

COMPOSER: Scalable and Robust Modular Policies for Snake Robots

Yuyou Zhang¹, Yaru Niu¹, Xingyu Liu¹ and Ding Zhao¹

Abstract—Snake robots have showcased remarkable compliance and adaptability in their interaction with environments, mirroring the traits of their natural counterparts. While their hyper-redundant and high-dimensional characteristics add to this adaptability, they also pose great challenges to robot control. Instead of perceiving the hyper-redundancy and flexibility of snake robots as mere challenges, there lies an unexplored potential in leveraging these traits to enhance robustness and generalizability at the control policy level. We seek to develop a control policy that effectively breaks down the high dimensionality of snake robots while harnessing their redundancy. In this work, we consider the snake robot as a modular robot and formulate the control of the snake robot as a cooperative Multi-Agent Reinforcement Learning (MARL) problem. Each segment of the snake robot functions as an individual agent. Specifically, we incorporate a self-attention mechanism to enhance the cooperative behavior between agents. A high-level imagination policy is proposed to provide additional rewards to guide the low-level control policy. We validate the proposed method COMPOSER with five snake robot tasks, including goal reaching, wall climbing, shape formation, tube crossing, and block pushing. COMPOSER achieves the highest success rate across all tasks when compared to a centralized baseline and four modular policy baselines. Additionally, we show enhanced robustness against module corruption and significantly superior zero-shot generalizability in our proposed method. The videos of this work are available on our project page: <https://sites.google.com/view/composer-snake/>.

I. INTRODUCTION

Snake robots have been extensively studied over decades. Their flexible characteristics and bioinspired structure have enabled compliant interaction with the environment, facilitating their usage across areas including medical applications [1], [2], extreme [3], [4] or confined [5], [6] environment exploration. The natural advantage of snake robots to work in unstructured environments is based on their material properties and morphology of their bodies [1], [7]–[9] rather than control strategy. In this work, we aim to improve the performance of snake robots by introducing a control strategy that leverages the inherent modularity of continuum robots. This modularity can be viewed from three perspectives: high dimensionality, scalability, and redundancy.

Snake robots, as continuum robots, possess a high-dimensional Degree of Freedom (DOF) and fall between the realms of soft robots and rigid-body robots. This high-dimensional DOF complicates the system’s kinematics and dynamics, posing a significant challenge to robot control. Model-based control methods [10]–[14] for snake robots

¹Yuyou Zhang, Yaru Niu, Xingyu Liu and Ding Zhao are with the Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA. {yuyouz, yarun, xingyul3, dingzhao}@andrew.cmu.edu

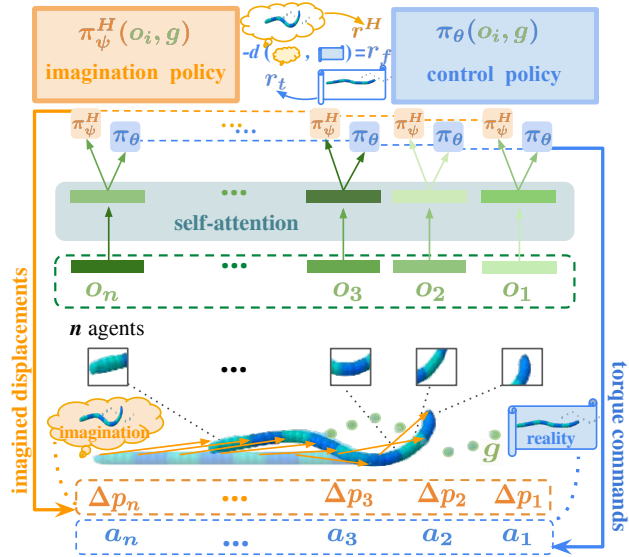


Fig. 1. Framework overview of COMPOSER. A snake robot with n joints is formulated as n agents. The modular control policy outputs individual torque commands, while the imagination policy forecasts an ideal displacement per step. The control policy is trained to both complete the task and adhere to the direction prescribed by the imagination policy.

suffer from this high dimensionality and usually struggle to approximate the dynamics model of the continuum body. Though reinforcement learning (RL) has shown remarkable model-free decision-making capabilities, learning an optimal policy in the presence of a high-dimensional space in snake robots can be costly [15]. Training a centralized policy, which requires exploring a high-dimensional space that grows exponentially with the snake robot DOF, can be highly inefficient.

Snake robots have inherent modularity as continuum robots. This modularity enables scalability and flexibility, simplifying the snake robot design and facilitating easier maintenance. Structural scalability and flexibility empower snake robots to adapt to various tasks and unstructured environments [6]. Current monolithic reinforcement learning [16] [17] processes joint observations as inputs and produces joint actions as outputs. This centralized policy lacks generalizability to match the scalability of the snake robot.

Modularity also gives snake robots hardware redundancy compared with traditional robots, which adds to system-level robustness against hardware failures, and allows for emergent behavior to adapt to perturbations [18]–[20]. However, the presence of hardware redundancy becomes futile without an effective control policy to achieve policy-level robustness in snake robots.

