

# Supervision and Evolution: Pretraining Neural Networks for Quadrupedal Locomotion

Jared M. Moore<sup>1</sup> and Anthony J. Clark<sup>2</sup>

<sup>1</sup>School of Computing,

Grand Valley State University, Allendale, MI, USA

<sup>2</sup>Computer Science, Pomona College, Claremont, CA, USA

jared.moore@gvsu.edu, anthony.clark@pomona.edu

## Abstract

Neural networks (NNs) are effective controllers for evolutionary robotics, imposing few limits on potential gaits. Morphology evolved with a controller enables brain and body to become tightly coupled. Typically, NN parameters (sometimes architectures) and animat bodies are randomly initialized at the start of evolution. In this paper, we pretrain NNs with supervised learning, bootstrapping NN outputs towards oscillating behaviors prior to evolution. We focus on quadrupedal gaits as they are well-studied in biology and several common gait patterns have been identified, named, and studied by the research community. We hypothesize that performance of evolved gaits will improve with pretraining compared to beginning evolution with randomly initialized NNs. Our results show that only some pretraining regimens outperform (in terms of distance traveled and viability) random initialization of NN parameters. Furthermore, some regimens introduce an initial bias that is difficult to overcome, resulting in better initial performance but worse performance in the long term.

## Introduction

Quadrupedal gaits follow specific rhythmic patterns coordinating the movement of multiple limbs and joints to realize locomotion. Body composition influences the type of gait such as hopping in kangaroos (Alexander and Vernon, 1975), galloping in horses (Alexander, 1988), and undulating swimming in salamanders (Ijspeert et al., 2005). Morphological characteristics, such as tendon elasticity (Alexander, 1988; Geyer et al., 2006) and flexible spines (Culha and Saranli, 2011), aid in performance and efficiency. Although animals refine their gaits throughout their lives—adjusting to changes in morphology due to growth and aging—the basic underlying movements are innate. Evolutionary robotics (ER) approaches, however, often begin with a bio-inspired morphology and randomly initialized controllers. Randomly generated NNs face the initial hurdle of mapping a periodic input signal to coordinated oscillating movements across many joints enabling effective locomotion.

In this paper, we pretrain NN controllers with a supervised learning algorithm and then continue optimization by coevolving morphology with control. Prior to evolution, supervised learning trains NN outputs to match periodic oscillating

signals hand-designed to simulate the motion of joints observed in several quadrupedal gaits. Pretraining enables the evolutionary process to begin with controllers predisposed to coordinated oscillatory motion. We begin with a comparison between NNs seeded with random parameters and NNs seeded with pretrained parameters in a many-objective task using Lexicase selection (Spector, 2012). For these initial experiments, we pretrain NNs to match simple sinusoidal patterns that do not directly correspond to a standard gait. Next, we explore evolving NNs pretrained on quadrupedal gaits observed in animals. Our key hypotheses are that 1) pretraining will result in higher performing quadrupedal gaits in terms of distance traveled and efficiency, and 2) pretraining speeds up the evolutionary process by alleviating the occurrence of NNs in early generations where input signals do not bring about oscillation in NN outputs.

Our results show that some pretraining regimens evolve effective locomotion more rapidly and outperform random initialization on some objectives. Effective locomotion evolves across all treatments, including the non-pretrained baseline, resulting in a variety of morphologies. Figure 1 shows a sample of evolved quadruped animats. Further, the mechanical feasibility of evolved gaits (i.e., how easily a gait could be implemented with physical parts and motors) for pretraining exceeds that of many randomly initialized NNs.

## Background and Related Work

Evolutionary robotics (ER) (Floreano et al., 2008) leverages digital simulation and high-performance computing evolving animats with genetic algorithms. ER approaches have demonstrated effective locomotion in quadrupeds (Clune et al., 2009), hexapods (Pretorius et al., 2019), and soft robots (Cheney et al., 2013), among others. Coevolving morphology and control leads to effective systems where brain and body are highly intertwined (Paul and Bongard, 2001; Hornby and Pollack, 2001). For locomotion, preevolving NNs, that is evolving the brain to elicit oscillation before embodiment, has elicited oscillatory behavior in legged animats (Stanton and Channon, 2015).




