Hyperswarms: A New Architecture for Amplifying Collective Intelligence

Gregg Willcox
Unanimous AI
San Francisco, CA, USA
gregg@unanimous.ai

Louis Rosenberg
Unanimous AI
San Francisco, CA, USA
louis@unanimous.ai

Colin Domnauer
Unanimous AI
San Francisco, CA, USA
colin@unanimous.ai

Hans Schumann
Unanimous AI
San Francisco, CA, USA
hans@unanimous.ai

***Abstract*—Artificial Swarm Intelligence (ASI) is a recently developed method that enables networked human groups to converge on more accurate group forecasts, estimations, and decisions. While ASI can significantly amplify collective intelligence, the process struggles when too large of a majority supports an inaccurate view, even if their average confidence is low. Thus, a major goal of ASI research is to increase resilience to low-confidence majorities. This paper introduces a new ASI structure called a Hyperswarm that enables a confident minority to more readily sway an unsure majority. The approach involves dividing a population P into a set of overlapping sub-groups (H_1, H_2... H_P) such that each member only interacts with members of their subgroup. And because each subgroup overlaps multiple other subgroups, the local interactions quickly propagate throughout the full population. In this paper we simulate hyperswarms, showing that a confident minority can intelligently overcome a less confident majority, even when 70% of participants initially harbor the majority view. In addition, we explore a variety of hyperswarm design parameters and derive guidelines for future development.**

Keywords—Swarm Intelligence, Artificial Swarm Intelligence, Collective Intelligence, Wisdom of Crowds, Stigmergy, Hive Minds, Collaboration, Decision-Making Systems, Small-World Networks.

I. Introduction

In the field of Collective Intelligence (CI), it is well known that a large group can significantly outperform its individual members. For well over a century, a wide variety of aggregation techniques have been explored for harnessing the intelligence of human populations to enable more accurate decisions [1-3]. Artificial Swarm Intelligence (ASI) is a recent technique that's been shown to significantly amplify the decision-making accuracy of networked human groups using algorithms modeled on natural swarms. Unlike votes, polls, surveys, or prediction markets, which treat each participant as a separable datapoint for statistical aggregation, the ASI process treats each individual as an active member of a real-time dynamic system, enabling the full group to efficiently converge on solutions as a unified intelligence [4,5].

For example, a recent study conducted at the Stanford University School of Medicine showed that small groups of radiologists, when connected by real-time ASI algorithms, could diagnose chest X-rays with 33% fewer errors than traditional methods of aggregating human input [6,7]. Researchers at Boeing and the U.S. Army recently showed that small groups of military pilots, when using ASI technology, could more effectively generate subjective insights about the design of cockpits than current methods [8]. Researchers at California Polytechnic published a study showing that networked business teams increased their accuracy on a standard subjective judgement test by over 25% when deliberating as real-time ASI swarms [9,10]. And researchers at Unanimous AI, Oxford University, and MIT showed that small groups of financial traders, when forecasting the price of oil, gold, and stocks, increased their predictive accuracy by over 25% when using ASI method [11,12].

While the power of swarm-based systems to amplify group intelligence has been validated across many disciplines, ASI can struggle in highly lopsided groups where an overwhelming majority holds a similar view. Of course, the most common reason for a large majority to harbor a singular view is that it's correct. In such cases, traditional swarming converges accurately. The problem arises when a large majority suspects the wrong answer because of prevailing misconceptions. In such situations, it is difficult for a small but confident minority to overcome a large, misinformed majority, even if the minority has significantly higher confidence. To address this problem, the hyperswarm technique was developed with the goal of enabling the sentiments of a confident minority to propagate with greater resilience in the face of a larger but less confident majority.

II. From Swarms to Hyperswarms

A central feature of swarm-based decision making is that all participants interact in real-time, simultaneously adjusting their input as they converge together on group solutions. At every time-step, swarming algorithms modulate the influence of each participant based on their real-time behaviors. For example, participants who find themselves supporting a popular view may display behaviors that reflect increasing confidence, while participants who support an unpopular view may display resistance or capitulation. Swarming algorithms use this diverse range of behaviors to predict the conviction of each member and adjust their influence over time. In this way, collective decisions emerge that can significantly outperform traditional collective intelligence techniques.

