2024

Dancing in the Dark: Evolving CPGs for Entrainment of Locomotion to Patterned Stimuli with Minimal Components

Emily A. Ertle
University of Vermont

Dancing in the Dark: Evolving CPGs for Entrainment of Locomotion to Patterned Stimuli with Minimal Components

An Honors Thesis Presented

by

Emily Abigail Ertle

to

The Faculty of The University of Vermont

May, 2024

Defense Date: April 25th, 2024
Thesis Examination Committee:

Josh Bongard, Ph.D., Advisor
Jeremy M. Barry, Ph.D.

Abstract

Legged locomotion presents a significant challenge in robotics. Many legged robots accomplish stable movement through models of "central pattern generators (CPGs)," a type of neural circuit which underlies biological rhythms from walking, flying and breathing to patterned cognitive and central nervous system activity. While current CPG models are effective solutions for moving from point A to point B, they have several important drawbacks. These include reliance on complex, specialized neuron models and specific neural topology, which make the system difficult to modify or improve. Artificial CPG design also sacrifices stability for adaptability, as their mechanism largely prevents gait variation. In this work, we used a multi-objective evolutionary algorithm to produce virtual robots able to rhythmically entrain—synchronize footstrikes—to a simple metronome. Robots had an "auditory neuron" to sense metronome strikes and the selection algorithm favored individuals which both traveled away from the origin and demonstrated strong rhythmic alignment. In this paper, we explore what conditions and methods might be conducive to evolving rhythmic entrainment in the spirit of minimal cognition. In addition, we demonstrate evolution of a functional CPG with only a small network of simple tanh neurons and make inferences about its mechanism.

This is dedicated to my family, friends, and all the teachers, professors, and mentors who have supported my curiosity.

Acknowledgements

I would like to thank Piper Welch and Dr. Josh Bongard for their many hours of help and support. Piper, I cannot thank you enough for always making yourself available. You helped me find simple solutions when things felt too complicated and always made me feel like my goals were possible. You are also a fountain of good advice. Dr. Bongard, thank you for your mentorship and for encouraging me to explore my own ideas. You provided the perfect balance of encouragement and skepticism for me to stay excited about my work but leave your lab a much better scientist and critical thinker.

Thank you to all the members of the Morphology, Evolution, and Cognition lab for sharing many great ideas and conversations. I have learned so much from you.

Thank you to Dr. Jeremy Barry for serving on my thesis committee. I am grateful for your support, your excellent feedback, and for helping me see my work from a new angle.

Finally, thank you to my friends and family. I truly could not have done this without all of your love and support.

Table of Contents

List of Figures

1 Introduction

Central pattern generators are defined as neuronal circuits that generate coordinated, rhythmic output in the absence of patterned sensory input [1, 2]. These circuits are common to vertebrates and invertebrates and generate an incredibly diverse array of behaviors, the majority of which can be characterized as either episodic (like locomotion) or continuous (like breathing) [1]. The dynamics of a given circuit are a function of both its connectivity pattern and the intrinsic membrane properties of its constituent neurons [2]. Further, these features on their own do not imply anything about the function of a circuit—rewiring the same group of neurons can produce a different behavior and two circuits with the same connectivity pattern can have different functions [3]. Instead, the features of each circuit have evolved independently to cater to an organism's neurobiology, behavior, and morphology [4], making each CPG a very elegant solution to a specific problem.

In vertebrates, CPG outputs are driven by supra-spinal networks activated by afferent (traveling from the periphery to the central nervous system) sensory inputs [5, 6]. Though CPGs, by definition, must be able to generate rhythmic output in the absence of sensory input [1], sensory feedback still plays a critical role in facilitating gait changes, frequency changes, and corrective responses to obstacles and uncertainties [2, 5, 7]. CPG function also includes modulation of behavioral strength and frequency, allowing for stability in response to perturbation or noise and modulation when desired [2, 5]. This combined open-loop (supraspinal interaction) and closed-loop (internal regulation) CPG control permits incredibly robust control of locomotion [5].

Models of CPGs have grown increasingly popular in the field of robotics, often to achieve stable locomotion [7]. CPGs are an effective solution for stable locomotion because they are designed to self-synchronize and thus have very stable oscillations [8]; this translates to success in navigating complex terrain and robustness against perturbations [7]. While this makes CPG models very effective solutions for moving from point A to point B, they have several important drawbacks. One of these is reliance on complex, specialized neuron models and specific neural topology, making the systems employing them difficult to modify or improve [9]. Artificial CPG design also sacrifices adaptability for stability, as their mechanism largely prevents gait variation and frequency modulation [10, 8].

In biological organisms, supraspinal networks and intrinsic neuron properties facilitate legged locomotion that balances stability and adaptability. Artificial CPGs are instead built from specialized neuron models called oscillators [8]. The most well-known and most commonly used of these is the Matsuoka oscillator [11], which can only entrain to a very limited range of frequencies around a pre-set intrinsic frequency [8]. In trying to reduce complex neuron behavior to an equation—or in this case a pair of nonlinear differential equations—critical functionality is lost. One author with 40 years of experience in invertebrate CPG analysis believes that due to the immense difficulty of describing and predicting multiple nonlinear interactions, "there probably is not [a correct computational model of CPG behavior]" [4]. Robots unable to adapt their gaits (or other rhythmic behaviors) to external stimuli in real-time have greatly reduced potential for human-robot interaction [8]. If CPG behavior may never be perfectly modeled, balancing stability and adaptability in robotic legged locomotion may require alternative solutions.

A study published this July by Szorkovszky et al. at the Centre for Interdisciplinary Studies in Rhythm, Time, and Motion in Oslo attempted to achieve this balance using an evolutionary algorithm to evolve CPGs for real-time adaptation to patterned stimuli [10]. Though this method successfully produced more adaptable CPGs, their system was highly complex. One reason for this complexity was their choice of neuron model: a custom Matsuoka oscillator with several added elements. Additionally, four-objective evolutionary optimization was used to find solutions, which involves selecting individuals on a four-dimensional plane and is both complicated and computationally expensive.

The field of evolutionary robotics attempts to model and understand natural embodied intelligence by exerting evolutionary pressure on robots simulated in a physics engine. Selection is facilitated by an evolutionary algorithm (EA), a type of AI search process that begins with random solutions and generates optimal ones through a cyclical process of evaluation, selection, and reproduction/mutation. EAs can be used to automate any aspect of robot design, such as neuron weights, neural network topology, or body morphology. In this way, an EA can occupy the role of a 'scientist' who, lacking human observer bias, approaches robot design from an otherwise inaccessible perspective—a neutral one. Maintaining neutrality while avoiding perverse instantiation can be a challenge. Perverse instantiation is a term use to describe behavior of an artificial agent that performs *exactly* the task it was asked to do, but in an unexpected way, akin to 'cheating'. If that challenge can be overcome, EAs permit the discovery of ideal, non-intuitive solutions suited to the constraints of a simulated organism and its environment by mimicking the process that yields the elegant but complex solutions found in nature.

