Can Expertise Impair the Quality of Crowdsourced Data?

Research-in-Progress

Shawn Ogunseye
Memorial University of Newfoundland
osogunseye@mun.ca

Jeffrey Parsons
Memorial University of Newfoundland
jeffreyp@mun.ca

Abstract

It is not uncommon for projects that collect crowdsourced data to be commissioned with incomplete knowledge of data contributors, data consumers, and/or the purposes for which the data collected are going to be used. Such unanticipated uses and users of data form the basis for open information environments (OIEs), and the information collected through systems designed to gather content from users have high quality when they are complete, accurate, current and provided in an appropriate format. However, as it is assumed that experts provide higher quality information, many types of OIEs have been designed for experts. In this paper, we question the appropriateness of this assumption in the context of citizen science systems – an exemplary category of OIE. We begin by arguing that experts are primarily efficient rule-based classifiers, which implies that they selectively focus only on attributes relevant to their classification task and ignore others. Drawing from existing literature, we posit that experts' focus on only diagnostic features of an entity leads to a learned inattention to non-diagnostic attributes. This may improve the accuracy of the information provided, but at the expense of its completeness, currency, format and ultimately the novelty (for unanticipated uses) of information provided. On the other hand, we predict that non-experts and amateurs may use rules to a lesser extent, resulting in less selective attention and leading them to provide more novel information with less trade-off of one dimension of information quality for another. We propose hypotheses derived from this view, and outline two experiments we have designed to test them across four dimensions of information quality. We conclude by discussing the potential implications of this work for the design of crowdsourcing platforms and the recruitment of experts, amateurs, or novice data contributors in studies of data quality in crowdsourcing settings.

Introduction

Crowdsourcing systems are an increasingly popular tool to collect data outside of traditional organizational settings (Castriotta & Guardo, 2011; Hosseini, Phalp, Taylor, & Ali, 2014; Tarrell et al., 2013; Tripathi & Tahmasbi, 2014). A popular advantage of crowdsourcing is the possibility to broaden the base of potential data contributors. At the same time, some crowdsourcing applications (e.g., citizen science) demonstrate a preference for expert contributors to ensure that information collected is of high quality (Budescu & Chen, 2014; Lukyanenko, Parsons, & Wiersma, 2016). This preference reflects the underlying assumption of a positive correlation between a contributor's level of expertise in a given domain and the quality of information the contributor provides. In other words, there is at least a tacit belief in many crowdsourcing applications that *the expert crowd is the better crowd*.

We argue that, to fully appreciate the ramifications of approaches to the design of distributed knowledge crowdsourcing systems, consideration must be given to not just the intended positive consequences of strategies implemented to elicit information from experts, but also the potential unintended negative consequences. More generally, this research is motivated by two factors.

First, there have been surprising findings in a number of studies involving distributed knowledge projects of no substantial difference between the task performance of experts and non-experts. For instance, examining three groups of participants (assigned into groups by their levels of expertise) in the detection of malicious data in a collection of "good" data, researchers found that, contrary to their expectations and the projections from prior research, participants' level of expertise had no significant influence on their performance (Biros, George, & Zmud, 2002). In another case, by re-conceptualizing and redesigning a crowdsourcing system to enable non-expert information contributors to provide content in ways familiar to them – via attributes of instances rather than classes based on inclusion rules – the knowledge requirement for making contributions was reduced and the quality of information contributed by non-experts improved (Lukyanenko, Parsons, & Wiersma, 2014; Parsons, Lukyanenko, & Wiersma, 2011). An even more recent study (Austen, Bindemann, Griffiths and Roberts, 2016) reported a citizen science experiment where experts and non-experts were tested on their capacity to report the sighting of a species of bumblebee. The study showed that, not only was overall identification accuracy low (at most about 57%), there was no significant difference in the performance of experts and non-experts. Consequently, the authors call for the investigation of "how experts and non-experts make observations" to be factored into study [IS] design (p. 1).

Second, there is an increasing interest in open information environments (OIEs) – informal communities in which "data are not necessarily created or used within, by, or for, a formal organization" (Parsons & Wand, 2014, p. 1). An OIE "creates opportunities to generate new information and use it in unexpected ways" (Parsons & Wand, 2014, p. 1). To achieve this, researchers and practitioners need to understand the extent to which expertise improves (or decreases) the IQ of crowdsourced data.

