

Heterogeneous complexification strategies robustly outperform homogeneous strategies for incremental evolution

Adam Stanton¹ and Alastair Channon¹

¹Research Institute for the Environment, Physical Sciences and Applied Mathematics
Keele University, ST5 5BG, UK
{a.stanton, a.d.channon}@keele.ac.uk

Abstract

The evolution of naturalistic, embodied agents and behaviours has been a long-standing goal of Artificial Life since the initial, impressive work of Karl Sims. Incremental evolution has been used extensively to improve the quality of evolutionary search in many complex, non-linear problem spaces. This work sets out to disambiguate the lexicon around incremental evolution, advocating the term *environmental complexification* to represent the complexification of the problem domain. We then go on to analyse various complexification strategies in a structured, complexifiable and yet simple environment: a 3D agent-based obstacle task. We divide the strategies conceptually into *homogeneous* and *heterogeneous*; homogeneous strategies expose successive generations of the population to a single or tightly clustered range of objective functions while heterogeneous strategies present many, covering the range of complexity. It was found that widely-used homogeneous complexification techniques, for example direct presentation of difficult tasks or linearly-increased difficulty, fail due to either loss-of-gradient or temporally-local over-fitting (analogous to *catastrophic forgetting* in neural systems). Heterogeneous methods of complexification (including oscillatory strategies) that eliminate these issues are devised and tested. The heterogeneous category outperforms the homogeneous in all metrics, establishing a much more robust approach to the evolution of naturalistic embodied agents.

Introduction

The evolution of naturalistic, embodied agents and behaviours has been a long-standing goal of Artificial Life since the initial, impressive work of Karl Sims. We are interested in evolving generalised behaviours rather than those that succeed in only a specific or narrow range of parameters: for example in agents able to climb over arbitrary obstacles rather than just those of a specific (maximal or other) height. It may appear desirable to evaluate each individual in each generation on all combinations of parameters for all behaviours, but this approach is infeasible as the number of combinations scales exponentially. We are therefore interested in evolutionary approaches in which each agent is evaluated on a small subset of parameters, in this paper on a single value for a single behaviour, and yet result in agents

able to perform over the full range of parameters, for example by having evolved generalised behaviours rather than ones that work only in a specific or narrow range of parameters.

Incremental evolution has been used extensively to improve the quality of evolutionary search in many complex, non-linear problem spaces. This work sets out to disambiguate the lexicon around incremental evolution, advocating the term *environmental complexification* to represent the complexification of the problem domain as described above. We then seek to identify and objectively compare the strengths and weaknesses of homogeneous and heterogeneous strategies for complexification of a problem domain when using incremental evolution. In homogeneous complexification strategies, for any short sequence of successive generations the population is exposed to a single or tightly clustered range of objective functions, while heterogeneous strategies present many, covering a range of complexity.

Incremental Learning in Evolutionary Systems Inman Harvey's SAGA paradigm, motivated by evolution in the natural world, set the stage for the computational use of incremental evolution by providing an evolutionary mechanism which allows an evolving species to maintain, at least theoretically, most if not all evolutionary pathways as potential candidates for exploration, no matter how converged the population has become to a single point in genotype space (Harvey, 1992, 1997). Once a SAGA algorithm is implemented, objective functions can be changed and the population can be expected to adapt to its new circumstances by traversing *neutral networks* in genotype space (Harvey, 1997, 2001). The requirements for the successful implementation of a SAGA-type incremental process are straightforward: inclusion of mutation as a genetic operator, smooth fitness landscapes and a redundant (high-dimensional) genotype to phenotype mapping which permits neutral networks - interconnected regions of equivalent fitness - to percolate through genotype space. Note that the term *incremental evolution* is used in a sense which implies continued change, development or acquisition of domain knowledge by the algorithm over time. Where there is a gradual increase in dif-

difficulty of objective function, we prefer the term *environmental complexification* as mentioned in (Mouret and Doncieux, 2009). The label *incremental evolution* is also applied where intermediate solutions are moved to a new objective domain; this case we also consider a flavour of environmental complexification, a point we explain in more detail below.

