

Stability-Driven CPG Design for More Effective Quality-Diversity in Evolutionary Robotics

Matthias Pex
Matthias.Pex@UGent.be
Ghent University – imec
Ghent, Belgium

Dries Marzougui
Dries.Marzougui@UGent.be
Ghent University – imec
Ghent, Belgium

Francis wyffels
Francis.wyffels@UGent.be
Ghent University – imec
Ghent, Belgium

Abstract

Evolutionary robotics commonly employs quality-diversity algorithms to discover diverse, high-performing behaviours. Meanwhile, central pattern generators provide a compact neural dynamical model for robotic locomotion control. Their differential-equation-based formulation induces smooth, rhythmic actions, while also reducing the dimensionality of the search space. However, the widely used Kuramoto-inspired central pattern generators exhibit nonlinear dynamics, for which convergence depends sensitively on network connectivity and coupling parameters. When convergence is slow or unstable, controller evaluations become noisy, which substantially degrades quality-diversity performance. This work introduces an analytical criterion to assess and guarantee the convergence properties of Kuramoto-based central pattern generators prior to optimisation or simulation. The method yields principled guidelines for selecting connectivity structures and coupling strengths based on the eigenvalues of the Laplacian of the connectivity graph. Experiments on a simulated brittle star robot demonstrate that adhering to this criterion, significantly improves the efficiency and discovery of coordinated gaits. The criterion (a) predicts stable convergence from network structure alone, (b) enhances both the breadth and quality of exploration, and (c) scales naturally to morphologies of arbitrary complexity.

CCS Concepts

• Computing methodologies → Evolutionary robotics; • Mathematics of computing → Evolutionary algorithms; Nonlinear equations.

Keywords

Evolutionary Robotics, Central Pattern Generators, Evolvability, Quality Diversity, Parameter Tuning, Kuramoto Model

ACM Reference Format:

Matthias Pex, Dries Marzougui, and Francis wyffels. 2026. Stability-Driven CPG Design for More Effective Quality-Diversity in Evolutionary Robotics. In *Genetic and Evolutionary Computation Conference (GECCO '26)*, July 13–17, 2026, San Jose, Costa Rica. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3795095.3805096>




Figure 1: Research overview. Comparison between CPG controllers for brittle star robots, optimised via MAP-Elites. The top pathway (dotted arrows) qualitatively shows the obtained diversity when no prior stability criterion is used, in this case resulting in a slowly converging, weakly coupled CPG. The bottom pathway (solid arrows) involves iterating over the controller connectivity by means of spectral graph theory to obtain faster, stably converging CPGs. We show that proper design of the controller can significantly improve the coverage of a MAP-Elites archive, hence impacting the discovery of highly coordinated behaviours.

1 Introduction

Locomotion is essential for a robot. It is achieved via a rhythmic activation of many different actuators. For meaningful, coordinated gaits, all those individual components need to synchronise together. Numerous controller designs have been designed to solve this problem, including a wide variety of neural connectionist models [54]. Such controllers produce signals through networks of many small components (neurons). The training typically happens via adjusting the topology and coupling between those nodes [37, 48].

A specific type of such neural systems are central pattern generators (CPGs), where every node is an oscillator that generates a rhythmic output [17, 55]. Such oscillators can represent fireflies in a flock, musicians in an orchestra, or muscles in a body [1, 40]. In order for the overarching system to function as intended, CPGs heavily depend on synchronisation between individual oscillators [1]. Thanks to their built-in rhythmic pattern generation, they have been used in a wide variety of robotic applications, both in simulation and in physical prototypes, including legged walking and undulating swimming motions [18, 34, 38, 45, 56, 57].



This work is licensed under a Creative Commons Attribution 4.0 International License. *GECCO '26, San Jose, Costa Rica*
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2487-9/2026/07
<https://doi.org/10.1145/3795095.3805096>

Getting such systems to synchronise is no straightforward assignment. Convergence may be either too slow for the system to ever stabilise, or so fast that instabilities arise. They can settle in one single coherent state, or get lost in a fractioned assembly of incoherent subsystems [20, 42]. The study of coupled, non-linear systems is immensely complex, often even more so when scaled up to large systems with many oscillators. Yet, inconsistent convergence can result in noisy fitness evaluations, and hence greatly diminish the quality of evolutionary optimisation algorithms.

In this research, we present clear design guidelines for systems of nonlinear, coupled oscillators. We propose a simple design criterion that is based on spectral graph theory, more specifically, by calculating the eigenvalues of the Laplacian of the connectivity graph. We prove the added value of the criterion through simulation and optimisation of a bio-inspired brittle star robot (Figure 1). MAP-Elites is used as an optimisation algorithm, since it favours diverse exploration of highly coordinated gaits [30, 39]. The eigenvalue-based criterion can be assessed before even running a single simulation. This gives other researchers a powerful tool for properly designing controllers for a wide range of control architectures and morphologies in evolutionary robotics.

2 Background and Related Work

2.1 Evolutionary Optimisation

Evolutionary optimisation is a very powerful, black box algorithm that does not put constraints on the parameters to optimise [24, 35]. This makes the technique ideal for optimising controllers that do not support differentiability, as is often the case for CPGs [5]. As a result, evolutionary robotics has gained increasing interest in the field of brain-body co-optimisation [13, 32]. Additionally, it enables straightforward curriculum learning approaches where the robot becomes increasingly more complex, mimicking phylogenetic evolution [7, 19]. Lastly, studying robots from an evolutionary perspective can provide insights into hypotheses regarding biological evolution, further closing the gap between biology and robotics [19, 25, 43].

Successful evolvability is by itself an extensive area of research, that studies aspects like the robustness of genotypes and phenotypes, the ruggedness of the fitness landscape, and the existence of neutral zones [2, 3, 46, 49, 50]. With respect to evolution as an optimisation algorithm, one of the most important prerequisites is determinism and reproducibility of the evaluations [15]. In other words, repeated evaluations of the same genotype should result in similar behavioural descriptors and fitnesses. This is important for consistent search and navigation of the fitness landscape. Additionally, a sufficiently smooth evolution landscape is implicitly assumed in techniques such as natural evolution strategies, which are based on fitness gradient estimations [52]. Rugged landscapes can cause fitness valleys, which evolution strategies cannot cross, or fitness peaks that cause convergence to greedy solutions.

