

The Effects of Peer Information on Problem-Solving in a Networked Group

Thomas N. Wisdom (tnwisdom@indiana.edu)

Department of Psychological and Brain Sciences, 1101 E. 10th Street
Bloomington, IN 47404 USA

Xianfeng Song (xsong@indiana.edu)

Department of Physics, 727 E. 3rd Street
Bloomington, IN 47405 USA

Robert L. Goldstone (rgoldsto@indiana.edu)

Department of Psychological and Brain Sciences, 1101 E. 10th Street
Bloomington, IN 47404 USA

Abstract

In this experiment, we implemented a problem-solving task in which groups of participants simultaneously play a simple puzzle game, with score feedback provided after each of 24 rounds. Each participant in a group is allowed to view and imitate the guesses of others during the game. Results show that when the utility of others' innovations is unambiguous, individuals base their own solutions on personal innovation and productively imitate other players' innovations early on, and that this tendency to imitate is proportional to the relative amount of information available from others. Average trends of innovation and imitation decreased across rounds as player guesses stabilized and scores increased. Mean scores and imitation increased with group size, while individual innovation decreased. Results are consistent with previously studied social learning strategies in several taxa.

Keywords: Social learning; innovation; imitation; problem solving; innovation diffusion.

Background

In uncertain situations, one can often obtain useful information from others that would be inaccurate or inefficient to procure on one's own. A particularly important use of this kind of socially-mediated information is "social learning," which has been defined by Boyd & Richerson (2005) as "the acquisition of behavior by observation or teaching from other conspecifics." Common examples of this include learning "rules of thumb" for accomplishing a task from a more experienced practitioner, or soliciting recommendations for restaurants or films from friends. Evidence has been presented for social learning in many nonhuman animals, including foraging choices in starlings (Templeton & Giraldeau, 1996), foraging sites and escape routes in guppies (Reader, Kendal & Laland, 2003), food preferences in various rodent species (Galef & Giraldeau, 2001), mate choices in black grouse (Höglund, Alatalo, Gibson, & Lundberg, 1995), and many others.

But resources must generally be spent in gathering or creating new information before it can be shared. Thus an apparent trade-off develops between individually learning or producing locally novel information (innovation) and social learning (imitation). Despite the negative connotations of

"conformity," it has been shown that the tendency to imitate others rather than innovate ("conformity bias") can be adaptive in uncertain environments (Boyd & Richerson, 1985; Kameda & Nakanishi, 2002). Cultural conventions are thought to be a form of large-scale imitation of behaviors that evolve along with their associated populations, subject to accompanying adaptive pressures (Boyd & Richerson, 2005).

Humans' rare talent among animals for direct and flexible imitation has been called "no-trial learning" (Bandura, 1965). This talent allows an imitator to add new behaviors to his or her repertoire without the costs of trial-and-error learning. However, the population-level benefits of imitation are not necessarily straightforward. A simulation by Rogers (1988) of agents in a nonstationary environment showed that if the avoidance of learning costs is the only benefit of imitation, then the addition of imitators to a population of individual learners does not improve the average fitness of the population.

Boyd & Richerson (1995) confirmed and extended this result, but showed conversely that imitation provides a net population fitness benefit if it makes individual learning more accurate or less costly (e.g. by allowing selective learning or cumulative improvements). Kameda & Nakanishi (2003) found that if agents can selectively switch between individual learning and imitation, a "cultural" population has a mean fitness advantage over an "acultural" population. Even though it operates at the individual level, imitation can generate substantial population-level effects when iterated over time and space.

Laland (2004) reviews studies of several strategies for the selection of imitation ("copying") over innovation ("asocial learning") in terms of *when* to copy, and *who* to copy. The *when* category includes the *copy when asocial learning is costly* strategy, which selects imitation when innovation entails excessive known costs in energy, time or risk, and *copy when uncertain*, which selects imitation when the environment or the payoff for innovation is uncertain. The *who* category includes the *copy successful individuals* strategy, which selects imitation of the action taken by successful observed models; and *copy if better*, which is similar to the latter but entails a comparison with one's

previous actions. Though many theoretical and animal accounts of these strategies exist, the empirical evidence for them in humans is still somewhat unclear.

The present experiment explores the dynamics of selective learning and cumulative improvements using a simple problem-solving task, in which guesses are composed of sets of discrete units, and each player can view and imitate the guesses of other players. Players receive round-based score feedback about the quality of their guesses, and the score function takes into account linear and interaction terms for the guess units to create a complex problem space. Players can see their own and each others' most recent guesses and scores, and may copy entire guesses or parts thereof from other players (see Figure 1).