Some of the earliest work which uses the environmental complexification approach directly is that of Gomez who (in addition to discrete, staged evolution over subtasks) gradually increased the speed of prey in a pursuit-avoidance simulation where neural networks were evolved to control simulated predators (Gomez and Miikkulainen, 1997). This work showed a very large performance gain by using the incremental approach. The work also identified an interesting adaptive approach where complexification is dependent upon agent performance at the current level of complexity. Mouret introduced a more general approach to rewarding sub-task performance than the hand-designed staged approach common until this point (Mouret and Doncieux, 2009). Complex agent behaviour was evolved incrementally in a two-dimensional task in (Robinson et al., 2007) where agents in a discrete world were evolved to navigate a hostile environment by avoiding and building bridges over increasingly challenging obstacles. Environmental complexification was used to evolve swarm robots in (Kadota et al., 2012), although the complexification chosen constituted arbitrary, discontinuous changes to the agents' environment and not a smooth transition over a range of difficulties. Notwithstanding, once again the incremental approach delivered a much higher rate of success in the given task (co-operatively foraging for food in a two-dimensional environment). Oh et al. evolved controllers for unmanned aerial vehicles first using a non-incremental strategy. This strategy was found to perform badly as more constraints were added into the objective function so an incremental, task-subdivision strategy was used instead (Oh and Suk, 2013).

Categorisation of Incremental Learning Techniques

Barlow identified two classes of incremental training schemes: functional incremental evolution and environmental incremental evolution (Barlow et al., 2004). In this definition, functional approaches parameterise fitness functions to increase the apparent difficulty of tasks toward the desired level of complexity whereas environmental approaches modify the environment around the evolving individuals without modifying the fitness function, with the same effect. Sub-categories of incremental evolution identified by Mouret in (Mouret and Doncieux, 2009) are *staged evolution*, *environmental complexification*, *fitness shaping* and *behavioural decomposition*. The most striking of these distinctions, common to both Barlow's and Mouret's work is environmental complexification; this category is of particular interest as semantically it can encompass all of the other categories identified and thus becomes synonymous with the

sense of incremental evolution where the problem is simplified and made progressively more difficult. Additionally, environmental complexification is the only category which adequately encompasses co-evolutionary systems (which can be seen as auto-complexification) which in turn are the natural precursors to open-ended evolutionary systems, the development of which is an active area of research in the Artificial Life domain.

Incremental Learning in Neural Systems The idea of incremental learning is not confined to evolutionary algorithms; neural network research has also considered this both as a problem (learning invariances piecewise) and as a solution (tackling complex problems) for networks generally, outside of any particular training scheme. In the standard approach of using neural networks, training and application are distinct phases: all training data are presented to the network and the system learns the invariances and abstractions in that data using some learning algorithm. Then, this trained network is put to work on unseen data. This method of presentation can make it difficult for the network to adapt to new, unseen data at a later time and cause networks to suffer the phenomenon of *catastrophic forgetting* (McCloskey and Cohen, 1989). In contrast to this, incremental learning algorithms are designed to allow the neural system to continually adapt to new information whilst maximising the information available in the network from previous training. This is an important concept for real world applications as often data is not available all at once and sometimes learning guides further exploration, meaning that learning is a continuous process rather than a discrete activity (Giraud-Carrier, 2000). One popular solution to the catastrophic interference problem found in these incremental learning schemes is to *rehearse* either already known data or pseudo-data representing the knowledge already in the network, interleaved with training on new information. See (French, 1999) for an overview and (Guajardo et al., 2010) for an example of recent work using this technique.

Complexification Strategies As noted above, previous work has successfully leveraged the power of incremental evolution through successive increases in environmental complexity. However, attention has been focused solely on the outcome and the particular strategies used to complexify the environment have not been examined in detail. We assert that a rigorous theoretical underpinning of complexification is necessary both for the practical application of incremental evolution and the further elucidation of the interplay of agent and environment in co-evolutionary settings which ultimately lead to unbounded evolutionary activity.

