69 Montrey, M., & Shultz, T. R. (2010). Evolution of social learning strategies. Proceedings of the Ninth IEEE International Conference on Development and Learning (pp. 95-100). Ann Arbor, MI: IEEE.

Evolution of Social Learning Strategies

Marcel R. Montrey and Thomas R. Shultz

***Abstract*—We study three types of learning with Bayesian agent-based modeling. First, we show that previous results obtained from learning chains can be generalized to a more realistic lattice world involving multiple social interactions. Learning based on the passing of posterior probabilities converges to the truth more quickly and reliably than does learning based on imitation and sampling from the environment; and the latter method gets closer to the truth than does pure imitation. The passing of posterior probability distributions can be viewed as teaching by explanation, and as an implementation of the cultural ratchet, which allows rapid progress without backsliding. We also show that evolution selects these learning strategies in proportion to their success. However, if the environment changes very rapidly, evolution favors the imitation-plus-reinforcement strategy over the more sophisticated posterior passing. Implications for developmental robotics, human uniqueness, and interactions between learning and evolution are discussed.**

***Index Terms*—Agent-based modeling, Bayesian learning, cultural ratchet, evolution.**

I. Introduction

The division between humans and other animals has often been framed in terms of communication. What separates humans from animal social groups may be cultural transmission, and the fidelity with which it allows information to be transmitted between individuals [1]. Sophisticated human communication, seemingly unavailable to other species, creates a ratchet effect, whereby the perpetuation of existing knowledge prevents backsliding and new discoveries enable further refinement. This allows individuals to benefit not only from their own experience, but also from the cumulative knowledge of countless peers and ancestors. In short, the wheel or the light bulb need only be invented once, and most revisions are likely to be improvements. The effectiveness of such learning strategies has implications for other domains as well. For instance, robots that are designed to learn from humans or each other may also benefit from an emphasis on theory learning and transmission, as opposed to the sharing of low-level data.

Manuscript received June 30, 2010. This work was supported in part by grants from the Natural Sciences and Engineering Research Council of Canada and the Dean's Excellence Fund of the Faculty of Science, McGill University to TRS.

Marcel R. Montrey is in the Cognitive Science Program at McGill University, Montreal, QC Canada (e-mail: marcel.montrey@mail.mcgill.ca).

Thomas R. Shultz is with the Department of Psychology and School of Computer Science, McGill University, Montreal, QC Canada (514-398-6139; e-mail: thomas.shultz@mcgill.ca).

Here we present computer simulations that extend previous research on learning strategies in simple linear chains to a more complex two-dimensional lattice structure. We find that more sophisticated strategies lead to more effective learning, and are selected by evolution in most cases, although a simpler strategy is favored if the environment changes too rapidly.

A. Cultural Transmission and Other Social Learning

Human cultural transmission stands apart from other social learning methods by being closely tied to the concept of a theory of mind [1]. In other kinds of social learning, an animal acquires information that is in some way influenced by the social environment. This can be as simple as a child following its parent to a food source. Cultural transmission differs in that it involves some attempt at grasping or adopting another's perspective. To amend the previous example, the parent might draw the child's attention to the fact that the food source is in close proximity to water. In this latter case, it is not simply the idea that *food is here* that is conveyed; it is also the understanding that *such food can be found near water, and there are good causal reasons for that association*. In other words, cultural transmission involves not just the communication of environmental data, but also of a theory which can be used to predict it [2]. The ratchet effect is therefore an exercise in theory-building, which relies on a deeper understanding of the problem in question to make gradual and beneficial changes to that understanding.

Fortunately, these ideas have proven amenable to both mathematical and empirical analysis. In particular, agent-based modeling and psychological experimentation support the notion that cultural transmission is unique in enabling the ratchet effect [3]. Like humans, the Bayesian agents employed in these simulations are rational in that they attempt to maximize their performance on the task that they are assigned.

