

Utah State University

DigitalCommons@USU

All Graduate Theses and Dissertations

Graduate Studies

8-2015

Detecting Experts on Quora: by their Activity, Quality of Answers, Linguistic Characteristics and Temporal Behaviors

Sumanth Kumar Reddy Nagarur Patil
Utah State University

Follow this and additional works at: <https://digitalcommons.usu.edu/etd>



Part of the [Computer Sciences Commons](#)

Recommended Citation

Nagarur Patil, Sumanth Kumar Reddy, "Detecting Experts on Quora: by their Activity, Quality of Answers, Linguistic Characteristics and Temporal Behaviors" (2015). *All Graduate Theses and Dissertations*. 4453.
<https://digitalcommons.usu.edu/etd/4453>

This Thesis is brought to you for free and open access by the Graduate Studies at DigitalCommons@USU. It has been accepted for inclusion in All Graduate Theses and Dissertations by an authorized administrator of DigitalCommons@USU. For more information, please contact digitalcommons@usu.edu.



DETECTING EXPERTS ON QUORA: BY THEIR ACTIVITY, QUALITY OF
ANSWERS, LINGUISTIC CHARACTERISTICS AND TEMPORAL BEHAVIORS

by

Sumanth Reddy

A thesis submitted in partial fulfillment
of the requirements for the degree

of

MASTER OF SCIENCE

in

Computer Science

Approved:

Dr. Kyumin Lee
Major Professor

Dr. Nick Flann
Committee Member

Dr. Tung Nguyen
Committee Member

Dr. Mark R. McLellan
Vice President for Research and
Dean of the School of Graduate Studies

UTAH STATE UNIVERSITY
Logan, Utah

2015

Copyright © Sumanth Reddy 2015

All Rights Reserved

ABSTRACT

Detecting Experts on Quora: By Their Activity, Quality of Answers, Linguistic
Characteristics and Temporal Behaviors

by

Sumanth Reddy, Master of Science

Utah State University, 2015

Major Professor: Dr. Kyumin Lee

Department: Computer Science

Quora is a fast growing social Q&A site where users create and answer questions, and identify the best answers by upvotes and downvotes with crowd wisdom. Unfortunately, little is known about properties of experts and non-experts and how to detect experts in general topics or a specific topic. To fill the gaps, in this paper we (i) analyze behaviors of experts and non-experts in five popular topics; (ii) propose user activity features, quality of answer features, linguistic features and temporal features to identify distinguishing patterns between experts and non-experts; and (iii) develop statistical models based on the features to automatically detect experts. Our experimental results show that our classifiers effectively identify experts in general topics and a specific topic, achieving up to 97% accuracy and 0.987 AUC.

(33 pages)

PUBLIC ABSTRACT

Detecting Experts on Quora: By Their Activity, Quality of Answers, Linguistic Characteristics and Temporal Behaviors

Sumanth Reddy

Question and answering sites are useful in sharing the knowledge by answering questions. It is a medium of sharing knowledge. Quora is the fastest emerging effective Q&A site, which is the best source of knowledge. Here you can ask a question, and get help in getting answers from people with firsthand experience, and blog about what you know. In this paper, we are investigating and identifying potential experts who are providing the best solutions to the questioner needs. We have considered several techniques in identifying user as an expert or non-expert. We have targeted the most followed topics in Quora and finally came up with five topics: Mathematics, Politics, Technology, Sports and Business. We then crawled the user profiles who are following these topics. Each topic dataset has many special features. Our research indicates that experts are quite different from normal users and tend to produce high quality answers to as many questions as possible to gain their reputation. After evaluation, we got a limited number of experts who have potential expertise in specific fields, achieving up to 97% accuracy and 0.987 AUC.

CONTENTS

	Page
ABSTRACT	iii
PUBLIC ABSTRACT	iv
LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER	
1 INTRODUCTION	1
2 DATASET	3
2.1 Dataset collection	3
2.2 Dataset Labelling	3
3 ANALYZING BEHAVIORS OF EXPERTS AND NON-EXPERTS	5
3.1 Analyzing Behaviors of Experts and Non-Experts: By Activity	5
3.2 Analyzing Behaviors of Experts and Non-Experts: By Linguistic Character- istics	7
4 DETECTION OF EXPERTS	9
4.1 Features	10
4.2 Feature Selection and Analysis	11
4.3 Experiments	12
5 DETECTING EXPERTS WITH TEMPORAL FEATURES	15
5.1 Data Collection	15
5.2 Analysis of temporal behaviors of experts and non-experts.	15
5.3 Temporal features	18
5.4 Experiments	19
6 DISCUSSION: ADDING FEATURES EXTRACTED FROM AN EXTERNAL SOURCE – TWITTER	20
7 RELATED WORK	21
8 CONCLUSION	23
REFERENCES	24

LIST OF TABLES

Table		Page
2.1	Dataset consisting of expert and non-expert profiles collected from five topics.	4
4.1	Features.	9
4.2	Classification results of the combined dataset.	13
4.3	Classification results of the five topics-based datasets.	13
5.1	Another dataset containing 22 user profile snapshots of each user.	17
5.2	Classification results without/with temporal features.	19