While Szorkovszky et al. demonstrated their method was sufficient to produce the desired behavior of real-time rhythmic entrainment (synchronization), this method was complex and computationally intensive. Further, its continued utilization of a CPG model to achieve locomotion reduces its potential use as a tool for understanding biological phenomena. Evolving a much simpler controller composed of tanh-activated neurons rather than oscillators for rhythmic entrainment would both reduce the complexity of the system and open opportunities for future study of real-world evolutionary pressures [12] related to biological CPGs and rhythmic processing. While many other studies have attempted to apply evolutionary or genetic algorithms to the problem of CPG design, adaptable locomotion, or locomotion with multiple defined gaits ([10], [9], [13], [14], [15]), literature searches revealed few solutions which did not involve a neuron or oscillator model and none which attempted to evolve a CPG simply by rewarding simple fundamental behaviors that result in their utilization. Here, we explore what conditions and methods might be conducive to evolving rhythmic entrainment in the spirit of minimal cognition. In addition, we demonstrate evolution of a functional CPG with only a small network of simple tanh neurons and make inferences about its underlying mechanism.

2 Methods

In order to determine the feasibility of evolving CPG-like behavior externally rather than internally, the relative success of different multi-objective optimization methods and fitness functions was explored. To better understand possible limitations of this approach, the performance of three body plans across evolution was compared.

Finally, one evolved individual was chosen for more in-depth testing and analysis. The Mann-Whitney U test was used for all statistical analyses [16]. The Holm-Bonferroni procedure was used to correct rejection criteria for the multiple pairwise comparisons [17].

2.1 The Robots

All three robots were given capsule-shaped legs to reduce the number of contact points between each leg and the ground to only one, making it easy to determine when a 'step' occurs. Doing so quickly improved the success of attempts to evolve rhythmic behavior and the ease with which observers can estimate rhythmicity from simulation audio. Robot proportions will be described relative to torso length, which was set to 1 meter (the standard unit of length in Unified Robot Description Format (URDF) files). The colors and markers used in Fig. 1 to indicate the limbs and joints discussed below will be listed beside their descriptions to increase clarity. All capsule-shaped body components have a radius of 1/5 (dark gray) the torso length (orange) and all joints are capable of 0.3 radians of rotation from their original position in both directions. All revolute joints, also known as hinge joints, for flexion and extension are marked by red circles and transverse revolute joints are marked by blue circles. With the exception of body plan comparison, the 1DOF (degree of freedom) radial quadruped is used in all experiments (Fig. 1a).

A. The 1DOF Radial Quadruped

The 1DOF radially-symmetric quadruped (Fig. 1a) has a cubic torso and four cylindrical legs, each with one upper (purple) and lower (green) component. The two

components of each leg have lengths equivalent to that of the torso. Each hip and knee contains a single revolute joint for flexion and extension.

B. The 2DOF Radial Quadruped

The 2DOF radially-symmetric quadruped (Fig. 1b) has a cubic torso and four cylindrical legs, each with two upper components and one lower component. The two components making up the upper part of the leg, from most to least central, have lengths of 1/5th (yellow) and 4/5ths (purple) that of the torso, respectively. The first, short component is connected to the torso by a transverse revolute joint. The joint connecting these two upper components as well as the knee joint connecting the long upper component to the sole lower component (green) are revolute joints for flexion and extension.

C. The 2DOF Bilateral Quadruped

The 2DOF bilaterally-symmetric quadruped (Fig. 1c) has a cubic torso and four cylindrical legs, each with two upper components and one lower component. The two components making up the upper part of the leg, from most to least central, have lengths of 1/5th (yellow) and 4/5ths (purple) that of the torso, respectively. The first, short component is connected to the torso by a transverse revolute joint. The joint connecting these two upper components as well as the knee joint connecting the long upper component to the sole lower component (green) are revolute joints for flexion and extension.

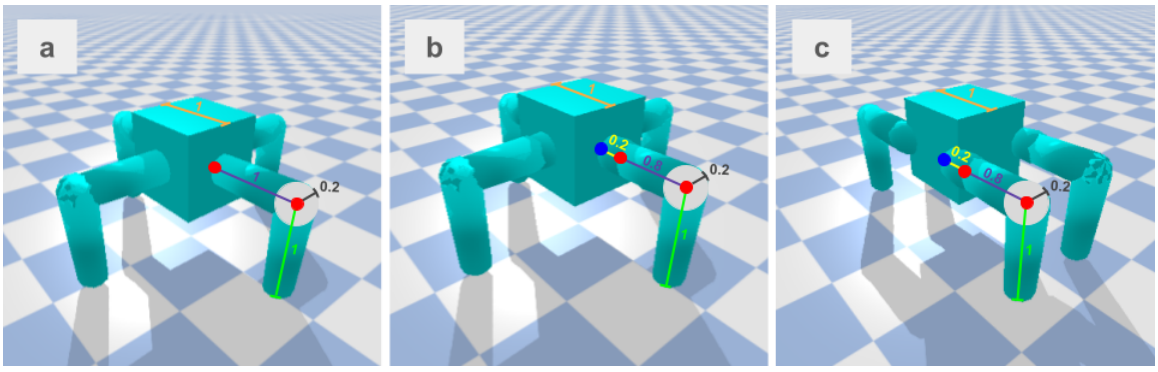

Figure 1: ***Experimental Body Plans.*** ***a:*** *The 1DOF radial quadruped.* ***b:*** *The 2DOF radial quadruped.* ***c:*** *The 2DOF bilateral quadruped. Analogous structures across the three body plans are labeled with lines of the same color. Revolute (hinge) joints for flexion and extension are marked by red circles and transverse revolute joints are marked by blue circles.*

2.2 The Controller

Each robot is controlled by a recurrent neural network with three layers (Fig. 2a). The first layer consists of one touch sensor for each leg in a given robot's motor plan, plus one *auditory neuron*, which accepts binary 'auditory' information. This simplified representation of auditory information is generated by a metronome built in to the simulation. The metronome emits *1s* at regular intervals and *0s* otherwise. These values are fed directly into the auditory neuron. For more information about these simulated auditory stimuli, see subsection 2.3. The sensory (input) layer is fully connected to a layer of three hidden neurons, which is fully self- and recurrently-connected as well as fully connected to the motor layer. Hidden neurons are neurons in an artificial neural network that are neither input nor output layers. The final output layer is composed of one motor neuron for each joint. The placement of each sensor and motor is shown for the quadruped in Fig. 2b. The weights of all outgoing connections from the auditory neuron to hidden neurons are maximized and frozen.

