

Testing the Variability Selection Hypothesis: The Adoption of Social Learning in Increasingly Variable Environments

James M. Borg¹ and Alastair Channon¹

¹Research Institute for the Environment, Physical Sciences and Applied Mathematics, Keele University, ST5 5BG, UK
{j.borg,a.d.channon}@epsam.keele.ac.uk

Abstract

The variability selection hypothesis predicts the adoption of *versatile* behaviors and survival strategies, in response to increasingly variable environments. In hominin evolution the most apparent adaptation for versatility is the adoption of social learning. The hypothesis that social learning will be adopted over other learning strategies, such as individual learning, when individuals are faced with increasingly variable environments is tested here using a genetic algorithm with steady state selection and constant population size. Individuals, constituted of binary string genotypes and phenotypes, are evaluated on their ability to match a target binary string, nominally known as the environment, with success being measured by the Hamming distance between the phenotype and environment. The state of any given locus in the environment is determined by a sine wave, the frequency of which increases as the simulation progresses thus providing increasing environmental variability. Populations exhibiting combinations of genetic evolution, individual learning and social learning are tested, with the learning rates of both individual and social learning allowed to evolve. We show that increasingly variable environments are sufficient but not necessary to provide an evolutionary advantage to those populations exhibiting the extra-genetic learning strategies, with social learning being favored over individual learning when populations are allowed to explore both strategies simultaneously. We also introduce a more biologically realistic model that allows for population collapse, and show that here the prior adoption of individual learning is a prerequisite for the successful adoption of social learning in increasingly variable environments.

Introduction

It is now widely accepted that the species *Homo sapiens*, to which all modern humans belong, evolved in Africa before leaving to populate the rest of world (Tattersall, 2009). In order to successfully populate new and challenging environments hominins must have developed versatile and robust behaviors and survival strategies, with the most apparent hominin adaptation for versatility being the adoption of extra-genetic learning strategies such as social learning (Tomasello, 1999). This leads us to ask what was it about the environments in which hominins evolved that enabled

them to adapt to be so versatile and ultimately so successful when moving into new and unfamiliar environments. In response to this question numerous authors have suggested a variety of theories and hypotheses regarding the relationship between hominin evolution and the environment (Potts, 1998a). In this work we do not seek to answer the question of how hominins became such expert social learners, we instead test one of the most prominent theories of hominin evolution and versatility, the Variability Selection Hypothesis (Potts, 1996, 1998a,b), using an artificial life simulation.

The Variability Selection Hypothesis

The variability selection hypothesis, as proposed by Richard Potts (Potts, 1996, 1998a,b), predicts the adoption of versatile behaviors and survival strategies, in response to increasingly variable environments. Over the past seven million years there have been a number of what Potts describes as “large disparities” in environmental conditions and a trend toward increasing climatic variation in and around known early hominin locations in eastern and southern Africa, such as the Turkana and Olduvai basins (Potts, 1998a). Evidence for such inter- and intra-generational changes have been found in a variety of climatic indicators including marine oxygen isotope levels (Potts, 1998a,b), providing insight into temperature changes, and ocean dust records (Potts, 1998a), providing evidence for dust plumes arising from strong seasonal rainfalls and prevailing wind patterns. Both of these indicators demonstrate an upward trend in environmental variability during the last seven million years in Africa, and around the world in general. Evidence from these, and other climatic indicators, shows that major shifts in the African climate correlate well with important early technological milestones and speciation events in hominin evolutionary history (Grove, 2011). Key hominin and hominid adaptations such as early bipedality and complex social behavior emerged during these periods of more pronounced environmental variability (Potts, 1998b). Though the climatic evidence for the variability selection hypothesis is impressive, the hypothesis has had very little theoretical work applied to it. Following the call from Potts (1998b) for

a mathematical framework to explore the variability selection hypothesis, and the work of Grove (2011) to that end, we here test the claim that increasing environmental variability is a sufficient selection pressure to elicit the adoption of social learning, in an artificial life simulation.

