

Emergent Collective Sensing in Human Groups

Peter M. Krafft (pkrafft@mit.edu)*, Robert X.D. Hawkins (rxdh@stanford.edu)†, Alex “Sandy” Pentland (pentland@mit.edu)‡, Noah D. Goodman (ngoodman@stanford.edu)†, Joshua B. Tenenbaum (jbt@mit.edu)*

*MIT Computer Science and Artificial Intelligence Laboratory, †Stanford Department of Psychology, ‡MIT Media Lab

Abstract

Despite its importance, human collective intelligence remains enigmatic. We know what features are predictive of collective intelligence in human groups, but we do not understand the specific mechanisms that lead to the emergence of this distributed information processing ability. In contrast, there is a well-developed literature of experiments that have exposed the mechanisms of collective intelligence in nonhuman animal species. We adapt a recent experiment designed to study collective sensing in groups of fish in order to better understand the mechanisms that may underly the emergence of collective intelligence in human groups. We find that humans in our experiments act at a high level like fish but with two additional behaviors: independent exploration and targeted copying. These distinctively human activities may partially explain the emergence of collective sensing in our task environment at group sizes and on times scales orders of magnitudes smaller than were observed in fish.

Keywords: collective intelligence; distributed cognition; social cognition; social computation; online experiments

Introduction

Many common examples of collective behavior illustrate apparent failures of collective intelligence. Mobs, market panics, and mass hysteria draw attention because of their perceived irrationality and drastic consequences. However, the successes of collective intelligence are as remarkable as the failures are devastating. The richness of human culture, the incredible pace of our technological developments, and the gradual progression of our scientific understanding of the universe stand out as both distinctively human and heavily reliant on the emergent behavior of the interactions of many individuals. Even at a less grandiose level, humans regularly agree to work together to accomplish tasks that no individual could accomplish alone via dynamic cooperative interactions that are hypothesized to be uniquely human (Tomasello, 2014). Yet little is known about the specific mechanisms underlying these synergistic processes of self-organization.

Many mathematical and computational models of collective behavior have been proposed. However, as a result of the logistical difficulties in conducting real-time human experiments involving multiple participants, and as a result of a broader lack of data analysis aimed at understanding collective behavior, the quantitative study of collective behavior has largely lacked an empirical basis. Recently, researchers have begun conducting carefully controlled laboratory experiments to test and refine models of collective behavior (Couzin, 2009; Goldstone & Gureckis, 2009). Yet many of these experiments, with some notable exceptions (Goldstone, Roberts, Mason, & Gureckis, 2008; Kearns, 2012), have been conducted using nonhuman animal subjects. We are therefore quickly developing a better understanding

of the collective behavior of ants (Pratt & Sumpter, 2006), bees (Seeley & Buhrman, 1999), cockroaches (Amé, Halloy, Rivault, Detrain, & Deneubourg, 2006), and fish (Ward, Sumpter, Couzin, Hart, & Krause, 2008), but our empirically-grounded quantitative understanding of human collective behavior remains limited.

In the present paper we harness recent technical advances in running real-time, networked experiments on the web (Hawkins, 2014) to develop and test a model of collective human behavior. We build on a recent experiment designed to study the collective behavior of a particular species of fish (Berdahl, Torney, Ioannou, Faria, & Couzin, 2013) that is one of the clearest illustrations of collective intelligence in a nonhuman animal group. In this previous experiment, the researchers studied a type of fish called the golden shiner that prefers to spend time in dark areas of the water, presumably to avoid predators. Aware of this natural propensity of the fish, the researchers projected time-varying spatially correlated light fields into a fish tank. The researchers then studied the effectiveness of the fish at finding the darker areas of the tank as a function of the number of fish participating in the task. The researchers found that average group performance increased significantly as a function of group size, and they identified two simple behavioral mechanisms driving this improvement: First, individual fish tended to move more slowly in darker areas. Second, individual fish also tended to turn towards conspecifics. The researchers argued that the combination of these mechanisms generated an emergent collective gradient sensing ability in groups of fish that had been absent in individual fish.

