

Simulation to Real

Learning Energy-Efficient Slithering Gaits for a Snake-Like Robot

By Zhenshan Bing , Long Cheng , Kai Huang , and Alois Knoll

© PHOTOCREDIT

To resemble the body flexibility of biological snakes, snake-like robots are designed as a chain of body modules, which gives them many degrees of freedom (DoF) on the one hand and leads to a challenging task to control them on the other.

Compared with conventional model-based control methods, reinforcement learning (RL)-based ones provide promising solutions to design agile and energy-efficient gaits for snake-like robots as RL-based methods can fully exploit the hyperredundant bodies of the robots. However, RL-based methods for snake-like robots have rarely been investigated even in simulations, let alone been deployed on real-world snake-like robots. In this work, we introduce a novel approach for designing energy-efficient gaits for a snake-like robot, which first learns a policy using an RL algorithm in simulation and then transfers it to the real-world testing, thereby leveraging a fast and economical gait-generation process. We evaluate our RL-based approach in both

simulations and real-world experiments to demonstrate that it can generate substantially more energy-efficient gaits than those generated by conventional model-based controllers.

Introduction

Inspired by the versatile locomotion skills of biological snakes, researchers have designed many kinds of snake-like robots to mimic the structural and functional properties of their counterparts in nature. With distinctive skills, snake-like robots are expected to be deployed in special task scenarios, where other kinds of robots are incapable to be used, such as disaster rescue, military surveillance, and pipeline maintenance in factories. As snake-like robots can only carry limited energy resources for field operations, it is essential to design energy-efficient gaits to reduce the impact of the power constraints. An energy-efficient movement, on the one hand, can maximize the service time of a robot while maintaining its performance. On the other hand, it can, in return, allow us to optimize the design of the robot by lowering its weight or adding other hardware [1]. However, it is difficult to design energy-efficient gaits for snake-like robots due to their redundant

Digital Object Identifier 10.1109/MRA.2022.3204237

Date of current version: 20 September 2022

long mechanisms and complex dynamic interactions with their surroundings.

The most commonly adopted method for controlling a snake-like robot is to model its discrete body as a continuous curve with sideslip constraints in the planar environment, which is also widely known as the *serpenoid curve* [2]. This serpenoid curve, as a mathematical formation of the lateral undulation, parameterizes the angle of each body module to fit the shape of a sinusoidal curve and drives the robot forward using the external force generated from the side constraints. This kinematic-based method is also extended to combine two parameterized sinusoidal waves that propagate along the lateral and dorsal planes for modular snake-like robots moving in 3D space [1], [3], [4]. This parameterized kinematic method, also known as the *gait equation*, allows for the emergence of complex behaviors from low-dimensional representations with only few key parameters, greatly expanding their maneuverability and simplifying user control. With this method, researchers developed several biological gaits for snake-like robots to move in indoor and outdoor environments. However, these simplified kinematic models also pose limitations on designing agile and energy-efficient gaits as they are confined to those abstracted gait parameters; let alone the parameter tuning process is inefficient and time consuming.

Designing a gait can be regarded as a parameter optimization problem, which aims at maximizing the objective functions, such as velocity and power. For snake-like robots, it becomes even more challenging, and the gait equation method mainly suffers for two reasons [5]. The extrinsic challenge comes from the complex dynamic interaction between the ground and the redundant mechanism with many DoF. The intrinsic challenge is how to synchronize and coordinate all the body joints to exhibit a proper motion pattern integrally, which is expected to be both robust and efficient.

Conventional control algorithms for snake-like robots require not only complex manual engineering to carefully tweak the analytical models of the robot but also prior knowledge of the gaits for a specific robotic system. From a control perspective, conventional control methods often require accurate models of the system and expertise on other domain knowledge about the environment, which is usually complex, inaccurate, and difficult to be acquired in the real world. From an implementation perspective, specialized methods designed to control snake-like robots usually comprise complex software architectures, which are lengthy and error prone. Therefore, it is promising to develop end-to-end learning methods that can greatly automate the gait-design process. Such end-to-end learning methods should be able to remove the need for accurate dynamic models or domain knowledge and can be applied to robotic systems without explicit system identification or too much manual engineering.

Instead of solving robotic tasks based on controllable kinematic or dynamic models, deep RL methods solve similar tasks by mimicking the natural learning process in the

trial-and-error paradigm. Agents can learn to master diverse and imaginative motor skills by only trying to maximize the reward signals. Although some robots can be directly trained in the real world to learn motor skills, it is more feasible to learn complex skills from the simulation first as it can greatly assist the development of RL algorithms and leads to a fast, cheap, and safe approach by alleviating the need for expensive and tedious real-world experiments.

On the basis of our previous work [6], this article aims to explore the gait-generation approach by learning a neural network (NN) controller via RL for snake-like robots and transfer learned knowledge from simulations to real-world implementations. We largely extend our work from [6] and study our research in five steps. First, instead of using a simplified snake robot, we design a new snake-robot prototype, both in the simulation environment and real world. Second, we design a reward signal that aims at encouraging energy-efficient locomotion. This reward function takes both the velocity and power consumption into consideration and integrates them after normalization. Third, instead of training only the controller using the proximal policy optimization (PPO) algorithm, we redesign the training pipeline by utilizing a dynamic randomization strategy, an informative prediction network, and the PPO algorithm to train the NN-based controller. The randomization includes the observation space and some of the crucial physical parameters. And the informative prediction network can predict unobservable features using the observable information via supervised learning. Fourth, to show the effectiveness of the proposed method in a more general context, we perform simulation tasks with different robot designs (joint number) and physical parameters (friction coefficient). We demonstrate that the proposed method can successfully leverage the dynamics of the robot in a general context and generate energy-efficient gaits in diverse scenarios. Finally, we compare gait results from the NN-based controller with results from the model-based method, both in simulation and real-world experiments.

Our main contributions are summarized as follows. First, different from prior works that used RL-based methods to adapt the parameters of a low-level gait generator, our NN-based controller takes full advantage of the flexible body of the snake-like robot and learns sophisticated gaits, simply from scratch. With the help of the reward signal and the iterative RL scheme, this method offers a new alternative for snake-like robots to study complicated motor skills without requiring any prior knowledge. Second, with the dynamic randomization strategy and informative prediction network, our NN-based controller successfully learns robust and adaptive behaviors in both simulation and the real world, where some physical parameters and sensory measurements are either difficult to obtain or inaccurate. Existing solutions that are used for sim-to-real transfer, such as domain randomization [7], can improve only the controller's performance on the basis of a known and yet inaccurate domain. Our informative prediction network can predict observations that cannot be measured in a self-supervised manner. Third, by comparing

the results, we demonstrate the superiority of our RL-based method against model-based ones in terms of energy efficiency.

