Subequivariant Reinforcement Learning in 3D Multi-Entity Physical Environments

Project Page: https://alpc91.github.io/SMERL/

Runfa Chen*, Ling Wang*, Yu Du, Tianrui Xue, Fuchun Sun, Jianwei Zhang, Wenbing Huang

Generalization of Embodied Agents

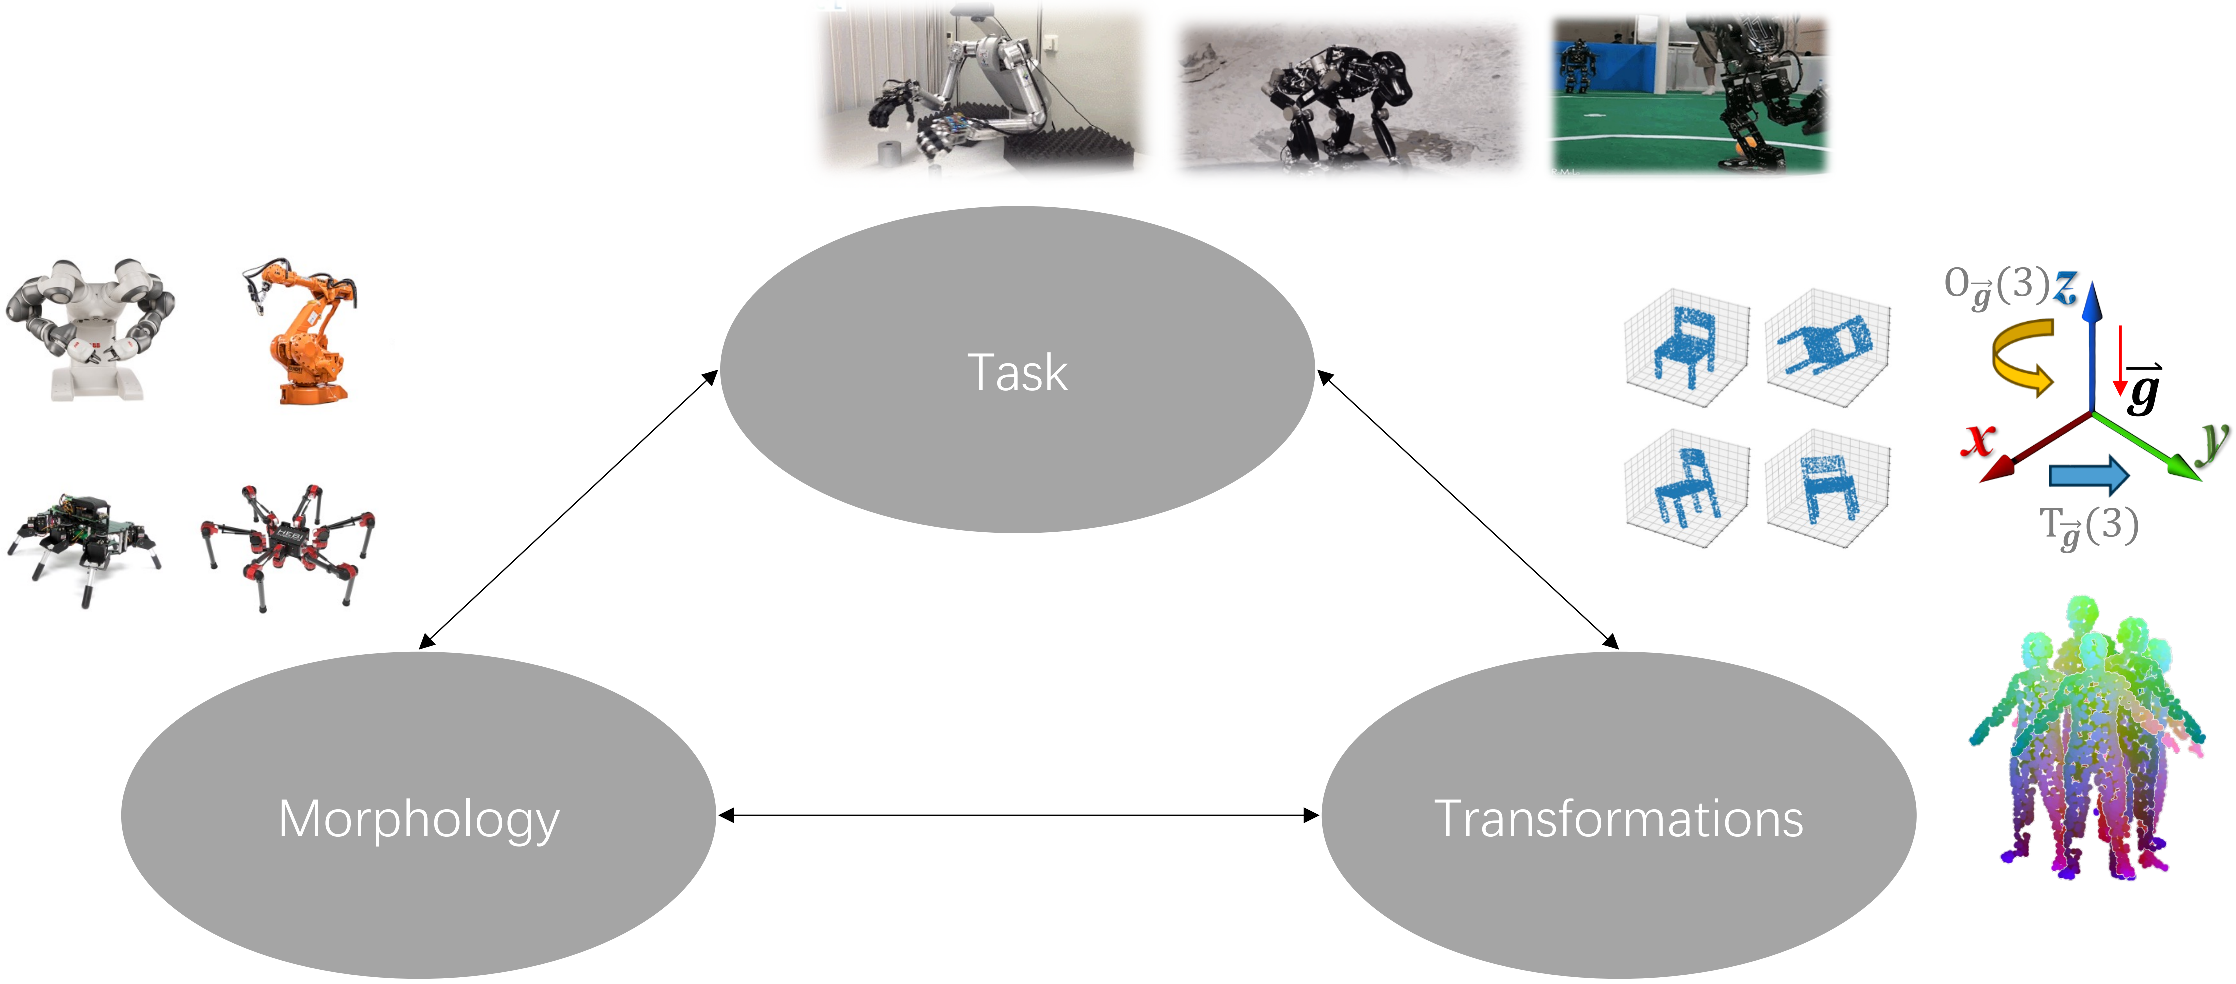

An intricate challenge is generalizing across configurations like transformations, morphologies, and tasks, which are interlinked and complicate the learning process.