In this work, we present cognitive principles based on classification theory and identify their consequences for understanding the impact of expertise on performance in crowdsourcing tasks. Subsequently, we describe a proposed laboratory experiment to investigate the impact of expertise on IQ in crowdsourced data. Finally, we discuss the potential implications of our experiment on research in the design of OIEs.

Theoretical Background

Research in cognitive psychology provides the theoretical foundation for this study. The literature on classification explains why we selectively attend to the attributes of objects and offers a suitable lens for exploring expertise and information quality.

Classification Theory

Classification (or categorization) is a fundamental human capability. We classify to make efficient use of our cognitive resources by organizing our existing knowledge about phenomena mainly through their similarities, which allows us to make predictions about new instances and events (Best, Yim, & Sloutsky, 2013; Parsons & Wand, 2008, 2013). Classes are therefore useful abstractions of the similarities of the classified phenomena. Classification theory and its relevance to information system (IS) analysis and design have been extensively discussed (e.g., Parsons, 1996; Wand, Monarchi, Parsons, & Woo, 1995). Classes are considered from the perspective of either the inclusion rules that form them (rule-based classification) or their instances (exemplars and prototypes) (Kloos & Sloutsky, 2008; Murphy, 2002).

In rule-based classification, it is necessary to pay selective attention to only relevant features crucial for identifying instances of the class, while irrelevant features (those not useful for predicting class membership) can be safely ignored. Although selective attention leads to efficient learning, especially when making connections between instances with very sparse similarities and dense dissimilarities, it has costs. The primary cost of selective attention is a learned inattention to features that are not "diagnostic" in the present context (Hoffman & Rehder, 2010). If these features, however, become diagnostic in another context, then the ability to make predictions and transfer knowledge is lost.

Even though rule-based classification has the cost of learned inattention to the non-diagnostic features of instances, it is dominant in IS and informs the design of interfaces, databases and programs (Lukyanenko, Parsons, & Samuel, 2015). In OIEs (e.g., citizen science projects), where unanticipated applications, uses and users of data may be paramount, instance-based designs are argued to be better because they encourage contributions that include non-diagnostic features (Parsons & Wand, 2014). Additionally, Lukyanenko et al. (2014) argued that designs based on instances rather than classes are more accommodating of less knowledgeable contributors.

Instance-based perception and reasoning about entities in our world is "a developmental default" (Gelman & Coley, 1990; Gelman & Markman, 1986, 1987; Kloos & Sloutsky, 2008). Rule-based classification is learned with development and distinguishes adults from children. Experiments conducted by Best, Yim, & Sloutsky (2013), comparing the ability of infants and adults to form inclusion rules and selectively attend to attributes of instances based on such rules, show that infants do not have the capacity for selective attention. Infants reason about classes by observing all the features of individual instances without any *a priori* class inclusion rules (Gelman, Collman, & Maccoby, 1986). We contend they are naturally comparable to non-experts in the domain of a distributed knowledge crowdsourcing project. Like infants, non-experts also lack *a priori* class-formation rules. Infants can, therefore, help us understand how non-experts and "expert amateurs" – people with incomplete knowledge – perceive instances (Keil, 2011; Kloos & Sloutsky, 2008).

We posit that a non-expert's exposure to an instance in a citizen science project is an example of a situation that activates the default ("infant") kind of reasoning about classes. Conversely, the tendency of adults to employ rule-based classification can help us understand expert contributors. Rule-based classification allows experts to focus on relevant features for identifying instances of classes, producing cognitive economy (efficiency of classification). Thus, we expect them to be less likely to attend to non-diagnostic attributes than will novices. These atypical data (irrelevant to the classification task) reported by non-experts leads to a richer dataset, emphasizing the "data completeness" component of IQ. It is valuable to OIEs, especially as interest in big data continues to grow, because it can lead to novel discoveries and unanticipated uses.

Research Model and Hypotheses

Based on the theoretical foundation above, we propose hypotheses and experiments that explore the effect of expertise on the completeness and overall quality of the information contributed by participants in a citizen science project.

We define three types of information that contributors may provide:

> Type A information – attributes of instances that match a classification rule;
> Type B information – attributes that are ascribed to an observed instance but are not part of a classification rule; and
> Type C information – attributes that are not inherent to the instance, but are a part of its surroundings and may directly or indirectly interact with the instance.