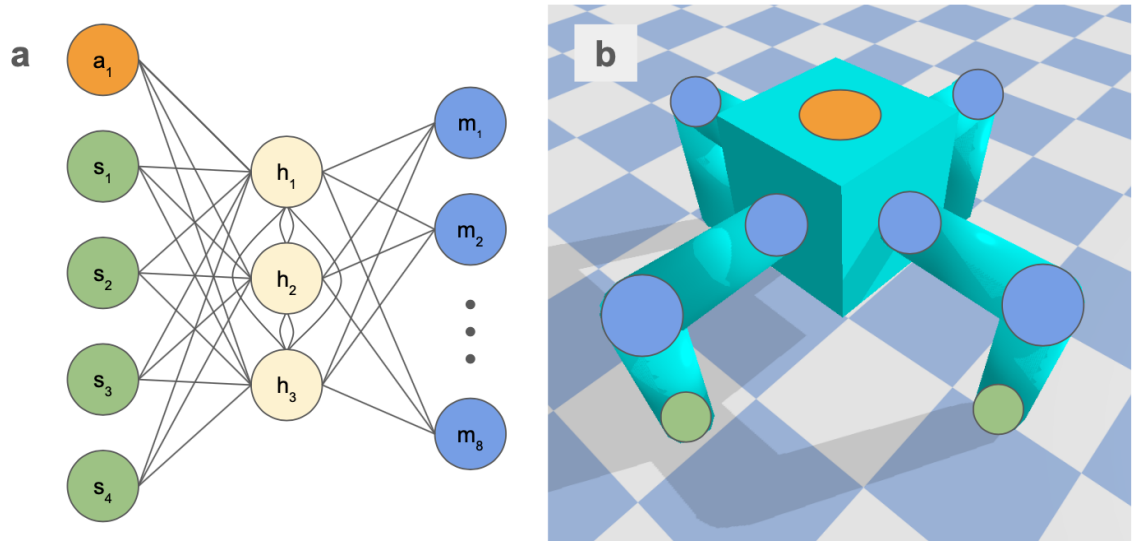

Figure 2: ***Quadruped Neural Network Architecture.*** ***a:*** *Each controller contains one auditory neuron (orange, a), one sensory neuron for each leg (green, s), three hidden neurons (yellow, h), and one motor neuron for each joint—two for each leg (blue, m. In the case of the quadruped, this is four sensory neurons and eight motor neurons.* ***b:*** *Sensor and motor locations in the quadruped morphology.*

2.3 External Rhythmic Stimuli

Often, rhythmic entrainment and creation of a regular gait involves providing a robot with an internal rhythmic stimulus. Typically, this is a neural oscillator model or a sine wave. Here, we instead attempt to evolve oscillatory neural activity by providing robots with a 'metronome'-like external stimulus, a series of isochronous binary impulses. This signal is used to update the value of each robot's auditory neuron. Over the course of one simulation, each robot is exposed to three different conditions, or stimulus frequencies. Stimuli are presented in frequencies of 2Hz, 1.33Hz, and 1.66Hz (in that order), each for 20 pulses. It should be noted that possible period lengths are limited by the length of time between discrete timesteps (in simulation time) in the physics engine. All experiments described here are updated with a step size of

0.025 seconds.

Ideally, these rates would be chosen by determining the natural gait of the body plan. However, because no central pattern generator or oscillator is employed in these robots they do not have pre-established natural gaits. Thus, trends in Froude number for quadrupedal animals (see Alexander [18]) were used to determine what periods would promote steady walking gaits. Froude number is calculated using the equation

$$\text{Froude Number} = \frac{v^2}{g \cdot h} \tag{1}$$

where v is velocity, g is gravitational acceleration ($9.8 m/s^2$), and h is hip height. Humans and most quadrupeds transition from walking to running at a Froude number of 0.5, so inter-stimulus periods are chosen to represent Froude numbers across the spectrum of the walking range (with goal distance per impulse, see G in Equation 9, held constant at the length of the quadruped's torso). Frequencies of 2Hz, 1.33Hz, and 1.66Hz correlate to Froude numbers of 0.4, 0.18, and 0.28, respectively.

2.4 Optimization Methods

In evolutionary robotics, evolutionary pressure is exerted on robots simulated in a physics engine. Selection is facilitated by an evolutionary algorithm (EA), a type of AI search process that begins with random solutions and generates optimal ones through a cyclical process of evaluation, selection, and reproduction/mutation. There are many different evolutionary algorithms, each with their own unique features. It is sometimes unclear which is ideal for a given application, thus the relative success of two algorithms was evaluated.

Age-Fitness Pareto Optimization

Age-Fitness Pareto Optimization (AFPO) is a common multi-objective evolutionary algorithm designed to prevent premature convergence, or search stagnation [19]. This is accomplished by minimizing solution age while maximizing solution fitness. The use of this method to select for more than one behavior requires sub-functions to be carefully combined into a single function. Two promising combined functions are tested. One includes all subfunctions (see subsection 2.5): rhythmicity (f_r), gait evenness (f_e), and movement (f_m)

$$F = f_r \cdot f_e \cdot f_m \tag{2}$$

while the other function excludes movement.

$$F = f_r \cdot f_e \tag{3}$$

Bi-Objective Pareto Optimization

Controllers are also evolved using a modified version of the Age-Fitness Pareto Optimization (AFPO) algorithm in which age is substituted for a second fitness metric. The two fitness metrics used for this method of selection were rhythmicity (f_r) and *gait evenness* (f_e) [19].

2.5 Fitness Functions

Each of the functions below plays a role in both shaping the desired behavior—marching to an external isochronous (evenly temporally spaced) auditory stimulus—

and preventing perverse instantiation. These functions are used to evaluate the 'fitness' of individuals in a population. During the *selection* step of an evolutionary algorithm, individuals are competed at random—meaning their fitnesses are compared to one another. Individuals with higher fitnesses are thus more likely to remain in the population and mutate/reproduce.