B. A Previous Agent-Based Model

Bayes' rule is a mathematical formulation intended to capture how probabilities change with new data. Each agent adopts what can be thought of as a theory of the world. In actuality, this is a set of probabilities $p(h)$ which describe the extent to which a series of hypotheses $h \in H$ are thought to be true. As new data d are acquired, the probability of a given hypothesis may change. For instance, consider an agent whose task is to predict whether or not it will rain tomorrow. The hypotheses h_1 and h_2 may be characterized as *it will rain tomorrow* and *it will not rain tomorrow*, respectively. The agent may begin by assuming that rain is highly unlikely: the $p(h_1)$ value it assigns is low, and $p(h_2)$ is high. But what happens if it does end up raining the next day? On the basis of this one day of experience, the assumption that rain is highly unlikely becomes implausible. The agent thus updates its $p(h_1)$ to a higher value, its $p(h_2)$ to a lower value, and it becomes more likely to predict rain in the future.

Formally, Bayes' rule is written as

$$p(h \mid d) = \frac{p(d \mid h)p(h)}{\sum_{h' \in H} p(d \mid h')p(h')} \qquad (1)$$

where $p(h|d)$ is the probability of a hypothesis h given data d. This is referred to as the posterior probability, and it varies directly with the likelihood of the data given the hypothesis $p(d|h)$ and the prior probability of that hypothesis $p(h)$. Intuitively, the more the data and the hypothesis disagree, the less credible the hypothesis becomes. The denominator is simply a normalizing term, which ensures that the sum of the posterior probabilities remains constant.

In this earlier work, a chain of Bayesian agents was tasked with solving a function approximation problem [3]. An agent would begin with some set of $p(h)$ values that comprise a prior theory about the function, and then interact with its predecessor using Bayes' rule. This process would then be repeated, with the successor acting as the predecessor for the next agent in the chain.

These chains of agents were divided into three conditions, each of which used a distinct learning strategy. The chain process was initialized by allowing the first agent to sample data directly from the environment, which produced data in accordance with the true hypothesis h^*. All agents also shared a prior distribution – a default $p(h)$ value for each hypothesis h. This can be thought of as a set of innate dispositions or prior assumptions about the problem, and is common to all agents.

C. Imitation

The *pure iterated learning* model was used to represent the simplest form of social learning: *imitation*. Here, the only data available to the i^{th} agent is the behavior of its predecessor. In terms of Bayes' rule, the probability that an agent assigns to a hypothesis is updated during an interaction using

$$p(h_i \mid d_i) \propto p(d_i \mid h_i)p(h_i) \qquad (2)$$

where $p(d_i|h_i)$ is the likelihood of the data produced by the predecessor given the hypothesis, and $p(h_i)$ is the prior probability of that hypothesis. Such an interaction is akin to a child observing where its parent searches for food. Because no new environmental data is ever collected, the probability assigned to any hypothesis $p(h|d)$ asymptotically approaches the prior $p(h)$.

While each agent learns something by observing its predecessor, the chain of agents ultimately falls back on its innate assumptions to fill in the blanks, since those provide the most consistent information. Given enough repetitions of this process, each agent's behavior becomes indistinguishable from that of an agent following its intuitions without reference to any external data. As in the popular *telephone game*, whatever knowledge was acquired from the environment by the first agent in the chain is rapidly lost.

D. Imitation-plus-Reinforcement

A slight improvement to the previous learning method is the *mixed data* strategy, here called *imitation-plus-reinforcement*. In this case, every interaction with an agent's predecessor is augmented by data from the environment. An intuitive example might be a child not only observing where its parent searches for food, but also finding some food on its own. In Bayesian terms, an interaction updates the probability an agent ascribes to a hypothesis using

$$p(h_i \mid d_i, d_i^*) \propto p(d_i \mid h_i)p(d_i^* \mid h_i)p(h_i) \qquad (3)$$

where the new term $p(d_i^*|h_i)$ is the likelihood of the data produced by the environment given the hypothesis. Although each agent learns something from the environment as well as from its predecessor, this learning still builds upon the agent's innate assumptions. More formally, while new data d and d^* are obtained during each interaction, $p(h)$ remains constant, making this agent chain incapable of cumulative cultural evolution. Instead, the prior's strong influence ensures that the average probability assigned to each hypothesis $p(h|d)$ approaches $\sum_{d^*} p(h|d^*)\, p(d^*|h^*)$, a point between the prior and the correct value.