To this end, our work is developed to break down the high dimensionality of snake robots while harnessing the

structural scalability and redundancy brought by the inherent modularity. We first form the snake robot control problem as a multi-agent reinforcement learning problem and propose a shared modular torque control policy. Then we propose a high-level imagination policy to provide additional reward for control policy training in long-horizon tasks. We further utilize a self-attention mechanism to enhance inter-module communication. Five snake robot tasks are designed, showcasing manipulation and locomotion skills. Extensive experiments are conducted on five tasks to demonstrate the effectiveness of the proposed modular policy COMPOSER, in terms of success rate in normal functionality, robustness against hardware malfunction, and generalizability. To the best of our knowledge, this is the first work that investigates the robustness and generalizability of a modular torque control policy for snake robots. The overview of our proposed method COMPOSER is shown in Fig. 1.

The contributions are summarized as follows.

- We propose the first communication-enhanced decentralized modular control policy for snake robots to accommodate their inherent structural modularity.
- We introduce an imagination policy for efficient morphology-aware planning in long-horizon tasks.
- Our proposed method, COMPOSER, has demonstrated superior performance compared to strong baselines in all five snake robot tasks and showed better robustness and zero-shot generalizability.

II. RELATED WORK

A. Snake and Continuum Robot Control

Early works of snake robots model-based control mainly focus on reproducing the innate behavior of real snakes. Various approaches have been explored, including generating torque proportional to the curvature derivative of the body curve [10], applying a serpenoid traveling wave propagated from the head to the tail [11], utilizing gaits for shape based control [12], [13]. However, simplifications made in these model-based approaches constrain the versatility of the modular snake robot, hurting compliance and flexible locomotion ability. Many recent works in snake robots control have exhibited reduced reliance on the model, including co-optimization of morphology and control [21], locomote with neural-network and CPG-based control [22], coach-based reinforcement learning [17], reinforcement learning for snake robot tracking control [23]. However, [17], [21], [23] use wheels to simplify snake robots and focus on in-plane control, leaving a gap in the exploration of complex tasks, such as manipulation and 3D locomotion in unstructured environments. Many prior studies have not fully leveraged the inherent modularity of snake robots. While Sartoretti et al. first explore snake robot modularity through the decentralized learning of shape-based locomotion parameters, its applicability is limited by reliance on a shape-based controller [24]. In contrast, our approach in this paper directly learns a control policy to generate actuator torque, which enables learning intricate manipulation and locomotion skills.

B. Modular Policy for Decentralized Control

Decentralized control is common in biological motor control in living creatures. For example, a majority of neurons in the octopus are found in the arms, which can independently control basic motions without input from the brain, allowing fast responses. Using modular policy for decentralized control has demonstrated efficiency, robustness, and generalizability [25]–[28]. Schilling et al. illustrate that within a decentralized control paradigm, reinforcement learning-based motor control not only attains high performance but also showcases enhanced robustness and improved generalizability [25]. Pigozzi et al. employ evolutionary computation to optimize a shared neural controller for Voxel-based Soft Robots [26]. Whitman et al. use modular policy for reconfigurable robots [27]. In [28], a shared modular policy is trained with messages passing. [12] shows that fewer individual degrees of freedom coupled means more efficient locomotion on irregular terrains. However, decentralized control has primarily been applied to snake robots within model-based control [10], [12], [24]. As far as we know, this is the first work investigating learning a modular torque control policy to exploit the modularity of snake robots.

C. Multi-agent Reinforcement Learning

On one hand, many works in MARL follow the Centralized Training Decentralized Execution (CTDE) framework, where decentralized policies are trained in a centralized fashion with extra information. MADDPG adopts actor-critic structures and learns a centralized critic [29]. Value-decomposition (VD) methods decompose the joint Q-function as a function of agents' local Q-functions [30], [31]. MAPPO demonstrates supreme performance in homogeneous multi-agent cooperative settings [32]. HATPRO and HAPPO achieve promising results in heterogeneous settings, without parameter sharing among agents [33]. On the other hand, some works employ centralized communication protocols during execution to share local information and augment the coordination among agents [34]–[36]. In this paper, we aim to utilize MARL to develop a modular policy tailored specifically for snake robots and continuum robots.

III. METHODOLOGY

We first introduce the Goal-conditioned CTDE paradigm for snake robot control as described in Section III-A. In Section III-B, we elaborate on training a modular torque control policy with additional reward from an imagination policy. Section III-C introduces a self-attention mechanism to enhance efficient cooperation between agents. The full algorithm including the joint training of the modular torque control policy and the imagination policy is summarized in Algorithm 1.

A. Multi-agent Reinforcement Learning for Modular Policy

In this paper, we formulate snake robot control as a fully cooperative MARL problem. Each actuator of the snake robot is considered an individual agent. We consider a multi-agent goal-conditioned partially observable Markov decision

process for a snake robot with n actuators, defined by $\langle \mathcal{S}, \mathcal{A}, O, R, P, \mathcal{G}, n, \gamma \rangle$. \mathcal{S} is the state space. $s \in \mathcal{S}$ describes states of all n agents. Given the same configuration of all actuators with the snake robot, each agent shares an identical observation space and action space. Local observation for agent i with s is $o_i = O(s; i)$. At each time step, agent i executes action a_i sampled from a distribution generated by policy $\pi_\theta(\cdot|o_i, g)$, which is parameterized by θ and shared by all agents. And all n agents receive shared reward $r_1, \dots, r_n = R(s, \mathbf{a}, g)$. $R(s, \mathbf{a}, g)$ denotes the reward function and $g \in \mathcal{G}$ is a goal shared among all agents. $\mathbf{a} = (a_1, a_2, \dots, a_n) \in \mathcal{A}$ is the joint action of n agents, which produces the next state s' with transition probability $P(s'|s, \mathbf{a})$.

