Do Crowds Go Stale? Exploring the Effects of Crowd Reuse on Data Diversity

Complete Research Paper

Shawn Ogunseye
Memorial University of Newfoundland
osogunseye@mun.ca

Jeffrey Parsons
Memorial University of Newfoundland
jeffreyp@mun.ca

Roman Lukyanenko
University of Saskatchewan
lukyanenko@edwards.usask.ca

Abstract

Crowdsourcing is increasingly used to engage people to contribute data for a variety of purposes to support decision-making and analysis. A common assumption in many crowdsourcing projects is that experience leads to better contributions. In this research, we demonstrate limits of this assumption. We argue that greater experience in contributing to a crowdsourcing project can lead to a narrowing in the kind of data a contributor provides, causing a decrease in the diversity of data provided. We test this proposition using data from two sources – comments submitted with contributions in a citizen science crowdsourcing project, and three years of online product reviews. Our analysis of comments provided by contributors shows that the length of comments decreases as the number of contributions increases. Also, we find that the number of attributes reported by contributors decreases as they gain experience. These finding support our prediction, suggesting that the diversity of data provided by contributors declines over time.

Keywords: Crowdsourcing, Data Quality, Data Diversity, Crowd-Reuse, Crowd Staleness

Introduction

Crowdsourcing involves outsourcing tasks traditionally carried out by employees, or others tightly linked to an organization or group, to an unknown and undefined group of people (Afuah & Tucci, 2012; Howe, 2006). More organizations, groups and individuals (crowdsourcers) are turning to crowdsourcing for tasks that include the collection and analysis of distributed data (Conrad and Hilchey, 2011; Cooper et al., 2007). For crowdsourcing projects, the quality of data collected is of utmost importance (Schenk and Guittard 2011).

One strategy to increase data quality is to control who gets to do a crowdsourcing task. To ensure that crowdsourced data is of high quality, crowdsourcers often decide who to recruit as contributors (Gura, 2013; Malone et al. 2010). A key issue in this regard is whether to recruit only those with prior experience in the data collection task, or to allow (or even encourage) participation by contributors with a range of experience or domain knowledge, including novices (Ogunseye and Parsons 2016; Lukyanenko, Parsons & Wiersma 2016).

Crowdsourcers must also concern themselves with how to recruit, giving rise to two important recruitment questions. For a new project, should participants be recruited from an existing crowd (e.g., from a similar project, within the same platform, such as Zooniverse) or should a new recruitment campaign be run? Should crowdsourcers actively recruit new participants through the lifecycle of their projects (i.e. continuous recruitment) or should recruitment be a singular event at the start of a project? We consider these issues as questions of the extent of *crowd reuse* - a crowdsourcing strategy that consists of drawing from the same crowd repeatedly for similar tasks across one or more crowdsourcing projects.

Given crowdsourcers' interest in ensuring high data quality, increasing crowd experience via crowd reuse is generally considered desirable (Galloway et al., 2006; Gura, 2013). Highly experienced contributors may have greater knowledge of the task and domain than amateurs or novices. This leads to a preference for experienced contributors over novices or amateurs, which in turn influences the recruitment of crowd members and the design of crowdsourcing systems (Wiggins et al. 2011, Austen et al., 2016, Lukyanenko et al. 2014a). Crowdsourcers often seek to retain their existing crowd while pursuing avenues to increase contributors' knowledge and participation (Nov et al., 2011; Rotman et al., 2012).

Data quality is traditionally measured by its intrinsic quality (e.g., accuracy), contextual quality (e.g., completeness, currency) and representational quality (e.g., comprehensibility) (Nelson, Todd, & Wixom, 2005; Wang & Strong, 1996). Research on the quality of crowdsourced data mainly focuses on the accuracy of data in attempts to address data quality issues (Budescu & Chen 2014; Wiggins, Newman, Stevenson, & Crowston 2011), paying less attention to other relevant dimensions of data quality such as completeness, comprehensibility and currency (Lukyanenko et al. 2016). However, these other dimensions are important as they determine a crowdsourcer's ability to verify, reuse, or repurpose contributed data.

