CrowdMuse: An Adaptive Crowd Brainstorming System

Victor Girotto[1] Erin Walker[1] Winslow Burleson[2]

[1]School of Computing, Informatics, and Decision Systems Engineering
Arizona State University
Tempe, AZ, USA
{victor.girotto, erin.a.walker}@asu.edu

[2]Rory Meyers College of Nursing
New York University
New York, NY, USA
wb50@nyu.edu

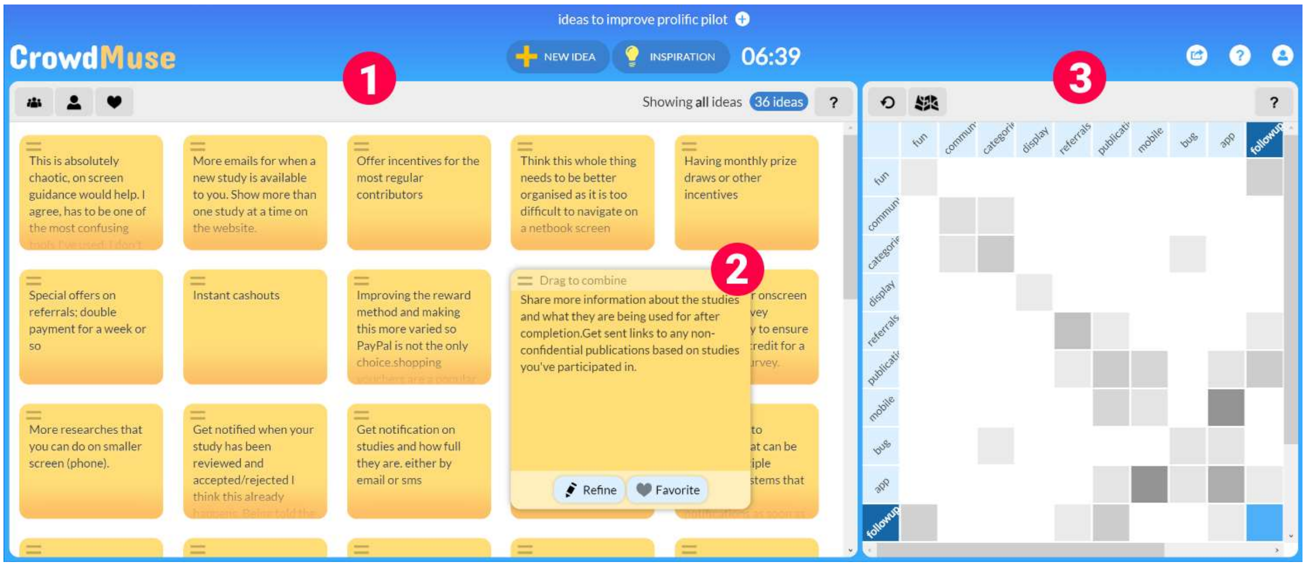

Figure 1 The CrowdMuse System and its two main views: the idea workspace (#1), on the left, allows users to view and manipulate ideas. Hovering over an idea expands it and shows other possible actions (#2); and the solution space (#3), on the right, provides an overview of the density of ideas developed for each tag.

ABSTRACT

Online crowds, with their large numbers and diversity, show great potential for creativity. Research has explored different ways of augmenting their creativity, particularly during large-scale brainstorming sessions. Traditionally, this comes in the form of showing ideators some form of inspiration to get them to explore more categories or generate more ideas. The mechanisms used to select which inspirations are shown to ideators thus far have been random or focused on characteristics of the inspirations rather than on ideators. This can hinder their effect, as creativity research has shown that ideators have unique cognitive structures and may therefore be better inspired by some ideas rather than others. We introduce CrowdMuse, an adaptive system for supporting large scale brainstorming. The system models ideators based on their past ideas and adapts the system views and inspiration mechanisms accordingly. An evaluation of this system could shed light on how to better individually support ideators.

