FreeMusco: Motion-Free Learning of Latent Control for Morphology-Adaptive Locomotion in Musculoskeletal Characters

Minkwan Kim, Yoonsang Lee*
Hanyang University

Motivation

- Existing motion capture datasets → we can generate a wide variety of realistic human motions.
- But what if **motion capture is difficult or infeasible** – such as for animals or imaginary creatures?

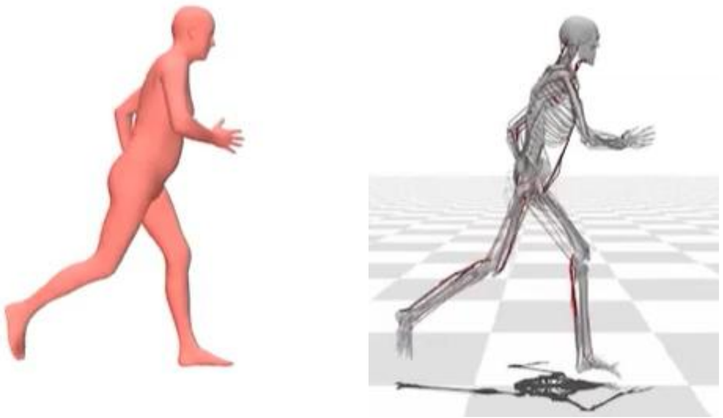

[Park et al. 2025]

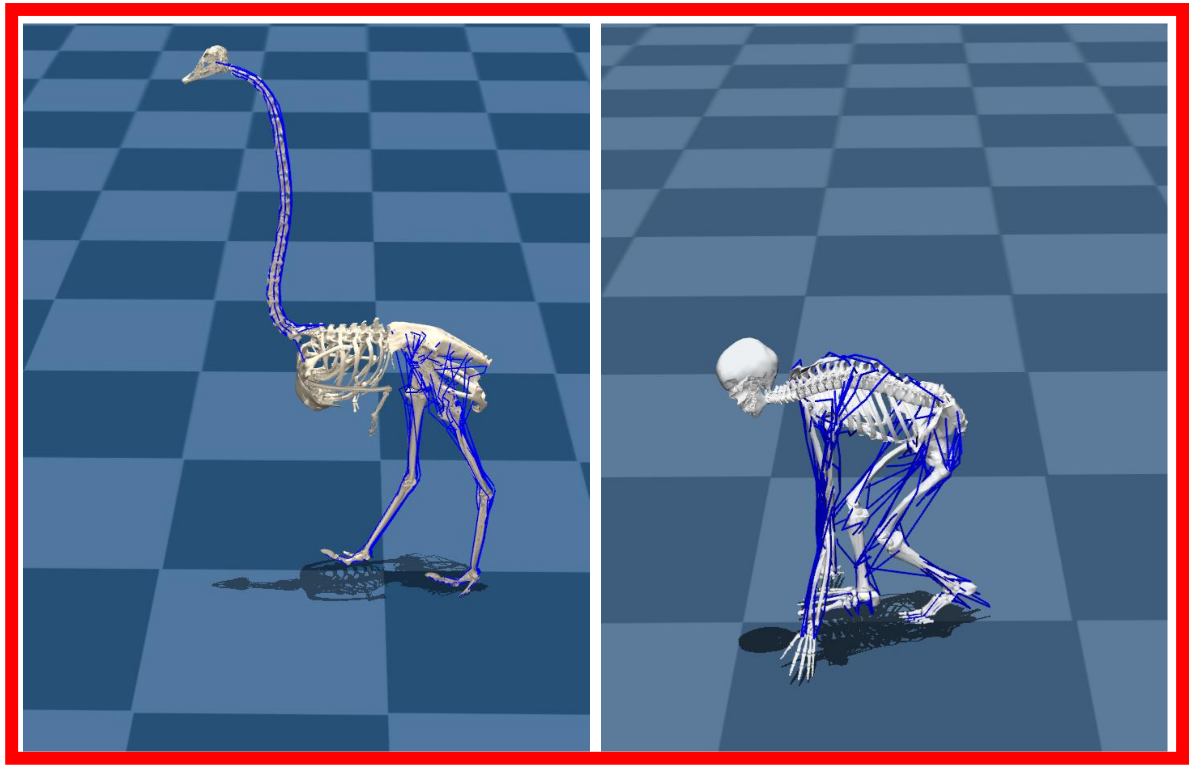

FreeMusco

- We propose a **motion-free** framework that learns a **latent space** of **morphology-adaptive locomotion** behaviors in various **musculoskeletal** characters.
- The model is trained without motion data, based only on **morphology and biomechanics**. It generalizes to **non-human** characters (e.g., Ostrich) and **synthetic** characters (e.g., Chimanoid).

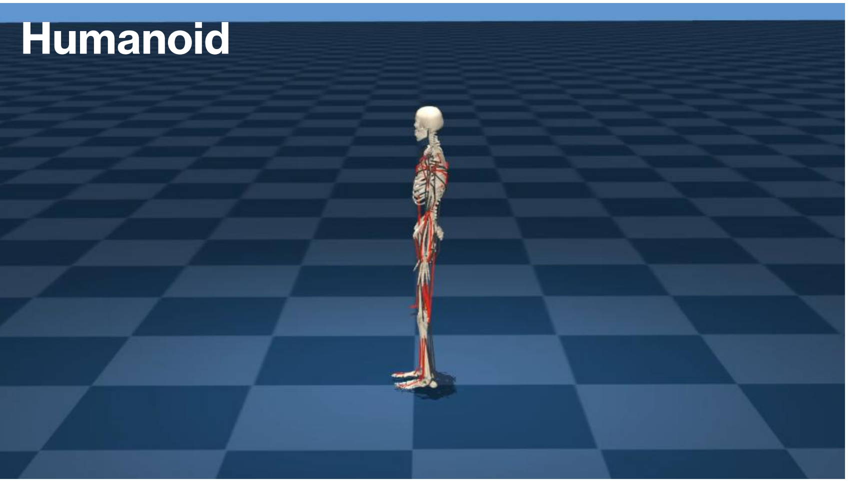

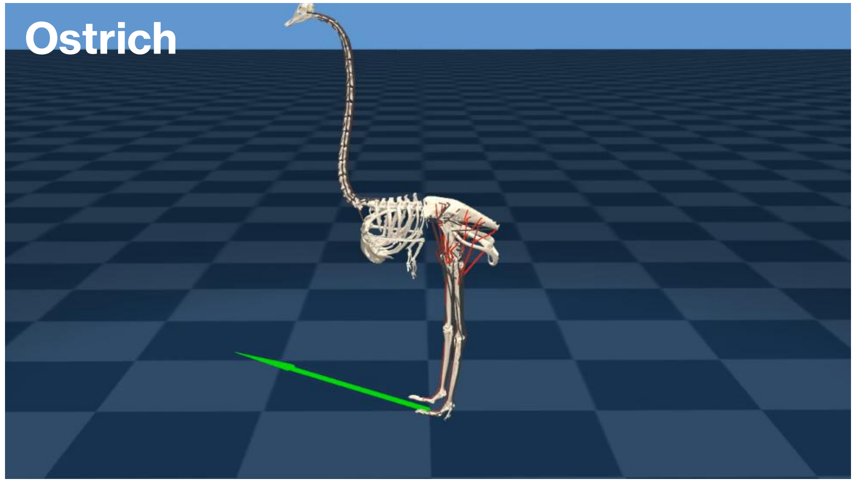

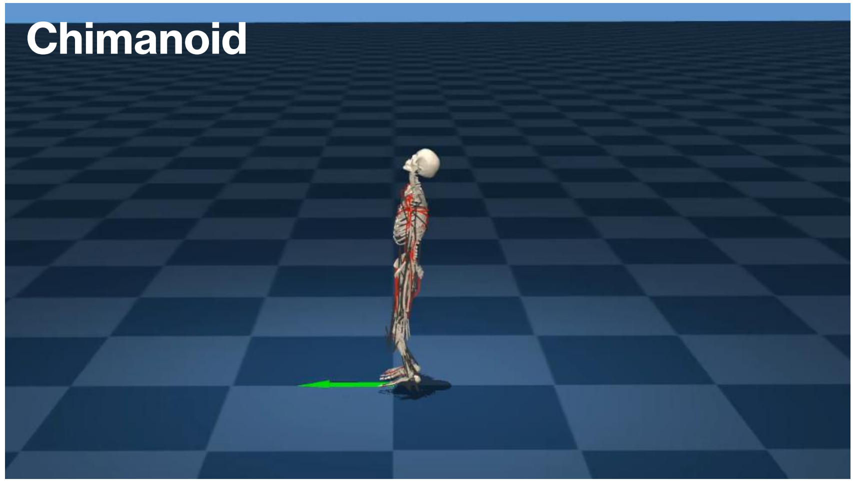

FreeMusco

- We propose a **motion-free** framework that learns a **latent space** of **morphology-adaptive locomotion** behaviors in various **musculoskeletal** characters.
- The model is trained without motion data, based only on **morphology and biomechanics**. It generalizes to **non-human** characters (e.g., Ostrich) and **synthetic** characters (e.g., Chimanoid).
- The learned latent space enables **high-level control**, like goal navigation and path following.

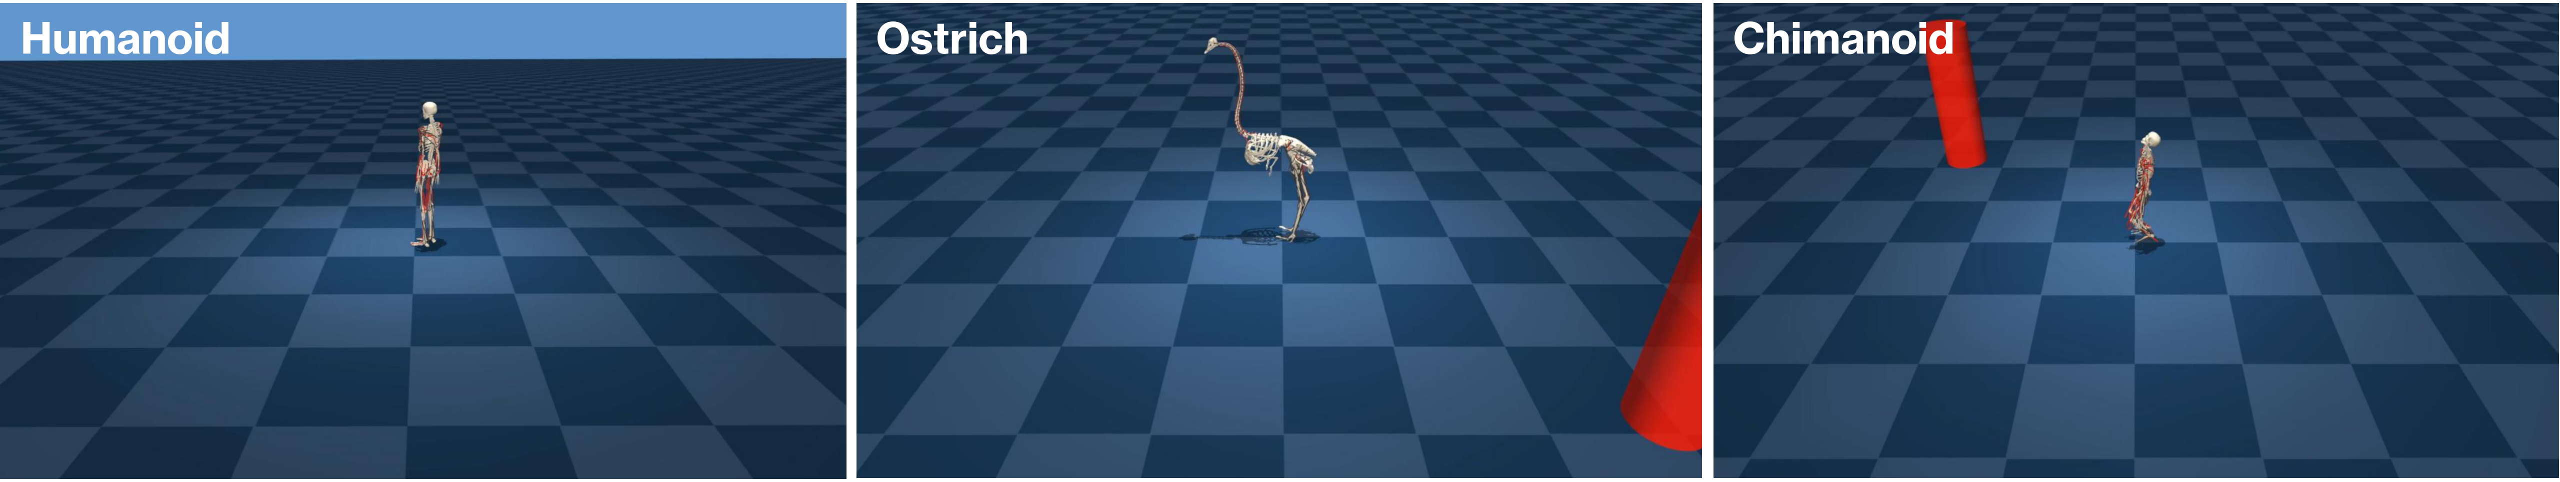

Prior Works

Prior Methods: Motion-Driven Character Control

[Lee et al. 2010]

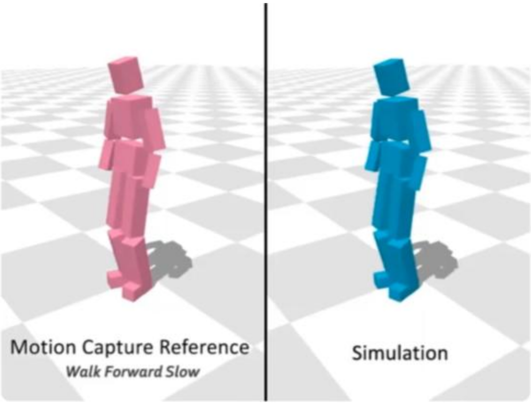

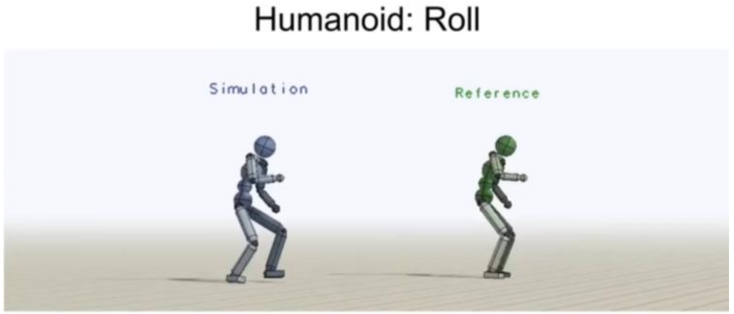

[Peng et al. 2018]

[Yao et al. 2022]

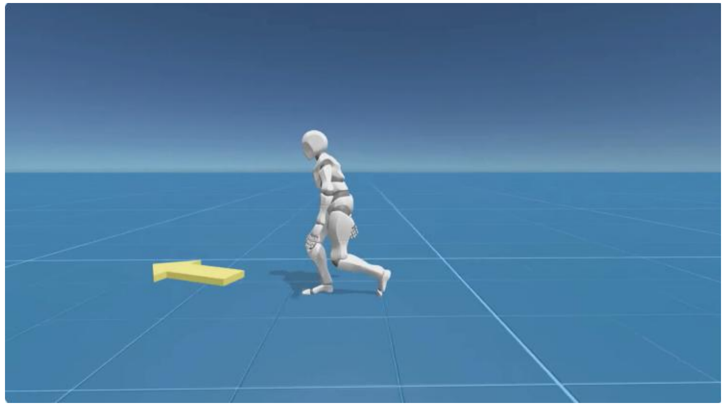

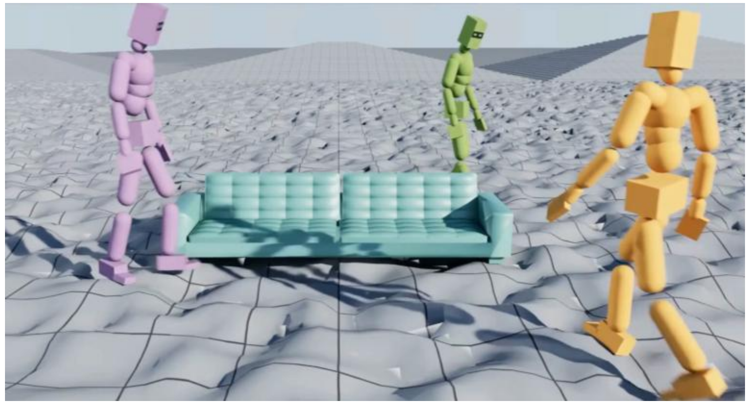

[Tessler et al. 2024]

- Advance of motion-driven character control:

  Mocap-based feedback control → DRL, imitation learning and generative models.

- Pros: Imitating diverse behaviors / **Cons: constrained by the distribution of demonstration data.**

Prior Methods: Motion-Free Character Control

[Hodgins et al. 1995]

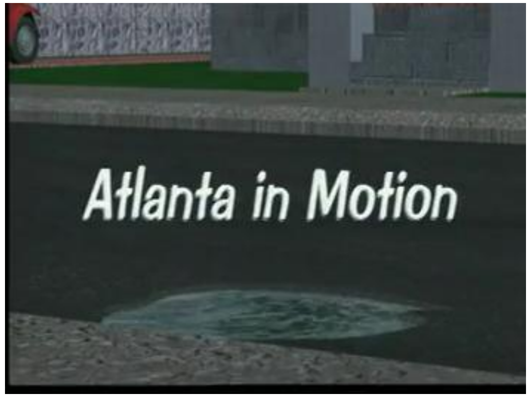

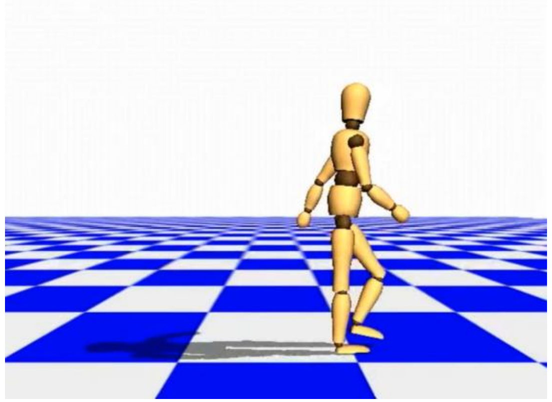

[Yin et al. 2007]

[Yu et al. 2018]

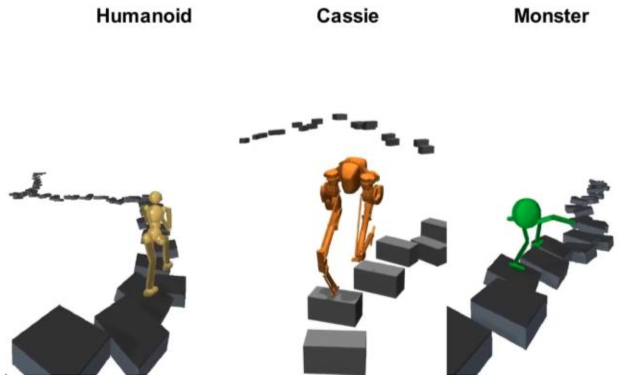

[Xie et al. 2020]

- Advance of motion-free character control:

  Handcrafted controllers & FSMs → RL frameworks (learn locomotion without reference data).

