# Know your limits! Optimize the robot's behavior through self-awareness



Esteve Valls Mascaró<sup>1</sup> and Dongheui Lee<sup>1,2</sup> evm7.github.io/Self-AWare

Fig. 1: **Optimizing robot behavior control from high-level task commands using self-awareness**. Given an instruction by an operator, the robot generates multiple potential behaviors to accomplish the task, evaluates them based on its capabilities and limits, and selects the most suitable one to execute. This work introduces a motion adaptation with a Self-AWare model (SAW) to anticipate how well a robot can follow a given reference by ranking multiple potential behaviors and choosing the optimal one. For instance, in a scenario with three potential actions—walking, running, and jumping—the robot assesses each option and determines that walking is the most appropriate. Consequently, the robot walks to the person and says, "Hi." In this image, the generated references are shown by orange robots, while the robot's behavior when attempting to follow them is depicted in blue. Note that the reference motions are fixed just for visualizations, and do not consider gravity.

Abstract-As humanoid robots transition from labs to realworld environments, it is essential to democratize robot control for non-expert users. Recent human-robot imitation algorithms focus on following a reference human motion with high precision, but they are susceptible to the quality of the reference motion and require the human operator to simplify its movements to match the robot's capabilities. Instead, we consider that the robot should understand and adapt the reference motion to its own abilities, facilitating the operator's task. For that, we introduce a deep-learning model that anticipates the robot's performance when imitating a given reference. Then, our system can generate multiple references given a highlevel task command, assign a score to each of them, and select the best reference to achieve the desired robot behavior. Our Self-AWare model (SAW) ranks potential robot behaviors based on various criteria, such as fall likelihood, adherence to the reference motion, and smoothness. We integrate advanced motion generation, robot control, and SAW in one unique system, ensuring optimal robot behavior for any task command. For instance, SAW can anticipate falls with 99.29% accuracy.

# I. INTRODUCTION

Imagine a world where bipedal robots walk among us, mimicking human movements and behaviors with remarkable precision. Such robots could revolutionize industries ranging from healthcare to disaster response, offering unparalleled assistance and efficiency. However, teaching robots how to robustly imitate human behaviors is very challenging, especially for bipedal robots. Recent works [1]-[7] have shown promising potential by training a robot to imitate a given human reference using goal-conditioned reinforcement learning (RL) [8]. However, exactly following a human reference motion might fall out of the robot's capabilities. To address this issue, previous works filtered out unfeasible motions and limited the commanded behaviors to upper body movements [1], [2] or one specific task [3]-[6]. However, this strategy presents two main problems: (a) Should the human operator adapt its reference behavior to the robot's abilities and constraints? (b) How can robots handle dynamic motions that are out of their expertise? This paper is motivated to answer those questions by doting the robots with self-awareness: the ability to understand any motion

<sup>&</sup>lt;sup>1</sup>Esteve Valls Mascaro and Dongheui Lee are with Autonomous Systems, Technische Universität Wien (TU Wien), Vienna, Austria (e-mail: {esteve.valls.mascaro, dongheui.lee}@tuwien.ac.at).

<sup>&</sup>lt;sup>2</sup>Dongheui Lee is also with the Institute of Robotics and Mechatronics (DLR), German Aerospace Center, Wessling, Germany.

command and adapt it according to their own limitations and capabilities.

As humans, we inherently possess self-awareness, which helps us recognize our physical and cognitive boundaries. For instance, a novice parkour enthusiast understands that attempting a high-risk maneuver without sufficient skill can lead to injury. This self-awareness guides them to modify their actions, opting for safer, more manageable moves until they gain more experience and confidence. Similarly, selfawareness in robots is equally important to ensure safety, efficiency, and high task performance. Without it, robots might blindly try to imitate human motions without considering their own physical constraints, leading to falls, collisions, or mechanical damage. In fact, self-awareness should also entail the robot's expertise: a robot might be able to jump with one leg, but still has not learned how. Therefore, during execution, the robot should ideally explore how to move like a human jumping with one leg and adapt its movement based on its own capabilities. By incorporating self-awareness, robots can assess their capabilities and make informed decisions, preventing actions beyond their abilities, as shown in Fig. 1. In this paper, we reformulate this adaptation problem by first generating potential movements and later selecting the most optimal according to the current situations and robot expertise and limitations.