2.2 Quality Diversity

Classical evolutionary optimisation favours solutions that perform well on prespecified objectives. Such techniques typically give only one solution and can suffer from premature convergence [39]. However, fitness with respect to predetermined objectives should not be

used to guide the search, but to recognise good solutions when they occur [13]. Quality diversity (QD) addresses this issue by favouring exploration in the behavioural space [9, 39]. This approach is known to provide a repertoire of diverse behaviours and controllers that are robust to various situations [11]. Additionally, rewarding novelty can provide stepping stones towards well-performing solutions in further regions of the optimisation landscape [26].

2.2.1 MAP-Elites. One of the most popular quality diversity algorithms is MAP-Elites [30]. This technique fills a grid with a select few behavioural descriptors, promoting diversity in a limited number of dimensions. It is a conceptually simple algorithm that performs well thanks to the fact that it tries to optimise many descriptor cells in parallel, where offspring of one cell can overflow to neighbouring cells [31, 47].

Keeping track of a repertoire of elites has proven very powerful in discovering diversity in the behavioural space, yet constructing and maintaining a repertoire has some disadvantages: (1) It is computationally intensive, since many evaluations will not outperform already present elites, and hence these evaluations do not impact the next generations of optimisation; (2) The introduction of new hyperparameters regarding the design and discretisation of the behavioural descriptors requires additional fine-tuning and optimisation; (3) Noisy evaluations can cause cells to be filled with ‘lucky elites’, which obtain a high fitness score or new behavioural descriptor values, but by means of a genotype that suffers from poor reproducibility and robustness [16].

2.2.2 Uncertain Quality Diversity. The subfield of Uncertain Quality Diversity focuses on the problem that noisy evaluation can generate unreliable archives and inferior evolvability. Such algorithms will also evaluate the fitness of candidate solutions by taking into account the reproducibility of the obtained performance and descriptor metrics [14, 16]. The major drawback of this technique is that it relies on repeated evaluation of genotypes. However, robotic simulations are computationally expensive and the MAP-Elites algorithm already suffers from computational inefficiency, as described in Section 2.2.1. Therefore, it is advisable to not just rely on algorithms that compensate for uncertain evaluations, but to actually prioritise the design of a stable system in the first place. Stable system design and robust optimisation algorithms are perfectly compatible, and one should not substitute the other.

2.3 Central Pattern Generators

Central pattern generators (CPGs) are small, autonomous neural circuits that generate rhythmic patterns without relying on sensory input [27]. They have been observed in many animals, ranging from leeches to humans, and produce rhythmic motor patterns for locomotion, breathing and circulation [12, 27, 36]. In humans, central pattern generators are typically located in the spinal cord, where they can easily be modulated by the brain to adjust the frequency. Even though they do not require sensory stimulation to generate rhythmic patterns, appropriate sensory feedback improves coordination, even in machine learning setups [6]. Thanks to these small clusters of neurons, certain vertebrates have shown to continue proper gait patterns, even without the presence of a brain [29].

CPGs have drawn the interest of robotics researchers thanks to their ability to reduce the complexity of the action space by intrinsically encoding the rhythmicity. Additionally, the differential-equation-based formulation induces smooth actions. Indeed, Ijspeert et al. (2007) have shown that it is possible to coordinate both walking and swimming in a salamander robot by simply modulating the frequency, amplitude, offset and relative phase biases of the oscillatory signals [18]. Despite their open-loop design, it is possible to chain them with multilayer perceptrons, which take the robot's observations and processes them into modulation signals for the CPG [6]. Doing so can simplify learning multi-skill locomotion [55]. Ijspeert and Spröwitz introduced the following differential equation CPG for robotics [18, 38]:

$$\dot{\phi}_i = \omega_i + \sum_j w_{ij} r_j \sin(\phi_j - \phi_i - \rho_{ij}), \quad (1)$$

$$\ddot{r}_i = a_r \left(\frac{a_r}{4} (R_i - r_i) - \dot{r}_i \right), \quad (2)$$

$$\ddot{x}_i = a_x \left(\frac{a_x}{4} (X_i - x_i) - \dot{x}_i \right), \quad (3)$$

$$\theta_i = x_i + r_i \cos(\phi_i). \quad (4)$$

Where θ_i is the oscillating output signal for every oscillator i and ϕ_i , r_i , x_i encode the state variables for phase, amplitude and offset respectively. Parameters ω_i , R_i and X_i are control parameters for the frequency, amplitude and offset steady-state values. Further, a_r and a_x are constant positive gains. For our research, mainly the weight couplings w_{ij} and desired phase biases ϕ_{ij} are of interest. That is because amplitude and bias consist of ordinary, second order linear differential equations that have been constructed to be critically damped. However, the phase biases represent a system of non-linear, coupled differential equations, strongly resembling the Kuramoto model with equation [23]:

$$\dot{\phi}_i = \omega_i + \frac{K}{N} \sum_{j=1}^N \sin(\phi_j - \phi_i), \quad (5)$$

where ϕ_i are the phase biases between oscillators, ω_i their natural frequencies, K the coupling strength, and N the number of oscillators.

Attention is required for the fact that Kuramoto oscillators have not been designed with the primary goal of robotics. They can virtually be used for any system that involves studying the process of synchronisation, from epileptic seizures to the chirping of crickets in unison [42]. Weakly coupled Kuramoto models benefit from the emergence of local interactions, which allows for standing waves and interesting spatio-temporal dynamics. As coupling increases, systems start to synchronise in unison [53]. Stronger coupling tends to make a system more responsive to modulation and induces faster convergence to a global, coherent state with fixed phase biases. Figure 2 shows the impact of Kuramoto dynamics on the convergence after modulation.

The Kuramoto model, despite its ingenious design and fruitful applications in research, is known as a puzzling system for which the stability is continuously studied in the context of new topologies and edge cases [41]. Spectral graph theory offers techniques for the analysis of the eigenvalues of the Jacobian of the weighted connectivity graph [10]. Jadbabaie et al. (2004) were the first to




Figure 2: The convergence problem in the Kuramoto model. The figure shows the connectivity graph of a weakly and a strongly coupled system in a cobweb-like CPG topology. The node pairs that have converged to their modulated phase bias after 200 solver steps are in grey, non-converged couplings are in dotted red. Networks of nodes where all phase biases have converged to their modulated value get the same colour.

apply this framework to Kuramoto models with arbitrary topologies to provide insights on how to tune the coupling strengths of such a system [20]. More specifically, they note that symmetric graphs with a positive definite adjacency matrix have only real eigenvalues for their Laplacian, that the number of zero-eigenvalues is equal to the number of connected components, and that the first non-zero eigenvalue is a measure for the connectedness, and as a result, the speed of convergence.