This experiment provides a beautiful example of a higher level of intelligence at the group level emerging from minimal intelligence at the individual level. However, while these simple mechanisms did appear to give rise to surprisingly effective group behavior, they only lead to substantial gains in performance for large groups of 50 or more fish. In contrast, we expect humans in a similar task to show significant gains with much smaller group sizes. In particular, we expect that humans should be able to make use of theory of mind, an ability to draw inferences about the underlying mental states of other players, to better utilize social information in a similar environment.

To elucidate these potential differences between humans and fish, we developed a version of the gradient-sensing task for human participants. Specifically, we recreated the environment used by Berdahl et al. (2013) as an online real-time multi-player game. In our experiment, participants controlled avatars in a virtual world. Every location in this world cor-

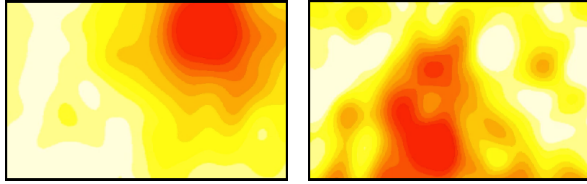


Figure 1: Example score fields from the low noise (left) and medium noise (right) conditions at particular points in time. Red areas indicate higher scoring areas.

responded to a score value that changed over time, and participants were awarded bonuses proportional to their cumulative scores in the game. The score of a player at a particular point in time was simply determined by the location of that player in the virtual world. Our incentives for participants to achieve high scores were designed to parallel the fishes' preferences for darker areas in their environment. The players either played alone or in groups of varying sizes. We used this virtual environment to investigate how the gradient-tracking performance of human groups changed as group size increased, and to attempt to identify behavioral mechanisms underlying collective sensing in human groups.

Methods

Participants We recruited 563 unique participants from Amazon Mechanical Turk to participate in our experiment. All participants were from the United States. After excluding 72 participants due to inactivity or latency, and 6 others for disconnecting in the first half of the game, we were left with usable data from 437 participants in 224 groups. These groups ranged in size from one to six individuals. Since we were only able to collect one group of size six, we ignored this group in our analysis.

Stimuli The game scores of the participants in our experiments were determined by underlying "score fields". These score fields consisted of 480×285 arrays of score values for each 125ms time interval in our game. We generated these score fields using the method reported by Berdahl et al. (2013). First, a "spotlight" of high value was created that moved in straight paths between uniformly randomly chosen locations. This spotlight was then combined with a field of spatially correlated noise. This procedure yields a complex landscape with many transient maxima and a single persistent time-varying global maximum.

We manipulated the weighting between the noise field and the spotlight to generate different task conditions. We used two weight values, corresponding to the "low" and "medium" noise levels reported by Berdahl et al. Examples of score fields are shown in Figure 1. 113 individuals (63 groups) were assigned to the low noise condition and 324 individuals (161 groups) were assigned to the medium noise condition. To decrease variability and increase statistical power, we generated only four distinct score fields per noise level, so multiple groups experienced the same fields. To discourage inactivity,

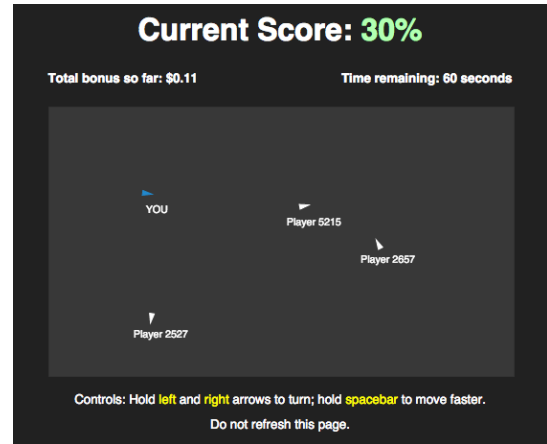


Figure 2: A screenshot of the interface that participants saw. The score displayed corresponds to the value of the score field at the location that the player's avatar is occupying.

players were awarded a score of zero, corresponding to zero bonus, if their avatars were touching a wall.

We attempted to give our participants perceptual and motor capabilities in this environment similar to the capabilities that Berdahl et al. observed in the fish in their experiments. In terms of perception, we restricted the information that participants received about the underlying score fields in the games. We allowed participants to see only the scores at their avatars' locations. The participants could *not* see the scores that other players were obtaining or the scores at any other locations besides their own. However, the positions, directions, and speeds of all other players were visible to each player. All of this information was updated in real-time every eighth of a second. A screenshot of the interface we used for the game is shown in Figure 2.