In fact, this area of developing intelligent embodied systems that understand their own limitations has already been explored in the past. Prior works have considered physical self-awareness by predicting whether a robot will fall in the near future [9]-[14]. The main goal of these works is to avoid falls by adjusting its motion before high disturbances are encountered. Both model-based [9], [10] and learning-based approaches [11]-[14] have been considered for the task of fall prediction. For instance, [13], [14] used deep-learning strategies to anticipate falls from the last observed robot states. However, all learning-based methods generated the training data using model-based controllers that were limited to simple movements, such as standing [14] or walking [13]. They collect the falling examples by applying external force perturbations to the robot's torso in a simulator. Consequently, their fall prediction approach was limited to very specific tasks (i.e., walking) and did not generalize to diverse human behaviors. On the contrary, this work proposes a system that understands the robot's capabilities when imitating any human reference provided, even highdynamic movements such as backflips or handstands. We envision a general approach that allows a robot to imitate any human and any behavior commanded if it knows how or adapt the movement to its own expertise.

Moreover, instead of just predicting a robot's potential fall, we anticipate how good a robot's imitation will be given a human reference. For that, we introduce a Self-AWare model, shorted as SAW, that infers a score to inform a robot on how smooth, safe, and faithful a robot's behavior will be if it decides to follow a human reference motion. Additionally, by integrating SAW with a real-time text-tomotion generator, such as MotionLCM [15], we can adapt the human's reference so that the robot imitates the high-level task commanded but improving the quality of the generated behavior. SAW can be used to generate and rank diverse potential robot behaviors and select the most adequate in advance, ensuring an optimal imitation. In conclusion, our efforts resulted in the following contributions:

- 1) A transformer-based model to anticipate the robot's performance when following any human motion given as reference.
- 2) A pipeline to optimize the robot's behavior, prevent falls, and enable non-expert users to control a bipedal robot using text or trajectory commands.

## II. RELATED WORKS

This section reviews the existing literature on robots imitating diverse human behavior, focusing on recent advancements and identifying current limitations. This analysis underscores the importance of integrating physical selfawareness in robots to facilitate their effective deployment in real world-applications.

### A. Imitation Learning

As robots are increasingly utilized in more complex and unstructured environments, manually preprogramming their behavior or defining it through reward functions is becoming exceedingly difficult [16], [17]. Instead, imitation learning (IL) provides an avenue for teaching robots a desired behavior by simply demonstrating it. During the learning process, robots are provided with a dataset of expert demonstrations, and the goal is either to replicate them by mapping the observed states to actions, known as behavior cloning (BC) [18], or to understand the underlying reward functions behind the expert's behavior with inverse RL (IRL) [19]. While IRL provides a more robust solution, its learning process is computationally more expensive and has lately struggled to scale to larger environments and replicate expert behaviors compared to BC [19]. In fact, BC has the advantage of being more efficient, as merely follows a traditional supervised learning task, but might suffer from a covariate shift problem [20]. This problem arises due to the bias induced during training where the state distribution is led by an expert, while in testing the state is induced by its action [21]. Different attempts [18], [22], [23] have focused on solving the covariate shift problem, either by incorporating the human expert in the loop [18] or by limiting the agent actions to be in the distribution covered by the expert demonstrations [23]. For instance, [23] learned to detect states that could lead to failures, and encourage BC to shift to known states. In fact, our work shares some common beliefs with [23] where anticipating potentially unstable states is crucial for ensuring optimal behaviors during BC, but we differ in nature. While [23] constraints the agent policy to be in the distribution of the expert demonstrations, we adapt the policy to the expertise of the agent itself. Our work builds upon behavior cloning but relaxes the imitation process, where the agent should perform a task optimally in its own domain. One similar motivation is proposed in [24], where the goal is to overcome discrepancies between agent embodiments without forcing a shared latent domain.

