上海期智研究院
SHANGHAI QI ZHI INSTITUTE

LIGHTSPEED STUDIOS

NEURAL INFORMATION PROCESSING SYSTEMS

Learning to Control Free-Form Soft Swimmers

Changyu Hu · Yanke Qu · Qiuan Yang · Xiaoyu Xiong · Kui Wu · Wei Li · Tao Du

Swimming is Complicated

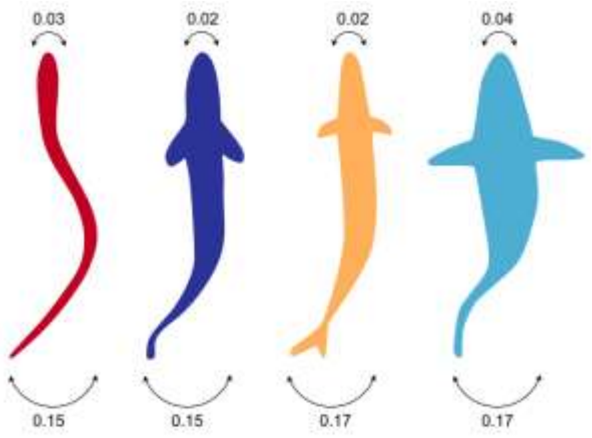

Fin/Body Undulation
[Santo et al. 2021]

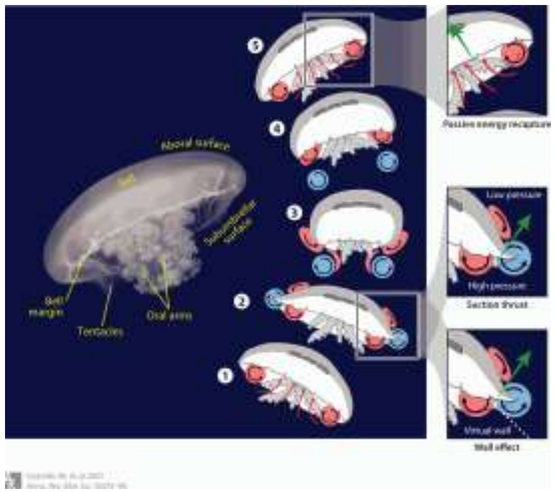

Jet Propulsion
[Costello et al. 2021]

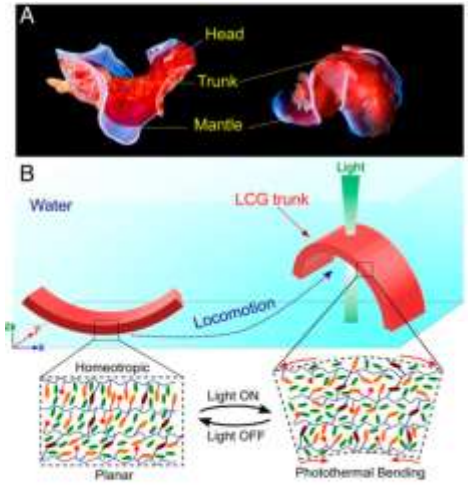

Undulatory Propulsion
[Shimoga et al. 2021]

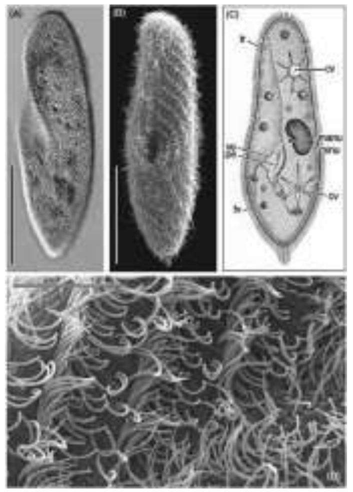

Ciliary Motion
[Valentine et al. 2022]

Problem

How to find efficient swimming morphologies / skills?

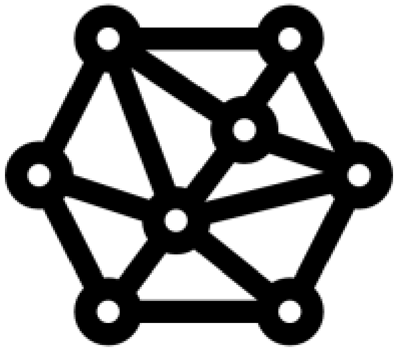

Shape & Control Modeling

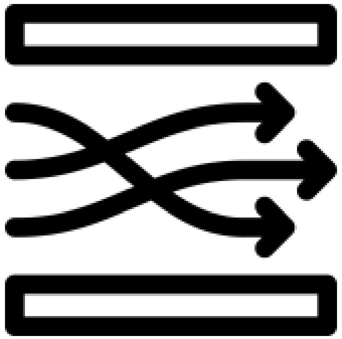

Interaction with Fluid Environment

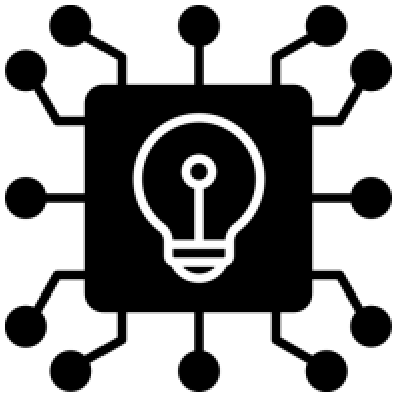

Policy Learning

Previous Research

Control

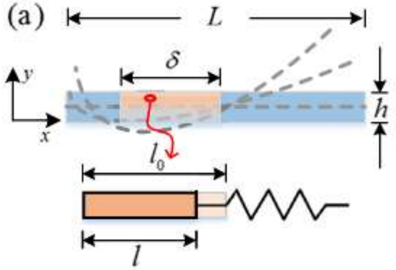

Expert
[Lin et al. 2019]

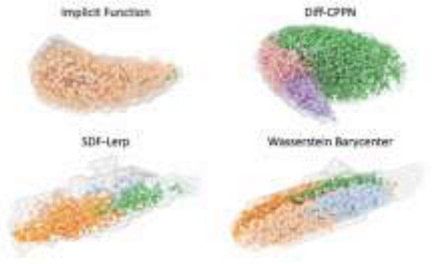

Clustering
[Wang et al. 2023]