The need to better understand the impact of repeated contributions by the same volunteers is important since the relationship between crowdsourcing expertise and quality is not a straightforward one. There are numerous reported examples where contributor experience and/or knowledge did not influence the accuracy of crowdsourced data (Austen et al. 2016; Kallimanis, Panitsa, & Dimopoulos 2017). e posit that quantity and diversity of contributions in a crowdsourcing project will decline when crowds are repeatedly exposed to the same project categories. This contention goes against the prevailing wisdom that more experience result in higher data quality, motivating the need to better understand how increased crowd knowledge gained through crowd reuse affects aspects of data quality beyond accuracy. Such insights will be useful to crowdsourcing organizations in their recruitment decisions (Ogunseye & Parsons 2017), and in the creation of more effective data collection designs sensitive to the nature of the crowds involved in their projects (see Lukyanenko, Parsons & Wiersma 2014a, 2014b). This is particularly important in crowdsourcing tasks centered on pooling complementary input from the crowd and fostering discoveries of new phenomena. Such crowdsourcing tasks include many tasks carried out in *citizen science crowdsourcing* – outsourcing scientific data collection and/or analysis to the general public using information technology (Lee, Crowston, Østerlund & Miller, 2017).

Considering the effort involved in recruiting, training and retaining volunteers, crowd reuse is a common practice. Specific examples in the domain of citizen science include:

(1) the eteRNA project, in which ordinary citizens help design RNA sequences that fold into particular shapes previously unknown to expert scientists (Bohannon 2016, Anderson-Lee et al. 2016); (2) Galaxy Zoo, in which participants classify galaxies in images taken by the Hubble telescope, and in which non-expert contributors have been instrumental in the discovery of an important astronomical phenomenon (See Lukyanenko et al. 2016, Clery 2011, Cardamone et al. 2009); and (3) eBird (ebird.org), a project that relies on thousands of dedicated volunteers that report millions of monthly observations (Callaghan and Gawlik 2015). Indeed, GalaxyZoo shares a common platform – Zooniverse.org - with other similar projects (e.g., Planet Hunters, Asteroid Zoo, Old Weather, Snapshot Serengeti) allowing the same crowds to participate in a variety of other tasks. Likewise, companies like Crowdflower and Amazon Mechanical Turk provide access to a massive pool of workers available on demand including for repeated use in the same or similar crowd sourcing tasks (Peer, Vosgerau, & Acquisti 2014; Paolacci & Chandler 2014).

Crowd Reuse and Data Diversity

In ongoing projects, committed contributors are learning and participating in a cycle within the project. The knowledge gained through participation (e.g., training, experience, self-study, or continued use of the crowdsourcing platform) affects future contributions in the same project (Jordan et al., 2011, Lukyanenko et al. 2014b). Further, crowds may self-aggregate by common interests and continually offer their services to projects they consider interesting (Gura, 2013; Newman et al., 2012). As considerable effort and cost might be needed to attract and retain volunteers, it is clear why crowdsourcers prefer a stable cohort of volunteers. Notwithstanding the benefits of crowd reuse, in this research we seek to examine potential limitations of relying on the same crowd, particularly for projects that engage crowds in discoveries or that evolve to encompass uses of data that were not anticipated when the project was designed and initiated.

Reuse of crowds in crowdsourcing increases task knowledge and manifests as experience and task expertise. For example, in an interview with *Science,* a molecular biologist – Arthur Olson – said of the success of ordinary citizens in predicting the 3D structure of protein through the FoldIt crowdsourcing platform: "I didn't know …that this game would actually create experts" (Bohannon 2010). Although expertise is generally viewed as desirable, we examine whether increased experience (regardless of domain knowledge) makes crowds "stale". We define ***crowd staleness*** as the extent to which a crowd provides less diverse data over time, as a consequence of increasing experience in reporting observations of the phenomena being studied. The theoretical basis for our claim is discussed next.