We aim to learn a shared modular torque control policy π_θ for the snake robot. Each agent makes decisions individually while all agents execute actions simultaneously. The coupling between agents of the snake robot requires precise coordination among all cooperative agents to successfully complete tasks. The modular torque control policy π_θ is trained to jointly maximize the expected return $J(\theta) = \mathbb{E}_{\mathbf{a}_t, s_t} [\sum_t \gamma^t R(s_t, \mathbf{a}_t, g)]$, with the discount factor γ . π_θ is optimized with the experiences of all agents simultaneously.

B. Imagination-guided Control

In addition to using a shared torque control policy to output torque commands for each agent, we propose to use a shared modular imagination policy, π_ψ^H , parameterized by ψ to predict per-step displacement for each agent. We use this imagination policy as an imaginary planner to decompose the long-horizon task into manageable steps and provide supplementary rewards for the torque control policy.

The imagination policy $\pi_\theta(\cdot|o_i, g)$ is conditioned on the local observation o_i and goal, and it generates distribution over anticipated displacement $\Delta p_i \in \mathbb{R}^3$. The displacement vector Δp_i points from agent k 's current position $p_i \in \mathbb{R}^3$ to an imagined next-step position $p'_i \in \mathbb{R}^3$, $p'_i = p_i + \Delta p_i$. $\Delta \mathbf{p}$ denotes the predicted snake whole body shape transition. Agents receive rewards by making good predictions so that the imagined next-step positions p'_i can be a reasonable sub-goal which decomposes the long-horizon task. The reward function for imagination policy is defined to capture the proximity of the imagined next-step positions to the goal completion state, $R^H(\mathbf{p}', g) = d(\mathbf{p}', g)$, $\mathbf{p}' = \{p'_i\}_{i=1}^n$ is the joint next-step positions predicted by n agents. For example, in the shape formation task of snake robots, we use the Wasserstein distance [37] between the imagined next-step position \mathbf{p}' and goal positions g as the inverse of the reward for imagination policy, $W(\mathbf{p}', g) = \inf_{\pi \in \mathcal{T}(\mathbf{p}', g)} \int_{\mathbb{R}^3 \times \mathbb{R}^3} |x - y| d\pi(x, y)$. The goal of the imagination policy is to maximize the accumulated reward $J(\psi) = \mathbb{E}_{\mathbf{p}'_t, s_t} [\sum_t \gamma^t R^H(\mathbf{p}'_t, g)]$.

The modular torque control policy receives rewards by both following the imagined displacement as prescribed by the imagination policy and completing the task. Agent i executes action $a_{i,t}$ induced by the torque control policy $\pi_\theta(\cdot|o_{i,t}, g)$, and reaches position $p_{i,t+1} \in \mathbb{R}^3$. The ‘‘follow-

the-imagination’’ reward r_f is defined as the sum of errors between actual position $p_{i,t+1}$ agent i reaches and the imagined next-step position $p'_{i,t}$.

$$r_f = \sum_{i=1}^n |p_{i,t+1} - p'_{i,t}| \quad (1)$$

The reward function for the torque control policy can be described as $R = r_t + r_f$, where r_t is a task-specific reward as shown in Table I in Section IV-A.

The shared imagination policy and the shared torque control policy are jointly trained following the framework of Multi-Agent PPO (MAPPO) [32]. For each policy, we train two separate neural networks, the actor network and the critic network, $\pi_\theta(\cdot|o_i, g)$ and $V_\phi(s)$ for the torque control policy, and $\pi_\psi^H(\cdot|o_i, g)$ and $V_\omega^H(s)$ for the imagination policy, respectively. The critic networks, V_ϕ and $V_\omega^H(s)$, map the joint state to value: $S \rightarrow \mathbb{R}$.

The actor network for the torque control policy π_θ is trained to maximize the objective

$$L(\theta) = \frac{1}{Bn} \sum_{k=1}^B \sum_{i=1}^n \min(r_{\theta,i}^{(k)} A_i^{(k)}, \text{clip}(r_{\theta,i}^{(k)}, 1 - \epsilon, 1 + \epsilon) A_i^{(k)}) \\ + \sigma \frac{1}{Bn} \sum_{i=1}^n \sum_{k=1}^n H[\pi_\theta(o_i^{(k)})]$$

where $r_{\theta,i}^{(k)} = \frac{\pi_\theta(a_i^{(k)}|o_i^{(k)})}{\pi_{\theta_{old}}(a_i^{(k)}|o_i^{(k)})}$, $A_i^{(k)}$ is the estimated advantage computed using the Generalizable Advantage Estimation (GAE) [38], H is the policy entropy, and σ is the entropy coefficient hyperparameter. The critic network V_ϕ is trained to minimize the loss function

$$L(\phi) = \frac{1}{B} \sum_{k=1}^B \max[(\text{clip}(V_\phi(s^{(k)}), V_{\phi_{old}}(s^{(k)}) - \epsilon, \\ V_{\phi_{old}}(s^{(k)}) + \epsilon) - \hat{R})^2, (V_\phi(s^{(k)}) - \hat{R})^2]$$

where \hat{R} is the discounted reward-to-go, B refers to the batch size, and n refers to the number of agents. The actor network and critic network for the imagination policy, π_ψ^H and V_ω^H are optimized similarly.

