


Learning to crawl: Benefits and limits of centralized versus distributed controlLuca Gagliardi^{*} and Agnese Seminara[†]*Machine Learning Genoa Center and Department of Civil Chemical and Environmental Engineering,
University of Genoa, villa Cambiaso, via Montallegro 1, 16145, Italy* (Received 11 June 2025; accepted 26 February 2026; published 3 April 2026)

We present a model of a crawler consisting of several suction units distributed along a straight line and connected by springs. The suction units are rudimentary proprioceptors-actuators, which sense binary states of compression versus elongation of the springs and can either adhere or remain idle. Muscular contraction is not controlled by the crawler but follows an endogenous, stereotyped wave. The crawler is tasked to learn patterns of adhesion that generate thrust in response to the wave of contraction. Using tabular Q -learning we demonstrate that crawling can be learned by trial and error and we ask what are the benefits and limitations of distributed versus centralized learning architectures. We find that by centralizing proprioceptive feedback and control, the crawler leverages long-range correlations in the dynamics and ride the endogenous wave smoothly. The ensuing benefits are measured in terms of both speed and robustness to failure, although they come at increased computational cost. At the opposite extreme, purely distributed feedback and control only leverages local information and yield a jerkier and slower crawling, although computationally cheap. Intermediate levels of centralization can negotiate fast and robust crawling while avoiding excessive computational burden, demonstrating the computational benefits of a hierarchical organization of crawling. Our model unveils the trade-offs among crawling speed, robustness to failure, computational cost, and information exchange that may shape biological solutions for crawling and could inspire the design of robotic crawlers.

DOI: [10.1103/9snq-xgqh](https://doi.org/10.1103/9snq-xgqh)**I. INTRODUCTION**

All organisms with a distributed nervous system need to coordinate their sensorimotor loop without descending commands from a central nervous system, typical of bilateral animals. Despite the lack of centralized control, many distributed nervous systems reliably manage effective locomotion involving the synchronization of many body parts.

These nervous systems showcase different architectures reflecting varying degrees of distribution versus centralization. In cnidarians, the nervous system consists only of diffuse nerve nets and is the most primitive of all animals'. Despite the lack of centralization, cnidaria are capable of complex behaviors. For example, *Hydra vulgaris* achieves highly non-trivial coordination of its body parts to perform somersaults [1]. Furthermore, associative learning in the form of operant conditioning can be induced by either visual or mechanical stimuli for the box jellyfish *Tripedalia cystophora* [2].

Echinoderms, such as starfish, also have a simple radial nervous system consisting of a nerve net of neurons. But while they have no brain, their nervous system is more structured

than in the cnidarians, as the nerves connect to a central nerve ring [3]. Starfishes crawl efficiently using hundreds of independent tubular feet which generate robust locomotion without central control. The nerve ring establishes the direction of motion by integrating sensory information detected by the arms [4,5]. Given the direction of motion, coordination of the tube feet can emerge purely from a passive mechanical coupling, as showed with a minimal model where behavior of the tube feet is prescribed and no information propagates along the nerve net [6,7].

Cephalopods, such as octopus, squid, and cuttlefish, have evolved a complex nervous system whose size is comparable to that of vertebrates and endows them with impressive cognitive capabilities. They do have a central brain; however, most of their nervous system is distributed throughout their arms which are covered in hundreds of suckers that can sense the environment as well as adhere to substrates. Like all soft-bodied animals, these organisms have to coordinate many degrees of freedom to achieve locomotion. For example, octopuses translocate on surfaces by leveraging their taste-by-touch sense, whereby suckers adhere to surfaces and sense them chemically upon contact [8]. The sensory information acquired by the periphery guides suckers' adhesion, yielding displacement of the organism along the surface. However, to what extent locomotion is controlled by the central nervous system versus by the ganglia distributed along the arms is not known [9,10]. Intriguingly, coordination in octopus may or may not depend on precise proprioceptive feedback, since the molecular mechanisms for proprioception and to what extent the brain stores positional information is not understood [11,12].

*Contact author: luca.gagliardi@edu.unige.it†Contact author: agnese.seminara@unige.it

This question is part of a large body of literature in neuromechanics that broadly asks how do animals and robots control various kinds of distributed locomotion to achieve specific goals including energy efficiency, maneuverability, stability, and speed (see, e.g., Refs. [13,14] and references therein). Approaches that range from a more centralized [15] to a purely distributed control [16] have been put forward for various locomotion mechanisms in bio-inspired robotics, reviewed, e.g., in Ref. [17]. Whether centralization increases locomotion speed is an open question (reviewed, e.g., in Ref. [18]). On the one hand, centralization likely requires additional processing which may be too slow to be implemented at fast speed. On the other hand, the need for added stability at high speed may require centralized control. The combination of central pattern generator (CPG) and local feedback was found beneficial for adaptation and versatility for amphibious centipedes [19], undulatory patterns in swimming and multileg terrestrial locomotion [20], and quadruped locomotion [21].

Here we focus on adhesion-driven crawling: We draw inspiration from octopus and focus on a toy model to ask what are the pros and cons of centralization versus distribution. The crawler is modeled as a collection of suckers distributed along a straight line and connected by springs and endowed with a rudimentary form of proprioception (extension versus elongation of the springs) and a rudimentary form of control of the suckers (adhere versus not adhere). Contraction of the springs is not controlled by the crawler but rather follows a stereotyped wave established by a CPG. If suckers were to adhere at random times, then the crawler center of mass would be unable to move on average: Is it possible to learn a pattern of adhesion that enables net translocation, purely from trial and error?

Learning is an underexplored aspect of crawling that is relevant to develop robotic crawlers in unpredictable environments, for example on irregular terrain or under other environmental perturbations. In these cases, programming the robot to crawl is unfeasible and learning from trial and error becomes a viable option.

To answer this question we develop a simulator of the crawler and let it learn to translocate by trial and error using the Q -learning reinforcement learning (RL) algorithm. We take the approach of fixing a widely used model of coupled oscillators, or template [22], where each unit has the ability to adhere. The mathematical model is similar to a large body of literature on the neuromechanics of crawling largely developed for larvae, worms, and insects (e.g., Refs. [23–25] and references within). Our crawler does not control muscle contraction but only adhesion with the substrate. A main difference with most literature (but see Ref. [26]) is that we do not program the crawler, but let it learn through the use of reinforcement learning, whereby we set the goal of translocating and the agent learns to crawl purely from trial and error. This approach allows us to assess how easy it is for a learning architecture to develop an efficient crawling and how computationally demanding learning is. Efficiency is defined as the combination of fast translocation and robustness to failure of individual suckers. Alternatively, one may target a specific dynamics and learn how to formulate a mathematical model that reproduces a specific attractor in the phase space [27]. We ask whether the sole control of

adhesion is sufficient to learn crawling with a distributed versus centralized architecture. Our results complement previous literature considering patterns of contraction—rather than adhesion—in fully centralized learning of crawling [26] and fully distributed learning of swimming [28], as well as works where crawling emerges from prescribed control patterns, rather than learning (see, e.g., Refs. [6,29]), and where the redundancies of preprogrammed gait design lead alone to locomotion robustness, for instance over rough terrains [30].

We compare the learnability, performance, and robustness of centralized versus distributed learning architectures. Despite the rudimentary proprioception and control, we find that the sole existence of a common reward—here the speed of the center of mass—is sufficient for independent suckers to learn crawling. When all suckers are simultaneously controlled by a single agent, crawling is both faster and more robust to failure, partially overcoming the crawler’s poor sensorimotor ability. However, the computational cost of the central control scales exponentially with the number of suckers and springs. We then introduce a shallow hierarchical architecture where control is centralized in few control centers. We find that this partial centralization achieves nearly optimal performance and robustness with a limited increase in computational burden.

In our setting, crawling is achieved by “surfing” the wave generated by the central pattern generator. The existence of a central pattern generator is commonly assumed and has been demonstrated in various systems (see, e.g., Ref. [31]). Note, however, that in the absence of a central pattern generator, crawling may still be achieved by proprioceptive feedback alone [32–34] or by leveraging synchronization of motor control with some natural frequency of the system (for example in the context of human inspired biped locomotion via imitation learning [35], limb oscillation via experimental analysis [36], or for flapping wings in flight [37]).

We hope that our work will serve as a benchmark to explore the optimal level of centralization in more realistic scenarios for natural and artificial crawling, where constraints on speed, robustness, and computational costs will all concur to define optimality.

II. MODEL AND LEARNING

A. The one-dimensional crawler model

The model, illustrated in Fig. 1, is strongly inspired by the one proposed in Ref. [29] and considers the crawler as a series of suckers (blocks) connected by springs along the body axis. The crawler is placed on a flat substrate and its dynamics is dictated by the springs and the interaction with the substrate:

$$m\ddot{x}_n(t) = -\zeta_n\dot{x}_n(t) + \kappa(x_{n+1}(t) - x_n(t) - l_n(t)) - \kappa(x_n(t) - x_{n-1}(t) - l_{n-1}(t)), \quad (1)$$

where x_n is the position of the n th sucker, with $n = 1, \dots, N_s$; N_s is the number of suckers; m is the mass of each sucker, ζ_n its friction coefficient and κ the spring constant; l_n is the rest length of the n th spring and it oscillates in time as prescribed by the CPG as described below. We assume friction is proportional to the sliding velocity of the suckers, modeling viscous friction with a wet substrate. Each sucker can increase its viscous coefficient ζ_n by adhering with the substrate. For

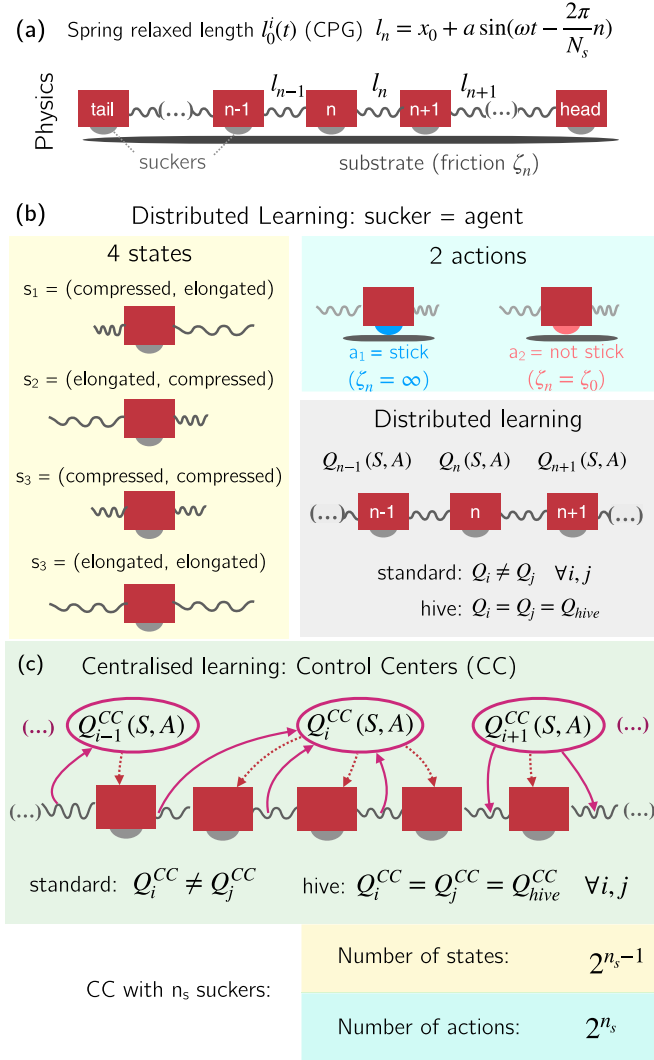


FIG. 1. (a) Sketch of the one-dimensional crawler model. (b) Sketch of the distributed architecture, including the 4 states available to each agent (individual sucker), the 2 actions the sucker controls and the 4×2 Q matrix for each sucker. (c) Centralized architecture: One or more control center (CC) control a subset of contiguous suckers. The states for each CC correspond to all combinations of compression and elongation states of the springs within the CC; the actions are all combinations of adhesion OR no-adhesion of suckers within the CC. For multiagency (whether the individual agents are suckers or CCs), the “hive update” corresponds to forcing all agents to agree, i.e., write and read the same Q matrix.

simplicity, we assume that each friction coefficient can select only two values: a constant value ζ_0 (no adhesion) or $\zeta_n \rightarrow \infty$ (the n th sucker adheres). The choice of infinite adhesion is one of the ingredients distinguishing our model from Ref. [29] where calculations are conducted under the assumption that individual suckers only slightly perturb their friction coefficient.

