Large-Scale Collaborative Innovation: Challenges, Visions and Approaches

Pao Siangliulue,[1] Joel Chan,[2] Kenneth C. Arnold,[1]
Bernd Huber,[1] Steven P. Dow,[2] Krzysztof Z. Gajos[1]
[1]Harvard School of Engineering and Applied Sciences, Cambridge, MA USA,
{paopow, kcarnold, bhb, kgajos}@seas.harvard.edu
[2]Carnegie Mellon University, Pittsburgh, PA USA, {joelchuc, spdow}@cs.cmu.edu

Abstract

Emerging online innovation platforms have enabled large groups of people to collaborate and generate ideas together in ways that were not possible before. However, these platforms also introduce new challenges in helping their members to generate diverse and high quality ideas. In this paper, we enumerate collaboration challenges in crowd innovation: finding inspiration for contributors from a large number of ideas, motivating crowd to contribute to improve group understanding of the problem and solution space, and coordinating collective effort to reduce redundancy and increase quality and breadth of generated ideas. We discuss possible solutions to this problem and present our recent work that addresses some of these challenges using techniques from human computation and machine learning.

Introduction

Innovation is a product of collective effort. In contrast to a popular belief, innovation does not come from a few "lone geniuses", but arises as combinations of ideas from a diverse set of viewpoints and experiences of many collaborating individuals (Hong and Page 2004; Krause et al. 2011; Pentland 2014). In order to make big progress such as creating new technology, making groundbreaking scientific discovery or solving challenging social problems, our society needs to support idea exchanges between people of various backgrounds and perspectives.

Online innovation platforms enable idea exchanges on the scale that was not possible before. These platforms allow people with diverse knowledge, experience and viewpoints to share their ideas and get inspired by ideas of others to solve problems together without the limitation of physical location or fields of expertise. Existing communities from various domains, like OpenIDEO (openideo.com) where people propose solutions to social problems, have already attracted large numbers of users, many of whom contribute ideas. The promise of these online communities is that participants will benefit from exposure to others' ideas and, thus inspired, will generate better and more diverse solutions than they would have otherwise.

However, large-scale online innovation platforms also pose new challenges that prevent the crowd to effectively generate ideas together. The large number of ideas makes it difficult for a community member to find non-redundant inspiring ideas. Community members have to sift through a lot of ideas with no effective way to discover the ideas they find inspiring.

Moreover, there is no mechanism to coordinate effort among individual contributors. Crowd innovation is an extreme instance of loosely coupled teamwork: the individual contributors work mostly alone, but in order to contribute productively they should have an awareness of the overall shape of the solution space explored so far. As no effective methods exist yet to provide crowd innovators with such awareness, the individual contributors either ignore the work of others or are overwhelmed by a sea of mostly mundane and redundant ideas, which form the majority of the idea corpus (Ward et al. 2002). While current research on crowd innovation builds largely on cognitive, social and emotional mechanisms specific to creativity, an urgent need exists to bring teamwork-related insights to crowd innovation.

Instead of each participant starting de novo and rediscovering the same few obvious ideas as everybody else, ideal systems would provide individual contributors with just the right awareness of what others have done and with just the right inspiration to help them contribute to promising, but still under-explored parts of the solution space. An ideal system would limit redundancy and instead appropriately coordinate the individual efforts to ensure both breadth and depth in the exploration of the solution space.

In this paper, we enumerate the challenges faced by creative crowd, share our recent and current work in addressing some of these challenges, and discuss future solutions that pull insights and methods from HCI, AI and social science.

Challenges for collaborative idea generation

Finding inspiring ideas

Online innovation platforms host large corpus of ideas that act as a medium for contributors to exchange ideas by seeing ideas of each other. While the size of large corpus of ideas presents more possibility of idea exchanges to produce new ideas, not all idea exchanges yield high quality or novel ideas. Instead, a large idea corpus forces ideators to sift through a large number of ideas before finding a variety of inspiring ones. To effectively generate diverse and high quality ideas, crowd innovation platforms should encourage idea exchanges that are most likely to improve the overall quality and diversity of the community. This challenge is similar to information sharing problem when resources are limited.

Exposing contributors to a large number of ideas can both benefit or harm their creative output. On the one hand, seeing a lot of uninspiring redundant ideas can fixate contributors to unfruitful ideation (Jansson and Smith 1991; Kohn and Smith 2011; Smith, Ward, and Schumacher 1993). On the other hand, if contributors see a set of ideas that are both diverse and of high quality, they are more likely to generate creative and diverse ideas (Marsh, Landau, and Hicks 1996; Siangliulue et al. 2015; Nijstad, Stroebe, and Lodewijkx 2002). Hence, a crowd innovation support system should carefully select ideas to avoid fixating contributors on bad ideas and promote good ideas.