As in a precursor to this study (Wisdom & Goldstone, 2007), the in-game interaction between players was intentionally impoverished so as to allow examination of imitation and innovation behavior unencumbered by more complex social interactions, such as direct communication among players. This paradigm can be considered a general case for specific group interactions like explicit competition or cooperation, in which task goals are more structured by incentive. The principal differences between the present study and the precursor study referenced above are (1) direct recording rather than estimation of occurrences of innovation, imitation, etc., (2) use of a more challenging, interesting, and general task, and (3) the addition of an “assistive memory” (last round / best round) display to reduce effects of participant memory limitations.

Though the task environment within each condition does not vary over time as in many studies of imitation, the problem space is designed to be sufficiently large and complex so as to provide substantial uncertainty about optimal guesses and strategies. In addition, the difficulty of the problem is manipulated across conditions by changing the size of the problem space, via the size of the set of units to be chosen from and the size of the subset that may be chosen and evaluated at one time. A range of group sizes is tested to investigate the effects of the availability of other players’ guess information on individual performance and strategies.

We expected to confirm the following results from a previous, similar experiment which used a different task (Wisdom & Goldstone, 2007): there should be an increase in the rate of imitation with group size, and a qualitative deficit in the performance of isolated individuals compared to those in groups. Furthermore, it was anticipated that participants would use the *copy when uncertain* strategy, so that there would be more imitation earlier in each game when the problem space is relatively unexplored, as well as in the more difficult condition. It was also expected that participants would use the *copy successful individuals* strategy, so that the participants with the highest scores would be copied most often. The two strategies above were selected because of their confirmed prevalence in nonhuman social learning, and their relative scarcity in the human social learning literature, in addition to their anticipated simplicity of detection using this task.

Method

Participants were recruited from the Indiana University Psychology Department undergraduate subject pool, and were given course credit for taking part in the study. Participants populated each session by signing up at will for scheduled experiments with a maximum capacity of 9 persons. 49 sessions (containing a total of 201 participants) were run. 9 sessions (containing a total of 39 participants) were discarded due to network or software problems. The remaining 162 participants were distributed as shown in Table 1.

Table 1: Distribution of participants across group sizes

| Group size | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
|----------------|---|----|----|----|----|----|----|----|----|
| # Sessions | 8 | 6 | 5 | 5 | 5 | 2 | 4 | 3 | 2 |
| # Participants | 8 | 12 | 15 | 20 | 25 | 12 | 28 | 24 | 18 |

We implemented the experimental game using custom software programmed in Java and Flash, which runs in a web browser. Each participant uses a mouse to interact with the game. All participant computers communicate with a game server, which updates score and team information for participants at the end of each round, and records data. In the game itself, participants attempt to maximize the number of points earned by their chosen subsets ("teams") from a set ("league") of creature icons over 24 rounds. The display includes an area for the participant's own current team, another that can be toggled to show the participant's previous round team or best-scoring team so far in the game (along with the associated score), a league area which shows all of the icons (team members) that can be chosen from, and indications of the current round in the game and the amount of time remaining in the current round.

In games with more than one participant, each participant's display also shows the teams and associated scores of all other participants in the previous round. Icons can be copied from any part of the display to the participant's current team by dragging and dropping them with the mouse, except for those already on the current team, which are faded in the display and non-clickable. Icons may be present on more than one participant’s team in the same round. The current team can be replaced entirely by another team by using the score box above the latter as a "handle" and dragging it to the current team area. A screenshot of the participant interface for a game with 5 participants is shown in Figure 1.

At the beginning of each session, players are given a hands-on demo of the game (including the various ways to move creatures to one's current team), and further informed about the mechanics of the game, including the following information. Each game consists of 24 rounds, and each round is 10 seconds long. Score feedback is given after each round: if the participant's score has improved from the




Figure 1: Example of experiment interface.

previous round, the display counts up to the new score and turns green, and if it has worsened, the display counts down to the new score and turns red. At the end of each game, the display shows the player's final score, along with a table of the scores of each player in each round of the game, which is sorted by average score. The player's own scores are highlighted to provide a notion of their relative performance without placing competitive emphasis on it. Players are instructed to do their best to find a high-scoring team in each round. At the beginning of each game, each player's team is a random selection of creature icons from the league.