#### B. Human Imitation in Bipedal Robots

Imitating human motion is no trivial task, especially for bipedal robots. Traditional methods in robot motion planning often rely on pre-defined movement patterns and control algorithms [25], [26] inspired by human movements. These methods, though effective in controlled environments, struggle to adapt to the unpredictability of real-world settings. The lack of flexibility in these approaches results in robots that can perform specific tasks but fail when faced with novel or dynamically changing situations.

Reinforcement Learning (RL) has emerged as a powerful technique to address these limitations of traditional methods, showing high performance in animated characters [27]-[29], quadruped robots, [30], [31] and recently in bipedal humanoids [1]–[7]. All these approaches tackled behavior cloning through goal-conditioned RL [8], where a robot is trained to replicate a given human reference while considering environmental and physical constraints. For instance, [3], [4] focused on training an RL agent to follow either velocity commands or jumps, respectively, and performed zero-shot transfer to a real Cassie robot. To gain more fine-grained control over the robot's behavior, [5], [6] designed extensive rewards to overfit a policy for specific tasks in a bipedal humanoid robot. However, the ultimate goal of robot imitation is to enable a humanoid to follow any human demonstration, independent of its complexity and diversity. To address this task, [1], [2] focused on only imitating the upper-body part to account for more expressive robot behaviors, but only controlled the robot legs to ensure motion feasibility. Instead, I-CTRL [7] focused on whole-body humanoid imitation and showcased high performance in four different bipedal robots, generalizing over 10.000 different motions. For that, I-CTRL constrained the exploration phase during learning so that the produced physics-based behavior only modified the retargeted human reference within a defined margin, ensuring a highly visual resemblance.

However, all the aforementioned works focused on blindly following a reference human motion with high precision, which could lead to falls or non-smooth behaviors when the intended reference behavior exceeds the robot's capabilities. In fact, previous works filter out complex human behaviors (i.e., remove high jumps, backflips, ...) [2], [6], [7], so that their generated behavior underperforms for complex reference motions, leading to falls. To overcome this issue, we propose to anticipate the resulting robot behavior of a trained policy when given a human reference. If a robot is aware that a commanded motion is out of its expertise and might lead to falls or poor imitation, the robot can decide to relax the reference motion constraints and further explore new reference behaviors that might lead to better performance. We described this ability as physical selfawareness, which refers to the robot's understanding of its own limitations and capabilities and enables the robot to make informed decisions.

### C. Self-Awareness in Robotics

Psychologists describe self-awareness as the ability to become the object of one's attention [32], which entails a person's knowledge of themselves. This capability significantly influences human motivations, decisions, and intentions [33], either at the cognitive or physical level. For instance, [33] states that when a human performs a movement, their action awareness is unconsciously monitoring discrepancies between the planned movement and the current state. If this error becomes significant, this awareness alerts a higher-level cognitive system to correct and replan this movement. Our work is inspired by this physical self-awareness and attempts to design a similar system for robots. Our developed Self-AWare model (SAW) continuously monitors discrepancies between the planned movements, the human references, and the current robot states. If SAW anticipates a potential discrepancy, such as a fall or a wrong robot imitation, our system proposes to replan the movement by consciously reasoning about the robot's limitations and capabilities.

In fact, preventing falls in robots, which can be considered as a subarea among physical self-awareness, has already been explored in the past. Early efforts focused on avoiding falls when small disturbances occurred using stabilizing controllers [34], [35]. However, when dealing with large disturbances, a robot requires more time to adjust its motion to avoid falling. Therefore, [9], [10] proposed to anticipate and prevent falls using model-based approaches. [9] used a stand space inverted pendulum model for fall prediction, while [10] modified the zero moment point (ZMP) for a simplified multi-rigid body model to predict falls in humans. Despite some advancements, these methods often oversimplify fall prediction by making assumptions that limit the robot's behavior to specific motions.

