Title

The Benefit of Imitating Particular Individuals

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN
1069-7977

Authors
Sakamoto, Yasuaki
Shi, Hongyuan

Publication Date
2010

Peer reviewed

The Benefit of Imitating Particular Individuals

Yasuaki Sakamoto (ysakamot@stevens.edu)
Hongyuan Shi (hshi@stevens.edu)
Center for Decision Technologies
Howe School of Technology Management, Stevens Institute of Technology
Hoboken, NJ 07030 USA

Abstract

We examined the benefits of different search strategies by testing four computational models. In one model, agents in a group always innovated. The other three models incorporated some mechanisms of imitation. In the second model, each agent imitated the best solution of a random other. In the third model, each agent followed preferential attachment and imitated the best solution of the agent that was asked by many agents. In the fourth model, each agent developed a familiarity with an agent based on how often it asked a certain agent, and imitated this agent. In two simulation studies, following the most popular or the most familiar agent resulted in a good compromise between efficiency and diversity in finding good solutions. People's desire to follow particular individuals may be a key to their adaptive behavior, allowing them to disseminate ideas efficiently while encouraging the exploration of new ideas.

Keywords: Innovation and imitation; computational modeling; social learning; search.

Introduction

How do we search for information? Some individuals like to innovate. Others like to imitate. We all engage in both. Because we are social beings, we often rely on other's behavior to shape our own behavior. By observing and imitating others, people can entertain solutions that they would not have even considered otherwise (Bandura, 1965). The creation of innovative solutions (Kraatz, 1998), the evolution of language (Smith et al., 2003), and the development of culture (Dennett, 1995) all result from the process of iterated learning, in which people learn from the previous outputs of others.

In the current work, we examine the benefits of different types of search strategies through computer simulation. We know that whereas too much innovation results in poor dissemination of good solutions, too much imitation results in under exploration of good solutions (Gureckis & Goldstone, 2006). A group of people needs to both innovate and imitate to prosper. But when should we innovate and when should we imitate? Who should we observe if we decide to imitate?

When people are unsure about the best solution, they use other's information as an indicator of what is best (Cialdini & Goldstein, 2004; Deutsch, & Gerard, 1995; Festinger, 1954; Sherif, 1935). People also adopt other's information due to their desire to be liked and to not appear deviant (Asch, 1956; Deutsch, & Gerard, 1995). This imitation behavior is consistent with the principle of preferential attachment (Barabási & Albert, 1999), in which people are attracted to already popular solutions. For example, people instantly get in line when they see a long line outside of a cupcake store, assuming that the store must be offering some really good cupcakes. If everyone imitates, however, it will be difficult for the group to find another cupcake store that also serves really good cupcakes. Thus, imitation leads to efficient problem solving when there is a single best solution. When there are multiple good solutions, however, imitation can lead the group to quickly converge to a single solution, under-exploring the others: some people need to explore other possibilities.

For studying innovation and imitation, we used a simple search game, inspired by a recent social learning tournament (http://www.intercult.su.se/cultaptation/tournament.php). In our game, five agents guessed an action value between 0 and 100, and received as feedback the number of points obtained from the guess. A function converted the guess to a payoff. The agents did not know the function and did not try to learn it. They simply stored the guessed action value that was associated with the highest payoff. The payoff distributions are displayed in Figures 1 and 2. In one case, the search space had a single peak at action 80 as shown in Figure 1. In another case, the search space had three peaks at action values 10, 50, and 80 as shown in Figure 2. Although the game may seem overly simple and artificial, it is analogous to many tasks we encounter every day (see Page, 2007).

The five agents, A_1, …, A_5, selected to either innovate (randomly select a value between 0 and 100) or imitate (receive another agent's value with the highest payoff) in turn. Four groups are simulated:

1. **Innovate**: Agents only innovated.
2. **Ask Random**: Agents imitated a randomly selected agent. The preference weight of A_i asking A_j, p_{ij}, was equal for all j.
3. **Ask Majority Preference**: Agents imitated another agent who was imitated by many others. That is p_{ij} was determined by the number of times A_j was asked by other agents, m_j. This group followed the principle of preferential attachment, and conformed to the majority's behavior.
4. **Ask Individual Preference**: Agents imitated another agent based on how often they asked a certain agent:

$$p_{ij} = P_1 + \frac{P_2 - P_1}{1 + \exp[-C(f_{ij} - F)]}$$

where P_1 = 0, P_2 = 10, C = 0.2, and F = 15. The ask history, f_{ij}, tracked the number of times A_i imitated A_j. Agents maintained a counter for every other agent it had interaction with. They followed the footstep of a particular agent they became familiar with.

