LineChange: An Analytic Framework for Automated Moderation of Crowdsourcing Systems

Brent D. Fegley(✉), Ryan Mullins, Ben Ford, and Chad Weiss

Aptima, Inc., Woburn, MA 01801, USA
bfegley@aptima.com

Abstract. If humans are more productive in collective problem-solving with a modicum of active help and guidance, then the potential of automated moderation of crowdsourcing systems has yet to be realized. Here, we present the conceptual design of an intelligent machine capable of (a) monitoring the temporal, structural, and emergent characteristics of participant behavior in a problem-solving process, and (b) modifying team structure and prompting participants for input at opportune or transitional moments in that process—by configuration, rule, or inference—to achieve collective goals and optimize output. The design is unique in treating teams as composable objects, in being scale-free, in relying on configuration and inference (not hard-coding), and in treating participant behaviors as sensory input.

Keywords: Crowdsourcing · Insourcing · Network generation
Automated facilitation · Decision support

1 Overview

Despite reported successes of human collective intelligence in collaborative problem-solving and decision-making, the power and utility of crowdsourcing remains encumbered by the nature and facilitation of its tasking. Crowdsourcing continues to be explored in a variety of contexts where focus tends toward highly structured, relatively simplistic, and parallelizable tasks, including, but not limited to, text editing [3] and dataset development [14]. For example, the Good Judgment Project [11] excels in part because of the elegance of its methods for eliciting and scoring forecasts from participants who work in parallel, and largely independently, on tasks that require an assessment of event probability, for which participants are personally accountable. In contrast, research shows that complex problem-solving by groups achieves more effective outcomes (collectively and interpersonally), especially when actively facilitated. However, expert human facilitators are rare; and the resources and coordination necessary to engage them represent significant hurdles to their use. (Note [1,16], among others.) Opportunity exists to address problems of increased social and technical

C. Stephanidis (Ed.): HCII Posters 2018, CCIS 850, pp. 401–408, 2018.
https://doi.org/10.1007/978-3-319-92270-6_57

complexity—especially those involving defeasible and incendiary arguments—through automated facilitation of ideation, collaboration, and consensus.

Here, we present work-in-progress to design a framework for automated moderation, facilitation, and intervention of participant activities within a hypothetical crowdsourced problem-solving process. We call the underlying analytic framework *LineChange*. We hypothesize that in problem-solving, like hockey, better outcomes may be obtained by varying the lineup of participants throughout the endeavor.[1] In an idealized implementation of this process, an intelligent machine (i.e., background service) monitors the temporal, structural, and emergent characteristics of participant behavior and then, at opportune or transitional moments, by configuration, rule, or inference, prompts participants for input or modifies team structure.

2 Related Work

If a goal of crowdsourced problem-solving is enlightened decision-making, then brainstorming, debate, and consensus are surely corresponding high-level objectives. These objectives have been discussed variously and extensively in the context of non-cooperative game theory, political science, and economics, as well as IS (information system)-centric research on group decision support systems (GSS), electronic brainstorming (EBS), and computer-supported cooperative work (CSCW; e.g., [2,5,9,13]).

Brainstorming (ideation) provides a useful way to understand problem-solving. Problems cannot be solved without some amount of creative thinking; but *quality* does not necessarily follow *quantity* in idea creation. Briggs and Reinig [4] identify six factors ("boundaries") that affect the quality of ideation: *mental ability*, a "function of both intelligence and domain-relevant expertise"; *solution space*, a group-level effect, the continuum of which extends from closed-ended tasks that have a finite number of solutions to open-ended tasks that have an infinite number of solutions; *problem understanding*, or the extent to which participants have accurate information about the problem and understand their task; *attention*, which is analogous to cognitive load; *goal congruence*, or "the degree to which participants perceive that working toward a group goal will be instrumental to the attainment of their salient private goals"; and *exhaustion* in both mental and physical contexts. *Dissent* should be added to this list as well, having been found to encourage "creative and divergent thinking" [12]. Thus, the factors defined by [4] should have currency in debate as well as brainstorming.

