

# GCNT: Graph-Based Transformer Policies for Morphology-Agnostic Reinforcement Learning

Yingbo Luo<sup>1</sup>, Meibao Yao<sup>1\*</sup> and Xueming Xiao<sup>2</sup>

<sup>1</sup>Jilin University

<sup>2</sup>Changchun University of Science and Technology

luoyb23@mails.jlu.edu.cn, meibaoyao@jlu.edu.cn, alexcapshow@gmail.com

## Abstract

Training a universal controller for robots with different morphologies is a promising research trend, since it can significantly enhance the robustness and resilience of the robotic system. However, diverse morphologies can yield different dimensions of state space and action space, making it difficult to comply with traditional policy networks. Existing methods address this issue by modularizing the robot configuration, while do not adequately extract and utilize the overall morphological information, which has been proven crucial for training a universal controller. To this end, we propose GCNT, a morphology-agnostic policy network based on improved Graph Convolutional Network (GCN) and Transformer. It exploits the fact that GCN and Transformer can handle arbitrary number of modules to achieve compatibility with diverse morphologies. Our key insight is that the GCN is able to efficiently extract morphology information of robots, while Transformer ensures that it is fully utilized by allowing each node of the robot to communicate this information directly. Experimental results show that our method can generate resilient locomotion behaviors for robots with different configurations, including zero-shot generalization to robot morphologies not seen during training. In particular, GCNT achieved the best performance on 8 tasks in the 2 standard benchmarks.

## 1 Introduction

In recent years, Deep Reinforcement Learning (DRL) has been widely used for continuous control tasks of mobile robots and has achieved satisfactory performance [Mnih *et al.*, 2015; Lillicrap *et al.*, 2015; van Hasselt *et al.*, 2016; Fujimoto *et al.*, 2018; Ye *et al.*, 2020; Xie *et al.*, 2022; Sun *et al.*, 2023]. However, in traditional reinforcement learning, for each robot configuration, a specific policy network for control needs to be trained from scratch, that is, one network corresponds to one configuration. Since training each policy network requires substantial time and resources

\*Corresponding Author

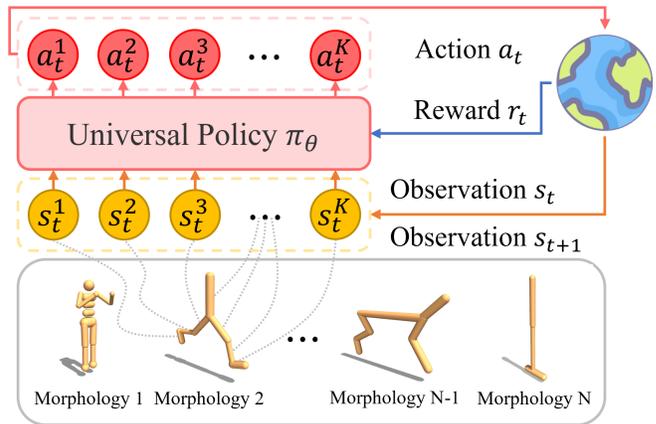


Figure 1: Learning universal controllers with various morphologies.

to interact with the environment, the overall cost is unacceptable when dealing with multiple robots. To address this issue, some scholars brought up the concept of the universal control network that allows all robot configurations to share a single policy network [Wang *et al.*, 2018; Pathak *et al.*, 2019; Huang *et al.*, 2020; Kurin *et al.*, 2021; Hong *et al.*, 2022; Chen *et al.*, 2023], that is, one policy network corresponds to multiple configurations. Such reinforcement learning models are called morphology-agnostic RL, inhomogeneous multi-task RL, or modular RL. Obviously, compared with traditional reinforcement learning, morphology-agnostic reinforcement learning has broader applications: First, it can significantly enhance the robustness and resilience of robots [Zhang *et al.*, 2017]. For example, a field exploration robot with a broken leg can still operate normally through the original policy network. Second, for unstructured terrain in the wild, we can focus on optimizing the configurations of robots that are best adapted to the current environment, without having to train a separate motion controller for each morphology [Wang *et al.*, 2019]. Third, as in the fields of Natural Language Processing (NLP) [Devlin *et al.*, 2019; Brown *et al.*, 2020] and Computer Vision (CV) [Dosovitskiy *et al.*, 2021; He *et al.*, 2022; Han *et al.*, 2022; Tian *et al.*, 2024], it can serve as a pre-trained prior for fine-tuning new morphologies [Gupta *et al.*, 2022]. That is, for brand new robot configurations, we can use the parameters of the universal policy

network as initialization and achieve satisfactory control performance through a small number of iterative training. This saves a lot of training time and resources. Last, by providing a unified representation for robots of different morphologies, it makes training large models in the field of robotic control possible [Furuta *et al.*, 2023].

Since robots of different configurations have different numbers of joints and sensors, they have different state spaces and action spaces [Zhao *et al.*, 2024a; Zhao *et al.*, 2024b]. However, the traditional MLP-based reinforcement learning model requires the input and output dimensions of the policy network be fixed, so they cannot control variably-configured robots. A natural idea is that any robot can be represented as a combination of multiple modules, and we can assume that the dimensions of the state and action space of each module are the same. If our policy network operates on a single module rather than the entire robot (that is, the policy network receives local observations from individual modules as inputs and produces local actions for them as outputs), and the local actions cooperate to form the overall actions of the robot, then the problem that a single network cannot adapt to different robot configurations is solved.