Social Learning

Social learning is not restricted to humans and their ancestors: it is a widely observed natural phenomenon, with many species using a variety of social learning mechanisms such as imitation, emulation, teaching and the use of public information to produce adaptive behaviors in dynamic and challenging environments (Laland, 2004; Reader and Biro, 2010; Whiten and van Schaik, 2007). It has been suggested that social learning enables animals to better track their environment by assimilating extra-genetic information from others during their lifetimes while avoiding potentially costly individual learning (Boyd and Richerson, 1995).

The effects and benefits of learning have been studied widely in simulation. According to Nolfi and Floreano (1999) learning may be seen as having several adaptive functions within an evolutionary perspective. These include allowing individuals to adapt to environmental change, enabling evolution to use information extracted from the environment, and guiding evolution. Famously Hinton and Nowlan (1987) demonstrated that by using individual learning, populations are able to solve “needle in a haystack” problems due to learning guiding evolutionary search. Best (1999) extended the work of Hinton and Nowlan (1987) by demonstrating that, given the same “needle in a haystack” problem, social learning outperforms individual learning. Further work using simulated robots (Acerbi and Nolfi, 2007), animats (Borg et al., 2011), autonomous robots (Acerbi et al., 2007), ungrounded neural networks (Curran and O’Riordan, 2007), and binary strings (Jones and Blackwell, 2011) has contributed further to our understanding of the evolutionary advantages provided by social learning.

Social Learning in Increasingly Variable Environments

Numerous models and simulations have demonstrated the adaptive advantages, and highlighted potential failings, of learning strategies in environments exhibiting some level of consistent variation (Borg et al., 2011; Boyd and Richerson, 1983, 1995; Grove, 2011; Jones and Blackwell, 2011; Whitehead and Richerson, 2009). In this work we test the hypothesis that increasing, rather than simply consistent, environmental variability is sufficient to elicit the adoption of social learning. To test this hypothesis populations of individuals, constituted of binary string genotypes and phenotypes, are evaluated on their ability to match a target binary string, nominally known as the environment, with success measured by the Hamming distance between the phenotype and environment. Three classes of environment are used.

1. Static environments in which an environment’s target string remains unchanged.
2. Consistently variable environments in which each locus of an environment’s target string switches on or off at regular, frequent, intervals.
3. Increasingly variable environments in which the frequency of change increases over the period of evolution.

For each class of environment, populations exhibiting combinations of genetic evolution, individual learning and social learning are evaluated, with the learning rates of both individual and social learning allowed to evolve. Mean population fitness is recorded for each combination of environment and learning strategy, with data also collected on the evolved rates of social and individual learning and the reproductive fitness of individuals exhibiting different learning rates when both extra-genetic learning strategies are combined.

Our expectations were as follows.

1. Social and individual learning strategies, both separately and in combination, will outperform genetic evolution on all environments.
2. When evolved simultaneously social learning will be favored over individual learning, with individuals exhibiting higher levels of social learning having a higher reproductive fitness, thus showing that social learning is adopted over individual learning in increasing and consistently variable environments.

The Model

The model used is a genetic algorithm with steady state selection, in which individuals, constituted of binary string genotypes and phenotypes of length L , are assessed on their ability to match a binary target string or, as we shall refer to it here, an environment denoted as E (also of length L). A phenotype is assessed by measuring the Hamming distance between it and the environment. A phenotype is initially a copy of the genotype but can acquire information through evolution and learning, which is discussed in more detail later. This may be achieved by one of four strategies.

1. Genetic Evolution - at reproduction random mutations occur with probability p_{mut} at each locus.
2. Individual Learning - at each epoch (iteration of the steady state genetic algorithm) every individual flips each of the bits in its phenotype with probability p_{ind} .
3. Social Learning - at each epoch every individual copies each locus from a random other individual’s phenotype with probability p_{soc} .
4. Individual and Social Learning (Combined) - at each epoch every individual engages in either individual learning or social learning, with equal probability, at each locus in the phenotype.

The learning rate (per locus probability of flipping or copying) is allowed to evolve independently for each individual. That is to say that a population wide learning rate is not set. Both p_{ind} and p_{soc} are floating point values bounded within the range $[0, 1]$.