To maximize the gain from a large corpus of ideas, the system should be able to help people find ideas that are diverse and inspiring. In other words, we need a synthesis that helps the system understand the whole knowledge space. One naive approach is to use human curators. However, previous innovations has shown that this is actually hard to do when the number of ideas grows large. During Google's Project 10^{100} \$10M innovation challenge, for example, the company had to recruit 3,000 of their employees to prune 150,000 ideas received from the crowd, pushing the project nine months behind schedule. In another setting, the change.gov web site was shut down prematurely because the staff could not process the incoming ideas fast enough.

In our recent work, we develop a method that helps synthesize the solution space using techniques in machine learning and human computation. By cleverly outsourcing microtask to external crowd, we can derive a synthesis of solution space that the system can use to help identify a set of inspiring and diverse ideas from others. When coupling this method with existing techniques that can scalably evaluate quality of ideas (Salganik and Levy 2012; Yu and Nickerson 2011; Tanaka, Sakamoto, and Kusumi 2011; Xu and Bailey 2012), we can use this synthesized solution space to help each contributor find a variety of high quality ideas for inspiration.

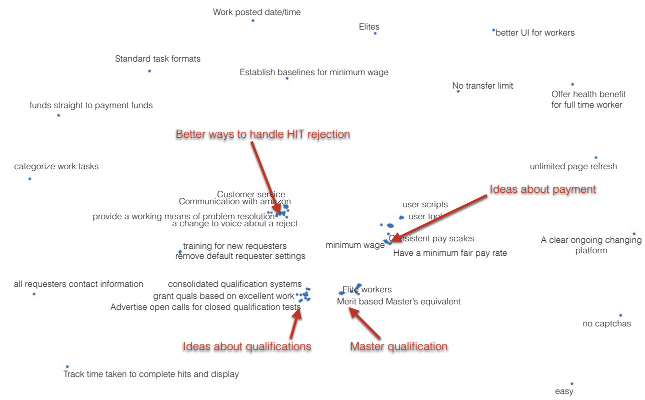

Figure 1: An example of idea map of for a new micro-task market platform generated from contributor's similarity between ideas using our algorithm. Similar ideas are placed close to each other while dissimilar ideas are placed far from each other. The red arrows point to emerged clusters of ideas around different topics.

Motivating contribution to improve group understanding of the solution space

A synthesis of solution space can guide individual idea exploration through aptly selected inspirations and give the crowd a big picture of various promising ideas to explore. This synthesis is usually generated by humans because automated topic analysis tools often fail to capture complex nuances of high-dimensional data that are meaningful to humans (Chang et al. 2009). This is a potential burden on the large community where the number of human judgement required grows with the number of ideas. Asking community members to do micro-tasks that provide information to create a synthesis (Chilton et al. 2013; André, Kittur, and Dow 2014; Siangliulue et al. 2015) can detract from their work, taking up time for idea generations. Moreover, most of these micro-tasks are tedious activities that people do not want to spend time on. While outsourcing the work can be a solution in some cases, some tasks require certain expertise or credentials from the human judges and, thus, cannot be outsourced to popular micro-task market.

We are exploring "organic" human computation approach, where we get all the potentially tedious human input as a byproduct of the activities people would like to do, and we use this organic input to manage information sharing in crowd innovation. For example, we observed that people usually arrange their ideas spatially to make senses of the problem and solution space during idea generation. We have prototyped an approach to harvest this computational signal in our idea generation system called IdeaHound 2. IdeaHound's user interface makes it convenient for users to spatially organize ideas on a large canvas. The system combines the spatial relationships created by all contributors into an aggregated low-dimensional embedding of all ideas, where distances between ideas correspond to their semantic differences. This is a coarse representation, but sufficient for IdeaHound to offer diverse inspirational examples to subsequent contributors (which in turn, inspires more diverse

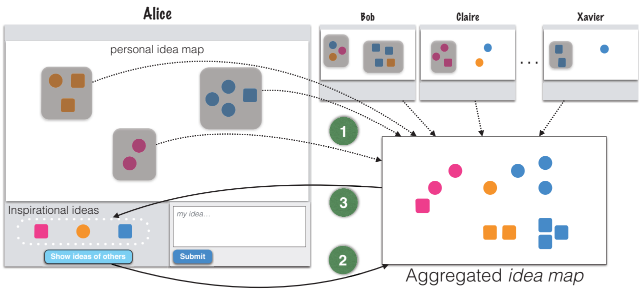

Figure 2: Data collection flow in IDEAHOUND: 1) IDEAHOUND collects information about semantic relationships between ideas from all users and use them to synthesize a global idea map with semantic relationships betweens all ideas. 2) A user can request to see a set of ideas by others for inspiration. 3) IDEAHOUND consults the idea map and selects a set of diverse ideas for the user.

ideas (Siangliulue et al. 2015)).

