Learning Behavior Representations Through Multi-Timescale Bootstrapping

Mehdi Azabou[1], Michael Mendelson[1], Maks Sorokin[1], Shantanu Thakoor[2], Nauman Ahad[1], Carolina Urzay[1], Eva L. Dyer[1]
[1] Georgia Institute of Technology [2] Deepmind

Motivation

Animal behavior unfolds over multiple timescales

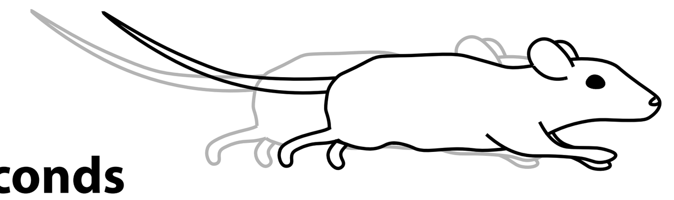

seconds
running, grooming or chasing

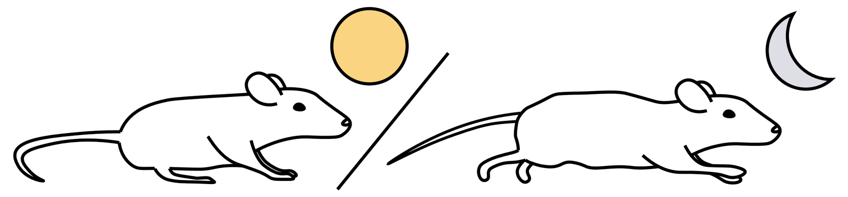

hours
mice are night active
changes based on time of day

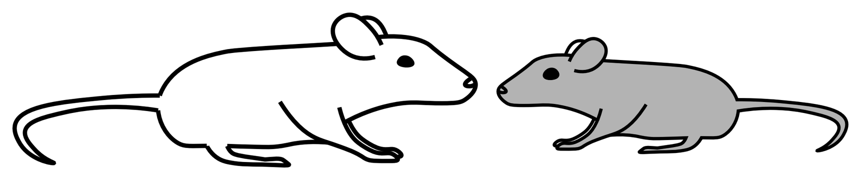

days - ∞
age, disease progression
strain, morphology

Simulated legged robots