3 Methodology

LineChange augments and facilitates complex, crowdsourced problem-solving, particularly efforts involving unstructured deliberation among participants,

[1] The National Hockey League formalized the term *line change* in its official rules; we co-opt the term here in name and spirit only. In our formulation, teams do not compete directly with one another but act independently, for collective benefit and in service of a single goal.

using parameterized rules based on a process model for problem-solving. One possible model is depicted in Fig. 1. This model is illustrative only; many different models are possible, including those that eschew debate or re-conceive problem-solving objectives entirely. Here, brainstorming and consensus bound a cycle of debate (i.e., argument) intended to help participants shrink the space of possible solutions. Arrows indicate control flow; activities with red backgrounds are participant-directed and time-limited, by configuration (not hard-coding). The model always produces a result, whether conclusive or not, and contains no short-circuits for "failure."

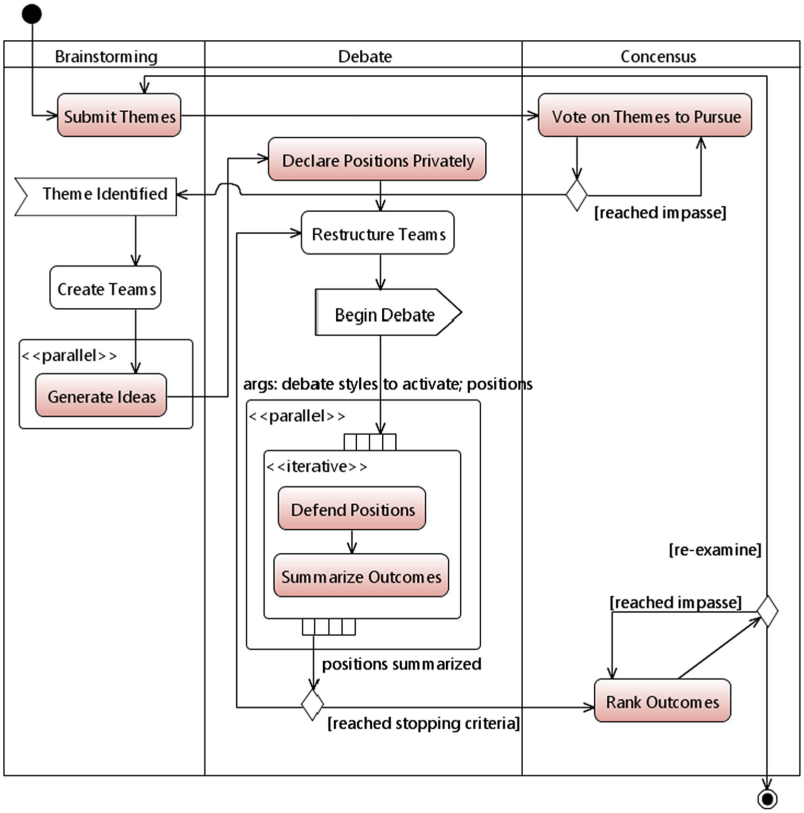

Fig. 1. Hypothetical process model for crowdsourced problem-solving. The process begins with brainstorming and ends with consensus but is facilitated by configurable rewiring of a supporting artificial network of participants.

LineChange is charged with (a) specifying how participants engaged in problem-solving should be organized at appropriate steps in the process, (b) monitoring participant use of computer input devices to trigger individual action or team progression from one activity to another, because events initiated by participant interaction with the system are the inputs of greatest interest, and (c) monitoring time bounds for participant-directed activities to signal delivery of messages or prompts to participants ahead of deadlines. In implementation, all of these responsibilities can be initialized by configuration and managed by inference. The first two may interact and deserve more explanation. The former is the subject of the wiring protocol; the latter, triggering mechanisms, which may be described and measured about the boundaries of Briggs and Reinig [4]. For reasons of space, only the wiring protocol will be discussed herein.