The naive strategy presents the most difficult task to the evolving species at every opportunity. This *straw man* is unlikely to be successful: it was the failing of this approach that spurred the development of alternative, progressive strate-

gies, such as linearly increasing difficulty as time passes. The linear approach has been used often to circumvent the bootstrapping problem, one of the first attempts occurring in (Gomez and Miikkulainen, 1997). This approach is a natural extension of human learning - start easy and then get harder - and the simplicity of implementation and broadness of potential application strengthen its appeal. Many task decomposition strategies can also be considered as implementations of this strategy (linear or at least monotonic increase in complexity), albeit discrete rather than continuous. Gomez also proposed an extension to the linear increase in task complexity where difficulty is only increased when the evolving species achieves a certain level of performance against the current objective function. This interesting strategy has not been developed in detail by others but we consider it a good candidate for analysis as it enforces gradient at every level of difficulty, potentially solving some or all of the issues described in the introduction to this work.

Although not often described in previous work, random presentation of different task complexities may also be useful and finally, drawing upon ideas from incremental learning in neural systems, we propose a strategy of repeated presentation of earlier, simpler tasks in an evolutionary setting. These strategies may have something to offer beyond linear or adaptive monotonic changes in task complexity.

Hypotheses

We anticipate that homogeneous complexification strategies, for example direct presentation of difficult tasks or linearly-increased complexity, will perform poorly due to either loss-of-gradient or temporally-local over-fitting (analogous to *catastrophic forgetting* in neural systems). Heterogeneous strategies are our proposed approach to overcoming forgetting, as an analogue of rehearsal, with smoothly changing heterogeneous strategies, such as oscillatory strategies, also overcoming the loss-of-gradient problem. For oscillatory strategies, the current range of difficulties is from zero to the amplitude of oscillation. A gradual increase of this range may be expected to show improved performance. At very low frequencies, such a strategy would degenerate to the homogeneous linear strategy, and at very high frequencies to the random strategy. Thus, we propose the following hypotheses:

H1: Homogeneous strategies will fail to achieve good coverage on the evaluation task.

H2: Heterogeneous strategies (with the possible exception of random) will achieve better coverage than homogeneous strategies.

H3: Heterogeneous strategies with a range of difficulties increasing over time will outperform heterogeneous strategies with constant range.

H4: A heterogeneous strategy using an oscillatory approach, as an analogue of rehearsal, will exhibit an optimal frequency for any particular problem.

Method

The general setup of our experiment is designed to test the above hypotheses in a task which provides a smooth fitness landscape and neutrality in genotype space. We have chosen the evolution of controllers for three-dimensional agents as the platform, tasked with learning how to walk and climb over an obstacle. The height of the obstacle represents the complexification parameter of the system; task difficulty varies somewhat as obstacle height varies but the ultimate objective for the agents is to deal with every possible obstacle - this is the most complex case. Thus, we can assess which of many possible complexification strategies (that is presentation of tasks of various difficulties) provide the strongest gradient for the evolutionary system to climb and the most robust final evolved agents.

A. The Physical Model In the tradition founded by (Sims, 1994) and continued by many others, we perform all experiments on agents in a three-dimensional virtual world consisting of collidable rigid bodies connected by powered constraints. Unlike Sims, our morphology is a fixed quadruped which is controlled by a feed-forward three-layer perceptron augmented by sinusoidal input. The cuboid quadruped torso (length 0.2m) is supported by four limbs, each comprising an upper and lower portion (length 0.075m). Constraints with two degrees of freedom limit the motion of torso and upper limb at the hips; constraints with one degree of freedom limit the motions of lower limb and upper limb at the knee. See figure 1 for a visual representation. The range of motion of each knee joint is limited from 0 to $\pi/2$ radians, so knees cannot bend outward. The maximum torque that can be applied at any constraint is 0.5 N m. The obstacle is situated 1m from the agent's origin and extends to infinity in x and for 0.05m in y. The height of the obstacle is varied as described elsewhere. The physical simulator used was ODE 0.12, using double-precision arithmetic, the standard big-matrix step function and a step-size of 0.02s. Coulomb friction was applied at contacts between the agent, the obstacle and the ground plane with $\mu = 2$.

B. The Control System The agent controller is modelled by a standard three-layer feed-forward neural network with 12 hidden nodes. Networks receive 4 real-valued inputs in addition to 12 joint-angle sensors. Inputs comprise two sinusoidal oscillators (sine and cosine, period 1 second), an input describing the target location in relation to agent position and orientation (difference between distance from target to each ear, divided by distance between ears) and an up-sensor which describes the orientation of the agent's head relative to the ground plane. Network updates are made synchronously with physics integration. Each hidden node activation is a weighted sum of its inputs with a hyperbolic tangent activation. Each output node activation is a weighted sum of hidden nodes with a logistic activation function.




