

# Supporting Collaborative Innovation at Scale

Pao Siangliulue

Harvard School of Engineering and Applied Sciences

Cambridge, MA USA

paopow@seas.harvard.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 finding inspiration from a large number of ideas, and coordinating the collective effort. In my dissertation, I address the challenges of large scale idea generation platforms by developing methods and systems for helping people make effective use of each other's ideas, and for orchestrating collective effort to reduce redundancy and increase the quality and breadth of generated ideas.

## Author Keywords

Creativity support system, Ideation, Collective intelligence.

## ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

## INTRODUCTION

Innovation is a product of collaborative effort. When people create together, each person brings unique knowledge, experiences and points of view, which can be combined to achieve amazing things such as creating new technology, making new scientific discovery and solving social problems. Online innovation platforms allow their contributors to seek inspirations from a large collection of ideas by other members and to generate novel ideas that none of the contributors could have come up with alone. OpenIDEO.com and Quirky are examples of such platforms, whose members have shown the potential of collective idea generation in solving challenging problems.

However, large-scale collaborative idea generation also brings new challenges. The large number of ideas makes it difficult for a contributor to find non-redundant inspiring ideas. Contributors have to look 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

contributors. Instead of working together, contributors generate ideas alone with no guidance on how to best contribute. They end up producing many instances of redundant common ideas while leaving many parts of the solution space unexplored [20]. To help a large community of contributors make full use of its size and diversity, we need new methods and tools that enable contributors to make effective use of each other's ideas, and work effectively together.

My dissertation focuses on developing intelligent systems that improve user experience and outcomes of large-scale collaborative idea generation by providing people with carefully selected inspiring ideas of others at appropriate time.

My work builds on studies of creativity from cognitive science. I have extended this body of work to further understand mechanisms that help people generate more creative and diverse ideas as they get inspired by ideas of others in an actual collaborative innovation platform. Building on these insights, I have developed automatic methods to provide ideators with inspirational ideas, leading to improved performance. Specifically, my work has thus far made the following contributions:

- I have demonstrated that presenting people with sets of *diverse* inspirational ideas results in people generating more diverse ideas than they would have if they were presented with a random selection of ideas [15]. To enable this intervention in real-world settings, I developed and evaluated an efficient, scalable and domain-independent method that can select a diverse set of ideas from a large set of ideas.
- I demonstrated the impact of timing of delivery of inspirational ideas on creative output. I showed that inspirational idea delivery mechanisms that attempt to deliver ideas when people are switching from generating ideas in one category to another can improve novelty and quantity of generated ideas [16].
- I am investigating mechanisms that select personalized set of inspirations to best inspire each contributor based on their solution path (ideas they have explored so far) and their cognitive state (focused exploration versus seeking new inspiration). Prior work suggests that providing a set of inspirations tailored to each individual might yield more creative ideas and improve the contributor's experience [12, 16].

I am now at the crucial point where I address technical challenges in building the system that embodies these findings.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

UIST'15 Adjunct, November 08–11, 2015, Charlotte, NC, USA.

ACM 978-1-4503-3780-9/15/11.

<http://dx.doi.org/10.1145/2815585.2815588>

I plan to demonstrate the viability of my vision in an innovation platform that is open to public, and to proceed in the following directions:

- To select a set of ideas that are inspiring to ideators, the system needs some knowledge about the quality of ideas and semantic relationship between ideas. In my prior work [15], the system acquired this knowledge from paid crowd workers, which is not always viable. I am building a system that uses a self-sustainable “organic” approach where the system subtly extracts information about the quality and the semantics of ideas from users’ natural interactions while they are generating ideas.
- In practice, idea generation spans a period of time. During this time, ideators might be incubating their ideas as they go about their other daily activities [3]. To support an ideation task that spans longer than one session, I plan to integrate the ideation task throughout the ideators’ day by sending prompts or ideas of others through mobile devices to remind or inspire them to come up with new ideas.
- One main problem in existing innovation platforms is that they elicit many redundant ideas that represent a small fraction of possible solutions while leaving some promising solutions unexplored. I plan to develop novel algorithms that help the community coordinate individuals’ efforts in parallel, reduce redundant contributions, and achieve a more thorough, diverse exploration of the solution space.

In the remainder of the paper, I first describe the theoretical findings of prior work and my own work. Then I propose how to integrate these findings into a real innovation platform.

## UNDERSTANDING HOW TO HELP PEOPLE GENERATE CREATIVE AND DIVERSE IDEAS

Idea exchanges on large innovation platforms often occur when contributors get inspired by each other’s ideas. But seeing ideas of others is not always helpful: depending on the choice of ideas and the timing of their delivery, the ideas that people see can either benefit or harm their creative output [12, 8, 6, 16]. An effective approach therefore should help a contributor find inspiring ideas from a large pool at an appropriate time.