However, there is a core problem that needs to be solved: how to better coordinate the discretized limbs of robots, so that the local actions of each module can be combined into the overall actions of the robot to accomplish various tasks. In other words, the key to solving the problem is how to determine the influence of the states of other modules on the current module. Previous studies generally opt for message passing mechanisms or Transformer architectures to address this issue. However, this approach is not ideal for the following two reasons. First, similar to the hidden state in Recurrent Neural Network (RNN) [Hochreiter and Schmidhuber, 1997; Guo *et al.*, 2020], message passing mechanisms face the problem of limited capacity. After messages pass through multiple nodes, the original information contained in them is greatly diminished. This results in ineffective communication between nodes that are far apart (e.g., the hand and foot nodes of a humanoid robot). Second, the purely Transformer-based architecture neglects the robot’s morphological information, which has been proven to be crucial for training a universal controller. Therefore, we propose GCNT, a morphology-agnostic robot control network architecture based on Graph Convolutional Network (GCN) [Kipf and Welling, 2017] and Transformer. Our key insight is that GCN is able to efficiently extract morphology information of robots, while Transformer ensures that it is fully utilized by allowing each node of the robot to communicate this information directly. GCNT converts the robot’s morphology into a corresponding graph, uses the local observations of individual modules as input, and generates local actions for each module as output, thereby achieving universal control for robots of different morphologies. In summary, our main contributions are as follows:

- We propose GCNT, a novel network architecture for morphology-agnostic reinforcement learning. It decomposes the problem into two independent processes: morphological information extraction and node communication, providing a new insight into robot control.

- To fully extract the morphological information of the robot, we designed complementary GCN module and Weisfeiler-Lehman module in GCNT to overcome the indexing inconsistency of traversal-based methods. Besides, we improve the structure of GCN by adding additional linear layers and residual connections to avoid gradient vanishing caused by network depth increasing.
- We conducted comparative experiments with 7 baselines in 2 standard benchmarks. Our method achieves leading performance in all 8 tasks and zero-shot generalization.

## 2 Problem Formulation

To enable the policy network to handle different state space and action space dimensions of various robots, we modularize the robots. In this manner we treat any morphology of a robot as a combination of multiple modules, where each module has identical state space and action space dimensions. The policy network takes the local state of each module as input and generates local actions for each module as output. Specifically, consider an agent with  $K$  limbs, where we regard each limb and its corresponding joint (actuator) as a module. At each discrete time step  $t$ , the agent’s state  $s_t$  is composed of the local states  $s_t^i$ , that is,  $s_t = \{s_t^1, s_t^2, \dots, s_t^K\}$  and the action output by the policy network is  $a_t = \{a_t^1, a_t^2, \dots, a_t^K\}$ , where  $a_t^i$  is the torque value for the corresponding actuator of the limb  $k$ . Note that for  $\forall i, j \in \{1, 2, \dots, K\}$ , we ensure  $\dim(s_t^i) = \dim(s_t^j)$  and  $\dim(a_t^i) = \dim(a_t^j)$ .

To represent the relationship among modules, we use an undirected graph  $\mathcal{G} := \langle \mathcal{V}, \mathcal{E} \rangle$  to describe the robot’s morphology. Each node  $v_i \in \mathcal{V}$  for  $i \in \{1, \dots, K\}$  represents a limb of the agent, and each undirected edge  $(v_i, v_j) \in \mathcal{E}$  indicates that limb  $v_i$  and limb  $v_j$  are directly connected through one or more joints. In practical applications,  $\mathcal{E}$  can be represented by an adjacency matrix  $A \in \{0, 1\}^{K \times K}$ , where  $A^{i,j} = 1$  if  $v_i$  and  $v_j$  are connected, otherwise  $A^{i,j} = 0$ .

In this work, we assume there are  $N$  agents with different morphologies in the configuration space. As shown in Figure 1, a universal policy network  $\pi_\theta$  is trained to control them for locomotion, where  $\theta$  represents the learnable parameters of the policy. We adopt a joint policy optimization approach to achieve morphology-agnostic reinforcement learning, which follows Markov Decision Process (MDP). For each discrete time step  $t$ , the policy network  $\pi_\theta$  generates appropriate composite action  $\{a_t^k\}_{k=1}^{K_n}$  based on the current state  $\{s_t^k\}_{k=1}^{K_n}$  of agent  $n$ . After the agent  $n$  executes the actions, the environment returns the next state  $\{s_{t+1}^k\}_{k=1}^{K_n}$  and the corresponding reward  $r_t^n$ . Since the aim of our policy network is to perform well for all morphologies of robots, we need to consider the reward returns of all agents. Thus, our objective function is:

$$F(\theta) = \mathbb{E}_{\pi_\theta} \sum_{n=1}^N \sum_{t=0}^{\infty} \left[ \gamma^t r_t^n \left( \{s_t^k\}_{k=1}^{K_n}, \{a_t^k\}_{k=1}^{K_n} \right) \right] \quad (1)$$

where  $a_t^k = \pi_\theta(s_t^k)$  and  $\gamma$  is the discount factor. We aim to find the optimal solution  $\theta^*$  that maximizes the expected cumulative return on discounts for all agents.

