

Stochastic Ontogenesis in Evolutionary Robotics

Adam Stanton

School of Computing and Mathematics, Faculty of Natural Sciences,
Keele University, Newcastle-under-Lyme, ST5 5BG, UK.
a.stanton@keele.ac.uk

Abstract

This paper investigates the hypothesis that noise in the genotype–phenotype mapping, here called *stochastic ontogenesis* (SO), is an important consideration in Evolutionary Robotics. This is examined in two ways: first, in the context of seeking to generalise controller performance in an incremental task domain in simulation, and second, in a preliminary study of its effectiveness as a mechanism for crossing the “reality gap” from simulation to physical robots. The performance of evolved neurocontrollers for a fixed-morphology simulated robot is evaluated in both the presence and absence of ontogenic noise, in a task requiring the development of a walking gait that accommodates a varying environment. When SO is applied, evolution of controllers is more effective (replicates achieve higher fitness) and more robust (fewer replicates fail) than evolution using a deterministic mapping. This result is found in a variety of incremental scenarios. For the preliminary study of the utility of SO for moving between simulation and reality, the capacity of evolved controllers to handle unforeseen environmental noise is tested by introducing a stochastic coefficient of friction and evaluating previous populations in the new problem domain. Controllers evolved with deterministic ontogenesis fail to accommodate the new source of noise and show reduced fitness. In contrast, those which experienced ontogenic noise during evolution are not significantly disrupted by the additional noise in the environment. It is argued that SO is a catch-all mechanism for increasing performance of Evolutionary Robotics designs and may have further more general implications for Evolutionary Computation.

Introduction

It is well-known in biology literature that identical genomes do not give rise to identical organisms. Studies of near-exact DNA copies show that during development (ontogenesis) small environmental variations cause relatively large phenotypic changes (Wong et al., 2005). In some cases there is a clear advantage conferred to organisms whose developmental programmes adapt to their environment as they progress, for example to accommodate differences of temperature or salinity. In other cases, random noise from either the environment or the local molecular machinery (Gordon et al., 2009) nudges the growth process into this or that dynamical trajectory, shaping the organism differently each time. This is one outcome of the presence of feedback relationships in

the regulatory architecture. The phenotypic discordance of monozygous human twins is a good example of these effects in action (Fogel et al., 1965). It is assumed by practitioners of Evolutionary Computation (EC) that the uncertainty of biological ontogenesis is an unavoidable consequence of the substrate; the possibility that some variation in the translation from genome to organism has been conserved due to its utility is often overlooked.

Evolutionary Robotics (ER) aims to use EC principles to build robots and robot controllers that operate effectively despite environmental perturbations and uncertain information (Nolfi and Floreano, 2000). Whilst many examples of research in this area focus exclusively on simulation, other attempts have been made to evolve robot controllers that are competent in real-world implementations. Sometimes the search takes place entirely in physical reality, but often simulations are used to accelerate the evaluation of solutions and thus the evolutionary process, albeit at a cost of accuracy and real-world performance, see for example Cliff et al. (1993), Zagal et al. (2004), and Koos et al. (2013).

This disparity between simulation and real robotics is known as the *reality gap*, and bridging it is an important area of research in ER (Mouret and Chatzilygeroudis, 2017). An evolutionary domain with real-world fidelity and faster-than-real-time evaluation would be a major achievement. Various attempts have been made to address this ambition: optimisation of simulator parameters based on real-world experiences (Bongard and Lipson, 2004); interleaving real-world and simulated evaluations (Goosen et al., 2007); learning to predict simulator accuracy from behaviour descriptors so that individuals with good predicted “transferability” can be selected for (Koos et al., 2013); and injecting noise into the simulation (Jakobi et al., 1995; Jakobi, 1997).

Jakobi proposed the *radical envelope-of-noise hypothesis*, a comprehensive account of how the principled addition of random noise to simulations can improve the performance of robots evolved in simulation when evaluated in real-world settings. Jakobi’s approach stipulates two important conditions that together aim to smooth the transition between simulation and reality. First, that every implementation aspect of the simulation must be randomly varied from trial to trial so that reliably fit controllers are *base-set exclusive*, mean-

ing that the evolved controllers are tuned only to the factors common to all simulators. Second, that every base-set aspect of the simulation is randomly varied from trial to trial so that reliably fit controllers are *base-set robust*—meaning that evolving species cannot exploit idiosyncratic features of the implementation of base-set aspects in any particular simulator.