This first non-zero eigenvalue λ_2 is often referred to as the algebraic connectivity or spectral norm, a measure for the reactivity and the speed of convergence of a system of coupled oscillators. In robotics simulations, a coherent system means that the system will converge to the steady state as defined by the modulation values. However, an incoherent system might not settle in that steady-state regime, meaning that they either converge too slowly or settle in metastable states. Where those states lie strongly depends on the initial conditions and cannot be controlled. However, incoherent states can have significant impact on the coordination, and hence, the performance and behavioural descriptors of an agent.

Similarly, the induced norm $\|L\|_2$ of the Laplacian of the adjacency matrix is a measure for the largest possible effect operations with the matrix can have on signals. Intuitively, this is the upper bound to prevent numerical instability in the simulations. On this ground, we use the spectral gap λ_2 and the induced norm $\|L\|_2$ to find design boundaries for the topology and coupling strength of the CPGs, in order to obtain responsive, yet stable systems. In other words, we aim to design systems that quickly converge to their coherent state, without being prone to numerical instabilities.

2.4 Brittle Stars

Brittle stars (Figure 3) are marine animals. They use long, flexible limbs for locomotion across the seafloor. From a robotics perspective, brittle stars are interesting agents to simulate. They are capable of both locomotion and grasping [33]. Their morphology can easily

be scaled to high-dimensional observation and action spaces by adding arms or increasing arm length. Despite their limited visual sensing, they can use their environment to push off against irregularities in the terrain [22]. Just like real brittle stars can dispose of limbs when attacked, the redundancy in their number of arms allows to experiment with damage resilience [21]. They have an innate distributed nervous system, which makes them a plausible substrate to test distributed control [28]. And lastly, with respect to their behaviour, their gait patterns have been carefully described regarding the synchronisation between arms and the roles the arms play in propulsion and stabilisation. Despite their pentaradial symmetry, their gait is functionally bilateral, with a central limb and a varying number of active rowing limbs [4].



Figure 3: *Ophiura ophiura*, a type of brittle star. © Hans Hillewaert, licensed under CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>), via Wikimedia Commons (https://commons.wikimedia.org/wiki/File:Ophiura_ophiura.jpg)

3 Methods

This section introduces the conceptual backbone for the experiments that are required to establish the design criterion, as well as the simulation setup and optimisation algorithm used for validation.

3.1 CPGs

3.1.1 Differential equations and solver. The differential equations (DEs), as proposed by Ijspeert et al. (2007), are used to govern the CPG (Eq. 1-4) [18]. The DEs will be solved numerically with the fourth order Runge-Kutta method. Phase ϕ_i , amplitude r_i , offset x_i and output θ_i are state variables of the CPG. All are initialised to zero, except the phases, which are sampled from uniform distribution $\mathcal{U}(-0.01, 0.01)$ to break symmetry. R_i , X_i and ρ_{ij} are considered the independent modulation parameters, defining the amplitude, offset and relative phase biases between all the oscillators in the steady-state regime. They are either sampled randomly or provided by the MAP-Elites optimiser. The weights w_{ij} between the oscillators are characteristic for the graph topology, and are obtained by multiplying the adjacency matrix with a weight coupling K . The remaining CPG and solver parameters get fixed values as highlighted in Table 1. The angular frequency is based on recordings of actual brittle stars. The amplitude and offset gains are adopted from the work of Spröwitz et al. (2008) [38]. The solver timestep has been determined by trial and error.

For assessing the convergence of the CPG, the model is modulated immediately after initialisation. Next, the solver iterates over 200 time steps. Afterwards, the relative fraction of converged couplings is assessed by comparing the relative phase biases $|\phi_i - \phi_j|$ after 200 time steps with the modulated phase bias ρ_{ij} . Figure 2

Table 1: CPG and Solver Parameters

Symbol	Value	Parameter
a_r	20	Amplitude gain
a_x	20	Offset gain
ω	2π	Angular frequency
dt	0.004s	Solver timestep

shows the result of such a convergence analysis. Further, a suggested measure for the speed of convergence could be to look at the time step at which 90% of all couplings has converged to their modulation value. Convergence of the amplitude and offset is governed by an ordinary, linear DE and is considered trivial for the scope of our research. For the modulation values, every oscillator gets their own amplitude R_i and offset X_i . However, if n is the number of oscillators, only $(n - 1)$ phase biases ρ_{ij} are defined. This defines a unique tree where every node has a parent node, with which it tries to obtain a relative phase bias. Phase biases are anti-symmetric, so that $\rho_{ij} = -\rho_{ji}$. Any additional graph edge (coupling) that is added to this tree will be calculated to ensure that there are no conflicting phase biases, which could hinder convergence [18]. In other words, summing the modulated phase biases when walking along a loop through the graph will always result in a net phase bias of zero.

3.1.2 Graph topologies. The graph topologies and corresponding adjacency matrices are limited to ring structures with long, arm-like appendices. Three types of graph topologies are defined:

- Base method: no additional edges are added on top of the basic star-like graph.
- Cobweb method: Additional edges are added to connect corresponding nodes across neighbouring arms.
- Fully connected method: Every node is connected to every other node.

For the experiment, a sweep is done across different morphology sizes, weight coupling strengths and connectivity methods, as listed in Table 2. Figure 4 shows the graph connectivities for our research.