In short, the behavior of an imitation-plus-reinforcement agent is a combination of its predecessor's behavior, its own sampling of the environment, and its own predispositions. Over time, the latter two factors become dominant, and each agent's behavior asymptotically approaches that of a naïve Bayesian agent which has sampled the environment once. The entire process leads to agents which are about as knowledgeable as the first agent in the chain.

E. Posterior Passing

The final learning strategy employed was called *posterior passing*. This was intended to represent cultural transmission. In addition to information from the predecessor and the environment, the agent also adopts its predecessor's posterior distribution as its prior. In effect, an agent's innate assumptions about the problem are supplanted by the previous agent's theory of the world. The probability of a hypothesis is thus determined by

$$p(h_i \mid d_i, d_i^*) \propto p(d_i \mid h_i)p(d_i^* \mid h_i)p(h_{i-1} \mid d_{i-1}) \qquad (4)$$

where $p(h_{i-1}|d_{i-1})$ is the preceding agent's posterior probability. Intuitively, an agent no longer learns from scratch, but rather from where its predecessor left off.

This is akin to possessing a theory of mind, and might again be exemplified by a parent leading a child to a food source and then calling attention to the fact that water is nearby. Not only does the child learn where the food is, but also that, in the parent's estimation, *water predicts food*. Thus, a sort of food-search theory is transmitted from parent to child. This faithful preservation and gradual improvement of knowledge yields a ratchet effect, with the probability assigned to each incorrect hypothesis asymptotically approaching 0, and the probability assigned to the true hypothesis asymptotically approaching 1.

A posterior-passing agent's learning is a combination of its predecessor's behavior and understanding of the problem, and

 Montrey, M., & Shultz, T. R. (2010). Evolution of social learning strategies. Proceedings of the Ninth IEEE International Conference on Development and Learning (pp. 95-100). Ann Arbor, MI: IEEE.

its own sampling of the environment. Given enough time, each agent's knowledge becomes entirely consistent with the true state of the world. In other words, the agent chain solves the problem. Such a cultural ratchet effect relies on a rather sophisticated and powerful style of communication.

F. *Some New Questions*

Again, these chain phenomena were demonstrated in both simulations and psychological experimentation [3]. The results opened the door to additional, closely-related questions. For instance, although an agent chain is a useful abstraction, such a model of social interaction is also quite unrealistic. Because social learning is rarely as simple as a series of one-to-one interactions comprising a linear chain, it is reasonable to ask whether these results would generalize to a multidimensional lattice where agents are able to interact with multiple predecessors. We were also interested in how agents adopting different learning strategies would behave if they were intermixed and then exposed to evolutionary forces. If an agent's ability to solve the assigned problem determines reproductive fitness, then which learning strategy would be favored, and to what extent? Moreover, how would a rapidly or slowly changing environment upset this balance? In short, using Bayesian agent-based modeling, we sought to analyze these three learning strategies in an evolutionary context.

II. METHODS

A. *The Problem*

In our simulations, Bayesian agents are faced with the problem of choosing a location to search for food. There are four locations L_1-L_4, only one of which can be searched. At each time step, three incorrect locations contain a modest amount of food, while one correct location L^* contains an abundance of food. An agent's theory of the world is the set of probabilities it assigns to four hypotheses h_1-h_4, each of which corresponds to one of the locations. Any given hypothesis h_i may be characterized as stating that *location L_i probably has the most food*. More formally, a hypothesis h_i assigns a probability of 0.55 to $L_i = L^*$ and a probability of 0.15 to $L_j = L^*$, for each $j \neq i$. At any given time step, L^* is determined by sampling one of the four hypotheses. The sampled hypothesis is called the true hypothesis h^*. Any agent assigning a probability $p(h^*)$ of 1 to the correct hypothesis, and a probability $p(h)$ of 0 to every other hypothesis can be said to have solved the problem, and to possess a perfect theory. In our simulations, h_1 was always the correct hypothesis. The sole exception was in cases where the true hypothesis was dynamic, and thus subject to regular change.