To address these limitations and build models that can better adapt to uncertainties, learning-based approaches were proposed [11]-[14]. [11], [12] hand-crafted robot state features to predict falls. Later, [13] inferred the likelihood of falling using a bidirectional Long-Short Term Memory (BiLSTM) network, which processed the evolution of robot sensor measurements such as the center of mass (CoM), the center of pressure (CoP) and the linear and angular momentum and its derivative. More recently, [14] adopted a 1D convolutional neural network (1D-CNN) with the same goal. However, these learning-based methods generated the training data using model-based controllers that were limited to simple movements, such as standing [14] and walking [13]. Consequently, their approach did not generalize to more diverse human behaviors a robot might want to imitate. On the contrary, we use I-CTRL [7] to generalize our SAW model to more complex and diverse motions, such as dancing, walking, running, and jumping. Moreover, and contrary to prior works [11]-[14] that only focus on the observed robot states to predict falls, we incorporate the reference human motion to imitate as input to our model. For instance, when a bipedal robot imitates a human performing a high jump, the robot's self-awareness can recognize that the



Fig. 2: Overview of our system for optimal robot behavior generation from high-level task commands with selfawareness. Upon a high-level instruction, given as trajectory or natural language description, we generate a human reference  $\mathbf{H}_{\mathbf{r}}$  for the robot to follow using MotionLCM [15]. Then, at each time horizon *t*, we assess the feasibility of the next future poses  $\mathbf{R}_{\mathbf{f}}^t$  to follow and optimize the robot behavior using the motion adapter module (shown in the gray center box). For that, MotionLCM is used first to generate potential new references with similar root trajectories and/or textual commands. Those references are retargeted to robot motions using ImitationNet, obtaining  $[\mathbf{\hat{R}}_{\mathbf{f},1}^t, \cdots, \mathbf{\hat{R}}_{\mathbf{f},\mathbf{n}}^t]$ . Our Self-AWare model (SAW, depicted in the light box on the right) ranks each edited references  $\mathbf{\hat{R}}_{\mathbf{f},\mathbf{i}}^t$  according to the robot's learned capabilities and the current robot states  $\mathbf{\hat{R}}_{\mathbf{o}}^t$ . Thus, SAW infers a score vector  $\hat{s}_i^t$  that describes how well the robot can follow a given reference *i* at time *t*, and that can be summarized using a weighted sum to a single scalar per reference. The motion adapter then chooses the best reference and forwards it to I-CTRL [7], which transfers this new reference to real robot commands  $\mathbf{\hat{R}}_{\mathbf{b}}^t$ .

skill is out of its limits and adapt the behavior to perform a simpler jump that is feasible with the current abilities. Note that this prevention formulation differs from prior works on quadruped robots [36] or small bipedal robots [37] that train a specific policy for fall recovery, where the robot has already fallen. In our case, by accurately knowing the robot's limits and anticipating potential failures in advance, we can replan the robot's behavior to be optimal while remaining as close as possible to the intended movement.

## III. METHODOLOGY

This section is structured as follows. First, we introduce the task of robot behavior generation and control from highlevel commands with self-awareness. Then, we present our motion adaptation system with the Self-AWare model (SAW) to tackle this task, which is illustrated in Fig. 2.

#### A. Problem Formulation

We envision the control of bipedal robots using high-level task commands (*c*), such as natural language descriptions or root trajectories. Instead of directly translating those commands to robot signals, we opted for using the human embodiment as a bridge and defining the task as an imitation problem. Therefore, our first subtask is to translate those high-level task commands *c* to an appropriate human motion  $\mathbf{H}_{\mathbf{r}}$ , which will serve as the reference for the robot to imitate.