Players controlled their avatars using the left and right arrow keys to turn (at a rate of 40° per second) and could hold the spacebar to accelerate. The avatars automatically moved forward at a constant velocity of 136 pixels per second whenever the spacebar was not depressed. The avatars instantaneously increased to a constant velocity of 456 pixels per second for the duration of time that the spacebar was held down. We chose these speed values to match the speeds that Berdahl et al. reported observing in their fish, and we also matched the playing area dimensions and game duration to the parameters of their experiments. Each participant played in a single continuous game lasting for 6 minutes.

Procedure After agreeing to participate in our experiment, participants were presented with a set of instructions. These instructions simply described the mechanics of the game. The participants were not informed about the nature of the underlying score fields and were not encouraged to work together. After successfully completing a comprehension test, participants were then redirected to a waiting room. In the waiting room participants would wait for up to 5 minutes or until a

pre-assigned number of other players joined the game. While in the waiting room, participants could familiarize themselves with the controls of the game. Players were not shown any score in the waiting room unless the participant was against a wall, in which case the displayed score would change from a dashed line to a red “0%”. We found no evidence for the amount time a player spent in the waiting room having any effect on individual performance in the game (linear regression slope $1.993e-06$, with 95% confidence interval $[-1e-05, 1.4e-05]$). As in the actual game, participants in the waiting room would be removed for inactivity if the player’s browser was active in another tab for more than 15 seconds or if the player’s avatar was unmoving against a wall for 30 seconds. We also removed players if their ping response latencies were greater than 125ms for more than 36 seconds. We paid participants 50 cents for reading our instructions, and the participants could receive a bonus of up to \$1.25 during the six minutes of gameplay. Final bonuses were computed to be the players’ cumulative scores divided by the total length of the game times the total possible bonus. Following the current convention on Mechanical Turk, each participant was also paid 12 cents per minute for any time spent in the waiting room, minus any time that player spent against a wall. These numbers were chosen so that the participants were expected to receive at least the U.S. federal minimum wage of \$7.25 per hour for the totality of their time active in the experiment.

We implemented this experiment using the MWERT framework (Hawkins, 2014). The MWERT framework uses a set of recent web technologies capable of handling the challenges of real-time, multi-player web experiments, including Node.js, the Socket.io module, and HTML5 canvases. Since MWERT was originally used for two-player games, we had to extend the MWERT framework in several ways to handle the challenges posed by hosting larger groups of players.

Results

We find that group size is positively related to group performance in this game in the low noise condition. However, we find that there was little effect of group size in the medium noise condition. Average performance as a function of group size in each of these conditions is shown in Figure 3. A linear regression on the individuals in the low noise condition produces a significant positive slope of 0.0238 and a 95% confidence interval (CI) of $[0.006, 0.041]$. A linear regression on the individuals in the medium noise condition produces a marginally significant positive slope of 0.0068, 95% CI $[-0.001, 0.015]$, and this trend is weakened substantially with the inclusion of the single 6-person group. Moreover, the marginally significant result in the medium noise condition is driven entirely by the effect of group size in one of the four distinct score fields we used. This particular score field displays a significant effect of group size with a positive slope of 0.0306, 95% CI: $[0.015, 0.046]$, while none of the others do. Qualitative inspection revealed that this particular score field seemed to share spatial properties more similar to

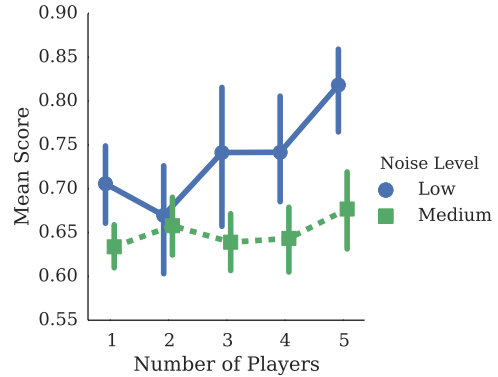


Figure 3: Mean performance as a function of group size in the low and medium noise levels. Error bars are 95% bootstrap confidence intervals using the group as the primary bootstrap unit. All points are averages over at least two groups. This plot excludes the single group we were able to collect of size six. Including this group weakens the trend in the medium noise condition.

