

Social Foraging in Groups of Search Agents with Human Intervention

Daniel S. Schloesser¹ (dschloesser@ucmerced.edu), Derek Hollenbeck² (dhollenbeck@ucmerced.edu),
& Christopher T. Kello¹ (ckello@ucmerced.edu)

Cognitive and Information Sciences¹, Mechanical Engineering², University of California, Merced
5200 N. Lake Road, Merced, CA 95343 USA

Abstract

Intelligent agents coordinate and cooperate flexibly when rules and dynamics of interaction can change over time and across different tasks and environmental conditions. *Loose coupling* emerges among agents when the rules of interaction are weak enough for agents to act independently or interdependently, and patterns of interaction vary as a function of conditions. Here, we examine collective foraging among simulated agents with and without human intervention. We find that loose coupling among search agents improved group foraging success, and that human players improved performance partly by subtle, indirect effects on group interactions. Analyses of movement patterns showed that loose coupling enabled collections of agents to self-organize and reorganize into a greater diversity of ad hoc groupings.

Keywords: Social foraging; Agent based modeling; Loose coupling

Foraging success often depends on cooperation. From lions to vultures to humans, cooperative behaviors have been key to their success and survival. The interactions needed for cooperation are complex and depend on various conditions in the environment, social dynamics, and individual differences. It would be useful, for example, if interactions among collective foragers could provide information about how long to spend at one resource before leaving to find another. In the present study, we investigate how simple, constrained interactions among simulated agents in a foraging game affect group performance, and how human players interact with simulated agents to improve group performance.

In social foraging, individual and group search behaviors influence the group outcome. Social Foraging Theory aims to explain how individual behaviors adapt to optimizing group foraging success and fitness (Giraldeau & Caraco, 2000). Social foraging is founded on the premise that evolutionary advantages can be gained by cooperating with other conspecifics. Cooperation occurs when group performance is better than any statistical gain expected to occur by mere aggregation of individuals (Giraldeau & Caraco, 2000). Established benefits of social foraging include greater likelihood of finding food resources (Beauchamp, 2005) as well as security and mating among other social benefits (Todd & Miller, 1999).

Some species can transition from searching socially and cooperatively, to searching independently. For example, humans are very capable of making this transition, and we need to be able to flexibly couple and uncouple from each other, often in various configurations of groupings. Different degrees and kinds of coupling are possible, and it seems

likely that variations in these parameters will result in variations from more individual to more collective search strategies. Such flexible couplings and configurations may serve to balance the degree of exploring new territory versus exploiting established territory as a group (Dreller, 1998; Seeley, 1983).

For example, Harel, Spiegel, Getz, and Nathan (2017) showed that griffon vultures (*Gyps fulvus*) switch between individual and social foraging strategies in response to social cues. They showed that cues beyond proximity allow one vulture to see whether others have found food. Vultures emit and perceive cues about from a carcass location that alter their foraging paths (Harel et al., 2017). The evidence showed that individual vultures who previously visited a carcass were statistically more likely to be followed when revisiting that same carcass location.

This pattern of social foraging was based on whether a vulture was an informed or uninformed forager. Perceivable information, like noticeable blood stains after feasting, serve as cues that inform nearby vultures to either exploit information to join the collective, or explore as an individual for fresh carcasses. Social foraging in this case means striking the right balance of choices to increase the overall rate of consumption and energy intake of the group. Flexibly switching between exploration and exploitation was shown to increase group search efficiency through interactions among vultures that shared information among them (Harel et al., 2017).

The cooperative behaviors of loosely coupled social foragers, like vultures, makes them less susceptible to predation, improves their chances against other scavenging species, and increased foraging success (Lamprecht, 1981). These benefits of social foraging have been shown in other birds of prey as well (Cortés-Avizanda et al., 2014; Harel et al., 2017). The recurrence of social foraging in the animal kingdom has led some researchers to study it in humans as convergent evidence and for the sake of understanding human social foraging per se. For example, in a simulation by Beauchamp (2005), the rate of food intake among social foragers was found to be less than that of individual foragers, but to compensate, social foragers were able to consume food patches faster than individual foragers. The net result was an increase in the mean food intake rate for social foragers which provides a potential safeguard against starvation (Beauchamp, 2005).