Figure 1: A variety of morphologies evolve across treatments. All quadrupedal animats have three torso segments and three segment legs. Limb lengths, torso dimensions, and joint ranges are some of the morphological aspects that evolve together with the NN controller.

Traditional neuroevolutionary algorithms (e.g., NEAT (Stanley and Miikkulainen, 2002)) are single objective, typically focusing on a performance metric like distance traveled. Evolving alongside a morphology is possible with modification, but the task remains single objective (Moore and McKinley, 2017). In living organisms, conversely, multiple factors are used to assess gait performance. For example, galloping is the fastest quadrupedal gait, but incurs high ground reaction forces, which can lead to degradation. Whereas, canters and trots reduce ground impact forces on the body but require more metabolic power output (McMahon, 1985). Observations in robotics have shown that flexible spines reduce the vertical center of mass (COM) movement improving efficiency (Ackerman and Seipel, 2013). More directly, Sellers et al. (2003) evolved bipedal walkers with metabolic efficiency as the fitness objective evolving effective, efficient locomotion. Drawing upon these observations, we evolve animats in a many-objective evolutionary algorithm incorporating traditional fitness objectives like distance traveled while also including objectives for efficiency.

Outside of neuroevolution, NNs are frequently pretrained or updated online. For example, Erhan et al. (2010) used an unsupervised learning process to improve the performance of gradient-based training. More recently, transfer learning (using a pretrained model on a new, different task) and fine-tuning (a training technique in which only some model parameters are updated) have been shown to dramatically decrease training time (Howard and Gugger, 2020). Reinforcement learning is a commonly used technique for developing walking gaits in real-world robots. Even state-of-the-art methods,

however, require hours of execution to develop a working gait (Haarnoja et al., 2019).

Lexicase selection (Spector, 2012) is a many-objective selection operator for evolutionary algorithms. Lexicase selection evaluates individuals one objective at a time, returning the best individual from a sample if it outperforms others in the sample on the objective under consideration. If two or more individuals are tied, another objective is randomly sampled and all tied individuals are evaluated again. If all objectives are exhausted, a random selection of the remaining individuals is performed. Epsilon-Lexicase ( $\epsilon$ -Lexicase) selection (La Cava et al., 2016) modifies the Lexicase selection mechanism for real-valued fitness metrics where “close” performance between individuals might mean there is not an appreciable performance difference. For example, all individuals within 95% of the best performer are considered to be tied with respect to the given objective. Two modifications to Lexicase facilitate use in real-valued fitness tasks and improve computational efficiency in ER. First,  $\epsilon$ -Lexicase selection is effective in ER where small differences in fitness do not translate to substantially different gait performance (Moore and McKinley, 2016; Moore and Stanton, 2018, 2019). Second, down-sampled Lexicase selection (Helmut and Spector, 2020) reduces the number of objectives considered per selection event by randomly sampling only a subset of the objectives under consideration during the selection process. This can significantly reduce computation time while not significantly hindering performance in practice.

## Methods

**Quadrupedal Animat** The quadrupedal animat consists of a three-segment torso and four three-segment legs. Torso segments are connected by 2 degree-of-freedom (DOF) joints allowing for rotations in the  $y$  (side-to-side) and  $z$  (up-and-down) axes. Legs are attached to the front and rear torso segments. The hip, knee, and ankle joints are also 2-DOF joints allowing for movement along the long axis of the animat as well as to move away from the midline. Each foot has a touch sensor providing information on whether a foot is in contact with the ground. Figure 1 highlights some morphologies that evolve in experiments presented in this paper.

The morphological component of an animat's genome comprises 69 real-value numbers. The first gene codes for spine type: rigid, flexible hinge, or actively controlled hinge. The second codes for the lowest leg joint: rigid, flexible slider, or actively controlled hinge. Joints are constrained by a maximum joint velocity gene, 14 genes for maximum exertable joint force, 28 genes specifying upper and lower joint limits, and 8 genes specifying flexibility of the spine and lowest limb joints. Note that some genes are not expressed in an animat depending on the type of joints coded for the spine and lowest leg joint in the genome. Four genes specify the initial rotation of the upper and mid-leg segments. Finally, five genes specify the torso dimensions (two for width and length of front and rear segments and one for mid-torso length) and six for the length of the upper, mid, and lower limb segments. Front and rear legs are grouped together in the genome enforcing left/right morphological symmetry.