B. *The Method*

An agent decides which location to search for food via a two-stage process. First, it selects a hypothesis. The probability of a hypothesis being adopted is directly proportional to how probable the agent believes that hypothesis to be. Next, the selected hypothesis is sampled to determine which location will actually be searched. For instance, if an agent assigns $p(h_1)$ a value of 0.9, then there is a 0.9 probability that h_1 will be selected as that agent's hypothesis. If h_1 is in fact selected, then the agent has a 0.55 chance of searching L_1 and a 0.15 chance of searching each of the other three locations.

Upon its creation, each agent is assigned a set of prior probabilities, with $p(h_1) = p(h_2) = p(h_3) = p(h_4) = 0.25$. That is to say, agents are innately predisposed to find each hypothesis to be equally probable.

C. *Learning*

Each simulation begins with all agents sampling the environment using Bayes' rule. At every subsequent time step, each agent is replaced by a successor, which interacts with its predecessors prior to being placed. Every such interaction is multidirectional and involves both the agent's immediate predecessor (parent) and all four of that predecessor's von Neumann neighbors. Each agent interacts only once, though this interaction yields an average of the five posteriors that would have resulted from interacting with each predecessor individually. This process permits an agent to integrate information from multiple sources without allowing the order of interactions to marginalize any of those influences. In effect, every agent's knowledge is molded by a small and localized community, with each of the five predecessors contributing equally. The precise nature of an interaction is determined by the agent's learning strategy. Imitators acquire only behavioral data, whereas imitation-plus-reinforcement agents add environmental data, and posterior-passing agents also take their predecessor's theory as their prior.

D. *The World*

All simulations are conducted using a 48 by 48 torus filled with 2304 agents. A torus is equivalent to a lattice where each edge is connected to the corresponding edge on the opposite side. Represented three-dimensionally, this yields a donut-shaped world. This configuration ensures that the world is continuous in all directions, giving each agent an equal number of neighbors – a common simplifying assumption in space-based simulations [4]. Each model is populated with agents of the appropriate type, then iterated over 1000 time steps, with each step representing a new generation of agents. All reproduction is asexual.

E. *Evolution and Learning*

In our later models, biological evolution is simulated by making an agent's reproductive success dependent on its ability to search the correct location. First, each agent is assigned a fitness value between 0 and 1. Then, for each position in the lattice, both the occupant and its von Neumann neighbors are considered. Of these five agents, one is randomly selected to create an offspring which will fill the current position at the next time step. The probability of being chosen is proportional to the agent's fitness, meaning that if one neighbor has a fitness of 1.0 and another has a fitness of

Montrey, M., & Shultz, T. R. (2010). Evolution of social learning strategies. Proceedings of the Ninth IEEE International Conference on Development and Learning (pp. 95-100). Ann Arbor, MI: IEEE.

0.5, then the former is twice as likely to produce an offspring in that location. As a result, each agent may potentially produce anywhere from 0 to 5 children.

An agent's fitness is determined by two factors. The first is the base fitness rate. This is the fitness value that each agent is created with, and is set to 0.5 in all cases. The other is whether or not the agent in question searches the correct location. If it does, it is rewarded by an increase in fitness of 0.5. These factors reflect the fact that even incorrect locations will yield a modicum of food and thus increase the potential for reproduction. Finding an overabundance by searching the correct location is highly advantageous however, and will double the agent's reproductive odds.

Note that a child will typically retain its parent's learning strategy. Each strategy can be thought of as corresponding to a particular allele of the same gene. As such, there exists the possibility of mutation, whereby a child adopts a learning strategy that differs from that of its parent. In our evolutionary models, a mutation rate of 0.005 was fixed.

