



# VIGOR: Visual Goal-In-Context Inference for Unified Humanoid Fall Safety


Osher Azulay   Zhengjie Xu   Andrew Scheffer   Stella X. Yu  
University of Michigan, Ann Arbor



1. Prone-to-stand recovery across fragmented support surfaces demands vision-guided contact selection and continuous whole-body rebalancing



2. Fall-mitigation recovery after a push towards stairs demands vision-guided anticipatory arm contact to arrest the fall and balance on a narrow step



3. Recovery behaviors across diverse terrain and initial conditions highlight the generality of vision-guided contact selection and whole-body rebalancing

Fig. 1. Vision-enabled unified fall safety for humanoids. A single learned policy integrates fall mitigation and stand-up recovery (sample runs in Rows 1–2) and transfers zero-shot to diverse terrains and conditions (poses over time shown as overlays in each image in Row 3).

**Abstract**—Reliable fall recovery is critical for humanoids operating in cluttered environments. Unlike quadrupeds or wheeled robots, humanoids experience high-energy impacts, complex whole-body contact, and large viewpoint changes during a fall, making recovery essential for continued operation.

Existing methods fragment fall safety into separate problems such as fall avoidance, impact mitigation, and stand-up recovery, or rely on end-to-end policies trained without vision through reinforcement learning or imitation learning, often on flat terrain. At a deeper level, fall safety is treated as monolithic data complexity, coupling pose, dynamics, and terrain and requiring exhaustive coverage, limiting scalability and generalization.

We present a unified fall safety approach that spans all phases of fall recovery. It builds on two insights: 1) Natural human fall and recovery poses are highly constrained and transferable from flat to complex terrain through alignment, and 2) Fast whole-body reactions require integrated perceptual-motor representations.

We train a privileged teacher using sparse human demonstrations on flat terrain and simulated complex terrains, and distill it into a deployable student that relies only on egocentric depth and proprioception. The student learns how to react by matching the teacher’s goal-in-context latent representation, which combines the next target pose with the local terrain, rather than separately encoding what it must perceive and how it must act.

Results in simulation and on a real Unitree G1 humanoid demonstrate robust, zero-shot fall safety across diverse non-flat environments without real-world fine-tuning. The project page is available at <https://vigor2026.github.io/>

## I. INTRODUCTION

Humanoids are intended to operate in the same cluttered and uneven environments as humans, yet they remain vulnerable when failures occur. Minor disturbances can quickly cascade into coupled rotations, impacts, and contact transitions, unfolding faster than conventional stabilization can correct. Unlike robots with inherently stable morphologies, humanoids must coordinate many joints and fleeting contacts under severe time pressure, making fall safety a whole-body problem in which early actions sharply limit later options.

Fall safety cannot be reduced to a post-impact reflex: How a robot falls determines which supports and contacts remain feasible, and whether standing up is possible at all. These choices depend on situational and terrain awareness, since contact feasibility and momentum redirection are shaped by local geometry invisible to proprioception alone. Fall safety thus requires the integration of terrain perception and coordinated whole-body control across the fall-recovery process.

Despite the intrinsic coupling between fall and recovery, existing approaches fragment fall safety along two axes and typically rely on blind, proprioceptive-only sensing.

First, many methods decompose fall-related behaviour into isolated subproblems: fall avoidance, impact mitigation, or




Fig. 2: **Factorized data generation yields sample-efficient imitation and scalable adaptation for humanoid fall safety learning.** Rather than treating **pose**, **time**, and **terrain** as a single monolithic data space requiring exhaustive coverage (left), we generate the same space by factorizing it into a small set of human pose trajectories from real-world demonstrations on flat terrain (middle) and independently varying terrain geometry in simulation (right), which can be arbitrarily complex.

stand-up recovery. Balance controllers focus on preventing loss of stability [7, 8], while stand-up controllers are typically invoked only after the robot comes to rest in a small set of predefined poses [23]. These components are designed largely in isolation and rely primarily on proprioceptive sensing, implicitly assuming flat ground and benign contact geometry. In practice, destabilization, impact, and recovery are tightly coupled: the robot’s orientation, contact sequence, and resting configuration are shaped by the surrounding terrain. Recovery therefore cannot be separated from *how the robot falls*, nor from *where and how contact occurs*.

Second, learning-based approaches that aim to address fall-related behaviors end-to-end often treat the problem as one of **monolithic data complexity**, and are typically trained without access to visual terrain perception. Reinforcement learning (RL) methods require hard reward engineering and extensive training, frequently leading to brittle or unnatural motions [40]. Imitation learning (IL) methods rely on dense trajectory demonstrations, often sourced from internet-scale human motion data, which transfer poorly across terrain geometry and contact conditions [42, 38, 43]. Both RL and IL methods treat environmental complexity as a monolithic data problem, entangling kinematics, dynamics, and terrain and requiring exhaustive coverage of their combinations.

In contrast, we advocate for a **factorized view of data complexity** (Fig. 2): *Our first insight is that natural human fall and recovery poses are far more constrained than they appear: Poses required on complex terrain often admit spatially aligned and physically compatible counterparts on flat terrain.* This observation allows pose and terrain variation to be treated as largely independent factors. As a result, a few human demonstrations collected on flat ground can serve as priors on kinematics and dynamics, while variation in terrain geometry, contact timing, and momentum redirection is handled through RL via interaction with the environment.

*Our second insight is that effective humanoid fall safety hinges on representing action goals in context.* Fast whole-body reactions during a fall require tight coupling between perception and action. Explicit prediction of terrain or dynamics [22, 44], or tracking reference trajectories during

execution [25, 18], is neither necessary nor sufficient: The humanoid must ultimately select actions conditioned jointly on its body state, the local terrain, and the next target pose. Inferring each of these factors in isolation demands solving difficult and error-prone subproblems, even though they must ultimately be considered together to form an action plan.

We therefore build the policy directly on a compact **visual goal-in-context latent representation**, which integrates the next target pose with the local terrain and body context in a single perceptual-motor space. Because contact feasibility and momentum redirection depend on terrain geometry that is not observable through proprioception alone, this latent is inferred from egocentric visual input together with short-term proprioceptive history. Rather than separately encoding the geometry the robot must perceive and the pose it must reach, the goal-in-context latent provides exactly the information required to formulate an action plan in situ.

Based on these insights, we introduce **VIGOR**, *Visual Goal-In-Context Inference for Unified Humanoid Fall Safety*. VIGOR treats fall safety as a single, unified task spanning fall avoidance, impact mitigation, and stand-up recovery. A privileged teacher policy is trained via RL using sparse human demonstrations on flat terrain together with access to local terrain geometry, yielding terrain-aware reactive behaviours. This knowledge is then distilled into a deployable student policy that operates using only egocentric depth observations and short-term proprioceptive history. By matching the teacher’s goal-in-context latent representation, the student learns to react across diverse fall scenarios without explicitly decomposing perception and control.

We evaluate VIGOR extensively in simulation and demonstrate successful zero-shot transfer to the real world across a wide range of fall and recovery conditions, validating a unified and vision-enabled approach to humanoid fall safety.

## II. RELATED WORK

### A. Humanoid Fall Mitigation and Standing-Up Control

Navigating real-world terrain poses significant stability risks for humanoid robots, whose high center-of-mass and complex contact dynamics make falls both likely and difficult to recover

from [15, 10, 29, 28, 12]. Prior work typically decomposes recovery into two disjoint stages: fall mitigation during descent and standing up after impact. Classical fall mitigation methods focus on maintaining balance or redirecting momentum to reduce impact forces [8, 7], while learning-based approaches similarly emphasize safe landing or damage reduction without explicitly addressing subsequent recovery [16]. Standing-up control, by contrast, is often treated as a separate problem and triggered only after the robot has settled into a small set of predefined post-impact configurations [33]. Although recent learning-based methods have improved robustness across diverse initial poses [35, 14, 13], they generally assume the fall has already concluded and do not reason about the dynamics leading to impact. Only a few works attempt to unify fall mitigation and standing-up within a single policy [11, 39], but these approaches typically rely on limited sensing or blind operation. In contrast, we learn a *unified, visually grounded* policy that explicitly couples fall mitigation and recovery under diverse terrain conditions.

### B. Visual Whole-Body Control

Methods in vision-based control have recently enabled humanoids to perform complex whole-body behaviors under partial observability [44, 6]. Visual input has been used to support loco-manipulation and contact-rich interaction [9, 20, 41], as well as to guide navigation and locomotion across challenging environments [22, 19, 5]. Other work explores emergent active perception and internal representations to maintain control consistency under changing viewpoints [24, 32]. Despite this progress, fall recovery presents a difficult perceptual regime where rapid body rotations, self-occlusion, and intermittent ground contact lead to narrow and unstable visual observations during critical control phases. To the best of our knowledge, our work is the first to utilize visual observations to facilitate more robust and adaptive humanoid fall recovery behaviors.

### C. Motion Priors and Style-Constrained RL

Reinforcement learning offers a flexible framework for humanoid control, but learning safe and coordinated whole-body behaviors remains challenging due to high-dimensional action spaces and complex dynamics [34, 4]. Prior work therefore relies on carefully shaped multi-term reward functions [37, 44] or constrains the control problem through reduced body representations or limited link actuation [19]. Human motion priors provide an alternative form of regularization, inducing structured behavior with minimal reward engineering.