In this paper I examine the hypothesis that noise in the genotype–phenotype mapping, here called *stochastic ontogenesis* (SO), exerts beneficial effects on an evolutionary system, in terms of its final fitness within individual runs and across replicates. I also present evidence that SO captures some of the essential characteristics of the radical envelope-of-noise hypothesis whilst significantly reducing the complexity and problem-specificity of implementing such a mechanism. To address these points, performance of simulated evolved robot neurocontrollers is evaluated in the presence of ontogenic noise. By comparing to a deterministic equivalent, it is shown in multiple contexts that evolution in the presence of this noise is more effective and produces more robust solutions. It is demonstrated that the improvement in performance is present despite fundamental differences in the underlying evolutionary schemes employed. The performance advantage of stochastic ontogenesis is lost only when carefully-chosen incremental environmental changes that introduce extrinsic, noise-like variation in the task are used. In this case, there is no disadvantage to the presence of stochastic ontogenesis.

The final part of the paper discusses the potential for stochastic ontogenesis to fulfil the role of an envelope of noise, delivering the generalising power of that approach but with a situation-agnostic encoding of the noise in the genotype–phenotype mapping, rather than explicitly in the environment. A preliminary empirical demonstration is presented by using a varying base-set aspect of the simulator (specifically, the coefficient of friction) as a proxy for a reality gap. The differing performance of controller species evolved according either to a stochastic or deterministic ontogenic scheme is observed in these worlds exhibiting previously unseen noise. Results show that individuals from the stochastic simulator have significantly better performance in the novel uncertain environment than those from the deterministic world.

Background

In addition to Jakobi’s work (Jakobi et al., 1995; Jakobi, 1997), other research in ER and the related field of Virtual Creatures has examined the effectiveness of noise in guiding evolutionary processes towards more robust solutions. Seth (1998) added noise to a fitness function to promote better generalisation, and Miconi and Channon (2005) used a co-evolutionary approach that effectively adds a similar noise component to their fitness function. Relatively few investigations have examined noise in the genotype to phenotype translation. Branke (2001) explored the use of noise by carefully selecting maximally-informative simulation variants to

accelerate robust evolution whilst incurring minimal computational overhead. The variants were defined based on the addition of noise between the genotype and phenotype. In this case however, the noise was carefully controlled in order to deliver the requisite extra exploration of the local solution space. Fernandez Leon (2011) injected a stochastic component into a dynamic leaky integrator network to increase its robustness and capacity to generalise. Innovative research in artificial cultural evolution has incorporated stochastic ontogenesis in order to capitalise on the extra exploration of the solution space afforded by such a modification. There it is called *transcription error* and used as a mechanism for “individual learning” (Borg et al., 2011). That work avoids considering stochastic ontogenesis as exerting a primary effect on the underlying evolutionary system (despite some results being suggestive of this effect) and in fact concludes the opposite: that it is only effective when combined with cultural (i.e. horizontal) information transmission.

Hypothesis

Noise has been applied to Evolutionary Robotics problems to increase robustness and performance, by manipulating simulation parameters to incorporate a stochastic component. In this paper, I explore the hypothesis that adding noise to a genotype when it is expressed as a phenotype (here called Stochastic Ontogenesis) will improve the performance of the evolutionary search (in terms of fitness of evolved species across a number of replicates) and improve robustness by acting as an alternative to the injection of noise into the environment.

Methods

The evolutionary robotics platform used in this work is a development of the environments presented in Stanton and Channon (2013), Moore and Stanton (2017), and Stanton (2017). The reader is invited to refer to those works for detailed descriptions of the platform and its development. However key detail, including any necessary exposition of algorithmic processes and parameter values, is summarised below for clarity and completeness.

Robot Platform and Simulation Environment Figure 1 shows the simulated quadrupedal robot (animat) used in this study. The robot has a cuboid torso and four legs placed at the corners. Each leg is divided into an upper and lower segment. The hip is a 2-degree of freedom (DOF) joint while the knee is a 1 DOF joint; see Table 1 for specific parameter values. The robot’s joints are actuated using a Proportional-Derivative (PD) control mechanism (Reil and Husbands, 2002) that takes a target angle as input and applies a torque to the joint according to Equation 1, where T is the applied torque, k_s and k_d are the spring and damper constants, θ_d is the desired angle, θ the current angle and $\dot{\theta}$ the angle change since the last timestep. As in previous work, $k_s = k_d = 0.5$ in this study.

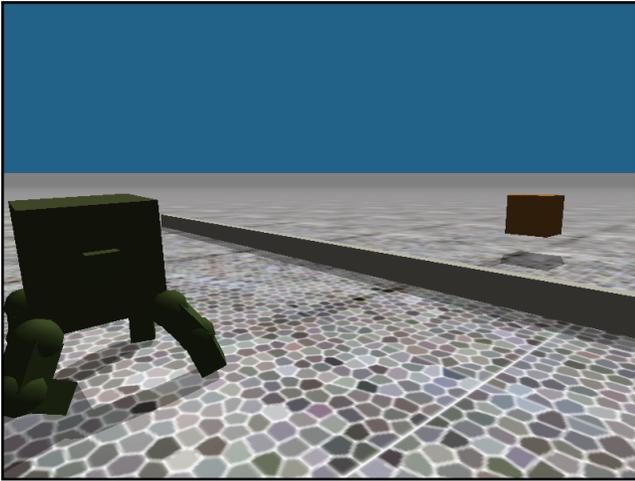


Figure 1: The quadrupedal animat and simulation environment in this study. The animat is tasked with crossing a wall and moving towards a target, represented by the box.