Theoretical Underpinnings

Classification (or categorization) is a fundamental human capability (Hoffman & Murphy, 2006). We learn by classification (Piaget & Cook, 1952) and continue to classify as a way to make efficient use of our cognitive resources (Goldstone & Kersten, 2003; Lakoff, 1982). By classifying, we organize our existing knowledge about phenomena mainly through their similarities, which allows us to make predictions about new instances and events (Best et al., 2013). The literature on classification posits that, due to our tendency to classify, as we gain experience we learn to identify instances as members of classes by paying selective attention to only relevant features crucial for identifying them, while irrelevant features (those not useful for predicting class membership) are safely ignored. Although such "selective attention" engenders efficient learning, it leads to a learned inattention to features that are not diagnostic for a class in a particular context (Colner & Rehder 2009; Hoffman & Rehder 2010), resulting in a reduction in the diversity of features we attend to.

The tendency toward selective attention and classification occurs naturally in humans as we acquire knowledge about entities in our world. In contrast, the absence of this tendency is "a developmental default" (Kloos & Sloutsky, 2008 p. 68; also see Gelman, 1988). Experiments by Best et al. (2013), comparing the ability of infants and adults to selectively attend to attributes of instances based on prior or current knowledge or experience, show that infants do not have the capacity for selective attention. Young children reason about classes by observing all of the features of individual instances. As they grow, they begin to pay selective attention to particular features determined to be relevant to the task at hand (Gelman & Markman, 1986; Plebanek & Sloutsky, 2017). Since knowledge naturally increases with experience, it is expected that contributors' knowledge about a crowdsourcing task will increase as they gain experience (Harnad, 2005). Experienced contributors are, therefore, expected to focus less on the non-diagnostic attributes of instances of phenomena, but concentrate on diagnostic attributes (attributes they have learned from prior knowledge and experience that help identify the instance). With increasing experience and a corresponding increase in learned inattention, contributors will provide less data for each observation manifesting as reduced length of comments and a reduction in the number of attributes reported in comments. We hypothesize accordingly:

> **H1. Data Quantity.** *The length of comments provided by contributors will decrease as they gain experience in participation in a crowdsourcing platform.*
>
> **H2. Data Diversity.** *The number of attributes used by contributors to describe a phenomenon will decrease as they gain experience in participation in a crowdsourcing platform.*

Empirical Evidence

To test the research hypotheses, we examined the relationship between contributor experience and the quantity and diversity of data they provide. First, we tested the relationship between experience and the quantity of data provided in a crowdsourcing context. Second, we examined the relationship between contributor experience and the number of attributes they report.

To increase the reliability and generalizability of our results, we analyzed data from two sources – both used in prior research and representing different crowdsourcing domains (He and McAuley 2016, Lukyanenko et al. 2017): (1) NLNature (www.nlnature.com);

and (2) Amazon.com. NLNature is an existing citizen science project to collect observations of flora and fauna in the Canadian province of Newfoundland and Labrador. Two of the authors participated in the design of the project. Participation was open to any resident of or visitor to the region, irrespective of level of expertise. The project encouraged participants to provide additional data in the comment section beyond the basic organism identification and description task they performed. These comments typically contained rich data to accompany an observation (see Table 1). Of 9148 submissions made to the crowdsourcing platform in this period of time, we found 6447 usable for this study as they were: (a) unique (not duplicates); and (b) were made by participants who contributed at least 30 times from January 1, 2014 to June 6 2017. We arbitrarily classified contributors who reported 30 or more observations in this period as "experienced" (in future work, we intend to vary the number of contributions considered as experienced to test whether our results are dependent on the chosen cutoff. In total, 39 contributors provided the comments we analyzed in this paper. We use the number of contributions as a measure of experience.

Review data from Amazon's "baby products" category (see source details in He & McAuley 2016) was also analyzed using the same criteria discussed above for a period of 3 years: 2010-2012. We found these years suitable for our analyses because before these years, none of the participants in the dataset had contributed 30 or more times. The dataset therefore self-mitigates the possibility that the reviewers were experienced reviewers in the domain on entering the time window we consider.

