

How Learning Strategies Can Promote an Evolving Population in Dynamic Environments

Nam Le*, Michael O'Neill†, Anthony Brabazon‡

Natural Computing Research and Applications Group

University College Dublin

Dublin, Ireland

Email: *nam.lehai@ucdconnect.ie, {†m.oneill, ‡anthony.brabazon}@ucd.ie

Abstract—Learning as a form of adaptation has been shown to benefit the evolutionary process through the Baldwin Effect, promoting the adaptivity of an evolving population. Learning generally can be classified into two types: *asocial learning*, e.g. trial-and-error; and *social learning*, e.g. imitation learning. Recent research has shown that a learning strategy (or learning rule) – a combination of social and asocial learning in a strategic way – can be more adaptive than both types of learning functioning alone, and that that finding should be scrutinised in different environmental variabilities. This paper investigates the adaptive advantage of social learning when combined with asocial learning under a number of environmental dynamics. An instance of dynamic optimisation problems is proposed to test the theory. Experimental results show that a combination of social and asocial learning has a significant impact on guiding evolutionary search in the dynamic landscapes. Indications for future work on dynamic optimisation are also presented.

Index Terms—Baldwin Effect, Social Learning, Dynamic Environment, Phenotypic Plasticity, Cultural Evolution

I. INTRODUCTION

Evolution and learning are two different forms of adaptation in which the organism can adapt their behaviour to environmental problems. The second form occurs when individuals exhibit plasticity in response to environmental conditions that may strengthen their survival. Individuals seek a behaviour that increases fitness. Therefore, it is plausible and rational for the individual to have some learning capabilities to prepare for the uncertain future, some sort of prediction or plastic abilities in different environments. In nature, when the environment changes too quickly that the slower evolutionary process cannot equip enough information for the population to survive, having the ability to learn during the lifetime is necessary to keep pace with the changing environment.

There exists an intriguing idea, called the **Baldwin Effect**, saying how learning influences the evolutionary process, enhancing the adaptivity and performance of the population. Hinton and Nowlan (henceforth H&N) presented a classic paper around 1987 [1] to demonstrate an instance of the Baldwin Effect in a computer simulation. Their initial success motivated a number of further studies [2], [3], [4], to name but a few. It is pretty surprising that the influence of learning on evolution has been rarely studied and employed in the field of Evolutionary Computation (EC), despite the fact that there exists not a small number of dynamic problems in the literature [5]. This might be due to the reason that some people in EC

tend to go *little too far* to regard evolution as learning based on the belief that evolutionary algorithms can be used to solve machine learning problems [?]. One recent paper amongst others on this topic showed that learning by trial-and-error can enhance the performance of an evolutionary algorithm when solving dynamic optimisation problems [6].

More interestingly, learning can be classified into two forms. *Asocial* (or individual) learning (IL) – learning by oneself through direct interaction with the environment, e.g. trial-and-error, and *social* learning (SL) – learning from others, e.g. imitation – are two alternative ways for an individual agent to acquire information from the environment at the phenotypic level. SL has been observed in organisms as diverse as primates, birds, fruit flies, and especially humans [7]. Social learning is a form of *information-parasitism* – can only learn from information produced by others. Research on social learning has shown that we should pay a serious attention to answer the question of whether an agent should rely on social or asocial learning in different environmental dynamics to make the population more adaptive [8]. Recent studies in computer simulations indicated that a learning strategy – a combination of social and asocial learning in a specific manner – can be more adaptive than asocial learning alone [9], [4].

The main aim of this paper is to investigate the effect learning strategies might have on the evolutionary process when dealing with rapid changing environments. We extend previous work on evolutionary optimisation [6], combining evolution with learning strategies to test how the combination of social and asocial learning performs. In the remainder of this paper, we briefly present research on learning and evolution. Social learning and related concepts are briefly introduced. We in turn describe the experiments we use in this paper. Results are analysed and discussed, then the conclusion and some future directions are proposed.