Related Work

Gait Optimization With Model-Based Methods

To mimic the body structure and motion functions of biological snakes, snake-like robots are usually designed as a chain of body modules. These modules are used to interact with the environment to propel the robot, thereby resulting in a highly complex system with which to be controlled. A biomorphic approach, i.e., a gait equation [1], [2], is widely used to control the locomotion of snake-like robots as it can easily mimic snake-like motions by shaping its body as a sinusoid curve. The gait equation can be regarded as a process to simplify parametric representations of snake-like trajectories. Crespi and Ijspeert adopted a heuristic optimization algorithm to rapidly adjust the travel speed of an amphibious snake robot [8]. Tesch et al. used the Bayesian optimization approach to regulate open-loop gait parameters for snake robots, which made the robot move faster and more reliably [9]. However, all these works were confined by the parameterized gait-generation system and have a very limited effect on further improving the energy efficiency of a gait. Moreover, this open-loop method is also time consuming and inefficient [10] for fine-tuning gait parameters to achieve expected control objectives.

Gait Optimization With RL-Based Methods

As an intelligent trial-and-error learning method, RL brings new solutions for free gait-generation tasks without knowing precise models or prior knowledge. RL technologies were used to fine-tune parameters for motor skills in a model-based manner, which can either be the controllable model of the system [11] or another abstracted gait generator such as the central pattern generator (CPG) [12]. Shi et al. [13] developed a simplified snake-like robot with two actuated joints and implemented a controller based on a deep Q-network algorithm. They demonstrated that their method could produce meaningful gaits. Similar to the concept in [14], Liu et al. proposed an RL-based controller to modulate the activity of a CPG controller for generating goal-reaching and tracking behaviors for a soft, snake-like robot in simulations [15]. Chatterjee et al. proposed a policy-improvement method to choose different control parameters for a snake-like robot with screw-drive units [16]. Sartoretti et al. presented an RL-based approach to adapt the shape parameters of a snake-like robot in response to external sensing [17].

Sim-to-Real Transfer

To narrow simulation-reality gaps, researchers have attempted to improve simulation fidelity by building up accurate models and refine them using real data. For instance, to match the performance of the real system, an actuator can be modeled and refined using its real data [18]. Researchers have also worked on improving the robustness of a learned policy to

variations of system properties and perception information, thereby enabling it to be feasible and adaptive in real-world systems [19]. The task to transfer a policy learned from simulation to the real world is treated as an instance of domain adaption, where the source domain (simulation) is modeled as close as possible to the target domain (real world) [7]. Robotic arms are the most common agents to deploy learned policy from simulations as the kinematics and dynamics can be accurately modeled. Christiano et al. controlled a robotic arm using learned policy from simulation, of which the inverse dynamics (from observations to actions) of that robot was learned from the data generated by the forward dynamics (from actions to observations) in simulations [20].

Domain randomization is another technique for sim-to-real transfer based on the introduction of higher variance in the domain-specific—but task-irrelevant—features of the training data. Peng et al. propose introducing variability in the dynamics of the simulation by sampling values for strategic features (e.g., friction, mass, and damping coefficients) during the training phase [21]. They argued that although modeling the simulation in accordance with the real environment is important for the sim-to-real policy transfer, a complete and accurate model is often impossible due to many reasons, such as unforeseen forces, left-out environmental characteristics, calibration issues, and so on. In the meantime, Wulfmeier et al. argued that an unfavorable domain in randomization can lead to poor performance of the policy that is transferred from simulation [22]. They proposed another idea to tackle systematic model discrepancies: aligning the distributions over visited states between the simulated and the real-world agent. They demonstrated their argument by performing experiments between two simulations with either different parameterizations or completely different simulation engines to create situations of misaligned and unknown system dynamics.

Robot and Model

Snake-Like Robot

Our planar snake-like robot adopts a modularized design manner, which consists of eight identical actuated body modules and one head module to slither forward (see Figure 1). Each module is connected to the adjacent module by an actuated joint, which can rotate 90° in both the left and right directions. Each body module is composed of an actuation system, the control system, housing components, and a pair of passive wheels to imitate the anisotropic friction property of the snake skin. The actuation system consists of a servo, gearbox, and an angular sensor to feed back the angular position of the output joint. The dc servo has a maximum torque of 12.8 Kg cm and drives a gearbox with a reduction factor of 3.71. The angular sensor is a Hall effect encoder that is used to measure the angular position of the output shaft of each module. The control system is a customized STM32 microcontroller. The STM32 runs three tasks: controlling the servo, reading the joint angle, and communicating with the other




Figure 1. The snake-like robot and its module dimensions. PCB: printed circuit board.

modules. A real-time operation system is executed on the control system, and the control loop is set as 20 Hz. Table 1 summarizes the technical specifications of the robot.

Simulation Model

We model the snake-like robot as close as the real-world counterpart and simulate it in Multi-Joint dynamics with Contact (MuJoCo). As the fundamental component of a motion’s behavior, the dynamic model plays a great role in narrowing the gap between the simulation and the real world. There are several key parameters that directly determine the accuracy of the dynamic model, such as the mass of the robot, inertia, friction, and other physical parameters.

Modeling Mass

The mass can be weighted exactly by measuring the module of the prototype, while the moment of inertia is estimated from its CAD assembly, in which each component is given its material.

Modeling Actuation

The simulated servo system consists of a position-controlled motor, feedback signal of the joint position, and gear ratio. For the sake of simplicity, by setting the gear ratio to one, we directly map the motor position to the output position of the servo system. The torque is limited to a range of $[-4.6, 4.6]$ in the Newton meter based on specifications of the servo.

Modeling Friction

To imitate this anisotropic friction property, each robot module is equipped with two passive wheels. Those passive wheels enable a minimum friction in the direction of rotation and a

Table 1. Technical details of the robot module.

Parameters	Quantity
Dimensions	Width: 70 mm; length: 100 mm; height: 54 mm
Mass	Body: 202 g; one wheel: 2.2 g
Servo	Max torque: 12.8 kg·cm; max speed: 0.07 s/60°
Gearbox	Gear ratio: 151:43
Communication	CAN bus
Sensors	AS5047D (resolution: 0.022°)
Power	Voltage: 24 V; max current: 10 A
Control rate	20 Hz
Max: maximum; CAN: controller area network.	

high friction in the lateral direction (see Figure 1). To obtain realistic movements, we directly measured the friction coefficients for these two directions in the real world.

Energy-Efficiency Metric

We aim to design gaits that can make snake-like robots move energy efficiently and in the meantime keep a steady speed of the movement. In this work, we define the power efficiency metric to evaluate the performance of all gaits.

The total power P of the robot can be calculated by adding up the power of all the N joints, which can be calculated as $P = \sum_{i=1}^N |\tau_i \dot{\phi}_i|$, where $\dot{\phi}_i$ is the velocity of joint i . The torque τ_j is the product of its applied force f_j and its gear constant parameter h_i (the length of the actuator). The model uses actuators with a limited force of f_{\max} as the maximum force

in both directions. With this property, the normalized power consumption \hat{P} is calculated as

$$\hat{P} = \frac{1}{N} \sum_{i=1}^N \frac{|f_i h_i \dot{\phi}_i|}{f_{\max} h_i \dot{\phi}_{\max}}.$$

This normalized power consumption \hat{P} will be further used for reward definition.

With the power P and the velocity v , several efficiency metrics can be calculated. The usual way is to calculate the cost of transport, which is the power P divided by the mass m , the gravity g and the velocity v :

$$\text{COT} = \frac{E}{mgd} = \frac{P}{mgv}. \quad (1)$$

This unit-less measurement is able to compare the efficiency between different mobile systems. As the same robot and environmental properties are compared, those constants only scale the results and would not provide any additional insight for the comparison. Thus, a simplified metric calculates the averaged power per velocity (APPV) as $\text{APPV} = P/v$.