3.1 Wiring Protocol

A phase entails a period of activity about a single objective and a particular arrangement of participants, such as "Declare Positions Privately" or the combination "Defend Positions" and "Summarize Outcomes" in the foregoing figure. Additionally, while the anonymity of participants is assumed for simplicity, nothing prohibits identifiability or a mixed mode in application.

We hypothesize that each phase of the problem-solving process may be facilitated by a particular arrangement (wiring, network) of participants. For example, brainstorming is facilitated by isolates or cliques (complete subgraphs), consensus by a complete graph (all participants are connected to one another and have global knowledge; note [8]). This is one approach, not a hard constraint. When not otherwise arranged as a collection of isolates or as a complete graph, participants may be configured as a set in a bipartite graph, the other disjoint set representing teams. In this case, edges would link participants to one or more teams.

Each phase of the process is structured about who can see what when. A team, as the term is used here, is an abstraction representing a group of participants with a flat structure (a structure without formal roles or hierarchies) that share resources (such as posts, comments, ratings, and summaries). For example, while brainstorming could be done individually as well as in teams, the content of a particular brainstorming phase would be accessible to members of the immediate team only. Edges between teams may be created to grow a team, in which case content previously accessible to some is now accessible to all members of the larger team. Self-organization is possible within teams, depending on task (such as in debate), but not across teams.

The wiring protocol determines how the bipartite graph is initialized and (re)configured during activity transitions. A few considerations follow in how the protocol could be enacted. First, a team may be *non-persistent*, in which case, unlike participants, teams would exist only for the duration of the activity (task). The wiring protocol would need to define how teams are lumped or split from one activity to another. Alternatively, a team may be *persistent*, in which case teams would persist from one activity to another. Characterizing teams becomes possible only if at least one team member remains. Edges could be added or dropped between teams to preserve original team structures (simplifying post hoc analysis), thus expanding or restricting the collective team resources to which individual participants have access. Second, participants could be allowed to vary in number from one activity to another (such as by transitioning to/from active or inactive status) depending on metrics of their performance or level of engagement, however those metrics might be defined. Edges between participants (not just between participants and teams) could be established to influence whether the participants are likely to remain paired (i.e., on the same team). Rewiring may be necessary if participants drop out, either actively (e.g., by deactivating their account) or passively (e.g., by not using the system again).

The wiring protocol could be extended in several ways. For example, a non-uniform probability distribution could be created over teams to allow selection of teams with different sizes. Additional parameters could include (a) the propensity of a participant to work with a previous collaborator (a control for increasing or decreasing team degree or team member swapping; see [7]) and (b) the probability that a given participant will be chosen as a boundary spanner (i.e., linked to more than one group, in the spirit of [6]). The foregoing could be predetermined or adaptive per phase, informed by metrics of participant behavior, for example, the ratio of ostensibly positive interactions over all interactions with another participant.

Example rewiring protocol (simplistic instantiation)
Setup

- Participants remain active throughout the problem-solving process.
- Teams, as well as edges between teams and participants, persist from one phase to another.
- The process model has three objectives (brainstorming, debate, and consensus), each with associated activities and with parameters held constant.

Parameters
N: (integer) number of eligible participants; K: (integer) number of separate subgraphs (i.e., teams); D: (integer) number of debate cycles; Q: (ratio) proportion of actual to possible pairwise participant connections at D.

Initialization

1. Set $K, where\ 1 \leq K \leq N$.
2. Assign N participants to K teams uniformly at random so that each team has N/K members, adjusting when necessary to ensure that no participant is left behind (integers only). At this point, all participants belong to a team, each with equal weight; and each team constitutes a subgraph.

Phase Transition

- debate → consensus: Ensure that all teams are connected to one another, creating a complete graph of them. Implication: participants achieve global knowledge.

- brainstorming → debate: Drop all edges between teams; but retain the edges between teams and participants. Implication: participants now have local knowledge only.
- debate → debate: Add edges selectively between teams that have none otherwise—a function of D, Q, the number of non-existent edges, and the index of the current cycle. Implication: exposure to new or different ideas or perspectives is controlled to minimize cognitive load and help narrow focus towards convergence on problem solutions.