The imitating agent always received another agent's current best solution. That is, the asked agent always returned the action value associated with the highest payoff that it previously guessed. In the current simulation, when asked agents returned worse solution than the existing one (i.e., the imitating agent had a better solution than the one asked), the agent innovated on the next round. Likewise, when asking someone does not result in good solution, humans often explore the environment by themselves. After innovating once, the agent tried to imitate again. Without this innovation round, always imitating can quickly converge to an action value regardless of its payoff.

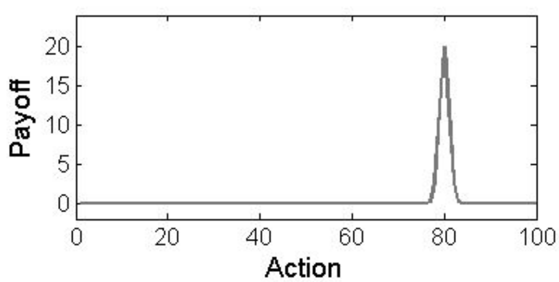

Figure 1: The distribution of payoff in Simulation 1 is shown. There is a single peak at action 80.

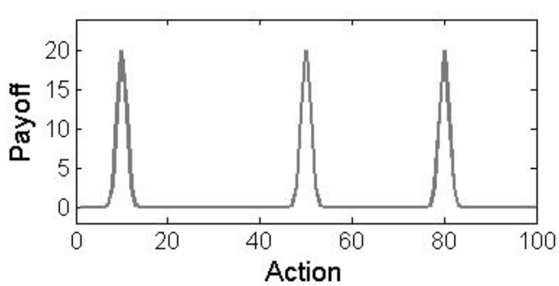

Figure 2: The distribution of payoff in Simulation 2 is shown. There are three peaks at action 10, 50, and 80.

Which group will result in all agents finding the value associated with the highest payoff most efficiently? Imitating others will help disseminate ideas. But which type of asking is best? Humans often conform to the group, similar to the Majority Preference model (e.g., Cialdini & Goldstein, 2004; Sakamoto et al., 2009). They also build familiarity for a particular other, and follow this individual (e.g., Sadlon, et al., 2009; Sakamoto et al., 2008). This following behavior allows us to make near-optimal decision in a limited amount of time in many different social circumstances (Gigerenzer et al., 2000). At the same time, vocal group members often sway the opinion of individuals, and thus the opinions produced by a group may only reflect those of a small subset of the group. Then, imitating others may not be advantageous when there are multiple good solutions to find. In this case, the Innovate group may be successful because the group has no social influence that converges their solutions. Previous work has focused on how given social network structures influence the dissemination of ideas (e.g., Mason & Goldstone, 2008). In the current work, the agent's behavior determines the kinds of social networks built and thus how information is spread within the group.

Simulation Study 1

In Simulation Study 1, the four groups of agents searched for the action with the highest payoff in a space with a single best solution as shown in Figure 1. The agents did not know what the maximum payoff was. Each group had five agents that all followed the same behavioral rule as described previously: innovate, ask random, ask majority preference, or ask individual preference. Each group had 500 cycles to search, each cycle consisting of an agent taking its action. We used 500 cycles so that each group would perform well at the end and we could see the entire course of evolution. Each group was simulated 30 times.

Figures 3 to 6 show the results from the four groups. The left most graph of each figure shows the evolution of total payoff (sum of all agents' payoffs) over the course of 500 cycles, averaged over 30 simulations. During the first 200 cycles, the Innovate model and the Majority Preference model lag behind the Random model and the Individual Preference model. The Innovate model is especially far behind the other models early on. After 300 cycles, the Innovate model is performing the worst, the Random model performing the best, and the two preference models in between. After 400 cycles, the two preference models catch up with the Random model, while the Innovate model is still behind the other models. At 500 cycles, every model has nearly all agents discovering the action with the highest payoff.