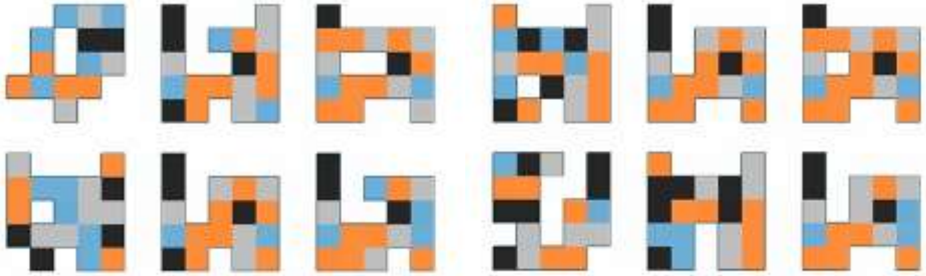

Per-Element
[Bhatia et al. 2021]

Simulation

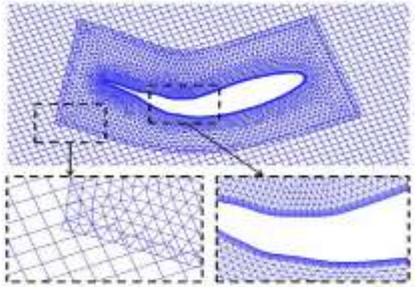

CFD
[Zhang et al. 2022]

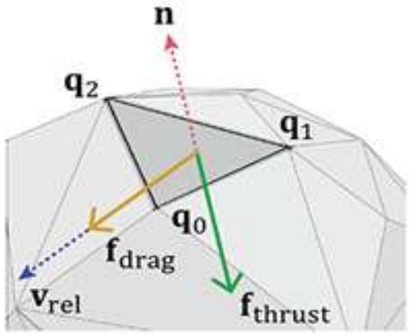

Lift & Drag
[Min et al. 2019]

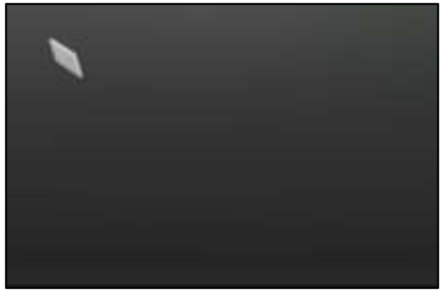

Flow Maps
[Chen et al. 2024]

Methodology

- Unified Representation for Shape & Control
- Physics-Based Simulation
- AI-Driven Controller Training

Methodology

- Unified Representation for Shape & Control
- Physics-Based Simulation
- AI-Driven Controller Training

Reduced Control Modeling

- Reduction on control space dimension
- Agnostic to morphology and robust to discretization
- Capable of representing diverse and characteristic deformations

