Imitation Tendencies of Local Search Schemes in Baldwinian Evolution

Shu Liu
Dept. of Electrical Engineering,
the Univ. of Tokyo,
Tokyo, Japan
liushu@iba.t.u-tokyo.ac.jp

Hitoshi Iba
Dept. of Electrical Engineering,
the Univ. of Tokyo,
Tokyo, Japan
iba@iba.t.u-tokyo.ac.jp

ABSTRACT

Baldwinian evolution is a type of hybridization of population-based global search and individual local search. The individuals take local refining processes, then in selection benefit from the improved fitness, but do not pass on the refined traits the data in to the offspring. The lost information of the refined phenotype implies that the inheritance encoded in genotypes is not directly benefit traits, but the traits having potential to achieve high fitness through the lifetime interaction with the environment. As the result, it is necessary to study how learning works comparing to the previous generation, in addition to how much it improves on the current population. The children may imitate what their parents performed and catch up with them, or alternatively, explore elsewhere and have no idea of where the parents arrived. In this paper, the trade-off is investigated, and it is revealed that in Baldwinian learning, the capability to follow the parents' footprints benefits. With higher imitation tendency, the evolving population can maintain a greater scale of learning potential, and the search results in better speed and convergence.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*Heuristic methods*

General Terms

Algorithms

Keywords

memetic algorithms, Baldwinian evolution, local search, learning potential, imitation tendency

1. INTRODUCTION

Memetic algorithms are a type of hybrid algorithms combining population-based evolution with local refinements. They can be divided into two categories: Lamarckian and Baldwinian. Lamarckian evolution assumes learned traits are fixed back to the genotypes and passed on to the offspring. On the other hand, in Baldwinian evolution selection is based on improved fitness, but the traits refined are not known to the offspring.

To apply memetic algorithms in real world problems, it is necessary the knowledge of how evolution proceeds and which factors influence the performance. So far the researchers mainly concentrated on Lamarckian evolution for its simpleness and effectiveness. For example there are works on optimizing the parameters such as learning frequency and intensity [16, 13, 1, 7, 11], and some on adapting them through the search [19, 17, 12, 18]. However, the results may not hold in Baldwinian evolution. For example, in Lamarckian search, it is common to choose only a random set from the population to take learning. However in Baldwinian scenario, if an individual is not allowed to take learning, the learning efforts paid by its parents, if any, become meaningless. The parents were selected for their refined traits, but the offspring is to be judged by the inherited initial traits.

There are also a number of studies on the mechanisms of Baldwinian evolution. The guiding effect of Baldwinian learning in the search was first verified by Hinton and Nowlan [8]. In the following years researches were taken on the characteristics in the process of Baldwinian evolution [23, 15, 22, 14], the interaction of learning, evolution and development [4, 5, 6], and the fact that Baldwinian learning can also be hiding [20, 21]. Furthermore there are works trying to combine Baldwinian and Lamarckian evolution [3, 9, 2]. However, as once claimed by Turney [23], Baldwinian evolution is very complex. By now there are still many mysteries in Baldwinian evolution.

In this paper we investigate further into the process of Baldwinian evolution, to find out what has effect on the search performance. Differing from many conventional studies on how learning performs in one generation, our viewpoint is taken on the comparison between the offspring's learning processes and their parents'. This work is inspired by Suzuki [22] who has reported climbing period and the effect of keeping track to the optimum. In our previous work [14], it was revealed the most learning efforts are paid to keep the inherited potential of achieving good fitness. Now we study what happens if the offspring follow or not follow the learning trajectories or directions of their parents', as it's already known they start from similar positions.

An image of Lamarckian and Baldwinian learning is shown in Fig.1. In Lamarckian learning, the children C_1 and C_2

GECCO'11, July 12–16, 2011, Dublin, Ireland.

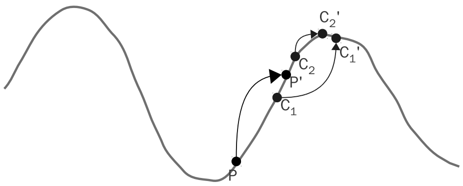

(a) Lamarckian learning

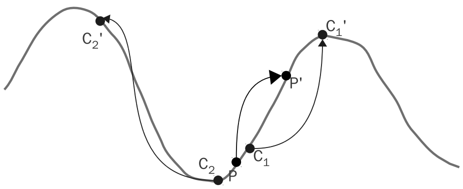

(b) Baldwinian learning

Figure 1: Illustration of two types of inheritance.

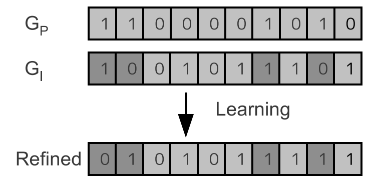

Figure 2: The genotype of an individual.