Table 2: Experimental variations in topology

Variable	Values	Variations
Nodes in the ring	3, 4, 5, 6, 7, 8, 9, 10	8
Nodes per arm	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10	11
Weight coupling K	0.05, 0.5, 5, 25, 50, 100, 500	7
Connectivity	base, cobweb, fully connected	3
Total		1848

3.1.3 Spectral Graph Theory. The weighted adjacency matrix W of the phase couplings is defined as

$$W = KA, \quad (6)$$

where A is the adjacency matrix of the graph-like CPG and K is the coupling strength. The Laplacian L of this weighted graph is then:

$$L = D - W, \quad (7)$$

where D denotes the degree matrix. This Laplacian matrix is symmetric positive definite and represents a connected graph, meaning




Figure 4: An illustration of the three different types of graph connectivities on a morphology with a ring size of 5 nodes and 5 nodes per arm

that all eigenvalues of the Laplacian are real, positive, and that the smallest eigenvalue λ_1 is zero. In that case, the Fiedler value, spectral gap, or algebraic connectivity is represented by the second smallest eigenvalue λ_2 , with bigger spectral gaps related to better synchronisability [20]. Additionally, for symmetric positive definite matrices, the largest eigenvalue corresponds to the induced norm $\|L\|_2$ of the Laplacian matrix, and is a measure of the largest action the system can exert, where large induced norms are associated with numerical instabilities in computation. Those spectral gap λ_2 and induced norm $\|L\|_2$ values will constitute the design criterion guidelines.

3.1.4 Statistical analysis. Every experimental variation as listed in Table 2 is repeated 100 times. The percentage of successfully converged trials per experimental condition is computed, i.e. the relative number of trials with 100% convergence of the couplings after 200 time steps. An experimental condition is labelled as successful when either 100% of the iterations were fully converged at time step 200 (strict criterion), or whether only 50% of the iterations were fully converged (loose criterion). Note that unsuccessful setups include both the CPGs that converged too slowly, as well as the unstable setups that diverged due to numerical instabilities.

3.2 Brittle Star Simulation

3.2.1 Morphology and Simulation. To validate the effect of the design criterion, controllers are applied to a custom brittle star simulator that we developed using the MuJoCo physics engine [44]. The morphologies will be constrained to brittle stars with five arms, with an example in Figure 5, similar to how they typically occur in nature. Every arm consists of several segments, that are connected to their parent segment with two hinge joints. Every hinge joint gets assigned to one unique oscillator in the CPG. The oscillators in the central ring of the graphs in Figure 4, are not coupled to actuators and merely serve as synchronisation nodes.

Three different morphologies will be assessed. A small morphology with two segments per arm, a medium morphology with 7 segments per arm, and a large morphology with 20 segments per arm. For every morphology, a CPG will be constructed that is weakly coupled, a CPG with a strong and stable coupling, and a strongly coupled but unstable CPG. Table 3 shows the experimental conditions for the brittle star simulations.

3.2.2 Optimisation space. MAP-Elites will be used as an optimisation algorithm, more specifically the QDax toolsuite implementation [8]. The optimisation space for n oscillators consists of the



Figure 5: Brittle star robot in MuJoCo with 5 arms and 7 segments per arm (medium size).

steady state modulation parameters, being n amplitudes R_i , n offsets X_i and $(n - 1)$ independent phase biases ρ_{ij} . The optimisation space is considered to be isotropic for the purpose of equal exploration in every dimension by adding isotropic noise. Specifically, this means that amplitude and offset sampled parameter will be clipped between $[-1, 1]$ and then rescaled to meaningful modulation parameter ranges, as listed in Table 4. Phase bias values are not constraint. The CPG outputs will again be clipped and rescaled to represent valid joint ranges in the simulator.

The MAP-Elites repertoire is initialised by uniformly sampling parameters in the range $[-1, 1]$. Subsequent generations are generated according to the IsoLine variation framework [47]. The relevant MAP-Elites parameters are displayed in Table 5.

3.2.3 Behavioural descriptors. Descriptors that promote exploration of highly coordinated behaviour are chosen. The first descriptor will be the bilateral-contralateral score D_1 . To compute this, a central limb is identified. This is the leading or trailing limb that is best aligned with the locomotion direction. Next, the angle of the four remaining limbs with respect to the locomotion direction is computed for every time step, with small angles when limb is reaching forward, and large angles when the limb is extending to the rear. For this calculation, the arm is represented by the vector from the base of the arm to the midway point in the arm, as this is most informative for the kinematics [28]. Whether arms are in sync is computed via the cosine similarity of the arm angular velocities across time between arms [51]. Hence, descriptor D_1 is given by:

$$D_1 = 0.25 \cdot (E_{\text{front}} + E_{\text{hind}} + E_{\text{right}} + E_{\text{left}}), \quad (8)$$

where E represents the cosine similarity between the front limbs, the hind limbs, the right side limbs and the left side limbs respectively. If all limbs move in phase, the bilateral-contralateral score will be 1. If the left front and right hind limb move in phase, and in perfect counter phase with the right front and the left hind limb, then the gait is contralateral with a bilateral-contralateral score of -1.

The second descriptor will be the assistive score D_2 , representing the relative number of arms that assist locomotion. For every time step, the resulting ground reaction force per arm is computed. The ground reaction forces of all arms are projected onto the locomotion direction to get the assistive forces, and those are summed across all time steps. The assistive force of every arm is normalised by the maximal assistive force occurring across the morphology. The resulting score will range from 1 (only one arm produces the assistive force) to 5 (all arms produce equal assistive force).

Table 3: An overview of the morphologies, connectivity methods and weight couplings K that will be verified with the brittle star simulator. Figure 6 shows that most values for the spectral gap λ_2 and the induced norm $\|L\|_2$ for the slow, fast/stable and unstable convergence CPGs indeed lie respectively beneath, within and above the ideal range according to the design criterion. Note that some do not lie within the ideal range (solid bounding box in Figure 6), but still within the acceptable range (dotted bounding box).

Morphology	Slow convergence	Fast, stable convergence	Unstable convergence
Small (4 osc/arm)	Base; $K = 5$	Cobweb; $K = 150$	Fully connected; $K = 200$
	$\lambda_2 = 0.45$; $\ L_2\ = 25$	$\lambda_2 = 57$; $\ L_2\ = 1085$	$\lambda_2 = 5000$; $\ L_2\ = 5000$
Medium (14 osc/arm)	Cobweb; $K = 5$	Fully connected; $K = 5$	Fully connected; $K = 50$
	$\lambda_2 = 0.22$; $\ L_2\ = 38$	$\lambda_2 = 375$; $\ L_2\ = 375$	$\lambda_2 = 3750$; $\ L_2\ = 3750$
Large (40 osc/arm)	Fully connected; $K = 0.05$	Fully connected; $K = 1$	Fully connected; $K = 10$
	$\lambda_2 = 10$; $\ L_2\ = 10$	$\lambda_2 = 205$; $\ L_2\ = 205$	$\lambda_2 = 2050$; $\ L_2\ = 2050$

