

---

# Success Factors for Effective Knowledge Sharing in Community-based Question-Answering

---

Fabio Calefato, Filippo Lanubile, Maria Raffaella Merolla, Nicole Novielli

Dipartimento di Informatica  
Università degli Studi di Bari 'Aldo Moro'  
via E. Orabona, 4 – 70125 Bari  
{fabio.calefato,filippo.lanubile,nicole.novielli}@uniba.it  
m.merolla1@studenti.uniba.it

## Structured Abstract

**Purpose** – Nowadays, people increasingly seek information and ask for help on Question and Answer (Q&A) sites. The enormous success of Stack Exchange<sup>1</sup>, a constantly growing network of Q&A sites, attests this increasing trend. The success of Q&A mainly depends on the will of their members to provide good quality answers to others' questions. We investigate the success factors of Q&A that is those factors that foster effective knowledge creation and sharing. In particular, we focus those factors that can be acted upon by contributors when writing a question.

**Design/methodology/approach** – Based on literature in the domain, we build an empirical model of the factors that predict the chance of getting a useful answer when asking a question on a Q&A site. The actionable factors in our model are grouped into three categories of features: *Presentation Quality*, *Time*, and *Affect*. We use a multivariate logistic regression framework for estimating the probability of success of a question based on our set of predictors, that is the metrics that operationalize affect, time and presentation quality. Stack Exchange makes user-contributed content freely available<sup>2</sup> under Creative Commons license, which we use in our empirical studies.

**Originality/value** – Previous research shows how the success of a question depends on its presentation quality (Treude *et al.* 2011, Asaduzzaman *et al.* 2013), on the day and time in which it is posted (Bosu *et al.* 2013), and on the asker's reputation (Althoff *et al.* 2014). The influence of affective factors is less evident. However, we argue that the path to effective question answering also involves consideration of emotions (Novielli *et al.*, 2014). Our ongoing research aims at filling this gap in literature by further investigating the role of affect in Stack Exchange.

**Practical implications** – The expected output of this ongoing research will be a user-driven netiquette for online Q&A sites. It will shed new light on how emotion expression facilitates or impairs effective knowledge sharing, leading to guidelines for fostering emotional awareness computer-mediated interactions. In details, we aim at defining new

---

<sup>1</sup> <http://stackexchange.com/>

<sup>2</sup> <https://archive.org/details/stackexchange>

empirically driven guidelines in order to inform the design of emotional interface, provide new insights on how to successfully exploit social Q&A sites in professional knowledge sharing, develop tools for embedding emotional intelligence into online Q&A communities to support both novice users and community managers.

**Keywords** – Community-based Question&Answering, Collaborative Software Development, Affective Computing.

**Paper type** – Academic Research Paper

## 1. Introduction

Today, people increasingly seek information and ask for help on Question and Answer (Q&A) sites. The enormous success of Stack Exchange, a growing network Q&A sites, attests this increasing trend. Stack Exchange originates from the success of Stack Overflow, a community of millions of programmers discussing issues about software developments launched in 2008. The Stack Exchange network now includes 139 Q&A sites<sup>1</sup>, each focusing on one specific topic, ranging from technology, finance to cooking, linguistics, and religion.

The emergent success of community-based Q&A largely depends on the will of their members to answer others' questions. Among driving factors of success of requests, there are both social aspects (Althoff *et al.* 2014), depending on people's status in the community, and technical aspects (Treude *et al.* 2011), dependent upon topics. Only recently, research has begun to investigate linguistic factors too, looking at how questions are formulated (Althoff *et al.* 2014; Mitra and Gilbert 2014).

In this paper, we investigate how an information seeker can increase the chance of getting a useful answer when asking a question on Stack Exchange. In particular, we focus on actionable factors that can be acted upon by users when writing a question. Accordingly, we exclude from our set of predictor variables factors, such as the topic of the question and the asker's reputation, which are not actionable despite having been found to influence both the probability of success of a request (Althoff *et al.* 2014, Bosu *et al.* 2013) and its quality (Ponzanelli *et al.* 2014). Hence, our research question is formulated as follows:

*RQ* – Which actionable factors in a question do affect its success in a Q&A site?