$$\theta^* = \arg \max_{\theta} F(\theta) \quad (2)$$

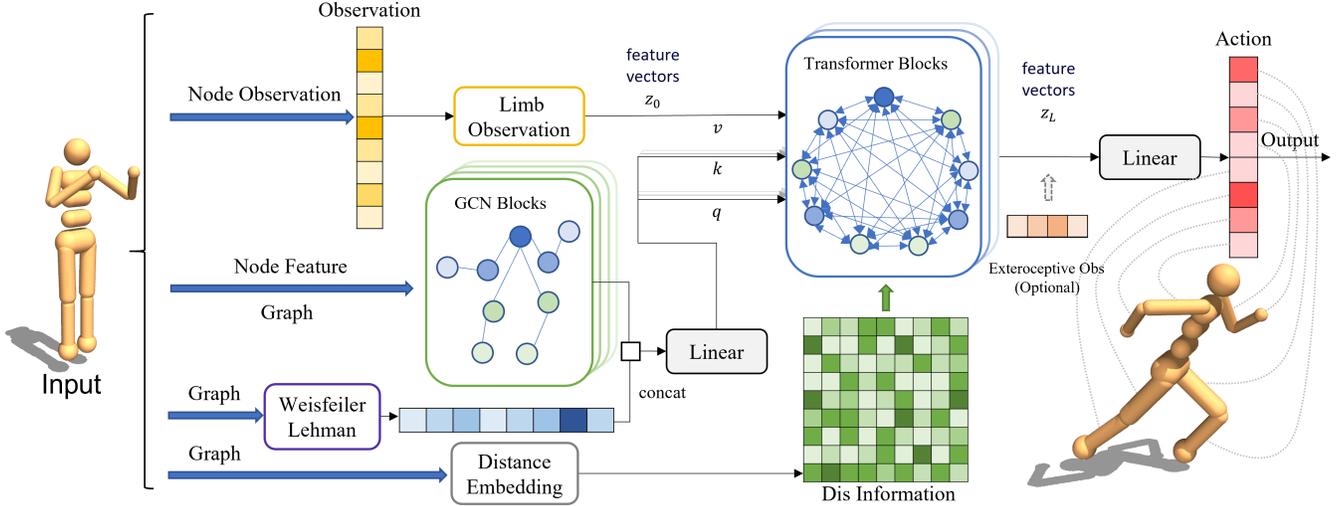


Figure 2: Overview of GCNT’s network architecture, mainly composed of five modules: the limb observation module processes state information, the GCN and Weisfeiler-Lehman modules extract morphological information, the learnable distance embedding module considers the impact of physical distances among limbs, and the Transformer module is used for communication between limbs.

### 3 Methodology

In this work, we present GCNT, a novel network architecture based on GCN and Transformer, for morphology-agnostic reinforcement learning in continuous control tasks. Our motivation is: GCN is able to aggregate information from neighboring nodes to the current node via graph convolution, which makes it naturally suitable for structural information extraction. Moreover, the Transformer allows each node to communicate directly and assess the importance of information through attention weights, thus preventing information loss during transmission. Therefore, we use these characteristics to solve the problem that discrete robot limbs cannot effectively cooperate with each other. More details of GCNT are shown below.

#### 3.1 Overview of GCNT

The network architecture of GCNT is shown in Figure 2. For robots of any morphology, the input to GCNT can be divided into four parts, represented by the four blue arrows in Figure 2. They go through the limb observation module, the GCN module, the Weisfeiler-Lehman module, and the learnable distance embedding module, respectively. The limb observation module includes the current state information of each limb, which is the main basis for generating actions. The GCN module is proposed to extract the robot’s local morphological information. The Weisfeiler-Lehman module is a complement to the GCN module and is used to extract the overall morphological information. The learnable distance embedding module determines the influence of limbs at different distances on the current limb. Afterwards, all information is aggregated in the Transformer module, and communication among limbs is completed through its attention mechanism. Finally, the results of communication are passed through a linear layer to obtain the actions that each limb should execute.

#### 3.2 Limb Observation Module

This module is mainly used to process limb observations. As mentioned, the dimension of each limb’s local observation are same. GCNT takes the local observation of each limb as input and generates local actions for each joint. In this work, the observation space of each limb in the two benchmarks remain the same as in previous work [Huang *et al.*, 2020; Kurin *et al.*, 2021; Hong *et al.*, 2022; Gupta *et al.*, 2022; Xiong *et al.*, 2023], which includes the motion state information and hardware attribute information of limbs and joints. The detailed information can be found in the Appendix D. All observations are uniformly passed through an MLP and mapped to a higher dimension for subsequent processing.

#### 3.3 GCN Module

In this paper, we propose the improved GCN module to extract the local morphology information of the robot. The mathematical principle behind GCN is based on spectral graph theory. It uses the normalized graph Laplacian for feature aggregation and approximates the convolution operation on the graph as a polynomial filter. This updates each node’s features as a weighted combination of its own features and those of its neighbors, so that the structure information in the graph can be captured effectively. The process of extracting information by GCN can be expressed as the following equation:

$$H^{(l+1)} = \sigma \left( \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right) \quad (3)$$

where  $\tilde{A} = A + I_K$ ,  $A$  is the adjacency matrix, and  $I_K$  is the identity matrix of size  $K$ ,  $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$ ,  $W^{(l)}$  is a layer-specific trainable weight matrix,  $\sigma(\cdot)$  is an activation function,  $H^{(l)} \in \mathbb{R}^{K \times D}$  is the feature matrix in the  $l^{th}$  layer, and  $H^{(0)} = X$ .  $X$  represents the original attributes of each node. In this work, we use one-hot encoding to represent different types of limbs and include them as the original attributes  $X$  of

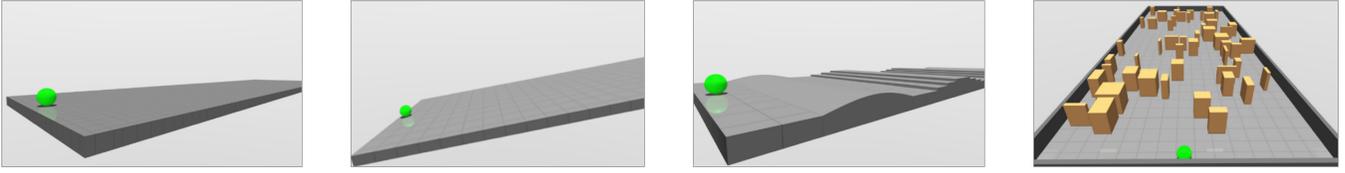


Figure 3: The four experiment environments in UNIMAL: Flat Terrain (FT), Incline, Mixed Terrain (MT), and Obstacles. Images credit to Gupta et al. [2021; 2022]