LIST OF FIGURES

Figure	Page
1.1 An example of Quora user profile.	2
3.1 Cumulative Distribution Functions of the number of followers, edits, words in answers and questions between experts (blue line with circles) and non-experts (red line with stars).	6
3.2 Three linguistic characteristics of experts (blue line with circles) and non-experts (red line with stars).	7
4.1 Top five features in the combined dataset.	11
4.2 Top five features in each of the five topic-based datasets.	12
5.1 Three types of experts grouped by their weekly change of the number of answers.	16
5.2 Weekly change in the average number of followers, edits and answers. . . .	17

CHAPTER 1

INTRODUCTION

Social Q&A sites are becoming more and more popular because people can post questions, get answers, and befriend with experts. The answers are evaluated by the number of upvotes and downvotes with crowd wisdom. These interactions naturally reveal the best answer for a question. Sometimes, new questions in these Q&A sites stimulate answerers to disseminate curated knowledge which may not be available in other websites or it may take time for a user to find, understand and summarize relevant information from other sites. For example, some raw information may be spread across several websites. It would take time for a user to search and understand these pages. Or the information may be not available on the Web. In this case, people may visit a social Q&A site and post a question, expecting experts would give them answers. As the popularity of social Q&A sites has increased, various platforms have emerged – from general-purpose social Q&A platforms such as Quora and Yahoo Answers to specialized platforms such as Stack Overflow (for programming) and Super User (for computer). As social Q&A sites have become popular with the number of users, people have desire to quickly identify experts in general topics or a specific topic. New users may not be familiar with the community, but they want to find experts who could give them relevant answers. Also, expert identification can be used for a expert recommendation system as a service of a social Q&A site. Unfortunately, little is known about properties of experts and non-experts, and how to detect experts in general topics or a specific topic. Hence, in this paper we choose Quora – a fast growing social Q&A site and the 200th most popular site to be the first to answer the following questions: Do experts and non-experts behave differently? Do they change their behaviors over time? Do answers of experts and non-experts contain their linguistic characteristics? Can we measure quality of answers? Based on this analysis and the corresponding observations, can we automati-

Matt Bodnick
Local/Global business & community issues
• Currently interested in business & community issues of local
• Co-founder, Education Partners
• Former board member, Insp. EducationPartners, Boston ...

Follow | 41.7k Facebook Twitter LinkedIn

FEEDS

All Activity

Questions	11,549
Answers	5,351
Posts	455

Followers 41,792

Following 5,778

Edits 267,368

American Sniper (2014 movie): In the movie story (American Sniper), why didn't Chris's wife leave him?... (continue)
Answer added.
In the context of the movie story only, I'd say:
• He is going to come bac...
#57398858 • 20 Jan, 2015 10:58 PM

Why does the media seem to portray the closure of the Google Glass Explo... (continue)
Answer added.
In large part, because Google Glass so far has been a really unexciting ...

Figure 1.1: An example of Quora user profile.

cally detect experts in general topics or a specific topic? Will adding temporal (dynamic) properties of experts and non-experts improve the success rate of expert detection?

To answer these questions, we make the following contributions in this paper:

- First, we collect user profiles from five popular topics on Quora, and analyze the properties of experts and non-experts.
- Second, we extract and analyze user activity features, quality of answer features, linguistic features based on the Linguistic Inquiry and Word Count (LIWC), and temporal features.
- Third, we develop statistical models based on the proposed features to detect experts in general topics and a specific topic. We evaluate what types of classifiers produce the best result. To our knowledge, this is the first study to focus primarily on Quora for analyzing behaviors of experts and non-experts and detecting experts in general topics and a specific topic.
- Finally, we study whether adding additional features extracted from a user's external accounts such as social media accounts would improve performance of expert detection.

CHAPTER 2

DATASET

2.1 Dataset collection

In order to analyze behaviors of experts and non-experts on Quora, the first step is to collect user information. Since there were no publicly available Quora dataset and no official APIs, we developed our own crawler which collected user information on Quora. Our crawling strategy is to first collect five popular topic pages each of which consists of a list of user profile URLs. Users on Quora can follow any topic that they like. Following a topic indicates that they are interested in the topic. We chose five topics such as Mathematics, Business, Politics, Sports and Technology. From each topic page, we randomly selected users, and the crawler collected these users' profiles consisting of information related to user activities and answers posted by the user. Figure 1.1 shows an example of a Quora user profile which consists of user activity related information such as the number of following, the number of followers and the number of answers, and a list of answers that the user have posted. By running our crawler, we collected 3,720 profiles of users who were interested in one of the five topics in October, 2013.

2.2 Dataset Labelling

Next step is to label the dataset to get the ground truth (i.e., which profile is an expert's profile or a non-expert's profile). In general, an expert is a person who has a comprehensive and authoritative knowledge of or skill in a particular area. An expert is supposed to post high quality answers and actively post answers (i.e., how many times a user posted answers). Quora provides a feature called *upvote*. Other users evaluate an answer and if they like it, they upvote it. The more number of upvotes an answer gets, the better the answer is. We

Table 2.1: Dataset consisting of expert and non-expert profiles collected from five topics.

