

Social Learning and Evolution in a Structured Environment

Daniel Jones¹ and Tim Blackwell¹

¹Department of Computing, Goldsmiths, University of London, UK
d.jones@gold.ac.uk, t.blackwell@gold.ac.uk

Abstract

We survey the relationships between evolution, individual learning and social transmission within well-mixed and structured environments. With a novel individual-based simulation, we determine the regimes under which each mode of learning dominates, in terms of the environment's relative complexity and its rate of change. We show that social learning can give rise to a particularly potent form of the "Baldwin effect", wherein an organism develops an innate trait having first acquired it socially. We demonstrate that social learning is of increased significance in a structured environment.

Introduction

To operate successfully in a Darwinian system, it is advantageous to possess maximal information about our environment. This is reflected in the functional information that all living creatures inherit via DNA, which codes for the set of functional characteristics most likely to benefit an organism in its future surroundings (Avery, 2003).

However, the environment which produced a parent is never quite the same as when its child is born. Ecological habitats are continually changing as their inhabitants consume and produce resources, with environments effectively co-evolving with their organisms. Inheritance is thus an intrinsically probabilistic process, which uses rules of thumb to provide the best possible solution given the expected habitat based on previous generations (Seth, 2007).

To optimally deal with uncertainty, all organisms exhibit some degree of *phenotypic plasticity* (West-Eberhard, 1989; Scheiner, 1993): the ability to alter behaviour or physiology in response to environmental conditions. By allowing some morphological decisions to be fixed later in an organism's lifetime, evolution can effectively defer decisions about functional specifics. This appears to be particularly prominent in fluctuating and heterogeneous environments, which are naturally less predictable (West-Eberhard, 1989).

Evolution, learning and culture

The most radical form of phenotypic plasticity is behavioural learning, which can respond rapidly and flexibly

to novel stimuli based on prior experience. Learning can be considered as giving foresight to the blind process of evolution, by enabling an organism to search the fitness landscape around the point determined by its genotype (Belew, 1990; Borenstein et al., 2006). As Maynard Smith (1987) observes,

"...finding the optimal [solution] in the absence of learning is like searching for a needle in a haystack. With learning, it is like searching for the needle when someone tells you when you are getting close." (Maynard Smith, 1987, p762).

In this paper, we are concerned with two forms of learning: *individual exploration*, which we shall define as trial-and-error learning solely between an individual and its (abiotic) environment; and *social learning*, in which an organism acquires traits by observing or mimicking the behaviours of others (Lefebvre and Palameta, 1988). Countless species engage in social learning (Galef and Laland, 2005; Laland, 2004a), through mechanisms such as mimicry, teaching, and goal emulation. We shall here deal with a general case in which a trait is exhibited after observing another organism as a model (the "exemplar").

When evolutionary systems are extended with lifetime behavioural plasticity, we should expect some interesting interactions to arise. One which came to the attention of the first generation of evolutionary theorists after Darwin (Baldwin, 1896; Morgan, 1896) is the "Baldwin effect", a term coined half a century later (Simpson, 1953) after one of its progenitors, ironically in an attempt to discredit the theory.

The general pattern encapsulated within the Baldwin effect is as follows:

1. A population arises in which some trait P becomes beneficial.
2. Some individuals arise which, through their phenotypic plasticity, are able to learn P .
3. In some of these individuals, the trait P becomes innate (*genetic assimilation*).

With the assumption that innate behaviours are less costly than those which are plastic, we would then expect selective pressure to lessen on these particular learning capabilities: if we can do it by nature, we no longer need to be able to learn it (West-Eberhard, 2003).

The status and prevalence of genetic assimilation within real-world ecosystems is as yet unresolved, and subject to some controversy (Pigliucci et al., 2006). Due to its onerous requirements – a species sufficiently advanced to partake in social learning, bred over a sufficient number of generations for a trait to become genetically incorporated – it is difficult to observe via *in vivo* studies, though Waddington’s (1953) “veinless” study elegantly demonstrates its biological plausibility. It is, therefore, a well-suited candidate for *in silico* experiments.