II. BACKGROUND

A. Evolution and Learning

In 1987, the British Cognitive Scientist Geoffrey Hinton and his colleague Kevin Nowlan at CMU presented a classic paper [1] to demonstrate an instance of the Baldwin effect in a computer simulation. Hinton and Nowlan used a Genetic Algorithm to evolve a population in an extreme landscape called *Needle-in-a-haystack*, showing that learning can help

evolution to search for the solution when evolution alone fails to do so. An interesting idea can be extracted from their work is that instead of genetically fixing the genotype, it is wiser to let just a portion of the genotype be genetically fixed, and the other be *plastic* that allows for changes through learning. It is these plastic individuals that promote the evolutionary process to search for the optimal solution, although the H&N landscape is static.

The model developed by Hinton and Nowlan, though simple, is interesting, opening up the trend followed by a number of studies investigating the interaction between learning and evolution. There have been several studies studying the Baldwin effect in the NK-fitness landscape [10]. Some notable papers include the work by Giles Mayley [11], and some others like [12]. Their results, again, demonstrated that the Baldwin Effect does occur, and learning helps evolutionary search overcome the difficulty of a rugged fitness landscape. (Please refer to [4] for more literature review of the Baldwin Effect in the computer)

B. Social Learning

SL has been studied in various disciplines, including Cognitive Biology, Evolutionary Psychology, Behavioral Ecology, Cognitive Science and Robotics. In general, SL covers several mechanisms through which individual organisms learn from others, such as stimulus enhancement, observational conditioning, imitation, and emulation (please refer to [13], [7] 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.

SL, at first glance, seems to be adaptive at a low cost when individual agents can acquire information from others without incurring the cost of trial-and-error learning. Thus, it is plausible to think that SL will result in more effective learning outcomes. Contrary to this belief, it has been found that agents should not learn socially all the time [14], [15], [9]. It is argued that individual learners produce new information about the environment, though at a cost. Social learners avoid this cost by copying the existing behaviors of others, but do not generate new information. Therefore, it is highly likely that social learners will copy outdated information when the environment changes, reducing the average fitness of the population.

Several theoretical models have been proposed to investigate how to use SL effectively [16], [17], [14]. It is said that social learning should be combined with individual learning in a strategic way in order to have an adaptive advantage. Social learning strategies consist of rules specifying the way an individual relies on social learning by answering three questions as follows:

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

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 recognise 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.

In addition to the *Who* question, the transmission from demonstrators to observers are classified into three types ([18]). The first is *vertical transmission* – transmission from parents to their children. The second is *oblique transmission* in which cultural traits will be passed to an individual from another individual from the previous generation but differs from its parent. The last is *horizontal transmission* in which an observer learns from a demonstrator in its current generation. In the scope of this paper, we only use *oblique transmission* in our experiments.

Evolutionary Algorithms (EAs) have been claimed a potential technique to solve dynamic optimisation problems [5]. EAs are a family of algorithms based on biological evolution metaphor. A number of parameter tweaking and operators have been proposed in the literature to enhance evolutionary search when dealing with dynamic optimisation. Though these engineering techniques show some good results in some way, relying on the evolutionary metaphor to cope with dynamic environments is not the optimal way if the optimal solution changes so fast that it takes many more generations for evolution alone to encode the environmental information back to the gene-like pool. Taking on this line of thought, individual learning was employed to benefit the evolutionary search, resulting in higher performance in a simple dynamic optimisation instance [6].

Recently, a model building on H&N's simulation presented in [4], in which the authors combine evolution, asocial and social learning. It was shown that social learning alone fails to search on Needle-in-a-haystack, but social learning when coupled with individual learning outperforms individual learning alone with respect to average fitness of the population. It has also been indicated in their work that the combination of social and asocial learning can increase the adaptivity of plasticity, that means the higher the level of plasticity, the higher the average fitness of the population. More importantly, social learning has also been shown to be a form of 'information-parasitism' which cannot produce new information about the environment and has to rely on information produced by others [9], [17], [8]. Social learning alone does little thing, yet when coupled with asocial learning by some strategy it has been shown to produce interesting effects, sometimes increasing the performance of an evolving population [9].

In this paper, we extend this research by combining individual and social learning to see if that combination can promote the evolutionary search at a greater scale. In the following section, we describe the experimental design and the algorithm we use to combine evolution, individual and social learning.

