

Intelligent Systems to Support Large-Scale Collective Creative Idea Generation

Pao Siangliulue

Harvard School of Engineering and Applied Sciences
Cambridge, MA USA
paopow@seas.harvard.edu

ABSTRACT

In recent years, it has become possible for large groups of people to collaborate and generate ideas together in ways that were not possible before. However, the large number of ideas and participants in this setting also pose new challenges in helping people find inspiration from a large pool of ideas, and coordinating the collective effort. My research aims to 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 orchestrate collective effort to reduce redundancy and increase the breadth of generated ideas.

INTRODUCTION

Emerging large online innovation platforms like OpenIDEO.com, My Starbucks Idea and Quirky promise an unprecedented torrent of ideas, harvested from minds of thousands. The strengths of these platforms lie not only on the sheer number of contributors but also on the diversity of thinking, expertise, and knowledge of all contributors as a whole. Idea exchanges in such settings—often in the form of inspiration—bring about novel ideas that none of the contributors would have thought of alone.

But the large-scale idea generation paradigm also introduces new challenges. First, the large number of ideas reduces the chance of a contributor finding non-redundant inspiring ideas. Contributors usually have to wade through a large number of ideas with no effective way to find inspiring sets of ideas. Second, there is no existing mechanism to coordinate effort among a large number of contributors in idea generation. Instead of working together, contributors generate ideas alone without knowing how to best contribute. This lack of coordination produces redundant common ideas and limits the breadth of exploration of new ideas [11].

To help a large community of ideators make full use of its size and diversity, we need new methods and tools that en-

able them to make effective use of each other's ideas, and orchestrate collective effort to explore diverse possibilities.

HELPING PEOPLE FIND INSPIRING IDEAS

Idea exchanges on large innovation platforms often occur when contributors get inspired by each other's ideas. The ideas people see and when they see them affect the ideas they generate [4, 8]. An effective approach therefore should help a contributor find inspiring ideas from a large pool at an appropriate time.

Selecting a diverse set of ideas

Prior research suggests that a set of inspiring ideas consists of diverse ideas [6]. However, there is no robust scalable method to identify a diverse set of ideas. Asking human experts to look at all ideas and keep track of how each idea relates to others does not scale. When the number of ideas rises to several hundreds, the task becomes difficult



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

In our recent work [7], we developed a crowd-powered method for 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?”). We built our method on prior work on multidimensional scaling and active similarity learning techniques [10, 9] to embed ideas in an idea space with as few human queries as possible. From the idea map, we can algorithmically extract diverse sets of examples at scale. In contrast to some existing methods [1], our approach is applicable to any creative artifact and not is limited to textual ones. We also demonstrated that people

generate diverse ideas when they see a set of diverse ideas selected by our algorithm.

Extracting information from user natural interactions

So far, we have outsourced human computation tasks necessary for idea map generation to micro-task marketplaces like Amazon Mechanical Turk. Instead of relying solely on outside labors, we look for solutions that enable users to contribute without disrupting their idea generation. User contribution could be more reliable. Plus, they do not incur extra financial cost. Prior work in online education has shown methods that extract information from learners by embedding the tasks into activities that learners find intrinsically motivating [3, 12]. Likewise, we can provide ways for users to bookmark ideas for their own references as a part of information discovery process [2]. By aggregating what each person's bookmarks are and how they organize them, we can derive an input to the algorithms used for selecting inspiring ideas.

Suggesting ideas that inspire individual contributors

We should provide ideas that fit individual's idea generation process. Different people find inspiration in different ideas. In a pilot study, we found that people are more receptive to inspirations that are different but still share some properties with their own ideas. This finding agrees with prior finding that people generate better ideas when they integrate ideas that are conceptually closer rather than farther from each other [1]. If we can identify ideas that are similar to the contributors' ideas so far yet different in an interesting way, we might inspire them to come up with more interesting ideas. One possible approach is to place generated ideas on an idea map and suggest ideas of certain distant from these ideas. Other approaches such as selecting inspirations based on user-generated labels would be worth exploring as well.

