Setting Group Priorities – Swarms vs Votes

Louis Rosenberg and David Baltaxe
Unanimous A.I.
2443 Fillmore Street, #116
San Francisco, CA. USA
david@unanimousai.com

***Abstract* — As established by the Condorcet Jury Theorem, the statistical average of a group-wise vote will generally outperform the accuracy of the individual participants. Because of this, many organizations use polls and surveys for critical decisions, such as setting group priorities. Unfortunately, the conditions required by the Condorcet Jury Theorem are very strict, demanding (a) that participants are fully independent when casting votes, with no cross-team influences or social biasing, (b) that all members of the team are skilled performers who render correct decisions more than 50% of the time, and (c) that the questions are binary, with members selecting between only two options. A major problem, therefore, is that real world teams engaged in authentic decisions, judgements, and estimations rarely satisfy the ideal conditions for statistical accuracy amplification. The present study explores the use of "human swarming" as an alternative to polls and surveys for real-world tasks such as the setting of group priorities. More specifically, this study tasked a group of 43 voting age Americans with prioritizing a set of political objectives by vote and by swarm, and then asked the members to rate their satisfaction with the resulting prioritizations. It was found that 68% of the participants rated the swarm-based result as a more accurate reflection of their personal priorities than the vote-based result. In addition, 74% of participants rated the swarm-based result as a more accurate reflection of the group's priorities than the vote-based result. With satisfaction being a core success measure for a prioritization task, it appears that real-time swarming may offer groups a significant benefit as compared to traditional polls and surveys.**

Keywords— Swarm Intelligence, Artificial Intelligence, Human Swarming, Wisdom of Crowds, Collective Intelligence

I. Introduction

From business teams to political parties, organizations often find it extremely challenging to prioritize their top objectives. As a consequence, priority-setting can easily become a high conflict endeavor within teams, especially when the group is diverse, including participants of varied background, discipline, or expertise. To make matters worse, conflict in priority-setting is not just unpleasant, it can be counterproductive, reducing the buy-in among participants in the final outcome. To mitigate such conflicts, many organizations have turned away from purely deliberative priority-setting methods in favor of more objective statistical means, using votes, polls, and surveys to derive average results that inform group-wise prioritization. This approach is often justified by historical research that shows the statistical average of group decisions, forecasts, and judgements, outperforming the accuracy of individual responses.[1]

Much of the rationale for treating groups as statistical rather than deliberative entities goes back to the Marquis de Condorcet, who worked to justify the shift from dictatorial monarchy to representative democracy during the turmoil of the French Revolution. His intent was to validate the "will of the people" as an intelligent and effective way to reach societal decisions, render judgements, and set political priorities. Memorialized as the *Condorcet Jury Theorem*, his work shows that so long as each member of a group provides a correct judgement more than 50% of the time, the statistical average of group members will outperform the individuals, the larger the group the greater the accuracy advantage. The theorem requires, however that all individuals provide their input independently, with no influence from other members. In other words, no deliberation, cross-pollination, or social biases – a purely statistical result that averages individuals in perfect isolation.[2]

But what if the individuals are not correct more than 50% of the time as required by the Condorcet Jury Theorem? In such cases, the statistical average of participants will underperform the accuracy of individuals, with the collective insights getting less accurate as the group size increases. This makes polling a risky endeavor for group decision-making as it can amplify poor judgement. Furthermore, is it realistic to model participants in real-world decision tasks as purely independent actors, as is formally required by the Jury Theorem? Probably not, for most members of a working team share similar biases and impose cross-team influences, not to mention the impact of a shared organizational culture. Clearly, the strict idealization of the Condorcet Jury Theorem faces real-world practicalities. In addition, while the use of statistical averages via vote, poll, or survey, has been shown to give improved results in idealized cases, there is no reason to believe that such methods yield the very best results. This inspires the research question – *is there a better way for groups to decide upon their common priorities?*

To find a more effective method for group prioritization, the present researchers looked to Mother Nature for guidance. That's because many natural species make collective decisions that greatly outperform the intellectual capacity of the individual organisms in the group. Referred to as *Swarm Intelligence* (SI), nature generally achieves this amplification by enabling groups to form closed-loop systems in which participants explore the decision-space in real-time synchrony and converge on optimal outcomes. One of the most studied examples of amplified Swarm Intelligence is among honeybee swarms, which have been shown to prioritize potential home sites and select the optimal destination 80% of the time.[3,4,5,6] But can humans use