- **Pros: could generalize to novel morphologies.**

Prior Methods: Musculoskeletal Character Control

[Wang et al. 2012]

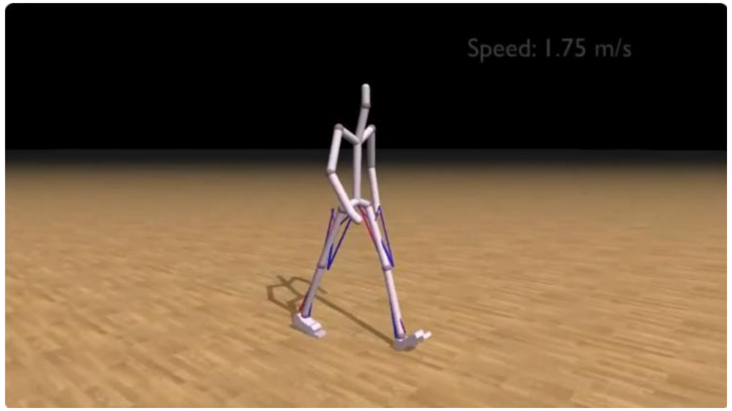

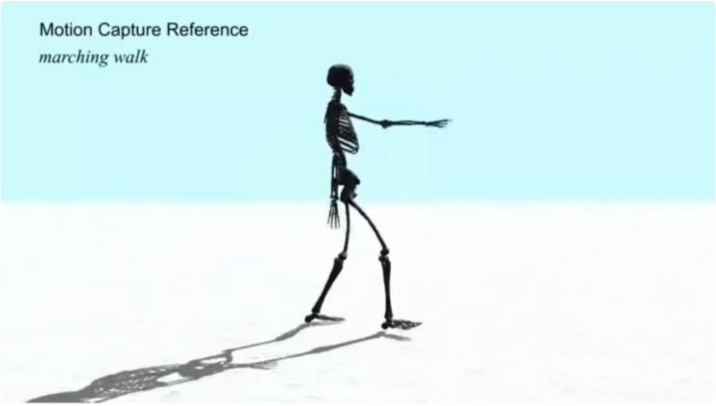

[Lee et al. 2014]

[Feng et al. 2023]

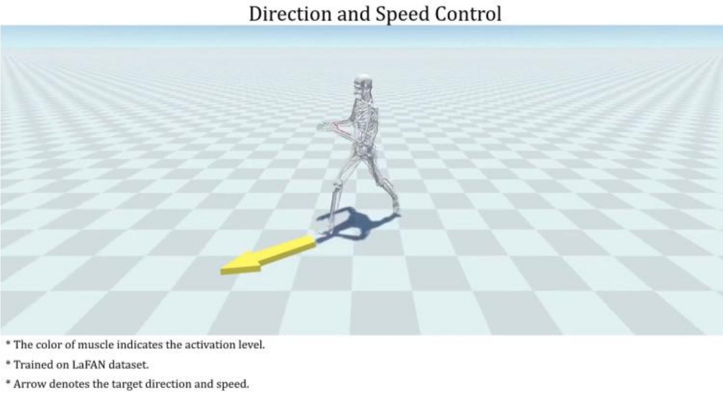

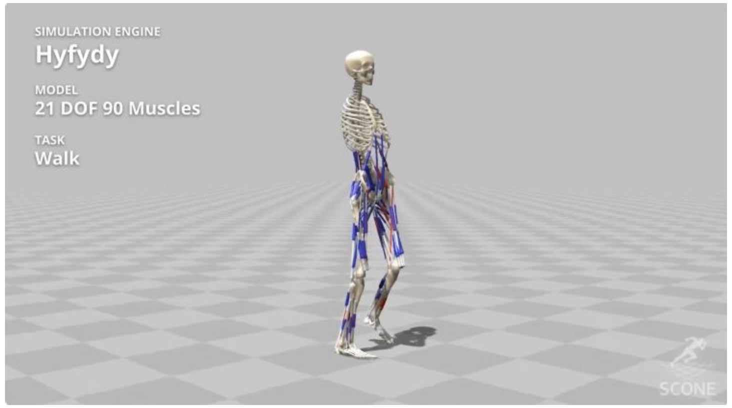

[Schumacher et al. 2025]

- Advance of musculoskeletal character control:
  Biological objectives & trajectory-optimization → learning-based methods (w/ or w/o motion data).
- We **revisit motion-free** method:
  Learning latent control policies for diverse & **morphology-adaptive locomotion.**

Method

- We adopt a ControlVAE [Yao et al. 2023]-likes architecture, but with **some modifications.**
- 1) Conditional VAE is guided by **goal signal** rather than reference motion.
- 2) World model predicts the **energy expenditure** as well. (Original: state-transition only)

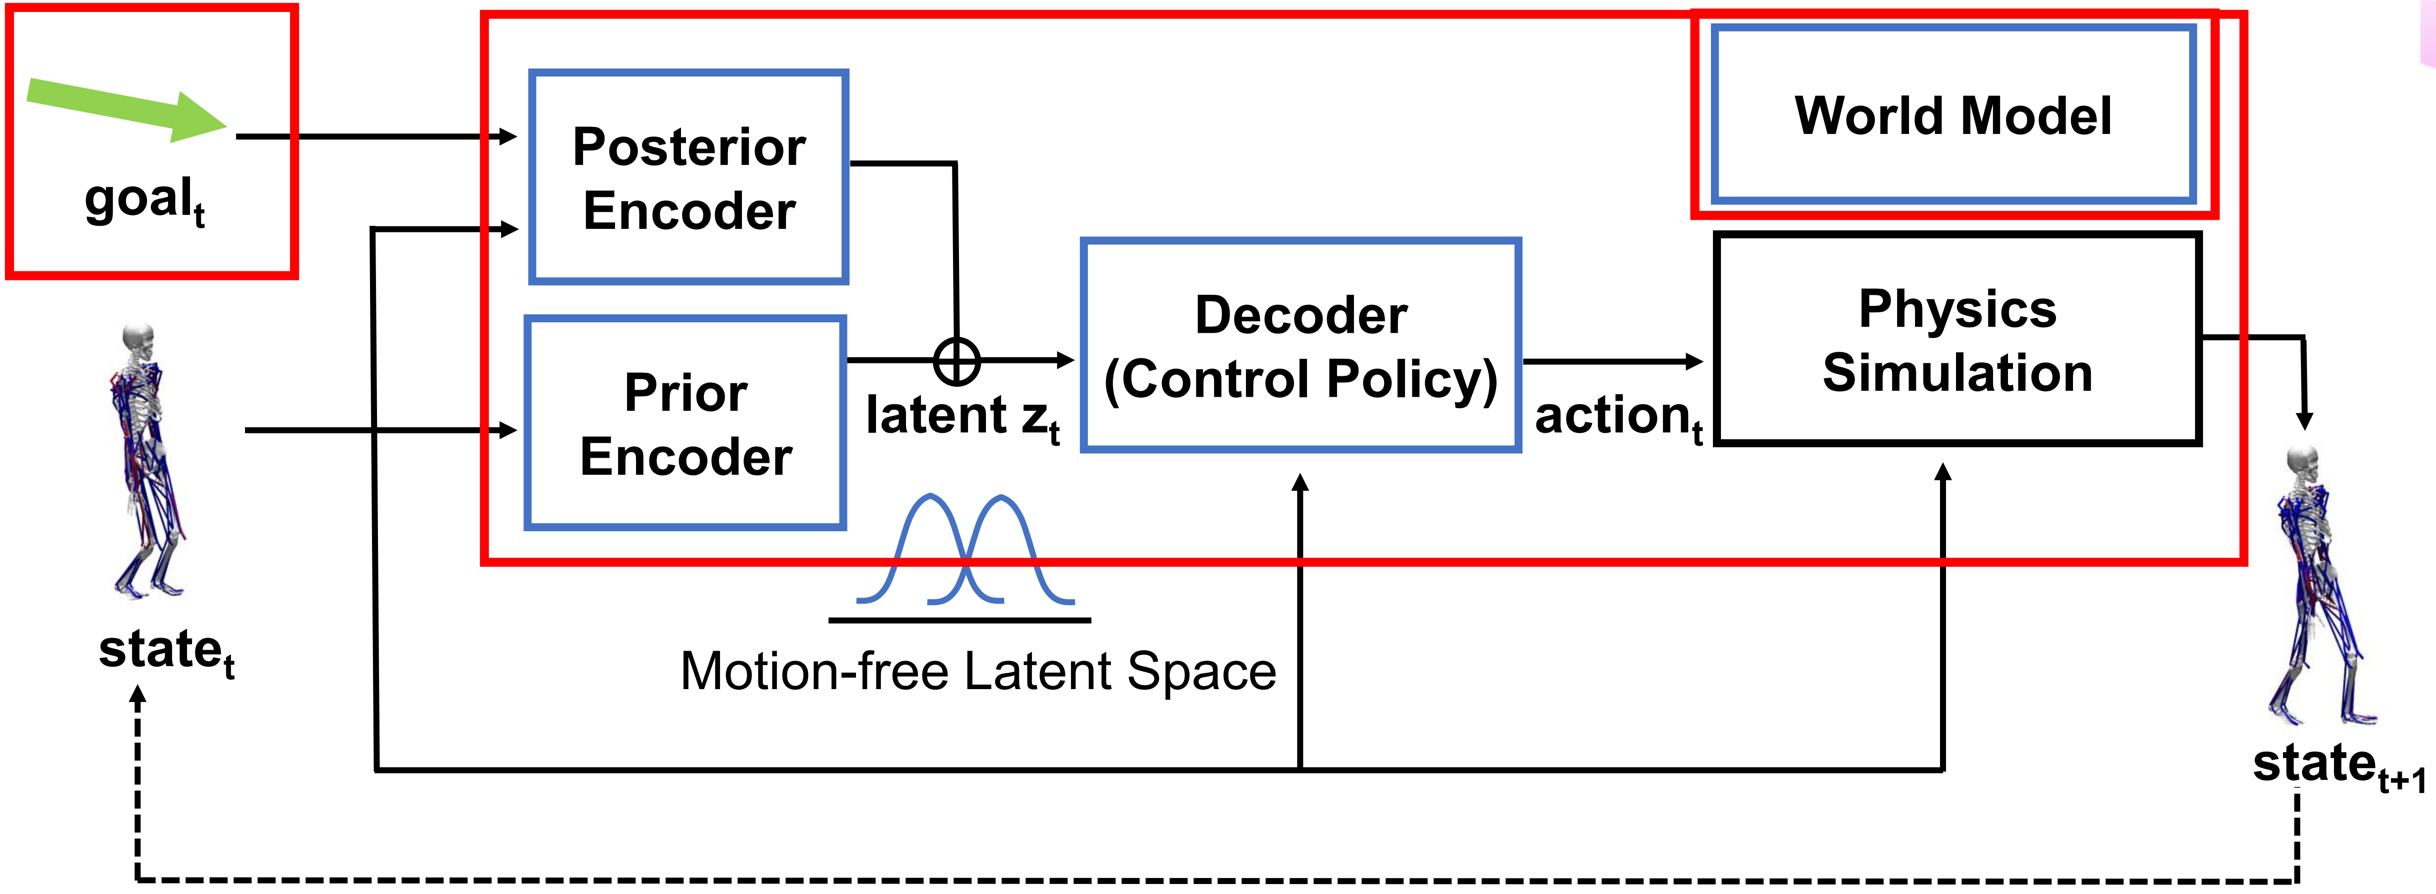

- **Goal:** target velocity, posture, energy, and facing direction (randomly assigned during training).
- **Character state:** position, velocities, linear and angular velocities (of all links).
- **Latent vector:** 64-dim vector.
- **Action:** muscle activation (control signals for muscle dynamics).

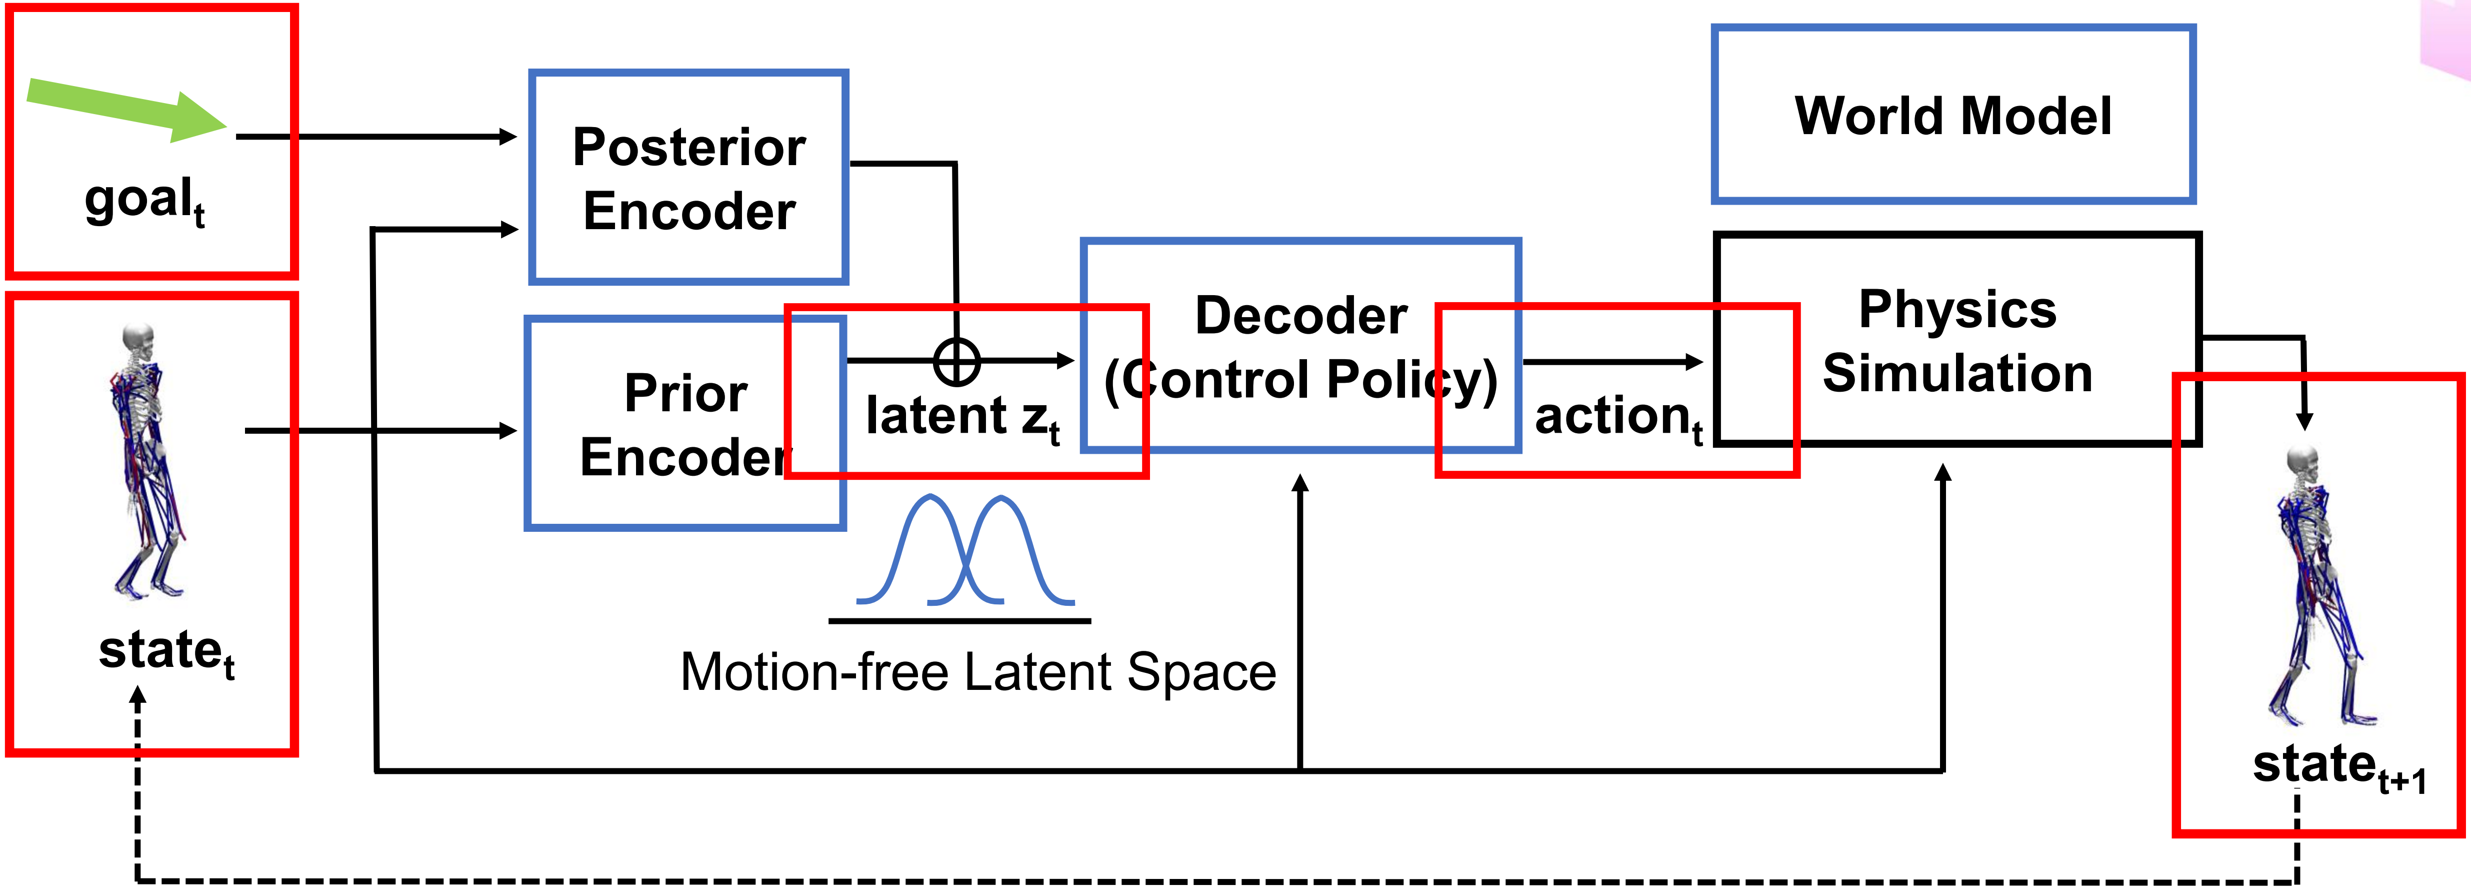

- The model is trained by repeating the following three stages in each iteration.

