The Baldwin effect reconsidered through the prism of social learning

Nam Le*, Michael O‘Neill‡, Anthony Brabazon§
Natural Computing Research and Applications Group
University College Dublin
Dublin, Ireland
Email: * namlehai90@gmail.com

Abstract—**Learning, through the Baldwin effect, has been shown to provide beneficial adaptation in evolving populations, which is not genetically inherited. More interestingly, learning can be classified into two categories. The first is asocial (individual) learning when individuals learn by directly interacting with their environment, e.g. trial-and-error. The second is social learning when individuals learn from others, e.g. imitation learning — a form of social learning. There has been a large body of research concerning the former. Drawing on a rich literature concerning social learning in animals, this paper reconsiders the Baldwin effect through the prism of social learning. The paper provides an introduction of social learning and related concepts. We also propose our own algorithm combining evolution and social learning and test this on a compact simulation. The results obtained are encouraging and indicative. Social learning in combination with asocial learning has been shown to benefit evolutionary search. Experimental results also indicate that social learning can also result in the Baldwin effect, increasing the average fitness of the population. The level of information fidelity is also investigated with respect to its influence on the effect of social learning by a series of comparative experiments. A number of suggestions for future work are also presented.**

I. INTRODUCTION

Evolution and learning are two different ways in which the behavior, and other traits, of organisms can change. Evolution is change at the genetic level of a population, in which organisms reproduce selectively subject to mechanisms,like mutation or sexual recombination or both, which maintain inter-individual variability. This causes changes in the population from one generation to the next. Learning, on the other hand, is change at the individual level. By interacting during its life with a specific environment an organism can change its behavior by incorporating, through its experience, aspects of the environment into its internal structure.

The orthodox view of evolution is that changes due to learning during life are not inherited and, more generally, that learning does not influence evolution. The basis for such a view is the physical separation between the germ cell line and the somatic cell line. Changes due to learning concern somatic cells whereas evolution is restricted to the germinal cells. Since the two types of cells are physically separated, it is not possible for changes in the somatic cells to have an influence on evolution. On the other hand, Baldwin [1], Waddington [2], and several others [3] have claimed that there is an interaction between learning and evolution and, more specifically, that learning can have an influence on evolution. This is called the **The Baldwin effect**.

Hinton and Nowlan presented a classic paper in 1986 [4] to demonstrate an instance of the Baldwin effect in the field of Evolutionary Computation. Their initial success motivated a number of further studies [5], [6], [7], [8], [9], [10], to name but a few.

Learning generally can be classified into two types, namely *social learning* and *individual (asocial) learning*. Individual or asocial learning can be understood as learning when the learner directly interacts with its environment, e.g. trial-and-error, without the presence of others. Social learning, on the other hand, can be interpreted as learning from others, e.g. imitation, through observation or interaction. The learning type used in the Hinton model is an instance of individual learning, where each individual finds the correct solution itself based on locally searching through its neighborhood.

Baldwin himself once mentioned the presence of imitation, a form of social learning in our language today, as an important aspect of how learning guides evolution [1]. Therefore, the main aim of this paper is to investigate how social learning guides evolution through the Baldwin effect. To achieve this, a literature of social learning with some related concepts will be briefly presented. We propose our own algorithm combining evolution with learning, comprising of social and asocial learning, to demonstrate the presence of the Baldwin effect. Based on this, we study how the relationship between asocial and social learning influences an evolutionary process. Furthermore, through social learning individuals with lower fitness tend to be closer to better individuals from whom they learn. Therefore it is plausible to think that social learning increases the average fitness of the population. We investigate the impact of social learning on average fitness in our simulation.

The remainder of the paper is organized as follows. Section II presents some main points in the Baldwin effect, the model used by Hinton and related researches. The literature on social learning is briefly mentioned in this section. Section III shows the experimental setup and pseudo-code that we use in our simulation, while the experimental results are analyzed and discussed in section IV. The final section is for conclusion and future researches.

II. BACKGROUND

A. Baldwin Effect

In 1896, the Baldwin Effect was proposed by both Baldwin [1] (published in Nature Magazine) and Lloyd-Morgan [3] (published in Science Magazine). In the paper *'A new factor in Biology'* [1], Baldwin used various phrases that warrant explanation to understand his idea. First, he coined the term *Organic selection* as a general way to refer to individual adaptation by learning that conforms to Darwinian processes. He adopted the term *'ontogenetic adaptation'* which includes pleasure and pain, reinforcement, and imitation. Ontogenetic adaptation was said to strengthen the organisms and make them more responsive to environmental contingencies. That also means that the more ontogenetically adapted an organism, the more adaptive it is, and the higher the chance it leaves offspring. It was also stated by Baldwin that in some species, ontogenetic adaptations are made more effective through *'Social heredity'*. By *social heredity* he meant any function or behavior socially acquired and transmitted through the interaction between animals. From that paper, it is also suggested that social heredity can be maintained in a population for a period of time, counting by generations, until phylogenetic modification turns them into congenital instincts. (please refer to [1] for more details about Baldwin‘s terminology usage).