C. Self-attention Mechanism

The coordinated behavior of a snake robot is heavily dependent on the cooperation of its highly coupled modules. We incorporate the self-attention [39] mechanism into the multi-agent training of the shared control policy, aiming to provide global information for each agent by augmenting inter-module communication.

The self-attention encoder takes joint observation sequence $\mathbf{o} = \{o_i\}_{i=1}^n \in \mathbb{R}^{n \times d}$ as input, where n is the sequence length, and d is the observation dimension. The self-attention encoder maps \mathbf{o} to new latent space observation sequence of equal length n as shown in Fig. 1. Attention is calculated as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (2)$$

Algorithm 1 Training of the control and imagination policy

```

1: Notation: number of agents  $n$ , number of episodes  $N_e$ ,
   episode length  $T$ 


---


2: for  $e$  in 1 to  $N_e$  do
3:   data buffer  $D \leftarrow \{\}$ 
4:   for  $t$  in 1 to  $T$  do
5:     for  $i$  in 1 to  $n$  do
6:       sample  $a_{i,t} \sim \pi(o_{i,t}, g; \theta)$ 
7:       sample  $\Delta p_{i,t} \sim \pi^H(o_{i,t}, g; \theta)$ 
8:        $v_{i,t} \leftarrow V(s_t, g; \phi)$ 
9:        $v_{i,t}^H \leftarrow V^H(s_t, g; \phi)$ 
10:      Execute  $\mathbf{a}_t, s_{t+1} \sim P(\cdot|s_t, \mathbf{a}_t)$ , predict  $\Delta \mathbf{p}_t$ ,
        receive rewards  $r_t, r_t^H$ , observe  $\mathbf{o}_{t+1}$  from  $s_{t+1}$ 
11:       $D \leftarrow D \cup \{(s_t, \mathbf{o}_t, \mathbf{a}_t, \Delta \mathbf{p}_t, r_t, r_t^H, s_{t+1}, \mathbf{o}_{t+1})\}$ 
12:      Compute advantage  $\hat{A}$  and reward-to-go  $\hat{R}$  on  $D$ 
13:      Adam update  $\theta, \phi, \psi, \omega$  with  $D$ 

```

Where Q, K, V are query, key, and value, $Q = XW_q \in \mathbb{R}^{n \times r}$, $K = XW_k \in \mathbb{R}^{n \times r}$, $V = XW_v \in \mathbb{R}^{n \times f}$, $W_q, W_k \in \mathbb{R}^{d \times r}$, $W_v \in \mathbb{R}^{d \times f}$ are linear transformations. By calculating the attention, the new local observation for agent i incorporates global information in the form of the weighted sum of n transformed local observations of n agents.

IV. EXPERIMENTS

The design of our COMPOSER framework is motivated by three hypotheses. First, compared to centralized policies, the modular policy employed by COMPOSER can improve learning efficiency. Second, the modular policy used by COMPOSER provides policy-level robustness against individual actuator malfunction. Third, by using modular policy, COMPOSER generalizes better to snake robots with different numbers of joints. To show this, we conducted extensive experiments across five snake robot tasks, evaluating success rates, robustness, and zero-shot generalizability to validate our hypotheses. We further illustrate the contributions of the imagination policy and the self-attention mechanism through ablation studies and visualization.

A. Experiment Setup

Environments. We introduce five snake robot tasks, Goal Reaching, Block Pushing, Shape Formation, Tube Crossing, and Wall Climbing as shown in Fig. 2. Our simulation environment is adapted from SomoGym [40]. The reward is calculated as the sum of instant rewards across an episode for each environment. Once the robot completes the task, which is determined by the task-specific success criterion as shown in Table I, or the maximum episode length is reached, the episode terminates.

The action space contains normalized torques for each actively controlled degree of freedom (DOF) of each actuator. The action a_i is represented by two scalar $[a_1, a_2]$ normalized to $[-1, 1]$, corresponding to two independently actuated bending axes in one actuator. In the multi-agent setting, each agent takes action only based on its own




Fig. 2. Snapshots of COMPOSER across five snake robot tasks.

TABLE I
TASK REWARD r_t AND SUCCESS CRITERION

Task	Reward r_t	Success criterion
Goal Reaching	$-d(\text{head}, \text{goal})$	$d(\text{head}, \text{goal}) < d_g$
Block Pushing	$-d(\text{head}, \text{object}) - d(\text{object}, \text{goal})$	$d(\text{object}, \text{goal}) < d_g$
Shape Formation	$-\sum_i^n d(\text{agent}_i, \text{goal}_i)$	$\sum_i^n d(\text{agent}_i, \text{goal}_i) < d_g$
Tube Crossing	$-d(\text{head}, \text{tube}) - d(\text{head}, \text{goal})$	$d(\text{head}, \text{goal}) < d_g$
Wall Climbing	head position	snake tail over the wall

observation including local positions, velocities, orientations, and applied torques. For monolithic baseline training, the control policy takes global observations as input and outputs the action including commands for n agents.

