



OPEN ACCESS

EDITED BY

Andy M. Tyrrell,
University of York, United Kingdom

REVIEWED BY

Edgar Buchanan,
University of York, United Kingdom
Anil Yaman,
VU Amsterdam, Netherlands
Léni Kenneth Le Goff,
Edinburgh Napier University, United
Kingdom

*CORRESPONDENCE

Alex Szorkovszky,
✉ alexansz@ifi.uio.no

RECEIVED 01 June 2023

ACCEPTED 18 August 2023

PUBLISHED 04 October 2023

CITATION

Szorkovszky A, Veenstra F and Glette K
(2023), From real-time adaptation to
social learning in robot ecosystems.
Front. Robot. AI 10:1232708.
doi: 10.3389/frobt.2023.1232708

COPYRIGHT

© 2023 Szorkovszky, Veenstra and
Glette. This is an open-access article
distributed under the terms of the
[Creative Commons Attribution License
\(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or
reproduction in other forums is
permitted, provided the original author(s)
and the copyright owner(s) are credited
and that the original publication in this
journal is cited, in accordance with
accepted academic practice. No use,
distribution or reproduction is permitted
which does not comply with these terms.

From real-time adaptation to social learning in robot ecosystems

Alex Szorkovszky^{1,2*}, Frank Veenstra^{1,2} and Kyrre Glette^{1,2}

¹RITMO Centre for Interdisciplinary Studies in Rhythm, Time and Motion, University of Oslo, Oslo, Norway, ²Department of Informatics, University of Oslo, Oslo, Norway

While evolutionary robotics can create novel morphologies and controllers that are well-adapted to their environments, learning is still the most efficient way to adapt to changes that occur on shorter time scales. Learning proposals for evolving robots to date have focused on new individuals either learning a controller from scratch, or building on the experience of direct ancestors and/or robots with similar configurations. Here we propose and demonstrate a novel means for social learning of gait patterns, based on sensorimotor synchronization. Using movement patterns of other robots as input can drive nonlinear decentralized controllers such as CPGs into new limit cycles, hence encouraging diversity of movement patterns. Stable autonomous controllers can then be locked in, which we demonstrate using a quasi-Hebbian feedback scheme. We propose that in an ecosystem of robots evolving in a heterogeneous environment, such a scheme may allow for the emergence of generalist task-solvers from a population of specialists.

KEYWORDS

social learning, evolutionary robotics, entrainment, central pattern generator, cultural evolution

1 Introduction

It is by now a mainstream opinion in robotics and artificial intelligence that a properly intelligent machine can only come about by continually interacting with its environment through some forms of sensory perception-action loops (Pfeifer and Bongard, 2006; Zador et al., 2023). Such situated cognition is a prevalent goal in evolutionary robotics, where robots come to adapt to their environments and exhibit morphological intelligence (Doncieux et al., 2015).

Taking evolutionary robotics to its logical conclusion, the Autonomous Robot Evolution project presents a radically bottom-up approach to design and fabrication of diverse populations of robots with high degrees of autonomy (Hale et al., 2019). The approach is neatly broken down into three components: fabrication of a robot from a genotype, learning in the physical world, and finally “mature life” in which tasks are performed, performance is evaluated, and the robot’s morphology and/or controller is passed to the next-generation. This cycle has been termed the Triangle of Life (Eiben et al., 2013). Of these three stages, learning is currently the least well developed.

There are several reasons for a robot to learn during its lifetime, and not only benefit from its successful ancestors’ genetic material. The “reality gap” refers to controllers evolved *in silico* not behaving the same in the real world due to imperfections in the simulation (Jakobi et al., 1995). Imperfections in manufacture also imply a noisy genotype-to-phenotype mapping. This learning may not even necessarily be a small tweaking of

parameters: a child robot resulting from mutations or crossover is likely to have a different number of sensor inputs and motor outputs to its parents, so the controller architecture often needs to be reconfigured or learned from scratch.

A controller archive is a natural solution to the problem of varying morphologies. This archives a high-fitness controller for each class of morphology (for example, a certain number of inputs and outputs) as a starting point for future child robots in this class (Le Goff et al., 2022). Controllers tuned during lifetime learning can then be passed down to compatible descendants. This approach resembles “quality diversity” schemes (Pugh et al., 2016) in that solutions occupying different parts of the space of solutions are preserved.

When considering a time-varying or heterogeneous environment, however, it is worth asking whether a certain learning and evolution scheme promotes specialists or generalists. Consider, for example, a task that can be solved by either stepping or hopping. If this is used to determine fitness, certain morphology-controller combinations will evolve to do either one or the other. Now, imagine a real-world task appears that requires some combination of stepping and hopping. The evolutionary solution in this case would most likely come from a child of one stepper and one hopper. As this is a new task, a new controller would need to be learned from scratch for the right combination of parents’ morphologies. Additionally, if the more complex task has been learned but is then absent for some time, catastrophic forgetting is likely to occur if it is not explicitly archived (McCloskey and Cohen, 1989).

One solution to generalist task-solving involves multi-objective evolutionary optimization (De Carlo et al., 2021), another common quality diversity technique. In this case the Pareto front will include both specialists closer to the edges favouring different fitness functions, and “jack-of-all-trades” solutions close to the middle. A learning stage can also be implemented to optimize specialized behaviours using multiple copies of the controller, which can be switched between (de Bruin et al., 2023). So, for example, a morphology that accommodates both stepping and hopping can learn a separate controller for each.

Here, we propose an alternative scheme, in which robots learn from each other instead of on their own. That is, we propose situated social learning of a variety of movement patterns from different “teacher” robots. We demonstrate a key component of the proposed learning method on a variety of controllers evolved using multi-objective optimization, as in de Bruin et al. (2023). In principle, the teacher and learner can exist in different regions of the Pareto front and have different morphologies. One advantage of this approach is that either specific behaviours can be human-defined as tasks, and selected upon, or behaviours can emerge spontaneously from the population if they are useful for survival. The latter is an example of open-ended evolution (Taylor, 2012), which by removing potential bounds on complexity intends to produce the “full generativity of nature” (Soros and Stanley, 2014).

A key to the success of the human species is precisely this kind of social learning, which can greatly enhance problem solving abilities (Herrmann et al., 2007) and the accumulation of knowledge and skills over time (Boyd et al., 2011). Not only does this accumulation take place “vertically,” from older to younger kin, but also “horizontally” across whole societies. Identifying the conditions in which genes and culture co-evolve, and the aspects of

cognition that make it possible, are primary goals of dual inheritance theory, or biocultural evolution (Boyd and Richerson, 1985). This differs from the “memetic” approach influential in computer science (Neri and Cotta, 2012) in that it is based on behaviours rather than information, and is hence a more suitable framework for situated agents.

Culture, in this sense of knowledge, practice and artifacts preserved via non-genetic means across generations, is not only confined to those preserved through syntactical language. It also includes gesture, dance, vocal calls, music and tool-use transmitted through action imitation. Recent studies have shown evidence of social learning of several of these behaviours in non-human animals, indicating those that higher forms may be built upon (Whiten, 2021; Aplin, 2022). These basic forms of social learning involve copying of another agent’s behaviour, followed by its transformation into autonomous behaviour. The propagation of behaviour from agent to agent in this way is termed “cultural transmission” (Mesoudi and Whiten, 2008).

