Journal of Electronic Research and Application, 2026, Volume 10, Issue 1
http://ojs.bbwpublisher.com/index.php/JERA
ISSN Online: 2208-3510
ISSN Print: 2208-3502

RoboMirror: Bridging Human Intent and Robot Capability Through Self-Reflective Motion Adaptation

Lin Zhu[1], Longliang Huang[1], Chuxiong Lin[2]*, Yujie Chen[1]

[1]Guizhou Equipment Manufacturing Polytechnic, Guizhou, China
[2]Beijing University of Technology, Beijing, China

**Corresponding author: Chuxiong Lin, syt20040909@126.com*

Abstract: Current humanoid robot control paradigms place the burden of feasibility assessment on human operators, who must carefully design commands within perceived robot limitations. This constraint significantly hinders practical deployment and limits the expressiveness of robot behaviors. This study proposed an inverting paradigm: rather than constraining operator inputs, robots should autonomously evaluate their capacity to execute commanded motions and intelligently adapt references to align with their physical constraints and learned skills. This study introduced the Performance Prediction Network (PPN), a transformer-based architecture that forecasts execution quality for arbitrary reference trajectories by analyzing both the commanded motion sequence and current robot state. Given a high-level task specification, our framework synthesizes multiple viable motion candidates and employs PPN to rank them across six dimensions: collision avoidance, kinematic feasibility, dynamic stability, trajectory smoothness, and goal satisfaction. This ranking enables autonomous selection of the most suitable reference motion before execution begins. Our complete system integrates motion generation, kinematic retargeting, and learned control policies with PPN-guided adaptation, creating a closed-loop framework where robots reason about their own limitations. Validated on 100,000 diverse human motions span walking, running, jumping, and acrobatic maneuvers, PPN achieves 99.14% accuracy in predicting imminent failures while maintaining low prediction error across all performance metrics. In deployment, our system successfully prevents 62% of anticipated falls by autonomously modifying commanded references, demonstrating that explicit capability modeling enables safer and more reliable humanoid control without sacrificing behavioral diversity.

Keywords: Deep reinforcement learning; Physical self-awareness; Safe motion planning; Failure prediction; Human-robot imitation

Online publication: February 12, 2026

1. Introduction

Humanoid robots are increasingly deployed in human-centered environments where commands are expressed through high-level social intent [1,2]. However, bridging the gap between intent-rich commands and a robot's physical capabilities remain a significant challenge [3–5]. Traditional control paradigms often impose a cognitive burden on operators, requiring them to internalize kinematic and stability limits to craft executable commands [6,7]. When reference motions slightly exceed a robot's feasible envelope, contemporary whole-body imitation systems often exhibit brittle behavior or failure, primarily due to the absence of self-evaluative mechanisms that reason about execution before action [8,9]. This study proposes a paradigm inversion: robots should proactively evaluate and adapt commanded behaviors to their own capabilities. This capability-aware view addresses three interconnected challenges: learning predictive models that generalize across diverse maneuvers, accounting for dynamic executability under real-world physic, and maintaining real-time latency for behavior selection [10–20]. Unlike approaches that restrict behavioral expressiveness to ensure safety, this study advocates for retaining open-ended commands while shifting the responsibility for feasibility to the robot via learned, predictive self-assessment [21–24]. Our framework instantiates this paradigm by forecasting execution quality to select optimal behavior candidates. At its core is the Performance Prediction Network (PPN), a transformer-based architecture that jointly encodes commanded reference trajectories and the robot's current state to predict multi-dimensional quality metrics, such as fall likelihood and tracking accuracy. This allows the system to rank multiple candidates, generated from text-to-motion models, and select the most viable trajectory before committing to control. The primary contributions of this work include:

(1) A capability-aware control paradigm
Inverts the feasibility burden, allowing robots to evaluate and adapt behaviors before execution.

(2) The performance prediction network (PPN)
A transformer model for forecasting multi- dimensional execution quality based on reference trajectories and robot state.

(3) An open-ended intent pipeline
Integrates high-level text-to-motion generation with learned, pre-execution ranking and selection.

(4) Real-time integration
A physics-based whole-body controller, enabling a low-latency assessment selection loop for responsive adaptation.