Topic	Experts	Non-Experts	Users
Business	74	698	772
Mathematics	94	683	777
Politics	114	742	856
Sports	82	533	615
Technology	68	632	700
Combined Dataset	432	3,288	3,720

labelled each user as an expert or non expert manually with close obseravtion into user profile. We spent 7 working days on the annotation. While labelling for an user’s expertise, we looked at:

- User reputation based on Follower count and Following count.
- Number of answers provided.
- Number of Upvotes user has gained while answering to an question.
- User’s Profession.
- Endorsement count an user has got for a specific topic by other users (This feature is limited and is available for few users).
- Quality of answers, etc.

Each topic-based data contained between 615 and 856 user profiles. Overall, there were 432 experts and 3,288 non-experts in the combined dataset.

CHAPTER 3

ANALYZING BEHAVIORS OF EXPERTS AND NON-EXPERTS

In the previous section, we presented the collected dataset consisting of expert profiles and non-expert profiles. In this section, we analyze behaviors of experts and non-experts on Quora. First we compare four activities of experts and non-experts.

3.1 Analyzing Behaviors of Experts and Non-Experts: By Activity

How many followers do experts and non-experts have? Quora provides following and follower features like Twitter does. A user can control the number of following, but cannot control the number of followers. An interesting research question is “will experts have a larger number of followers than non-experts?” Figure 3.1(a) presents a cumulative distribution function (CDF) of the number of followers between experts and non-experts. The number of followers of experts is greater than the number of followers of non-experts. Since we are analyzing users on a social Q&A site, we conjecture that users tend to follow experts who have posted high quality answers. This is interesting phenomena compared with following celebrities on social media sites like Twitter and Facebook.

How many edits have experts and non-experts made? A user profile contains the number of edits which means how many times a user edited postings (e.g., editing answers, editing questions). This number would indicate how active a user is on Quora. Figure 3.1(b) shows that experts have made a larger number of edits than non-experts, indicating that experts are users who were more active than non-experts. This is an interesting observation. Naturally following questions are “Have experts posted longer answers than non-experts?” and “how many questions have experts and non-experts posted?”




Figure 3.1: Cumulative Distribution Functions of the number of followers, edits, words in answers and questions between experts (blue line with circles) and non-experts (red line with stars).

Have experts posted longer answers than non-experts? To answer this question, we counted the average number of words in answers created by experts and non-experts. Figure 3.1(c) shows CDFs of the average number of words. Until reaching to 0.9 in y -axis (i.e., 90% of experts and non-experts), experts have posted longer answers than non-experts. But, some non-experts (the above 0.9 in y -axis value) have posted longer answers.

How many questions have experts and non-experts posted? Figure 3.1(d) shows CDFs of the number questions that experts and non-experts have posted. People may think experts would be only interested in answering questions. But, surprisingly experts have posted a larger number of questions than non-experts. We conjecture that some experts may be knowledgeable in a specific topic, but may be not knowledgeable in other topics, so they may post many questions related to other topics.

So far, we have analyzed four activities of experts and non-experts. Next, we study the linguistic characteristics of answers posted by experts and non-experts.




Figure 3.2: Three linguistic characteristics of experts (blue line with circles) and non-experts (red line with stars).

3.2 Analyzing Behaviors of Experts and Non-Experts: By Linguistic Characteristics

Do experts create answers with different language use? To answer this question, we used the Linguistic Inquiry and Word Count (LIWC) dictionary, which is a standard approach for mapping text to psychologically-meaningful categories [1]. LIWC-2001 defines 68 different categories, each of which contains several dozens to hundreds of words. Given each user's tweets, we measured his linguistic characteristics in the 68 categories by computing his score of each category based on LIWC dictionary. First we counted the total number of words in his tweets (N). Next we counted the number of words in his tweets overlapped with the words in each category i on LIWC dictionary (C_i). Then, we computed his score of a category i as C_i/N .

Figure 3.2 shows linguistic characteristics of experts and non-experts in three categories such as Home (e.g., house, kitchen, lawn), TV/Movies (e.g., TV, sitcom, cinema) and Touch

(e.g., touch, hold, felt). Most non-experts have not used any Home, TV/Movies and Touch related words in their answers. But, some non-experts have used these words many times. Interesting, 94% experts have used words in these categories with less than 0.005 score.

In summary, we have analyzed behaviors of experts and non-experts by their activity and linguistic characteristics. We observed that their behaviors were clearly different. These observations motivated us to do further study.

CHAPTER 4

DETECTION OF EXPERTS

In the previous section, we have analyzed behaviors of experts and non-experts and found that their activities and linguistic characteristics are different. Based on the observations, in this section we propose features, measure distinguishing power of the features, and then develop and test expert classifiers (see Table 4.1).

Table 4.1: Features.