Table 4: CPG parameter ranges

Symbol	Range	Parameter
R_i	$[0, 1]$	Amplitude
X_i	$[-1, 1]$	Offset
ρ_{ij}	$[-\infty, \infty]$	Offset range
θ_i	$[-\frac{\pi}{6}, \frac{\pi}{6}]$	Joint angle

Table 5: MAP-Elites parameters

Symbol	Value	Parameter
σ_{iso}	0.05	Standard deviation for isometric noise
σ_{line}	0.10	Standard deviation for line noise
D_1	$[-1, 1]$	Bilateral-contralateral score range
D_2	$[1, 5]$	Assistive score range
n_{pop}	200	Population size
n_{gen}	150	Number of generations

3.2.4 Fitness function. The fitness function consists of two main components: a reward and a cost term. The fitness is meant to be maximised. The reward is the distance d travelled by the end of the episode. The cost represents the energy consumption, estimated by the sum of all the torques $\tau_{i,t}$ produced by the actuators i over all time steps t . The coefficient α has been set to a custom value of 10 via trial and error to obtain gaits that value both distance travelled and smooth behaviour:

$$\text{fitness} = \frac{d}{1 + \alpha \sum_{i,t} \tau_{i,t}}. \quad (9)$$

3.2.5 Statistical analysis. For every full optimisation the following metrics are stored for every generation:

- (1) Coverage: the fraction of cells in the MAP-Elites archive that contain an elite.
- (2) QD-Score: the quality-diversity score is the sum of all the fitnesses that are stored in the archive.

Every full optimisation is repeated 20 times. The evolution of the above metrics can be plotted across generations. Via bootstrapping with 1000 repetitions per generation, the 95% CI can be constructed for every experimental condition. The CPGs with a strong, stable coupling, thus satisfying the design criterion, are compared with

both the CPGs with a weak coupling and the CPGs with an unstable coupling. Whether the curves of the coverage and QD-score are significantly higher for CPGs that satisfy the design criterion, is assessed via a one-sided t-test ($p=0.05$) on the area under the curve (AUC).

4 Results

This section focusses on three main questions:

- Does adding more couplings between oscillators provide a reliable way to improve the speed of convergence to a global, coherent state? (Section 4.1)
- Does there exist a range for the spectral gap λ_2 and the induced norm $\|L\|_2$ for which convergence to a coherent state is guaranteed? (Section 4.2)
- How does obeying this design criterion impact the quality and diversity of behaviour in evolutionary robotics? (Section 4.3)

The source code of the experiments is available at <https://github.com/matthiaspex/CPG-QD.git>.

4.1 Impact of connectivity on convergence

The first goal is to gain insights into how the connectivity approach impacts the convergence across a wide range of morphologies. Table 6 highlights the importance of the number of couplings that are present. The base method has significantly less couplings per node, the cobweb method slightly more, and the fully connected method has the maximum number of possible couplings. Apparently, the more nodes are coupled, the bigger the total number of nodes can get while still allowing convergence. Even stronger results are found for the fully connected graphs, where no morphological sizes have been found for which convergence was impossible.

4.2 Eigenvalue-based design criterion

The second goal is to define a proper convergence criterion. A prerequisite of this, is that it should be possible to verify the criterion prior to running optimisations and robotic simulations. Thanks to spectral graph theory, we know that the spectral gap λ_2 and the induced norm $\|L\|_2$ are respectively measures for the speed of convergence and the susceptibility for numerical instabilities. Both can be calculated from the weighted adjacency matrix. Figure 6

Table 6: Overview of how the connectivity influences the successful convergence at time step 200. As soon as a weight coupling K exists for which the chosen connectivity fully converges, the technique is considered successful for that morphology. The base method only produces successful results for the smallest morphologies. The cobweb method works for slightly larger morphologies. Only the fully connected graph connectivity works for every morphological size.

Method	Success rate	Smallest successful morphology size	Biggest successful morphology size
Base	5%	3 nodes	6 nodes
Cobweb	41%	3 nodes	49 nodes
Fully connected	100%	3 nodes	110 nodes




Figure 6: A representation of the bounding boxes for spectral gap λ_2 and induced norm $\|L\|_2$. Every dot in the scatter plot corresponds to a unique graph topology and coupling weight K , each of which representing 100 repetitions with a different seed. The bounding boxes are drawn to only contain dots that had, respectively, 100% and $\geq 50\%$ of the trials that obtained full convergence after 200 time steps of solving the DE. Increasing coupling weight K increases λ_2 and $\|L\|_2$, whereas increasing the number of couplings per node closes the gap between λ_2 and $\|L\|_2$ until they become equal.

shows that there is indeed a clear bounding box for λ_2 and $\|L\|_2$ for which the convergence is relatively quick and stable.

Both Figure 6 and Table 3 show that a more interconnected graph has a closer gap between λ_2 and $\|L\|_2$. This means that a system with more couplings per node, becomes very easily tunable. That is because the weight coupling K can still be freely chosen to shift the eigenvalues. Hence, the following statement applies to fully connected graphs:

$$\exists K : \lambda_2 = \|L\|_2 \in [80, 1200]. \quad (10)$$

4.3 Effect of criterion on evolutionary robotics

The final goal is to see whether the application of the convergence criterion brings additional benefits in the context of evolutionary optimisation and QD in robotics simulations. Figure 7 shows that

the $\lambda_2/\|L\|_2$ -criterion causes significant improvements in both coverage and QD-score for every morphology. Additionally, Figure 8 shows that those improvements in coverage are specifically in the edges of the archive, which characterise highly coordinated gaits (perfect bilateral or contralateral scores), as well as rare gaits that only use 1 arm, or all 5 arms equally to generate propulsive force. On the other hand, the centre of the grid contains uncoordinated, chaotic gaits. Those are the gaits that have also been discovered by the CPG controllers that do not obey the design criterion.