A behaviour, in our case, is communicated as a periodic rhythm via a series of impulses (for example, indicating swing-to-stance transitions). The “learner robot” first synchronizes to the input, a process known as rhythmic entrainment (Schachner et al., 2009), and then self-synchronizes to preserve the resulting motion pattern. Insofar as the frequency or pattern of ground contact indicates a behaviour, this form of communication allows specific behaviours to not be restricted to a particular area of morphological space.

We will first review current and potential generative approaches to social sensorimotor learning, and then demonstrate a scheme to achieve this in a diverse population of central pattern generator (CPG) based robots. The first ingredient of this process, CPGs with the ability to spontaneously entrain to periodic stimuli, has recently been achieved (Szorkovszky et al., 2023a). We will first demonstrate that Hebbian-like spike timing-dependent plasticity can be used to lock-in movement patterns achieved through synchronization. Using autocorrelation functions, we characterize both the diversification of movement patterns, as well as the cultural transmission of patterns from teacher to learner. Finally, we propose how this approach can be incorporated into a learning scheme for evolving robots.

2 Related work

2.1 Imitation and sensorimotor learning

A number of subfields of robotics and computational neuroscience have already successfully modelled aspects of social sensorimotor learning. For robot arms, learning from demonstration is a common technique, where operator training data are generalized into smooth dynamical systems with stable fixed points or limit cycles at desired positions in absolute or relative space (Ravichandar et al., 2020). Using extra dimensions in the dynamical system can even allow multiple overlapping limit cycles, such as clockwise and anticlockwise circles (Khoramshahi and Billard, 2019). Less common is robotic imitation and coordination of gestures from visual signals (Billard and Arbib, 2002; Arbib et al., 2014). However, much progress has been made in this area, and it

has recently been proposed as a plausible means for open-ended evolution, particularly with conditions that encourage spontaneity (Ikegami et al., 2021).

Social and collective behaviours are also commonly studied using wheeled robots such as e-pucks. Often this is through directly copying controller parameters (Heinerman et al., 2015; Bredeche and Fontbonne, 2022). A notable exception is Winfield and Erbas (2011) in which robots attempted to copy each others' trajectories based on their visual perception of their neighbours. This results in "noisy imitation" and hence increased variation in behaviours. Another interesting application of collective robotics is the use of synchronization to identify faulty neighbours (Christensen et al., 2009).

Another field of intense research is in vocal learning: the production and imitation of speech sounds with biophysical models. These studies use a variety of machine learning methods to maximize the match between perceived and produced sounds (Pagliarini et al., 2020), including reward-modulated spike timing dependent plasticity (Warlaumont and Finnegan, 2016). Early work in this area showed that a discrete set of vocal sounds can emerge in a self-organized fashion from mutual interaction (Oudeyer, 2005).

More general motor pattern learning is also an active topic in computational neuroscience. Here, large reservoirs of recurrent spiking neurons are commonly used to model the learning of arbitrary patterns in time, generally with feedback control of chaotic outputs (Sussillo and Abbott, 2009; Laje and Buonomano, 2013). This method takes advantage of such networks' universalizability (Maass and Markram, 2004) and capacity for multifunctionality (Flynn et al., 2021).

2.2 CPG-based entrainment

For locomotion, robotic systems have been created that can adapt frequencies of movement to intrinsic body mechanics. In Buchli et al. (2006), CPG controller parameters were continuously modified according to phase-error feedback. In this case, once the feedback is turned off, the learned behaviour can be set in place by its new parameters. However, the need to employ a phase variable and to calculate an error signal limits the potential complexity of inputs.

Neuron-based CPGs, which due to their nonlinearity are faster and more flexible in their adaptation, have also been used in feedback loops to adapt to mechanical resonances in real time (Williamson, 1998; Iwasaki and Zheng, 2006). In the case of locomotion, force feedbacks can enable adaptation to physical environments, even with interconnections between CPG modules disabled (Thandiackal et al., 2021).

It has recently been demonstrated that self-organization of a locomotion CPG without feedback is sufficient for real-time entrainment to external rhythms, such as those transmitted through audio (Szorkovszky et al., 2023a). This can be seen as an embodied version of the "dynamic attending" approach to beat perception, which was proposed using abstracted nonlinear oscillators (Large and Jones, 1999). More broadly, this falls within the approach of exploiting self-organized nonlinear dynamics in order to generate complex behaviours (Steingrube et al., 2010; Husbands et al., 2021).

3 Materials and methods

3.1 Virtual robots

We begin with the same modular CPG and recurrent filter design as in Szorkovszky et al. (2023a). Each limb module contains three Matsuoka neurons modified to have input-dependent frequency. Two motor neurons drive the joints for each limb, while one interneuron is used for inter-limb coupling (see Figures 1A, B). The use of highly nonlinear CPGs and recurrent networks was intended to encourage self-organization and avoid fixed gait periods.

A constant input is applied to all neurons in a module, each with its own coefficient, to model the brain-stem modulation of the gait. In addition, external inputs are fed to the interneurons through a single-layer recurrent neural network (RNN) (see Figure 1C) consisting of six neurons of the same type as the CPG.

We used two virtual robot body layouts, simulated in Unity using the ML agents package (Griminger et al., 2020). The first body type is a short-legged quadruped, based on the Open Dynamic quadruped (Juliani et al., 2018), and previously studied in Szorkovszky et al. (2023a). For this body, each motor neuron drives one joint angle (see Supplementary Material). Two control parameters were used to evaluate a controller's flexibility: the first was the brain-stem drive, which typically affects gait period and amplitude. For the quadruped, an offset angle of the upper joint was also used to modulate the forward-backward direction of motion via the centre of mass.

The virtual hexapod body design is based on the 18-DOF robot used in Allard et al. (2022). Motor neuron output A was connected to the horizontal coxa joint, while output B was connected to both vertical joints, with two independent coefficients included in the genotype (see Supplementary Material). A simple coefficient for the coxa joint (-1 to 1) was used to modulate the direction of motion.

3.2 Evolution

The multi-objective evolutionary algorithm NSGA-III (Deb and Jain, 2013) was used to simultaneously optimize CPGs for flexibility in speed and direction, as well as stability, resulting in a diverse range of CPGs spanning the Pareto front. The genotype included CPG parameters, interconnection weights and tilt feedback weights [with the same ranges as in Szorkovszky et al. (2023a) for both morphologies] as well as central joint angles, angle limits and joint amplitude coefficients (see Supplementary Material of this paper for ranges).