Group	Feature
AF	the number of edits
AF	the number of posted questions
AF	the number of followers
AF	the percentage of bidirectional friends: $\frac{ following \cap followers }{ following }$
QAF	the average number of words in posted answers
QAF	the average number of uppercase words in posted answers
QAF	subjectivity of answers: the average number of subjective words in posted answers
QAF	Average Upvotes: the average number of upvotes for the answers provided by an user
QAF	entropy of answers
QAF	readability of answers
LF	68 LIWC features, which are Total Pronouns, 1st Person Singular, 1st Person Plural, 1st Person, 2nd Person, 3rd Person, Negation, Assent, Articles, Prepositions, Numbers, Affect, Positive Emotions, Positive Feelings, Optimism, Negative Emotions, Anxiety, Anger, Sadness, Cognitive Processes, Causation, Insight, Discrepancy, Inhibition, Tentative, Certainty, Sensory Processes, Seeing, Hearing, Touch, Social Processes, Communication, Other References to People, Friends, Family, Humans, Time, Past Tense Verb, Present Tense Verb, Future, Space, Up, Down, Inclusive, Exclusive, Motion, Occupation, School, Job/Work, Achievement, Leisure, Home, Sports, TV/Movies, Music, Money, Metaphysical States, Religion, Death, Physical States, Body States, Sexual, Eating, Sleeping, Grooming, Swearing, Nonfluencies, and Fillers

4.1 Features

To build an expert classifier, we need to convert user profile information to meaningful feature values. Based on our previous analysis and observations, we propose 78 features as shown in Table 4.1 and grouped them into the following three categories:

- **Activity Features (AF):** These features measure a user’s activities on Quora. They consist of the number of edits, the number of questions, the number of followers and the number of bidirectional friends. The first three features indicate activity levels of the user. The last feature measures how many bidirectional friends the user has.
- **Quality of Answer Features (QAF):** These features are extracted from a user’s aggregated answers. The features consist of the average number of words, the average number of uppercase words, the average number of subjective words in answers, entropy of answers and readability of answers and average number of Upvotes for whole answers. To count the number of subjective words, we used the Subjectivity Lexicon [2] which consists of 8,222 subjective words collected from various information sources.

We measured the complexity of answers by the entropy of the words in the answers:

$$entropy(a_j) = - \sum_{i=1}^k P(x_i) \log P(x_i) \quad (4.1)$$

, where k is the number of distinct words in answers, and $P(x_i)$ is $\frac{\text{frequency of a word } i}{\text{total number of words } n \text{ in answers}}$.

A low entropy score indicates that answers contain a few words or repetitive words. A high entropy score indicates that a user’s answers contain various words and are complex. In other words, the user with a high entropy score has knowledge to use various words, and know how to present complex or complicated ideas.

The readability of aggregated answers was measured by the following SMOG formula:

$$1.043 \sqrt{|polysyllables| \times \frac{30}{|sentences|}} + 3.1291 \quad (4.2)$$




Figure 4.1: Top five features in the combined dataset.

The SMOG grade estimates the years of education needed to understand a piece of writing [3].

- **Linguistic Features (LF):** Researchers have found that word usage in one’s writings is related to one’s personality or linguistic characteristics [4, 5]. By using LIWC [1], we measured 68 linguistic features as shown in Table 4.1. Each feature is a word category which contains up to hundreds of words selected by psychologists. We extract these features from answers posted by a user. Detailed information regarding how we calculated these features was described in the previous section.

4.2 Feature Selection and Analysis

Before building classifiers, we conduct feature selection to make sure only using features having positive distinguishing power between experts and non-experts. To measure discriminative power of our proposed 78 features, we computed the χ^2 value [6] of each of the features. Our χ^2 test results showed that all features had positive discriminative power, though with different relative strengths.

Next, we measured the Mean Decrease Accuracy (MDA) of Random Forests which is another method to measure importance features. The larger its mean decrease accuracy is, the more important a feature is. We measured MDA of all features in the combined dataset in Table 2.1. Figure 4.1 shows top five important features – the number of edits, the number of followers, entropy of answers, TV/Movies and Anger in LIWC.

While we were measuring MDA of features in the combined dataset, an interesting question was raised. “Does each topic-based dataset have a different set of top features?” To answer this question, we measured MDA of all features in each topic-based dataset presented in Table 2.1. Figure 4.2 shows the experimental results. Interestingly, important features varied across the topic-based datasets. A commonly important feature was the number of edits. Some LIWC features like TV/Movies and Home in Sports, and Family and Assent in Politics were considered as important features. We conjecture that while experts answered sports related questions, they might express what they watched (e.g., NFL games) on TV with some feelings like some sadness and joy because these experts wanted to deliver detailed information and feelings regarding the sports.