Table 1. Sample contributions made to the citizen science project

Source	Identified Organism	Comment
NLNature	Mink (Neovison vison)	*may have killed a snowshoe hare found nearby that didn't appear to eaten as a mammal such as a coyote would have*
	Fox	*fox eating something buried under a small mound of grass*
Amazon.com	Baby Toy	*This is colorful, soft, and makes lots of fun sounds babies love. I would recommend it to anyone for their baby.*

The results of the analysis of the two datasets are shown in Figures 1 and 2. The figures reveal that, regardless of the initial level of knowledge of the contributor, and when (prior or during participation) the knowledge was acquired, there is a decrease in the length of comment as experience increases (by about one half word per comment for the NL Nature data and nearly six words per comment for the 2010 Amazon data). This is consistent with our theoretical arguments. Table 2 summarizes the results for the Amazon data for each of the three years. The results offer strong support for Hypothesis 1, based on data from two different crowdsourcing contexts.

To examine the possibility that the number of properties used by contributors to describe a phenomenon decreases as contributors gain experience in participation, we analyzed data from Amazon.com. From the Amazon reviews, thirty pairs of comments were randomly selected, each pair from one contributor.

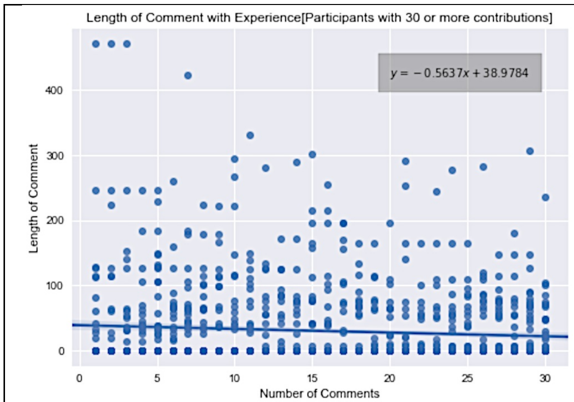

Figure 1: NLNature

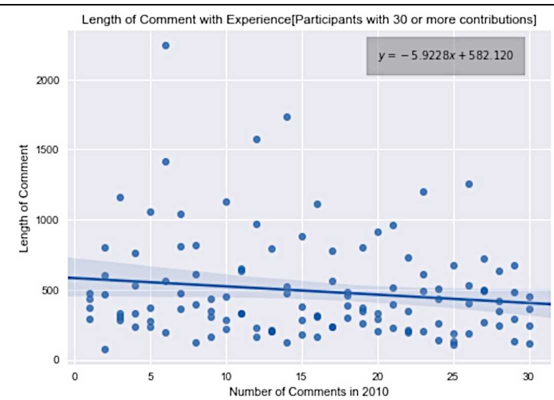

Figure 2: Amazon Reviews 2010 [p<0.01]

(Similar results were obtained for years 2011 and 2012)

Table 2: Summary of Results of Amazon Data Reviewed

Year	Participants (used)	Comments (used)	p-value	R^2	Mean comment length
2010	1142 (143)	5173 (120)	9.15e-17	0.442	692
2011	2094 (416)	11499 (270)	1.43e-14	0.198	671
2012	3095 (906)	17904(690)	3.87-e55	0.299	615

For the contributors selected, one comment was from the first three comments made in that year and another comment from their last three. Two of the authors independently coded each comment in this sample to determine the total number of attributes of the reported entity the comment contained. The third author resolved coding disagreements among the coders. The number of attributes is used as a measure of the diversity of the contributions. The result and summary statistics are shown in Figure 3 below. The regression analysis performed on the coded sample indicates that experience accounts for over 20% of the variance in the total number of attributes reported by crowds.