2. Related works

2.1. Imitation learning

Imitation learning (IL) enables robots to acquire skills from demonstrations without manual reward engineering [25–29]. While Behavior Cloning (BC) offers computational efficiency, it suffers from co- shift and compounding errors [27,28,30–32]. Inverse Reinforcement Learning (IRL) provides robustness by inferring reward functions but at a higher computational cost [27, 33–36]. Recent address failure by constraining policies to expert manifolds. Unlike these approaches that enforce strict adherence to demonstrations, our work builds on BC but relaxes imitation constraints; this study posit that agents should execute tasks optimally within their specific embodiment limits rather than perfectly replicating human motion. See **Figure 1**.

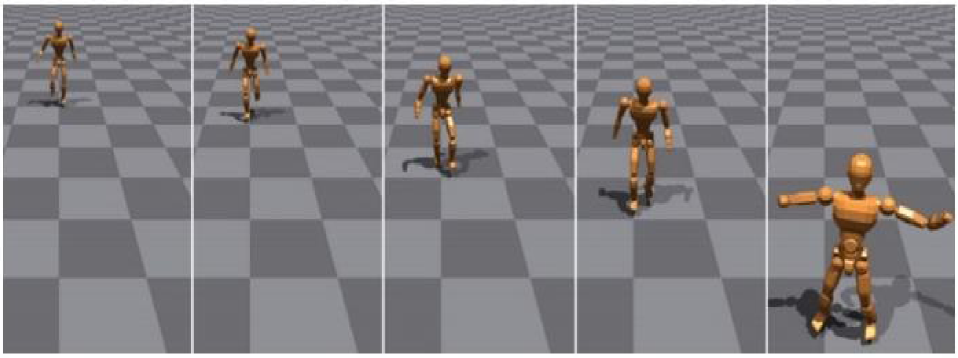

Figure 1. When completing a task, the robot engages in self-reflection to select the optimal plan. For example, when need to reach a certain location, it chooses to walk, thereby avoiding the risks associated with running or jumping.

2.2. Humanoid control

Traditional humanoid control relies on predefined patterns and model-based methods, which often struggle in unpredictable environments [22,37–40]. Reinforcement learning (RL) has emerged as a 72 robust alternative for bipedal locomotion, achieving zero-shot sim-to-real transfer on platforms like Cassie [8,41–44]. Advanced systems such as I-CTRL have extended whole-body imitation to over 7410,000 motions by constraining exploration to ensure visual resemblance. However, most existing 75 systems blindly pursue high reference fidelity, leading to failure when commanded motions exceed 76 the robot's capabilities. Current mitigation strategies often involve filtering complex behaviors (e.g., 7 acrobatics), which limits expressiveness. This study proposes a "capability assessment" mechanism: robots 78 should anticipate execution outcomes and autonomously relax reference constraints when risks are 79 detected.

2.3. Self-awareness

In robotics, physical self-awareness involves monitoring discrepancies between planned movements and current states. Our performance Prediction Network (PPN) draws inspiration from this by continuously analyzing the gap between human references and robot states. Fall prevention is a critical subset of this capability. While early model-based methods used simplified inverted pendulum or ZMP models, they lack generalization to dynamic motions [14,38]. Recent learning-based methods using LSTMs or 1D-CNNs address these limits but remain restricted to simple movements like walking [14,38,45–47]. In contrast, this study leverages I-CTRL to train PPN on a diverse spectrum of human movements. By incorporating the intended reference motion as an input, not just the current state, our system can proactively adapt a high-risk command (e.g., a high jump) into a feasible one before failure occurs [8,44,48–50].

3. Methodology

This study presents a capability-aware motion adaptation pipeline that maps high-level commands c to safe robot motions Rpby interposing human embodiment and predictive self-assessment. The system operates in three stages.

(1) Synthesizing intent consistent human references H, via MotionLCM(f)

(2) Retargeting references to robot space R, via ImitationNet (9h2r)

(3) Refining trajectories into physics consistent motion RP via I-CTRL(gr2p), see **Figure 2**.