Theoretical studies of social learning

A large body of theoretical work has been developed at the confluence between evolution, learning and cultural transmission (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1985; Wakano et al., 2004). The watershed work was a computational model by Hinton and Nowlan (1987, henceforth ‘H&N’), who extended binary genetic algorithms with an undefined third value, whose outcome is determined by lifetime learning. Though intentionally simplistic, this model effectively demonstrated the “needle in a haystack” function of learning as a dowsing rod to guide evolution towards discontinuous fitness peaks.

Belew (1990) and Best (1999) have extended H&N with differing forms of cultural transmission, both in a well-mixed environment, incorporating oblique and horizontal forms of social exchange. While Belew models cultural exchange as a bias towards higher fitness, we will follow Best as treating it as a more neutral form of behavioural mimicry, in which an organism may imitate deleterious as well as adaptive behaviours.

Models of social transmission within a spatial environment include work by Boyd and Richerson (1988), Lowen (1996) and Borenstein (2003). A consensus view has emerged that sociality is of benefit within structured environments. We wish to extend these analyses to survey the regimes under which each mode of learning excels, and whether unforeseen mixed strategies may come to the fore given a heterogeneous, individual-based model.

We also wish to model the scenario conjectured by Papineau (2005), who posits that the Baldwin effect may become significantly more prominent when bolstered with social learning. This can be roughly encapsulated by the inequality:

$$p(G) \ll p(L) \ll p(S) \quad (1)$$

Where $p(G)$ is the probability of exhibiting a trait innately, $p(L)$ is the probability of learning it through ex-

ploration, and $p(S)$ is the probability of acquiring the trait through social learning. Quite simply, wherein it is effectively impossible to acquire a functional trait P through evolution – perhaps because it is comprised of multiple sub-traits, which are jointly necessary to reap a fitness benefit – this process may be somewhat more likely when lifetime learning is possible, and even moreso when social learning enables organisms to share traits.

This argument, though intuitively sound, is thus far based on heuristic assumptions. The following model is intended to quantitatively explore situations in which a social Baldwin effect can take place, and particularly those in which combination strategies can arise: evolved individuals can exhibit both individual and social learning in proportion. We are furthermore interested in how these phenomena interact in a context which is explicitly spatial, a combination which has not yet received significant attention.

Model specification

We will now describe the components of the individual-based model used to explore these ideas¹. An *environment* E consists of a B -bit string, representing a ‘target’ task: $E \in \{0, 1\}^B$. The current environmental state can therefore be considered as a vertex on an B -dimensional hypercube.

It is inhabited by a population of N agents, each of which has the following properties:

- $b_{evo}, b_{exp}, b_{soc} \in [0, 1]$ – *behavioural traits* determining the propensity towards evolutionary instinct, individual exploration, and social learning. These are collectively normalised to sum to unity.
- $g \in \{0, 1\}^B$ – *genotype*, a B -bit string corresponding to the capability to fulfil the environment’s target task.
- $p \in \{0, 1\}^B$ – *phenotype*, a B -bit string initially equal to g , but subject to modification through individual and social learning. If p is equal to E then the agent’s fitness is maximised.
- Ω – current metabolic state, initialised to a constant Ω_0 .

An agent’s current phenotype determines how well it complies with the environment’s demands, based on its Hamming distance from E . Its metabolic state determines the extent to which it has ‘grown’ throughout its lifetime.

Actions and learning

Every timestep, each agent selects a behavioural mode according to a weighted random of $\{b_{evo}, b_{exp}, b_{soc}\}$:

- b_{evo} – act according to the agent’s current phenotype
- b_{exp} – act according to the agent’s current phenotype, with β bits toggled at random

¹For all subsequent parameter values, see the *Methods* section.

- b_{soc} – act according to the agent’s current phenotype, with β bits copied from a neighbour using roulette wheel selection weighted by Ω . With a probability p_{noise} , each of these bits may be copied erroneously (that is, toggled from $0 \rightarrow 1$ or $1 \rightarrow 0$). This models the inaccuracy present in real-world imitative learning: a behaviour may be only partially observed, or reproduced incorrectly.