While ASI can amplify the accuracy of group decisions in many contexts, it can fail when a large majority supports a singular view, even if the majority has low confidence. This is because the very presence of a large majority reduces the range of behaviors that participants are likely to express, as members of the majority have no reason to show uncertainty when their views are validated by so many others. Thus, members of a large majority who harbor low confidence are not driven to reveal their low confidence, further strengthening the majority. As a consequence, traditional swarms may converge on the majority

978-1-6654-0066-4/21/$31.00 ©2021 IEEE

opinion with little opportunity for a confident minority to influence the group.

To solve this problem, the hyperswarm method was developed in which each member of a population is exposed only to a random subset a subset of other members, thereby driving a more diverse range of behaviors. This is achieved by dividing the full population **P** into a series of overlapping subgroups (**H_1, H_2... H_P**) such that each participant is only exposed to the changing sentiments of their subgroup. And because the subgroups are defined with overlapping membership, interactions quickly propagate across the full population. In this way, a hyperswarm is a closed-loop system that can converge on a unified solution but does so with greater diversity of behavior than traditional swarms [14].

The hyperswarm method is particularly useful when a large majority supports a similar view, as traditional swarming fails to reveal hidden uncertainty within the majority. By dividing the population into a set of overlapping subgroups of size S, the hyperswarm structure creates a series of overlapping swarms, each with unique distributions of opinions. By statistical chance, some subgroups will have balanced or inverted distributions as compared to the full population. This will motivate behaviors within some subgroups that reveal uncertainty that would otherwise be hidden. And because of the overlapping structure, shifting views within the majority can propagate across the full population, enabling a confident minority to sway an uncertain majority.

To make the hyperswarm concept more concrete, let's consider a twenty member population (P=20) split into 20 overlapping subgroups H, each of size S=7. As shown in Figure 1, this hyperswarm structure is modeled after a simple ring network in which each participant P_1, P_2... P_{20} is exposed to the real-time sentiments of S-1 other members through the series of unique but overlapping subgroups.

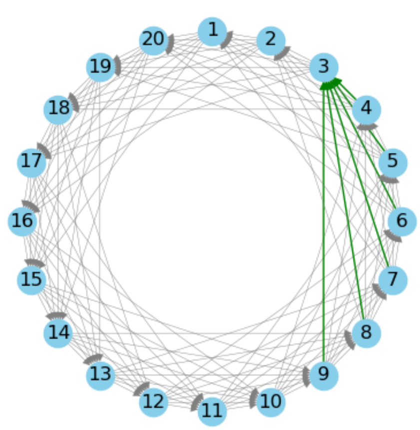

Fig. 1. HyperSwarm Influence Diagram for P=20, S=7. Green arrows show that member P3 is exposed to six other members. Grey arrows show all other connections among subgroups

Looking closely at Figure 1, we see that member **P_1** is exposed to subgroup **H_1** of members **P_2, P_3... P_7**. Shifting over one step, member **P_2** is exposed to subgroup **H_2** of members **P_3, P_4... P_8**. This continues around the ring structure and wraps around at **P_{20}**. In this way, each member **P_n** is exposed only to the real-time sentiments of a unique subgroup **H_n**. This means that each member **P_n** will react not to full population, but instead to their unique subgroup **H_n**. In other words, all members interact in synchrony but each is driven by potentially different sentiment distributions. And because all subgroups overlap, sentiments are quickly propagated throughout the full population, enabling all members to influence all others. Hence, a hyperswarm is a single closed-loop dynamic system in which behaviors are inspired locally but propagate globally.

III. Simulation Study

To explore the hypothesis that hyperswarms enable confident minorities to sway unsure majorities, a simulation was created to test the hyperswarm structure on a theoretical question with two possible answers, **A** and **B**. The simulation was seeded with a population of size **P** of which the majority fraction **M** initially supports answer **A** and a minority fraction **(1-M)** supports answer **B**. Each simulated member **P_n** of the population is assigned a confidence level **C_n** in their initial answer and is exposed to a unique subgroup of other members **H_n**. At each timestep, every member is assigned a simulated probability of switching answers depending on their simulated confidence in that answer and whether their answer is shared by a majority of other members in the unique subgroup **H_n** that they are exposed to at that timestep.

The simulation randomly generates a seed population of **P** members: **M** fraction of whom (e.g. 40%) support the minority answer initially and **(1-M)** fraction of whom support the majority answer initially. Each majority-respondent and minority-respondent are assigned a randomly-sampled confidence from a distribution which has average confidence **C_A** and **C_B** respectively, where **$C_B > C_A$** such that the minority is more confident than the majority. The distribution of confidences for groups **A** and **B** is a skewed normal distribution with means at **C_A**, **C_B** and a skewness of **5** and **-5** respectively, shown in Figure 2 for **C_A=0.4**, **C_B=0.6**. In general, any distribution could be used in which the minority population has on average higher confidence than the majority population: the authors have used simple normal curves and Dirac delta functions as distributions in this simulation and have observed essentially identical results.