The middle three histograms in each figure show the frequency of total payoff for the 30 simulations. At 50 and 100 cycles, the disadvantage of the Innovate model is apparent: no simulation resulted in total payoff of 80 or 100. At 500 cycles almost all 30 simulations for each group result in every agent knowing the best action.

The color map on the right side of each figure shows the evolution of each agent's payoff averaged over 30 simulations. The Innovate model is darker in general, indicating that it took longer to find good solutions than the other models. In addition, the Innovate model has darker horizontal band, indicating that some agents had hard time innovating a good solution. In contrast, the other three models incorporating imitation disseminated the best action efficiently. The Random model was especially quick at disseminating good solutions.

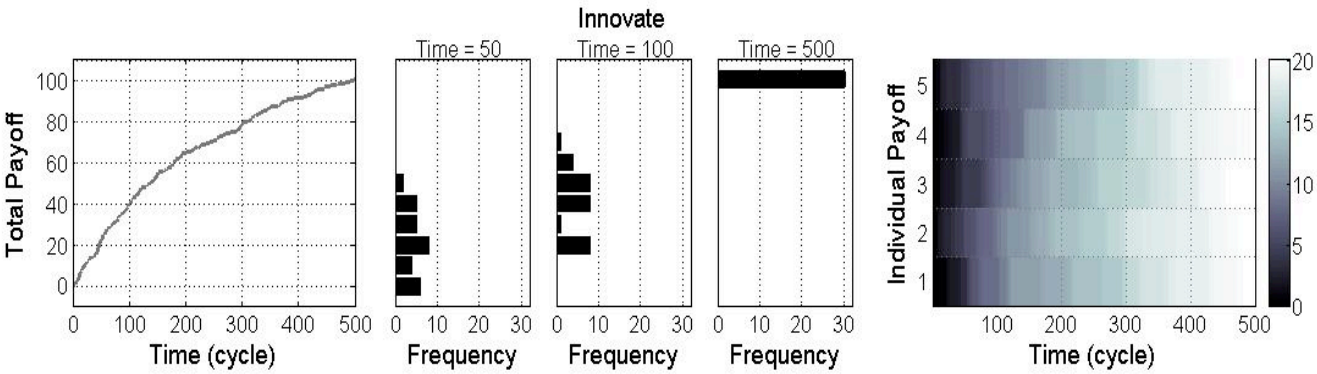

Figure 3: The results from the Innovate model in Simulation Study 1 are shown. The left most graph shows the evolution of total payoff (sum of all agents' payoffs) over the course of 500 trials, averaged across 30 simulations. The middle three histograms show the frequency of total payment for the 30 simulations at 50, 100, and 500 cycles. The color map on the right side shows the evolution of each agent's payoff averaged over 30 simulations.