If b_{exp} or b_{soc} is employed *and* the resultant action gives a higher payoff than the agent’s own current phenotype, the corresponding bits in p are replaced by the new action: discovering (or imitating) a successful new trait results in its being incorporated into the agent’s roster. This reflects phenotypic plasticity, where β is the limiting factor on the rate at which new skills can be acquired.

In the case of b_{soc} , weighting the exemplar by their Ω value reflects a tendency towards mimicking those organisms which are perceived as being fittest. This is described by Laland (2004b) as a “*copy-successful-individuals*” strategy, as observed in avian, chimpanzee and bat societies.

The agent’s metabolism is then modified according to the following update rule:

$$\Delta\Omega = \left(1 - \frac{H(p, E)}{B}\right)^{\alpha^{-1}} \quad (2)$$

where H denotes the Hamming distance between two bit strings. The exponential of α is used to determine the fitness differential between perfect and almost-perfect task fulfilment: a lower value of α means that payoffs fall more rapidly with distance. With $\alpha = 1$, scaling is linear in distance.

In general, if an agent’s g matches precisely the tasks specified in E_i , its metabolism will increase by the maximal value of 1. If g is precisely the complement of E_i , its metabolism will increase by 0.

Taken as a population mean, the metabolic rate $\Delta\Omega$ can be considered as a measure of *fitness*, as it is directly proportional to reproductive rate. We will subsequently use the terms interchangeably.

Reproduction

When an agent’s metabolism Ω reaches the value $2\Omega_0$, the agent reproduces asexually. Its offspring has an identical genotype, subject to each bit of g mutating with small probability p_{mut} . Behavioural traits b_{evo} , b_{exp} , b_{soc} are modified by a zero-mean Gaussian noise function, standard deviation μ , and clipped to $[0, 1]$. These are again collectively normalised to unity. The child replaces a member of the population selected uniformly randomly, and its parent’s Ω value is reset to Ω_0 .

Sexual recombination was considered as a reproductive strategy. Kauffman (1993) observes that recombination is an effective method of finding ‘middle ground’ locations between points on a complex fitness landscape. However,




Figure 1: Distribution of behaviours in a static environment, averaged over 25 simulation outcomes.

given our single-peaked landscape, we focus on clonal reproduction for the sake of simplicity. A number of recombinative trials indicated that the results would not be qualitatively different.

Results

The results of this model are presented in incremental form, with processes introduced gradually. The motivation behind this approach is to understand pairwise interactions between adaptive mechanisms. By doing so, we hope to fully understand the causal basis behind the emergent phenomena.

Static environment

We initialise the environment’s task to 1^B for clarity (following Hinton and Nowlan (1987)). Behavioural traits are initialised to uniformly random values, and the population left to evolve.

The changing distribution of behavioural traits over time is shown in Figure 1, as averaged over multiple iterations (see Methods). At step 0, the frequency of each is $\frac{1}{3}$, indicating the initial uniformly random distribution of behavioural modes.

The dynamics can subsequently be divided into three phases. (i) Between steps **1–4000**, the population is dominated by social learners, with a generally low level of genotypic fitness meaning that a costlier but fitter social learning is preferable. (ii) From steps **4000–9000**, the trait has been assimilated into the genotype of the majority, and so innate b_{evo} agents outcompete their costlier plastic rivals. (iii) Beyond step **9000**, a stable optimum is reached.

Sharpening these costs by reducing payoff scaling factor α_0 results in a more rapid convergence to a predominantly b_{evo} population. This also reduces the effectiveness of lifetime learning, of course, which introduces a penalty in fluctuating environments.

This is a clear example of the Baldwin effect. Phenotypically plastic individuals first outcompete their peers (*i*) as they scramble to higher fitness through learning and social exchange, and are subsequently replaced (*ii*, *iii*) by innate mutants, who do not bear the costs of exploration.

A repeated trial wherein all agents begin with a genotype of 1^B reveals, as anticipated, that they continue to maintain a stable state with only low levels of social and individual learning.