4 An Example

To illustrate how *LineChange* might facilitate problem-solving, we will now explore the organizational challenge of company reorganization due to an influx of new employees following merger or acquisition, with particular focus on roster selection following a decision about hierarchical structure.

The essential task is relatively straightforward: closed-category card sorting [15], where the categories are the leaves in the organizational chart (including an "I don't know" category) and where the cards are the personnel that need to be sorted. The challenges in this task are twofold: (a) coordinating the elicitation of solutions to mitigate groupthink (i.e., bias towards majority opinion); and (b) resolving conflicting sorts (e.g., one person is equally valued in two or more parts of the organization).

LineChange might be engaged as follows, given the workflow of Fig. 1. First, a few assumptions: a participant in the workflow is anonymous unless that participant explicitly chooses to reveal his or her identity; participants include the personnel being sorted; at least one "theme" (a concise statement of the problem to be solved) submitted by a participant would include a call for card-sorting with a recommendation to make card-sorting tractable by creating disjoint subsets of the personnel for sorting; participants use a voting system to select the foregoing submission; themes subsidiary to the one submitted are constructed, each containing a disjoint subset of personnel, each queued for submission to the brainstorming objective; and consensus is reached by majority vote using the Borda count method (see [10]).

For the brainstorming objective, by configuration, isolates (that is, participants working in virtual isolation of one another rather than teams) carry out idea generation. When the configured time limit for this activity has ended, the results of card-sorting efforts are fed to the debate objective where they become privately declared positions. Isolates are then organized into teams.

Debate begins (a) with the anonymized exposure of team members' private positions to the rest of their team, not to all teams, and (b) using one or more of several possible debate interfaces (again, by configuration; pro vs. con debate, threaded discussion, argument mapping) to compare and narrow down the set of possibilities. *LineChange* prompts all teams to begin summarizing positions in anticipation of the time-bound set for a round of debate. At the end of a round, the summarized positions of all teams are gathered and, if stopping criteria have not been met, another round of debate commences with a team rewiring. For simplicity, we imagine that *LineChange* is configured to rewire teams by connecting nodes representing teams in the bipartite graph. By this method, team size increases with each round of debate, thus limiting participant exposure to too many new positions at once.

Criteria that stop debate for all teams could entail an overall time-bound for debate or a target mixing rate. Once such criteria have been met, the problem-solving process ends with all participants ranking the summarized positions of the last teams engaged in debate. This process continues to iterate until all cards have been sorted into agreement.

Consensus is determined by voting using the Borda count method. Conceivably, the workflow just described may end in stalemate, where agreement about categorization cannot be reached. In such cases, a preferential method of voting could be used, such as that used by the Australian electoral system; voting recurs after the least popular choice is dropped. Regardless, a call can be made for re-examination in the form of a new theme submission.

5 Epilogue

LineChange is an analytic framework that makes active moderation, facilitation, and intervention part of the problem-solving process. In practice, it could function as an automated service, listening to events representing participant interaction with a crowdsourcing system and triggering actions intended to elicit different behaviors as a process unfolds.

LineChange distinguishes itself from other approaches, in part, because (a) it treats teams as composable objects (they may vary in size or number or cease to exist from one phase of a problem-solving process to the next, depending on the wiring protocol); (b) the framework is scale-free in being unaffected by time (process duration), space (proximity of participants to one another), and size (of the participant pool); (c) the framework is configurable in depending on a parameterized set of rules; and (d) participants provide first-order input as sensors.

Acknowledgements. This research was performed in connection with contract N68335-18-C-0040 with the U.S. Office of Naval Research. We would like to thank Dr. Yiling Chen and Dr. Predrag Neskovic for their contributions to this work as thought partners.