Variable Environments

Populations are tested on one of the three environmental setups introduced earlier, two of which exhibit some level of variability. Variability is dictated by a sine wave. At initialization each locus l in the environment is assigned a random value f , which is used to determine the binary value of the environmental locus at each epoch (1).

$$E^l = \sin((f^l \times epoch) \times (\pi/180)) \begin{cases} <0 \rightarrow 0 \\ >0 \rightarrow 1 \end{cases} \quad (1)$$

The range of values f may be initially set to is determined by which environment the population is being tested on:

1. No Variability (static): $f = 0$
2. Consistent Variability: $f \in N(1.8, \frac{1.8^2}{2})$
3. Increasing Variability: $f \in N(0.018, \frac{0.018^2}{2})$

Values of $f \approx 1.8$ equate to approximately one change per 100 epochs, with 100 epochs being considered to be one generation of the algorithm (where $L = 100$). A value of $f \approx 0.018$ equates to approximately one change per 10000 epochs, or one hundred generations. One change per generation is referred to as *high* frequency variability, one change per ten generations as *medium* frequency, and one change per one hundred generations as *low* frequency. As each environmental locus has a unique initial value of f , the sine wave dictating the value at each locus will be different, thus avoiding uniform environmental change.

For increasing variability tests the f values increase over time. The f value for any environmental locus (E^l) during increasing tests is determined by the initial f value at that locus (f^0), the maximum f value ($f^{max} = 1.8$), the current epoch and the number of epochs the evaluation is permitted to run for (2).

$$f^{epoch} = f^0 + (f^{max} - f^0) \times \left(\frac{epoch}{epoch^{max}} \right) \quad (2)$$

Evolution and Learning

Each test is populated by N individuals, each constituted of the following:

- $g \in \{0, 1\}^L$ - *genotype*, an L -bit string
- $h \in \{0, 1\}^L$ - *phenotype*, an L -bit string initially equal to g but subject to learning. The individual's fitness is L minus the Hamming distance between h and E .

- $p_{ind} \in [0, 1]$ - *individual learning rate*, set initially to 0. In populations allowed to learn in this manner p_{ind} may evolve via mutation.

- $p_{soc} \in [0, 1]$ - *social learning rate*, set initially to 0. In populations allowed to learn in this manner p_{soc} may evolve via mutation.

These properties are broadly consistent with the properties used by Jones and Blackwell (2011). However, unlike Jones and Blackwell (2011) the learning rates are not normalized to sum to unity, instead each rate may evolve to a maximum value of 1.

At each epoch two individuals are selected at random from the population for tournament selection. Reproduction then takes place between the tournament winning individual (the one with the higher fitness) and a random individual from the population, the progeny of this reproduction replacing the tournament loser. Reproduction consists of both recombination and mutation. Recombination is by way of single point crossover, where a random position $l \in [0, L - 1]$ is selected. Bits 0 to l being taken from one of the parents and bits $l + 1$ to $L - 1$ from the other, with the order of the parents determined at random at each reproduction. Mutation occurs at each locus in the child's genotype, with probability $p_{mut} = 1/L$ of the bit at that locus being flipped. Following reproduction g is copied without error to h which from this point in the child individual's lifetime is used for fitness evaluation and learning. In learning populations parental values of p_{ind} and p_{soc} are also inherited (depending on the learning strategy implemented for the population). The child inherits one of its parents' learning rates at random, with the learning rate then being mutated by the addition of Gaussian random noise (mean 0, standard deviation 0.01).

Learning comes in two distinct strategies: individual and social. At each epoch all individuals from a learning population are afforded the opportunity to learn. Individual learning takes the same form as mutation at reproduction, with each locus in h bit-flipping with probability p_{ind} . Social learning on the other hand is a little more involved: for each locus in h there is a probability p_{soc} of copying the tournament winning individual's equivalent locus. Copying the tournament winning individual in social learning strategies may be seen as akin to the "*copy-successful-individuals*" strategy outlined by Laland (2004) and implemented (though in a slightly different manner) by Jones and Blackwell (2011). In those populations exhibiting both individual and social learning in combination, which of the two learning strategies to use is chosen at random (50:50) for each locus of each individual, and applied with the appropriate learning rate. Individuals are also afforded the opportunity to unlearn any learned information. Each individual maintains a copy of their phenotype from before learning; if after learning their fitness is less than it was during the last epoch, their previous phenotype is restored.