Static environment with restricted strategies

The above experiment was repeated with a fixed trait mutation factor of $\mu = 0$ and initial behavioural traits restricted to specific combinations: either pure evolutionary learning (b_{evo}), or evolution plus learning ($b_{evo} + b_{exp}$), or evolution plus social learning ($b_{evo} + b_{soc}$), or all three traits in combination.




Figure 2: Convergence rates with four different strategies: b_{evo} , $b_{evo} + b_{exp}$, $b_{evo} + b_{soc}$, and $b_{evo} + b_{exp} + b_{soc}$.

Figure 2 depicts the relative effectiveness of each strategy in a static environment, plotting the global mean fitness (that is, $\Delta\Omega$) over a number of generations. The key indicators of success are the convergence rate and the value to which the population converges.

All four strategies eventually converge around the same peak of 0.8. The times taken to do so, however, are markedly different. Notably, evolution plus learning takes substantially more time to converge than pure evolution alone, and continues to trail throughout the simulation. This confirms the findings of Borenstein et al (2008) that, in a static, unimodal fitness landscape, individual learning actually serves to slow convergence rates.

With social learning, convergence times are markedly more rapid, reaching a mean fitness of 0.5 in less than half the time as evolution or evolution plus learning.

Static environment with single perturbation

Here, the scenario was repeated as per the Static Environment case, with with an environmental perturbation induced at step 10000: each of its bits were flipped according to a

probability $p = 0.5$. As indicated in Figure 3, this change results in a temporary increase in social and exploratory learners, bringing up phenotypic fitness through plasticity whilst evolution takes time to work out the necessary series of mutations.




Figure 3: A single perturbation occurs at $t = 10000$. Subsequently, agents are selected for increased social and exploratory learning tendencies.




Figure 4: Genotypic and phenotypic fitness after perturbations.

This further demonstrates Baldwin-like phenomena, and moreover with a *social* focus: whilst a small proportion of individuals respond to environmental change by switching to individual exploration, the predominant trend is to rely on social learning, observing the behaviour of others to maximise fitness.

Fluctuating environment

We now extend the above by introducing irregular environmental fluctuations. Each time step, a single bit of the environmental task may be toggled, according to a small probability p_{switch} . A value of $p_{switch} = 0.01$ reflects an ex-




Figure 5: With a regularly fluctuating environment ($p_{switch} = 0.01$), a social learning strategy is more frequently adopted.




Figure 6: With a higher reliance on phenotypic plasticity, genetic selection pressure is lower, and so genotypic constitution drifts.

pected period of 100 timesteps between fluctuations. With an initial metabolism $\Omega_0 = 10$ and a typical $\Delta\Omega = 0.5$, the environment could be expected to fluctuate once every 5 agent-lifespans.

The optimal combination of strategies is markedly different than in a fixed environment (Figures 5 and 6). Social learning dominates, reflecting the benefit of a faster adaptive rate with changing fitness targets.

Convergence patterns are also markedly different (Figure 7). In a rapidly changing environment ($p_{switch} = 0.005$), no strategy attains a mean fitness of above 0.7: even with the ability to mimic successful peers, it is difficult to maintain a high performance level in the face of continuous change. Social learning is frontrunner once more, with $b_{evo} + b_{exp}$ significantly outperforming pure evolution. This reflects the advantage in random trials when an organism's genome is lagging behind the rate of change of its environ-




Figure 7: Convergence rates are markedly different within a fluctuating environment ($p_{switch} = 0.005$).

ment.

Environmental rate and complexity

To gain fuller insight into the relative strengths of individual, social and exploratory learning in fluctuating environments, we carried out an array of simulations over a range of rates of change ($p_{switch} \in [5 \times 10^{-6}, 0.5]$) and environmental complexities ($B \in [1, 2048]$). Each permutation of p_{switch} and B was executed for 10^5 timesteps, and a snapshot taken of the final distribution of behavioural traits. These are mapped in Figure 8, with the dominance of each trait demonstrated by its share of the pie chart at the given (*complexity*, *rate*) combination.