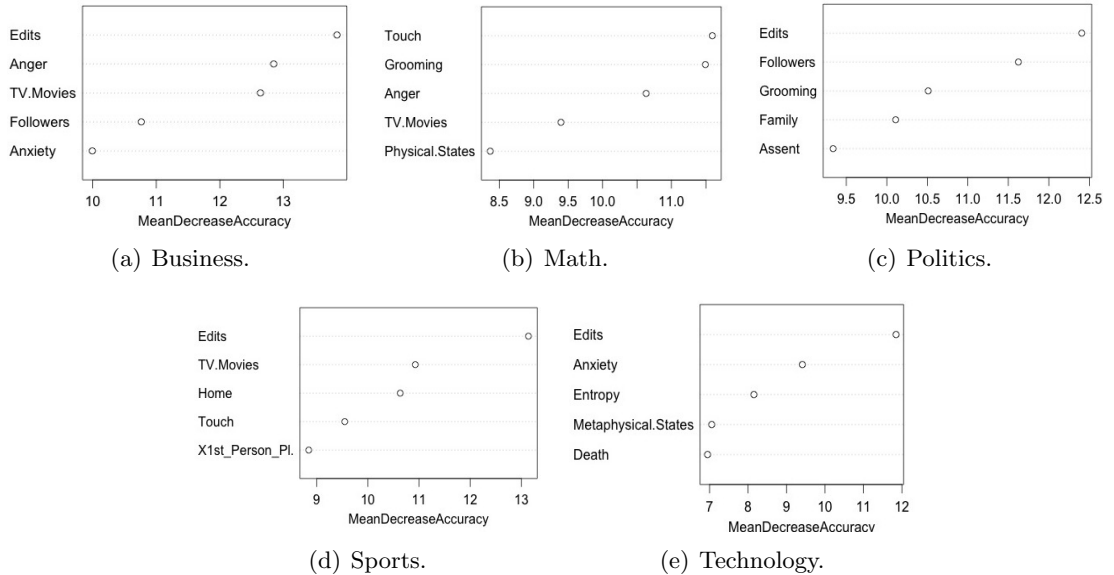


Figure 4.2: Top five features in each of the five topic-based datasets.

4.3 Experiments

So far, we have learned that all of the proposed features have positive discriminative powers, and each topic-based dataset has had a different order of important features. Based on this analysis and observation, now we turn to develop classifiers to see whether

Table 4.2: Classification results of the combined dataset.

Classifier	AUC	Accuracy (%)
J48	0.843	94.00
Random Forest	0.979	95.94
SMO	0.509	88.52

automatically detect experts in the *combined dataset* (i.e., containing general topics – multiple topic-based datasets) is possible. Further, we develop topic-specific classifiers to test whether we can detect experts in each *topic-based dataset*.

Combined Dataset. We chose three classification algorithms such as Random Forest, J48 and SMO (SVM) to compare how their classification performances are different and which one is the best. We used Weka [7], a machine learning toolkit consisting of implementations of these algorithms. First, we converted the *combined dataset* consisting of profiles of 432 experts and 3,288 non-experts in Table 2.1 to feature values, and ran 10-fold cross-validation for each classification algorithm. Table 4.2 shows classification results of the combined dataset by measuring area under the ROC curve (AUC) and accuracy of each classification algorithm. Random Forest classifier outperformed J48 and SMO classifiers, achieving 0.979 AUC and 95.94% accuracy.

Table 4.3: Classification results of the five topics-based datasets.

Topic	Classifier	AUC	Accuracy (%)
Business	J48	0.803	95.07
	Random Forest	0.976	96.37
	SMO	0.52	90.80
Mathematics	J48	0.777	91.89
	Random Forest	0.972	94.20
	SMO	0.5	87.90
Politics	J48	0.83	93.22
	Random Forest	0.983	96.61
	SMO	0.512	86.79
Sports	J48	0.811	92.52
	Random Forest	0.987	96.91
	SMO	0.529	87.15
Technology	J48	0.767	91.71
	Random Forest	0.951	95.42
	SMO	0.499	90.14

Topic-based Datasets. As we described in our data collection strategy in Section “Dataset” we intentionally collected five topic-based datasets – Technology, Politics, Sports, Mathematics and Business. Users in each dataset were interested in the topic and followed the topic. For example, experts were interested in the topic and wanted to observe what kind of questions had been posted in this topic thread. An interesting research question is “can we detect a topic-specific experts by using the classification approach”? As we observed in the previous subsection, importance of features varied in each topic-based dataset. Developing topic-specific classifiers would make sense. To answer the research question, we developed three classifiers in each topic-based dataset. Table 4.3 shows 15 classifiers’ experimental results after running 10-fold cross-validation. Overall, Random Forest classifier outperformed J48 and SMO classifiers in all five topic-based datasets. Random Forest classifiers achieved 96.37%, 94.20%, 96.61%, 96.91% and 95.42% in Business, Mathematics, Politics, Sports and Technology, respectively. Especially, topic-specific classifiers for Sports and Technology achieved higher accuracies than the classifier built based on the combined dataset. These results show that building statistical models in each topic is possible, and the models work well in detecting topic-specific experts.

In summary, we have thoroughly analyzed our proposed features, and developed two types of classifiers – (i) a universal classifier to detect experts in general topics (containing multiple topics); and (ii) topic-based classifiers to detect topic-specific experts. The both types of classifiers worked well, and achieved up to 96% accuracy and 0.979 AUC.

CHAPTER 5

DETECTING EXPERTS WITH TEMPORAL FEATURES

In previous section, we developed classifiers based on static features extracted from a snapshot of user profiles. In this section, we are interested in studying temporal behaviors of experts and non-experts. Do they have clearly different temporal patterns? If yes, can we use these temporal patterns to improve performance of expert classifiers?