To make it clear in today's language, we summarize the Baldwin effect as follows:

i. If an individual is capable of learning an adaptive behavior, it will be given selective advantage over the rest of the population, leading to the increase of that individual‘s genes in the subsequent generations.
ii. Learning could be social or asocial. Assume that one adaptive behavior is found in the population, if social learning is permitted it will propagate that adaptive behavior through the population very quickly.
iii. Learning involves a cost. It is the cost of learning that makes selective pressure favor individuals who can learn with minimum cost compared to others in the same population. Thus, if the cost of social learning is less than that of individual learning, individuals with more social learning abilities will be favored.
iv. Over many generations, the learning capability becomes close to being innate, or genetically specified.

B. Hinton and Nowlan‘s model

The work by Hinton and Nowlan was proposed more than 30 years ago and has received many citations. We discuss their model in detail for the purpose of clarity and future usage in our experimental design.

Hinton and Nowlan used a genetic algorithm [11] as a model of evolution to search for the wiring of a neural network in hopes of making an individual learn effectively.

In their setting, each individual agent has a *genotype* — a string of twenty characters, which encodes the neural network connections as its *phenotype* with corresponding twenty possible wirings. Each position in a genotype, or *locus*, can have three alternative values: '0', '1', and '?'. Each locus is randomly initialized with 25% chance of being assigned a '0', 25% chance of being '1', and 50% chance a '?'. The value '1', also called **allele**, specifies that the neural connection exists, whereas the allele '0' says that the connectivity is absent. The allele '?' allows for lifetime learning (or plasticity). Each agent will have 1000 rounds of learning during its lifetime. On each round, an individual agent is allowed to do individual learning by changing its allele '?' to either '0' or '1' as the expressed value.

In their model, a neural network can be considered successful if it has all of the neural connections correctly specified. Thus there is exactly one right answer in the space of 2^{20} possible settings. Without loss of generality, we can assume that the correct solution is when all the connections are present; the correct solution should be found with all loci specified by the allele '1', or *111...1*. Only the correct answer has a high value (fitness value), and the incorrect ones (with '$0s$' in the genotype) have the same fitness of 1. For example, the agent with the genotype like this 1110?...1 will have the fitness of 1 because it contains one zero in its genotype. It is crucial to understand the difficulty of the fitness landscape to be searched. It is a 'Needle-in-a-haystack' (see figure 1) because there is only one correct solution out of a massive search space. This means there is no gradient towards the optimal solution. In such a problem like this, a Genetic Algorithm alone has no information to search for solutions and is analogous to random search in this landscape.

There are three types of individual agents with different fitness measures. The genotype-phenotype mapping is one-to-one and at birth, each individual has its phenotype string identical to its genotype string. If an agent is born with any '0' in its genotype (hence its phenotype at birth), its fitness is set to 1 — the lowest possible fitness value (note, in this model the higher the fitness value, the better the individual); conversely, if its initial genotype is comprised of all ones, its fitness is set to the maximal value of 20. The other type of individual with its genotype comprising of only '1' and '?' (we call it *potential individual*) will be involved in life-time learning. The fitness of that a potential individual agent x_i is calculated as:

$$f(x_i) = 1 + \frac{19(1000 - n)}{1000} \tag{1}$$

in which n is the number of trials required to find the correct combination of alleles - the all-one string. It can be inferred from the fitness function that the more trials an agent needs, the lower the fitness it will get.

The allowance for lifetime learning actually shows its advantage over evolution alone. In Figure 1, we can observe that learning helps smooth the landscape by creating two 'shoulders‘ in proximity of the 'needle'. It can be understood that not only the exact correct configuration of alleles but also combinations which have correct alleles in part and the remaining plastic (the question mark) alleles have an average fitness higher than 0.

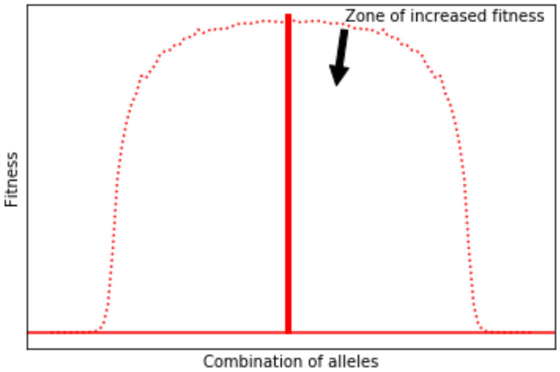

Fig. 1: The fitness landscape by Hinton and Nowlan, with and without learning. Without learning, the landscape is flat, with a thin spike indicating the only correct solution to be found (thick vertical line). When learning is enabled, the landscape has a shoulder around the spike consisting of the allele combinations that have some correct values in part with the remaining values unspecified (learnable or plastic).

C. Other works

The model developed by Hinton and Nowlan, though simple, is interesting, opening up the trend followed by a number of research papers investigating the interaction between learning and evolution. Following the framework of Hinton and Nowlan, there have been a number of other papers studying the Baldwin effect in the NK-fitness landscape, which was developed by Stuart Kauffman [12] to model 'tunably rugged' fitness landscapes. Problems within that kind of landscape are shown to fall in NP-completeness category [12]. Several notable studies of the Baldwin effect in the NK-model include works by Giles Mayley [13], [14], and some others [15], [16]. Their results, again, demonstrated that the Baldwin effect does occur and the allowance for lifetime learning, in the form of individual learning, helps evolutionary search overcome the difficulty of a rugged fitness landscape.