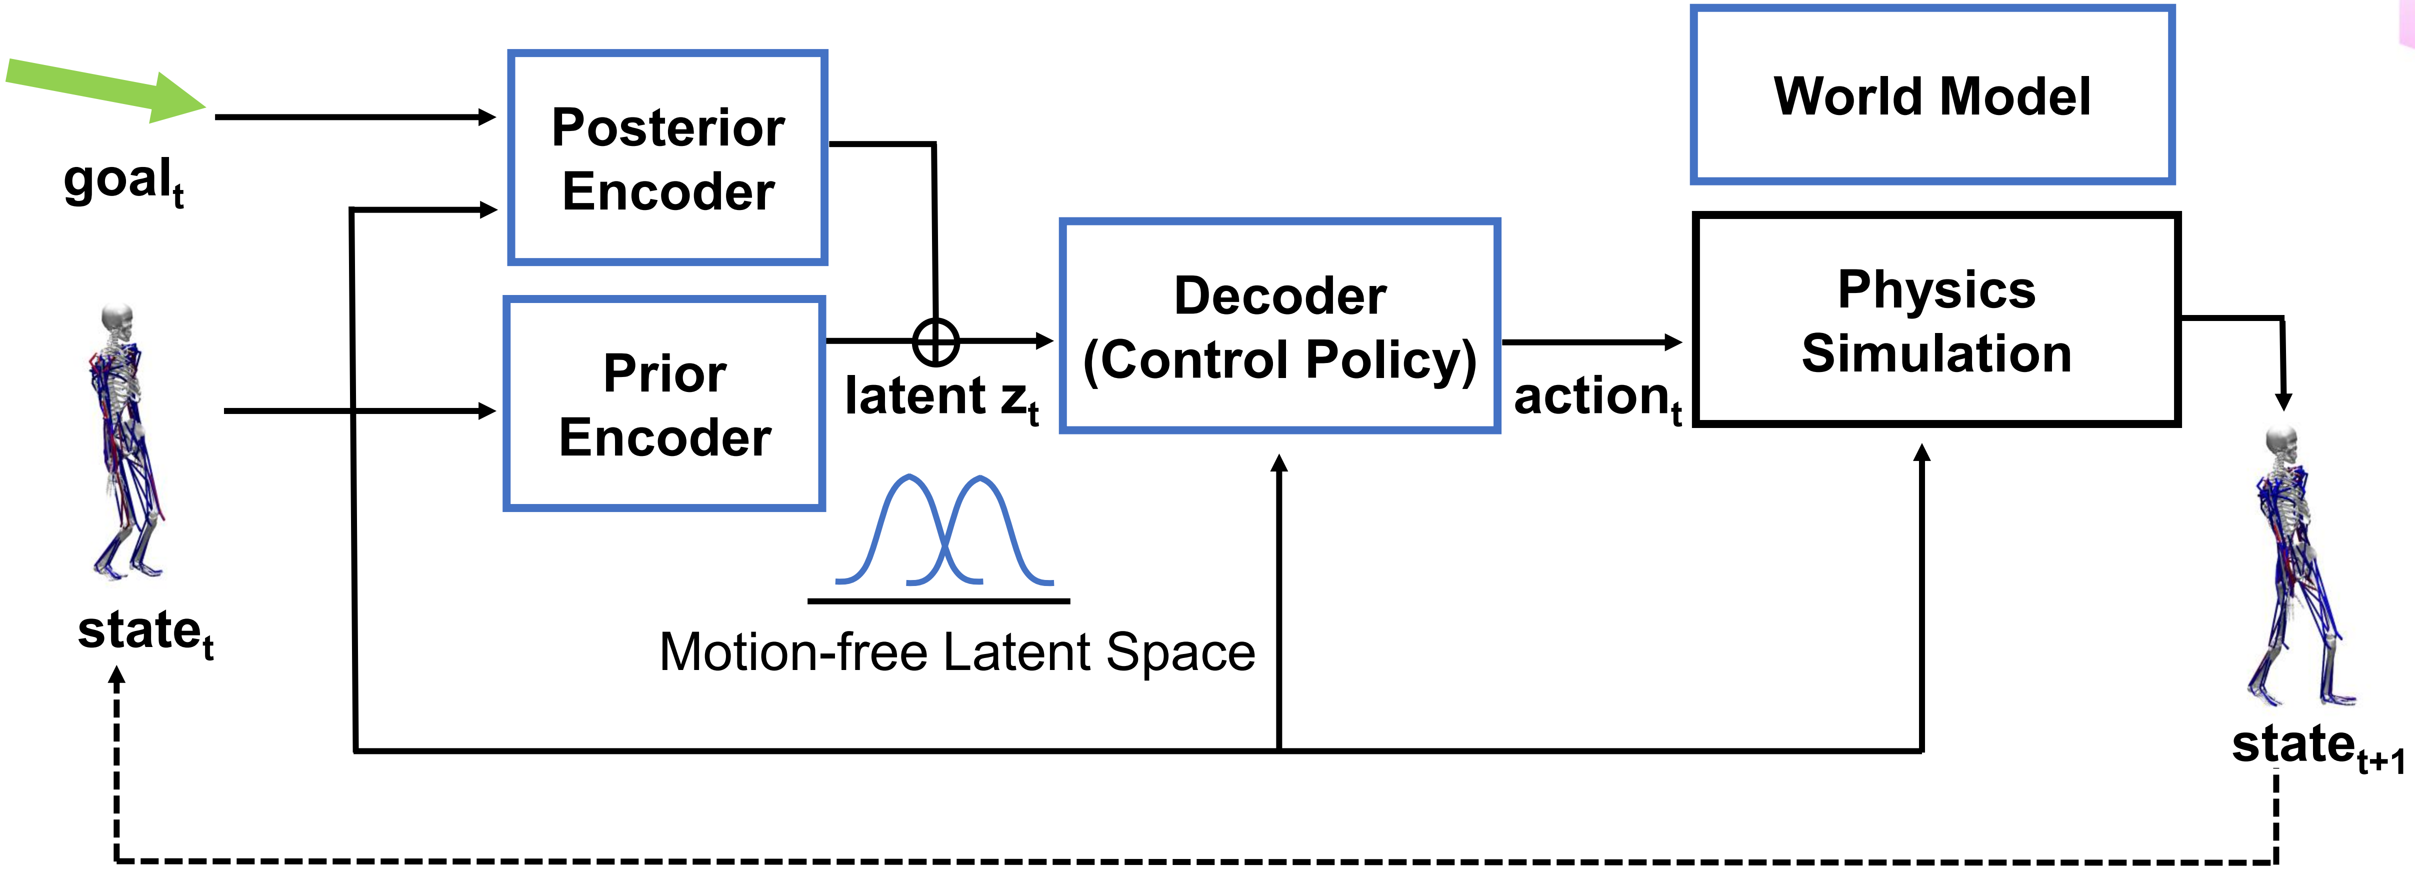

- 1st stage: Generating **simulation trajectories (Goal, action, state transitions, energy expenditure)**.

→ Stored in trajectory buffer and **used in 2nd and 3rd stages for updating each network.**

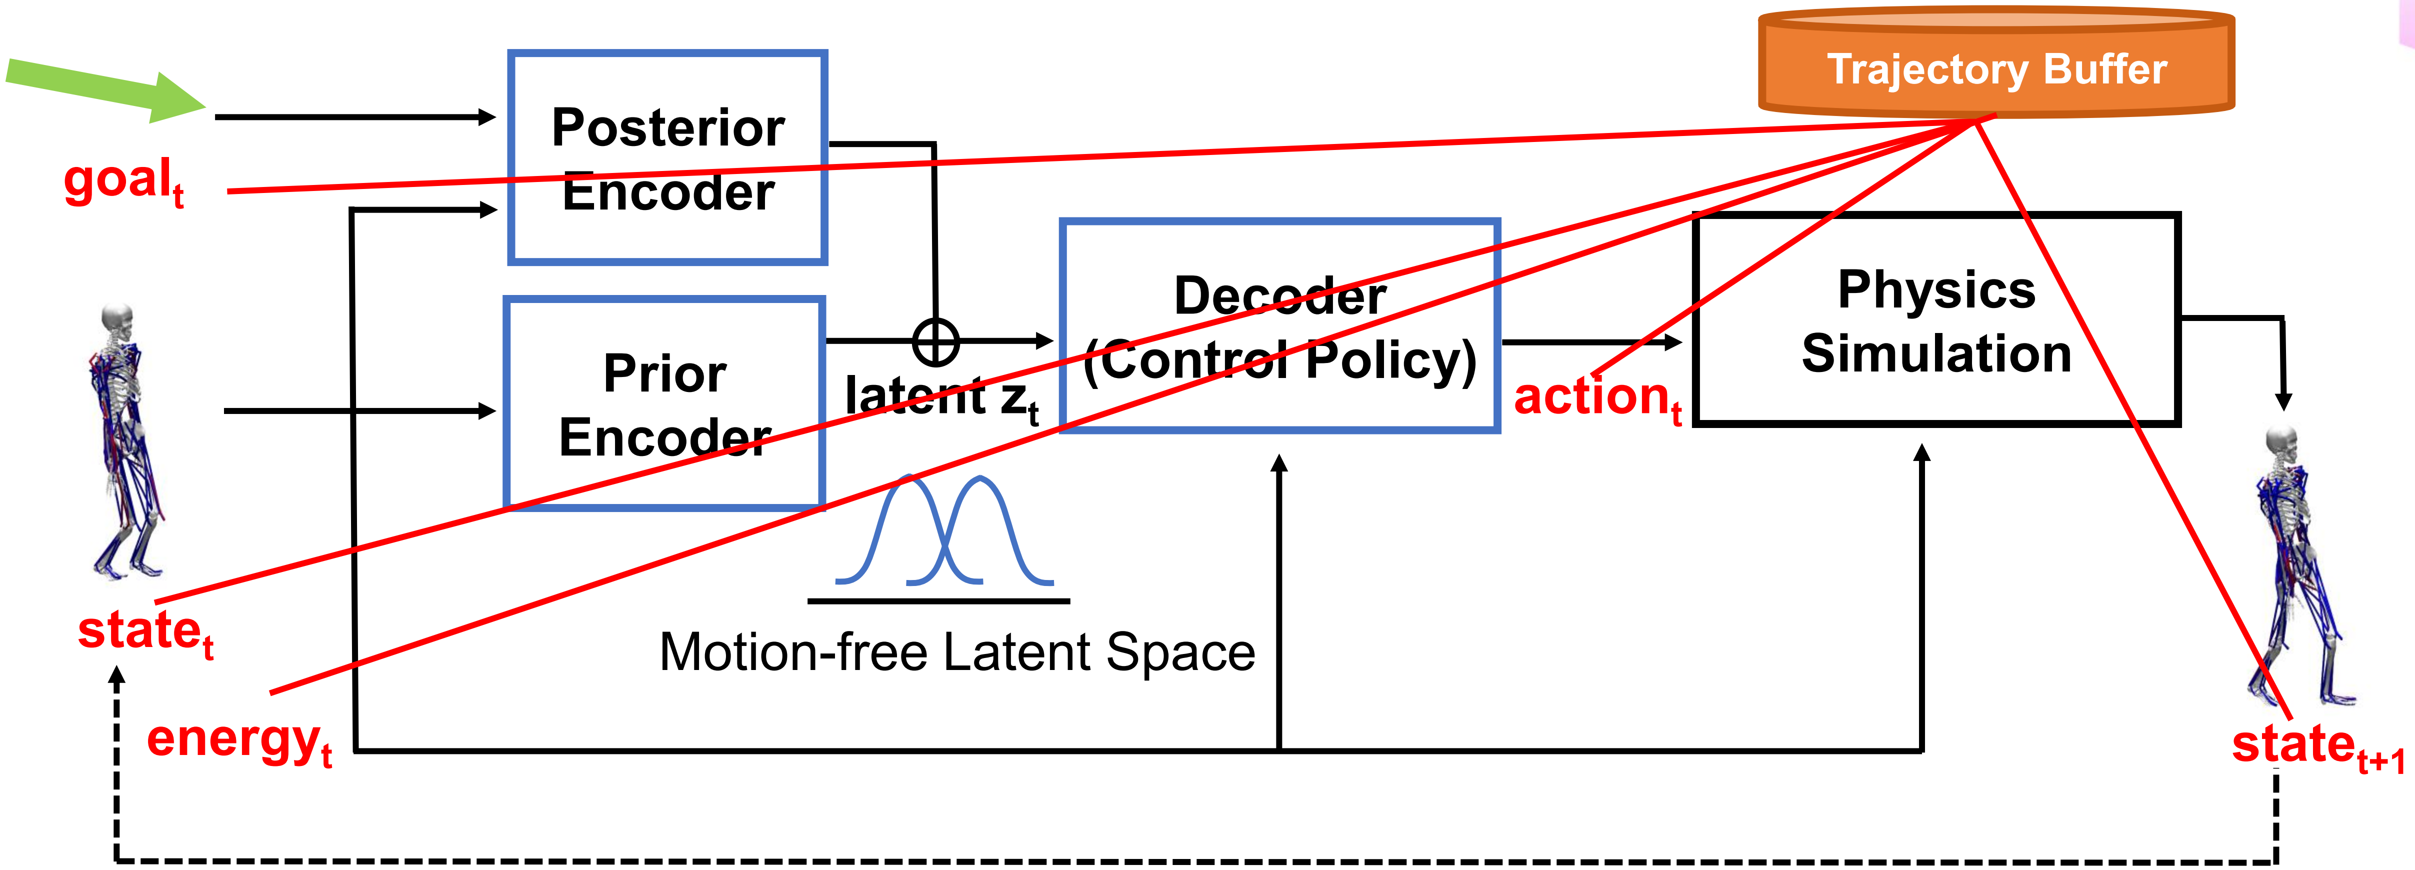

- 2nd stage: **Updating world model** (to predict ground truth **state transitions** and **energy expenditure).**

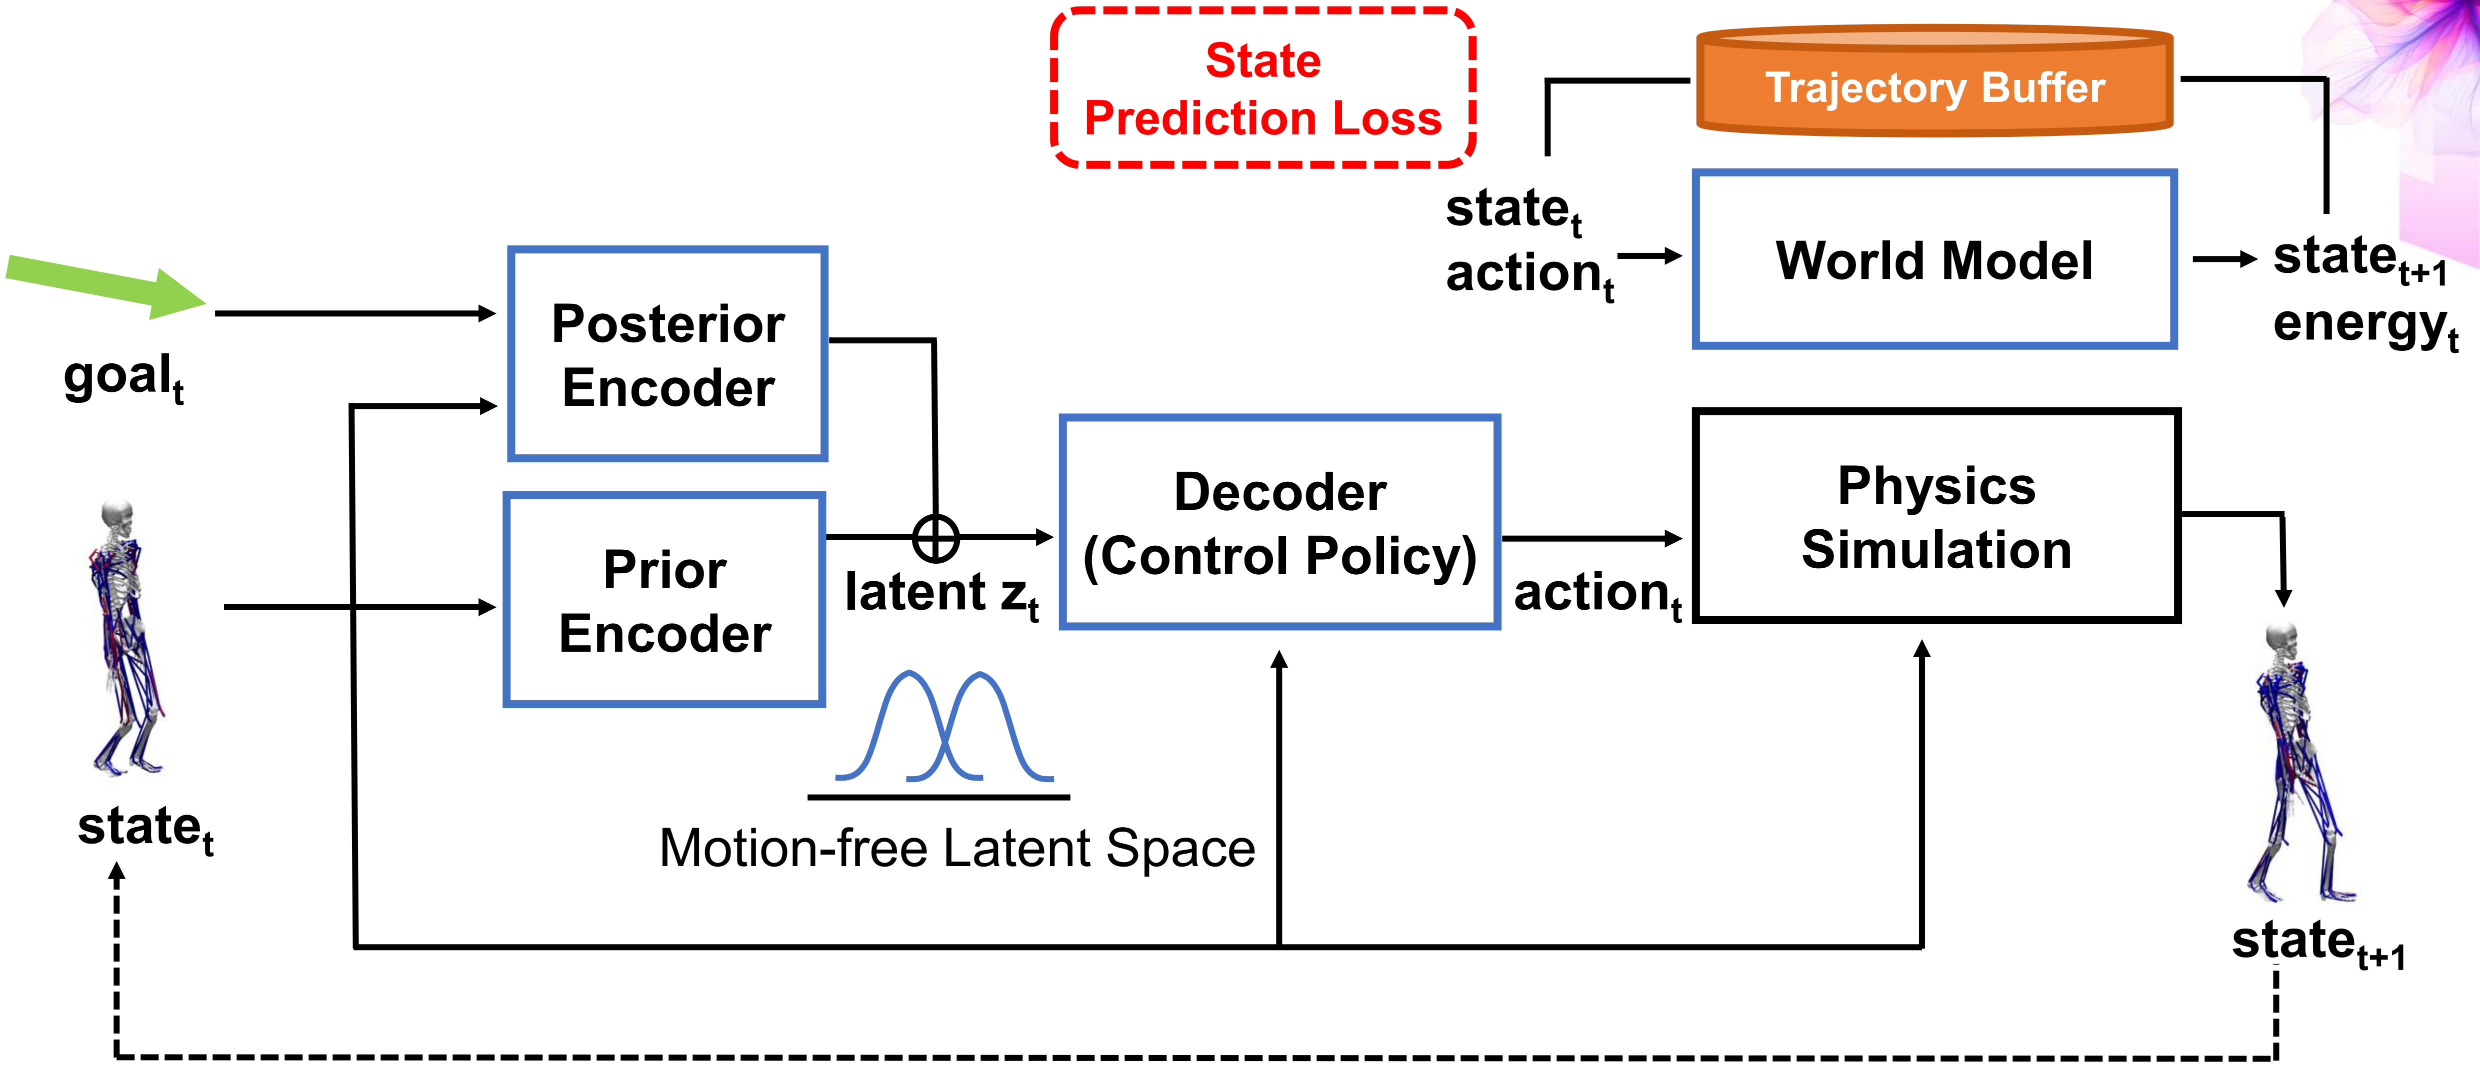

- 3rd stage: **Encoder and decoder** are updated with proposed **locomotion objective loss (using world model).**