This preliminary work demonstrates that organic human computation can overcome resource limitations, but also leaves open key questions to explore. For example, what other synthesis information—such as labeled attributes, analogies, or constraints within the solution space—can we harvest organically during ideation? How can we amplify human judgements with judicious use of machine learning? We plan to further explore different solutions for these challenges through iterative design, grounded in careful evaluation of the human computation tasks' impact on ideators' cognitive resources and motivation.

Coordinating crowd effort

Beyond increasing individual creativity, successful collaborative ideation involves coordinated exploration of the solution space. The system should help contributors avoid duplicating the efforts of others or going down solution paths known to be dead ends. It should support and encourage parallel explorations of promising, but non-obvious solution paths. How can we influence the individual creative efforts such that they contribute to a coordinated and systematic exploration of the solution space? What is the most effective method for directing an individual's creative effort toward a particularly promising part of the solution space? Could we induce "productive" design fixation by carefully curating the inspirational examples we present to the contributor? Or should we instead explicitly suggest an attribute of the solution space to explore? Or should we present the contributor with a glanceable representation of the evolving solution space in hopes that ideators will gravitate towards promising, but relatively unexplored solution paths?

One possible approach is to automatically identify the parts of the solution space that have not been explored as much as others and direct more people to generate ideas in those unexplored regions of the space. An algorithm can help the community decide whether the contributors should focus on generating ideas related to particular topics or exploring new ones. This mechanism can be integrated into a system as a form of an explicit todo list (Zhang et al. 2012) or as a subtle nudge through judicious selection of inspirational examples. The most effective mechanisms will quite likely depend on the stage of ideation process (early exploration or later refinement), and the ideator's cognitive state. To effectively intervene, we need to understand how the form of a nudge (e.g., showing examples or attributes)—in combination with key contextual factors (e.g., stage of ideation process, cognitive state of individual, etc.)—influences key ideation process and outcome measures (e.g., diversity or depth of ideas generated) and interacts with ideators' limited attentional resources and motivation to participate. We need to conduct more empirical studies to understand how to optimally help people in different states. The findings can then guide us on designing systems that empower motivated individuals and direct their participation towards activities that will improve the overall quality and diversity of ideas.

Another mechanism is to provide an overview of the community's idea generation along with a way for contributors to communicate so that contributors do not interfere with each other. The overview of the community's idea exploration can give people a sense of promising parts of solution space that are underexplored. A contributor can state the parts they want to generate ideas for and the system will update this information to other contributors to avoid duplicate work. These indirect communications happen in a lightweight fashion with an overview synthesis that keeps everyone informed of the progress and current actions of the group.

So far we have discussed idea generation for singular problems where the proposed ideas address the whole problem. However, real-world problems usually involves many interdependent sub-problems. How can we utilize the strength of crowd so that they generate multiple good solutions for each sub-problem in parallel? The challenges are 1) dividing the problem into manageable sub-problems and 2) integrating appropriate solutions to sub-problems to make up complete solutions.

The first challenge—dividing problems into smaller ones—lies in finding the right division of problems. One can divide a problem into parts before generating ideas for each based on the existing knowledge of the problem domain. However, the resulting division can be arbitrary and end up introducing complexity to the process (Alexander 1964). For example, in generating ideas for a web application, it is possible to divide sub-problems by implementation (e.g. aesthetic, input methods) instead of a better division by functionality (e.g. login experience, shopping cart). An ideal sub-problem division should produce a set of sub-problems that are small enough for an individual and are relatively independent of each other to minimize the communication cost between sub-components. We are exploring an approach that detects an appropriate division of problem from the crowd's initial idea generation. At the beginning of idea generation, people usually generate incomplete ideas (Goel 1995). These incomplete ideas can be complemented by other incomplete ones. The complimenting ideas show natural sub-components that make up the whole solutions.

The second challenge—integrating appropriate solutions—lies in communication between people working on different sub-problems. The system should ensure that, the generated ideas for different components are complement to each other. This process would involve communication between sub-problems that depend on each other. For example, in a challenge to find solutions for water transportation in a developing country, a contributor who designs a water container should be aware of the ideas for transportation mechanisms generated by others. However, communication is expensive and there are different mechanisms to exchange information.