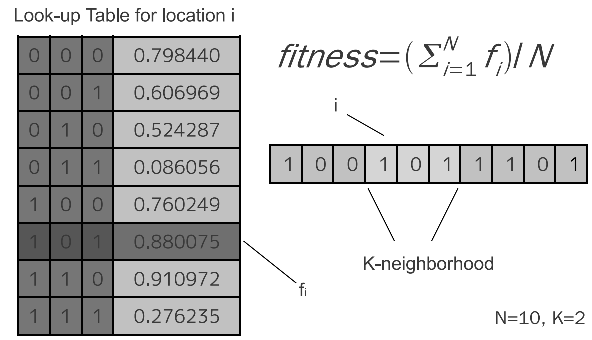

Figure 3: Illustration of partial fitness in NK model with N=10, K=2.

fully inherited the refined traits of the parent in P', and what they achieve in their own learning processes ($C_1 \to C_1'$ and $C_2 \to C_2'$) are mostly beyond the learning destination of the parent ($P \to P'$). In Baldwinian learning, the children are reproduced close to P, the start point of the parent, where there is a distance from the learning destination P'. The children have to again discover routes to improve fitness by themselves, whether the same to or different from the parent's. The inheritance is no longer direct benefit traits, but the potential to achieve good fitness by learning.

Basically there are two possible patterns to realize the potential. The individual may follow the parent's footprints and go further ($C_1 \to C_1'$), and have most of its learning efforts imitating or repeating what the parent has done once. It is also possible to explore somewhere else and improve the fitness ($C_2 \to C_2'$), which leaves the route of the parent and may be either a better direction or a worse one.

Even with the same search capabilities on the same individuals or populations, different learning schemes have different tendencies of imitating or exploring, thus lead to different search performances. The refined traits are not passed on directly, however, individuals in different generations are taking the same learning scheme, and the common procedure may provide some clues of the parents' routes. With more such clues the children, which are reproduced similar to the parents' initial shapes, are expected to have more chance to step on or get close to the footprints of the parents, and vice versa. With different levels of the balance, different individuals survive in selection, the population is tailored into different shapes through evolving, and the search performs differently.

Considering the trade-off between imitating and wandering elsewhere, an idea associated is the topic of exploitation versus exploration. However, the nature of Baldwinian evolution makes the story in another way. Without leaving the attained learning route to the offspring realizable, neither exploitation nor exploration could make sense. Even when an individual imitates the learning of its parent(s), it is just keeping track to the refining destination once attained, and makes a little further improvement. From this point of view it can be expected that with a lower imitation tendency, the achievements of ancestors will be more difficult to repeat, and the search may result in worse speed and convergence.

In this paper, we examined the effects of learning scheme imitation tendencies in Baldwinian evolution. It is revealed more imitation enables the population to keep more learning potential, and may improve search speed and convergence. In the second section we introduce the NK model as our benchmark, and present the three learning schemes to compare, which have similar local search performances on the same populations. In section 3, experiments examining learning potential scales and search performances are presented, on three different landscapes separately. Section 4 includes discussions on the results of the experiments, also some looking ahead topics. Finally the conclusion of the paper is presented.

2. MODEL

2.1 NK Model with Plasticity

In our examination, NK model is used as the benchmark. The model is simple, and allows us to design the needed learning schemes easily. It is a broadly used artificial model to simulate general discrete optimization problems proposed by Kauffman [10]. Some researchers use this model to examine general rules in evolution [15].

The genotype of an individual includes two N-bit strings, as shown in Fig.2. G_I records the start point of learning, in other words the initial traits. G_P points out which bits in G_I are allowed to be modified in lifetime learning, as known as the plasticity. In the learning process, the plastic bits (the positions having 1 in G_P) of G_I are modified to try new strings and pursue higher fitness. The survival depends on refined fitness, while the children receive G_I and G_P.

To calculate the fitness of G_I or its refinements, a series of look-up tables are needed, which determine the landscape.

Table 1: Experiment Parameters

N	20
K	0/5/15
Population Size	500
Crossover Rate	0.50
Mutation Rate	0.05

As shown in Fig.3, for any location i, calculate a partial fitness f_i. The value is determined by the bits in the according K-neighborhood, those from the i-th bit to the $(i+K)$-th bit, and is found in the prepared i-th look-up table, indexed by the 2^{K+1} possible 0/1 combinations. After extracting all the N partial fitness values from the N tables, the fitness of the whole string is calculated as $f = \frac{1}{N}\sum_{i=1}^{N} f_i$.

From these explanations it is clear that N determines the size of the search space, and K adjusts the epistasis and complexity level. The experiment parameters are included in Table 1. Furthermore, conventional 2-tournament selection, two-point crossover and one-bit mutation are employed. Experiments are taken on three landscapes with K values 0, 5 and 15, to include various levels of epistasis and complexity. For each landscape the N look-up tables are generated randomly once, and on this landscape various tests and their iterations take place.

2.2 Local Search Schemes

With the settings mentioned above, we compared three local search schemes with similar search capabilities but different imitation tendencies. They are all bit-wise trial climbing schemes: From a learning individual with current phenotype string s, a trial changes one of the plastic bits to form a new string s'. Evaluate s', compare to the fitness of s, and keep the better string from the two as the result of this trial and the base of next. Starting from string G_I, iterative trials are taken till the given number of trials is reached, which we call it the budget. The final fitness is considered the fitness of the learning individual.