Each participant group plays 8 games, of which half have a large league and team size (48 and 6, respectively), and half smaller (24 and 5). These two parameter settings are intended to vary the level of difficulty of the game, with the former being more difficult because the score distribution and combinatorics make high-scoring teams rarer than in the latter case, even though higher scores are possible.

In each game, each icon is associated with a certain positive number of points, and several special pairs of icons are associated with separate score bonuses or penalties. The score for a team is computed by summing the individual point values for each icon and then adding or subtracting the value of any special pairs present. The pairs do not overlap, and the distribution is designed to be challenging: pairs which give large positive bonuses are distributed among icons with small individual point values, and pairs which give large negative penalties are found among icons with large individual point values (see Figure 2).

Individual point values per icon range from 1 to 8 points, and pair interaction values range from -20 to 20 points, so that the possible score ranges for the large and small league and team size combinations are [-6,60] and [-6,51], respectively. Participants are not given information about the maximum score, the score distribution, and the identity of the interaction pairs, though they could conceivably be deduced during play. The icons' display position and




Figure 2: Distribution of individual and pair interaction point values for a League Size of 24.

associations with the point distribution are shuffled randomly for each game, so that their appearance and placement in the display will not give clues as to their point values during the course of an experiment session.

In each round, the following data are automatically recorded for each player: the icons on the current team at the end of the round, the source of each icon, and the resulting score. The source information indicates whether each icon is unchanged from the previous round (Incumbent), copied from the player's own previous round team after initially being removed from the team (Previous Round), copied from the player's own best-scoring team so far (Best Round), chosen from the league display (League), or copied from another player's team (Copy). The latter includes the persistent identifier of the copied player, to allow further analyses of imitation decisions. Though any creature icon may always be copied from the League display, it is assumed that this action is not substituted for copying from another intended source with any significant frequency, due to the additional effort required.

Results

Dependent variables were averaged across participants to give measures for the group's aggregate activity in a session.

Score

Scores for each game were normalized to the range [0,1] according to the score range possible with the associated league and team size. The average score for all participants in all rounds was .597. Curve fit estimation for a linear regression of score data against game round showed a significant positive relationship ($F(1,958)=1015, p<.0001$) (see Figure 3). The average improvement in score over the course of a game was .261, and there was a significant but comparatively small average increase (.051) in scores across game order within a session ($F(1,318)=11.9, p<.001$). Trends in score over rounds did not significantly differ with group size. Scores increased significantly with group size ($F(1,38)=70.5, p<.0001$), with an average score difference of .101 between isolated participants and those in the largest group size of 9 (see Figure 4). In similar analyses against league size, scores decreased significantly from the smaller to the larger league size ($F(1,78)=31.7, p<.0001$), from .607 to .54, an average difference of .067.




Figure 3: Average score vs. round.

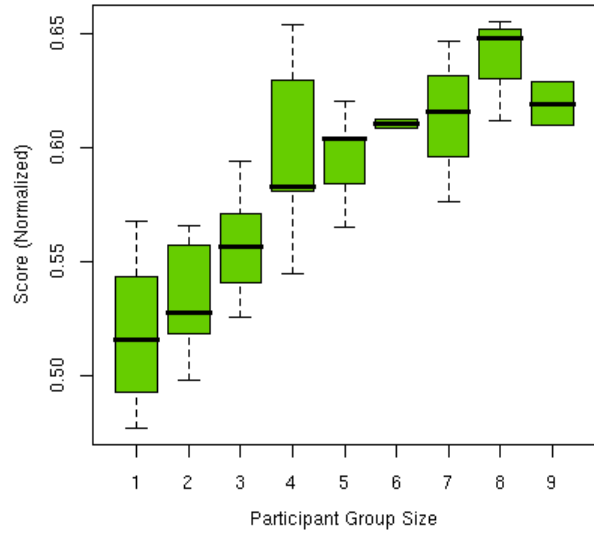


Figure 4: Average score vs. group size.




Figure 5: Average Choice Source vs. round.




Figure 6: Average Choice Source vs. group size.