In another study, Liu and Passino (2004) created a collective foraging model based on an attract-repel rule where agents within the group sought to find a “comfortable”

position in relation to their nearest neighbors. The environment had varying resource gradients that foragers could follow to find food, and there was uncertainty in agents' ability to detect and follow gradients, including their own positions and velocities. The authors measured cohesion between the agents by the average position and velocity of the swarm. Closer positions and more similar velocities corresponded with cohesive movement. Results showed that a balance of attraction and repulsion engendered loose coupling and coordinated behaviors of foraging agents as a group, and these effects grew with group size. Results also showed that the detrimental effects of noise on individual foraging behaviors were averaged out in the collective groups. In more recent work, Copenhagen, Quint, and Gopinathan (2016) implemented cohesion among simulated agents using the *Lennard-Jones potential*, which is a single parameterized equation that governs the degree to which agents are attracted to or repelled from each other as a function of their distance apart (details in Method section).

Here we report an agent-based simulation of foraging agents in which we manipulate the degree of cohesion and collective behavior using the Lennard-Jones potential plus a formalization of the influence of directional alignment among search agents. Our simulation is based on scavenging vultures who search for food somewhat independently, while also utilizing each other to detect scarce food sources (Cortés-Avizanda, Jovani, Donazar, & Grimm, 2014; Harel, Spiegel, Getz, & Nathan, 2017). Our research question is how loose coupling affects group performance. We include a condition of human intervention where one of the agents is controlled by a human player instead of automatic rules of movement and interaction. We expected human intervention to improve group performance, and we tested whether humans could beneficially influence the behaviors of simulated agents indirectly through their rules of interaction.

Method

Collective Foraging Game and Participants

Sixty undergraduate participants volunteered as participants for course credit. Participants played a cooperative search game based on a previous social foraging model (Vicsek & Zafeiris, 2012). The game was implemented in NetLogo and each player controlled an on-screen avatar with the use of a computer mouse, along with nine other avatars controlled by one of four different search algorithms, detailed below. Thus, human intervention was always defined as one human player with nine automated search agents. In addition to human intervention, we also examined a fully automated condition of 10 search agents.

The game was played by searching the 2D space for “sheep” represented by gold star targets, which were not visible beyond a certain radius around each agent, including the human. Targets were placed in the game space one at a time, and the goal was to find each one as fast as possible, and “consume” it as fast as possible, before moving on to the next. Consumption occurred simply by staying over the

target, so the only way to consume faster was for multiple search agents to converge on the target.

The search space was a 200x200 grid of pixels with periodic boundary conditions. The grid was empty except for one target at time, located at random. Agents could not “see” targets until they came within a 22.5-pixel radius. Automated agents immediately headed toward each target upon detection, and upon arrival, each agent consumed one unit of target “resource” per time step, for 500-time steps (each time step was 3.5 ms long). Multiple agents could land on a target and together consume it faster. When a target was completely consumed, it disappeared, and a new target appeared at a random location.

Rules for Automatic Search Agents

In the absence of target detection, automated search agents moved based on some combination of four possible rules. One rule present in all conditions was *visual chaining* as implemented by Cortés-Avizanda et al. (2014): if one agent saw a second agent approaching a target, the first agent headed towards the second one. This rule favored collective foraging because it could effectively extend an individual's “field of view” to include the radii of others in view. The three other possible rules added forces that drove search agents to converge, align, or wander.

The first rule was the Lennard-Jones potential (Copenhagen et al., 2016; Spears & Spears, 2012) which came into effect when agents could see each other, and caused agents to converge if in the periphery, and separate as they came close to “colliding”, although no collision rules were implemented. The second rule caused agents to align their movement directions when nearby, and the last added variability to movements through a random correlated walk (RCW). The governing equations for the cohesion, alignment, and RCW parameters are as follows:

$$\text{Cohesion } (\vec{d}_{AJ}): \quad \vec{d}_{LJ} = -\sum_i \left[\left(\frac{s}{r} \right)^4 - \left(\frac{s}{r} \right)^3 \right] \hat{r}_i, \quad r_i \leq d_v$$

$$\text{Alignment } (\vec{d}_A): \quad \vec{d}_A = \sum_j \vec{d}_j(t - \delta t), \quad |\vec{d}| < d_v$$