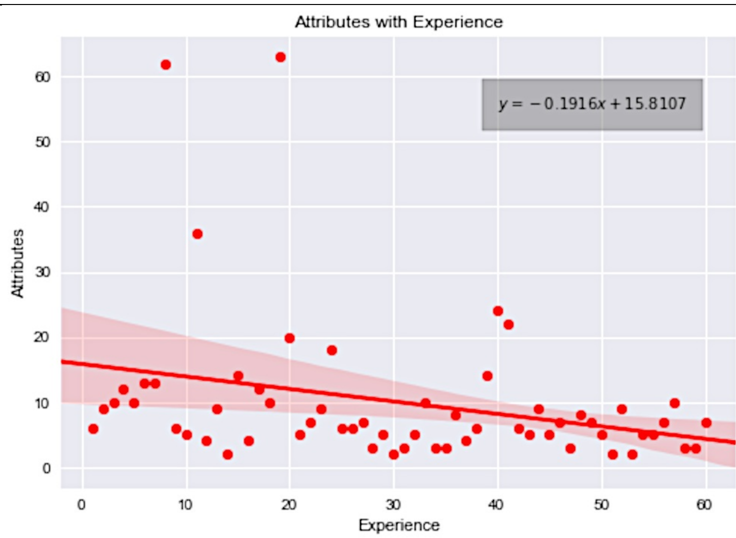

Figure 3: Change in number of attributes with increasing experience

```
R-squared:                     0.215
P (F-statistic):               0.000170
```

The statistically significant negative slope confirms the decline in the reporting of attributes predicted in Hypothesis 2.

Discussion

In this paper, we seek to answer the question: does the diversity of crowdsourced data decline as crowds gain knowledge and experience? To answer this question, we provide arguments for how increasing crowd knowledge may lead to crowd staleness, diminishing the tendency for crowd members to provide diverse data. To test the validity of our argument, we used data from two real world data sources to examine the diversity in data provided by a stable crowd over a period, investigating if the tendency of crowd members to provide diverse data declines with experience. Overall, we find that diversity, measured by the length of contributed comments and the number of attributes reported in comments, declines with the number of contributions made regardless of the level of prior knowledge of the contributor (i.e. whether they are experts or novices). Even though our research is preliminary, our results suggest that crowds can become stale, implying a need to continually recruit to make up for declining diversity due to staleness.

Finally, while we do not test the relationship between crowd reuse and data accuracy, we believe crowd staleness could also negatively affect accuracy. Studies demonstrate that crowds seek to report data in the format they deem appropriate for the crowdsourcer, which may result in lower accuracy if people guess or select choices at random (Parsons et al. 2011, Lukyanenko et al. 2014a). For example, in biology-related projects, contributors may feel the need to report observed organisms as biological species, and thus may provide potentially inaccurate species identification labels. In contrast, most volunteers are able to describe the observed organisms using a variety of attributes or generic classes (e.g., bird), but if they perceive this to be of no use to the crowdsourcer (an awareness volunteers may reach with some experience with the project, see Lukyanenko et al. 2014a), they may avoid doing so to the detriment of observational accuracy and data diversity. Future work is needed to apply our theoretical propositions to the relationship between crowd reuse/crowd staleness and data accuracy.

Conclusion

Online crowdsourcing is an increasingly popular tool to collect data from diverse and distributed crowds. Notwithstanding the many advantages of crowdsourcing, we do not yet understand the most effective means of engaging crowds. In this paper, we challenge a common and tacit assumption that engaging the same volunteers – crowd reuse – brings only benefits to crowdsourcing projects. Indeed, as we show, as crowd members gain experience, the diversity of data they provide may decline. As a starting point to understanding the issue of crowd staleness, this work has focused on the analyses of comment length and number of attributes included in a comment as surrogates for data diversity. Future studies are needed to more fully explore the characteristics of diversity that might be affected by crowd staleness and to propose and test measures to mitigate any loss of diversity.

References

Afuah, A., & Tucci, C. L. 2012. "Crowdsourcing as a solution to distant search," *Academy of Management Review*, (373), 355-375.

Anderson-Lee, J., Fisker, E., Kosaraju, V., Wu, M., Kong, J., Lee, J., … Das, R. 2016). "Principles for Predicting RNA Secondary Structure Design Difficulty," *Journal of Molecular Biology*, *428*(5, Part A), 748–757.

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*.