The actionable factors in our model are broadly grouped into three categories of features: *presentation quality*, *time*, and *affect*. Previous research clearly shows how the success of a question depends on its presentation quality (Treude *et al.* 2011, Asaduzzaman *et al.* 2013) as well as on the day and time in which it is posted (Bosu *et al.* 2013). The influence of affective factors is less evident. Although displaying emotions is common in face-to-face interaction, people might not be prepared for effectively dealing with the barriers of social media to non-verbal communication, as also discussed by Stack Exchange users who complain about harsh comments from expert community members<sup>2</sup>. Moreover, expression of gratitude, urgency and reciprocity has been shown to be a factor

---

<sup>1</sup> As of April 2015.

<sup>2</sup> <http://meta.stackexchange.com/questions/179003/stack-exchange-is-too-harsh-to-new-users-please-help-them-improve-low-quality-po/179009#179009>

of success for altruistic requests on online social communities (Althoff *et al.* 2014). Accordingly, we also consider textual cues for affective states among the potential factors of success for a question.

This paper represents a first step towards the long-term goal of understanding the role of emotions in Stack Exchange (Novielli *et al.* 2014). While previous research has mostly focused on Stack Overflow, the Stack Exchange network now comprises many other sites belonging to categories other than technology. In our study, we consider data from four Stack Exchange websites, from different categories, in the attempt of drawing conclusions that are not specific to programmers' behaviors and are more generalizable to community-based Q&A overall.

The remainder of this paper is organized as follows. In Section 2, we describe Stack Exchange and the dataset considered in this study. In Section 3, all the actionable factors and related metrics are presented. Preliminary results are reported in Section 4 and discussed in Section 5. Finally, in Section 6 we draw conclusions, discuss the main limitation of this study, and anticipate the next steps for future work.

## 2. Stack Exchange

In this section, we shortly describe how Stack Exchange works and then, we provide details about the dataset used in our study.

### 2.1. Background

The Stack Exchange network comprises today over a hundred of different Q&A websites organized by categories, namely *Technology*, *Culture/Recreation*, *Life/Arts*, *Science*, *Business* and *Professional*. All these sites were born after the huge success of Stack Overflow, the original site dedicated to computer programming questions. Each site in Stack Exchange consists of pages containing one question posted by a user and an arbitrary number of answers submitted by other users. Questions can have one accepted answer, chosen by the original asker, if that solves the problem. We consider as 'successful' those questions for which an accepted answer has been provided.

Questions, answers and users are subject to a reputation award process. Questions and answers can receive upvotes and downvotes from users who want to either underline the usefulness of content or express disagreement. In the true spirit of *gamification*, through their contributions, users earn extra rights, reputation points and badges, which reflect users' skills as well as their status in the community. In turn, this reputation system motivates users to generate high quality content.

Stack Exchange sites are self-moderated by community members that unlocked moderation rights by earning sufficient reputation points. To keep quality high, moderators can remove questions or answers because inappropriate or irrelevant.

## 2.2. The Dataset

Stack Exchange makes user-contributed content freely available<sup>1</sup> under Creative Commons license, with the exception of users' personal data. For this study, we downloaded the dump of May 2014.

We decided to include in our study the site with the highest number of questions per day for four Stack Exchange categories, as reported in Table 1. We excluded questions that were removed or closed by moderators, as well as those questions that were edited after the original post. Since modifying a question may have indirect side effects on the quality evaluation, previous research has adopted the same approach as ours (Ponzanelli *et al.* 2014). The final dataset resulting from this pre-processing phase ended up containing almost 300k questions.

Table 1. Overview of the dataset after pre-processing.

Category	Website	Active users	Total # of questions	# questions per day	# questions after pre-processing	% successful	Creation Date
Culture / Recreation	Arqade (videogaming)	50,894	42,320	38.0	17,303	76%	2009
Business	Bitcoin	12,832	6,894	5.7	3,659	56%	2011
Science	Mathematics	121,666	269,819	533.0	272,697	55%	2010
Life/Arts	Science Fiction & Fantasy	20,477	13,495	22.0	4,768	72%	2011
Total Questions Included in the Study					298,427		

## 3. Success Factors for Questions

In this section, we present the predictor variables, grouped as factors, which compose our framework of analysis.