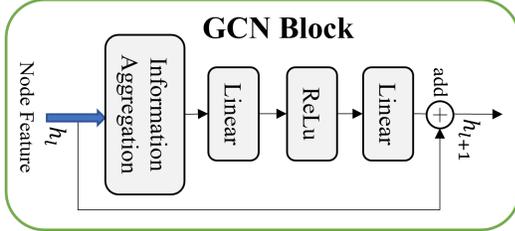


Figure 4: Internal structure of the GCN block.

each node in the GCN operations. Notably, we made changes to the structure of GCN, as shown in Figure 4. To enhance the extraction performance, we added additional linear layers to the network structure and we incorporated residual connections similar to those in ResNet [He *et al.*, 2016] to prevent vanishing gradient caused by the excessive depth of the GCNT network [Ying *et al.*, 2021]. The computation process of the improved GCN is as follows:

$$H^{(l+1)} = \sigma \left( \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W_1^{(l)} \right) W_2^{(l)} + H^{(l)} \quad (4)$$

### 3.4 Weisfeiler-Lehman Module

The Weisfeiler-Lehman (WL) algorithm is widely used for graph isomorphism testing and graph classification tasks [Weisfeiler and Leman, 1968]. In this work, we use the WL algorithm to capture the overall morphological information of the robots. We believe that the color string obtained through the WL algorithm can reflect the global structural information of the graph to a certain extent. The more similar the morphology of two robots is, the more similar the results obtained by the WL algorithm will be, and their control policy should also be more similar. Therefore, we use the WL module as a complement to the GCN module to capture the overall morphological information of the robot. In practical applications, we use the type of each limb as the initial color of each node and adopt a fixed number of iterations for all robot morphology. After the iteration ends, we use a vector to represent the occurrence times of each color during the iteration and use this vector as the overall morphological information, concatenating it with the local morphological information generated by the GCN module for each node.

### 3.5 Learnable Distance Embedding Module & Transformer Module

The learnable distance embedding module is proposed to work with the Transformer module, considering the influence of distance in the communication between different limbs.

For a given robot configuration represented by an undirected graph  $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ , we denote the distance value of two directly connected nodes as 1 and use the Floyd algorithm to calculate the shortest distance  $D$  among all nodes in the graph, as shown by the green matrix in Figure 2. We input this into the Transformer module, as shown in Figure 5. We propose a learnable distance embedding layer to map the limb distance information to each head of the multi-head attention layer. Specifically,  $R^{i,j} = g_\phi(D^{i,j})$ , where  $g_\phi(\cdot)$  is a mapping function with parameters  $\phi$ , and  $R^{i,j}$  is a vector with each element corresponding to a head in the multi-head attention. When calculating the attention weights,  $R^{i,j}$  is added to the attention score and then passed through the Softmax function to reflect different impact of limbs at different distances on the current limb. This process can be expressed by Equation 5, where the subscript( $h$ ) denotes the  $h$ -th head and  $R_{(h)}^{i,j}$  is a  $h$ -th entry of  $R^{i,j}$ .

$$\alpha_{(h)}^{i,j} = \frac{q_{(h)}^i k_{(h)}^j}{\sqrt{d_{(h)}}} + R_{(h)}^{i,j} \quad (5)$$

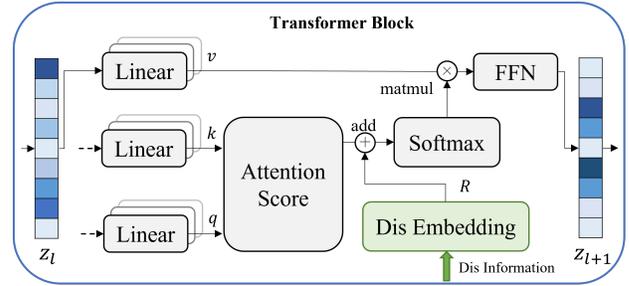


Figure 5: Internal structure of the Transformer block.

### 3.6 Optimization Method

In order to optimize the parameter  $\theta$  in GCNT, we use an actor-critic architecture based on the deterministic policy gradient algorithm to find the optimal value of the objective function in Equation 1, which is a standard practice for continuous control tasks [Lillicrap *et al.*, 2015]. The actor and critic have the same network architecture, which remains the same as Figure 2. Specifically, we use the TD3 algorithm [Fujimoto *et al.*, 2018] in benchmark 1 and the PPO algorithm [Schulman *et al.*, 2017] in benchmark 2 to ensure a fair comparison with previous works.