III. RESULTS

A. Moving to a Lattice

Our first three simulations involved populating lattices with agents of a single type. We produced an imitation model, an imitation-plus-reinforcement model, and a posterior-passing model. Results in each case are averaged over 5000 independent runs.

In the imitation case, agents converged asymptotically on a theory of the world that matched their prior expectations. Specifically, each hypothesis' posterior probability $p(h|d)$ rapidly approached its prior probability $p(h)$. This can be seen in Fig. 1, which shows the mean posterior probabilities as a function of time.

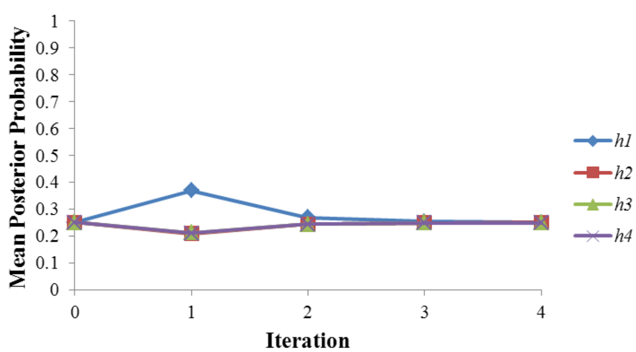

Figure 1. Mean posterior probabilities $p(h|d)$ with imitation learning. These agents' theory of the world quickly converged to their prior expectations. The change in probabilities from time step 0 to time step 1 represents the initialization procedure, where agents sample the environment directly.

In the case of the imitation-plus-reinforcement agents, all posterior probabilities $p(h|d)$ approached $\sum_{d^*} p(h|d^*)\, p(d^*|h^*)$, a set of values which lie partway between the agents' prior expectations and the truth. In other words, each agent's behavior remained that of a naïve Bayesian agent sampling the environment for the first time, as shown in Fig. 2.

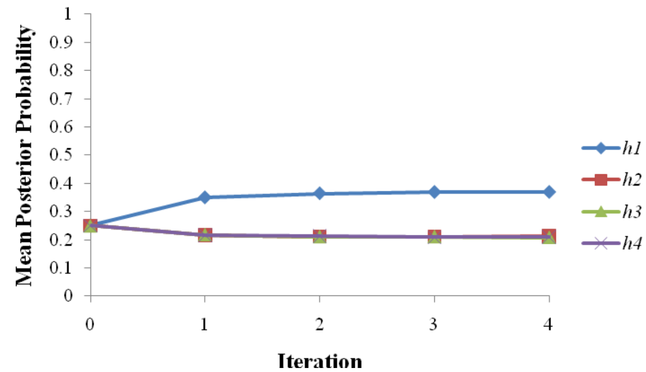

Figure 2. Mean posterior probabilities $p(h|d)$ with imitation-plus-reinforcement. A theory partway between the agents' prior expectations and the truth was maintained following the initial sampling of the environment.

Finally, in the posterior passing case, agents gradually adopted the true hypothesis h^* to the exclusion of all others. More formally, $p(h^*|d)$ approached 1, while $p(h|d)$ for all other hypotheses trended towards 0. This is illustrated in Fig. 3, where the mean posterior probability is plotted for the first 20 cycles.

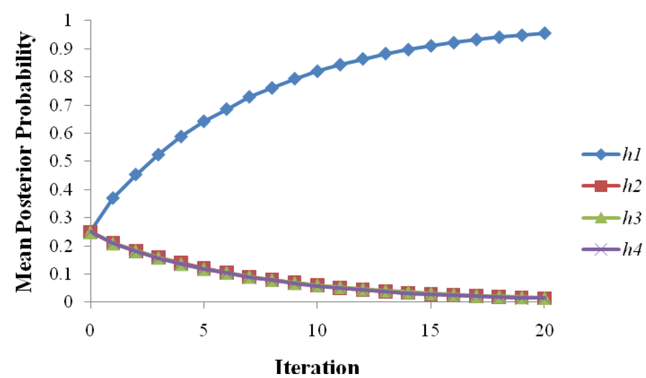

Figure 3. Mean posterior probabilities $p(h|d)$ with posterior passing. These probabilities asymptotically approached their ideal values, signaling a correct solution.