### 3.1. Affect

Affective computing techniques (Picard, 2000) have already been successfully applied for analyzing written interactions on social media (Thelwall, 2012). We include affective lexicon as a source of information to build our framework. Specifically, we consider metrics describing the overall *sentiment* of a post and the use of simple linguistic cues to express *gratitude*.

**Sentiment.** By sentiment analysis, we refer to the task of detecting the overall positive/negative polarity of a text (Pang and Lee 2008). Kucuktunc *et al.* (2012) perform a large-scale sentiment analysis study on Yahoo! Answers. Their work shows how best answers tend to a neutral sentiment. Consistently with the study by Kucuktunc *et al.*, we

<sup>1</sup> <https://archive.org/details/stackexchange>

capture the sentiment of the questions in our dataset using SentiStrength<sup>1</sup>, a state-of-the-art tool (Thelwall *et al.* 2010) already employed in social computing (Guzman *et al.* 2013). SentiStrength is capable of dealing with informal language, including abbreviations, intensifiers and emoticons. Based on the assumption that a text, even if short, can convey mixed sentiment, SentiStrength outputs both positive and negative sentiment scores for all the input sentences. Positive sentiment scores range from +1 (neutral) to +5 (extremely positive). Analogously, negative sentiment scores range from -1 (neutral) to -5 (extremely negative). We mapped both scores into the [0,4] interval.

**Gratitude.** Expressing gratitude in a question has been found to positively correlate with success of requests (Althoff *et al.* 2014) because it can be seen as a clue of positive disposition towards the future answerer. Therefore, we decide to include the presence of lexical cues of gratitude in the set of our affective predictors. In particular, we are interested in verifying whether paying gratitude forward makes a request more likely to be successful. We model gratitude as a binary feature indicating the presence of lexical cues of appreciation (i.e., ‘*thanks*’, ‘*thank you*’).

### 3.2. Time

Previous research has demonstrated that time is a critical factor for the success of a question in technical Q&A (Bosu *et al.* 2013). In fact, high and low efficiency hours exist in a community, as well as off-peak hours. Therefore, we include as predictors both the question posting **Time** (GMT), assuming values in {*Morning*, *Afternoon*, *Evening*, *Night*}, and the **Day** of the week, coded with values in {*weekend*, *weekday*}.

### 3.3. Presentation Quality

Writing an easy-to-read and -understand question by providing contextual information is regarded as the strongest indicator of the quality of a question and, hence, a good predictor of its success. Unclear and vague questions have been shown to remain unanswered in technical Q&A (Asaduzzaman *et al.* 2013). Moreover, Stack Overflow uses a set of simple textual metrics to feed the queue of low quality posts that need to be analyzed by a selected subset of users who unlocked moderation privileges holding sufficient reputation score (Ponzanelli *et al.* 2014), including length, uppercase ratio, URLs count, and textual similarity between title and body. Therefore, we define the following metrics.

**Title and Body Length**, which we calculate as number of words. Too short questions tend to remain unanswered, failing to clarify their meaning to community members (Asaduzzaman *et al.* 2013). On the other hand, overly long questions might be associated to a greater effort in providing an answer and might discourage potential helpers.

**Number of URLs**, which we treat as a numeric feature indicating how many links are in the question.

**Number of Tags** associated to a question. When properly employed, tags are useful to attract the potential answerers based on the topic of interest, since Stack Exchange allows users to search and browse questions based on the user-specified tags.

---

<sup>1</sup> <http://www.sentistrength.wlv.ac.uk>

**Uppercase Ratio**, calculated as the ratio between the total uppercase characters over the total characters in a question. Use of uppercase is employed as a textual cue in pre-filtering of low quality posts (Ponzanelli *et al.*, 2014).

**Code snippet**. A yes/no feature indicating whether a question embeds an excerpt of code, which the asker has provided as an example. Code snippets are considered key for success in ‘code review’ questions (Treude *et al.* 2011). We use this variable for all the sites in our dataset except for ‘Science Fiction’, to which it does not apply.