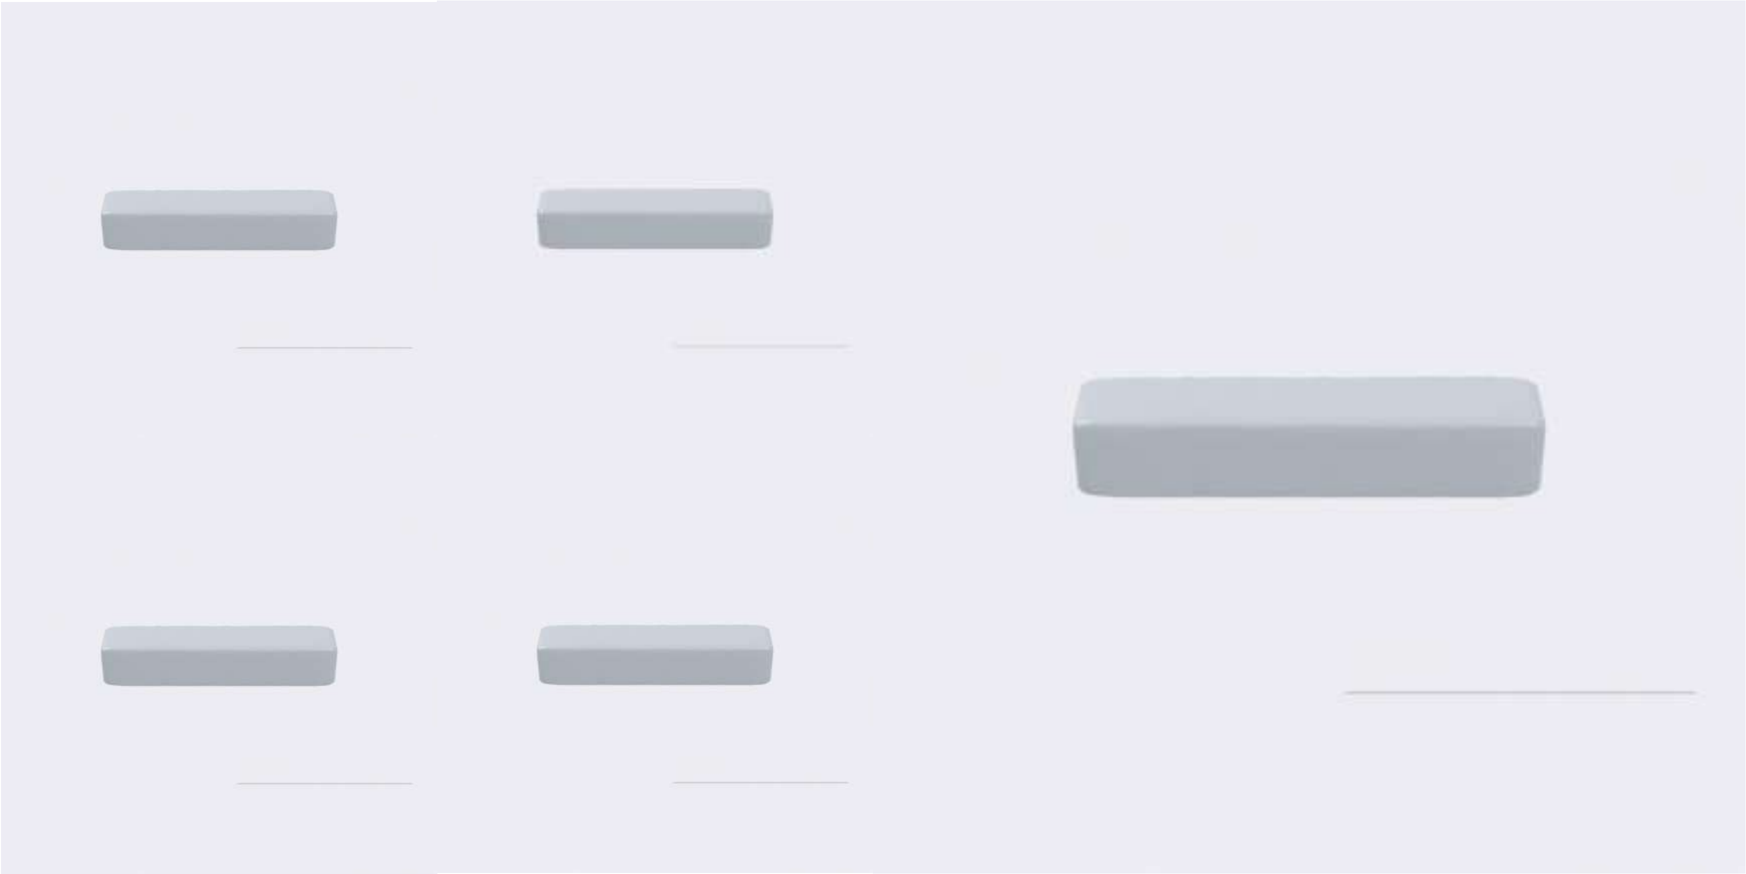

Farthest Point Sampling

1. High Dimensional Volume Mesh

Farthest Point Sampling

1. High Dimensional Volume Mesh

2. Iterative Sampling via Geodesic Distance

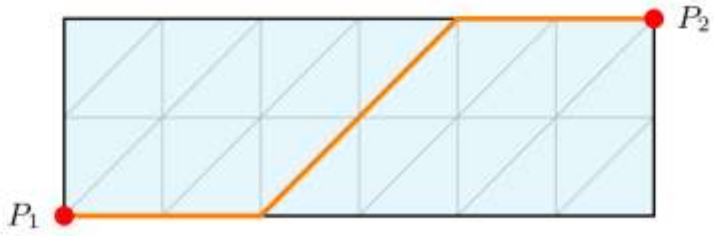

Path represents the shortest distance between points traveling through the **interior edges** of the volume mesh.

Farthest Point Sampling

1. High Dimensional Volume Mesh

2. Iterative Sampling via Geodesic Distance

Path represents the shortest distance between points traveling through the **interior edges** of the volume mesh.

Farthest Point Sampling

1. High Dimensional Volume Mesh

2. Iterative Sampling via Geodesic Distance

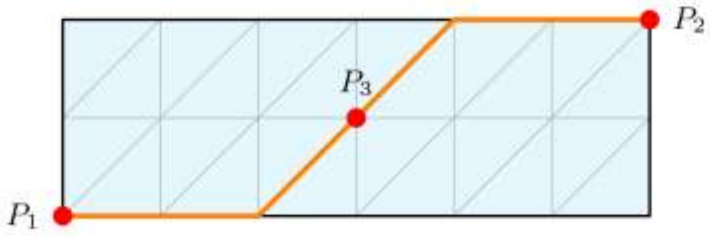

Path represents the shortest distance between points traveling through the **interior edges** of the volume mesh.

3. Influence Weights Precomputation

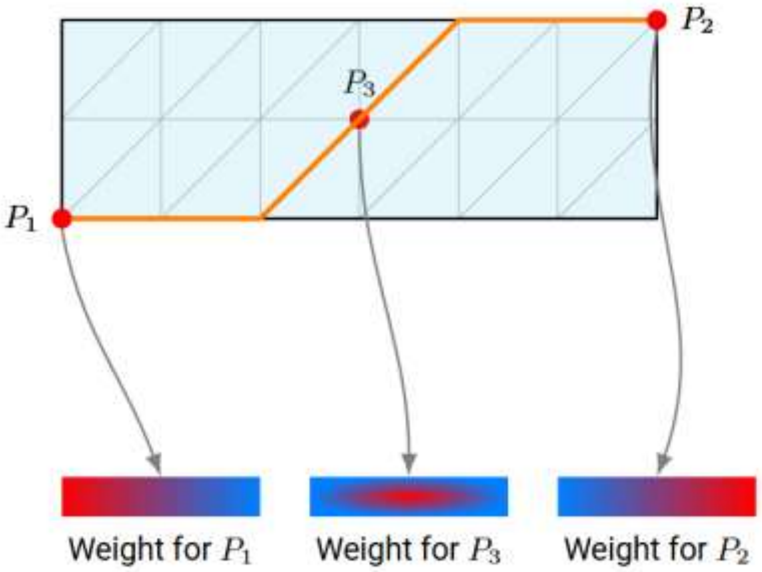

Kinematic Displacement (LBS)