$$\text{RCW } (\vec{d}_N): \quad \vec{d}_N = U[0,180] - U[0,180]$$

All agents were in constant motion which meant that the movement rules only governed turning to change directions. U was a random turn from a uniform distribution between 0 and 180 degrees, the d variables are directional headings, position is measured at time t and δt , s is the separation between agents, and the r variables are distances between pairs of agents. Rules were combined additively, $\vec{d} = \vec{d}_A + \vec{d}_{AJ} + \vec{d}_N$, and simulation conditions were defined by including or not the alignment and cohesion rules (the RCW rule was always in effect). This resulted in four different movement conditions: Random (both off), Alignment, Cohesion, and Combined (both on). All simulated agents were uniformly assigned to one of the four conditions for each run of the game, and parameters were tuned to create loose coupling between agents in the combined condition.

See Figure 1 for general trajectory examples of each movement type.

The human player controlled the movement of one agent by placing the mouse at a desired location. Automated agents had no effect on the human player, but the human player affected automated agents by the rules reviewed above. Foraging performance with a human player was compared with simulation-only conditions in which all ten agents were fully automated. The conditions were otherwise identical.




Figure 1: Example movement trajectories for 2000-time steps for each movement condition. Top Left: Random condition. Top Right: Alignment condition. Bottom Left: Cohesion condition. Bottom Right: Combined condition.

Procedure and Measures

The orders of movement conditions were randomized, and all participants completed all four movement conditions. Agents started each search session in random positions, each session ran for 13,500 time steps, which was ~7-8 minutes in real time when a human player was included. All agents moved at a constant velocity of 1 pixel per simulation time step. Performance was measured simply in terms of the number of targets found and consumed in each session. This measure is composed of the times taken to find each target, and times taken to consume them. A “trial” was defined as the onset of each new target and lasted until it was completely consumed.

Movement patterns were analyzed in terms of the frequencies with which agents acted collectively with other agents, as defined by being within view of each other. Each search agent, including the human, could have from zero to nine other agents in view on each time step while searching for a target, and a target could be consumed by one to ten agents on each time step of target consumption. We analyze the frequencies with which each possible configuration occurred over time, as a function of movement condition, and with and without the inclusion of a human player.

Results

Figure 2 shows mean levels of performance as a function of condition and human intervention, where performance includes times to detect targets, times to consume them, and numbers of targets acquired.



Figure 2: Right: Ten simulated agents ($N=60$ per condition); Left: User and nine simulated agents ($N=60$ per condition); Red: Averaged time to detect; Teal: Averaged time to consume; Asterisk: Mean targets found.

Differences between conditions were tested using a 2 (Mode: experiment vs. simulation) $\times$ 2 (Cohesion: on or off) $\times$ 2 (Alignment: on or off) repeated-measures analysis of variance (ANOVA) with numbers of targets acquired as the dependent measure. We found significant main effects for Mode, $F(1, 59) = 302.45$, $p < 0.001$, $\eta_p^2 = 0.069$, Cohesion, $F(1, 59) = 8.58$, $p = 0.004$, $\eta_p^2 = 0.502$, and Alignment, $F(1, 59) = 6.32$, $p = 0.012$, $\eta_p^2 = 0.927$. Additionally, results indicated a significant two-way interaction between Cohesion and Alignment, $F(1, 59) = 6.18$, $p = 0.013$, $\eta_p^2 = 0.614$, and a significant three-way interaction between Mode, Cohesion, and Alignment, $F(1, 59) = 8.73$, $p = 0.003$, $\eta_p^2 = 0.253$. Performance was best in the combined condition for both modes, where loose coupling was hypothesized to occur, and performance was always improved by a human player. The agents benefitted from human intervention the most in the combined condition, suggesting that human players were best able to beneficially influence search agents when they were loosely coupled.