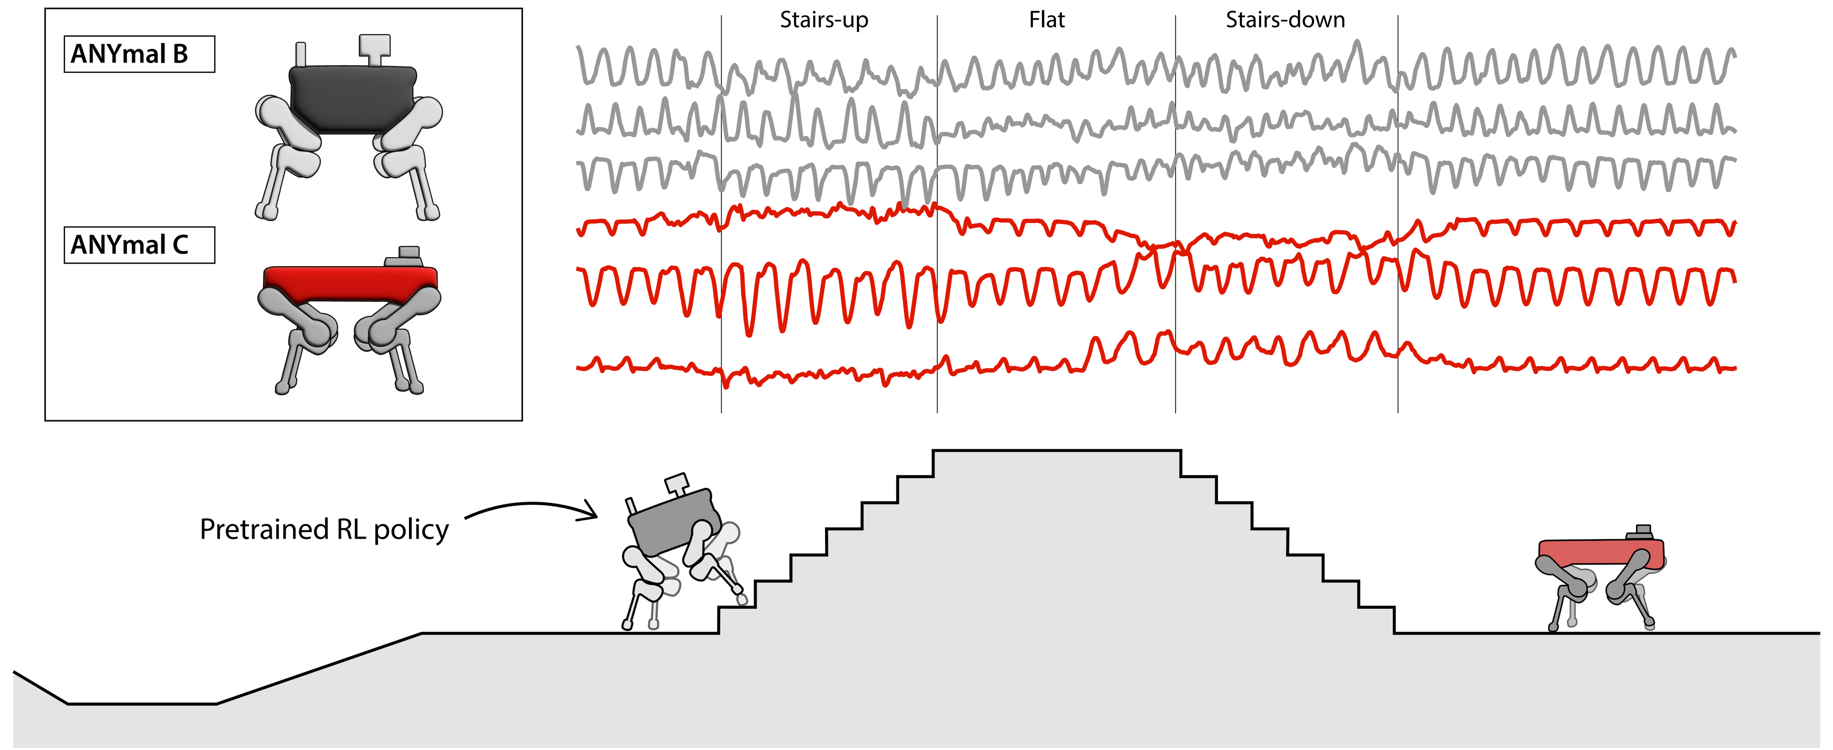

Simulation: NVIDIA's Isacc Gym simulator + Pretrained RL agents
We simulate multiple robots walking on procedurally generated terrains

Data: We track 12 keypoints representing the robot's joints

Two factors modulate the behavior:
- Terrain type: Flat, Stairs or Slope (**Short-term**)
- Robot type: Robot's morphology influences movement (**Long-term**/Global)

Model	F1-score	
	Terrain type	Robot type
Short-term + Long-term	**0.73**	0.98
Short-term only	0.50	0.86
Long-term only	0.62	**0.99**

BAMS: Bootstraping Across Multiple Scales

Pull the embeddings of neighboring timepoints closer to each other.