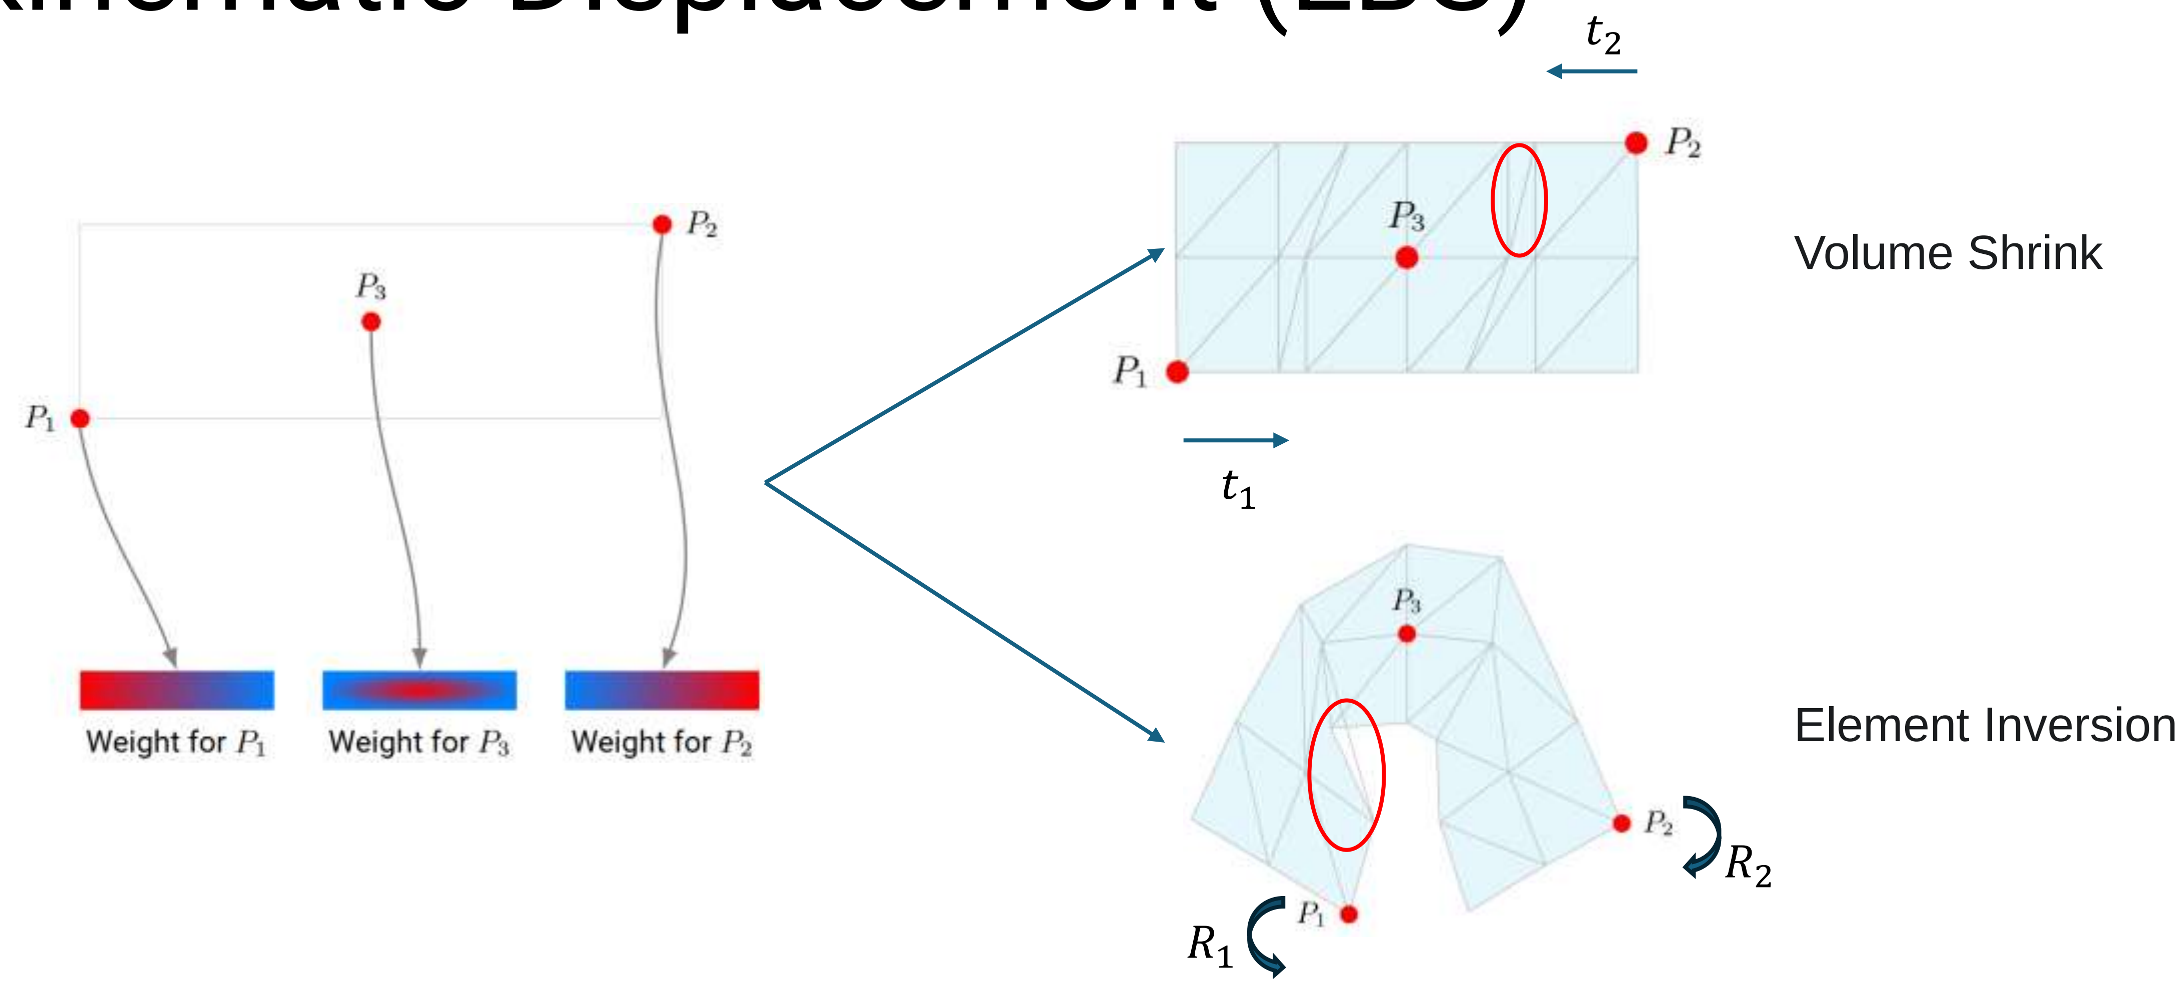

$$\boldsymbol{u}_k = \sum_{i=0}^{m} w_i(\boldsymbol{X}, \boldsymbol{p}_i)(\boldsymbol{R}_i(\boldsymbol{x}_k - \boldsymbol{c}) + \boldsymbol{t}_i + \boldsymbol{c})$$