Baselines. We compare our method with five baselines to investigate the effectiveness of COMPOSER. We apply the same hyper-parameters of baseline algorithms from their original papers. PPO is a strong on-policy RL method. We use PPO with zero padding as a monolithic trained baseline to learn a centralized policy. Following the setup in [16], we zero-pad the states and actions to the maximum dimension across all agents with different lengths. MAT is an on-policy few-shot learner on unseen tasks regardless of changes in the number of agents [41]. MAT_dec is a CTDE-variant of MAT [41]. SMP is a modular policy to generate locomotion behaviors for different skeletal structures via message passing [28]. MAPPO is a strong multi-agent reinforcement learning baseline with competitive sample-efficiency [32].

B. Results Comparison and Analysis

We evaluate our COMPOSER framework in the following three aspects: learning efficiency in terms of task success rate, robustness against hardware failure, and zero-shot generalization to previously unseen robots.

Success Rate. In Fig. 3, we illustrate the performance of our proposed method and the baselines across five snake robot tasks. The results are averaged over three seeds, and the standard errors are provided. In almost all tasks, COMPOSER outperforms monolithic trained zero-padding PPO and other multi-agent RL baselines including MAPPO, MAT, MAT_dec, and SMP. Monolithic trained PPO shows clear defects indicating the difficulties inherent in exploring a high-dimensional action space for snake robots. This observation validates the effectiveness of modular policy to tackle this high-dimensional challenge with decentralized control. For modular policy baselines, MAPPO is the strongest one while our method still outperforms it by a large margin,




Fig. 3. Learning curves of different methods in the tasks of (a) *Goal Reaching*, (b) *Block Pushing*, (c) *Wall Climbing*, (d) *Shape Formation*, and (e) *Tunnel Crossing*. The evaluation rewards under desired goals are averaged over three seeds, and the shaded region represents the standard error.

TABLE II

EVALUATIONS OF ROBUSTNESS AGAINST HARDWARE FAILURES WITH DIFFERENT CORRUPTIONS IN EIGHT-AGENT SNAKE *Goal Reaching* TASKS

Method	0 fault			1 fault action			2 fault action			1 fault observation		
	Success Rate \uparrow	Distance \downarrow	Step \downarrow	Success Rate \uparrow	Distance \downarrow	Step \downarrow	Success Rate \uparrow	Distance \downarrow	Step \downarrow	Success Rate \uparrow	Distance \downarrow	Step \downarrow
PPO	0.51 \pm 0.15	3.14 \pm 0.37	102.42 \pm 6.74	0.28 \pm 0.084	5.14 \pm 0.61	111.62 \pm 3.77	0.17 \pm 0.06	6.50 \pm 0.77	116.67 \pm 2.31	0.35 \pm 0.085	4.66 \pm 0.099	108.94 \pm 2.97
MAPPO	0.76 \pm 0.08	2.35 \pm 0.09	93.01 \pm 3.45	0.52 \pm 0.093	3.76 \pm 0.54	103.13 \pm 3.81	0.28 \pm 0.07	5.30 \pm 0.79	112.36 \pm 3.98	0.44 \pm 0.14	5.20 \pm 1.26	106.29 \pm 5.03
COMPOSER (ours)	0.90 \pm 0.028	2.08 \pm 0.059	82.78 \pm 1.94	0.55 \pm 0.073	3.67 \pm 0.35	100.07 \pm 1.07	0.33 \pm 0.058	5.19 \pm 0.38	110.62 \pm 1.35	0.39 \pm 0.033	4.61 \pm 0.37	105.07 \pm 1.83

in terms of lower success rates and sample efficiency, showing the effectiveness of the self-attention mechanism and imagination policy. MAT features a complex model structure with both an encoder and a decoder, resulting in slow learning growth compared to COMPOSER. MAT_dec, being fully decentralized without the encoder for global observation embedding and the decoder that grants agents access to preceding agents’ actions in MAT, significantly underperformed in comparison to our proposed method. SMP does not manage to learn a successful control policy across all five tasks, which aligns with results reported in [40], that TD3 has also shown no relevant improvement in nearly all training.

Robustness against Hardware Failures. In Table II, we demonstrate the performance of different policies trained on a normally functioning snake robot for *Goal Reaching* when tested in scenarios involving partial hardware failures. This partial hardware failure includes agent action corruption and agent observation corruption. In the case of action corruption, faulty agents are assigned actions with a value of zero. In the case of observation corruption, the observation of faulty agents is assigned a value of zero. At the beginning of each episode, faulty agents are randomly selected based on the specified faulty probability, $p = n_{\text{fault}}/n$, where n_{fault} is the number of corrupted agents, and n is the total agent number. Corrupted agents fail to follow the prescribed policy.

In Table II, we evaluate different policies on four scenarios, including zero corrupted agent, one corrupted action, two corrupted actions, and one corrupted observation. The corresponding model is evaluated on 100 episodes, and results include success rate, final distance to the goal, and episode length. The upper arrow means higher is better, and the down arrow means lower is better. The best performance of each task is marked in bold. We compare COMPOSER with monolithic trained baseline PPO and the strongest modular policy baseline MAPPO. COMPOSER demonstrates greater fault tolerance by having a shorter distance to the goal, fewer episode steps across all scenarios, and a higher success rate

across action corruption scenarios, when compared to PPO and MAPPO, as shown in Table II. In the case of one corrupted observation, MAPPO demonstrates a slightly higher success rate. This outcome could be attributed to the fact that the corruption of a single agent’s local observation might lead to confusion in the observations of other agents, given the presence of the attention mechanism. In contrast, MAPPO is less affected by the corruption of local observations due to the lack of inter-agent communication.