Experimentation and Results

Experimentation was initially conducted on the static, consistently variable and increasingly variable environments. Forty initially random populations of size $N = 100$ were tested for each environmental setup: ten populations per learning strategy. Each environment, of size $L = 100$, was initially identical in its binary composition, as was the random number seed from which the initial f values were derived. Each population was run for 100000 epochs (1000 generations), with the population being sampled every 100th epoch (once per generation). The data presented here takes the mean performance of each of the ten populations per learning strategy at every generation.

A set of further tests were also conducted to assess in which conditions of environmental variability populations were likely to collapse. These tests were conducted in two differing setups. In both setups N was maintained at 100 but before standard tournament selection took place all individuals with a fitness less than $L/2$ were killed, these individuals being deemed to be unfit. If at this point the new population size $N' \leq N \times 0.1$ the population is considered to have collapsed and evolution is terminated. If the population does not collapse, tournament selection takes place to replace one surviving individual, and the population is then re-populated to $N = 100$ by the progeny of randomly selected other surviving individuals. The first test setup was conducted for a maximum of 100000 epochs, with populations reaching this epoch being considered as surviving populations.

The second population collapse test setup differs from the first in three distinct ways: tests were simulated for 200000 epochs; only populations exhibiting the individual and social learning strategies combined were tested; and social learning was prohibited from being used or evolving for the first half of each experiment.

Static Environments

As can be seen from figure 1(a), under static conditions both social learning and individual and social learning combined perform much better than genetic evolution and individual learning. These results are broadly consistent with those of Jones and Blackwell (2011) who also found social explorations to be advantageous and individual learning to sub-optimal in static environments. However, unlike Jones and Blackwell (2011), in these tests individual learning does not outperform no-learning (genetic evolution alone) over the entire simulation. This result is a little surprising given Hinton and Nowlan (1987), which demonstrates that individual learning should be able to better guide evolution than random mutation alone. It also seems that individual learning is not highly expressed when used in isolation. Figure 4 shows that under unchanging environmental conditions individual learning does not achieve a maximum p_{ind} of above 0.2, this value being lower than in all other environmental conditions and significantly lower than p_{soc} , which in static




Figure 1: Static Environment Tests: (a) Mean fitness of each learning strategy, (b) Mean fitness of individual and social learning with the evolved learning rates, (c) Reproductive fitness of combined learning rates.

environments achieves a value in excess of 0.7. Individual learning is also marginalized when expressed in combination with social learning. Figure 1(b) shows that when evolved together social learning outstrips individual learning by some distance, with individual learning becoming almost unused after an initial spike before 1000 epochs. Interestingly, for static environments the maximum value of p_{soc} achieved is larger when individual and social learning are found together, than when social learning is evolved in isola-

tion, implying that social learning requires individual learning to be fully expressed. As hypothesized social learning is adopted over individual learning, this adoption also being reflected by the reproductive fitness of individuals exhibiting the combined learning strategy as shown in figure 1(c). Individuals exhibiting intermediate values for p_{soc} and low values (below 0.1) of p_{ind} are shown to be more reproductively fit by contributing to a larger number of reproductions over the evaluation period.

Consistently Variable Environments

As shown in figure 2(a), under consistently variable conditions, where f is maintained at 1.8, the extra-genetic learning strategies all outperform no-learning (genetic evolution alone). In high variability environments non-learners find it difficult to track changes in the environment using mutation and recombination alone, causing populations of non-learners to average out at a fitness of $L/2$: no better than random. Of the extra-genetic learning strategies the combined strategy far outperforms individual and social learning alone. Individual learning when exhibited in isolation tends to find a stable value very quickly, but is unable to improve upon it. Social learning on the other hand rapidly (though also rather noisily) finds highly optimal solutions. However, the ever increasing reliance on social learning, as demonstrated by a maximum learning rate of above 0.9 (see figure 4), causes social learners' fitness to decrease to a value equal to that of individual learners, suggesting that overly conformist learning strategies are no better than trial-and-error personal innovations at tracking high levels of environmental change. By combining individual and social learning the negative aspects of both strategies in isolation seem to vanish: fitness does not stabilize at a sub-optimal value early on and fitness does not decrease over time. This suggests that the conformist bias imposed by social learning is in some way tempered by non-social innovation. However, as we can see in figure 1(b and c) social learning is largely adopted over individual learning, with p_{ind} being sidelined to values well below 0.1 and highly reproductive individuals exhibiting high levels of social learning and low levels of individual learning. The initial spike in individual learning seen early in the combined strategy, while p_{soc} is also low, may indicate that the vast majority of innovation is introduced into the population before it becomes overly conformist. It is also interesting to note that the spike in p_{ind} correlates well with the noisiest fitness period. Once enough innovation is introduced into the population innovation appears to be sidelined, although maintained at a low level, and individuals become increasingly reliant on social learning.

