# 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

*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) – which combines social and asocial learning in a strategic way – can be more adaptive than either type of learning functioning alone. However, this general finding leaves open the important question as to how best to combine these forms of learning in differing environmental conditions. This paper investigates this question under a variety of environmental dynamics and also provides some indications for future work.**

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

## I. INTRODUCTION

Evolution and learning are complementary forms of adaptation by which an organism can adapt its behaviour in response to environmental challenges. While evolution can 'hard wire' patterns of behaviour, learning permits a significantly higher degree of behavioural plasticity in response to environmental conditions which in turn can enhance the chances of survival of an organism. In particular, when the external environment is highly dynamic, learning can permit much faster behavioural adaptation than occurs via a multi-generational evolutionary process. Indeed, evolution and lifetime learning are closely intertwined, as a capability for lifetime learning can only arise as a result of an evolutionary process. Perhaps less evidently, the linkage also goes the other way and learning can significantly influence the evolutionary process and therefore enhance the adaptivity of a species over time.

This phenomenon is termed the **Baldwin Effect** and was demonstrated in a classic simulation paper by Hinton and Nowlan (henceforth H&N) [1]. This paper stimulated a number of important follow-on studies including [2], [3], [4]. However, it is noticeable that apart from some recent papers by [5], [6], the influence of learning on evolution has subsequently been little studied in the field of Evolutionary Computation (EC) [7], despite the fact that many problem domains addressed by EC are inherently dynamic [8].

Delving a little deeper into lifetime learning, this category can be subdivided into *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. Each is a plausible way 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 [9]. Social learning can be considered as a form of 'information-parasitism' as an individual can only socially learn from information produced by others. A key open question in the learning literature is how best to combine IL and SL under different environmental dynamics [10]. A rule for combining of each type of learning is termed a 'learning strategy' and such strategies can be more adaptive than either IL or SL alone [5], [11].

The main aim of this paper is to investigate the effect that different forms of learning strategies can have on the evolutionary process when dealing with rapidly 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 contribution, we initially overview some prior research on learning and evolution and some important concepts from social learning are introduced. We then describe our experimental design, provide our results and discuss the same. Finally, some directions for future work are proposed.

## II. BACKGROUND

### A. Evolution and Learning

In 1987, the British Cognitive Scientist Geoffrey Hinton and his colleague Steven J. Nowlan at CMU presented a classic paper [1] which demonstrated an instance of the Baldwin effect using a computer simulation. Hinton and Nowlan used a Genetic Algorithm to evolve a population in a Needle-in-a-haystack landscape showing that learning can help evolution to search for a solution when evolutionary search alone is ineffective. An interesting idea which can be extracted from their work is that instead of genetically fixing the structure of the genotype ab initio, permitting elements of the genotype to remain *plastic*, and therefore to be 'learnable' during the evolutionary process provides a much easier route for the population to eventually uncover determine the solution to the difficult optimisation problem. Crucially, the H&N landscape – although difficult – is static.

The model developed by Hinton and Nowlan, opened up an initial surge of interest in the investigation of the in-

978-1-7281-2153-6/19/$31.00 ©2019 IEEE

teraction between learning and evolution. Exemplars of this work include studies examining the Baldwin effect in an NK-fitness landscape [12], and [13]. Their results, again, demonstrated that the Baldwin Effect does occur, and learning helps evolutionary search overcome the difficulty of a rugged fitness landscape.

### 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 [9], [14], [15] for more detail 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 organism (physical or digital) in order to complete a novel task.

SL, at first glance, seems to be a low-cost, adaptive, mechanism as 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 [16], [17], [11]. It is argued that individual learners produce new information about the environment, though at a cost. In contrast, social learners avoid this cost by copying the existing behaviors of others, but do not themselves generate any new information about the environment. 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 [18], [19], [16]. It is said that social learning should be combined with individual learning in a strategic way in order to produce an adaptive advantage. Social learning strategies consist of rules specifying the way an individual relies on social learning by answering three questions:

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 ([20]). 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 a (non-parent) member of the previous generation. The final mechanism, *horizontal transmission*, occurs when an observer learns from a demonstrator in its current generation. In the scope of this paper, we only use *oblique transmission* in our experiments, leaving the examination of other mechanisms for future work.