**Controller** Controllers consist of a fully-connected feed forward NN with 33 inputs and 28 outputs. Inputs to the NN include: one periodic oscillating signal with an evolved frequency, four touch sensors (one per foot), and 28 joint angle sensors (one for each DOF). There are 28 outputs, one for each DOF on the quadrupedal animat. The NN is implemented in PyTorch (Paszke et al., 2019) with four hidden layers of 16 nodes each. The hidden layers have sigmoid activation functions while the output layer is a tanh allowing for output values in the range of -1 to 1. An animat's controller evolves by mutating NN weights as well as the frequency of the periodic oscillating signal provided to the NN.

**Neural Network Pretraining** NNs are trained prior to evolution with the Adam algorithm (Kingma and Ba, 2017) with a batch size of 16, learning rate of  $1 \times 10^{-3}$  and Mean-Squared Error loss function. Training proceeds for 1,000 epochs. The training goal is for each output of the neural network to match a sine wave with frequency of 1.0 sampled 30,000 times over a 10 second period. This frequency matches the update frequency of the physics simulation. Some treatments alter the phase offset of the sine wave for different outputs, those are detailed next.

**Treatments** Treatments *RandInPT*, *ZeroInPT*, and *JointFbPT* examine different possible simulated input combinations during pretraining and are compared against a baseline with no pretraining (*NoPT*). In *NoPT*, NN parameters are randomly initialized (i.e., no pretraining) and evolved along with the morphology of the animat. *RandInPT* (denoting pretraining with randomized inputs to all inputs aside from the oscillating input) pretrains the NN to match a series of oscillating outputs assigned to each joint. The target oscillating output of the knees is one phase offset of the hips resulting in a pretrained gait where the knees and hips move opposite of each other. A periodic oscillating input signal with frequency of 1.0 is sent to the first input of the neural network. The four touch sensors receive random inputs of either 0 or 1 while the 28 joint position sensors receive random inputs in the range of  $-1.0$  to  $1.0$ . *ZeroInPT* has the same pretraining strategy as *RandInPT* except all inputs but the oscillating input are sent zeroes. Pretraining with all zeros would indicate that all sensors (touch and joints) are reading as zero. In *JointFbPT*, each joint's desired oscillation is offset  $1/28 * phase$  of the next and previous joint—simulating that the joint sensors report that each joint is in motion as directed. *JointFbPT* emulates robot joints that precisely follow NN commands, simulating conditions of smooth locomotion with all joints moving through their range and feeding back into the NN.

Treatments *AKneesPT*, *DiagCoordPT*, *S3LegsPT*, and *3GaitsPT* are pretrained to match specific gait patterns instead of a simple sinusoidal pattern. *AKneesPT* is trained with all four legs moving symmetrically but the knees being out of phase of the hips and ankles. It is similar to *RandInPT* but now the desired joint positions are sent back to the NN as inputs. In *DiagCoordPT* each leg's knee is out of phase from the hips but the leg movements are diagonally coordinated. The front-left and rear-right legs move forward while the front-right and rear-left legs move backward. *S3LegsPT*'s gait maintains the knee out of phase of its respective hip with three of the legs moving in phase, while the front-right leg moves out of phase. *3GaitsPT* doesn't pretrain a specific gait, instead the population is seeded randomly with the three gaits in Treatments *AKneesPT*, *DiagCoordPT*, and *S3LegsPT*.

**Evolutionary Algorithm and Objectives** Animats evolve across 20 replicate runs, each with a unique starting seed, over 2,000 generations with a population size of 120. Child animats are formed from two parents through two-point crossover between two parents with crossover rate of 50%, otherwise asexual reproduction occurs. Mutation is applied with a 4% chance of changing a gene's value. Parents are selected using  $\epsilon$ -Lexicase selection with an  $\epsilon$  of 5% and four individuals randomly sampled from the population per selection event. Downsampling is applied with four objectives sampled randomly from seven possible objectives at each generation. Tiebreaks are settled by a random selection of

the remaining individuals that have been evaluated as equal on all objectives under consideration.