Figure 8: Dominant learning modes at equilibrium, varying dimensionality and rate of change of environment.




Figure 9: Ring of environmental ‘cells’. Tasks are numbered according to Gray code: note that adjacent cells have 1 bit difference.

At low rates of change and in simple environments, the population demonstrates a significantly greater mean growth rate, with a clear prevalence of b_{evo} . As *either* rate or complexity increase, strategies become more mixed, with a trend towards social learning at median values of each. A greater amount of noise in the results suggests that selection pressures per are weaker, leading to more vulnerability to stochastic variation.

At very high rates of change or complexity, a sudden increase of b_{exp} dominance is evident. This is relatively simple to interpret in the former case: if the environment is changing faster than information can percolate through a social group, then even social learning is inferior to individual trial and error.

The benefit of learning in very complex environments is less clear; even in a virtually static environment ($p_{switch} = 5 \times 10^{-6}$), exploration exceeds innate strategies for $B = 2048$. Analysis reveals that the fitness ($\Delta\Omega$) in these regimes is uniformly low: given the rapid fitness falloff due to α , neither evolution nor social learning are fit to find suitable values. With such a large parameter space, the optimal resort is simply bit-wise trial and error.

Structured environment

We now extend the model by introducing a form of spatial structure. The single environment is replaced by a 1-dimensional ring of L environmental ‘cells’, each with a distinct population and set of tasks (Figure 9). Inhabitants of each cell can only interact with each other. As before, the size of the total metapopulation remains constant at N .

Each environmental cell has a single neighbour on each side, with the rightmost cell wrapping around to the leftmost. To introduce correlation between the task structure of neighbouring cells, integer sequences were produced us-

ing Gray code, a base-2 numeric encoding in which any two adjacent integers have a Hamming distance of 1. A further property of Gray code sequences is that they are cyclical, with the first and final integer of any 2^N -length sequence also one bit apart. It is possible, therefore, to produce integer rings with pairwise Hamming distance of 1.

During a timestep, agents may move from their current location to a neighbouring cell with a small probability p_{move} . Evolutionary and individual learning are unaffected; social learning, however, is now restricted to exemplars within the agent’s current location.

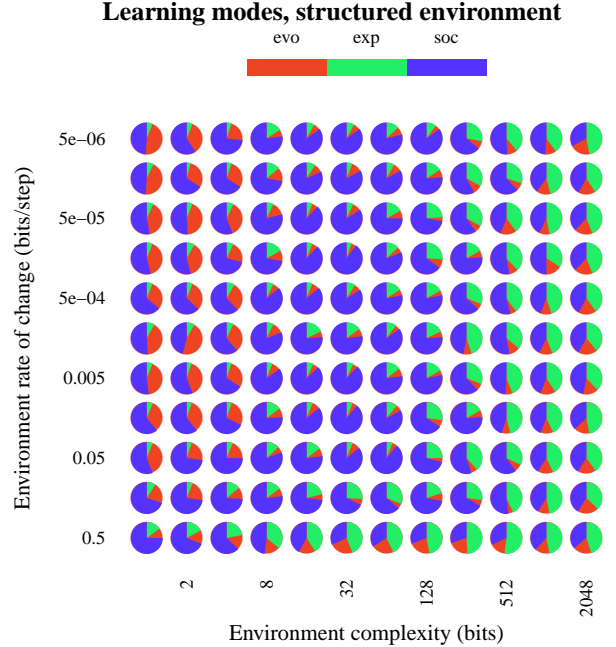


Figure 10: Dominant learning behaviours in a structured environment with migration.

From Figure 10, we can see that the overall distribution of learning patterns is similar: in static, simple environments, innate behaviour is commonplace, moving towards social learning in more complex and fluctuating contexts.

With rapid fluctuations, exploratory learning still excels, but it appears to have slightly less prevalence in environments with a large B value. This appears to be due to what we will call the “*local specialist*” effect: in a well-mixed, complex environment, there are a large range of behaviours to mimic, drawn from a large variety of sources. Even if we select our exemplar wisely, we may still mimic the wrong behaviour, as they too will be employing random search to test out new tasks.

