

AnyMorph: Learning Transferable Policies By Inferring Agent Morphology



Brandon Trabucco ¹, Mariano Phielipp ^{* 2}, Glen Berseth ^{* 3}

¹ Carnegie Mellon University, ² Intel AI, ³ Mila, ^{*} Equal Advising



Machine
Learning
Department


Carnegie Mellon University
School of Computer Science

intel[®] ai




Controlling Many Morphologies With One Policy


- **Agent-Agnostic RL** aims to learn **multitask policies** that generalize to new agent morphologies.




Why Is Controlling Many Morphologies Hard?




Why Is Controlling Many Morphologies Hard?



Why Is Controlling Many Morphologies Hard?




Why Is Controlling Many Morphologies Hard?




- Morphologies can have **incompatible** designs and abilities.

Why Is Controlling Many Morphologies Hard?



- Morphologies can have **incompatible** designs and abilities.
- The policy **must infer** which morphology its controlling.

Why Is Controlling Many Morphologies Hard?

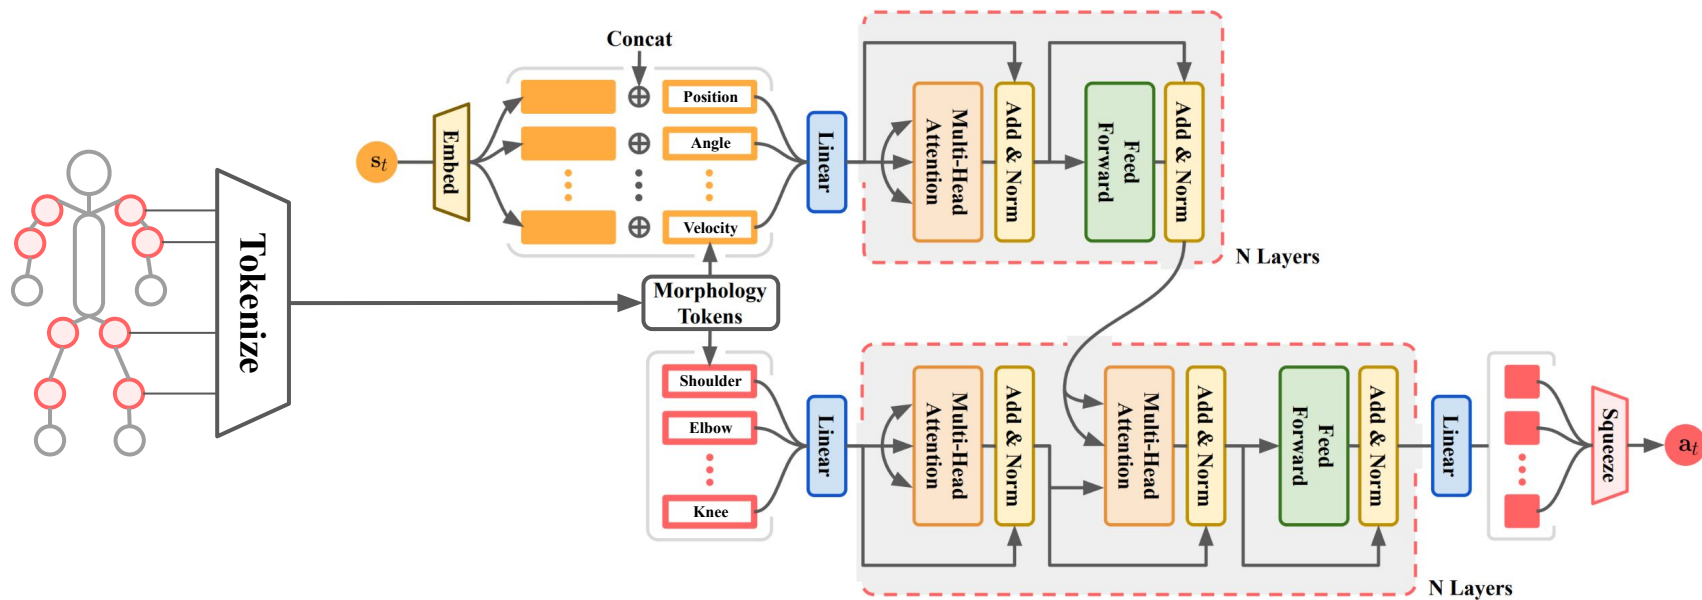


- Morphologies can have **incompatible** designs and abilities.
- The policy **must infer** which morphology its controlling.