There has also been a large body of research in agent-based simulation studying this topic. Most use the evolutionary connectionist approach to test if neural network learning facilitates an evolutionary process, often represented by a genetic algorithm. Some notable papers on this line of research can be named as [5], [6], especially several works by Stefano Nolfi, Domenico Parisi and their colleagues on Evolutionary Robotics [7], [8], [9], [10], to name but a few. All of these papers attempted to confirm the existence of the Baldwin effect, by showing how learning interacts with evolution making the system perform better than with evolution alone.

However, most research papers on this topic only study the relationship between individual learning and the evolutionary process. We think that it is worth investigating more about social learning [17] and the Baldwin effect, how social learning benefits the evolutionary process, what the Baldwin effect looks like through the course of social learning. A brief literature on social learning and our algorithm with social learning in combination with evolution are presented in the following sections.

D. Social Learning

Social learning arises in various fields of research, including evolutionary biology, psychology, anthropology, behavioral ecology, evolutionary game theory, cognitive science and robotics. Social learning research could be considered a rapidly growing subfield of animal cognition studies: Social learning has been observed in organisms as diverse as primates, birds, fruit flies, and especially humans [18]. Although the use of social learning is widespread, understanding when and how individuals learn from others is a significant challenge [19].

In general, Social Learning in animals can occur via: local enhancement, opportunity provision, stimulus enhancement and observational conditioning, imitation, and emulation (please refer to [17], [20] for the definition of these mechanisms).

In this study we focus on one of these mechanisms, namely imitation learning. In this instance of learning, the observer directly copies the behavior of the observed animal in order to complete a novel task. Implementing imitation learning requires specifying [17], [21], [18]:

i. When an individual should copy;
ii. From whom they should learn; and
iii. What information should be copied.

The question of *when to copy* covers the decision as to when to seek social information. *Whom to copy* may depend on factors such as the social structure of the population and the ability of the individual to recognize whether other individuals are obtaining higher payoffs. Possibilities include the copying of the most successful individual, copying of kin, or adherence to a social norm by copying the majority. *What to copy* considers which behavior or more specifically what part of that behavior to copy.

E. Information fidelity

Another factor we will investigate in this study is *information fidelity* in social transmission. Information fidelity considers how accurately the information is copied from one individual to another. Information fidelity has been studied to show how it affects the social learning process, resulting in the diversity of information in the population [22]. In Science Magazine 2017 the Harvard Anthropologist Joe Henrich [23] argued that a small change in information fidelity could lead to a big difference in human evolution. We incorporate fidelity of transmission in our experiments in later sections.

III. EXPERIMENTAL SETUP

In this section, we present experimental setups used in our paper. It is often said that evolutionary search finds it hard to search in 'Needle-in-a-haystack' landscape. Furthermore, Hinton and Nowlan ran evolutionary search alone and it was claimed that genetic algorithm fails to search in their landscape [4]. From that, we conduct three experiments with the same parameter settings with the original work by Hinton and Nowlan. The settings are shown in Table I.

TABLE I: Parameter setting

Parameter	Value
Genome length	20
Replacement	Generational
Generations	50
Population size	1000
Selection	Fitness-Proportionate selection
Reproduction	Sexual reproduction
Fitness function	Equation 1
Maximal learning trials	1000
Fidelity	1 (by default)

A. Setup I: The replication of the Hinton and Nowlan model

The first experimental setup we use is the replication of Hinton‘s model, when evolutionary search is combined with asocial learning (denoted EVO+IL). As described in section II-B, asocial learning used here is the local search process by which each ‘?‘ allele will guess its value to be ‘0‘ or ‘1‘ in each learning trial (please refer back to section II-B or to the original paper by Hinton [4] for more details). The evolutionary algorithm used in Hinton‘s model is a genetic algorithm without mutation, crossover is the only genetic operator allowed (See table I).

B. Setup 2: Evolution with social learning

In order to implement social learning, first we propose the imitation procedure, with pseudo-code described in algorithm 1 below. This presents the process by which an individual observer imitates the phenotype of its demonstrator, which is controlled by a $fidelity$ parameter. The imitative process starts by extracting the positions of question marks in the phenotype of the observer. For each question mark position, the observer will decide whether to copy exactly the $trait$ or a mutated version of that $trait$ from the demonstrator based on the parameter $fidelity$. Because there are only two expressed values namely 0 and 1, the mutated version of one value is equal to the subtraction of this value from one (line 12 in algorithm 1). By default, the $fidelity$ is set to 1, that means imitative process will copy exactly the values from the demonstrator to the observer. The imitative process is visualized in Figure 2b.