Head dimension	$0.2 \times 0.2 \times 0.2$
Leg section dimension	$0.075 \times 0.05 \times 0.05$
Head mass	2.0
Leg section mass	0.5
Hip joint axis 1	vertical, range $[-\frac{\pi}{4}, \frac{\pi}{4}]$
Hip joint axis 2	horizontal, range $[0, \frac{\pi}{2}]$
Knee joint axis 1	horizontal, range $[0, \frac{\pi}{2}]$
Maximum torque	0.125

Table 1: Physical parameters of robot.

$$T = k_s \times (\theta_d - \theta) - k_d \dot{\theta} \quad (1)$$

Simulations are conducted in the Open Dynamics Engine (ODE) (Smith et al., 2005), a real-time rigid body physics engine. ODE version 0.15.2 was used. ODE computes the interaction between the different rigid components of the robot as well as interaction with the environment; Table 2 lists key parameters used.

Task Environment Animats are evaluated on their ability to cross a wall and move towards a target position. Wall heights vary, from 0.01 up to a maximum value of 1.0 in 0.01 increments. Fitness values for animats are negative values, with a maximum fitness of 0.0 corresponding to an animat reaching the target.

Controller The controller is a feed-forward ANN with fixed, fully-connected topology. Inputs comprise a number of signals from the environment as well as spontaneous activity generated by sinusoidal functions. These inputs are shown in Table 3. The ANN has 12 hidden nodes and 12 outputs mapping to desired angles for each joint at the next

Timestep	0.02 seconds
Gravity	-1.2
Friction model	Pyramid; $\mu_f = 2.0$
Global ERP	0.2
Global CFM	5.0×10^{-5}
Wall dimension	$0.05 \times 5.0 \times h$; $h_{max} = 0.1$
Wall position	$x = 1$; $y = 0$
Target position	$x = 2$; $y = 0$
Start location	$x = 0$; $y = 0$
Simulation time	20 seconds

Table 2: ODE simulation parameters.

timestep within the ranges described in Table 1. tanh function is used for the transfer function in the hidden layer, and the logistic function for outputs. The ANN is updated in lock-step with the dynamics simulation, and inputs are propagated completely through the network at each update.

1	$\sin(2\pi t)$
2	$\cos(2\pi t)$
3	balance: $\arccos(\mathbf{H}[10])$
4	$(\vec{H}_l - \vec{H}_r) \div H_{width}$
5-12	hip joint angles
13-16	knee joint angles

Table 3: ANN controller inputs, where \mathbf{H} is the ODE rotation matrix of the robot’s head, $|\vec{H}_x|$ is the distance from each side of the robot’s head to the target and H_{width} is the width of the head.

Evolutionary Algorithm A generational EA was employed in this study. A genome specifies a set of floating-point weights for the ANN controllers. Populations comprise 50 individuals, each randomly initialized with values drawn from $\mathcal{U}(-1.0, 1.0)$ based on a starting seed corresponding to the replicate number. Fitness-based selection is used where the population is ranked and the lower half is replaced with mutated copies of the upper half, recombined with a random individual from the upper half using single-point crossover. The mutation rate used was $\frac{2}{N}$, where N is the length of the genome. If applied, the mutation is drawn from $\mathcal{N}(0.0, 1.0)$, and allele values are unbounded. Morphologies of the animats remain fixed throughout evolution. An animat is evaluated based on its Euclidean distance from a target at the end of a simulation. High fitnesses are only attained if the animat climbs over an intervening wall of varying height.

Treatment analysis Each treatment is replicated 20-fold for statistical validity. In each replicate, after the Evolutionary Algorithm has proceeded for 5000 generations ac-

Scenario	Changes to the environment (height of the wall) at each generation
A	The environment varies between 0 and h_{max} according to a uniform random distribution.
B	The environment varies between 0 and h_{max} linearly between 0 and 4000 generations, and then remains static at h_{max} .
C	The environment varies sinusoidally between 0 and h_{max} every 100 generations.