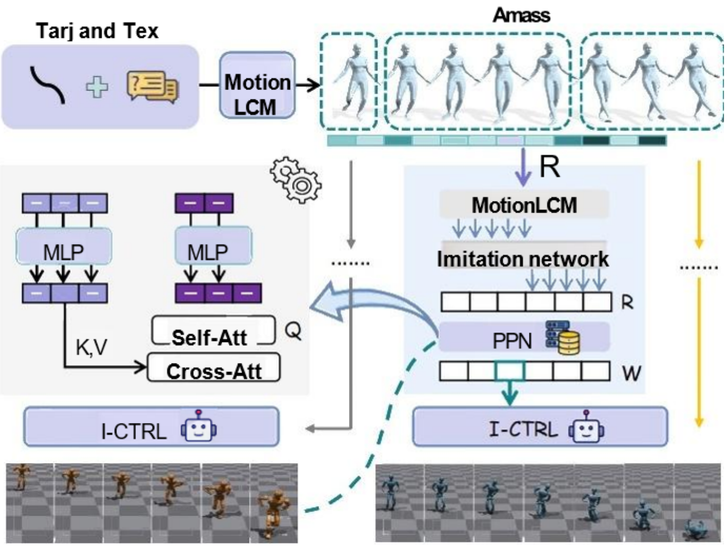

Figure 2. Overview of the motion adaptation system. Given a command, the system generates diverse motion candidates and ranks them using a PPN based on the robot's physical capabilities and current state. The highest-ranked motion is executed by the low-level controller.

3.1. Problem formulation

This study formulates humanoid control as a motion adaptation problem. Given a task c, this study generates motion Rp that respects kinematic and dynamic constraints through the mapping.

$$c \xrightarrow{f} \mathbf{H}_r \xrightarrow{g_{h2r}} \mathbf{R}_r \xrightarrow{g_{r2p}} \mathbf{R}_p \tag{1}$$

Human motion is respected as $\mathbf{H}_r = \{\mathbf{h}_r^t\}_{t=1}^T \in \mathbb{R}^{T\times J\times 3}$, while robot reference R, and physics-based motion Rp include root states (p*,θt) and joint configurations (q6, qt). This study's key insights were to discover an adapted reference H, that maximizes quality Q while maintaining safety S above a threshold Tsafe

$$\hat{\mathbf{H}}_r = \arg\max_{\mathbf{H}'_r} \mathcal{Q}(g(\mathbf{H}'_r), c)\, s.t. S(g(H)) > Tsafe \tag{2}$$

3.2. Motion adaptation framework

The system operates in a receding horizon. At each step t, this study consider observed states $\mathbf{R}^t\mathbf{o}$ and the future reference poses $\mathbf{H}^t\mathbf{f}$. Candidate generation: This generate and diverse candidate $\{\mathbf{H}_{f,i}^t\}_{i=1}^n$ using MotionLCM to provide multiple behavioral alternatives (e.g. walking vs running). Selection: Each candidate is retargeted via gh2r and evaluated by the Performance Prediction Network (PPN): The optimal index i* was selected by prioritizing safety lexicographically, then maximizing quality $\mathbf{w}^\top \mathbf{s}^t\mathbf{i}$.

$$\hat{\mathbf{S}}^t\mathbf{i} = \mathrm{PPN}(\hat{\mathbf{R}}_{f,i}^t, \mathbf{R}^t\mathbf{o}) \tag{3}$$

3.3. Performance prediction network

PPN is a transformer based architecture that forecasts execution quality Encoding: Reference motion $\hat{\mathbf{R}}_{f,i}^t$ and observed states $\mathbf{R}^t\mathbf{o}$ are encoded via MLPs into $\mathbf{E}^t\mathbf{f}$ and $\mathbf{E}^t\mathbf{o}$. This study appends a [cls] token to $\mathbf{E}^t\mathbf{o}$ and apply self-attention to capture temporal dependencies. Conditioning: Cross attention allows observed states to attend to reference features:

$$\mathbf{E}^{t}c = \mathrm{CrossAttn}\,(\mathrm{Query} = \hat{\mathbf{E}}^{t}c\ ,\ \mathrm{Key} = \mathbf{E}^{t}f\ ,\ \mathrm{Value} = \mathbf{E}^{t}f) \tag{4}$$

Score prediction: The context vector $\mathbf{zc} = \mathbf{E}^{t}c^{[0]}$ was project to predict $\hat{\mathbf{S}}^{t}\mathbf{i} = [dfall\ ,\ \hat{A}_q\ ,\ \hat{A}_q^{-}\ ,\ \hat{A}_q^{--},\ \hat{A}_p\ ,\ \hat{A}\theta\ ]^{T}$, quantifying fall probability, alignment errors, and smoothness. Training: The objective is $L = Lfall + \lambda\ 1\ Lalign + \lambda 2\ Lsmooth$, using binary cross entropy for Lfall and MSE for alignment and smoothness.