Reference-guided RL methods such as DeepMimic [26] and AMP [27] demonstrate strong whole-body control on flat terrain, but often rely on dense trajectory tracking or periodic motion assumptions, which become limiting in highly non-periodic, contact-rich behaviors such as fall recovery. Recent work relaxes dense tracking through higher-level structural guidance, including keyframe-based objectives [42, 36], style-constrained learning [38], or hybrid formulations that combine imitation on flat terrain with residual learning and task-driven

adaptation on uneven terrain [43]. These methods allow deviation from demonstrations when needed, but still depend on sufficient motion–terrain coverage, and fall recovery remains challenging due to its non-periodic nature and strong coupling to terrain and contact geometry. We therefore treat human fall-recovery demonstrations as *sparse structural priors* rather than full trajectory targets.

## III. METHOD

We propose VIGOR, a unified framework for humanoid stand-up, fall mitigation, and recovery in unstructured environments, as illustrated in Fig. 3. Unlike pipelines that treat falling, impact mitigation, and standing as separate modules, VIGOR learns the full process in a single control policy conditioned on onboard egocentric depth sensing. Following a destabilizing disturbance, the robot may enter arbitrary configurations involving high-energy impacts, whole-body contact, and large changes in body orientation. At each timestep  $t$ , the robot receives proprioceptive measurements  $\mathbf{o}_t$  and an egocentric depth image  $\mathbf{I}_t$  from a head-mounted camera, and outputs joint-space control targets  $\mathbf{a}_t$ . The policy must continuously regulate contact and body motion during the fall and produce a stand-up behavior that returns the robot to an upright configuration. The system contains two components: (1) a privileged teacher that observes sparse keyframe targets  $\mathbf{k}_t$ , the current observation  $\mathbf{o}_t$ , and local terrain samples  $\mathbf{h}_t$  to provide high-level recovery structure; and (2) a deployable student that reconstructs the teacher’s *goal-in-context* latent using egocentric depth and short-term observation history.

### A. Motion Collection and Sparse Keyframe Extraction

We obtain fall–recovery demonstrations from monocular human videos recorded on flat ground, covering forward, sideways, and backward recovery behaviors, shown in Appendix Fig. 12. Each video is processed using the VideoMimic pipeline [2] to reconstruct full-body 3D motion and fit an SMPL model [21], which is then retargeted to the Unitree G1 humanoid via a kinematics-aware mapping with conservative joint-limit constraints to avoid retargeting artifacts. Following reconstruction and retargeting, we retain a total of nine high-quality recovery motion sequences from three different body sizes. Rather than using full trajectories as strict imitation targets, we extract a sparse set of uniformly sampled keyframes to provide coarse temporal structure for the motion. These keyframes serve as *high-level structural priors* for the privileged teacher, guiding learning without over-constraining recovery behavior.

During training, reference keyframes are coarsely projected onto the terrain following geometric alignment inspired by [36]. Concretely, we apply a vertical projection

$$\Delta z = \max_{i=1, \dots, N_{\text{links}}} (h(\mathbf{p}_i^{\text{ref}}) - z_i^{\text{ref}}),$$

where  $h(\cdot)$  queries the terrain height at the reference link position  $\mathbf{p}_i^{\text{ref}}$  and  $z_i^{\text{ref}}$  is its vertical coordinate; all reference poses are shifted by  $\Delta z$  to ensure clearance above the terrain. Reference poses are initialized at the terrain center

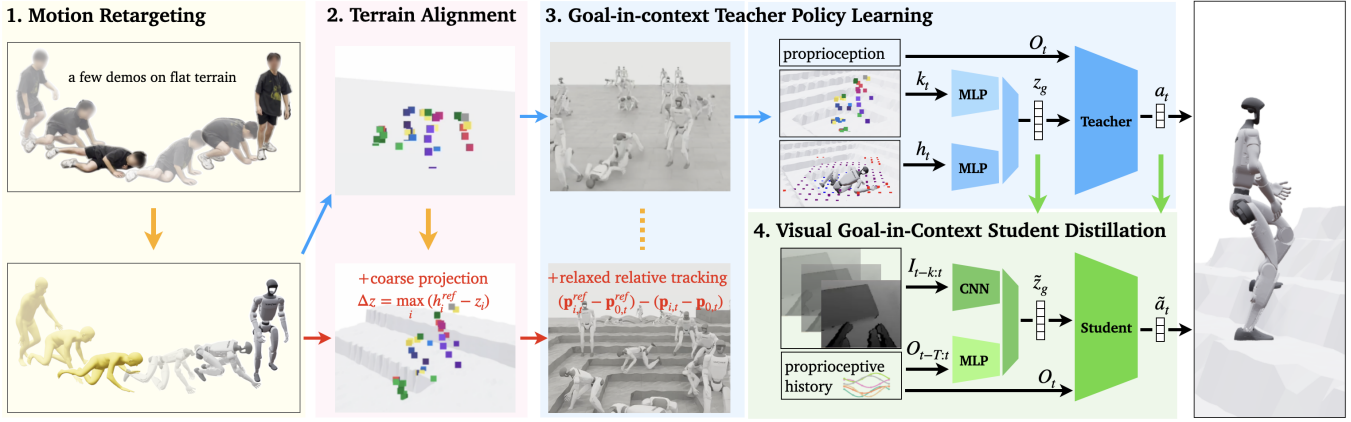


Fig. 3: **Overview of VIGOR.** **1)** Motion retargeting: human fall–recovery demonstrations are kinematically retargeted to the robot. **2)** Terrain alignment: reference poses are used directly on flat terrain and coarsely projected onto uneven terrain to provide sparse tracking targets. **3)** Goal-in-context teacher policy learning: a privileged teacher policy is trained with RL to acquire a goal-in-context representation that encodes the immediate recovery target pose together with local terrain information. **4)** Visual goal-in-context student distillation: a student policy distills the teacher’s terrain-aware recovery behavior from egocentric depth and short-term proprioceptive history for deployment.

(origin), with small  $x$ – $y$  perturbations constrained relative to the terrain frame. We further randomize reference conditioning and initialization by sampling different reference trajectories and different starting points along each trajectory to improve coverage of recovery configurations. Additional details are provided in the Appendix Section D.

### B. Privileged Goal-in-Context Teacher

We train a privileged fall-mitigation and recovery teacher policy  $\pi_\theta$  using PPO [31]. Training proceeds in two stages: the policy is first trained on flat terrain to acquire basic fall mitigation and stand-up behaviors, and then continued on randomized non-flat terrains to handle diverse contact geometry and perturbations. The teacher receives the observation  $\mathbf{o}_t = (\mathbf{o}_t^{\text{prop}}, \mathbf{k}_t, \mathbf{h}_t)$ , where  $\mathbf{o}_t^{\text{prop}}$  contains proprioceptive features,  $\mathbf{k}_t$  encodes sparse keyframe targets, and  $\mathbf{h}_t$  is a privileged terrain scan containing local height information. The reference and terrain signals are fused into a single *goal-in-context* latent  $\mathbf{z}_g = g(\mathbf{k}_t, \mathbf{h}_t)$ , which summarizes the immediate recovery target pose together with local terrain information. The teacher actor conditions on this latent and proprioception to produce joint targets:  $\mathbf{a}_t = \pi_\theta(\mathbf{z}_g, \mathbf{o}_t^{\text{prop}})$ .

While the teacher is trained in a tracking-based manner similar to prior motion imitation policies [2, 26], the reference signals are intentionally sparse and enforced only at a coarse, relative level. Together with access to privileged terrain observations, this under-specification allows RL to resolve contact timing, body placement, and terrain-dependent execution details beyond what is prescribed by the demonstrations.

### C. Reward Design

We design a composite reward function to support fall mitigation, stabilization, and recovery within a single policy:

$$r_t = r_t^{\text{imit}} + r_t^{\text{reg}} + r_t^{\text{post}}$$

where  $r_t^{\text{imit}}$  is a DeepMimic-style tracking term [26],  $r_t^{\text{reg}}$  aggregates standard motion-regularization penalties, and  $r_t^{\text{post}}$  provides post-recovery stabilization by rewarding upright hold and suppressing residual motion once standing. All reward terms are summarized in Table I.

**Motion imitation.** We project reference poses, recorded on flat ground, onto the simulated terrain before computing tracking errors, using a simple geometric alignment procedure inspired by MaskedMimic [36]. This alignment preserves the demonstrated motion structure but is not physically exact on uneven terrain, leaving small residual offsets introduced by the projection itself. Tracking full world-frame positions  $\|\mathbf{p}_{i,t}^{\text{ref}} - \mathbf{p}_{i,t}\|$  would therefore penalize projection residuals and terrain-dependent deviations that are necessary for recovery. Instead, we define imitation targets in a root-relative canonical frame by comparing  $(\mathbf{p}_{i,t}^{\text{ref}} - \mathbf{p}_{0,t}^{\text{ref}}) - (\mathbf{p}_{i,t} - \mathbf{p}_{0,t})$ , which preserves relative body configuration while remaining insensitive to projection artifacts, allowing demonstrations to act as structural priors rather than exact pose targets.

**Regularization and safety.** The term  $r_t^{\text{reg}}$  aggregates standard regularization penalties, including joint limit violations, joint velocities, accelerations, momentum change, action smoothness, undesired contacts, and contacts near terrain edges. These terms stabilize learning and reduce unsafe behaviors without prescribing a specific recovery strategy.

**Post-recovery stabilization.** Once the robot reaches a stable upright configuration,  $r_t^{\text{post}}$  encourages stable upright posture by suppressing residual base motion and joint velocities.

### D. Egocentric Student Policy

At deployment time the robot has no access to privileged terrain or reference motion. The student therefore learns to infer the teacher’s goal-in-context latent using only egocentric depth and short-term proprioceptive history. The same set of

TABLE I: **Reward terms used for fall-recovery learning.** Tracking rewards use a Gaussian kernel  $f(d; \sigma) = \exp(-d^2/\sigma)$  applied to pose, velocity, and joint-space errors. Regularization and safety terms follow standard formulations widely adopted in prior humanoid control work [39, 18, 37].