B. Evolution and Learning Interactions

Rather than studying each learning strategy independently, our next model intermixed agents of all three types. Such heterogeneity in and of itself significantly influences an agent's effectiveness. With access to the behavior of more knowledgeable agents, pure imitation becomes a considerably more successful strategy, allowing these agents to make genuine, though limited progress in solving the problem. Imitation-plus-reinforcement agents obtain similar benefits from interacting with their posterior-passing counterparts, whereas the posterior-passing agents' efficacy suffers slightly, due to their exposure to peers whose theories are less accurate than their own.

The original proportions of the three agent types were equal, and the initial spatial distribution was randomly determined. Finally, natural selection and evolution were introduced.

Montrey, M., & Shultz, T. R. (2010). Evolution of social learning strategies. Proceedings of the Ninth IEEE International Conference on Development and Learning (pp. 95-100). Ann Arbor, MI: IEEE.

Unlike in our previous models, our interest here is the frequency of each learning strategy. By observing the changes in their relative proportions over time, we can see which strategy might be favored by evolution, and to what extent. To this end, the mean proportions of each agent type over 1000 evolutionary cycles are plotted in Fig. 4. These results are averaged over 250 independent runs.

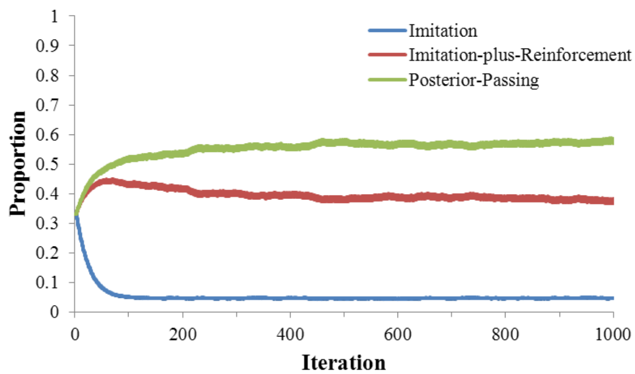

Figure 4. Mean proportions of each agent type over 1000 cycles. Line thickness represents standard error. Posterior-passing agents quickly occupy the majority of the world.

In order to probe the effects of a changing environment on evolutionary trends, our final model introduced a new parameter. The rate of change of the environment reflects the number of iterations that must pass before a new correct hypothesis h^* is selected. We found that the shorter the period between such changes, the more successful imitation-plus-reinforcement agents became relative to posterior-passing ones. This trend is illustrated in Fig. 5.

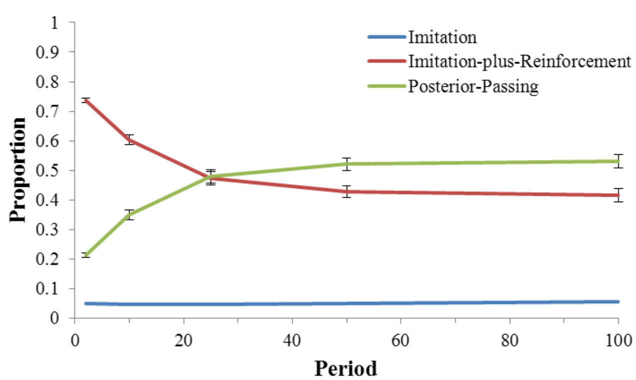

Figure 5. Mean proportions of agents of each type, with standard error, after 1000 cycles as a function of rate of environmental change. The faster the true hypothesis h^* changes, the more successful imitation-plus-reinforcement agents become relative to posterior-passing ones.