### Selecting a diverse set of inspirations

Prior laboratory studies have demonstrated that presenting people with diverse inspirational examples help people generate more diverse ideas themselves [12]. However, prior to my work, there was no robust, scalable and domain-independent method to dynamically identify a diverse set of ideas. Existing crowdsourced approaches designed to organize large collections of items are designed mainly for batch processing and not for the settings where new ideas keep trickling in [1].

I developed a method powered by human computation and machine learning for incrementally constructing an abstract spatial “idea map” (Figure 1) from simple human’s judgement on similarity between ideas (“Is idea A more similar to B or C?”) [15]. I developed this method based on prior work on multidimensional scaling and active similarity learning techniques [18, 17] to embed ideas in an idea space with



**Figure 1. An idea map of birthday messages to a 50-year-old female firefighter. Similar ideas are placed close to each other and dissimilar ideas are kept far apart. See emergent clusters of ideas around different themes and sentiments.**

as few human queries as possible. From the idea map, a system can algorithmically extract diverse sets of examples by selecting a set of ideas that are situated far apart from one another on the map. In contrast to some existing methods which use text mining algorithms to infer similarity between ideas in text form [4], my approach is domain independent. I also demonstrated that people generate more diverse ideas when they saw a set of diverse ideas selected by the algorithm compared to when they saw randomly selected ideas or no ideas at all [15].

### Timing of inspiration delivery

The timing of example delivery, not just the choice of examples themselves, can impact creative outcomes [7]. Prior work on a theory of idea generation called SIAM (Search for Ideas in Associative Memory) [11] and the Prepared Mind theory of insight [14] suggested that new inspirations should be offered when people finish exploring ideas in one semantic category, as seeing ideas of others can direct them towards new parts of the solution space. Showing ideas of others while people are still exploring their current topics; however, can cut their train of thoughts short and reduce the quantity and quality of generated ideas [11].

In my recent work [16], I explored two timing mechanisms that attempt to offer ideas of others when people just finished exploring ideas in one category: (1) participants can request ideas on demand whenever they want, and (2) the system automatically infers when participants are stuck from their idle time and provides ideas at that moment. The results of evaluation of the two mechanisms show that people who requested examples themselves generated the most novel ideas and people who received ideas automatically when they were idle produced the largest quantity of ideas.

### Providing personalized inspirations for individuals

While providing a diverse set of ideas at appropriate time is a sensible strategy in general, it does not take into account the differences between individuals. Prior work suggests the prior experience and the ideas they have already explored can impact how much individuals can benefit from ideas of others [20, 11]. In this case, inspiration delivery could be personalized to a user’s current solution path. With an algorithm

that adaptively shows ideas of others that are appropriately related to a user’s current ideas, the system can either direct the person in new directions or suggest deep exploration within a category for highly creative ideas [11, 10, 13].

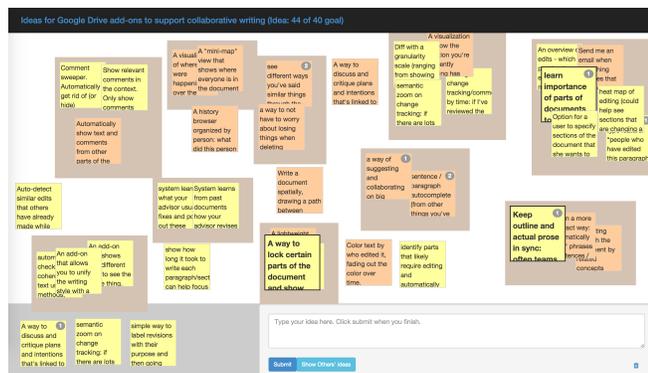
I am developing mechanisms for real-time semantic analysis of participants’ solution paths, and I plan to conduct experiments to test whether personalized inspiration could further help people benefit from inspirational examples. The three main technical challenges for this approach are (1) the system needs to be able to identify the semantic relationships between ideas dynamically so that it knows the users’ paths of exploration relative to other ideas, (2) the system should be able to infer the users’ cognitive states—whether the users are receptive to ideas of others that are similar or different from their own ideas—to select appropriate suggestions, and (3) for each cognitive state, the system should have a mechanism to decide what kinds of examples should be shown.

## SYSTEM SUPPORT FOR INNOVATION AT SCALE

### Organic interactions

In my initial work [15], the system learned about relationships between ideas from outsourced microwork. This approach is domain-independent and scalable, but it has two disadvantages. First, paid crowd workers are not intrinsically motivated to produce high quality work. Second and more importantly, many types of creative tasks require special knowledge and thus cannot be outsourced.