Environments of Increasing Variability

Unlike in consistently noisy environments, all populations exhibiting extra-genetic learning strategies find it difficult to maintain high levels of fitness when confronted with increas-




Figure 2: Consistently Variable Environment Tests: a) Mean fitness of each learning strategy, b) Mean fitness of individual and social learning with the evolved learning rates, c) Reproductive fitness of combined learning rates.

ing levels of variability (see figure 3(a)). As the environment becomes more noisy individual learning rates begin to increase, possibly to reintroduce an element of personal innovation to the population, which has become stagnant due to the high levels of conformist learning imposed by large quantities of social learning during times of minimal variability. The reproductive fitness of individuals, as seen in figure 3(c), is also interesting, as reproductively successful




Figure 3: Increasingly Variable Environment Tests: a) Mean fitness of each learning strategy, b) Mean fitness of individual and social learning with the evolved learning rates, c) Reproductive fitness of combined learning rates.

individuals tend to exhibit high levels of social learning and increased levels of individual learning, when compared to the reproductive fitnesses of individuals in consistently variable or static environments. It is also interesting to note the comparisons between maximum learning rates for social and individual learning on increasingly variable environments (see figure 4): despite individual learning being a necessary component of the combined strategy, it is not

exhibited to as high a degree as when found alone; conversely social learning is always exhibited at higher levels when accompanied by individual learning. This again suggests that, while social learning is adopted over individual learning, individual learning is necessary for social learning to be used to greatest effect (Acerbi and Nolfi, 2007; Acerbi et al., 2007). Evidence from all stages of environmental variability seem to tell a similar story, though to different degrees: social learning is widely adopted over individual learning when found together, with all extra-genetic learning strategies performing better than random on all tests. Extra-genetic learning strategies are also exhibited at higher levels in noisy environments than in static environments. The evidence presented does suggest that increasing variability is sufficient to cause the adoption of versatile survival strategies such as learning, with social learning being the learning strategy of choice.

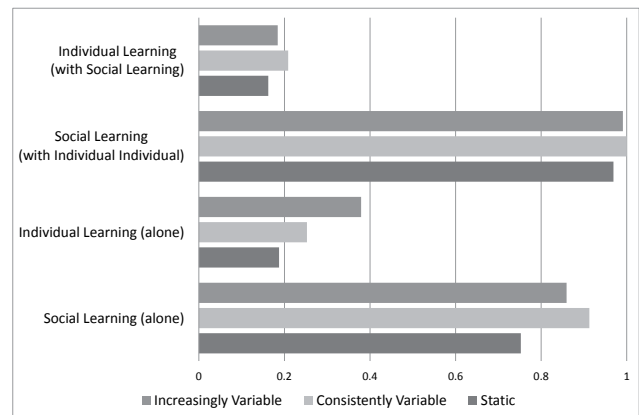


Figure 4: Maximum learning rates exhibited over all environmental test cases for all learning strategies.

Population Collapse in Variable Environments (Consistent and Increasing)

One of the pitfalls of the kind of genetic algorithm used so far is that even when populations exhibit low levels of evolutionary proficiency, they still survive; of course this is not the case in nature. To explore whether or not the learning strategies implemented in this model are really robust we have also implemented a set of tests where populations may become extinct. The first tests follow the test setups above, with populations exhibiting different learning strategies being tested on environments with consistent and increasing variability. Populations falling below $N \times 0.1$ individuals are considered as being collapsed.