The only difference of the three schemes is how to determine the bit to change in each trial. Note the plastic bit positions in an individual (1s' positions in G_P) as $i_1 < i_2 < ... < i_m, 1 \leq i_j \leq N$, the schemes are as follows:

1) Change the plastic bits one by one in sequential trials, according to the position sequence in the genotype: $i_1, i_2, ..., i_m, i_1, i_2, ...$

2) Change the plastic bits one by one, according to a random permutation generated before the learning of the individual. Note the permutation of 1 to m as $p(1), p(2), ..., p(m)$, the changed bits are: $i_{p(1)}, i_{p(2)}, ..., i_{p(m)}, i_{p(1)}, i_{p(2)}, ...$

3) In each trial choose a random bit from all the plastic bits. A possible series of changed bits is: $i_4, i_2, i_1, i_2, ...$

Here we use the budget limit as the termination condition of learning, for it is easy to control and measure, and convenient to compare between generations. It is also common to finish learning when it makes little or no further improvement. Both are reasonable but have difficulties in determine the best time to stop. There is always risk of insufficient local search or redundant computational cost. Furthermore, in Baldwinian evolution with the same learning scheme, an individual pays almost the same effort as the parent to realize the received potential and catch up. Any online adaptation of learning scheme or learning intensity needs to be smooth enough, or they will encounter additional inefficiency. That is the reason we have tests on multiple budgets and use fixed budgets in each run.

It should also be noted that in our experiment, the learning cost is not employed in fitness function, as we are only interested in the search phase of Baldwinian evolution. In this phase it is claimed that enhancement of local search dominates [8, 3], and the punishment of learning cost is considered and set trivial before genetic assimilation. Our further consideration is, the cost punishment in fitness function substantially changes the target of evolution. It is no longer "to find individuals achieving good fitness through learning", but "to find individuals achieving good fitness and cost little through learning". Here comes two objectives, and the trade-off coefficient between is determined arbitrarily in conventional works. It needs more careful investigations.

3. EXPERIMENT

In the first sub-section we confirm that the three learning schemes bring similar improvements on the same populations. Then in other three sub-sections, corresponding to the three landscapes, we examine how different learning tendencies shape the population, and how they influence the search performance.

3.1 Search Capabilities of Learning Schemes

The three learning schemes are similar in operations. In this experiment, they are confirmed quite similar in search capabilities, when applied to the same population evolved without bias of either scheme. We evolve a population with no-learning evolution till the average fitness exceeds 0.70, then apply the three learning schemes on the population with various budgets, and check how much the average fitness is improved since previous generation.

As shown in Fig.4, X axis shows the budget, which means how many trials are taken in the learning. Y axis shows the average fitness improvement through the learning processes. The three data lines distinguish the schemes applied. Each data bar shows the result of 50 sub-sections, results on different landscapes are presented separately. random iteration runs with the scheme and budget, including the average and the error range. Considering these both aspects, on all landscapes and with various budgets, learning scheme 1 and 2 perform almost the same on the same population. Learning scheme 3 is relatively poor in search comparing to the other two, but not too much. Comparing different landscapes, as the epistasis and complexity grow, the improvement of learning decreases, and differences of the schemes become smaller.

However, the tendencies of imitation are quite different. Considering starting from similar or even the same initial genotypes, with some learning schemes the destination of learning are likely to be similar, but others not. In scheme 1, the order of the changed bits is fixed to the positions in genotypes, thus some priority information can be passed on to the offspring: a bit close to the beginning of the string is always taken in trials in the early stages of learning. The same genotype (G_I+G_P) performs always the same in this learning, and similar individuals are likely to reach similar or the same destination. In scheme 2, the order is generated randomly every time the learning begins. The same genotype performs differently as permutation varies, and similar individuals are less likely to arrive at the same place. In

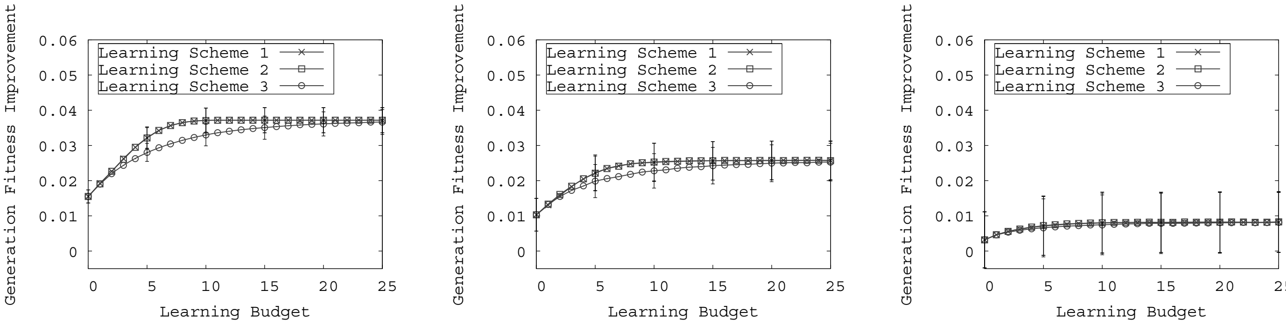

(a) K=0 (b) K=5 (c) K=15

Figure 4: Compare learning schemes on the same population.