Conference'10, Month 1–2, 2010, City, State, Country.

DOI: http://dx.doi.org/10.1145/12345.67890

CCS Concepts

• Information systems→ Collaborative and social computing systems and tools • Human-centered computing→ Collaborative and social computing systems and tools

Keywords

Creativity; Brainstorming; Crowd; Adaptive systems;

1. INTRODUCTION & RELATED WORK

Online crowds show great potential for creativity, due in great part to their large numbers and diversity [4,5]. However, simply recruiting large numbers of ideators is not enough to ensure a creative output. Research has attempted to support crowd brainstorming through different kinds of *inspirations*, such as employing expert facilitators [3], showing ideators a diverse set of ideas [8], or by leveraging analogies [9]. These attempts share a focus on the inspirations themselves rather than on the ideators.

Creativity research has shown that ideators leverage their own cognitive structures during brainstorming [2,7]. This means that each ideator is more likely to focus on some areas (i.e. idea categories) rather than others and may vary in their strategy (focus on breadth vs. focus on depth). An inspiration strategy that doesn't take the ideator into consideration may be missing out on leveraging their unique strengths for idea generation. For example, if an ideator is more familiar with ideas in category A than those in B, showing ideas in category B may not effectively inspire him or her to come up with new ideas. The overarching research question for this research then is: "*How can we adapt inspirations to ideators in order to maximize their effect?*"

2. THE CROWDMUSE SYSTEM

To explore this question, we built the CrowdMuse system[1]. The system models ideators based on the categories of ideas they develop and adapts its views and inspiration mechanisms based on that model. Figure 1 depicts the system's main interface. It is comprised of two main views. The first is the idea workspace (#1), whose purpose is to allow users to explore and manipulate existing ideas. It has options for loading all ideas, only the current user's ideas, or all ideas a user marked as favorites. Ideas can be expanded by hovering the mouse over them (#2).

The second is the solution space (Figure 1, #3), whose purpose is to provide an overview of which categories have been thoroughly explored and, conversely, those which are yet to be explored. It is an $n \times n$ matrix in which the rows and columns correspond to the idea categories developed so far. The color of the cell indicates how many ideas have been developed at the intersection of two categories—the darker the cell, the more ideas have been developed. Clicking a cell will open all ideas at that category intersection in the idea workspace.

Ideas can be added by clicking the "new idea" button at the top of the UI. When adding a new idea, the user is prompted to pick at most two categories for the idea (based on previously used categories), or to suggest new ones. To the right of the new idea button is an inspiration button, which when clicked presents three ideas along with a small microtask on each (e.g. "rate the idea's originality and usefulness"). The microtasks are used to increase the attention to ideas, following our previous work [6].

2.1 Adaptations

The system's purpose is to enhance idea generation by prioritizing categories that could be inspiring to an ideator. This is done through two forms of adaptation: explicit inspirations, and through an adaptation of the solution space. Explicit inspirations are shown when the user clicks the inspiration button, at which point the system selects and shows three ideas from other ideators (along with the microtask for each). This intervention is designed to strongly draw the ideator's attention to the ideas [6], and therefore it will prioritize the categories that are least likely to be distracting, that is, those that would not interrupt a train of thought [7]. In practice, this means either selecting the user's *current category* (i.e. the category used in his or her last idea), or *adjacent categories*, that is, a category that the ideator has previously visited after the current category. The system may also suggest *inferred categories*, explained below. This selection is done probabilistically, based on how often the user switches between categories (we call this the *category switching ratio*).

The system also performs an adaptation of the solution space by ordering its rows and columns. This reordering happens every time the user submits a new idea. Since the goal for the solution space is to give users an overview of all developed ideas, the purpose for adapting this view is to guide users' attentions to the most relevant categories when exploring the solution space. The categories are prioritized in the following order: 1) current category followed by adjacent categories. This is similar to the prioritization used in the explicit inspirations; 2) *inferred new categories*. These are categories that have not yet been visited by this ideator, but that are likely to be useful for him or her, based on other ideators who have a similar behavior (see the next section); 3) other previously visited categories, but that are not adjacent; 4) any other category that has not yet been visited.