Table 4: Incremental Scenarios

cording to the schedule outlined above, the final population is evaluated to assess its capacity to achieve the wall-crossing task. 100 evaluations take place for each individual in each population. Each individual’s proximity to the target after 20 simulated seconds is recorded for each wall height from 0 up to (but not including) h_{max} (where h_{max} is the maximum height), in increments of $0.01 \times h_{max}$. For each replicate in each treatment, the individual with the highest mean fitness across the 100 evaluations is chosen as the representative individual. The performances of the 20 “best” individuals across all replicates for each treatment are then analysed to build a picture of the relative capability of each treatment to address the wall-climbing problem. After this preprocessing, treatments are presented as box-whisker plots and per-treatment heatmaps of raw evaluation performance showing replicates and evaluation-stage fitness. Treatments are compared for significant differences using the non-parametric Mann-Whitney U test null hypothesis (that it is equally likely that a randomly selected value from one sample will be less than or greater than a randomly selected value from the second sample). Where a significant statistical difference is found and discussed, effect size is presented with the Common Language Effect Size methodology (CLES).

Stochastic Ontogenesis To inject randomness into the genotype–phenotype transformation, Gaussian noise was added to the evolved weights of the neural network at construction time, leaving the individual’s genome unchanged in the EA. The mean μ of this noise is zero and the spread σ^2 is varied according to the scheme outlined below. When $\sigma^2 = 0$, ontogenesis is deterministic. This process occurs each time an individual is generated from its genome, meaning network weights are modified a) for each evolutionary trial, and b) for each of the 100 tests in the evaluation stage.

Treatments Treatments compare stochastic ontogenesis with deterministic ontogenesis (DO) in three different incremental evolutionary scenarios (see Table 4; refer to Stanton and Channon (2013) for expanded details on these scenarios). Where not otherwise specified, Scenario A is used. In the deterministic ontogenic model, neural weights are

read directly from the genome and used in the robot’s controller without modification (Treatment I). In the stochastic ontogenic model, neural weights read from the genome are subject to Gaussian noise with $\sigma^2 = 0.05$, $\sigma^2 = 0.1$, $\sigma^2 = 0.15$, and $\sigma^2 = 0.2$ (Treatments IIa, IIb, IIc, and II d). Additionally, a distinction is made between evolutionary time and evaluation time. Two additional treatments are examined: first, where only the evaluation phase incorporates noise (during the evolutionary phase, the mapping between genotype and phenotype is deterministic; Treatment III). Second, the reverse, where the evolutionary phase incorporates ontogenic noise but the evaluation phase does not (Treatment IV). Finally, the other combinations of scenario and stochasticity are considered: Scenario B, deterministic versus stochastic (Treatments V and VI), and Scenario C, deterministic versus stochastic (Treatments VII and VIII). Treatments III, IV, VI and VIII all use $\sigma^2 = 0.1$.

Additional treatments In order to establish the environment-agnostic effect that stochastic ontogenesis exerts, two additional treatments were conducted with the objective of varying the environment and at the same time eliminating noise in the genotype–phenotype mapping. In this way, the degree to which stochastic solutions generalise beyond their immediate simulation environment can be assessed, and a judgement made as to whether this method of introducing noise is effective in the ambition to make evolving species, in the parlance of Jakobi, base-set robust. To this end, the simulation was adapted to include noise in the friction coefficient that governs the robots’ interaction with the ground, effectively making the environment stochastic. When applied, the friction parameter of the simulator was drawn from a normal distribution $\mathcal{N}(2.0, 1.0)$. In standard environments, this parameter is fixed at 2.0—see Table 2. Treatments IXa (Scenario A; deterministic ontogenesis in a stochastic environment), IXb (Scenario A; stochastic ontogenesis in a stochastic environment), Xa (Scenario C; deterministic ontogenesis in a stochastic environment), and Xb (Scenario C; stochastic ontogenesis in a stochastic environment) are thus undertaken and compared. Treatments IXb and Xb both use $\sigma^2 = 0.1$.

Results

- When stochastic ontogenesis is used in scenario A, there is a significant ($p < 1.0 \times 10^{-100}$) increase in performance for all examined σ^2 compared to the deterministic mapping, with a very large effect size of 0.752 when $\sigma^2 = 0.1$ (c.f. Figure 2, I vs. IIa-d and Figure 3, I vs IIb).
- The increased performance is not an artefact of the evaluation mechanism. When individuals evolved in a deterministic ontogenic environment are evaluated in a stochastic world, the worst performance of all is observed (Figures 2 and 3, III). In contrast, when individuals evolved using stochastic ontogenesis are evaluated using deterministic ontogenesis, the highest performance of all is observed (Figures 2 and 3, IV).