Table 2: Fitness improvement of learning schemes with fixed learning budget, on landscape K=0

	Budget	5	10	20
vs before Learning	Scheme 1	0.046891 (±0.003227)	0.068673 (±0.003760)	0.069930 (±0.003838)
	Scheme 2	0.040463 (±0.003097)	0.068623 (±0.004322)	0.070024 (±0.005128)
	Scheme 3	0.024141 (±0.002321)	0.044181 (±0.003300)	0.061050 (±0.004799)
vs Previous Generation	No Learning	0.015492 (±0.001858)		
	Scheme 1	0.015488 (±0.001897)	0.017195 (±0.001450)	0.017373 (±0.001460)
	Scheme 2	0.014235 (±0.001503)	0.017354 (±0.001412)	0.017863 (±0.001367)
	Scheme 3	0.013780 (±0.001652)	0.014528 (±0.002164)	0.016777 (±0.001695)

Table 3: Search performance of learning schemes, on landscape K=0

	Budget	5	10	20
Converging Fitness	No Learning	0.793938 (±0.000205)		
	Scheme 1	0.794080 (±0.000190)	0.794220 (±0.000163)	0.794399 (±0.000146)
	Scheme 2	0.793963 (±0.000201)	0.794157 (±0.000182)	0.794380 (±0.000171)
	Scheme 3	0.793871 (±0.000245)	0.793908 (±0.000250)	0.793955 (±0.000244)
Ev. Calls to Reach 99%	No Learning	6840 (±410)		
	Scheme 1	35700 (±2443)	56650 (±3191)	105840 (±5547)
	Scheme 2	39120 (±2496)	57530 (±2984)	105210 (±5404)
	Scheme 3	43080 (±2554)	74360 (±5240)	126000 (±7648)

scheme 3, a bit may be chosen again for a trial soon after used once, not waiting for all the plastic bits have been taken once. Even less certainty is passed through generations. The differences are tested in details in the next experiments.

3.2 No Epistasis Landscape

The first landscape has K=0. There is no epistasis and the landscape is very simple. For each position of G_I the bit has two different partial fitness values according to its 0/1 code. Just find the better choice from 0/1 and cover all bits, the search reaches the global optimum.

First we have an experiment measuring learning improvements, to examine how the population is shaped according to the learning schemes. A population is evolved with a fixed learning budget applied to all individuals, till the average fitness exceeds 0.70. Then apply the same learning scheme with various budgets to the population, and record how much the average fitness is improved comparing to the refined fitness of previous generation. It shows the scale of learning potential kept with the learning scheme and budget.

The result is shown in Fig.5 and Table 2. They are averages of 50 runs. In Fig.5, X axis shows the differences between current and previous budgets. For example, the population reaching fitness 0.70 is evolved with fixed budget b_0, and in current generation the budget is changed to b, then the X coordinate is $b - b_0$. Y axis shows the differences of refined fitness between two generations, which can measure how much the search proceeds. Table 2 shows the fitness improvement when current budget reaches the previous, corresponding to the point with $X = 0$ in Fig.5, and also includes results comparing to the fitness before learning. The numbers in brackets with ± signs are deviations of the independent iterations.

On this landscape, with any of the learning schemes, individuals have to take almost the same budget to catch up with their parents, and can hardly go further. However, different imitation tendencies lead to different scales of learning potential kept. As shown in Table 2, comparing to initial fitness of current generation, the fitness increments through learning are different. Considering the deviations, the differences are significant. Scheme 1 and 2 perform similarly. Scheme 3 holds a smaller scale of potential, and improves less through generations. Further statistics of budget=20 shows that scheme 1 and 2 usually achieve greater learning potential than scheme 3 (49/50 and 46/50).

The second test is on the search performance: speed and fitness. To compare the schemes more fairly, for each learning scheme we assign three different budgets 5, 10 and 20, and test separately. In the experiment we record the computational cost measured in evaluation calls, and the average fitness of the population.

The result is shown in Fig.6 and Table 3. The data are averages of 50 runs. In Fig.6, X axis shows the evaluation call numbers, and Y axis is the average fitness. Table 3 shows convergence fitness, the evaluation calls used to reach 99% of the fitness, and their deviations. The numbers in brackets with ± signs are deviations of the independent iterations.

The global optimum fitness of this landscape is 0.794861, found by enumeration. From the table we see with all meth-

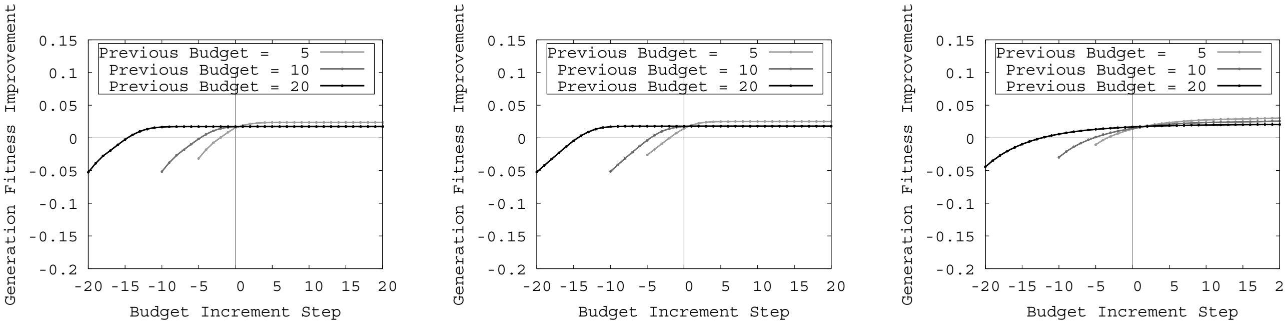

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 5: Learning improvement on K=0 landscape.