III. EXPERIMENTAL DESIGN

A. The Dynamic String Match Problem

The chosen problem domain is the String Match problem, in which we have to match the target string. The String Match domain is not a rare problem, and it can be seen in a wide range of contexts, such as in Immune Systems (both natural and artificial), in Antivirus or Intrusion Detection systems. All these systems need some sort of string matching mechanism to match their dictionaries to an incoming signal to see if that signal is abnormal or not. In the scope of this paper, we propose a simple instance of the String Match problem, in which the target string contains only binary characters (0 or 1). The target string changes over time forming a dynamic problem. For simplicity, we restrict the length of the target string to 20.

Suppose we are evolving a population to solve this string match problem. The fitness of an individual is calculated as the proportion of matched characters between the individual and the target over the number of characters (the length of the target string). Mathematically, the fitness of an individual x_i (with phenotype p_i) is computed as one of the following two ways:

$$f(x_i) = \frac{\text{Number of matched characters}}{\text{Length of the target string}} = 1 - \frac{\text{dist}(p_i, \text{target})}{\text{length}(\text{target})} \quad (1)$$

with $\text{dist}(p_i, \text{target})$ is the hamming distance between the phenotype p_i and the *target*. Based on this fitness function, an individual with higher proportion of matching has a higher fitness; and the higher the fitness value, the better the individual. The optimal individual has the fitness of 1, whereas the worst possible fitness value is 0.

Without loss of generality, suppose we have to match the original string *111...11* (20 ones). The target will change based on two parameters: the *frequency* and *magnitude* of change. The first parameter tells us after many generations the target will move to another point in the landscape, while the latter helps determine the likelihood of change for each element of the target. Assume that at a generation g the target is all-one (20 bits of one), *frequency* = 10 and *magnitude* = 0.1 (10%). This informs us that after 10 generations or at generation $g+10$ the target $t = 111...1$ (20 bits of 1) is likely to be changed. The magnitude of 0.1 tells us that there are, on average, $20 \times 10\% = 2$ bits in the target that are likely to be modified. For each bit in the target sequence, a random number is generated and then compared with the *magnitude*: if the random value is less than 0.1, the current bit is mutated to its subtraction from 1 (1 becomes 0, and vice versa). Suppose the new target at generation $g + 10$ is $t_1 = 001...1$ (two first bits are changed).

B. Experimental Setup

1) Experimental Setup I: Evolution + Asocial Learning:

The first setup evolves a population of individuals with learning capabilities. The evolutionary process is implemented as an evolutionary algorithm. In this simulation, we also allow

lifetime learning in the form of asocial learning, in addition to evolutionary algorithm, to update the phenotype of the individual. To allow for lifetime learning we used the same encoding scheme as in [4], [1], [9]: Instead of being fully specified, the genotype now is composed of three alleles ‘0’, ‘1’, and ‘?’. The allele ‘?’ allows for lifetime learning (or plasticity). Each agent will have 100 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. Thus, the behaviour of an individual agent is partly specified by its genetic composition, and partly by what it learns in the course of its lifetime.

At each generation, two individuals are selected from the population as parents to produce one child. The newly born child is mutated based on a small probability, then being added into a new population. This process repeats until the new population is filled up and replaces the old population of parents, i.e. generational replacement without elitism.

When an individual learns, it updates its phenotypic behaviour, and hence its fitness. However, instead of being implemented as a blind random search as in previous work [4], we devise a new learning algorithm as a *hill-climbing* process. The learning algorithm adopted by every individual is presented as Algorithm 1 below.

Algorithm 1 Learning

```

1: function LEARNING(ind)
2:   best_fitness = ind.fitness
3:   best_phenotype = ind.phenotype
4:   while ind.learning_attempts < max_attempts do
5:     ind.learning_attempts + = 1
6:     Flip all question marks to get a new phenotype
7:     best_fitness =
       compute_fitness(ind.phenotype)
8:     if best_fitness > ind.fitness then
9:       ind.fitness = best_fitness
10:      best_phenotype = ind.phenotype
11:    end if
12:  end while
13:  ind.phenotype = best_phenotype
14: end function

```

The above algorithm is relatively self-explanatory. When an individual expresses a new phenotypic behaviour, it checks if the new behaviour is more adaptive than the current before deciding to replace the current phenotype by the new one. This process helps each agent keeps its best behaviour as the current phenotype.

After lifetime learning, the population goes through the evolutionary process as follows. At each generation, two individuals are selected from the population as parents to produce one child. This process repeats until the new population is filled up and replaces the old population of parents, i.e. generational replacement without elitism. No mutation is employed in the current work as in previous studies [9], [4].