IV. DISCUSSION

A. *Summary of Results*

By employing Bayesian agent-based modeling, we showed that previous results with agent chain simulations of social learning generalize to a more complex and realistic two-dimensional lattice. By intermixing agents of various types, we also obtained data suggesting that cultural transmission, as opposed to other, simpler social learning, is favored by evolution. This conclusion comes with a caveat however, since we also found that the more straightforward imitate-and-sample strategy is superior under rapid environmental change.

B. *The Weakness of Pure Imitation*

Across our simulations, pure imitation proved to be an ineffective learning strategy. Merely copying the behavior of others is not enough to compete against peers with direct access to environmental data. The strength and consistency of the prior distribution ensures that these agents cannot converge on the correct hypothesis, even when interacting with more accurate agents. Simply put, imitation alone is too limited to allow mastery of this particular problem.

C. *Better Learning Strategies*

Of greater interest are the imitation-plus-reinforcement and posterior-passing strategies. Because their learning mechanism is so powerful, one might expect posterior-passing agents to easily dominate. However, our results suggest that this is not always the case. In most instances, the imitation-plus-reinforcement agents' ability to duplicate the behavior of their posterior-passing neighbors enables improved performance. Conversely, posterior-passing agents are somewhat impeded by their interactions with less knowledgeable agents, which further narrows the gap between these two strategies.

Moreover, for very dynamic problems, a stronger learning mechanism can be a liability. If rapid environmental change modifies the rules according to which food is distributed, then the high fidelity with which a now-incorrect theory is transmitted will often leave an agent in a worse position than if it had no outside information at all. In such cases, the imitation-plus-reinforcement strategy is superior. Because their learning is shallower, these agents react more quickly to new environmental conditions.

While the ratchet effect may lead to superior learning that even bystanders can benefit from, posterior passing is not a universally optimal strategy. In highly dynamic environments, adaptability offers an advantage. In such cases, complex problems are better approached with a more superficial understanding, which allows for both adequate behavior and faster recovery when that understanding fails.

D. *Should Developmental Robotics Consider Evolution?*

Developmental robotics is typically concerned with how a robot's control system changes through learning and development. Such concerns are usually seen as distinct from the sister subfield of evolutionary robotics, which employs populations of robots that evolve over time according to artificial natural selection and even genetic crossover [5].

However, it is worth noting that biologists view development as the key link between genotype and phenotype and one of the most important challenges in contemporary biology [6]. If developmental robotics wants to know which abilities should be built into robots as a basis for future

Montrey, M., & Shultz, T. R. (2010). Evolution of social learning strategies. Proceedings of the Ninth IEEE International Conference on Development and Learning (pp. 95-100). Ann Arbor, MI: IEEE.

learning and development, our results suggest that the role of evolution should not be ignored. In nature, learning strategies are complex, and do not arise in a vacuum. They probably evolve, and in turn affect the course of evolution by changing the environment in which natural selection occurs. For instance, as a social environment alters, previously mediocre learning strategies may become more effective, allowing exploitation of their other advantages. Paradoxically, intermixing stronger learners with weaker ones may thus lead to more resilient groups, despite the decline in individual ability. Though dissimilar from the machine learning concept of boosting [7] in that group performance is also diminished, such phenomena illustrate the surprising benefits of introducing theory sharing into a more biologically plausible heterogeneous population. Whether adopted exclusively, or within a small segment of a larger group, cultural transmission can be a powerful tool.

One might also imagine that developmental robotics could potentially offer insights into how development links genotype to phenotype. It is often more effective to study evolution and learning in simulation experiments than in nature, from sparse evolutionary records. Like many aspects of robotics, such approaches will likely be more feasible in software than in physical robots, but this may well change in the future.

E. What Was the Human Spark?

Much recent effort has been directed to the question of what makes humans unique [8]. What was the evolutionary spark that set humans so far apart from other species? Along with others [1, 2, 8], we suggest that the ability to transmit information in a ratchet-like fashion is a reasonable candidate for such a spark. Until very recently, the cultural ratchet effect has been outside the realm of rigorous scientific study. Now, rather suddenly, we have a formal scientific treatment of the cultural ratchet [3], in the form of passing, from agent to agent, posterior probability distributions for hypotheses and theories that explain how phenomena in the world seem to work.