4. Experiments

This study validated our capability aware motion adaption framework using 85,000 human motion sequences and 255,000 robot trajectories. Our evaluation focuses on

(1) PPN accuracy across multiple time horizons
(2) Adaptation effectiveness in preventing failures
(3) Architectural ablation studies

4.1. Experimental setup

Dataset and Platform. This study generated 85,000 sequences from 8,500 textual prompts using MotionLCM, covering locomotion, dynamic actions, and expressive gestures. Robot executions were simulated using the JVRC-1 model (23 DOF) in IsaacGym via the I-CTRL controller. The test set was balanced with 50% fall samples (e.g. jumps > 45 cm or rapid turns) to ensure discriminative power. Simulations ran at 60 Hz with Kp = 100 and Kd = 10. Metrics. This study evaluate Fall Prediction Accuracy, Alignment MSE (A_q , A_q^{-} , A_p , $A\theta$), and Smoothness MSE ($\hat{A}_q^{--}$). Adaptation is measured by Fall Prevention Rate, Task Completion rate, and trajectory deviation.

4.2. Performance prediction accuracy

Table 1 summarizes the performance of PPN and its variants. Our full model achieves a 99.14% fall prediction accuracy at a 1s horizon.

Table 1. Performance prediction accuracy on test set

Model	Tf	Fall Acc. ↑	A_q^{--}↓	Aq ↓	A_q^{-}↓	Ap ↓	Aθ ↓
w/o Rf	1s	96.85	0.108	0.0289	8.147	0.1253	0.0521
w/o Ro	1s	98.73	0.094	0.0417	6.382	0.1142	0.0298
w/o Cross-Attn	1s	98.91	0.086	0.0183	5.874	0.1067	0.0264
PPN (Ours)	1s	99.14	0.078	0.0159	5.138	1.004E-1	8.90E-3
PPN (Ours)	3s	98.93	0.051	0.0142	4.417	0.1158	0.0287

Analysis. Ablation results confirm that removing reference motion (**R**f) causes the most degradation, with fall accuracy dropping by 2.29% and joint errors increasing by 81.8%. Removing observed states (**R**o) primarily impact root pose accuracy. Compared to simple concatenation (w/o Cross-Attn), the cross-attention mechanism reduces orientation error by 197%, validating its efficacy in modeling state-reference interactions.

4.3. Motion adaptation effectiveness

On a test set of 420 failing commands, our framework achieved a 58.3% fall prevention rate, with 87.6% of adapted motions successfully completing the semantic task (**Table 2**).

Table 2. Motion adaptation performance on failing commands

Metric	Value
Fall prevention rate	58.3%
Task completion rate	87.6 %
Random selection fall prev.	19.8 %
Avg. adaptation time	0.21 s

Quantitative & Qualitative. PPN-guided selection provides a 2.9 × improvement over random selection. Qualitative analysis (**Figure 3**) demonstrates intelligent constraint relaxation: in the karate kicks task, the system reduces kick height by 30% to ensure stability while maintaining the dynamic character. In defend-punch, it shortens the lunge distance to preserve the center of mass while executing the strike.

(a) Human defends him self

(b) Kick

(c) Stand with one leg

(d) Hammering very strong

Figure 3. Qualitative examples of motion adaptation across diverse scenarios. Each row shows: (left) original failing motion, (middle) adapted motion selected by PPN, (right) comparison of root trajectories. Our framework intelligently modifies motion characteristics while preserving task semantics.

4.4. Efficiency and failure analysis

The full pipeline averages 0.21 s supporting 1–5 Hz real time planning. Failures are primarily due to insufficient candidate diversity (48%) and prediction errors (23%)

5. Conclusion

The PPN demonstrates that self-evaluation mechanisms can preemptively identify failures with 99.14% accuracy, allowing robots to transcend rigid constraints. While current results are simulation based, future work must address the sim to real gap, specifically sensor noise and actuator latency. The 58.3% prevention rate indicates significant potential for improvement by expanding candidate motion libraries and refining transition smoothness between original and adapted trajectories.

Disclosure statement

The authors declare no conflict of interest.

References

[1] Dautenhahn K, 2007, Socially Intelligent Robots: Dimensions of Human–Robot Interaction. Philosophical Transactions of the Royal Society B, 362: 679–704.

[2] Thomaz A, Cakmak M, 2016, Learning About Objects with Human Teachers. Human–Robot Interaction, 5: 1–42.

[3] Khatib O, Sentis L, Park J, et al., 2008, Whole-Body Dynamic Behavior and Control of Human-Like Robots. International Journal of Humanoid Robotics, 5: 29–43.