Figure 1: Visualisation of physical environment. Agent, obstacle and target location are shown.

C. The Evolutionary Algorithm Individual genotypes specify floating-point weights for the neural control system. Initial values for the first generation are drawn from a uniform distribution $x \in [-1, 1]$. In each run, the evolutionary simulation is progressed for 5000 generations using a population of 50 individuals. Individuals are evaluated for 20 simulated seconds and the objective function is defined as the reduction in distance in the x-y plane to a target position situated on the other side of the obstacle. At each new generation, individuals are scored according to the objective function and ranked in order of fitness. The lower half of the population is replaced with mutated, crossed-over variants of the upper half. Single-point crossover is implemented at a random point on the genotype and crosses the current parent individual with another random individual from the best half of the population (possibly itself). Mutation occurs on average twice per genotype and consists of adding a value drawn from a Gaussian distribution with $\sigma = 1$ and $\mu = 0$.

D. The Experimental Setup Sixteen possible strategies for environmental complexification have been identified and tested. Each of these strategies modifies the height of the obstacle in the environment for the current generation of the species. In every case the maximum height of the obstacle, τ , is 0.1m. The height function h for generation G and wavelength λ is defined for each strategy as follows:

1. Direct presentation of environment with complexity τ at every generation: $h(G) = \tau$ (Strategy 1)
2. Presentation of a randomly complex environment at each generation, with complexity drawn from a uniform distribution between 0 and τ : $h(G) = \text{random}(0, \tau)$ (Strategy 2)

3. Gradual complexification of the environment, with complexity interpolated linearly between 0 and τ from generation 0 to generation 4000 and fixed at τ from generation 4001 to 5000 (Strategy 7):

$$h(G) = \begin{cases} \frac{\tau G}{4000}, & G < 4000, \\ \tau & \text{otherwise} \end{cases}$$

4. Oscillating complexification of the environment ($\lambda = 50, 100, 200, 400$ generations), with complexity following a sinusoidal increase and decrease over wavelength λ with maximum amplitude τ (Strategies 3, 4, 5 and 6):

$$h(G, \lambda) = \tau \frac{1 + \sin(\frac{2\pi G}{\lambda} - \frac{\pi}{2})}{2}$$

5. Oscillating complexification of the environment as above, with maximum amplitude interpolated linearly between 0 and τ from generation 0 to generation 4000 and fixed at τ from generation 4001 to 5000 (Strategies 8, 9, 10 and 11):

$$h(G, \lambda) = \begin{cases} \frac{\tau G}{4000} \frac{1 + \sin(\frac{2\pi G}{\lambda} - \frac{\pi}{2})}{2}, & G < 4000, \\ \tau \frac{1 + \sin(\frac{2\pi G}{\lambda} - \frac{\pi}{2})}{2} & \text{otherwise} \end{cases}$$

6. Adaptive modification of 3, where h is instead increased by 1% of τ when the population's average fitness has increased or remained the same and otherwise decreased by 1% of τ , while being kept in the range $[0, \tau]$. (Strategy 12)
7. Adaptive modification of 5 where the value multiplying the oscillator is instead increased by 1% of τ when the population's average fitness has increased or remained the same and otherwise decreased by 1% of τ , while being kept in the range $[0, \tau]$. (Strategies 13, 14, 15 and 16).

Results

Table 1 shows that *no homogeneous complexification strategy (direct, linear or adaptive) was able to achieve success on all task difficulties*, in any experimental run. In contrast, *all heterogeneous strategies did*. The adaptive oscillating ($\lambda=50$) strategy achieved 100% success in 20% of runs and 95% success in 48% of runs.