Dynamic Correction

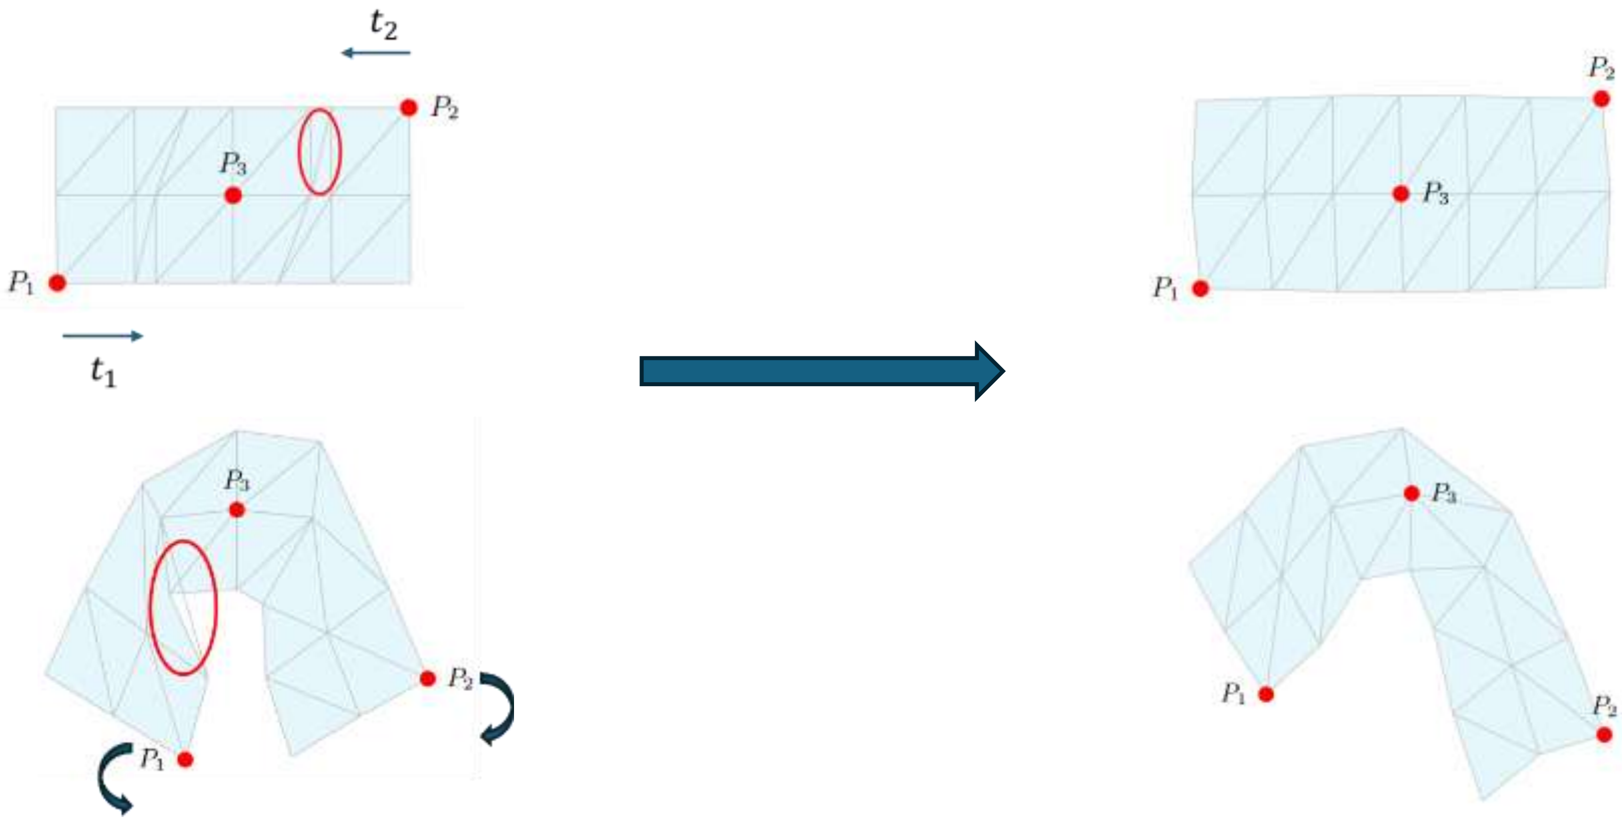

$$\underset{\boldsymbol{x}_d}{\operatorname{argmin}}\, E = \underbrace{\Psi(\boldsymbol{x}_d)}_{E_{elastic}} + \alpha \underbrace{\sum_i \|\boldsymbol{x}_d - \boldsymbol{x}_{lbs}\|_2^2}_{E_{LBS}}$$

Methodology

- Unified Representation for Shape & Control
- Physics-Based Simulation
- AI-Driven Controller Training

Efficient and Accurate Simulator

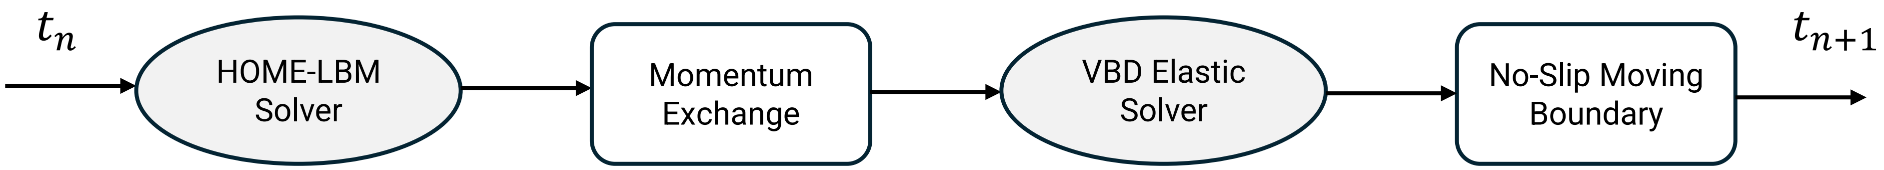

- SOTA fluid and elastic solver
  - Fully GPU-parallelized
- Weak two-way coupling scheme
  - Carefully optimized for time consuming parts (e.g. intersection detection)

	Fluid Domain Size	Elastic Body DOFs	Time per Step
2D Fish	1200x400	~300 Triangles	2.7ms
3D Clownfish	512x128x128	~3.5k Tetrahedra	18.3ms

Methodology

- Unified Representation for Shape & Control
- Physics-Based Simulation
- AI-Driven Controller Training