Spatial Intelligent

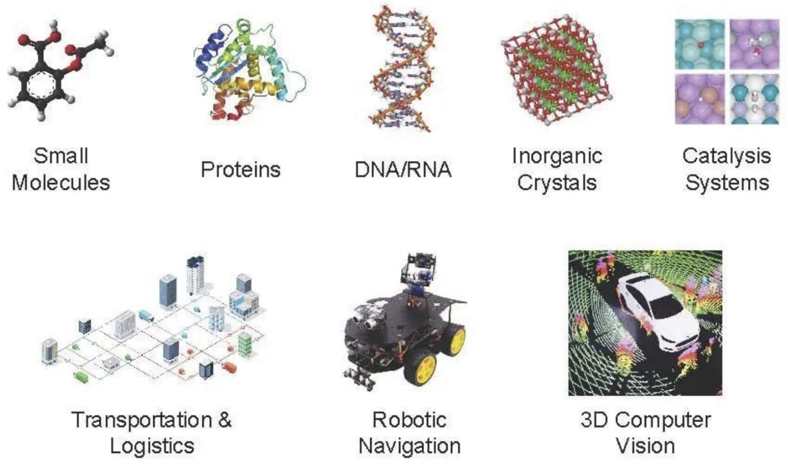

Geometric Structure and Systems
[Joshi et.al. Geometric Graph Neural Networks. NeurIPS 2022]

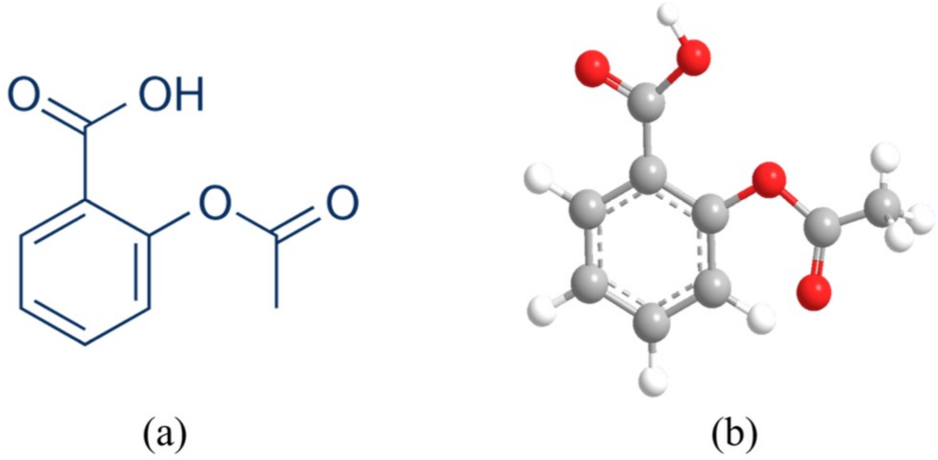

The examples in the figure correspond to Aspirin (acetylsalicylic acid) molecules:(a) topology graph and (b) geometric graph.

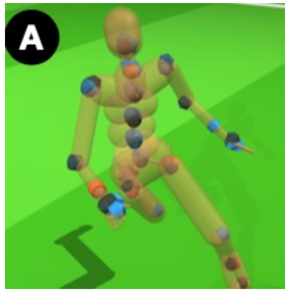

3D Geometric Graph!

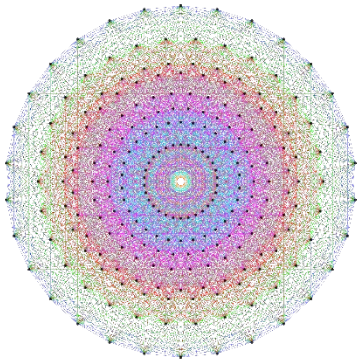

Symmetry!

Geometric Deep Learning

Geometric Symmetry The symmetrical structure in 3D environments is $\mathrm{E}(3)$, which is a 3-dimensional Euclidean *group* that consists of rotations, reflections, and translations.

Definition 2.1 (Group): A group G is a set of transformations with a binary operation " $\cdot$ " satisfying these properties: " $\cdot$ " is closed under associative composition, there exists an identity element, and each element must have an inverse.

Symmetrical structure enforced on the model is formally described by the concept of *equivariance*.

Definition 2.2 (Equivariance): Suppose $\vec{\boldsymbol{Z}}$ to be 3D geometric vectors (positions, velocities, etc) that are steerable by a group G, and $\boldsymbol{h}$ non-steerable features. The function f is G-equivariant, if for any transformation $g \in G$, $f(g \cdot \vec{\boldsymbol{Z}}, \boldsymbol{h}) = g \cdot f(\vec{\boldsymbol{Z}}, \boldsymbol{h})$, $\forall \vec{\boldsymbol{Z}} \in \mathbb{R}^{3\times m}, \boldsymbol{h} \in \mathbb{R}^d$. Similarly, f is invariant if $f(g \cdot \vec{\boldsymbol{Z}}, \boldsymbol{h}) = f(\vec{\boldsymbol{Z}}, \boldsymbol{h})$.

Specifically, the $\mathrm{E}(3)$ operation " $\cdot$ " is instantiated as $g \cdot \vec{\boldsymbol{Z}} := \boldsymbol{O}\vec{\boldsymbol{Z}}$ for the orthogonal group that consists of rotations and reflections where $\boldsymbol{O} \in \mathrm{O}(3) := \{\boldsymbol{O} \in \mathbb{R}^{3\times 3} | \boldsymbol{O}^\top \boldsymbol{O} = \boldsymbol{I}\}$, and is additionally implemented as the translation $g \cdot \vec{\boldsymbol{x}} := \vec{\boldsymbol{x}} + \vec{\boldsymbol{t}}$ for the 3D coordinate vector where $\vec{\boldsymbol{t}} \in \mathrm{T}(3) := \{\vec{\boldsymbol{t}} \in \mathbb{R}^3\}$. To align with the principles of classical physics under the influence of gravity, we introduce a relaxation of the group constraint. Specifically, we consider equivariance within the subgroup of $\mathrm{E}(3)$ induced by gravity $\vec{\boldsymbol{g}} \in \mathbb{R}^3$, defined as $\mathrm{O}_{\vec{\boldsymbol{g}}}(3) := \{\boldsymbol{O} \in \mathbb{R}^{3\times 3} | \boldsymbol{O}^\top \boldsymbol{O} = \boldsymbol{I}, \boldsymbol{O}\vec{\boldsymbol{g}} = \vec{\boldsymbol{g}}\}$ and $\mathrm{T}_{\vec{\boldsymbol{g}}}(3) := \{\vec{\boldsymbol{t}} \in \mathbb{R}^3 | \vec{\boldsymbol{t}}\vec{\boldsymbol{g}} = \vec{\boldsymbol{0}}\}$. By this means, the $\mathrm{E}_{\vec{\boldsymbol{g}}}(3)$-equivariance is only restrained to the translations/rotations/reflections along the direction of $\vec{\boldsymbol{g}}$ .We term *subequivariance* primarily referring to $\mathrm{E}_{\vec{\boldsymbol{g}}}(3)$-equivariance.