2) Experimental Setup II: Evolution + Learning Strategies:

In the second setup, we evolve populations of strategic individuals - individuals that can perform both SL and IL based on a learning rule. In order to implement social learning, first we propose the imitation procedure, with pseudo-code described in algorithm 2 below. This presents the process by which an individual observer imitates the phenotype of its demonstrator. 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 copy exactly the *trait* from the demonstrator.

Algorithm 2 IMITATION

```

1: function IMITATION(observer, demon)
2:   questions = [] comment: question mark array
3:   for position i ∈ observer.pheno do
4:     if i = ? then
5:       questions.add(i)
6:       observer.learning_attempt += 1
7:     end if
8:   end for
9:   for i ∈ questions do
10:    observer.pheno(i) = demon.pheno(i)
11:  end for
12: end function

```

The population now has just one type of individual - strategic individuals that can learn both asocial and socially. We specify the learning strategy for every individual agent as follows: At each generation, an agent first *observes* and learns socially from its demonstrator, then learns asocially on its own until the maximum learning attempt is reached. This scenario can be interpreted as a novice first copies from an expert, then sharpens what it has learned to further its own quality. In the scope of this study, we test on with *oblique transmission* – the individual agent learns from one individual in the previous generation. This also means that there is no social learning at the initial generation. In our current study, all the individuals at each generation have the same demonstrator – the best individual agent from in terms of fitness belonging to the previous generation. After lifetime learning process for each agent, the population goes through selection and reproduction as in EVO+IL.

We run our experiments through 9 different combinations of *frequency* and *magnitude*. It can be understood that the lower the *frequency* value, the faster the target will change; the bigger the value of *magnitude*, the bigger the change of the target. The environment becomes more dynamic or harder to cope with by faster changing and bigger magnitude of change, and vice versa. We also compare the two populations when the environment is static, the target is kept stable over generations. It is interesting that in this sense our problem becomes the

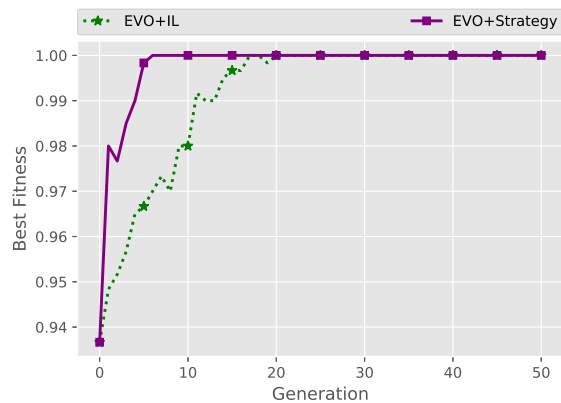


Fig. 1: Best fitness – when the environment is static

canonical *one-max* problem – the trivial problem solvable by Genetic Algorithms.

TABLE I: Parameter setting

Parameter	Value
Original target	111...1 (20 bits of 1s)
Genome length	20
Replacement	Generational
Generations	51
Elitism	No
Population size	100
Selection	Fitness-Proportionate selection
Reproduction	Sexual reproduction
Fitness function	Equation 1
Max learning attempts	50
Frequency	5, 10, 20
Magnitude	0.05, 0.1, 0.2

IV. RESULTS, ANALYSIS AND EXPLANATION

In this section we present the comparison between the two experimental setups, evolving populations with and without learning, in terms of both best fitness and average fitness of the population. All results are averaged over 30 independent runs.

A. When the environment is static

As noted above, our problem becomes *onemax*-like (as the initial target string is comprised of all ones) when the environment is static – EVO+IL was shown to find the correct solution easily in the previous study [6]. It can be seen in Figure 1 and 2 there is no clear difference between the performance of EVO+IL and EVO+Strategy as the evolution progresses.

Despite that the EVO+Strategy population still shows some small advantage in terms of the best fitness since it can find the solution much quicker than the EVO+IL.