Reward	Definition	Weight
<i>Motion imitation</i>		
RB pos. track.	$\exp(-\ \Delta x\ ^2/\sigma)$	1.25
RB rot. track.	$\exp(-\ \Delta \theta\ ^2/\sigma)$	0.50
RB lin. vel. track.	$\exp(-\ \Delta v\ ^2/\sigma)$	0.125
RB ang. vel. track.	$\exp(-\ \Delta \omega\ ^2/\sigma)$	0.125
Joint pos. track.	$\exp(-\ \Delta q\ ^2/\sigma)$	0.50
Joint vel. track.	$\exp(-\ \Delta \dot{q}\ ^2/\sigma)$	0.125
<i>Regularization &amp; safety</i>		
Torque penalty	$\ \tau\ ^2$	$-1.0 \times 10^{-6}$
Torque limit	$\max(0,  \tau  - \tau_{\max})$	-0.1
Joint pos. limit	$\max(0, q - q_{\max}) + \max(0, q_{\min} - q)$	-10.0
Joint vel. limit	$\max(0,  \dot{q}  - \dot{q}_{\max})$	-5.0
Joint acc.	$\ \ddot{q}\ ^2$	$-2.5 \times 10^{-7}$
Momentum change	$\ \Delta p\ $	$-5.0 \times 10^{-3}$
Body yank	$\ \Delta F\ ^2$	$-2.0 \times 10^{-6}$
Joint vel. penalty	$\ \dot{q}\ ^2$	$-1.0 \times 10^{-4}$
Action rate	$\ a_t - a_{t-1}\ ^2$	-0.1
Undesired contacts	$r_{\text{contacts}}$ [18]	-0.1
Support at edge	$r_{\text{foothold}}$ [37]	-1.0
<i>Post-recovery stabilization</i>		
Head height	$\exp(-\ \Delta h\ ^2/\sigma)$	0.25
Base linear velocity	$\ \mathbf{v}_{\text{base}}\ ^2$	-0.01
Base angular velocity	$\ \boldsymbol{\omega}_{\text{base}}\ ^2$	-0.05

sparse demonstrations used to train the teacher defines the space of recovery behaviors available to the student, which learns to infer an appropriate target pose from history rather than relying on explicit reference signals.

The student receives a short history of egocentric depth images  $\mathbf{I}_{t-k:t}$  and proprioceptive observations  $\mathbf{o}_{t-T:t}^{\text{prop}}$ . A perceptual encoder maps the stacked depth images to a feature  $\mathbf{f}_t^{\text{img}}$ , while a temporal encoder summarizes the proprioceptive history into  $\mathbf{f}_t^{\text{hist}}$ . These are fused to produce a predicted goal latent  $\tilde{\mathbf{z}}_g$ , which encodes the control objective as the target pose conditioned on terrain. The student actor outputs joint-space actions conditioned on the predicted latent and the current proprioceptive observation,  $\tilde{\mathbf{a}}_t = \pi_\phi(\tilde{\mathbf{z}}_g, \mathbf{o}_t^{\text{prop}})$ .

Training is supervised by the teacher in both the latent and action spaces. The latent-matching loss encourages the student to reconstruct the teacher’s goal-in-context representation, while behavioral cloning aligns the student’s actions with the teacher’s:

$$\mathcal{L}_{\text{goal}} = \|\tilde{\mathbf{z}}_g - \mathbf{z}_g\|^2, \quad \mathcal{L}_{\text{BC}} = \|\tilde{\mathbf{a}}_t - \mathbf{a}_t\|^2$$

A DAGger-style [30] mixing schedule gradually replaces teacher actions with student actions during rollouts while continuing to provide supervision in both the latent and action spaces, enabling deployment without privileged information.

Notably, the student is distilled from unstable initial states, biasing it toward reactive stabilization rather than deliberate falling. As a result, the policy maintains a stable posture in nominal conditions and engages in fall mitigation and recovery




Fig. 4: **Terrains used for training.** From top to bottom: rough, waves, slope, inverted slope, stairs, and inverted stairs. The figure shows three representative difficulty levels per terrain for visualization.

under disturbance.

#### E. Domain Randomization

To improve robustness and sim-to-real transfer, we randomize both dynamics and perception at the start of each episode and throughout rollout. On the dynamics side, we sample randomize friction, restitution, as well as the robot’s initial pose, reference motion clip and phase, yaw, base height, and joint states; we additionally apply stochastic external pushes and occasional joint torque dropouts to mimic partial actuator failures. On the perception side, following prior works [3, 41], we perturb depth observations using depth clipping and non-linear remapping, multiplicative noise, spatial and temporal dropout, synthetic occlusions, and small random camera pose jitter. These perturbations encourage policies that rely on stable geometric structure rather than brittle simulator-specific visual or dynamical artifacts; full ranges and ablations are provided in the Appendix Section F.

## IV. EXPERIMENTS

### A. Implementation Details

We summarize the main components of our simulation, training, and deployment setup below; full implementation details, hyperparameters, and architectural specifications are provided in the Appendix Section C.

**Simulation Setup:** We use HumanoidVerse [17] and extend it with terrain generation and image rendering, enabling the full pipeline to run in both *IsaacGym* and *IsaacLab* using identical control and learning modules. Environments use a 23-DoF Unitree G1 model with a head-mounted depth camera, executed at 50 Hz for up to 7.5 s per episode.

**Terrain Setup.** Training proceeds in two phases. The policy is first trained on flat terrain to acquire core stand-up and

impact-mitigation behaviors, and is then continued on randomized non-flat terrains. Three representative difficulty bands are shown in Fig. 4; the full training distribution spans fifteen continuous difficulty levels.

**Training:** The teacher is trained using PPO [31] with 4,096 parallel environments. Due to rendering cost, the student policy is trained with egocentric depth observations on 512 environments. Training is performed on a single RTX 4090 workstation for IsaacLab experiments, and on an NVIDIA A40 GPU for IsaacGym experiments.

**Neural Architectures:** All network components rely on lightweight MLP backbones with ELU activations. Depth is processed by a compact CNN, and short-term proprioceptive history is encoded by a small temporal convolutional module.

**Real-World Setup:** Hardware experiments are performed on a Unitree G1 equipped with an Intel RealSense head mounted depth camera. Proprioception streams at 500 Hz and depth at 30 Hz. The student policy runs at 50 Hz and outputs joint-space position targets to the low-level PD controller. Depth preprocessing mirrors simulation, and deployment is fully zero-shot without any real-world fine-tuning.

### B. Simulated Experiments

We evaluate policies under two initialization regimes: *Stand-Up* and *Fall-Recovery* (dynamic falls with nonzero base velocity). Unless otherwise stated, results are reported on mixed terrains in IsaacGym; IsaacSim results are provided in the Appendix. For Stand-Up, episodes are initialized by sampling around the lowest-configuration keyframe of each demonstration with added noise. For Fall-Recovery, episodes are initialized near the onset of falling with random perturbations.

*a) Metrics:* We evaluate recovery using the following metrics (details in Appendix G), averaged over 300 trials (3 seeds). **Success (Succ)** measures upright stabilization within 7.5 s. **Safe success (Succ<sub>safe</sub>)** excludes trials where the head approaches within 5 cm of the terrain. **Time** is time-to-recovery. **Tracking error (Track.)** is the root-mean-squared deviation from the reference during non-stationary phases. **Energy** measures mechanical power consumption. **Displacement (Disp.)** quantifies cumulative pelvis drift over the episode.

*b) Baselines:* We compare VIGOR against representative prior methods for humanoid stand-up and recovery under identical initialization and terrain distributions. **HOST** [14] learns stand-up behaviors for different starting configurations separately directly through RL, using curriculum scheduling and multiple critics, while **FIRM** [39] leverages motion-level structure by conditioning a goal diffusion model to guide recovery behaviors. Both methods were originally designed and evaluated on flat terrain, and we assess their generalization under our terrain setup. We evaluate HOST only on stand-up tasks, as it was not trained for fall-recovery scenarios. For each terrain setup, we use the closest available HOST starting configuration when available; otherwise, we use the flat-terrain policy. Since no existing visual baseline for humanoid fall recovery is available, we instead study the role of perception through targeted student ablations.

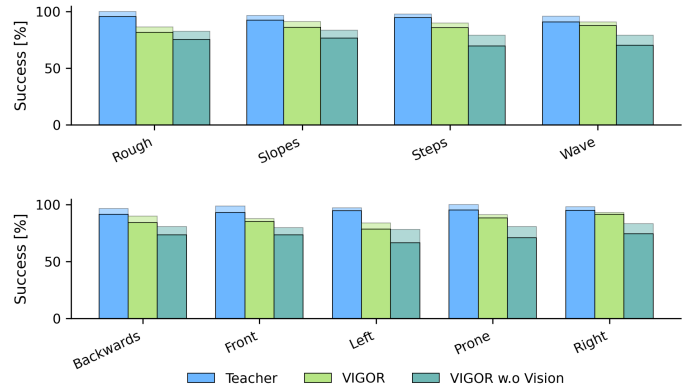


Fig. 5: **Simulation performance, grouped by terrain and motion type.** Top: success rate by terrain family. Bottom: success rate by initial fall direction, aggregated over terrains. The semi-transparent segment indicates unsafe successes.

