

Specialization and selective social attention establishes the balance between individual and social learning

Charley M. Wu^{1,2} (charley.wu@uni-tuebingen.de), Mark K. Ho³, Benjamin Kahl², Christina Leuker^{2,4}, Björn Meder^{2,5} & Ralf H.J.M. Kurvers²

¹ Human and Machine Cognition Lab, University of Tübingen, Tübingen, Germany

² Max Planck Institute for Human Development, Berlin Germany

³ Department of Computer Science, Princeton University, Princeton, USA

⁴ Robert Koch Institute, Berlin, Germany ⁵ Health and Medical University Potsdam, Potsdam, Germany

Abstract

A key question individuals face in any social learning environment is when to innovate alone and when to imitate others. Previous simulation results have found that the best performing groups exhibit an intermediate balance, yet it is still largely unknown how individuals collectively negotiate this balance. We use an immersive collective foraging experiment, implemented in the Minecraft game engine, facilitating unprecedented access to spatial trajectories and visual field data. The virtual environment imposes a limited field of view, creating a natural trade-off between allocating visual attention towards individual innovation or to look towards peers for social imitation. By analyzing foraging patterns, social interactions (visual and spatial), and social influence, we shine new light on how groups collectively adapt to the fluctuating demands of the environment through specialization and selective imitation, rather than homogeneity and indiscriminate copying of others.

Keywords: social learning; spatial foraging; collective intelligence; visual attention; producer-scrounger; Rogers' paradox

Introduction

Social information is a valuable resource, yet it often comes at the cost of lost opportunities for individual innovation. Previous research has shown that the best collective and individual outcomes occur when there is an intermediate balance of individual and social learning (Miu, Gulley, Laland, & Rendell, 2020; Ehn & Laland, 2012). However, it is not yet known how groups collectively negotiate this balance.

Here, we use a collective spatial search task to study the social dynamics of human behavior, under different reward environments designed to alter the effectiveness of social learning (Barkoczi, Analytis, & Wu, 2016; Rendell et al., 2010). Our task is implemented in an immersive virtual environment, where a limited field of view creates an attention allocation problem (Fig 1a). Visual attention can either be allocated to focus on individual exploration or to look towards peers for social imitation. We study how the predictability of resources alters the balance between social and individual learning. Employing high-resolution spatial tracking and automated transcription of visual field data, we provide an unprecedented level of detail into the interaction dynamics that negotiate this balance.

Between environments, we find that participants flexibly adapt their level of social learning, relying more on social information when profitable. In environments where predictable rewards favor social learning, participants reach a mixed equilibrium through specialization and selective attention, rather than all individuals homogeneously converging

on the same intermediate strategy. Participants become specialized as either the target or the source of social attention. This produces an asymmetric social attention structure where the most observed individuals are those who specialize in individual learning. Thus, social learners selectively attend to strongly individual learners, who innovate to discover new reward patches, rather than copying other social learners.

Individual vs. social learning

In social environments, individuals can either explore on their own or imitate the successes of others. This trade-off is encapsulated in the *producer-scrounger dilemma* (Barnard & Sibly, 1981; Kurvers et al., 2009), where social imitation (i.e., scrounging) comes at the cost of reduced (attentional) resources for individual search (i.e., producing). Attempted solutions to this dilemma generally have *frequency-dependent payoffs*, meaning one's performance depends on the strategies used by other individuals in the population. While social learning is cheap and effective among mostly individual learners, it fails amidst an abundance of imitators, causing a collapse in both individual and group performance. This phenomenon is known as Rogers' (1988) paradox.

However, more recent work has shown that Rogers' paradox can disappear in a variety of more realistic setting. For instance, when agents can flexibly switch between individual and social learning (Boyd & Richerson, 1995; Enquist & Ghirlanda, 2007; Kameda & Nakanishi, 2002), when rewards and agents are embedded in a spatial structure (Beauchamp, 2008; Rendell et al., 2010; Kobayashi & Ohtsuki, 2014) or when social information is not copied verbatim, but adapted in the process of transmission (Ehn & Laland, 2012; Boyd & Richerson, 1995). Yet the majority of these results are derived from simulations, with little experimental data about how people flexibly arbitrate between social and individual learning in realistic environments.