Algorithm 1 IMITATION

```
 1: function IMITATION($observer, demon, fidelity = 1$)
 2:     $questions = []$ comment: question mark position array
 3:     for position $i \in observer.pheno$ do
 4:         if $i = ?$ then
 5:             $questions.add(i)$
 6:         end if
 7:     end for
 8:     for $i \in questions$ do
 9:         if $rand() < fidelity$ then
10:             $observer.pheno(i) = demon.pheno(i)$
11:         else
12:             $observer.pheno(i) = 1 - demon.pheno(i)$
13:         end if
14:     end for
15: end function
```

Having that imitation procedure in hand, we conduct our own experiments in two different setup.

1) Setup 2.1: Evolution with social learning alone:

Algorithm 2 presents the process in which evolution is combined with only social learning in place of asocial learning as in Hinton‘s model (denoted by EVO+SL). The demonstrator is set to be the best individual in the population in terms of fitness. In this algorithm, we can see that the best individual is found based on its original fitness at birth. Remember that only potential individuals can learn (i.e. individuals containing ‘?’ and without any ‘0’). Every individual in the potential set will imitate the phenotype of the demonstrator based on the Imitation algorithm 1. After social learning, the population operates an evolutionary process as in Hinton‘s model.

Algorithm 2 EVO + SL

```
function EVO+SL($pop, fidelity = 1$)
    comment: Do life-time learning
    $potentials = pop.findPotentials()$
    $demon = pop.best()$ comment: extract the best individual
    for $ind \in potentials$ do
        $Imitation(ind, demon, fidelity)$ comment: do imitation
    end for
    comment: Evolve the population
    Do selection
    Do reproduction
    Do replacement
end function
```

2) Setup 2.2: Evolution with both social learning in combination with asocial learning:

Lastly, we devise Algorithm 3 that includes evolution and social learning in combination with individual learning (denoted by EVO+IL+SL). Similar to setup 2.1 above, social learners here will imitate the most successful individual. The only difference with this setting is that we allow individual learning to search for the correct solution when the best individual is not the correct one (line 4-13 in algorithm 3). It is noted that once the correct solution is found, the algorithm stops the individual learning process, and switches all remaining potential individuals to learn socially from that successful individual.

Algorithm 3 EVO + IL + SL

```
 1: function EVO+IL+SL($pop, fidelity = 1$)
        comment: Do life-time learning
 2:     $potentials = pop.findPotentials()$
 3:     $demon = pop.best()$
 4:     if $demon.isCorrect() == False$ then
 5:         for $ind \in potentials$ do
 6:             $ind.individual\_learning()$
 7:             $potentials.remove(ind)$
 8:             if $ind.isCorrect()$ then
 9:                 $demon = ind$
10:             end if
11:             break
12:         end for
13:     end if
14:     for $ind \in potentials$ do
15:         $Imitation(ind, demon, fidelity)$
16:     end for
        comment: Evolve the population
17:     Do selection
18:     Do reproduction
19:     Do replacement
20: end function
```

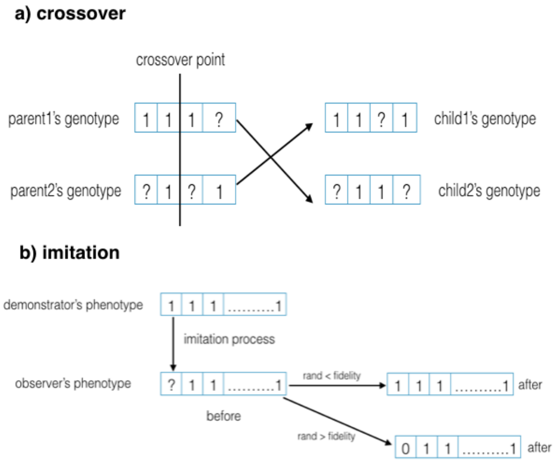

Fig. 2: a) Crossover happens at genotypic level, while b) imitation learning occurs at phenotypic level, only changing the phenotype of the observer/learner.

From this pseudo-code, three imitative strategies mentioned in section II-D are specified:

i. When to copy: Copy when the successful individual is found in the population
ii. From whom to copy: Copy the most successful individual
iii. What to copy: Copy the corresponding value from the demonstrator‘s phenotype.

Please note that, unlike the so-called memetic algorithm [24] learning here only happens at the phenotypic level, what an individual learns does not change its genotype. The recombination operators work on the genotypic level, so children may inherit question marks from their parents. Please see Figure 2 for the visualization of how imitation and recombination work at different levels.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. *Results*

Hinton and his colleague calculated the proportion of loci occupied by ‘0s’, ‘1s’, ‘?s’ as the measurement of how well the simulation performs. According to the experimental setting, the percentage of ‘0’ and ‘1’ are the same at 0.25, while the proportion of ‘?’ is 0.5, in the initial generation. The motivation for this measurement is to see if evolution in combination with learning can drop the number of ‘$0s$’ towards zero, and increase the number of the adaptive allele ‘1’ to be as large as possible. Thus, the Baldwin effect shows up when the incorrect allele ‘0’ disappears and the adaptive allele increases its proportion over time, indicating that the population is more adaptive as time goes by. We adopt the same measurement in our simulation, plus two more criteria. The first is the average fitness of the population. The second is the average number of learning attempts — the total number of learning trials performed by individual learners in the population.