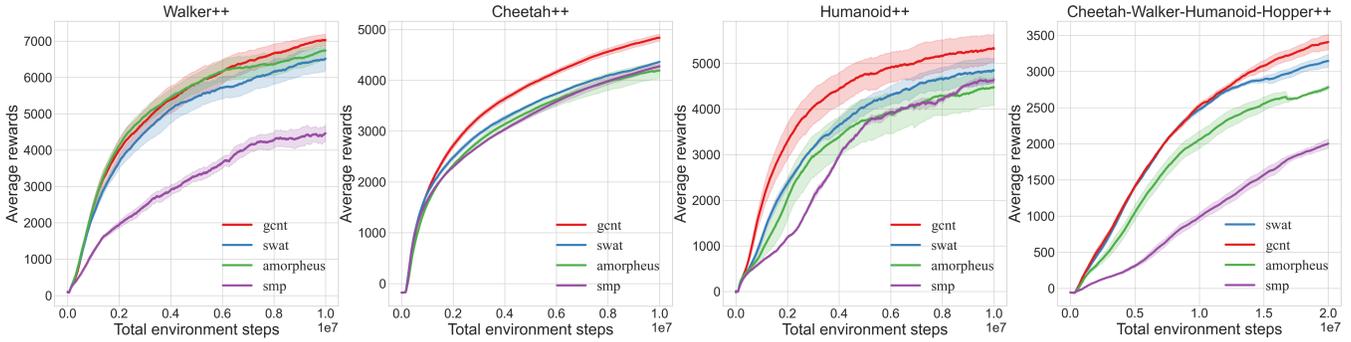


Figure 6: Average rewards across all agents in environments of Walker++, Cheetah++, Humanoid++, CWHH++ in SMPENV.

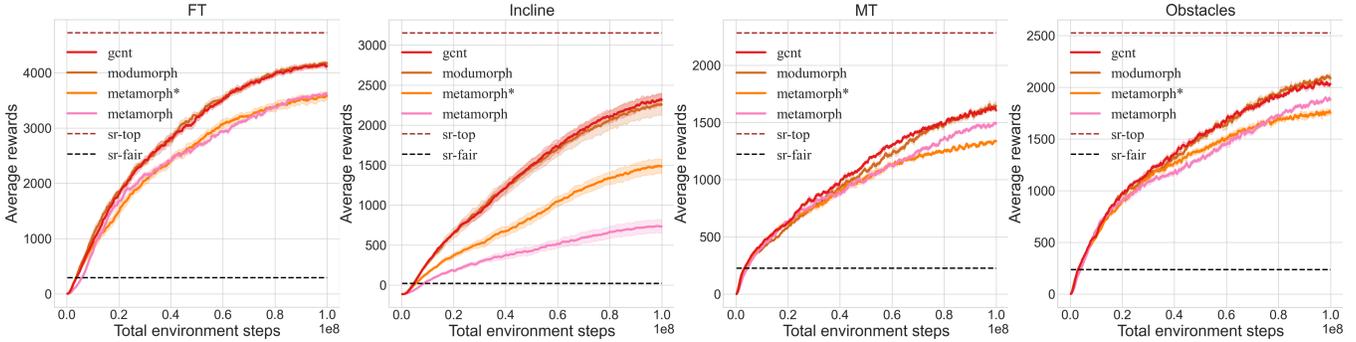


Figure 7: Average rewards across all agents in environments of FT, Incline, MT and Obstacles in UNIMAL.

## 4 Experimental Results

We conducted experiments across 8 different scenarios on two standard benchmarks and compared with 7 baselines, ensuring thorough testing. The benchmarks are SMPENV [Huang *et al.*, 2020] and UNIMAL [Gupta *et al.*, 2021]. All tests are based on the MuJoCo [Todorov *et al.*, 2012].

In SMPENV, we set up four test environments: (1) Walker++, (2) Cheetah++, (3) Humanoid++, (4) Cheetah-Walker-Humanoid-Hopper++ (CWHH++). These four environments can be divided into two categories, in-domain and cross-domain. For the first three test environments, robots of different morphologies in the environment have some similarities. Each test environment contains complete morphologies and variants with missing limbs of the same category of robots. For example, Humanoid++ includes both the complete humanoid\_2d.9\_full and the humanoid\_2d.7\_left\_leg, which is missing the left leg. While the last CWHH++ environment is a comprehensive test environment. The morphologies and the movement manner of the robots in the environment vary significantly. For example, the walker needs to move by alternately coordinating its left and right legs, while the hopper can only move by jumping with a single leg.

In UNIMAL, we conducted experiments in four environments. As shown in Figure 3, they are: (1) Flat Terrain (FT), (2) Incline, (3) Mixed Terrain (MT), and (4) Obstacles. In all environments, the goal of the robot is to pursue the farthest possible travel distance. Additionally, the terrain in MT and the pillars in Obstacles are randomized in each round. So the last two environments provide a height map centered around

the robot as an extra external observation input, allowing the agent to perceive and respond to different terrains or obstacles. Finally, it is worth noting that in UNIMAL we need to control 100 robots with different morphologies. So UNIMAL has greater complexity than SMPENV.

### 4.1 Training Results in SMPENV

To demonstrate the effectiveness of our method, GCNT, we compare it with three baselines: the message passing method SMP [Huang *et al.*, 2020], the Transformer-based method AMORPHEUS [Kurin *et al.*, 2021], and the Structure-aware Transformer method SWAT [Hong *et al.*, 2022]. We run all experiments with 3 random seeds to report the mean and the standard error. For fairness, we use the TD3 algorithm to train the policy network for all methods. The results are shown in Figure 6, where the solid line represents the mean and the shaded area represents the standard error.

It can be observed that our method, GCNT, has the highest sample efficiency and achieves the best performance in all test environments. The experimental results validate our analysis. SMP, which uses the message-passing mechanism, suffers from severe information loss during multi-hop transmission, making it difficult for two distant limbs to coordinate. Thus it performed the worst in all environments. AMORPHEUS did not utilize morphological information. While it performed well for simpler robot configurations like Walker++ (each robot consisting of only a torso and two legs), it showed poor performance for more complex robot configurations. SWAT incorporated morphological information into the Transformer, but the information extracted based on

traversal and graph methods does not always have a positive effect. For example, in the Walker++ environment, SWAT’s performance is inferior to that of AMORPHEUS, which uses a pure Transformer architecture. This indicates that the morphological information extracted by SWAT actually hinders the training of the policy network. In contrast, the morphological information extracted by our GCNT is useful in all scenarios. Moreover, GCNT consistently outperforms SWAT, demonstrating that our method has a stronger ability to extract morphological information.