978-1-5090-3502-1/16/$31.00 ©2016 IEEE

similar real-time swarming methods to reach optimized group decisions? Prior research into human swarming has shown that by enabling groups of online users to combine their knowledge, wisdom, insights, and opinions in real-time swarms, enhanced predictions and forecasts can be made.[7,8,9,10,11]

Prior research, however, does not address priority setting, which inspires the question: Can real-time swarming be used by groups to converge upon preferred sets of priorities as compared to traditional polls, votes, and surveys? To answer this question, researchers used the UNU swarm intelligence platform to compare priority-setting among diverse groups by vote and by swarm. More specifically, researchers assembled a group of 43 voting-age Americans of mixed party affiliation and tasked them with evaluating and prioritizing a set of political objectives that the government should focus on. The group was required to order the set of priorities, from most important to least important, in two ways: (i) by ranking individual preferences on a traditional online survey, which would then be mathematically combined to set priorities and (ii) by working together as online swarm, setting the priorities in real-time synchrony.

II. Enabling Human Swarms

To enable real-time decisions among groups of networked users, the UNU online platform was employed. UNU allows users to login simultaneously from all around the world and participate in closed-loop swarms. As shown in Figure 1, users answer questions by collectively moving a graphical puck to select among a set of alternatives. The puck is modeled as a physical system with a defined mass, damping and friction. Users provide input by manipulating a graphical magnet with a mouse or touchscreen. By positioning their magnet, users impart their personal intent as a force vector on the puck. The input from each user is not a discrete vote, but a stream of vectors that varies freely over time. Because the full population of users can adjust their intent at every time-step, the puck moves in response to the dynamics of the full system. This enables a real-time negotiation among the members of the swarm, the group collectively exploring the decision-space and converging on the most agreeable answer.[7]

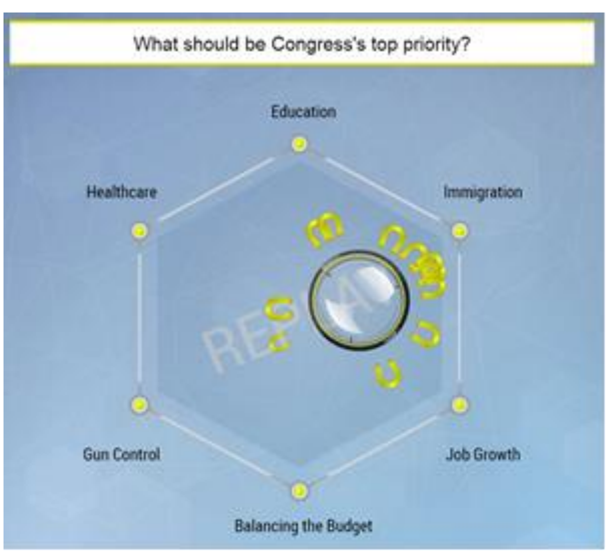

Fig 1. A human swarm comprised of user-controlled magnets.

It's important to note that users don't only vary the direction of their input, but also the magnitude by adjusting the distance between the magnet and the puck. This enables users to convey not only which choice they prefer most at a given time-step, but also their level of conviction in that choice. In addition, real-time predictive algorithms infer variations in user conviction based on the frequency of choice changes over time.

III. Swarms vs Votes

To compare the effectiveness of swarming and voting in the setting of group priories, 43 voting age Americans reviewed a list of 24 popular political objectives that have been debated during the 2016 Presidential and Congressional campaigns. From that full list, participants were asked to identify and rank which of the objectives they believed should be the top five priorities for the new President and Congress in 2017. This is a challenging task for any group, but to ensure high conflict in the prioritization process, the pool of participants were selected as a mix of Republican, Democrat, and Independent leaning voters.

In the first phase of the study, each participant completed an online survey to identify and rank their top five priorities. The surveys were performed independently and participants had no opportunity to communicate with one another about their selections. In the second phase of the experiment, the participants worked together as a unified real-time swarm (using the UNU swarming platform) to collectively rank their top five priorities. In this way, the 43 participants produced two different sets of priorities – one set generated individually by ranked survey and combined statistically, and one set generated by the group working collectively as a real-time swarm.