Best, C. A., Yim, H., and Sloutsky, V. M. 2013. "The cost of selective attention in category learning: Developmental differences between adults and infants," *Journal of experimental child psychology*, (116:2), pp. 105–119.

Bohannon J., 2010. "Video Game Helps Solve Protein Structures", www.sciencemag.org

Bohannon J., 2016. "For RNA paper based on a computer game, authorship creates an identity crisis", www.sciencemag.org

Budescu, D. V. & Chen, E. 2014. "Identifying expertise to extract the wisdom of crowds", *Manag. Sci.* **61,** 267–280.

Callaghan, C.T., and Dale E. G. 2015. "Efficacy of eBird data as an aid in conservation planning and monitoring," *Journal of Field Ornithology 86,* no. 4 (2015): 298-304.

Cardamone, C., Schawinski, K., Sarzi, M., Bamford, S.P., Bennert, N., Urry, C.M., Lintott, C., Keel, W.C., Parejko, J., Nichol, R.C. and Thomas, D., 2009. Galaxy Zoo Green

Peas: discovery of a class of compact extremely star-forming galaxies. *Monthly Notices of the Royal Astronomical Society*, *399*(3), pp.1191-1205.

Clery, D. 2011. "Galaxy Zoo volunteers share pain and glory of research," *Science* 333**,** 173–175

Collins, H., & Evans, R. 2007. "Rethinking expertise," *University of Chicago Press.* Chicago IL.

Colner, B. and Rehder, B. 2009. "A new theory of classification and feature inference learning: An exemplar fragment model," in *Proceedings of the 31st Annual Conference of the Cognitive Science Society, pp. 371–376*

Conrad, C. C., and Hilchey, K. G. 2011. "A review of citizen science and community-based environmental monitoring: issues and opportunities," *Environmental monitoring and assessment,* (176:1), pp. 273–291.

Cooper, C., Dickinson, J., Phillips, T., and Bonney, R. 2007. "Citizen science as a tool for conservation in residential ecosystems," *Ecology and Society,* (12:2)

Erickson, L., Petrick, I., & Trauth, E. 2012. "Hanging with the right crowd: Matching crowdsourcing need to crowd characteristics. *Proceedings of the Eighteenth Americas Conference on Information Systems,* Seattle, Washington, August 9 12, 2012

Galloway, A. W., Tudor, M. T., and Haegen, W. M. V. 2006. "The reliability of citizen science: a case study of Oregon white oak stand surveys," *Wildlife Society Bulletin,* (34:5), pp. 1425–1429.

Gelman, S. A., and Markman, E. M. 1986. "Categories and induction in young children," *Cognition*, (23:3), pp. 183–209.

Goldstone, R. L. and Kersten, A. 2003. "Concepts and categorization," *Handbook of psychology.*

Gura, T. 2013. "Citizen science: Amateur experts," *Nature,* (496:7444), pp. 259–261 19

Harnad, S. 2005. "To cognize is to categorize: Cognition is categorization," *Handbook of categorization in cognitive science*, pp. 20–45.

He, R. and McAuley, J., 2016. "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," *In Proceedings of the 25th International Conference on World Wide Web* (pp. 507-517)*. International World Wide Web Conferences Steering Committee.*

Hoffman, A. B., and Murphy, G. L. 2006. "Category dimensionality and feature knowledge: When more features are learned as easily as fewer.," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* (32:2), p. 301.

Hoffman, A. B., and Rehder, B. 2010. "The costs of supervised classification: The effect of learning task on conceptual flexibility," *Journal of Experimental Psychology: General,* (139:2), p. 319.

Jordan, R. C., Gray, S. A., Howe, D. V., Brooks, W. R., and Ehrenfeld, J. G. 2011. "Knowledge gain and behavioral change in citizen-science programs," *Conservation Biology,* (25:6), pp. 1148–1154.

Kallimanis, A. S., M. Panitsa, and P. Dimopoulos 2017. "Quality of non-expert citizen science data collected for habitat type conservation status assessment in Natura 2000 protected areas." *Scientific Reports* 7, no. 1, 8873.

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.