Baseline Controller

In this section, we first introduce the widely used method for generating gaits for snake-like robots, which is the gait equation controller. Then we use the grid-search algorithm to explore energy-efficient gaits at different velocities by brutally searching the parameter space with fine steps.

Gait Equation Controller

As the most widely adopted method, the gait equation controls the locomotion gaits of a snake-like robot by shaping its back curve as a moving sinusoid wave that propagates along the body. This method or similar ideas (e.g., CPG-based methods) have been used for decades by many kinds of snake-like robots to generate different gaits [2]. We utilize the gait equation from [9], which is modeled as

$$\phi(n, t) = \left(\frac{n}{N}x + y\right) \cdot A \cdot \sin(\omega t + \lambda n). \quad (2)$$

$\phi(n, t)$ presents the joint-angle value at time t , where n is the joint index and N is the total joint number. λ and ω are the

spatial and temporal frequency of the movement, respectively. The spatial frequency represents the cycle number of the wave, and the temporal frequency represents the traveling speed of the wave. A is the serpentine amplitude and x and y are the constants for shaping the body curve.

Grid-Search Optimization

We take the grid-search algorithm to explore the parameter space, which can lead to optimum gaits as long as the searching intervals are small enough. We use the grid-search method to generate a variety of gaits and determine those parameter combinations with the best power efficiency at different velocities. The grid-search method generates a Cartesian product from the parameters in Table 2, resulting in 10,080 parameter sets. Then, each motion parameter set gets tested by running 1,000 steps in the simulation environment. For each run, the first 200 time steps are ignored and the remaining 800 time steps are evaluated for collecting experimental data. It is because it has been observed that the snake robot needs approximately 200 time steps to accelerate and then moves at a steady speed.

NN-Based Controller

Observation Space

The full observation space consists of all the joint position ϕ_j , the joint velocity $\dot{\phi}_j$, the actuator torque τ_j , the head module velocity v_1 , and the target velocity v_t . The joint position ϕ_j and its joint velocity $\dot{\phi}_j$ are required to learn the locomotion and represent the proprioceptive awareness of the robot. The head module velocity v_1 helps to sense its global velocity, which offers better movement awareness. To learn an energy-efficient gait, the sense of energy consumption is necessary. Therefore, the actuator torque of each joint τ_j is provided and can be interpreted in combination with $\dot{\phi}_j$ to determine the total power usage. The specified target velocity v_t is passed to the environment and can be changed, which is required to control velocity of the robot.

Action Space

The action space has 8 DoF with finite-continuous values in the range of $[-1.5, 1.5]$, which linearly translates to a corresponding joint angle ϕ in the range of $[-90^\circ, 90^\circ]$. Each action represents eight actuator angle positions of the servo motors. In the real-world setup, the joint can turn around with the motor speed set as $0.07 \text{ s}/60^\circ$ and the gear ratio as $151:43$. The control frequency for the snake-like robot is set as 20 Hz. Thus, the maximum turning angle in 0.05 s is set as 12° in the simulation. The simulation time step is also set as 0.05 s. It should be noted that 20 Hz is selected to ensure real-time calculation of the NN-based controller.

Reward Function

The objective of our task is to learn an energy-efficient gait for a variety of specified velocities. Therefore, the energy consumption and the difference between the actual model

Table 2. The gait parameters used for the grid search.

Parameters	Descriptions	Values
ω	Temporal frequency	0.2, 0.4, 0.6, 0.8, 1, 1.2, 1.4, 1.6, 1.8, 2
A	Amplitude ($^\circ$)	30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80, 85
y	Linear reduction	0.1, 0.2, 0.3, 0.4
x	Linear reduction	$(1-y)$
λ	Spatial frequency ($^\circ$)	40, 45, 50, 55, 60, 65, 70, 75, 80, 85, 90, 95, 100, 105, 110, 115, 120, 125, 130, 135, 140

velocity and the target velocity are the main criteria to find a successful behavior. The challenge is to combine the variables into one numerical reward and weigh them for each time step. Therefore, we first split the power efficiency and velocity criteria into two normalized reward function components.

First, a normalized reward is defined to maintain the specified velocity. The objective is to reach and maintain the target velocity v_t by comparing it with the head velocity v_1 . The following function represents the velocity reward:

$$r_v(t) = \left(1 - \frac{|v_t - v_1|}{a_1}\right)^{\frac{1}{a_2}}. \quad (3)$$

$a_1 = 0.2$ influences the spread of the reward curve by defining the x -axis intersections with $x = v_t \pm a_1$. $a_2 = 0.2$ affects the changes of the curve's gradient. If $|v_t - v_1| = 0.0$, the velocity reward r_v reaches the maximum value of one.

Second, the normalized value of the total power usage \hat{P} in the "Energy-Efficiency Matrix" section is used to determine the power efficiency reward component r_p , which is represented by

$$r_p(t) = r_{\max} \left|1 - \hat{P}\right|^{b_1^{-2}}. \quad (4)$$

Here, r_{\max} is the maximum reward value, and $b_1 = 0.6$ is the slope of the curve. Power efficiency is influenced by the desired target velocity. Thus, the normalized r_{\max} represents this influence by limiting the maximum value of r_p .

Finally, rewards from the velocity r_v and the power efficiency r_p are combined to form the overall reward r :

$$r(t) = \left(1 - \frac{|v_t - v|}{a_1}\right)^{\frac{1}{a_2}} \left|1 - \hat{P}\right|^{b_1^{-2}}. \quad (5)$$

This equation replaces the r_{\max} in (4) with r_v . With that, the maximal power efficiency depends on the absolute value of the difference between the desired velocity and the robot velocity and total power consumption.

Network Architecture

Given the input (observation \mathbf{o}_i^t) and output (action \mathbf{a}_i^t), we now elaborate the policy network mapping \mathbf{o}_i^t to \mathbf{a}_i^t . We design a fully connected two-hidden-layer NN as a nonlinear function approximator to the policy π_θ . The input layer has the same dimension as the action space \mathbf{o}_i^t . Both hidden layers each have 200 neurons and are each followed by a rectified linear unit layer. The final layer outputs joint position commands for the robot. To train the network, the PPO algorithm adapted from [23] is used.

We train our policy network on a computer with an i7-7700 CPU and a Nvidia GTX 1080 GPU. A total of six million time steps (roughly 7,500 episodes) are used for training. The maximum number of time steps in an episode is 2,000. With the environment settings of 50 ms per time step, the training takes approximately 42 h in total simulation time and 2 h in wall-clock time for the policy to converge.

Sim-to-Real Transfer

We present two approaches that are designed for robust transfer of the policy learned in simulation to the real world, namely, the parameter randomization and informative network.

Parameter Randomization

Parameter randomization is an effective approach to improve the robustness of the system, especially taking the gap between simulation and real-world experiments into consideration. By randomizing the physical parameters and observations during training, the learned policy can be more robust for deployment in real-world experiments.