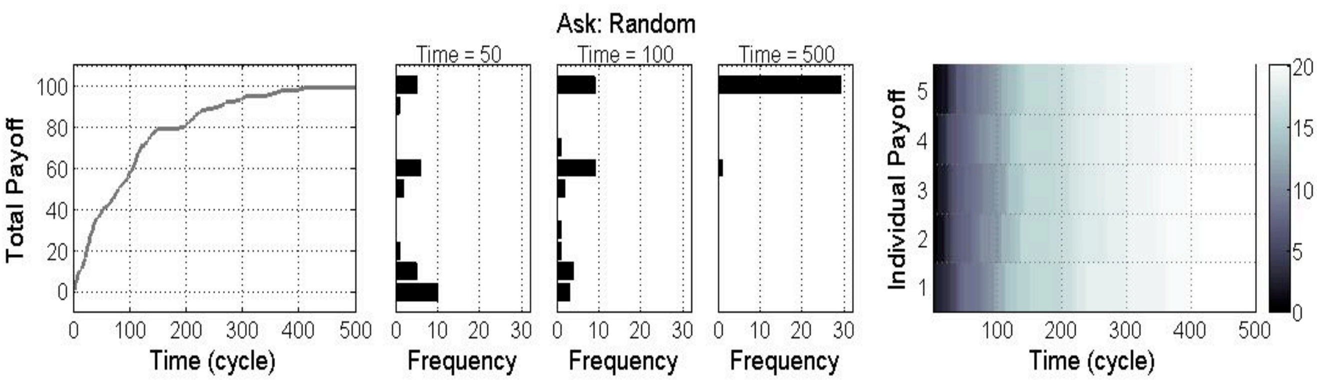

Figure 4: The results from the Random model in Simulation Study 1 are shown. The left most graph shows the evolution of total payoff (sum of all agents' payoffs) over the course of 500 trials, averaged across 30 simulations. The middle three histograms show the frequency of total payment for the 30 simulations at 50, 100, and 500 cycles. The color map on the right side shows the evolution of each agent's payoff averaged over 30 simulations.

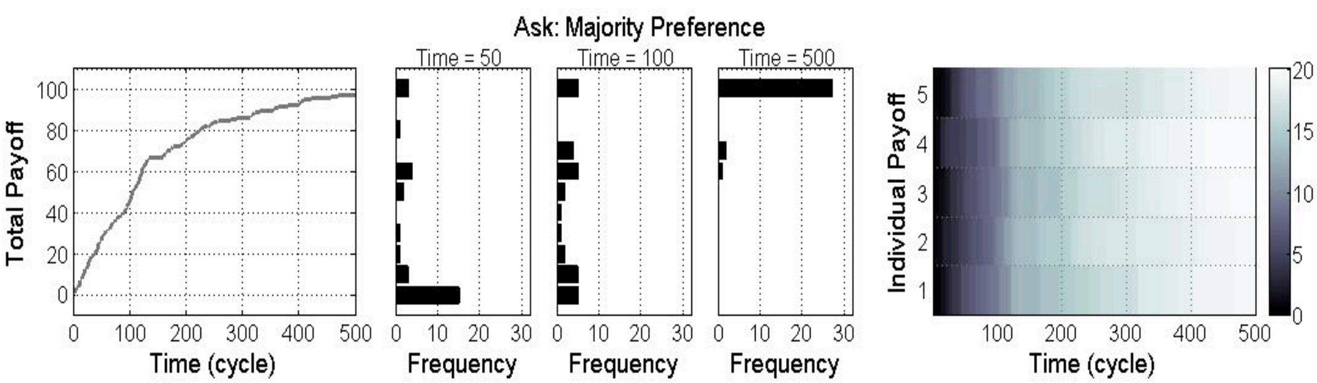

Figure 5: The results from the Majority Preference model in Simulation Study 1 are shown. The left most graph shows the evolution of total payoff (sum of all agents' payoffs) over the course of 500 trials, averaged across 30 simulations. The middle three histograms show the frequency of total payment for the 30 simulations at 50, 100, and 500 cycles. The color map on the right side shows the evolution of each agent's payoff averaged over 30 simulations.

