Hierarchical Decomposition of Terrain Adaptation for Quadrupedal Locomotion

From Hierarchical Decomposition to Role-Specialized Adaptation

Ji-Hyeon Yoo, Kartikeya Singh, Karthik Dantu

jihyeony@buffalo.edu

Motivation

① Heterogeneous terrain creates multiple failure sources

- The same locomotion task can lead to different outcomes across terrain

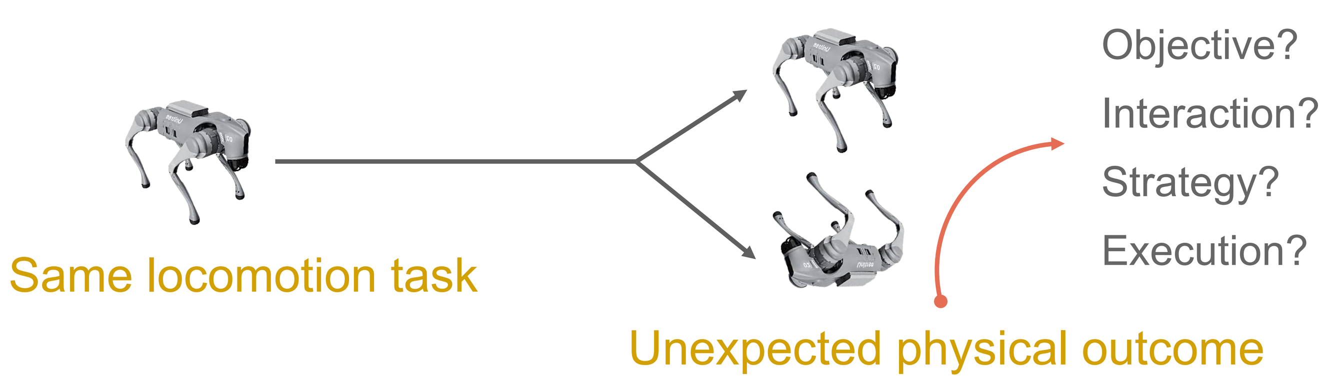

② Terrain adaptation can occur at multiple levels of the locomotion system.

- Existing approaches adapt different components, such as objectives, motion strategies, or low-level execution.
- **Which component should adapt to a given terrain-induced failure?**

③ Reward conditioning does not guarantee the desired physical outcome

- Increasing a reward weight does not necessarily improve the physical quantity with the same label.

Contribution

Hierarchical Decomposition

- We separate terrain adaptation into functional roles: **Objective, Interaction, Strategy, Execution**

Explicit Interface Semantics

- Distinguish η, β, locomotion strategy Θ, and J to localize mismatches across the hierarchy

Interface-Level Diagnosis

- Show that objective conditioning changes strategy **without guaranteeing the same physical outcome**