Unified RL Training Pipeline

- Unified state space representation

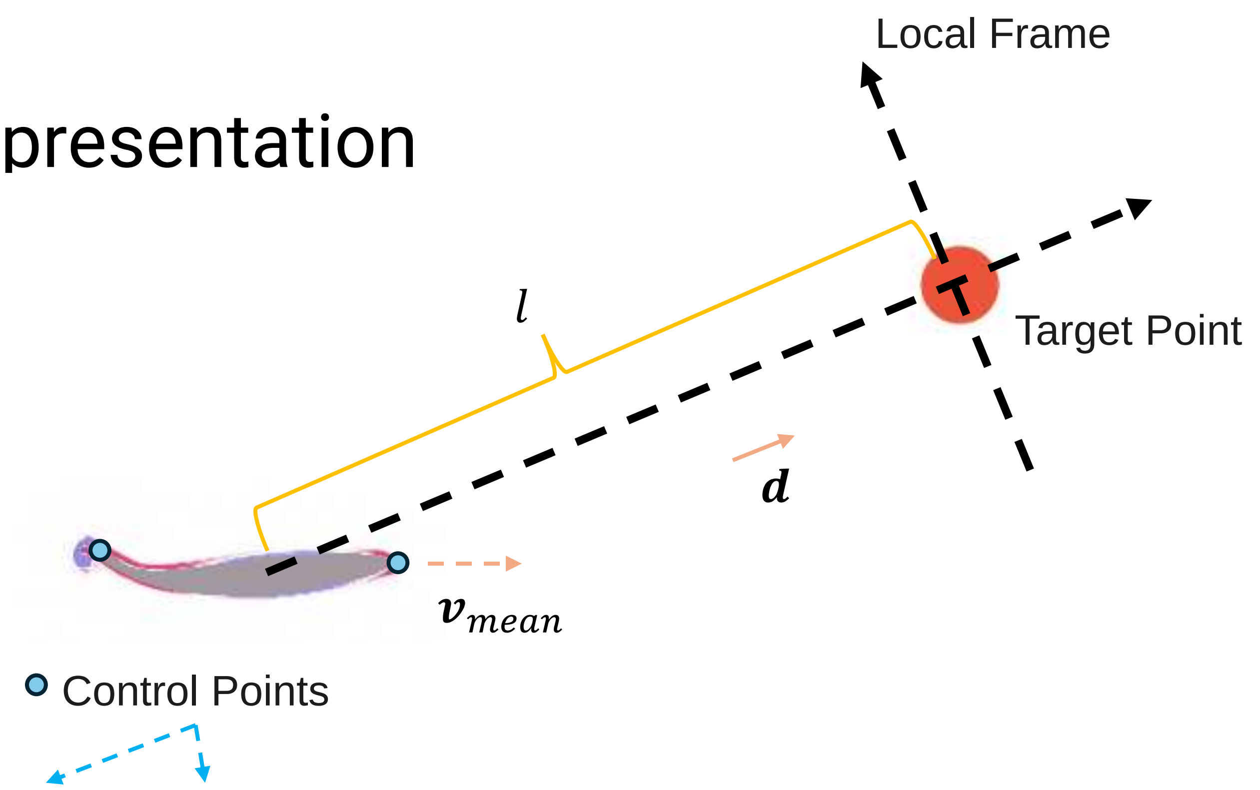

$$s = \{\boldsymbol{x}_{local}, \boldsymbol{v}_{local}, \boldsymbol{v}_{mean}, \boldsymbol{d}, l, \boldsymbol{a}_{last}\}$$

- Capture the local deformation patterns of the soft swimmer efficiently
- Ensure the learned policy is rotation- and translation-invariant by construction

Unified RL Training Pipeline

- Unified reward function
  - Use a target point to control the locomotion of swimmers

$$R = R_{\text{task}} + \lambda_{\text{smooth}} p_{\text{smooth}} + \lambda_{\text{reg}} p_{\text{reg}},$$
$$R_{\text{task}} = \boldsymbol{v}_{\text{mean}} \cdot \boldsymbol{d},$$
$$p_{\text{smooth}} = -||\boldsymbol{a} - \boldsymbol{a}_{\text{last}}||_2^2/(6m),$$
$$p_{\text{reg}} = -||\boldsymbol{a}||_2^2/(6m).$$