Hoffman and Rehder (2010) showed the need to differentiate supervised classification – engendered by some form of explicit training (e.g., by a teacher) with sufficient feedback to improve the classifier's skill – from unsupervised classification – classification formed without explicit training (self-taught). They argued that the latter "may involve less rule-based processing" and consequently, more attentiveness to attributes. As a result, their definition of expert classifiers excludes "amateur experts" or people who are deeply interested in a domain. Such amateur classifiers may or may not have accurate knowledge. They also explained that "expert classification involves the same sort of attention optimization that characterizes supervised classification learning" (p. 336), which is due to extensive training and the type of task. Studies that have considered amateurs when studying domain knowledge are scarce both in psychology and business domains.

We predict that the level of domain knowledge possessed by a contributor will affect his/her tendency to communicate at the class level. Domain knowledge therefore affects the format[1] of information contributed.

> H1: Experts will contribute more information at a higher level of abstraction than will Amateurs; Amateurs in turn will contribute more information at a higher level of abstraction than Novices.

Through a series of experiments comparing adults who learned categories using either inclusion rules or instances, Hoffman and Rehder's study emphasized that the use of classification is accompanied by the

[1] Format has been described by Nelson, Todd & Wixom (2005) as the way information is presented for understandability. We adapt this definition in this user generated content context to mean the way in which contributors present information to the consumer based on their understanding of the domain.

cost of limited knowledge (resulting from focusing narrowly) and of learned inattention. Their experiments show that when an individual learns to ignore certain attributes of an instance, he/she continues to ignore these attributes when they are present in novel instances (associated or not to previous instances), even if the ignored attributes have become critical to good performance. This is consistent with findings from Lukyanenko, Parsons, & Wiersma, (2014) study's, as the use of an instance-based design increased the use of their system and also the number of novel sightings reported by novices. Other studies have also reported that users of rule-based classification perform poorly at tasks that require identifying a novel instance of a class (Best et al., 2013; Kloos & Sloutsky, 2008; Sloutsky, 2003). We predict that selective attention will particularly affect the accuracy, currency and completeness dimensions of IQ.

For currency[2], we predict that high domain knowledge will cause participants to ignore atypical information that describes the current state of the object being observed. In terms of the amount of Type C information, i.e. information that have no direct bearing to the instance being reported but are present in its environment:

> H2: Novices will contribute more Type C information than will Amateurs; Amateurs in turn will contribute more Type C information than Experts.

For accuracy[3], in terms of noticing and reporting type A information:

> H3a: Experts will contribute more Type A information than will Amateurs; Amateurs in turn will contribute more Type A information than Novices.
>
> Regarding detecting and reporting qualitative changes (like length, color, etc.) or quantitative changes (number) to type A attributes, we expect that:
>
> H3b: Experts will notice and report more qualitative & quantitative changes to Type A attributes than will Amateurs; Amateurs in turn will notice and report more qualitative and quantitative changes to Type A attributes than will Novices.
>
> For completeness, we predict that domain knowledge will negatively affect the amount of information provided.
>
> H4: Novices will contribute more Type B information than will Amateurs; Amateurs in turn will contribute more Type B information than Experts.

We are interested in all kinds of contributors who may participate in OIEs. Unlike previous research that has approached domain knowledge from a binary perspective – Expert or Non-expert – we consider the different levels of knowledge that may be possessed by contributors who do not fall into both categories i.e. amateurs, and even within both categories, allowing us to study the effect of domain knowledge (expertise) on a wide range of OIE contributors.

Confidence (specifically overconfidence) is argued to be a primary source of bias (see Kahneman & Tversky, 1977; Tversky & Kahneman, 1975) that affects experts more than non-experts (Spence and Brucks 1997) and is the cause of experts' poor performance in classification-like tasks. We contend that this may be an oversimplification of the fact, and while this may account for some of the results found in similar experiments, it is our position that selective attention due to classification is the main source of polarization for contributors. To test this view, we propose the following:

---

[2] Currency as defined in Nelson, Todd, & Wixom (2005), is the degree to which information precisely reflects the current situation of the world that it represents. We distinguish currency from completeness – defined as the tendency for information to capture all possible states relevant to the user – as currency covers atypical type C information, while completeness refers to type B non-diagnostic information that may be directly relevant to the user (consumer).

[3] Accuracy was also defined as the tendency for information to be correct, clear, meaningful and believable. The correct application of the inclusion rule and assessment of diagnostic information of attributes would lead to the most accurate classification.