**Number of Images**. A numeric feature indicating how many images are provided in the question.

**Topic Alignment**. Coded as a binary variable, it indicates whether the title and the body refer to the same topic. When accessing the web interface of Stack Exchange sites, people visualize a list of snippets including the title and the incipit of questions. Hence, it is crucial to define a title that clearly specifies the problem being addressed in the post. To verify the title-body alignment, we annotate our dataset using topic labels for both bodies and titles of questions. In order to do so, we exploit topic modeling and, in particular, Latent Dirichlet Association (LDA) (Blei *et al.* 2003). LDA builds a model of topics in a corpus of documents in an unsupervised fashion, based on word frequencies and co-occurrences. The number of topics is a customizable parameter, according to the desired granularity of the topic taxonomy. For each site, we experimented with a number of topics varying from 10 to 40, finally choosing 15 as the optimal number for a meaningful granularity. We extracted topics separately for each site in our dataset, using Mallet (McCallum 2002), which provides an implementation of LDA. We pre-process questions to discard code snippets, remove HTML tags and apply the Porter stemming algorithm (Porter, 1997) to reduce data sparsity.

#### 4. Preliminary Results

We applied a logistic regression framework for estimating the probability of success of a question based on our set of predictors, that is, the measures grouped by affect, time, and presentation quality, as outlined in Section 3.

For each predictor, the logistic regression outputs three values, namely estimate, odds ratio and statistical significance. The sign of the estimate coefficient indicates the (positive/negative) impact of the predictor on the success of a question. The odds ratio weighs the magnitude of this impact. Finally, the statistical significance, which is based on a likelihood ratio test traditionally employed in logistic regression frameworks (Peng *et al.* 2002), determines whether a predictor has a significant explanatory value.

In Table 2, we report the results of the logistic regression for all predictors and sites in our dataset. As for affective predictors, we observe an overall significant impact on the probability of success for all sites. In particular, a negative sentiment significantly reduces the probability of success for Science Fiction. On the contrary, it has a positive impact on the success of questions posted to Mathematics and Bitcoin sites. Similarly, a positive sentiment correlates to unsuccessful questions in Arqade and Bitcoin. Finally, gratitude expressions negatively correlate to success in all sites except for Mathematics.

As for temporal metrics, both time and day of the week significantly correlate with success in Mathematics. For this site, we observe a higher probability of success during the weekend. Also, the majority of successful requests for Mathematics are posted during evening and nights (GMT).

As for metrics related to presentation quality, clearness of the title is positively correlated with successful questions for Mathematics unlike references to additional resources (i.e., URLs and images). Except for Mathematics, the uppercase ratio appears as the most predictive metric (larger coefficient estimates). It is generally associated to unsuccessful requests but we observe a very high correlation with success in the Science Fiction community. As for body length, it negatively impacts successful requests also in Arqade and Bitcoin. On the contrary, including images here significantly increases the probability of success (second larger parameter estimate). Finally, we observe a negative impact for code snippet in the Bitcoin community where, conversely, including URLs greatly increases the chance of success.

Table 2. Results of logistic regression for Affect, Time, and Presentation Quality

	Arqade			Science Fiction			Mathematics			Bitcoin		
	Estim.	OR	*	Estim.	OR	*	Estim.	OR	*	Estim.	OR	*
(Intercept)	1.715	5.55	*	1.153	3.17	*	0.010	1.01		0.034	1.03	
<b>Affect</b>												
Positive Sent.	-0.101	0.90	*	-0.044	0.96		0.007	1.01		-0.098	0.91	*
Negative Sent.	0.020	1.02		-0.068	0.93	*	0.026	1.03	*	0.140	1.15	*
Gratitude	-0.477	0.62	*	-0.702	0.50	*	0.042	1.04	*	-0.368	0.69	*
<b>Time</b>												
Weekend	-0.169	0.84	*	-0.158	0.85	*	0.079	1.08	*	0.067	1.07	
GMT h.: Night	-0.032	0.97		0.034	1.03		0.032	1.03	*	0.048	1.05	
Afternoon	0.019	1.02		-0.008	0.99		-0.001	1.00		0.110	1.12	
Evening	-0.009	0.99		-0.020	0.98		0.066	1.07	*	0.006	1.01	
<b>Presentation Quality</b>												
Code Snippet	-0.156	0.86		--	--		0.063	1.07		-0.316	0.73	*
Body Length	-0.002	1.00	*	0.000	1.00		0.000	1.00		-0.005	1.00	*
Title Length	-0.011	0.99	*	-0.033	0.97	*	0.017	1.02	*	-0.011	0.99	
Num. of Images	0.200	1.22	*	0.140	1.15		-0.051	0.95	*	0.092	1.10	
Upp. Chars Ratio	-2.403	0.09	*	5.273	194.94	*	0.617	1.85		-4.067	0.02	*
Num. of Tags	-0.140	0.87	*	0.041	1.04		-0.014	0.99	*	0.084	1.09	*
URL	0.038	1.04		0.011	1.01		-0.044	0.96	*	0.229	1.26	*
Topic Alignment	0.069	1.07		-0.010	0.99		0.039	1.04	*	-0.127	0.88	
Statistical significance: * $p < 0.05$												