One simply short explanation for this behaviour is that when the environment is stable, the problem is easy enough (the *one-max* in this case) so that even evolution alone is sufficient to encode the information of the environment for the population (can find the solution). Moreover, asocial learning was also




Fig. 2: Average fitness – when the environment is static

showed to quicken solution seeking procedure through the hill-climbing process [9]. Therefore, adding social learning before asocial learning may have some advantage in early generations, but does not bring much benefit in terms of finding the fixed target solution as time goes by. In the following section, we provide an explanation in a greater detail which is for dynamic environments but also captures the static scenario.

B. When the environment is changing

The main focus of our current work is how learning strategy performs when the environment is dynamic. First, we look at both the best and the average fitness of the population as measurements of how well each simulation performs. All plots are grouped together, sharing the same labels for x-axis and y-axis as well as the annotation. Each row and column shows the increasing level of difficulty of the problem from left to right and from top to bottom.

A similar trend can be simply seen in Figure 3 and Figure 4 that generally there is a fall in both the best fitness and the average fitness of all environmental circumstances at the generation when the environment changes. This is understandable because when the environment changes and becomes harder, adaptive behaviors from previous generations are no longer adaptive in the current generation, reducing performance of the population. By looking at the behavior of each corresponding line through each row or each column of Figure 4, we can see another shared behavior that the more difficult the environment, the lower the average fitness of the population.

Another shared behaviour can be observed is that the difference between EVO+IL and EVO+Strategy populations becomes bigger, with respect to both best and average fitness, when the environment becomes harder over time in the direction of both frequency and magnitude of change.

Specifically, first we look at the best fitness of the population as a measurement of how well each simulation performs. When the frequency of change is 20, the environment changes at a slow rate between generations. In this case, the two populations behave relatively similar in the first 2 instances of

the magnitude, and show some divergence at the end of the run when the magnitude of change is 0.2 – with %20 of the target being changed. Looking at the minimum magnitude of change at 0.05, there is little difference between the two populations, even when the environment changes after every 2 generations. Despite that over trend, if we look more closely into the best fitness it can be seen that the EVO+Strategy population can reach the best fitness more quickly than the EVO+IL.

The magnitude of change shows more effects on the performance of the population. When the magnitude of change is bigger (0.2, and especially 0.5 in our experiment), the EVO+Strategy population demonstrates a better performance than the EVO+IL population. This is even clearer when the rate of change between generations is faster.

All of these observations here can be explained by the fact that, in any kind of environment, the individual learning process still does its job – upgrading the current phenotype of an agent towards the target step-by-step through a hill-climbing process. Here it is plausible to think that if by some way an agent has a better base for individual learning, the result after learning would be better. This is exactly the story we are talking here and the nature of our learning strategy mentioned above. Each strategic agent, after being born, first copies behaviour from the best individual in the previous generation, then update its behaviour itself through individual learning. This scenario can be interpreted in casual language as follows: A *novice* copies a good enough skill from an expert, then makes its own effort to upgrade and sharpen that skill. Through the imitation process, the learning agent has more chance to have a better phenotype before the individual learning process takes place. Therefore, strategic agents – those that can learn from others and on their own – have more advantage over agents that can only learn individually.

One key thing to be extracted here is that agents in EVO+IL learn individual based on what evolution provides to the agent – this is individual learning from scratch with innate base. More than this, agents in EVO+Strategy learn individually based on what has been learned, or found, individually by agents from previous generations. This creates a form of cumulative learning process – updating information based on what has been found so far in the history.

Therefore, in earlier generations when the evolutionary process plus individual learning from scratch would need more time to find the best solution, copying from the best seems to give some initial advantage. This is because the imitation process provides a better base for individual learning, compared to individual learning from innate information only when the environmental information is not encoded enough. Over generations, however, when the environment slightly changes (or static), even the evolutionary process alone still has time to encode new information, not to mention the fact that individual learning will upgrade the phenotype of each agent after being born and find the solution (as done in previous work [6]). Since the target is slightly changed, all the individual agents in EVO+IL still have a chance to move closer and closer to the target (the solution). If an agent is born close enough to