Individuals are evaluated on seven fitness objectives: (1) forward distance traveled, (2) euclidean distance traveled, (3) distance per unit of power, (4) vertical center of mass displacement, (5) time until flipped, (6) number of leg direction switches, and (7) number of touches by non-toe body segments. For pure locomotive performance, distance traveled (objective 1) is the measure typically used in ER experiments. However, the additional objectives help bootstrap initial locomotion (objectives 2 and 4), support efficient movement (objectives 3 and 4), and encourage stable locomotion (objectives 4 and 5), while discouraging less practical gaits (objectives 6 and 7).

## Results and Discussion

**Pretraining Input Strategies** Figure 2 plots the best fitness across replicates for the non-pretrained baseline (*NoPT*) as well as the first three pretraining treatments (pretraining without explicitly simulating known gaits). Pretrained NNs have an early generation performance advantage as *RandInPT*, *ZeroInPT*, and *JointFbPT* exhibit rapid increases in distance traveled for the first 500 generations. At generation 500, pairwise performance differences in distance traveled are significant for *NoPT* / *RandInPT*, *NoPT* / *ZeroInPT*, and *NoPT* / *JointFbPT* ( $p < 0.01$ ). (We use the Wilcoxon rank-sum test with Bonferroni correction for all p-values reported in this paper.) By generation 1,000, differences in distances traveled by the best individuals in *NoPT* and *RandInPT*, *ZeroInPT*, and *JointFbPT* are no longer significant.




Figure 2: Distance traveled of the best individual per replicate, per generation across the non-pretrained baseline, and the initial three pretraining strategies not based on quadruped gait patterns. Shaded areas represent the 95% confidence intervals.

Figure 3 plots the farthest traveling individual per replicate after 2,000 generations. After the initial performance boost provided by pretraining, the randomly initialized NNs evolve similar performance. Random initialization is as effective as *RandInPT* and *ZeroInPT*, and competitive with *JointFbPT*. Random initialization (*NoPT*) is the typical strategy

employed for gait evolution in ER, having proven effective in neuroevolution (Clune et al., 2009; Moore et al., 2015). These initial results show that, in terms of gait performance, pretraining aids in the early stages of evolution.




Figure 3: Distance traveled of the best individual per replicate at generation 1999. All treatments evolve statistically similar performance on this fitness metric.

Figure 4 plots the efficiency of the farthest traveling individual over time. Here we note that pretraining also appears to aid in evolving efficient locomotion whereas random initialization incurs an initial lower efficiency before ultimately becoming competitive with the three pretraining treatments. *JointFbPT* is significantly more efficient than *NoPT* up to generation 850, but this advantage is no longer present in the best individual per replicate thereafter.




Figure 4: Locomotive efficiency of the best individual per generation per replicate for the first four treatments over evolution. Shaded areas represent 95% confidence intervals.

In terms of distance traveled and efficiency, the four treatments are similar at generation 1999, however, practicality of evolved gaits is also critical. Figure 5 plots the number of leg direction switches over time for the farthest traveling individual per replicate. Here we see a significant difference ( $p < 0.001$ ) between *NoPT* and the pretraining replicates. While some number of direction switches are required for any locomotion driven by a periodic oscillating signal, *NoPT* far exceeds any other treatment.

Even with the number of direction switches added as a minimization fitness objective in Lexicase selection, some




Figure 5: Number of leg direction switches for the best individual per generation per replicate over evolution. Shaded areas represent 95% confidence intervals. High numbers of direction switches indicate vibrating locomotion instead of walking/running gaits.

replicates in *NoPT* converge towards vibrating locomotion rather than classic quadrupedal gaits. Indeed, 10/20 evolved gaits for *NoPT* exhibit significantly more vibrating locomotive patterns than any of the other treatments. Figure 6 depicts the strategy for half of the evolved animats in *NoPT*. This gait produces locomotion by rapidly oscillating legs inducing vibration rather than smooth walking/trotting gaits. Treatments *RandInPT*, *ZeroInPT*, and *JointFbPT* all exhibit more traditional quadruped gaits albeit with the varied morphologies in Figure 1. Two representative gaits from *JointFbPT* are shown in Figure 7.




Figure 6: A high frequency vibrating gait evolves in 10/20 replicates in *NoPT*. This gait relies on rapid direction changes in the legs rather than smooth oscillating gaits. Here, the blur in the legs shows the range of motion of the quadruped animat as it moves across the world. Note the tight range of motion. This animat executes a total of 4802 direction switches in the legs over the 10 second simulation.