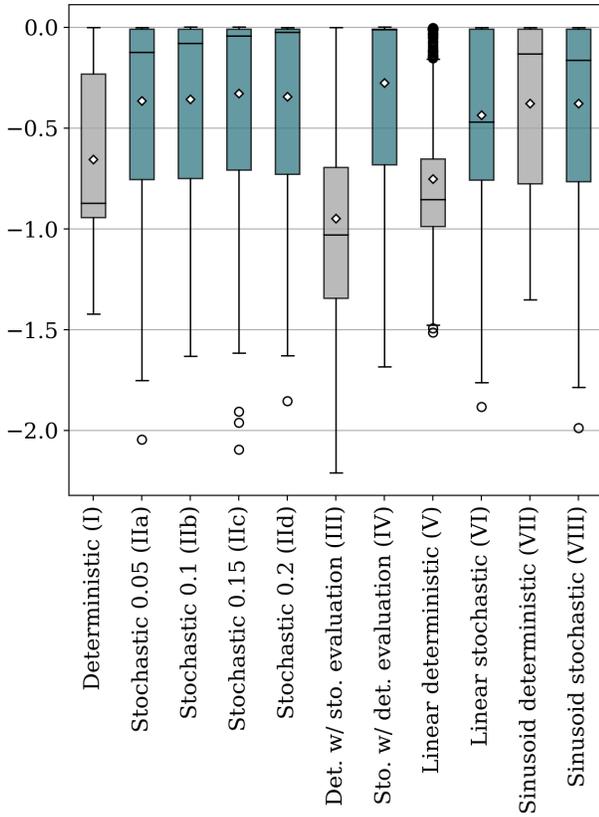


Figure 2: Box-whisker plot (where diamonds indicate means and whiskers indicate $1.5 \times IQR$, or population extremes where there are no outliers) showing differences between treatments that involve a stochastic ontogenetic element during evolution, evaluation, neither, or both. Adding noise during evolution significantly improves performance on the given task, whereas adding it only at evaluation time significantly lowers the quality of results.

- The advantage of stochastic ontogeny is not limited to a particular incremental evolutionary scheme. It is shown that stochastic ontogeny significantly improves performance in both the uniform random presentation scheme (Scenario A; Figures 2 and 3, I and IIa-d) and the linear presentation scheme (Scenario B; Figure 2, V and VI, $p < 1.0 \times 10^{-100}$).
- Species evolved with SO show an improvement in the absolute fitness achieved by the highest-scoring replicates, as well as a general improvement in the quality of all replicates (Figure 6).
- Neither a positive nor negative effect is observed when stochastic ontogenesis is applied in the sinusoidal presentation scheme (Scenario C; Figure 2, VII and VIII).
- When evolved species are evaluated in an uncertain environment (where friction changes for each evaluation), in-

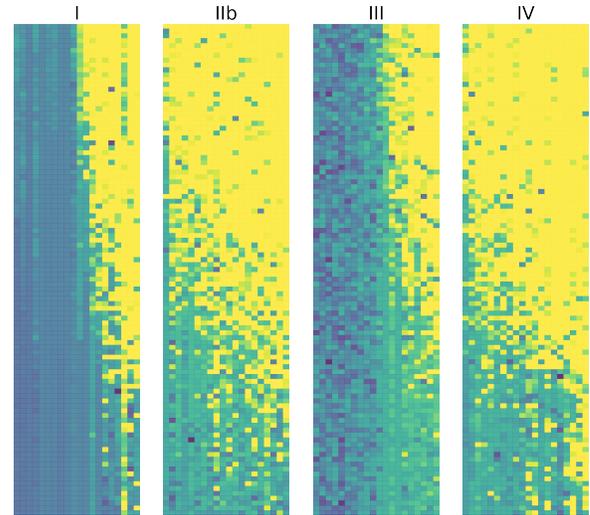


Figure 3: Heatmaps showing fitnesses at evaluation time for: (I) deterministic ontogenesis during evolution and at evaluation; (IIb) stochastic ontogenesis during evolution and at evaluation; (III) stochastic ontogenesis only at evaluation time; and (IV) stochastic ontogenesis only during evolution. Blue areas indicate poor performance; yellow areas indicate strong performance. Replicates are presented in increasing order of successful (reach target) evaluations across the horizontal axis; wall heights are presented on the vertical axis. It is notable that performance is extremely suppressed in the case where noise is introduced only during the evaluation phase (III).

dividuals evolved with stochastic ontogeny perform significantly better than those evolved with deterministic ontogeny (Figure 4, IXb vs IXa; $p < 1.0 \times 10^{-100}$, $ES = 0.759$).

- There is no significant difference between evaluations in deterministic or stochastic environments for individuals evolved with stochastic ontogeny (Figure 4, IXb vs IIb), but those evolved with deterministic ontogeny show a significant decrease in performance in these worlds (Figure 4, IXa vs I; $p < 1.0 \times 10^{-25}$, $ES = 0.355$).
- When species evolved with deterministic ontogeny and a sinusoidal presentation (Scenario C) are evaluated in a simulation with an unpredictable friction coefficient, performance decreases (Figure 4, Xa vs. Figure 2, VII; $p < 1.0 \times 10^{-25}$, $ES = 0.390$). In contrast, where SO is used, performance is maintained and no significant difference is observed (Figure 4, Xb vs Figure 2, VIII).
- Aggregate performance changes in the simulation with an unpredictable friction coefficient are driven by an increase in the number of successful runs (more replicates achieve some competency in the task), and an increase in the maximum performance achieved by individual runs (Figure

5). In the Scenario C version of this treatment, the improvement is driven more by the increase in successful runs and less by the increase in maximum performance (Figure 5, Xa vs Xb).