Evolutionary Algorithms (EAs) have been widely applied for the purposes of solving dynamic optimisation problems [8]. Prima facie, this would appear to be a reasonable approach as biological evolution is situated in a dynamic environment and has proven itself capable of generating an almost infinite variety of organisms, each inhabiting a dynamic ecological niche. Many mechanisms have been used in developing EA variants for dynamic problem environments with simple approaches implementing varying approaches for the maintenance of populational diversity, or the rapid adaptation of the population when environmental change is detected.

Though these techniques can produce good results in specific problem instances, it does not concord with the multiple layer learning found in most biological organisms. Taking on this line of thought, some studies have combined individual learning with evolutionary search, resulting in higher performance in a simple dynamic optimisation instance [6].

Recently, a model building on H&N's simulation was presented in [5], in which the authors combine evolution, asocial and social learning. It was shown that social learning alone fails when presented with a Needle-in-a-haystack landscape, but social learning when coupled with individual learning successfully outperformed individual learning alone with respect to average fitness of the population. The study also indicated that the combination of social and asocial learning can increase the adaptivity of plasticity. Specifically, the higher the level of plasticity, the higher the average fitness of the population.

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 the objective is to match a target string. The String Match problem bears close parallel with many real-life problem domains, and it can be seen in a wide range of contexts, such as in Immune Systems (both natural and artificial), and 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 determine whether a 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). Crucially, 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 of binary genotypes 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:

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

where $dist(p_i, target)$ is the Hamming distance between the phenotype $p_i$ and the *target*. Based on this fitness function, an individual with higher proportion of matching bits 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*** (i.e., 20 ones). The target will change based on two parameters: the *frequency* and *magnitude* of change. The first parameter tells us after how many generations the target will move to another point in the landscape, while the latter helps determine the likelihood of change for each element (bit) of the target. Assume that at 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 [5], [1], [11]: 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 with a small probability, and is then added into a new population. This process repeats until the new population is completed upon which it replaces the old population of parents, i.e., the process consists of 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 [5], 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

```
function LEARNING($ind$)
    $best\_fitness = ind.fitness$
    $best\_phenotype = ind.phenotype$
    while $ind.learning\_attempts < max\_attempts$ do
        $ind.learning\_attempts += 1$
        Flip all question marks to get a new phenotype
        $best\_fitness = compute\_fitness(ind.phenotype)$
        if $best\_fitness > ind.fitness$ then
            $ind.fitness = best\_fitness$
            $best\_phenotype = ind.phenotype$
        end if
    end while
    $ind.phenotype = best\_phenotype$
end function
```

The above algorithm is relatively self-explanatory. When an individual expresses a new phenotypic behaviour, it checks whether the new behaviour is more adaptive than the current before deciding whether the current phenotype is replaced by the new one. This process helps each agent keep its best behaviour as its 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. No mutation is employed in the current work as in previous studies [11], [5].

#### 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.

The population now has just one type of individual - strategic individuals that can learn both asocial and socially.

**Algorithm 2** IMITATION

```
function IMITATION($observer, demon$)
    $questions = []$ comment: question mark array
    for position $i \in observer.pheno$ do
        if $i = ?$ then
            $questions.add(i)$
            observer.learning_attempt += 1
        end if
    end for
    for $i \in questions$ do
        $observer.pheno(i) = demon.pheno(i)$
    end for
end function
```

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

TABLE I: Parameter setting

| Parameter | Value |
|---|---|
| Original target | 111...1 (20 bits of 1$s$) |
| 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 |