H5: Novices will contribute more Type B and C information than will Amateurs; Amateurs in turn will contribute more Type B and C information than Experts even when experts lack confidence in their general task performance.

Proposed Experimental Design

In this paper, and like extant literature, we define expertise as *high domain knowledge* (Stein, 1992). Accordingly, we adopt the idea of performance-based measurement and definition of expertise (Clemen, 2008; Davis-Stober, Budescu, Dana, & Broomell, 2014; Lin & Cheng, 2009):

1. *Novices:* These subjects will not receive any training and have been screened to ensure that they do not have any domain knowledge required to classify the objects to be presented.
2. *Experts:* The participants are trained and tested to ensure that they have a firm understanding of the classification rules required for the task. In line with previous research, we assess participants' level of expertise based on their performance on related tasks.
3. *Amateurs:* The members of this group will receive a general introduction to the object of study, with information pertaining to all its parts and dimensions, including a generic version (imprecise) of the information provided to the experts. They will be allowed to observe the artifact and form their own concepts of the artifact.

Stimuli

The stimuli will be an artificial artifact to prevent participants from using their existing knowledge. Apart from the "Normal" artifact, the attributes (e.g. color, length of parts) of the stimuli will be manipulated to create different variations of "Deficient "and "Excess" artifacts. A Normal artifact will imply an artifact that perfectly matches the inclusion rule. They will contain all the type A information that is required to classify the artifact and a default quantity of type B and C information. A "Deficient" artifact will be an artifact that contain all the type A information in the inclusion rule, but will possess less (in number or size) of type B than the "Normal". An "Excess" artifact will be an artifact that contains a modified version of the type A information needed to successfully classify it. It may also possess more (in number or size) of type B information. Both the Deficient and Excess artifacts will contain variations of type C information. Simplified examples to illustrate the possible variations are presented in Figure 1 below.

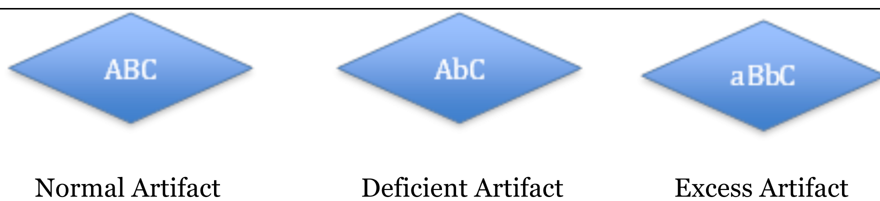

Figure 1: Possible variations of attribute types provided in the stimuli

Procedure

The experiment will be conducted in a citizen science lab setting. Undergraduate students of the business school in a mid-size Canadian university will be invited to participate in the experiment for a chance to win a small prize. Business students will be selected because there is a low probability that they will have pre-existing knowledge of the objects to be tested. First, one third of participants will be randomly selected to form our novice group. The others will be trained and tested on their ability to classify the stimulus. From these trained participants and based on their performance we will create two extra (equal) groups, experts and amateurs. The experts and amateur participants will be involved in the second experiment.

Experiment 1 - the experiment itself will be a 3*3 design as each individual will be randomly assigned to one of three groups: experts, amateurs and novices. Each participant will also receive 3 tests: reporting, comparison and detection. The order of tasks will be randomly assigned to control for learning effects.

We will compare the ability of members of each group to notice and report the features of an artificial stimulus as affected by their level of expertise. By asking open questions, we give each participant the opportunity to report sightings of the stimuli by providing its name, describing its features or doing both.

For the reporting test, each participant will receive three tests with distractions interspersed between them. Even though the order of the tests will be randomized, the first test (from the experimenter's standpoint) will involve reporting the sighting of the artifact. A normal artifact will be presented with sufficient type B and C information.

For the comparison test, participants will be presented with "distractor" artifacts, before being presented with a deficient artifact. For the Novices, in this case, they will be shown a Normal artifact for 40 seconds, after which it will be taken away (see Lukyanenko et al. 2014) and after some distractions, will be asked to compare the Normal artifact with a Deficient artifact during the reporting task.

For the detection test, expert and amateur participants will be presented with an Excess artifact and the participants will be asked to report it, testing to see if they will detect the changes in diagnostic attributes. The novice participants will be asked to compare the new artifact with the normal artifact they previously observed.