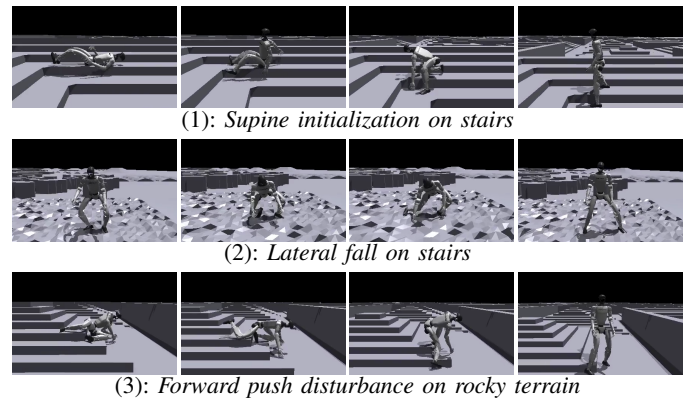


Fig. 6: **Recovery scenario examples.** Each row shows a different initial condition and terrain, visualized with four key frames from left to right.

*c) Teacher Ablations:* We evaluate teacher-side ablations to isolate the contribution of privileged structure and supervision. **noKeypoints** removes keypoint-based observations. **DofKeypoints** replaces spatial keypoint positions with joint angle observations as the target observation. **AbsTrack** trains the teacher using absolute pose tracking objectives instead of relative tracking. **NoScandots** removes access to privileged terrain information, while **Teacher** denotes the full privileged teacher with complete terrain access and supervision.

*d) Student Ablations:* We study student-side ablations of the VIGOR student policy by selectively removing components while keeping training conditions fixed. **w.o Shared** disables shared latent supervision, removing goal-level distillation. **w.o Vision** removes egocentric depth input, yielding a proprioception-only student. **w.o History** removes temporal observation history, restricting the student to single-step observations. Lastly, **VIGOR** includes egocentric depth, temporal history, and shared latent supervision.

### C. Simulated Results

Table II reports average recovery performance across all simulated terrains for both Stand-Up and Fall-Recovery. De-

TABLE II: **Average simulated recovery performance** under Stand-Up and Fall-Recovery initializations across all terrains. **best** and **second-best** indicate the top two methods per column. **VIGOR significantly outperforms baselines on both stand-up and fall recovery.**

Stand Up	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
HOST [14]	15.2 ( $\pm 3.7$ )	12.8 ( $\pm 4.1$ )	<b>2.7</b> ( $\pm 0.7$ )	–	480.3 ( $\pm 60.2$ )	3.2 ( $\pm 1.4$ )
FIRM [39]	30.8 ( $\pm 5.2$ )	21.4 ( $\pm 4.3$ )	3.1 ( $\pm 1.9$ )	20.4 ( $\pm 2.3$ )	490.1 ( $\pm 103.7$ )	2.8 ( $\pm 1.2$ )
VIGOR (Ours)	89.5 ( $\pm 3.0$ )	86.7 ( $\pm 3.1$ )	4.9 ( $\pm 2.1$ )	15.1 ( $\pm 5.4$ )	<b>305.6</b> ( $\pm 124.8$ )	1.7 ( $\pm 0.7$ )
Teacher	<b>97.7</b> ( $\pm 1.5$ )	<b>93.0</b> ( $\pm 2.5$ )	4.0 ( $\pm 1.3$ )	<b>10.6</b> ( $\pm 4.6$ )	315.6 ( $\pm 135.3$ )	<b>1.5</b> ( $\pm 0.5$ )
Fall Recovery	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
FIRM [39]	20.2 ( $\pm 6.3$ )	15.3 ( $\pm 3.2$ )	5.7 ( $\pm 2.2$ )	26.4 ( $\pm 4.3$ )	320.5 ( $\pm 99.3$ )	2.9 ( $\pm 1.2$ )
VIGOR (Ours)	90.5 ( $\pm 2.0$ )	89.3 ( $\pm 3.4$ )	5.4 ( $\pm 1.9$ )	14.1 ( $\pm 4.4$ )	287.5 ( $\pm 125.8$ )	<b>1.8</b> ( $\pm 1.0$ )
Teacher	<b>98.0</b> ( $\pm 1.4$ )	<b>94.6</b> ( $\pm 2.3$ )	<b>5.3</b> ( $\pm 1.2$ )	<b>10.5</b> ( $\pm 3.4$ )	265.8 ( $\pm 96.5$ )	<b>1.8</b> ( $\pm 0.6$ )

TABLE III: **Teacher ablations.** Ablations of the privileged teacher under both Stand-Up and Fall-Recovery initializations across all terrains.

Teacher	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
Stand Up	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
noKeyoints	55.3 ( $\pm 4.9$ )	52.0 ( $\pm 5.0$ )	7.1 ( $\pm 2.9$ )	15.9 ( $\pm 4.5$ )	<b>198.7</b> ( $\pm 123.2$ )	<b>1.3</b> ( $\pm 0.6$ )
DofKeyoints	86.0 ( $\pm 3.4$ )	78.0 ( $\pm 4.6$ )	7.3 ( $\pm 1.6$ )	15.1 ( $\pm 7.7$ )	216.5 ( $\pm 92.6$ )	2.4 ( $\pm 1.2$ )
AbsTrack	74.1 ( $\pm 4.3$ )	71.8 ( $\pm 4.5$ )	5.5 ( $\pm 2.7$ )	18.6 ( $\pm 1.1$ )	281.6 ( $\pm 145.1$ )	1.5 ( $\pm 0.5$ )
NoScandots	90.7 ( $\pm 2.9$ )	75.7 ( $\pm 4.2$ )	4.9 ( $\pm 1.9$ )	10.8 ( $\pm 3.8$ )	297.0 ( $\pm 103.1$ )	1.8 ( $\pm 0.8$ )
Teacher	<b>97.7</b> ( $\pm 1.5$ )	<b>93.0</b> ( $\pm 2.5$ )	<b>4.0</b> ( $\pm 1.3$ )	<b>10.6</b> ( $\pm 4.6$ )	315.6 ( $\pm 135.3$ )	1.5 ( $\pm 0.5$ )
Fall Recovery	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
noKeyoints	77.4 ( $\pm 4.2$ )	76.0 ( $\pm 4.2$ )	6.6 ( $\pm 2.1$ )	20.1 ( $\pm 4.4$ )	<b>149.8</b> ( $\pm 112.5$ )	<b>1.4</b> ( $\pm 0.9$ )
DofKeyoints	83.3 ( $\pm 3.7$ )	70.0 ( $\pm 4.5$ )	7.3 ( $\pm 1.2$ )	15.6 ( $\pm 7.2$ )	193.0 ( $\pm 72.6$ )	2.9 ( $\pm 1.0$ )
AbsTrack	67.0 ( $\pm 4.7$ )	64.3 ( $\pm 4.7$ )	5.5 ( $\pm 2.4$ )	22.0 ( $\pm 0.7$ )	246.7 ( $\pm 136.6$ )	2.2 ( $\pm 0.6$ )
NoScandots	88.0 ( $\pm 3.2$ )	75.0 ( $\pm 4.3$ )	6.7 ( $\pm 1.6$ )	12.1 ( $\pm 4.3$ )	280.3 ( $\pm 88.3$ )	2.7 ( $\pm 0.9$ )
Teacher	<b>98.0</b> ( $\pm 1.4$ )	<b>94.6</b> ( $\pm 2.3$ )	<b>5.3</b> ( $\pm 1.2$ )	<b>10.5</b> ( $\pm 3.4$ )	265.8 ( $\pm 96.5$ )	1.8 ( $\pm 0.6$ )

tailed breakdowns by terrain family and fall direction are shown in Fig. 5, while representative recovery rollouts under diverse initializations and contact sequences are visualized in Fig. 6.

**Baselines.** Prior methods exhibit limited robustness under the unified terrain setting. Both **HOST** and **FIRM**, which were originally designed for flat terrain, achieve substantially lower success and safe success rates compared to our approach, with higher tracking error, energy consumption, and base displacement. In contrast, **VIGOR** shows large gains in both Stand-Up and Fall-Recovery, indicating improved robustness across diverse terrain geometries. Notably, Fall-Recovery generally requires less energy than Stand-Up across methods, consistent with its shorter stabilization horizon, whereas Stand-Up involves longer multi-contact transitions and higher mechanical cost.

**Teacher Ablations.** Table III highlights the role of structural priors in the teacher. Removing keypoints (**noKeyoints**) causes the largest degradation in both Stand-Up and Fall-Recovery, substantially reducing success and safe success, indicating that explicit spatial structure is central to coordinated recovery. Replacing spatial keypoints with joint-angle targets (**DofKeyoints**) partially recovers performance but remains below the full model, suggesting that geometric body

TABLE IV: **Student ablations.** Ablations of the VIGOR under Stand-Up and Fall-Recovery initializations across all terrains. The realization gap,  $\|\tilde{\mathbf{z}}_t^{\text{goal}} - \mathbf{z}_t^{\text{goal}}\|^2$ , ( $\times 10^{-2}$ ) is 8.8 ( $\pm 4.6$ ) for **w.o Vision** and 6.2 ( $\pm 3.6$ ) for **w.o History**.