This motivates our use of an immersive virtual environment (Fig. 1a), where a limited field of view constrains how participants allocate visual resources. In this spatially-explicit environment, visibility is naturally defined in terms of distance and orientation between participants (Fig. 1b). Thus, the environment is only partially observable at any point in time, with opportunity costs when learning either socially or individually. Visual attention allocated towards individual search detracts from social learning, and vice versa. Addi-

tionally, the spatial distribution of participants influences the ease of exploiting social information, since imitation incurs travel costs depending on the distance to the target.

Methods

Participants and materials

Participants ($N = 44$) were recruited from the REDACTED recruitment pool¹ and selected to be between the ages of 18 and 30 to minimize generational differences in exposure to first person computer games (mean age: 26.5 ± 4.6 SD; 25 Females, 1 non-binary). The study was approved by the Institutional Review Board of REDACTED and participants signed an informed consent form prior to participation. Participants earned a base payment of €12 plus a bonus of €0.03 per reward, earning on average €17.21 \pm 0.88.

The experiment was conducted in a computer lab using a modified Minecraft server, where participants controlled an avatar from a first person perspective (Fig. 1a). Each round of the experiment took place on a 60x60 field (bounded by a fence), containing 400 resource blocks, laid out in a 20x20 grid with a two block gap between each block (Fig. 1b). Each resource block (either watermelon or pumpkin, depending on environment) could be destroyed by continually hitting it (holding down left mouse button) for 2.25 seconds, yielding a binary outcome of either reward or no reward.

Rewards were indicated by a blue splash effect (Fig. 1a), visible from any position on the map. Only resource blocks (watermelon or pumpkin) were capable of being destroyed in the experiment and were not renewed within the round. Additionally, blocks did not possess any *a priori* visual features indicating whether or not they contained a reward. However, rewards in smooth environments were partly predictable, since observing a reward predicts other rewards nearby (Fig. 1c). Participants were incentivized to collect as many rewards as possible, which were translated into a bonus payment at the end of the experiment.

Design and procedure

Participants completed the task in randomly assigned groups of four. Groups completed the task in the same room, and were made aware they were interacting with these same people. After an in-game tutorial and two practice rounds (see below), participants completed 16 rounds of the task, each lasting two minutes. We used a 2x2 within-subject design, manipulating the reward environment (random vs. smooth) and search condition (solo vs. group), with each combination completed in four sequential rounds (Fig. 1d).

The reward environment of a given round was made salient by the use of either pumpkin or watermelon blocks (Fig. 1c), which was randomly assigned and counter-balanced across groups. Both environments had the same number of rewards (25% of blocks), but with rewards either randomly or

¹The original target sample size was 128 participants, but we were forced to pause lab-based experiments since March of 2020 due to COVID-19. These current analyses thus focus on within-subject and within-session comparisons for the highest statistical power.