The Symmetry in 3D Multi-Entity Physical Environments

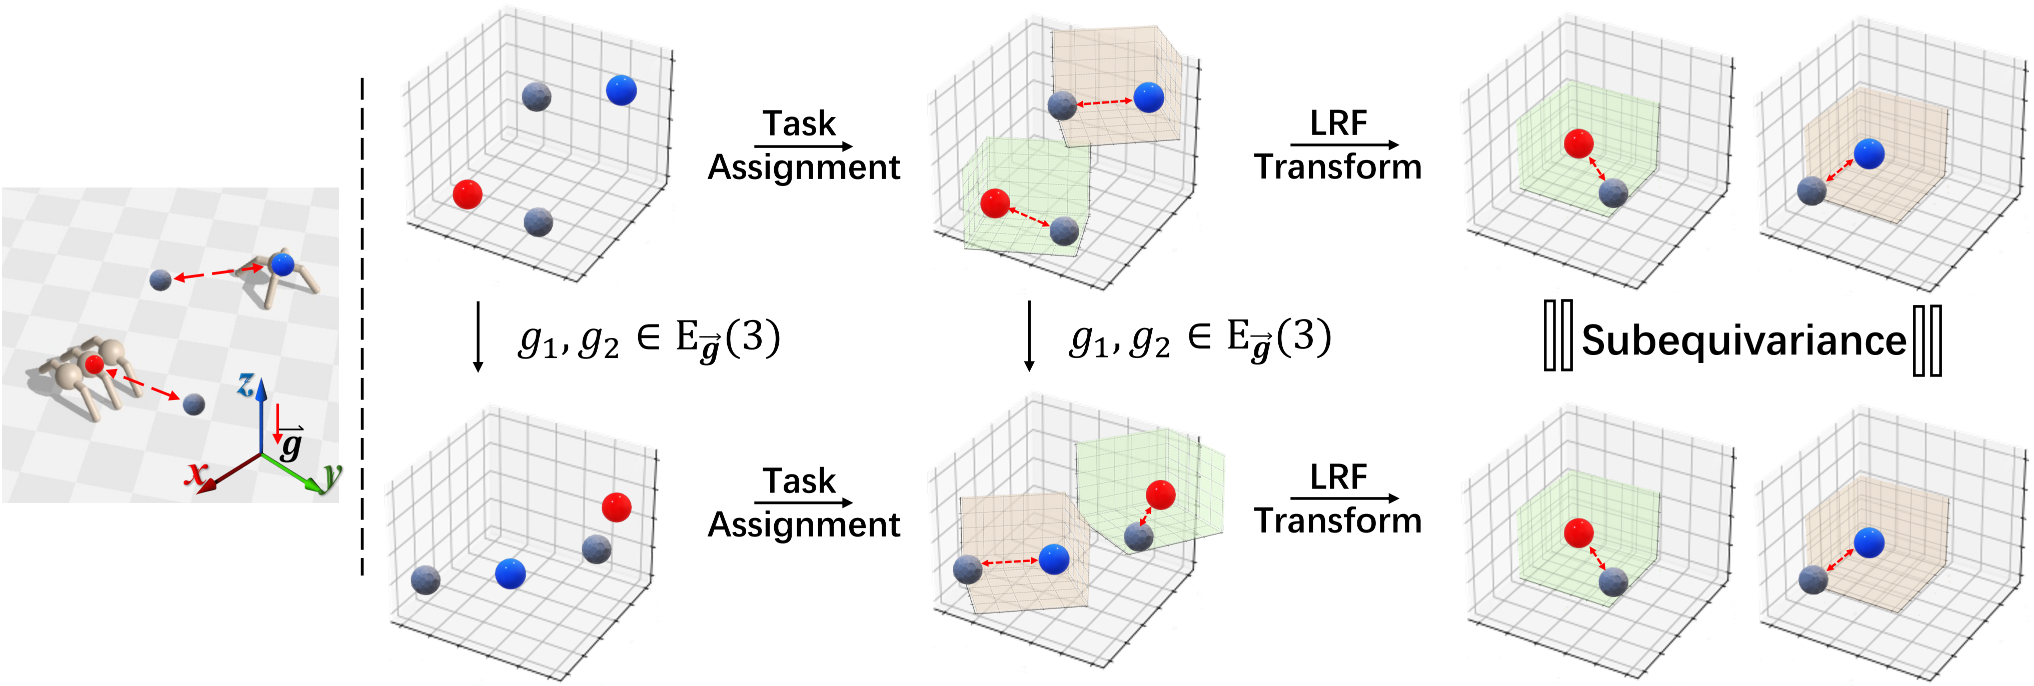

In particular, multi-entity systems, which include agents, objects, present considerable challenges compared to single-entity scenarios, partly due to exponential expansion of global transformations as the number of entities increases.

Subequivariant Hierarchical Neural Networks (SHNN)

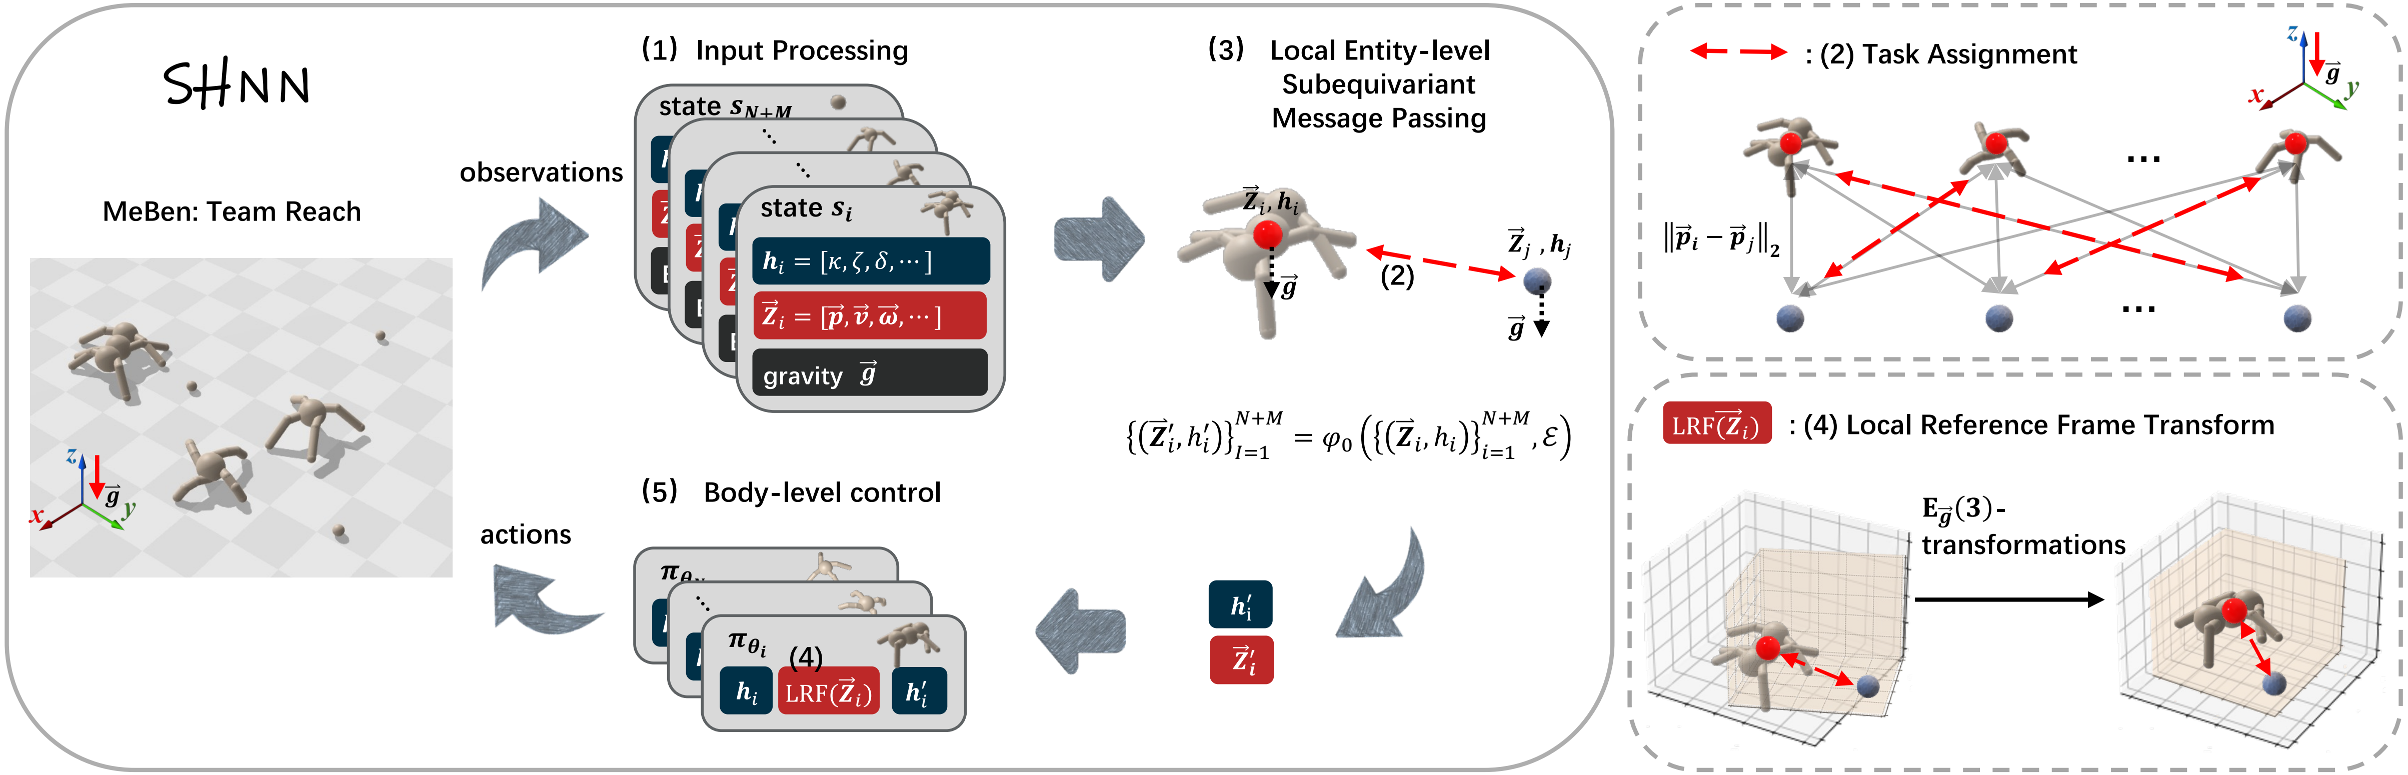

The right side illustrates our key innovation: the dynamic task assignment leveraging bipartite graph matching, and the construction of an $\mathrm{E}_{\vec{g}}(3)$-equivariant local reference frame for each entity to address local transformations.