Physical Parameter Randomization in Simulation

The physical parameters given in the "Robot and Model" section are either measured or estimated, which may lead to inaccuracy. And different settings of physical parameters will directly impact performance of the generated gaits. When the agent is trained in a stable environment with fixed dynamic parameters, it usually leads to an overfitting controller, which will not work properly in the real-world environment. We randomly sample physical parameters, as listed in Table 3.

Observation-Space Randomization

In a real-world setup, the joint angles cannot be perfectly measured due to several uncertain factors, such as the noise of the encoder and the data errors or delay caused by the communication. To eliminate the impact of this inaccuracy, we add a Gaussian noise $\mathcal{N}(\mu, \sigma)$ on every joint-angle position in the observation space during training, where $\mu = 0$ rad and $\sigma = 0.05$ rad.

Informative Prediction Network

The observation space used in the simulation is only partially observable or measurable in the real world. To solve this problem, we propose an informative prediction network to predict the joint torque, joint velocity, and head velocity under the framework of supervised learning. We present three steps to construct such an informative prediction network (see Figure 2).

Table 3. Technical details of the robot module. During simulations, these physical parameters are randomly selected from the randomization range of the baseline value of each parameter.

Parameters	Baselines	Unit	Randomization Ranges
Ground friction	0.6	—	90 ~ 100%
Mass	0.206	kg	90 ~ 100%
Armature inertia	0.01	kg·m ²	80 ~ 120%
Motor damping	0.3	N·s/m	90 ~ 110%

Compact Observation Space

In the real world, some of the observations are difficult to acquire, such as the joint velocity $\dot{\phi}_j$ and the joint torque τ_j . We can calculate the joint velocity by differentiating two consecutive joint positions at each time step. However, this is infeasible due to the noise of the joint position. To solve this issue, we also design a compact observation space that contains only the joint position ϕ_j and the target velocity v_t . The informative network is used to infer the corresponding joint velocity $\dot{\phi}_j$, joint torque τ_j , and head velocity v_1 according to the given joint position ϕ_j . The input of the network is the joint position ϕ_i . The informative network is trained using supervised learning, where the dataset is directly sampled from simulation.

Data Generation

With a well-trained NN-based controller, the agent is controlled to slither in the simulation environment with a varying target velocity to sample training data. Each episode contains 400 time steps, the target velocity starts from 0.05 m/s and slightly increases $1.33e^{-3}$ m/s at each episode until 0.25 m/s. One logged data point contains the target velocity, joint position of the timestamp, and values of the last action of the NN-based controller. The generated data are shuffled and split, where 80% of the dataset is used for training and the other 20% is used for testing.

Informative Network Model

Three observation features, namely, the joint velocity $\dot{\phi}_j$, joint torque τ_j , and head velocity v_1 , used for training the NN-based controller in the simulation are not available in

the real world. Because each feature has its own characteristic, we design three separate prediction models to predict them with the same architecture. For one prediction network, the input consists of the action values and the joint positions for each joint in the last two time steps, together with the target velocity.

As explained in [24], dynamics of the actuators are independent of each other and, accordingly, they trained their models for predicting actuator forces separately. However, for the snake-like robot, we believe that the action of one joint may affect other joints because the snake robot is a class of serially connected, active-cord mechanisms. Therefore, we adopted two approaches for sampling the input data: the jointwise and nonjointwise models. For a jointwise model, each value of the feature is strongly related to one's specific joint index, e.g., the torque of the first joint is predicted using only the data from the first joint. The jointwise prediction models split the input features of the input layer by joint and process them separately. The processed values are concatenated at the output layer. Oppositely, the nonjointwise models process values from all the existing joints together at once as input and predict output features for all the joints at once. The main assumption of this is that a target feature depends on the entire kinematics of the snake-like robot.

In this work, prediction models with two layers of multi-layer perceptron long short-term memory (LSTM), or a single Temporal Convolutional Network layer (which consists of multiple convolution layers) were tested. For the length of input sequences (the number of past time steps), two, four, and eight were used for training to observe how the input-sequence length affects model performance. The

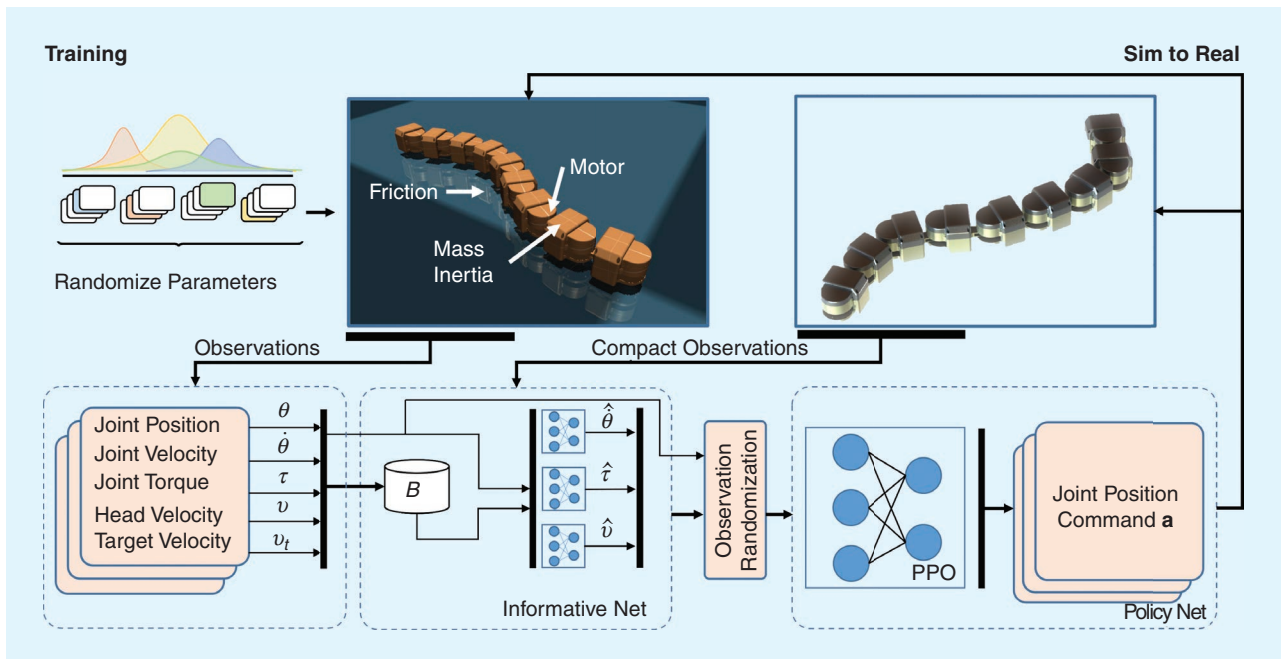


Figure 2. The training pipeline for the sim-to-real transfer of the NN-based controller. There are three steps in this pipeline. First, the physical parameters are randomized in each episode during training. Second, the observation space is used to train an informative NN to predict the observations that cannot be acquired directly from real-world environments. Third, the randomized observations are fed to the NN-based controller, which is trained using the PPO algorithm. Sim to Real: Simulation to Real.

mean-square error was used as the loss function for training models, which is commonly used for regression tasks. We skip the training and testing details for choosing proper models or the length of input sequences and present the final setup for each feature in Table 4. To predict torque, we choose multilayer perceptron as the architecture of the prediction network, which is trained via the jointwise style. The length of the sequential input is two. To predict the joint velocity, we choose LSTM as the network architecture, which is trained via the nonjointwise style. The input length is also two. To predict the head velocity, we choose LSTM as the network architecture as well. The input length of this prediction network is eight. As the head velocity is one overall property of the robotic system, the jointwise or nonjointwise style is not applicable here. The inputs for predicting these three states are the joint position, action, and target velocity.