Rhythmicity

Further freedom in gait selection was provided by only counting the step closest to the metronome strike in the rhythmicity score, thus permitting robots to strike the ground with the tracked leg at twice the stimulus frequency without punishment (a decrease in fitness).

$$f_r = \frac{1}{R} \cdot \sum_{c=1}^{C} \sum_{i=1}^{I} \max_{0 \le t \le p_c - 1} \left(\frac{p_c}{2} \cdot \cos\left(\frac{2\pi}{p_c} \cdot t \right) + \frac{p_c}{2} \right) \tag{4}$$

Above, C is the total number of rhythmic conditions presented, p_c is the period of the current condition (c), I is the number of impulses presented for each condition, and t is the current timestep. R, the maximum rhythmicity score, can be found using the equation below.

$$R = \sum_{c=1}^{C} p_c \cdot I \tag{5}$$

Conversion of condition frequency in Hz to its period in discrete timesteps was performed using the equation

$$p = (frequency \cdot S)^{-1} \tag{6}$$

where p is the period between impulses and S is the time in seconds between discrete timesteps in the physics engine.

Gait Evenness

Here, we will define *gait evenness* as the amount of variation between the number of steps taken by each leg. This function is included because only one leg's rhythmicity is evaluated to avoid prescribing a specific gait. If the behavior of all legs were factored into the calculation, it would be necessary to either select for pronking, a unique quadrupedal gait in which all four legs connect with and lift from the ground simultaneously, or choose a desired offset for each leg to mimic a more common quadrupedal gait. However, gait evenness was also selected for to prevent the robot from achieving rhythmicity by stepping with only one or two legs. Gait evenness was calculated using the following equations

$$f_e = \sqrt{e+1} \tag{7}$$

$$e = \frac{1}{U} \cdot \sum_{u=1}^{U} (k_{tracked} - k_u)^2 \tag{8}$$

where U is the number of untracked legs and k is the number of footstrikes recorded for a given leg. The choice to square $k_{tracked} - k_u$ rather than take the absolute value of the expression was made to select against having any legs be significant outliers. Taking the square root of e creates a more realistic gradient for selection to follow, increasing the tractability of the problem. Without taking the square root of the function, f_e climbs from 0.1 to 1 over a range of $[k_{tracked} - 3, k_{tracked}]$. With it, the slope decreases, allowing the same growth over a more reasonable range of $[k_{tracked} - 10, k_{tracked}]$. Adding 1 before taking the square root of e both normalizes the values of f_e and avoids undefined values.

Movement

The movement fitness subfunction rewarded solutions one point each time a goal distance (G) was traveled and the square root of the ratio of actual distance traveled (d_i) to the goal distance otherwise. The final result of this calculation was normalized over the maximum movement score, as seen in the equation below.

$$f_m = \frac{1}{C \cdot I} \cdot \sum_{i=1}^{I \cdot C} \left(\begin{cases} 1 & d_i \geq G \\ \sqrt{\frac{d_i}{G}} & otherwise \end{cases} \right) \tag{9}$$

2.6 In-Depth Individual Analysis

An individual is selected for in-depth analysis based on visual assessment of footprint graphs, fitness graphs, and graphical 3D simulation results. Ideal solutions will mimic common features of CPG-controlled biological locomotor entrainment including consistent gait for the span of each individual stimulus frequency, adaptability across tempos, continued motion throughout the course of the simulation, and succesful synchronization of one leg with the stimulus frequency.

3 Results

3.1 Shaping Evolutionary Pressure

Bi-objective Pareto Optimization was found to produce larger and more rapid and gains in rhythmicity, as Pareto fronts were dominated by scores above 0.8 after

100 generations of simulation of 100 individuals. However, gait evenness scores lagged, often resulting in best solutions which did not exhibit the desired behavior—entrainment of locomotion to an external isochronous stimulus.

Employment of Age-Fitness Pareto Optimization (AFPO) yielded more balanced improvement across sub-objectives. Throughout the process of tuning simulation hyperparameters, runs with the same population size and number of generations yielded the desired behavior more frequently when AFPO was used to select for the product of several sub-functions as opposed to selecting for each as an independent optimization objective.

After selection of AFPO as the primary optimization strategy, it was necessary to determine what combination of objectives as a single function would be most likely to yield the desired behavior. Exploratory runs indicated rhythmicity and gait evenness objectives were necessary for evolution of locomotor entrainment. The former because it was the primary goal and the latter because all experiments in which the sub-function was excluded or modified resulted in perverse instantiation. The details of this result are explored in greater depth in the discussion section (see subsection 4.1). At both the start and end of evolution, selection without the movement subfunction yielded significantly lower rhythmicity and movement performance than with the subfunction ($p < 0.05$). Comparison of these two methods for 10 replicates with 95% confidence intervals is visible in Supplementary Fig. 1.

3.2 Body Plan Comparison and Analysis

At the beginning of the simulation, there was no significant difference between the performance of any of the body plans (Fig. 3). At the end of 10 generations, the