5 Discussion and Limitations

Our research shows that convergence can be assessed by calculating the spectral gap λ_2 and the induced norm $\|L\|_2$ of the Laplacian of the weighted adjacency matrix of a CPG with Kuramoto-like synchronisation. This in turn ensures significant improvements regarding both exploitation and exploration in evolutionary robotics. Additionally, it shows that fully connected graphs will always be tunable thanks to the fact that λ_2 and $\|L\|_2$ become equal, and scale with coupling weight K . However, one should not simply adopt the specific numerical ranges that have been described above, as the exact boundaries for convergence and numerical stability depend on the specific setup. For instance, bigger natural frequencies can increase the derivative values of the oscillators, making the setup more vulnerable for numerical instabilities. Further, decreasing the solver time step could push the boundaries of faster convergence without instabilities.

The results in this research argue in favour of fully connected CPG topologies, which might feel counter-intuitive since it blurs any spatial relationships between actuators. Nevertheless, our research also shows that connectivities with about two (base method) to four (cobweb method) couplings per node, can also satisfy the design criterion, but only for relatively small morphologies. A famous example is the salamander robot by Ijspeert et al. (2007), which closely resembles the cobweb-like topology for a small to medium number of nodes [18]. We invite further research to investigate the sensitivity of the exploration for different robots and tasks.

The proposed design criterion mainly serves the purpose of obtaining reproducibility upon repeated genotype evaluations, and the benefits it brings within the field of evolutionary optimisations. One could also value different design criteria when fast convergence to a global coherent state is subordinate. Examples include designing a physical robot or a biologically plausible nervous system, where spatial proximity of actuators should be respected. Additionally, some applications might even value slow convergence, which does not enforce a single, global, coherent state. Such research might even seek to purposely violate the design criterion. That is because




Figure 7: Comparison of the means and 95% confidence intervals across different morphologies. The green curves represent CPG controllers which satisfy the design criterion, and the orange and red curves represent CPGs for which λ_2 and $\|L\|_2$ lie respectively below and above the prescribed ranges from Figure 6.




Figure 8: Two randomly chosen MAP-Elites repertoires for brittle star simulations for the medium morphology with either a weakly coupled, slow converging CPG, or a strongly coupled, quick converging, stable CPG, that has been designed with the spectral graph analysis criterion. The colour scale represents the fitness. Empty cells highlight that no controllers were able to produce the behaviour corresponding to those descriptor values.

weakly coupled Kuramoto systems allow investigation of the emergence of synchronisation, instead of enforcing coherency, which is for instance of interest in neuroscience (epilepsy) and chemistry (chemical oscillators) [42]. In evolutionary robotics however, tightly coupled systems apparently greatly improve global coordination and reproducibility of genotype evaluation.

6 Conclusion

CPGs greatly simplify the control and optimisation of locomotion in robotics. A popular approach to CPGs is to use the Kuramoto model for synchronisation. The problem with such CPGs is that they do not always converge to a single, global coherent state. This is especially the case when the coupling is too weak, or so strong that the system becomes unstable. However, evolutionary optimisation relies heavily on reproducibility of the fitness evaluation for a given genotype. In this work, we present a clear design criterion to guarantee both fast and stable convergence. This criterion is based on the calculation of the spectral gap λ_2 and the induced norm $\|L\|_2$ of the connectivity graph. The strength of the criterion is that it can be evaluated even before running a single robotic simulation or genetic optimisation. Moreover, we prove that obeying this criterion significantly improves the diversity and coordination of gait patterns for a brittle star robot, and that it scales to any morphological size. This work provides other researchers with a powerful tool for designing controllers that will improve the quality and consistency of their optimisation in evolutionary robotics.

Acknowledgments

We would like to thank Dominique Adriaens and Joni Dambre for their valuable insights. This work was supported by the Research Foundation Flanders (FWO) (1154126N) and Ghent University's Special Research Fund (BOF) (BOF21/DOC/015)