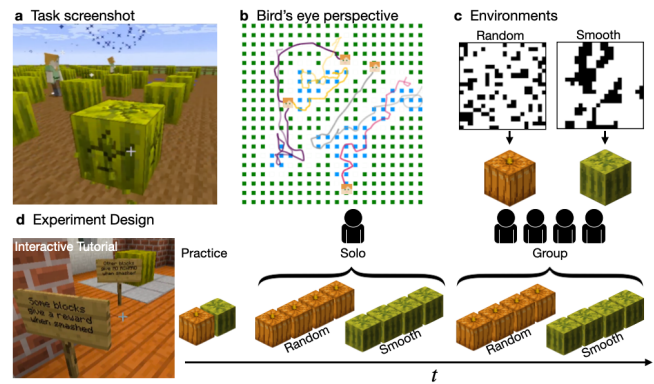


Figure 1: Minecraft experiment. **a**) A first person screenshot (cropped) from a social round, where a reward discovered by another player (top) is visible as a blue splash in the air. **b**) A bird's eye recreation of the experiment, showing the full 60x60 field and participant trajectories (colored lines). The blue blocks correspond to discovered rewards, while gaps in the regular grid of green blocks indicates no reward. **c**) Examples of smooth and random reward environments, which were mapped to either melons or pumpkins (counter-balanced across sessions). Each pixel indicates either reward (black) or no reward (white). Gaps between blocks are omitted. **d**) Experimental design. Participants first completed an interactive tutorial, followed by one training round of each environment. The main task was a 2x2 within-participant design, manipulating reward environment (smooth vs. random; mapped to melons or pumpkins) and search condition (solo vs. group). Each of the four configurations was completed in four consecutive rounds of the same type. Each melon/pumpkin block indicates one round. Round order was pseudo-randomized and counterbalanced between sessions.

smoothly distributed. Random environments were generated by uniformly sampling 25% of blocks (without replacement). Smooth environments were designed to contain clustered reward distributions, which varied smoothly over space. We first sampled a bivariate reward function from a Gaussian Process prior, where we used a radial basis function kernel with the lengthscale parameter set to 4 (similar to Wu, Schulz, Garvert, Meder, & Schuck, 2020). Sampled reward functions were then binarized, such that the top quartile (25%) of block locations were set to contain rewards. In the tutorial, participants were given verbal descriptions of each reward condition, saw two fully-revealed illustrations of each environment class from a bird's-eye perspective (Fig 1c), and interactively destroyed a 3x3 patch of both environments.

The search conditions were made salient by having participants either stand on an isolated teleportation platform (solo) or on a common teleportation platform with the other participants (group) in order to start the round. In the solo condition, participants searched on identical replications of the same environments but without interactions with one another. In the group condition, participants searched on the same environment and could compete with and imitate one another.

After receiving verbal instructions, participants completed

a tutorial to familiarize themselves with the controls, how to destroy blocks, the difference between smooth and random reward distributions, and the overall task structure (Fig. 1d). They then completed two practice rounds in both smooth and random environments, which were identical to the solo condition of the main task (but without contributing to bonus payments). Each round lasted two minutes, with the end of the round corresponding to the sun setting below the horizon. This served as an approximate in-game timer for each round, and was communicated to participants in the tutorial. A three second countdown timer was also shown on the screen. At the end of the round, participants were given an on-screen announcement indicating the number of rewards they earned and notifying them of the reward environment and search condition for the next round. Participants were then teleported into a lobby (either separate lobbies for solo or a communal lobby for group rounds), and were required to all stand on a “teleportation” block to indicate readiness for the next round. Then players were teleported into a random position in the next environment, facing a random direction. For the entire experiment the sound was turned off, participants could not see each other’s screens, and task-irrelevant controls (e.g., crafting menu) were disabled.

Data collection

Experimental data was collected using a custom data logging module programmed in Java, separated into player logs and map logs. Player logs contained each player’s $[x, y]$ spatial positions together with the pitch and yaw components of their visual orientation sampled at 20hz (every 0.05s). Map logs contained information about each player’s interactions with resource blocks, also sampled at 20hz. Together, the player and map logs allowed us to completely reconstruct the visual fields from all participants using a simulation programmed in the Unity game engine. This allowed us to automatically transcribe visibility information. We simulated each player’s field of view with all entities and other players rendered using a unique RGB color mask. From each frame of the simulations, we analyzed the pixels of the resulting image to determine which players and entities were visible at a given time.

Results

We first focus on task performance before analyzing social interactions using visual and spatial data. Lastly, we model social influence by detecting pull events in foraging patterns.