Multi-entity Benchmark (MEBEN)

Table 1. Comparison of Morphology-based Environment Setup.

Aspect	SGRL	MxT-Bench	MEBEN
Multi-Morphology	✓	✓	✓
Multi-Agent	×	×	✓
Diverse-Task	×	✓	✓
Supported-Symmetry	✓	×	✓
Accelerated-Hardware	×	✓	✓

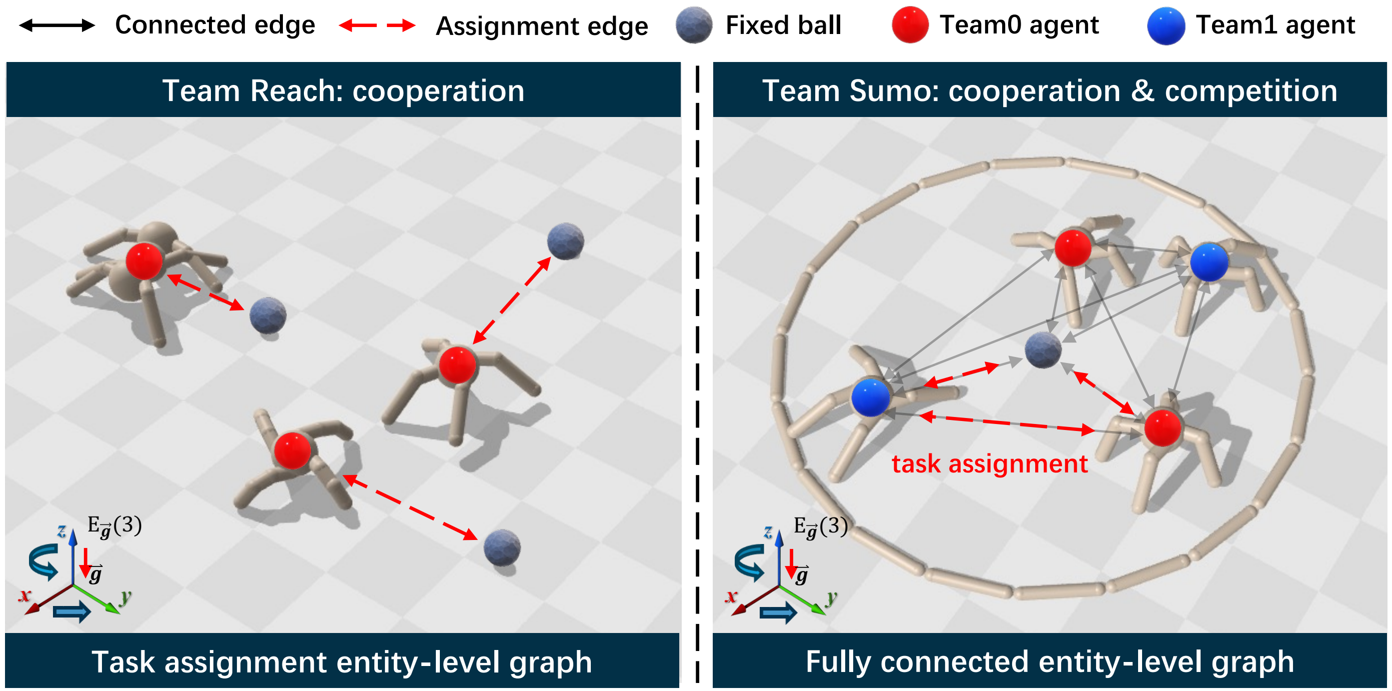

Team Reach (left) where agents cooperate to collectively reach all fixed balls, and Team Sumo (right) where agents engage in both cooperation and competition to push opponents away from the fixed ball.