Choosing appropriate time to show inspirations

Once we find a set of inspiring ideas, we have to select the right moment to present them. Recent work has demonstrated that not just the choice of examples, but also the timing of their delivery can impact creative outcomes [5]. We have explored two timing mechanisms: (1) participants can request ideas on demand, and (2) a system automatically infers when participants are stuck and provides ideas at that moment. Our results show that people who requested examples themselves generated the most novel ideas and people who received ideas automatically when idle produced the most ideas [8]. These results can inform the design of ideation support systems that aim to help people generate many high quality ideas.

COORDINATING EXPLORATION OF IDEA SPACE

As the number of contributors increases, the more crucial and difficult it becomes to coordinate the community efforts. Without coordination or some guidance from the system, people may not know how to best contribute. They could end up generating redundant ideas that are very similar to one another instead of exploring a broader range of ideas in the idea space. However, too much communication overhead can be overwhelming.

We are still drafting the plan to explore different mechanisms that reduce idea redundancy and widen the breadth of

ideas generated while minimizing communication overhead. I would love to hear suggestions on which of various directions I should take for my thesis. One possible mechanism is to identify which parts of the idea space have not been explored and direct more people to generate ideas in those unexplored regions of the space. This mechanism can be integrated into a system as a form of an explicit todo list [13] or as a subtle nudge through inspiration suggestion. Another approach is to provide people with an overview of the state community's idea generation along with an effective way for contributors to inform others of their plan so that they don't interfere with each other.

CONCLUSION

Large ideation platforms promise a huge opportunity for innovation, but are limited by lack of appropriate methods and tools that helps people discover inspiring ideas and coordinate their effort. By synthesizing knowledge human computation and machine learning, we can create an intelligent system that addresses the existing challenges and improves collective idea generation at scale.

REFERENCES

1. Chan, J., Dow, S., and Schunn, C. Conceptual distance matters when building on others' ideas in crowd-collaborative innovation platforms. In *Proc. CSCW'14*, ACM (2014), 141–144.
2. Kerne, A., Koh, E., Dworaczyk, B., Mistrot, J. M., Choi, H., Smith, S. M., Graeber, R., Caruso, D., Webb, A., Hill, R., et al. combinformation: a mixed-initiative system for representing collections as compositions of image and text surrogates. In *Proc. JCDL'06*, ACM (2006), 11–20.
3. Kim, J., Guo, P. J., Cai, C. J., Li, S.-W. D., Gajos, K. Z., and Miller, R. C. Data-driven interaction techniques for improving navigation of educational videos. In *Proc. UIST'14*, ACM (2014), 563–572.
4. Kohn, N. W., and Smith, S. M. Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology* 25, 3 (2011), 359–371.
5. Kulkarni, C., Dow, S. P., and Klemmer, S. R. Early and repeated exposure to examples improves creative work. In *Design Thinking Research*. Springer, 2014, 49–62.
6. Nijstad, B. A., Stroebe, W., and Lodewijkx, H. F. M. Cognitive stimulation and interference in groups: Exposure effects in an idea generation task. *Journal of Experimental Social Psychology* 38, 6 (2002), 535–544.
7. Siangliulue, P., Arnold, K. C., Gajos, K. Z., and Dow, S. P. Toward collaborative ideation at scale—leveraging ideas from others to generate more creative and diverse ideas. In *Proc. CSCW'15* (2015).
8. Siangliulue, P., Chan, J., Gajos, K. Z., and Dow, S. P. Providing timely examples improves the quantity and quality of generated ideas. In *Proc. Creativity and Cognition'15* (2015). To appear.
9. Tamuz, O., Liu, C., Belongie, S., Shamir, O., and Kalai, A. T. Adaptively Learning the Crowd Kernel. *arXiv.org* (May 2011).
10. van der Maaten, L., and Weinberger, K. Stochastic triplet embedding. *Machine Learning for Signal Processing (MLSP), 2012 IEEE International Workshop on* (2012), 1–6.
11. Ward, T. B., Patterson, M. J., Sifonis, C. M., Dodds, R. A., and Saunders, K. N. The role of graded category structure in imaginative thought. *Memory & Cognition* 30, 2 (2002), 199–216.
12. Weir, S. A. Learnersourcing subgoal labels for how-to videos. In *CHI'14 Extended Abstracts*, ACM (2014), 945–950.
13. Zhang, H., Law, E., Miller, R., Gajos, K., Parkes, D., and Horvitz, E. Human computation tasks with global constraints. In *Proc. CHI'12*, ACM (2012), 217–226.