



Motivation

Could we learn a single vision-based policy for agile parkour?




Challenges

Vision Locomotion:

- The action is noisy and laggy.
- The camera has artifacts, latency and jitter.

Agile Parkour:

- Extreme motion needs precise control.
- The heading should be adjusted by robot itself.
- Walking in different styles is still under exploration.



Unified Rewards

- Velocity Tracking Reward:** Encourage to track heading

$$r_{tracking} = \min(\langle \mathbf{v}, \hat{\mathbf{d}}_w \rangle, v_{cmd}) \quad \hat{\mathbf{d}}_w = \frac{\mathbf{p} - \mathbf{x}}{\|\mathbf{p} - \mathbf{x}\|} \text{ is target direction.}$$

- Clearance Reward:** Penalize dangerous footprints

$$r_{clearance} = - \sum_{i=0}^4 c_i \cdot M[p_i]$$

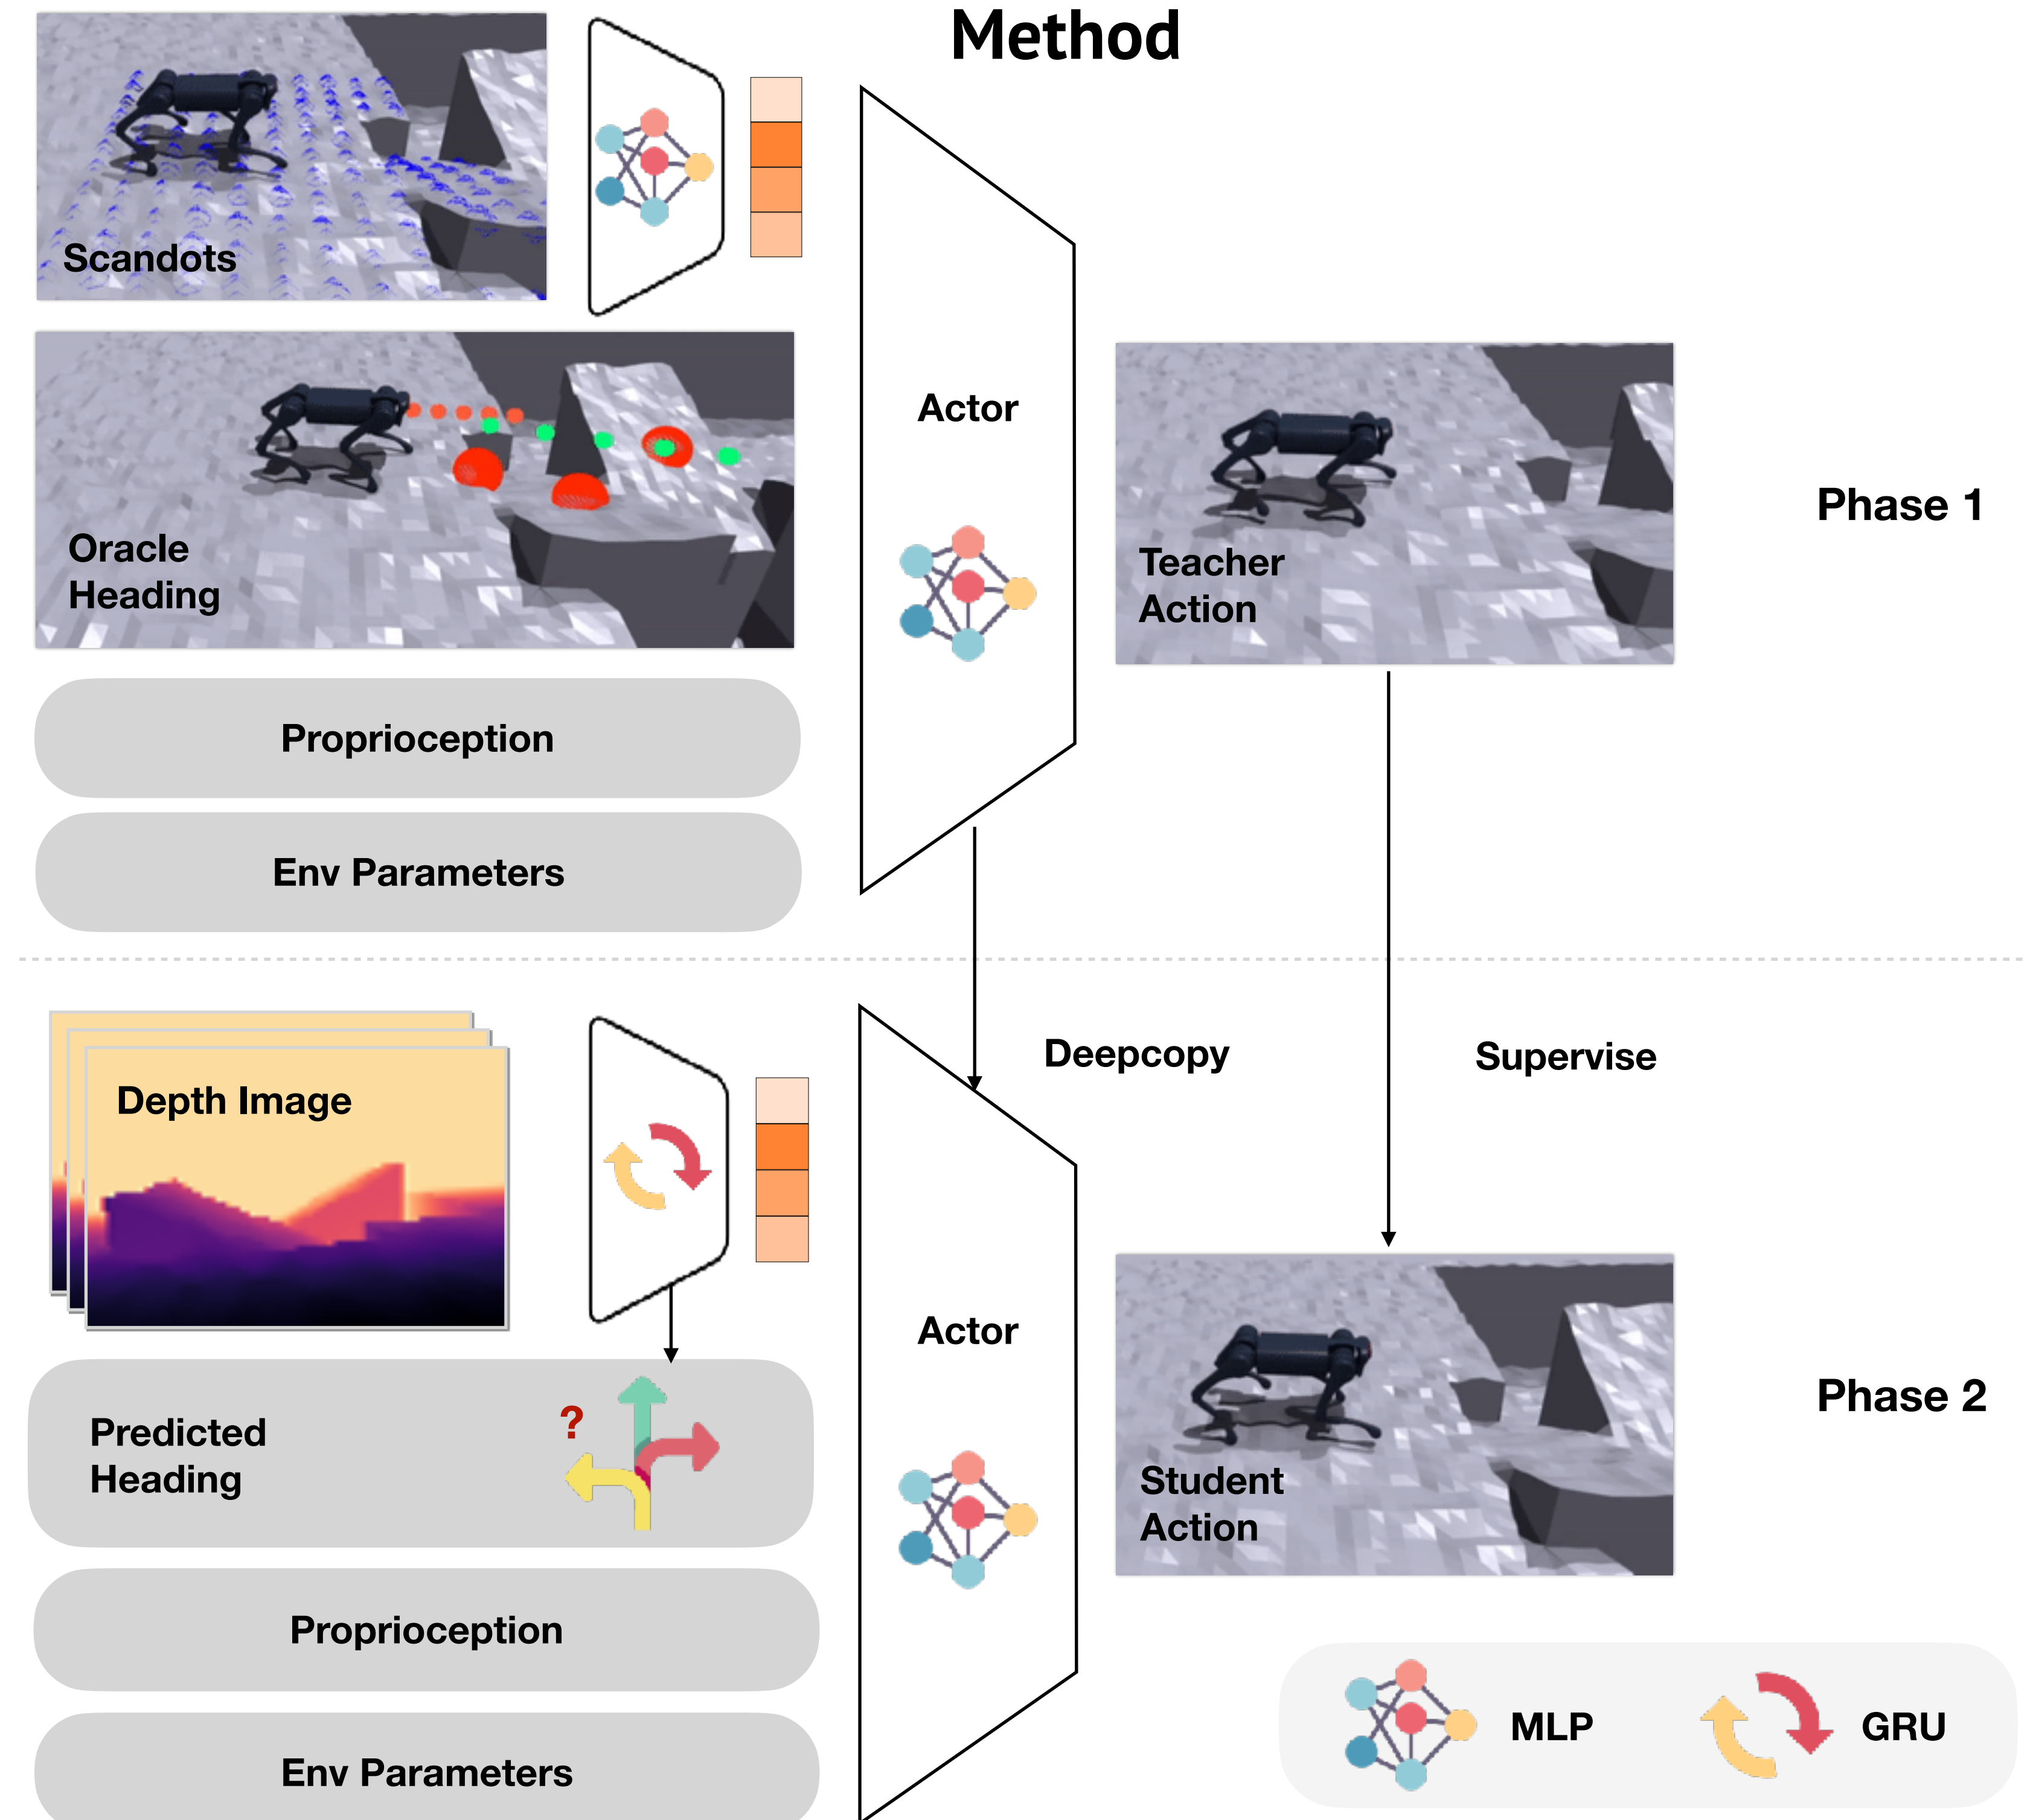
where M is a boolean function which is 1 iff the robot lies within 5cm of an edge.

- Stylized Reward:** Encourage handstand walk

$$r_{stylized} = W \cdot [0.5 \cdot \langle \hat{\mathbf{v}}_{fwd}, \hat{\mathbf{c}} \rangle + 0.5]^2$$

is sampled randomly in $\{0, 1\}$ in training and controlled via joystick in deployments.

Method



Results