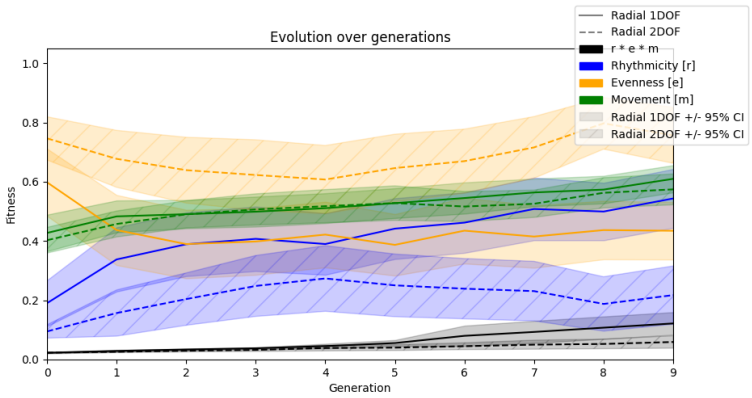

(a) Radial 1DOF v. Radial 2DOF

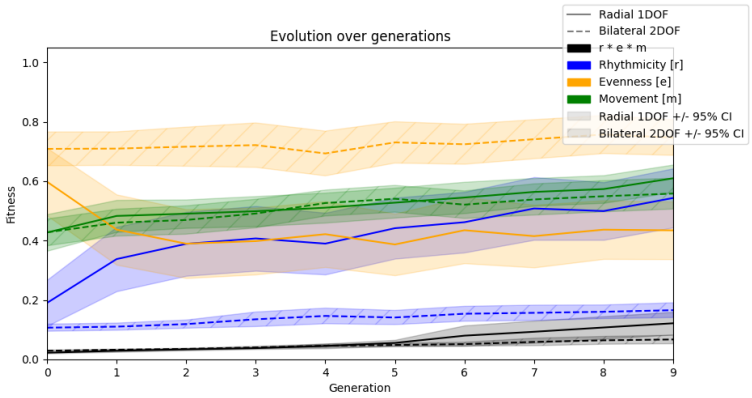

(b) Radial 1DOF v. Bilateral 2DOF

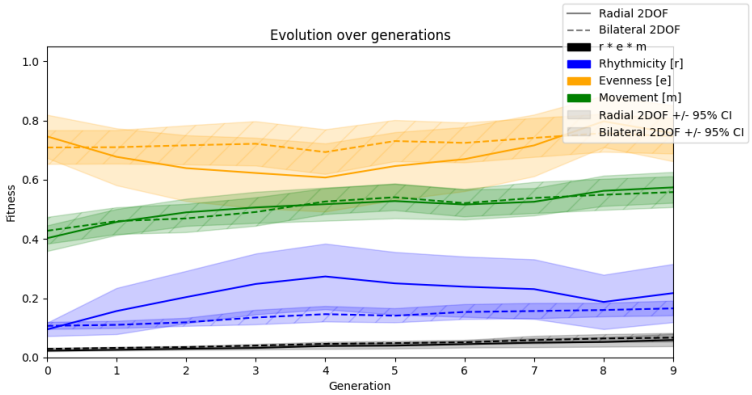

(c) Radial 2DOF v. Bilateral 2DOF

Figure 3: ***Comparison of 3 body plans.*** *Performance over evolution of 100 individuals for 10 generations with 95% confidence intervals. 30 replicates were performed for each body plan.* ***a:*** *Radial 1DOF v. Radial 2DOF.* ***b:*** *Radial 1DOF v. Bilateral 2DOF.* ***c:*** *Radial 2DOF v. Bilateral 2DOF.*

radial 1DOF quadruped performed significantly better overall and on rhythmicity than the radial 2DOF quadruped ($p < 0.01$) and the bilateral 2DOF quadruped ($p < 0.05, p < 0.01$). The difference between the two 2DOF quadrupeds was not statistically significant.

3.3 Selected Individual Analysis

Graphs detailing the evolution of the selected individual are visible in (Fig. 4). This individual was evolved from a population of 100 individuals for 100 generations using AFPO with the fitness function $F = f_r \cdot f_e \cdot f_m$ (see Equation 3). Footprint graphs for the selected individual's performance on frequencies seen during training (Fig. 5a) is shown alongside a graph of auditory and hidden neuron activity throughout the simulation (Fig. 6b). Below the neuron activity graph is a footprint graph from a test simulation of the same brain with stimulus frequencies of 1Hz, 1.25Hz, and 2.5Hz, none of which were supplied during evolution (Fig. 5b). It should also be noted that these test frequencies vary significantly from each other—more so than those provided during evolution (2Hz, 1.33Hz, and 1.66Hz). They also are presented in increasing order, whereas the sequence presented in the evolutionary set is non-monotonic.

4 Discussion

4.1 Shaping Evolutionary Pressure

One limitation of using this technique to evolve rhythmic entrainment is that it is not apparent from any single fitness value what the exact behavior of the robot

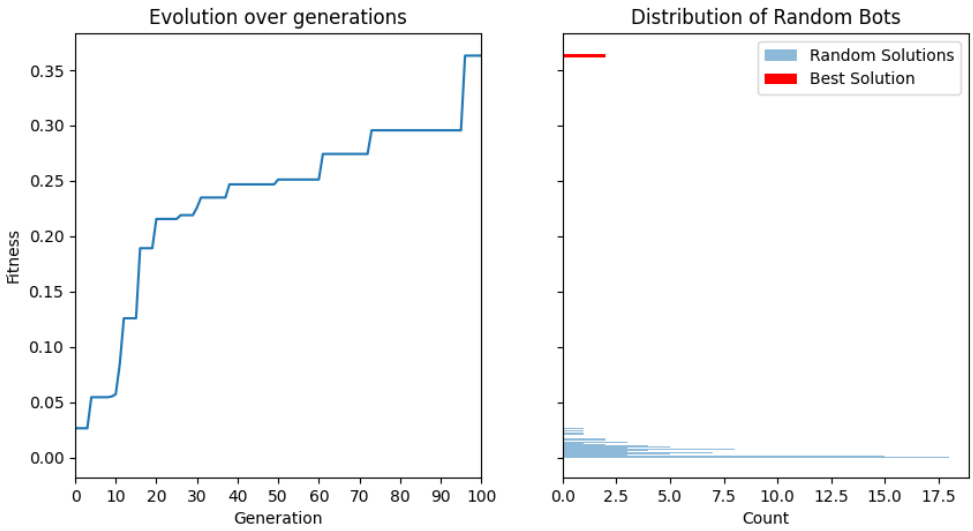

(a)

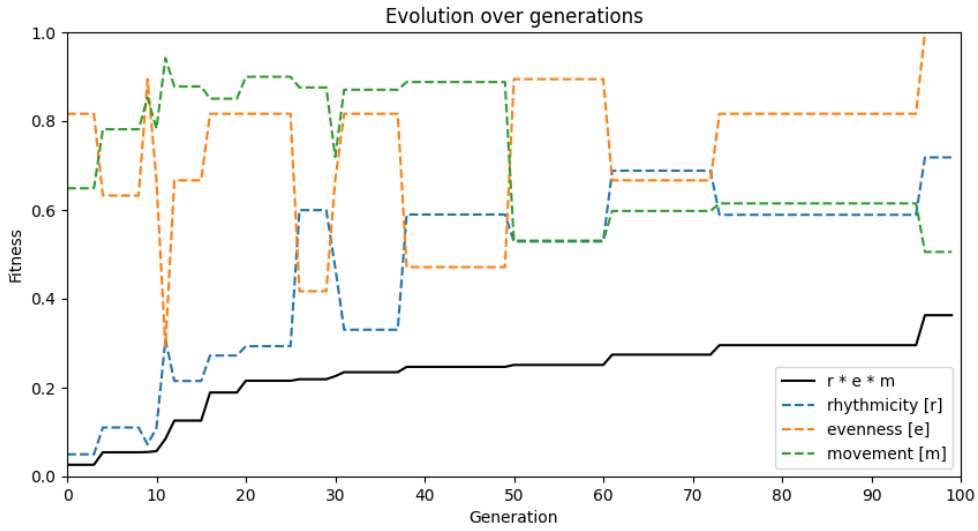

(b)