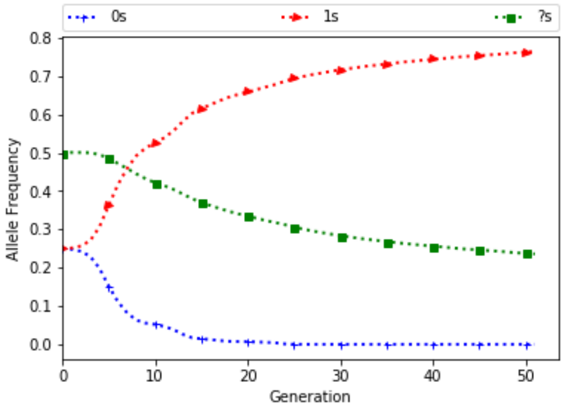

Fig. 3: Replication of the result by Hinton and Nowlan (EVO+IL)

We carry out 30 independent runs and get the average of all measurements above. The results are presented and discussed in the following sections.

1) Setup 1: EVO+IL:

From Figure 3, we can see that the trend of the graph is pretty similar to what was observed in the original simulation of Hinton and Nowlan. By similar, we mean that we observe the same behavior of the graph as a whole, not the behavior of any single line. To be more specific, the proportion of ‘0’ drops to zero after around 24 generations, the percentage of ‘1’ increases while the number of question mark decreases until the end of the simulation. Therefore, we still reach the same conclusion as in the Hinton and Nowlan simulation that learning facilitates evolutionary process in searching a needle-in-a-haystack landscape. Furthermore, it is shown that the Baldwin-like effect occurs due to the fact that from time to time, more alleles of ‘1s’ dominate the population. That means, with more ‘1s’ and less ‘0s’, learning the correct behavior (correct configuration of all-one) costs less time after generations.

Obtaining the same behavior as stated in Hinton‘s simulation makes us believe that the results obtained from our own algorithm with social learning is reliably comparable.

2) Setup 2.1: EVO+SL:

It is shown in figure 4 that without individual learning, social learning fails to guide evolution in a ‘needle-in-a-haystack’ landscape. The Baldwin effect does not show up in this case. Figure 4 shows that frequency of all three alleles keeps relatively constant. The algorithm fails to decrease the proportion of ‘$0s$’ to zero compared to the success of the original result by Hinton and Nowlan.

3) Setup 2.2: EVO+IL+SL:

When running our experiment setup 2.2 (Algorithm 3) we observe that the proportion of ‘0’ drops to 0 quickly. We put the result obtained from our algorithm in the same graph with that obtained from Hinton‘s simulation for comparison. All

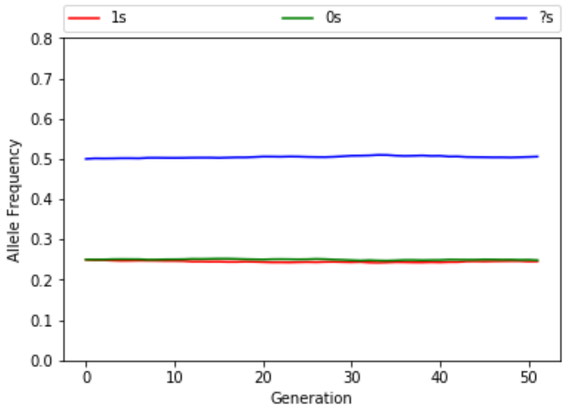

Fig. 4: Frequency of alleles, Evo + SL

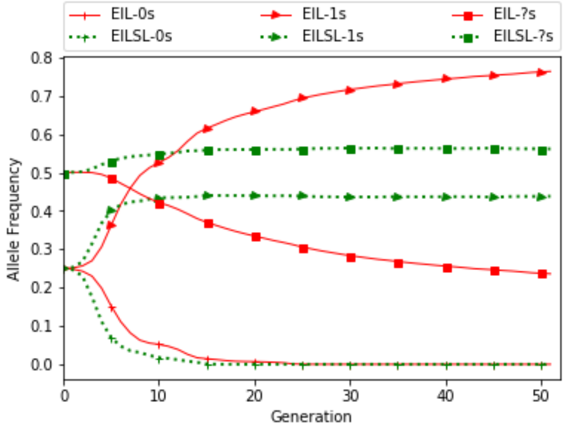

Fig. 5: Frequency of alleles (Fidelity = 1.0 by default)

the lines produced by our algorithm are plotted in green, while those produced by our replication of Hinton‘s simulation are in red. EIL in each figure points to the results by Hinton‘s simulation, using individual learning only. $EILSL$ refers to our results, meaning our algorithm 3.

It is demonstrated in Figure 5 that the frequency of ‘$0s$’ quickly goes down to zero at about generation 15 and remains steady until the end of the simulation. This decrease of ‘$0s$’ in our simulation occurs earlier than in Hinton‘s. The frequency of ‘$1s$’ goes up to reach about 0.43 at the sixth generation, and levels off, though having some small fluctuation. The same trend can be seen in the proportion of allele ‘?‘ - the plastic allele goes up a little to reach about 0.58, and almost remains constant from generation 6.