Experiment 2 – This experiment is aimed at testing the effect of overconfidence on the performance of participants (hypothesis 5). Some members of the Expert and Amateur groups who have been randomly selected to take part in the second experiment will be provided information that claims that the use of inclusion rules may not be the most effective strategy for classification. The information will support the claim against classification (a strategy used in Burton-Jones & Meso 2008). The goal is to reduce their *confidence* in the inclusion rule and in their current strategy. They will be required to repeat all the tests under this "reduced confidence" condition.

Conclusions and Potential Implications

The studies proposed in this paper can shed light on the relevance of expertise to OIEs, especially as it relates to realizing high information quality. If these hypotheses are supported, the findings will help researchers and practitioners better understand the costs and benefits of seeking out expert contributors, and may serve as a precursor to research on how to maximize these benefits while mitigating the costs. It may also further underscore the need for inclusive designs in crowdsourcing systems and OIEs, where the potential uses of information may not be fully known when a system is developed, or may evolve over time. Additionally, the hypotheses presented here, grounded in classification theory, may help us understand and explain prior research findings where expertise has been shown to be irrelevant to performance. We expect the implication of this research to go beyond traditional crowdsourcing platforms to other platforms that collect user-generated content, such as personal health records, leading to improved designs. The research will also contribute to the literature and theories in psychology, expertise, and information system design.

REFERENCES

1. Austen, G. E., Bindemann, M., Griffiths, R. A., & Roberts, D. L. (2016). Species identification by experts and non-experts: comparing images from field guides. *Scientific Reports*, *6*, 33634.
2. Best, C., Yim, H., & Sloutsky, V. (2013). The cost of selective attention in category learning: developmental differences between adults and infants. *Journal of Experimental Child Psychology*.
3. Biros, D., George, J., & Zmud, R. (2002). Inducing sensitivity to deception in order to improve decision making performance: A field study. *MIS Quarterly*.
4. Budescu, D., & Chen, E. (2014). Identifying expertise to extract the wisdom of crowds. *Management*

Science.
5. Burton-Jones, A., & Meso, P. N. (2008). The effects of decomposition quality and multiple forms of information on novices' understanding of a domain from a conceptual model. *JAIS*,*9*(12), 748.
6. Castriotta, M., & Guardo, M. Di. (2011). Open Innovation and Crowdsourcing: The Case of Mulino Bianco. *Information Technology and Innovation*.
7. Clemen, R. (2008). Improving and measuring the effectiveness of decision analysis: Linking decision analysis and behavioral decision research. *Decision Modeling and Behavior in Complex and uncertain environments* (pp. 3-31). Springer New York.
8. Davis-Stober, C., Budescu, D., Dana, J., & Broomell, S. (2014). When is a crowd wise? *Decision*.
9. Escoffier, N., & McKelvey, B. (2015). The Wisdom of Crowds in the Movie Industry: Towards New Solutions to Reduce Uncertainties. *International Journal of Arts Management*.
10. Gelman, S., & Coley, J. (1990). The importance of knowing a dodo is a bird: Categories and inferences in 2-year-old children. *Developmental Psychology*.
11. Gelman, S., Collman, P., & Maccoby, E. (1986). Inferring properties from categories versus inferring categories from properties: The case of gender. *Child Development*.
12. Gelman, S., & Markman, E. (1986). Categories and induction in young children. *Cognition*.
13. Gelman, S., & Markman, E. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*.
14. Goldstone, R. L., Kersten, A., & Carvalho, P. F. (2013). Concepts and Categorization. In *In Weiner, I.B, Healey, A.J., & Proctor, R.W. (Eds.) Handbook of Psychology, Volume 4, Experimental Psychology, 2nd Edition (pp. 607-630). New York, NY: Wiley*.
15. Hoffman, A. B., & Rehder, B. (2010). The costs of supervised classification: The effect of learning task on conceptual flexibility. *Journal of Experimental Psychology: General*, *139*(2), 319–340. http://doi.org/10.1037/a0019042
16. Hosseini, M., Phalp, K., Taylor, J., & Ali, R. (2014). Towards crowdsourcing for requirements engineering.
17. Kahneman, D., & Tversky, A. (1977). Intuitive prediction: Biases and corrective procedures. Retrieved from http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA047747
18. Keil, F. C. (2011). Science Starts Early. *Science*, *331*(6020), 1022–1023. http://doi.org/10.1126/science.1195221
19. Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General*, *137*(1), 52–72. http://doi.org/10.1037/0096-3445.137.1.52
20. Lin, S., & Cheng, C. (2009). The reliability of aggregated probability judgments obtained through Cooke's classical model. *Journal of Modelling in Management*.
21. Lukyanenko, R., Parsons, J., & Samuel, B. (2015). Do We Need an Instance-Based Conceptual Modeling Grammar ? Do We Need an Instance-Based Conceptual Modeling, (MAY). http://doi.org/10.13140/RG.2.1.4613.9365
22. Lukyanenko, R., Parsons, J., & Wiersma, Y. (2016). Participatory Design for User-generated Content: Understanding the challenges and moving forward. *Journal of Information* …. Retrieved from http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1443&context=sjis
23. Lukyanenko, R., Parsons, J., & Wiersma, Y. F. (2016). Emerging problems of data quality in citizen science. *Conservation Biology*, n/a–n/a. http://doi.org/10.1111/cobi.12706
24. Lukyanenko, R., Parsons, J., & Wiersma, Y. (2014). The impact of conceptual modeling on dataset completeness: a field experiment.
25. Lukyanenko, R., Parsons, J., & Wiersma, Y. F. (2014). The IQ of the Crowd: Understanding and Improving Information Quality in Structured User-Generated Content. *Information Systems Research*, *25*(4), 669–689. http://doi.org/10.1287/isre.2014.0537
26. Lukyanenko, R., Parsons, J., & Wiersma, Y. F. (2016). Emerging problems of data quality in citizen science. *Conservation Biology*, n/a–n/a. http://doi.org/10.1111/cobi.12706
27. Murphy, G. (2002). The big book of concepts. *MIT Press*.
28. Murphy, G., & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review*. Retrieved from http://psycnet.apa.org/journals/rev/92/3/289/
29. Nelson, R. R., Todd, P. A., & Wixom, B. H. (2005). Antecedents of information and system quality: an empirical examination within the context of data warehousing. *Journal of Management Information Systems*, *21*(4), 199–235.
30. Parsons, J. (1996). An Information Model Based on Classification Theory. *Management Science*,