Task performance

Using a two-way within subjects ANOVA, we found that participants acquired more rewards in smooth environments ($F(1, 43) = 154.0, \eta^2 = .47, p < .001$) and in the solo condition ($F(1, 43) = 8.54, \eta^2 = .02, p = .006$), with an additional interaction of smooth:solo ($F(1, 43) = 6.42, \eta^2 = .01, p = .015$; Fig. 2a). This improved performance in the smooth:solo condition was mediated by learning over rounds (Fig. 2b). We fit a Bayesian mixed effects regression using environment, search condition, and round number to predict reward,

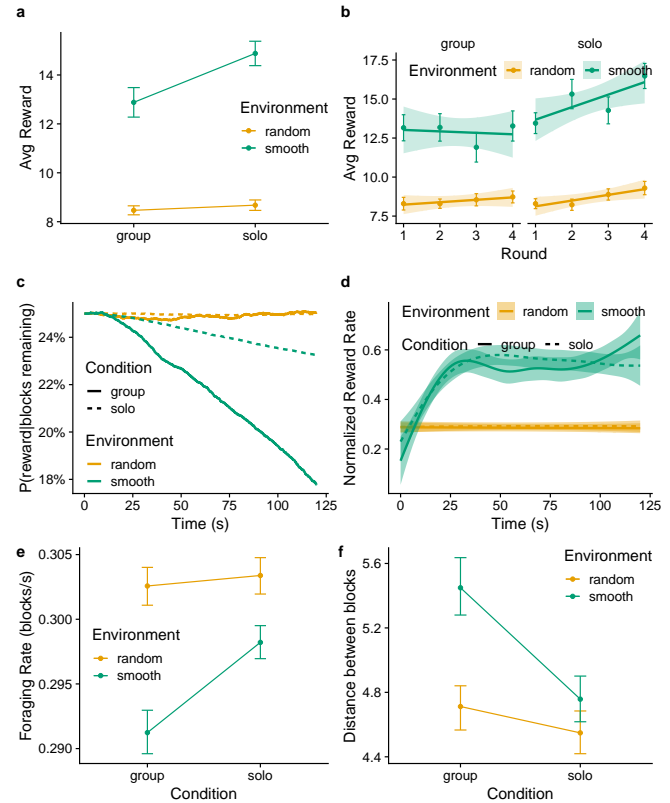


Figure 2: Behavioral results. **a)** Average reward across conditions. Dots indicate group means with error bars showing the standard error. **b)** Average over rounds, where each dot and error bar show the mean and standard error, and the lines and ribbons show the mean and 95% CI of a linear regression. **c)** The probability that a randomly sampled block contains a reward at any given time. This illustrates the higher depletion of rewards in more predictable environments (i.e., smooth) and when more participants are competing for the same finite number of resources (i.e., group). **d)** Normalized reward rate shows the average rate of rewards, normalized by the expectation of rewards (panel c). Lines and ribbons show the mean and 95% CI of a generalized additive model. **e)** Foraging rate defined in terms of blocks destroyed per second. **f)** The average distance between destroyed blocks.

while treating participants as random effects. The only reliable interaction with round was found in the smooth:solo condition ($b = 0.07$, 95% Highest Posterior Density (HPD): $[0.01, 0.12]$, $p(b > 0) = .99$). Thus, participants improved their ability to detect rewards in smooth environments over successive rounds when searching alone, but not in groups.

However, these reward differences were substantially influenced by the dynamics of reward depletion (Fig. 2c), where more predictable rewards (i.e., smooth) and more participants searching for the same finite number of rewards (i.e., groups) both contributed to a faster decay in the baseline probability that one of the remaining blocks contained a reward. Thus, we computed the normalized reward rate (Fig. 2d), by normalizing the instantaneous reward rate by the current ex-

pected reward rate $P(\text{reward}|\text{blocksRemaining})$. This measure shows how well participants perform relative to a dynamically changing random baseline. We found higher performance in smooth environments ($F(1,43) = 182.5$, $\eta^2 = .53$, $p < .001$), but no difference between group or solo conditions ($F(1,43) = 0.33$, $\eta^2 = .001$, $p = .57$). We also no longer found any significant interaction in the group:smooth condition ($F(1,43) = 0.03$, $\eta^2 < .001$, $p = .863$). Thus, when accounting for differences in expected rewards due to depletion, participants performed equivalently in the solo and group conditions.