Swimming forward: right

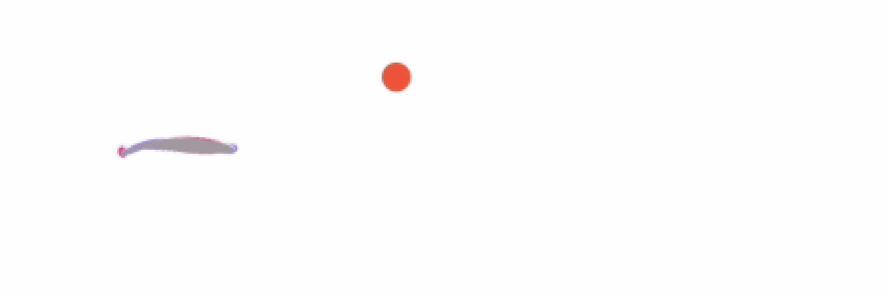

Navigation: randomly generated

Flow Resistance: center of solid

Collection of Morphologies

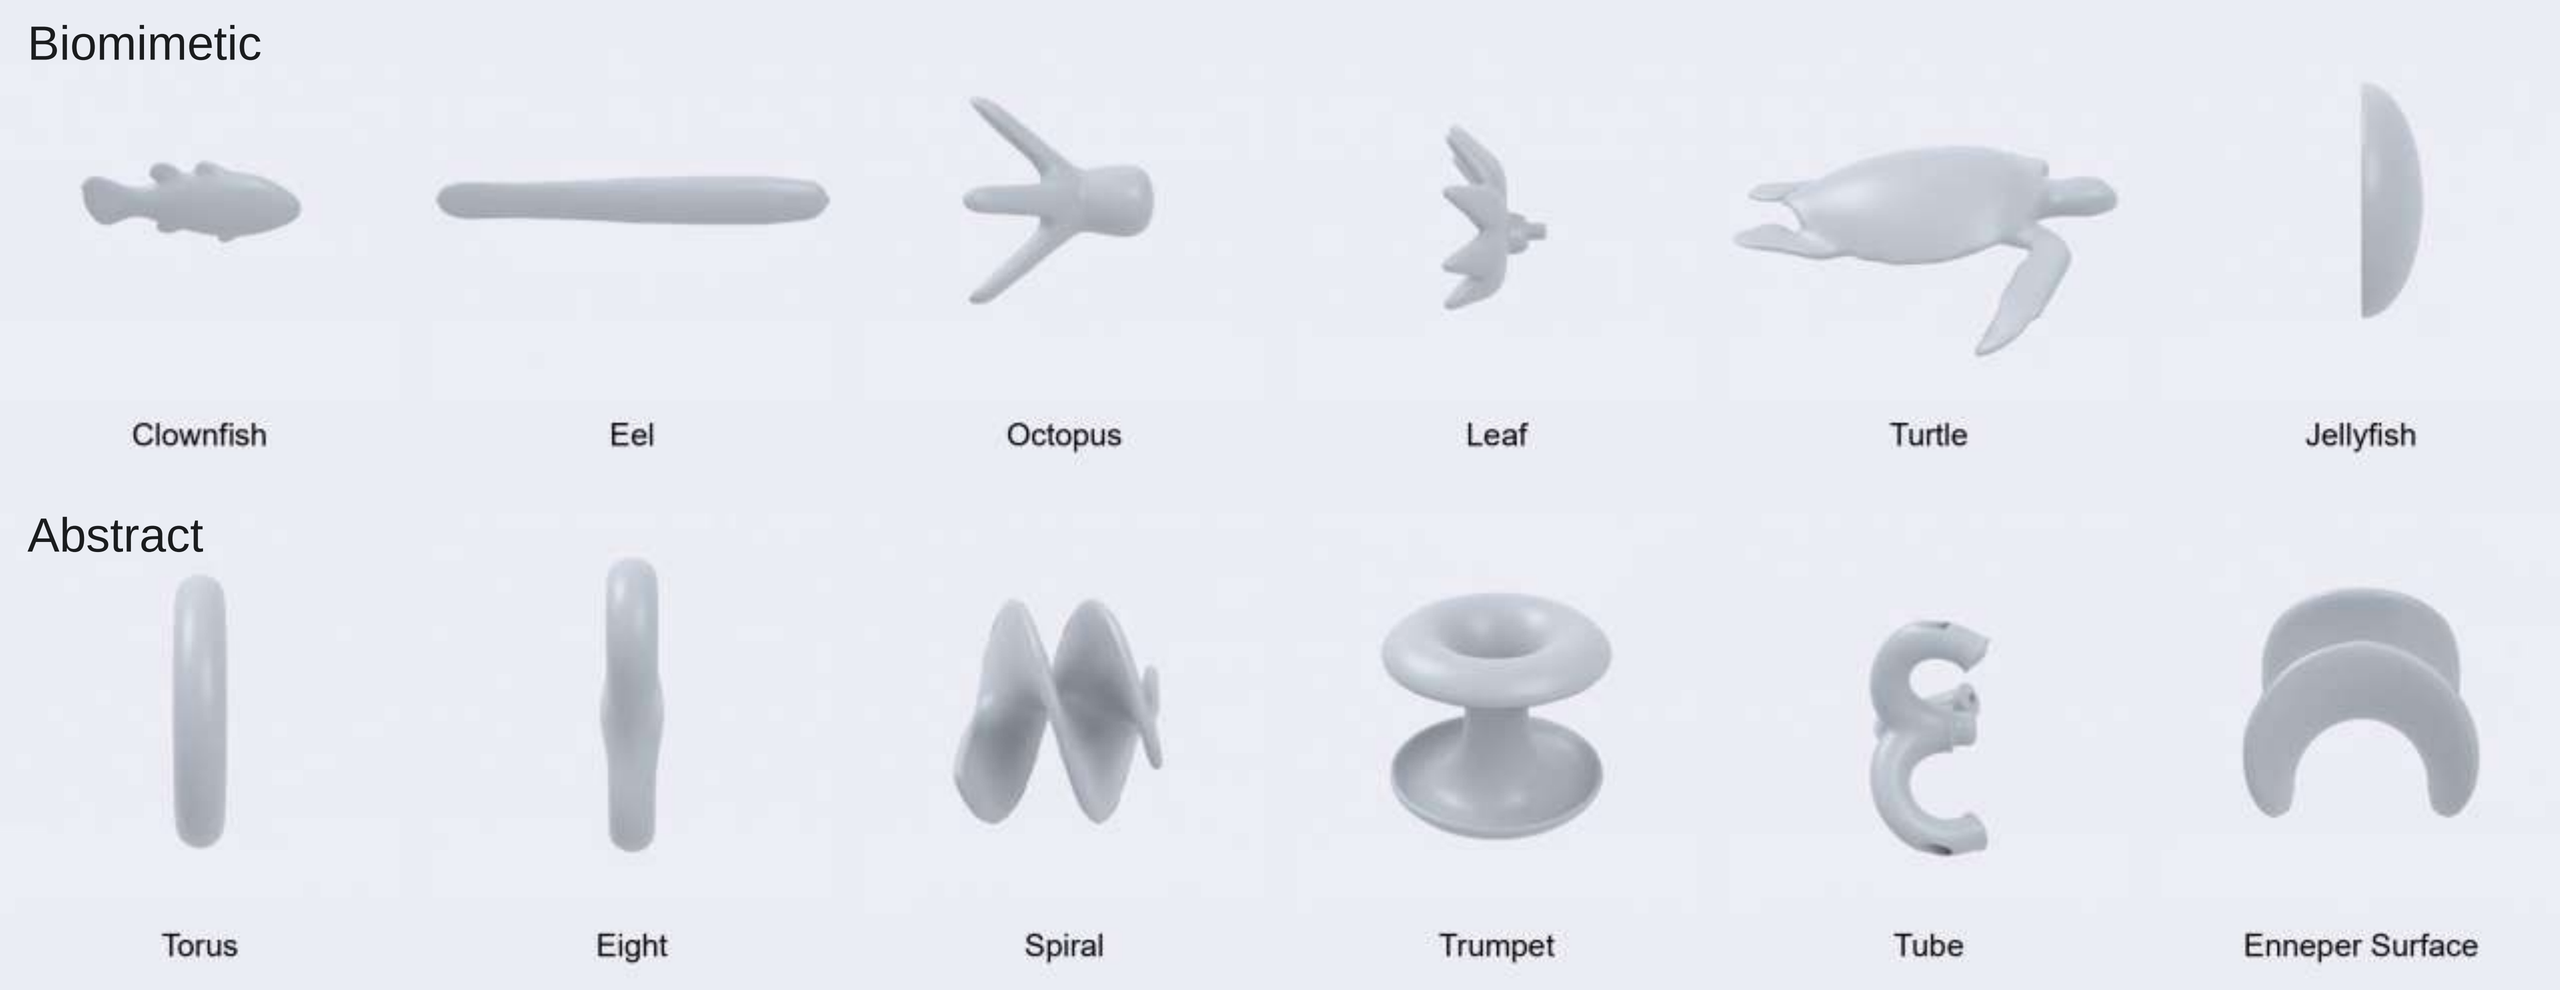

Main Results: Forward Swimming

Clownfish

Torus

Enneper Surface

Eight

Comparison with Baselines

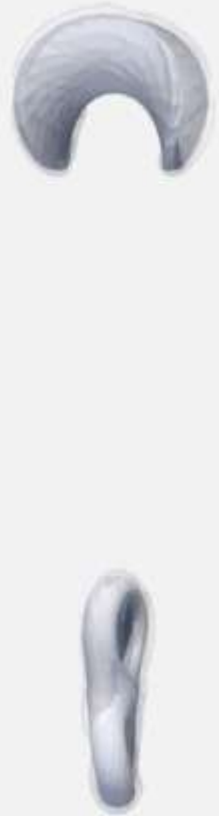

Ours

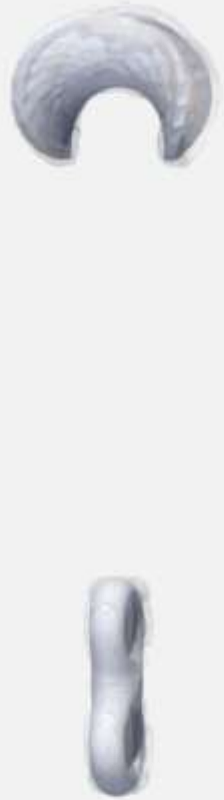

Clustering
(Adapted from Wang et al. 2023)

Main Results: Energy-efficient Swimming

w/o energy penalty

with energy penalty

Training Stability

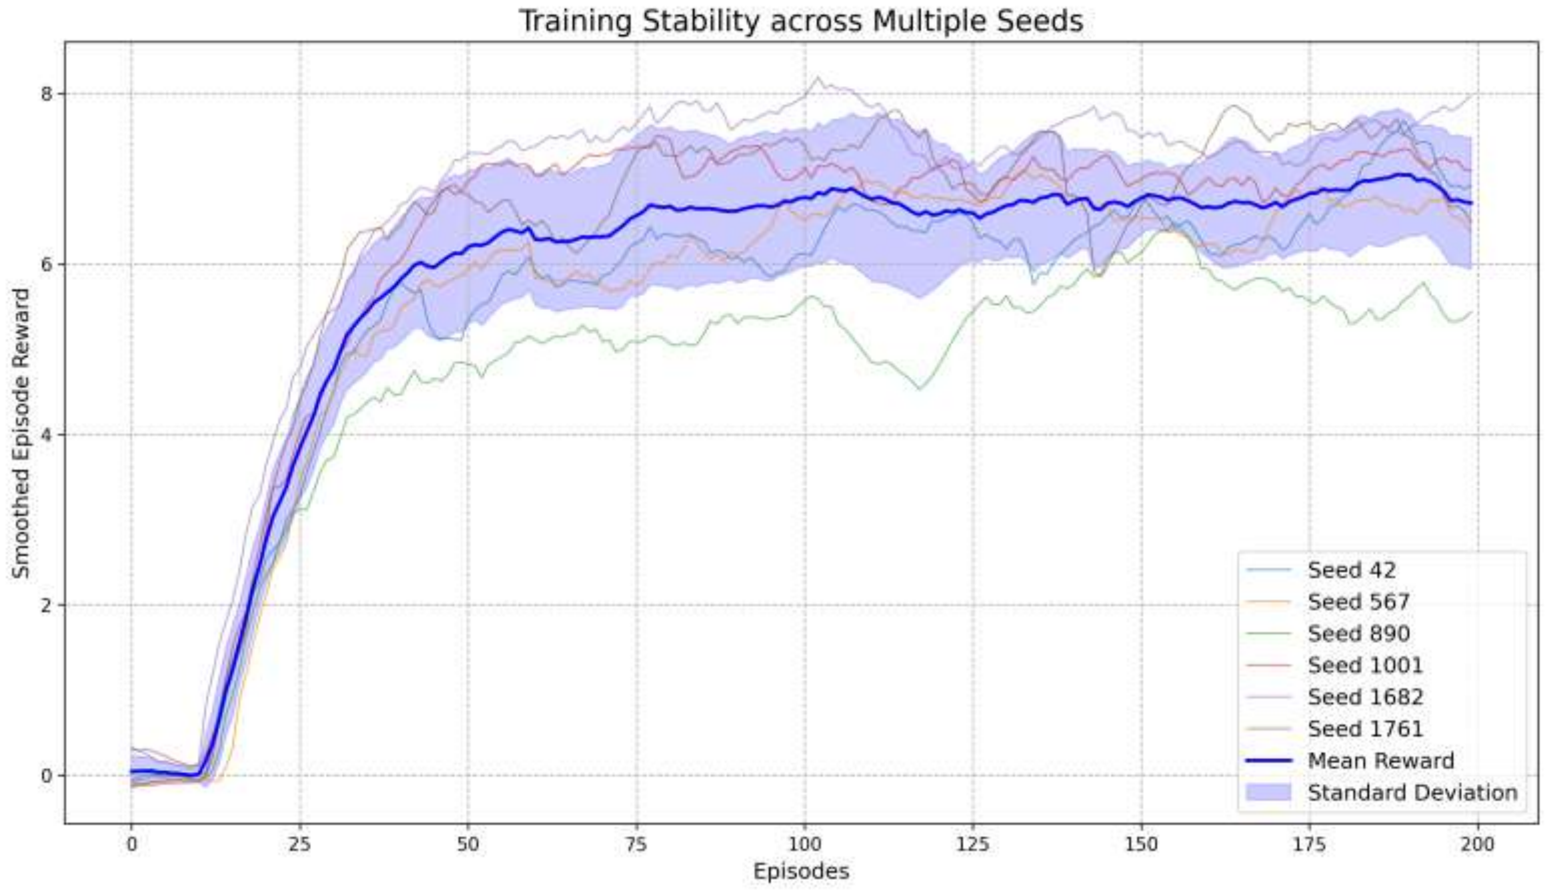

Conclusions

- We introduced a unified framework automating controller learning via:
  - Morphology-agnostic control representation
  - High-fidelity, GPU-accelerated simulator
- We demonstrated emergence of effective gaits for diverse morphologies, significantly outperforming baselines on unconventional shapes.

Thank You for Watching

- Acknowledgements

  - We would like to thank Dr. Chao Yu for her valuable advice on reinforcement learning training. Tao Du acknowledges the research funding support from Tsinghua University and Shanghai Qi Zhi Institute, and Wei Li benefits from SJTU's startup funds.

- Contact

  - hucy24@mails.tsinghua.edu.cn

Webpage