Instead, I am designing interactions that encourage ideators to generate useful information about ideas for the system while engaging in meaningful activities for their own idea generation. For example, people usually spatially organize ideas and group ideas together to get a better sense of possible solutions [9]. These grouping activities reveal users’ perceptions of similarity among ideas and can be used to generate an idea map without the need to resort to external crowds.



**Figure 2.** A prototype of an interface with organic interactions. The users can organize ideas on the whiteboards from which the system extracts information about relationships between ideas.

I am building an intelligent system that supports and leverages such “organic” interactions (Figure 2). The current prototype is a web application that allows users to create projects, invite people to contribute to their projects, submit ideas and request to see others’ ideas. The users can organize ideas

on their personal “whiteboards”, bookmark ideas that they want to easily access later and remove ideas that they no longer want to see from the board. These interactions give the system information about the quality of ideas (bookmarking, deleting from the work space) and how they are related to each other (spatial organization on the whiteboard) which the system uses to help select inspiring ideas when users request them.

Pilot user testing with small groups over several days showed that users naturally organize ideas into groups and bookmark ideas they consider interesting. I also learned that participants sometimes had other uses of grouping ideas other than organizing semantically similar ideas into clusters. Some created a separate group for their own ideas, while some had a pile for ideas that they do not want to organize yet. Such groups do not give meaningful information about semantic relation between ideas and introduce noise to the algorithm that selects ideas. I am working on a way to automatically identify which clusters are semantically meaningful.

Will such organic activities provide enough input for the system to reason reliably about the semantic relationships among ideas? To answer this question, I am designing a study to compare the idea maps generated organically, as ideators naturally go about their creative activities, to those generated using my previously validated method that relied on paid crowd workers performing explicit assessment tasks.

The pilot studies also revealed opportunities for other interaction designs. There are two in particular that I plan to explore. The first one is derived from my observation that participants sometimes submit rationales for their ideas on top of the idea itself. According to prior work, seeing the rationale is less likely to cause fixation than seeing the examples [21]. If the users can provide rationales or tag the interesting aspects of ideas, the system can intelligently provide these rationales as an alternative to showing the ideas themselves.

The second opportunity is based on participants’ comments that it took some time for them to reorient themselves to their workspace in the following sessions. The system can allow users to put labels or notes on groups of ideas. These labels can help inform the system whether the grouping is meaningful and on which dimension it is based in case there are several ways to group the available ideas. Such information might help the system provide better explanation on how it selects to people. The explanation can increase trust in the recommendation by the system [19] and help ideators pinpoint which part of the inspirational ideas to focus on.

### Making incubation more effective

Prior work in cognitive science suggests that idea generation is a long process that involves incubation periods where ideators take a break from actively thinking about the problems [3]. While people are not trying to come up with ideas, there is a solution searching process running in the background. To keep this process running, people need to be aware that they are returning to generate ideas later [5]. A system can help remind people of the task while they are not actively thinking about ideas by sending prompts or showing

inspiring ideas by others. However, the system also needs to be careful in choosing which reminders to send based on the context of ideators to avoid interfering with other tasks the users try to accomplish [2].

Mobile devices grant access to information about users' context such as location, activity and schedule. This information can help a system decide when and how to instigate interactions with the users. For example, knowing a user's schedule and location, a system can offer related inspirations or prompting her to record her ideas at a certain point in time. I plan to expand my prototype to include a mobile application that provides inspirational ideas of others and prompts users based on their contexts. This version of collective ideation platform can also better inform how people generate ideas in different situations for future creativity research as well.

### Coordinating collaborative exploration of idea space

As the number of contributors increases, the more crucial and difficult it becomes to coordinate the community efforts. Without coordination, people are likely to generate redundant ideas that are very similar to one another instead of exploring a broader range of ideas [20]. However, the community also wants to avoid too much communication overhead that can take time away from generating ideas.

I plan to explore mechanisms that reduce idea redundancy in large creative communities and widen the breadth of ideas generated. One possible mechanism is to automatically identify the parts of the idea 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 [22] or as a subtle nudge through judicious selection of inspirational examples. Another approach is to provide people with an overview of the community's idea generation along with a way for contributors to communicate so that they do not interfere with each other.

### CONCLUSION

Large ideation platforms promise a huge opportunity for innovation, but are limited by lack of appropriate methods and tools to help people discover inspiring ideas and coordinate their efforts. By synthesizing knowledge and approaches from cognitive science, human computation and machine learning, we can create an intelligent system that addresses the existing challenges and improves the experience of ideators and quality of generated ideas at scale.