B. Analyses of our models: EVO+SL and EVO+IL+SL

In EVO+SL, we have observed that social learning alone fails to guide evolutionary search a ‘needle-in-a-haystack’ landscape. An explanation for this is that the performance of social learners depends on how good their demonstrators are. Evolutionary search itself fails to find the correct solution. All individuals have the lowest fitness of 1, and social learners have no better individuals to learn from in hopes of increasing their fitness. Thus, the social learning process does not change the fitness of each social learner and all individuals keep their fitness unchanged, standing at the value of 1 after going through their lifetime learning. The plot of average fitness of the population in Figure 6 (the blue line) shows a flat-line standing at 1, the same as with evolutionary search alone. In this type of hard landscape, social learning fails to guide evolution.

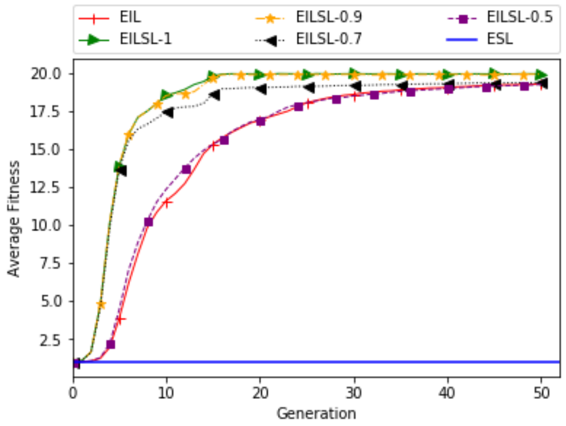

Fig. 6: Comparative average fitness. The solid red for individual learning alone. The green, orange, black, purple lines for our algorithm 3 with fidelity = 1, 0.9, 0.7, 0.5, respectively. The solid blue line for evolution with social learning only.

Conversely, results obtained from EVO+IL+SL show that social learning in combination with asocial learning guides the evolutionary search, even more quickly than does asocial learning alone (Figure 5). It can also be seen in Figure 6 that EVO+IL+SL helps increase the average fitness more quickly and reaches higher value than that in the simulation by Hinton and Nowlan. Furthermore, the learning attempts performed in EVO+IL+SL is much less than that in EVO+IL, as shown in figure 7.

One notable point here is the percentage of ‘$?s$’ is still higher than than that of ‘1’ in EVO+IL+SL. This behavior is different from what we observed in EVO+IL model. The explanation for the observation here is that once the correct solution is found, the imitation procedure will copy exactly the correct value ‘1’ from the solution to potential learners but with much less learning attempts because the learners just need to copy and do not need to find the correct value themselves. This is because from the fitness function in equation 1 that a lower learning cost results in a higher fitness for the learner. Therefore, the average fitness of the population in our Evo+IL+SL is higher than that in EVO+IL. That also indicates that having more plastic alleles, specifying the ability to learn socially, is more adaptive in the future, hence the dominance of ‘$?s$‘. This corresponds to the interpretation of the Baldwin effect as we have described in section II.

We argue that the default $fidelity = 1$ makes it much easier for social learners to copy the correct solution with the much less cost. What is obtained here indicates that the information

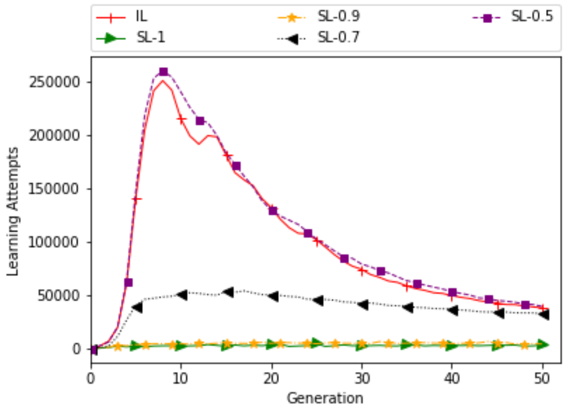

Fig. 7: Comparative learning attempts. The solid red for individual learning alone. The green, orange, black, purple lines for our algorithm 3 with fidelity = 1, 0.9, 0.7, 0.5, respectively.

fidelity could have an influence on the effect of social learning on evolution. We also run EVO+IL+S with different levels of fidelity in order to test this argument. We chose to test on three other levels of fidelity at 0.9, 0.7, and 0.5.

One interesting thing is that when $fidelity = 0.5$ the imitation process as shown in Algorithm 1 performs pretty much the same as a random guessing. This is because a plasticity '?' now, on average, has 50 percent of being correct as '1', or incorrect as '0'. Thus, it is highly expected that the behavior of social learning when $fidelity = 0.5$ is quite similar to that of asocial learning alone as in Hinton's simulation. Experimental results are analyzed in the next section.

C. EVO+IL+SL with fidelity = 0.9, 0.7, 0.5

By looking at Figures 8, 9, and 10, we see that EVO+IL+SL can lead to the decrease of $0s$ in a faster pace than EVO+IL in the two first cases. The proportion of '0' in EVO+IL+SL always drops to zero more quickly than that in EVO+IL. Moreover, it can be observed that the higher the fidelity, the higher the dominance of plastic allele (question mark '?'), the less the amount of '1', the less the average fitness (see figure 6, and vice versa.