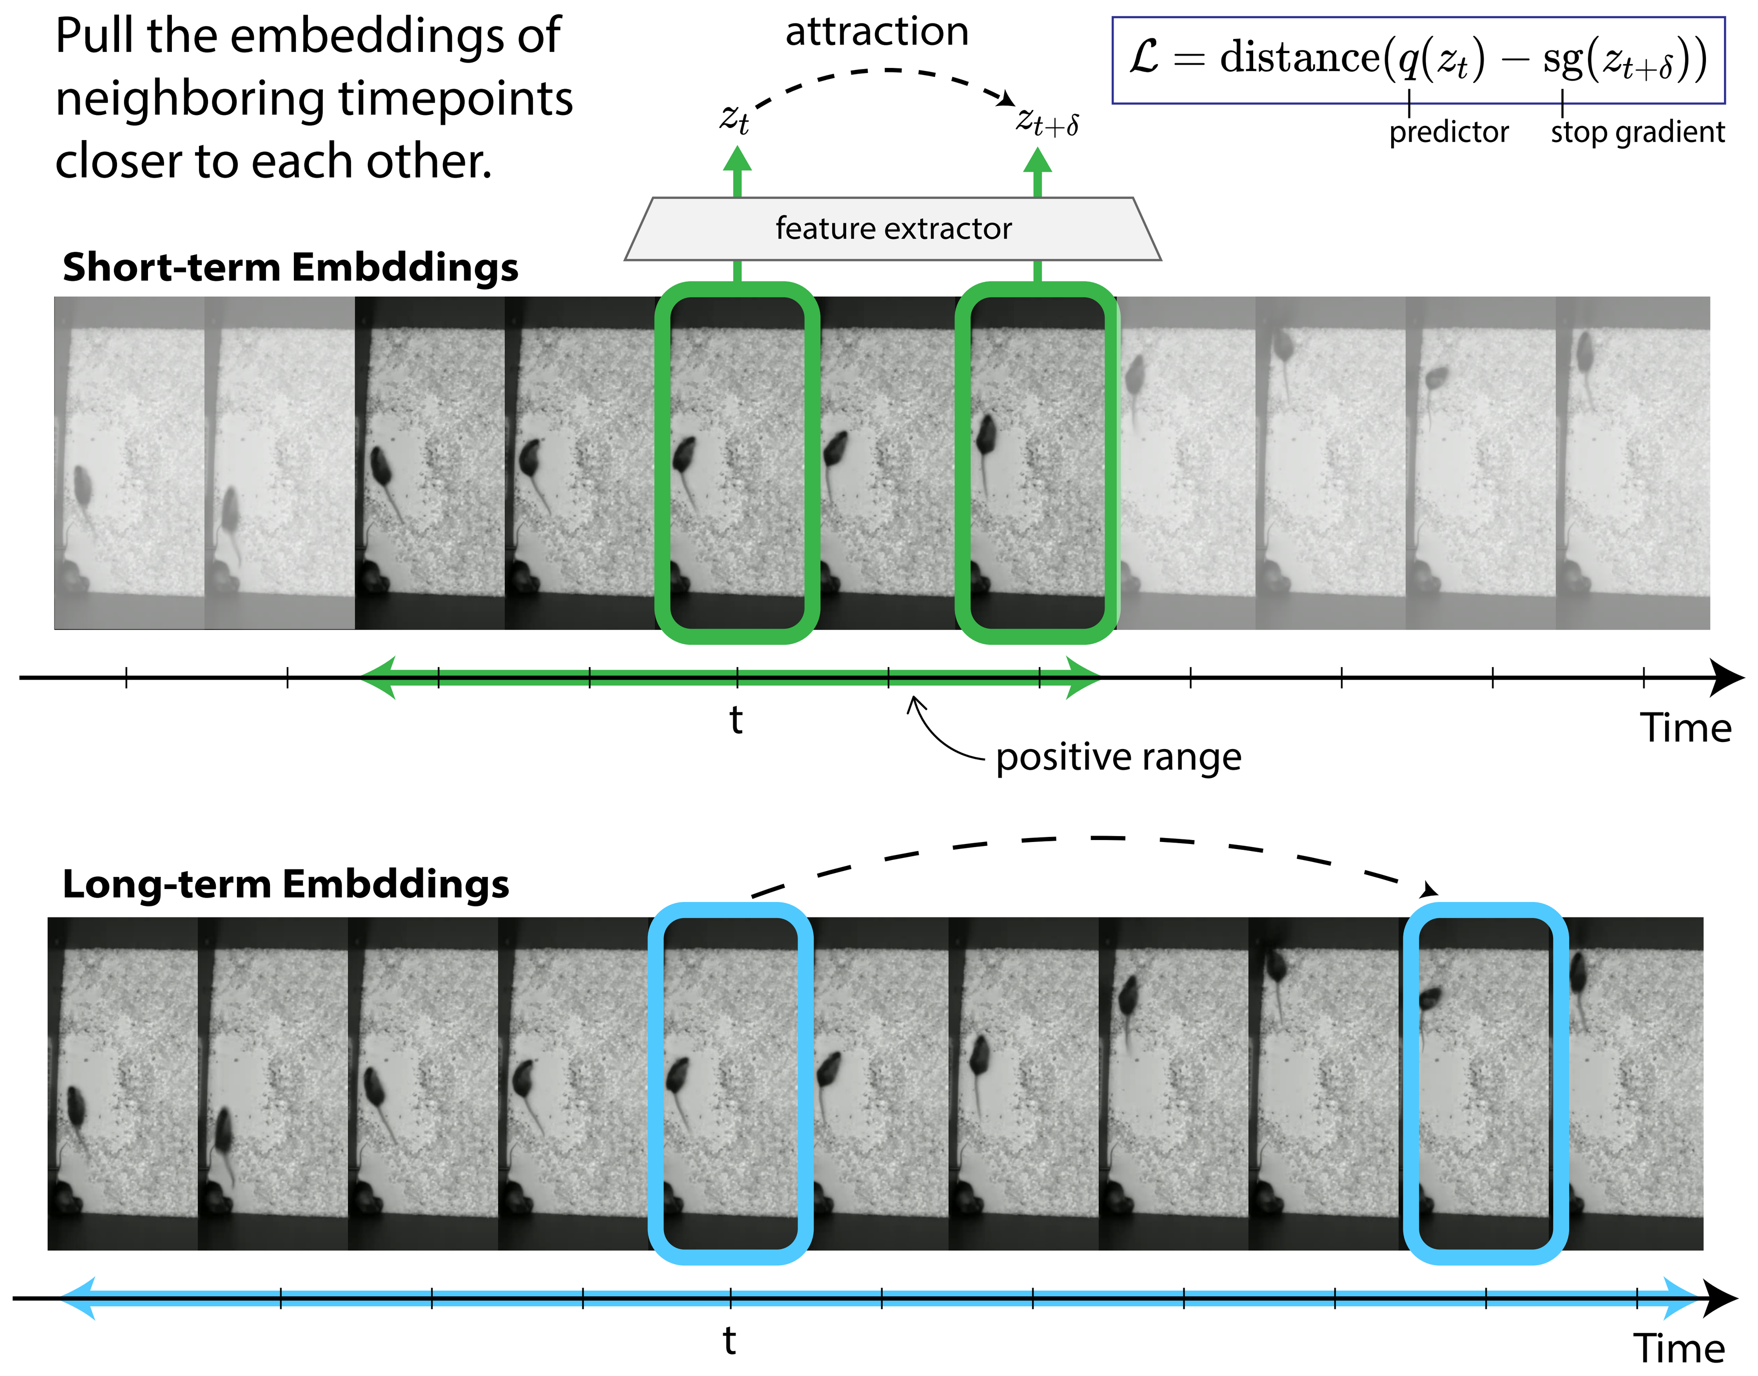

Temporal Pyramid Pooling Module
Multiple TCNs with different receptive fields:
- Recent past encoder (**sub-sec**)
- Short-term encoder (**seconds**)
- Long-term encoder (**minutes - hours**)

Building block: Temporal Convolutional Networks. TCNs have tunable receptive fields (receptive field = observable history)