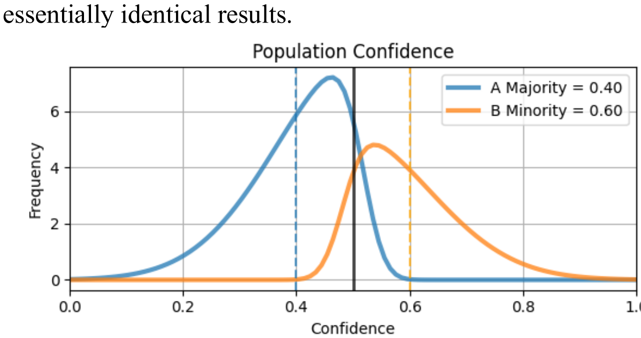

Fig. 2. Distribution of individual confidences

Once seeded, the simulation employs a set of simple behavioral rules to approximate how real participants would act based on their initial answer, confidence level, and the distribution of opinions they are exposed to in their random subgroup. These simulated rules allow us to estimate what percent of participants are likely to switch answers (or resist switching) when they find themselves within the minority of their unique subgroup. These rules were defined as follows:

Rule 1: If a user is in a local majority, they don't switch. *(i.e. if they're in the majority, there's no motive to switch).*

Rule 2: If a user is in a local minority and has not yet switched from their original view, they switch with probability = (1 - confidence). *(i.e. higher confidence in their original view means lower probability they switch).*

Rule 3: If a user already switched but finds themselves in a local minority, they switch back with probability = confidence. *(i.e. higher confidence in their original view means higher probability that they switch back).*

A single simulation trial that evolves according to these rules is shown below in Figure 3. It was initialized with a simulated population P=25 such that 60% of the population initially supported answer A and 40% initially supported answer B (M=0.4). Furthermore, in this simulation supporters of A have an average confidence CA=0.4, supporters of B have an average confidence CB=0.7, and the resulting confidence differential is CD=0.3. Each individual's subgroup $\mathbf{H_n}$ contains 6 other individuals (S=7).

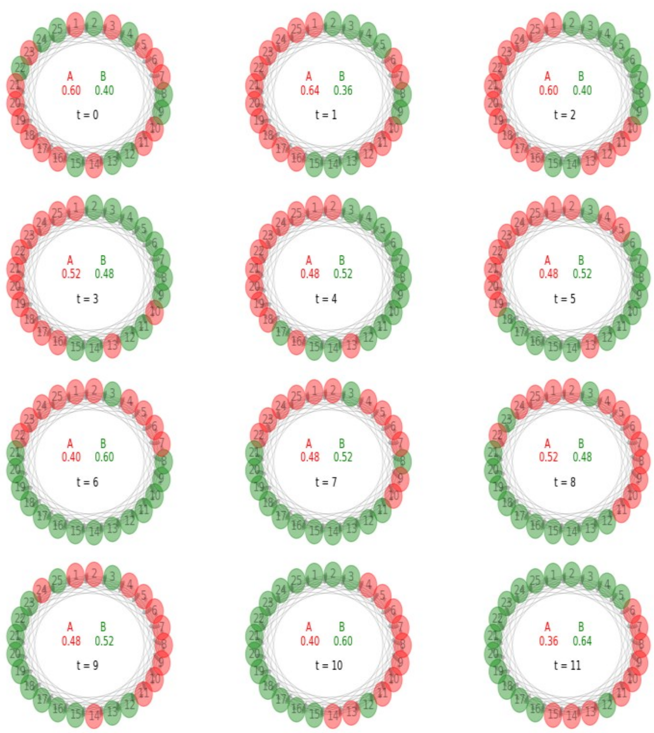

Figure 3: Hyperswarm evolution over 12 time-steps.