5.1 Data Collection

To answer these questions, we collected another dataset presented in Table 5.1 in November 2014. Our data collection strategy is that first we randomly selected 786 users. Then we collected their profiles once per day during 22 consecutive days. In other words, each day we got one snapshot of each user, in total we collected 22 snapshots of each user. By using the labeling method in normal datasets, we got the ground truth. Finally, the dataset consisted of user profiles (22 user profile snapshots of each user) of 114 experts and 672 non-experts

5.2 Analysis of temporal behaviors of experts and non-experts.

Next, we analyze temporal behaviors of experts and non-experts. First, to understand what types of experts exist on Quora in terms of different temporal behaviors, we calculated two variables – (i) the average value of weekly change of the number of answers; and (ii) standard deviation of weekly change of the number of answers. We calculated these two variable values for each expert. Then, we grouped the experts to three categories as shown in Figure 5.1:

- **Fluctuating Experts (18.4%):** The fluctuating experts have posted different number of new answers each week. For example, these experts posted less number of answers




Figure 5.1: Three types of experts grouped by their weekly change of the number of answers.

in the first week. Then they posted more in the second week and posted a little less number of answers. We found 21 (18.4%) fluctuating experts in the dataset.

- **Stable (and Active) Experts (54.4%):** The stable experts are very active experts who post almost similar number of answers every week. They constantly provided answers and played a prominent role on disseminating knowledge on Quora. We found 62 (54.4%) stable experts in the dataset.
- **Idle Experts (27.2%):** The idle experts have posted very less number of answers constantly every week. They posted high quality answers, but posted very less number of answers recently. We conjecture that they used to be active and posted high quality answers, but might lose passion on posting more answers on Quora. Q&A service providers should think of how to motivate them to become active (stable) experts again. We found 31 (27.2%) idle experts in the dataset.

Second, we analyze how average number of followers, edits and answers of experts and non-experts had been changed over time. To do this, we measured weekly change of


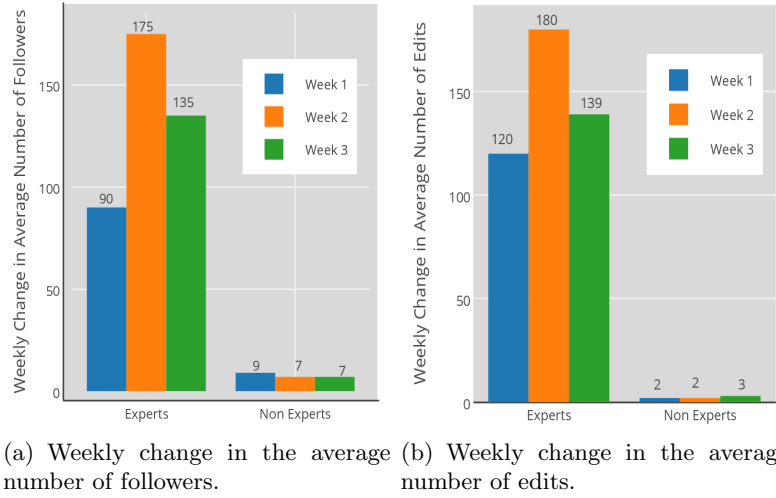


Figure 5.2: Weekly change in the average number of followers, edits and answers.

Table 5.1: Another dataset containing 22 user profile snapshots of each user.

Experts	Non-Experts	Users
114	672	786

followers, edits and answers. Figure 5.2 shows weekly change results of experts and non-experts. Change in average number of followers between experts and non-experts was very different. Experts had a larger number of weekly change in average number of followers than non-experts. More number of other users followed experts than non-experts. We conjecture that other users tend to follow experts to get useful information. Weekly change

in average number of edits and answers also followed similar patterns with weekly change in average number of followers. Experts were more active than non-experts by making a larger number of edits. Experts were also posted more number of answers than non-experts during the period of 22 days. Of course, we observed that some non-experts had increased the number of answers over time. We conjecture that these non-experts have a high change to be experts in the future. The temporal data analysis presents that experts and non-experts have different temporal behaviors. With the positive observations, next we extract temporal features toward building an expert classifier.

5.3 Temporal features

In each user profile snapshot, we extracted 5 variable values – the number of following, followers, edits, questions and answers. By doing this for 22 user profile snapshots of each user, we got 22 time-series values of each variable (e.g., number of following, number of followers, number of edits). Then, we computed following temporal features for each variable:

- **Average daily change (in total five features):** From 22 time-series values of each variable, we calculated daily change (increase/decrease) between two consecutive days. Then we averaged these values. Finally, we got five average daily change features each of which was calculated from each variable.
- **Standard deviation of daily change (in total five features):** We measured standard deviation of 21 daily changes of each variable.
- **Probability of average daily change on the day of a week (in total 35 features):** These features capture the day’s average change. In this context, the day of a week means Monday, Tuesday, Wednesday, Thursday, Friday, Saturday or Sunday. From each variable, we calculated seven features (in total, 35 features).

Overall, we extracted 45 temporal features.

Table 5.2: Classification results without/with temporal features.