Participants' foraging patterns were also influenced by the environment and social search conditions. Figure 2e shows the average foraging rate, which is defined as the number of blocks destroyed per second. Participants were more selective and foraged slower in smooth environments ($F(1,43) = 5.15$, $\eta^2 = .02$, $p = .028$) and in the group condition ($F(1,43) = 7.39$, $\eta^2 = .01$, $p = .009$). There was also an interaction between search condition and environment ($F(1,43) = 8.20$, $\eta^2 = .003$, $p = .006$), where participants were especially slower when combining the predictably smooth rewards with group dynamics.

These adaptive patterns of foraging selectivity are also present when analyzing the distance between destroyed blocks (Fig. 2f). Participants travelled further between foraged blocks in smooth environments ($F(1,43) = 14.3$, $\eta^2 = .07$, $p < .001$) and in the group condition ($F(1,43) = 34.6$, $p < .001$, $\eta^2 = .07$). We again see the same interaction between search condition and environment ($F(1,43) = 15.1$, $p < .001$, $\eta^2 = .03$), where participants in the group:smooth condition especially foraged over the longest distances.

Social Interactions

Next, we focus on the social interactions between participants in the group rounds. By recreating all experimental data in the Unity game engine (see methods), we were able to programmatically annotate all field of view (FOV) data. This allowed us to determine, when any given participant was visible to any other participant at all points in time. We then computed the average number of visible peers as a proxy for *social attention* (Fig. 3a), which is higher when more peers were visible and for longer durations. Consistent with the fact that social information had no value in random environments (due to unpredictable rewards), we found higher social attention in smooth environments (paired t-test: $t(43) = 2.7$, $d = 0.5$, $p = .011$). There was also a marginal interaction effect of round, where social attention tended to decrease over rounds in the random environment (Bayesian mixed effects regression: $b = -0.01$, 95% CI: $[-0.02, 0.01]$, $p(b < 0) = .86$; Fig. 3b). Thus, participants observed other participants less in random environments (where social information had no value), with the magnitude of this difference increasing over successive rounds.

In order to better understand the pairwise interactions between participants' visual attention, we constructed visibility networks for each round (Fig. 3c), where the directed edges are weighted based on the amount of time one participant