Time to Detect

We conducted the same repeated-measures ANOVA using time to detect as the dependent measure. There was a significant main effects for Cohesion, $F(1, 59) = 7.66$, $p = 0.008$, $\eta_p^2 = 0.162$, Alignment, $F(1, 59) = 119.18$, $p < 0.001$, $\eta_p^2 = 0.312$, and Mode, $F(1, 59) = 169.94$, $p < 0.001$, $\eta_p^2 = 0.418$. With significant two-way interactions between Cohesion and Alignment, $F(1, 59) = 38.33$, $p < 0.001$, $\eta_p^2 = 0.181$, and Mode and Alignment, $F(1, 59) = 27.59$, $p < 0.001$, $\eta_p^2 = 0.105$. There was also a significant three-way interaction between Mode, Cohesion, and Alignment, $F(1, 59) = 18.18$, $p < 0.001$, $\eta_p^2 = 0.071$. The combined condition again showed the best performance, and performance again improve with human intervention. However, time to detect benefitted the most from humans in the alignment condition, and later analyses will show that this benefit came from individual foraging separately from the overly aligned and unresponsive group of autonomous agents.

Time to Consume

Next we conducted the same repeated-measures ANOVA with time to consume as the dependent measure. We found a significant main effect of Cohesion, $F(1, 59) = 234.72, p < 0.001, \eta_p^2 = 0.185$, Alignment, $F(1, 59) = 2500.57, p < 0.001, \eta_p^2 = 0.927$, and Mode, $F(1, 59) = 17.61, p < 0.001, \eta_p^2 = 0.069$. With significant two-way interactions between Cohesion and Alignment, $F(1, 59) = 432.32, p < 0.001, \eta_p^2 = 0.614$, Mode and Cohesion, $F(1, 59) = 55.82, p < 0.001, \eta_p^2 = 0.191$, and between Mode and Alignment, $F(1, 59) = 86.15, p < 0.001, \eta_p^2 = 0.267$. Followed by a significant three-way interaction between Mode, Cohesion, and Alignment, $F(1, 59) = 80, p < 0.001, \eta_p^2 = 0.253$. The pattern of results was the same time to detect, except that human intervention caused an *increase* in the average time to consume, indicating that individual foraging away from the “pack” increased time to detect but at the cost of slightly longer times to consume.

Diversity of Search and Consumption Patterns

We used entropy to measure the diversity of differently sized groupings that occurred during times of search versus times of consumption. Our hypothesis was that performance was best in the combined condition because search agents were loosely coupled and thereby able to change between modes of flying together versus flying separately. Entropy served to quantify the diversity of groupings, computed over the probability of a given agent being in one of ten possible grouping states on each time step. We calculated entropy as $-\sum [p(x_i) \log(p(x_i))]$ where x_i is the number of agents in view and p is the associated probability.

Figure 3 shows a normalized histogram (i.e. probabilities) of the numbers of agents in view while searching for a target. The first 50-time steps at the start of each trial were excluded to avoid transitory effects from the previous trial. Figure 3 features the autonomous search agents with human intervention, but the human player was removed in order to focus on measuring coordination of the automated search agents, and the degree to which this coordination is influenced by human intervention. Figure 4 shows the same histogram but for agents without human intervention.




Figure 3: Normalized histogram of the number of agents in view during the time to detect period is respective to the nine autonomous agents minus the human agent.

First and foremost, the histograms show a more even distribution in the combined condition, whereby search agents ranged from individual to collective search and consumption in group sizes ranging 1 to 5+. The random and cohesion conditions showed less diversity, and the alignment condition showed the least. The latter occurred because agents flocked entirely together most of the time. The effect of human intervention is subtle, so we turn to the entropy analyses for a more precise comparison (see Figure 5).




Figure 4: Normalized histogram of the number of agents in view during the time to detect period respective to nine autonomous agents without human intervention.

Again, we conducted the same 2 (Mode: experiment vs. simulation) x 2 (Cohesion: on or off) x 2 (Alignment: on or off) repeated-measures ANOVA as before, but with entropy as the dependent measure, where entropy was computed for each session in each condition. We found a significant main effect of Cohesion, $F(1, 59) = 6082.08, p < 0.001, \eta_p^2 = 0.964$, and a significant two-way interaction between Cohesion and Alignment, $F(1, 59) = 117.11, p < 0.001, \eta_p^2 = 0.958$. We also found a significant two-way interaction of Mode and Alignment, $F(1, 59) = 9.47, p = 0.002, \eta_p^2 = 0.039$. The three-way interaction between Mode, Cohesion, and Alignment was non-significant, $F(1, 59) = 0.019, p = 0.891, \eta_p^2 = 0.001$.