Figure 4: ***Selected Individual Evolution.*** *Evolution of a population of 100 radial 1DOF quadrupeds for 100 generations.* ***a:*** *On the left, $f_r \cdot f_m \cdot f_e$ (see Equation 3) on the right, fitness of the best evolved solution v. 100 random solutions.* ***b:*** *Values of subfunctions (rhythmicity, gait evenness, and movement) over the course of evolution as well as the overall fitness function (Equation 3).*

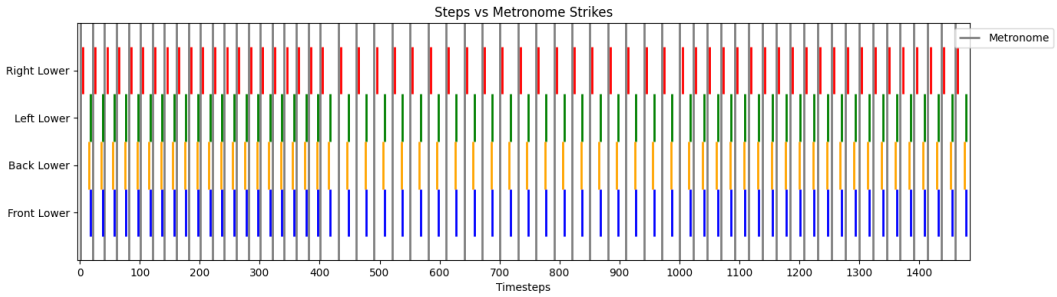

(a) Entrainment to frequencies used during evolution.

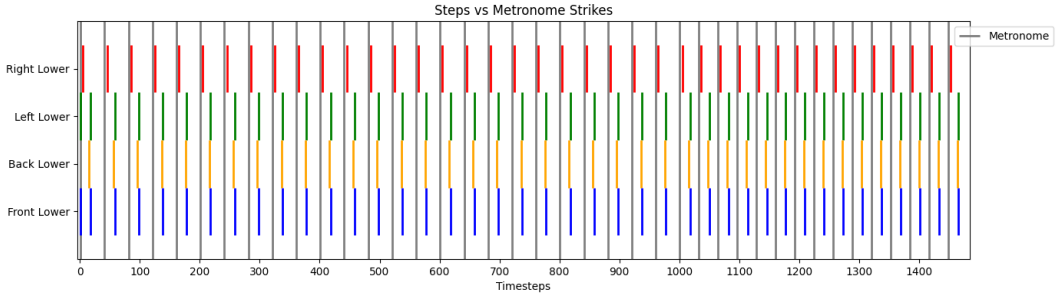

(b) Entrainment to novel test frequencies.

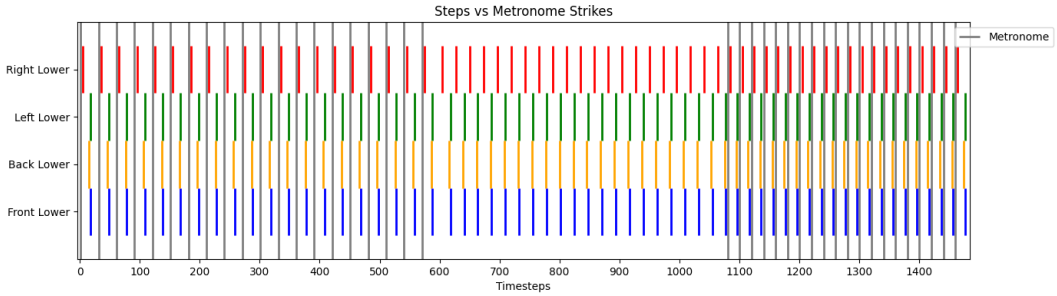

(c) Entrainment in the absence of external input.

Figure 5: ***Selected Individual Behavior.*** ***a:*** *Footprint graph of the selected individual on training frequencies of 2Hz, 1.33Hz, and 1.66Hz.* ***b:*** *Footprint graph of the selected individual on test frequencies of 1Hz, 1.25Hz, and 2.5Hz, none of which were included in evolutionary runs. c: Footprint graph of the selected individual on frequencies of 1.33Hz, followed by a period of no external input, then another period of 2Hz input.*

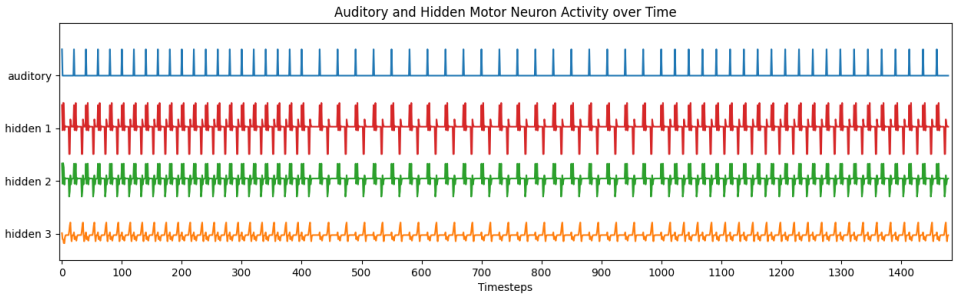

(a) Auditory and hidden neuron activity of the selected individual.

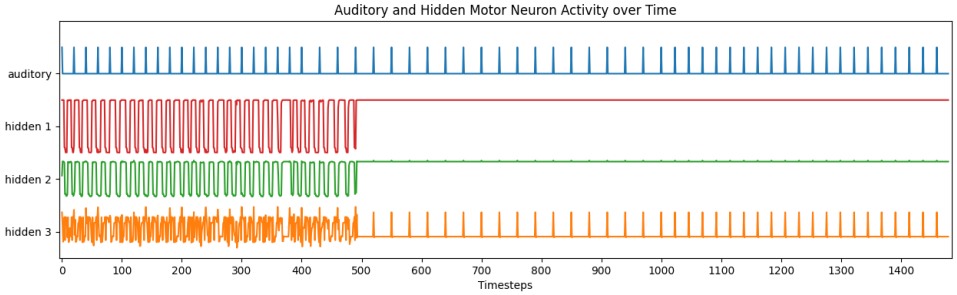

(b) Auditory and hidden neuron activity of a random individual with identical parameters.