Consistently variable environments were produced with four levels of variability;

1. No variability (static): $f = 0$
2. Low variability: $f \in N(0.018, \frac{0.018^2}{2})$

- 3. Medium variability: $f \in N(0.18, \frac{0.18^2}{2})$
- 4. High variability: $f \in N(1.8, \frac{1.8^2}{2})$

Learning Strategy	Static	Low	Medium	High
Genetic	100%	100%	100%	0%
Individual	100%	100%	100%	50%
Social	100%	100%	90%	0%
Individual & Social	100%	100%	100%	0%

Table 1: Consistently Variable Environments: % of populations surviving.

The percentages of populations surviving until the end of evaluation are reported in table 1. As may be expected, populations are unable to survive highly variable environments as the increased chance of death makes it all but impossible to re-adapt to new environments. However, individual learning does seem to be more robust than all other strategies, achieving a 50% survival rate on high frequency environments. It may be the case that higher rates of individual learning, though risky, are better able to deal with sudden environmental shifts. Social learning on the other hand begins to struggle in environments exhibiting medium amounts of variability. As with our earlier tests it may simply be the case that conformism spreads through the population, increasing the likelihood of population collapse. Combining individual and social learning alleviates the problem to some extent.

Increasingly variable environments were produced at three initial levels of variability: static, low and medium. In these environments variability increase throughout evolution, to a level of high variability.

Learning Strategy	Static	Low	Medium
Genetic	0%	0%	0%
Individual	100%	100%	100%
Social	0%	0%	0%
Individual & Social	0%	0%	0%

Table 2: Increasingly Variable Environments: % of populations surviving.

Unlike in consistently variable environments all learning strategies, excluding individual learning alone, result in populations that are unable to survive in any increasingly variable environment (see table 2). It seems social learning completely undermines individual learning when combined, perhaps owing to over-conformism in times of lower variability stagnating the population’s pool of knowledge to the point that the increase in individual learning, usually seen later in increasingly variable environments (see figure 3(b)) is insufficient to redeem the population’s fortunes.

As indicated by tables 1 and 2, individual learning is the only learning strategy robust enough deal with increasing

and high levels of environmental variability. However, in early tests the combined strategy of both individual and social learning was seen to be adaptive in all environmental settings. To investigate whether individual learning is necessary for the successful introduction of social learning we implemented a final set of tests. In these, individual learning was allowed to evolve in isolation for 100000 epochs before the introduction of social learning alongside it for a further 100000 epochs. These tests provide a greater challenge for populations as they are required to survive for twice the evaluation period previously tested. However, this increase in evaluation time does reduce the rate at which environmental variability increases during increasing-variability tests.

As table 3 shows, the evolution of individual learning prior to social learning does provide some benefits in increasingly variable environments, but only when beginning from medium levels of variability ($f = N(0.18, \frac{0.18^2}{2})$). It may be that noisier environments provide a greater selection pressure for high levels of innovation, which in turn introduces a larger pool of knowledge for social learning to access; or that the lower rate of increase in variability is significant. Further tests will need to be conducted to analyze the precise learning rates, reproductive fitnesses and death rates exhibited in these “goldilocks” conditions.

Variability	Static	Low	Medium	High
Consistent	100%	100%	100%	0%
Increasing	0%	0%	100%	N/A

Table 3: Individual and Social Learning: % of populations surviving when individual learning is allowed to evolve before the introduction of social learning.

Conclusions and Future Work

Reader and Laland (2002) have demonstrated that personal innovations (individual learning) and social learning co-vary across species. The above results go some way to explaining why social learning was adopted most strongly when combined with individual learning. It seems that individual learning is necessary for effective social learning. This may also be a mechanism of avoiding population collapse. Whilst social learning alone can maintain adaptive knowledge in the population, over-reliance on it can just as easily reinforce sub-optimal or incorrect knowledge when the environment is highly stochastic, potentially causing the population to collapse (Whitehead and Richerson, 2009). By maintaining a level of personal innovation alongside social learning, populations can maintain non-conformist local search whilst ensuring that useful innovations are transmitted over generations (Acerbi and Nolfi, 2007). However, in environments of lower variability conformist social learning ensures a high level of individual fitness. Individual learning on

the other hand may impose unnecessary local search which could cause individuals to lose useful adaptations if high levels of individual learning are maintained. The data presented here suggests that when environments are in minimally variable states individual learning plays a smaller role than it does in more variable environments. It is also found to be the case that mortality is greatly increased in environments of high or increasing variability when social learning is exhibited unless individual innovation is allowed to develop in isolation (Acerbi et al., 2007).