For that, we adopt a pre-trained MotionLCM [15] as a textto-motion generation model f, where  $(f : c \mapsto \mathbf{H_r})$ . Our second subtask is to translate this human reference  $\mathbf{H_r}$  to a physics-based robot movement  $\mathbf{R_p}$ , such as  $(g : \mathbf{H_r} \mapsto \mathbf{R_p})$ . In fact, we attain the description of g as presented in [7], where first a human-to-robot retargeting module  $g_{h2r}$ translates a human reference  $\mathbf{H_r}$  to a robot reference  $\mathbf{R_r}$ , only considering visual resemblance between the embodiments, and then a pre-trained RL module  $g_{r2p}$  refines  $\mathbf{R_r}$  to ensure plausibility under the real-world physics, generating  $\mathbf{H_p}$ . Following [7], we define  $g_{h2r}$  as ImitationNet [38] and  $g_{r2p}$ as I-CTRL [7].

However, due to the differences between kinematics and dynamics between humans and robots, as well as the error accumulation from ensembling multiple modules (f,  $g_{h2r}$  and  $g_{r2p}$ ), the resulting  $\mathbf{R}_{\mathbf{p}}$  might largely deviate from  $\mathbf{H}_{\mathbf{r}}$ . Thus, we design a Self-AWare model (SAW) that learns how g performs under different human references  $\mathbf{H}_{\mathbf{r}}$ , and slightly adapts those references motions  $\hat{\mathbf{H}}_{\mathbf{r}}$  to achieve an optimal robot behavior that still satisfies the task command c.

Note that, given the use of multiple pre-trained models (*f* is MotionLCM,  $g_{h2r}$  is ImitationNet and  $g_{r2p}$  is I-CTRL),  $\mathbf{H}_{\mathbf{r}}, \mathbf{R}_{\mathbf{r}}$ , and  $\mathbf{R}_{\mathbf{p}}$  are represented differently. A human motion  $\mathbf{H}_{\mathbf{r}} \in \mathbb{R}^{T \times J \times 3}$  is represented as a sequence of *T* human poses, where each pose is defined as *J* joints in Cartesian representation. On the contrary,  $\mathbf{R}_{\mathbf{r}} \in \mathbb{R}^{T \times D_r}$  includes the

root position  $\mathbf{p}^t \in \mathbb{R}^3$  and orientation  $\theta^t \in \mathbb{R}^4$  in quaternion, as well as the robot joint angles  $\mathbf{q}^t \in \mathbb{R}^S$ , so that D = 3 + 4 + S. Similarly,  $\mathbf{R}_{\mathbf{r}} \in \mathbb{R}^{T \times D_r + S}$  also includes the robot joint velocities  $\dot{\mathbf{q}}^t \in \mathbb{R}^S$ .

### B. Motion Adaptation

We consider the motion adaptation task as finding a new human reference  $\hat{\mathbf{H}}_{\mathbf{r}}$  that ensures optimal robot behavior  $\mathbf{R}_{\mathbf{p}}$  while still resembling the original reference  $\mathbf{H}_{\mathbf{r}}$  and the task command *c*. Our reference adaptation block utilizes the last  $T_o$  observed robot states  $\mathbf{R}_{\mathbf{p}}^{t-\mathbf{T}_0:t} = [\mathbf{r}_p^{t-T_o}, \cdots, \mathbf{r}_p^t]$  to ensure that the new references are feasible with the current robot state, simplified as  $\mathbf{R}_{\mathbf{o}}^t$ , and the subsequent  $T_f$  human reference poses  $\mathbf{H}_{\mathbf{r}}^{t:t+\mathbf{T}_f} = [\mathbf{h}_r^t, \cdots, \mathbf{h}_r^{t+T_f}]$  which we aim to adapt, simplified as  $\mathbf{H}_{\mathbf{f}}^t$ , where *t* represents the current time.