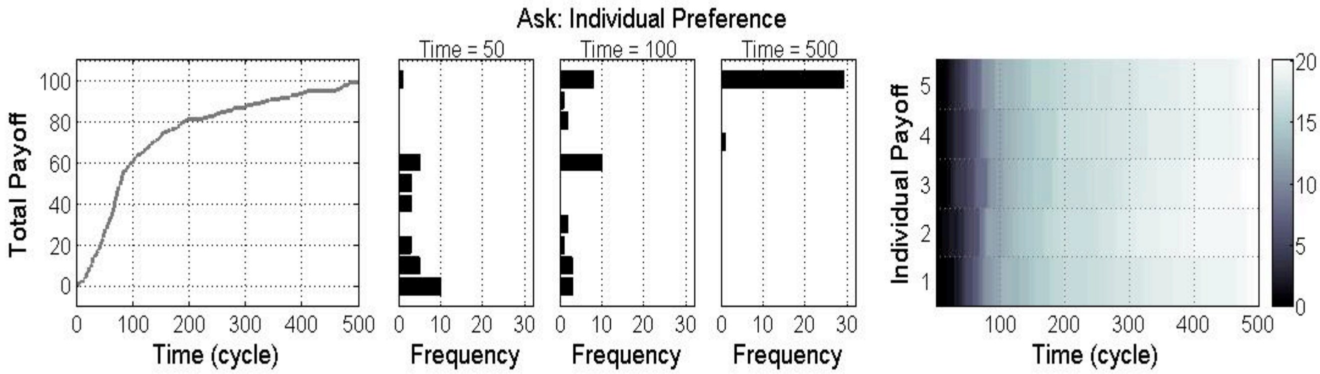

Figure 6: The results from the Individual Preference model in Simulation Study 1 are shown. The left most graph shows the evolution of total payoff (sum of all agents' payoffs) over the course of 500 trials, averaged across 30 simulations. The middle three histograms show the frequency of total payment for the 30 simulations at 50, 100, and 500 cycles. The color map on the right side shows the evolution of each agent's payoff averaged over 30 simulations.

The results from Simulation Study 1 show that for a single peak search space, asking random others can be especially beneficial when the time to search is limited. Every group member innovating can slow the team performance down. If there is a reasonable amount of time, the Majority Preference model and the Individual Preference model work fine. The success of the Random model suggests that we should sometimes observe different, random others, instead of always following the same individuals.

Simulation Study 2

In Simulation Study 2, the search space had three best solutions as shown in Figure 2. In this case, imitating can limit the number of good solutions the group discovers by causing all agents to conform to a single good solution. In contrast, the group can collectively find different solutions if group members innovate.

The procedure for Simulation Study 2 was the same as that for Simulation Study 1. The same four models were evaluated using a diversity metric and a normalized search speed for finding good solutions. The diversity metric was defined as the percentage of the group finding two or more best actions in 30 simulations. The normalized search speed, $\hat{v}_e$, is a relative metric defined by the time required to achieve 70% of the optimal result for a group, T_e. If a constant S quantifies the solution space, behavior model k has an observed average exploration speed, v_e:

$$v_e(k) = \frac{S}{T_e(k)}$$

Then the normalized search speed for model k, $\hat{v}_e(k)$, is:

$$\hat{v}_e(k) = \frac{v_e(k)}{\min_j v_e(j)} = \frac{S/T_e(k)}{\min_j S/T_e(j)} = \frac{\max_j T_e(j)}{T_e(k)}$$

	Number of Solutions 1	2	3	Diversity Metric
Innovate	0%	30%	70%	**100%**
Ask: Random	96.7%	3.33%	0%	**3.33%**
Ask: Majority Preference	70%	30%	0%	**30%**
Ask: Individual Preference	73.3%	26.7%	0%	**26.7%**

Table 1: The results from Simulation Study 2 are shown. The diversity metric shows the percentage of finding two or more best solutions in 30 simulations. The Innovate model was able to find two best solutions 9 times (30%) and three best solutions 21 times (70%), resulting in 100% diversity score. In contrast, the Random model resulted in finding only one good solution in 29 of 30 simulations. The performances of the Majority and Individual Preference models were in between those of the Innovate and Random models.

Table 1 displays the simulation results for the payoff distribution with three peaks. As predicted, the Innovate model was able to find multiple best solutions, resulting in a high diversity score. In contrast, the Random model resulted in the discovery of only one good solution in almost all 30 simulations (96.7%). The Majority Preference model and the Individual Preference Model were in between the Innovate model and the Random model. Although the two preference models could not find all three best solutions, they were able to find two best solutions in some cases, much more frequently than the Random model.