### 4.2 Training Results in UNIMAL

In UNIMAL, we compare the method with four baselines: the depth-first traversal-based method MetaMorph [Gupta *et al.*, 2022], the variant MetaMorph\* [Xiong *et al.*, 2023], the hypernetwork-based method ModuMorph [Xiong *et al.*, 2023], and the methods SR-fair and SR-top, which train a separate policy network for each robot. Both SR-fair and SR-top use an MLP with 3 hidden layers and 256 hidden units in each layer. The difference is that SR-fair trains each robot with the same budget as used in multi-robot training (1M steps), while SR-top trains each robot until convergence (10M steps). SR-top can be viewed as an approximate upper bound. All methods use the PPO algorithm for training (using 4 seeds) and the experimental results are shown in Figure 7. As can be seen, all methods exhibit higher sample efficiency than SR-fair, and our method, GCNT, achieves leading performance across all scenarios. It is worth noting that although ModuMorph has similar effects to our method, the number of parameters of ModuMorph is 2-3 times that of our method because it uses a Hypernetwork to generate different encoders and decoders for each limb. More importantly, it has worse generalization, as detailed in the zero-shot experiments. Additionally, our method can also integrate Hypernetworks to generate different encoders and decoders for each body part to enhance performance. The corresponding experimental results can be found in the Appendix G.

	gcnt	swat	amorpheus	smp
w3	171±4	417±34	<b>651±36</b>	142±6
w6	<b>2933±116</b>	935±22	1463±110	1424±76
h7	<b>3509±9</b>	2464±60	883±53	1056±27
h8	<b>3332±27</b>	958±34	1598±88	304±4
c3	96±8	-157±2	87±6	<b>211±8</b>
c5	<b>2440±18</b>	986±12	2245±17	2149±7
c6	<b>4780±21</b>	3252±7	4322±9	4726±8

Table 1: Zero-shot generalization to unseen morphologies. The w, h, and c in the first column are abbreviations for walker, humanoid, and cheetah, respectively.

### 4.3 Zero-Shot Generalization to Unseen Morphologies

The goal of learning a universal controller is not limited to training on specified morphologies but also includes zero-shot generalization to new morphologies without any further fine-tuning. Table 1 shows the results of the zero-shot generalization to unseen morphologies in SMPENV. Each method

was trained on the training sets of Walker++, Humanoid++, and Cheetah++ and evaluated on their corresponding test sets. We evaluated the average performance and the standard error over 3 random seeds, with each seed evaluated on 100 roll-outs. It can be seen that GCNT achieved the highest average score in 5 out of the 7 tests. The cases where GCNT’s generalization was not as good as other methods occurred in simpler configurations with fewer limbs (walker\_3\_main and cheetah\_3\_balanced both have only 3 limbs). However, as the number of limbs increased, GCNT’s performance improved. This further validates that GCNT can effectively extract and utilize morphological information of robots. To sum up, we believe that GCNT’s performance in zero-shot generalization is superior to other methods.

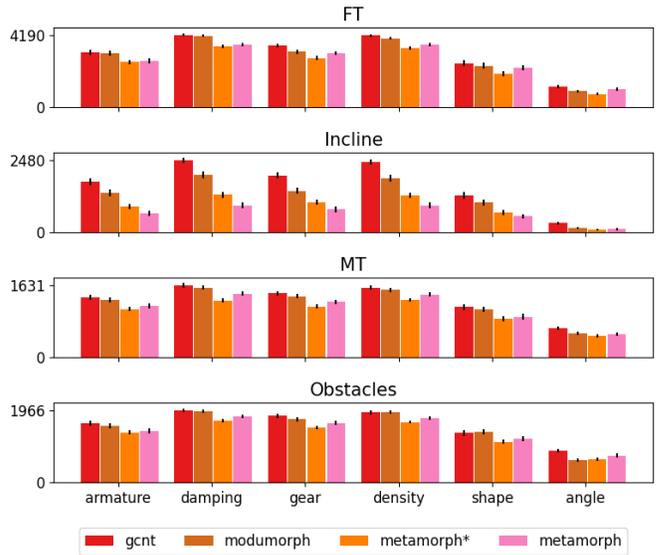


Figure 8: Zero-shot generalization to kinematics and dynamics variations. Error bars denote 95% bootstrapped confidence interval.

### 4.4 Zero-Shot Generalization to Kinematics and Dynamics Variations

We also conducted generalization experiments on dynamics and kinematics in UNIMAL. In this test set, the topological structure of each robot remains the same as in the training set, while the attributes of the limbs and joints are changed (including armature, density, damping, gear, module shape, and joint angle). We performed generalization experiments on 400 new robots, and the corresponding results are shown in Figure 8. As can be seen, our method, GCNT, exhibits the strongest generalization ability to both kinematics and dynamics.

### 4.5 Ablation Study

To demonstrate the role of each module in GCNT, we conducted the ablation study. Specifically, we removed the GCN module, Weisfeiler-Lehman module, and learnable distance embedding module separately. All methods were evaluated on Humanoid++ because it has the most limbs and the most

complex configuration, providing a complete test for validating the effectiveness of each module. The specific experimental results are shown in Figure 9. It can be seen that removing each module led to a performance drop, but the network still outperformed the pure Transformer AMORPHEUS.