For the tail, $n = 0$, and the head, $n = N_s$, the dynamics is driven by a single adjacent spring:

$$m\ddot{x}_0(t) = -\zeta_0\dot{x}_0(t) + \kappa(x_1(t) - x_0(t) - l_0(t)), \quad (2)$$

$$m\ddot{x}_{N_s}(t) = -\zeta_0\dot{x}_{N_s}(t) - \kappa(x_{N_s}(t) - x_{N_s-1}(t) - l_{N_s-1}(t)). \quad (3)$$

B. Central pattern generator

We assume an elongation-contraction wave is generated by a CPG and propagates from the tail to the head of the crawler:

$$l_n = x_0 + a \sin(\omega t - kn), \quad (4)$$

where l_n is the rest length of the n th spring, defined as the one to the right of the n th sucker; $k = 2\pi/N$ with N the periodicity; ω the angular velocity; a a constant and uniform amplitude; and x_0 the natural length of the unperturbed spring. The unperturbed rest length of the springs, x_0 , fixes the spatial scaling of the model. The spatial periodicity N in the CPG oscillation is in principle an arbitrary integer. For simplicity, we fix $N = N_s$, so that the wavelength $\lambda = x_0N$ corresponds to the length of the crawler L . Arbitrary values of the periodicity are considered in Appendix B, where we show how N alters the importance of each sucker when a crawling strategy is played. Note that when all suckers are controlled by a single agent, it is possible to learn the contraction patterns for coordinated crawling without assuming a CPG, as shown in Ref. [26]. We assume the CPG is unaffected by the suckers, relevant to real locomotion in the case where proprioceptive feedback has low gain [38].

C. Scaling and parametrization

The above equations can be nondimensionalized by rescaling all length scales with the unperturbed rest length x_0 , $\bar{x} = x/x_0$, $\bar{l} = l/x_0$, and time with the frequency of the CPG, $\bar{t} = \omega t$. In the overdamped limit, inertia is negligible compared to the other terms, $m\omega^2/\kappa \ll 1$, and the dynamics reduces to kinematics:

$$\begin{aligned} \dot{\bar{x}}_n(t) &= B(\bar{x}_{n+1}(t) + \bar{x}_{n-1}(t) - 2\bar{x}_n(t) - \bar{l}_n(t) + \bar{l}_{n-1}(t)) \\ \dot{\bar{x}}_0(t) &= B(\bar{x}_1(t) - \bar{x}_0(t) - \bar{l}_0(t)) \\ \dot{\bar{x}}_{N_s}(t) &= B(\bar{x}_{N_s-1}(t) - \bar{x}_{N_s}(t) + \bar{l}_{N_s-1}(t)) \\ \bar{l}_n &= 1 + \bar{a} \sin(\bar{t} - 2\pi n/N), \end{aligned} \quad (5)$$

where the nondimensional group $B_n = \kappa/(\zeta_n\omega)$ is finite for suckers that do not adhere ($B_n = B_0$) and it vanishes for suckers that adhere ($B_n = 0$, i.e., when the n th sucker adheres its velocity drops instantaneously to zero, $\dot{\bar{x}}_n = 0$). The problem is thus dictated by four nondimensional parameters:

$$\begin{aligned} B_0 &= \kappa/(\zeta_0\omega) = 10 \\ N &= N_s \text{ except for Fig. 10} \\ N_s &= L/x_0 \text{ varying} \\ a/x_0 &= 0.25. \end{aligned} \quad (6)$$

The model equations (5) with the nondimensional parameters (6), the initial condition $\dot{\bar{x}}_n = 0$ and $x_n = nx_0$, and a specific scheduling for the control ζ_n (or its nondimensional form B_n) fully prescribe the dynamics. As detailed in Ref. [29], without coordinated adhesion of each sucker or block, the crawler just oscillates along its center of mass and the average velocity is zero. In contrast, when ζ_n is controlled appropriately, the crawler undergoes net translocation. In the following, we will compare the net translocation velocity v under different strategies for the control of ζ_n . To compare

crawlers of different lengths, we will express velocities in units of the typical velocity $x_0\omega$: $\bar{v} = v/(x_0\omega)$.

1. Expected qualitative behavior from the idealized continuous infinite crawler

In Ref. [29] the authors find an analytical solution for the velocity of the center of mass of a continuous ($N_s \rightarrow \infty$, $x_0 \rightarrow 0$, and $Nx_0 = \lambda$) one-dimensional infinitely long crawler, assuming individual blocks can only perturb slightly the friction coefficient (rather than fully adhere like in our case). For the sake of clarity we summarize here their results. The n th friction coefficient is perturbed according to

$$\zeta_n(t) = \zeta_0(1 + \epsilon p(s, t)), \quad (7)$$

where $s = nx_0$ is the continuous position along the crawler; $\epsilon \ll 1$ is a small parameter and only a single sucker can increase its viscous friction at a time, i.e., $p(s, t)$ is only nonzero for a single sucker at a time. In Ref. [29] the authors show that to ensure fastest translocation $p(s, t)$ should follow a wave traveling with the CPG Eq. (4), with a phase shift of $3\pi/2$. Under these optimal conditions, the center of mass translocates with a velocity:

$$v \propto \frac{\tilde{\omega}}{\sqrt{1 + \tilde{\omega}^4}} \quad (8)$$

where $\tilde{\omega} = \omega\zeta_0 N^2 / (4\pi^2\kappa)$. Note that this equation has a maximum at $\tilde{\omega} = 1$, which occurs when

$$N_{\max} = 2\pi\sqrt{\kappa/(\omega\zeta_0)}.$$

This maximum is achieved when the wavelength of the CPG, $\lambda = Nx_0$, matches a diffusion length scale $\lambda_D = \sqrt{D/\omega}$, where $D = kx_0^2/\zeta_0$ can be derived analytically as the continuous model reduces to an effective diffusive equation. While our finite crawler cannot be treated analytically as in Ref. [29], the diffusive length scale is still dimensionally correct, and we expect that qualitatively the crawler finds a maximum speed when $N_{\max} \sim \sqrt{\kappa/(\omega\zeta_0)}$. We do indeed observe that the crawler translocates with a speed that depends on the number of suckers, with a maximum at an intermediate location with about $N = 12$ suckers (Fig. 3). The location of the maximum varies with the parameters as prescribed by the scaling above (data not shown), suggesting that the physics of this optimal condition is robust to details of the model. In Appendix A we discuss and compare extensively our model and simulator to Ref. [29].

Note that living crawlers may evolve to match this constraint by, e.g., developing specific secretions that will modify their friction coefficient ζ_0 , or by modifying the frequency ω or wavelength λ of their CPG, or the stiffness of their muscles κ . In the following, we do not focus on the mechanical parameters that define the crawler but rather on how to learn the controls that enable it to move.

D. Learning architectures

In the RL paradigm, the learning problem is framed in terms of environment and agent. The agent, i.e., the learner, interacts continuously with the environment selecting at each time step an *action* (stick or not stick to the substrate) to which the environment responds with a numerical *reward*

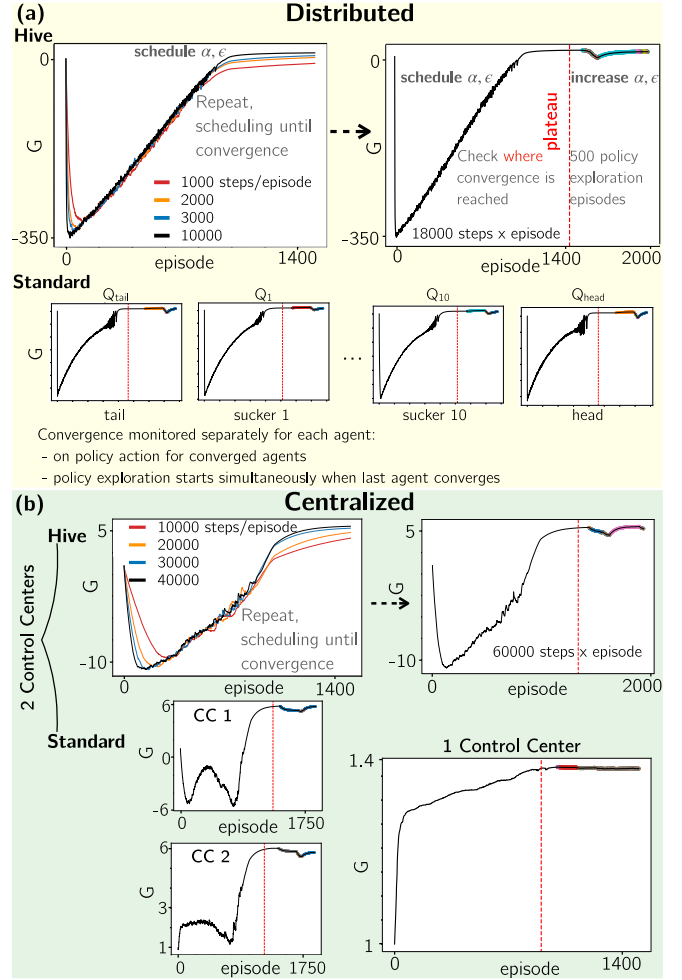


FIG. 2. Illustration of the training protocol on a 12 suckers crawler for various learning architectures. (a) Average return $G = 1/|\mathcal{S}| \sum_s \max_a(Q(s, a))$ as a function of learning episode for distributed learning architectures with the same Q matrix for all agents (hive, top) or one Q matrix per agent (standard, bottom). Top left: G initially increases and then slows down; the process is repeated increasing the number of steps per episode (color lines) until learning reaches a flat plateau (black). Right: At the plateau, identified automatically (red dashed line), ϵ and α are increased again within an exploration phase that stores a set of 250 policies (each color represents a different policy). (b) Same as (a), for centralized architectures with two control centers hive (top) and standard (bottom) and for a single control center (bottom, right). See Table I for details on the values of the parameters.