Particularly, when $fidelity = 0.5$, all the measurements, from allele frequency through average fitness to learning attempts, EVO+IL+SL performs similarly to EVO+IL. It is indicated that the results obtained are not only as expected, but also consistent with what we have argued so far.

An explanation for this is, again, the understanding of the Baldwin effect through learning (social or asocial). When the fidelity is high, a potential individual learning by imitation spends less learning effort than it does by trial-and-error. This leads to the fact that having more plastic alleles shows advantages in terms of fitness, resulting in higher average fitness. The selection process will favor that kind of plastic allele over others. When the fidelity decreases, an observer has more chances of not copying correct values from the demonstrator. This means some plasticity '?' results in higher chance of being incorrect (having the value of 0), leading to the situation that having more plasticity '?' means having more chance of being incorrect. Furthermore, that also means that each plastic value in this case requires more learning effort to find the correct value of 1. Thus, having fewer number of '?' reduces the learning cost. Again, the selection process will favor a correct individual with less learning cost, the allele '?' will be less favored when the fidelity is lower.

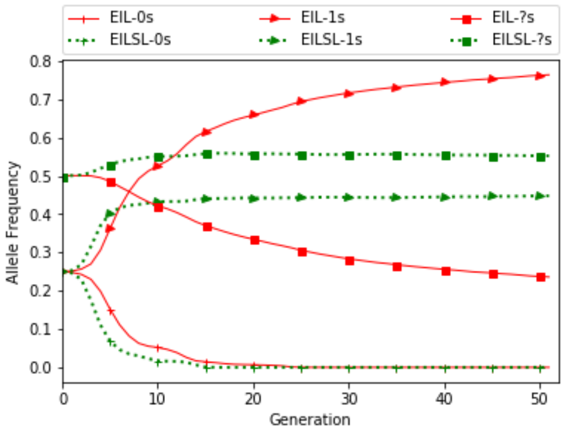

Fig. 8: Frequency of alleles (Fidelity = 0.9)

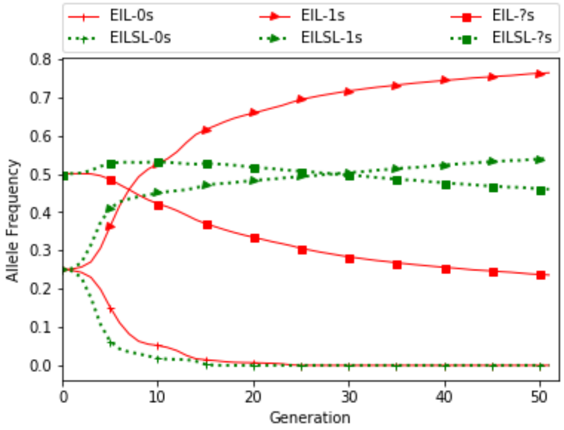

Fig. 9: Frequency of Alleles (Fidelity = 0.7)

From all analyses above, we can conclude that information fidelity plays an important role in the effectiveness of social learning, at least for the problem of interest in this paper.

V. CONCLUSION & FUTURE WORKS

We have set out to understand the role social learning may have on the evolutionary process, leading to the Baldwin effect. Experimental results have shown that social learning alone fails to guide evolution to search a 'Needle-in-a-haystack' landscape. When coupled with asocial learning, however, social learning can reduce the number of *maladaptive* alleles at a faster pace than with asocial learning alone. Our EVO+IL+SL has been shown to increase the average fitness

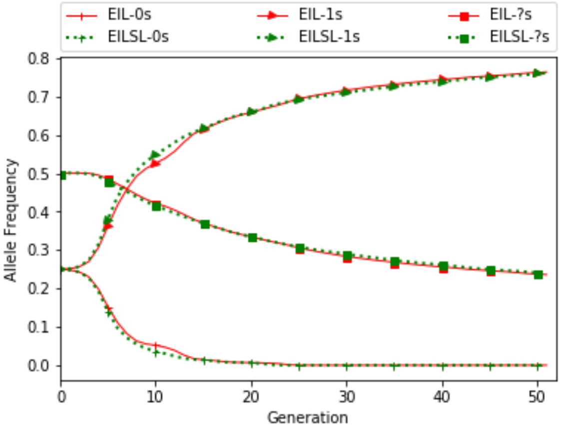

Fig. 10: Frequency of Alleles (Fidelity = 0.5)

of the population with much less learning attempts compared to evolution with asocial learning alone.

We have also shown that information fidelity plays an important role in the performance of social learning by a series of experimental results. It has been suggested in our paper that social and asocial learning should be combined in a strategic manner to improve the population. Social learning should be scrutinized in future research to consolidate the conclusion about the real effect of social learning on evolution.

We have admitted ourselves that the experimental setting used in this paper is the same in Hinton‘s simulation [4] for the purpose of fair comparison. We propose to carry out more experiments with different parameter settings. In the future, we will setup our algorithm when evolution is installed with both asexual and sexual reproduction to see the effect of social learning. We will run the simulation for more generations to observe the effect on allele frequency.