Each evaluation was in three stages of 10 s each: 1) negative direction parameter with decreasing brain-stem drive, targeting fast backwards motion; 2) positive direction parameter with low and constant brain-stem drive, targeting steady forwards motion; and 3) positive direction parameter with increasing brain-stem drive, targeting fast forward motion. The fitnesses were calculated as.

$$F_1 = -y_1 - \left(\frac{x_1}{x_0}\right)^2 \quad (1)$$

$$F_2 = 2y_0y_2 - y_2^2 - \left(\frac{x_1}{x_0}\right)^2 \quad (2)$$

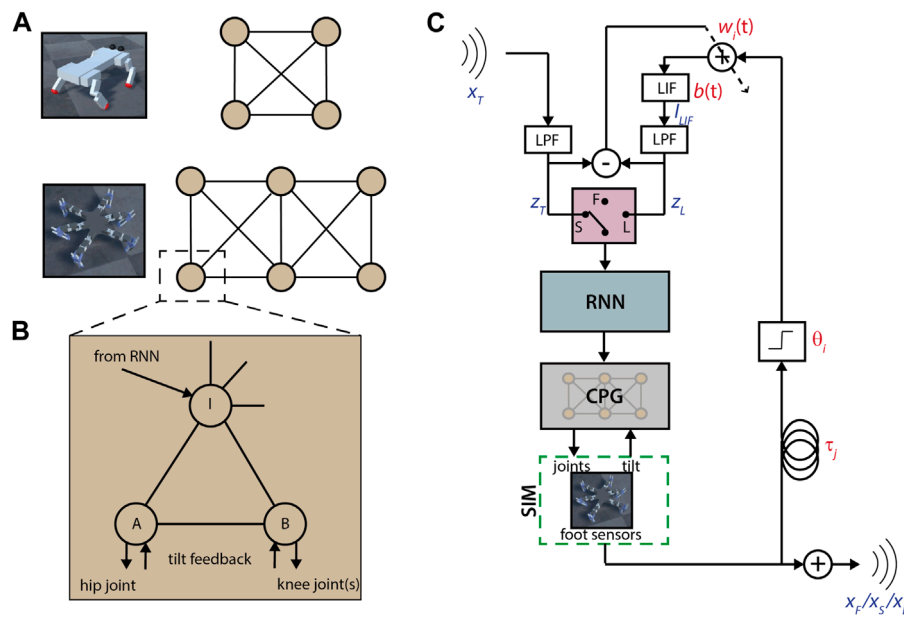


FIGURE 1 Robot controllers and learning scheme. **(A)** Quadruped and hexapod CPG layouts. Each circle is a module, connections are to the module’s interneuron. All connections between modules are inhibitory. **(B)** Diagram of a single CPG module. Each circle is a modified Matsuoka neuron. Connections with arrows are unidirectional, otherwise bidirectional. All connections can be inhibitory or excitatory. **(C)** Diagram of the learning scheme. Selection between free (F), synchronization (S) and learned feedback (L) stages is depicted as a switch. Variables shown in red are parameters, variables in blue are time series. The input signal x_T is the free output x_r of another agent. LPF: low pass filter; LIF: leaky integrate-and-fire neuron; RNN: recurrent neural network of 6 fully connected modified Matsuoka neurons with inhibition, thresholded so there is no output in the absence of input. During synchronization, $b(t) = 1$ and during the feedback stage, each $w_j(t)$ is fixed.

$$F_3 = y_3 - \left(\frac{x_1}{x_0}\right)^2 \tag{3}$$

$$F_4 = \frac{y_0^2 H_{tot}}{1 + t_{tot}} \tag{4}$$

where y_i and x_i were the parallel and perpendicular motion, respectively, in the i th stage. The perpendicular terms, with constant $x_0 = \sqrt{5}$ m, were introduced to discourage turning. F_2 is maximized at $y_2 = y_0 = 2.5$ m, while F_1 and F_3 are unbounded. The fourth fitness F_4 targets stability across the entire evaluation, where H_{tot} is the mean height (normalized and limited for a maximum of 1), and t_{tot} is the root mean square body tilt.

Five runs were performed for each morphology, with each run using a population of 168 individuals evolving over 200 generations. After each run, four controllers were selected, each preferentially weighting one of the four fitness functions:

$$F_m^* = zF_m + \sum_{k=1}^4 F_k, \tag{5}$$

where z was incremented in intervals of one until the maximum of each F_m^* was unique, or until a limit of $z = 100$ was reached.

For both morphologies, two of the post-evolution selections using Eq. 5 did not converge, meaning only three unique CPGs were output instead of four. Therefore, only 18 out of a possible 20 CPGs were selected for each, making a total of 36. For each of these CPGs, recurrent filter layers were evolved for maximum entrainment to a repetitive impulse pattern with a range of periods (Szorkovszky et al., 2023a).

3.3 Learning scheme

Using these 36 virtual robots, exhibiting a range of gait styles, we now consider firstly whether they can entrain to each others’ movement patterns, and secondly whether they can learn its entrained motion pattern. Entrainment ability has been demonstrated for various external periodic inputs in the quadruped morphology (Szorkovszky et al., 2023a; b). However, due to the open-loop nature of this entrainment, the robot’s original movement pattern reappears shortly after the input stops.

While reservoir-based methods can successfully learn to replicate time-series inputs using spike-timing dependent plasticity, as detailed in Section 2.1, these require large numbers of neurons to work effectively (Ganguli et al., 2008; Sussillo and Abbott, 2009). Instead, we approximate the outcome of a larger reinforcing feedback network by using time-delayed impulses from foot sensors to approximate the rhythmic input (see Figure 1C). These feedback connections are separate from the CPG and RNN modules, which are kept fixed and define the free motion and stimulus response, respectively. Therefore, the robot can return to its free motion by removing the feedback, or even switch between different learned motor patterns by interchanging the learned feedback parameters.

We consider every possible pairing of one “teacher,” which transmits its rhythmic step pattern, and one “learner” which attempts to entrain to the pattern and learn it. That is, each robot attempts to learn a new pattern from every other robot. The procedure for each teacher-learner pair consists of three stages,

elaborated in the following sections. First, the learner synchronizes in an open-loop fashion to impulses from the teacher's steps. Secondly, feedback parameters are learned to approximate the teacher's input while still in a synchronized state. Finally, the teacher input is replaced with the feedback signal, and the learner then continues its behaviour autonomously, in a closed-loop "self-synchronized" state.

Foot sensors were added in simulation to record all instances of swing-to-stance transitions. Each of the 36 controllers was initially run for 60 s (4,000 frames at $\Delta t = 15$ ms) to record its free motion, which is also its "teacher" output. This was the sum of impulses for each foot at each time step. The amplitude for each foot was one-half of the impulse amplitude used during evolution. In addition, the natural period of the robot's free motion was calculated using autocorrelation of the CPG outputs.

Of the 36 controllers, five were discarded due to foot-dragging behaviour, resulting in an average of less than one step every two periods, and hence very low output levels. All teacher-learner combinations from the remaining 31 robots were run in the following three stages, each of 60 s duration.

3.4 Synchronization stage 1: feedback delay learning

The impulses in the teacher output were run through an exponential low-pass filter, using the decay rate in the learner's genotype. This decay rate was optimized during evolution for entrainment to rhythmic impulse patterns. This input (z_T) was then passed through the recurrent filter layer to the CPG.

After 60 s, a two-time cross-correlation was performed between each leg's impulse output and the low-pass filtered input. All peaks were then identified with height greater than zero, with a distance of more than 1/20th of the learner's CPG period from any higher peak, and with a lag of less than the learner's CPG period.

The average number of foot sensor outputs per cycle was also calculated, and this was used as a threshold θ_i for limb i . This is because a synchronized system containing n teacher impulses and m_i foot steps per cycle for limb i is expected to generate nm_i cross-correlation peaks for each limb. Therefore, if m_i is greater than one, generating one feedback impulse per step will produce more than the n impulses in the input.