Figure 2 shows a complete view for each strategy, with $\lambda=50$ selected for each oscillating strategy and each strategy's 100 runs sorted along the horizontal axis by proportion of successful evaluations (shown on the vertical axis). Note that we are primarily interested in the upper portion of this graph, that is in those populations able to complete the task at most obstacle heights. The adaptive strategy generated fewer populations than the linear strategy, successful on fewer than 50% of evaluations (over the full range of obstacle heights) but of greater interest is that it generated only a comparable number of populations successful on more than 90% of evaluations. The random strategy, whilst better than

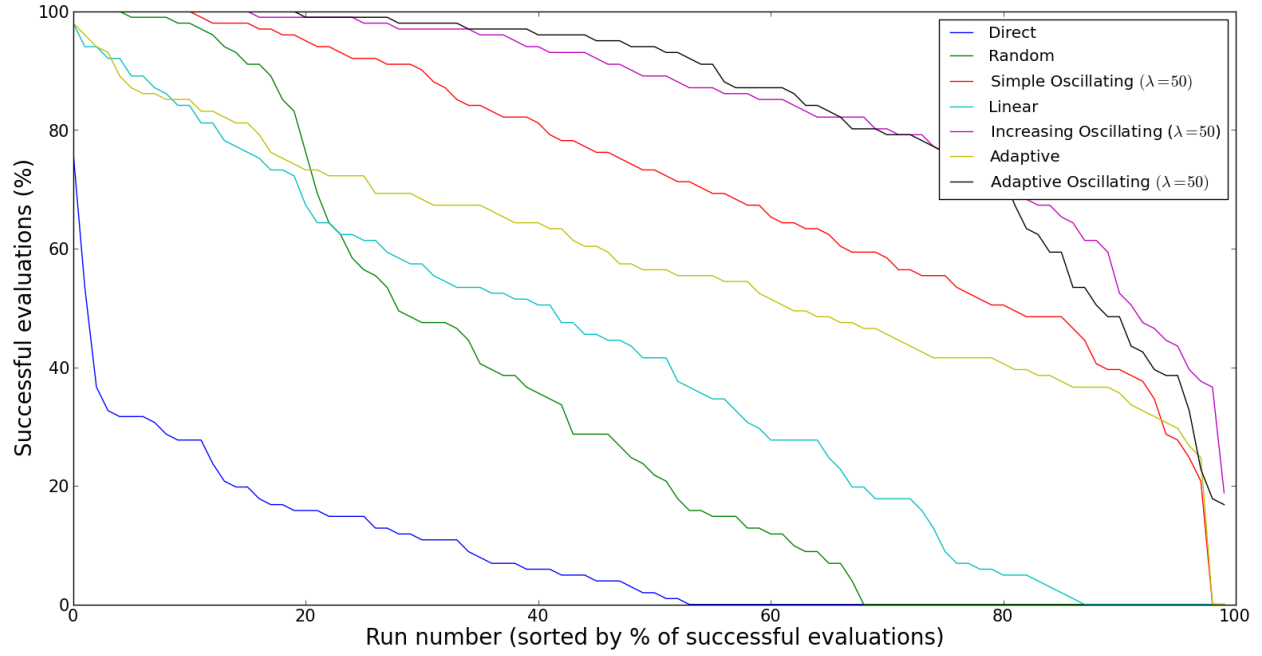


Figure 2: Performance of various strategies, 100 runs per strategy sorted best to worst.

Number	Strategy	% runs with success of at least:	
		95%	100%
Homogeneous Strategies			
1	Direct	0%	0%
7	Linear	1%	0%
12	Adaptive	2%	0%
Heterogeneous Strategies			
2	Random	13%	5%
3	Simple Oscillating ($\lambda=50$)	21%	11%
4	Simple Oscillating ($\lambda=100$)	16%	7%
5	Simple Oscillating ($\lambda=200$)	17%	8%
6	Simple Oscillating ($\lambda=400$)	10%	2%
8	Increasing Oscillating ($\lambda=50$)	39%	16%
9	Increasing Oscillating ($\lambda=100$)	30%	14%
10	Increasing Oscillating ($\lambda=200$)	30%	12%
11	Increasing Oscillating ($\lambda=400$)	29%	10%
13	Adaptive Oscillating ($\lambda=50$)	48%	20%
14	Adaptive Oscillating ($\lambda=100$)	44%	16%
15	Adaptive Oscillating ($\lambda=200$)	26%	9%
16	Adaptive Oscillating ($\lambda=400$)	31%	9%

Table 1: Number of runs achieving success on 95% and 100% of obstacle heights.