the low noise score fields, which may explain the strength of the effect in that particular score field. Overall these results indicate that larger groups do tend to perform systemically better on our task than those in smaller groups, at least in the low noise condition.¹

In order to understand the factors that may have contributed to the increases in performance achieved by larger groups in the low noise condition, we examine the behavior of the players in our games. We assume a simple state-based representation of player behavior. We then attempt to identify how participants choose to occupy particular behavioral states at each point in time, and we examine the relationship between the players’ decisions to occupy particular states and the performance of those players. Specifically, we assume that at any particular point in time a player is either “exploring”, “exploiting”, or “copying” (see Rendell et al., 2010, for a similar classification). Conceptually, a player is exploring if that player is looking for a good location to exploit, a player is exploiting if that player has found a location where the player wants to remain, and a player is copying if that player is intending to move to the location of another player.

We empirically determine the state of each player at each point in time using a set of hand-tuned filters. All of these filters depend only on information that is observable to any player in the game (i.e., the filters do not depend directly on the scores of any individuals), and hence we can use the inferred states of players as proxies for what other players might infer as the states of those players. Also, since the states are not defined in terms of scores, we can meaningfully quantify the relationship between state and performance.

¹Results were similar using a mixed-effects regression including group and score field as random effects, and also revealed larger variability due to score field in the “medium” noise condition than the “low” noise condition.

We now define the three states: exploiting, copying, and exploring. Exploiting a particular location in the environment is not completely trivial for players since the avatars always move at least at a slow constant velocity. In order to attempt to stay in a single location, a player can either meander around a particular location or can persistently hold down one of the arrow keys while moving at a slow speed, which creates a tight circular motion around a particular location. We call this second activity “spinning” because of its distinctive appearance. We then classify a player as exploiting if the player is spinning for 500ms or if the player moves at the slow speed for 3 seconds and has not traveled more than two thirds of the possible distance that the player could have traveled in that time. The second condition is supposed to capture the meandering behavior of individuals who have not discovered how to spin. Copying behavior is more difficult to identify, but appears to often be characterized by fast directed movements towards other players. We thus classify a player as copying if the player is moving in a straight line at the fast speed towards any particular other player consistently for 500ms. We classify a player as moving towards another player if the second player is within 60° on either side of the first player’s straight-line trajectory. Finally, we classify a player as exploring if the player is neither exploiting nor copying. Thus a player will be classified as exploring if that player is either moving slowly but not staying in the same general location, if the player is moving quickly but not towards any particular person, or if the player is moving quickly and turning.

We use these filters to analyze how players behave in our game. First, we compute the probability of a player being in a particular state conditional on the current score that the player is receiving. We find that the probability of a player occupying a particular state is closely related to that player’s score. Specifically, players in higher scoring locations are more likely to be exploiting than exploring or copying, but the probability that a player is exploring or copying increases as the player’s score decreases. These results, which are visualized in Figure 4, suggest that players are choosing their states relatively rationally. Players will tend to remain in good areas and will leave bad areas quickly either by exploring independently or by copying other individuals.

Second, we find substantial variation in the types of copying behavior that different individuals display. Some individuals appear to focus their copying behavior on other players who tend to have higher scores, whereas other individuals appear to be less discriminating in their copying behavior. Moreover, as shown in Figure 5, groups that contain individuals who focus their copying behavior on higher scoring individuals achieve significantly higher performance in our task (slope: 0.2639, 95% CI: [0.145, 0.383]). This result, though subject to the confounding of correlation and causation, could be explained by theory of mind assisting in individual and group performance. A player who is able to accurately infer whether another player is receiving a high score may be able to achieve higher performance on our task by leveraging these

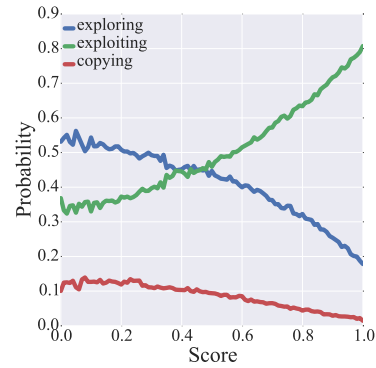


Figure 4: The probability of an individual being in a particular behavioral state as a function of the individual’s score.