3.5 Synchronization stage 2: feedback weight learning

The second synchronization stage is simply a continuation of the first, but now the foot sensors with the calculated delays are fed into a leaky integrate-and-fire (LIF) neuron to combine them into a single feedback signal $z_L(t)$. Weights for the delayed impulses are learned continuously in order to match $z_L(t)$ to the input $z_T(t)$ as closely as possible.

The impulse signals per limb S_i are created using the delays $\tau_1 \dots \tau_p$, where p is the number of cross-correlation peaks. and thresholds θ_i

determined from the previous stage:

$$S_i(t) = I \left(\sum_{j=1}^p s_i(t - \tau_j) \geq \theta_i \right), \quad (6)$$

where $s_i(t)$ is 1 if the i th foot sensor is triggered during a three frame window centred at time t , and 0 otherwise, and I is an indicator function. A learner output $z_L(t)$ is then generated on-the-fly using the following update equation:

$$\Delta z_L(t) = I_{\text{LIF}}(t) - \gamma z_L(t), \quad (7)$$

where γ is the learner robot's low pass filter decay parameter. We use a standard leaky integrate-and-fire output current I_{LIF} and membrane voltage V_{LIF} , where

$$I_{\text{LIF}}(t) = \begin{cases} I_{\text{out}} & \text{if } V_{\text{LIF}}(t-1) > b(t) \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

and

$$V_{\text{LIF}}(t) = \begin{cases} 0 & \text{if } V_{\text{LIF}}(t-1) > b(t) \\ \sum_i h(w_i(t)) S_i(t) + (1 - \Gamma) V_{\text{LIF}}(t-1) & \text{otherwise.} \end{cases} \quad (9)$$

Here, $h(x)$ is a rectified linear unit, meaning that the inputs are strictly excitatory. The firing threshold $b(t)$ is set to a constant value of one during the learning stage. For the results presented in this study, the LIF decay was set to $\Gamma = 10 \text{ s}^{-1}$ unless specified otherwise.

At the same time, the feedback weights are updated according to:

$$\Delta w_i(t) = \begin{cases} \frac{a}{m} (z_T(t) - z_L(t)) \sum_{t'=0}^T e^{-\Gamma t' \Delta t} S_i(t-t') & \text{if } w_i(t-1) < w_{\text{max}} \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

where $a = 0.01$, m is the number of feet, $T = 2/(\Gamma \Delta t)$ rounded to the nearest integer, w_{max} is the maximum weight (set to 1.5 for this study), and the beginning weights $w_i(0)$ are zero for all i .

3.6 Feedback stage

In the final stage, the teacher input is replaced with the feedback from the feet governed by Eq. 7 with the feedback weights w_i fixed at their final learned values. Since stability cannot be guaranteed in the closed-loop case, we use a homeostatic feedback to stabilize the output level [Husbands et al. \(2021\)](#). We therefore adjust the LIF firing threshold $b(t)$ if the feedback level is not close to the time-averaged teacher input $E_t[z_T]$:

$$\Delta b(t) = \begin{cases} \eta (\bar{z}_L - E_t[z_T])/k & \text{if } \bar{z}_L < E_t[z_T]/k \\ \eta (\bar{z}_L - kE_t[z_T]) & \text{if } \bar{z}_L > kE_t[z_T] \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

where \bar{z}_L is a moving average of the last 200 frames (3 s), $b(0) = 1$, $\eta > 0$ and $k > 0$. Hence, if too many foot sensors are triggered (for example, due to noise or imperfect feedback), producing excess input, the threshold is raised in order to lower the rate of firing of the LIF neuron. Likewise, if the robot is not stepping enough to produce the correct feedback level, the threshold is lowered. In this study, we use $k = 1.5$, so that there is a tolerance of 50% in the feedback level, and an increment $\eta = 2 \times 10^{-4}$.

3.7 Analysis

Analysis of the controllers' similarity of output patterns was done using autocorrelation functions. The differences between movement patterns for the free, synchronized and learned (closed-loop feedback) trials were calculated using:

$$d(X, Y) = \frac{1}{T_T} \int_0^{T_T} (R_{XX}(\tau) - R_{YY}(\tau))^2 d\tau \quad (12)$$

where X and Y are two low-pass filtered time series, T_T is the CPG period of the teacher, and $R_{xx}(\tau)$ is the autocorrelation of time-series x at lag τ , calculated over the last 40 s of each stage. This function, rather than the cross-correlation was used as it is relatively insensitive to relative phase drifts. Low-pass filters were used due to the inherent noisiness of correlation functions when one or both time series consist of short spikes.

The period of motion was also determined from the autocorrelation functions of the joint outputs. For each limb, a complex signal was made with a real part corresponding to the hip joint, and imaginary part corresponding to the knee joint, using the last 40 s of each stage. The real parts of the autocorrelation functions were then averaged, and the lag at the highest peak at $\tau > 0.1$ s was used as the period, while its height was used as a measure of stability.

4 Results

We focused on two abilities. First is the ability of the learner to substantially modify its movement pattern temporarily through synchronization to a teacher, and/or in a lasting way through applied feedback. This ability implies a large synchronized-to-free difference $d(x_S, x_F)$ and learned-to-free difference $d(x_L, x_F)$, respectively. These were calculated for every teacher-learner pair.

Another important ability is to retain the input signal in the feedback, and then to transmit it with some fidelity. This can

be quantified by the difference between the feedback signal and the input $d(z_L, z_T)$ during the closed-loop feedback stage, and the difference between the teacher input and the final learner output that would be transmitted further $d(x_L, z_T)$. Both of these are low for well performing pairs.

Synchronization of learner to teacher was largely learner-dependent, and did not show substantial preferences for the same morphology (see Figure 2A) or evolutionary run (see Supplementary Material for unsorted teacher-learner pair plots). Average within-teacher variability of the synchronized-to-free difference $d(x_S, x_F)$ was 72% greater than average within-learner variability (standard deviation of 0.16 compared to 0.093). Although some learners rarely succeeded to proceed to the feedback stage (most often due to reduced motion), those that did synchronize successfully tended to also lock in new motion patterns in the feedback stage, as shown by Figure 2B. Here, the within-learner variability was more than twice the within-teacher variability (standard deviation of 0.176 compared to 0.087). In other words, some gaits are more able to be modified than others, while the characteristics of the teacher's gait appear relatively unimportant.

Another way to show diversification is to examine movement characteristics such as speed and rotation. An example of a learner's movement profile modified by the learned feedback is shown in Figure 3. In this example, the learned motion can be seen to be outside the normal range of motion patterns controlled by the brain-stem drive. Switching between the intrinsic and learned motion is achieved by simply turning on the feedback.

The synchronized motion, as expected, often matched the period of the teacher, or a multiple thereof, as shown by Figure 4A. By comparison, robots with learned feedback tended to drift away from these periods, most often to a shorter one, and approached the input-free distribution.