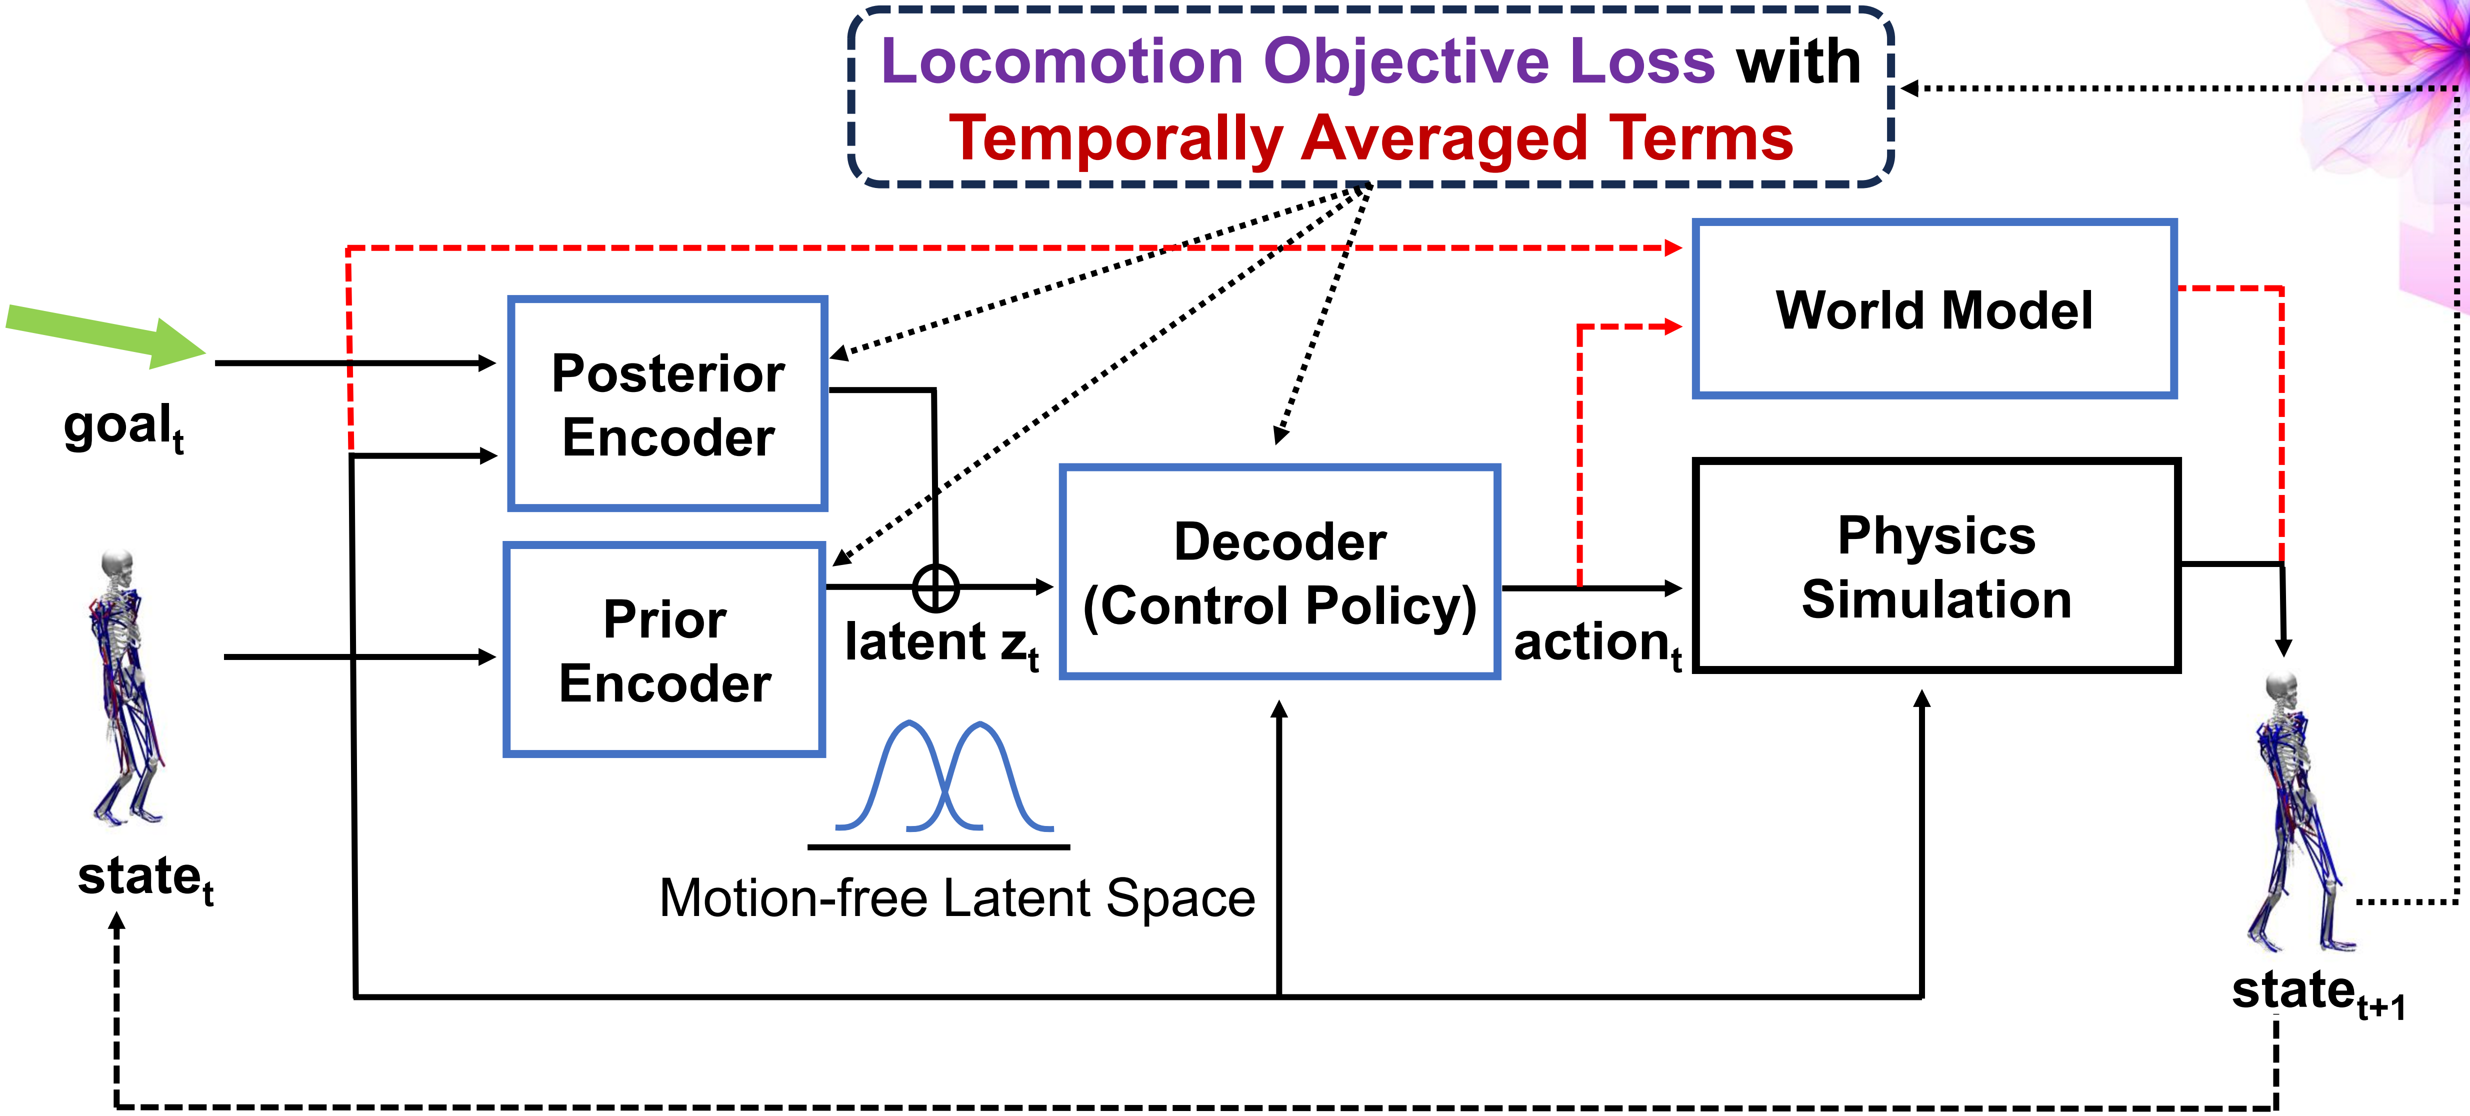

Loss Function

$$L_{objective} = \underbrace{L_{vel} + L_{dir}}_{\text{Control Objective}} + \underbrace{L_{height} + L_{up}}_{\text{Balancing Objective}} + \underbrace{L_{pose} + L_{energy}}_{\text{Biomechanical Objective}}$$

Loss Formulation: Per-step **vs Temporally Averaged**

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

- **Per-step** loss has proven effective in imitation-based frameworks, because **reference trajectories naturally exhibit rhythmic variation.**
- However, out motion-free setting lacking such patterns.

$$L_{step}(\{\bar{x}_t\}, \{x_t\}) = \frac{1}{T_p} \cdot \sum_{t=0}^{T_p - 1} \gamma^t \cdot \|\bar{x}_t - x_t\|$$

Loss Formulation: Per-step **vs Temporally Averaged**

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

- We introduce the **temporally averaged loss** to promote biologically plausible locomotion by accounting for **natural oscillations in movement.**
- This loss compares averages of the **simulated** and **target** states over a **short temporal window (32 steps).**

$$L_{avg}(\{\bar{x}_t\}, \{x_t\}) = \left\| \frac{1}{T_p} \cdot \sum_{t=0}^{T_p-1} \gamma^t \cdot \bar{x}_t - \frac{1}{T_p} \cdot \sum_{t=0}^{T_p-1} \gamma^t \cdot x_t \right\|_1$$

$$L_{step}(\{\bar{x}_t\}, \{x_t\}) = \frac{1}{T_p} \cdot \sum_{t=0}^{T_p-1} \gamma^t \cdot \|\bar{x}_t - x_t\|$$

Loss Function (Control)

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

$L_{vel} = L_{avg}(\{\overline{vel_t}\}, \{vel_t\})$

- Character **averaged** root (pelvis) speed → target speed ([0, 4.5] m/s during training).

$L_{dir} = L_{step}(\{\overline{dir_t}\}, \{dir_t\})$

- Character root (pelvis) facing direction → target direction (360° during training).

Loss Function (Balancing)

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

$$L_{height} = L_{step}(\{\overline{height_t}\}, \{height_t\})$$

- Character root (pelvis) height ≥ target height **(to avoid falling).**

$$L_{up} = L_{avg}(\{\overline{up_t}\}, \{up_t\})$$

- Character **averaged** root (pelvis) up direction → global up axis **(to maintain an upright posture).**

Loss Function (Biomechanical)

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

$$L_{pose} = L_{avg}(\{\overline{p_t}\}, \{p_t\})$$

- Character's **averaged** posture → target posture **(default: rest pose).**

$$L_{energy} = L_{step}(\{\overline{e}_t\}, \{e_t\})$$

- Character's metabolic energy consumption → target energy value **(default: zero energy).**

Results

- FreeMusco Framework Trained for Various Locomotion Behaviors

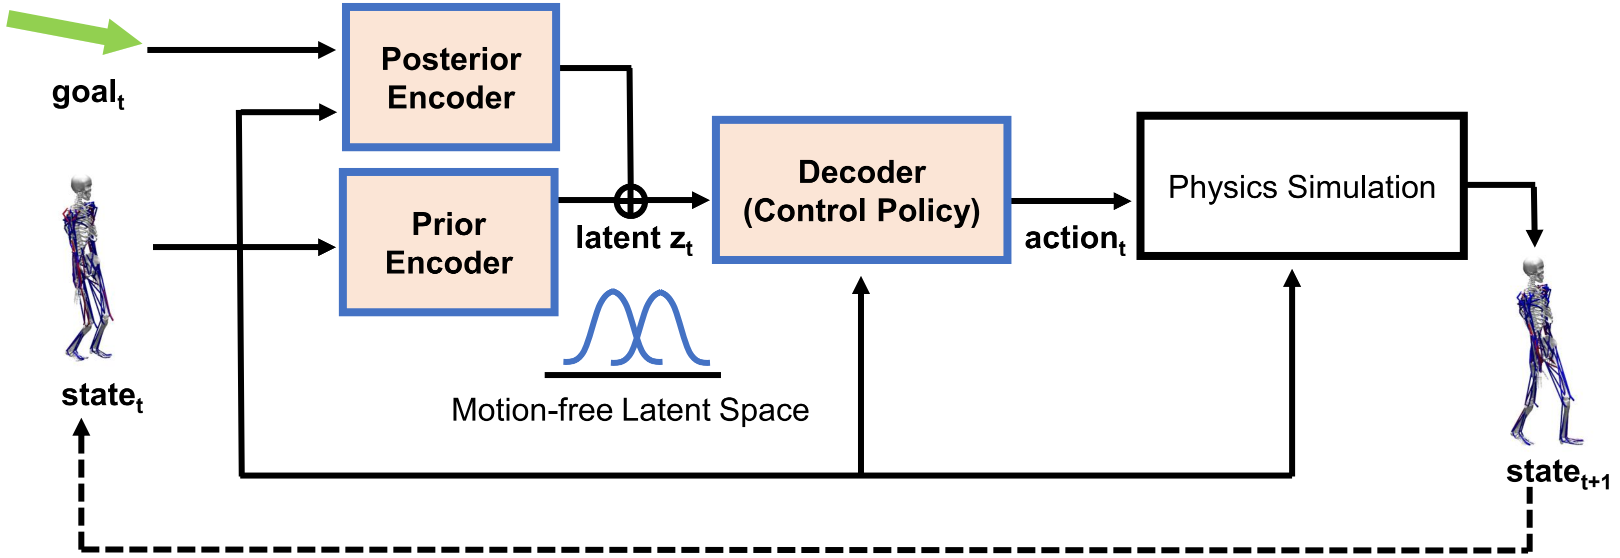

Humanoid Locomotion

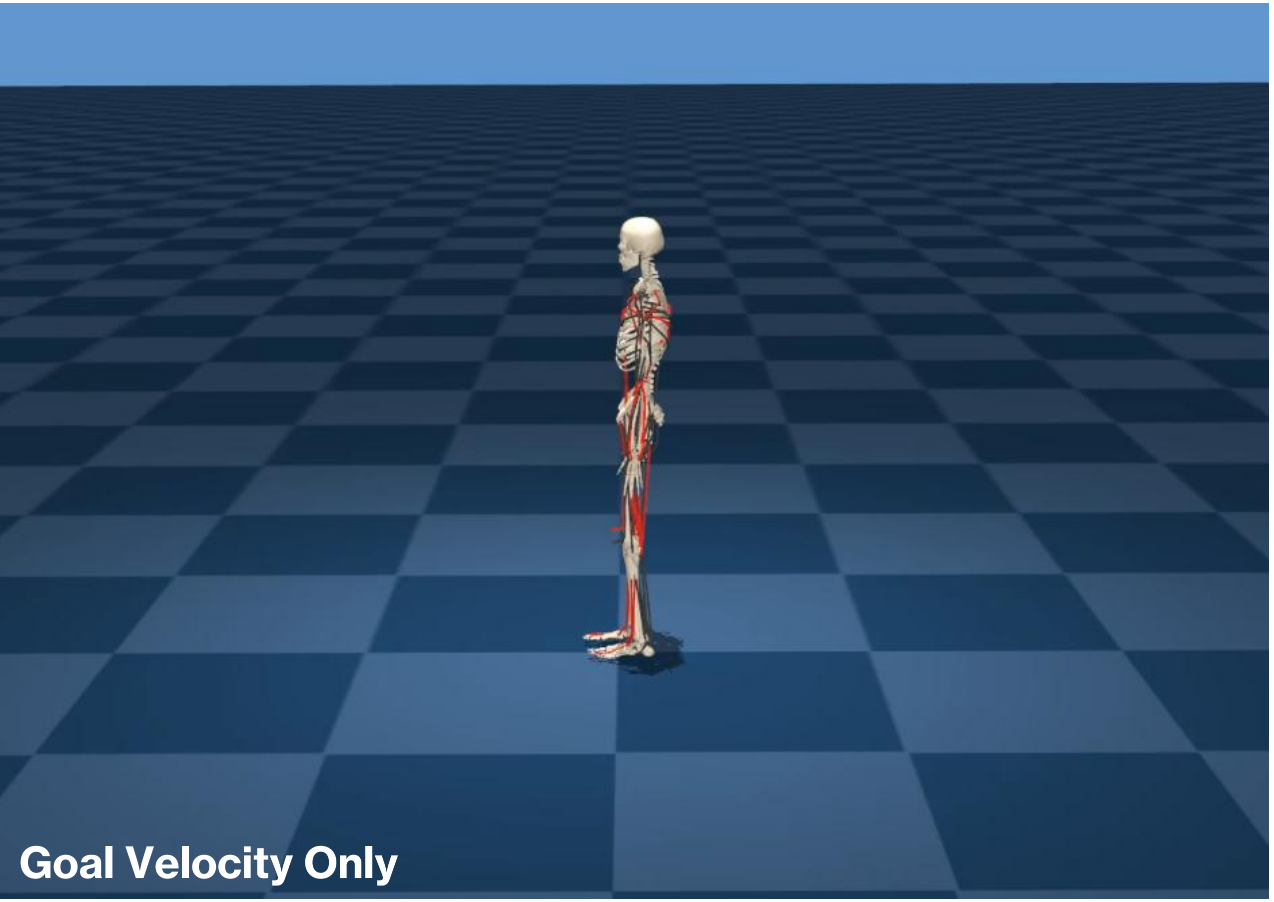

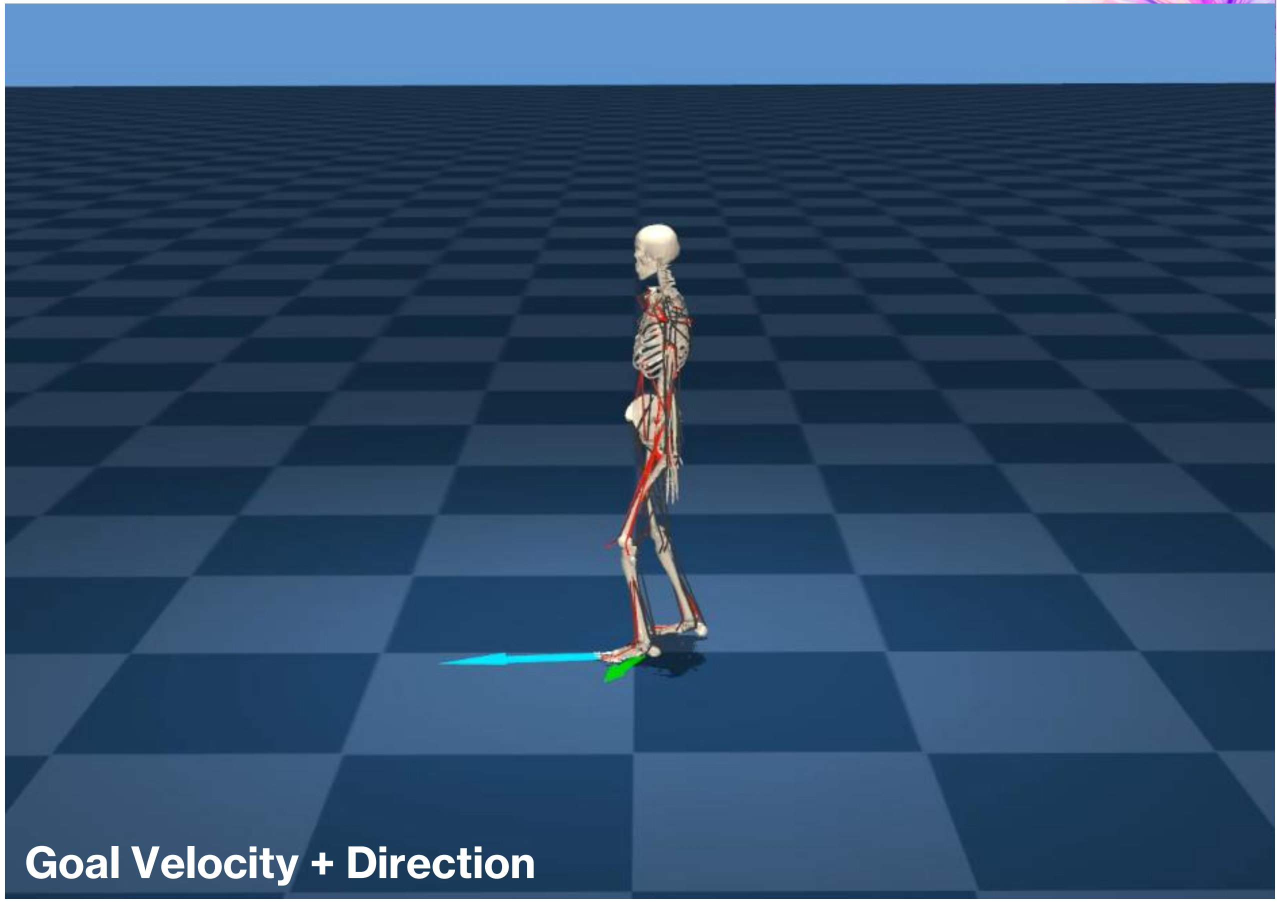

- We can control moving **velocity** and facing **direction** of the Humanoid.