Traditional servo motors are not capable of rapidly changing direction as would be required for vibrating gaits evolved in *NoPT*. Moreover, in all but the smoothest, high friction environments, vibrating gaits would not be able to handle ob-

stacles placed in an animat's path. Figure 8 plots the distance traveled for the farthest traveling individual per replicate in *NoPT* grouped by their gait pattern. Vibrating gaits are significantly lower than regular gaits in terms of distance traveled and efficiency while also having almost double the number of leg switches. Given the disparity in performance between vibrating and regular gaits, we hypothesize that vibrating gaits are an evolutionary trap preventing the evolution of regular oscillating locomotion. Even though the number of leg switches objective is included, it is not enough to escape the area of the search space where vibrating gaits are prevalent.

Initial performance of gaits at generation 0 are still relatively low for the pretrained NNs as the controller and morphology have not yet had time to coevolve. Bodies of the animats are randomly initialized including limb length and joint range of motion. Here, we study what happens if NNs are pretrained for regular locomotion but the morphologies are not predefined. Performance could perhaps improve in initial generations by starting with a predefined quadruped morphology and evolving from that fixed start; we leave this for future study. The relative effectiveness, and realizability of the gaits in the pretraining treatments indicate that pretraining the NN, even with a randomly generated morphology is effective. Oscillatory behaviors are evident in initial generations whereas the baseline randomly initialized NNs saturate and assume fixed postures. The new sensor information alters the pretrained behavior of the brains.

**Pretraining Gait Patterns** Treatments *AKneesPT*, *DiagCoordPT*, *S3LegsPT*, and *3GaitsPT* pretrain NNs based on gait patterns observed in quadrupeds. Figure 9 plots the fitness of the best individual over time across replicates for *AKneesPT*, *DiagCoordPT*, *S3LegsPT*, and *3GaitsPT* as well as *NoPT* and *JointFbPT* for comparison. Results of these treatments are similar to those observed in the initial pretraining treatments with an initial rapid increase in distance traveled. *AKneesPT*, *S3LegsPT*, and *3GaitsPT* are significantly better than *NoPT* up to generation 500, after which only *S3LegsPT* is significantly better up to generation 1,250.

Figure 10 plots the performance of the best individual per replicate after the final generation. Although distance traveled is not significantly different for the baseline versus pretraining, the type of gait that evolves is qualitatively different in observed behavior. Figure 11 plots the number of leg direction switches across the treatments. *NoPT* has significantly more direction switches than all gait pretraining treatments due to the prevalent vibrating locomotion pattern noted earlier. *3GaitsPT* does have two outliers exhibiting vibrating locomotion, but the rest resemble traditional quadruped gaits. We hypothesize that the outliers in *3GaitsPT* evolve as the populations are seeded with three different pretrained gaits. Crossover is not restricted to animats with similar gaits, therefore, it is likely that this treatment sees a destructive series of crossovers breaking the pretrained behaviors in the NNs.




Figure 7: Two gaits that evolve in *JointFbPT*. (Top) Predominant gait that evolves across treatments in this study. The rear legs drive movement while the front legs maintain stability as the rear legs move forward. This gait maintains a relatively stable center of mass minimizing the COM vertical movement. (Bottom) A four legged slow gallop with front and rear legs moving opposite of each other.




Figure 8: Distance traveled for the farthest traveling individual per replicate in *NoPT*. Animats with regular gaits are significantly better than those that exhibit vibrating gaits.




Figure 9: Distance traveled for the best individual per generation per replicate of the final four treatments as compared to the randomly generated NNs and *JointFbPT*. Shaded areas represent the 95% confidence intervals.

The pretrained NNs for *AKneesPT*, *DiagCoordPT*, and *S3LegsPT* are based on four legged gait patterns observed in nature. This introduces a potential bias in the populations as all individuals are likely predisposed to variations of these gait patterns due to pretraining. From an exploratory perspective this is limiting, but from an exploitative/engineering perspective it might help to focus the search around a specific gait for a given treatment. While observing evolved gaits in




Figure 10: The farthest traveling individual per replicate after evolution for the final four treatments do not significantly outperform *NoPT* and *JointFbPT*.