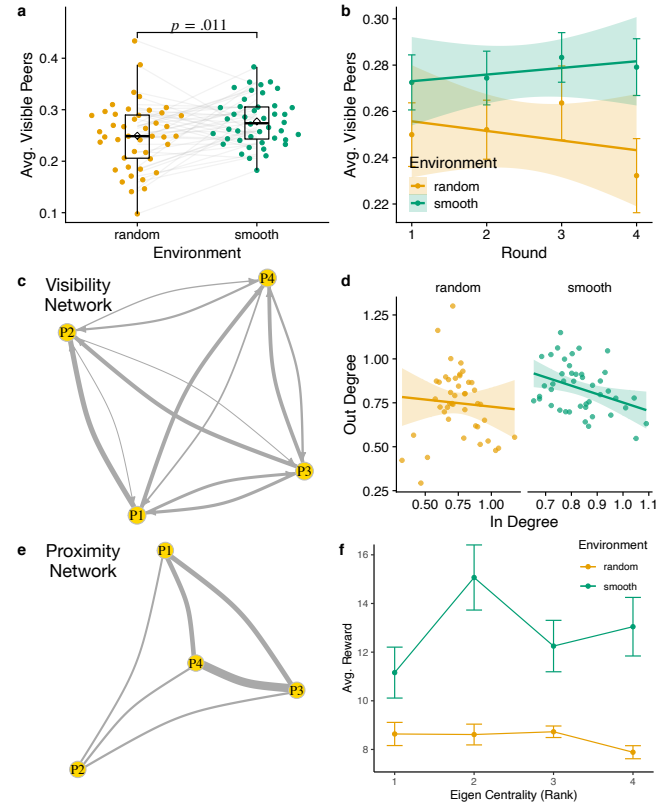


Figure 3: Social interaction results (group rounds only). **a**) The average number of visible peers (at any point in time). Each connected dot is a participant, with an overlaid Tukey boxplot providing group-level statistics. The diamonds show the group means. The p -value is for a paired t -test. **b**) Changes in social visibility over consecutive rounds. Dots and error bars indicate the aggregate mean and standard error, while the line and ribbon show the mean and 95% CI of a linear regression. **c**) An example of a visibility network, where each node is a participant and the directed edges are weighted by the proportion of time (within a round) that the target participant is visible to the observer. **d**) A comparison of the in- and out-degree of each participant (dot) with the lines indicating a linear regression. **e**) An example of a spatial network, where the undirected edges are weighted by the average spatial proximity between participants. **f**) Correspondence between the rank Eigen centrality (how connected one is to other connected nodes) of each participant with their average reward. Dots and error bars show the aggregate mean and standard error.

observed another. Based on these graphs, we computed the weighted in- and out-degree for each participant (Fig. 3d). Higher in-degree corresponds to being observed more frequently (i.e., celebrity factor), while higher out-degree corresponds to observing others more frequently (i.e., paparazzi factor). In smooth environments, we find a negative correlation between in- and out-degree ($r = -.37$, $p = .014$), but no correlation in random environments ($r = -.07$, $p = .652$). We obtain similar results using rank correlation (smooth:

$r_\tau = -.21$, $p = .044$; random: $r_\tau = -.11$, $p = .275$). This suggests that participants were more specialized in smooth environments, with celebrities who were frequently observed but rarely observed others, and paparazzi who observed others intensely but were seldom observed themselves.

We also built proximity networks where the undirected edges were weighted by the spatial proximity between participants (stronger edge weights for closer average distance; Fig. 3e). Based on these spatial relationships to other participants, we computed the Eigenvector centrality (EC) of each participant as a holistic measure of spatial proximity to other participants. Figure 3f shows the relationship between the rank EC (computed within each group) and average rewards. In random environments there was a negative but unreliable relationship between spatial proximity and reward ($r_\tau = -.19$, $p = .081$). But in smooth environments, we find a non-linear relationship, where intermediate levels of proximity produced the highest performance. Participants with a rank EC of 2 performed better than their peers with lower EC who were the most distant from the group ($t(20) = 2.3$, $p = .032$, $d = 1.0$), but with no significant difference to peers with higher EC ($t(31) = 1.7$, $p = .107$, $d = 0.6$), who were more central. These results suggest that an intermediate level of spatial proximity may have facilitated better performance in smooth environments, although there was no significant disadvantage to being more central. Thus, in smooth environments where social information was predictive of other rewards in the same area, it was better to be more central—despite increased competition for rewards—than to be at the outskirts, where imitation was more expensive due to increased travel costs.

Social Influence

In order to more directly measure social influence, we adapted methods developed to analyze the movement patterns of geo-tracked Baboons in the wild (Strandburg-Peshkin, Farine, Couzin, & Crofoot, 2015). This allows us to detect discrete “pull” events over arbitrary time scales, where the movement patterns of one participant (leader) pulls in another (follower) to imitate and forage in the same vicinity (Fig. 4).

We first computed the pairwise distance between all participants (Fig. 4a) and defined candidate pull events from min-max-min sequences. These candidate sequences were then filtered based on *strength*, *disparity*, *leadership*, and *duration* in order to be considered a successful pull.