Student	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
Stand Up	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
w.o Shared	60.0 ( $\pm 4.9$ )	54.3 ( $\pm 4.9$ )	6.9 ( $\pm 2.8$ )	17.4 ( $\pm 4.9$ )	<b>250.3</b> ( $\pm 124.9$ )	2.1 ( $\pm 1.3$ )
w.o Vision	77.4 ( $\pm 4.1$ )	62.7 ( $\pm 4.8$ )	5.6 ( $\pm 2.6$ )	19.1 ( $\pm 5.4$ )	267.8 ( $\pm 138.8$ )	<b>2.0</b> ( $\pm 0.7$ )
w.o History	86.8 ( $\pm 3.3$ )	81.2 ( $\pm 3.9$ )	5.3 ( $\pm 2.2$ )	16.6 ( $\pm 4.9$ )	298.5 ( $\pm 115.1$ )	2.0 ( $\pm 1.0$ )
VIGOR	<b>89.5</b> ( $\pm 3.0$ )	<b>86.7</b> ( $\pm 3.1$ )	<b>4.9</b> ( $\pm 2.1$ )	<b>15.1</b> ( $\pm 5.4$ )	305.6 ( $\pm 124.8$ )	<b>1.7</b> ( $\pm 0.7$ )
Fall Recovery	Succ $\uparrow$	Succ $_{\text{safe}}\uparrow$	Time $\downarrow$	Track $\downarrow$	Energy $\downarrow$	Disp $\downarrow$
w.o Shared	57.7 ( $\pm 4.9$ )	57.7 ( $\pm 4.9$ )	7.2 ( $\pm 2.6$ )	19.0 ( $\pm 3.7$ )	<b>240.0</b> ( $\pm 121.9$ )	2.5 ( $\pm 1.5$ )
w.o Vision	84.3 ( $\pm 3.6$ )	82.3 ( $\pm 3.8$ )	5.6 ( $\pm 2.2$ )	18.6 ( $\pm 4.5$ )	273.5 ( $\pm 129.3$ )	2.5 ( $\pm 1.8$ )
w.o History	<b>91.0</b> ( $\pm 2.0$ )	<b>89.5</b> ( $\pm 1.0$ )	<b>5.4</b> ( $\pm 1.9$ )	<b>18.2</b> ( $\pm 4.2$ )	293.5 ( $\pm 125.3$ )	2.3 ( $\pm 1.0$ )
VIGOR	90.5 ( $\pm 2.0$ )	89.3 ( $\pm 3.4$ )	<b>5.4</b> ( $\pm 1.9$ )	14.1 ( $\pm 4.4$ )	287.5 ( $\pm 125.8$ )	<b>1.8</b> ( $\pm 1.0$ )

configuration provides stronger guidance than joint supervision alone. Using absolute pose tracking (**AbsTrack**) markedly increases tracking error and reduces success, showing that world-frame supervision degrades terrain transfer. In contrast, removing privileged terrain samples (**NoScandots**) preserves much of the raw success but leads to a pronounced drop in safe success, indicating that terrain observations primarily contribute to safety rather than to core motion coordination. The full privileged **Teacher** consistently achieves the best overall performance.

**Student Ablations.** Table IV reveals distinct roles for each student component. Removing shared latent supervision (**w.o Shared**) causes the largest degradation in both Stand-Up and Fall-Recovery, substantially reducing success and safe success. This confirms that distillation of the goal-in-context representation is the primary driver of coordinated recovery. Eliminating egocentric vision (**w.o Vision**) preserves moderate raw success but consistently lowers safe success, especially in Stand-Up, indicating that depth primarily improves terrain-aware stabilization rather than basic recovery capability. Removing temporal history (**w.o History**) slightly reduces Stand-Up performance but has no negative effect on Fall-Recovery, where it even achieves marginally higher success. This suggests that Stand-Up benefits from temporal context, while Fall-Recovery is dominated by rapid reactive control. Overall, the full **VIGOR** model achieves the most balanced behavior, combining high success, stronger safety, lower tracking error, and reduced displacement. Notably, improvements in success are generally accompanied by lower displacement and energy, suggesting more efficient stabilization rather than more aggressive control.

#### D. Real-World Experiments

We evaluate the student policy on a real Unitree G1 without any task-specific tuning or parameter changes. Experiments are conducted under two regimes: *Stand-Up* and *Fall-Recovery*, each spanning multiple terrains and initialization conditions.

a) *Fall-Recovery*: Fall-recovery experiments are conducted on three terrain types: *flat ground*, a *raised platform*, and *stairs*. On flat ground, external pushes are applied from

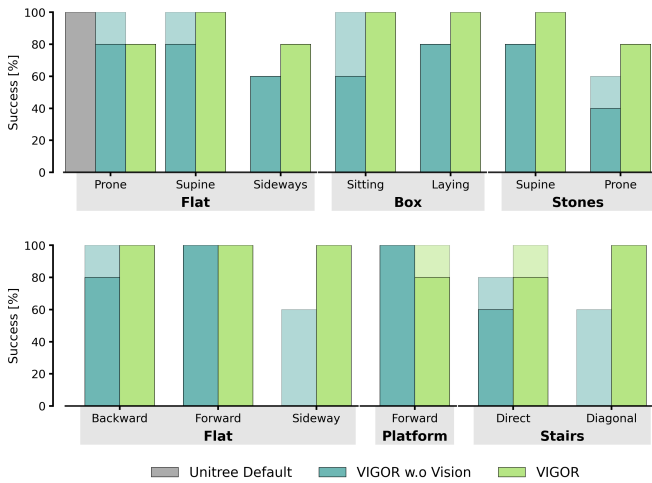


Fig. 7: **Real-world recovery performance across surfaces and initial configurations.** (Top) Stand-up. (Bottom) Fall recovery across push directions; bars show success over five trials per condition, with lighter segments indicating unsafe successes. **Vision improves terrain-aware reactions.**

three directions—backward, forward, and sideways—to induce diverse falling motions. On the platform terrain, the robot is pushed forward off the platform edge, resulting in a forward fall with a height change prior to impact. On stairs, two push configurations are evaluated: a direct push toward the stairs and a diagonal push across the stair direction, inducing asymmetric contact sequences during descent and recovery. Because external perturbations alone do not always produce uncontrolled falls on hardware, the learned policy is disabled during the initial disturbance and activated only once the robot reaches a predefined fall angle. This angle is detected using the projected gravity vector estimated from the onboard IMU, ensuring that control is engaged only after a genuine falling state with nonzero base velocity is observed.

*b) Stand-Up:* Stand-up experiments are conducted on three terrain types: *flat ground*, a *box obstacle*, and *stones*. On flat ground, the robot is initialized in prone, supine, and sideways configurations, with randomized joint configurations while in contact with the ground. On the box terrain, two initializations are used: a seated configuration with the robot’s back adjacent to the box, and a laying configuration with the torso on the box, both with randomized joint angles. On the stones terrain, the robot is initialized in prone and supine configurations, reflecting uneven contact geometry under torso and limbs. Across all stand-up experiments, the policy must generate a sequence of whole-body actions that brings the robot to a stable upright posture and maintains balance.

### E. Real-World Results

Fig. 7 reports real-world performance on a Unitree G1 under *Stand-Up* (top) and *Fall-Recovery* (bottom), grouped by terrain and initialization or disturbance type. Representative images for each evaluation setting are provided in Appendix.

**Stand-Up.** On flat terrain, both policies achieve similar overall success, with small differences across specific initial

poses. As terrain complexity increases (box and stones), differences become clearer: while raw success remains comparable in many cases, the blind policy exhibits lower safe success under uneven contacts, whereas VIGOR maintains consistently higher safety across initializations.

**Fall-Recovery.** For forward and backward pushes on flat terrain, both policies recover successfully. However, under more challenging disturbances, such as sideways pushes on flat ground and stair interactions, the blind policy shows marked drops in safe success. In contrast, VIGOR maintains high success and markedly better safety, with the advantage becoming even more pronounced in complex setups that require terrain-aware reactions, such as diagonal stair pushes.

**Effect of Vision.** Vision does not universally improve raw success in simple scenarios, and in some flat stand-up cases the blind policy performs similarly. The primary benefit of egocentric depth appears in enabling terrain-aware reactions and safer contact-aware stabilization: under asymmetric contacts and stair geometries, vision significantly reduces unsafe recoveries while preserving high overall recovery rates.

## V. LIMITATIONS

Our current formulation emphasizes robust reactive fall recovery after loss of balance, providing a strong foundation for safe post-disturbance behavior. While the policy already exhibits stabilizing responses at fall onset, it is not yet jointly trained with locomotion or long-horizon navigation. Extending the framework to unify fall avoidance and recovery within continuous locomotion is a promising direction for future work. More broadly, our approach leverages sparse keyframe demonstrations as high-level structural priors for training the privileged teacher, enabling flexible recovery without over-constraining motion. However, the use of sparse keyframe supervision may limit motion fidelity in fine-grained contact transitions. On the student side, recovery is encoded through a compact goal-in-context latent that efficiently captures target pose and local terrain, offering a scalable representation that could be further enriched for highly multimodal scenarios.

## VI. CONCLUSION

We present VIGOR, a unified framework for humanoid fall safety that integrates fall mitigation and recovery within a single vision-conditioned policy. By factorizing data complexity into sparse human motion priors and independently varying terrain, and by representing action goals through a compact goal-in-context latent, VIGOR enables RL to resolve contact timing and terrain-dependent execution without dense demonstrations. A privileged teacher learns terrain-aware recovery strategies using sparse keyframes and terrain access, which are distilled into a deployable student operating solely from egocentric perception. Simulation and real-world results demonstrate robust recovery across diverse terrains and initial conditions, validating structured teacher-student learning as a practical path to zero-shot sim-to-real humanoid fall safety.