**Acknowledgments** This work was funded in part by gifts from Google and Adobe.

### REFERENCES

1. Paul André, Aniket Kittur, and Steven P. Dow. 2014. Crowd Synthesis: Extracting Categories and Clusters from Complex Data. In *Proc. CSCW'14*. ACM, 989–998.
2. Brian P. Bailey and Shamsi T Iqbal. 2008. Understanding changes in mental workload during execution of goal-directed tasks and its

- application for interruption management. *ACM TOCHI* 14, 4 (2008), 21.
3. Benjamin Baird, Jonathan Smallwood, Michael D Mrazek, Julia WY Kam, Michael S Franklin, and Jonathan W Schooler. 2012. Inspired by distraction mind wandering facilitates creative incubation. *Psychological Science* (2012).
4. Joel Chan, Steven Dow, and Christian Schunn. 2014. Conceptual distance matters when building on others' ideas in crowd-collaborative innovation platforms. In *Proc. CSCW'14*. ACM, 141–144.
5. Jason Gallate, Cara Wong, Sophie Ellwood, RW Roring, and Allan Snyder. 2012. Creative people use nonconscious processes to their advantage. *Creativity Research Journal* 24, 2-3 (2012), 146–151.
6. Nicholas W Kohn and Steven M Smith. 2011. Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology* 25, 3 (2011), 359–371.
7. Chinmay Kulkarni, Steven P Dow, and Scott R Klemmer. 2014. Early and repeated exposure to examples improves creative work. In *Design Thinking Research*. Springer, 49–62.
8. Richard L Marsh, Joshua D Landau, and Jason L Hicks. 1996. How examples may (and may not) constrain creativity. *Memory & Cognition* 24, 5 (1996), 669–680.
9. Bill Moggridge and Bill Atkinson. 2007. *Designing interactions*. Vol. 14. MIT press Cambridge.
10. Bernard A Nijstad, Carsten KW De Dreu, Eric F Rietzschel, and Matthijs Baas. 2010. The dual pathway to creativity model: Creative ideation as a function of flexibility and persistence. *European Review of Social Psychology* 21, 1 (2010), 34–77.
11. Bernard A Nijstad and Wolfgang Stroebe. 2006. How the group affects the mind: A cognitive model of idea generation in groups. *Personality and social psychology review* 10, 3 (2006), 186–213.
12. Bernard A Nijstad, Wolfgang Stroebe, and Hein FM Lodewijckx. 2002. Cognitive stimulation and interference in groups: Exposure effects in an idea generation task. *Journal of experimental social psychology* 38, 6 (2002), 535–544.
13. Eric F Rietzschel, Bernard A Nijstad, and Wolfgang Stroebe. 2007. Relative accessibility of domain knowledge and creativity: The effects of knowledge activation on the quantity and originality of generated ideas. *Journal of Experimental Social Psychology* 43, 6 (2007), 933–946.
14. Colleen M Seifert, David E Meyer, Natalie Davidson, Andrea L Patalano, and Ilan Yaniv. 1994. Demystification of cognitive insight: Opportunistic assimilation and the prepared-mind hypothesis. (1994).
15. Pao Siangliulue, Kenneth C. Arnold, Krzysztof Z. Gajos, and Steven P. Dow. 2015a. Toward Collaborative Ideation at Scale—Leveraging Ideas from Others to Generate More Creative and Diverse Ideas. In *Proc. CSCW'15*.
16. Pao Siangliulue, Joel Chan, Krzysztof Z. Gajos, and Steven P. Dow. 2015b. Providing Timely Examples Improves the Quantity and Quality of Generated Ideas. In *Proc. Creativity and Cognition'15*.
17. Omer Tamuz, Ce Liu, Serge Belongie, Ohad Shamir, and Adam Tauman Kalai. 2011. Adaptively Learning the Crowd Kernel. *arXiv.org* (May 2011).
18. L van der Maaten and K Weinberger. 2012. Stochastic triplet embedding. *Machine Learning for Signal Processing (MLSP), 2012 IEEE International Workshop on* (2012), 1–6.
19. Weiquan Wang and Izak Benbasat. 2007. Recommendation agents for electronic commerce: Effects of explanation facilities on trusting beliefs. *Journal of Management Information Systems* 23, 4 (2007), 217–246.
20. Thomas B Ward, Merryl J Patterson, Cynthia M Sifonis, Rebecca A Dodds, and Katherine N Saunders. 2002. The role of graded category structure in imaginative thought. *Memory & Cognition* 30, 2 (2002), 199–216.
21. Lixiu Yu, Aniket Kittur, and Robert E Kraut. 2014. Distributed analogical idea generation: inventing with crowds. In *Proc. CHI'14*. ACM, 1245–1254.
22. Haoqi Zhang, Edith Law, Rob Miller, Krzysztof Gajos, David Parkes, and Eric Horvitz. 2012. Human computation tasks with global constraints. In *Proc. CHI'12*. ACM, 217–226.