Q How do we specify the morphology to the policy?

AnyMorph: A Simple Encoding For Any Morphology

Key Idea: Encode Morphology Via A **Sequence Of Discrete Tokens**.

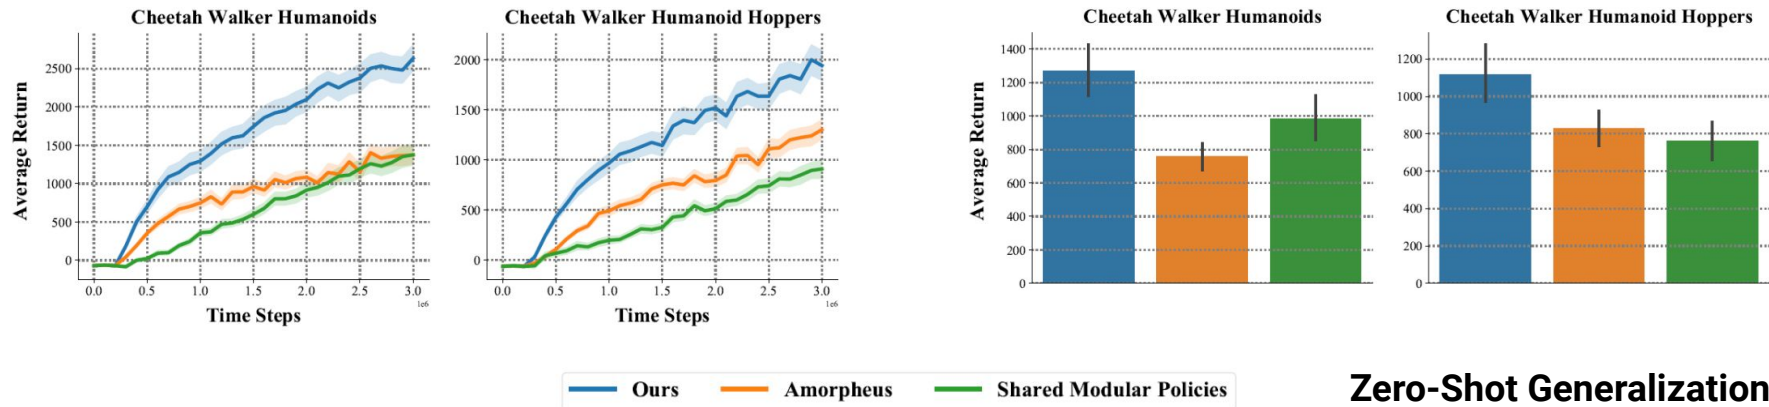


AnyMorph Improves Sample Efficiency



- Outperforms prior work despite needing **less information** about the agent.
- Leads to an improvement of **+53%** when trained on 32 morphologies.



AnyMorph Improves Zero-Shot Generalization



Zero-Shot Generalization

- Outperforms prior work despite needing **less information** about the agent.
- Leads to an improvement of **+53%** when trained on 32 morphologies.
- Improves zero-shot generalization by **+32%** on the largest task.

AnyMorph Improves Robustness To Breakage



Automatically discovers which sensors can be ignored!

- **Maintains 50% original performance** when 50% of sensors are corrupted.
- Prior work **drops to 13% original performance** in the same conditions.

Concluding Remarks

- We propose a simple yet flexible encoding for any morphology.
- Improves zero-shot generalization by **+16%** overall.
- Robust when the agent's sensors break.

Open Problems & Future Work:

- Translating AnyMorph to visual tasks.
- Scaling up the number of morphologies.
- Prompt Engineering for morphology tokens.

Meet The Authors:



Brandon Trabucco



Mariano Phielipp



Glen Berseth

Find out more at: <https://sites.google.com/btrabucco.com/anymorph-icml2022>

Thanks For Listening!