Figure 11: The gaits that evolve in the pretraining treatments exhibit regular periodic oscillating locomotion except for two replicates in *3GaitsPT* which evolve vibrating gaits.

the best individuals both (a) gaits related to the pretrained NN pattern and (b) other gaits evolve. Starting from an initial gait does not completely prevent evolution from discovering other gait patterns.

In terms of performance, the farthest traveling individuals from *AKneesPT*, *DiagCoordPT*, and *S3LegsPT* do not outperform the other treatments. Although *S3LegsPT* evolves the

farthest traveling individual across all treatments, *AKneesPT* and *DiagCoordPT* are among the lowest performing treatments. The mixed results across these three treatments suggest that pretraining from biological gaits might not be the best approach for robotic systems or easing the transition between pretraining and evolution might be necessary. *JointF-bPT* has higher performance than *AKneesPT* and *DiagCoordPT* yet its pretraining was closer to a randomly generated pattern. It may be that to fully exploit natural gaits, the initial morphology of the animats must more closely resemble quadrupeds that exhibit those gaits in nature. However, we leave this to future investigation.

## Conclusions and Future Work

Living animals do not learn how to walk from scratch. Innate reflexes, preflexes, and instincts finely tuned to morphologies, enable many animals to walk within hours of birth. Taking this cue from nature, here we evolve walking animats with pretrained NNs such that they do not have to evolve simple joint motions from random initialization. Different pretraining regimens are explored including random noise, all zeroes, and simulated joint feedback to NN inputs. Although pretraining doesn't produce significantly higher performance in this study across all seven fitness objectives, the animats do evolve effective locomotion in fewer generations and the quality of the gaits is objectively more realistic. Without pretraining, many evolved animats exhibit a vibrating locomotion switching leg directions thousands of times in a 10 second simulation. This behavior would likely damage traditional robot actuators. The pretraining configuration also does not appear to significantly effect results as performance is similar across all pretraining treatments. Training oscillating output from the periodic input is apparently the main factor in the performance of evolved animats rather than pretraining to a specific quadruped gait pattern.

In future work we plan to examine improvements to the pretraining process, explore alternate NN architectures, and further investigate evolutionary dynamics of Lexicase selection. First, we note that morphologies are randomly initialized and the NNs are pretrained with random inputs for footfalls. Kinematic simulations might aid in pretraining by more accurately simulating foot-ground contact and joint angles. The morphologies have also been randomly initialized, performance might be improved if the quadruped bodies were initialized with parameters used during kinematic modeling. Second, NN topology is fixed across treatments, only weights are optimized toward oscillating behaviors during pretraining. Neuroevolutionary approaches like NEAT have shown the effectiveness of evolving topology along with weights. We will investigate if tuning network topology in addition to pretraining weights increases performance and efficiency. Third, fitness objectives in this paper are intended to improve both performance in terms of distance traveled as well as the efficiency of the evolved gaits. Some of these objectives

might hinder performance to increase efficiency. Changing the objectives, or selectively removing some from consideration at different points in the evolutionary process might result in more effective locomotion. We plan to investigate the impact of different combinations of objectives on performance of evolved individuals and scheduling of objectives with in Lexicase.

## Acknowledgements

The authors would like to thank William La Cava, Thomas Helmuth, Edward Pantridge, Emily Dolson, and Lee Spector for their conversations about various aspects of Lexicase selection. We would also like to thank Adam Stanton for his thoughts on evolving body and brain of robotic systems along with his contributions to applying Lexicase selection to evolutionary robotics. Thanks to Pomona College and Grand Valley State University for supporting this research.

## References

- Ackerman, J. and Seipel, J. (2013). Energy efficiency of legged robot locomotion with elastically suspended loads. *IEEE Transactions on Robotics*, 29(2):321–330.
- Alexander, R. and Vernon, A. (1975). The mechanics of hopping by kangaroos (macropodidae). *Journal of Zoology*, 177(2):265–303.
- Alexander, R. M. (1988). Why mammals gallop. *American Zoologist*, 28(1):237–245.
- Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, pages 167–174, Amsterdam, The Netherlands. ACM.
- Clune, J., Beckmann, B. E., Ofria, C., and Pennock, R. T. (2009). Evolving coordinated quadruped gaits with the HyperNEAT generative encoding. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 2764–2771, Trondheim, Norway. IEEE.
- Culha, U. and Saranlı, U. (2011). Quadrupedal bounding with an actuated spinal joint. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1392–1397, Shanghai, China.
- Erhan, D., Courville, A., Bengio, Y., and Vincent, P. (2010). Why does unsupervised pre-training help deep learning? In Teh, Y. W. and Titterton, M., editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 201–208, Chia Laguna Resort, Sardinia, Italy. JMLR Workshop and Conference Proceedings.
- Floreano, D., Husbands, P., and Nolli, S. (2008). Evolutionary Robotics. In *Handbook of Robotics*. Springer Verlag, Berlin.