References

- [1] Juan A. Acebrón, L. L. Bonilla, Conrad J. Pérez Vicente, Félix Ritort, and Renato Spigler. 2005. The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics* 77, 1 (April 2005), 137–185. doi:10.1103/RevModPhys.77.137
- [2] Jacobo Aguirre. 2022. Life finds a way. *Nature Ecology & Evolution* 6, 11 (Nov. 2022), 1599–1600. doi:10.1038/s41559-022-01877-x
- [3] Torsten Asselmeyer, Werner Ebeling, and Helge Rosé. 1996. Smoothing representation of fitness landscapes — the genotype-phenotype map of evolution. *Biosystems* 39, 1 (Jan. 1996), 63–76. doi:10.1016/0303-2647(95)01601-5
- [4] Henry C. Astley. 2012. Getting around when you're round: quantitative analysis of the locomotion of the blunt-spined brittle star, *Ophiocoma echinata*. *Journal of Experimental Biology* 215, 11 (June 2012), 1923–1929. doi:10.1242/jeb.068460
- [5] Charles Audet and Michael Kokkolaras. 2016. Blackbox and derivative-free optimization: theory, algorithms and applications. *Optimization and Engineering* 17, 1 (March 2016), 1–2. doi:10.1007/s11081-016-9307-4
- [6] Guillaume Bellegarda and Auke Ijspeert. 2022. CPG-RL: Learning Central Pattern Generators for Quadruped Locomotion. *IEEE Robotics and Automation Letters* 7, 4 (Oct. 2022), 12547–12554. doi:10.1109/LRA.2022.3218167
- [7] Josh Bongard. 2011. Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences* 108, 4 (Jan. 2011), 1234–1239. doi:10.1073/pnas.1015390108
- [8] Felix Chalumeau, Bryan Lim, Raphael Boige, Maxime Allard, Luca Grillotti, Manon Flageat, Valentin Macé, Guillaume Richard, Arthur Flajolet, Thomas Pierrot, and others. 2024. Qdax: A library for quality-diversity and population-based algorithms with hardware acceleration. *Journal of Machine Learning Research* 25, 108 (2024), 1–16.
- [9] Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret. 2021. Quality-Diversity Optimization: A Novel Branch of Stochastic Optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*, Panos M. Pardalos, Varvara Rasskazova, and Michael N. Vrahatis (Eds.). Springer International Publishing, Cham, 109–135. doi:10.1007/978-3-030-66515-9_4
- [10] Fan R. K. Chung. 1997. *Spectral Graph Theory*. American Mathematical Soc. Google-Books-ID: 4IK8DgAAQBAJ.
- [11] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (May 2015), 503–507. doi:10.1038/nature14422
- [12] V. Dietz. 2003. Spinal cord pattern generators for locomotion. *Clinical Neurophysiology* 114, 8 (Aug. 2003), 1379–1389. doi:10.1016/S1388-2457(03)00120-2
- [13] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E. (Gusz) Eiben. 2015. Evolutionary Robotics: What, Why, and Where to. *Frontiers in Robotics and AI* 2 (March 2015). doi:10.3389/frobt.2015.00004
- [14] Manon Flageat, Johann Huber, François Helenon, Stephane Doncieux, and Antoine Cully. 2025. Extract-QD Framework: A Generic Approach for Quality-Diversity in Noisy, Stochastic or Uncertain Domains. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '25)*. Association for Computing Machinery, New York, NY, USA, 140–148. doi:10.1145/3712256.3726404
- [15] Manon Flageat, Hannah Janmohamed, Bryan Lim, and Antoine Cully. 2025. Exploring the Performance-Reproducibility Trade-off in Quality-Diversity. *IEEE Transactions on Evolutionary Computation* (2025), 1–1. doi:10.1109/TEVC.2025.3548438
- [16] Luca Grillotti, Manon Flageat, Bryan Lim, and Antoine Cully. 2023. Don't Bet on Luck Alone: Enhancing Behavioral Reproducibility of Quality-Diversity Solutions in Uncertain Domains. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '23)*. Association for Computing Machinery, New York, NY, USA, 156–164. doi:10.1145/3583131.3590498
- [17] Auke Jan Ijspeert. 2008. Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks* 21, 4 (May 2008), 642–653. doi:10.1016/j.neunet.2008.03.014
- [18] Auke Jan Ijspeert, Alessandro Crespi, Dimitri Ryczko, and Jean-Marie Cabelguen. 2007. From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model. *Science* 315, 5817 (March 2007), 1416–1420. doi:10.1126/science.1138353
- [19] Michael Ishida, Fidji Berio, Valentina Di Santo, Neil H. Shubin, and Fumiya Iida. 2024. Paleoinspired robotics as an experimental approach to the history of life. *Science Robotics* 9, 95 (Oct. 2024), eadn1125. doi:10.1126/scirobotics.adn1125
- [20] A. Jadbabaie, N. Moté, and M. Barahona. 2004. On the stability of the Kuramoto model of coupled nonlinear oscillators. In *Proceedings of the 2004 American Control Conference*, Vol. 5. 4296–4301 vol.5. doi:10.23919/ACC.2004.1383983 ISSN: 0743-1619.
- [21] Takeshi Kano, Daichi Kanauchi, Hitoshi Aonuma, Elizabeth G. Clark, and Akio Ishiguro. 2019. Decentralized Control Mechanism for Determination of Moving Direction in Brittle Stars With Penta-Radially Symmetric Body. *Frontiers in Neurobotics* 13 (Aug. 2019). doi:10.3389/fnbot.2019.00066
- [22] Takeshi Kano, Eiki Sato, Tatsuya Ono, Hitoshi Aonuma, Yoshiya Matsuzaka, and Akio Ishiguro. 2017. A brittle star-like robot capable of immediately adapting to unexpected physical damage. *Royal Society Open Science* 4, 12 (Dec. 2017), 171200. doi:10.1098/rsos.171200
- [23] Yoshiki Kuramoto. 2003. *Chemical oscillations, waves, and turbulence*. Courier Corporation.
- [24] Jeffrey Larson, Matt Menickelly, and Stefan M. Wild. 2019. Derivative-free optimization methods. *Acta Numerica* 28 (May 2019), 287–404. doi:10.1017/S0962492919000060
- [25] George V. Lauder. 2022. Robotics as a Comparative Method in Ecology and Evolutionary Biology. *Integrative and Comparative Biology* 62, 3 (Sept. 2022), 721–734. doi:10.1093/icb/icac016
- [26] Joel Lehman and Kenneth O. Stanley. 2011. Novelty Search and the Problem with Objectives. In *Genetic Programming Theory and Practice IX*, Rick Riolo, Ekaterina Vladislavleva, and Jason H. Moore (Eds.). Springer, New York, NY, 37–56. doi:10.1007/978-1-4614-1770-5_3
- [27] Eve Marder and Dirk Bucher. 2001. Central pattern generators and the control of rhythmic movements. *Current Biology* 11, 23 (Nov. 2001), R986–R996. doi:10.1016/S0960-9822(01)00581-4
- [28] Yoshiya Matsuzaka, Eiki Sato, Takeshi Kano, Hitoshi Aonuma, and Akio Ishiguro. 2017. Non-centralized and functionally localized nervous system of ophiuroids: evidence from topical anesthetic experiments. *Biology Open* 6, 4 (April 2017), 425–438. doi:10.1242/bio.019836
- [29] S. Miller, J. van der Burg, and F. G. A. van der Meché. 1975. Coordination of movements of the hindlimbs and forelimbs in different forms of locomotion in normal and decerebrate cats. *Brain Research* 91, 2 (June 1975), 217–237. doi:10.1016/0006-8993(75)90544-2
- [30] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. doi:10.48550/arXiv.1504.04909 arXiv:1504.04909 [cs].
- [31] Anh Mai Nguyen, Jason Yosinski, and Jeff Clune. 2015. Innovation Engines: Automated Creativity and Improved Stochastic Optimization via Deep Learning. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation (GECCO '15)*. Association for Computing Machinery, New York, NY, USA, 959–966. doi:10.1145/2739480.2754703
- [32] Stefano Nolfi, Josh Bongard, Phil Husbands, and Dario Floreano. 2016. Evolutionary Robotics. In *Springer Handbook of Robotics*, Bruno Siciliano and Oussama Khatib (Eds.). Springer International Publishing, Cham, 2035–2068. doi:10.1007/978-3-319-32552-1_76
- [33] Jaleigh Q. Pier and Jansen Smith. 2022. Ophiuroidea. <https://www.digitalatlasofancientlife.org/learn/echinodermata/ophiuroidea/>
- [34] Ricardo Pérez-López, Andrés Espinal, Marco Sotelo-Figueroa, Erick I. Guerra-Hernandez, Patricia Batres-Mendoza, and Horacio Rostro-Gonzalez. 2024. Evolutionary Deployment of Central Pattern Generators for Legged Robots Using Nengo. In *New Directions on Hybrid Intelligent Systems Based on Neural Networks, Fuzzy Logic, and Optimization Algorithms*, Patricia Melin and Oscar Castillo (Eds.). Springer Nature Switzerland, Cham, 141–155. doi:10.1007/978-3-031-53713-4_12
- [35] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. doi:10.48550/arXiv.1703.03864 arXiv:1703.03864 [stat].
- [36] Allen I. Selverston. 2010. Invertebrate central pattern generator circuits. *Philosophical Transactions of the Royal Society B: Biological Sciences* 365, 1551 (Aug. 2010), 2329–2345. doi:10.1098/rstb.2009.0270
- [37] Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer Hicks, Sergey Levine, Christopher G. Atkeson, and Scott L. Delp. 2021. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *Journal of NeuroEngineering and Rehabilitation* 18, 1 (Aug. 2021), 126. doi:10.1186/s12984-021-00919-y
- [38] Alexander Sproewitz, Rico Moeckel, Jérôme Maye, and Auke Jan Ijspeert. 2008. Learning to Move in Modular Robots using Central Pattern Generators and Online Optimization. *The International Journal of Robotics Research* 27, 3–4 (March 2008), 423–443. doi:10.1177/0278364907088401
- [39] Michiel Stock, Daan Van Hauwermeiren, Bernard De Baets, Steff Taelman, Dries Marzougui, and Maxime Van Haeverbeke. 2025. Quality-Diversity Methods for the Modern Data Scientist. *WREs Computational Statistics* 17, 4 (2025), e70047. doi:10.1002/wics.70047 eprint: <https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wics.70047>
- [40] Steven Strogatz. 2003. *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press.
- [41] Steven H. Strogatz. 2000. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena* 143, 1 (Sept. 2000), 1–20. doi:10.1016/S0167-2789(00)00094-4