and a representation of its *state* [39]. When treating the one-dimensional crawler model as an RL problem, since the task is to achieve unidirectional and fast motion, the reward at each time step is given by the instantaneous center-of-mass velocity, $R_t = v_{c.m.}(t)$, when the velocity is positive and $R_t = -1$ if $v_{c.m.}(t) < 0$. After several trials, we found this choice to be effective to promote policies maximizing the crawling speed. In order to keep the problem as simple as possible and minimize the requirements on proprioception, we use binary states associated to the spring's tension: 1 = elongated and 0 = compressed. As previously mentioned, actions are also bi-

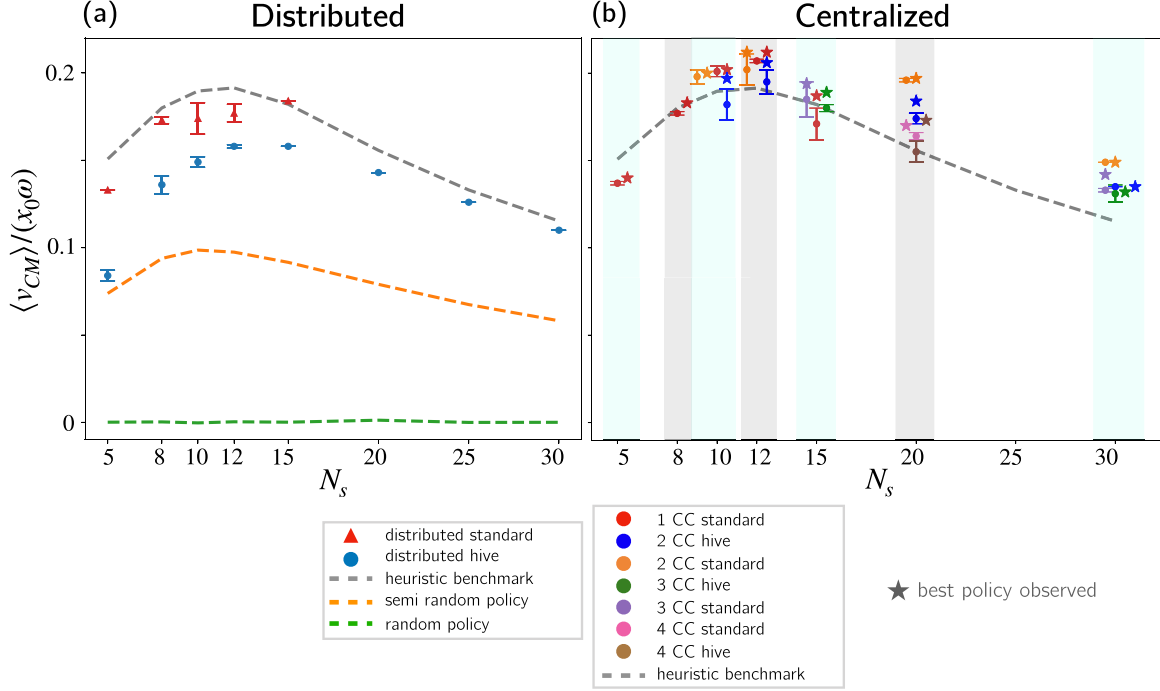


FIG. 3. Average velocity of the crawler center of mass versus number of suckers, for different learning architectures. Markers and errorbars correspond to average and standard deviation of performance over the 250 best learned policies for distributed architecture (left) and centralized architectures (right). Colored bands in panel (b) group results obtained for the same value of N_s (which have been horizontally scattered for clarity, so symbols do not overlap).

nary and related to each sucker being adhering or not: $\zeta_n = \infty$ or $\zeta_n = \zeta_0$, respectively. Note that the instantaneous reward R_t is a random variable, as the net speed of the center of mass is the result of the coupled spring dynamics which is only partially sensed and controlled by the crawler. We thus seek to maximize the cumulative translocation over many steps in the future, as usual in RL. To do this, we introduce the long-term reward G_t , defined as $G_t = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$, where R_{t+k+1} is the instantaneous reward $k+1$ steps in the future with respect to t . We use the standard formulation with discount rate $\gamma < 1$, which is a parameter discounting rewards further in the future ($\gamma = 0$ truncates the sum to its first term, i.e., $G_t = R_{t+1}$ is the instantaneous reward, whereas $\gamma = 1$ weighs each contribution equally for all times). We use $\gamma = 0.99$, corresponding to truncating the sum to an effective horizon of 100 steps. For continuing tasks, another possible formulation of the reward is that of *average reward* [39,40], which does not require discount and where the reward becomes also an estimate. Preliminary results, however, do not indicate any visible benefits of this formulation in our problem.

Learning seeks to find a *policy* of actions, $\pi(a|s)$, i.e., a rule that prescribes how to choose actions a at each given state s that maximizes the expected future reward. Optimal policies can be searched via several algorithms. We here use standard tabular Q -learning, where a *quality matrix*, or Q matrix, is learned with an iterative process. The Q matrix is the expected cumulative reward conditioned to being in state s and choosing action a , i.e., $Q(s, a) = E(G_t | S_t = s, A_t = a)$, while continuing with policy π . Q is estimated iteratively and each estimate is used to build a policy consisting in selecting at each state the action that maximizes the empirical return with

highest probability $(1 - \epsilon_t)$ and choosing a random action with a small probability $\epsilon_t/|\mathcal{A}|$. Thus for a given agent its Q matrix follows the update rule

$$\begin{aligned}
 Q_{t+1}(S, A) &= Q_t(S_t, A_t) + \alpha_t \left[R_{t+1} + \gamma \max_a Q_t(S_{t+1}, a) - Q_t(S_t, A_t) \right] \\
 \pi_{t+1}(a|s) &= \begin{cases} \arg \max_a Q_{t+1}(a, s) \text{ w.p. } 1 - \epsilon_t \\ \text{any } a \in \mathcal{A} \text{ w.p. } \epsilon_t/|\mathcal{A}| \end{cases} \quad (9)
 \end{aligned}$$

where R_t , S_t , and A_t are the reward, state, and action of the agent at integration step t . The parameters α_t (learning rate) and ϵ_t (exploration) vary during learning in a process known as *scheduling*, discussed in the next section. Therefore, the dimensionality of the Q matrix is given by $|\mathcal{S}| \times |\mathcal{A}|$, with \mathcal{S} and \mathcal{A} the ensembles of states and action, respectively. As illustrated in Fig. 1, agency can be defined by the learning architecture in several ways. Each architecture carries a specific definition of states and actions which dictates the dimensionality of the Q matrix, and hence the computational cost. In particular, states are binary words describing all possible combinations of elongation and compression of all $N_s - 1$ springs contained within an agent. Spring k is compressed when it is shorter than its rest length, i.e., $\bar{x}_{k+1} - \bar{x}_k < \bar{l}_k$, where \bar{x}_k and \bar{x}_{k+1} and \bar{l}_k are defined by Eqs. (5). Finally, we assume that each agent can act on temporal scales comparable to the CPG oscillation frequency, so that responses are instantaneous. To investigate the role of centralization versus distribution,

we explore two families of learning architectures, which we illustrate next.

1. Distributed control

Each sucker is considered an independent agent [Fig. 1(b)]. Each agent, excluding tail and head, has four states which are all the combinations of the spring states of the neighboring springs and two actions. Tail and head of the crawler, have only two states since for them one spring is missing. Therefore, each intermediate agent-sucker has a 4×2 -dimensional Q matrix, and tail and head a $2 \times 2Q$ matrix.

2. Centralized control

We introduce centralization by defining an agent as an assembly of several suckers and springs which we call control center (CC), as illustrated in Fig. 1(c). The number of CCs controlling the crawler depends on how many suckers are included in each CC and can range from several CCs down to a single CC, which corresponds to the highest degree of centralization. Centralized architectures control several suckers and springs simultaneously and hence have access to a richer representation of the dynamical state of the crawler and can realize many more policies than those realizable in the distributed setting. The state of the CC is given by the compression and elongation state of each spring located between the first and last sucker within the CC; while the action of the CC consists of specifying the binary adhesion–no-adhesion instruction for each of the suckers within the CC. Since the CC has access to all combinations of compression-elongation and adhesion–no-adhesion states and actions, the state-action space scales exponentially as $|\mathcal{S}| = 2^{n_s-1}$ and $|\mathcal{A}| = 2^{n_s}$, with n_s the number of suckers controlled by the given CC (for n_s suckers under a CC there are $n_s - 1$ connected springs). Note that to facilitate comparisons and have the same dimensionality of the Q matrix for each CC we chose to ignore the spring between two adjacent CCs. This is the only consistent choice for hive updates, as if the spring at the interface of two control centers was included in the state, the information would be interpreted differently by the two adjacent control centers. To make a fair comparison, we keep the same choice even when each control center learns its own policy (nonhive). We do not expect considerable differences because the control centers already perform as well as the single agent, to within the error bar.

3. The hive assumption

Whenever multiple agents control crawling, a simplifying assumption may be adopted imposing that all agents learn the same policy. This is the so-called *hive* assumption or population-based training [41,42], where a population of agents is used to find a shared solution. In all our multiagent learning schemes—including distributed and centralized with multiple CCs—we compare standard learning to hive learning, as illustrated in Fig. 1. In practice, each agent writes and reads the same Q matrix. The distributed hive Q matrix becomes a 8×2 tensor, where in addition to the four states associated to the internal sucker there are two extra elements of the matrix for the tail, (null,elongated) and (null, compressed),

and for the head, (elongated,null) and (compressed,null).¹ For the hive CCs there is no need to modify the Q matrix between hive and standard learning, since all CCs have Q matrices that are identical in shape, as discussed above. Clearly, the hive assumption speeds up training as it allows to parallelize the learning process. However, it also affects significantly the learned policies, as all agents are forced to learn the same policy, despite being placed at different locations relative to the propagating wave. In other words, the hive update imposes a consensus mechanism and all agent must “agree” on the same policy. In fact in the hive policy we find that different agents will learn distinct policies, and the hive update reduces the number of policies that are learned, as discussed under Results.