We define the motion adaptation similarly to a brainstorming process, where we first generate a set of *n* modified reference motions  $[\hat{\mathbf{H}}_{\mathbf{f},\mathbf{1}}^t, \cdots, \hat{\mathbf{H}}_{\mathbf{f},\mathbf{n}}^t]$ , and then we rank those potential motions to select the most optimal  $\hat{\mathbf{H}}_{\mathbf{f},\mathbf{i}}$  according to the robot feasibility. For that, we employ MotionLCM [15] which allows joint and trajectory-level editing and can ensure that the edited reference motions start from the current pose  $\mathbf{h}_r^t$  and closely approximate to  $\mathbf{h}_r^{t+T_f}$  and its root trajectory. An example of this behavior is shown in Fig. 3.

Later, we convert those edited references to robot references  $[\hat{\mathbf{R}}_{\mathbf{f},1}^t, \cdots, \hat{\mathbf{R}}_{\mathbf{f},n}^t]$  using ImitationNet [38]. Finally, each edited reference  $\hat{\mathbf{R}}_{\mathbf{f},i}^t$  generated is then processed by SAW alongside the observed robot states  $\mathbf{R}_{\mathbf{0}}^t$ , inferring a score  $\hat{s}_i^t$  that informs on how optimal a reference  $\hat{\mathbf{R}}_{\mathbf{f},i}^t$  is. By calculating a score  $\hat{s}_i^t$  for each edited reference  $\mathbf{R}_{\mathbf{0}}^t$ , we can rank and select the most suitable option that prioritizes avoiding falls and then optimize the quality of robot behavior. This chosen reference is followed until the SAW module identifies a better alternative for I-CTRL to guide the robot's actions. An overview of this process is shown in Fig. 2.

#### C. Self-AWare model (SAW)

This section provides a thorough description of our deeplearning-based Self-AWare model (SAW) for evaluating reference motions for optimal robot control. SAW is designed based on two main ideas: a robot's behavior might lead to failure as the current robot state is already not adequate or if the reference is unfeasible for the current situation. Therefore, we first encode the reference motion  $\hat{\mathbf{R}}_{fi}^{t}$  and the observed robot states  $\mathbf{R}_{o}^{t}$  using individual multi-layer perceptrons (MLPs), which leads to  $\mathbf{E}_{f,i}^t$  and  $\mathbf{E}_o^t$  respectively. Next, we make use of the attention mechanisms [39] to summarize both sequences. For that, we append a learnable class token *cls* [40] to the  $\mathbf{E}_{\mathbf{0}}^{\mathbf{t}}$  to aggregate all motion information during the transformer process. We add a sinusoidal positional embedding to  $\mathbf{E}_{\mathbf{0}}^{\mathbf{t}}$  and forward it to a self-attention transformer model to learn the temporal relationships of the observed robot states, embedded as  $\hat{E}_{o}^{t}$ . Finally, we condition this representation with the planned reference motion  $\mathbf{E}_{fi}^{t}$  using cross-attention. Following [40], we extract the appended cls token from the cross-attention output which represents



Fig. 3: Example of our Top-5 ranked edited motions given commanded human reference using our motion adaptation. Here, the motion in orange is the original reference  $\mathbf{H}_r$ , while in blue, we show Top-5 edited references  $[\hat{\mathbf{H}}_{f,1}^t, \cdots, \hat{\mathbf{H}}_{f,5}^t]$  at different times. Note that, as SAW predicts a high likelihood of falling for the 'standing with one leg' for a long time, it tends to place the leg down and start 'walking', as instructed by the high-level task command.

the expected robot behavior in the future if imitating the provided reference motion.  $\widehat{cls}$  is then projected to a score vector  $\hat{s}_i^t$  that assesses factors such as the likelihood of falling  $(\widehat{fall})$ , the smoothness of the generated robot's behavior  $(\hat{A}_{\vec{q}})$ , and the alignment of joint angles  $(\hat{A}_q)$ , joint velocities  $(\hat{A}_{\vec{q}})$ , and root position  $(\hat{A}_p)$  and orientation  $(\hat{A}_{\theta})$  with the reference. We define alignment as the mean square error between the reference state and the generated robot state, and smoothness as low robot joint accelerations (no jittering). Thus, a predicted score  $s = [\widehat{fall}, \hat{A}_q, \hat{A}_{\dot{q}}, \hat{A}_p, \hat{A}_{\theta}] \in \mathbb{R}^6$ .