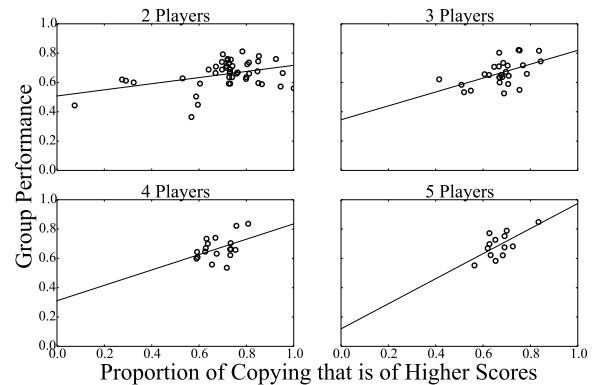


Figure 5: Average group performance as a function of the fraction of copying in the group that consists of “intelligent copying”—copying of an individual with a higher score. Lines are individually fitted regression lines.

inferences to more effectively copy others.

Behavioral Model

The trends we observe suggest a potential set of behavioral mechanisms that effective human groups may use in our task. We propose that each player in an effective group chooses a state based on the following rules:

1. If the player is in a good area, the player will remain in that area exploiting.
2. If the player is not in a good area and the player perceives another person as possibly having a higher score, the player may choose to copy that person.
3. Otherwise the player will explore independently.

According to this model, players in bad locations improve their scores by copying exploiting individuals instead of wasting time by copying low scoring players or wasting time by exploring many poor quality areas. The model also has interesting emergent collective properties. When any individual

finds a good area, that player will attract the other players to that location by exploiting. Then, when all the players are together in a group exploiting a particular area, one of the players will start to lose bonus as the score field shifts. This player will then either move closer to the others who are still exploiting or will shift to an exploring state. If that player starts exploring but doesn't find any good locations, the player will return to the group if the group is still exploiting. If that player does find a new good area, though, the player will start exploiting that area. The rest of the group will then follow after the highest scoring region shifts to where the exploiting player is. This mechanism creates a kind of gradual crawling that effectively tracks the moving score field. Thus, by using this mechanism players are improving both their own performances directly and also that of the entire group by participating in this process of emergent collective sensing. An example of this process occurring in participant gameplay is shown in Figure 6.

Discussion

In our experiment, we observed that humans were able to achieve increases in performance at much smaller group sizes than fish. Fish exhibited mild improvements in group performance at groups of 16 and more substantial improvements at groups of 64 and 128. However, we see significant improvements in human performance at just five players. This difference may be at least partially explained by the differences in the mechanism that humans appear to use in this task as compared to fish.

Interestingly, the mechanism we identify in humans is similar to that of fish in some ways, but it is also distinct in important ways. Similar to the behavior of humans in choosing appropriate states based on current score, fish modulated their speeds based on the level of darkness that they were experiencing. Fish moved slower in their preferred darker areas and faster in lighter areas. Similar to the copying behavior we observe, fish had a tendency for turning towards other fish. However, Berdahl et al.'s model of the behavior of their fish did not require any reference to the kind of discerning social awareness that we see in humans. Whereas fish appear to equally weight information from all nearby conspecifics, effective humans appear to modulate their copying behavior based on the inferred scores of other players. The strategic use of independent exploration (a form of asocial learning) was also key to the mechanism enabling human success. These key differences support recent work in social learning (Wisdom, Song, & Goldstone, 2013; McElreath et al., 2008), which find an impressive flexibility in the strategic deployment of imitation in humans. Of course, it is difficult to compare human performance directly to that of fish given the differences between the perceptual and motor abilities of fish in an actual fish tank and the abilities of the participants in our simulated environment. Nevertheless, our comparison hints at a superior capacity for distributed cognition in humans, possibly enabled by our ability for theory of mind.