E. Learning protocol

Reinforcement learning theory and algorithms are extensively covered in specialized books and articles (e.g., Ref. [39]). As briefly summarized before, in tabular Q -learning during training we iteratively update the Q matrix using Eq. (9). To keep a certain degree of exploration during training we adopt the standard procedure of following an ϵ -greedy policy, that is, with a finite probability ϵ a random action is chosen instead of the one prescribed by the current policy (i.e., the action maximizing the current estimate of the value). During training, both the exploration ϵ and learning rate α in Eq. (9) are scheduled to decrease linearly, to ensure convergence to a final estimate of Q . Note that convergence to the optimal Q matrix is only guaranteed asymptotically, and hence different trainings will lead to different suboptimal policies, which we explore systematically in the following. In this work the scheduling process starts at a value of $\epsilon_{\max} = 0.9$ and $\alpha_{\max} = 0.1$ and is scheduled down to $\epsilon_{\min} = 0.001$ and $\alpha_{\min} = 0.001$. Because here we seek fast unidirectional motion, the task is *continuing* and there is no termination. With the term *episode* we intend a series of updates characterized by the same value of the parameters α and ϵ . Note that this definition is different from the usual definition used in episodic tasks, where the end of the episode corresponds to the agent reaching a specific terminal state.

Note that although the equations of motion are deterministic, our learning problem is stochastic. There are two sources of stochasticity: (i) the coarse-grained state space and (ii) the definition of agents. Both (i) and (ii) imply that the reward signal is a stochastic variable. Indeed, the speed of the crawler depends on the full dynamical state of the system, defined by the (continuous) speed, position, and adhesion of all suckers and on the relaxed length of all springs. However, as per (i), our crawler does not observe the full dynamical state of the system but rather its coarse projection onto discrete binary states of compression versus elongation of all the springs. Thus given the actions of all suckers, the center of mass is not fully determined yielding a stochastic reward $R(t)$. Moreover, in all multiagent architectures each agent is unaware of other

¹Note that in this case the problem is effectively hive only for the internal suckers, while tail and head remain genuinely independent (since their states are unique) as in standard multiagent RL.

TABLE I. Approximate value of the exploration parameters used for the various architectures trained, including the number of integration steps per episode needed to obtain convergence (consider also Fig. 2). The precise value is case specific and was tuned manually by considering the trade-off between satisfactory levels of exploration and convergence.

Architecture	ϵ	α	Minimum training steps \times episode
Distributed (standard)	$\lesssim 10^{-2}$	$\sim 10^{-3}$	$\sim 50\text{k steps}^a$
Distributed (hive)	$\sim 10^{-2}$	$\sim 10^{-3}$	$\sim 18\text{k step}$
Centralized (small ^b CC)	$\sim 10^{-2}$	$\sim 10^{-2}$	$\lesssim 100\text{k steps}$, about half for hive
Centralized (large ^c CC)	$\sim 10^{-2}$	$\gtrsim 10^{-1}$	from $\gtrsim 100\text{k}$ to $\sim 1000\text{k}$ steps, about half for hive
Fully centralized (single CC)	$\sim 10^{-1}$	robust up to $\gtrsim 10^{-1}$	from $\gtrsim 100\text{k}$ to $\sim 1000\text{k}$ steps

^aDifficult to converge and dependent on the number of suckers (agents).

^bFive or six suckers per control center.

^cTen or 15 suckers per control center.

agents' actions, as per (ii), and hence once again the actions of any individual agent do not individually determine the speed of the center of mass (i.e., the reward). This latter source of noise in the reward is bypassed by the fully centralized architecture with a single control center (1CC).

To monitor progress during learning we consider the average value G , with $G = 1/|S| \sum_s \max_a(Q(s, a))$. Clearly, when G increases, the Q matrix is being updated and learning is in progress, whereas when G reaches a plateau, the learning process has come to an end. As illustrated by Fig. 2, we set an automatic protocol consisting in repeating the learning procedure until a plateau in G is detected: When the plateau is not encountered after 1500 episodes, we repeat the process changing the scheduling by increasing the number of steps per episode. When several Q matrices are learned simultaneously (for instance, in standard distributed control) the convergence is monitored separately for each agent. Due to the stochastic nature of the problem, we noticed that repeating the learning procedure yields different policies. Therefore, we design a protocol to account for the stochasticity of this learning process. To efficiently store several suboptimal policies, without repeating the entire training procedure, we introduced the following strategy: When convergence is detected, the learning parameters are raised to a small and constant value and 500 extra episodes are run letting Q evolve accordingly. At each episode the corresponding policy is stored. As reported in Table I, the value of α and ϵ for this "policy exploration" phase are set differently according to the learning architecture considered with the purpose of exploring (sub) optimal policies without disrupting significantly the plateau that was reached by G . The value of α and ϵ that is needed to remain in the plateau reflects how robust the optimal policy is to perturbations, and appears correlated to the degree of stochasticity previously discussed. Centralized architectures appear to need stronger exploration of the plateau to learn distinct suboptimal policies. We find a fairly stable plateau region in the exploration phase of the fully centralized ICC architecture even for large values of α and ϵ . This is due to the high degeneracy of this architecture, where many different policies perform similarly, as discussed later. Finally, we find that the hive update introduces some degree of stabilization of the learning process since we observed the need of less integration steps (as expected) and a smaller sensitivity to scheduling and exploration parameters. The source code for all different architectures is available in Ref. [43]. Next we

discuss results obtained by repeating the protocol described above for several centralizes and distributed architectures and a number of suckers that ranges from $N_s = 5$ to $N_s = 30$. We reserve a particularly in-depth analysis to the case of $N_s = 12$.

III. RESULTS

Using the learning protocol described above, we store 500 optimized policies for each architecture and quantify performance as the average speed of the center of mass of the crawler for the best 250 policies. We obtained velocities by playing each policy over a trajectory of 20 000 steps or about 300 full cycles of the CPG traveling wave across the full crawler. In computing the average we accounted for the multiplicity of each policy, so that policies which are found several times will be more represented as well as quasidegenerate policies realizing similar crawling velocities.

We find that performance for all architectures depends on the number of suckers N_s (Fig. 3), qualitatively consistent with the infinite continuous crawler case previously discussed. With our choice of parameter, performance peaks for $N_s \approx 12$ and then exhibits a long-range decay beyond the peak (as discussed under Methods, the number of suckers is identical to the spatial periodicity of the CPG wave, N). To investigate the limits and properties of multiagency and centralization in cooperative or coordinated actions we will focus largely on the peak. The value of the parameters we chose correspond to the peak occurring near $N_s = 12$, which is computationally treatable for all learning architectures from the most distributed to the most centralized. Whether and how natural crawlers orchestrate the central pattern generation and the suckers to operate near the peak is a fascinating question for further research.

We find that the tail and head of our finite crawler play a key role, thus distinguishing our problem from the infinite crawler for which an asymptotic solution is available. However, the analytical optimal policy is a useful heuristic for the internal suckers as discussed below. We thus construct a reference heuristic policy prescribing that (i) internal suckers follow the analytical optimal policy Eq. (7) from Ref. [29] (see Supplementary Video 1 [44]) and (ii) tail and head anchor in response to negative forces, i.e., the head adheres when it is dragged backward by the elongated spring to its left and the tail adheres when it is pushed backward by the compressed spring to its right. It is important to observe that to

implement the heuristic strategy, the simulator prescribes adhesion according to a phase shift, thus measuring time and space. In contrast, because we use binary compressed and elongation spring states, our agents (the suckers, or the CCs) have no access to the CPG’s wave phase. Thus, our simple crawlers cannot implement the heuristic policy, due to their limited states and controls. Intriguingly, our poor state representation can be compensated by a more centralized control which recovers a smooth traveling wave by controlling multiple suckers at a time, as we will see in the following.

A. Distributed control

Let us first focus on distributed forms of control, where each sucker is an agent [Fig. 3(a)]. When suckers play random actions, the center of mass does not move on average, as expected (random policy, green dashed line). If tail and head adhere in response to negative forces, even if the internal suckers perform a random action, then the center of mass does move, showing the importance of the crawler’s tail and head (semirandom policy, orange dashed line). If tail and head adhere in response to negative forces and the internal suckers adhere according to the analytical results described above, performance further increases (heuristics, gray dashed line). Although the analytically optimal policy of the internal suckers is obtained in the infinite continuous crawler limit, the clear improvement of the policy relative to the semirandom policy suggests that away from the boundary, the infinite crawler may be a useful approximation of our finite discrete crawler. This is consistent with the fact that our crawler translocates with a speed that varies with the number of suckers, qualitatively similar to the infinite crawler in Ref. [29]. We report this heuristics as a qualitative reference to guide the eye. Note that experiments on leech [45] support the role of head and tail postulated here. When we let agents act according to a policy that has been learned through trial and error, performance is intermediate between the semi random policy and the heuristics [Fig. 3(a), blue circles and red triangles]. We find that standard learning is harder than hive learning, but it achieves better performance and nearly matches the heuristics described above (e.g., for 15 suckers). As a reminder, in hive learning all suckers learn the same policy, whereas in standard learning each sucker learns its own individual policy.

Almost invariably trainings with hive architectures learn the same policy π_H , regardless of N_s . According to π_H , internal suckers adhere in response to the $\leftarrow | \leftarrow$ state while the tail and head—which are connected to a single spring—adhere in response to the states $\times | \leftarrow$ and $\leftarrow | \times$, respectively. The interpretation of the hive policy π_H is simple: Each sucker adheres in response to negative forces, i.e., when all adjacent springs push it backward. For $N_s \geq 15$, all our trainings converge to the same policy π_H , with no suboptimal policies ever learned [hence the absence of the errorbar in Fig. 3(a)]. The same behavior was found when training distributed hive policies for even larger crawlers with $N_s > 30$ (results not shown) further confirming that π_H is the best consensus policy. Departures from this policy emerge only for small N_s , where the specific dynamics of few suckers might affect the global consensus. For $N_s < 15$ trainings converge to a small number of policies; the best learned policy is π_H for $N_s = 5$ and 12,

and it is a slightly different policy $\hat{\pi}_H$ only for $N_s = 8$ and 10. π_H and $\hat{\pi}_H$ are almost identical, except internal suckers also adhere in response to the $\rightarrow | \rightarrow$ state, and the tail never adheres. For small number of suckers, $\hat{\pi}_H$ is marginally faster than π_H (e.g., for 10 suckers speed is ≈ 0.017 versus ≈ 0.015). Thus collective hive learning is highly transferrable, in that the hive policy can be applied to crawlers with any number of suckers with no need for retraining, consistent with previous results obtained with genetic algorithms [28]. Adopting the hive policy results in a clear traveling wave of adhesion, with only one sucker adhering at a time (see animation of the hive policy in Supplementary Video 2 [44]). Note that, as already remarked above, the individual suckers do not have access to their phase shift, and as a result, this architecture cannot learn a policy that implements our heuristic benchmark. Indeed, this heuristics requires a finer proprioceptive signal, e.g., inferring the phase from the precise intensity of the restoring forces, rather than its sign.

Intriguingly, the standard learning, where each sucker acts according to its own policy, realizes a more complex dynamics with multiple suckers adhering simultaneously and no visible traveling wave of adhesion. In doing so it achieves a more complex group behavior and better performance. Note that, although standard learning achieves better performances, we find that training is more computationally expensive and in our hands was limited to $N_s = 15$: Beyond this limit the training did not converge. This is likely due to the intrinsic level of stochasticity: All agents learn simultaneously, but rewards are coupled, as the motion of the crawler depends on the full dynamical state of the system. Thus good actions may be hard to learn as they may not be individually good enough to yield a positive reward. Moreover, in all cases where standard learning was successful, we noticed an extreme sensitivity to the scheduling parameters, with convergence only using about three times the number of steps per episode relative to the hive solution. In contrast, the distributed hive policies were easily learned.