Random Forest Classifier	AUC	Accuracy (%)
without Temporal Features	0.982	95.92
with Temporal Features	0.986	96.56

5.4 Experiments

Next, we are interested in testing whether adding these temporal features to the existing static feature set improves performance of expert detection. To answer this research question, we conducted two experiments – evaluate (i) performance of classifiers based on the existing static features without temporal features; and (ii) performance of classifiers based on the existing static features with temporal features. We ran 10-fold cross-validation for each classifier. Table 5.2 shows experimental result of Random Forest classifier without temporal features. Random Forest classifier *without* temporal features achieved 0.982 AUC and 95.92% accuracy. Then, we added 45 temporal features to the existing feature set and developed another Random Forest classifier. As shown in Table 5.2, Random Forest classifier *with* temporal feature achieved 0.986 AUC and 96.56% accuracy. Based on these experiments, Random Forest classifier *with* temporal features outperformed Random Forest classifier *without* temporal features of 0.64

CHAPTER 6

DISCUSSION: ADDING FEATURES EXTRACTED FROM AN EXTERNAL SOURCE – TWITTER

Some people have accounts in multiple sites such as social Q&A sites like Quora and social media sites such as Twitter and Facebook. Sometimes people link their own accounts each other. We came up interesting research questions. How many people link their external accounts to their Quora profiles? Will collecting a user’s information in external sources like social media sites and extracting features from these external sources help improving performance of expert detection?

To answer these research questions, we analyzed what percent of Quora users in Table 2.1 linked their accounts in external sites to their Quora profile pages. We found that 60% of the users linked URLs of their Twitter and Facebook profile pages to their Quora profile pages. Then, we collected their Twitter account information such as user profile, recent 200 tweets, a list of following and a list of followers. From the Twitter data, we extracted profile features like the number of following, the number of followers and the number of posted tweets, linguistic features (based on LIWC) extracted from the recent 200 tweets, and the user’s Klout score to measure her influence on Twitter network [8]. We added Twitter features to the Quora feature set, and built Random Forest classifier.

Experimental result showed that adding Twitter features did not improve performance of expert detection. While Quora is a place where people share their knowledge with detailed information regarding specific questions, Twitter is a place where people share personal thoughts, breaking news, sentiments regarding products or politics, and etc in a brief format. Because of these reasons, we conjecture that adding features extracted from Twitter did not improve the performance of the expert classifier.

CHAPTER 7

RELATED WORK

Social Q&A sites have been used by many people for years. As people have shared their curiosities, questions, information and knowledge with others in the social Q&A sites, researchers have conducted research to solve various research problems in these sites. In this section, we summarize existing research work by three research problems.

First research problem is to understand what kind of social Q&A sites exists and what people do in these Q&A sites. Harper et al. [9] have categorized Q&A sites to three types – digital reference services, ask an expert service, community Q&A sites. They further analyzed how these three types of Q&A sites are different in terms of responsiveness about questions. Furtado et al. [10] have analyzed contributors’ activity on a stack exchange Q&A platform by clustering contributors’ profiles. They have measured the quality and quantity of contribution of these users. Wang et al. [11] suggested that user-topic model produced user interest in searching and answering questions on Quora.

Second research problem is to measure quality of an answer. Su et al. [12] examined the quality of answers in Q&A sites. Jeon et al. [13] built a model for detecting the quality of answer based on features derived from the particular answer being examined. Zhu et al. [14] examined the quality of answers in Q&A sites. Zhou et al. [15] proposed a joint learning method to measure quality of an answer. Toba et al. [16] examined a question type to select a right answer. Paul et al [17] studied how people evaluate quality of an answer.

Third research problem is to measure expertise of users or detect experts based on a question. Kao et al. [18] proposed a hybrid approach to effectively find expertise of users in different categories of the target question in Q&A sites. They used user’s reputation, subject relevance and their authority of a category in detecting experts. Bouguessa et al. [19] model the expertise of users based on the number of best answers in Yahoo Answers. Zhang et

al. [20] proposed Z-score measure to calculate the expertise level of users in Q&A sites. Some researchers measured expertise of users by analyzing their link structure using PageRank and HITS [21–23]. Other researchers studied how to detect experts based on a question in Q&A sites. Pal et al. [24] and Pal et al. [25] proposed a probabilistic model that captures the selection preferences of users based on the questions they choose. Liu et al. [26] proposed a hybrid approach to find experts for the category of a target question. Luo et al. [27] studied to recommend answerers in an enterprise social Q&A system.

Compared with the previous research work, our work is the first study to focus primarily on Quora for analyzing behaviors of experts and non-experts and detecting experts in general topics and a specific topic. We proposed user activity features, quality of answer features, linguistic features and temporal features, and developed statistical models to detect experts. This research will complement the existing research work.

To our knowledge, this is the first study to focus primarily on Quora for analyzing behaviors of experts and non-experts and detecting experts in general topics and a specific topic.

CHAPTER 8

CONCLUSION

In this paper, we have presented analysis of behaviors of experts and non-experts on Quora and identified four types of features such as user activity features, quality of answer features, linguistic features based on LIWC and temporal features. Then, we measured what features are top features in general topics and a specific topic. Based on this analysis and the observations, we proposed and developed statistical classification models to automatically identify experts. Our experimental results showed that these models effectively detected experts in general topics with 0.979 AUC and 95.94% accuracy, and in a specific topic with up to 0.987 AUC and 97% accuracy. We also studied whether adding temporal features would improve performance of expert detection. The experimental results showed that adding temporal features further improved performance of expert detection by additionally increasing 0.64% accuracy.