In a structured environment, conversely, we have smaller number of local neighbours to mimic. With the roulette-wheel mechanism used to select exemplars, a smaller population also means a higher likelihood of selecting a highly-

ranked target. Combined with the fact that selection pressure still operates by removing the weakest agents of the *global* population, this means that positive behaviours are disseminated and adopted rapidly within individual cells, giving rise to social ‘specialist’ cliques.

Structured environment with migration

In this scenario, we remove environmental fluctuations, and instead vary p_{move} : the rate at which migration occurs. Figure 11 depicts this new distribution, with its Y-axis representing the rate of migration, over the same range as the fluctuation rate was previously plotted. The only variation that an agent will experience in its environment is when it moves from cell to neighbouring cell, so this can effectively be considered analogous to our previous environmental fluctuations.

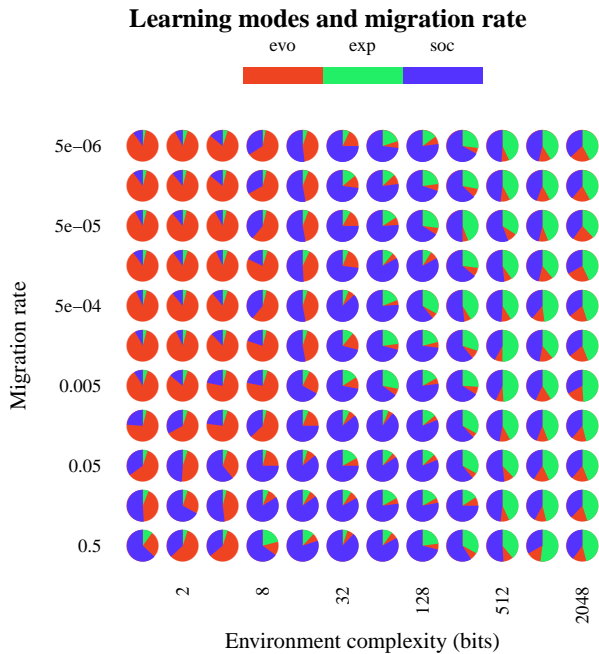


Figure 11: Dominant learning behaviours by migration rate and environmental complexity.

Learning strategies do not appear to correlate significantly with movement rates, despite the fact that movement between cells *does* effectively change the environment that an agent experiences. However, a significant difference takes place at high movement rates. Rather than resorting to individual trial and error, agents make greater use of social learning. This may be interpreted as a more focused version of the local specialist effect; in a static environment with frequent migration, we would expect the rapid dissemination of local knowledge to become of paramount importance.

In other words, if an agent is commonly moving from environment to environment, the most effective way to obtain

information about novel functions is to mimic the locals. This has both logical and biological plausibility.

Discussion

We have seen that three discrete regimes appear within varying classes of environment, each favouring different forms of learning. Within static environments, innate behaviour excels; within rapidly-changing environment, exploratory behaviour comes to the fore. Social behaviour, conversely, fills the gap between the two.

Beyond this, social transmission serves to inform and drive subsequent evolutionary behaviour, with what Papineau (2005) terms a “social Baldwin effect”. Our results suggest that this may play a pivotal role in the aftermath of major environmental changes – which, in ecosystems wherein organisms act as background to other organisms, may also correspond to the aftermath of major *ecological* changes.

In a structured environment, we have seen that successful behaviours are disseminated rapidly, due to reliance on smaller, focused groups of ‘specialists’ in each location. With greater environmental complexity, these local effects are amplified yet further.

Methods

Simulation results are averaged over 25 iterations to minimise stochastic fluctuations. Default variable values are given below.