B. Centralized control

Centralized architectures achieve even better performances, at the cost of increased computational burden (Fig. 3, right). Indeed, the number of entries in each Q matrix grows exponentially with the number of suckers controlled by the corresponding CC. Given that $n_s \leq N_s$ is the number of suckers controlled by each control center, there are N_s/n_s CCS whose quality matrix has dimensions $|\mathcal{S}| \times |\mathcal{A}| = 2^{n_s-1} \times 2^{n_s}$. Consistently, learning requires considerably longer training (Table I) relative to distributed architectures, where training requires to optimize $N_s Q$ matrices whose dimensions are independent from the crawler size $|\mathcal{S}| \times |\mathcal{A}| = 4 \times 2$.

At the same time, a higher degree of centralization achieves better performance. We find that a single control center dictating the actions of all N_s suckers performs best for $N_s = 10$ and 12, suggesting that long-range correlations across the whole crawler are useful if well orchestrated. However, training of the single CC is extremely complex; it failed beyond $N_s = 15$, because of the sheer size of the quality matrix, and for $N_s = 15$ it under performs slightly relative to the 3 CCs. Interestingly, multiple CCs achieve almost the same performance as a single control center and can be trained even with a much larger

number of suckers. For example, the reduction in complexity for $N_s = 12$ is quite impressive: The Q matrix for a single CC has approximately 10^7 entries, whereas it shrinks to about 2000 entries with two control centers. As for distributed learning, hive learning where each CC is forced to learn the same policy is faster than standard learning but performs slightly more poorly, although the margin is not as impressive as for the distributed control. Next, let us illustrate in more detail the differences between these architectures, focusing on the 12 sucker crawler. For illustrative purposes, the optimal policies for the various learning architectures, both distributed and centralized, are shown for a 12 suckers crawler in Supplementary Video 3 [44].

C. Robustness

We showed in the previous sections that moving from pure multiagency to some form of centralization improves performance. However, the benefits are not always sizable, e.g., 1CC improves $\sim 17\%$ over distributed standard and $\sim 30\%$ with respect to distributed hive, when considering 12 suckers. At the same time, computational cost grows dramatically and optimal policies become hard to learn. Are there any other benefits to centralized control?

We next show that centralized architectures are considerably more robust to failure of individual suckers. In Fig. 4, we focus on the best policy obtained for each architecture for the 12 suckers crawler and test its robustness to failure. To this end, we extract one or more random suckers at each integration step and let them perform a random action (adhere or not with the same probability). All other suckers follow the optimal policy, and no extra learning is allowed. Centralized architectures are more robust to failure: Their performance drop more slowly with the number of failing suckers Fig. 4(a), with a single random failure eroding less than 10% performance versus about 20% for distributed architectures, Fig. 4(b). We then asked how robust these policies are to failure of the most crucial sucker. When we make a single sucker fail consistently over time, we find that the head plays the most prominent role in all of our architectures, as its failure contributes the largest drop in performance Fig. 4(c). In fact, the head is always the dominant sucker when the CPG wavelength, λ , is equal or larger than the crawler length and not otherwise (Appendix C). Failure of the head is particularly acute for distributed architectures (60% drop in performance for fully distributed architectures) against much less dramatic drops in centralized architectures (the most robust being the 2CC with a $<20\%$ drop).

To further quantify the distinct policies we measure the frequency that each sucker experiences the four local compression and elongation states of its adjacent springs, as the policy is played over time. The four local states are represented with the symbols $\rightarrow | \rightarrow$; $\leftarrow | \rightarrow$; $\rightarrow | \leftarrow$ and $\leftarrow | \leftarrow$. The local states are visited with frequencies that vary across the different architectures [Fig. 4(e) top row], but the pattern is relatively preserved. This suggests that the underlying physics constrains the spring dynamics considerably and may partly overshadow differences in the control architecture. In contrast, the average action played by each sucker in its local compression or elongation state varies dramatically across

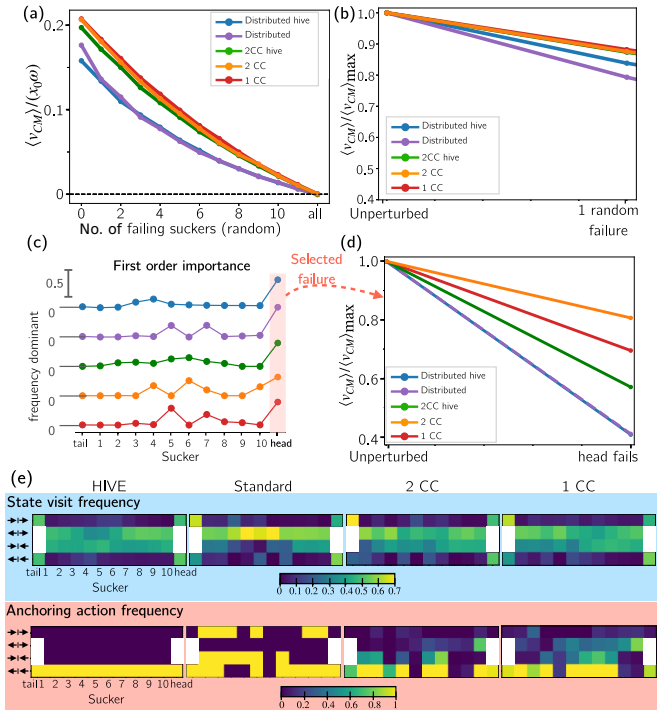


FIG. 4. Robustness analysis of the best policy for 12 suckers crawlers with various learning architectures. (a) Performance (velocity of the crawler) versus number of failing suckers. (b) Performance of the optimal policy with 1 failing sucker relative to optimal performance with no failing suckers. (c) Frequency each sucker is the dominant contributor, measured by the largest drop in performance when the sucker fails: The head is the most important sucker. (d) Performance drop due to failure of the head. (e) Statistical behavior of a 12 sucker crawler for different distributed and centralized architectures. Top row: For each sucker (x axis), and each compression state of its adjacent springs (y axis), the color represents the frequency that state was visited while performing the optimal policy. Bottom row: Fraction of time a sucker adheres, color coded from purple (always adhere) to yellow (never adhere) for each of the 12 suckers (x axis) and each of the compression states of their adjacent springs (y axis). The blank entries for head and tail represent that these suckers only have access to one spring, hence two compression states (rather than four).

different control architectures [Fig. 4(e) bottom row]. In distributed architectures (left panels), suckers play an action that depends only on their local compression and elongation state, and hence colors are binary (adhere or not adhere). In contrast, in centralized architecture, a sucker acts in response to the global (albeit binary) state of compression or elongation and not only in response to its local state. Intermediate colors in the right panels signify that the local compression or elongation state is not enough to determine whether the sucker will adhere or not. Because centralized policies are dictated by global compression or elongation states, they may be more tolerant to failures of individual suckers.

D. Redundance

To further elaborate on why centralized policies are more robust, we characterize redundancy. Within a given learning

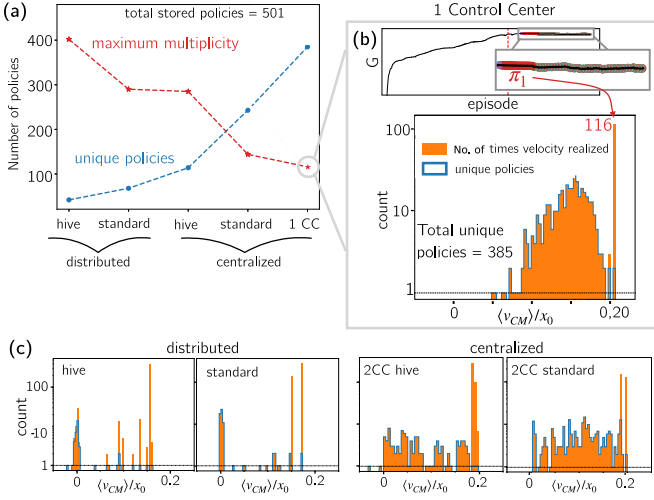


FIG. 5. Distribution of speed for the policies learned by the 12 suckers crawler. (a) Number of distinct policies learned by the 12-sucker crawler with different architectures (blue); number of times the most popular policy was learned (red); (b) detail of policies learned by the fully centralized agent; top: Plateau in G , showing the optimal policy, labeled as π_1 , is selected at several iterations of the exploration phase of the learning protocol (same as bottom right panel in Fig. 2(b), color change in the plateau marks policy change); bottom: histogram of the velocities across all 250 policies. Orange: Number of times a crawler velocity is realized by a learned policy. Blue: Number of unique policies associated to the velocity (within the bin). Bins where the blue histogram is lower than the orange histogram show identical policies which are selected multiple times in the learning protocol. (c) Same as (b), for the remaining learning architectures.

architecture, are many suboptimal policies learned and are they considerably worse than the best policy? Centralized learning converges several times to the optimal policy but also to several fair suboptimal policies. For instance, the 1CC architecture converges 116 times of 500 to the optimal policy π_1 but also converges to another 384 suboptimal policies whose performance is only slightly poorer than the optimal performance (Fig. 5 top row). Conversely, distributed learning converges even more often to the best policy (large peaks in the histograms in Fig. 5 bottom row, corresponding to the best policy) and to fewer considerably worse suboptimal policies (group of bins near velocity zero). These results suggest that the higher degree of complexity of centralized control offers numerous ways to achieve fast crawling strategies making it harder for the algorithm to spot *the* best one but at the same time more likely to learn *a* good policy.