Humanoid Locomotion

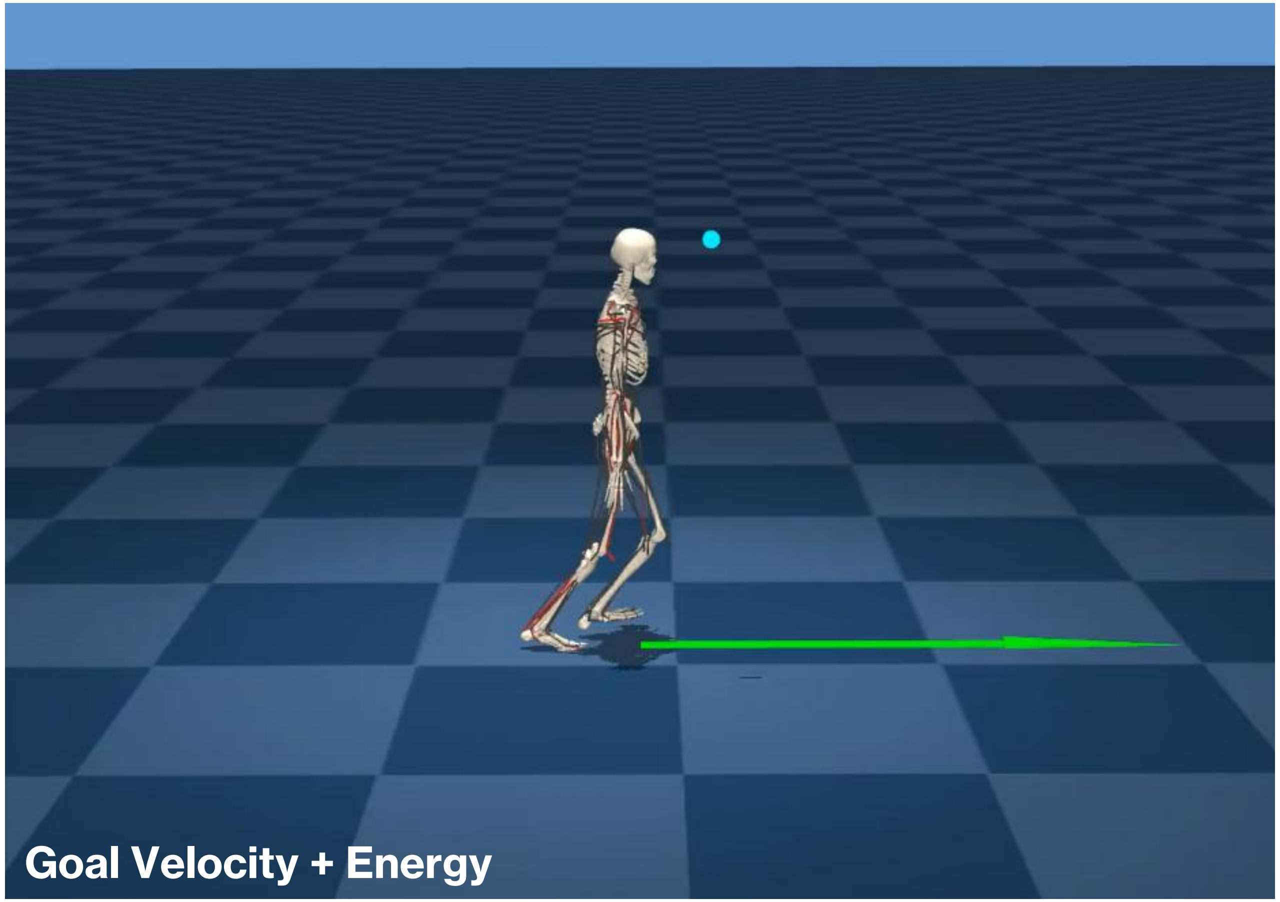

- **Low** target **energy** → **walking.**
- **High** target **energy** → transitions to **running** (activating most of the muscles).

Humanoid Locomotion

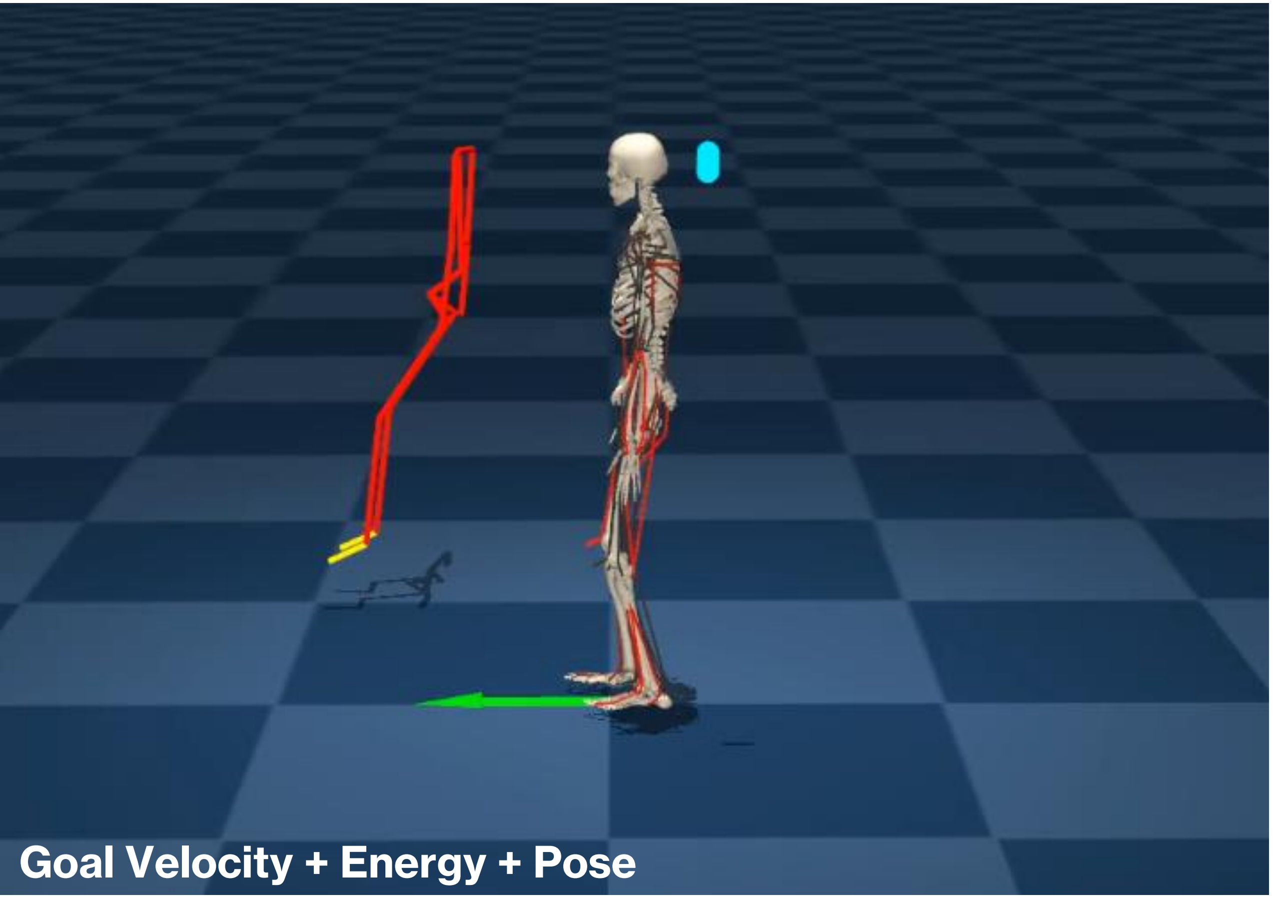

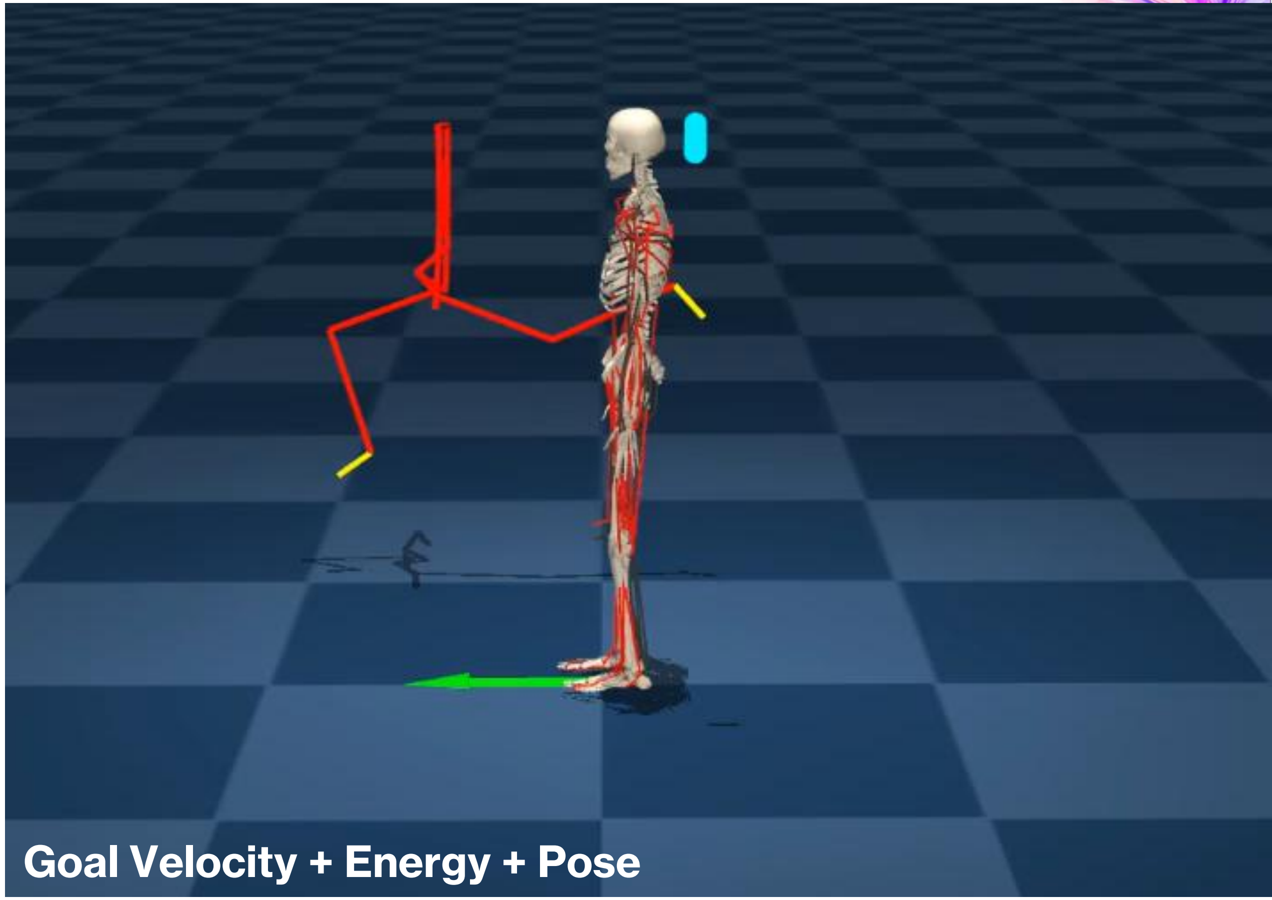

- Target **posture** and **energy** expenditure are **controllable at runtime.**
- The character can adapt to both **symmetric** and **asymmetric** lower-body postures.

Ostrich Locomotion

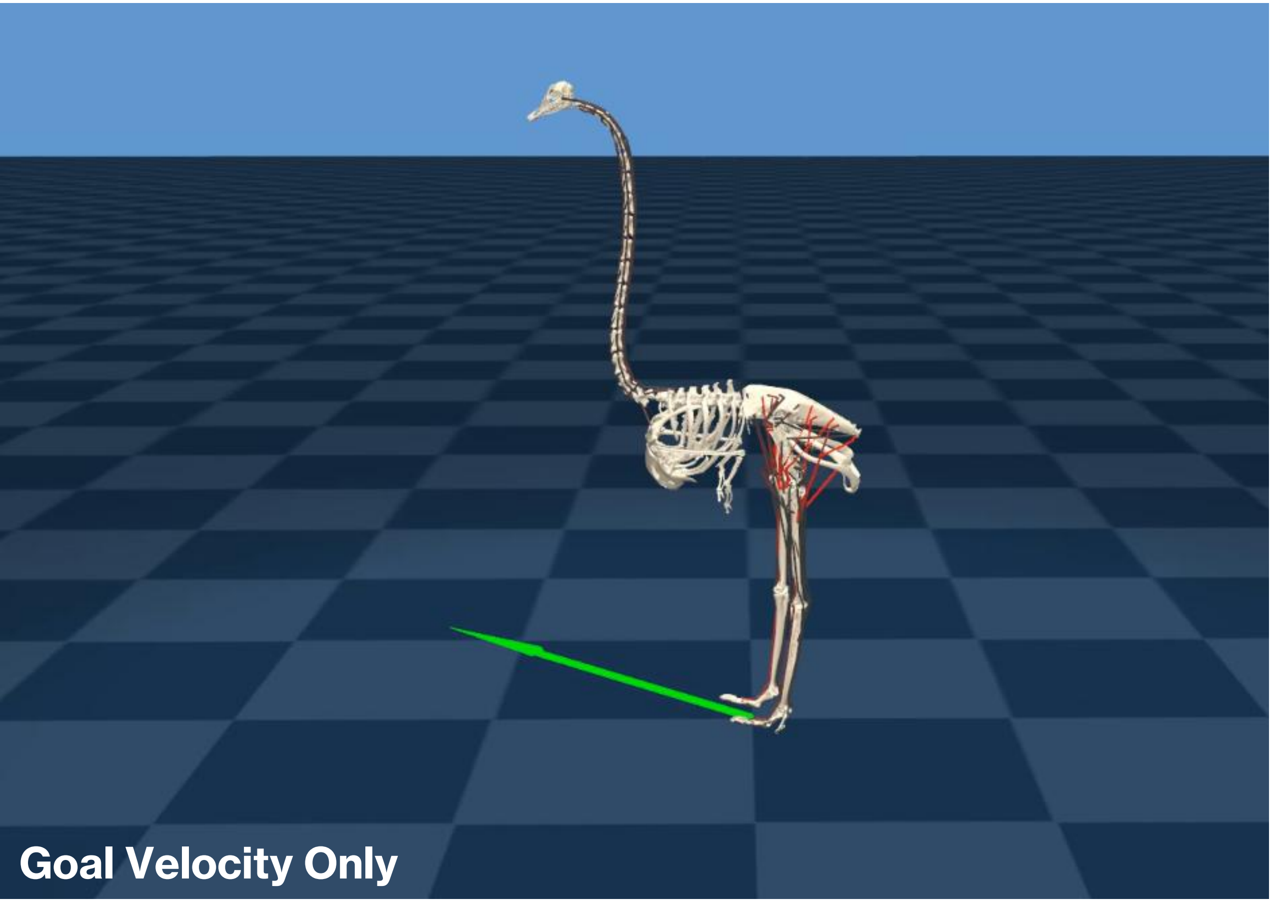

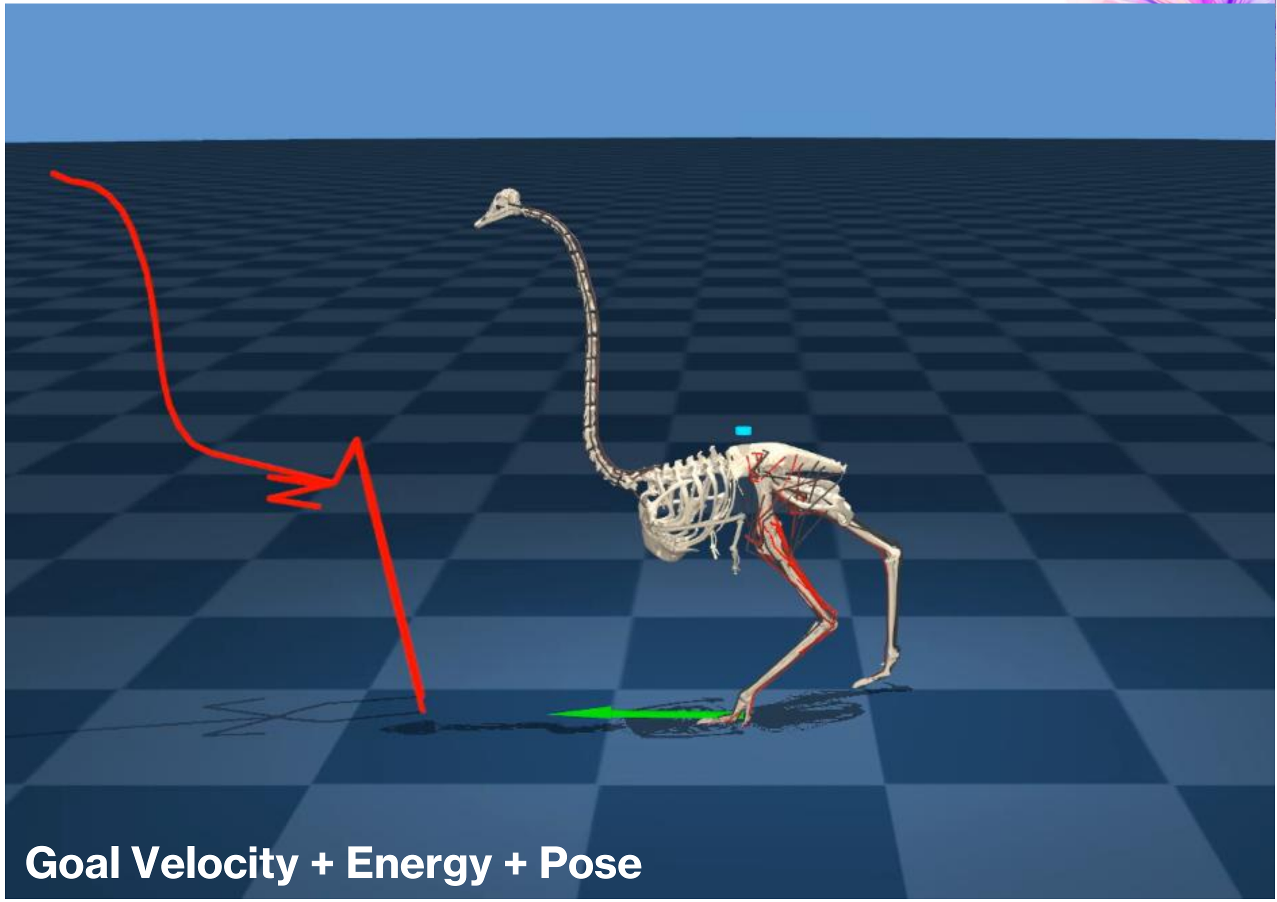

- Our motion-free approach enables learning locomotion for **non-human characters** like an ostrich.

Chimanoid Locomotion

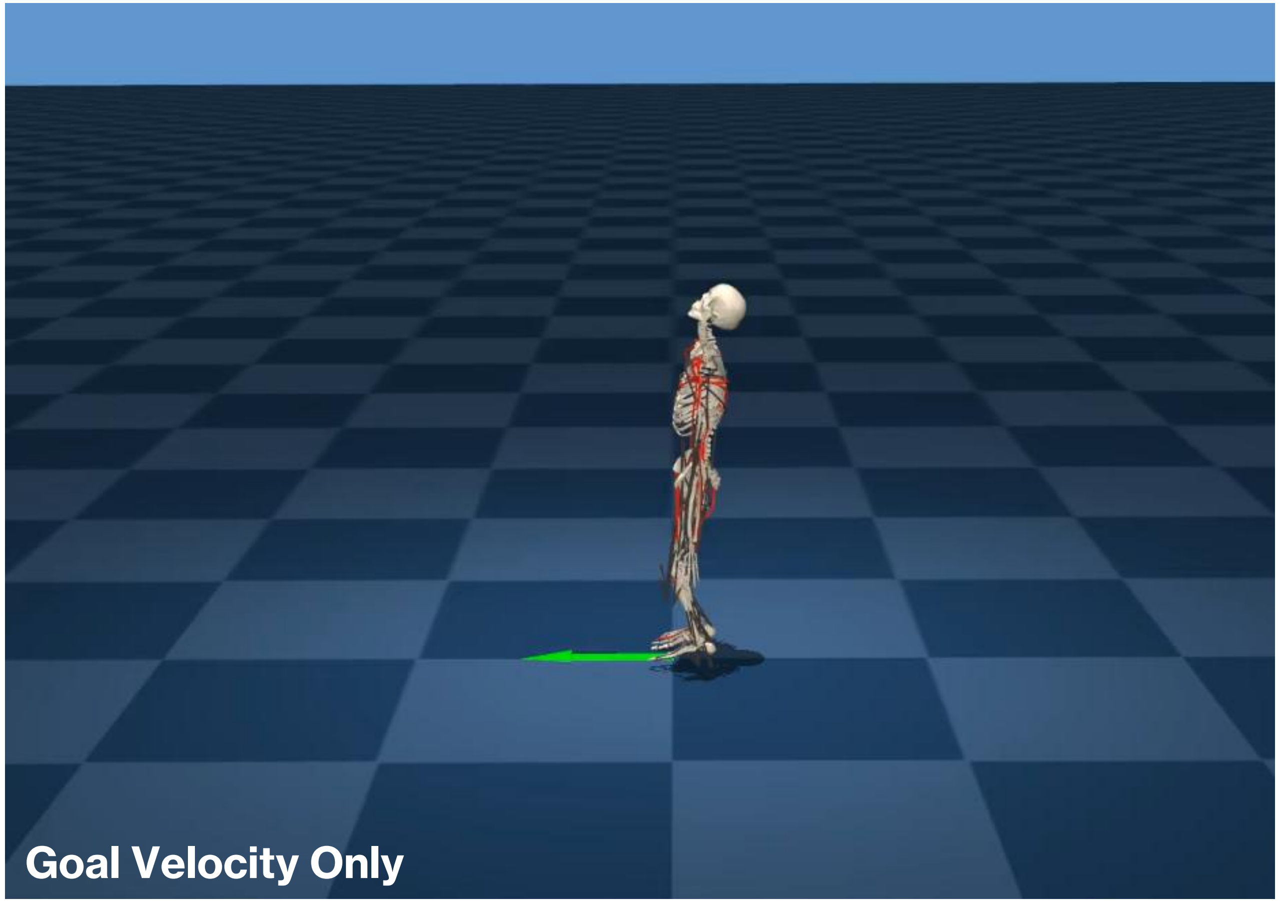

- Chimanoid learns a **natural** **quadrupedal** **gait:**

→ Our motion-free framework can generate **locomotion behaviors adapted to character's morphology**, without relying on motion priors.