Choice Source

The first two rounds of choice source data were excluded from round-based analysis due to artifacts related to the beginning of the game (neighbor score data not provided in the first round, etc.) Regression analyses of the remaining data showed that across rounds, there were significant increases in the average proportions of the Incumbent ($F(1,878)=5.8, p=.016$) and Best Round ($F(1,878)=26.6, p<.0001$) choices in each team, while significant decreases were observed in the League ($F(1,878)=61.6, p<.0001$) and ($F(1,878)=18.5, p<.0001$) choices (see Figure 5). As group size increased, the proportion of the choice increased ($F(1,38)=75.6, p<.0001$), while accompanying decreases were observed in proportions of the Best Round ($F(1,38)=15.4, p<.001$) and League ($F(1,38)=24, p<.0001$) choices (see Figure 6).

The proportion of the Previous Round choice source was negligible on average (approximately .01 or less). No significant relationships were found between the use of any choice source and League Size, or game order within sessions. The overall imitation rate (the proportion of rounds for all participants in which the Copy choice proportion was greater than zero) was .293. When imitating, participants copied a participant with the highest score 82.6% of the time, and copied participants with higher scores than their own 92.4% of the time.

League Coverage

In order to measure the heterogeneity of the solutions within a group, we calculated the League Coverage, defined as the proportion of icons in the league represented on one or more participants' teams during a given round, normalized by the participant group size.




Figure 7: Normalized League Coverage vs. Round, by Participant Group Size.

Regression analyses showed that average League Coverage decreased with game round ($F(1,958)=19.3$, $p<.0001$) and with increasing league size ($F(1,78)=20.6$, $p<.0001$). League Coverage also decreased with increasing group size ($F(1,38)=194$, $p<.0001$), and trends over rounds were more strongly negative for increasing group size (see Figure 7).

Discussion

Overall, the results show changes in strategy and performance, both over time and across group sizes. The positive relationship of score with game round confirms that participants were improving during the game. Despite the fact that a higher maximum score is possible in the larger league size condition, the significant drop in average scores in this condition confirms that it was indeed more difficult, as intended. No qualitative differences were found in the performances of isolated and grouped individuals.

The Incumbent, Best Round, League, and Copy choice sources will hereafter be discussed respectively as *retention* of guess information, *retrieval* of information from a previous best guess, *innovation*, and *imitation*. The increases over rounds in the proportions of retained and retrieved information and accompanying decreases in innovation and imitation show that individual solutions were established over the course of the game, rather than being produced by random behavior or exclusively by copying. Larger amounts of innovation and imitation occurred early on, and less in later rounds when better guesses were found and smaller changes were made. This is consistent with the *copy when uncertain* strategy in that imitation occurs more often early on in each game when participants have less experience with the current problem space. Though innovation decreased over rounds as well, the increases in retrieval and retention show that individual learning was taking place, as they function as a kind of

memory for successful (or at least satisfactory) guesses for each participant. This also shows that the improvements in score within games were due to cumulative improvements by participants employing the *uncertainty-reduction* capability of social learning (Kameda & Nakanishi, 2002).

Since the ratio of innovation to imitation did not change significantly across rounds, it may be that they were treated as complements for one another rather than substitutes. This may have been due to the relatively low cost to innovation in terms of risk. Though it was possible to pursue fruitless innovation and thus lower one's score relative to neighbors, there was always the possibility of copying the highest-scoring participant and "catching up" with the rest of the group.

The increase in imitation with larger group sizes implies that copying is more attractive when there are more neighbors from which to copy, perhaps simply because there is a greater probability of encountering a neighbor with a score greater than one's own. A different kind of consistency with the *copy when uncertain* strategy is shown here: imitation is favored when the payoff for innovation is relatively uncertain, compared to the abundant unambiguous information available about the content and utility of neighbors' guesses. The accompanying decreases in innovation and retrieval in larger group sizes imply that individual learning accordingly becomes less attractive than social learning when there are more neighbors to imitate. A related interpretation is that having more neighbors may simply make it more likely that one of them will imitate and improve upon one's own guesses more productively than oneself, and thus that recourse to a previous best guess is not as useful.

In any case, the correlation of score increase with increasing group size implies that the above strategy changes are advantageous under these conditions. This in turn implies that the adaptive value of imitation in this context is due to its facilitation of selective learning and the generation of cumulative improvements (Boyd & Richerson, 1995; Kameda & Nakanishi, 2003).

The use of the *copy-successful-individuals* strategy was confirmed by the tendency of imitators to copy the highest-scoring individuals. Furthermore, equally successful individuals do not simply copy each other at random, since nearly all imitation results in score increases for the imitator, confirming the use of the *copy-if-better* strategy.

It is slightly curious that the difference in difficulty between the two League Size conditions did not produce an accompanying difference in guessing strategies. In particular, this is in conflict with the findings of Baron, Vandello, & Brunzman (1996), which indicate that copying increases with task difficulty.