#### IV. EXPERIMENTS

### A. Dataset generation

To train the SAW model effectively, we automatically generate a dataset that accounts for the robot's ability to accurately mimic real-world motions. This involved synthesizing 100,000 human motion references using MotionLCM, incorporating 10,000 diverse textual annotations, with motion durations ranging from 3 to 12 seconds. For each reference, we simulated three unique JVRC-1 robot behaviors [41] using I-CTRL within the IsaacGym simulator, culminating in a comprehensive dataset of 300,000 robot behaviors. To ensure a fair evaluation, in our validation and testing dataset, the robot falls in half of the motions. Our SAW module observes 0.5 seconds in the past to score the reference of the next 1, 2, or 3 seconds.

#### B. Metrics

To evaluate our SAW module we make use of simple mean-square error metrics between the predicted and groundtruth alignment scores  $(\hat{A}_q, \hat{A}_q, \hat{A}_p, \hat{A}_\theta)$ . We also use accuracy as a metric to assess the falling prediction ability of SAW. Note that a robot is considered to have fallen when the root

	Future Horizon	Fall Accuracy. (%) $\uparrow$	Smoothness Error $(A_{\ddot{q}}) \downarrow$	$A_q$ Error $\downarrow$	$A_{\dot{q}}$ Error $\downarrow$	$A_p$ Error $\downarrow$	$A_{\theta}$ Error $\downarrow$
w/o $\mathbf{R}_{f}$	1	97.70	0.093	0.0252	7.232	0.1136	0.0477
w/o <b>R</b> <sub>o</sub>	1	99.16	0.080	0.0396	5.760	0.1008	0.0225
w/o C.A.	1	99.14	0.074	0.0162	5.320	0.0944	0.0231
SAW*	1	99.28	0.076	0.0156	5.248	0.0992	0.0207
SAW	1	99.29	0.070	0.0144	4.592	0.0928	0.0066
SAW	2	99.17	0.057	0.0144	4.016	0.0976	0.0228
SAW	3	99.20	0.043	0.0126	4.000	0.1024	0.0261

TABLE I: Quantitative evaluation of SAW for score prediction. The first section showcases the benefits of SAW over other variants in the architecture. The second section showcases the robustness of SAW over longer future horizons.

height of the robot is lower than a predefined threshold, following the definition from [7].

## C. Robot

In this work, we demonstrate the benefit of self-awareness on the JVRC-1 robot [41]. This robot has 23 DoFs with a height of 140 centimeters and a weight of 62.2 kilograms.

#### D. Quantitative and Qualitative Evaluation

Due to the absence of existing benchmarks for score prediction, we conducted an ablation study to assess (a) different variations of our proposed architecture, and (b) different future horizons in which we anticipate the robot's behavior. The results in Table I showcase the high accuracy of SAW in predicting if a robot will fall ( $\approx 99\%$ ), as well as the high precision in determining the quality of the generated motion across multiple horizons.

First of all, we observe that when only considering the observed robot states (w/o  $\mathbf{R}_{f}$ ), SAW underperforms when predicting the alignment of the robot with the reference, as expected. Likewise, when SAW only considers the reference motion (w/o  $\mathbf{R}_{o}$ ), it can not ground the expected behavior to its current state, thus also failing to correctly predict an alignment. Later on, we evaluate two modifications on how to integrate  $\mathbf{R}_f$  and  $\mathbf{R}_o$  for the score prediction. First, we consider one independent self-attention block per sequence and compute the score from the concatenation of the cls token from both sequences (w/o C.A.). Secondly, we inverted the sequences in the computation, such as the self-attention is computed over the reference and later conditioned on  $\mathbf{E}_{o}$ . Both variants achieve higher alignment errors and lower fall prediction accuracy. In general, we observe that the observed motion  $\mathbf{R}_o$  has a higher influence on the future robot behavior quality rather than  $\mathbf{R}_{f}$ , as also shown in SAW w/o  $\mathbf{R}_{f}$ .