To enable learning of different velocities, the parameter v_t is changed by iterating over 0.05, 0.1, 0.15, 0.2, and 0.25 for each episode while training. It should be noted that in theory, we can also select target velocity from a uniform distribution to train the network. But empirical results show that the training is much more successful and stable when using discretized target velocities. Meanwhile, to simplify the beginning of the learning process, the first 100 episodes are trained with a fixed target velocity of 0.1 m/s.

Simulation Results and Comparisons

Baseline Performances

The power consumption and corresponding velocity results from the grid-search algorithm are shown in Figure 3 as a point cloud of parameter sets using dot markers. The lowest points at different velocities in the point cloud have the highest energy efficiency. To reduce the number of equation parameters, we select five temporal frequencies to show the pattern of gaits generated by the gait equation as temporal frequency is the most direct factor that impacts velocity. Second, instead of connecting the scatter points with lines, we show the confidence region of the gaits generated with each frequency. We can observe that low frequency (0.25 Hz) leads to very slow gaits, while high frequency (2 Hz) leads to energy-intensive gaits. The gaits with 1.5 Hz are more promising for generating energy-efficient gaits. We also demonstrate that the grid-search method suffers from being inefficient to search proper parameters for energy-efficient gaits as most of the parameter sets in the searching space are distributed in the low-velocity area (0 – 0.1 m/s). Note that we show all the gaits from the grid-search algorithm as we want to depict the inefficiency of the baseline method.

NN-Based Controller Performance

In this study, target velocities in the range of [0.025, 0.25] m/s with a step interval of 0.005 m/s are used for the evaluation. The simulation results demonstrate that the NN-based controller has succeeded in learning a series of gaits from scratch without knowing any prior locomotion skills.

First, the NN-based controller can perform very accurate locomotion gaits in terms of velocity even though the NN-based controller is trained with only five target velocities (0.05, 0.1, 0.15, 0.2, and 0.25 m/s). As shown in Figure 4, the

Table 4. Training configurations for prediction. The length of the input sequence means the number of past time steps.

ANN	Torque	Joint velocity	Head velocity
	MLP	LSTM	
Input feature	Joint position, action, and target velocity		
Input length	2	2	8
Nonjointwise/ jointwise	jointwise	nonjointwise	—

ANN: artificial neural network; MLP: multilayer perceptron.




Figure 3. This scatter plot directly shows the energy-consumption results of the controllers at a range of velocities in the simulation. The red points represent the performance of the NN-based controller. For the gait equation controller, we select five temporal frequencies w , namely, 0.25, 0.5, 1, 1.5, and 2 Hz. We then run the grid-search algorithm with each frequency and plot the scatters and its confidence region.

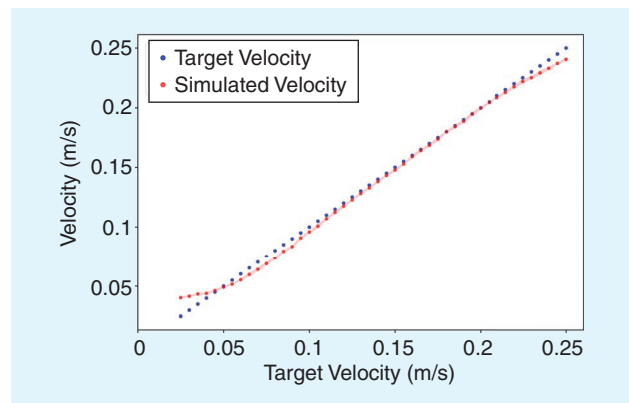


Figure 4. The actual velocities (red dots) measured from simulations and the corresponding target velocities (blue dots). The error interval of the red line represents the standard deviation over the five random seeds.

targeted velocities are represented by blue solid dots, and the measured velocities from the simulation are marked by red solid dots. We can observe that the measured velocities almost match the target velocities from 0.05 to 0.20 m/s with only very small errors close to zero. In the low-velocity range (less than 0.05 m/s) and high-velocity range (higher than 0.20 m/s), measured velocities differ from target velocities. This is because under extreme conditions, the NN-based controller is difficult to maintain target velocities while keeping optimization of the power consumption. In the low-velocity range, the NN-based controller tends to move faster as a minor modification of the motion will lead to a speed change. In the high-velocity range, the NN-based controller is limited by physical configuration of the snake-like robot to reach target velocities.

Second, the power consumption results of the learned gait are marked with red points in Figure 3. The data depict a linear relationship between travel velocity and power consumption, which is in line with the physical law $Power = Force \times Speed$. There are only a few points with higher power consumption when the velocity is below 0.08 m/s. Importantly, this also reveals the adaptability of the learning approach for generating gaits in a range of velocities. It can also be observed that the mean velocities do not exactly align with the specified interval of 0.005, especially at higher target velocities. The reason for this is the difficulty to achieve the exact ratio between holding the right velocity and performing the corresponding power-efficient locomotion.

Comparison

We first have a look at the energy metric based on averaged power per velocity. After putting the power efficiency data of

the gait equation controller (see Figure 3) and the NN-based controller together, we can clearly conclude that the NN-based controller has a much better power efficiency at a range of velocities. As the velocity grows, the advantage of the NN-based controller for saving energy is even more obvious.

To show that the proposed method can generate energy-efficient gaits in a more general context, we also evaluated both controllers in different environments (friction coefficient) and with different robot designs (the number of body modules) in simulation. The results are visualized in Figure 5. In Figure 5(a), we change the robot's design by adding or reducing body modules to test the proposed method. For a snake robot with a joint number range from six to nine (the default number is eight), the proposed method shows consistent performance superiority to generate many energy-efficient gaits across different traveling speeds. In Figure 5(b), we change the friction coefficient of the environment to test effectiveness of the proposed method. On one hand, results clearly show that the NN-based controller can generate gaits that consume less energy than the gaits generated from an equation-based controller. On the other, the NN-based controller also shows a much more stable performance in generating energy-efficient gaits, while the equation-based method exhibits fluctuating energy-consumption results.