Evaluations in Diverse Environments

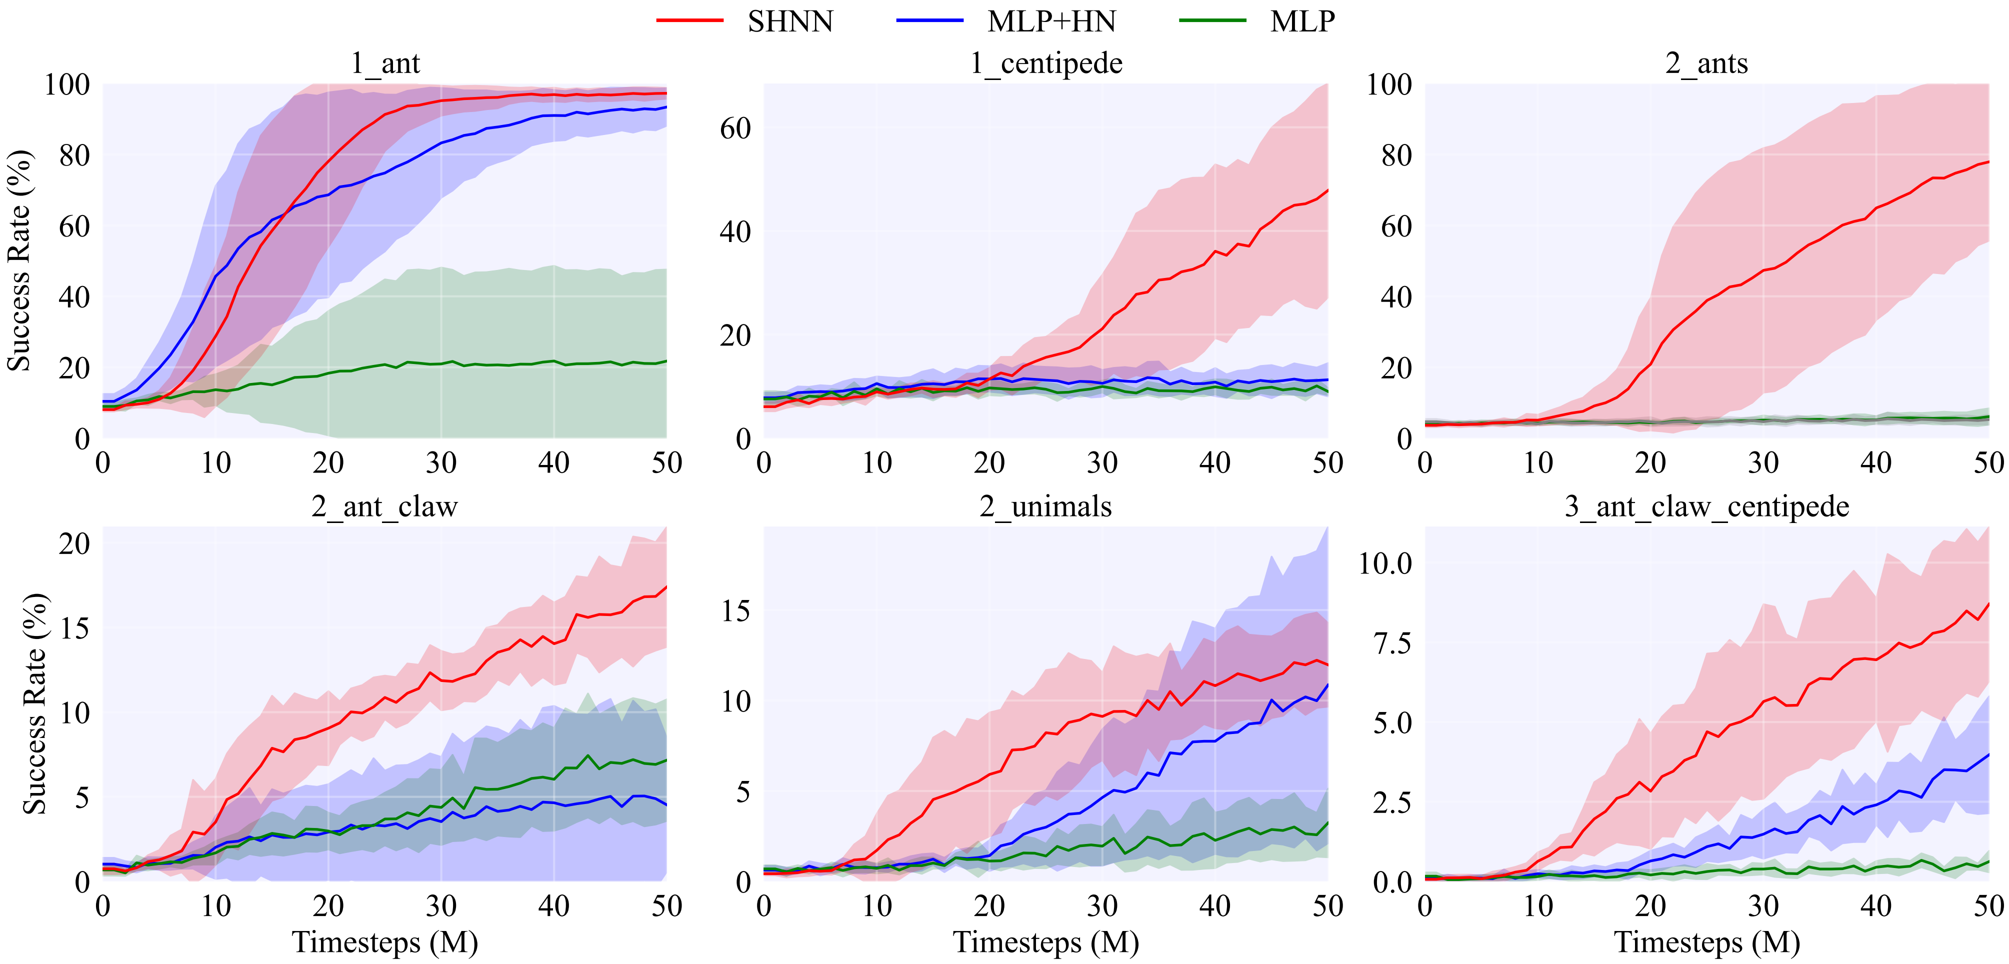

Team Reach

Evaluations in Diverse Environments

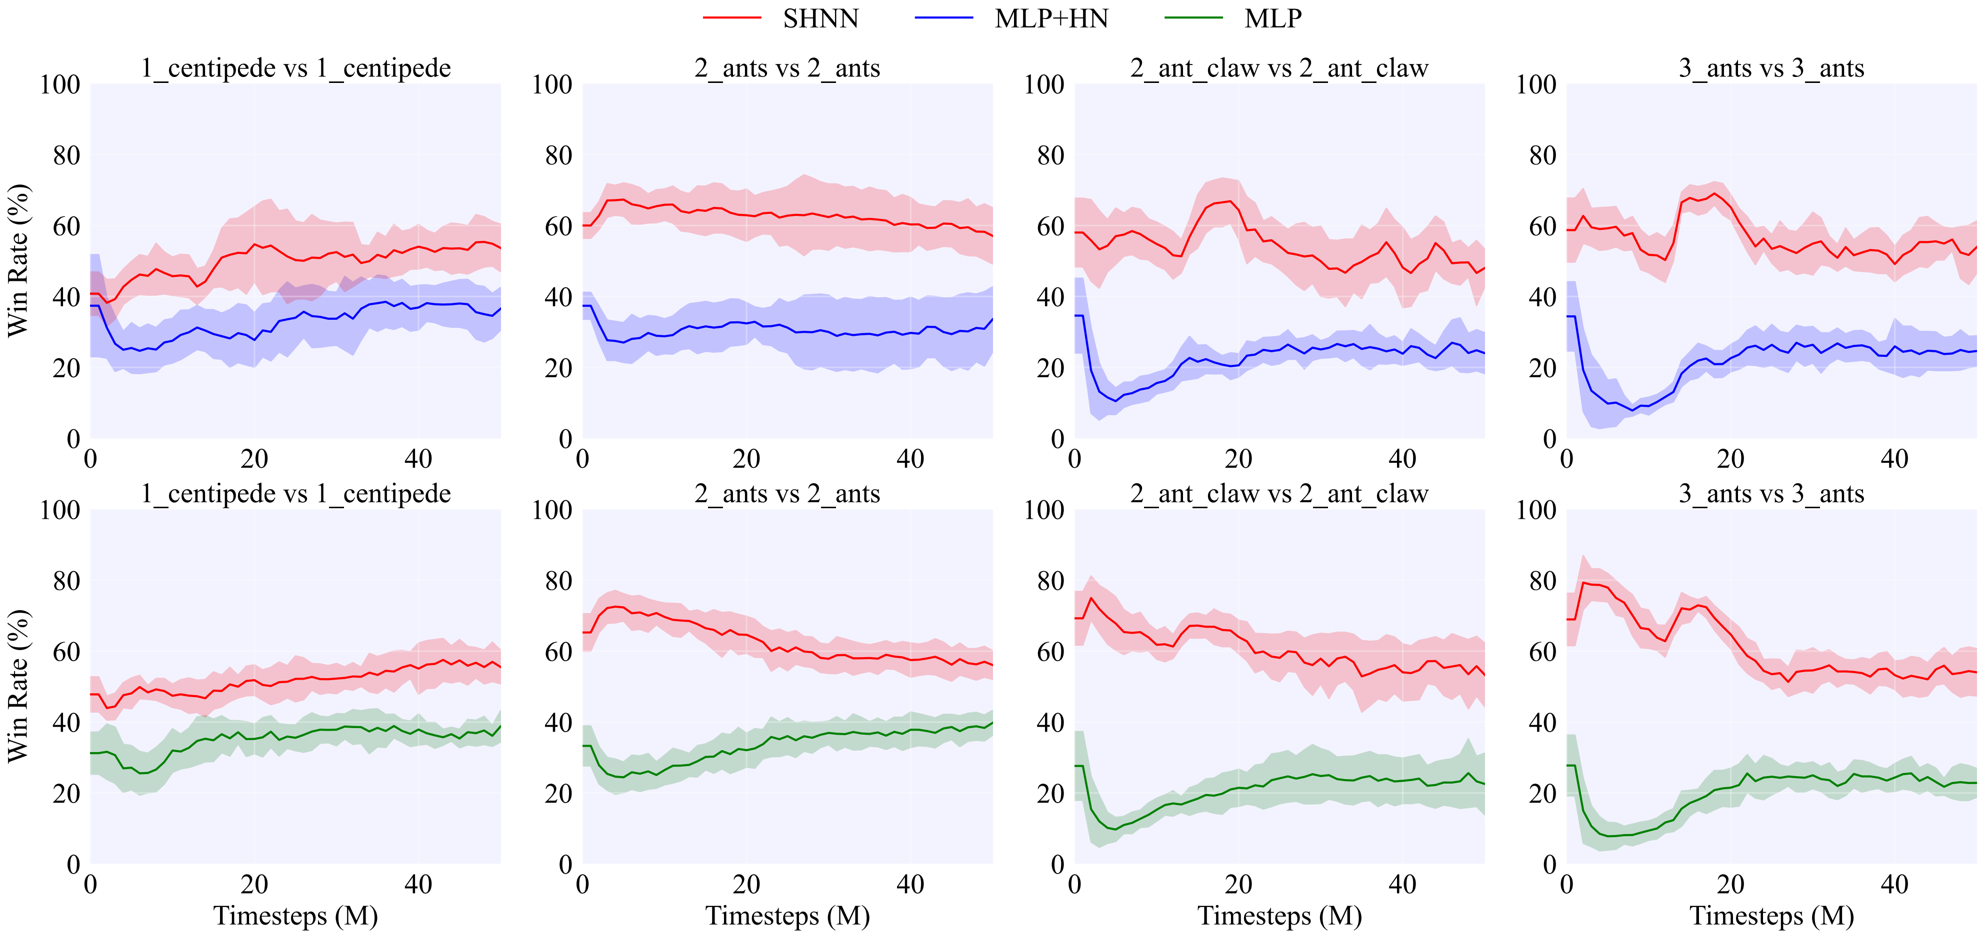

Team Sumo

Extended Evaluations on Transformer

Methods	1_ant	1_centipede	2_ants	2_ant_claw	2_unimals	2_ant_claw_centipede
MLP+HN	93.39 ± 5.25	11.28 ± 3.21	5.25 ± 1.399	4.52 ± 3.93	10.86 ± 8.74	3.98 ± 1.83
SHNN	**97.26** ± 1.51	**47.82** ± 20.62	**77.93** ± 22.22	**17.40** ± 3.54	**11.97** ± 2.31	**8.70** ± 2.42
Transformer+HN	5.47 ± 2.84	5.55 ± 2.99	2.47 ± 0.73	0.51 ± 0.19	0.25 ± 0.10	2.26 ± 1.92
SHTransformer	**63.61** ± 39.57	**11.52** ± 2.39	**11.37** ± 15.50	**1.17** ± 0.70	**0.26** ± 0.15	**7.12** ± 2.29