## REFERENCES

- [1] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *6th Annual Conference on Robot Learning*, 2022. URL <https://openreview.net/forum?id=Re3NjSwf0WF>.
- [2] Arthur Allshire, Hongsuk Choi, Junyi Zhang, David McAllister, Anthony Zhang, Chung Min Kim, Trevor Darrell, Pieter Abbeel, Jitendra Malik, and Angjoo Kanazawa. Visual imitation enables contextual humanoid control. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2025.
- [3] Osher Azulay, Dhruv Metha Ramesh, Nimrod Curtis, and Avishai Sintov. Visuotactile-based learning for insertion with compliant hands. *IEEE Robotics and Automation Letters*, 10(4):4053–4060, 2025.
- [4] Lingfan Bao, Joseph Humphreys, Tianhu Peng, and Chengxu Zhou. Deep reinforcement learning for bipedal locomotion: A brief survey. *arXiv preprint arXiv:2404.17070*, 2024. URL <https://arxiv.org/abs/2404.17070>. Last revised January 7, 2026.
- [5] Sirui Chen, Yufei Ye, Zi-ang Cao, Jennifer Lew, Pei Xu, and C Karen Liu. Hand-eye autonomous delivery: Learning humanoid navigation, locomotion and reaching. *arXiv preprint arXiv:2508.03068*, 2025.
- [6] Helei Duan, Bikram Pandit, Mohitvishnu S. Gadde, Bart J. van Marum, Jeremy Dao, Chanho Kim, and Alan Fern. Learning vision-based bipedal locomotion for challenging terrain. *arXiv preprint arXiv:2309.14594*, 2023. URL <https://arxiv.org/abs/2309.14594>.
- [7] Johannes Engelsberger, Twan Koolen, Sylvain Bertrand, Jerry Pratt, Christian Ott, and Alin Albu-Schäffer. Trajectory generation for continuous leg forces during double support and heel-to-toe shift based on divergent component of motion. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4022–4029. IEEE, 2014.
- [8] Diego Ferigo, Raffaello Camoriano, Paolo Maria Viceconte, Daniele Calandriello, Silvio Traversaro, Lorenzo Rosasco, and Daniele Pucci. On the emergence of whole-body strategies from humanoid robot push-recovery learning. *IEEE Robotics and Automation Letters*, 6(4): 8561–8568, 2021.
- [9] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: Learning a unified policy for manipulation and locomotion. In *6th Annual Conference on Robot Learning*, 2022. URL <https://openreview.net/forum?id=zldI4UpuG7v>.
- [10] Kiyoshi Fujiwara, Fumio Kanehiro, Shuuji Kajita, Kenji Kaneko, Kazuhito Yokoi, and Hirohisa Hirukawa. Ukemi: Falling motion control to minimize damage to biped humanoid robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 2521–2526. IEEE, 2002.
- [11] Clément Gaspard, Marc Duclausaud, Grégoire Passault, Mélodie Daniel, and Olivier Ly. Frasa: An end-to-end reinforcement learning agent for fall recovery and stand up of humanoid robots. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 15994–16000. IEEE, 2025.
- [12] Zhaoyuan Gu, Junheng Li, Wenlan Shen, Wenhao Yu, Zhaoming Xie, Stephen McCrory, Xianyi Cheng, Abdulaziz Shamsah, Robert Griffin, C. Karen Liu, Abderrahmane Kheddar, Xue Bin Peng, Yuke Zhu, Guanya Shi, Quan Nguyen, Gordon Cheng, Huijun Gao, and Ye Zhao. Humanoid locomotion and manipulation: Current progress and challenges in control, planning, and learning. *arXiv preprint arXiv:2501.02116*, 2025. doi: 10.48550/arXiv.2501.02116.
- [13] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Learning Getting-Up Policies for Real-World Humanoid Robots. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.063.
- [14] Tao Huang, Junli Ren, Huayi Wang, Zirui Wang, Qingwei Ben, Muning Wen, Xiao Chen, Jianan Li, and Jiangmiao Pang. Learning Humanoid Standing-up Control across Diverse Postures. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.064.
- [15] Eric Krotkov, Douglas Hackett, Larry Jackel, Michael Perschbacher, James Pippine, Jesse Strauss, Gill Pratt, and Christopher Orlowski. The darpa robotics challenge finals: Results and perspectives. In *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*, volume 121 of *Springer Tracts in Advanced Robotics*, pages 1–26. Springer, Cham, 2018. doi: 10.1007/978-3-319-74666-1\_1.
- [16] Visak CV Kumar, Sehoon Ha, and C Karen Liu. Learning a unified control policy for safe falling. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3940–3947. IEEE, 2017.
- [17] CMU LeCAR Lab. Humanoidverse: A multi-simulator framework for humanoid robot sim-to-real learning. <https://github.com/LeCAR-Lab/HumanoidVerse>, 2025.
- [18] Qiayuan Liao, Takara E Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C Karen Liu. Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion. *arXiv preprint arXiv:2508.08241*, 2025.
- [19] Kwan-Yee Lin and Stella X Yu. Let humanoids hike! integrative skill development on complex trails. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22498–22507, 2025.
- [20] Minghuan Liu, Zixuan Chen, Xuxin Cheng, Yandong Ji, Ri-Zhao Qiu, Ruihan Yang, and Xiaolong Wang. Visual whole-body control for legged loco-manipulation. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=cT2N3p1AcE>.
- [21] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned

- multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, October 2015.
- [22] Antonio Loquercio, Ashish Kumar, and Jitendra Malik. Learning visual locomotion with cross-modal supervision. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 7295–7302. IEEE, 2023.
- [23] Dingsheng Luo, Yaoxiang Ding, Zidong Cao, and Xihong Wu. A multi-stage approach for efficiently learning humanoid robot stand-up behavior. In *2014 IEEE international conference on mechatronics and automation*, pages 884–889. IEEE, 2014.
- [24] Zhengyi Luo, Chen Tessler, Toru Lin, Ye Yuan, Tairan He, Wenli Xiao, Yunrong Guo, Gal Chechik, Kris Kitani, Linxi Fan, et al. Emergent active perception and dexterity of simulated humanoids from visual reinforcement learning. *arXiv preprint arXiv:2505.12278*, 2025.
- [25] James Ni, Zekai Wang, Wei Lin, Amir Bar, Yann LeCun, Trevor Darrell, Jitendra Malik, and Roei Herzig. From generated human videos to physically plausible robot trajectories. *arXiv preprint arXiv:2512.05094*, 2025.
- [26] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.
- [27] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (ToG)*, 40(4):1–20, 2021.
- [28] Jerry E. Pratt, John Carff, Sergey V. Drakunov, and Ambarish Goswami. Capture point: A step toward humanoid push recovery. In *Proceedings of the 6th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2006)*, pages 140–145, Genoa, Italy, 2006. doi: 10.1109/ICHR.2006.321385.
- [29] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Malik Jitendra, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 2024? doi: 10.1126/scirobotics.adi9579.
- [30] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [32] Remo Steiner, Alexander Millane, David Tingdahl, Clemens Volk, Vikram Ramasamy, Xinjie Yao, Peter Du, Soha Pouya, and Shiwei Sheng. mindmap: Spatial memory in deep feature maps for 3d action policies, 2025. URL <https://arxiv.org/abs/2509.20297>.
- [33] Jörg Stückler, Johannes Schwenk, and Sven Behnke. Getting back on two feet: Reliable standing-up routines for a humanoid robot. In *IAS*, pages 676–685, 2006.
- [34] William Suliman, Egor Davydenko, Ekaterina Terina Chaikovskaia, and Roman Gorbachev. Reinforcement learning-based footstep control for humanoid robots on complex terrain. *IEEE Access*, PP:1–1, 2025. doi: 10.1109/ACCESS.2025.3622091. URL <https://ieeexplore.ieee.org/document/10364844>.
- [35] Tianxin Tao, Matthew Wilson, Ruiyu Gou, and Michiel Van De Panne. Learning to get up. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–10, 2022.
- [36] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 2024.
- [37] Huayi Wang, Zirui Wang, Junli Ren, Qingwei Ben, Tao Huang, Weinan Zhang, and Jiangmiao Pang. Beamdojo: Learning agile humanoid locomotion on sparse footholds. *arXiv preprint arXiv:2502.10363*, 2025.
- [38] Kehan Wen, Chenhao Li, Junzhe He, and Marco Hutter. Constrained style learning from imperfect demonstrations under task optimality. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=TFbT7kHD89>.
- [39] Zhengjie Xu, Ye Li, Kwan-yeo Lin, and Stella X Yu. Unified humanoid fall-safety policy from a few demonstrations. *arXiv preprint arXiv:2511.07407*, 2025.
- [40] Chuanyu Yang, Can Pu, Guiyang Xin, Jie Zhang, and Zhibin Li. Learning complex motor skills for legged robot fall recovery. *IEEE Robotics and Automation Letters*, 8(7):4307–4314, 2023.
- [41] Shaofeng Yin, Yanjie Ze, Hong-Xing Yu, C. Karen Liu, and Jiajun Wu. Visualmimic: Visual humanoid locomanipulation via motion tracking and generation. *arXiv preprint arXiv:2509.20322*, 2025.
- [42] Fatemeh Zargarbashi, Jin Cheng, Dongho Kang, Robert Sumner, and Stelian Coros. Robotkeyframing: Learning locomotion with high-level objectives via mixture of dense and sparse rewards. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=wcbRhPnOei>.
- [43] Zewei Zhang, Chenhao Li, Takahiro Miki, and Marco Hutter. Motion priors reimaged: Adapting flat-terrain skills for complex quadruped mobility. In *9th Annual Conference on Robot Learning*, 2025. URL <https://openreview.net/forum?id=JXBm4Xfrvj>.
- [44] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=fs7ia3FqUM>.

## APPENDIX

This Appendix provides implementation details and extended results that complement the main paper and support reproducibility; for further intuition and visual examples, we refer the reader to the supplementary video.