[4] Mainprice J, Hayne R, Berenson D, 2015, Predicting Human Reaching Motion in Collaborative Tasks Using Inverse Optimal Control and Iterative Replanning. IEEE International Conference on Robotics and Automation: 885–892.

[5] Lasota P, Fong T, Shah J, 2017, A Survey of Methods for Safe Human–Robot Interaction. Foundations and Trends in Robotics, 5: 261–349.

[6] Javdani S, Admoni H, Pellegrinelli S, et al., 2018, Shared Autonomy via Hindsight Optimization for Teleoperation and Teaming. International Journal of Robotics Research, 37: 717–742.

[7] Dragan A, Lee K, Srinivasa S, 2013, Legibility and Predictability of Robot Motion. IEEE International Conference on Human–Robot Interaction: 301–308.

[8] Peng X, Abbeel P, Levine S, et al., 2018, DeepMimic: Example-Guided Deep Reinforcement Learning of Physics-Based Character Skills. ACM Transactions on Graphics, 37: 1–14.

[9] Atkeson C, Schaal S, 1997, Robot Learning from Demonstration. International Conference on Machine Learning: 12–20.

[10] Hwangbo J, Lee J, Dosovitskiy A, et al., 2019, Learning Agile and Dynamic Motor Skills for Legged Robots. Science Robotics, 4: eaau5872.

[11] Mordatch I, Todorov E, Popović Z, 2012, Discovery of Complex Behaviors through Contact-Invariant Optimization. ACM Transactions on Graphics, 31: 1–8.

[12] Pratt J, Carff J, Drakunov S, et al., 2006, Capture Point: A Step Toward Humanoid Push Recovery. IEEE-RAS International Conference on Humanoid Robots: 200–207.

[13] Stephens B, Atkeson C, 2010, Push Recovery by Stepping for Humanoid Robots with Force-Controlled Joints. IEEE-RAS International Conference on Humanoid Robots: 52–59.

[14] Koolen T, de Boer T, Rebula J, et al., 2012, Capturability-Based Analysis and Control of Legged Locomotion, Part 1: Theory and Application. International Journal of Robotics Research, 31: 1094–1113.

[15] Bretl T, Lall S, 2008, Testing Static Equilibrium for Legged Robots. IEEE Transactions on Robotics, 24: 794–807.

[16] Hauser K, Bretl T, Latombe J, et al., 2008, Motion Planning for Legged Robots on Varied Terrain. International Journal

of Robotics Research, 27: 1325–1349.

[17] Dai H, Valenzuela A, Tedrake R, 2014, Whole-Body Motion Planning with Centroidal Dynamics and Full Kinematics. IEEE-RAS International Conference on Humanoid Robots: 295–302.

[18] Zucker M, Ratliff N, Dragan A, et al., 2013, CHOMP: Covariant Hamiltonian Optimization for Motion Planning. International Journal of Robotics Research, 32: 1164–1193.

[19] Kalakrishnan M, Chitta S, Theodorou E, et al., 2011, STOMP: Stochastic Trajectory Optimization for Motion Planning. IEEE International Conference on Robotics and Automation: 4569–4574.

[20] Tedrake R, Manchester I, Tobenkin M, et al., 2010, LQR-Trees: Feedback Motion Planning via Sums-of-Squares Verification. International Journal of Robotics Research, 29: 1038–1052.

[21] Ott C, Roa M, Hirzinger G, 2008, Posture and Balance Control for Biped Robots Based on Contact Force Optimization. IEEE-RAS International Conference on Humanoid Robots: 26–33.

[22] Sentis L, Khatib O, 2005, Synthesis of Whole-Body Behaviors through Hierarchical Control of Behavioral Primitives. International Journal of Humanoid Robotics, 2: 505–518.

[23] Peng X, Guo Y, Halper L, et al., 2022, ASE: Large-Scale Reusable Adversarial Skill Embeddings for Physically Simulated Characters. ACM Transactions on Graphics, 41: 1–17.

[24] Rempe D, Birdal T, Zhao Y, et al., 2021, HuMoR: 3D Human Motion Model for Robust Pose Estimation. IEEE International Conference on Computer Vision: 11488–11499.

[25] Schaal S, 1999, Is Imitation Learning the Route to Humanoid Robots? Trends in Cognitive Sciences, 3: 233–242.