In the final phase of the study, participants were surveyed again and asked to individually reflect upon the two sets of priorities that were generated by the group, indicating (a) which set better reflected their personal views, and (b) which set better reflected the views of the full population. Participants were also asked to reflect on the process itself and indicate which methodology was more enjoyable.

IV. Results

As described above, a group of 43 voting age Americans, with mixed party affiliation, collectively produced two ordered sets of political priorities from a master list of 24 options. As provided in Figure 2 below, **List A** shows the top five priorities produced by the group working together as a unified swarm, while **List B** shows the top five priorities produced by aggregating the rankings provided on the individual surveys.

List A	List B
1. Provide Universal Healthcare	1. Provide Universal Healthcare
2. Create Jobs	2. Create jobs
3. Repair crumbling infrastructure	3. Eliminate poverty
4. Ensure fair elections	4. Defeat ISIS
5. Reduce college costs and student debts	5. Reduce wealth inequality between rich and poor

Fig 2. Ranked priorities produced by (A) swarm and (B) vote.

As shown in Figure 2, the sets of top-five priorities from the swarm and the survey had significant similarities and important differences. The key similarity is that first and second priorities on the lists – *Provide Universal Healthcare* and *Create Jobs* – were the same for both approaches. The next three priorities, however, were completely different for the two methodologies.

It is interesting to observe that priorities 3, 4 and 5 in **List A** (from the swarm) – *Repair crumbling infrastructure, Ensure fair elections*, and *Reduce college costs and student debt* – reflect concrete issues that could have immediate and direct impact on respondents' lives. In contrast, priorities 3, 4 and 5 in **List B** (from the survey) – *Eliminate poverty, Defeat ISIS*, and *Reduce wealth inequality between rich and poor* – address longer-term issues and are more removed from the day to day lives of participants. In fact, respondents commented that issues such as "eliminating poverty" were not realistic goals for any President and Congress to tackle, and yet it was highly ranked in the survey results. This suggests that when filling out the survey (which is an abstract individual exercise), respondents may have felt a personal need to express abstract altruistic goals, while participating in the collaborative swarm, where every ranking was a real-time exercise in group negotiation and compromise, users provided responses that were more grounded and realistic.

The findings raise the question of whether there is a bias towards "altruism" associated with surveys, as the individuals may feel they are being personally judged and therefore may be more inclined to answer the way "they think they're supposed to" as opposed to how they truly feel. Referred to generally as the Hawthorne Effect, this conforms with prior research that suggests altruistic bias can distort the participants true feelings when providing individual responses.[12] This raises an important question – does swarming mitigate this problem by having participants respond together as a synchronous group? To explore this, Part III of the research asked the 43 participants to reflect on each set of priorities.

In Part III, participants were asked to review both sets of priorities and independently complete an online questionnaire. Participants were asked to indicate which set of issues best represented their personal political priorities. As shown in Figure 3 below, **66%** of the respondents favored the list that resulted from the swarm, compared with 34% that favored the results of the survey.

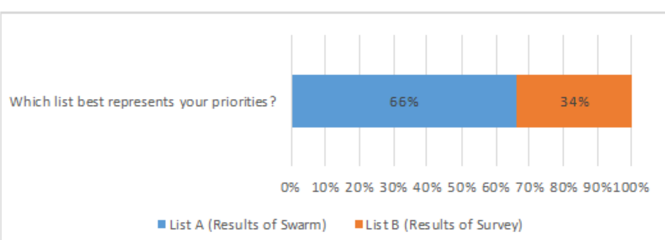

Fig 3. "Which list best represents your priorities", among those that expressed a preference (n=36)

Participants were also asked to reflect on which process (swarm or survey) they found to better represent their view of the group's overall priorities. As shown in Figure 4, **74%** of the respondents believed the swarm better represented the priorities of the group, with 26% that believed the results of the survey were more representative. This result suggest improved buy-in among the participants as three out of four participants believe the swarming process yielded a more accurate reflection of the group's collective will.