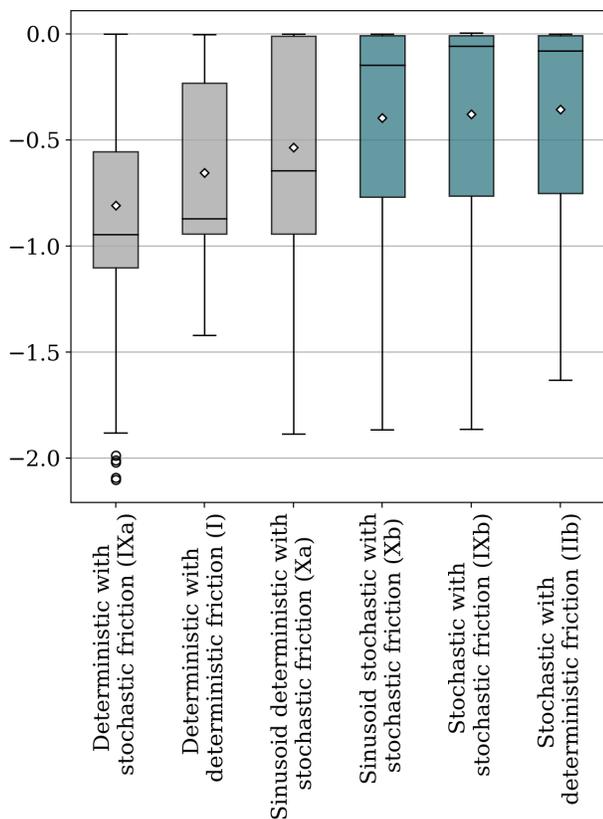


Figure 4: Box-whisker plot (where diamonds indicate means and whiskers indicate $1.5 \times \text{IQR}$, or population extremes where there are no outliers) showing differences between treatments when the physical environment is noisy, manifest in this case in the form of unpredictable friction coefficients (μ_f). Populations evolved with deterministic ontogenesis show decreased performance, including populations evolved in a sinusoidal environment which have equivalent performance in the predictable world to those evolved with stochastic ontogenesis. In contrast, populations evolved with stochastic ontogenesis are able to accommodate the additional environmental noise and do not exhibit decreased performance.

Conclusions

The main conclusion drawn from this work is that the performance of the evolutionary system is significantly improved when SO is applied in the translation from evolving genotype to expressed phenotype. This improvement arises from an increase in the performance of high-scoring replicates (as shown in Figures 2 and 4), and also from a general increase

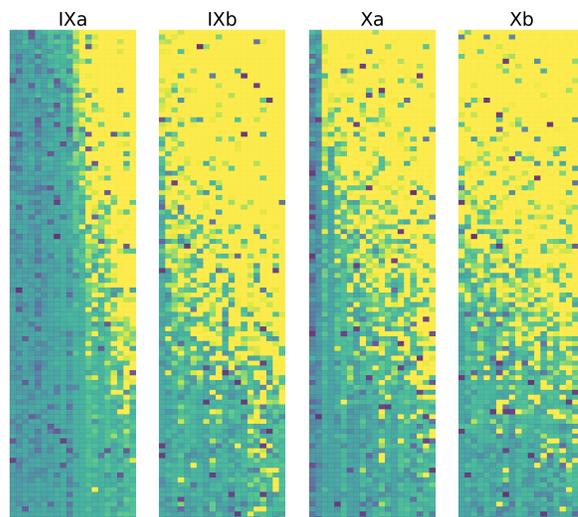


Figure 5: Heatmaps showing fitnesses at evaluation time in an uncertain world for: (IXa) deterministic ontogenesis; (IXb) stochastic ontogenesis; (Xa) deterministic ontogenesis in incremental Scenario C; and (Xb) stochastic ontogenesis in incremental Scenario C. Blue areas indicate poor performance; yellow areas indicate strong performance. Replicates are presented in increasing order of successful (reach target) evaluations across the horizontal axis; wall heights are presented on the vertical axis. Treatment Xa shows reduced performance for some replicates (darker columns on left-hand side); this is evidence of compromise due to the uncertain world. Treatment Xb, where species have evolved with SO, show robustness to this change.

in the number of runs performing well (shown in Figure 6). The possibility that there is an artefact in the method of evaluation arising from the inclusion of SO has been eliminated by looking at performance when DO is used during evolution and SO in evaluation. Since in this case performance is comparatively poor, the improvements in fitness with SO point to a general change in evolutionary pressure, from favouring brittle solutions early on (that cause suboptimal species to arise and persist), towards solutions that colonise broader areas of the evolutionary space and thus fare better in a broader range of task configurations. Stochasticity in the genotype–phenotype mapping is modulating the evolutionary system itself.