Each user in this simulation (labeled 1, 2…. 25) is colored red if they answer **A** at each timestep, and green if they answer **B**. Initially, the distribution of views is randomly distributed across network structure, but as the simulation progresses and information is shared, local regions form that favor answer A (red) or B (green). These local regions of support propagate around the hyperswarm, spreading clockwise due to the asymmetry of the subgroup connectivity structure in this simulation. Since **A**-responding participants have lower confidence and are therefore more likely to switch answers than **B**-responding participants, the fraction of respondents choosing **B** grows over the course of the swarming process. In the particular trial shown, the effect was so pronounced that the distribution of support flipped from 60% of the population initially favoring answer **A** (@ t=0) to 64% of the population favoring answer **B** at the end of the swarming process (@ t=12)

Using these rules, the simulation was run 1000 times, each with a randomly selected seed population of 60% favoring A and 40% favoring B, with $C_A < C_B$. The objective was to determine the probability that a population with a significant majority initially supporting A will switch its majority support to B over the duration of the swarming process. In this way, the simulation allows us to estimate the effectiveness of the proposed hyperswarm structure and assess the conditions under which the hyperswarm technique enables a higher confidence minority to overcome a lower confidence majority.

To provide a baseline for comparison, the simulation was also run using a traditional flat structure in which all members are exposed to all others in real-time. To emulate the flat swarm, we ran the simulation with a population P=25 and subgroup size S=25 (so all members are exposed to all others). We then set the population parameters to a define 60% majority supporting answer A and a 40% minority supporting B, with average confidence in B higher than average confidence in A (i.e. M=0.4, C_A=0.3, C_B=0.7, CD=.4). Running 1000 trials of this configuration, we found that the flat structure never converged on the minority opinion. That's because none of the simulated users were ever motivated to switch from the majority A answer (regardless of confidence level) as everyone was influenced by the 60% global majority supporting A. Obeying rule 1 above, only minority-responding participants were ever motivated to switch (depending on exposure and confidence). The evolution of this system is shown in figure 4a.

Next, we tested the hyperswarm structure with subgroups of size S=7, also with P=25, M=0.4, C_A=0.4, C_B=0.7 and CD=0.3, and found substantially improved performance. As shown in Figure 4b, the 5000 simulation trials confirm that after 8 simulated time steps, 56% of the randomly generated populations switched their majority preference from A to B.

This suggests that the hyperswarm structure, when ideally simulated and seeded with initial conditions within certain bounds, can enable a more confident minority to overcome a less confident majority. The question remains, could other connection structures better enable a confident minority to overcome an unsure majority?

In other words, how does the connection structure of an ASI hyperswarm impact the probability that the views of confident minority can overcome a substantially larger but less confident majority? To investigate this, we created five different connection structures and explored their impact on the minority group's ability to overcome a less confident majority in the face of varying conditions.

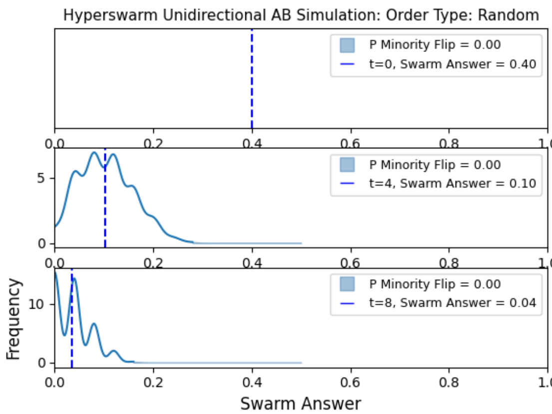

Figure 4a: Fully Connected Swarm never flips a majority.

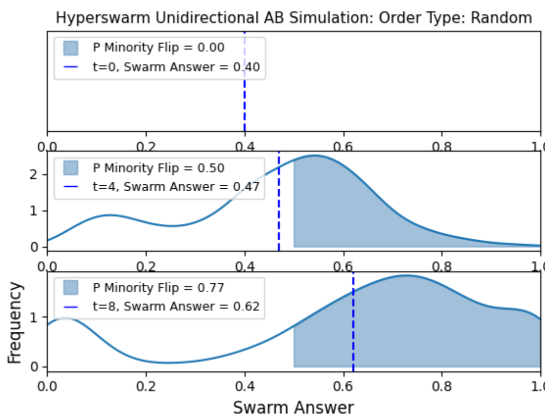

Figure 4b: After 8 time steps, hyperswarms enable 77% of simulated populations flip to support the initial minority view.

IV. Connection Structures Influence Outcomes

The four connection structures we study are: (i) Unidirectional, (ii) Bidirectional, (iii) Random, and (iv) Unidirectional Small World. An example of each is given in Figure 5 for a population size of P=10 and a subgroup size of S=5.

The Unidirectional structure is the same as the structure in Figure 1, and studied in [14], where member $\mathbf{P_j}$ is exposed to members $\mathbf{H_j=\{P_{j+1}, P_{j+2}, \ldots P_{j+S-1}\}}$. In the Bidirectional structure, member $\mathbf{P_j}$ is exposed to members $\mathbf{H_j=\{P_{j-(S-1)/2}, \ldots P_{j-2}, P_{j-1}, P_{j+1}, P_{j+2}, \ldots P_{j-(S-1)/2}\}}$, where S is constrained to be odd to ensure the structure is symmetric.

The Unidirectional Small World structure is created using the Watts-Strogatz approach [13] by first creating a Unidirectional connection structure, and then rewiring 25% of connections randomly—without replicating existing connections and ensuring that each member still is exposed to S-1 other members. This structure was created to examine whether minimizing the path length between hyperswarm members would speed up the rate at which confident minority opinions can propagate through the population.

The random structure is initialized purely randomly, again such that each user is exposed to exactly S-1 other members. Next, we explore how the connection structure between participants impacts the likelihood that a confident minority can overcome a weak majority across a wide variety of initial conditions, including the impact of population size **P**, subgroup size **S**, the population fraction that supports the minority view **M**, and the difference in average confidence between members of the minority and majority **CD**.

To do so, we ran this simulation with P=10, S=5, M=0.4 and CD=0.4 for each of the connection structures. The fraction of hyperswarms in which the initial minority answer becomes the global majority answer by timestep 5 and timestep 20 is reported in Table 1 for each of these cases. We expected that the Bidirectional or Small World models would perform best, since they enable two-way communication that has the potential to propagate across the network quickly. However, we observed that the Unidirectional structure outperformed the other connection structures by a large margin—flipping 48% of majorities by the 5th timestep and 54% of majorities by the 20th timestep.

Table 1: Probability of Minority Flip by connection structure.

	Probability of Flip	
Timestep	T=5	T=20
Unidirectional	**48%**	**54%**
Bidirectional	17%	18%
Random	29%	30%
Unidirectional Small World	36%	36%

To examine the sensitivity of these results to the parameters that we chose for this simulation, we reran the simulation for a range of parameter values—first, by varying the population size **P** and subgroup size **S** (Figure 6), and then by varying the initial minority fraction **M** and confidence difference between groups **CD** (Figure 7).

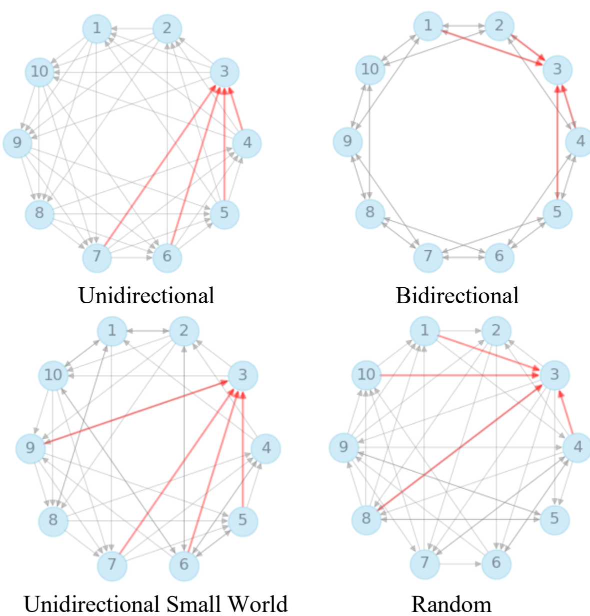

Figure 5: Hyperswarm Influence Diagram for P=10, S=5, for each of the five studied connection structures. Red arrows show that member P_3 is exposed to four other members. Grey arrows show all other connections among subgroups.

Figure 6 shows that the Unidirectional structure outperforms the Unidirectional Small World structure for a wide range of population sizes P and subgroup sizes S. This result holds for the Bidirectional, and Random structures to an even greater degree. An important case shown on these heatmaps is when the subgroup size approaches the population size (P=S), which is equivalent in all cases to a global connection structure. In these global connection structures, all individuals can see all other individuals, and the global majority is known to all participants, so there's never a chance at flipping the minority answer. In this way, reducing the subgroup size and creating a hyperswarm enables the group to aggregate their insights by their confidence, rather than their initial answer distribution.

Figure 7 shows us that the Unidirectional structure again outperforms the Unidirectional Small World structure for a wide range of initial minority fractions **M** and confidence differences **CD** between the minority and majority populations. These results hold for the Bidirectional and Random structures also. There are two general trends that we can discover through Figure 7: first, across all these connection structures, the lower the initial fraction of members responding **B** (the initial minority answer), the lower the likelihood that **B** will overcome **A** regardless of the difference in confidence between the two groups. This is to be expected—the fewer people support **B**, the fewer subgroups will have **B** majorities. Second, the higher the group's average confidence in **B** relative to the average confidence in **A**, the more likely the minority (**B**) will overcome the majority (**A**). Both trends are to be expected and fit with the author's intuition about swarm-like systems in real groups.