2.2 User Modelling

The adaptations described above are powered by an underlying user model. This model is inferred based on a user's behavior within the system: whenever users add an idea, they are asked to choose one or two categories for their idea. This selection is done through a list of existing categories. The system then uses the chosen categories to update the user's model. The model's purpose is to answer the question: *which categories of ideas should the system show next?*

The model attempts to answer this question by keeping track of different types of information. The first is by keeping track of basic descriptive information. This includes the user's current categories (i.e., the categories of their last idea), as well as the frequency of category switching ratio (calculated as the ratio between sequential idea pairs of different categories over all idea pairs). This is in line with previous idea generation models, which foresee some ideators being more likely to stay within the same category, while other may more likely frequently switch categories [2].

The system also generates a *transition graph* based on the sequential categories used by one ideator. This is a directed, weighted graph, in which each node represents a category. When the user adds an idea, the system creates an edge between the categories for the latest idea and the one added before it. The weight of the edge increases is that transition is repeated. This is how the system determines the *adjacent categories*.

Finally, the system also keeps track of a user's category vector. This vector contains the frequency of ideas per category for each user. Drawing from collaborative filtering techniques [1], this vector is used to *infer new potential categories* based on the most similar users. In other words, this functions as a category recommender system for each user. This allows the system to intelligently suggest categories other than those that have already been explored by a given user.

3. CONCLUSION

Ideators have differences among themselves in which idea categories they may be more fluent. Current approaches for improving crowd brainstorming have not yet explored these differences. To this purpose, we built the CrowdMuse system, which keeps track of an ideation model, and adapts the system based on that model. Evaluation of the efficacy of this system can provide us with a better understanding of how to support idea generation across individual differences.

Our hypothesis is that through these adaptations, users can explore categories in greater depth (by showing them more ideas in their current category), as well as greater breadth by exposing them categories that have been useful to similar ideators. We expect the explicit inspirations to have greater effect than simple view adaptations. We are currently running pilots to test the system, both face-to-face and online through crowdsourcing platforms. The purpose for these pilots is to ensure that the system is usable, while collecting data that will be used during the actual studies. The actual study, to be ran online, will provide us with better understanding of whether this approach works, or how it could be improved.

[1] http://www.crowdmuse.io

4. REFERENCES

1. J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez. 2013. Recommender systems survey. *Knowledge-Based Systems* 46: 109–132. https://doi.org/10.1016/j.knosys.2013.03.012
2. Vincent Brown, Michael Tumeo, Timothy S. Larey, and Paul B. Paulus. 1998. Modeling Cognitive Interactions During Group Brainstorming. *Small Group Research* 29, 4: 495–526.
3. Joel Chan, Steven Dang, and Steven P. Dow. 2016. Improving Crowd Innovation with Expert Facilitation.
4. Alan R. Dennis and Mike L. Williams. 2003. Electronic Brainstorming: Theory, Research, and Future Directions. In *Group creativity: Innovation through collaboration*. Oxford University Press.
5. Gerhard Fischer. 2005. Distances and diversity: sources for social creativity. 128. https://doi.org/10.1145/1056224.1056243
6. Victor Girotto, Erin Walker, and Winslow Burleson. 2017. The Effect of Peripheral Micro-tasks on Crowd Ideation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
7. 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: 186–213.
8. Pao Siangliulue, Kenneth C. Arnold, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Toward Collaborative Ideation at Scale: Leveraging Ideas from Others to Generate More Creative and Diverse Ideas. 937–945. https://doi.org/10.1145/2675133.2675239
9. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Distributed analogical idea generation: inventing with crowds. 1245–1254. https://doi.org/10.1145/2556288.2557371