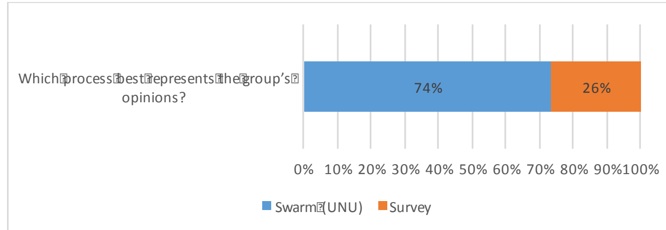

Fig 4. "Which process best represents the group's opinions?" among those that expressed a preference (n=34)

Lastly, the participants were asked to reflect on the process itself and indicate which method they found to be more enjoyable – prioritizing by survey, or prioritizing by swarm. As shown in Figure 5, **65%** of the respondents found the swarming process to be more enjoyable, while 35% preferred the survey. These results echo other research that indicates that swarming is a more pleasant process than taking surveys. This is an important result, for one of the primary logistical barriers to collecting data by survey is user aversion to the process.

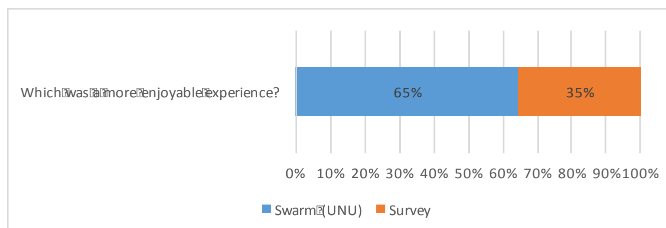

Fig 5. "Which was a more enjoyable experience?" among those that expressed a preference (n=34)

V. Discussion and Conclusions

As reflected by the results above, this study suggests that human swarming may be a more effective methodology for setting priorities among diverse groups than traditional polling. When participants compared the output of their swarm with the aggregate results of their survey responses, a significant majority reported that the swarm better represented both their personal priorities and their perceived opinions of the broader group. Two thirds of the subjects also found that participating in the unified swarm was more enjoyable than taking the survey.

With surveys and other forms of polling widely used by business organizations, market researchers, and news outlets to gauge the sentiments of the public, the benefits of swarming may have many applications. Surveys aggregate individual opinions as isolated snapshots, highlighting differences within the group rather than explicitly eliciting common ground. Surveys may also encourage participants to mask their true feelings vs what they believe they "should say". In contrast, the swarming process immerses respondents in a group decision dynamic that is specifically aimed at converging on common ground and results in clearer representations of overall group intent. Swarming may also mitigate the Hawthorne Effect by enabling respondents to feel part of a synchronous group rather than an exposed individual who risks being personally judged. And finally, swarming is perceived to be more enjoyable than surveys and is therefore more likely to get repeat engagements.

ACKNOWLEDGMENT

This work was directly supported by Unanimous A.I., the maker of the UNU platform for real-time human swarming. For more information about UNU, visit http://UNU.ai.

REFERENCES

[1] Armstrong, J.S. (2001), Principles of forecasting: a handbook for researchers and practitioners, Kluwer Academic Publishing, pages 417-439.

[2] Bottom, W., P., Ladha, Krishna., Miller, Gary J. (2002) "Propagation of Individual Bias through Group Judgement," Journal Of Risk Uncertainty, 25: 152-154.

[3] Seeley, Thomas D., Visscher, P. Kirk. Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making. Behavioral Ecology and Sociobiology 54 (5) 511-520.

[4] I.D. Couzin, Collective Cognition in Animal Groups, Trends Cogn. Sci. 13,36 (2008).

[5] J.A.R. Marchall, R. Bogacz, A. Dornhaus, R. Planque, T.Kovacs, N.R. Franks, On optimal decision making in brains and social insect colonies, J.R. Soc Interface 6,1065 (2009).

[6] Seeley, Thomas D. Honeybee Democracy. Princeton University Press, 2010.

[7] Rosenberg, L.B., "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[8] Rosenberg, Louis. "Artificial Swarm Intelligence vs Human Experts", Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE.

[9] Palmer, Daniel W., et al. "Emergent Diagnoses from a Collective of Radiologists: Algorithmic versus Social Consensus Strategies." Swarm Intelligence. Springer International Publishing, 2014. 222-229.

[10] Eberhart, Russell, Daniel Palmer, and Marc Kirschenbaum. "Beyond computational intelligence: blended intelligence." Swarm/Human Blended Intelligence Workshop (SHBI), 2015. IEEE, 2015.

[11] L. B. Rosenberg, "Human swarming, a real-time method for parallel distributed intelligence," Swarm/Human Blended Intelligence Workshop (SHBI), 2015, Cleveland, OH, 2015, pp. 1-7.

[12] Bardsley, N. (2008). Dictator game giving: altruism or artefact?. *Experimental Economics*, *11*(2), 122-133.