Figure 3: Aggregate success rate over all obstacle heights for various strategies, sorted by median success rate. Each evolved population was evaluated on the task at heights 0%, 1%, ... 100%. (See Table 1 for description of the numerical labels.)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1		↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑
2	←		↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑	↑
3	←	←				←	←	↑	↑	↑		←	↑	↑		
4	←	←					←	↑	↑	↑		←	↑	↑		
5	←	←				←	←	↑	↑	↑		←	↑	↑		
6	←	←	↑	↑	↑		←	↑	↑	↑		↑	↑	↑	↑	↑
7	←	←	↑	↑	↑	↑		↑	↑	↑	↑	↑	↑	↑	↑	↑
8	←	←	←	←	←	←	←			←	←	←			←	←
9	←	←	←	←	←	←	←				←	←	↑			
10	←	←	←	←	←	←	←	↑				←	↑	↑		
11	←	←				←	←	↑	↑			←	↑	↑		
12	←	←	↑	↑	↑		←	↑	↑	↑	↑		↑	↑	↑	↑
13	←	←	←	←	←	←	←		←	←	←	←			←	←
14	←	←				←	←			←	←	←			←	←
15	←	←				←	↑					←	↑	↑		
16	←	←				←	↑					←	↑	↑		

Table 2: Significance table showing one-tailed statistical relationship between strategies, $p < 0.05$. Arrows indicate statistical dominance of one strategy over another (significantly higher median value). Statistical test used was the Mann-Whitney U-test.

all homogeneous strategies, is by far the worst method of the heterogeneous strategies. In turn, the simple oscillating strategy is outperformed by the increasing oscillating and adaptive oscillating strategies.

Figure 3 shows a box plot of successful evaluations (%) for each strategy (with whiskers to 1.5 interquartile ranges below and above the lower and upper quartiles), complete with a range of wavelengths (50, 100, 200 and 400 generations). Mann-Whitney U tests were performed to examine significant differences in median number of evaluative successes between strategies and within strategies (by varying wavelength). In table 2, a left arrow indicates that the strategy corresponding to the row number has a significantly higher ($p < 0.05$) median success rate than the strategy corresponding to the column number (and an up arrow vice versa), shown particularly clearly by strategies 8 and 13. Within each of the increasing and adaptive oscillating strategies, the median number of successful evaluations was found to be significantly higher ($p < 0.05$) for strategies with wavelengths of 50 to 100 generations when compared to the same strategy with four times the wavelength or higher. Within the simple oscillating strategy, the long wavelength (400 generations) produced a significantly lower median ($p < 0.05$) than shorter wavelengths (50, 100, 200 generations).

Strategies which oscillate showed the best performance. We found no significant difference in median between the increasing and adaptive oscillating strategies at equal wavelengths. We found that either a linear or an adaptive increase in maximum amplitude over the training time performed significantly better than simple oscillation. For both

increasing oscillating and adaptive oscillating the two lower wavelengths (50 and 100 generations) showed a significantly higher ($p < 0.05$) median number of successful evaluations than the simple oscillating strategies at all wavelengths.

On average, the adaptive strategy performed significantly better than the direct, linear and random strategies, and significantly worse than every oscillatory strategy (except the simple oscillating strategy at wavelength 400 for which we found no significant difference).

The linear strategy resulted in a significantly higher median number of successful evaluations than the direct and random strategies (even though the random strategy produced more highly fit populations from many more runs) and a significantly lower median than all other strategies.

On average, the random strategy performed significantly worse than all other strategies except for the direct method, which was significantly worse than all other strategies.

In order to determine whether the poor results of the linear strategy is due to either evolutionary loss or failure to gain we determined the proportion of successful evaluations at each obstacle height throughout the evolutionary progress, for each strategy. All strategies achieved 8% success at all obstacle heights, with the exceptions of direct (for which obstacle height is always 100%) and adaptive (low coverage at high obstacle height). The linear strategy achieved more successful evaluations than the simple oscillating strategy at all wavelengths during the evolutionary phase, indicating that its ultimate failure is due to evolutionary loss rather than a failure to gain. Only 10% of the final populations from linear runs contained an individual able to walk to the target with no obstacle, compared to at least 69% for the increasing and adaptive oscillating strategies. Figure 4 shows the performance of the linear strategy and the three oscillatory strategies of wavelength 50 generations, against obstacle height, during evolution.