It is also shown that the effects of SO transcend specific instances of the evolutionary system. The increased performance of species evolved in both Scenario A (where sub-task presentation is random) and Scenario B (where sub-task presentation changes monotonically) points to a general mechanism at work when SO is applied. The *lack* of observable effects in Scenario C implies that this incremental evolutionary technique has captured some of the advantages of SO. However, since this technique is the best from a broader sweep of parameters in this problem (i.e. it is a

strongly task-specific formulation), the SO method is still advantageous since it is situation-agnostic. In other words, the naive SO algorithm performs as well as the task-specific Scenario C treatment, so should be preferred.

Last, SO species handle unseen environmental uncertainty better than DO species. Where DO species evolved in Scenario C perform well in the initial problem, their advantage is lost in the more difficult Uncertain World problem, used here as a proxy for crossing the reality gap. Since species evolved with SO do not show the same performance drop, there is an implication that SO is a general mechanism for handling uncertainty. In other words, the important result that the sinusoidal treatment performs poorly in the deterministic ontogeny/stochastic environment world but well in the deterministic ontogeny/deterministic environment and stochastic ontogeny/stochastic environment implies that although the sinusoidal environment has some effect on the evolutionary system, much less robustness is incorporated into the final populations than with the SO scheme. The increased robustness is evident in Figure 5, Xa vs Xb, where absolute performance does not show significant improvement, but the overall number of successful replicates (where some progress is made in the task) increases.

To put this result in terms of Jakobi’s radical envelope-of-noise hypothesis, Stochastic Ontogenesis appears to be a workable alternative to this scheme, at least for base-set robustness. A noisy coefficient of friction has been used as a proxy for a difference between simulation and the real world. Although only a preliminary demonstration in simulation, it has been shown that species evolved with SO are able to perform well even in environments in which they have had no prior experience at all: environments where a base-set aspect, friction, is varied as if moving from simulation to reality. SO populations show robustness to this variation; indeed, this robustness stems from SO. Unlike Jakobi’s scheme, the environmental noise was not present during evolution so SO in that respect fulfils the role of the envelope of noise. It has not been conclusively shown that it is an alternative for base-set exclusivity, nor that the results found in simulation will translate effectively to a real reality-gap problem. However, this interesting question will be the focus of immediate future work.

In other further work, the extent to which this result will generalise across different neuroevolutionary problems, both inside and beyond Evolutionary Robotics, should be the first point of investigation. In addition, the relationship between the complexity (or sensitivity) of the evolutionary problem in question and the degree to which ontogenic noise can be applied should be elucidated. No decrease in performance was observed in the present work for even the most noisy genotype–phenotype mapping (where $\sigma^2 = 0.2$); understanding the limits in particular problem instances will shed light on the complexity of the task itself.

Finally, since it is relatively cost-free to add an SO mechanism into existing neuroevolutionary platforms, researchers may wish to consider exploring this dimension alongside

others in their own work.