### A. Contribution Overview

This work makes the following contributions:

- **Unified, visually grounded fall safety.** To the best of our knowledge, we present the first learning-based humanoid fall-safety framework that unifies fall mitigation and stand-up recovery within a single policy with *visual awareness*, enabling safer behavior under complex and uncertain environments.
- **Factorized data formulation for fall recovery.** We propose a factorized view of *fall-recovery* data complexity that decouples human pose structure from terrain variation, enabling sample-efficient learning by combining a small number of flat-ground human demonstrations with large-scale terrain randomization in simulation.
- **Visual goal-in-context distillation.** We introduce a compact goal-in-context latent that jointly encodes the next target pose (guidance [26, 25, 39]), local terrain geometry (environment awareness [1, 22]), and body state, distilled from a privileged terrain-aware teacher and deployed in a student policy using only egocentric depth and short-term proprioceptive history.
- **Comprehensive evaluation and real-world transfer.** We evaluate the approach across diverse simulated fall recovery scenarios and demonstrate zero-shot transfer to a real humanoid robot.

### B. Future Work: Proactive Fall Avoidance

Our method, VIGOR, is particularly well suited to the relatively short-horizon tasks of humanoid fall-recovery and standing-up. Notably, the framework is completely decoupled from the locomotion policy that initially caused the fall. As a result, VIGOR emphasizes reactive fall response and recovery rather than proactive fall avoidance during locomotion. Investigating the interplay between robust humanoid locomotion and fall avoidance is an exciting direction for future work.

### C. Additional Implementation Details

Training hyperparameters for both the teacher (PPO) and student (DAgger) policies are summarized in Table V. Policies operate at 50 Hz, while physics simulation runs at 200 Hz using four substeps per control step. Egocentric depth images are rendered at 30 Hz to match the real-world sensing rate. Policy actions correspond to joint-space offsets, which are scaled and clipped before being applied as position targets to a low-level PD controller. Specifically, the policy output  $a_t$  is mapped to desired joint positions as

$$q_t^{\text{des}} = q^{\text{default}} + c \cdot \text{clip}(a_t, -a_{\text{clip}}, a_{\text{clip}}),$$

where  $q^{\text{default}}$  denotes the default joint configuration,  $c$  is a fixed action scaling coefficient, and  $a_{\text{clip}}$  bounds the action

TABLE V: Hyperparameters used for PPO and DAgger.

PPO (Teacher)	
Number of GPUs (Sim)	1 × RTX 4090
Number of GPUs (Gym)	NVIDIA A40
Number of environments	4096
Learning epochs	5
Steps per environment	24
Minibatch size	19576
Discount factor ( $\gamma$ )	0.99
GAE parameter ( $\lambda$ )	0.95
Clipping parameter ( $\epsilon$ )	0.2
Entropy coefficient	0.005
Optimizer	Adam
Learning rate	$1 \times 10^{-3}$
Learning rate schedule	Adaptive
Target KL divergence	0.01
DAgger (Student)	
Number of GPUs	1 × RTX 4090
Number of environments	512
Optimizer	Adam
Steps per environment	4
Learning rate	$1 \times 10^{-4}$

magnitude. These desired joint positions are tracked using a PD controller,

$$\tau_t = K_p (q_t^{\text{des}} - q_t) - K_d \dot{q}_t,$$

which outputs joint torques applied to the simulator.

Episodes are not terminated early upon failure; instead, the policy is allowed to continue executing until the episode horizon. This design encourages learning recovery behaviors from a diverse set of failure states, including prolonged contact, partial collapses, and unstable intermediate configurations.

### D. Demos, Retargeting Constraints, and Reference Processing

Our model is trained using only nine flat-ground indoor fall-recovery demonstrations (Fig. 12). Despite this coarse pose guidance, VIGOR learns robust fall-recovery behaviors that generalize to complex terrains (such as stairs and waves) in simulation, enabling strong vision-based humanoid fall safety in both challenging simulated and real-world settings.

These human demonstrations of fall–recovery motions involve high-impact and multi-contact transitions. Direct retargeting can introduce kinematic artifacts such as excessive pelvis or hip rotations and inconsistent contact timing. To mitigate these effects, we apply conservative joint-limit constraints during retargeting, primarily on the pelvis and hip joints, to maintain physically plausible configurations while preserving high-level recovery intent. From each retargeted sequence, we extract sparse keyframes by uniform temporal subsampling (typically  $\sim 5$  Hz), which are used as structural priors rather than strict trajectory targets. During training, environments randomly sample both the reference sequence and temporal phase to initialize diverse recovery configurations, including early falling stages, and apply simple global transformations. When operating on uneven terrain, reference poses are vertically shifted using the coarse projection described in the main text to avoid penetration and maintain clearance above the

local ground surface. To expose the policy to uncontrolled fall dynamics, some episodes begin with a brief *free-fall* interval in which a random subset of joint torques is suppressed; once control is enabled, we optionally re-synchronize the reference phase by matching the current base height to a small set of candidate keyframes.

### E. Observation and Architecture Details

Table VI summarizes the observations used by the teacher actor, student actor, and critic.

TABLE VI: Observations used by the teacher policy, student policy, and critic.

Input	Dim.	Teacher	Student	Critic
Base lin. vel.	3	×	×	✓
Base ang. vel.	3	✓	✓	✓
Projected gravity	3	✓	✓	✓
Joint positions	23	✓	✓	✓
Joint velocities	23	✓	✓	✓
Previous action	23	✓	✓	✓
Terrain heights	132	✓	×	✓
Obs. history	10 × 75	×	✓	×
<b>Motion</b>				
Local ref. body	3 × 31	✓	×	✓
Diff. local body	3 × 31	×	×	✓
<b>Perception</b>				
Depth img	64 × 64	×	✓	×
Depth img history	3 × 64 × 64	×	✓	×

### F. Domain Randomization and Noise

To improve robustness and sim-to-real transfer, we apply extensive domain randomization to dynamics and perception during training (Table VII).

TABLE VII: Domain randomization and feature noise used during training.

Parameter	Range / Value
<b>Dynamics Perturbations</b>	
Random push interval	$\mathcal{U}[15, 20]$ s
Random push (xy)	$\mathcal{U}[-1, 1]$ m/s
Ground friction	$\mathcal{U}[0.5, 1.25]$
Base CoM offset	$\mathcal{U}[-0.05, 0.05]$ m
Link mass	$\mathcal{U}[0.9, 1.1]$
PD gains ( $K_p, K_d$ )	$\mathcal{U}[0.9, 1.1]$
Control delay	{0, 1, 2} steps
<b>Policy Input Noise</b>	
Base linear velocity	$\mathcal{N}(0, 0.1)$ m/s
Base angular velocity	$\mathcal{N}(0, 0.2)$ rad/s
Projected gravity	$\mathcal{N}(0, 0.05)$ m/s <sup>2</sup>
Joint position	$\mathcal{N}(0, 0.01)$ rad
Joint velocity	$\mathcal{N}(0, 1.5)$ rad/s
Terrain heights	$\mathcal{N}(0, 0.1)$ m
<b>Visual Perturbations</b>	
Depth erasing	$p = 0.4$ , area [0.02, 0.10], ratio [0.3, 3.3]
Gaussian blur	kernel 5, $\sigma \in [0.01, 0.1]$
Gaussian noise (depth)	$\mathcal{N}(0, 0.005)$
Image rotation	$\pm 5^\circ$
Camera position jitter	$\pm 0.01$ m
Camera pitch jitter	$\pm 0.087$ rad ( $\approx \pm 5^\circ$ )
Depth clipping	[0.1, 2.0] m

### G. Evaluation Metrics

a) *Success (Succ.)*: An episode is successful if the robot stably reaches a reference standing configuration, with all tracked links and the relative head height within a fixed tolerance for a sustained duration.

b) *Safe success (Succ<sub>safe</sub>)*: Successful episodes with no unsafe head proximity events, where  $h_t^{\text{head}}$  and  $h_t^{\text{ground}}$  denote the head and local ground heights at time  $t$ :

$$\frac{1}{T} \sum_t \mathbf{1}\{h_t^{\text{head}} - h_t^{\text{ground}} < 5\} = 0.$$

c) *Time (Time)*: Episode duration, where  $T$  is the number of timesteps and  $\Delta t$  is the control timestep:

$$\text{Time} = T \Delta t.$$

d) *Tracking error (Track.)*: Root-mean-square error between root-relative link positions  $\mathbf{p}_{j,t}$  and reference positions  $\mathbf{p}_{j,t}^{\text{ref}}$  where  $j = 1, \dots, M$  indexes links:

$$\text{Track.} = \sqrt{\frac{1}{TM} \sum_{t=1}^T \sum_{j=1}^M \|\mathbf{p}_{j,t} - \mathbf{p}_{j,t}^{\text{ref}}\|_2^2}.$$

e) *Energy (Energy)*: Mean absolute mechanical power, where  $\tau_t$  and  $\dot{q}_t$  are joint torques and velocities:

$$\text{Energy} = \frac{1}{T \Delta t} \sum_t |\tau_t^\top \dot{q}_t|.$$

f) *Base displacement (Disp.)*: Horizontal base displacement, where  $\mathbf{p}_{t,xy}^{\text{base}}$  denotes the base position in the horizontal plane:

$$\text{Disp.} = \|\mathbf{p}_{T,xy}^{\text{base}} - \mathbf{p}_{0,xy}^{\text{base}}\|_2.$$

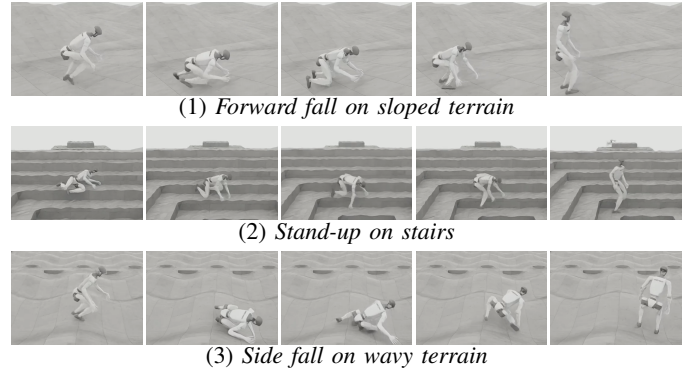


Fig. 8: **Recovery scenario examples.** Each row shows a different terrain and initial condition, visualized over key frames from left to right.