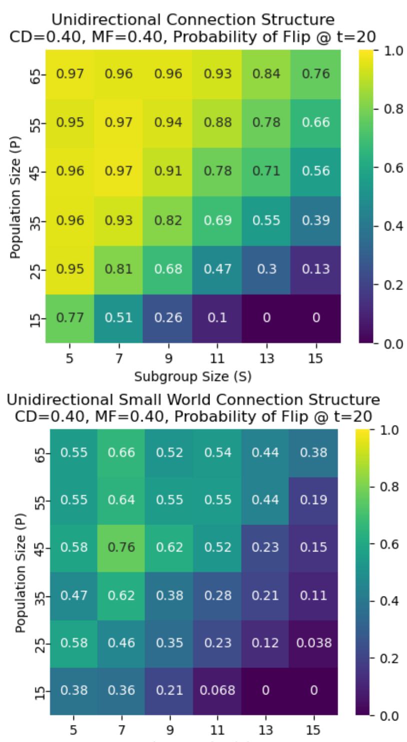

Figure 6: Probability of minority flip after 20 timesteps for the Unidirectional and Unidirectional Small World connection structures across a range of population and subgroup sizes.

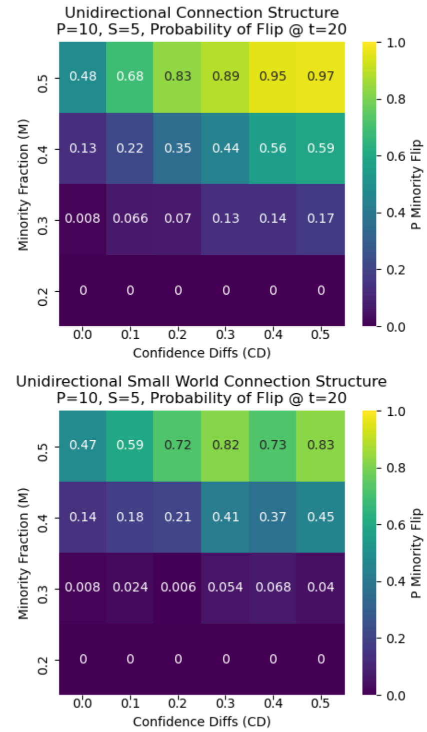

Figure 7: Probability of minority flip vs minority fraction **M** and confidence difference **CD** after 20 timesteps for the Unidirectional and Unidirectional Small World connection structures.

To gain better intuition into why the Unidirectional structure performs so well, we created the visualization shown in Figure 8 in which each row represents a time-step of the evolving distribution of opinions across the full population. At time **t=0** we initialize each simulation with the same randomly-generated seed population. Each simulated user supports either answer **A** (**shown in purple**) or answer **B** (**shown in yellow**) and harbors a randomly assigned confidence level in their initial view. We then visualize the evolution of the system across 20 timesteps, showing the changing distribution of views over time.

In all four graphs, we use P=50, S=7, M=0.4, CD=0.4; the initial majority answer is **A** (purple), while the initial minority answer is **B** (yellow). To compare the Unidirectional structure with the other structures, we used the same seed population in four visualizations (a), (b), (c), and (d) shown above but varied the connectivity model such that (a) is the Unidirectional structure, (b) is the Bidirectional structure, (c) is the Unidirectional small world structure, and (d) is the Random connectivity structure.

Qualitatively, the Unidirectional structure has two wavefronts that both propagate leftwards around the population—a fast wavefront where **A**-responding participants are exposed to a near-unanimous response of **B** and quickly end up converting, and a slower wavefront where **B**-responding participants are exposed to a unanimous response of **A** and slowly decide to convert. In the end, **B** wins because people convert to **B** faster than they convert to **A**. This is only possible because participants are exposed to a different set of people than they broadcast to: the flow of information is unidirectional.

In the Bidirectional structure, we observe no such information flow: instead, the Bidirectional structure creates clusters of uniform answers that do not propagate in any direction. This occurs because individuals on all edges of these clusters see themselves as a part of the majority of their subgroup—their subgroup is evenly split and their answer tips the tie into a majority. This preserves the diversity of initial answers—a unanimous decision is unlikely—and does not lead to consensus forming for the highest-confidence group answer.

Finally, in the Unidirectional Small World and Random structures, each subgroup is exposed to a greater variety of opinions from around the global population, leading to more nonlocal information propagation and, paradoxically, a slower global transition to the minority answer.

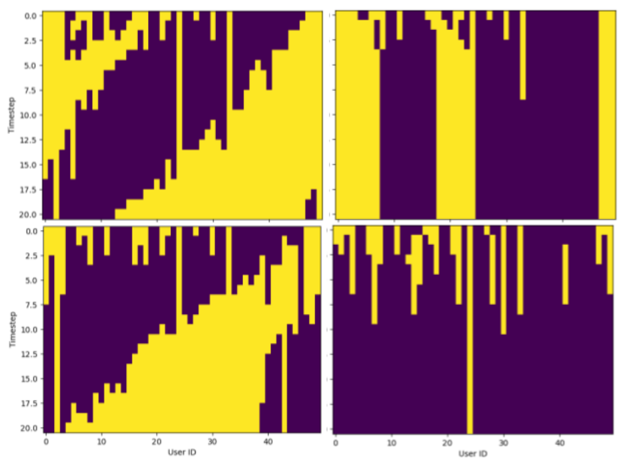

Figure 8: Heatmaps of user responses for the same initialization by connection structure (clockwise from top left): Unidirectional, Bidirectional, Unidirectional Small World, and Random.

V. Connection Structure Influences Subgroup Agreement

Why does nonlocal information propagation have such an impact on the final outcomes of this system?

To better understand this, we examine how each structure exposes members to beliefs that agree or conflict with their own over time. We ran 100 hyperswarms with P=50, S=7, CD=0.4, and M=0.4 as before, and calculate the fraction of participants who agree with the majority of their subgroup at each timestep in figure 9.

As we expect, there are minimal differences initially between connection structures: upon initialization at T=0, about two thirds of users agree with their subgroup majority regardless of the connection structure used. Over time, however, the Unidirectional structure exposes more individuals to conflicting options than any of the other network structures. This creates more opportunities for each individual to be challenged by opposing views, enabling the group to more often reach the answer it is collectively most confident in, rather than simply converging on the answer that most people initially support.

We also include a Bidirectional Small World structure here, to show that the effect of the 25% of random connections have the impact of bringing the Experienced Agreement curve closer to that of a random connection structure.

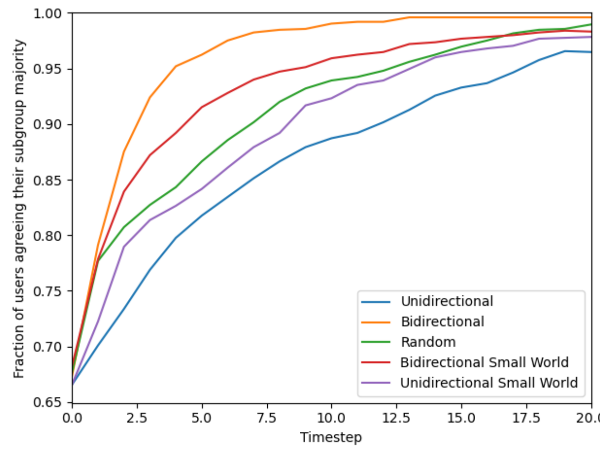

Figure 9: Fraction of participants who agree with the majority sentiment of their Hyperswarm subgroup by connection structure.

VI. Discussion

In this paper we simulated hyperswarm dynamics, showing that a confident majority can sway a less confident minority in real time, even when 70% of the participants initially harbor the majority view. We also demonstrated that this is not the case for traditional swarms in which participants can see all other participants.

We then explored the impact of the chosen connection structure on these results by comparing the likelihood that a confident minority could flip a less confident majority when connected by Unidirectional, Bidirectional, Random, and Small World structures. We found that Unidirectional structures enabled confident minorities to prevail across a substantially wider range of trials than did Bidirectional structures, whether they had symmetric connections, small world connections, or random connections.

While this result may seem counter-intuitive, it does conform with many swarming structures found in nature such as schooling fish and flocking birds wherein each member is influenced primarily by members directly ahead or adjacent, but not by members behind them. In other words, biological Swarm Intelligence has evolved using Unidirectional structures with great success.