Although feedback learning was less reliable than synchronization at copying the teacher's gait period, learned controllers were more stable than the synchronized controllers,

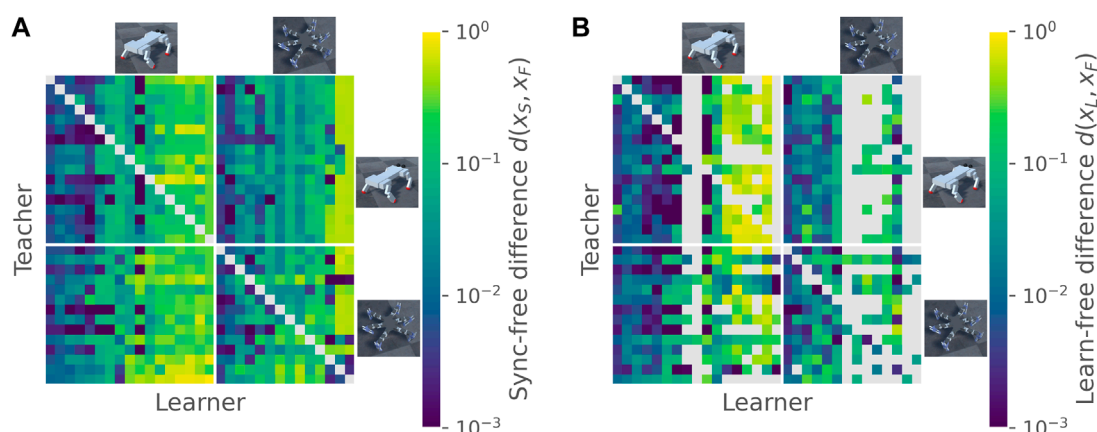


FIGURE 2

Diversification of gait pattern. Panel (A) shows the difference between synchronized and free gaits for all pairs of teachers and learners, ordered by morphology. Lower values indicate a smaller difference. Within each morphology, individuals are ordered by their mean synchronization difference as a learner. Diagonals are left blank since individuals were not tested against themselves in teacher-learner pairs. Panel (B) shows the difference between feedback-learned and free gaits, ordered as in (A). Non-diagonal blank elements indicate that no cross-correlation peaks were found during the period learning stage, and hence the feedback stage was not run.

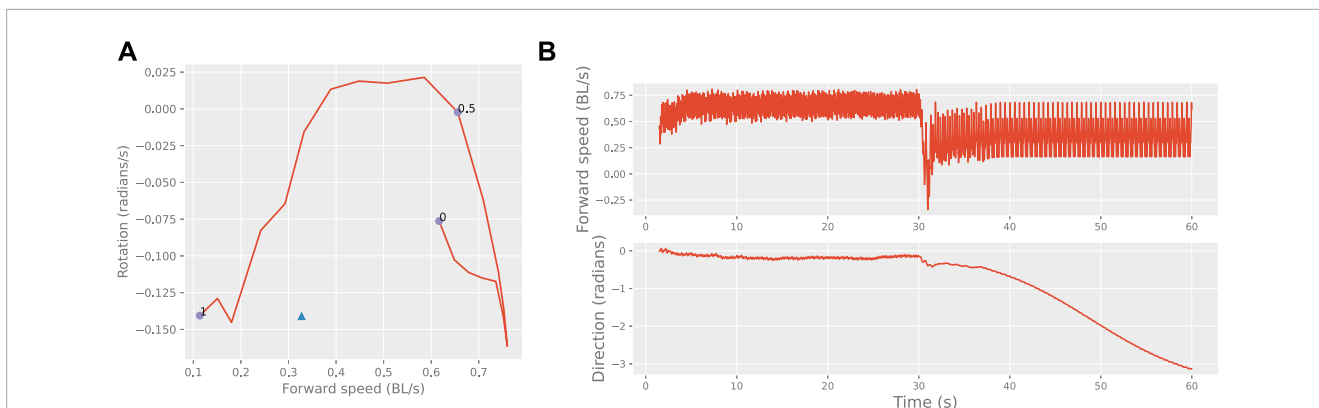


FIGURE 3 Behavioural switching. Panel (A) shows the mean movement characteristics of one learner during free motion as a function of the brain-stem input (line and circles), and the modified movement characteristics under self-synchronization (triangle). BL: body lengths. Panel (B) shows the learner’s switching from its intrinsic movement pattern (brain stem input 0.5) to its learned pattern upon the application of feedback parameters after 30 s.

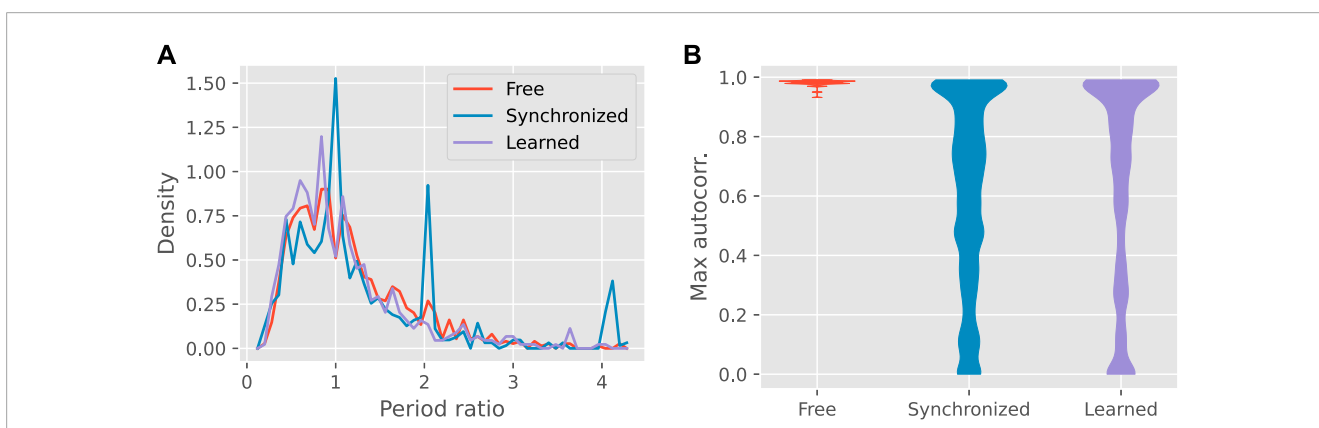


FIGURE 4 Panel (A) shows the frequency distribution of the ratio of the output to input period over all teacher-learner pairs, for the three input conditions. (B) Violin plots of the overall maximum of the autocorrelation function, showing the distributions over all teacher-learner pairs. The distribution for learned gaits only includes pairs for which feedback was actually applied.

as shown by the height of the autocorrelation peak (see Figure 4B). Compared to synchronized gaits, learned gaits were more likely to have autocorrelation peaks near zero or one. The overall median for learned gaits was 0.76 compared to 0.66 for synchronized gaits (Mann-Whitney U-test: $p < 0.001$).

Upon both synchronization and feedback learning, agents that deviated more from their free gait were less stable, as shown by Figure 5. Some, however, had gait patterns closer to the input than their own free gait, shown by points on the lower-right of the plots ($d(x_S, x_T) < d(x_S, x_F)$, and $d(x_L, x_T) < d(x_L, x_F)$, respectively). A significant proportion of controllers in each stage were closer to the input pattern than the free pattern, as shown at the bottom right of each panel, which we call the “success rate.” Under synchronization, the learner was closer to the input for 43.5% of pairs. After feedback learning, the success rate was 23.7% among pairs where the feedback stage was run (total 15.7%). Hence, for some teacher-learner pairs, the combination of synchronization and feedback led to successful cultural transmission (see Supplementary Material for teacher-learner matrix plots). However, in other cases, the feedback

led to a motion pattern unrelated to both the learner’s free motion and the teacher’s motion, as shown by the points in the top-right corner of Figure 5B. The low autocorrelation in this area implies that the feedback is leading to chaotic behaviour, and that the stabilization method can therefore be further optimized.