Strength $S_{i,j}$ defines the absolute change in dyadic distance relative to absolute distance:

$$S_{i,j} = \frac{|d_{i,j}(t_2) - d_{i,j}(t_1)| |d_{i,j}(t_3) - d_{i,j}(t_2)|}{(d_{i,j}(t_1) + d_{i,j}(t_2)) (d_{i,j}(t_2) + d_{i,j}(t_3))}, \quad (1)$$

where $d_{i,j}(t_k)$ is the dyadic distance between participants i and j at time $k \in [1, 2, 3]$ (corresponding to the timepoints of the min-max-min sequence). We required pull events to have a minimum strength of $S_{i,j} > .1$, such that they correspond to meaningful changes in spatial proximity rather than minor “jitters” at long distance.




Figure 4: Pull events. **a)** Candidate pull events were selected from min-max-min sequences (dashed lines) of the pairwise distances between participants. These candidate sequences were then filtered by strength, disparity, leadership, and minimum duration (see text). **b)** An example of a successful pull. The orange and pink lines indicate the trajectories of the leader and follower (respectively), where the timepoints [1, 2, 3] correspond to the min-max-min sequence in panel a (dashed lines). The colored blocks illustrate the foraged blocks up until t_3 . **c)** The average number of successful pull events in each session (connected dots). The Tukey boxplots illustrate the aggregate statistics, with the diamonds indicating the aggregate mean. The p -value is for a paired t -test.

Disparity $\delta_{i,j}$ defines the extent to which one participant moves more than the other in each segment, relative to the total distance moved by both participants:

$$\delta_{i,j} = \frac{|\Delta x_i(t_1, t_2) - \Delta x_j(t_1, t_2)| |\Delta x_i(t_2, t_3) - \Delta x_j(t_2, t_3)|}{(\Delta x_i(t_1, t_2) + \Delta x_j(t_1, t_2)) (\Delta x_i(t_2, t_3) + \Delta x_j(t_2, t_3))}, \quad (2)$$

where $\Delta x_i(t_1, t_2)$ is the displacement between t_1 and t_2 . We filtered pull events to have a minimum disparity of $\delta_{i,j} > .1$, such that changes in spatial proximity were asymmetrically driven by one of the interaction partners.

Leadership is a simple binary filter requiring that the participant who moved more in the first segment (t_1 to t_2) moved less in the second segment (t_2 to t_3). We refer to the participant who moved the most in the first segment $\max_{a \in (i,j)} \Delta x_a(t_1, t_2)$ as the *leader* and the participant who moved the most in the second segment $\max_{b \in (i,j)} \Delta x_b(t_2, t_3)$ as the *follower*. Thus, successful pulls are defined as $a \neq b$, where the leader and follower are separate participants.

Duration was the final filter, where we required pulls to be at least 3 seconds in duration (since it takes 2.25s to destroy a block). After all filters were applied, the average pull duration was $15s \pm 0.73$ (SEM).

Altogether, we detected 135 successful pull events from the group rounds in our data. Figure 4a-b shows an example of a successful pull in a smooth environment. At t_1 , both the leader (orange) and follower (pink) are in similar locations, but begin to move in opposite directions. At t_2 , the leader has found a new patch at the center of the map while the follower has been unsuccessful at the south-east corner. Between t_2 and t_3 , the follower turns around and starts moving towards the leader and begins foraging in the same proximity.

When comparing the influence of environment on the frequency of pull events, we found a higher frequency of pulls

in smooth than random environments ($t(10) = 4.0$, $p = .003$, $d = 1.9$; Fig. 4c). This effect was also amplified over successive rounds. We fit a Bayesian mixed effects Poisson regression model to predict the number of pull events in each round, using the reward environment and round number as predictors, and treating session as a random intercept. Pull events increased over rounds in smooth environments ($b = 0.42$, 95% HPD: $[-0.02, 0.89]$, $p(b > 0) = .97$), whereas they tended to decrease over rounds in random environments ($b = -0.27$, 95% HPD: $[-0.76, 0.21]$, $p(b < 0) = .86$). This adaptation of pull frequency was consistent with the fact that social information was predictive of rewards in smooths environment, but had no value in random environments.