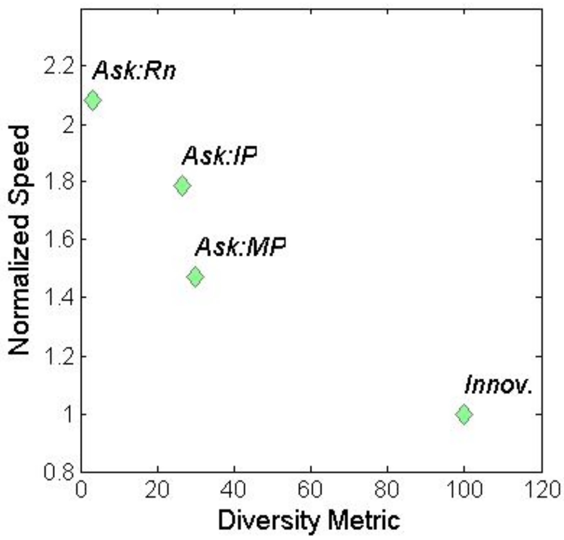

Figure 7: Each group's normalized search speed is shown as a function of its diversity metric. The normalized speed axis shows how quickly the group achieves a high total payoff (higher speed means faster). The diversity metric shows the percentage of finding two or more best actions in 30 simulations. The Innovate model (Innov) results in a high diversity measure but is slow to have all agents finding a good solution, indicated by low normalized speed. The Random model (Ask: Rn) leads to a high normalized speed, but this group converges to a single solution too quickly and thus results in low diversity of good solutions. The Majority Preference model (Ask: MP) and the Individual Preference Model (Ask: IP) were in between the Innovate model and the Random model.

Figure 7 shows each group's normalized search speed as a function of its diversity metric. The normalized speed axis shows how quickly the group achieves a high total payoff. This axis shows that we have essentially replicated Simulation Study 1 in terms of speed of finding a good solution as a group: Random, Individual Preference, Majority Preference, and Innovate, from fastest to slowest.

The diversity metric is the new measure relevant to the multiple best solutions in Simulation Study 2. All agents innovating results in a high diversity measure but, as Simulation Study 1 found, is slow to have all agents finding a good solution. The opposite of the Innovate model is the Random model. Asking random others leads to efficient dissemination of a solution and thus a high normalized speed, but makes the group converges to a single solution too quickly. Thus the Random model under explore the search space. The Majority Preference model and the Individual Preference model are quite efficient in disseminating a solution relative to the Innovate model. At the same time the two preference models have the time to explore the space. This is because always asking a particular individual has a higher chance of resulting in incidental innovation in the next round than asking a random other. When the asked agent does not have a good solution, the asking agent will be dissatisfied and innovate on the next round. When imitating a particular other, the imitating agent will likely keep asking the same agent. If this asked agent does not have a good solution, the asking agent will have many opportunities to innovate. When imitating a random other, the imitating agent will ask different agents at different cycles. There is less chance that the imitating agent always asks another agent with a poor solution in the Random ask model than in the two preference models. Thus the random imitation does not result in innovation as often as the other types of imitation.

Discussion

In the current study, we examined the benefits of different search strategies through computer simulation. We tested four models. In the Innovate model, each of the five agents in the group innovated on each cycle. The other three models incorporated some mechanisms of imitation. In the Random model, each agent imitated the best solution of a random other on each cycle. In the Majority Preference model, each agent imitated the best solution of the agent that was asked by many agents. This group followed the principle of preferential attachment, and conformed to the majority's behavior. In the Individual Preference model, each agent tracked how often it imitated the other agent, and imitated another agent based on how often it asked a certain agent. In this group, agents developed familiarity with a particular agent and followed this agent. We tested these four models in two kinds of search space: single best solution and three best solutions. In the current simulation, when imitating did not result in a better solution than the existing one, the agent innovated on the next time cycle and then resumed the imitation on the following time cycle.

The results from Simulation Study 1 showed that for a single peak search space, asking random others could be especially beneficial if the time to search was limited. In contrast, every group member innovating could take a long time for all the group members to find a good solution. The Majority Preference model and the Individual Preference model found good solutions in a reasonable amount of time.

In Simulation Study 2, the four models were tested under the three-peak environment. All agents innovating resulted in the group finding multiple good solutions, but, as Simulation Study 1 found, was slow to have all agents finding a good solution. In contrast, asking random others led to efficient dissemination of a solution, but the group converged to a single solution too quickly, and thus the Random model under explored the search space. Majority Preference model and the Individual Preference model had the time to explore the space and were still quite efficient in disseminating a solution relative to the Innovate model.