The NN-based controller improves upon the traditional kinematic-based one in two ways. On one hand, the traditional gait equation controller is based on kinematics and describes the gait movement with no influence of physical forces such as friction or damping. It only extracts several critical parameters to represent the gait, while the interaction between the snake-like robot and the environment is highly

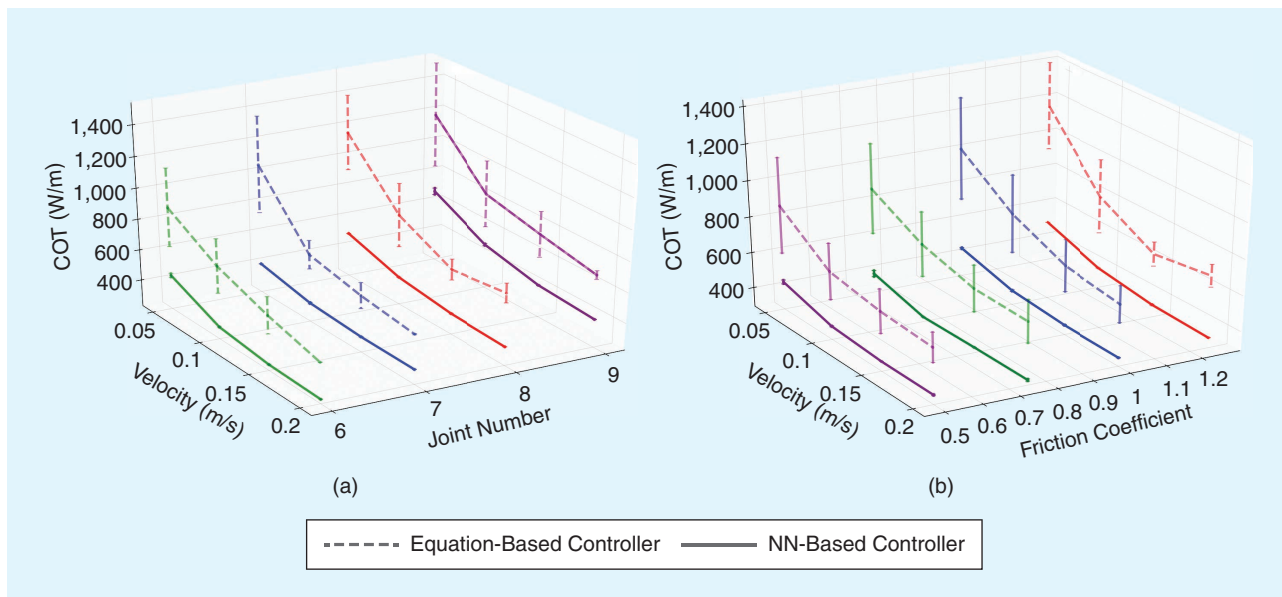


Figure 5. The simulation results from generalized scenarios. (a) The cost-of-transport (COT) performance of gaits generated from different robot design (body modules). (b) The COT performance of gaits generated from environments with different friction coefficients. The solid line shows the gaits generated from the NN-based controller. The dashed line shows the gaits generated from the gait equation controller. The error bars show the standard deviation of the energy consumption at each velocity.

complex. Although the parameterized gait equation simplifies the control task, it inevitably brings difficulties for designing more sophisticated gaits despite using traditional optimization technologies (the grid-search algorithm), especially when the parameter space will grow exponentially with the increasing joint numbers. Conversely, the RL method shows its effectiveness in its ability to solve this kind of complex control problem as it is trained directly in the dynamic environment. It is able to generate undulation gaits that not only imitate the movement of real snakes but also explore its limitations while continuing to improve its behavior under its hardware constraints.

Prototype Experiments

Sim-to-Real Transfer Framework

To make it feasible to transfer a learned policy from simulation to the real world, we should design a framework to easily connect the NN-based controller to the simulation environment or the real world with the same configuration. As shown in Figure 6, the NN-based controller has an interface used to switch between the simulation and real-world environments. For running both the simulation and real-world environments, the compact observations are the joint positions $\hat{\phi}$. The compact observations are first fed into the informative network to predict the full observations \mathbf{o}_t^f . Then, the NN-based controller outputs the actions to be executed either in the simulation or the real world. For real-world experiments, the NN-based controller is computed on a desktop computer and the actions are sent to an Arduino DUE via Robotic Operating System (ROS) messages. Then, the Arduino DUE serves as a host to transfer the ROS messages into CAN messages and sends them to the STM32 controller on each module via a CAN bus.

Experiment Setup

The power consumption can be calculated using its running current and voltage. In our experiments, we use a 24-V constant voltage power supply to actuate the robot. The current can be acquired by measuring the voltage of a high-load resistor (50 W/0.1 Ω), which is stringed into the circuit. Thus, power consumption of the prototype snake-like robot can be calculated as $P = (U_p/R_p)(24 - U_p)$, where $R_p = 0.1 \Omega$ and $U_p = 24 \text{ V}$. The voltage U_p is measured using an oscilloscope, which samples at 1,000 Hz.

Results

Due to time and hardware durability, we tested only five target velocities for the NN-based controller in our prototype experiments, namely, 0.05, 0.10, 0.15, 0.2, and 0.25 m/s. We controlled the robot to run for 10 s to measure power consumption and actual velocity. The actual velocities for the gaits generated by the NN-based controller were 0.03, 0.05, 0.08, 0.12, 0.15 m/s, approximately, which were all less than the target velocities. This is because of the reality gap between the simulated environment and the real world. For the same reason, the performance of the gait equation controller is also impacted by this reality gap. To fairly compare energy efficiency, we chose parameters that lead to gaits traveling at these five actual velocities with the best energy efficiency in the simulation.

First, we present montages of the slithering gaits at 0.15 m/s from the gait equation controller and the NN-based controller (see Figure 7). To show a full cycle of the movement, seven video frames (25 frames per second) are selected to visualize the pattern of the gaits generated by both controllers. For the gait generated by the gait equation controller, frames 89–119 show one full cycle of the gait. For the gait learned by the NN-based controller, roughly two motion cycles are

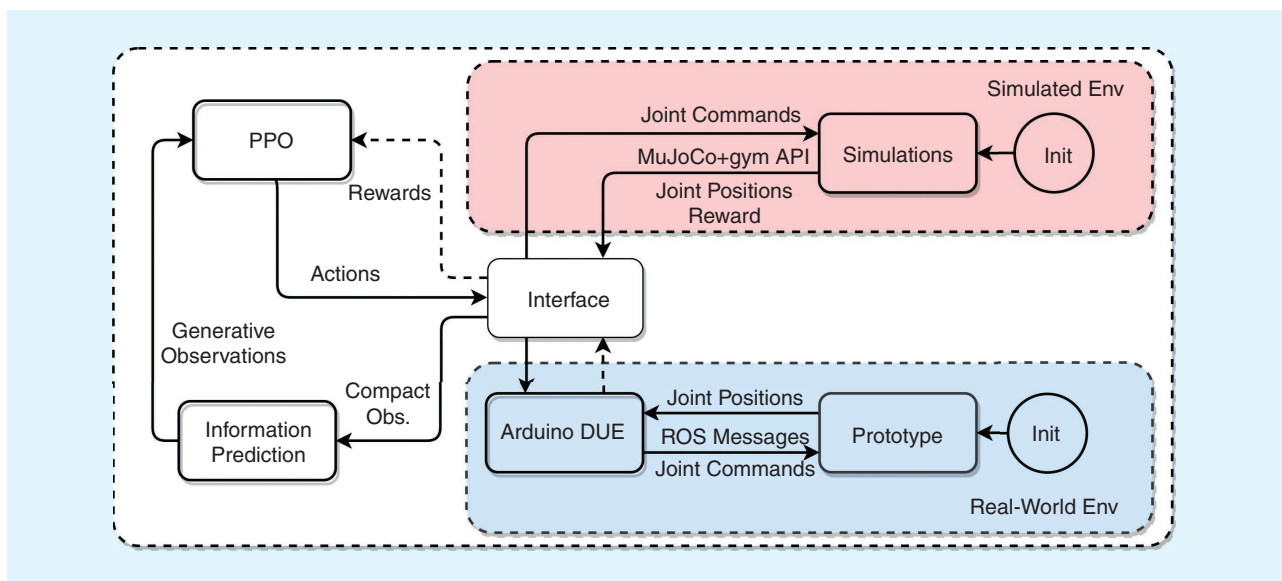


Figure 6. The software framework for transferring learned policy from simulations to the real-world controller. The left part of the figure visualizes how the NN-based controller is trained, and the right part illustrates the concept of sim-to-real transfer. Env.: environment; Obs.: observations; ROS: Robotic Operating System; API: application programming interface.