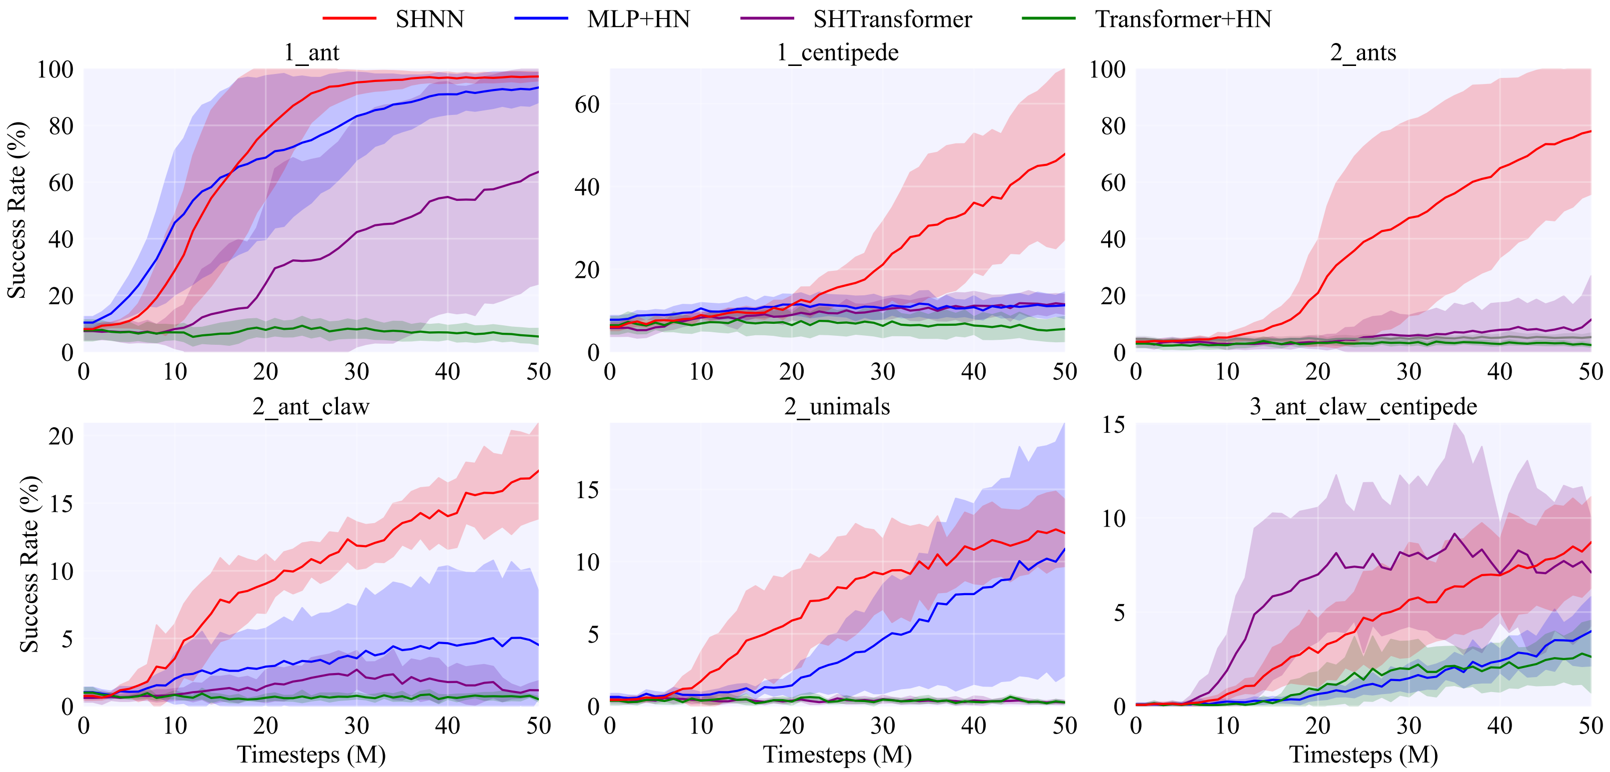

Ablations on Assignment

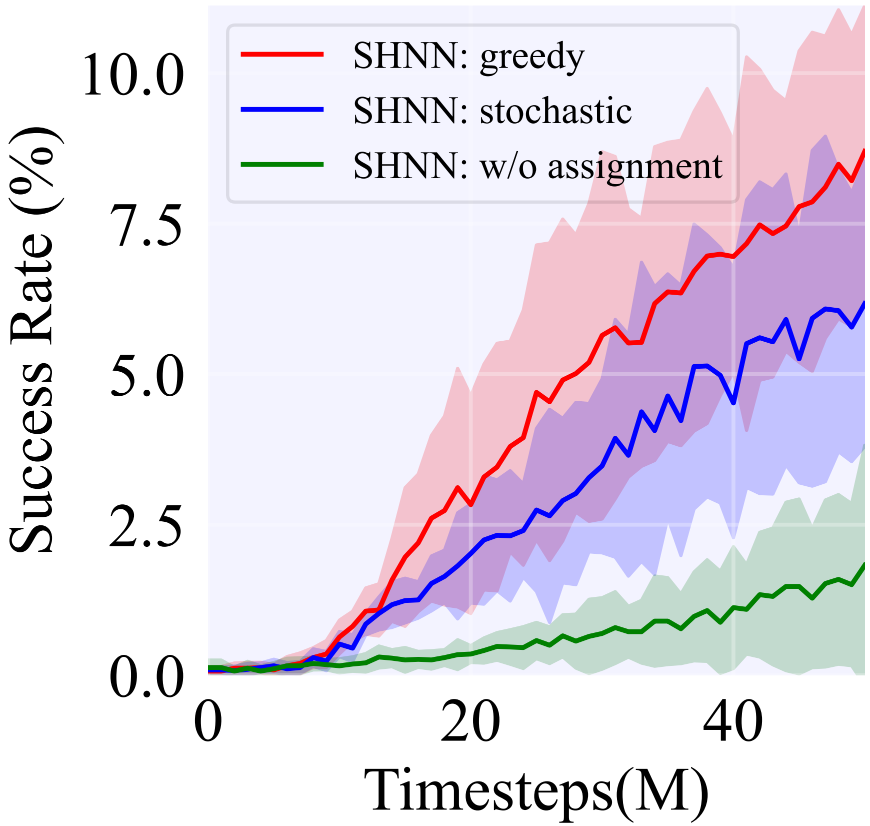

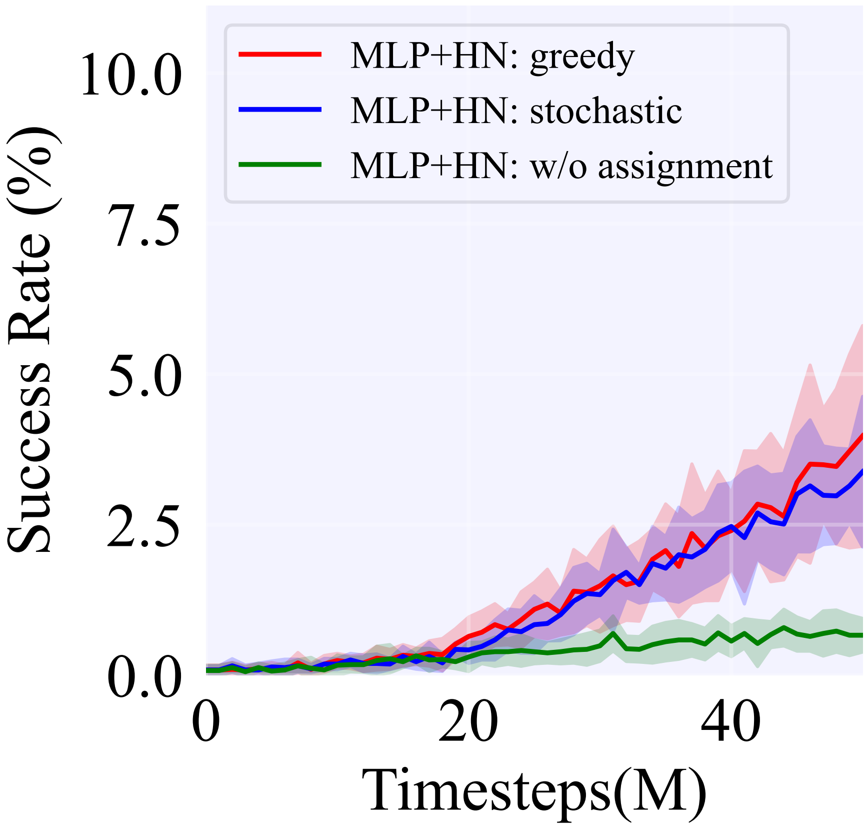

Ablations on Equivariance

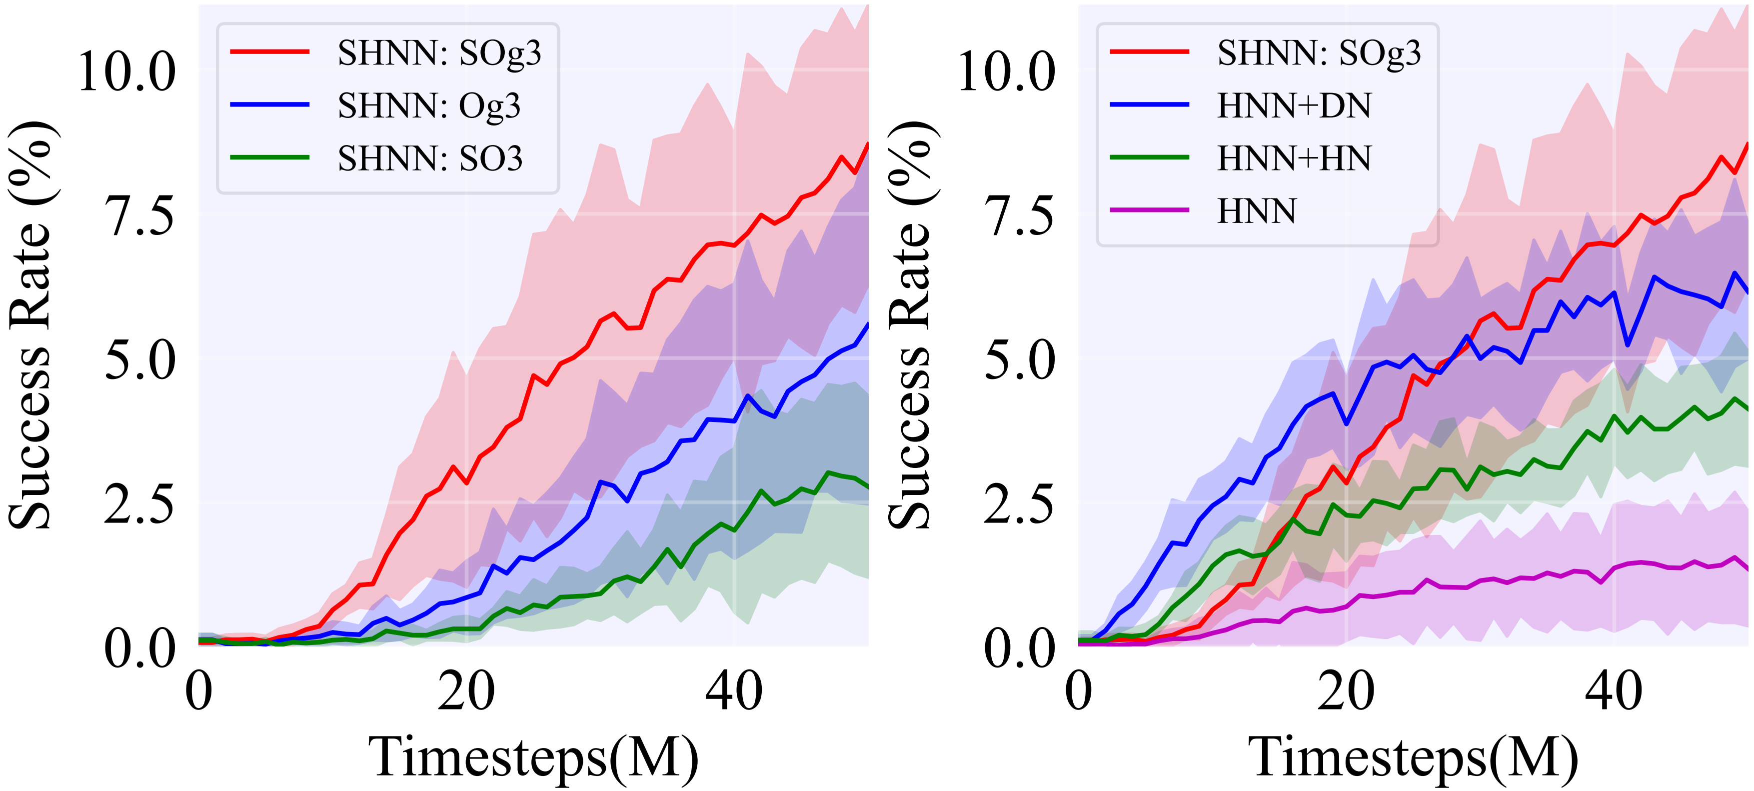

Analyses of Morphology-shared Policy

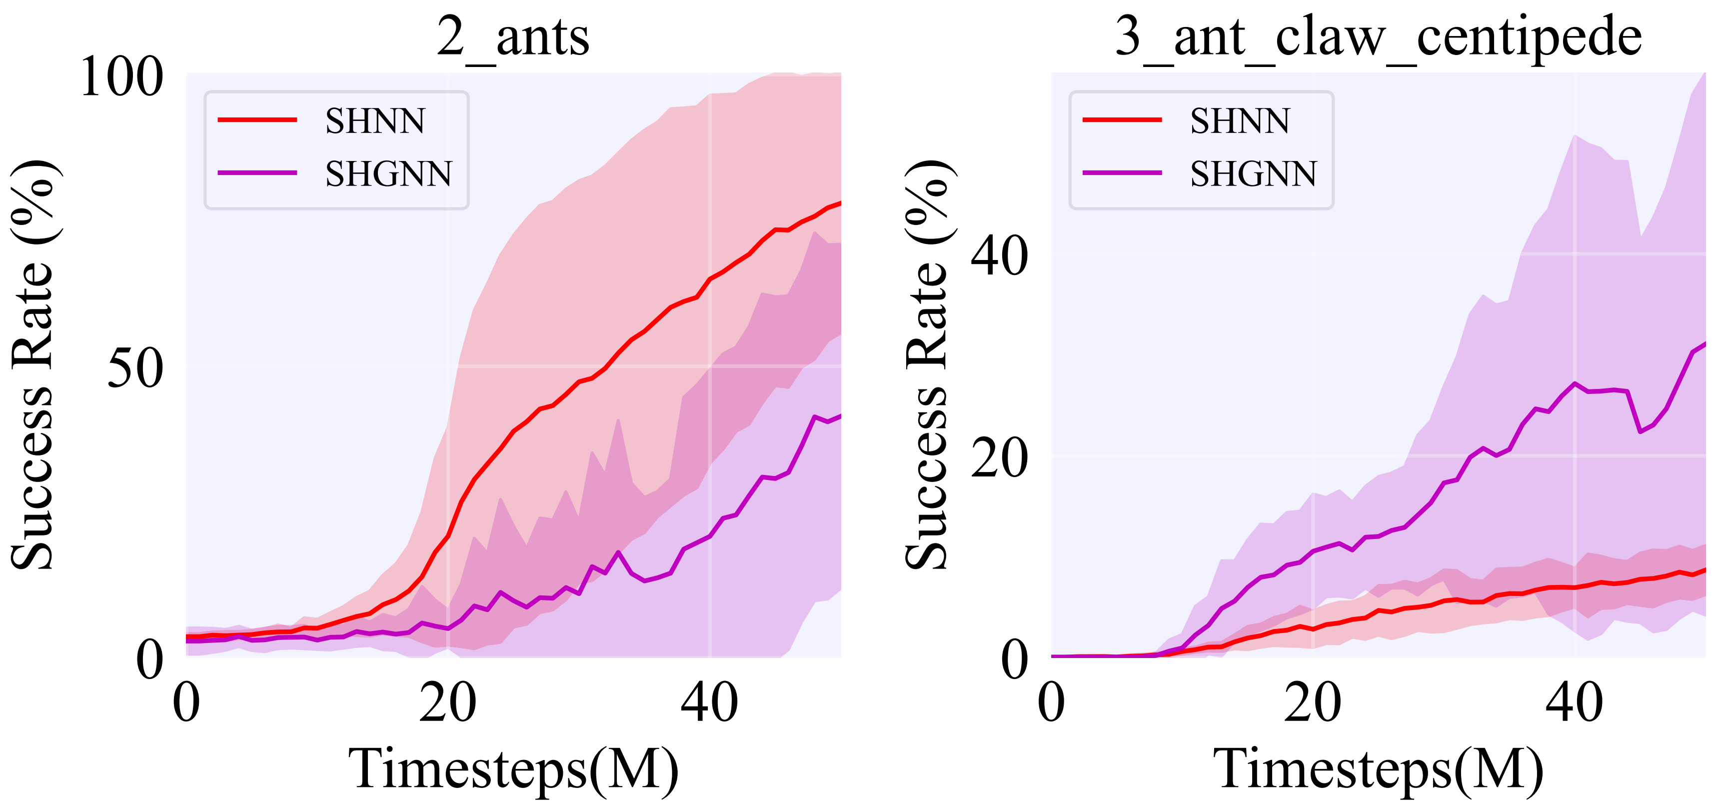

Importance of Local Symmetry

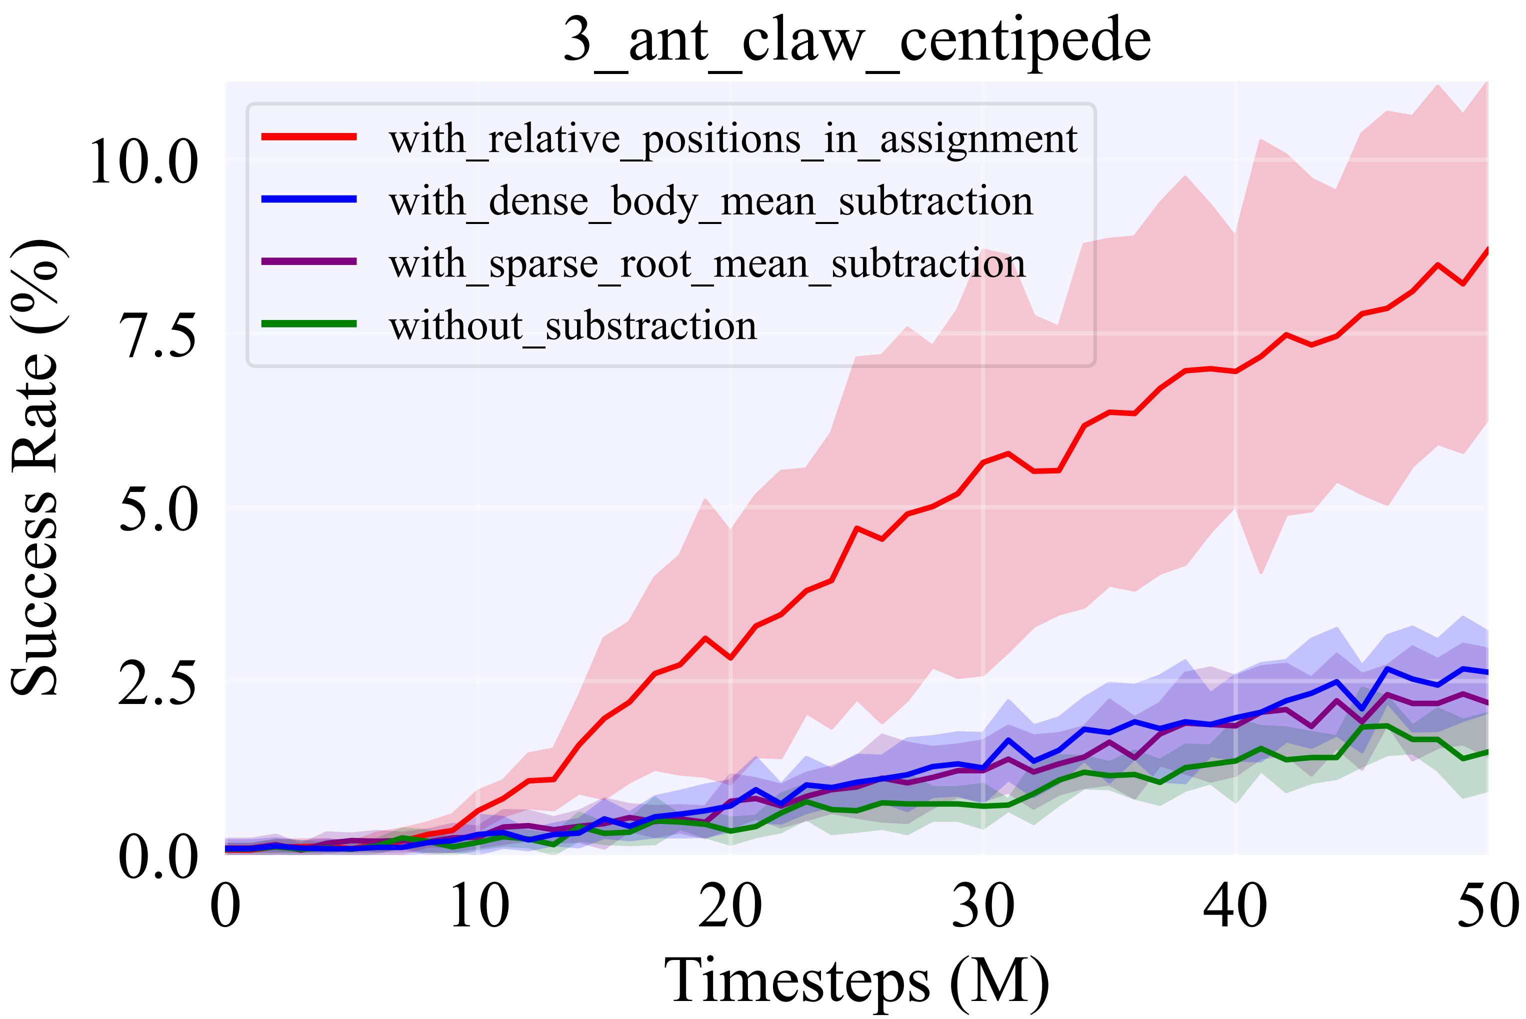

Model Comparison: Parameters and Training Time

Model	**Parameters** (M)	**Training Wall Time** (h)
MLP+HN	**1.759**	0.299±0.003
Transformer+HN	0.416	1.545±0.001
SHNN	0.772	1.302±0.012
SHTransformer	0.711	2.487±0.014
SHGNN	0.613	**5.538**±1.411

Take-away

- To effectively optimize the policy in 3D multi-entity physical environments, we propose SHNN, a novel framework that offers a superior plug-in alternative to hand-crafted LRFs. It decouples local transformations from the overall structure and compresses the state space by leveraging local physical geometric symmetry, particularly in gravity-affected environments.

- We introduce MEBEN, a collection of subequivariant morphology-based MARL environments, designed for in-depth exploration of multi-entity interactions within physical geometric symmetry constraints. These environments, including a diverse range of inter-entity transformations, facilitate both cooperative and competitive dynamics.

- We demonstrate the effectiveness of SHNN in the proposed 3D multi-entity physical environments, including Team Reach and Team Sumo. Our extensive ablations and comparative analyses also reveal the efficacy of the proposed ideas.

Thanks!

For more information, welcome to visit our website:

https://alpc91.github.io/SMERL/