Figure 5: Mean entropy values by Condition and Mode (Simulation and Experiment)

We conducted one-sample t-tests comparing the difference value after subtracting the global mean search entropy values

of the simulation searchers for each condition, respectively. These effects are largely driven by the influence the movement rules had on the coupling between the non-human autonomous agents, irrespective of anything the human agent was doing. Interestingly, we did find significant effect that can be attributed to the influence the human agent's coupling had on the other non-human agents.




Figure 6: Search Entropy difference values for each condition. The mean search entropy value was subtracted from the average search entropy value for each group of nine autonomous searching agents.

We found a significant departure from zero in the Cohesion condition, $t(59) = -2.533$, $p = 0.014$, and Random condition, $t(59) = -2.178$, $p = 0.033$. After correcting the significance value for running four independent one-sample t-tests the significance threshold is $p = 0.125$. Therefore, any indirect effect the human search had on the autonomous searchers during the experiment is marginally significant in the Cohesion and Random conditions. We find these results to be interesting because the human searcher indirectly influenced the searching behaviors of the autonomous searchers beyond that of the simulation autonomous searchers when the searching behaviors of the autonomous searchers is not as complex as it is in the Combined condition. Despite not finding any significant difference in the Combined condition, both sets of autonomous searchers obeyed the same movement rules dependent upon the condition. These effects indicate that human players on average *decreased* the diversity of search agent groupings. They displayed some indirect effect of coupling beyond the agents without human intervention while searching for targets.

Time to Consume

For the time to consume entropy analysis we calculated the number of agents consuming the target per time step during the consuming period of the task. This analysis differs from the time to detect analysis in that all agents in both modes were included. These analyses serve to examine how differences in coupling between conditions and across modes influenced the rate of consumption of the targets found. See Figure 7 and 8 for the respective normalized plots of the number of agents consuming a target. Figure 9 shows the corresponding mean entropy values.




Figure 7: Normalized histogram of the number of agents consuming the target during the time to consume period for autonomous agents with human intervention.




Figure 8: Normalized histogram of the number of agents consuming the target during the time to consume period without human intervention.

We conducted the same repeated-measures ANOVA as before with the entropy values for time to consume distributions as the dependent measure. We observed a significant main effect of Cohesion, $F(1, 59) = 667.78$, $p < 0.001$, $\eta_p^2 = 0.74$, Alignment, $F(1, 59) = 282.25$, $p < 0.001$, $\eta_p^2 = 0.541$, and Mode, $F(1, 59) = 85.78$, $p < 0.001$, $\eta_p^2 = 0.267$. These main effects were qualified by significant two-way interactions between Cohesion and Alignment, $F(1, 59) = 222.64$, $p < 0.001$, $\eta_p^2 = 0.586$, Mode and Cohesion, $F(1, 59) = 15.99$, $p < 0.001$, $\eta_p^2 = 0.063$, and marginally significant two-way interaction between Mode and Alignment, $F(1, 59) = 3.17$, $p = 0.076$, $\eta_p^2 = 0.013$. There is a significant three-way interaction between Mode, Cohesion, and Alignment, $F(1, 59) = 18.91$, $p < 0.001$, $\eta_p^2 = 0.074$. Due to the human agent's ability to flexibly switch between independent and collective search strategies, the overall entropy values for that mode increased because the human user can leave the group to consume targets independently. This created a more variable probability distribution, which increased the overall entropy for that mode of our study (see Figure 8).




Figure 9: Entropy values by Condition and Mode (Simulation and Experiment)

For consuming entropy we found a significant departure from zero in all conditions but the Combined condition: Combined, $t(59) = 1.617, p = 0.111$; Cohesion, $t(59) = 5.903, p < 0.001$; Alignment, $t(59) = 13.49, p < 0.001$; and Random; $t(59) = 4.305, p < 0.001$. See below in Figure 10 for the consuming entropy difference scores. As can be seen in Figure 10, the consuming entropy was higher for all the conditions but the Combined condition. We take this to indicate the human intervention increased the overall entropy of consuming targets across the conditions. The increased entropy provides an indication that there is more variability in the balance between searching and consuming targets based on the intervention of the human working alongside the other autonomous searching agents.