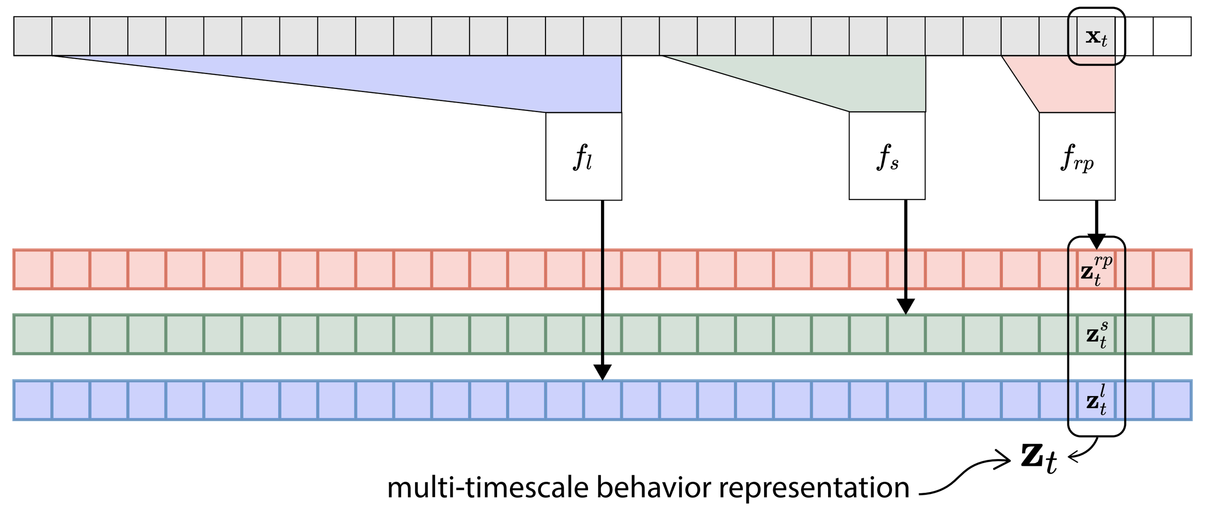

Objective 1: Latent Predictive Learning
- Use different positive ranges for short-term embeddings and long-term embeddings

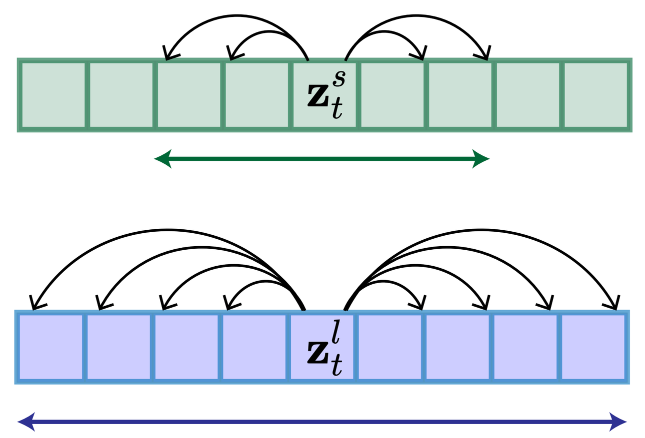

- Bootstrap the embeddings in each latent space, inducing different rates of smoothness

Objective 2: Pretext task
- Given the multi-timescale embedding, predict future actions using a multi-step predictor

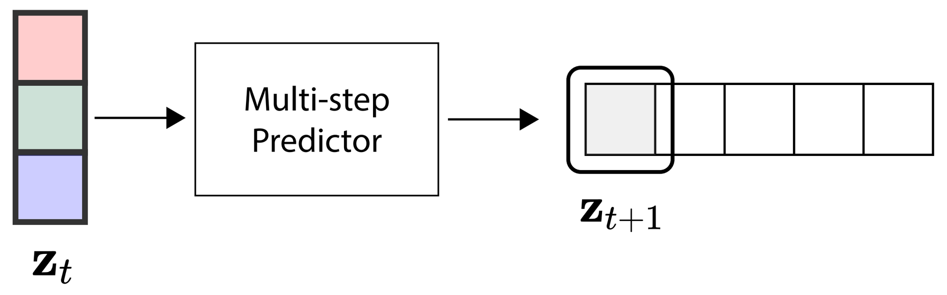

MABe Mouse Triplets Behavior Dataset

Dataset: Mouse Triplets - Multi Agent Behavior Challenge 2022
Task: Unsupervised learning of behavior embeddings from tracking data
Evaluation: To evaluate the representational quality of the learned embeddings, we perform linear readout of 13 sets of labels, including time of day, chasing behavior, and mouse strain