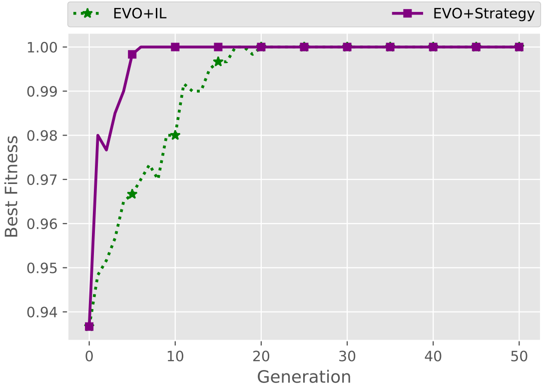

Fig. 1: Best fitness – when the environment is static

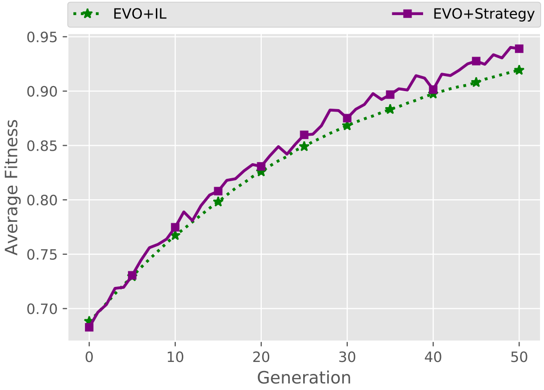

Fig. 2: Average fitness – when the environment is static

## 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 showed to quicken the solution-seeking procedure through the hill-climbing process [11]. 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. Initially, 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-xis 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 expected because as the environment changes and becomes harder, adaptive behaviors from previous generations are no longer adaptive in the current generation, reducing the 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 that 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 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 an 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 cummulative 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 genrations, 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 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 looses 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

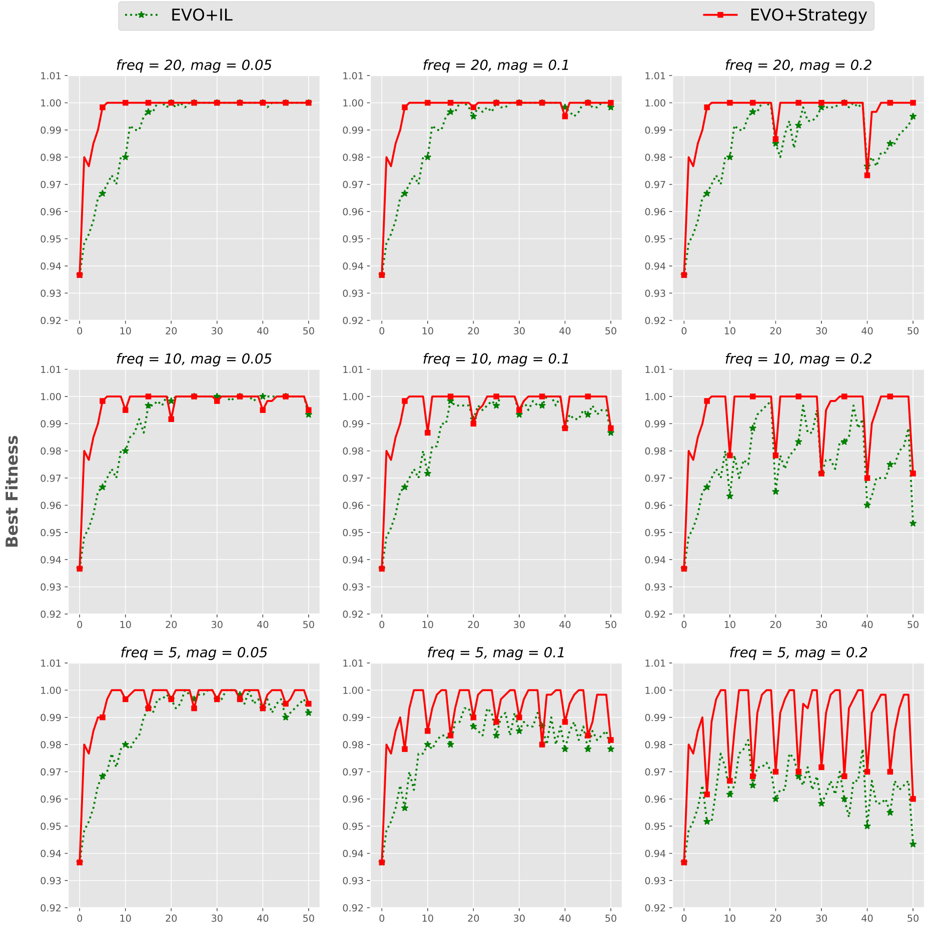

Fig. 3: Comparative Best Fitness

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.