An example of successful synchronization and learning is shown in Figure 6A. While the shape of the autocorrelation was largely kept upon feedback learning, there was a slight change in period that increased the learn-to-target difference.

The main parameter that was tuned was the LIF decay rate Γ , which determines the window in which impulses from different limbs can be combined. As shown in Figure 6B, there is a trade-off between stability and flexibility, with a shorter window (higher Γ) increasing the success rate of learning but decreasing the average stability of the final gaits. This illustrates the importance of accurate timing in the feedback, as an impulse from one foot that reaches the LIF neuron early or late will either reduce the accuracy of the input pattern (long window) or destabilize the gait completely (short window).

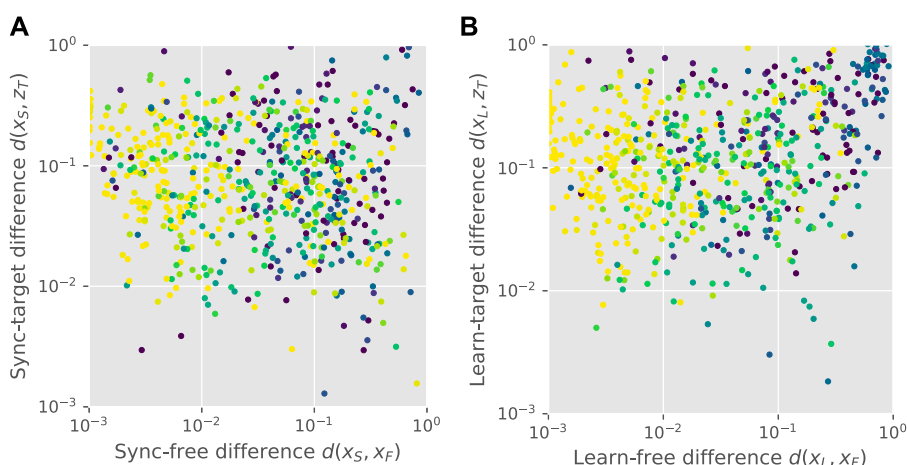


FIGURE 5 Cultural transmission. Panel (A) shows a scatter plot of the sync-free difference against the sync-target difference for all pairs, with colour indicating the autocorrelation peak height, ranging from zero (dark blue) to one (light yellow). Points on the lower-right correspond to synchronized gaits that are closer to the teacher’s gait than the learner’s original gait. Panel (B) shows the equivalent plot for learned gaits.

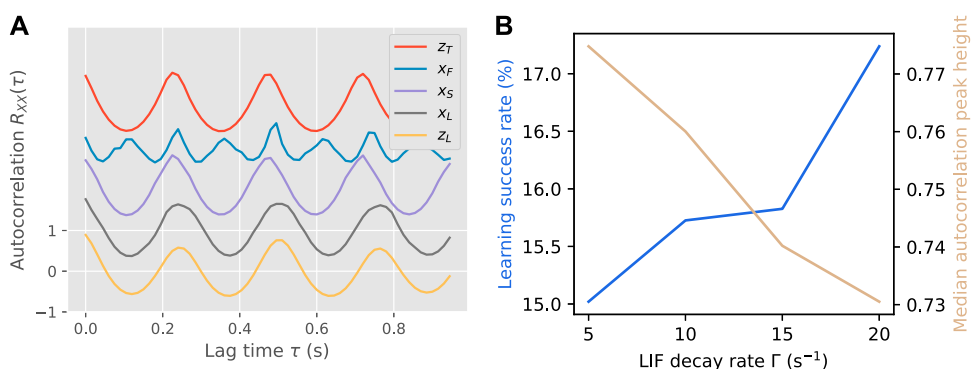


FIGURE 6 Panel (A) shows an example of autocorrelation functions modified by synchronization and learning, for the same teacher-learner pair as in Figure 3. Each autocorrelation function is offset from the previous by 1 for visibility. For this pair, $d(x_S, x_F) = 0.279, d(x_L, x_F) = 0.240, d(x_S, z_T) = 0.002, d(x_L, z_T) = 0.145$. Panel (B) shows the learning success rate and median autocorrelation peak height over all teacher-learner pairs as a function of the LIF decay rate Γ .

5 Discussion

We have demonstrated a proof of principle of a simple social learning scheme for robot gaits. Useful behaviours can be imitated by only communicating a series of foot contact events, such as via audible footsteps. Depending on fidelity, the imperfect copying that we demonstrate (Winfield and Erbas, 2011) can generate behavioural diversification and/or cultural transmission, which can be seen as population-level processes of exploration and exploitation, respectively.

Our scheme is a potential starting point towards robot to robot imitation in evolving populations where morphologies can differ drastically. Alternatives involving visual processing are computationally intensive and rely on an understanding of an agent’s own body and that it is imitating. New ways of communicating

behaviour, either implicitly or explicitly, using rhythmic signals, would be a valuable continuation on this path.

The ability to entrain to a stimulus is a prerequisite of the demonstrated technique. However, the CPG architecture used here is modular, so can be straightforwardly implemented in the modular framework typically used in an evolutionary robotics setting Hale et al. (2019).

Our approach allows for multiple behaviours to be learned and switched between. Therefore, a learning environment could be composed of several elementary tasks, and individuals learn new tasks from teachers who have mastered them. We found that individuals differed substantially in their ability to learn new motion patterns, as quantified by the average difference between original patterns and those learned from various teachers. By incorporating the performance from numerous tasks into a fitness

function, good generalist learners could therefore be selected by evolution.

Another possibility is for individuals to learn a behaviour-environment mapping, so that they automatically decide which motion pattern to use in an environment, and new individuals learn both the mapping and motion pattern from their neighbours. This unsupervised scheme avoids imposing artificial divisions in the learning environment and defining numerous tasks. Such an absence of designed goals could help to satisfy the theoretical requirements for the goal of open-ended evolution (Soros and Stanley, 2014). This lack of oversight, however, comes with risks of maladaptive behaviour spreading quickly, and may require precautionary safeguards (Eiben et al., 2021).

We found that on average, closed-loop feedback was more stable than the open-loop synchronization, despite pronounced delays in the feedback. Delayed feedback has been successfully used to generate a variety of motion patterns from discrete-time chaotic oscillators (Steingrube et al., 2010), and our results show this may also be a promising avenue for continuous-time CPGs with impulse feedback. The integrate-and-fire neuron with self-adapting threshold was a key factor in the stability of the learned gaits. The stability could further be increased using a more detailed spike-timing dependent plasticity (Kempster et al., 1999), which may be useful for adapting to different terrains.

Culture, like any human behaviour, is not trivial to generate in an artificial setting. However, there are clues pointing to the bootstraps that it is lifted by. It is believed that copying specifically via imitation of actions, as opposed to emulation of outcomes, is crucial for sustaining cultural transmission of complex behaviours (Whiten et al., 2009). To better understand how to sustain complexity, iterated learning over many individuals can be tested (Ravignani et al., 2016).