REFERENCES

- [1] J. Pennebaker *et al.*, *Linguistic Inquiry and Word Count*. Erlbaum Publishers, 2001.
- [2] E. Riloff and J. Wiebe, “Learning extraction patterns for subjective expressions,” in *In Proc. Conf. on Empirical Methods in Natural Language Processing*, 2003.
- [3] H. G. McLaughlin, “SMOG grading - a new readability formula,” *J. Reading*, pp. 639–646, May 1969.
- [4] L. A. Fast and D. C. Funder, “Personality as manifest in word use: correlations with self-report, acquaintance report, and behavior,” *J. of Personality and Social Psychology*, vol. 94, no. 2, pp. 334–346, Feb. 2008.
- [5] A. J. Gill *et al.*, “What are they blogging about? personality, topic and motivation in blogs.” in *in Proc. of the Third Int. Conf. on Weblogs and Social Media, ICWSM*, 2009.
- [6] Y. Yang and J. O. Pedersen, “A comparative study on feature selection in text categorization,” in *Proc. Fourteenth Int. Conf. on Machine Learning*, ser. ICML ’97, 1997.
- [7] M. Hall *et al.*, “The weka data mining software: an update,” *SIGKDD Explor. Newsl.*, vol. 11, no. 1, pp. 10–18, Nov. 2009.
- [8] Klout, “See how it works. - klout,” <http://klout.com/corp/how-it-works>, 2013.
- [9] F. M. Harper *et al.*, “Predictors of answer quality in online q&a sites,” in *In Proc. twenty-sixth annual SIGCHI conf. on Human factors in computing systems*, 2008.

- [10] A. Furtado *et al.*, “A case study of contributor behavior in Q&A site and tags: the importance of prominent profiles in community productivity,” *J. of the Brazilian Computer Society*, vol. 20, no. 1, p. 5, 2014.
- [11] G. Wang *et al.*, “Wisdom in the social crowd: An analysis of quora,” in *In Proc. 16th Int. Conf. on World Wide Web*, 2013.
- [12] Q. Su *et al.*, “Internet-scale collection of human-reviewed data,” in *In Proc. 16th Int. Conf. on World Wide Web*, 2007.
- [13] J. Jeon *et al.*, “A framework to predict the quality of answers with non-textual features,” in *Proc. 29th Annual Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*, 2006.
- [14] Z. Zhu *et al.*, “A multi-dimensional model for assessing the quality of answers in social q&a sites,” in *Int. Conf. on Information Quality*, 2009.
- [15] G. Zhou *et al.*, “Joint relevance and answer quality learning for question routing in community qa,” in *Proc. 21st ACM Int. Conf. on Information and knowledge management*, 2012.
- [16] H. Toba *et al.*, “Discovering high quality answers in community question answering archives using a hierarchy of classifiers,” *Inf. Sci.*, vol. 261, pp. 101–115, Mar. 2014.
- [17] S. A. Paul *et al.*, “Who is authoritative? understanding reputation mechanisms in quora,” in *In J. of Collective Intelligence*, 2012.
- [18] W.-C. Kao *et al.*, “Expert finding in question-answering websites: A novel hybrid approach,” in *Proc. 2010 ACM Symposium on Applied Computing*, ser. SAC ’10, 2010.
- [19] M. Bouguessa *et al.*, “Identifying authoritative actors in question-answering forums: the case of yahoo! answers,” in *Proc. 14th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, ser. KDD ’08, 2008.

- [20] J. Zhang *et al.*, “Expertise networks in online communities: Structure and algorithms,” in *In Proc. 16th Int. Conf. on World Wide Web*, 2007.
- [21] P. Jurczyk and E. Agichtein, “Discovering authorities in question answer communities by using link analysis,” in *Proc. Sixteenth ACM Conf. Information and Knowledge Management*, ser. CIKM ’07, 2007.
- [22] G. Zhou *et al.*, “An empirical study of topic-sensitive probabilistic model for expert finding in question answer communities,” *Knowledge-Based Systems*, vol. 66, no. 0, pp. 136 – 145, 2014.
- [23] D. Schall and F. Skopik, “An analysis of the structure and dynamics of large-scale q/a communities,” in *Proc. 15th Int. Conf. on Advances in databases and information systems (ADBIS)*, 2011.
- [24] A. Pal *et al.*, “Early detection of potential experts in question answering communities,” in *Proc. 19th Int. Conf. on User Modeling, Adaption, and Personalization*, ser. UMAP’11, 2011.
- [25] A. Pal *et al.*, “Exploring question selection bias to identify experts and potential experts in community question answering,” *ACM Trans. Inf. Syst.*, vol. 30, no. 2, pp. 10:1–10:28, May 2012.
- [26] D.-R. Liu *et al.*, “Integrating expert profile, reputation and link analysis for expert finding in question-answering websites,” *Inf. Process. Manage.*, vol. 49, no. 1, pp. 312–329, Jan. 2013.
- [27] L. Luo *et al.*, “Who have got answers?: Growing the pool of answerers in a smart enterprise social qa system,” in *In Proc. 19th Int. Conf. on Intelligent User Interfaces*, ser. IUI ’14, 2014.