## V. CONCLUSION, FURTHER DISCUSSION, AND FUTURE WORK

In this paper we have set out to better understand the role of learning from others in an evolving population under different environmental dynamics. By employing a simple instance of dynamic optimisation, our experimental results illustrate that when the environment is stable or only slightly changing, adding social learning does not bring much benefit to individual learning in helping an evolving population to adapt. However, when the environment becomes more dynamic (i.e., harder), imitation learning shows a clear benefit and facilitates individual learning, promoting the adaptation of the evolving population better than is the case under individual learning alone.

The dynamic problem used in this paper is concisely described but captures a wide variety of dynamic settings. While an individual study can of course only directly speak

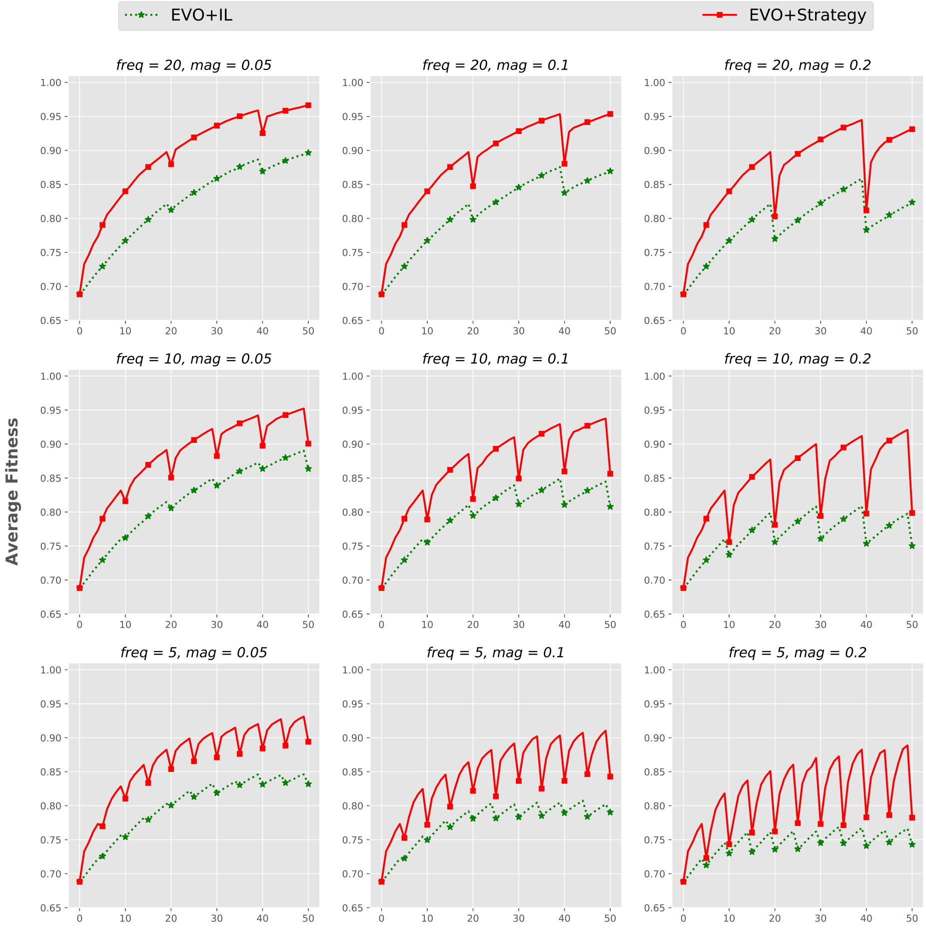

Fig. 4: Comparative Average Fitness

to the problem instances tested, the results in this study are consistent with a more general claim that a combination of social and individual learning can strengthen the performance of evolutionary algorithms when dealing with dynamic optimisation problems. The ability to learn can help individuals to efficiently track and follow the changing target. Indeed, the beneficial effects of social learning extend beyond their ‘first order’ impact, as an ability to usefully imitate, or learn from others, also provides better foundations for future individual learning to subsequently further improve the socially-learnt information. Future work will look at this aspect and also examine different dynamic optimisation problems.