Perhaps an even more interesting difference that emerged between humans and fish has to do with the time scale over which the collective intelligence mechanism evolved. For fish, the ability to gain from group performance in these collective sensing tasks is likely based on innate behaviors, selected over many generations of fish facing exactly this problem over their whole lifespans. In contrast, some of our human groups, facing this particular problem for the first time, appear to have discovered reasonable collective sensing strategies in just a matter of minutes.

Beyond the recent literature on collective intelligence in nonhuman animal groups, there has been a long line of work studying the factors that predict the performance of human groups in various scenarios (Kerr & Tindale, 2004). Our findings are consistent with previous work suggesting that having a larger group is beneficial in complex, uncertain environments (Stewart, 2006). Unlike much of this previous work, however, we focus here on the possibility in larger groups of new emergent group abilities and behaviors, and on the mechanisms leading to these emergent properties.

Our work therefore may shed light on one of the pressing puzzles of human collective intelligence and human distributed cognition. What are the specific mechanisms by which humans establish effective coordinated distributed information processing agents that can accomplish more than any individual alone, and how do our abilities play a role in these mechanisms? The perspective of group behavior as distributed processing (Hutchins, 1995) suggests the importance of communication for collective intelligence because of the importance of communication in distributed systems. Moreover, theory of mind—an enabler of implicit communication—has been shown to be predictive of collective intelligence (Woolley, Chabris, Pentland, Hashmi, & Malone, 2010; Engel, Woolley, Jing, Chabris, & Malone, 2014). While our work does not have a powerful enough experimental design to be definitive, our work at least further suggests that one of the roles that theory of mind plays in the emergence of collective intelligence is facilitating implicit communication that allows for coordination on good collective actions. Moreover, our work also suggests that the benefit of a group's coordinating on good actions could be more than simply the benefit to each individual independently. By combining a natural human tendency for independent exploration with a discerning social awareness, humans appear to be able to fluctuate between exploiting known good actions, independently exploring new options, and intelligently copying the promising choices of other individuals. A simultaneous combination of these activities by a cohesive group appears to lead to a collective memory of recently good actions from individuals who continue to exploit, and a collective movement towards actions that promise to be good in the near future driven by independently exploring individuals. The reactive distributed sensing ability that appears to emerge from this process may confer a unique benefit to working together in tightly knit groups.