Figure 10: Consumption Entropy difference values for each condition. The mean consumption entropy value was subtracted from the average consumption entropy value for each group of nine autonomous searching agents.

In general, results followed those for time to detect distributions. The combined condition again showed the greatest entropy and hence diversity of patterning, as with entropy values for time to detect distributions. This result further supports the hypothesis that better performance in the combined condition was supported by loose coupling. The cohesion and random conditions showed moderate levels of diversity, and the alignment condition showed the least.

The biggest difference between time to detect and time to consume patterns was that human intervention *increased* the mean entropy values for time to consume across all four conditions, and especially in the alignment condition. All the analyses taken together, it appears that humans spent a substantial portion of time switching between individual and collective search strategies. Individual periods decreased the diversity of search patterns but increased the diversity of consumption patterns.

DISCUSSION

In the present study, we found that group performance was best in the combined condition when cohesion and alignment were both active. The combined condition exhibited loose coupling among the human agent and the non-human agents. This loose coupling allowed for wider coverage of the area while maintaining connection to other agents compared to other conditions where agents were usually either all together or all independent of each other. The human agents uniformly improved performance across all movement conditions, and part of this improvement was via the effects of human intervention on search agent movement patterns.

Pairing the time to detect and time to consume entropy results provides a more complete picture of how human players influenced the non-human agents, and how they uniformly improved performance across movement conditions. The indirect influence of the human player on search agents was subtle and requires further investigation to understand how their movements, which were the only means of interaction, impacted the group.

Our collective foraging game was simple which allowed us to attribute results to specific manipulations, and possibly generalize our results to other more realistic search conditions. For example, loose coupling may be beneficial in teams of robotic search agents, and human operators may be able to influence team coordination simply through their movements, as in the present study. Our results may apply to other real-world problems, such as search and rescue missions, surveillance tasks, and gas leak detection. A significant motivation for applying social foraging to cognitive engineering problems is that many species have been shown to address these problems efficiently using strategies of social foraging. Our results contribute to the growing body of literature on the complementary benefits individual and social foraging, and the conditions under which different strategies are most appropriate.

References

- Beauchamp, G. (2005). Does group foraging promote efficient exploitation of resources? *Oikos*, 111(2), 403–407.
- Copenhagen, K., Quint, D. A., & Gopinathan, A. (2016). Self-organized sorting limits behavioral variability in swarms. *Scientific Reports*, 6.
- Cortés-Avizanda, A., Jovani, R., Donázar, J. A., & Grimm, V. (2014). Bird sky networks: How do avian scavengers use social information to find carrion?

- Ecology*, 95(7), 1799–1808.
- Dreller, C. (1998). Division of labor between scouts and recruits: Genetic influence and mechanisms. *Behavioral Ecology and Sociobiology*, 43(3), 191–196.
- Giraldeau, L. A., & Caraco, T. (2018). *Social foraging theory*. Princeton University Press.
- Harel, R., Spiegel, O., Getz, W. M., & Nathan, R. (2017). Social foraging and individual consistency in following behaviour: Testing the information centre hypothesis in free-ranging vultures. *Proceedings of the Royal Society B: Biological Sciences*, 284(1852).
- Lamprecht, J. (1981). The function of social hunting in larger terrestrial carnivores. *Mammal Review*, 11(4), 169–179.
- Liu, Y., & Passino, K. M. (2004). Stable social foraging swarms in a noisy environment. *IEEE Transactions on Automatic Control*, 49(1), 30–44.
- Seeley, T. D. (1983). Division of labor between scouts and recruits in honeybee foraging. *Behavioral Ecology and Sociobiology*, 12(3), 253–259.
- Spears, W. M., & Spears, D. F. (2012). Physicomimetics: Physics-based swarm intelligence. In *Physicomimetics: Physics-Based Swarm Intelligence* (Vol. 97836422228).
- Todd, P. M., & Miller, G. F. (1999). From pride and prejudice to persuasion. In *Simple heuristics that make us smart* (pp. 287–308).
- Vicsek, T., & Zafeiris, A. (2012). Collective motion. *Physics Reports*, Vol. 517, pp. 71–140.