### H. Additional Simulated Results

This section reports additional quantitative results for our method evaluated in IsaacSim, covering both stand-up and fall initializations. These results are included for completeness and provide a controlled comparison across policy variants. In real-world deployment, we found that policies trained using

TABLE VIII: Overall recovery performance in IsaacSim across all terrains under **Stand-Up** and **Fall** initializations. Realizability Gap is the mean squared error between the student-predicted goal latent and the teacher goal latent.

Stand-Up							
Method	Real. Gap		Succ <sub>safe</sub> ↑	Time ↓	Track. ↓	Energy ↓	Disp. ↓
	( $\times 10^{-2}$ ) ↓	Succ. ↑					
<b>w.o Vision</b>	1.3 ( $\pm 0.5$ )	75.0 ( $\pm 4.3$ )	52.7 ( $\pm 5.0$ )	5.4 ( $\pm 2.5$ )	16.8 ( $\pm 5.0$ )	311.5 ( $\pm 150.6$ )	2.2 ( $\pm 1.0$ )
<b>VIGOR</b>	0.9 ( $\pm 0.2$ )	91.5 ( $\pm 1.5$ )	84.1 ( $\pm 3.9$ )	5.7 ( $\pm 1.0$ )	13.6 ( $\pm 4.6$ )	220.8 ( $\pm 95.7$ )	2.2 ( $\pm 1.0$ )
<b>Teacher</b>	–	97.6 ( $\pm 1.5$ )	93.5 ( $\pm 3.9$ )	4.0 ( $\pm 1.0$ )	8.2 ( $\pm 3.0$ )	276.8 ( $\pm 113.0$ )	1.3 ( $\pm 0.4$ )
Fall Recovery							
Method	Real. Gap		Succ <sub>safe</sub> ↑	Time ↓	Track. ↓	Energy ↓	Disp. ↓
	( $\times 10^{-2}$ ) ↓	Succ. ↑					
<b>w.o Vision</b>	1.4 ( $\pm 0.9$ )	77.3 ( $\pm 4.1$ )	75.7 ( $\pm 4.2$ )	6.2 ( $\pm 2.4$ )	17.3 ( $\pm 3.9$ )	301.0 ( $\pm 137.5$ )	2.6 ( $\pm 1.4$ )
<b>VIGOR</b>	0.4 ( $\pm 0.8$ )	91.0 ( $\pm 1.9$ )	85.7 ( $\pm 2.2$ )	6.5 ( $\pm 1.1$ )	14.7 ( $\pm 3.6$ )	209.2 ( $\pm 86.2$ )	2.3 ( $\pm 0.1$ )
<b>Teacher</b>	–	98.8 ( $\pm 0.1$ )	90.7 ( $\pm 0.4$ )	5.3 ( $\pm 1.2$ )	9.4 ( $\pm 3.1$ )	241.8 ( $\pm 94.17$ )	1.8 ( $\pm 0.4$ )

TABLE IX: Real-world stand-up performance across surfaces and initial configurations.

Model	Metric	Flat		Box		Stones		Total	
		Face Up	Face Down	Side	Sitting	Face Down	Face up	All Flat	All Non-Flat
<b>Unitree Default</b>	Succ $\uparrow$	5/5	–	–	–	–	–	–	–
	Succ <sub>safe</sub> $\uparrow$	5/5	–	–	–	–	–	–	–
<b>w.o Vision</b>	Succ $\uparrow$	5/5	5/5	3/5	5/5	4/5	4/5	3/5	13/15
	Succ <sub>safe</sub> $\uparrow$	4/5	4/5	3/5	3/5	4/5	4/5	2/5	11/15
<b>VIGOR</b>	Succ $\uparrow$	4/5	5/5	4/5	5/5	5/5	5/5	4/5	13/15
	Succ <sub>safe</sub> $\uparrow$	4/5	5/5	4/5	5/5	5/5	5/5	4/5	13/15

TABLE X: Real-world fall-recovery performance across surfaces and push categories.

Model	Metric	Flat			Off Platform	On Stairs		Total
		Backward	Forward	Sideway	Forward	Direct	Diagonal	
<b>w.o Vision</b>	Succ $\uparrow$	5/5	5/5	3/3	5/5	4/5	3/3	25/26
	Succ <sub>safe</sub> $\uparrow$	4/5	5/5	0/3	5/5	3/5	0/3	13/26
<b>VIGOR</b>	Succ $\uparrow$	5/5	5/5	5/5	5/5	5/5	5/5	30/30
	Succ <sub>safe</sub> $\uparrow$	5/5	5/5	5/5	4/5	4/5	5/5	28/30

Note: For *w.o Vision*, only three diagonal-stairs and sideway push trials were conducted to avoid further hardware damage.

this setup transferred more reliably and exhibited improved stability compared to the IsaacGym version. Table VIII summarizes the overall performance under both *Stand-Up* and *Fall* initializations.



Fig. 9: **Sim-to-real camera mapping**. Left: simulation rendering. Right: real image captured by the G1 head-mounted camera.

### I. Real-World Sensing and Control

On the real robot, depth images are acquired from an Intel RealSense sensor at a native resolution of  $640 \times 480$  and a frame rate of 30 Hz. Prior to being provided to the policy, each depth frame is center-cropped and resized to  $64 \times 64$ . This preprocessing was adopted after several experiments in which physical contact caused partial degradation near the image boundaries, while the central region remained usable. Proprioceptive measurements are streamed directly from the motor system via Unitree Python SDK. The control policy runs at 50 Hz and outputs joint position targets for all actuated joints. These targets are sent to the robot at the same rate. Low-level torque control is handled internally by the motors using built-in controllers running at 500 Hz.

### J. Real-World Results

This section expands on Fig. 7 in the main text and the results reported in Sec. E. Tables X and IX summarize the Fig. 7 results for clearer presentation. Each policy was evaluated over five runs. For **w.o Vision**, only three diagonal-stair and sideways-push trials were conducted to avoid further hardware

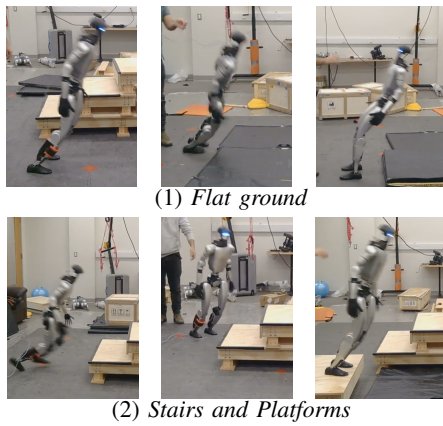


Fig. 10: **Tested fall-recovery scenarios:** (1) From left to right: forward, sideways, and backward pushes on flat ground. (2) From left to right: direct stair push, diagonal stair push, and forward push off a platform.

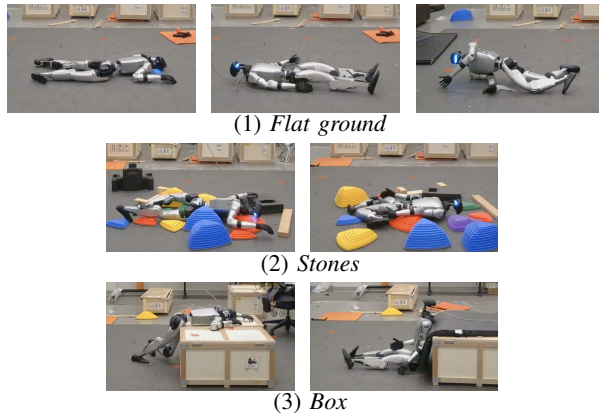


Fig. 11: **Tested stand-up scenarios:** Top: from left to right, prone, supine, and sideways stand-up on flat ground. Middle: prone and supine stand-up on stones. Bottom: seated next to a box and lying on a box.

damage. Figures 10 and 11 illustrate the tested scenarios. We define a successful trial as one in which the robot reaches and maintains a stable standing configuration for at least 7.5 s. Safe success further requires that no head collision with the environment occurs during the trial. As shown in Table IX, the **Unitree default** controller succeeds only for face-up stand-up on flat terrain and does not generalize to other initial configurations or non-flat surfaces. It is therefore omitted from the fall-recovery evaluation in Table X, highlighting the limited robustness of hand-engineered stand-up routines.

Demo 1: fall forward



Demo 2: fall forward



Demo 3: fall backward



Demo 4: fall backwards left



Demo 5: fall backwards right



Demo 6: fall right



Demo 7: fall right



Demo 8: fall left



Demo 9: stand up from prone



Fig. 12: **Human Fall Recovery Demonstrations.** Each row shows one of the nine human demonstration used to shape the learned recovery behaviors of VIGOR. Each demonstration is visualized with six key frames read from left to right. Our model is trained solely on these nine flat-ground indoor fall-recovery demonstrations. Despite this coarse pose guidance, it learns robust fall-recovery behaviors that generalize to complex terrains (such as stairs and waves) in simulation, enabling strong vision-based humanoid fall safety in both challenging simulated and real-world settings.