## 5. Discussion

Overall, the results of the logistic regression show that success factors vary depending on the community being analysed. This might depend on the domain of interest of each site

or depend on the cultural background of the community, causing the adoption of a different netiquette for each website. In the following, we analyse the impact of each factor on the success of a question.

**Affect.** Previous research has provided divergent evidence on the role of sentiment in the success of requests (Mitra and Gilbert, 2014; Kucuktunc *et al.* 2012; Althoff *et al.* 2014), thus suggesting that different communities have different perception of emotionally loaded posts. Our findings reflect this diversity, highlighting the needs for further investigation in this direction.

More in detail, previous research on the success of questions in Q&A sites reports about the impact of sentiment on both the success of a post and its perceived quality. The large-scale sentiment analysis on *Yahoo! Answers* by Kucuktunc *et al.* (2012) shows how best answers tend to a neutral sentiment. Similarly, Bazelli *et al.* (2013) discover that developers do not appreciate questions showing more extroversion. Our results confirm these findings for Science Fiction, for which negative sentiment is associated to unsuccessful questions. Similarly, Arqade and Bitcoin users did not appreciate questions with positive emotional load. Unlike previous evidence provided by Althoff *et al.* (2014), gratitude seems to have a negative impact on the probability of success. The only exception is represented by the users of Mathematics, which seems to appreciate the exhibition of gratitude in a request for help.

**Time.** As predicted by literature (Bosu *et al.* 2013), the day and time of a question has a large impact on the probability of all sites, except for Bitcoin. We speculate this might be due to the experts mainly reside in the US, as proposed by previous research on Stack Overflow (Bosu *et al.* 2013), hence they are available in the time slices corresponding to afternoon and evening in GMT zone. However, information about the user nationality and residence are not available and therefore we cannot verify this hypothesis.

**Presentation Quality.** Metrics describing the comprehensibility of a question seem to significantly contribute in predicting its success. This confirms previous findings that clear questions make people more willing to help (Asaduzzaman *et al.* 2013). Moreover, Althoff *et al.* (2014) found that clearly explaining the needs, through the narrative and pictures, has an impact on the success as well as linguistic indications of gratitude, generalized reciprocity, as well as the status of an asker in the community. We found that the presence of code snippets is the second most effective predictor for Bitcoin. While being consistent with literature (Treude *et al.* 2011), the effectiveness of code snippets cannot be generalized beyond those domains for which this kind of information makes sense.

Explicative titles whose content is aligned with the content of the question are significantly correlated with success in Mathematics. Similarly, shorter titles are significantly correlated with success for Arqade and Science Fiction. However, the contribution of this feature is of minor importance (e.g., see the odds ratio close to 1). Moreover, we observe that the higher number of tags the lower the probability of obtaining useful answers (as observed in Arqade and Mathematics). All these indicators suggest a potential asker to privilege conciseness in title definition and specificity in question tagging, in order to have more chances to attract potential contributors and making them more willing to help.