The same ideas can also be tested on different problem domains of Evolutionary Computation, such as Genetic Programming [21] in which the phenotype is represented as a computer program, or Evolutionary Robotics in which Evolutionary Deep Learning (Neuroevolution) [22], [2] techniques are employed to evolve the brain of agents living in a simulated world. As occurs in higher-order animals in the biological world, it is plausible that combining evolution, individual and social learning will assist in the creation of better simulated or indeed embodied learning agents.

With respect to philosophical and scientific understanding, what has been presented and explained in this paper is generally consistent with the nature of knowledge in the biological world, especially in the case of humans. The human agents do not rely totally on instinctive behaviour arising from an evolutionary process but can also learn. More than this, the learning process in humans is cumulative – we learn from others and then sharpen what we have learned by individual learning and all scientific advances rely on what has gone before. Knowledge is transmitted between and within generations through a number of channels, including verbal and written media and via formal educational processes. This is termed Human Cultural Evolution [20] with social learning playing a key role in the knowledge transmission process. Aspects of these processes can be explored via computer simulations in order to gain insights into the process of cultural and technological advance. As noted by the Cognitive Scientist Daniel Dennett [23], ‘Artificial Life can be considered another way of doing Natural Philosophy’.

## 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] N. Le, "Evolving self-taught neural networks: The baldwin effect and the emergence of intelligence," in *10th Symposium on AI & Games, in AISB Convention 2019*, Falmouth, UK, 16-18 Apri 2019.

[3] 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.

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

[5] N. Le, M. ONeill, and A. Brabazon, "The baldwin effect reconsidered through the prism of social learning," in *2018 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, jul 2018. [Online]. Available: https://doi.org/10.1109/cec.2018.8477654

[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] A. Brabazon, M. O'Neill, and S. McGarraghy, *Natural Computing Algorithms*, 1st ed. Springer Publishing Company, Incorporated, 2015.

[8] J. Branke, *Evolutionary Optimization in Dynamic Environments*. Springer US, 2002.

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

[10] 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.

[11] 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.

[12] 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.

[13] N. Le, M. O'Neill, and A. Brabazon, "Evolutionary consequences of learning strategies in a dynamic rugged landscape," in *Proceedings of the Genetic and Evolutionary Computation Conference*, ser. GECCO '19. Prague, Czech Republic: ACM, 2019 forthcoming.

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

[15] A. Brabazon and S. McGarraghy, *Foraging-Inspired Optimisation Algorithms*. Springer International Publishing, 2018.

[16] K. N. Laland, "Social learning strategies." *Learning and Behavior*, vol. 32, pp. 4–14, 2004.

[17] A. R. Rogers, "Does biology constrain culture?" *American Anthropologist*, vol. 90, no. 4, pp. 819–831, 1988.

[18] M. W. Feldman, K. Aoki, and J. Kumm, "Individual versus social learning: Evolutionary analysis in a fluctuating environment," Santa Fe Institute, Working Papers, 1996.

[19] 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.

[20] P. J. R. Robert Boyd, *Culture and the Evolutionary Process*. Chicago: University of Chicago Press., 1985.

[21] N. Le, H. N. Xuan, A. Brabazon, and T. P. Thi, "Complexity measures in genetic programming learning: A brief review," in *2016 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, Jul. 2016. [Online]. Available: https://doi.org/10.1109/cec.2016.7744087

[22] N. Le, M. O'Neill, and A. Brabazon, "The evolution of self-taught neural networks in a multi-agent environment," in *22nd International Conference on the Applications of Evolutionary Computation*. Leipzig, Germany: Springer, 24-26 Apri 2019.

[23] C. G. Langton, *Artificial Life: An Overview*. Cambridge, MA, USA: MIT Press, 1995.