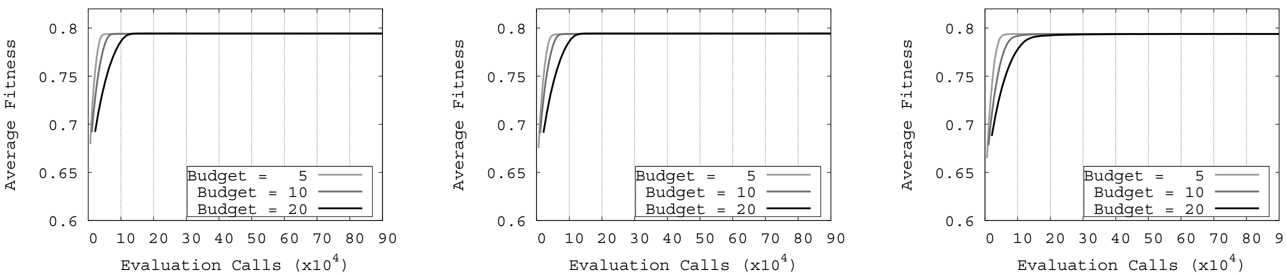

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 6: Search performance on K=0 landscape.

Table 4: Fitness improvement of learning schemes with fixed learning budget, on landscape K=5

	Budget	5	10	20
vs before Learning	Scheme 1	0.102299 (±0.016521)	0.138046 (±0.015239)	0.154631 (±0.015365)
	Scheme 2	0.070663 (±0.020377)	0.120990 (±0.012895)	0.143240 (±0.011780)
	Scheme 3	0.030695 (±0.007967)	0.065227 (±0.011316)	0.110085 (±0.011123)
vs Previous Generation	No Learning	0.010298 (±0.004638)		
	Scheme 1	0.009380 (±0.002965)	0.009794 (±0.002825)	0.009612 (±0.002988)
	Scheme 2	0.009106 (±0.002895)	0.009198 (±0.002565)	0.009011 (±0.002980)
	Scheme 3	0.009230 (±0.003879)	0.008564 (±0.002428)	0.008872 (±0.002731)

Table 5: Search performance of learning schemes, on landscape K=5

	Budget	5	10	20
Converging Fitness	No Learning	0.787684 (±0.011294)		
	Scheme 1	0.798414 (±0.002960)	0.798978 (±0.003703)	0.799183 (±0.002938)
	Scheme 2	0.794109 (±0.007212)	0.798162 (±0.003513)	0.798776 (±0.002472)
	Scheme 3	0.795399 (±0.006534)	0.795821 (±0.005731)	0.796702 (±0.004016)
Ev. Calls to Reach 99%	No Learning	14800 (±3231)		
	Scheme 1	80220 (±20739)	142450 (±84018)	237930 (±63093)
	Scheme 2	81420 (±17545)	141130 (±32312)	279930 (±107113)
	Scheme 3	88080 (±23013)	182820 (±71180)	295050 (±55733)

ods search converges close to the optimum fast. Considering deviations and differences of averages, Scheme 1 performs quite similar to scheme 2, and prior to scheme 3 not significantly, in both speed and fitness. Further statistics of budget=20 shows that scheme 1 and 2 usually attain better fitness (both 48/50) and speed (49/50 and 48/50).

3.3 Low Epistasis Landscape

This landscape has K=5. It has some epistasis, however still fairly easy. On this landscape, it becomes difficult to imitate parents.

Fig.7 and Table 4 are the results of learning improvement examination. The three schemes achieve similar improvements between generations, and lower than no learning evolution. The differences on the scale of learning potential are much more distinct on this landscape, according to both average and deviation data. It is apparent that scheme 3 keeps the least potential, and scheme 1 holds more than scheme 2. Further statistics of budget=20 shows that for learning potential scales, scheme 1 is often prior to 2 (37/50), and scheme 2 always higher than 3 (50/50).

The examination of search performance is shown in Fig.8 and Table 5. The global optimum fitness of the landscape is 0.803160, found by enumeration. The evolution without learning converges at 0.787684, not close to the optimum, though much faster than Baldwinian evolutions. All the Baldwinian searches converge close to the optimum, with no big difference in fitness. For convergence speed, scheme 1 is relatively faster than scheme 3, and scheme 2 is between

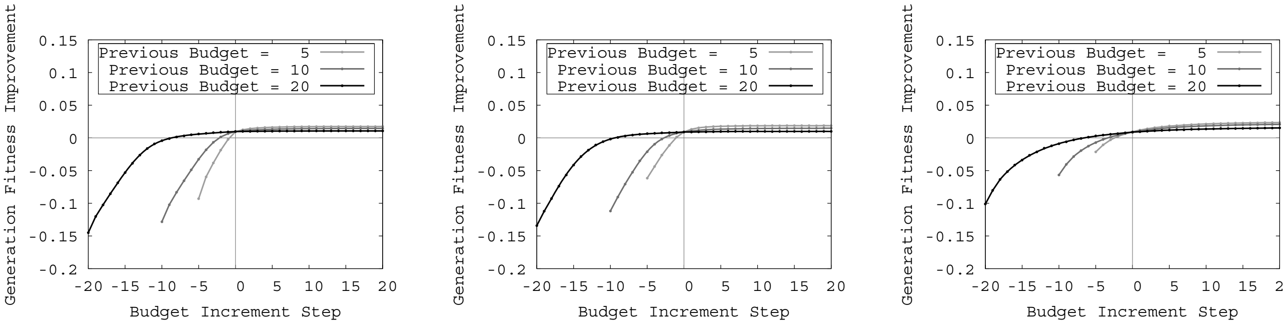

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 7: Learning improvement on K=5 landscape.