Upper character ratio is the most simple and predictive feature and it is mainly associated with unsuccessful questions (i.e., in Arqade and in Bitcoin). The massive use of uppercase in generally perceived as an indicator of low quality of a post. In fact, Stack Exchange includes this metric in the set of indicators used for feeding the review queue of low quality posts (Ponzanelli *et al.* 2014). Surprisingly, uppercase ratio is strongly



associated to the probability of success in Science Fiction. In searching for an answer we discovered that it is common for users of this site to discuss about actors and TV (e.g., “[...] *And I mean the Joseph Curwen from H. P. Lovecraft's 'The Case of Charles Dexter Ward' [...]*”). Therefore, while literature has provided evidence that simple textual and readability metrics can be effectively exploited for pre-filtering of poor quality of posts (Ponzanelli et al. 2014), the results of this study cast doubt on their generalizability and suggest further study to assess their validity across the communities in Stack Exchange.

## 6. Conclusions

In this paper, we have reported the results of an empirical study involving four Stack Exchange websites, in which we investigate the impact of affect, time, and presentation quality on the success of questions. In particular, our study provides evidence that may serve as a hint for users aiming at increasing the chance of getting help from the studied communities.

We acknowledge several limitations of our study. To address the problem of external validity (i.e., the ability to generalize the results outside the scope of the study), we included in our dataset Q&A sites from different categories of Stack Exchange. Sites were selected based on the size and level of activity of the community. However, Stack Exchange counts more than a hundred of very diverse sites and the selected dataset might not be representative of the whole network. Furthermore, in this study we choose to exclude factors that have been proven to effectively predict the success of requests such as the topic and the reputation of the askers. While consistent with our goal of investigating only actionable factors, this choice is a limitation with respect to goal of building a general model of success predictors in Q&A sites.

In our future work, we plan to both address the limitations of the present study and investigate new research questions suggested by our preliminary results. In particular, to address the issue of the generalizability of our findings, we plan to replicate our study by considering a larger number of Stack Exchange sites. In addition, we intend to further investigate the role of affect in Stack Exchange by considering more fine-grained emotion analysis, since different emotions might be relevant for different domains.

## Acknowledgements

The computational work has been executed on the IT resources made available by two projects financed by the MIUR (Italian Ministry for Education, University and Re-search) in the "PON Ricerca e Competitività 2007-2013" Program: ReCaS (Azione I - Interventi di rafforzamento strutturale, PONa3\_00052, Avviso 254/Ric) and PRISMA (Asse II - Sostegno all'innovazione, PON04a2\_A).

## References

Althoff, T., Danescu-Niculescu-Mizil, C., and Jurafsky, D. 2014. How to Ask for a Favor: A Case Study on the Success of Altruistic Requests. *In Proc. of the 8th International AAAI Conf. on Weblogs and Social Media (ICWSM 2014)*.