shown. As we can see, the body curve of the gait equation gait is shaped as a sinusoid wave. The body curve of the RL gait is slightly more slender than the body curve of the parameterized gait. This slider-body curve can propel the robot more smoothly and avoid joint power consumption as much as possible, which is in accordance with the simulation results. We can also observe that the learned gait moves at a higher frequency than the parameterized gait and moves straight toward the forward direction, while the parameterized gait tends to deviate from the forward direction.

Second, we show results of the power consumption of the gaits from both the gait equation controller and the NN-based controller, which are presented in Figure 8(a). All of the power consumption data were measured in 10 s for fair comparison. The solid line represents the exponentially weighted moving average of the power consumption, and the error bar depicts standard deviation of the raw power consumption measured by the oscilloscope. Note that we first process raw data by the removing outliers that are outside the range of the five standard deviations. Note that the outliers are caused by the noisy measurement of the current signal. From the results of the gait equation controller, we can observe that the power consumption increases with movement velocity, which ranges from approximately 21.5 to 28 W. As the velocity of the gait generated from the gait equation controller is directly related to the frequency of the sinusoid wave, it moves much slower at a low speed and faster at

a high speed, thereby consuming more power. We can see that the NN-based controller consumes more power at a low speed. This is because the controller generates a concertina gait, which only moves parts of the body to inch forward and those static joints require the locked-motor current, thereby consuming more power. Another potential reason is that the NN-based controller has difficulties maintaining low and accurate velocities. The small adjustments of each joint will impact the overall velocity observably and thus consumes more energy. Compared to the power consumption of the gait equation controller, we find that the NN-based controller can increase energy efficiency by 5 to 10%, especially at intermediate and high velocities. For both controllers at one velocity, power consumption shows an obvious fluctuation, which can be interpreted as the periodic movement of the robot. Figure 8(b) depicts the power-profile comparison of the simulation and real-world results. We can find that there is still a performance gap between the simulation and

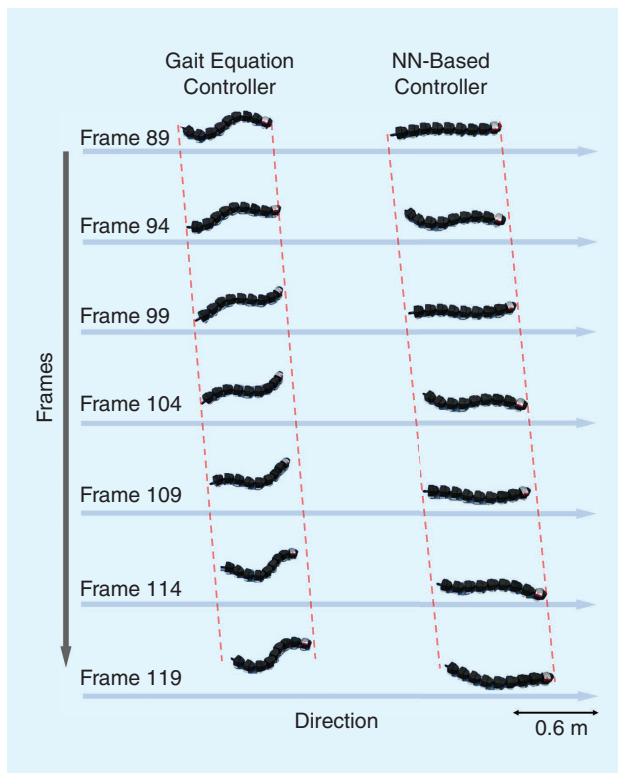
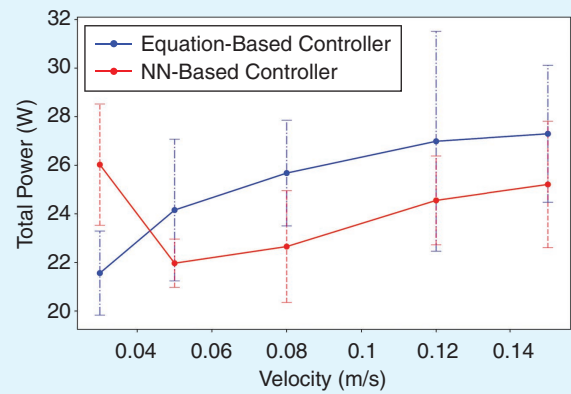
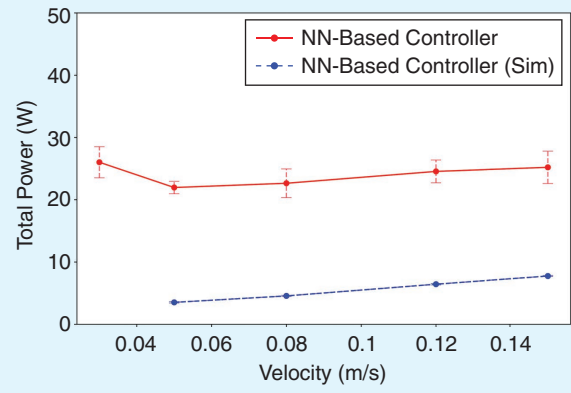


Figure 7. Montages of the snake-like robot performing slithering gait at 0.15 m/s, controlled with the NN-based controller and the gait equation controller. The video was recorded at 25 frames per second.



(a)



(b)

Figure 8. (a) A real-world power comparison of slithering gaits from the gait equation and NN-based controllers. The target velocities are [0.05, 0.10, 0.15, 0.2, 0.25] m/s. The actual velocities are approximately [0.03, 0.05, 0.08, 0.12, 0.15] m/s. The solid points represent the averaged power over 10 s, and the error bars depict the standard deviation. Note that the standard deviation is calculated based on the data over the duration of the experiment (10 s). (b) Power comparison of the NN-based controller in simulation (Sim) and the real world. In simulation, the NN-based controller cannot achieve a gait of 0.03 m/s.

real-world experiment due to some potential reasons. For instance, it is challenging to simulate the environment with perfect fidelity, which therefore leads to the performance drop. Inaccurate sensory measurements or noisy observation stats can also cause the performance drop.