- [42] Steven H. Strogatz and Renato E. Mirollo. 1991. Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics* 63, 3 (May 1991), 613–635. doi:10.1007/BF01029202
- [43] Marco Tamborini. 2021. The Material Turn in the Study of Form: From Bio-Inspired Robots to Robotics-Inspired Morphology. *Perspectives on Science* 29, 5 (Oct. 2021), 643–665. doi:10.1162/posc_a_00388
- [44] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 5026–5033. doi:10.1109/IROS.2012.6386109
- [45] Gabriel Urbain, Alexander Vandesompele, Francis Wyffels, and Joni Dambre. 2018. Calibration Method to Improve Transfer from Simulation to Quadruped Robots. In *From Animals to Animats 15*, Poramate Manoonpong, Jørgen Christian Larsen, Xiaofeng Xiong, John Hallam, and Jochen Triesch (Eds.). Springer International Publishing, Cham, 102–113. doi:10.1007/978-3-319-97628-0_9
- [46] Vesselin K. Vassilev, Terence C. Fogarty, and Julian F. Miller. 2003. Smoothness, Ruggedness and Neutrality of Fitness Landscapes: from Theory to Application. In *Advances in Evolutionary Computing: Theory and Applications*, Ashish Ghosh and Shigeyoshi Tsutsui (Eds.). Springer, Berlin, Heidelberg, 3–44. doi:10.1007/978-3-642-18965-4_1
- [47] Vassilis Vassiliades and Jean-Baptiste Mouret. 2018. Discovering the elite hyper-volume by leveraging interspecies correlation. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '18)*. Association for Computing Machinery, New York, NY, USA, 149–156. doi:10.1145/3205455.3205602
- [48] Sharmila Venugopal. 2022. Connectionist Models of CPG Networks. In *Encyclopedia of Computational Neuroscience*. Springer, New York, NY, 957–959. doi:10.1007/978-1-0716-1006-0_122
- [49] Andreas Wagner. 2007. Robustness and evolvability: a paradox resolved. *Proceedings of the Royal Society B: Biological Sciences* 275, 1630 (Oct. 2007), 91–100. doi:10.1098/rspb.2007.1137
- [50] Günter P. Wagner and Lee Altenberg. 1996. PERSPECTIVE: COMPLEX ADAPTATIONS AND THE EVOLUTION OF EVOLVABILITY. *Evolution* 50, 3 (June 1996), 967–976. doi:10.1111/j.1558-5646.1996.tb02339.x
- [51] Daiki Wakita, Katsushi Kagaya, and Hitoshi Aonuma. 2020. A general model of locomotion of brittle stars with a variable number of arms. *Journal of The Royal Society Interface* 17, 162 (Jan. 2020), 20190374. doi:10.1098/rsif.2019.0374
- [52] Daan Wierstra, Tom Schaul, Jan Peters, and Juergen Schmidhuber. 2008. Natural Evolution Strategies. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, Hong Kong, China, 3381–3387. doi:10.1109/CEC.2008.4631255
- [53] Arthur T. Winfree. 1967. Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology* 16, 1 (July 1967), 15–42. doi:10.1016/0022-5193(67)90051-3
- [54] Joe Wright and Ivan Jordanov. 2015. Intelligent Approaches in Locomotion - A Review. *Journal of Intelligent & Robotic Systems* 80, 2 (Nov. 2015), 255–277. doi:10.1007/s10846-014-0149-z
- [55] Chuanyu Yang, Can Pu, Yuan Zou, Tianqi Wei, Cong Wang, and Zhibin Li. 2025. Bio-inspired neural networks with central pattern generators for learning multi-skill locomotion. *Scientific Reports* 15, 1 (March 2025), 10165. doi:10.1038/s41598-025-94408-0
- [56] Ya-Tang Yang, Jhih-Ci Li, Feng-Yu Wang, Jheng-Lin Lin, Lai Jhih-Hao, and Jyun-Wei Shih. 2024. Myriapod robot locomotion control using central pattern generator based on Kuramoto oscillator model". In *2024 IEEE International Conference on Cyborg and Bionic Systems (CBS)*. 160–164. doi:10.1109/CBS61689.2024.10860652
- [57] Ya-Tang Yang, Jyun-Wei Shih, Feng-Yu Wang, and Jhih-Ci Li. 2025. Fish Robot Locomotion Control Using Central Pattern Generator Based on Kuramoto Oscillator Model. In *Technologies and Applications of Artificial Intelligence*, Wei-Ta Chu, Chih-Ya Shen, and Hong-Han Shuai (Eds.). Springer Nature, Singapore, 231–239. doi:10.1007/978-981-96-4596-1_17

Received 26 January 2026; revised 10 April 2026