Zero-shot Generalizability to Snake Robots with Different Number of Joints. To demonstrate zero-shot generalizability, we train the policy on a snake robot with 8 agents and directly apply the trained policy on previously unseen snake robots with 9 agents, 10 agents, and 11 agents. As shown in Table III, our modular policy demonstrates notable generalizability in comparison to PPO and MAPPO, particularly when applied to a 10-agent snake robot, where PPO and MAPPO almost always fail, but our proposed method has a 33% success rate. This shows the potential of our method to be applied to a scalable modular snake robot.

C. Ablation Study: Attention and Imagination

As shown in Fig. 4, we evaluate the performance of COMPOSER alongside its variants: one devoid of the imagination policy (COMPOSER_w/oI) and another excluding the self-attention mechanism (COMPOSER_w/oA) across five tasks. While COMPOSER_w/oI and COMPOSER_w/oA achieved a success rate similar to COMPOSER in simpler tasks such as *Goal Reaching* and *Block Pushing*, they fell notably short in tasks of increased complexity. In more challenging tasks, such as *Shape formation* which requires desired deformation, and *Wall climbing* and *Shape Formation* which necessitate effective interaction with the environment, the ablations demonstrate markedly worse performance. COMPOSER enjoys a higher success rate by leveraging both the attention mechanism for enhanced coordination and the imagination policy for more efficient planning.

TABLE III

EVALUATIONS OF ZERO-SHOT GENERALIZATION TO SNAKE ROBOTS WITH DIFFERENT LENGTHS ON *Goal Reaching* TASK

Method	8 agent (32 links)			9 agent (36 links)			10 agent (40 links)			11 agent (44 links)		
	Success rate \uparrow	Distance \downarrow	Step \downarrow	Success rate \uparrow	Distance \downarrow	Step \downarrow	Success rate \uparrow	Distance \downarrow	Step \downarrow	Success rate \uparrow	Distance \downarrow	Step \downarrow
PPO	0.51 \pm 0.15	3.14 \pm 0.37	102.42 \pm 6.74	0.25 \pm 0.01	5.23 \pm 0.26	112.79 \pm 1.51	0.067 \pm 0.047	7.10 \pm 0.59	120.43 \pm 2.52	0.033 \pm 0.023	8.51 \pm 1.73	122.13 \pm 1.32
MAPPO	0.76 \pm 0.08	2.35 \pm 0.09	93.01 \pm 3.45	0.47 \pm 0.05	3.61 \pm 0.21	107.99 \pm 1.58	0.10 \pm 0.07	7.14 \pm 1.42	119.41 \pm 2.68	0.023 \pm 0.02	11.4 \pm 1.70	122.45 \pm 0.71
COMPOSER	0.90 \pm 0.028	2.08 \pm 0.059	82.78 \pm 1.94	0.57 \pm 0.20	2.83 \pm 0.40	102.24 \pm 4.28	0.35 \pm 0.18	4.20 \pm 1.49	111.25 \pm 6.00	0.13 \pm 0.09	7.26 \pm 2.72	118.87 \pm 3.15




Fig. 4. The comparison between COMPOSER and its ablations, COMPOSER without imagination policy (COMPOSER_w/oI) and COMPOSER without self-attention mechanism (COMPOSER_w/oA). From left to right, success rate learning curves are plotted for five tasks, (a) *Goal Reaching*, (b) *Block Pushing*, (c) *Wall Climbing*, (d) *Shape Formation*, and (e) *Tunnel Crossing*.

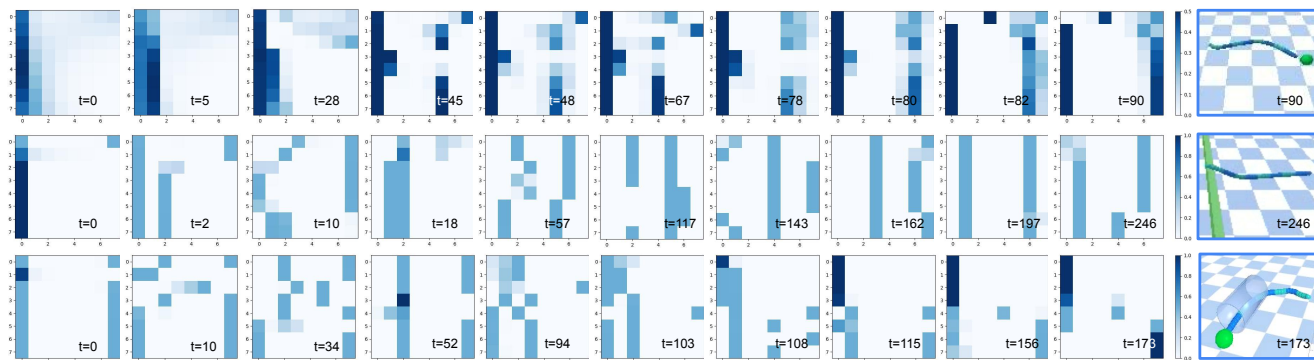


Fig. 5. Visualization of attention matrices $A \in \mathbb{R}^{8 \times 8}$ for an 8-agent snake. A_{ij} is the attention score for agent i attending to agent j . The higher the attention score is, the more agent i attends to agent j . From left to right, we show transitions of the attention matrix throughout one successful episode for three tasks, (a) *Goal Reaching*, (b) *Wall Climbing*, and (c) *Tunnel Crossing*.