Taken together, these results suggest that following a particular other, whether the most popular one or the most familiar one, results in a good compromise between speed and diversity in finding good solutions. It is interesting that these models that incorporate characteristics found in humans are most robust in the sense that they work well in different environments, although they may not be optimal in a single environment. Perhaps people's desire to follow particular others is a key to adaptive behavior, allowing people to disseminate ideas efficiently while still encouraging the innovation of new ideas.

Future work should explore more complex models, in which the group can have a mix of innovators and imitators. Individual differences can be useful when the group tends to converge too quickly. When group members converge quickly to an optimal solution, responding to a new situation becomes a problem (Resnick, 1994). For example, if all team members responded to an immediate threat in area X (which happens in the real world), it may take a while for everyone to respond to a new alert in area Y. Analogously, a group may fail to respond to a new and better solution when the group converges to a good solution too quickly. A simple way to avoid such failure to adapt to better solution is to include individuals with different abilities in a team (Sakamoto & Nickerson, 2007). By making some individuals innovate more often than others, we can encourage some learners to focus on disseminating solutions, and others to explore the space for new situations. These models incorporating individual difference can be robust to changing environments, such as when the payoff distribution shifts from time to time. Future work should include these variables, such as changing environment and individual difference, to make the simulation world closer to the world we live in. Future work should also compare these models against people.

In conclusion, the current simulation studies showed that people's natural tendency to follow particular others may have survived for a good reason: It leads to reasonable performances in a reasonable amount of time in different environments. If the dimension to optimize is well defined, one may tailor the search strategy. For example, if the time is not an issue, a group of agents that all innovate can find a diverse set of good solutions as a group. If there is a need to disseminate information widely and quickly, then asking random others will be the way to search the space. If one does not know what to optimize, following the particular others will result in a reasonable performance.

References

Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs*, *70* (Whole no. 416).

Bandura, A. (1965). Behavioral modification through modeling procedures. In L. Krasner & L. P. Ulmann (Eds.), *Research in behavior modification: New development and implications* (pp. 310–340). New York: Rinehart and Winston.

Barabási, A. L., & Albert, R. (1999, October 15). Emergence of scaling in random networks. *Science*, *286*, 509–512.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*, 591–621.

Dennett, D. C. (1995). *Darwin's dangerous idea*. New York: Touchstone.

Deutsch, M., & Gerard, H. B. (1995). A study of normative and informational social influences upon individual judgment. Journal of Abnormal Social Psychology, 51, 629–636.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, *7*, 117–140.

Gigerenzer, G, Todd, P. M., and the ABC Research Group (2000). Simple Heuristics that Make Us Smart. New York: Oxford.

Gureckis, T. M., & Goldstone, R. L. (2006). Thinking in groups. In S. Harnad & I. Dror (Eds.), *Distributed cognition: Special issue of pragmatics & cognition*, *14* (pp. 293–311). Amsterdam, The Netherlands: John Benjamins.

Kraatz, M. S. (1998). Learning by association? Interorganizational networks and adaptation to environmental change. *Academy of Management Journal*, *41*, 621–643.

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. Journal of Experimental Psychology: General, 137, 422-433.

Page, S. E. (2007). The Difference. Princeton University Press.

Resnick, M. (1994). Turtles, termites and traffic jams: Exploration in massively parallel microworlds. Cambridge, MA: MIT Press.

Sadlon, E., Sakamoto, Y., Dever, H. J., Nickerson, J. V. (2008). The Karma of Digg: Reciprocity in Online Social Networks. In *Proceedings of the 18th Annual Workshop on Information Technologies and Systems*.

Sakamoto, Y., & Nickerson, J. V. (2007). Social behavior in a team of autonomous sensors. In *Proceedings of the Intelligence and Security Informatics Conference*.

Sakamoto, Y., Sadlon, E., & Nickerson, J. V. (2008). Bellwethers and the emergence of trends in online communities. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*.

Sakamoto, Y., Ma, J., & Nickerson, J. V. (2009). 2377 people like this article: The influence of others' decisions on yours. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society*.

Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, *27*, 1–60.

Smith, K., Kirby, S., and Brighton, H. (2003). Iterated learning: A framework for the emergence of language. Artificial Life, 9, 371–386.