Why is this? Based on the simulation results herein, we propose both low-level and high-level explanations. At a low level, we compared the switching behaviors of individuals in different connection structures and found that the Unidirectional structure propagates information around the hyperswarm in waves, such that the higher-confidence answer in the population propagates faster and ultimately converts members more rapidly than the lower-confidence answer. At a high level, we demonstrated that the Unidirectional structure maintains a higher level of discord in the hyperswarm over time: fewer people are exposed to subgroup majorities that agree with them at each timestep. We argued that this leads to a global decision-making process that tends to select the answer in the population that has the highest average confidence, rather than the answer that people most agree upon at the outset.

Our aim in this work is to shed light onto what collective intelligence structures enable groups to reach optimal decisions. Often, a confident minority of people in a group can accurately select an optimal answer, while a less-informed majority of people select a suboptimal answer. In classical swarms, the less confident majority often overpowers the confident minority; our results show that this effect can be mitigated with the appropriate design of hyperswarms and suggests that local, unidirectional information flow is important to hyperswarm construction.

While these simulations are promising, the results are limited in that we only investigated a single behavioral rule—one in which a user who finds themselves agreeing with their subgroup majority never switches—so future work will investigate other behavioral rules where this is not the case. There are, of course, many other network structures that could be explored, including network-like and multidimensional models (e.g. toroidal networks rather than ring networks). Further, future work will seek to validate that hyperswarm structures enable these effects in real human populations.

REFERENCES

[1] De Condorcet, Nicolas. Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. Cambridge University Press, 2014.

[2] Boland, P. J. (1989). Majority systems and the Condorcet Jury theorem. Statistician 38, 181. doi:10.2307/2348873

[3] Larrick, Richard P., and Jack B.Soll. "Intuitions about combining opinions: Mis-appreciation of the averaging principle." Management science 52.1 (2006): 111-127.

[4] Rosenberg, Louis. Artificial Swarm Intelligence, a human-in-the-loop approach to A.I., Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI 2016), Phoenix, Arizona. Pg 4381-4382. AAAI Press, 2016.

[5] Rosenberg, Louis, "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659 ECAL 2015, York, UK. MIT Press 2015, ISBN 978-0-262-33027-5

[6] Halabi, Safwan., et. Al. "Radiology SWARM: Novel Crowdsourcing Tool for CheXNet Algorithm Validation", SiiM Conference on Machine Intelligence in Medical Imaging, 2018.

[7] Rosenberg, L, Willcox, G., Halabi, S., Lungren, M, Baltaxe, D. and Lyons, M. "Artificial Swarm Intelligence employed to Amplify Diagnostic Accuracy in Radiology," 2018 IEEE 9th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), Vancouver, BC, 2018.

[8] Befort, K., Baltaxe, D., Proffitt, C., & Durbin, D. (2018). "Artificial Swarm Intelligence Technology Enables Better Subjective Rating Judgment in Pilots Compared to Traditional Data Collection Methods." Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 62(1), 2033–2036.

[9] Askay, D., Metcalf, L., Rosenberg, L., Willcox, D., "Enhancing Group Social Perceptiveness through a Swarm-based Decision-Making Platform." Proceedings of 52nd Hawaii International Conference on System Sciences (HICSS-52), IEEE 2019.

[10] Metcalf L, Askay DA, Rosenberg LB. Keeping Humans in the Loop: Pooling Knowledge through Artificial Swarm Intelligence to Improve Business Decision Making. California Management Review. 2019;61(4):84-109.

[11] Rosenberg, Louis. Pescetelli, Nicollo, Willcox, Gregg., "Artificial Swarm Intelligence amplifies accuracy when predicting financial markets," 2017 IEEE 8th Annual Ubiquitous Computing, Electronics and Mobile Communication Conference (UEMCON), New York City, NY, 2017, pp. 58-62.

[12] H. Schumann, G. Willcox, L. Rosenberg and N. Pescetelli, ""Human Swarming" Amplifies Accuracy and ROI when Forecasting Financial Markets," 2019 IEEE International Conference on Humanized Computing and Communication (HCC), Laguna Hills, CA, USA, 2019, pp. 77-82, doi: 10.1109/HCC46620.2019.00019.

[13] Watts, D. J.; Strogatz, S. H. (1998). "Collective dynamics of 'small-world' networks" (PDF). Nature. 393 (6684): 440–442. Bibcode:1998Natur.393..440W. doi:10.1038/30918. PMID 9623998

[14] Rosenberg, L., Domnauer, C., Willcox, G., Schumann, H. "From Swarms to Hyperswarms: A New Methodology for Amplifying Group Intelligence" ACM Collective Intelligence Conference 2021, Copenhagen, Denmark – June 29, 2021.