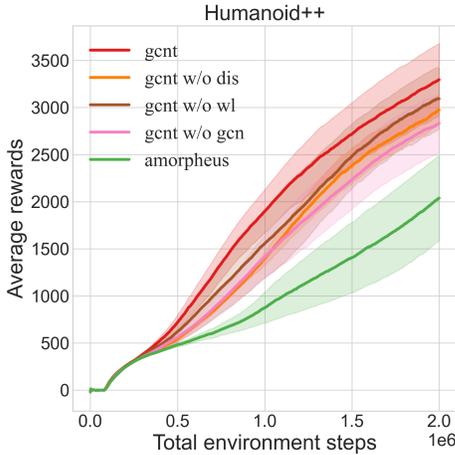


Figure 9: The results of ablation study of GCNT.

### 4.6 Morphological Information Embedding

Previous work has shown that the morphological information of a robot is crucial for training a universal controller. However, previous methods use depth-first traversal for extraction [Hong *et al.*, 2022; Gupta *et al.*, 2022], which does not guarantee that limbs of the same type have the same index in different robots, as shown in the Appendix H. To observe whether GCNT extracts the information effectively, we visualize the structural information learned in the CWHH++ environment. Specifically, we use the t-SNE algorithm [Van der Maaten and Hinton, 2008] to map the 128-dimensional structural vectors extracted by GCNT and SWAT into 2 dimensions. The corresponding results are shown in Figures 10, where the color of the points represents the type of limb, and the shape of the points represents the robot configuration to which it belongs. It can be seen that for GCNT, limbs with similar functions in different robot formation are mapped closely, while for SWAT, limbs with similar functions in different robot formation are mapped very scattered. This shows that SWAT’s traversal-based method has different performance on different robot configurations and cannot effectively extract the robot’s structural information. In contrast, the method based on GCN blocks and Weisfeiler-Lehman blocks in GCNT is effective for all robot configurations. This also explains why GCNT performs better in locomotion tasks and zero-shot generalization experiments.

### 5 Related Work

In general, research on morphology-agnostic reinforcement learning can be categorized as three stages. In the first stage, the core of related work is the use of message-passing mechanisms. Wang *et al.* [2018] proposed NerveNet, which is based on Graph Neural Network (GNN) [Sanchez-Lengeling

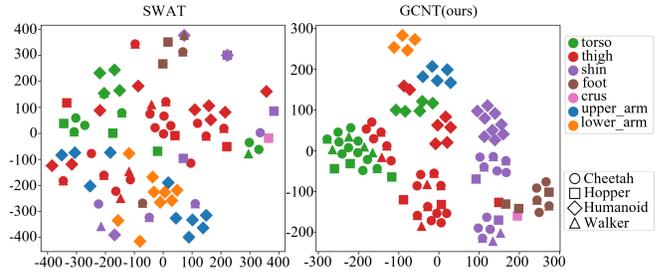


Figure 10: Morphological information embedding learned by GCNT and SWAT from different robot configurations.

*et al.*, 2021]. Pathak *et al.* [2019] and Huang *et al.* [2020] converted the robot’s configurations into corresponding graphs, and proposed Dynamic Graph Network (DGN) and Shared Modularization Policy (SMP) as the policy networks shared among robot modules. In the second stage, the focus of research shifted to the utilization of Transformers. To address the issue of critical information loss during multi-hop communication in message-passing mechanisms. Kurin *et al.* [2021] propose a Transformer [Vaswani *et al.*, 2017] policy network AMORPHEUS. It utilizes a multi-head self-attention mechanism, allowing direct communication among robot modules without the need for multiple hops through different nodes. Notably, the work by Kurin *et al.* [2021] initiated a trend in research on Transformer-based policy networks, and most subsequent work has been based on this architecture. In the third stage, researchers began to integrate the robot’s morphological information into the Transformer. Hong *et al.* [2022] proposed SWAT, a Transformer-based policy network that incorporates traversal-based positional embeddings in node attributes and graph-based relational embeddings in the multi-head attention mechanism to encode this information. Gupta *et al.* [2022] and Xiong *et al.* [2023] tried to use morphological information on larger robot benchmarks and proposed MetaMorph and ModuMorph respectively. Patel and Song [2024] leveraged embodied graph connectivity for structural bias in attention mechanism.

Although the aforementioned research has made some progress in morphology-agnostic reinforcement learning tasks, there are still many shortcomings. For instance, it does not ensure efficient communication between nodes or adequately extract the morphological information of robots. These are areas worth further exploration in the future.

### 6 Conclusion

In this work, we propose GCNT, a modular network architecture based on GCN and Transformer, for morphology-agnostic reinforcement learning. Through the effective extraction of robot morphological information by the GCN module and the direct communication among different limbs by the Transformer module, GCNT can provide resilient locomotion control for robots of any morphology. Compared to previous work, our method achieves leading performance in all eight environments and zero-shot generalization experiments under Mujoco platform.

## Acknowledgments

This work was sponsored by the National Natural Science Foundation of China (NSFC) through grant No. 52472448, and Deep Earth Probe and Mineral Resources Exploration-National Science and Technology Major Project through grants No. 2024ZD1000800 and No. 2024ZD1000804. We thank the above mentioned funds for their financial support.