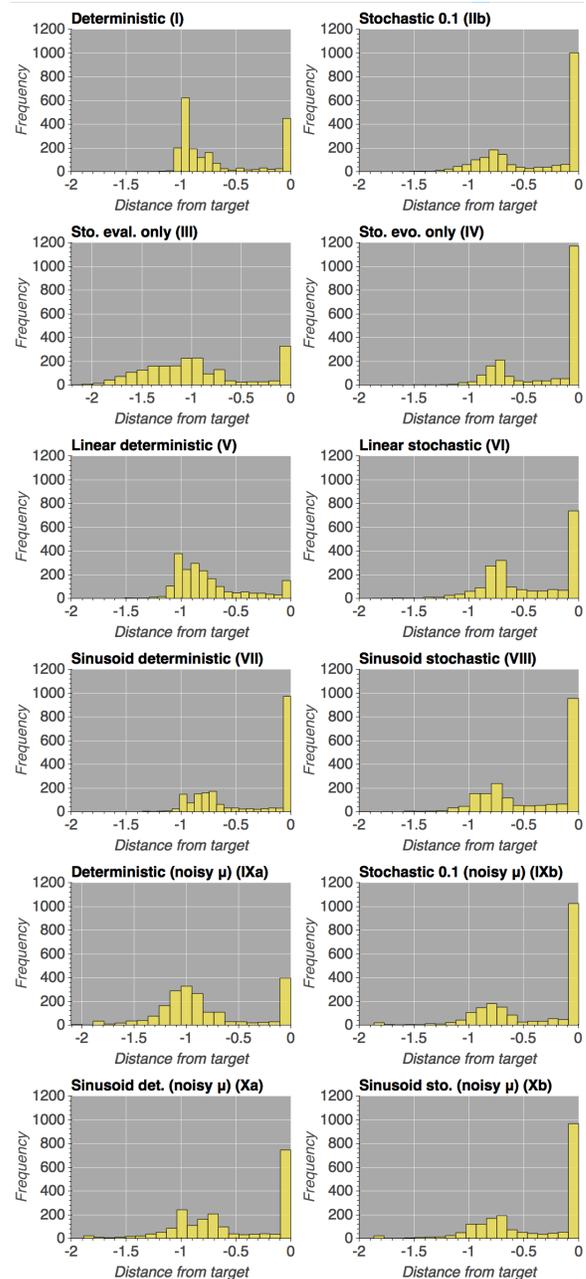


Figure 6: Histograms showing distribution of final fitnesses of individuals for treatments that involve ontogenic noise during evolution, evaluation, neither, or both. Treatments using deterministic ontogenesis are shown in the left-hand column; treatments using stochastic ontogenesis are shown in the right-hand column. A clear shift towards higher fitnesses (distances from target) in all runs, as well as many more high-scoring runs, can be seen in the stochastic treatments. A clear suppression of fitness is observed in treatments with uncertain friction coefficients (noisy μ) and deterministic ontogenesis.

Acknowledgements

Computational resources for this work were provided by the Evolutionary Systems Research Group, the School of Computing and Mathematics, and the Faculty of Natural Sciences, all at Keele University. I would like to thank Peter Andras, James Borg, Alastair Channon, Jennifer Cook, Charles Day, and Jared Moore for helpful discussions and suggestions during preparation of the paper.

References

- Bongard, J. and Lipson, H. (2004). Once more unto the breach: Co-evolving a robot and its simulator. In *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, pages 57–62.
- Borg, J. M., Channon, A., Day, C., et al. (2011). Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In *Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, pages 101–108. MIT Press.
- Branke, J. (2001). Reducing the sampling variance when searching for robust solutions. In *Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation, GECCO'01*, pages 235–242, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Cliff, D., Husbands, P., and Harvey, I. (1993). Explorations in Evolutionary Robotics. *Adaptive Behavior*, 2(1):73–110.
- Fernandez Leon, J. A. (2011). *Behavioural robustness and the distributed mechanisms hypothesis*. PhD thesis, University of Sussex.
- Fogel, B. J., Nitowsky, H. M., and Gruenwald, P. (1965). Discordant abnormalities in monozygotic twins. *The Journal of pediatrics*, 66(1):64–72.
- Goosen, A., Brule, R., Janssen, J., and Haselager, W. (2007). Interleaving simulated and physical environments improves evolution of robot control structures.
- Gordon, A. J. E., Halliday, J. A., Blankschien, M. D., Burns, P. A., Yatagai, F., and Herman, C. (2009). Transcriptional infidelity promotes heritable phenotypic change in a bistable gene network. *PLOS Biology*, (2).
- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive behavior*, 6(2):325–368.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer.
- Koos, S., Mouret, J.-B., and Doncieux, S. (2013). The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*, 17(1):122–145.
- Miconi, T. and Channon, A. (2005). Analysing co-evolution among artificial 3D creatures. In *International Conference on Artificial Evolution (Evolution Artificielle)*, pages 167–178. Springer.
- Moore, J. and Stanton, A. (2017). Lexicase selection outperforms previous strategies for incremental evolution of virtual creature controllers. In *Proceedings of the 14th European Conference on Artificial Life*, Lyon, France. MIT Press.
- Mouret, J.-B. and Chatzilygeroudis, K. (2017). 20 years of reality gap: a few thoughts about simulators in evolutionary robotics. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1121–1124. ACM.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. The MIT Press.
- Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168.
- Seth, A. K. (1998). Noise and the pursuit of complexity: A study in evolutionary robotics. In *European Workshop on Evolutionary Robotics*, pages 123–136. Springer.
- Smith, R. et al. (2005). Open Dynamics Engine.
- Stanton, A. and Channon, A. (2013). Heterogeneous complexification strategies robustly outperform homogeneous strategies for incremental evolution. In *Proceedings of the 12th European Conference on Artificial Life*, pages 973–980, Taormina, Italy.
- Stanton, A. J. (2017). *Simultaneous incremental neuroevolution of motor control, navigation and object manipulation in 3D virtual creatures*. PhD thesis, Keele University.
- Wong, A. H., Gottesman, I. I., and Petronis, A. (2005). Phenotypic differences in genetically identical organisms: the epigenetic perspective. *Human molecular genetics*, 14(suppl_1):R11–R18.
- Zagal, J. C., Ruiz-del Solar, J., and Vallejos, P. (2004). Back to reality: Crossing the reality gap in evolutionary robotics. *IFAC Proceedings Volumes*, 37(8):834–839.