As in figure 3, figure 5 shows the number of successful evaluations for each strategy, but drawn only from those runs able to reach the target with no obstacle (that is eliminating those runs which experienced the greatest evolutionary loss). It shows that in these cases, linear performance has a range comparable to the simple oscillatory strategies and a median comparable to the increasing and adaptive oscillating strategies.

To investigate the dependency of success rate on oscillatory frequency we evaluated the simple, increasing and adaptive oscillating strategies across a range of wavelengths from 2 to 10000 generations; figure 6 demonstrates this relationship. As wavelength approaches zero, the proportion of successful evaluations approaches that of random. As wavelength approaches total evolutionary time (number of generations), the proportion of successful evaluations approaches that of linear. Between these points, it can be seen that for each strategy there is an optimal wavelength (for the current algorithm, around 50-100 generations).

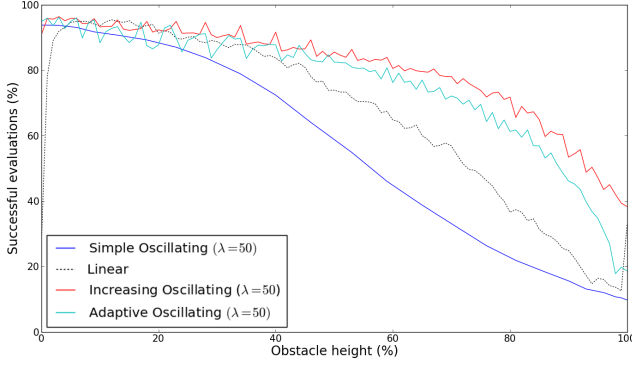


Figure 4: Strategy performance against obstacle height during evolution.




Figure 5: Success rate over all obstacle heights for various strategies (only aggregates runs which solved the task at zero-height). (See Table 1 for description of numerical labels; order preserved from Figure 3.)




Figure 6: Strategy performance (% success) against wavelength for oscillating strategies.

Discussion

It is clear from the results presented above that there is a strong distinction between the homogeneous and heterogeneous strategies. No homogeneous strategy achieved 100% coverage of the evaluation task in any run (Table 1) whereas all heterogeneous strategies did. Within the homogeneous category, the trivial, direct method of presentation was by far the least successful (Figures 2 and 3). The linear strategy was more successful but the best strategy in this category was the adaptive strategy. The poor performance of the homogeneous category can be explained by *evolutionary forgetting*: these strategies have either lost evolutionary gradient and drifted away from any early successes (linear) or over-specialised on later parts of the problem (adaptive).

The heterogeneous strategies performed better than the homogeneous group: the most successful strategies we explored all made multiple presentations of easier tasks at later stages of the evolutionary run, at the expense of fewer consecutive presentations of very similar tasks. These strategies all performed well at the hardest task and had the best generalisation performance over the whole range of tasks, suggesting that our hypothesis has merit.

The random strategy is the least successful strategy in this category. This may be due to the same problem of gradient loss as in the homogeneous group. As found in the homogeneous group, the linear and adaptive modifications of the oscillating strategy showed the best performance of all; the slow increase in task difficulty maintains a strong evolutionary gradient and the cyclical nature of task presentation consolidates earlier gains and causes the evolving population to prefer generalised solutions abstracted over the whole problem domain.

This consolidation is dependent on the frequency of representation of earlier, or easier, parts of the task. When investigating this frequency, it can be seen that a clear optimum exists in the frequency domain where cyclical strategies are able to maximise this consolidation without losing gradient. This optimum is likely to be problem-specific and a range of values should be explored for any given task. However, in the limit of wavelength, that is at very low and very high frequencies, it can be seen that the performance of the evolving populations begins to approximate, for low and high frequencies respectively, the linear and random strategies. This offers an abstract insight into the underlying mechanism at work - the maintenance of selective pressure and whole-task capability. As these components reduce in effectiveness due to the change in wavelength, so the oscillating strategies degenerate into the simpler strategies described above. The successful cases are those where environmental change is fast enough to induce a generalisation in the agent's approach to the task but slow enough to prevent catastrophic loss of gradient when evaluating partial solutions.