References

1. Anson, R., Bostrom, R., Wynne, B.: An experiment assessing group support system and facilitator effects on meeting outcomes. Manag. Sci. **41**(2), 189–208 (1995). https://doi.org/10.1287/mnsc.41.2.189
2. Austen-Smith, D., Feddersen, T.: Deliberation and voting rules. In: Austen-Smith, D., Duggan, J. (eds.) Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks. Studies in Choice and Welfare, pp. 269–316. Springer, Heidelberg (2005). https://doi.org/10.1007/3-540-27295-X_11

3. Bernstein, M.S., Little, G., Miller, R.C., Hartmann, B., Ackerman, M.S., Karger, D.R., Crowell, D., Panovich, K.: Soylent: a word processor with a crowd inside. Commun. ACM **58**(8), 85–94 (2015). https://doi.org/10.1145/2791285
4. Briggs, R.O., Reinig, B.A.: Bounded ideation theory. J. Manag. Inf. Syst. **27**(1), 123–144 (2010). https://doi.org/10.2753/MIS0742-1222270106
5. Dennis, A.R., Williams, M.L.: Electronic brainstorming: theory, research, and future directions. In: Paulus, P.B., Nijstad, B.A. (eds.) Group Creativity: Innovation Through Collaboration, pp. 160–180. Oxford University Press, Oxford (2003). https://doi.org/10.1093/acprof:oso/9780195147308.003.0008
6. Granovetter, M.: The strength of weak ties: a network theory revisited. Sociol. Theory **1**, 201–233 (1983). http://www.jstor.org/stable/202051
7. Guimer, R., Uzzi, B., Spiro, J., Amaral, L.A.N.: Team assembly mechanisms determine collaboration network structure and team performance. Science **308**(5722), 697–702 (2005). http://science.sciencemag.org/content/308/5722/697
8. Judd, S., Kearns, M., Vorobeychik, Y.: Behavioral dynamics and influence in networked coloring and consensus. Proc. Nat. Acad. Sci. **107**(34), 14978–14982 (2010). http://www.pnas.org/content/107/34/14978.abstract
9. Kim, T., Chang, A., Holland, L., Pentland, A.S.: Meeting mediator: enhancing group collaboration using sociometric feedback. In: Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work, CSCW 2008, pp. 457–466. ACM (2008). https://doi.org/10.1145/1460563.1460636
10. Mao, A., Procaccia, A.D., Chen, Y.: Better human computation through principled voting. In: AAAI (2013)
11. Mellers, B., Stone, E., Murray, T., Minster, A., Rohrbaugh, N., Bishop, M., Chen, E., Baker, J., Hou, Y., Horowitz, M., Ungar, L., Tetlock, P.: Identifying and cultivating superforecasters as a method of improving probabilistic predictions. Perspect. Psychol. Sci. **10**(3), 267–281 (2015). https://doi.org/10.1177/1745691615577794
12. Nemeth, C.J., Nemeth-Brown, B.: Better than individuals? The potential benefits of dissent and diversity for group creativity. In: Paulus, P.B., Nijstad, B.A. (eds.) Group Creativity: Innovation Through Collaboration, pp. 63–84. Oxford University Press, Oxford (2003). https://doi.org/10.1093/acprof:oso/9780195147308.003.0004
13. Ottaviani, M., Srensen, P.: Information aggregation in debate: who should speak first? J. Public Econ. **81**(3), 393–421 (2001). https://doi.org/10.1016/S0047-2727(00)00119-5
14. Post, M., Callison-Burch, C., Osborne, M.: Constructing parallel corpora for six Indian languages via crowdsourcing. In: Proceedings of the Seventh Workshop on Statistical Machine Translation, pp. 401–409. Association for Computational Linguistics, Montréal, June 2012. http://www.aclweb.org/anthology/W12-3152
15. Spencer, D.: Card Sorting: Designing Usable Categories. Rosenfeld Media, New York (2009)
16. Wong, Z., Aiken, M.: Automated facilitation of electronic meetings. Inf. Manag. **41**(2), 125–134 (2003). https://doi.org/10.1016/S0378-7206(03)00042-9