The decrease in guess heterogeneity among participants (League Coverage) within games, and the accompanying increase in average score, confirm the intuition that participant guesses will tend to converge toward similar higher-quality guesses over the course of a game. The increasingly negative round-based trends in guess

heterogeneity with increasing group sizes show that the convergence of guesses occurs more quickly and to a greater extent with larger groups. This is in accordance with the above-noted increase in copying with larger group sizes. The fact that average final scores are less than 70% of the maximum possible score implies that, especially in larger groups, participants are settling on good but suboptimal solutions due to insufficient search of the multimodal, “rocky” problem space. This result agrees with the findings of Mason, Jones, and Goldstone (2005), that fully-connected groups (like the ones in this experiment) performed relatively poorly on a multimodal problem space, while more sparsely-connected groups (small-world networks) found optimal solutions more reliably, though more slowly.

The use of differing network structures was not examined in this experiment, nor were other potentially interesting manipulations such as adding noise or explicit costs to information gained by individual or social learning; varying the “rockiness” of the problem space; changing the problem space over time; or allowing more complex forms of social communication between participants. The development of selective learning strategies and cumulative improvements are likely to be strongly affected by such variations, and this paradigm can be extended to pursue them.

Tomasello (1994) contends that human capabilities for selective and cumulative learning constitute a “ratchet effect” that allows culture to develop stably across generations, and that this effect may be unique to humans. Though this study presents a greatly simplified environment for such learning, it confirmed and extended several previous theoretical and empirical results in the field of social learning. Further understanding of the structure and dynamics of these phenomena is vital to the study of technological and cultural development.

Acknowledgements

The authors would like to thank Todd Gureckis for his help in designing the experiment. This work is funded by Department of Education, Institute of Education Sciences grant R305H050116, and National Science Foundation REC grant 0527920.

References

Bandura, A. (1965). Vicarious processes: a case of no-trial learning. In L. Berkowitz (Ed.) *Advances in Experimental Social Psychology, Vol. II*. New York: Academic Press.

Baron, R.S., Vandello, J.A., & Brunsman, B. (1996). The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology, 71*(5), 915-927.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Boyd, R., Richerson, P.J. (1995). Why Does Culture Increase Human Adaptability? *Ethology and Sociobiology, 16*, 125-143.

Boyd, R. & Richerson, P.J. (2005). *The origin and evolution of cultures*. New York: Oxford University Press.

Galef, B.G. Jr., & Giraldeau, L.A. (2001). Social influences on foraging in vertebrates: causal mechanisms and adaptive functions. *Animal Behaviour, 61*(1), 3-15.

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior, 19*, 215 – 241.

Höglund, J., Alatalo, R.V., Gibson, R.M., & Lundberg, A. (1995). Mate-choice copying in black grouse. *Animal Behaviour, 49*(6), 1627-1633.

Kameda, T., & Nakanishi, D. (2002). Cost-benefit analysis of social/cultural learning in a non-stationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior, 23*, 373-393.

Kameda, T., Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers’s question revisited. *Evolution and Human Behavior, 24*, 242 – 260.

Laland, K.N. (2004). Social learning strategies. *Learning & Behavior, 32*(1), 4-14.

Mason, W. A., Jones, A., & Goldstone, R. L. (2005). Propagation of innovations in networked groups. *Proceedings of the Twenty-seventh Annual Conference of the Cognitive Science Society*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Reader, S.M., Kendal, J.R., & Laland, K.N. (2003). Social learning of foraging sites and escape routes in wild Trinidadian guppies. *Animal Behaviour, 66*(4), 729-739.

Rogers, A.R. (1988). Does biology constrain culture? *American Anthropologist, 90*, 819-831.

Rogers, E.M. (2003). *Diffusion of Innovations* (5th ed.). New York: Free Press.

Templeton, J.J., & Giraldeau, L.-A. (1996). Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behavioral Ecology and Sociobiology, 38*, 105-113.

Tomasello, M. (1994). The question of chimpanzee culture. In R. Wrangham, W. McGrew, F. de Waal, & P. Heltne (Eds.), *Chimpanzee cultures*. Cambridge, MA: Harvard University Press.

Wisdom, T.N., & Goldstone, R.L. (2007). Group Problem Solving Behavior in a Networked Puzzle Game. *Proceedings of the Twenty-ninth Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.