## References

- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020.
- [Chen *et al.*, 2023] Runfa Chen, Jiaqi Han, Fuchun Sun, and Wenbing Huang. Subequivariant graph reinforcement learning in 3D environments. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 4545–4565. PMLR, 23–29 Jul 2023.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [Fujimoto *et al.*, 2018] Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1587–1596. PMLR, 10–15 Jul 2018.
- [Furuta *et al.*, 2023] Hiroki Furuta, Yusuke Iwasawa, Yutaka Matsuo, and Shixiang Shane Gu. A system for morphology-task generalization via unified representation and behavior distillation. In *International Conference on Learning Representations*, 2023.
- [Guo *et al.*, 2020] Qing Guo, Zhu Sun, Jie Zhang, and Yin-Leng Theng. An attentional recurrent neural network for personalized next location recommendation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):83–90, Apr. 2020.
- [Gupta *et al.*, 2021] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *Nature communications*, 12(1):5721, 2021.
- [Gupta *et al.*, 2022] Agrim Gupta, Linxi Fan, Surya Ganguli, and Li Fei-Fei. Metamorph: Learning universal controllers with transformers. In *International Conference on Learning Representations*, 2022.
- [Han *et al.*, 2022] Kai Han, Yunhe Wang, Jianyuan Guo, Yehui Tang, and Enhua Wu. Vision gnn: An image is worth graph of nodes. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 8291–8303. Curran Associates, Inc., 2022.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [He *et al.*, 2022] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [Hong *et al.*, 2022] Sunghoon Hong, Deunsol Yoon, and Kee-Eung Kim. Structure-aware transformer policy for inhomogeneous multi-task reinforcement learning. In *International Conference on Learning Representations*, 2022.
- [Huang *et al.*, 2020] Wenlong Huang, Igor Mordatch, and Deepak Pathak. One policy to control them all: Shared modular policies for agent-agnostic control. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4455–4464. PMLR, 13–18 Jul 2020.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [Kurin *et al.*, 2021] Vitaly Kurin, Maximilian Igl, Tim Rocktäschel, Wendelin Boehmer, and Shimon Whiteson. My body is a cage: the role of morphology in graph-based incompatible control. In *International Conference on Learning Representations*, 2021.

- [Lillicrap *et al.*, 2015] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *CoRR*, abs/1509.02971, 2015.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Belle-mare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [Patel and Song, 2024] Austin Patel and Shuran Song. GET-Zero: Graph embodiment transformer for zero-shot embodiment generalization, 2024.
- [Pathak *et al.*, 2019] Deepak Pathak, Christopher Lu, Trevor Darrell, Phillip Isola, and Alexei A Efros. Learning to control self-assembling morphologies: A study of generalization via modularity. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [Sanchez-Lengeling *et al.*, 2021] Benjamin Sanchez-Lengeling, Emily Reif, Adam Pearce, and Alexander B. Wiltschko. A gentle introduction to graph neural networks. *Distill*, 2021. <https://distill.pub/2021/gnn-intro>.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Sun *et al.*, 2023] Jieqiang Sun, Meibao Yao, Xueming Xiao, Zhibing Xie, and Bo Zheng. Co-optimization of morphology and behavior of modular robots via hierarchical deep reinforcement learning. In *Robotics: Science and Systems*, 2023.
- [Tian *et al.*, 2024] Pengzhi Tian, Meibao Yao, Xueming Xiao, Bo Zheng, Tao Cao, Yurong Xi, Haiqiang Liu, and Hutao Cui. 3d semantic terrain reconstruction of monocular close-up images of martian terrains. *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [Todorov *et al.*, 2012] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- [Van der Maaten and Hinton, 2008] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [van Hasselt *et al.*, 2016] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), Mar. 2016.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [Wang *et al.*, 2018] Tingwu Wang, Renjie Liao, Jimmy Ba, and Sanja Fidler. Nervenet: Learning structured policy with graph neural networks. In *International Conference on Learning Representations*, 2018.
- [Wang *et al.*, 2019] Tingwu Wang, Yuhao Zhou, Sanja Fidler, and Jimmy Ba. Neural graph evolution: Automatic robot design. In *International Conference on Learning Representations*, 2019.
- [Weisfeiler and Leman, 1968] Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra which appears therein. *nti, Series*, 2(9):12–16, 1968.
- [Xie *et al.*, 2022] Zhibing Xie, Hua Shang, Xueming Xiao, and Meibao Yao. Mapless navigation of modular mobile robots using deep reinforcement learning. In *2022 IEEE International Conference on Industrial Technology (ICIT)*, pages 1–6. IEEE, 2022.
- [Xiong *et al.*, 2023] Zheng Xiong, Jacob Beck, and Shimon Whiteson. Universal morphology control via contextual modulation. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 38286–38300. PMLR, 23–29 Jul 2023.
- [Ye *et al.*, 2020] Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang, Xipeng Wu, Qingwei Guo, Qiaobo Chen, Yinyuting Yin, Hao Zhang, Tengfei Shi, Liang Wang, Qiang Fu, Wei Yang, and Lanxiao Huang. Mastering complex control in moba games with deep reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):6672–6679, Apr. 2020.
- [Ying *et al.*, 2021] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do transformers really perform badly for graph representation? In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.
- [Zhang *et al.*, 2017] Tan Zhang, Wenjun Zhang, and Madan M. Gupta. Resilient robots: Concept, review, and future directions. *Robotics*, 6(4), 2017.
- [Zhao *et al.*, 2024a] Huanfeng Zhao, Meibao Yao, Xueming Xiao, and Bo Zheng. Dynainsremove: A real-time dynamic instance-aware static 3d lidar mapping framework for dynamic environment. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1803–1809. IEEE, 2024.
- [Zhao *et al.*, 2024b] Huanfeng Zhao, Meibao Yao, Yan Zhao, Yao Jiang, Hongyan Zhang, Xueming Xiao, and Ke Gao. M2cs: A multimodal and campus-scapes dataset for dynamic slam and moving object perception. *Journal of Field Robotics*, 2024.