[26] Argall B, Chernova S, Veloso M, et al., 2009, A Survey of Robot Learning from Demonstration. Robotics and Autonomous Systems, 57: 469–483.

[27] Abbeel P, Ng A, 2004, Apprenticeship Learning via Inverse Reinforcement Learning. International Conference on Machine Learning: 1–8.

[28] Pomerleau D, 1988, ALVINN: An Autonomous Land Vehicle in a Neural Network. Advances in Neural Information Processing Systems, 1: 305–313.

[29] Ho J, Ermon S, 2016, Generative Adversarial Imitation Learning. Advances in Neural Information Processing Systems, 29: 4565–4573.

[30] Ross S, Gordon G, Bagnell J, 2011, A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. International Conference on Artificial Intelligence and Statistics, 15: 627–635.

[31] Mehta S, Ciftci Y, Ramachandran B, et al., 2024, Stable-BC: Controlling Covariate Shift with Stable Behavior Cloning. arXiv preprint arXiv:2408.06246.

[32] Park J, Kim Y, Song K, et al., 2024, Mitigating Covariate Shift in Behavioral Cloning via Robust Stationary Distribution Correction. Advances in Neural Information Processing Systems, 37.

[33] Ng A, Russell S, 2000, Algorithms for Inverse Reinforcement Learning. International Conference on Machine Learning: 663–670.

[34] Ziebart B, Maas A, Bagnell J, et al., 2008, Maximum Entropy Inverse Reinforcement Learning. AAAI Conference on Artificial Intelligence: 1433–1438.

[35] Wulfmeier M, Ondruska P, Posner I, 2015, Maximum Entropy Deep Inverse Reinforcement Learning. arXiv preprint arXiv:1507.04888.

[36] Finn C, Levine S, Abbeel P, 2016, Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization. International Conference on Machine Learning, 48: 49–58.

[37] Kajita S, Kanehiro F, Kaneko K, et al., 2003, Biped Walking Pattern Generation Using Preview Control of Zero-

Moment Point. IEEE International Conference on Robotics and Automation, 3: 1620–1626.

[38] Vukobratovic M, Borovac B, 2004, Zero-Moment Point—Thirty-Five Years of Its Life. International Journal of Humanoid Robotics, 1: 157–173.

[39] Kuffner J, Kagami S, Nishiwaki K, et al., 2002, Dynamically Stable Motion Planning for Humanoid Robots. Autonomous Robots, 12: 105–118.

[40] Kuffner J, LaValle S, 2000, RRT-Connect: An Efficient Approach to Single-Query Path Planning. IEEE International Conference on Robotics and Automation: 995–1001.

[41] Li Z, Cheng X, Peng X, et al., 2021, Reinforcement Learning for Robust Parameterized Locomotion Control of Bipedal Robots. IEEE International Conference on Robotics and Automation: 2811–2817.

[42] Lee J, Hwangbo J, Wellhausen L, et al., 2020, Learning Quadrupedal Locomotion over Challenging Terrain. Science Robotics, 5: eabc5986.

[43] Xie Z, Clary P, Dao J, et al., 2020, Learning Locomotion Skills for Cassie: Iterative Design and Sim-to-Real. Conference on Robot Learning, 100: 317–329.

[44] Radosavovic I, Xiao T, Zhang B, et al., 2024, Real-World Humanoid Locomotion with Reinforcement Learning. Science Robotics, 9: edi9579.

[45] Renner R, Behnke S, 2006, Instability Detection and Fall Avoidance for a Humanoid Using Attitude Sensors and Reflexes. IEEE/RSJ International Conference on Intelligent Robots and Systems: 2967–2973.

[46] Zhong S, Gao J, Li M, et al., 2025, Fall Analysis and Prediction for Humanoids. Robotics and Autonomous Systems, 185: 104995.

[47] Yang T, Zhang W, Yu Z, et al., 2018, Falling Prediction and Recovery Control for a Humanoid Robot. IEEE-RAS International Conference on Humanoid Robots: 481–487.

[48] Mastalli C, Merkt W, Xin G, et al., 2024, Know Your Limits! Optimize the Robot's Behavior through Self-Awareness. arXiv preprint arXiv:2409.10308.

[49] Cheng X, Shi K, Agarwal A, et al., 2024, Extreme Parkour with Legged Robots. arXiv preprint arXiv:2309.14341.

[50] Khamassi M, Lallée S, Enel P, et al., 2018, Toward Self-Aware Robots. Frontiers in Robotics and AI, 5: 88.