- Asaduzzaman, M.; Mashiyat, A.S.; Roy, C.K.; Schneider, K.A. 2013. Answering questions about unanswered questions of Stack Overflow, *In Proc. of the 10th IEEE Working Conf. on Mining Software Repositories (MSR 2013)*, 97,100.
- Blei, D., Ng, A., and Jordan, M. 2003. *Latent dirichlet allocation*. The Journal of Machine Learning Research 3, 993-1022.
- Barua, A., Thomas, S. W., and Hassan, A. E. 2012. What are developers talking about? An analysis of topics and trends in Stack Overflow. *Empirical Software Engineering*, June 2014, Volume 19, Issue 3, pp 619-654.
- Bazelli, B., Hindle, A., and Stroulia, E. 2013. On the Personality Traits of StackOverflow Users. *In Proceedings of the 2013 IEEE International Conference on Software Maintenance (ICSM '13)*, IEEE Computer Society, Washington, DC, USA, 460-463.
- Bosu, A., Corley, C.S., Heaton, D., Chatterji, D., Carver, J.C., and Kraft, N.A. 2013. Building Reputation in StackOverflow: An Empirical Investigation. 2013 *In Proc. of the 10th IEEE Working Conf. on Mining Software Repositories (MSR 2013)*, 89, 92.
- Guzman, E., Azócar, D., and Li, Y. 2014. Sentiment analysis of commit comments in GitHub: an empirical study. *In Proc. of the 11th Working Conf. on Mining Software Repositories (MSR 2014)*. ACM, New York, NY, USA, 352-355.
- Kucuktunc, O., Cambazoglu, B.B., Weber, I., and Ferhatosmanoglu, H. 2012. A large-scale sentiment analysis for Yahoo! answers. *In Proc. of the fifth ACM international conf. on Web search and data mining (WSDM '12)*. ACM, New York, NY, USA, 633-642.
- McCallum, A. K. 2002. MALLET: A Machine Learning for Language Toolkit. <http://mallet.cs.umass.edu>. 2002. Accessed October 2014.
- Mitra, T. and Gilbert, E. 2014. The language that gets people to give: phrases that predict success on Kickstarter. *In Proc. of the 17th ACM conf. on Computer supported cooperative work & social computing (CSCW '14)*. ACM, New York, NY, USA, 49-61.
- Novielli, N., Calefato, F., and Lanubile, F. 2014 Towards Discovering the Role of Emotions in Stack Overflow. *In Proc. 6th Int'l Workshop on Social Software Engineering (SSE'14)*, Nov. 16, 2014, Hong Kong, China, pp. 33-36, DOI:10.1145/2661685.2661689
- Pang, B. and Lee, L. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1-2), pp. 1-135, 2008
- Peng, C.J., Lee, K.L., and Ingersoll, G.M. 2002. An Introduction to Logistic Regression Analysis and Reporting. *The Journal of Educational Research*, 96(1), 3-14.
- Picard, R. W. 2000. *Affective Computing*. MIT Press.
- Ponzanelli, L., Mocci, A., Bacchelli, A., and Lanza, M. 2014. Improving Low Quality Stack Overflow Post Detection. *In 30th International Conference on Software Maintenance and Evolution (ICSME'14) Industry Track*, pp. 541-544. 2014
- Porter, M.F. 1997. An algorithm for suffix stripping. *In Readings in information retrieval*. Morgan Kaufmann, San Francisco, 313-316
- Thelwall, M., Buckley, K., and Paltoglou, G. 2012. Sentiment Strength Detection for the Social Web. *Journal of the American Society for Information Science and Technology*, 63(1):163-173.
- Thelwall, M., Buckley, K., Paltoglou, G., Cai, D., and Kappas, A. 2010. Sentiment strength detection in short informal texts. *Journal of the American Society for Information Science and Technology*, 61(12), 2544-2558.
- Treude, C., Barzilay, O., and Storey, M. 2011. How do programmers ask and answer questions on the web? *In Proc. of the 33rd International Conf. on Software Engineering (ICSE '11)*. ACM, New York, NY, USA, 804-807.

## Authors' Biographical Notes

**Fabio Calefato** received the MSc and PhD degrees in computer science from the University of Bari, Italy. Currently, he is as a postdoctoral research assistant at the Collaborative Development Group (COLLAB) at the University of Bari. His research interests are focused on studying computer-mediated communication, collaboration in global software engineering, social software engineering and community-driven knowledge sharing.

**Filippo Lanubile** is an Associate Professor of computer science at the University of Bari where he leads the COLLAB Group. He has the National Scientific Qualification for the Full Professor position. His research interests lie in the areas of software engineering and social computing. He is a recipient of an IBM Eclipse Innovation Award (2006), an IBM Faculty Award (2008) and the Microsoft Research Software Engineering Innovation Foundation Award (2011). Filippo Lanubile has been the Program Chair and General Chair of the IEEE Int. Conf. on Global Software Engineering (ICGSE), respectively on 2008 and 2013. He is a member of the Editorial Board of the Empirical Software Engineering (Springer) journal..

**Nicole Novielli** is a postdoctoral research assistant at COLLAB. Her research interests are in the areas of Human-Computer Interaction and Affective Computing, with a focus on the role of human factors and emotions in natural language dialogues and in computer-supported collaboration. She obtained a PhD in Computer Science from the University of Bari in 2010, discussing a thesis on 'Lexical Semantics of Dialogue Act'.

**Maria Raffaella Merolla** is a MSc student. She is developing her master thesis project at COLLAB.

## Acknowledgements

This work fulfills the research objectives of the PON 02\_00563\_3470993 project "VINCENTE - A Virtual collective INtelligenCe ENvironment to develop sustainable Technology Entrepreneurship ecosystems" funded by the Italian Ministry of University and Research (MIUR).