Figure 6: ***Selected Individual Neuron Activity.*** *Plots of neuron values over the course of a single simulation. The values of the auditory neuron (top) spike at each point the auditory stimulus is presented.* ***a:*** *Neurons of the selected evolved individual.* ***b:*** *Plot of neuron values for the individual with the highest fitness of 100 random individuals with the same parameters as the selected individual.*

will be. However, it seems there is some benefit to leaving room for interpretation in fitness functions when trying to evolve highly complex nonlinear, adaptable behaviors. During exploratory simulations, it quickly became clear that with greater specificity comes greater risk of perverse instantiation and reduced problem tractability. This led to adoption of multiple relaxed fitness functions that are ineffective alone but collectively work to shape behavior. For example, the fitness function below tested early in the research process reliably rewards robots who step only at the same time as the stimulus.

$$f_r = \sum_{c=1}^{C} \sum_{t=0}^{(p_c \cdot I)-1} \left(\begin{cases} p_c \cdot \cos\left(\frac{2\pi}{p_c} \cdot t\right) & \text{if a footstrike occurred at } t \\ 0 & otherwise \end{cases} \right) \tag{10}$$

However, it is still prone to perverse instantiation, as robots can gain even more points by double-stepping close to each impulse. It also prevents natural gait adaptation from occurring—if the stimulus frequency is too slow to maintain a consistent gait and a gait transition occurs, transition to a gait cycle at exactly twice the frequency of the stimulus will score zero points. Further, robots who stand still will be selected for over robots who walk at the wrong times (and would thus incur a negative score).

The fitness function ultimately used to select for rhythmicity supplies a reward equivalent to the value of cosine only for the step that is temporally closest to the stimulus impulse (Equation 4). Without complementary support from another function—critically, gait evenness (Equation 7)—solutions often step as many times as possible without punishment to stop them and obtain very high rhythmicity scores. This explains why bi-objective Pareto optimization tends to be relatively unsuccessful. The Pareto front quickly becomes filled with solutions with subtly-varying high rhyth-

micity scores—likely solutions who step as frequently as possible—and population diversity becomes low in a manner that restricts the possible evolution of other objectives. This flaw was not shared by the alternative strategy tested, AFPO, which appeared to promote co-evolution of multiple behaviors. This result was not surprising, as AFPO was designed specifically to increase solution diversity and therefore reduce premature solution convergence.

4.2 Body Plan Comparison and Analysis

Increased success of body plans with fewer degrees of freedom early in the evolutionary process is consistent with our current understanding of motor learning and synchronization (Fig. 3). Bernstein's theory of graded skill acquisition suggests early stages of motor task learning are characterized by the "freezing" of degrees of freedom [20], suggesting this reduction increases ease of performance. However, later in learning, performance is benefited from the freeing of these joints, permitting increased efficiency and flexibility as movements are fine-tuned. Several studies have also found increased utilization of degrees of freedom translates to movements that are more accurately timed [21, 22]. Thus, it is possible that continuing simulation for more generations would reveal superior performance of quadrupeds with more degrees of freedom in their hips.

4.3 Selected Individual Analysis

The individual selected for in-depth analysis had a relatively low overall fitness ($f_r \cdot f_e \cdot f_m$) of 0.35 and a rhythmicity score of only 0.6 (Fig. 4a). However, upon inspection

of the footprint graph, it was clear the robot instead had a near perfect rhythmicity score—though the leg which had aligned with the auditory stimulus was not the leg being evaluated. Additionally, the individual appeared to have a clearly developed gait that was sustained across changes to the stimulus frequency with little to no variation between gait cycles (Fig. 5a). These are common features of locomotion in biological quadrupeds and are facilitated by central pattern generators in the spinal cord.

To determine the adaptability of the controller, it was tested without alteration on its ability to facilitate entrainment to three new frequencies of 1Hz, 1.25Hz, and 2.5Hz. The results suggest that the controller's performance remained largely consistent across frequencies (Fig. 5b). Interestingly, the only possible exception to this is the 2.5Hz frequency, although the behavior of the robot when this frequency becomes active appears more like a gait transition than a failure to entrain, as the robot adopts a different cyclic pattern of movement based on the rate of the impulses. Gait transition as the speed of locomotion approaches a limit of 'comfort' is biologically realistic, allowing animals to minimize the energy cost of travel [18]. In fact, Alexander's formal definition of gait is formed around this concept. He asserts "A gait is a pattern of locomotion characteristic of a limited range of speeds, described by quantities of which one or more change discontinuously at transitions to other gaits" [18]. Thus, we can say this controller demonstrates at least two gaits and that the point of transition between the two occurs somewhere between 2Hz and 2.5Hz. We can also say that this controller appears to be adaptable, with a range of entrainment spanning at least from 1Hz to 2.5Hz.

Central pattern generators are defined as neuronal circuits that generate coordi-

nated, rhythmic output in the absence of patterned sensory input [1, 2]. Thus, to determine whether the controller could be considered a central pattern generator, we tested the same controller in the absence of external input Fig. 5c. After beginning the simulation with a 1.33Hz stimulus, external impulses were temporarily suspended. The footprint graph produced demonstrates the robot almost immediately changes its pace and settles into locomotion with the same gait at a slightly higher frequency. When the stimulus returns at a frequency of 2Hz, the robot again matches the input frequency. **We can therefore claim that the controller evolved is a central pattern generator by definition.**

Finally, to gain insight into what neural mechanisms underlie this behavior, we plotted the spiking activity of the auditory neuron and values of the three hidden neurons across the span of the simulation (Fig. 6b). While it is difficult to make claims about the exact mechanisms of control, there are several features of the graph that bear resemblance to their biological equivalents. First, hidden neurons appear to spike at regular intervals, followed by a period of negativity before returning to baseline. Similarly, the action potentials of biological neurons begin with depolarization followed by a rebound period of hyperpolarization as ion concentrations normalize. During this time, neurons are temporarily inactivated and another action potential cannot be generated. It is this property of neurons that is (most notably) exploited by central pattern generators to produce oscillatory output in the absence of oscillatory sensory input. Rhythm generation is an emergent property of synaptically-coupled neurons that are non-rhythmic in isolation. Through reciprocal inhibition, a mechanism which elegantly harnesses intrinsic properties of neuronal membranes, non-rhythmic neurons can instead be made to exhibit the same firing patterns as

rhythmic pacemaker neurons [2]. The synchronization between the first two hidden neurons and their apparent alternation with the third hidden neuron also supports the conclusion that the controller evolved to solve for locomotor entrainment with a version of reciprocal inhibition. This result demonstrates the potential merits of a minimalistic evolutionary approach to certain challenging problems in robotics—life finds a way, so perhaps our machines also can.

5 Acknowledgements

This work was supported by The University of Vermont College of Engineering and Mathematical Sciences REU program.

Bibliography

[1] Ari Berkowitz. Expanding our horizons: central pattern generation in the context of complex activity sequences. *Journal of Experimental Biology*, 222(20):jeb192054, October 2019.