Additionally, we evaluated the performance of SAW over longer future horizons (from 1 second to 3 seconds) as reported in Table I, which showcases that we can still predict falls with high accuracy with low error in the robot performance behavior. We observed that the anticipation of the smoothness and joint angle and velocity alignment is independent of the future horizon, but rather depends strongly on the reference motion to follow. On the contrary, it is harder to anticipate the root alignment the longer the future horizon.

Finally, we evaluated the overall system when the original robot's behavior (without self-awareness) was falling. Fig. 4

showcases the benefit of SAW when commanding motions out of the robot expertise. In general, our results demonstrate that our motion adaptation system can prevent 62% falls while still following the same root trajectory.

### E. Limitations and future work

Despite our SAW module's high performance in predicting the quality of the robot's motion when following a given reference ( $\approx 99\%$ ), the end-to-end system only prevents 62% of falls during tasks. We recognize that simply stopping the robot's motion and avoiding complex motions when a fall is anticipated could address these issues. However, this approach would prevent the robot from following the original commands of the human operator which was part of the scope of the motion adapter. Our analysis indicates that the primary cause of the system's limited fall prevention is the noisy reference motions generated by MotionLCM when conditioned on prior poses. These noisy references can lead to abrupt changes that I-CTRL cannot effectively predict. Additionally, we currently generate only 15 reference motions and rank the robot's choices among them. Increasing this number to 30 or 50 would provide greater diversity, giving SAW better options to choose from, thereby potentially improving fall prevention.

## V. CONCLUSION

Controlling any robot with high-level task commands often requires operators who understand the capabilities of the robots and limit the command accordingly. Instead, we propose to provide robots with self-awareness so that they can autonomously adjust the commanded instruction to their own ability. Our end-to-end system converts natural language and trajectory instructions into a human motion reference for the robot to follow. To mitigate the risk of poor robot performance due to strict adherence to commands, we introduce the Self-AWare model (SAW). Our SAW module is able to anticipate a fall with  $\approx 99\%$  accuracy at different future horizons, and rank different references with high precision. Finally, we extend our self-awareness model with a motion adapter system that enables the robot to intelligently select the optimal reference motion according to the current robot states and capabilities, enhancing performance and preventing falls.

# A human defends itself and then proceeds to throw a punch.



Fig. 4: Robot behavior comparison between using self-awareness (in dark and light blue) or not (in yellow and brown). Given a textual command, we use MotionLCM [15] to generate a human reference, which we retarget to a JVRC-1 kinematics using ImitationNet [38] to obtain a reference robot motion. Then, I-CTRL [7] is used to refine this robot reference (in yellow) to conform with real-world physics (in orange). However, if our SAW model anticipates poor imitation (i.e., falls), we adapt the reference motion to prevent wrong behaviors. For that, we use MotionLCM+ImitationNet to generate various new robot reference motions (in dark blue) and SAW selects the more appropriate, which we then command to I-CTRL to optimize the robot behavior (in light blue) to follow the initial commands (described here by text). Note that  $t_k$  indicates different time instants of the motion sequences.

#### ACKNOWLEDGMENT

This work is funded by Marie Sklodowska-Curie Action Horizon 2020 (Grant agreement No. 955778) for the EU projects of PERSEO and INVERSE (GA nr. 101136067).

#### REFERENCES

- [1] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024.
- [2] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid realtime whole-body teleoperation. arXiv:2403.04436, 2024.
- [3] Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [4] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Robust and versatile bipedal jumping control through multi-task reinforcement learning. arXiv preprint arXiv:2302.09450, 2023.
- [5] Qiang Zhang, Peter Cui, David Yan, