Crowdsourcing Literature Reviews in New Domains

Michael Weiss

> " *This is a needle in a haystack problem where* "
> *the appearance of the needle is unknown.*
>
> Lin et al. (2014)

Conducting a literature review in new domains presents unique challenges. The literature in a new domain is typically broad, fragmented, and growing quickly. Because little is known about the new domain, the literature review cannot be guided by established classifications of knowledge, unlike in an existing domain. Rather, it will be driven by evidence that challenges and extends existing knowledge. In a way, exploring a new domain means looking for anomalies in the evidence that cannot be explained by what is already known. This article summarizes lessons from conducting two literature reviews in new domains in the area of cybersecurity. It then presents a design for using leader-driven crowdsourcing to collect evidence and synthesize it into insights in a new domain. The article will be relevant to those who are exploring a new domain, in particular students, researchers, and members of R&D projects in industry.

Introduction

A standard approach to exploring a domain is to conduct a literature review. However, conducting a literature review in a *new* domain presents unique challenges. Whereas in an existing domain, researchers can use established classifications of knowledge to guide their search for and interpretation of the literature, this is not the case for a new domain that lacks such classifications. In a new domain, the literature is typically broad, fragmented, and, at the same time, growing quickly. The task of the researcher is to make sense of evidence when it does not fit existing models and classifications. Encountering such evidence forces them to extend existing knowledge.

This article first examines the characteristics of new domains and summarizes lessons from conducting two literature reviews in new domains. It then reviews the goals and types of literature reviews and the typical structure of a systematic narrative literature review. Third, it introduces crowdsourcing as a technique for leveraging groups of people to solve complex tasks and examines the problems crowdsourcing can solve. The article then presents a design for crowdsourcing the creation of literature reviews to collect evidence and synthesize it into insights in a new domain. The article closes with the identification of challenges and open questions when using this new approach.

Exploring New Domains

Exploring a new domain can be conceptualized as looking for anomalies in the evidence that cannot be explained by what is already known, and subsequently building models and classifications that incorporate this evidence. A particular challenge in exploring a new domain is that the very criteria for searching the domain are co-evolving with our understanding of the domain. At the outset of the literature review, there are few established criteria for what the researchers should be looking for, something that Lin and colleagues (2014), in their study on crowdsourcing the search for Genghis Khan's tomb, refer to as a "needle in a haystack problem where the appearance of the needle is unknown".

Crowdsourcing Literature Reviews in New Domains

Michael Weiss

Searching for unknowns

Very little is known about what the key concepts in the new domain are. Thus, researchers should not limit their search criteria to what can be ''expected'' based on existing literature. As observed by Attenberg, Ipeirotis, and Provost (2011), organizations make decisions based on explicit or implicit models of the world. While it is important to understand where these models have limitations and can be improved, it is often not clear when these models are limited. In other words, we often don't know what we don't know.

From requirements engineering, we also know that ignorance of a domain often has advantages (Berry, 1995). It allows a requirements analyst to uncover unstated assumptions that domain experts have come to accept. Experts have tacit knowledge of a domain (aspects of the domain they take for granted), whereas an ignorant "newbie" in the domain would have to think about those aspects explicitly and evaluate them from first principles (Mehrotra & Berry, 2012). We consider this outsider's perspective as the "newbie's advantage".

Attenberg and colleagues (2011) also recognize the advantage of a non-expert's perspective. They found that non-experts can easily find holes in decision models that pass "standard" tests used by experts. These holes in an organization's decision model correspond to situations where the model is confident but wrong (these are "unknown unknowns"), not where the model is uncertain ("known unknowns"). From these observations, we conclude that, in a new domain, researchers should especially be looking for areas in the existing knowledge that are supposedly firmly established. This is where the biggest blind spots may lie.

Lessons from two literature reviews in new domains

The author had an opportunity to observe two teams conducting literature reviews in new domains within the area of cybersecurity (see the Acknowledgements). The teams consisted of experienced researchers and graduate students and early-career researchers. All team members had prior experience writing traditional literature reviews. The key observations were:

1. *Fragmentation and size of domain:* There were not yet established classifications of the knowledge in the new domains and the knowledge appeared fragmented. This observation was more apparent in one of the reviews, which lacked a reference point for starting the literature search.

2. *Evolving search criteria:* Questions drive the search for evidence and the search criteria evolve with the understanding of the domain. Competing interpretations require adjustments to the search criteria.

3. *Output of literature review:* The intent of the literature review is to obtain a sense of the future evolution of the domain, and to identify gaps and challenges. While, in some sense, every literature reviews strives to achieve those goals, a literature review in a new domain will put more emphasis on these aspects. Our understanding of a new domain starts with gaps in and challenges to the existing literature.

4. *Grounded in examples:* The review is grounded in examples of the phenomenon investigated.

5. *Non-traditional sources of literature:* Because the domain is still evolving, other sources than traditional conference and journal papers need to be considered (e.g., online presentations and news articles).

6. *Diversity:* In these two cases, the team consisted of generalists and specialists. The generalists in the team had a broad background in technology and innovation, whereas the specialists had expertise in cybersecurity. However, none of them had specific expertise in the emerging domains.

7. *Modularity:* The search and interpretation of literature was chunked into independent pieces. This observation is more applicable to one of the reviews, where scoping the domain into subdomains helped focus the review process.

8. *Leader-driven scoping and synthesis:* Questions (scoping) and synthesis of the answers were driven by one individual (an experienced researcher), who took a lead role in the literature review process.

Table 1 summarizes the evidence for these observations, which the author solicited by email from the team members. The team members were presented with an initial version of the eight observations above and asked to comment on them. The quotations are provided as they appeared in the emails, except for correcting obvious spelling or grammatical errors.

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

Table 1. Evidence collected from the authors of two literature reviews in new domains

Observation	Evidence
1. Fragmentation and size of domain	"With [Review 2] we had Machine Learning and its structures to start with. So, that classification exists. But, especially with the new evolving techniques, papers were rare and scattered."
	"[Review 1] didn't really have a starting structure at all, and we had to search around for the meaning of code reuse."
2. Evolving search criteria	"[We] used survey papers, where possible, to start the exploration and capture a broad context that would be refined."
	"[We] started by looking for recent papers (post 2010) based on evolving sets of keywords as we explored the domain."
	"I'm trying to show keyword search evolution in three phases. I would consider the first process was an initial validation of keywords to use for search, the second phase was a top-down approach for finding literature, and the third phase was a bottom-up approach to search.
3. Output of literature review	"Overall, the intent of the article is to contribute an understanding of the recent literature and a sense of future directions."
	"Certainly trying to capture trends, but also to identify gaps, challenges, etc."
4. Grounded in examples	"The review is grounded in examples of the phenomenon investigated. [Review 1] collects examples of code reuse covered by the media."
	"[We] extracted snippets from papers making points that we deemed of interest as pertaining to the focus questions from the client [funding the literature review]. [It] helped to ground the review as we moved forward in a domain [where] we were not experts."
5. Non-traditional sources of literature	"For [Review 1], we had to search outside Google Scholar for definitions and tutorials including YouTube videos."
6. Diversity	"I agree that team diversity is important – especially in a novel domain. In this case, I think it could be argued that we had no specialists on the team. I guess it depends upon the abstraction. Am I a specialist because I have insights into cybersecurity?"
7. Modularity	"True for [Review 2]; slightly less true for [Review 1]. [Review 2] resulted in five modules being written on specific subject matters (that, in a way, reflected the overall categorizations within Machine Learning)."
	"[Researcher B] modularized/scoped the research into subdomain topics that I thought included foundational elements of machine learning (i.e., feature extraction, clustering, datasets)."
8. Leader-driven scoping and synthesis	"I think in spirit this is true. But, in a way, it was the contract that set up the questions/direction. [Researcher A], however, certainly was key in determining the process used to perform the literature review."
	"In both [Review 1] and [Review 2], the project leader had to define the depth of the search as well as the level of details in the synthesis."
	"I concur; ultimately someone had to make a decision on scope."

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

Goals and Types of Literature Reviews

A literature review aims to summarize the current knowledge on a given topic based on previously published research. The authors of a literature review search through the literature, retrieve sources of information, and synthesize the findings of those sources into one paper (Green et al., 2006). We can classify literature reviews in terms of their goals and the ways in which the literature review is conducted. Baumeister and Leary (1997) identify five possible goals of a literature review. Starting with the most ambitious goal, these are: developing theory, evaluating theory, surveying the state of knowledge on a particular topic, identifying a gap or a problem, and, in some cases, providing a historical account of the development of theory and research on a particular topic.

Green et al. (2006) differentiate three broad categories of literature reviews:

1. Narrative literature reviews synthesize the findings of literature retrieved from searches of databases, manual searches, and authoritative texts. They are helpful when presenting a broad perspective on a topic. However, they are usually less systematic and comprehensive than other types of literature reviews and may be biased to one researcher's perspective. Editorials, commentaries, and overview articles are all examples of narrative literature reviews.

2. Qualitative systematic literature reviews are based on a detailed search of the literature. They are driven by a focused question or purpose. A systematic literature review aims to decrease the amount of bias that can occur when evidence is extracted from the literature by establishing systematic criteria for selecting literature to include in the survey and including multiple authors in the review. Results of the review are typically compiled in evidence tables.

3. Quantitative systematic literature reviews synthesize the results of the reviewed literature in a statistical manner. A quantitative literature review is also known as a meta-analysis.

It is also possible to create a taxonomy of literature reviews by combining the goals and types of literature reviews (Pare et al., 2015). Other authors such as Grant & Booth (2009) have created more detailed classifications of literature reviews. The literature reviews conducted by the two teams above can best be characterized as a systematic narrative literature review. They are more systematic than a narrative literature review, but do not meet all the formal requirements of a qualitative systematic literature review. This type of literature review is the focus of our paper.

Structure of a Systematic Narrative Literature Review

Green, Johnson, and Adams (2006) describes a (systematic) narrative literature review as a "best-evidence synthesis". A best-evidence synthesis contains the following elements:

1. *Focus:* The authors should state the purpose or focus of the literature review.

2. *Relevance:* The authors also need to make a case for the relevance of the review.

3. *Glossary:* The literature review should define any unusual terminology.

4. *Sources of information:* The authors of the literature review need to report on the electronic databases searched and the keywords used to search for papers.

5. *Search terms:* To limit the number of papers that need to reviewed, the authors should turn the main concepts of the domain under exploration into search terms.

6. *Selection criteria:* The literature review should describe on what grounds papers were included or excluded. Such criteria help avoid bias in the selection of the papers.

7. *Synthesis:* The information obtained from the literature should be organized into common themes or streams. Tables are a good way of categorizing the evidence collected. A goal of the synthesis is to identify agreements, disagreements, and gaps in the literature.

8. *Limitations:* The authors should identify weak points of the review and areas for future work.

9. *Conclusion:* The conclusion should relate back to the purpose and summarize the major findings of the literature review and identify the contributions to knowledge made.

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

Crowdsourcing

Crowdsourcing is a technique for leveraging a group of people (the crowd) to solve complex tasks. In crowdsourcing, there are two types of users: requesters and members of the crowd (Bigham et al., 2015). Requesters are the people or organizations who define a problem or task, and aggregate the partial solutions produced by the crowd. Crowd members are people who contribute. Crowdsourcing is a special type of co-creation, a practice where developers and stakeholders collaborate to create a product or service (Pater, 2009). However, unlike Pater (2009) we do not limit crowdsourcing to a scenario where anyone can join the crowd, but also include the case where crowd members are selected based on participation criteria, such as, for their expertise or collaboration history.

Types of crowdsourcing
Crowdsourcing systems differ in terms of the incentives of requesters and crowd members, the complexity of the tasks, the amount of time crowd members spend on tasks, the level of collaboration between crowd members, and in terms of whether the work is done as part of "standard" work or not. Bigham, Bernstein, and Adar, (2015) distinguish between three types of crowdsourcing: directed crowdsourcing, collaborative crowdsourcing, and passive crowdsourcing.

1. In directed crowdsourcing, a single requester recruits the members of the crowd to pursue a specific goal. In this type of crowdsourcing, the members of the crowd generally act independently. A good example of directed crowdsourcing is Amazon's Mechanical Turk platform (mturk.com), in which workers get paid for performing specified tasks for the requester. In directed crowdsourcing, large tasks are often decomposed into so-called microtasks.

2. In collaborative crowdsourcing, the crowd self-determines their organization and work. In this type of crowdsourcing, members of the crowd are usually intrinsically motivated to participate, that is, they share in interest in accomplishing a joint task such as the creation of an online encyclopedia as in the case of Wikipedia, or identifying features on satellite images such as shapes that may indicate the location of a tomb (Lin et al., 2014).

3. In passive crowdsourcing, the crowd produces a useful outcome as part of their regular behaviour. Instead of directing the activity of the crowd, the requester is simply collecting traces of the crowd's behaviour and drawing inferences from them. An example of passive crowdsourcing is tracking messages on Twitter to predict a political outcome (iHub Research, 2013).

An interesting hybrid between directed crowdsourcing and collaborative crowdsourcing is leader-driven crowdsourcing. In this type of crowdsourcing, a leader maintains a high-level vision of the task and directs other crowd members (contributors) to make specific contributions towards this task. An example of leader-driven crowdsourcing is the collaborative writing system called Ensemble (Kim et al., 2014). We will build on the concept of leader-driven crowdsourcing in the proposed design below.

Benefits of crowdsourcing
Crowdsourcing is beneficial for a number of reasons, including:

1. *Time:* By distributing a task across a large group, crowdsourcing can reduce the time it takes to complete the task, given a clear division of the task into subtasks (Brown et al., 2014).

2. *Validation criteria:* Lacking a pre-existing reference for what constitutes an anomaly in the new domain, consensus can be used as a training mechanism for the crowd (Lin et al., 2014).

3. *Diversity:* A crowd can provide access to a diversity of perspectives (André et al., 2014).

4. *Domain knowledge required:* When appropriately structured, complex problems can be solved by crowds with little to no pre-existing domain knowledge (Bigham et al., 2015).

5. *Scale:* When a task is distributed among the members of a crowd, much larger tasks can be addressed such as large-scale surveys of datasets (Lin et al., 2014).

Other work on crowdsourcing literature reviews
The application of crowdsourcing to exploring new domains has not been widely studied yet. Most applications are in the medical domain, for example finding papers that mention certain diseases or drugs (Good et al., 2015) or searching for treatments (Elliot et al., 2014), and in education, for example learning new concepts (Luther et al., 2015). Although most of the early work on crowdsourcing has focused on datasets in domains that

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

most users are familiar with, such as images or travel advice, recent research has developed techniques that can deal with more complex qualitative datasets in unfamiliar domains, such as synthesizing textual data that requires domain-specific knowledge (André et al., 2014).

A search on Google Scholar for combinations of the keywords "crowdsourcing" and "literature review", "new domains", or "unfamiliar domains" only found two examples of crowdsourcing used to conduct a literature review, both in the medical domain. In the first, Brown and Allison (2014) describe a process for evaluating the literature that involves decomposing a research question of interest into microtasks that can be distributed to members of the crowd. In the second, Elliot, Thomas, and Owens (2014) describe an ongoing initiative for crowdsourced screening of citations (Embase, 2016).

One lesson from Brown and Allison (2014) is that quality checks are essential not only to guarantee the validity of the results. Quality checks are also required to demonstrate the competence of the members of the crowd to conduct a literature review. Such competence can be demonstrated through "pre-flight" qualification tests that are administered as an entry criterion, before allowing workers to participate in the crowd. A second lesson is that it is important to decide on the scope of the literature review to ensure that the output of the literature review only includes sources relevant to the question.

Design for Crowdsourcing Literature Reviews

In this section, we describe a design for a leader-driven crowdsourcing platform that can be used to collect evidence and synthesize it into insights in a new domain. First, we identify the design principles that guide the design of the platform. Then, we present a conceptual model for crowdsourcing literature reviews. It includes both the structure of the artefacts produced by the crowd and the roles and responsibilities of the members of the crowd.

Design principles
The design of the proposed crowdsourcing platform builds on the lessons learned from the two manually conducted literature reviews in new domains and on recent advances in crowdsourcing. These lessons lead to seven design principles:

1. *Scoping and synthesis:* put a leader in charge to decide on which questions should be examined (scoping) and to synthesize the answers into new insights.

2. *Chunking:* partition the literature review task into focused microtasks that can be executed without having to consider the literature review as a whole.

3. *Diversity:* crowdsourcing benefits from having a diverse membership with different perspectives. Initially, it is assumed that crowd members cannot self-select to participate. The model, thereby, corresponds to the club of experts model of co-creation (Pater, 2009).

4. *Scaffolding:* embed expertise into the design of the tools to magnify worker efforts.

5. *Incremental points of reference:* show answers from other participants.

6. *Consensus building:* create a consensus among the crowd members through commenting, voting, and tagging.

7. *Incentives:* build on the complementary motivations of leaders (to receive feedback) and contributors (to be recognized for their expertise).

Table 2 provides known uses in the crowdsourcing literature for each design principle.

Conceptual model
The design of the crowdsourcing platform proposed here draws on previous work on collaborative writing systems (Kim et al., 2014) and crowd-based clustering of documents (André et al., 2014). In a leader-driven crowdsourcing approach to collaborative writing (Kim et al., 2014), there are two types of participants: leaders and contributors. Leaders constrain and specify the nature of the contributions: the lead author of a literature review sets the scope of the literature review and guides the synthesis. Other crowd members (contributors) are recruited to focus on specific writing tasks.

As shown in Figure 1, we conceptualize a literature review as a story or narrative. Each narrative consists of a series of chunks that we call scenes (Kim et al., 2014). Each scene is anchored around a writing goal (such as providing an overview of the literature review, defining key features of the topic of the review, identifying examples illustrating the topic, or identifying gaps in the literature). Each writing goal is associated with a prompt that helps focus the work of the contributors. Answers to prompts are collected in drafts. Drafts can be commented and voted on, as well as categorized. It

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

Table 2. Known uses of the design principles in the crowdsourcing literature

Design Principle	Known Use
Scoping and synthesis	In the Ensemble system, leaders decide what contributors should write by creating writing prompts and select drafts from the submitted drafts or write their own (Kim et al., 2014).
Chunking	In the Ensemble system, the writing tasks are distributed across different roles: leaders, moderators, and contributors (Kim et al., 2014).
	Lin and colleagues (2014) split satellite images into tiles and asked crowd members to tag individual tiles.
Diversity	In the Ensemble system, leaders reported that they found the perspectives of others beneficial for writing their narratives (Kim et al., 2014).
Scaffolding	Contributors write drafts for scenes in response to prompts created by leaders. When contributors write a draft, they can see all the drafts from other contributors for the same scene. This visibility provides context for their task (Kim et al., 2014).
	Scaffolding gives crowd members the right information, at the right location, and at the right time, to help them accomplish a given activity (Owens, 2013).
Incremental points of reference	In another study conducted by Kim and colleagues (2015), contributors are asked to create questions about social media posts. For each post, they are shown sample questions created by other contributors.
	In Lin et al. (2014), crowd members are shown all tags other crowd members have assigned to a geographical location. This visibility allows crowd members to compare their own decisions against those of other peer participants.
Consensus building	In André et al. (2014), crowd members iteratively categorize text fragments. When a crowd member is asked to categorize a fragment, they see how other fragments have been categorized. They can then decide to put the new fragment into an existing category or create a new one.
	In Kim et al. (2014), consensus is built by voting: leaders and contributors can vote on the drafts created by others.
Incentives	Leaders and collaborators have complementary motivations: leaders want to receive feedback on their narratives, whereas collaborators view their expertise as valuable input to leaders (Kim et al., 2014).

is up to the leader to choose the best draft for each scene, so as to produce a final version of the narrative.

Cognitive science research on writing has identified that the writing process can be viewed as a series of rhetorical problems (Flowers & Hayes, 1981). For each narrative, there is a top-level rhetorical problem (which includes the constraints given to the writers, and the goals the writers create for themselves), which is then decomposed into subproblems that drive the creation of the narrative. For example, if the lead author of a literature review needs input on examples illustrating the topic of the review, they can ask the contributors for specific contributions with a prompt. In this way, the lead author can maintain a high-level vision of the literature review, while providing contributors with enough context of the overall flow of the literature review and direction towards specific tasks to complete.

Table 3 lists the roles of the participants in a crowdsourced literature review process and their responsibilities. Note that, although terms such as scene, prompt, and draft are still generic, we expect to identify catalogs of scenes and prompts specific to the creation of literature reviews once an initial prototype of the proposed platform has been developed and can be subjected to

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

Table 3. Roles and responsibilities in the crowdsourced literature review process.

Role	Narratives	Scenes	Drafts
Leader	Creates narratives	Creates scenes	Creates drafts
		Creates prompts in scenes	Synthesizes drafts into final draft for a narrative
		Determines sequence of scenes in a narrative	Comments on drafts
			Votes on drafts
			Categorizes drafts
Contributor			Creates drafts
			Comments on drafts
			Votes on drafts
			Categorizes drafts

systematic user testing. For example, it is already apparent from the experience with the two manually created literature reviews that the platform will need to support different types of drafts.

In one case, a prompt may ask contributors to produce alternatives to the leader's draft. For example, the leader could ask contributors for a definition of software lineage for malware. In this case the drafts are strictly alternative versions of a scene. In another case, a prompt could ask for a list of instances that together form the answer to the question. For example, a leader might ask for examples of code reuse attacks and for contributors to categorize them (André et al., 2014). As contributors collect and categorize the examples, they produce a taxonomy of code reuse attacks that could serve as a basis for further exploration. In this case, all or a subset of the drafts should be included in the review.

Conclusion

In this article, we proposed the design of a platform for crowdsourcing literature reviews in new domains. In particular, our focus was on creating systematic narrative literature reviews. Benefits expected from crowdsourcing literature reviews include:

1. Reducing the time it takes to complete a review

2. Being able to rely on emergent validation criteria given that a new domain lacks a pre-existing reference
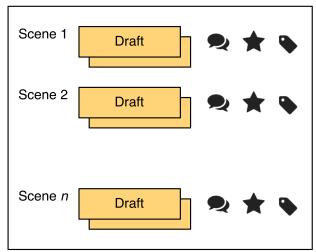
Narrative

Figure 1. Conceptualization of a literature review as a narrative, the components of a literature review (scenes and drafts), and the actions that can be performed on each of the components (commenting, voting, and categorization)

Crowdsourcing Literature Reviews in New Domains
Michael Weiss

for what constitute anomalies that may indicate gaps in current knowledge

3. Leveraging the diversity of perspectives of crowd members

4. Limiting the level of specific domain knowledge required to create a literature review in a new domain

Challenges for crowdsourcing literature reviews that we foresee include:

1. How to encourage participation (what kind of incentives need to be provided)

2. How to ensure the quality of the reviews produced (what aspects of the crowdsourcing process should be instrumented)

3. How to further support the synthesis stage of the review (what role can advanced techniques such as visualization and text mining techniques)

A prototype of the platform is currently being implemented by a team of developers at VENUS Cybersecurity Corporation (venuscyber.com). Systematic user testing of the platform and resulting extensions to the platform are left for future work.

Acknowledgements

About the Authors

Michael Weiss holds a faculty appointment in the Department of Systems and Computer Engineering at Carleton University in Ottawa, Canada, and is a member of the Technology Innovation Management program. His research interests include open source, ecosystems, mashups, patterns, and social network analysis. Michael has published on the evolution of open source business, mashups, platforms, and technology entrepreneurship.

References

André, P., Kittur, A., & Dow, S. P. 2014. Crowd Synthesis: Extracting Categories and Clusters from Complex Data. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW):* 989–998. New York: ACM.
http://dx.doi.org/10.1145/2531602.2531653

Attenberg, J., Ipeirotis, P. G., & Provost, F. J. 2011. Beat the Machine: Challenging Workers to Find the Unknown Unknowns. *Human Computation: Papers from the 2011 AAAI Workshop (WS-11-11):* 2–7.

Baumeister, R. F., & Leary, M. R. 1997. Writing Narrative Literature Reviews. *Review of General Psychology,* 1(3): 311–320.
http://dx.doi.org/10.1037/1089-2680.1.3.311

Berry, D. 1995. The Importance of Ignorance in Requirements Engineering. *Journal of Systems and Software,* 28(1): 179–184.
http://dx.doi.org/10.1016/0164-1212(94)00054-Q

Bigham, J. P., Bernstein, M. S., & Adar, E. 2015. Human-Computer Interaction and Collective Intelligence. In T. Malone & M. Bernstein (Eds.). H*andbook of Collective Intelligence:* 57–83, MIT Press.
http://cci.mit.edu/CIchapterlinks.html

Brown, A. W., & Allison, D. B. 2014. Using Crowdsourcing to Evaluate Published Scientific Literature: Methods and Example. *PloS ONE,* 9(7): e100647.
http://dx.doi.org/10.1371/journal.pone.0100647

Elliot, J., Thomas, J., & Owens, N., Dooley, G., Riis, J., Wallace, B., Thomas, J., Noel-Storr, A., Rada, G., Struthers, C., Howe, T., MacLehose, H., Brandt, L., Kunnamo, I., & Mavergames, C. 2014. Editorial: #CochraneTech: Technology and the Future of Systematic Reviews. *Cochrane Library,* 2014(9).
http://dx.doi.org/10.1002/14651858.ED000091

Embase. 2016. Embase Screening. *The Cochrane Collaboration.* Accessed on February 7, 2016:
http://screening.metaxis.com/EMBASE/login.php

Flower, L., & Hayes, J. R. 1981. A Cognitive Process Theory of Writing. *College Composition and Communication,* 32(4): 365–387.
http://www.jstor.org/stable/356600

Grant, M. J., & Booth, A. 2009. A Typology of Reviews: An Analysis of 14 Review Types and Associated Methodologies. *Health Information & Libraries Journal,* 26(2): 91–108.
http://dx.doi.org/10.1111/j.1471-1842.2009.00848.x

Crowdsourcing Literature Reviews in New Domains

Michael Weiss

Green, B. N., Johnson, C. D., & Adams, A. 2006. Writing Narrative Literature Reviews for Peer-Reviewed Journals: Secrets of the Trade. *Journal of Chiropractic Medicine,* 5(3): 101–117. http://dx.doi.org/10.1016%2FS0899-3467(07)60142-6

iHub Research. 2013. *Viability, Verification, Validity: 3Vs of Crowdsourcing.* Nairobi, Kenya: iHub Research. http://www.ihub.co.ke/ihubresearch/jb_VsReportpdf2013-8-29-07-38-56.pdf

Kim, J., Cheng, J., & Bernstein, M. S. 2014. Ensemble: Exploring Complementary Strengths of Leaders and Crowds in Creative Collaboration. *Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW):* 745–755. New York: ACM. http://dx.doi.org/10.1145/2531602.2531638

Lin, A. Y. M., Huynh, A., Lanckriet, G., & Barrington, L. 2014. Crowdsourcing the Unknown: The Satellite Search for Genghis Khan. *PloS ONE,* 9(12): e114046. http://dx.doi.org/10.1371/journal.pone.0114046

Luther, K., Hahn, N., Dow, S. P., & Kittur, A. 2015. Crowdlines: Supporting Synthesis of Diverse Information Sources through Crowdsourced Outlines. *Third AAAI Conference on Human Computation and Crowdsourcing,* November 8-11, San Diego, CA. Palo Alto, CA: The AAAI Press.

Mehrotra, G. and Berry, D.M. 2012. *Role of Domain Ignorance in Software Development: Software Development Tasks Benefiting from Domain Ignorance.* Technical Report, Cheriton School of Computer Science, University of Waterloo, Canada. https://cs.uwaterloo.ca/~dberry/FTP_SITE/tech.reports/Mehrotra BerrySurvey.pdf

Owens, T. 2013. Digital Cultural Heritage and the Crowd. Curator: *The Museum Journal,* 56(1): 121–130. http://dx.doi.org/10.1111/cura.12012

Paré, G., Trudel, M. C., Jaana, M., & Kitsiou, S. 2015. Synthesizing Information Systems Knowledge: A Typology of Literature Reviews. I*nformation & Management,* 52(2): 183–199. http://dx.doi.org/10.1016/j.im.2014.08.008

Pater, M. 2009. Co-Creation's 5 Guiding Principles. *Fronteer Strategy:* April 9, 2009. Accessed February 1, 2016: http://fronteerstrategy.blogspot.ca/2009/04/co-creations-5-guiding-principles-or.html