Discussion

Using a collective spatial foraging experiment implemented in an immersive virtual environment, we were able to bring together an unprecedented combination of behavioral data, including spatial trajectories, visual field data, and their complex social interactions. By analyzing foraging patterns, social interactions (visual and spatial), and social influence, we were able to shine new light on how groups collectively negotiate the balance between social and individual learning.

By manipulating the structure of rewards in the environment, we were able to study how groups adapt their search strategies to the value of social information. Smooth environments had more predictable rewards, such that observing when another player finds a reward provided actionable information about where to search next. In contrast, random environments had no predictable pattern of rewards, and thus time spent observing other players came only at the cost of lost opportunities for individual foraging without any benefits. Accordingly, participants observed each other less in random environments, and were less susceptible to social influence, as captured by a lower frequency of pull events.

Furthermore, our visibility analysis indicates that groups achieved a balance between individual and social learning through a mixed rather than a homogeneous equilibrium. In smooth environments, participants specialized as either the target of social attention or source of it. This asymmetric social attention structure may avoid the perils of maladaptive information cascades (Toyokawa, Whalen, & Laland, 2019; Tump, Pleskac, & Kurvers, 2020), since attention was selectively directed towards participants who rely more on individual learning.

In sum, we have only begun to fully leverage the richness of this spatial-visual dataset, where future work using computational models and more data will allow us to better understand the interaction dynamics that shape our social learning strategies.

References

Barkoczi, D., Analytis, P. P., & Wu, C. M. (2016). Collective search on rugged landscapes: a crossenvironmental analysis. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 918–923).

- Barnard, C. J., & Sibly, R. M. (1981). Producers and scroungers: a general model and its application to captive flocks of house sparrows. *Animal Behaviour*, 29(2), 543–550.
- Beauchamp, G. (2008). A spatial model of producing and scrounging. *Animal Behaviour*, 76(6), 1935–1942.
- Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16, 125–125.
- Ehn, M., & Laland, K. (2012). Adaptive strategies for cumulative cultural learning. *Journal of Theoretical Biology*, 301, 103–111.
- Enquist, M., & Ghirlanda, S. (2007). Evolution of social learning does not explain the origin of human cumulative culture. *Journal of Theoretical Biology*, 246(1), 129–135.
- Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, 23(5), 373–393.
- Kobayashi, Y., & Ohtsuki, H. (2014). Evolution of social versus individual learning in a subdivided population revisited: comparative analysis of three coexistence mechanisms using the inclusive-fitness method. *Theoretical Population Biology*, 92, 78–87.
- Kurvers, R. H., Prins, H. H., van Wieren, S. E., van Oers, K., Nolet, B. A., & Ydenberg, R. C. (2009). The effect of personality on social foraging: shy barnacle geese scrounge more. *Proceedings of the Royal Society B: Biological Sciences*, 277(1681), 601–608.
- Miu, E., Gulley, N., Laland, K. N., & Rendell, L. (2020). Flexible learning, rather than inveterate innovation or copying, drives cumulative knowledge gain. *Science Advances*, 6(23), eaaz0286.
- 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.
- Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, 90(4), 819–831.
- Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D., & Crofoot, M. C. (2015). Shared decision-making drives collective movement in wild baboons. *Science*, 348(6241), 1358–1361.
- Toyokawa, W., Whalen, A., & Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 3(2), 183–193.
- Tump, A. N., Pleskac, T. J., & Kurvers, R. H. (2020). Wise or mad crowds? The cognitive mechanisms underlying information cascades. *Science Advances*, 6(29), eabb0266.
- Wu, C. M., Schulz, E., Garvert, M. M., Meder, B., & Schuck, N. W. (2020). Similarities and differences in spatial and non-spatial cognitive maps. *PLOS Computational Biology*, 16, e1008149.