E. Orchestrating a traveling wave

The results described so far suggest that centralized control achieves better performance at the cost of exploring a higher-dimensional space of possible states and actions. Why is centralization beneficial? Our crawler is composed of individual suckers that only have the ability to sense a binary compressed-elongated state of their adjacent springs. This makes for a system of poor proprioceptors that do not have access to, e.g., their spatial location nor to a clock. We find

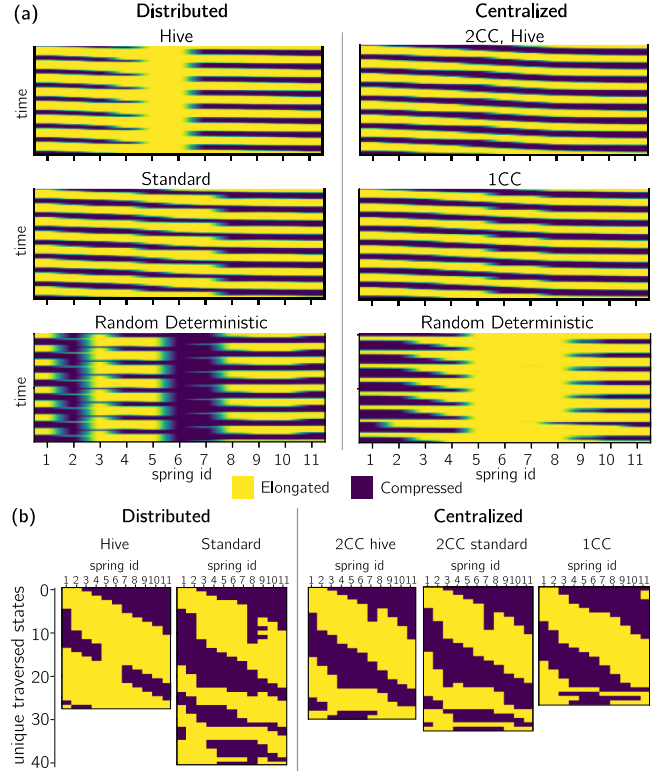


FIG. 6. Spring states observed when playing the best policy for the different architectures considered in the 12 suckers crawler, compared to a random deterministic policy. (a) Kymographs of states experienced by each spring (x axis) during integration: Each row corresponds to a different time; time increases on the y axis from top to bottom. (b) Unique binary states of compression-elongation of the 11 springs experienced as the crawler adopts specific policies. States appear in an arbitrary order that does not account for the time spent in each state nor the temporal order these states are encountered (e.g., state 0 is not visited more often than state 1, nor it is visited right before state 1).

that by orchestrating the action of many poor proprioceptors, centralized control is able to ride the traveling wave generated by the CPG, despite the lack of a direct measure of space or time. To characterize this intuition, we visualize kymographs of the state of compression of the springs generated by the dynamics as a crawling policy is enacted [Fig. 6(a)]. Note that the state of compression of the springs is not deterministically controlled by the adhesion pattern of the suckers. We find that all optimal policies generate periodic states of compression as seen by repeating patterns representing traveling waves of compression with speed ≈ 0.19108 , perfectly matching the underlying central pattern generator phase velocity, $v_p/x_0 = \omega/k \approx 0.19099$, with $\omega = 0.1$, $k = 2\pi/N_s$, and $N_s = 12$. This is in contrast to random deterministic policies, which generate random states of compression of the springs and cannot reproduce the CPG wave [bottom panel of Fig. 6(a)]. Interestingly, for distributed architectures in the hive case the bands abruptly end near the center of the crawler, signifying that when the compression wave hits springs 5 and 6 it stops: A clump of springs near the center of the crawler remains constantly elongated. In

contrast, for centralized architectures the wave never stops but rather travels along the crawler all the way to the end, so that the CPG movement is effectively transferred to the crawling. Thus by orchestrating adhesion of multiple suckers at a time, centralized architectures realize a smoother traveling wave of compression. Consistent with the observation that centralized policies generate a smooth traveling wave of compression, the dynamics does not generate all possible combinations of compression and elongation of all springs, which would be 2^{N_s-1} . Instead, only a limited number of compression states is visited. More precisely, for the optimal fully centralized policy (1CC), the 12 springs experience only 26 states of compression, about 1% of the available 2^{N_s-1} combinations of binary states associated to the $N_s - 1$ springs. When the crawler enacts any other policy, the coupled dynamics generates more states, indicative of stumbling blocks in wave propagation. This is despite the fact that centralized policies respond to the global state of compression of all springs simultaneously and are thus much more complex. Thus, somewhat counter intuitively, more complex policies with more degrees of freedom generate fewer states than simpler policies with fewer degrees of freedom.

IV. DISCUSSION AND PERSPECTIVES

Here we presented a toy model of a one-dimensional crawler composed of blocks connected by springs. Each block is a rudimentary actuator representing a sucker making a single binary decision at each time: adhere versus not-adhere. The suckers are also rudimentary proprioceptors, as they measure the binary state of compression versus elongation of their adjacent springs. Positional information and speed of the individual suckers is not measured nor inferred. We assume that a central pattern generator, or CPG, periodically modifies the rest length of the springs, according to a traveling wave. The CPG is dictated by an internal clock, independent from proprioception. This may be a useful model for the octopus, where proprioceptive feedback has been proposed to be implemented with an embodied mechanism rather than a neural process based on somatotopy [46]. Using this toy model we have demonstrated that this group of rudimentary proprioceptors or actuators can learn to interact with the substrate and exploit the wave generated by the CPG to achieve a net unidirectional motion of the crawler. Particularly, we focused on the pros and cons of distributed versus centralized forms of learning and control.

Using tabular Q -learning, we found that locomotion can be learned in a pure multiagent setting, where each sucker is considered an independent agent and with no explicit communication among agents, consistent with previous literature [28]. Although the agents do not communicate explicitly, the result of their actions is indirectly coupled via the reward. Indeed, the reward is the net translocation which results from the dynamics of the system of oscillators, hence to maximize reward agents effectively learn how to act cooperatively. A purely distributed form of learning and control is cost effective in that it requires no communication and thus scales well with the number of agents and the degrees of freedom in the system.

We find that introducing a certain degree of centralization, in the form of control centers encompassing several suckers, provides two benefits. On the one hand, the increased complexity of centralized policies can be effectively exploited to couple the actions of many suckers at a time. These long-range correlations enable centralized policies to more effectively ride the wave of compression provided by the CPG and ultimately yield faster crawling. On the other hand, as multiple suckers are activated at a time, failure of an individual sucker results in less dramatic loss of performance, making centralized policies more robust to failure.

Centralization of course comes at the cost of increased computational cost and requires two-way communication of proprioceptive input from the suckers to the control center and of motor control from the control center to the suckers. We find that intermediate architectures where multiple control centers orchestrate the actions of a subset of suckers achieve the best of both worlds, providing virtually identical benefits of full centralization but limiting the number of degrees of freedom to a manageable size. Ultimately, partial or full centralization overcome the limits of binary proprioception and activation by realizing smooth and continuous dynamics, despite binary input and output. Further work is needed to establish the optimal degree of centralization, emerging from the trade-offs between robustness, speed and computational cost, which will be relevant to understand the control architecture of biological crawlers and to design efficient robotic crawlers (see, e.g., Refs. [47–50]).

Although the physics underlying crawling is clearly deterministic, the learning task is stochastic due to partial observability. Indeed, the deterministic dynamics of the crawler is dictated by the instantaneous positions and speeds of all suckers and rest length of all springs, but these are unknown to the agent. On top of this, in all but the fully centralized architecture, the individual agents are unaware of other agent's actions and states of compression, which yields another source of stochasticity. In the context of partial observability, stochastic policies may provide better performance (e.g., adopting policy gradient methods) [39,51]. More refined forms of Q -learning may also be beneficial here, for example double Q -learning which was reported to be more robust when dealing with stochastic Markov decision processes [52] or other modern RL techniques, for instance based on deep-RL which couples neural networks to Q -learning [39,41,53]. Some of these techniques may prove more effective than our simple approach based on tabular Q -learning. However, they remain outside the scope of the present paper where we focused on disentangling the benefits and limits of distributed versus centralized learning architectures. For example, a complex neural network naturally couples inputs and outputs and may thus not be the natural choice to parse the role of centralization versus distribution.

In contrast, previous biologically inspired work with simple models of crawling (or templates [22]) focused on either a fully centralized [26] or fully distributed (hive) control [28]. In Ref. [26], a biologically plausible model of crawling is proposed which also uses rudimentary forms of proprioception and control: Input is limited to the identity of the most compressed body segment and actions correspond the designation of the sole active neuron in the neural network

controlling muscular contraction, while adhesion is not controlled. The authors show with a fully centralized Q -learning architecture that the central pattern generator can be learned rather than imposed as we do here. The benefits of distributed versus centralized control for animal locomotion have been studied in several other contexts and algorithms and generally pertain added stability in uneven terrain, adaptation, and versatility; see, e.g., Refs. [18–21].

In Ref. [28], swimming of an agent is achieved through training of a genetic algorithm with a fully distributed architecture, where individual actuators are all forced to learn the same policy—which we have called hive update here. Again, the actuators are not suckers but rather beads which control the restoring force between their neighbors. The hive update was found to generalize well to a different number of beads, with no need for retraining. This is consistent with our finding, where we also observe that the hive update almost invariably learns the same policy, independent on the number of suckers composing the crawler. However, while we focus on binary sensing and control, Ref. [28] considers actuators with full proprioception, including the position and speeds of their neighbors. An advantage of considering binary proprioception is its simple implementation and interpretation. For example, our hive update can be understood in simple terms: Suckers learn to only adhere when they are pushed and pulled backward from both sides.

Here we have focused on the learning process that enables suckers to achieve crawling by combining mechanical contractions and interaction with the substrate. In our model, thrust is achieved by letting suckers control patterns of adhesion to the solid substrate, whereas contraction results from the stereotypic wave of compression and the dynamics. In Refs. [26,28] switch the controls: The actuators learn patterns of contraction, and interaction with their solid and fluid environment respectively is embedded in the dynamics. Finally, in bio-inspired models of sea star locomotion [6,7], thrust is achieved by fully prescribing—rather than learning—actions. Synchronization here emerges through the dynamical interaction of the substrate with the tube feet, which are physically connected to the body, with a mechanism reminiscent of the Kuramoto model of synchronization [54,55].

Here we parsed the pros and cons of distribution versus centralization for learning the internal coordination among body parts, aimed at translocation. The benefits of centralization for speed and robustness point to a potential origin of selective pressure towards more centralized forms of nervous systems. To what extent early nervous systems allowing the Cambrian explosion may have emerged from the need for internal coordination versus from the need of processing complex sensory information is currently debated; see, e.g., Refs. [56–58]. Our results are dictated purely by the need to internally coordinate body parts to generate thrust, but the role of centralization may become even more prominent when sensation and locomotion are coupled to aim toward a specific target. Indeed, this requires that each unit senses not only the internal state of compression of the organism but also sensory signals from the environment, from, e.g., vision, audition, taste, and olfaction. Our approach may be further extended to sensory navigation by enriching the state definition with

sensory cues from the environment and rewarding not only speed but also the successful reaching of a sensory target.

Our model is inspired by octopuses and other cephalopods, organisms with a distributed nervous system whose proprioception remains elusive [11,59]. Control centers coordinating adhesion of a few suckers with rudimentary proprioception achieves fast translocation without incurring into the computational cost of a fully centralized brain. Our approach may be further developed to quantify performance, robustness, and computational cost as a function of the number of control centers to select an “optimal” number of control centers for crawling, aimed at both fast locomotion as well as sensory navigation toward a target. Control centers may be adapted to model how ganglia distributed along the arms of an octopus control locomotion. To what extent the robustness, performance, and computational cost of crawling shape the connectivity of the ganglia distributed along the octopus arms’ and the organization of their taste-by-touch sensory system [8,60,61] is an exciting avenue for further research.

ACKNOWLEDGMENTS

This research was supported by grants to Agnese Seminara from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 101002724 RIDING) and the National Institutes of Health (NIH) under Award No. R01DC018789. This work represents only the view of the authors: the European Research Council Executive Agency and the other funding agencies are not responsible for any use that may be made of the information it contains.

DATA AVAILABILITY

The data that support the findings of this article are openly available [62], embargo periods may apply.