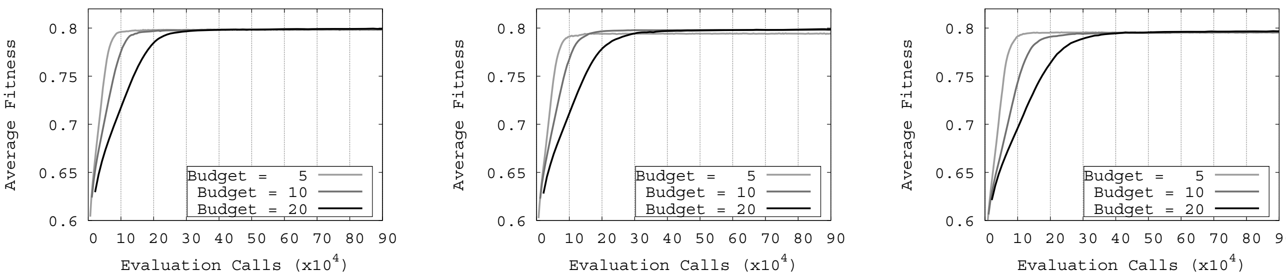

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 8: Search performance on K=5 landscape.

the two. Further statistics shows that for fitness and speed, scheme 1 is slightly better than 2 (30/50 and 31/50), and scheme 2 is better than 3 (37/50 and 33/50). Comparing scheme 1 and 3, scheme 1 usually wins (44/50 and 40/50).

3.4 High Epistasis Landscape

This landscape has K=15. When one bit is changed, 16 partial fitness values will change. The epistasis is high, thus crossover or mutation frequently breaks the selected learning routes. In this situation, greater budgets achieve little further learning capability, and the realization of inherited potential becomes very difficult. As the result, increment of budget not always brings greater learning potential. Sometimes the imitation tendency even decreases, such as scheme 1 shown in Fig.9 and Table 6. Without determined trial order, scheme 2 and 3 suffer less from this phenomenon, and further budget let them keep a greater scale of learning potential. It implies that the level of imitation tendency is not always positively correlated with budget or learning capability.

Learning improvements of three schemes are included in Fig.9 and Table 6. The learning potential scales vary apparently, considering averages and deviations. Though scheme 1 suffers from the disadvantage of broken learning routes, it keeps more potential than scheme 2, and scheme 3 is still the least. Further statistics of budget=20 shows that for learning potential scales, scheme 1 is always prior to 2 (50/50), and scheme 2 usually higher than 3 (46/50).

Fig.10 and Table 7 show the performance of the search. The global optimum fitness of the landscape is 0.801496, found by enumeration. All methods fail to converge close to the optimum, and the lowest is no learning evolution, though still the fastest. Different learning schemes bring similar convergence times, but different fitness levels. Scheme 1 has

Table 6: Fitness improvement of learning schemes Learningwith fixed learning budget, on landscape K=15

	Budget	5	10	20
vs before Learning	Scheme 1	0.160329 (±0.055076)	0.139534 (±0.018094)	0.137678 (±0.007222)
	Scheme 2	0.063902 (±0.041059)	0.099236 (±0.012206)	0.114280 (±0.005558)
	Scheme 3	0.028494 (±0.021138)	0.077244 (±0.011407)	0.103274 (±0.005545)
vs Previous Generation	No Learning	0.003169 (±0.007954)		
	Scheme 1	0.007173 (±0.008703)	0.002722 (±0.004311)	0.001201 (±0.002764)
	Scheme 2	0.005126 (±0.006681)	0.003498 (±0.004571)	0.000733 (±0.002355)
	Scheme 3	0.005552 (±0.006894)	0.003281 (±0.005199)	0.001404 (±0.002793)

the highest fitness, and scheme 3 the lowest. For convergence fitness scheme 1 is often better than 2 (37/50), and scheme 2 is better than 3 (32/50). Comparing scheme 1 and 3, scheme 1 usually wins (48/50).