D. Qualitative Visualization

In Fig. 5, we visualize the heatmap of the attention matrix $A = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})$, $A \in \mathbb{R}^{8 \times 8}$, for an 8-agent snake. A_{ij} is the attention score for agent i attending to agent j . The first agent ($i = 0$), denotes the snake head, and the last agent ($i = 7$) denotes the snake tail. As shown in Fig. 5, our attention mechanism effectively captures different patterns for different tasks, and the attention dynamically moves to different body segments over time. This shows that our attention mechanism has the potential to encompass both task-related details and the underlying dynamics of the snake robot. In *Goal Reaching* (the first row in Fig. 5), the attention matrix consistently focuses on the snake head, which directly determines the shared reward and the success condition of the task. We also observe that a portion of the attention transitions from the snake head to the snake tail from $t = 0$ to $t = 90$, akin to the sinusoidal actuation pattern that generates locomotion through anisotropic friction [42]. In *Wall Climbing* (the second row in Fig. 5), the agents attend to different body segments that are interacting with the wall over time. In *Tunnel Crossing* (the third row in Fig. 5), the attention is distributed over more segments than the first two

tasks at the early stage, because its whole body has to adapt to the tunnel shape. Then, the agents attend more to the head part as the head of the snake gets close to its targeted position. Notably, the colors across different agents are often nearly uniformly distributed within the same column of the heatmap, implying that n agents achieve a shared consensus on global information through self-attention communication.

V. CONCLUSION

Controlling snake robots for dexterous and compliant behaviors is extremely challenging due to complex dynamics. To address this challenge, we introduce a scalable and robust modular policy that leverages the inherent modularity of snake robots. In this work, we formulate the snake robot control framework within a cooperative multi-agent reinforcement learning setting. Extensive experiments are conducted across five tasks, showcasing the modular policy's notable efficiency in learning snake robot control, its robustness against hardware failures, and its generalizability for previously unseen robots. Furthermore, we plan to extend the application of the trained modular policy to physical continuum robots and soft robots.