Emergent Gait Strategies Across Morphologies

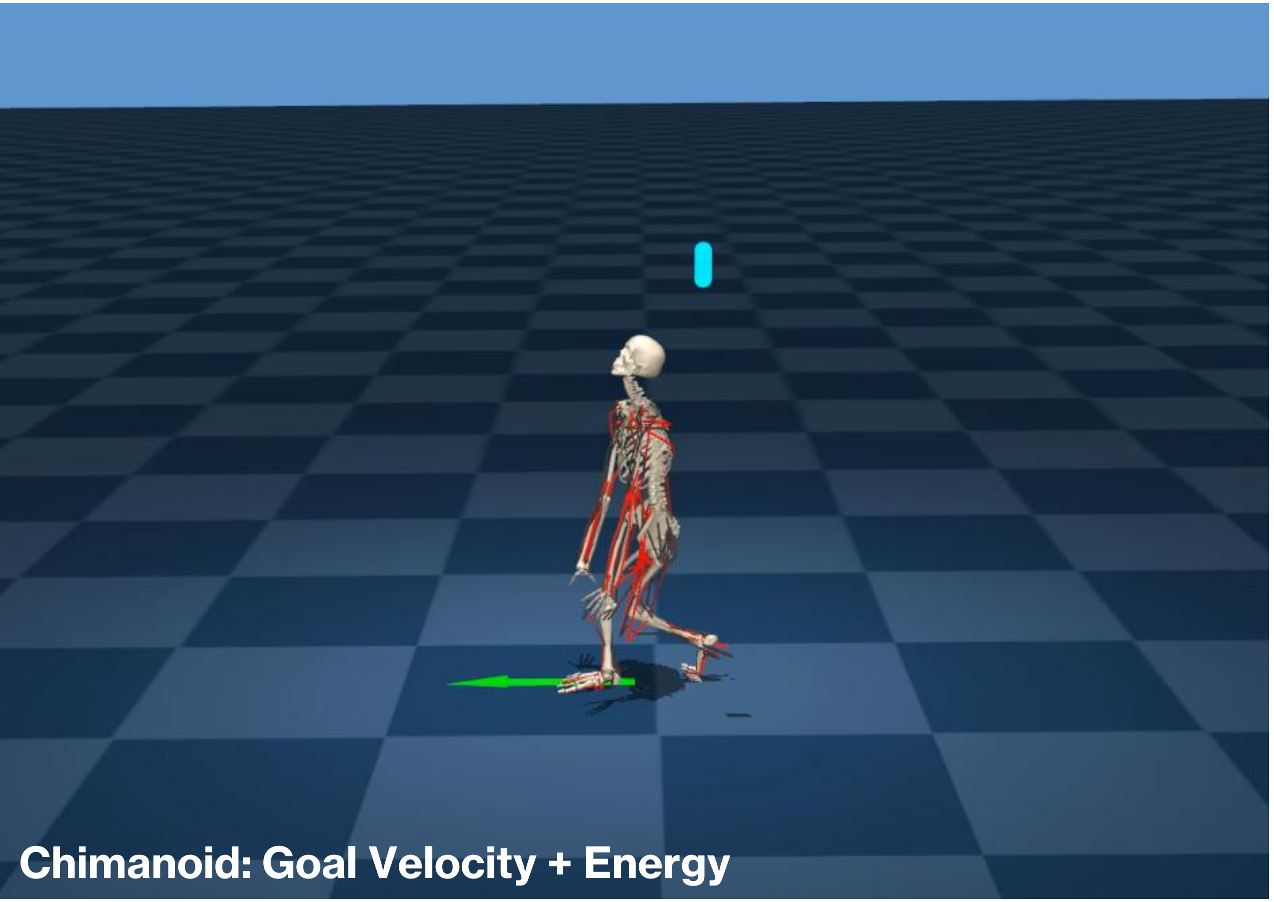

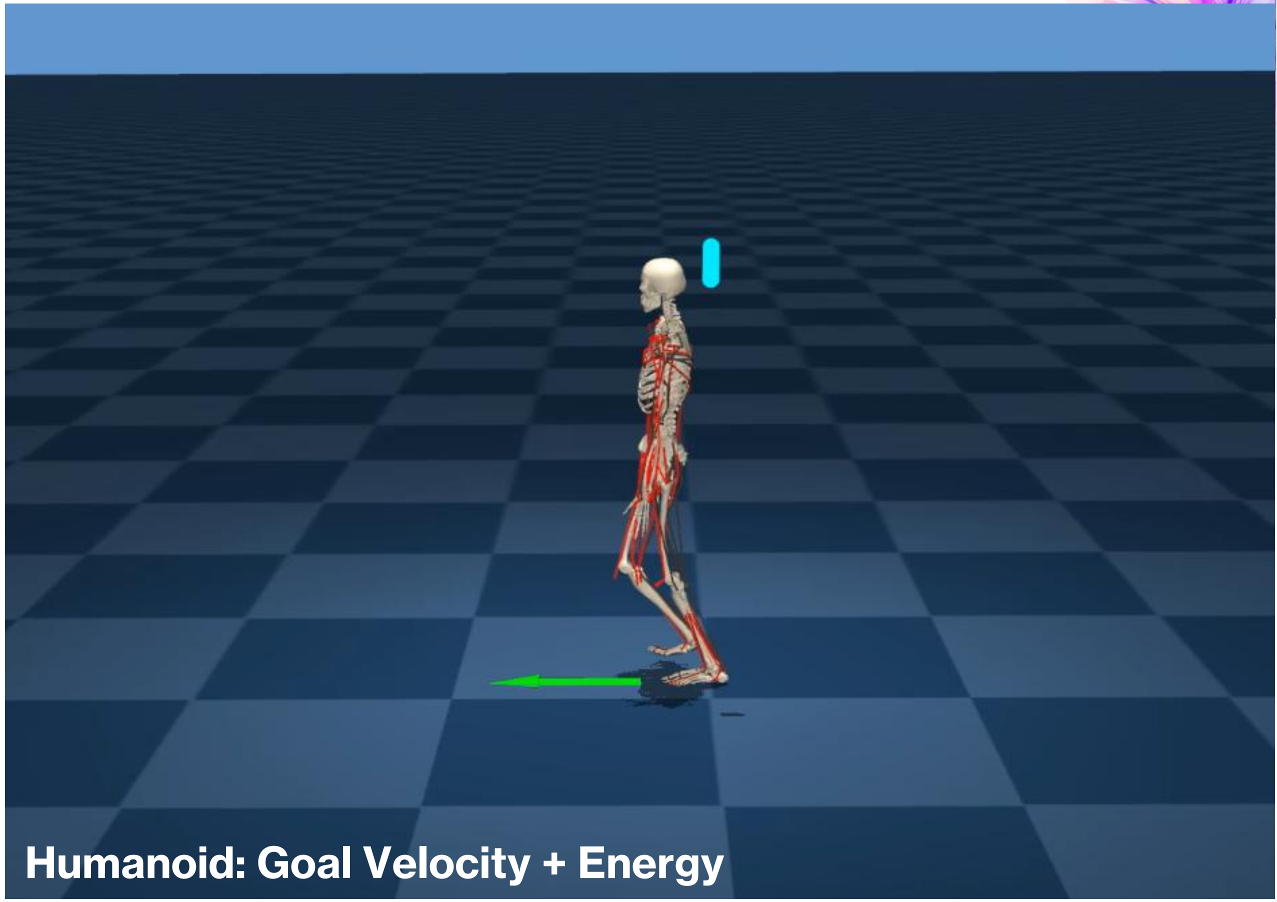

- **Chimanoid: bipedal** gait (high energy) → **quadrupedal** gait **(as energy decreases).**
- **Humanoid: never** adopts quadrupedal gaits.

→ **Energy-efficient gait depend on morphology** and can naturally emerge in our motion-free framework.

Random Sampling From Latent Space

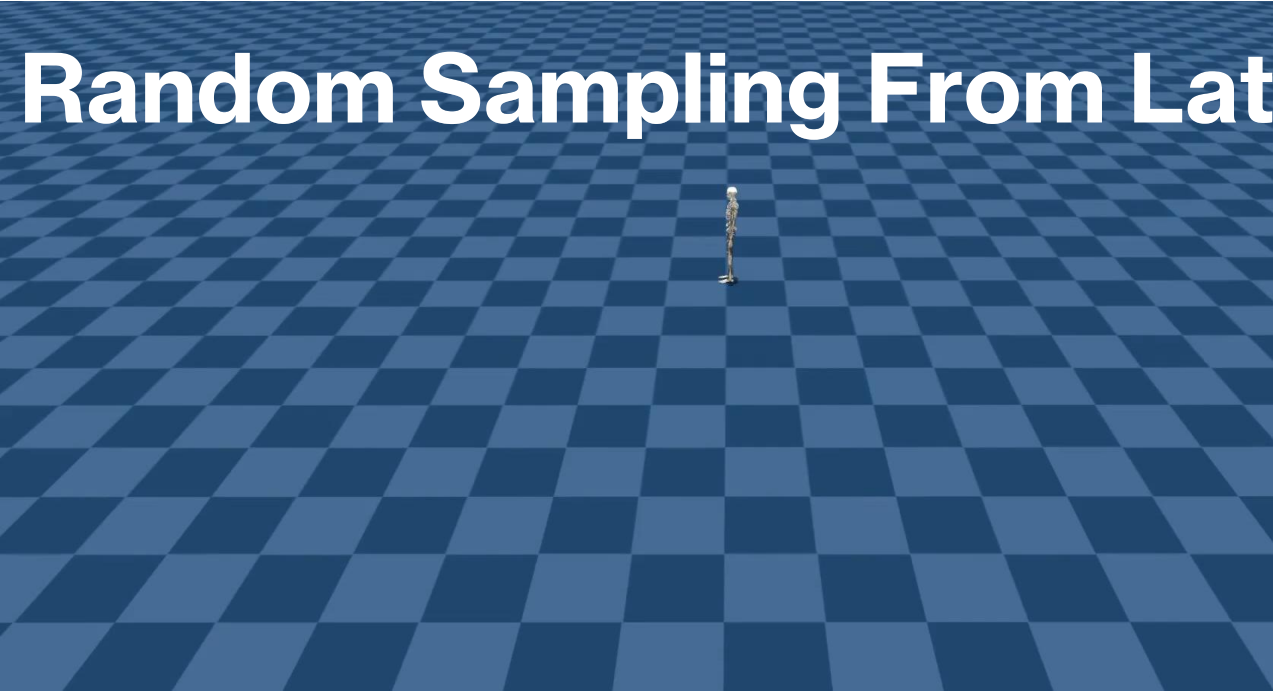

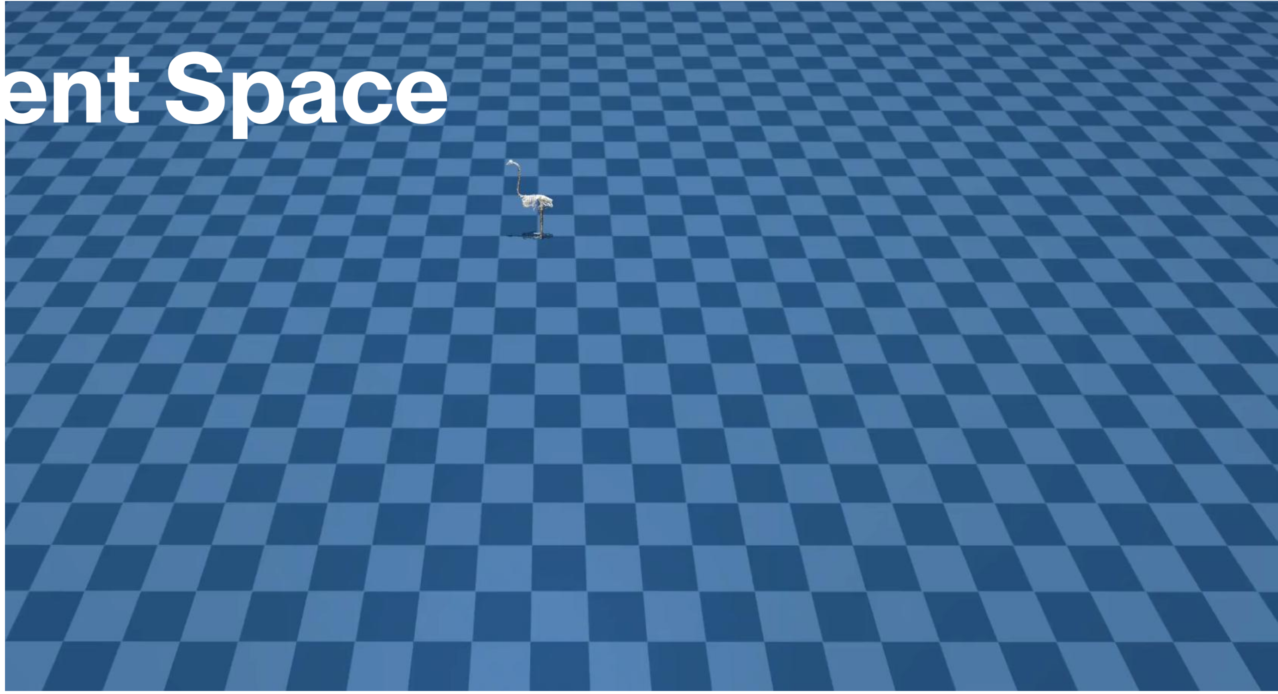

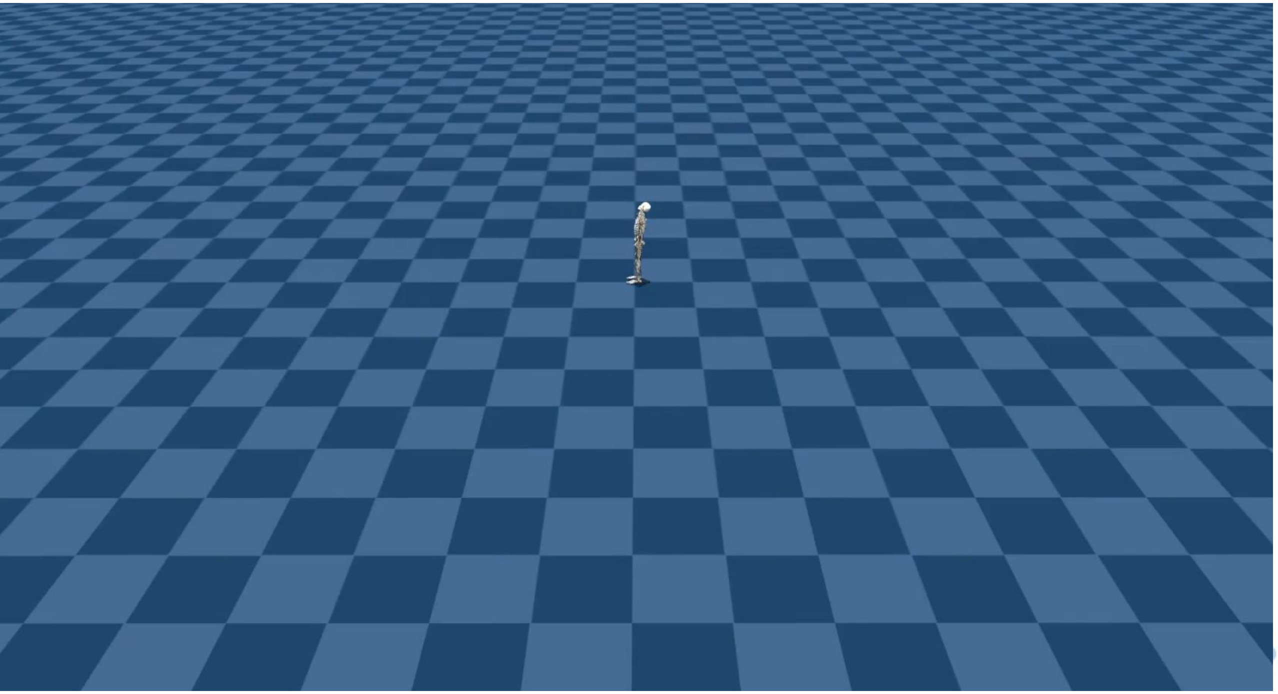

- Learned latent space (unit hypersphere, $||z|| = 1$) enabling random sampling to generate diverse locomotion.