Our work can also inform research into the function of rhythmic entrainment. Due to the fact that animals that can entrain to a beat are often skilled at vocal mimicry, it has been widely theorized that these processes are built on the same neural substrate (Schachner et al., 2009). Forms of entrainment have also been linked to faculties for temporal prediction (Patel and Iversen, 2014), turn-taking (Takahashi et al., 2013), separating other agents from objects (Premack, 1990), and more advanced social capabilities such as joint attention (Knoblich and Sebanz, 2008).

In general, it is important to understand the link between rhythmic movement and cognition. Neuromorphic models of locomotion are potentially crucial for a bottom-up development of intelligence as they are dynamical systems that can exhibit attractor states, a proposed solution to the symbol grounding problem (Pfeifer and Bongard, 2006). Opening these systems to inputs from other networks and sensory data from the physical environment leads to “open dynamical systems,” which constantly adapt in their attractor landscapes (Hotton and Yoshimi, 2011; Beer and Williams, 2015). We believe that our proposal fruitfully extends this idea to social environments. It remains to be seen whether agents influencing each other through their basic motion can lead to the emergence of new forms of perception and action.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/aszorko/COROBOREES/tree/Paper4>.

Author contributions

AS and KG devised the study. AS and FV created the virtual robot simulation. AS ran the experiment and analyzed the results. AS wrote the first draft of the manuscript with subsequent contributions from FV and KG. All authors contributed to the article and approved the submitted version.

Funding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101030688, and is partially supported by the Research Council of Norway through its Centres of Excellence scheme, project number 262762.

Acknowledgments

We would like to thank the organizers and other participants of the ARE workshop in York, United Kingdom, October 2022 for feedback on the talk that gave rise to this study, and the reviewers for helpful comments.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2023.1232708/full#supplementary-material>

References

- Allard, M., Smith, S. C., Chatzilygeroudis, K., Lim, B., and Cully, A. (2022). *Online damage recovery for physical robots with hierarchical quality-diversity*. arXiv preprint arXiv:2210.09918.
- Aplin, L. (2022). Culture in birds. *Curr. Biol.* 32, R1136–R1140. doi:10.1016/j.cub.2022.08.070
- Arbib, M., Ganesh, V., and Gasser, B. (2014). Dyadic brain modelling, mirror systems and the ontogenetic ritualization of ape gesture. *Philosophical Trans. R. Soc. B Biol. Sci.* 369, 20130414. doi:10.1098/rstb.2013.0414
- Beer, R. D., and Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive Sci.* 39, 1–38. doi:10.1111/cogs.12142
- Billard, A., and Arbib, M. (2002). “Mirror neurons and the neural basis for learning by imitation: Computational modeling,” in *Mirror neurons and the evolution of brain and language*. Editors M. I. Stamenov, and V. Gallese (John Benjamins Publishing Company)344–352.
- Boyd, R., and Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Boyd, R., Richerson, P. J., and Henrich, J. (2011). The cultural niche: why social learning is essential for human adaptation. *Proc. Natl. Acad. Sci.* 108, 10918–10925. doi:10.1073/pnas.1100290108
- Brederche, N., and Fontbonne, N. (2022). Social learning in swarm robotics. *Philosophical Trans. R. Soc. B* 377, 20200309. doi:10.1098/rstb.2020.0309
- Buchli, J., Iida, F., and Ijspeert, A. J. (2006). Finding resonance: adaptive frequency oscillators for dynamic legged locomotion. In *2006 IEEE/RSJ international conference on intelligent robots and systems (IEEE)*, 3903–3909.
- Christensen, A. L., Ogrady, R., and Dorigo, M. (2009). From fireflies to fault-tolerant swarms of robots. *IEEE Trans. Evol. Comput.* 13, 754–766. doi:10.1109/tevc.2009.2017516
- de Bruin, E., Hatzky, J., Hosseinkhani Kargar, B., and Eiben, A. (2023). “A multi-brain approach for multiple tasks in evolvable robots,” in *International conference on the applications of evolutionary computation (part of EvoStar)* (Springer), 129–144.
- De Carlo, M., Ferrante, E., Ellers, J., Meynen, G., and Eiben, A. (2021). “The impact of different tasks on evolved robot morphologies,” in *Proceedings of the genetic and evolutionary computation conference companion*, 91–92.
- Deb, K., and Jain, H. (2013). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE Trans. Evol. Comput.* 18, 577–601. doi:10.1109/tevc.2013.2281535
- Doncieux, S., Brederche, N., Mouret, J.-B., and Eiben, A. E. (2015). Evolutionary robotics: what, why, and where to. *Front. Robotics AI* 2, 4. doi:10.3389/frobt.2015.00004
- Eiben, A. E., Brederche, N., Hoogendoorn, M., Stradner, J., Timmis, J., Tyrrell, A., et al. (2013). The triangle of life: evolving robots in real-time and real-space. in *European Conference on Artificial Life (ECAL-2013)*, 1–8.
- Eiben, Á. E., Ellers, J., Meynen, G., and Nyholm, S. (2021). Robot evolution: ethical concerns. *Front. Robotics AI* 8, 744590. doi:10.3389/frobt.2021.744590
- Flynn, A., Tsachouridis, V. A., and Amann, A. (2021). Multifunctionality in a reservoir computer. *Chaos Interdiscip. J. Nonlinear Sci.* 31, 013125. doi:10.1063/5.0019974
- Ganguli, S., Huh, D., and Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proc. Natl. Acad. Sci.* 105, 18970–18975. doi:10.1073/pnas.0804451105
- Grimminger, F., Meduri, A., Khadiv, M., Viereck, J., Wüthrich, M., Naveau, M., et al. (2020). An open torque-controlled modular robot architecture for legged locomotion research. *IEEE Robotics Automation Lett.* 5, 3650–3657. doi:10.1109/LRA.2020.2976639
- Hale, M. F., Buchanan, E., Winfield, A. F., Timmis, J., Hart, E., Eiben, A. E., et al. (2019). “The are robot fabricator: how to (re) produce robots that can evolve in the real world,” in *Artificial life conference proceedings* (MIT Press), 95–102.
- Heinerman, J., Drupsteen, D., and Eiben, A. E. (2015). “Three-fold adaptivity in groups of robots: the effect of social learning,” in *Proceedings of the 2015 annual conference on genetic and evolutionary computation*, 177–183.
- Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B., and Tomasello, M. (2007). Humans have evolved specialized skills of social cognition: the cultural intelligence hypothesis. *science* 317, 1360–1366. doi:10.1126/science.1146282
- Hotton, S., and Yoshimi, J. (2011). Extending dynamical systems theory to model embodied cognition. *Cognitive Sci.* 35, 444–479. doi:10.1111/j.1551-6709.2010.01151.x
- Husbands, P., Shim, Y., Garvie, M., Dewar, A., Domcsek, N., Graham, P., et al. (2021). Recent advances in evolutionary and bio-inspired adaptive robotics: exploiting embodied dynamics. *Appl. Intell.* 51, 6467–6496. doi:10.1007/s10489-021-02275-9
- Ikegami, T., Masumori, A., and Maruyama, N. (2021). “Can mutual imitation generate open-ended evolution,” in *Proceedings of Artificial Life 2021 workshop on OEE*.
- Iwasaki, T., and Zheng, M. (2006). Sensory feedback mechanism underlying entrainment of central pattern generator to mechanical resonance. *Biol. Cybern.* 94, 245–261. doi:10.1007/s00422-005-0047-3
- Jakobi, N., Husbands, P., and Harvey, I. (1995). “Noise and the reality gap: the use of simulation in evolutionary robotics,” in *Advances in artificial life: Third European conference on artificial life granada, Spain, june 4–6, 1995 proceedings 3* (Springer), 704–720.
- Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., et al. (2018). *Unity: a general platform for intelligent agents*. arXiv preprint arXiv:1809.02627.
- Kempler, R., Gerstner, W., and Van Hemmen, J. L. (1999). Hebbian learning and spiking neurons. *Phys. Rev. E* 59, 4498–4514. doi:10.1103/physreve.59.4498
- Khoramshahi, M., and Billard, A. (2019). A dynamical system approach to task-adaptation in physical human–robot interaction. *Aut. Robots* 43, 927–946. doi:10.1007/s10514-018-9764-z
- Knoblich, G., and Sebanz, N. (2008). Evolving intentions for social interaction: from entrainment to joint action. *Philosophical Trans. R. Soc. B Biol. Sci.* 363, 2021–2031. doi:10.1098/rstb.2008.0006
- Laje, R., and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* 16, 925–933. doi:10.1038/nn.3405
- Large, E. W., and Jones, M. R. (1999). The dynamics of attending: how people track time-varying events. *Psychol. Rev.* 106, 119–159. doi:10.1037/0033-295x.106.1.119
- Le Goff, L. K., Buchanan, E., Hart, E., Eiben, A. E., Li, W., De Carlo, M., et al. (2022). Morpho-evolution with learning using a controller archive as an inheritance mechanism. *IEEE Trans. Cognitive Dev. Syst.* 15, 507–517. doi:10.1109/tcds.2022.3148543
- Maass, W., and Markram, H. (2004). On the computational power of circuits of spiking neurons. *J. Comput. Syst. Sci.* 69, 593–616. doi:10.1016/j.jcss.2004.04.001
- McCloskey, M., and Cohen, N. J. (1989). “Catastrophic interference in connectionist networks: the sequential learning problem,” in *Psychology of learning and motivation* (Elsevier), 24, 109–165.
- Mesoudi, A., and Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Trans. R. Soc. B Biol. Sci.* 363, 3489–3501. doi:10.1098/rstb.2008.0129
- Neri, F., and Cotta, C. (2012). Memetic algorithms and memetic computing optimization: a literature review. *Swarm Evol. Comput.* 2, 1–14. doi:10.1016/j.swevo.2011.11.003
- Oudeyer, P.-Y. (2005). The self-organization of speech sounds. *J. Theor. Biol.* 233, 435–449. doi:10.1016/j.jtbi.2004.10.025
- Pagliarini, S., Leblois, A., and Hinaut, X. (2020). Vocal imitation in sensorimotor learning models: a comparative review. *IEEE Trans. Cognitive Dev. Syst.* 13, 326–342. doi:10.1109/tcds.2020.3041179
- Patel, A. D., and Iversen, J. R. (2014). The evolutionary neuroscience of musical beat perception: the action simulation for auditory prediction (asap) hypothesis. *Front. Syst. Neurosci.* 8, 57. doi:10.3389/fnsys.2014.00057
- Pfeifer, R., and Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. MIT press.
- Premack, D. (1990). The infant’s theory of self-propelled objects. *Cognition* 36, 1–16. doi:10.1016/0010-0277(90)90051-k
- Pugh, J. K., Soros, L. B., and Stanley, K. O. (2016). Quality diversity: a new frontier for evolutionary computation. *Front. Robotics AI* 3, 40. doi:10.3389/frobt.2016.00040
- Ravichandar, H., Polydoros, A., Chernova, S., and Billard, A. (2020). Recent advances in robot learning from demonstration. *Annu. Rev. Control, Robotics, Aut. Syst.* 3, 297–330. doi:10.1146/annurev-control-100819-063206
- Ravnigani, A., Delgado, T., and Kirby, S. (2016). Musical evolution in the lab exhibits rhythmic universals. *Nat. Hum. Behav.* 1, 0007. doi:10.1038/s41562-016-0007
- Schachner, A., Brady, T. F., Pepperberg, I. M., and Hauser, M. D. (2009). Spontaneous motor entrainment to music in multiple vocal mimicking species. *Curr. Biol.* 19, 831–836. doi:10.1016/j.cub.2009.03.061
- Soros, L., and Stanley, K. (2014). “Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria,” in *Life 14: the fourteenth international conference on the synthesis and simulation of living systems* (MIT Press), 793–800.
- Steingrube, S., Timme, M., Wörgötter, F., and Manoonpong, P. (2010). Self-organized adaptation of a simple neural circuit enables complex robot behaviour. *Nat. Phys.* 6, 224–230. doi:10.1038/nphys1508
- Sussillo, D., and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557. doi:10.1016/j.neuron.2009.07.018
- Szorkovszky, A., Veenstra, F., and Glette, K. (2023a). Central pattern generators evolved for real-time adaptation to rhythmic stimuli. *Bioinspiration Biomimetics* 18, 046020. doi:10.1088/1748-3190/ace017
- Szorkovszky, A., Veenstra, F., Lartillot, O., Jensenius, A. R., and Glette, K. (2023b). “Embodied tempo tracking with a virtual quadruped,” in *Proceedings of the 20th sound and music computing conference*, 283–288.