Conclusion

Designing energy-efficient gaits for snake-like robots remains a challenging task as they come with redundant DoF and have complicated interactions with the environment. The most widely used method for generating gaits for snake-like robots is the gait equation, which mimics the snake's body shape using sinusoid-like curves. In this article, we presented a novel gait-design method based on RL. Compared with the optimized gaits generated by the gait equation, the learned gaits show better energy efficiencies at medium and high velocities. Although the learned gaits perform similarly or slightly worse than the parameterized gaits at low speeds, the proposed method shows great efficiency in discovering energy-efficient gaits automatically and maximizing performance by removing the limitations of predefined models. Our work contributes to and serves as an exploration for designing sophisticated moving patterns for snake-like robots. Future work will aim at designing gaits using RL for snake-like robots without passive wheels.

Acknowledgment

This project/research has received funding from the European Union's Horizon 2020 Framework Program for Research and Innovation under Specific Grant Agreement (SGA) Number 945539 (Human Brain Project SGA3). This work was also supported by the National Natural Science Foundation of China (grant 61902442), Pazhou Lab PZL2021KF0020 and Shenzhen Basic Research Grant JCYJ20180508152434975. The corresponding author is Long Cheng.

References

[1] M. Tesch et al., "Parameterized and scripted gaits for modular snake robots," *Adv. Robot.*, vol. 23, no. 9, pp. 1131–1158, 2009, doi: 10.1163/156855309X452566.

[2] S. Hirose, *Biologically Inspired Robots: Snake-Like Locomotors and Manipulators*, vol. 1093. London, U.K.: Oxford Univ. Press, 1993.

[3] Z. Bing, L. Cheng, G. Chen, F. Röhrbein, K. Huang, and A. Knoll, "Towards autonomous locomotion: CPG-based control of smooth 3d slithering gait transition of a snake-like robot," *Bioinspiration Biomimetics*, vol. 12, no. 3, p. 35,001, Apr. 2017, doi: 10.1088/1748-3190/aa644c.

[4] Z. Bing, L. Cheng, K. Huang, M. Zhou, and A. Knoll, "CPG-based control of smooth transition for body shape and locomotion speed of a snake-like robot," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pp. 4146–4153, doi: 10.1109/ICRA.2017.7989476.

[5] P. Liljebäck, K. Y. Pettersen, Ø. Stavadahl, and J. T. Gravdahl, *Snake Robots: Modelling, Mechatronics, and Control*. London: Springer Science & Business Media, 2012.

[6] Z. Bing, C. Lemke, Z. Jiang, K. Huang, and A. Knoll, "Energy-efficient slithering gait exploration for a snake-like robot based on reinforcement learning," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Macao,

China: International Joint Conferences on Artificial Intelligence Organization, Aug. 2019, pp. 1–7.

[7] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Brisbane, QLD, Australia, May 2018, pp. 3803–3810, doi: 10.1109/ICRA.2018.8460528.

[8] A. Crespi and A. J. Ijspeert, "Online optimization of swimming and crawling in an amphibious snake robot," *IEEE Trans. Robot.*, vol. 24, no. 1, pp. 75–87, Feb. 2008, doi: 10.1109/TRO.2008.915426.

[9] M. Tesch, J. Schneider, and H. Choset, "Using response surfaces and expected improvement to optimize snake robot gait parameters," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Sep. 2011, pp. 1069–1074, doi: 10.1109/IROS.2011.6095076.

[10] S. Ouyang and W. Wei, "Flexible adaptive control of snake-like robot based on LSTM and gait," *J. Phys., Conf. Ser.*, vol. 1487, no. 1, p. 012049, 2020, doi: 10.1088/1742-6596/1487/1/012049.

[11] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 7559–7566, doi: 10.1109/ICRA.2018.8463189.

[12] Y. Nakamura, T. Mori, M-a Sato, and S. Ishii, "Reinforcement learning for a biped robot based on a CPG-actor-critic method," *Neural Netw.*, vol. 20, no. 6, pp. 723–735, 2007, doi: 10.1016/j.neunet.2007.01.002.

[13] J. Shi, T. Dear, and S. D. Kelly, "Deep reinforcement learning for snake robot locomotion," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 9688–9695, 2020, doi: 10.1016/j.ifacol.2020.12.2619.

[14] S. Fukunaga, Y. Nakamura, K. Aso, and S. Ishii, "Reinforcement learning for a snake-like robot controlled by a central pattern generator," in *Proc. IEEE Conf. Robot., Autom. Mechatronics*, 2004, vol. 2, pp. 909–914, doi: 10.1109/RAMECH.2004.1438039.

[15] X. Liu, R. Gasoto, Z. Jiang, C. Onal, and J. Fu, "Learning to locomote with artificial neural-network and CPG-based control in a soft snake robot," in *Proc. 2020 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, pp. 7758–7765, doi: 10.1109/IROS45743.2020.9340763.

[16] S. Chatterjee et al., "Reinforcement learning approach to generate goal-directed locomotion of a snake-like robot with screw-drive units," in *Proc. 2014 23rd Int. Conf. Robot. Alpe-Adria-Danube Region (RAAD)*, pp. 1–7, doi: 10.1109/RAAD.2014.7002234.

[17] G. Sartoretti, Y. Shi, W. Paivine, M. Travers, and H. Choset, "Distributed learning for the decentralized control of articulated mobile robots," in *Proc. 2018 IEEE Int. Conf. Robot. Autom. (ICRA)*, pp. 3789–3794, doi: 10.1109/ICRA.2018.8460802.

[18] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, "Learning to walk via deep reinforcement learning," in *Proc. Robot., Sci. Syst.*, Freiburg, Jun. 2019, doi: 10.15607/RSS.2019.XV.011.

[19] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Sci. Robot.*, vol. 7, no. 62, p. eabk2822, 2022, doi: 10.1126/scirobotics.abk2822.

[20] P. Christiano et al., "Transfer from Simulation to real world through learning deep inverse dynamics model," Oct. 2016, *arXiv:1610.03518 [cs]*.

[21] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *Proc. 2018 IEEE Int. Conf. Robot. Autom. (ICRA)*, pp. 3803–3810, doi: 10.1109/ICRA.2018.8460528.

[22] M. Wulfmeier, I. Posner, and P. Abbeel, "Mutual alignment transfer learning," in *Proc. Conf. Robot Learn.*, 2017, pp. 281–290.

[23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.

[24] J. Hwangbo et al., "Learning agile and dynamic motor skills for legged robots," *Sci. Robot.*, vol. 4, no. 26, p. eaau5872, Jan. 2019, doi: 10.1126/scirobotics.aau5872.

Zhenshan Bing, Department of Informatics, Technical University of Munich, Munich, 85748, Germany. Email: bing@in.tum.de.

Long Cheng, College of Computer Science and Artificial Intelligence, Wenzhou University, Wenzhou, 325035, China, and Shenzhen Institute, Sun Yat-sen University, Guangzhou 543001, China. Email: chenglong@wzu.edu.cn.

Kai Huang, School of Data and Computer Science, Sun Yat-sen University, Guangzhou 543001, China, and Pazhou Lab, Guangzhou, 543001, China. Email: huangk36@mail.sysu.edu.cn.

Alois Knoll, Department of Informatics, Technical University of Munich, Munich 85748, Germany. Email knoll@in.tum.de.