Results

- High-Level Control Policies Trained for Downstream Tasks

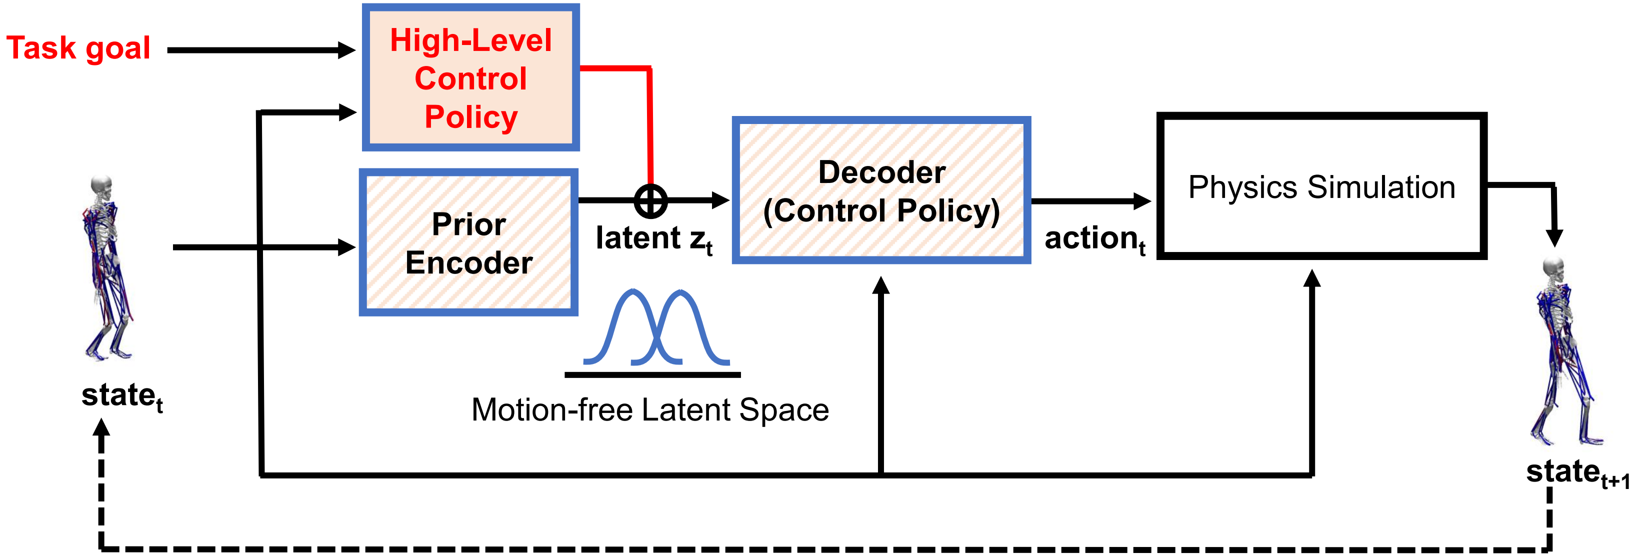

Point Goal Navigation

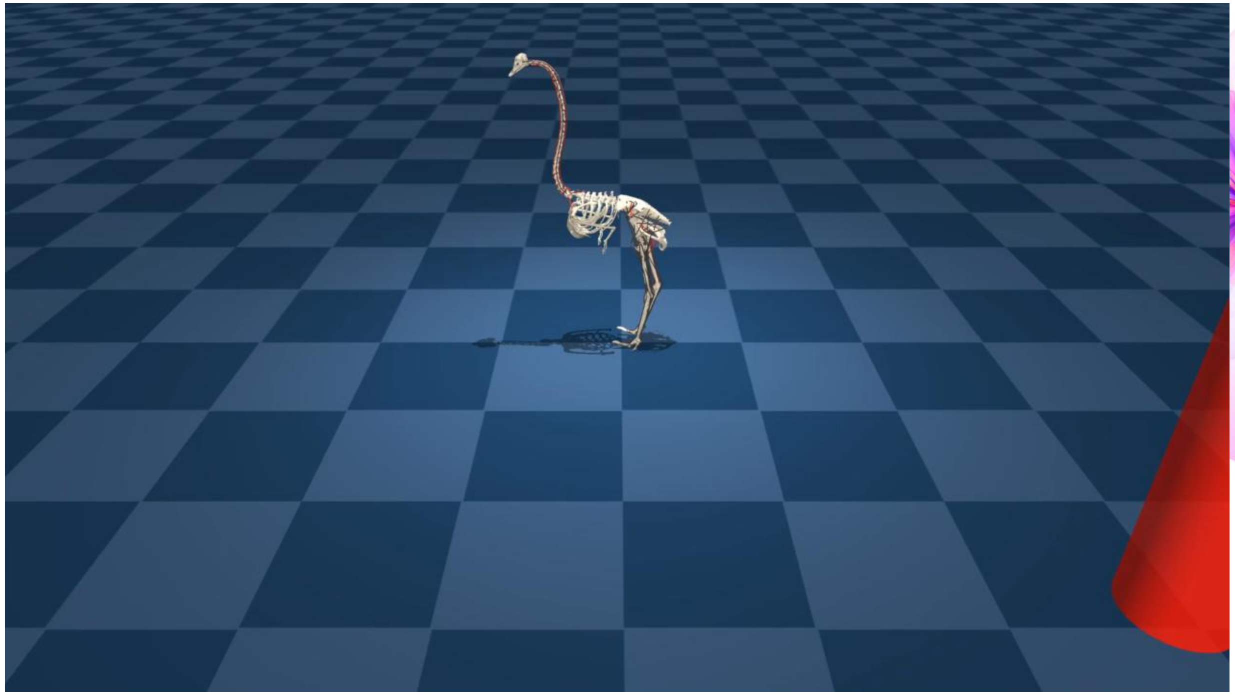

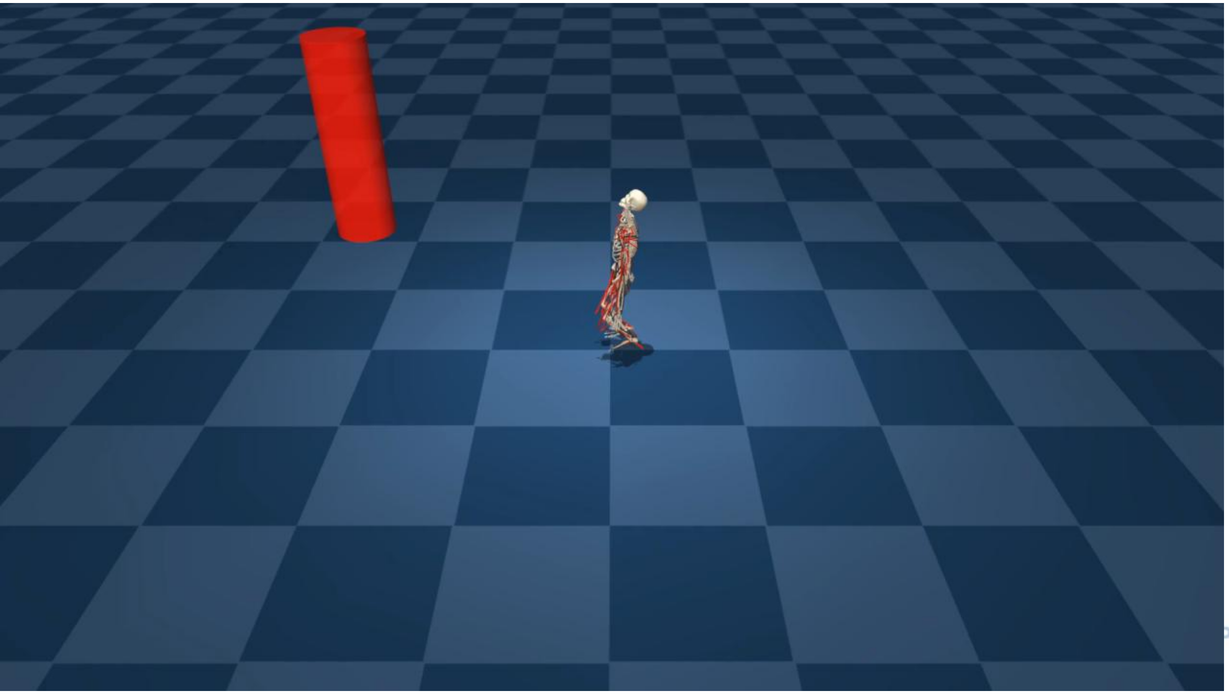

- **Goal input:** target position (relative).
- **Reward:** distance to target.

Path Following (Ribbon-shape)

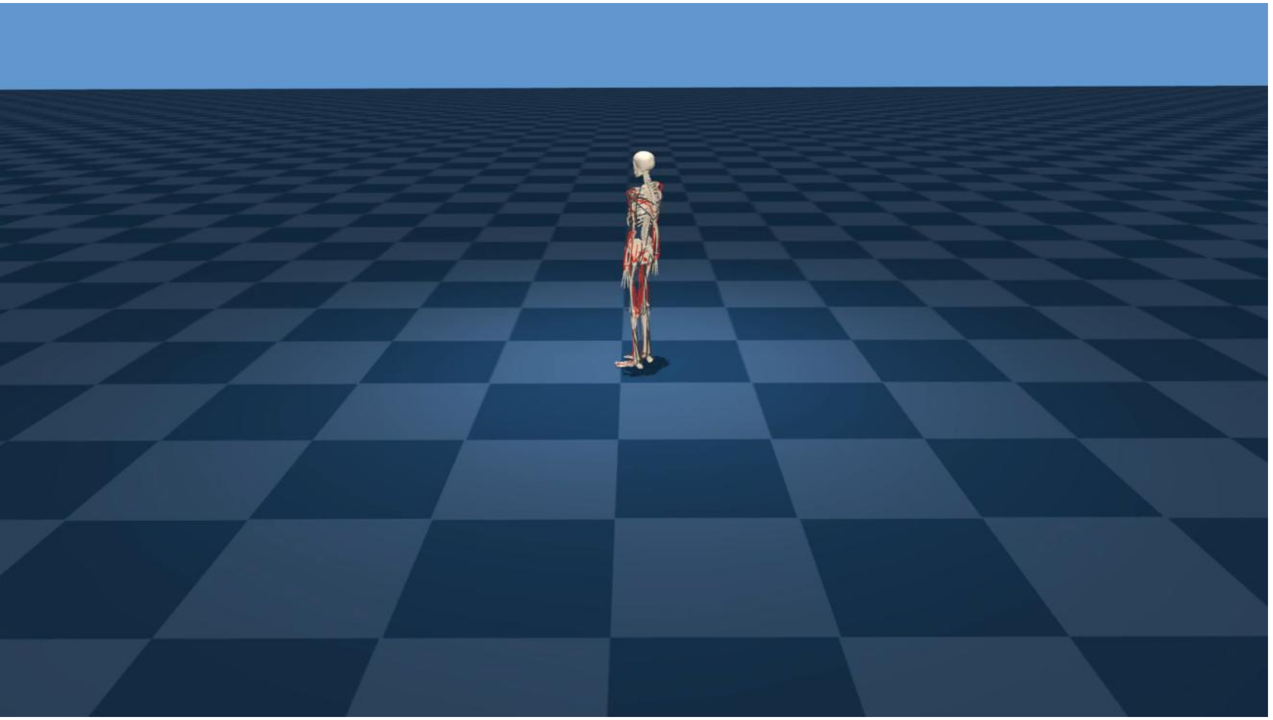

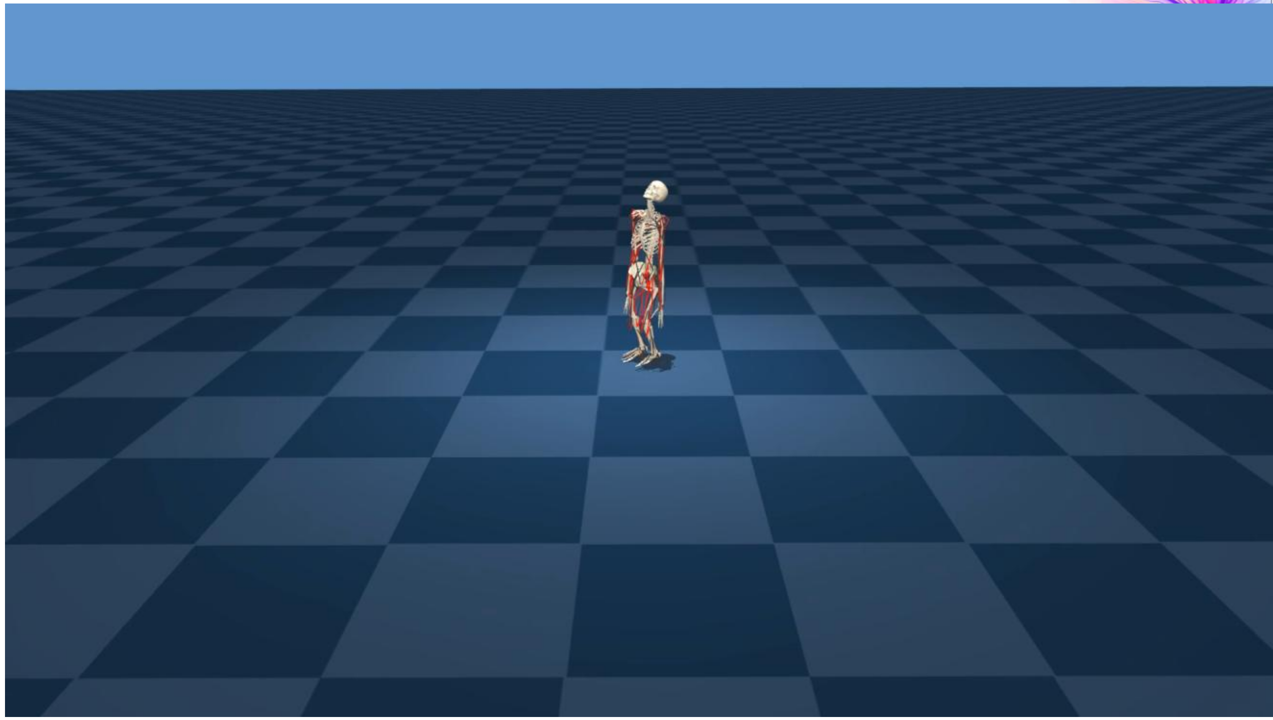

- **Goal input:** four consecutive target positions (relative).
- **Reward:** distance to the next target.

Effect of Temporally Averaged Loss Formulation

Ablation for L_{pose}

$$L_{objective} = \underbrace{L_{vel} + L_{dir}}_{\text{Control}} + \underbrace{L_{height} + L_{up}}_{\text{Balancing}} + \underbrace{L_{pose} + L_{energy}}_{\text{Biomechanical}}$$

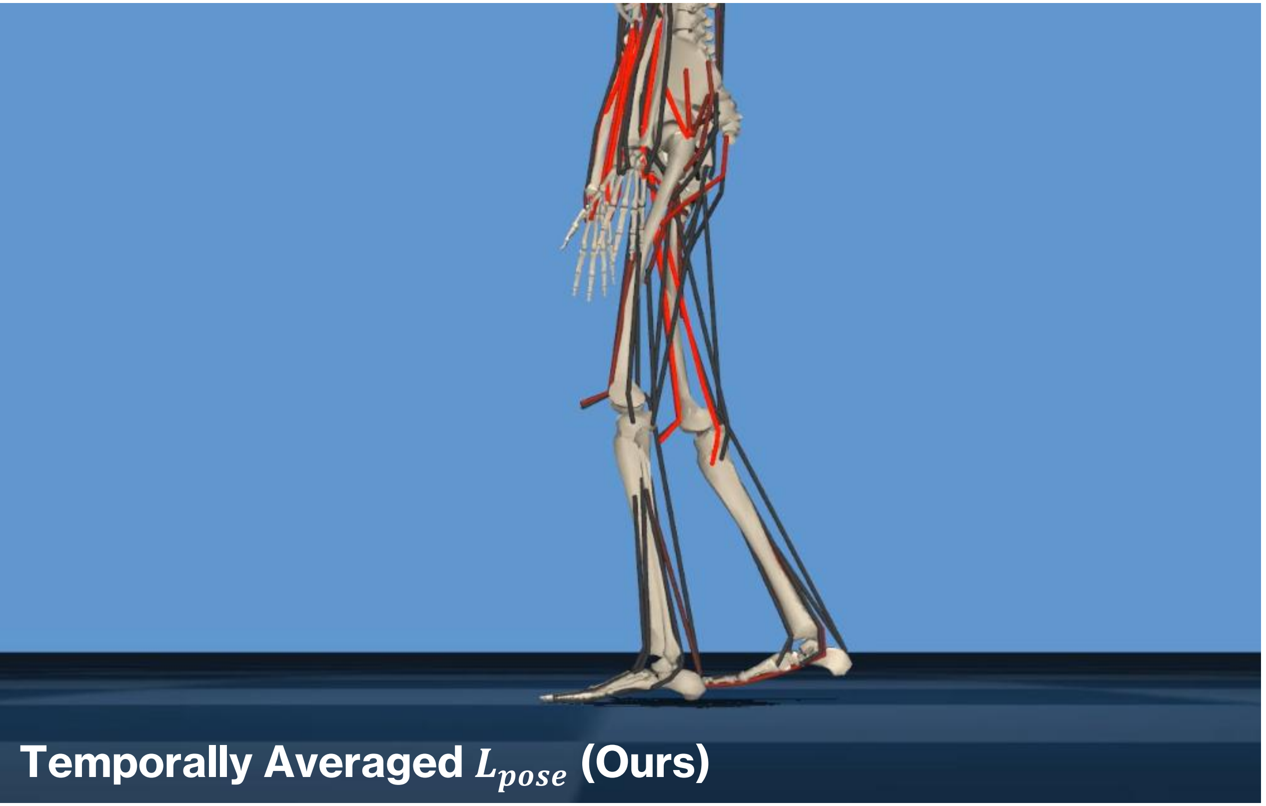

Temporally Averaged L_{pose} (Ours)

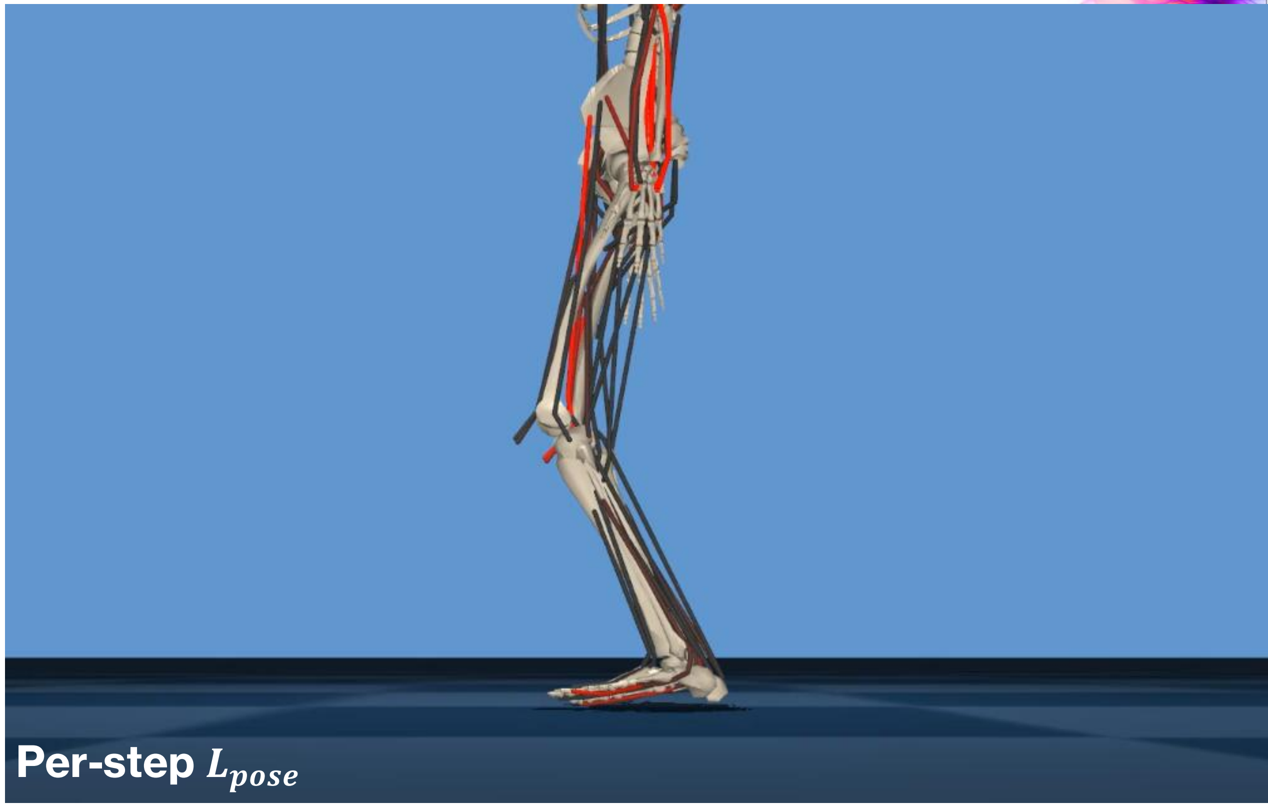

Per-step L_{pose}

Per-step L_{pose}

- Rigid frame-level pose matching.
- Slightly crouched, short-stepped gait with suppressed pelvic rotation.