Variable	Value	Comments
N	256	Population size
L	32	Number of spatial locations
B	32	Number of bits per task
Ω_0	10	Initial metabolic state
α	0.1	Rate of fitness dropoff based on task proximity
β	1	Maximum number of bits learned per timestep
μ	0.01	s.d. of mutation as applied to $b_{evo}, b_{exp}, b_{soc}$
p_{switch}	0.01	Probability of a single environmental fluctuation
p_{noise}	0.25	Probability of incorrect observation during mimicking
p_{mut}	0.01	Probability of sustaining a mutation per gene
p_{move}	0.1	Probability of migrating to a neighbouring cell

Acknowledgements

The authors would like to thank Christophe Rhodes, Tim Taylor and Mark d’Inverno for their insightful comments.

This work was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/P503418/1.

References

- Avery, J. (2003). *Information Theory and Evolution*. World Scientific.
- Baldwin, J. M. (1896). A new factor in evolution. *The American Naturalist*, 30(354):441–451.
- Belew, R. K. (1990). Evolution, learning and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4:11–49.
- Best, M. L. (1999). How culture can guide evolution: An inquiry into Gene/Meme enhancement and opposition. *Adaptive Behavior*, 7(3-4):289–306.
- Borenstein, E., Feldman, M. W., and Aoki, K. (2008). Evolution of learning in fluctuating environments: when selection favors both social and exploratory individual learning. *Evolution; international journal of organic evolution*, 62(3):586–602.
- Borenstein, E., Kendal, J., and Feldman, M. (2006). Cultural niche construction in a metapopulation. *Theoretical Population Biology*, 70(1):92–104.
- Borenstein, E. and Ruppin, E. (2003). Enhancing autonomous agents' evolution with learning by imitation. *Journal of Artificial Intelligence and the Simulation of Behavior*, 1(4):335–347.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the Evolutionary Process*. The University of Chicago Press.
- Boyd, R. and Richerson, P. J. (1988). *An Evolutionary Model of Social Learning: The Effects of Spatial and Temporal Variation*. Erlbaum, Hillsdale, NJ.
- Cavalli-Sforza, L. L. and Feldman, M. W. (1981). *Cultural Transmission and Evolution: A Quantitative Approach*. Princeton University Press.
- Galef, B. G. and Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioScience*, 55(6):489–499.
- Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
- Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, USA, 1 edition.
- Laland, K. N. (2004a). Extending the extended phenotype. *Biology and Philosophy*, 19(3):313–325.
- Laland, K. N. (2004b). Social learning strategies. *Learning & Behavior*, 32(1):4–14.
- Lefebvre, L. and Palameta, B. (1988). Mechanisms, ecology, and population diffusion of socially-learned, food-finding behavior in feral pigeons. In Zentall, T. R. and Galef, B. G., editors, *Social Learning: Psychological and Biological Perspectives*, pages 141–164. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Lowen, C. B. and Dunbar, R. I. M. (1996). Cultural evolution in a structured population. In *ESRC Centre on Economics Learning and Social Evolution*.
- Maynard Smith, J. (1987). When learning guides evolution. *Nature*, 329:761–762.
- Morgan, C. L. (1896). On modification and variation. *Science*, 4:733–740.
- Papineau, D. (2005). Social learning and the baldwin effect.
- Pigliucci, M., Murren, C. J., and Schlichting, C. D. (2006). Phenotypic plasticity and evolution by genetic assimilation. *The Journal of experimental biology*, 209(Pt 12):2362–2367.
- Scheiner, S. M. (1993). Genetics and evolution of phenotypic plasticity. *Annual Review of Ecology and Systematics*, 24:35–68.
- Seth, A. K. (2007). The ecology of action selection: insights from artificial life. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 362(1485):1545–1558.
- Simpson, G. G. (1953). The baldwin effect. *Evolution*, 7(2).
- Waddington, C. H. (1953). Genetic assimilation of an acquired character. *Evolution*, 7(2):118–126.
- Wakano, J., Aoki, K., and Feldman, M. (2004). Evolution of social learning: a mathematical analysis. *Theoretical Population Biology*, 66(3):249–258.
- West-Eberhard, M. J. (1989). Phenotypic plasticity and the origins of diversity. *Annual Review of Ecology and Systematics*, 20.
- West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. Oxford University Press, USA, 1 edition.