While some people find it helpful to directly communicate with others to work on complement solutions in parallel, others just need to know the scope of available solutions for other sub-problems that would allow them to generate ideas that complement other components. An ideal system should be able to identify appropriate means of communication and adapt accordingly. Direct communication requires more effort, time and attention but can help team establish common understanding quicker. A system can help contributors form small teams based on their experience, interests, and current components they are working on. Indirect communication is a light-weighted channel where the system infers relevant information to send to people working on related sub-problems. We are investigating mechanisms of information exchanges in a loosely coupled fashion in a large scale idea generation context. One possible approach is to extract shared dependent attributes of concern between components and present them to related individuals. A contributor will only see only a range of attributes that are related to the sub-problem they are solving instead of having to navigate to all available attributes that do not affect their solution.

Conclusion

Crowd innovation is an emerging paradigm that leverages quantity and diversity of contributors to generate ideas in a way that was not possible before. The rising challenges of crowd innovation can be addressed with insights and methods from various fields including HCI, AI and social science. This paper enumerates some of these challenges, namely, helping people find inspiring ideas from a big collection, motivating people to contribute to improve collective understanding of the solution space, and coordinating effort of contributors. We share our recent and current work that address these challenges. However, more importantly, we would like this paper to be a starting point for discussion from experts in various fields on how to address this problem together.

Acknowledgements

This work was funded in part by a Sloan Research Fellowship, gifts from Google and Adobe, a Nuance Foundation grant and awards from the National Science Foundation (IIS-1208382, IIS-1217096, and IIS-1122206).

References

Alexander, C. 1964. *Notes on the Synthesis of Form*, volume 5. Harvard University Press.

André, P.; Kittur, A.; and Dow, S. P. 2014. Crowd synthesis: Extracting categories and clusters from complex data. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work*, CSCW '14, 989–998. New York, NY, USA: ACM.

Chang, J.; Gerrish, S.; Wang, C.; Boyd-Graber, J. L.; and Blei, D. M. 2009. Reading tea leaves: How humans interpret topic models. In *Advances in neural information processing systems*, 288–296.

Chilton, L. B.; Little, G.; Edge, D.; Weld, D. S.; and Landay, J. A. 2013. Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the 2013 ACM annual conference on Human factors in computing systems*, 1999–2008. ACM.

Goel, V. 1995. *Sketches of thought*. MIt Press.

Hong, L., and Page, S. E. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America* 101(46):16385–16389.

Jansson, D. G., and Smith, S. M. 1991. Design fixation. *Design Studies* 12(1):3–11.

Kohn, N. W., and Smith, S. M. 2011. Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology* 25(3):359–371.

Krause, S.; James, R.; Faria, J. J.; Ruxton, G. D.; and Krause, J. 2011. Swarm intelligence in humans: diversity can trump ability. *Animal Behaviour* 81(5):941–948.

Marsh, R. L.; Landau, J. D.; and Hicks, J. L. 1996. How examples may (and may not) constrain creativity. *Memory & Cognition* 24(5):669–680.

Nijstad, B. A.; Stroebe, W.; and Lodewijkx, H. F. 2002. Cognitive stimulation and interference in groups: Exposure effects in an idea generation task. *Journal of experimental social psychology* 38(6):535–544.

Pentland, A. 2014. *Social Physics: How Good Ideas Spread-The Lessons from a New Science*. Penguin.

Salganik, M. J., and Levy, K. E. 2012. Wiki surveys: Open and quantifiable social data collection. *arXiv preprint arXiv:1202.0500*.

Siangliulue, P.; Arnold, K. C.; Gajos, K. Z.; and Dow, S. P. 2015. Toward collaborative ideation at scale—leveraging ideas from others to generate more creative and diverse ideas. In *Proceedings of CSCW'15*.

Smith, S. M.; Ward, T. B.; and Schumacher, J. S. 1993. Constraining effects of examples in a creative generation task. *Memory & Cognition* 21(6):837–845.

Tanaka, Y.; Sakamoto, Y.; and Kusumi, T. 2011. conceptual combination versus critical combination: devising creative solutions using the sequential application of crowds. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*.

Ward, T. B.; Patterson, M. J.; Sifonis, C. M.; Dodds, R. A.; and Saunders, K. N. 2002. The role of graded category structure in imaginative thought. *Memory & Cognition* 30(2):199–216.

Xu, A., and Bailey, B. 2012. A reference-based scoring model for increasing the findability of promising ideas in innovation pipelines. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, 1183–1186. ACM.

Yu, L., and Nickerson, J. 2011. Cooks or cobblers?: crowd creativity through combination. *CHI '11: Proceedings of the 2011 annual conference on Human factors in computing systems*.

Zhang, H.; Law, E.; Miller, R.; Gajos, K.; Parkes, D.; and Horvitz, E. 2012. Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 217–226. ACM.