Our initial hypothesis (developed in order to test Potts's variability selection hypothesis), that when individual and social learning rates are evolved simultaneously, both increasing and consistently variable environments are sufficient for the adoption of social learning over individual learning, holds true here, though with two main caveats: individual learning is required for successful social learning, and population collapse may only be avoided when individual learning is allowed to pre-evolve in already noisy environments before the introduction of social learning. Both of these caveats require further investigation in steady state genetic algorithms, neural networks (Curran and O'Riordan, 2007) and grounded animat simulations (Borg et al., 2011).

The way noise is implemented also requires further investigation. Sine waves, though used elsewhere to produce environmental variation (Grove, 2011), are not the only pattern of environmental variability found in nature. Further tests could include empirically derived data sets (Grove, 2011) or *red noise* (Whitehead and Richerson, 2009).

References

- Acerbi, A., Marocco, D., and Nolfi, S. (2007). Social facilitation on the development of foraging behaviors in a population of autonomous robots. In Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., and Coutinho, A., editors, *Advances in Artificial Life*, volume 4648 of *Lecture Notes in Computer Science*, pages 625–634. Springer Berlin / Heidelberg.
- Acerbi, A. and Nolfi, S. (2007). Social learning and cultural evolution in embodied and situated agents. In *Proceedings of the First IEEE Symposium on Artificial Life, ALIFE'07*, pages 333–340. IEEE Press.
- Best, M. L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7(3–4):289–306.
- Borg, J. M., Channon, A., and Day, C. (2011). Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In *Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems, ECAL'11*, pages 102–109. MIT Press.
- Boyd, R. and Richerson, P. J. (1983). The cultural transmission of acquired variation: Effects on genetic fitness. *Journal of Theoretical Biology*, 100(4):567–596.
- Boyd, R. and Richerson, P. J. (1995). Why does culture increase human adaptability. *Ethology and Sociobiology*, 16:125–143.
- Curran, D. and O'Riordan, C. (2007). The effects of cultural learning in populations of neural networks. *Artificial Life*, 13(1):45–67.
- Grove, M. (2011). Speciation, diversity, and mode 1 technologies: The impact of variability selection. *Journal of Human Evolution*, 61(3):306–319.
- Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
- Jones, D. and Blackwell, T. (2011). Social learning and evolution in a structured environment. In *Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems, ECAL'11*, pages 380–387. MIT Press.
- Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, 32(1):4–14.
- Nolfi, S. and Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, 7(1):89–113.
- Potts, R. (1996). Evolution and climate variability. *Science*, 273(5277):922–923.
- Potts, R. (1998a). Environment hypotheses of hominin evolution. *Yearbook of Physical Anthropology*, 41:93–136.
- Potts, R. (1998b). Variability selection in hominid evolution. *Evolutionary Anthropology*, 7(3):81–96.
- Reader, S. and Biro, D. (2010). Experimental identification of social learning in wild animals. *Learning and Behavior*, 38(3):265–283.
- Reader, S. M. and Laland, K. N. (2002). Social intelligence, innovation, and enhanced brain size in primates. *PNAS*, 99(7):4436–4441.
- Tattersall, I. (2009). Human origins: Out of africa. *PNAS*, 106(38):16018–16021.
- Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology*, 28:509–529.
- Whitehead, H. and Richerson, P. J. (2009). The evolution of conformist social learning can cause population collapse in realistically variable environments. *Evolution and Human Behavior*, 30(4):261–273.
- Whiten, A. and van Schaik, C. (2007). The evolution of animal cultures and social intelligence. *Phil. Trans. R. Soc. B*, 362(1480):603–620.