Our research builds on this insight by studying cultural transmission in the context of spatial realism and natural selection. Here, evolution selects a social learning strategy based on posterior passing, except when the environment changes so rapidly that it is more efficient for many agents to get by with copying and sampling. Our results also offer insight into another aspect of posterior passing: its ability to indirectly benefit bystanders.

Although other mammals conform to Bayes' rule in conditioning experiments [9], it seems likely that the effective communication of theories would require a powerful and flexible language. So, there may well turn out to be a cluster of human abilities, involving social inclinations and skills, theory of mind, and language, that enable the cultural ratchet.

Much more remains to be discovered about these phenomena. How robust are these effects and what parameter variations govern their appearance? Can the computation and passing of posterior probabilities be implemented in a more realistic way, perhaps with neural networks? To what extent can learning and development alter the course of evolution?

F. Is Cultural Transmission Rational?

Because Bayes' rule can be mathematically derived, it is typically considered to be an optimal and rational solution to induction, inference, and learning [9]. However, concern has been raised about whether cognition is always so rational [10]. Now that our current results suggest that evolution generally favors Bayesian learning, a number of puzzling new questions emerge about rationality.

Biological evolution itself is not particularly optimal in its solutions. It often produces pretty good contraptions, rather than excellent designs [11], and serious tradeoffs that, while enabling organisms to reproduce, eventually kill them [6]. Curious then that our simulations show that evolution favors a mathematically provable inference-and-learning device like Bayes' rule. Perhaps significantly, such reasoning is not universally adopted by our agent population.

Like other evolutionary simulations [12], ours does not actually *create* Bayes' rule from scratch – rather it selects a designed Bayesian allele over other, less sophisticated, also-designed alleles. Bayesian inference and learning in psychology simulations use this same select-from-available-designs strategy [10]. More interesting simulations may eventually show us how evolution creates solutions out of less, and what role those solutions play within communities, as well as at the individual level.

REFERENCES

[1] M. Tomasello, A. Kruger, and H. Ratner, "Cultural learning," *Behavioral and Brain Sciences,* vol. 16, pp. 495-552, 1993.

[2] R. Boyd and P. J. Richerson, "Rationality, imitation, and tradition," in *Nonlinear dynamics and evolutionary economics*, R. H. Day and P. Chen, Eds. Oxford: Oxford University Press, 1993.

[3] A. Beppu and T. L. Griffiths, "Iterated learning and the cultural ratchet.," in *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, N. A. Taatgen and H. v. Rijn, Eds. Austin, TX: Cognitive Science Society, 2009, pp. 2089-2094.

[4] M. A. Nowak, *Evolutionary dynamics*. Cambridge, MA: Harvard University Press, 2006.

[5] S. Nolfi and D. Floreano, *Evolutionary robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.

[6] S. C. Stearns and R. F. Hoekstra, *Evolution: an introduction*. Oxford: Oxford University Press, 2005.

[7] R. E. Schapire, "The boosting approach to machine learning: An overview," in *MSRI Workshop on Nonlinear Estimation and Classification*, 2001.

[8] M. Tomasello, "How are humans unique?," in *New York Times Magazine*, 2008.

[9] J. B. Tenenbaum, T. L. Griffiths, and C. Kemp, "Theory-based Bayesian models of inductive learning and reasoning," *Trends in Cognitive Sciences,* vol. 10, pp. 309-318, 2006.

[10] T. R. Shultz, "The Bayesian revolution approaches psychological development," *Developmental Science,* vol. 10, pp. 357-364, 2007.

[11] S. J. Gould, *The panda's thumb*. New York: Norton, 1980.

[12] T. R. Shultz, M. Hartshorn, and A. Kaznatcheev, "Why is ethnocentrism more common than humanitarianism?," in *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, N. A. Taatgen and H. v. Rijn, Eds. Austin, TX: Cognitive Science Society, 2009, pp. 2100-2105.