42(10), 1437–1453. http://doi.org/10.1287/mnsc.42.10.1437
31. Parsons, J., Lukyanenko, R., & Wiersma, Y. (2011). Easier citizen science is better. *Nature*, *471*(7336), 37. http://doi.org/10.1038/471037a
32. Parsons, J., & Wand, Y. (2008). Using cognitive principles to guide classification in information systems modeling. *MIS Quarterly*.
33. Parsons, J., & Wand, Y. (2013). Extending Classification Principles from Information Modeling to Other Disciplines. *Journal of the Association for Information*.
34. Parsons, J., & Wand, Y. (2014). A Foundation for Open Information Environments.
35. Schultze, U., & Leidner, D. (2002). Studying knowledge management in information systems research: discourses and theoretical assumptions. *MIS Quarterly*.
36. Sloutsky, V. (2003). The role of similarity in the development of categorization. *Trends in Cognitive Sciences*.
37. Smith, E., & Medin, D. (1981). *Categories and concepts*. Retrieved from http://www.cs.indiana.edu/~port/teach/sem08/Smith.Medin.1983.ch1.2.3.pdf
38. Spence, M. T., & Brucks, M. (1997). The moderating effects of problem characteristics on experts' and novices' judgments. *Journal of marketing Research*, 233-247.
39. Stein, E. (1992). A method to identify candidates for knowledge acquisition. *Journal of Management Information Systems*.
40. Tarrell, A., Tahmasbi, N., Kocsis, D., Pedersen, J., Tripathi, A., Onook, O., … DeVreede, G. (2013). Crowdsourcing: A snapshot of published research. *Proceedings of the Nineteenth Americas Conference on Information Systems, Chicago, Illinois, August 15-17, 2013*.
41. Tripathi, A., & Tahmasbi, N. (2014). Crowdsourcing typology: a review of is research and organizations. *Proceedings of the Ninth Midwest Association for Information Systems Conference, Ames, Iowa May 15-16, 2014*.
42. Tversky, A., & Kahneman, D. (1975). Judgment under uncertainty: Heuristics and biases. *Utility, Probability, and Human Decision Making*. Retrieved from http://link.springer.com/chapter/10.1007/978-94-010-1834-0_8
43. Wand, Y., Monarchi, D., Parsons, J., & Woo, C. (1995). Theoretical foundations for conceptual modelling in information systems development. *Decision Support Systems*.