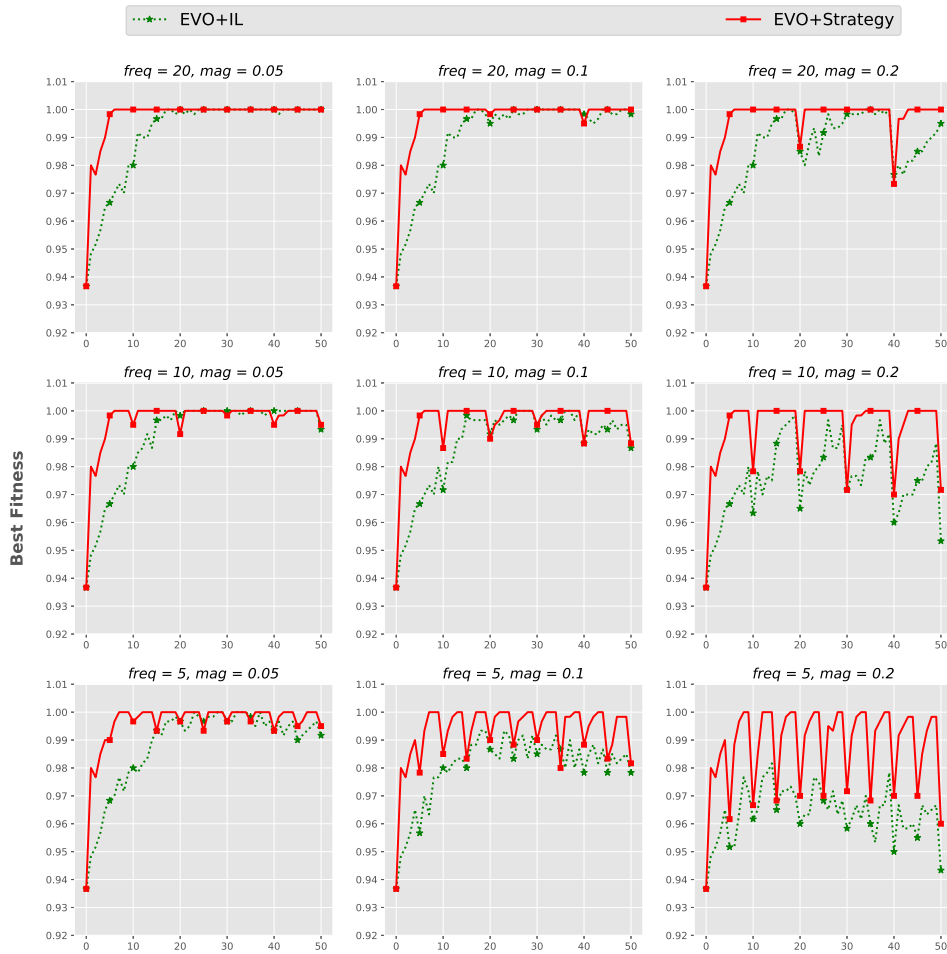


Fig. 3: Comparative Best Fitness

a specific target, and the target is fixed or changed just by a small amount, then the agent still can learn individually and move towards the target easily. There is no huge benefit to copy from any *expert* (the demonstrator). This is why adding social learning before asocial learning does not bring much benefit in these scenarios.

When the environment becomes more difficult to cope with, the target is changed by a sufficiently big amount, the individual learning process still updates the phenotypic information to match the target solution. However, the evolutionary process loses more information over generations. Therefore, learning from others, or more precisely, learning from what others have found previously by individual learning is more advantageous than learning from innate information on one's own. This is why EVO+Strategy shows better performance in terms of the best fitness in all cases. More specifically, the EVO+Strategy still can find the solution whereas the EVO+IL cannot when the environment becomes harder (the cases when the magnitude of change is 0.1 or 0.2, and the frequency of change is 10 or 5).

A similar trend can be observed when comparing the average fitness of the two populations and the explanation is

the same as above. More importantly, the difference between the two evolving populations is clearer in terms of average fitness, even when the environment changes slightly, compared to that of the best fitness.

The explanation presented above shows a more important effect in questioning and elucidating why the average fitness of EVO+Strategy is better than that of EVO+IL. Learning *good behaviours* from others which can provide a better base for individual learning to take place, compared to the based provided by the evolution alone. This cumulative learning process can strengthen the behaviour of the population over time. With the imitation process before individual learning, every individuals in the learning population would be likely to have better *substrates* to learn, updating their behaviour to match with the target, increasing their fitness, hence the average fitness of the whole population. This applies for the whole population, thus the difference in average fitness is a bit clearer between the two evolving populations, with and without imitation learning. Learning from others presents a huge effect on the population as a whole, more than on one single individual agent.