- Takahashi, D. Y., Narayanan, D. Z., and Ghazanfar, A. A. (2013). Coupled oscillator dynamics of vocal turn-taking in monkeys. *Curr. Biol.* 23, 2162–2168. doi:10.1016/j.cub.2013.09.005
- Taylor, T. (2012). “Exploring the concept of open-ended evolution,” in *Proceedings of the 13th international conference on artificial life* (Cambridge, MA): MIT Press), 540–541.
- Thandiackal, R., Melo, K., Paez, L., Herault, J., Kano, T., Akiyama, K., et al. (2021). Emergence of robust self-organized undulatory swimming based on local hydrodynamic force sensing. *Sci. Robotics* 6, eabf6354. doi:10.1126/scirobotics.abf6354
- Warlaumont, A. S., and Finnegan, M. K. (2016). Learning to produce syllabic speech sounds via reward-modulated neural plasticity. *PLoS one* 11, e0145096. doi:10.1371/journal.pone.0145096
- Whiten, A., McGuigan, N., Marshall-Pescini, S., and Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Trans. R. Soc. B Biol. Sci.* 364, 2417–2428. doi:10.1098/rstb.2009.0069
- Whiten, A. (2021). The burgeoning reach of animal culture. *Science* 372, eabe6514. doi:10.1126/science.abe6514
- Williamson, M. M. (1998). Neural control of rhythmic arm movements. *Neural Netw.* 11, 1379–1394. doi:10.1016/s0893-6080(98)00048-3
- Winfield, A. F., and Erbas, M. D. (2011). On embodied memetic evolution and the emergence of behavioural traditions in robots. *Memetic Comput.* 3, 261–270. doi:10.1007/s12293-011-0063-x
- Zador, A., Escola, S., Richards, B., Ölveczky, B., Bengio, Y., Boahen, K., et al. (2023). Catalyzing next-generation artificial intelligence through neuroAI. *Nat. Commun.* 14, 1597. doi:10.1038/s41467-023-37180-x