- Geyer, H., Seyfarth, A., and Blickhan, R. (2006). Compliant leg behaviour explains basic dynamics of walking and running. *Proceedings of the Royal Society B: Biological Sciences*, 273(1603):2861–2867.
- Haarnoja, T., Ha, S., Zhou, A., Tan, J., Tucker, G., and Levine, S. (2019). Learning to walk via deep reinforcement learning.
- Helmuth, T. and Spector, L. (2020). Explaining and exploiting the advantages of down-sampled lexibase selection. In *The 2020 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, pages 341–349, Montreal, Quebec, Canada.
- Hornby, G. S. and Pollack, J. B. (2001). Body-brain co-evolution using L-systems as a generative encoding. In *Proceedings of the 2001 ACM Genetic and Evolutionary Computation Conference*, pages 868–875, San Francisco, California, USA. Morgan Kaufmann.
- Howard, J. and Guggen, S. (2020). *Deep Learning for Coders with Fastai and PyTorch*. O'Reilly.
- Ijspeert, A. J., Crespi, A., and Cabelguen, J.-M. (2005). Simulation and robotics studies of salamander locomotion. *Neuroinformatics*, 3(3):171–195.
- Kingma, D. P. and Ba, J. (2017). Adam: A method for stochastic optimization.
- La Cava, W., Spector, L., and Danai, K. (2016). Epsilon-lexibase selection for regression. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, pages 741–748, Denver, Colorado, USA. ACM.
- McMahon, T. A. (1985). The role of compliance in mammalian running gaits. *Journal of Experimental Biology*, 115(1):263–282.
- Moore, J. M., McGowan, C. P., and McKinley, P. K. (2015). Evaluating the effect of a flexible spine on the evolution of quadrupedal gaits. In *Proceedings of the 13th European Conference on Artificial Life*, pages 166–173, York, UK.
- Moore, J. M. and McKinley, P. K. (2016). *A Comparison of Multiobjective Algorithms in Evolving Quadrupedal Gaits*, pages 157–169. Springer International Publishing, Aberystwyth, UK.
- Moore, J. M. and McKinley, P. K. (2017). Evolution of joint-level control for quadrupedal locomotion. *Artificial Life*, 23(1):58–79.
- Moore, J. M. and Stanton, A. (2018). Tiebreaks and diversity: Isolating effects in lexibase selection. In *Proceedings of the 16th International Conference on the Simulation and Synthesis of Living Systems*, pages 590–597, Tokyo, Japan. ACM.
- Moore, J. M. and Stanton, A. (2019). The limits of lexibase selection in an evolutionary robotics task. In *The 2019 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, pages 551–558, Newcastle-Upon-Tyne, UK. MIT Press, MIT Press.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Paul, C. and Bongard, J. C. (2001). The road less travelled: Morphology in the optimization of biped robot locomotion. In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 226 – 232, Maui, Hawaii, USA.
- Pretorius, C. J., du Plessis, M. C., and Gonsalves, J. W. (2019). Evolutionary robotics applied to hexapod locomotion: a comparative study of simulation techniques. *Journal of Intelligent & Robotic Systems*, pages 1–23.
- Sellers, W. I., Dennis, L. A., and Crompton, R. H. (2003). Predicting the metabolic energy costs of bipedalism using evolutionary robotics. *Journal of Experimental Biology*, 206(7):1127–1136.
- Spector, L. (2012). Assessment of problem modality by differential performance of Lexibase selection in genetic programming: A preliminary report. In *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*, pages 401–408, Philadelphia, Pennsylvania, USA. ACM.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.
- Stanton, A. and Channon, A. (2015). Incremental neuroevolution of reactive and deliberative 3d agents. In *Proceedings of the 13th European Conference on Artificial Life*, pages 341–348, York, UK.