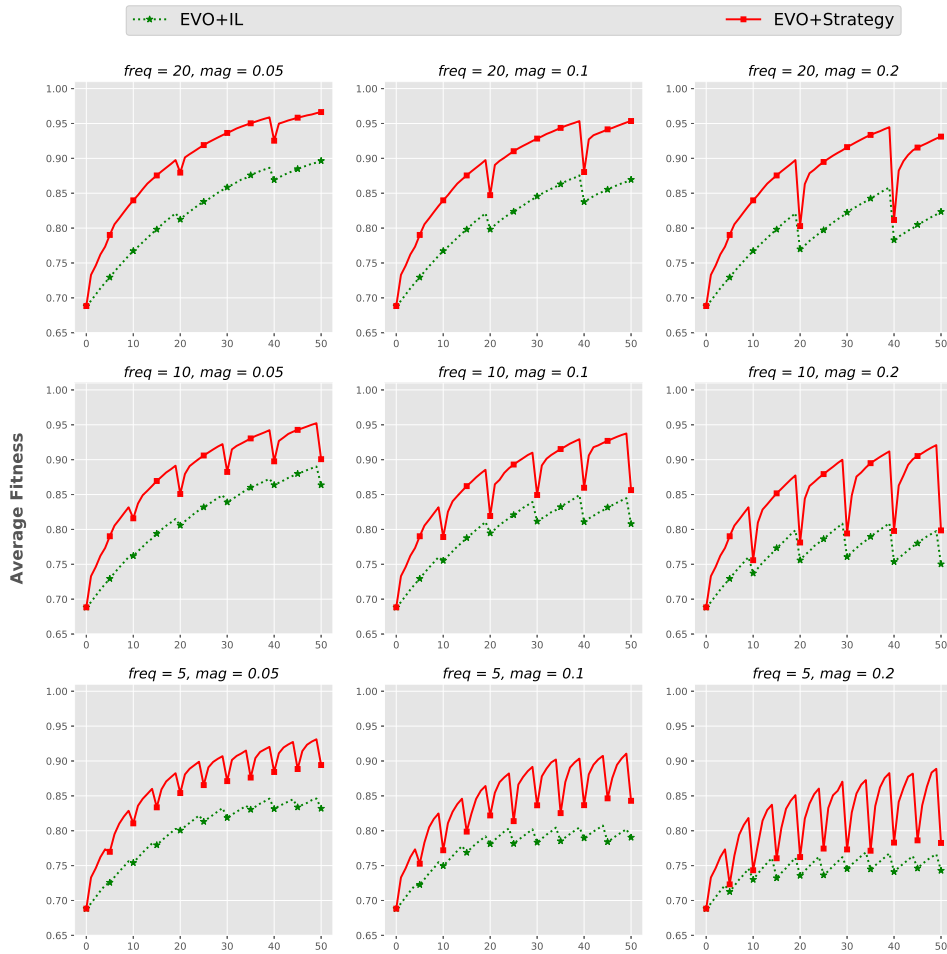


Fig. 4: Comparative Average Fitness

V. CONCLUSION, FURTHER DISCUSSION, AND FUTURE WORK

In this paper, we have set out to understand the role of learning from others in evolving population under different environmental dynamics. By a simple instance of dynamic optimisation, experimental results have empirically shown that when the environment is stable enough or relatively slightly changed, adding social learning does not bring much benefit to individual learning to help the evolving population. When the environment becomes harder, imitation learning has been shown to benefit and facilitate individual learning, promoting the evolving population better than individual learning alone. When the environment is stable and easy enough for asocial learning to cope with, adding social learning shows little advantage. Generally, copying *good behaviour* from others and then learn on one's own is more advantageous than learning individually from scratch.

The dynamic problem used in this paper is simple but the results obtained are indicative for future work. We hypothesise that a combination of social and individual learning can help with evolutionary algorithms when dealing with dynamic optimisation problems. The ability to learn can help individuals to

track and follow the changing target. More than this, the ability to imitate, or learn from others, is expected to provide better foundations for individual learning to take place. Future work will take on this idea of research, and test on with different dynamic optimisation problems.