Social learning strategy is a promising area of research to elaborate on with this work. A number of combination between ‘when’, ‘who’, and ‘what’ strategies are worth studying in order to find the best learning strategy for the problem of interest.

Another agenda for future research can be investigating more social learning strategies with the Baldwin effect on different landscapes like NK-landscape [14], [13], with more ruggedness in the search space. Agent-based modeling and robotics are also suggested for future research on how social learning strategies guide evolution in these tasks. Several successes in these fields with asocial learning [9], [7] motivate more research on social learning. We also propose to run the simulation on different dynamic environments to see the interrelationship between evolution, individual and social learning.

ACKNOWLEDGMENTS

This research is based upon works supported by the Science Foundation Ireland under Grant No. 13/IA/1850.

REFERENCES

[1] J. M. Baldwin, “A new factor in evolution,” *The American Naturalist*, vol. 30, no. 354, pp. 441–451, 1896.

[2] C. H. Waddington, “Canalization of development and the inheritance of acquired characters,” *Nature*, vol. 150, no. 3811, pp. 563–565, nov 1942.

[3] C. L. Morgan, “On modification and variation,” *Science*, vol. 4, no. 99, pp. 733–740, nov 1896.

[4] G. E. Hinton and S. J. Nowlan, “How learning can guide evolution,” *Complex Systems*, vol. 1, pp. 495–502, 1987.

[5] D. Ackley and M. Littman, “Interactions between learning and evolution,” in *Artificial Life II, SFI Studies in the Sciences of Complexity*, C. G. Langton, C. Taylor, C. D. Farmer, and R. S., Eds. Reading, MA, USA: Addison-Wesley, 1992, vol. X, pp. 487–509.

[6] J. R. Watson, N. Geard, and J. Wiles, “Stability and task complexity: A neural network model of evolution and learning,” in *IN*. MIT Press, 2002, pp. 153–156.

[7] S. Nolfi, “How learning and evolution interact: The case of a learning task which differs from the evolutionary task,” *Adaptive Behavior*, vol. 7, no. 2, pp. 231–236, 1999.

[8] S. Nolfi and D. Floreano, “Learning and evolution,” *Autonomous Robots*, vol. 7, no. 1, pp. 89–113, Jul 1999.

[9] S. Nolfi, D. Parisi, and J. L. Elman, “Learning and evolution in neural networks,” *Adaptive Behavior*, vol. 3, no. 1, pp. 5–28, 1994.

[10] D. Parisi, D. Parisi, S. Nolfi, and S. Nolfi, “The influence of learning on evolution,” in *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison Wesley, 1994, pp. 419–428.

[11] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press, 1975, second edition, 1992.

[12] S. A. Kauffman and E. D. Weinberger, “The NK model of rugged fitness landscapes and its application to maturation of the immune response,” *Journal of Theoretical Biology*, vol. 141, no. 2, pp. 211–245, nov 1989.

[13] G. Mayley, “Landscapes, learning costs, and genetic assimilation,” *Evolutionary Computation*, vol. 4, no. 3, pp. 213–234, sep 1996.

[14] ——, “Guiding or hiding: Explorations into the effects of learning on the rate of evolution.” in *In Proceedings of the Fourth European Conference on Artificial Life*. MIT Press, 1997, pp. 135–144.

[15] R. Suzuki and T. Arita, “Repeated occurrences of the baldwin effect can guide evolution on rugged fitness landscapes,” in *2007 IEEE Symposium on Artificial Life*. IEEE, apr 2007.

[16] L. Bull, “On the baldwin effect,” *Artif. Life*, vol. 5, no. 3, pp. 241–246, Jun. 1999.

[17] W. Hoppitt and K. N. Laland, *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press, 2013.

[18] L. Rendell, L. Fogarty, W. J. Hoppitt, T. J. Morgan, M. M. Webster, and K. N. Laland, “Cognitive culture: theoretical and empirical insights into social learning strategies,” *Trends in Cognitive Sciences*, vol. 15, no. 2, pp. 68–76, feb 2011.

[19] N. Le, M. O’Neill, D. Fagan, and A. Brabazon, “Social grammatical evolution with imitation learning for real-valued function estimation,” in *2017 IEEE Congress on Evolutionary Computation (CEC)*, J. A. Lozano, Ed. Donostia, San Sebastian, Spain: IEEE, 5-8 June 2017, pp. 1572–1578.

[20] C. M. Heyes, “Social learning in animals: Categories and mechanisms,” *Biological Reviews*, vol. 69, no. 2, pp. 207–231, may 1994.

[21] L. K, “Social learning strategies.” *Learning and Behavior*, vol. 32, pp. 4–14, 2004.

[22] N. Claidière and D. Sperber, “The natural selection of fidelity in social learning,” *Communicative & Integrative Biology*, vol. 3, no. 4, pp. 350–351, jul 2010.

[23] J. Henrich, “High fidelity,” *Science*, vol. 356, no. 6340, pp. 810–810, may 2017.

[24] F. Neri, C. Cotta, and P. Moscato, Eds., *Handbook of Memetic Algorithms*. Springer Berlin Heidelberg, 2012.