[2] Eve Marder and Dirk Bucher. Central pattern generators and the control of rhythmic movements. *Current Biology*, 11(23):R986–R996, November 2001.

[3] Paul S. Katz. Evolution of central pattern generators and rhythmic behaviours. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1685):20150057, January 2016.

[4] Allen I. Selverston. Invertebrate central pattern generator circuits. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1551):2329–2345, August 2010.

[5] Matteo Lodi, Andrey Shilnikov, and Marco Storace. Design of Synthetic Central Pattern Generators Producing Desired Quadruped Gaits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 65(3):1028–1039, March 2018. Conference Name: IEEE Transactions on Circuits and Systems I: Regular Papers.

[6] Serge Rossignol, Réjean Dubuc, and Jean-Pierre Gossard. Dynamic Sensorimotor Interactions in Locomotion. *Physiological Reviews*, 86(1):89–154, January 2006. Publisher: American Physiological Society.

[7] Auke Jan Ijspeert. Central pattern generators for locomotion control in animals and robots: a review. *Neural Networks: The Official Journal of the International Neural Network Society*, 21(4):642–653, May 2008.

[8] Melanie Jouaiti and Patrick Hénaff. Comparative study of forced oscillators for the adaptive generation of rhythmic movements in robot controllers. *Biological Cybernetics*, 113(5-6):547–560, December 2019.

[9] Andres Espinal, Horacio Rostro-Gonzalez, Martin Carpio, Erick I. Guerra-Hernandez, Manuel Ornelas-Rodriguez, and Marco Sotelo-Figueroa. Design of Spiking Central Pattern Generators for Multiple Locomotion Gaits in Hexapod Robots by Christiansen Grammar Evolution. *Frontiers in Neurorobotics*, 10, July 2016. Publisher: Frontiers.

[10] Alex Szorkovszky, Frank Veenstra, and Kyrre Glette. Central pattern generators evolved for real-time adaptation to rhythmic stimuli. *Bioinspiration & Biomimetics*, 18(4):046020, July 2023. arXiv:2210.08102 [nlin].

[11] Kiyotoshi Matsuoka. Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biological cybernetics*, 52:367–76, February 1985.

[12] Josh Bongard. Using robots to investigate the evolution of adaptive behavior. *Current Opinion in Behavioral Sciences*, 6:168–173, December 2015.

[13] Atılım Güneş Baydin. Evolution of central pattern generators for the control of a five-link bipedal walking mechanism. *Paladyn, Journal of Behavioral Robotics*, 3(1):45–53, March 2012. Publisher: De Gruyter Open Access Section: Paladyn.

[14] Arndt von Twickel and Frank Pasemann. Evolved Neural Reflex-Oscillators for Walking Machines. volume 3561, pages 376–385, June 2005.

[15] Randall D. Beer and John C. Gallagher. Evolving Dynamical Neural Networks for Adaptive Behavior. *Adaptive Behavior*, 1(1):91–122, June 1992. Publisher: SAGE Publications Ltd STM.

[16] H. B. Mann and D. R. Whitney. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1):50–60, March 1947. Publisher: Institute of Mathematical Statistics.

[17] Sture Holm. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979. Publisher: [Board of the Foundation of the Scandinavian Journal of Statistics, Wiley].

[18] Robert McNeill Alexander. *Principles of animal locomotion*. Princeton Paperbacks. Princeton University Press, Princeton, 2006.

[19] Michael Schmidt and Hod Lipson. Age-Fitness Pareto Optimization. In Rick Riolo, Trent McConaghy, and Ekaterina Vladislavleva, editors, *Genetic Programming Theory and Practice VIII*, pages 129–146. Springer, New York, NY, 2011.

[20] N. Bernstein. *The Co-ordination and Regulation of Movements*. Pergamon Press, 1967. Google-Books-ID: mUhzjwEACAAJ.

[21] Julius Verrel, Steven Pologe, Wayne Manselle, Ulman Lindenberger, and Marjorie Woollacott. Coordination of degrees of freedom and stabilization of task variables in a complex motor skill: expertise-related differences in cello bowing. *Experimental Brain Research*, 224(3):323–334, February 2013.

[22] Fiona Manning, Jennifer Harris, and Michael Schutz. Temporal prediction abilities are mediated by motor effector and rhythmic expertise. *Experimental Brain Research*, 235, March 2017.

6 Supplementary Materials

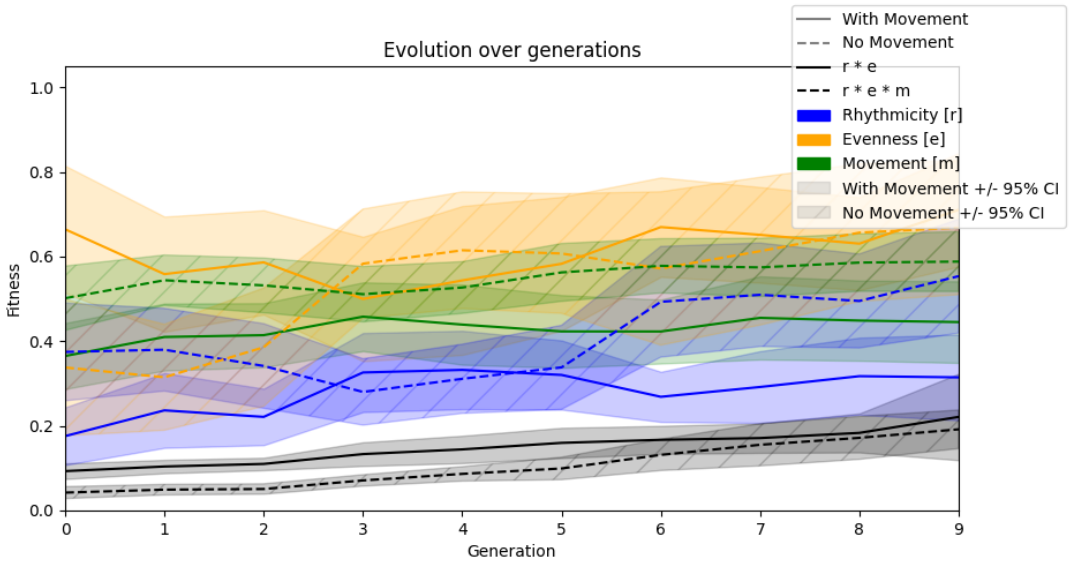

Supplementary Figure 1: ***Comparison of two multi-objective fitness functions.*** *Evolution of 100 individuals for 10 generations with 95% confidence intervals. 10 replicates were performed for each fitness function.*