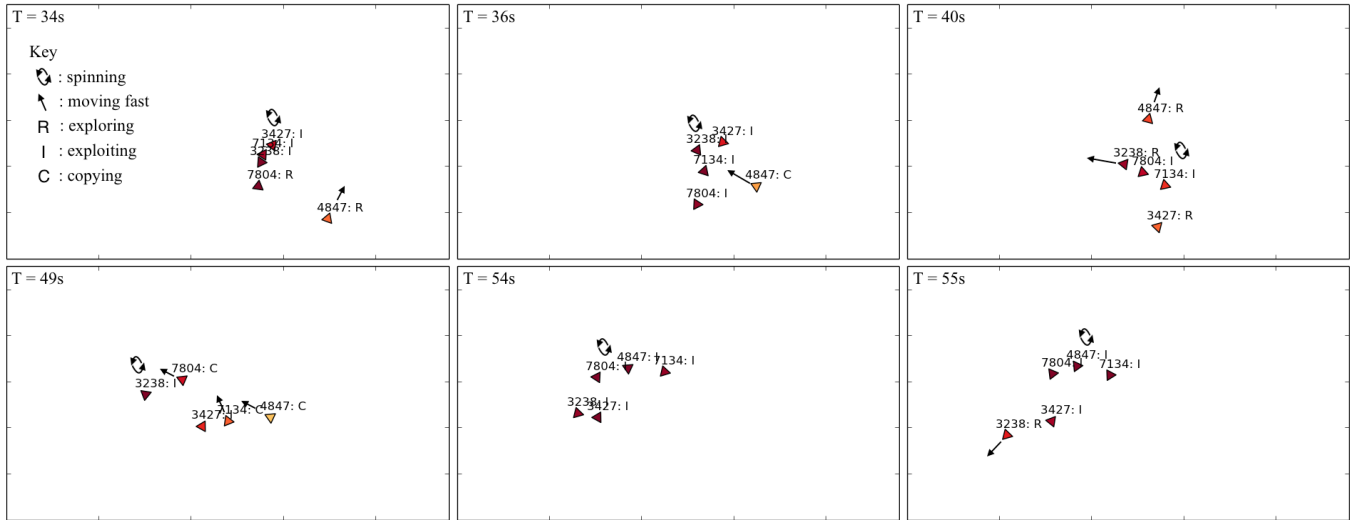


Figure 6: Reconstructions of actual gameplay in a five-person group illustrating both failed exploration leading to intelligent copying and successful exploration leading to collective movement. Colors indicate the individuals' scores, with red being higher and orange/yellow being lower. The player labels indicate both player IDs and also the player states our feature extraction procedure inferred. Other annotations are provided to give a sense for the game dynamics. At 34 seconds, in the first panel, most of the group has converged on exploiting a particular area while one individual is exploring independently. To the right, at 36 seconds, the exploring individual appears to have failed to find a good location and ceases exploring by copying the group. At 40 seconds, the final panel in the first row, the score field has shifted and some of the group begins exploring while others continue to exploit. By 49 seconds, the first panel in the second row, one of the exploring individuals found a good location, and other players have begun to move towards that individual. At 54 seconds, the entire group is exploiting the new area. In the final panel, at 55 seconds, the background has shifted enough again that one of the individuals begins to explore.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 to PK and Grant No. DGE-114747 to RXDH. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This material is based upon work supported by the Center for Minds, Brains and Machines (CBMM), funded by NSF STC award CCF-1231216.

References

- Amé, J.-M., Halloy, J., Rivault, C., Detrain, C., & Deneubourg, J. L. (2006). Collegial decision making based on social amplification leads to optimal group formation. *Proceedings of the National Academy of Sciences*, *103*(15), 5835–5840.
- Berdahl, A., Torney, C. J., Ioannou, C. C., Faria, J. J., & Couzin, I. D. (2013). Emergent Sensing of Complex Environments by Mobile Animal Groups. *Science*, *339*(6119).
- Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, *13*(1), 36–43.
- Engel, D., Woolley, A. W., Jing, L. X., Chabris, C. F., & Malone, T. W. (2014). Reading the Mind in the Eyes or Reading between the Lines? Theory of Mind Predicts Collective Intelligence Equally Well Online and Face-To-Face. *PLoS ONE*, *9*(12), e115212.
- Goldstone, R. L., & Gureckis, T. M. (2009). Collective Behavior. *Topics in Cognitive Science*, *1*(3), 412–438.
- Goldstone, R. L., Roberts, M. E., Mason, W., & Gureckis, T. (2008). Collective search in concrete and abstract spaces. In *Decision Modeling and Behavior in Complex and Uncertain Environments* (pp. 277–308). Springer.
- Hawkins, R. X. D. (2014). Conducting real-time multiplayer experiments on the web. *Behavior Research Methods*.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.
- Kearns, M. (2012). Experiments in social computation. *Communications of the ACM*, *55*(10), 56–67.
- Kerr, N. L., & Tindale, R. S. (2004). Group Performance and Decision Making. *Annual Review of Psychology*, *55*(1), 623–655.
- McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3515–3528.
- Pratt, S. C., & Sumpter, D. J. (2006). A tunable algorithm for collective decision-making. *Proceedings of the National Academy of Sciences*, *103*(43), 15906–15910.
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., . . . Laland, K. N. (2010). Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science*, *328*(5975), 208–213.
- Seeley, T. D., & Buhrman, S. C. (1999). Group decision making in swarms of honey bees. *Behavioral Ecology and Sociobiology*, *45*(1), 19–31.
- Stewart, G. L. (2006). A Meta-Analytic Review of Relationships Between Team Design Features and Team Performance. *Journal of Management*, *32*(1), 29–55.
- Tomasello, M. (2014). *A Natural History of Human Thinking*. Harvard University Press.
- Ward, A. J., Sumpter, D. J., Couzin, I. D., Hart, P. J., & Krause, J. (2008). Quorum decision-making facilitates information transfer in fish shoals. *Proceedings of the National Academy of Sciences*, *105*(19), 6948–6953.
- Wisdom, T. N., Song, X., & Goldstone, R. L. (2013). Social Learning Strategies in Networked Groups. *Cognitive Science*, *37*(8), 1383–1425.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science*, *330*(6004), 686–688.