4. DISCUSSION

In the experiments, three learning schemes are examined. They have similar search capabilities, if performed on the same populations. However, the learning behavior of individuals varies with the learning scheme, when comparing to their ancestors learned and evolved with that scheme. As observed in the results, the different imitation tendencies make the populations and searches distinct. When the scheme brings more chances to repeat parents' achievements,

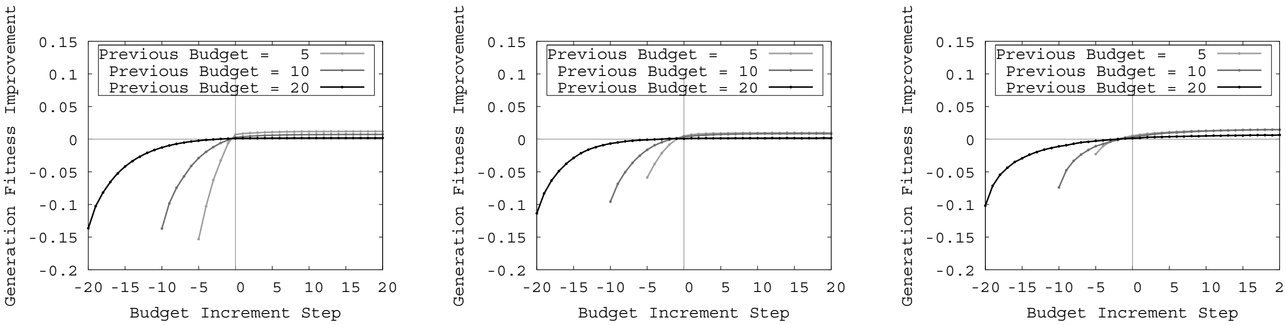

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 9: Learning improvement on K=15 landscape.

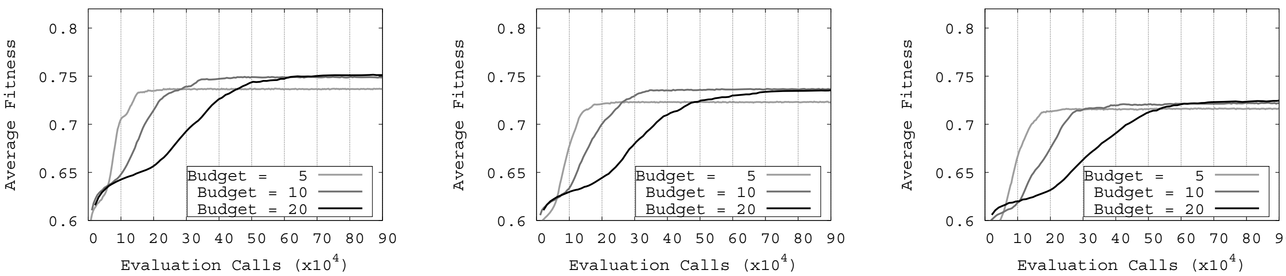

(a) Learning Scheme 1 (b) Learning Scheme 2 (c) Learning Scheme 3

Figure 10: Search performance on K=15 landscape.

Table 7: Search performance of learning schemes, on landscape K=15

	Budget	5	10	20
Converging Fitness	No Learning	0.706966 (±0.020083)		
	Scheme 1	0.736789 (±0.019725)	0.748969 (±0.020085)	0.751067 (±0.019072)
	Scheme 2	0.723160 (±0.018781)	0.736710 (±0.019307)	0.735063 (±0.018210)
	Scheme 3	0.716105 (±0.023496)	0.721917 (±0.018423)	0.724330 (±0.013875)
Ev. Calls to Reach 99%	No Learning	38970 (±126332)		
	Scheme 1	123480 (±35223)	246840 (±73754)	434910 (±100401)
	Scheme 2	127260 (±36937)	252780 (±111422)	458850 (±137939)
	Scheme 3	140100 (±66803)	269500 (±115074)	459480 (±119729)

the offspring follow the footprints more effectively, thus allow the population to keep greater scales of learning potential during the evolution. This also means clues of local refinement directions are passed through generations with less loss, and enables more efficient search, both in speed and fitness. The converging time and fitness data could support this explanation.

The task to compare learning processes between generations is probably unique in Baldwinian evolution. In Lamarckian evolution, children's learning efforts are always beyond their parents', and previous successes can not guarantee the way ahead. In Baldwinian evolution, the realization of the inheritance is no longer a deterministic translation from genotype to phenotype, but with in addition a learning process containing uncertainties. Individuals start from similar situations as their parents', repeat the history or happen to explore the unknown.

The story can be extended to another understanding. As a necessary procedure to achieve final fitness used in selection, the learning scheme is also an entity the individuals have to fit, similar to the environment. If an individual always goes in the same direction in learning, its offspring will not lose track to the direction. On the other hand if every individual has multiple random choices, who can survive through generations must be those surrounded by good fitness areas, as they have to achieve good fitness in all the directions they may go. To be strict, in either situation the true goal of evolution is not "the solution with best fitness". It is replaced by "the base point to achieve good fitness through the given learning scheme". The request of achieving high fitness in all possible directions thus may limit the search.

What the result means in application is: to improve the performance of Baldwinian search, it is better to make sure that individuals can usually follow their parents' learning footsteps. It is something before the trade-off of exploration and exploitation, and looks strange to some extent. If the individuals should follow the parents, why not Lamarckism? It is different. Lamarckian evolution completely fixes the learning products. The offspring do not have chances to go elsewhere for exploration, to correct possible mistakes the ancestors happened to make. It is necessary to design Baldwinism-Lamarckism combinations to avoid disadvantages of the both.

This paper contributes in knowledge accumulation of Baldwinian evolution, yet the discussion on learning behavior between generations is not enough. This time we mainly concentrated on the different learning schemes. However, it is also related to the variation part (crossover, mutation, etc.), as children are not copies of parents. The imitation, if happens, starts with the traits varies from parents with the given operators. This is even further related to the landscape and calls for more studies.