REFERENCES

- [1] M. Runciman, A. Darzi, and G. P. Mylonas, "Soft robotics in minimally invasive surgery," *Soft robotics*, vol. 6, no. 4, pp. 423–443, 2019.
- [2] H. Wang, R. Zhang, W. Chen, X. Wang, and R. Pfeifer, "A cable-driven soft robot surgical system for cardiotoracic endoscopic surgery: preclinical tests in animals," *Surgical endoscopy*, vol. 31, pp. 3152–3158, 2017.
- [3] G. Li, X. Chen, F. Zhou, Y. Liang, Y. Xiao, X. Cao, Z. Zhang, M. Zhang, B. Wu, S. Yin *et al.*, "Self-powered soft robot in the mariana trench," *Nature*, vol. 591, no. 7848, pp. 66–71, 2021.
- [4] D. A. Schreiber, F. Richter, A. Bilan, P. V. Gavrilov, H. M. Lam, C. H. Price, K. C. Carpenter, and M. C. Yip, "Arcsnake: An archimedes' screw-propelled, reconfigurable serpentine robot for complex environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7029–7034.
- [5] C. Tang, B. Du, S. Jiang, Q. Shao, X. Dong, X.-J. Liu, and H. Zhao, "A pipeline inspection robot for navigating tubular environments in the sub-centimeter scale," *Science Robotics*, vol. 7, no. 66, p. eabm8597, 2022.
- [6] J. Whitman, N. Zevallos, M. Travers, and H. Choset, "Snake robot urban search after the 2017 mexico city earthquake," in *Proceedings of IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR '18)*, August 2018.
- [7] T. Wang, J. Whitman, M. Travers, and H. Choset, "Directional compliance in obstacle-aided navigation for snake robots," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 2458–2463.
- [8] T. Wang, C. Pierce, V. Kojouharov, B. Chong, K. Diaz, H. Lu, and D. I. Goldman, "Mechanical intelligence simplifies control in terrestrial limbless locomotion," *Science Robotics*, vol. 8, no. 85, p. eadi2243, 2023. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adi2243>
- [9] S. Yi, K. Sycara, and Z. Temel, "Reconfigurable Robot Control Using Flexible Coupling Mechanisms," in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [10] T. Kano, T. Sato, R. Kobayashi, and A. Ishiguro, "Decentralized control of scaffold-assisted serpentine locomotion that exploits body softness," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5129–5134.
- [11] Q. Fu and C. Li, "Robotic modelling of snake traversing large, smooth obstacles reveals stability benefits of body compliance," *Royal Society open science*, vol. 7, no. 2, p. 191192, 2020.
- [12] M. Travers, J. Whitman, and H. Choset, "Shape-based coordination in locomotion control," *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1253–1268, 2018.
- [13] D. Rollinson, K. V. Alwala, N. Zevallos, and H. Choset, "Torque control strategies for snake robots," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 1093–1099.
- [14] F. Xu, Y. Zhang, J. Sun, and H. Wang, "Adaptive visual servoing shape control of a soft robot manipulator using bezier curve features," *IEEE/ASME Transactions on Mechatronics*, vol. 28, no. 2, pp. 945–955, 2022.
- [15] A. Plaat, W. Kusters, and M. Preuss, "Deep model-based reinforcement learning for high-dimensional problems, a survey," *arXiv preprint arXiv:2008.05598*, 2020.
- [16] T. Chen, A. Murali, and A. Gupta, "Hardware conditioned policies for multi-robot transfer learning," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [17] Y. Jia and S. Ma, "A coach-based bayesian reinforcement learning method for snake robot control," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2319–2326, 2021.
- [18] M. Yim, K. Roufas, D. Duff, Y. Zhang, C. Eldershaw, and S. Homans, "Modular reconfigurable robots in space applications," *Autonomous Robots*, vol. 14, pp. 225–237, 2003.
- [19] O. Itani and E. Shammas, "Motion planning for redundant multi-bodied planar kinematic snake robots," *Nonlinear Dynamics*, vol. 104, no. 4, pp. 3845–3860, 2021.
- [20] I. Erkmén, A. M. Erkmén, F. Matsuno, R. Chatterjee, and T. Kamegawa, "Snake robots to the rescue!" *IEEE Robotics & Automation Magazine*, vol. 9, no. 3, pp. 17–25, 2002.
- [21] J. Sun, M. Yao, X. Xiao, Z. Xie, and B. Zheng, "Co-optimization of morphology and behavior of modular robots via hierarchical deep reinforcement learning," in *Robotics: Science and Systems (RSS) 2023*, 2023.
- [22] 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 *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 7758–7765.
- [23] Z. Bing, C. Lemke, F. O. Morin, Z. Jiang, L. Cheng, K. Huang, and A. Knoll, "Perception-action coupling target tracking control for a snake robot via reinforcement learning," *Frontiers in Neurobotics*, vol. 14, p. 591128, 2020.
- [24] G. Sartoretti, W. Paivine, Y. Shi, Y. Wu, and H. Choset, "Distributed learning of decentralized control policies for articulated mobile robots," *IEEE Transactions on Robotics*, vol. 35, no. 5, pp. 1109–1122, 2019.
- [25] M. Schilling, A. Melnik, F. W. Ohl, H. J. Ritter, and B. Hammer, "Decentralized control and local information for robust and adaptive decentralized deep reinforcement learning," *Neural Networks*, vol. 144, pp. 699–725, 2021.
- [26] F. Pigozzi, Y. Tang, E. Medvet, and D. Ha, "Evolving modular soft robots without explicit inter-module communication using local self-attention," in *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 148–157.
- [27] J. Whitman, M. Travers, and H. Choset, "Learning modular robot control policies," *IEEE Transactions on Robotics*, 2023.
- [28] W. Huang, I. Mordatch, and D. Pathak, "One policy to control them all: Shared modular policies for agent-agnostic control," in *International Conference on Machine Learning*. PMLR, 2020, pp. 4455–4464.
- [29] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," *Neural Information Processing Systems (NIPS)*, 2017.
- [30] P. Sunehag, G. Lever, A. Grusl, W. M. Czarnecki, V. Zambaldi, M. Jaderberg, M. Lanctot, N. Sonnerat, J. Z. Leibo, K. Tuyls *et al.*, "Value-decomposition networks for cooperative multi-agent learning based on team reward," in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2018, pp. 2085–2087.
- [31] T. Rashid, M. Samvelyan, C. S. De Witt, G. Farquhar, J. Foerster, and S. Whiteson, "Monotonic value function factorisation for deep multi-agent reinforcement learning," *arXiv preprint arXiv:2003.08839*, 2020.
- [32] C. Yu, A. Velu, E. Viniitsky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, "The surprising effectiveness of ppo in cooperative multi-agent games," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 611–24 624, 2022.
- [33] J. G. Kuba, R. Chen, M. Wen, Y. Wen, F. Sun, J. Wang, and Y. Yang, "Trust region policy optimisation in multi-agent reinforcement learning," *arXiv preprint arXiv:2109.11251*, 2021.
- [34] S. Sukhbaatar, R. Fergus *et al.*, "Learning multiagent communication with backpropagation," *Advances in neural information processing systems*, vol. 29, 2016.
- [35] A. Singh, T. Jain, and S. Sukhbaatar, "Learning when to communicate at scale in multiagent cooperative and competitive tasks," *arXiv preprint arXiv:1812.09755*, 2018.
- [36] Y. Niu, R. R. Paleja, and M. C. Gombolay, "Multi-agent graph-attention communication and teaming," in *AAMAS*, 2021, pp. 964–973.
- [37] C. Villani *et al.*, *Optimal transport: old and new*. Springer, 2009, vol. 338.
- [38] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.
- [39] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [40] M. A. Graule, T. P. McCarthy, C. B. Teeple, J. Werfel, and R. J. Wood, "Somogym: A toolkit for developing and evaluating controllers and reinforcement learning algorithms for soft robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4071–4078, 2022.
- [41] M. Wen, J. Kuba, R. Lin, W. Zhang, Y. Wen, J. Wang, and Y. Yang, "Multi-agent reinforcement learning is a sequence modeling problem," *Advances in Neural Information Processing Systems*, vol. 35, pp. 16 509–16 521, 2022.
- [42] D. L. Hu, J. Nirody, T. Scott, and M. J. Shelley, "The mechanics of slithering locomotion," *Proceedings of the National Academy of Sciences*, vol. 106, no. 25, pp. 10 081–10 085, 2009.