The same idea can also be tested on different problem domains such as Evolutionary Robotics in which Evolutionary Deep Learning (Neuroevolution) techniques are employed to evolve the brain of agents living in a simulated world. Combining evolution, individual and social learning is expected to give more advantage to the learning agents when dealing with increasing complexity of the simulated world. Neuroevolution to solve machine learning problems will also be considered. Future work will investigate the same idea on these domains of interest.

In terms of philosophical and scientific understanding, what has been presented and explained in this paper is generally fit with the nature of knowledge in biological world, especially in humans. The human agents do not rely totally on instinctive behaviour which was created by the human evolutionary process. Humans do rely on learning to cope with their ever-changing world. More than this, the learning process

in humans is cumulative – we learn from others and then sharpen what we have learned by individual learning. This is the way a number of discoveries have been made in the history of mankind. Knowledge is transmitted between and within generations through a number of channels, both verbal and written media. This is what we call Human Cultural Evolution [18] in which social learning is the key to mediate the knowledge transmission process. By computer simulations, we can observe what would emerge from natural-like phenomena in the computer yet with some over-simplifications. What people do with Artificial Life can be considered another way of doing Natural Philosophy (as the Cognitive Scientist Daniel Dennett once said [19]). By combining evolution, individual and social learning in the computer, we are simply the creation of knowledge through the evolutionary history but at a much simpler level compared to the complex history of humans. In future work, it will be more interesting to study how learning from others, or cultural learning in humans, change the evolutionary pathway of the species. ALife simulations, especially in multi-agent systems, can play a role in the understanding of the interaction between evolution, learning, and culture.

ACKNOWLEDGMENT

This research is funded by the Science Foundation Ireland under Grant No. 13/IA/1850.

REFERENCES

- [1] G. E. Hinton and S. J. Nowlan, “How learning can guide evolution,” *Complex Systems*, vol. 1, pp. 495–502, 1987.
- [2] 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.
- [3] S. Nolfi, D. Parisi, and J. L. Elman, “Learning and evolution in neural networks,” *Adaptive Behavior*, vol. 3, no. 1, pp. 5–28, 1994.
- [4] N. Le, M. O’Neill, and A. Brabazon, “The baldwin effect reconsidered through the prism of social learning,” in *IEEE Congress on Evolutionary Computation, CEC 2018*. Rio, Brasil: IEEE Press, 8-13 Jul forthcoming.
- [5] J. Branke, *Evolutionary Optimization in Dynamic Environments*. Springer US, 2002.
- [6] N. Le, A. Brabazon, and M. O’Neill, “How the “baldwin effect” can guide evolution in dynamic environments,” in *Theory and Practice of Natural Computing*. Springer International Publishing, 2018, pp. 164–175.
- [7] C. M. Heyes, “Social learning in animals: Categories and mechanisms,” *Biological Reviews*, vol. 69, no. 2, pp. 207–231, may 1994.
- [8] 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.
- [9] N. Le, M. O’Neill, and A. Brabazon, “Adaptive advantage of learning strategies: A study through dynamic landscape,” in *Parallel Problem Solving from Nature – PPSN XV*. Springer International Publishing, 2018, pp. 387–398.
- [10] 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.
- [11] G. Mayley, “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.
- [12] L. Bull, “On the baldwin effect,” *Artif. Life*, vol. 5, no. 3, pp. 241–246, Jun. 1999.
- [13] W. Hoppitt and K. N. Laland, *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press, 2013.
- [14] K. N. Laland, “Social learning strategies,” *Learning and Behavior*, vol. 32, pp. 4–14, 2004.
- [15] A. R. Rogers, “Does biology constrain culture?” *American Anthropologist*, vol. 90, no. 4, pp. 819–831, 1988.
- [16] M. W. Feldman, K. Aoki, and J. Kumm, “Individual versus social learning: Evolutionary analysis in a fluctuating environment,” Santa Fe Institute, Working Papers, 1996.
- [17] J. Y. Wakano, K. Aoki, and M. W. Feldman, “Evolution of social learning: a mathematical analysis,” *Theoretical Population Biology*, vol. 66, no. 3, pp. 249 – 258, 2004.
- [18] P. J. R. Robert Boyd, *Culture and the Evolutionary Process*. Chicago: University of Chicago Press., 1985.
- [19] C. G. Langton, *Artificial Life: An Overview*. Cambridge, MA, USA: MIT Press, 1995.