APPENDIX A: VALIDATION

1. Periodic crawler

As discussed in Ref. [29], in the continuous limit the equilibrium spring position (central pattern generator) is $l(s, t) = x_0 + a \sin(\omega t - ks)$, with $s = nx_0$, $k_c = 2\pi/\lambda$, and $\lambda = Nx_0$ with N the periodicity. It can be proven [29] that without any perturbation, i.e., no adhesion, that is uniform and constant friction across all the crawler, ζ_0 , the continuous equation of motion for a sucker at location s is

$$u^0(s, t) = C \cos(\omega t - k_c s - \theta) \quad (\text{A1})$$

with $\theta = \tan^{-1} \tilde{\omega}$, $\tilde{\omega} = \zeta_0 \omega / (k^2 \kappa)$, with κ the springs elastic constant, and $C = a/N \cos \theta$. We can approach the system described in Ref. [29], where they assume an infinite crawler, by considering a discrete crawler with periodic boundary conditions and check how the unperturbed solution (no sucker adhering) compares to the analytical prediction of Eq. (A1). As illustrated in Fig. 7, periodic boundary conditions are implemented by considering a virtual replica of the tail connected to the last right-hand sucker of the crawler. The comparison is made in Fig. 8 where we track the head position



FIG. 7. Sketch of the periodic crawler model.

of the simulated periodic crawler with no adhesion (so no net movement of the center of mass, $v_{c.m.} = 0$). The comparison is extremely good with what predicted from theory.

When the crawler is perturbed by the adhesion pulse of Eq. (7), an asymptotic solution is available from Ref. [29] (Fig. 9). The implementation of the optimal crawling strategy [Eq. (7)] in the periodic crawler discussed here has the same qualitative behavior predicted by the theory, Eq. (8), with a different location and magnitude of the maximum. These discrepancies likely arise due to our approximation of the infinite crawler with the periodic crawler.

However, as also shown under Results for the finite crawler, we have a fairly good qualitative comparison with the theoretical behavior, namely the existence of a maximum in N ($= N_s$, the number of suckers for our parametrization) and a decay for $N \rightarrow \infty$. The position of the maximum is not exactly the one predicted by the analytical theory but similar (with $\omega = 0.1$, $N_{\max}^{\text{finite}} \sim 12 < \tilde{N}_{\max} = (2\pi)/\sqrt{\omega} = 20$). Interestingly, as illustrated by Fig. 3, the functional shape of the dependency with N_s is preserved even when crawling policies allow for the adhesion of multiple suckers at a time which was excluded from the analytical treatment of Ref. [29].

Damped limit

For validation purposes, we can drop the overdamped assumption and solve the full damped equation for the finite crawler Eq. (1). The overdamped regime is already approached when inertia balances the restoring force of the




FIG. 8. For validation purposes, we here track the movement of the head of a 10 suckers crawler in overdamped dynamics and with periodic boundary conditions (infinite crawler ideal model) and compare it to the analytical solution of Eq. (A1). The position are reported in the reference frame of the time-averaged position and in normalized units $\bar{x} = x/x_0$. The amplitude of the oscillation is here $a = x_0/3$.

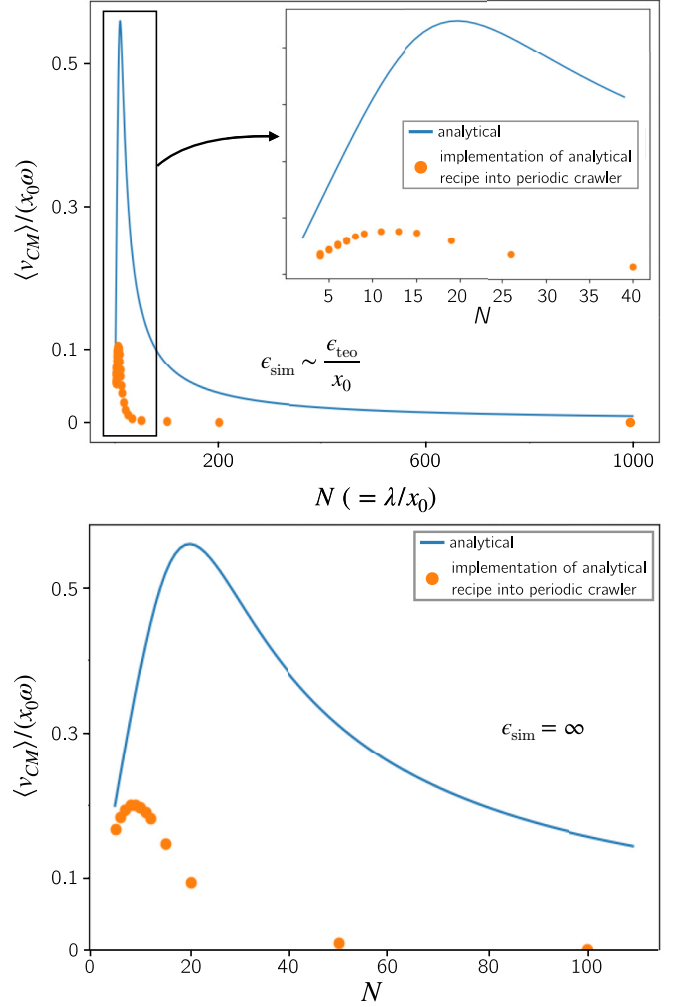


FIG. 9. Comparison between Eq. 4.9 of Ref. [29] (blue lines), and the implementation of Eq. (7) in the simulator (orange dots) for the periodic crawler model. (a) A more honest comparison with the perturbation regime used in Ref. [29] which accounts for the small adhesion limit, using $\epsilon = 1$ in the simulator. (b) With infinite adhesion in the simulator as the one used generally in this work.

spring, i.e., $m\omega^2/\kappa = 1$. Preliminary results also suggest that policies learned in the damped regime are not significantly different from the overdamped ones. Therefore, we expect that the conclusions reached in this work should still hold in the damped scenario. However, further verifications would be needed to assert this claim unequivocally.

APPENDIX B: EFFECT OF CPG PERIODICITY OVER SUCKER IMPORTANCE

In this work in general we assume for simplicity that a single wavelength is contained in the crawler, $\lambda = L$ (that is, $N = N_s$). As discussed in the main text, we found that when playing a crawler policy the head is on average the most impactful sucker upon failures, whatever the considered learning architecture. To establish if this is a general property and how it correlates to the physics of the system, we analyzed for the 12 suckers system how this property is affected when

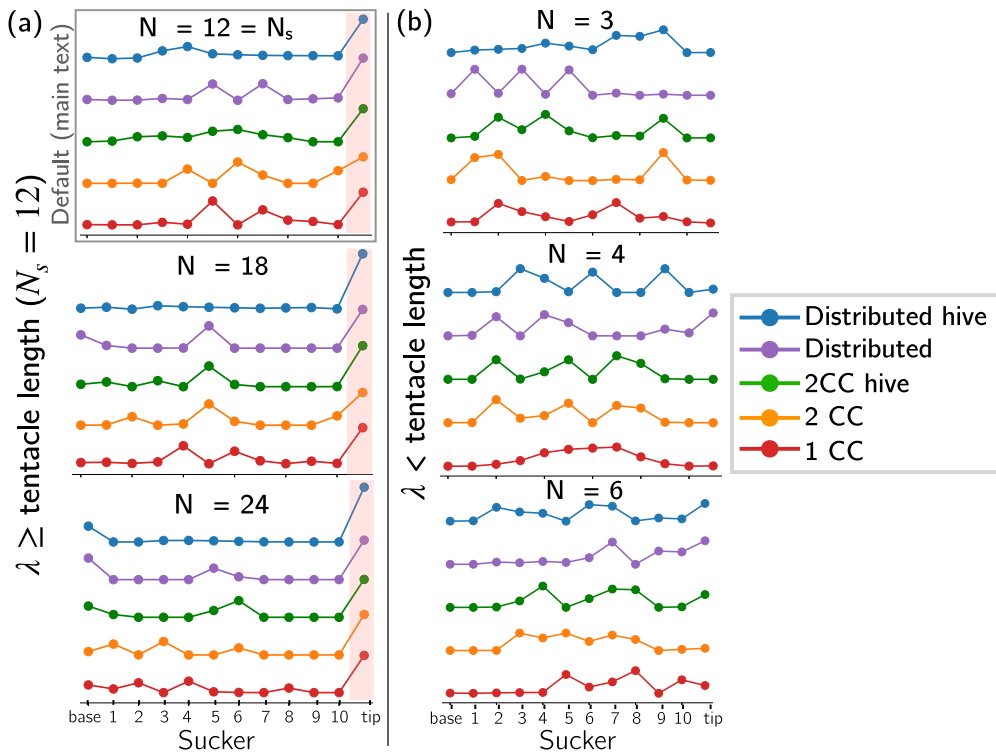


FIG. 10. Effect of different wave forms of the CPG over the first order sucker importance for a $N_s = 12$ suckers crawler. We can note that when the wavelength is equal or larger than the crawler length, $N \geq N_s$, the head is always the most impactful sucker for all the considered learning architectures. This invariance does not hold anymore when $\lambda < L$.

dropping the assumption that the CPG wavelength, λ , is identical to the crawler length, L . The results, shown in Fig. 10, indicate that the head is always the most determinant sucker only when the CPG wavelength is $\lambda \geq L$. However, when $\lambda < L$ we cannot find a single sucker to be the most important for all the considered architecture, but its

identity varies with the considered policy (architecture) and wavelength. We do not have yet a definite explanation of this phenomenon, but this results already indicates a nontrivial interplay between the underlying physic of crawling, namely the (discrete) placement of each sucker with respect to the CPG, and crawling policies.

[1] W. N. Frost, Neuroethology: Generating complex behavioral sequences with distributed nerve nets, *Curr. Biol.* **33**, R398 (2023).

[2] J. Bielecki, S. K. Dam Nielsen, G. Nachman, and A. Garm, Associative learning in the box jellyfish *Tripedalia cystophora*, *Curr. Biol.* **33**, 4150 (2023).

[3] C. A. Freas and K. Cheng, Neuroecology beyond the brain: Learning in Echinodermata, *Learn. Behav.* **50**, 20 (2022).

[4] T. Kano, D. Kanauchi, H. Aonuma, E. G. Clark, and A. Ishiguro, Decentralized control mechanism for determination of moving direction in brittle stars with penta-radially symmetric body, *Front. Neurobot.* **13**, 66 (2019).

[5] E. G. Clark, D. Kanauchi, T. Kano, H. Aonuma, D. E. G. Briggs, and A. Ishiguro, The function of the ophiuroid nerve ring: How a decentralized nervous system controls coordinated locomotion, *J. Exp. Biol.* **222**, 192104 (2018).

[6] T. Po, E. Kanso, and M. J. McHenry, Cooperative transport in sea star locomotion, *Curr. Biol.* **34**, 2551 (2024).

[7] S. Heydari, A. Johnson, O. Ellers, M. J. McHenry, and E. Kanso, Sea star inspired crawling and bouncing, *J. R. Soc. Interface* **17**, 20190700 (2020).