Lakoff, G. 1982. "Categories", *Linguistics in the morning calm*

Lee, T.K. Crowston, K., Østerlund, C. & Miller, G. 2017. "Recruiting messages matter: Message strategies to attract citizen scientists". in *CSCW 2017 - Companion of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing.* Association for Computing Machinery, Inc, pp. 227-230, 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW 2017, Portland, United States, 25-1 March.

Lukyanenko, R., Parsons, J., and Wiersma, Y. F. 2014. "The IQ of the crowd: understanding and improving information quality in structured user-generated content," *Information Systems Research,* (25:4), pp. 669–689.

Lukyanenko, R., Parsons, J., & Wiersma, Y. (2014). The impact of conceptual modeling on dataset completeness: a field experiment. ICIS 2014.

Lukyanenko, R., Parsons, J., & Wiersma, Y. F. 2016. Emerging problems of data quality in citizen science. *Conservation Biology*, *30*(3), 447-449.

Lukyanenko, R., Parsons J, Wiersma Y., Wachinger G., Huber B. & Meldt R. (2017). Representing Crowd Knowledge: Guidelines for Conceptual Modeling of User-generated Content. Journal of the Association for Information Systems (JAIS). 18 (4), pp. 297-339.

Malone, T. W., Laubacher, R., and Dellarocas, C. 2010. "The collective intelligence genome," *MIT Sloan Management Review,* (51:3), p. 21.

Nelson, R. R., Todd, P. A., and 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), pp. 199–235.

Newman, G., Wiggins, A., Crall, A., Graham, E., Newman, S., and Crowston, K. 2012. "The future of citizen science: emerging technologies and shifting paradigms," *Frontiers in Ecology and the Environment*, (10:6), pp. 298–304.

Nov, O., Arazy, O., and Anderson, D. 2011. "Dusting for science: motivation and participation of digital citizen science volunteers," in *Proceedings of the 2011 iConference ,* ACM, pp. 68–74

Ogunseye, S., and Parsons, J. 2016. "Can Expertise Impair the Quality of Crowdsourced Data?," *Proceedings of the 15th AIS SIGSAND Symposium,* Lubbock, Texas

Ogunseye, S., and Parsons, J. 2017. "What Makes a Good Crowd? Rethinking the Relationship between Recruitment Strategies and Data Quality in Crowdsourcing." *Proceedings of the 16th AIS SIGSAND Symposium,* Cincinnati, Ohio

Parsons, J., Lukyanenko, R. and Wiersma, Y., 2011. "Easier citizen science is better." *Nature*, *471*(7336), pp.37-37.

Peer, E., Vosgerau, J., & Acquisti, A. 2014. "Reputation as a sufficient condition for data quality on Amazon Mechanical Turk," *Behavior research methods*, 46(4), 1023-1031.

Paolacci, G., & Chandler, J. 2014. "Inside the Turk: Understanding Mechanical Turk as a participant pool," *Current Directions in Psychological Science,* 23(3), 184-188.

Piaget, J., and Cook, M. 1952. "The origins of intelligence in children," *International Universities Press* New York*,* (Vol. 8)

Plebanek, D. J., and Sloutsky, V. M. 2017. "Costs of Selective Attention: When Children Notice What Adults Miss," *Psychological Science*, p. 956797617693005

Rotman, D., Preece, J., Hammock, J., Procita, K., Hansen, D., Parr, C., Lewis, D., and Jacobs, D. 2012. "Dynamic changes in motivation in collaborative citizen-science projects," *in Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work ,* ACM, pp. 217–226

Schenk, E., & Guittard, C. 2011. "Towards a characterization of crowdsourcing practices," *Journal of Innovation Economics & Management*, (1), 93–107.

Wang, R. Y., and Strong, D. M. 1996. "Beyond accuracy: What data quality means to data consumers," Journal of management information systems, (12:4), pp. 5–33.

Wiggins, A., Newman, G., Stevenson, R. D. & Crowston, K. 2011. "Mechanisms for Data Quality and Validation in Citizen Science," in *IEEE Seventh International Conference on e-Science Workshops* 14–19 (2011).