Model	F1-score	MSE	Sequence-level subtasks				Frame-level subtasks								
			T1*	T2*	T3	T13	T4	T5	T6	T7	T8	T9	T10	T11	T12
# 1	**30.3**	**0.09296**	0.09019	**0.09523**	**82.20**	69.40	1.90	1.24	**71.62**	55.52	**30.20**	0.40	1.63	1.10	**20.45**
# 2	28.3	0.09289	0.09057	**0.09513**	67.20	66.90	**2.70**	**6.60**	71.47	**54.67**	20.30	**0.68**	**3.31**	**2.37**	18.54
# 3 BAMS (Ours)	28.4	0.09298	**0.09037**	0.09513	67.10	**69.50**	2.16	2.31	66.42	53.28	30.18	0.45	1.65	1.14	19.14
PCA baseline	7.99	0.09430	0.09415	0.09449	33.83	4.13	0.00	0.00	12.69	0.08	0.00	0.00	0.00	0.00	0.00

BAMS achieves competitive results. Our analysis shows that some factors are exclusively captured in either the short-term or the long-term embedding, or require both timescales to be revealed

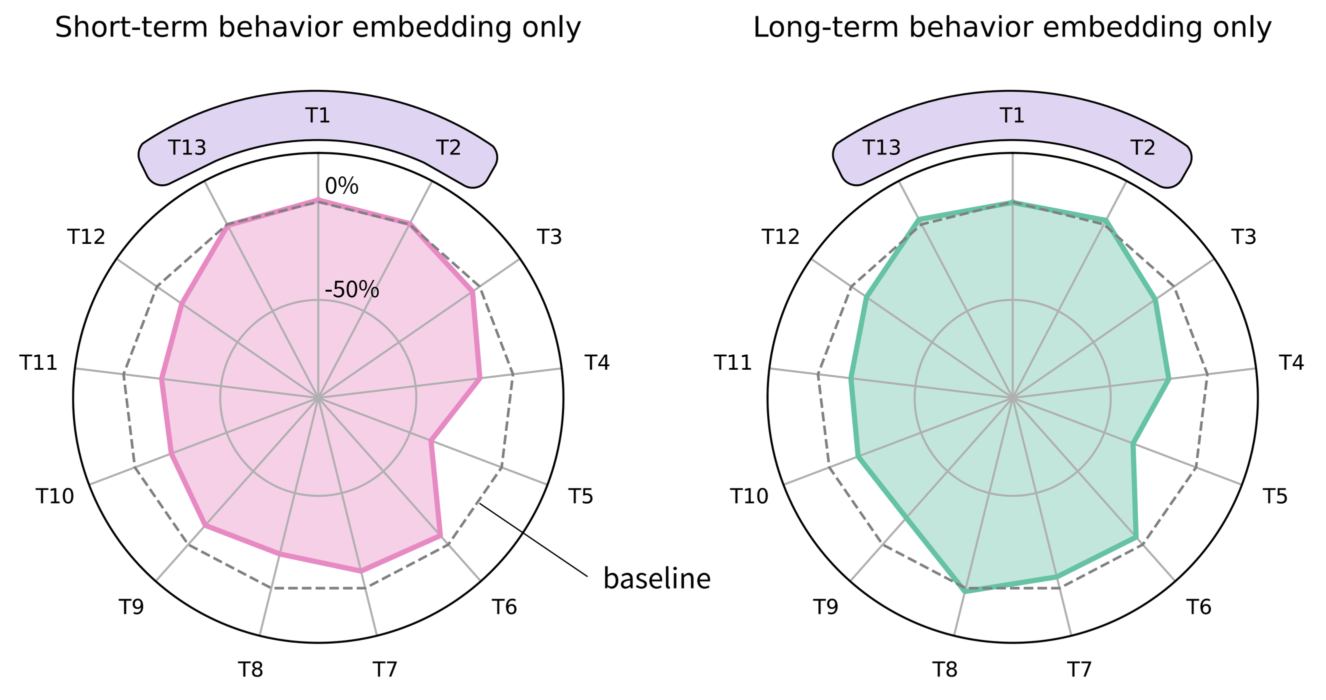

Relative decrease of accuracy from baseline