[8] C. A. Allard, W. A. Valencia-Montoya, and N. W. Bellono, Cephalopod chemotactile sensation, *Curr. Biol.* **33**, R1081 (2023).

[9] J. A. Mather and L. Dickel, Cephalopod complex cognition, *Curr. Opin. Behav. Sci.* **16**, 131 (2017).

[10] G. Levy, T. Flash, and B. Hochner, Arm coordination in octopus crawling involves unique motor control strategies, *Curr. Biol.* **25**, 1195 (2015).

[11] T. Gutnick, R. A. Byrne, B. Hochner, and M. Kuba, *Octopus vulgaris* uses visual information to determine the location of its arm, *Curr. Biol.* **21**, 460 (2011).

[12] L. Zullo, G. Sumbre, C. Agnisola, T. Flash, and B. Hochner, Nonsomatotopic organization of the higher motor centers in octopus, *Curr. Biol.* **19**, 1632 (2009).

[13] M. Dickinson, C. Farley, R. Full, M. Koehl, R. Kram, and S. Lehman, How animals move: An integrative view, *Science* **288**, 100 (2000).

[14] K. Nishikawa, A. A. Biewener, P. Aerts, A. N. Ahn, H. J. Chiel, M. A. Daley, T. L. Daniel, R. J. Full, M. E. Hale, T. L. Hedrick, *et al.*, Neuromechanics: An integrative approach for understanding motor control, *Integr. Comp. Biol.* **47**, 16 (2007).

- [15] A. J. Ijspeert, Central pattern generators for locomotion control in animals and robots: A review, *Neural Netw.* **21**, 642 (2008).
- [16] T. Umedachi, K. Takeda, T. Nakagaki, R. Kobayashi, and A. Ishiguro, Fully decentralized control of a soft-bodied robot inspired by true slime mold, *Biol. Cybern.* **102**, 261 (2010).
- [17] J. C. Spagna, D. I. Goldman, P.-C. Lin, D. E. Koditschek, and R. J. Full, Distributed mechanical feedback in arthropods and robots simplifies control of rapid running on challenging terrain, *Bioinspir. Biomimet.* **2**, 9 (2007).
- [18] I. D. Neveln, A. Tirumalai, and S. Sponberg, Information-based centralization of locomotion in animals and robots, *Nat. Commun.* **10**, 3655 (2019).
- [19] K. Yasui, T. Kano, E. M. Standen, H. Aonuma, A. J. Ijspeert, and A. Ishiguro, Decoding the essential interplay between central and peripheral control in adaptive locomotion of amphibious centipedes, *Sci. Rep.* **9**, 18288 (2019).
- [20] K. Yasui, A. Gupta, Q. Fu, S. Suzuki, J. Hainer, L. Paez, K. Lutek, J. Arreguit, T. Kano, E. M. Standen, *et al.*, Multisensory feedback makes swimming circuits robust against spinal transection and enables terrestrial crawling in elongate fish, *Proc. Natl. Acad. Sci. USA* **122**, e2422248122 (2025).
- [21] S. Suzuki, T. Kano, A. J. Ijspeert, and A. Ishiguro, Spontaneous gait transitions of sprawling quadruped locomotion by sensory-driven body–limb coordination mechanisms, *Front. Neurobot.* **15**, 645731 (2021).
- [22] R. J. Full and D. E. Koditschek, Templates and anchors: Neuro-mechanical hypotheses of legged locomotion on land, *J. Exp. Biol.* **202**, 3325 (1999).
- [23] A. Ayali, A. Borgmann, A. Büschges, E. Couzin-Fuchs, S. Daun-Gruhn, and P. Holmes, The comparative investigation of the stick insect and cockroach models in the study of insect locomotion, *Curr. Opin. Insect Sci.* **12**, 1 (2015).
- [24] D. E. Koditschek, R. J. Full, and M. Buehler, Mechanical aspects of legged locomotion control, *Arthropod Struct. Dev.* **33**, 251 (2004).
- [25] K. Pearson, Ö. Ekeberg, and A. Büschges, Assessing sensory function in locomotor systems using neuro-mechanical simulations, *Trends Neurosci.* **29**, 625 (2006).
- [26] S. Mishra, W. M. Van Rees, and L. Mahadevan, Coordinated crawling via reinforcement learning, *J. R. Soc. Interface* **17**, 20200198 (2020).
- [27] A. Jan Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, Dynamical movement primitives: Learning attractor models for motor behaviors, *Neural Comput.* **25**, 328 (2013).
- [28] B. Hartl, M. Levin, and A. Zöttl, Neuroevolution of decentralized decision-making in N -bead swimmers leads to scalable and robust collective locomotion, *Commun. Phys.* **8**, 194 (2025).
- [29] Y. Tanaka, K. Ito, T. Nakagaki, and R. Kobayashi, Mechanics of peristaltic locomotion and role of anchoring, *J. R. Soc. Interface* **9**, 222 (2012).
- [30] B. Chong, J. He, D. Soto, T. Wang, D. Irvine, G. Blekherman, and D. I. Goldman, Multilegged matter transport: A framework for locomotion on noisy landscapes, *Science* **380**, 509 (2023).
- [31] S. Grillner, Biological pattern generation: The cellular and computational logic of networks in motion, *Neuron* **52**, 751 (2006).
- [32] Q. Wen, M. D. Po, E. Hulme, S. Chen, X. Liu, S. Kwok, M. Gershow, A. Leifer, V. Butler, C. Fang-Yen, *et al.*, Proprioceptive coupling within motor neurons drives *C. elegans* forward locomotion, *Neuron* **76**, 750 (2012).
- [33] J. H. Boyle, S. Berri, and N. Cohen, Gait modulation in *C. elegans*: An integrated neuromechanical model, *Front. Comput. Neurosci.* **6**, 1 (2012).
- [34] C. Pehlevan, P. Paoletti, and L. Mahadevan, Integrative neuromechanics of crawling in *D. melanogaster* larvae, *eLife* **5**, e11031 (2016).
- [35] J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, and M. Kawato, Learning from demonstration and adaptation of biped locomotion, *Rob. Auton. Syst.* **47**, 79 (2004).
- [36] N. G. Hatsopoulos and W. H. W. Jr., Resonance tuning in rhythmic arm movements, *J. Mot. Behav.* **28**, 3 (1996).
- [37] J. Lynch, J. Gau, S. Sponberg, and N. Gravish, Dimensional analysis of spring-wing systems reveals performance metrics for resonant flapping-wing flight, *J. R. Soc. Interface* **18**, 20200888 (2021).
- [38] J. C. Tuthill and E. Azim, Proprioception, *Curr. Biol.* **28**, R194 (2018).
- [39] R. S. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., Adaptive Computation and Machine Learning (The MIT Press, Cambridge, MA, 2020).
- [40] S. Mahadevan, Average reward reinforcement learning: Foundations, algorithms, and empirical results, *Mach. Learn.* **22**, 159 (1996).
- [41] A. Plaat, *Deep Reinforcement Learning* (Springer Nature, Singapore, 2022).
- [42] D. L. Young and C. Eccles, in *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, edited by T. Pham, L. Solomon, and K. Rainey (SPIE, Bellingham, WA, 2020), p. 44.
- [43] L. Gagliardi, Git repo: Control a 1D crawler (2025), https://github.com/PimLb/control_1D_crawler.git.
- [44] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/9snq-xgqh> for animations of different crawling policies.
- [45] W. Stern-Tomlinson, M. Nusbaum, L. Perez, and W. J. Kristan, A kinematic study of crawling behavior in the leech, *Hirudo medicinalis*, *J. Comp. Physiol. A* **158**, 593 (1986).
- [46] B. Hochner, L. Zullo, T. Shomrat, G. Levy, and N. Nesher, Embodied mechanisms of motor control in the octopus, *Curr. Biol.* **33**, R1119 (2023).
- [47] S. Chen, Y. Cao, M. Sarparast, H. Yuan, L. Dong, X. Tan, and C. Cao, Soft crawling robots: Design, actuation, and locomotion, *Adv. Mater. Technol.* **5**, 1900837 (2020).
- [48] T. Sato, W. Watanabe, and A. Ishiguro, in *Proceedings of the IEEE International Conference on Robotics and Automation* (IEEE, Los Alamitos, CA, 2010), pp. 709–714.
- [49] B. Mazzolai, A. Mondini, F. Tramacere, G. Ricconi, A. Sadeghi, G. Giordano, E. Del Dottore, M. Scaccia, M. Zampato, and S. Carminati, Octopus-inspired soft arm with suction cups for enhanced grasping tasks in confined environments, *Adv. Intell. Syst.* **1**, 1900041 (2019).
- [50] S. Sareh, K. Althoefer, M. Li, Y. Noh, F. Tramacere, P. Sareh, B. Mazzolai, and M. Kovac, Anchoring like octopus: biologically inspired soft artificial sucker, *J. R. Soc. Interface* **14** (2017).
- [51] J. Peters and S. Schaal, Reinforcement learning of motor skills with policy gradients, *Neural Netw.* **21**, 682 (2008).
- [52] H. V. Hasselt, Double Q-learning, *Advances in Neural Information Processing Systems* **23**, 1 (2010).

- [53] M. Kim, J.-S. Kim, M.-S. Choi, and J.-H. Park, Adaptive discount factor for deep reinforcement learning in continuing tasks with uncertainty, *Sensors* **22**, 7266 (2022).
- [54] Y. Kuramoto, in *International Symposium on Mathematical Problems in Theoretical Physics*, edited by H. Araki (Springer-Verlag, Berlin, 1975), Vol. 39, pp. 420–422.
- [55] S. H. Strogatz, From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators, *Physica D* **143**, 1 (2000).
- [56] G. Jékely, F. Keijzer, and P. Godfrey-Smith, An option space for early neural evolution, *Philos. Trans. R. Soc. B* **370**, 20150181 (2015).
- [57] D. Arendt, M. A. Tosches, and H. Marlow, From nerve net to nerve ring, nerve cord and brain—Evolution of the nervous system, *Nat. Rev. Neurosci.* **17**, 61 (2016).
- [58] F. Keijzer and A. Arnellos, The animal sensorimotor organization: A challenge for the environmental complexity thesis, *Biol. Philos.* **32**, 421 (2017).
- [59] M. J. Wells, Proprioception and visual discrimination of orientation in *Octopus*, *J. Exp. Biol.* **37**, 489 (1960).
- [60] H. Bagheri, A. Hu, S. Cummings, C. Roy, R. Casleton, A. Wan, N. Erjavic, S. Berman, M. M. Peet, D. M. Aukes, *et al.*, New insights on the control and function of octopus suckers, *Adv. Intell. Syst.* **2**, 1900154 (2020).
- [61] L. Van Giesen, P. B. Kilian, C. A. Allard, and N. W. Bellono, Molecular basis of chemotactile sensation in octopus, *Cell* **183**, 594 (2020).
- [62] L. Gagliardi and A. Seminara, Codes and scripts to reproduce results from present publication (2025), https://github.com/PimLb/control_1D_crawler.