Conclusions

The points made in the discussion section support our hypotheses. The homogeneous strategies showed weak performance on the evaluation task, with no strategy achieving full coverage in any run. Conversely the heterogeneous strategies, including surprisingly the random strategy, all achieved full coverage in some runs. Those heterogeneous strategies with a range of difficulties increasing over time (increasing and adaptive oscillating) outperformed the simple (constant range) oscillating strategies, showing a much higher proportion of successful runs. Finally, we demonstrated that oscillating strategies do exhibit an optimal frequency.

Complexification strategies for incremental evolution offer a powerful mechanism for adaptive problem solving. However, this power comes at a price: it is easy to lose information learned earlier in the process. In order to fully exploit this power appropriate complexification strategies have to be realised in order to drive populations along desirable adaptive pathways. There are many options for formulating these strategies: much previous work has involved, in one manner or another, a simplification of the objective function and then a progressive complexification as time passes. In this work we found that many strategies encounter loss-of-gradient or over-fitting problems. We present a solution in the form of heterogeneous complexification strategies which combine solutions to those problems to deliver robust populations. Our approach can be translated to many scenarios where progressive complexification is used to guide an incremental evolutionary process; further exploration of the limitations and advantages of heterogeneous complexification within different problem domains would be useful in order to generalise these conclusions. Additionally, the oscillating strategies exhibited an optimal wavelength for representation. It is unclear whether this optimum is task-dependent or whether there is an underlying principle and optimal wavelength for this type of training; this question also merits further work.

Finally, we would advise that in general while a random presentation of subtasks or objective difficulty levels is preferable to a linear increase, as a minimum guideline an increasing heterogeneous complexification strategy should be used. This rehearsive, cyclical approach to presentation not only maintains evolutionary gradients but also promotes generalisation amongst the evolving populations from subtask-specific adaptation to performance across the super-task.

References

- Barlow, G. J., Oh, C. K., and Grant, E. (2004). Incremental evolution of autonomous controllers for unmanned aerial vehicles using multi-objective genetic programming. In *Proceedings of the 2004 IEEE Conference on Cybernetics and Intelligent Systems*, volume 2, pages 688–693.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135.
- Giraud-Carrier, C. (2000). A note on the utility of incremental learning. *AI Communications*, 13(4):215–223.
- Gomez, F. and Miikkulainen, R. (1997). Incremental evolution of complex general behavior. *Adaptive Behavior*, 5(3-4):317–342.
- Guajardo, J. A., Weber, R., and Miranda, J. (2010). A model updating strategy for predicting time series with seasonal patterns. *Applied Soft Computing*, 10(1):276–283.
- Harvey, I. (1992). Species adaptation genetic algorithms: A basis for a continuing SAGA. In *Proceedings of the First European Conference on Artificial Life*, pages 346–354.
- Harvey, I. (1997). Artificial evolution for real problems. In *Proceedings of the 5th International Symposium on Evolutionary Robotics*, pages 127–149.
- Harvey, I. (2001). Artificial evolution: A continuing SAGA. In *Proceedings of the 8th International Symposium on Evolutionary Robotics*, pages 94–109.
- Kadota, T., Yasuda, T., Matsumura, Y., and Ohkura, K. (2012). An incremental approach to an evolutionary robotic swarm. In *Proceedings of the 2012 IEEE/SICE International Symposium on System Integration*, pages 458–463.
- McCloskey, M. and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165.
- Mouret, J.-B. and Doncieux, S. (2009). Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *Proceedings of the 2009 IEEE Congress on Evolutionary Computation*, pages 1161–1168.
- Oh, S.-H. and Suk, J. (2013). Evolutionary controller design for area search using multiple UAVs with minimum altitude maneuver. *Journal of Mechanical Science and Technology*, 27(2):541–548.
- Robinson, E., Ellis, T., and Channon, A. D. (2007). Neuroevolution of agents capable of reactive and deliberative behaviours in novel and dynamic environments. In *Proceedings of the 9th European Conference on Artificial Life*, pages 345–354.
- Sims, K. (1994). Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22.