Ablation for L_{up}

$$L_{objective} = L_{vel} + L_{dir} + L_{height} + L_{up} + L_{pose} + L_{energy}$$

Control | Balancing | Biomechanical

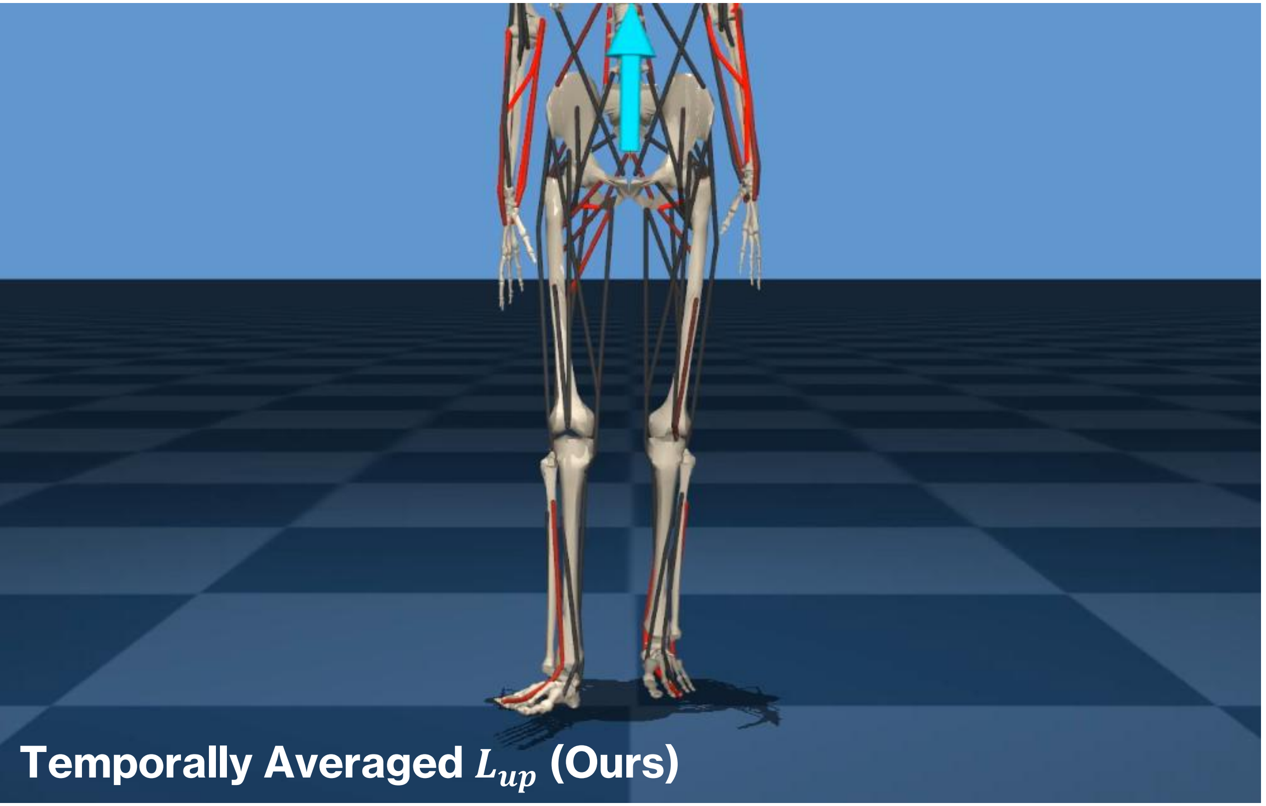

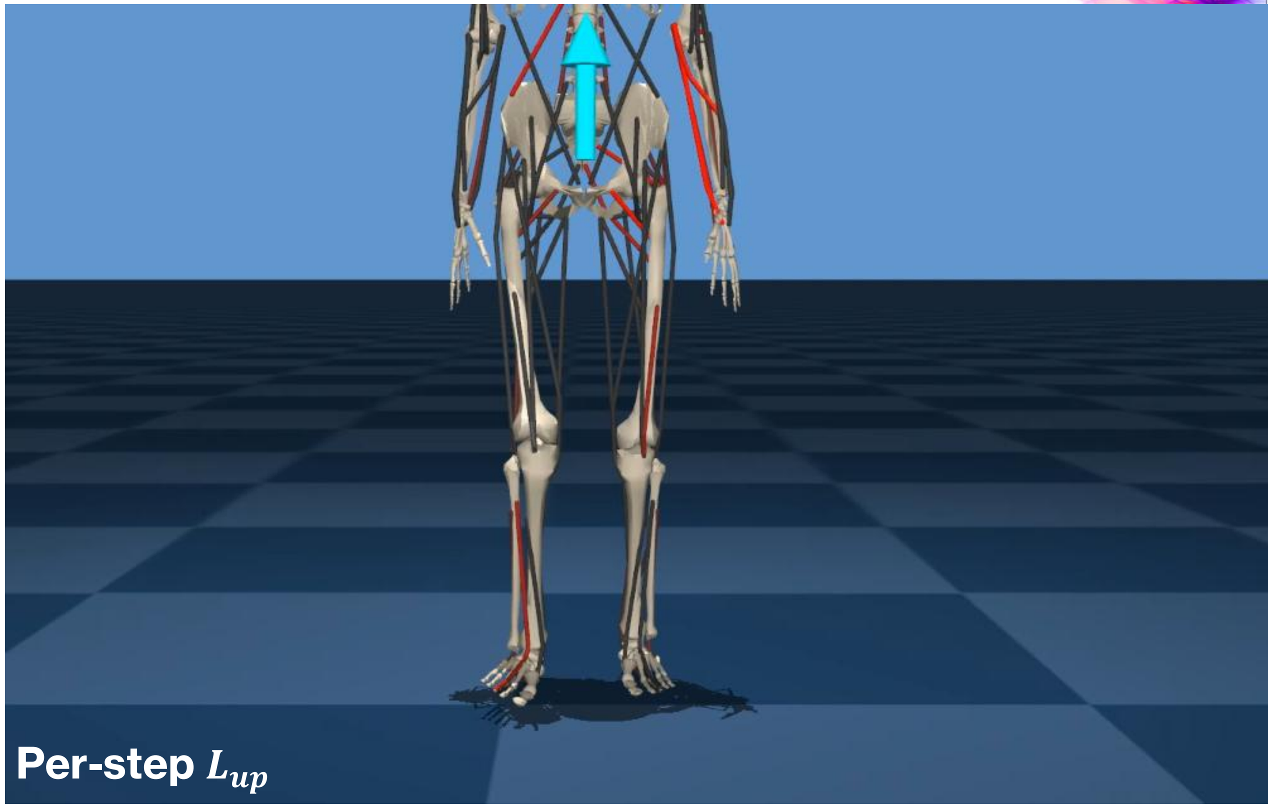

Per-step L_{up}

- Strongly constrains pelvic dynamics.
- Pelvic rotation appears nearly rigid without natural oscillation.

Comparison with Torque-Actuated Humanoid

Comparison with Torque-Actuated Humanoid

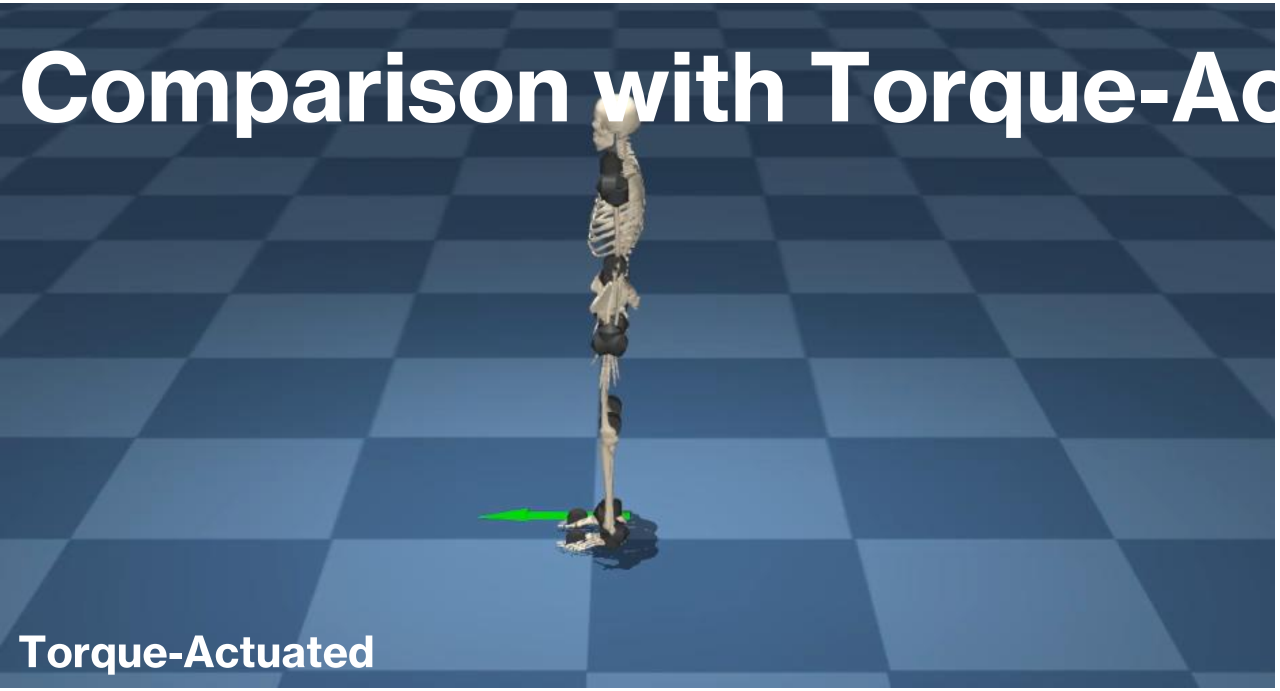

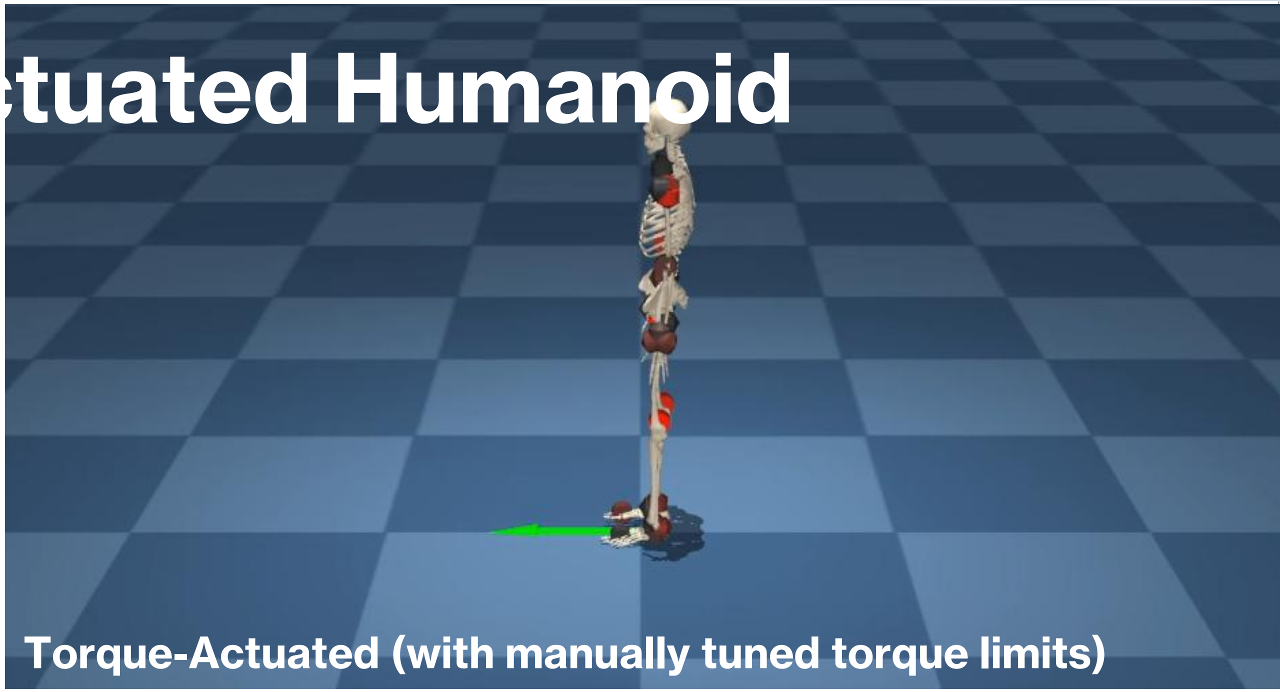

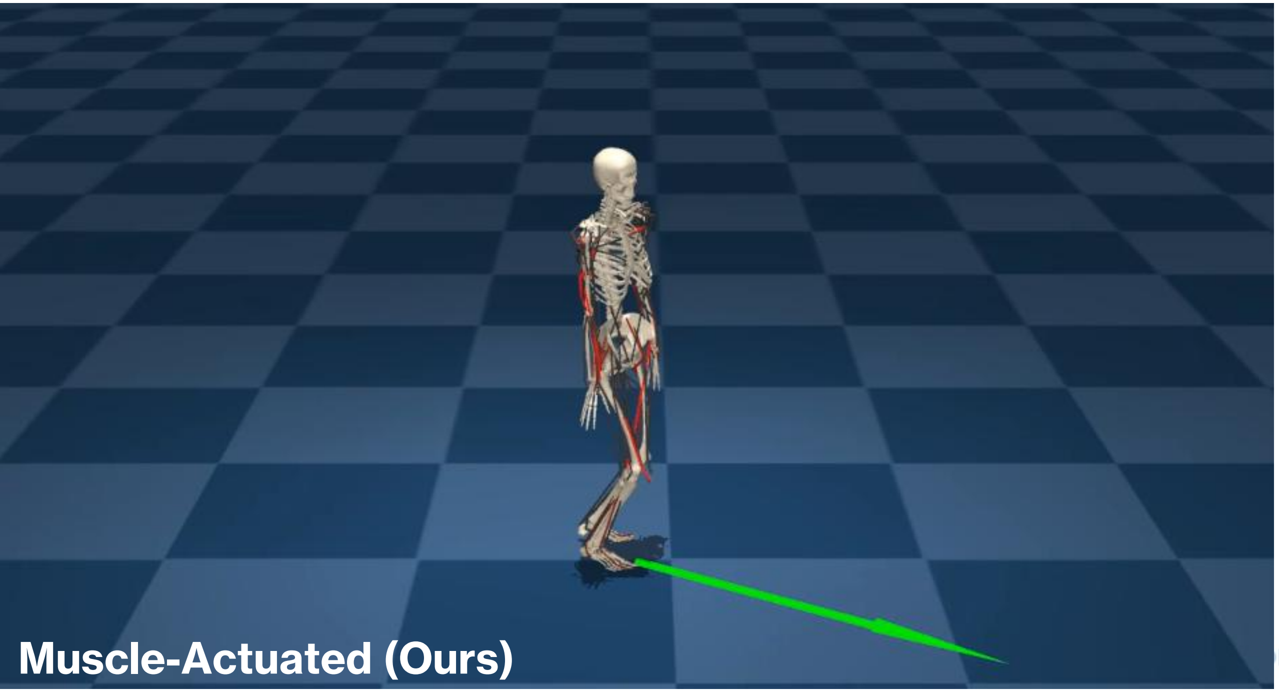

- **(Top-Left)** Torque-actuated humanoid failed to learn even basic balancing behavior.

- **(Top-Right)** Even with manual tuning of joint torque limits, it only learned an unnatural locomotion.

Conclusion & Discussion

Discussion

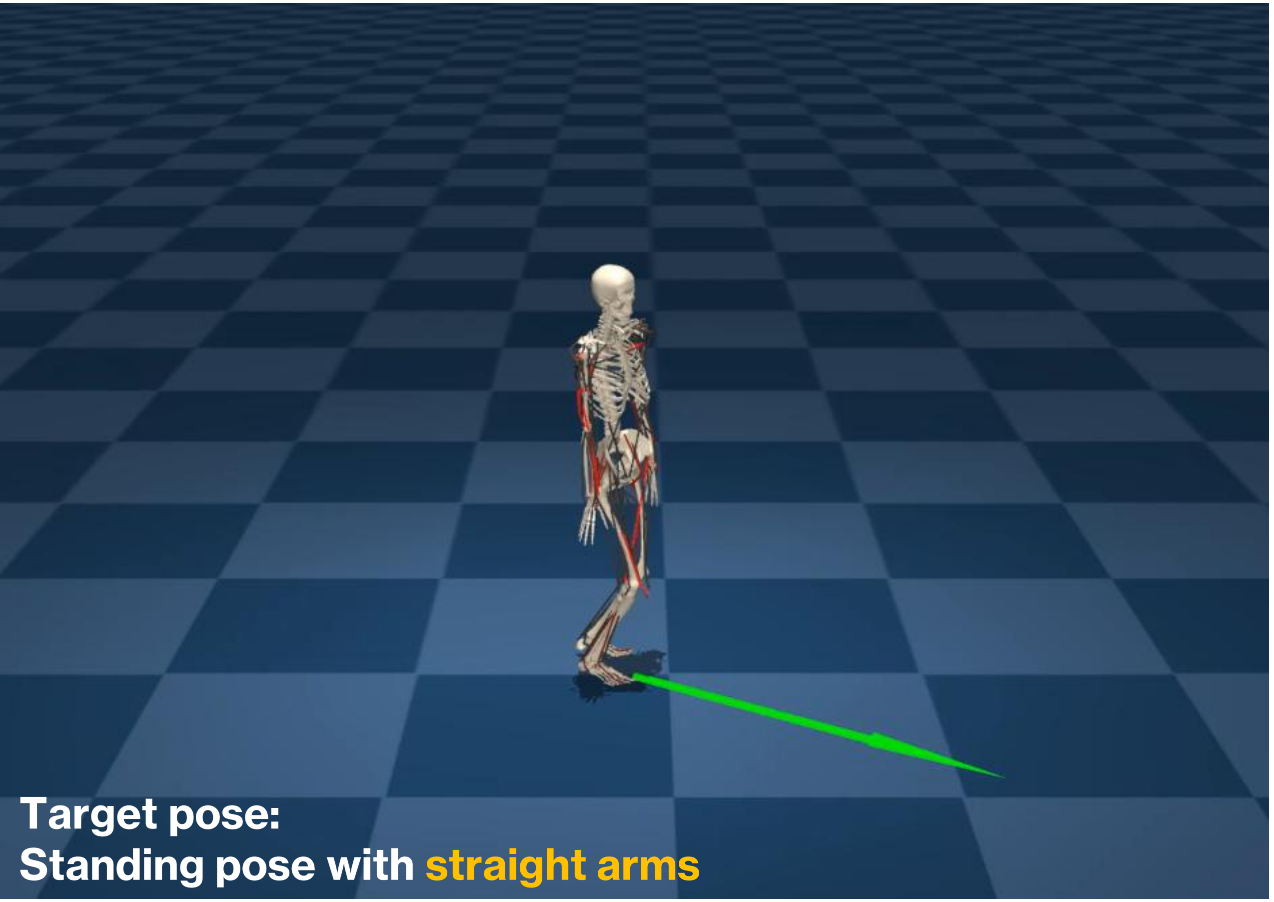

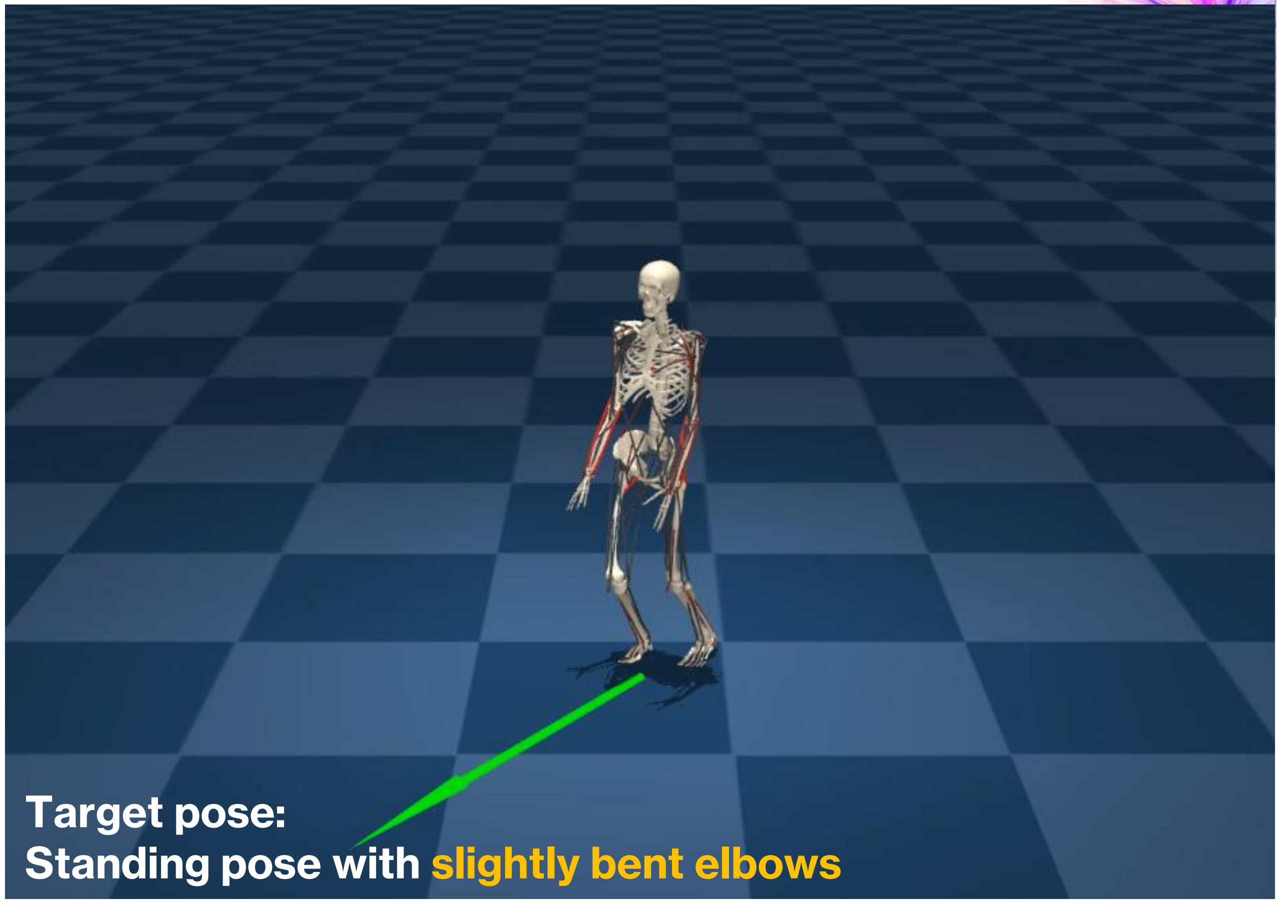

Learning with target poses **featuring slightly bend elbows.**

→ **Mitigates the straight and stiffness arm artifact.**

Contributions

- **Motion-free latent control learning:** Learn locomotion and latent-based control from musculoskeletal models without motion capture.
- **Cross-character generalization / Morphology-adaptive:** Apply to humanoid, non-humanoid, and synthetic characters with morphologically adaptive behaviors.
- **Locomotion objective loss / Temporally averaged formulation:** Integrate control, balance, biomechanics, and temporal averaging to induce natural gait cycles.
- **Diverse behavior modulation:** Randomize targets and energy during training for flexible control of form and intensity.

Future work

- **Reducing manual tuning efforts & finding optimal target postures.**


FreeMusco: Motion-Free Learning of Latent Control for
Morphology-Adaptive Locomotion in Musculoskeletal Characters
Minkwan Kim, Yoonsang Lee* (Hanyang University)
Come to our table to try the demo!