of the 7th ICWSM (2013)
16. Zhang, H., Horvitz, E., Chen, Y., Parkes, D.C.: Task routing for prediction tasks. In: Proc. of the 11th AAMAS, pp. 889–896. AAMAS (2012)
17. Zhou, Y., Cong, G., Cui, B., Jensen, C.S., Yao, J.: Routing questions to the right users in online communities. In: Proc. of the 25th ICDE, pp. 700–711. IEEE (2009)