Methodology

TRACER Functional Hierarchy

** Extended in ongoing work beyond the workshop study.*

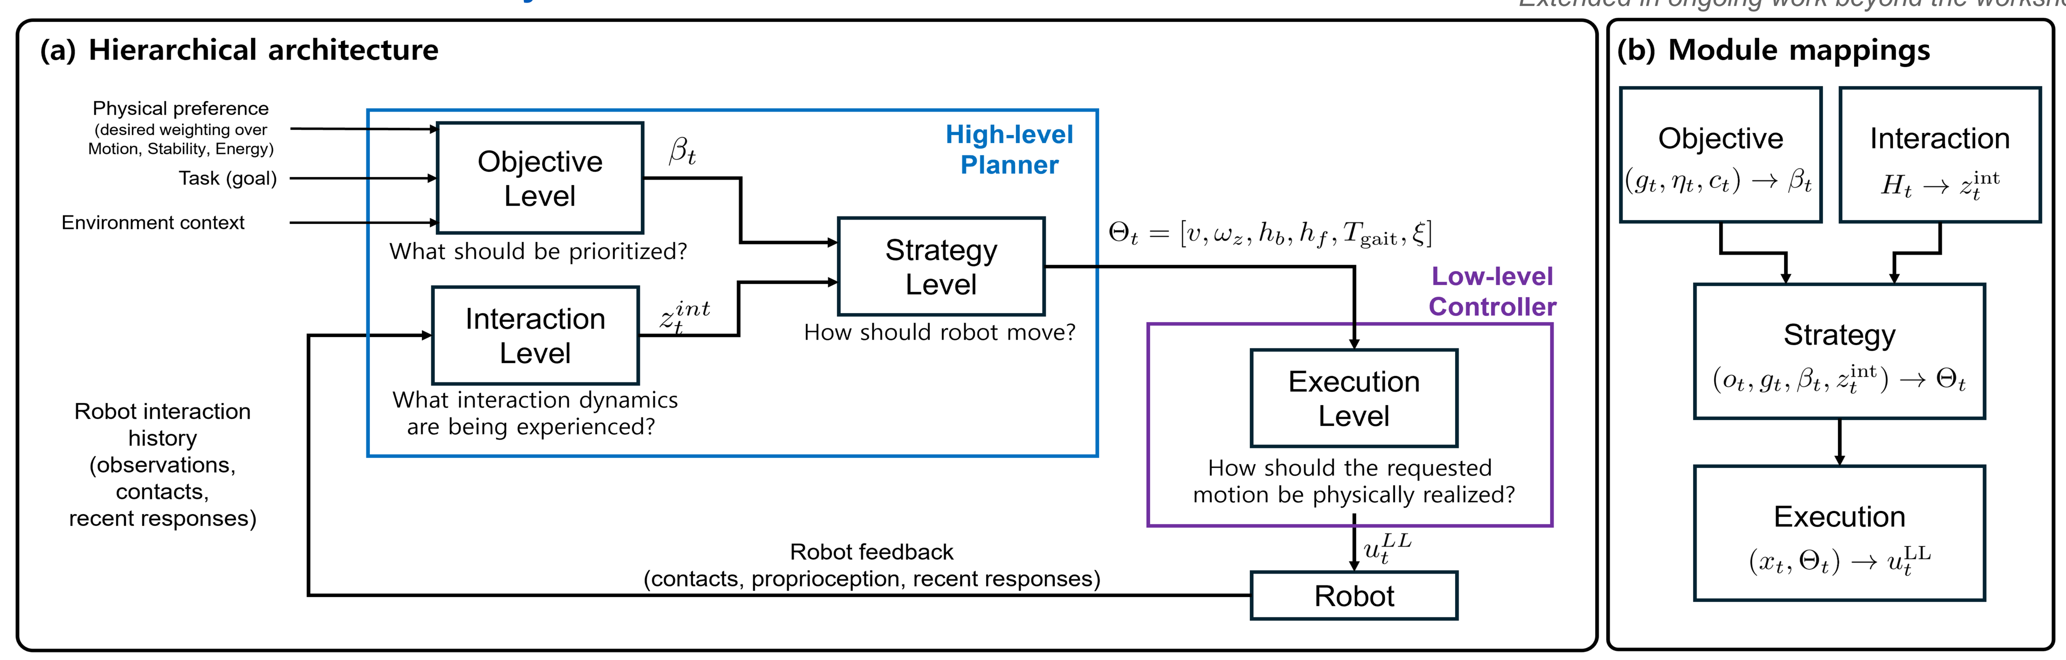

TRACER decomposes heterogeneous-terrain adaptation into functional roles: **Objective and interaction** provide distinct upstream information, **Strategy** converts them into interpretable locomotion references, and **execution** physically realizes the requested motion.

A. Objective – What should be prioritized?

Calibrates *desired physical preference* η and *context* into an internal conditioning coordinate β:

$$(g_t, \eta_t, c_t) \rightarrow \beta_t$$

η_t is also retained as an explicit semantic input to the strategy policy*.

B. Interaction – What dynamics are being experienced?

$$H_t \rightarrow z_t^{\text{int}}$$

Summarizes recent command–response history into an interaction representation z_t^{int}.

Not instantiated in the initial workshop study.

C. Strategy – How should the robot move?

Generates interpretable high-level locomotion references:

$$\Theta_t = \pi_{\text{HL}}(o_t, g_t, \beta_t, z_t^{\text{int}})$$

The full interface may contain *velocity, body height, swing clearance, gait period,* and *duty factor;* the current study uses only :

$$\Theta_t^{3D} = [v_x, \omega_z, h_{\text{body}}]$$

Semantic action path ($\eta \rightarrow \Theta$)*: adds direct preference-aware semantic authority

D. Execution – How is the strategy physically realized?

A fixed model-based low-level controller tracks the requested strategy under *robot dynamics* and *contact constraints*:

$$u_t^{LL} = C_{LL}(x_t, \Theta_t)$$

Findings & Implications

Preference-Conditioned Physical Outcomes

Task Preference Produces Distinct Locomotion Operating Points

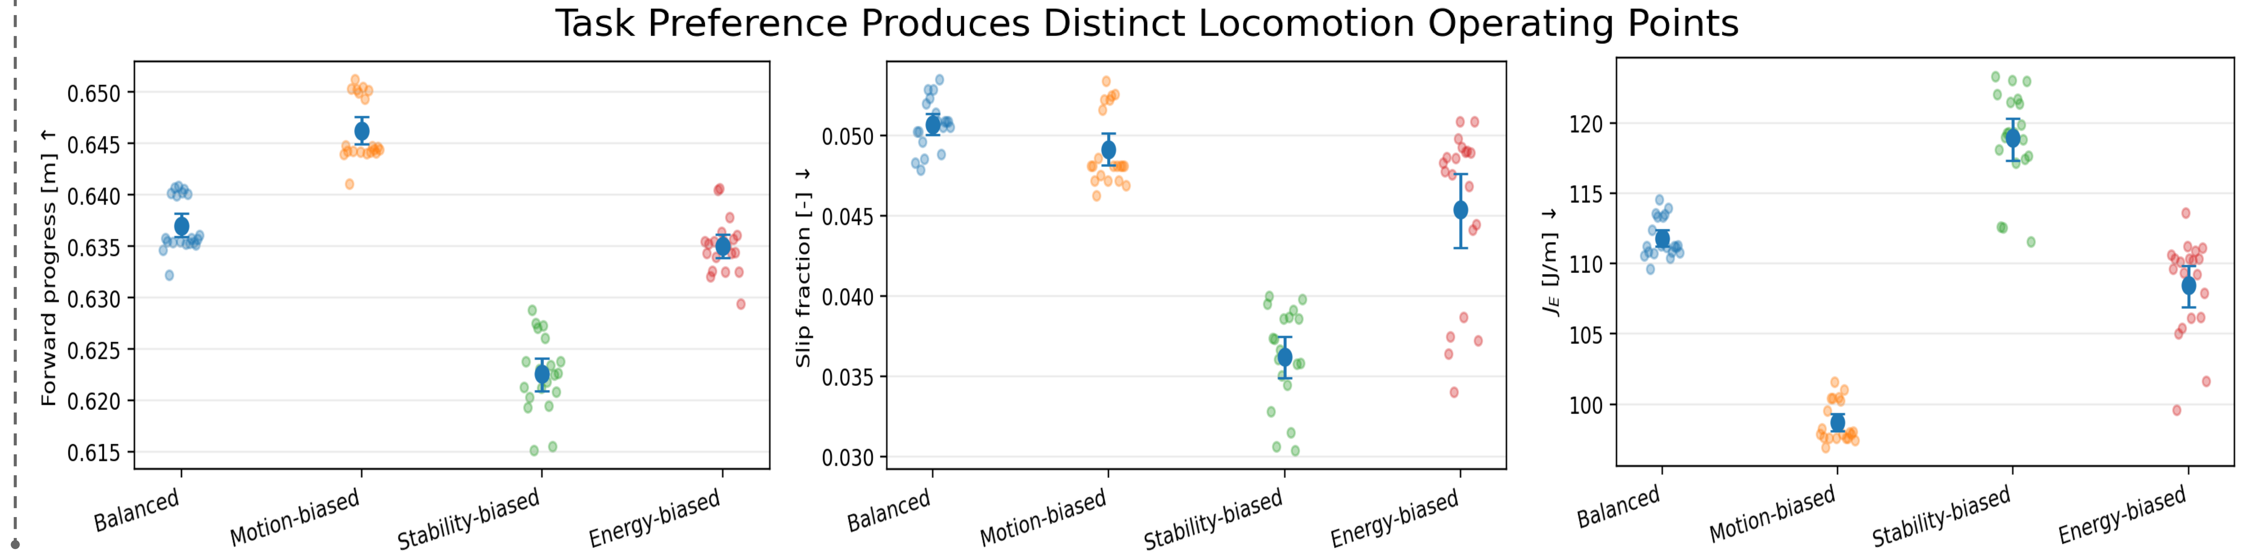

Motion, Stability, and Energy preferences shift the robot toward distinct physical operating points.

Controlled Preference Sweep

Preference	Target Physical effect	vs. Balanced	Consistency
Motion	Forward Progress ↑	+1.45 %	20 / 20
Stability	Slip ↓	-28.5 %	20 / 20
Energy	J_E ↓	-2.99 %	19 / 20

Interpretable table

Interpretation table

Observation	What it tells us
Objective conditioning changes strategy	Internal conditioning is behaviorally meaningful
Physical effects remain coupled	Reward Semantics ≠ Physical Semantics
Stability is harder to isolate	Current strategy needs richer contact / gait authority

These findings motivated the next TRACER stage

Current Work

- **From Interface diagnosis to Role-Specialized Adaptation**

Observed limitations		Ongoing Extensions
Physical preference ≠ reward semantics		Preference calibration
Interaction not explicitly modeled	→	Interaction adaptation
Limited strategy authority		Structural gait authority