Looking ahead even further, there is the problem of what to encode and inherit in Baldwinian evolution. Basically the genotype for initial phenotype, as the start point of learning, has to be encoded. Then plasticity of genes is usually mentioned in conventional studies. However it is still simple and vague. For example in this paper, the three schemes take the same genotype encoding and similar meanings of the codes, but applying different learning schemes, the evolution goes different directions. Inheritances can have, and are containing further information concerning the learning scheme. Similar ideas are already mentioned in some recent works such as Downing's [6]. In solving real world computational problems, surely we can try to encode further information into the genotypes and have the contents selected, to improve the performance.

5. CONCLUSION

This paper presents an study on how individuals realize learning potential inherited from parents in Baldwinian evolution. It is revealed that evolution pushes the population to fit the applied learning scheme. Learning schemes with higher tendencies of imitating parents enable the population to better receive inheritances, thus lead to greater potential scales, and furthermore advantages in converging speed and final fitness. The results help to further understand Baldwinian evolution, and can be a reference for attempts of applications. It is worth studying more precise descriptions of the issue, and how it is related to other aspects of Baldwinian evolution.

6. REFERENCES

[1] N. K. Bambha, S. S. Bhattacharyya, J. Teich, and E. Zitzler. Systematic integration of parameterized local search in evolutionary algorithm. *Evolutionary Computation, IEEE Transactions on*, 8(2):137–155, 2004.

[2] C. G. Carrier. Unifying learning with evolution through baldwinian evolution and lamarckism: A case study. In *in Proc. Symp. Computational Intelligence and Learning (CoIL-2000)*, pages 36–41, 2000.

[3] K. L. Downing. Deisgning neutral networks using genetic algorithms with graph generation system. *Genetic Programming and Evolvable Machines*, 2(3):259–288, 2001.

[4] K. L. Downing. Development and the baldwin effect. *Artificial Life*, 10(1):39–63, 2004.

[5] K. L. Downing. Computational explorations of the baldwin effect. In *Proceedings of the First Norwegian Artificial Intelligence Symposium*, pages 41–52, 2009.

[6] K. L. Downing. The baldwin effect in developing neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 555–562, 2010.

[7] D. E. Goldberg and S. Voessner. Optimizing global-local search hybrids. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 220–228, 1999.

[8] G. Hinton and S. Nowlan. How learning can guide evolution. *Complex Systems*, 1:495–502, 1987.

[9] H. Ishibuchi, S. Kaige, and K. Narukawa. Comparison between lamarckian and baldwinian repair on multiobjective 0/1 knapsack problems. In *in Evolutionary Multi-Criterion Optimization, Third International Conference, EMO 2005, ser. Lecture Notes in Computer Science*, pages 370–385. Springer, 2005.

[10] S. Kauffman. *The Origin of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.

[11] N. Krasnogor and J. Smith. A memetic algorithm with self-adaptive local search: Tsp as a case study. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, pages 987–994, 2000.

[12] N. Krasnogor and J. Smith. A tutorial for competent memetic algorithms: Model, taxonomy, and design issues. *Evolutionary Computation, IEEE Transactions on*, 9(5):474–488, 2005.

[13] K. W. C. Ku, M. W. Mak, and W. C. Siu. A study of the lamarckian evolution of recurrent neural networks. *Evolutionary Computation, IEEE Transactions on*, 4(1):31–42, 2000.

[14] S. Liu and H. Iba. A study on the computational efficiency of baldwinian evolution. In *Proceedings of World Congress on nature and Biologically Inspired Computing 2010 (NaBIC2010)*, pages 474–479, 2010.

[15] G. Mayley. Landscapes, learning costs and genetic assimilation. *Evolutionary Computation*, 4:213–234, 1996.

[16] Q. Nguyen, Y. Ong, and N. Krasnogor. A study on the design issues of memetic algorithm. In *Evolutionary Computation, 2007. CEC 2007. IEEE Congress on*, pages 2390–2397, 2007.

[17] Q. H. Nguyen, Y.-S. Ong, and M. H. Lim. A probabilistic memetic framework. *Evolutionary Computation, IEEE Transactions on*, 13(3):604–623, 2009.

[18] N. Noman and H. Iba. Accelerating differential evolution using an adaptive local search. *Evolutionary Computation, IEEE Transactions on*.

[19] Y.-S. Ong, M.-H. Lim, N. Zhu, and K.-W. Wong. Classification of adaptive memetic algorithms: a comparative study. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 36(1):141–152, 2006.

[20] I. Paenke, T. J. Kawecki, and B. Sendhoff. On the influence of lifetime learning on selection pressure. In *Proceedings of Artificial Life X*, pages 500–506, 2006.

[21] I. Paenke, T. J. Kawecki, and B. Sendhoff. The influence of learning on evolution: A mathematical framework. *Artificial Life*, 15(2):227–245, 2009.

[22] R. Suzuki and T. Arita. The dynamic changes in roles of learning through the baldwin effect. *Artificial Life*, 13(1):31–43, 2007.

[23] P. Turney. Myths and legends of the baldwin effect. In *Proceedings of the Workshop on Evolutionary Computing and Machine Learning at the 13th International Conference on Machine Learning (ICML-96)*, pages 135–142, 1996.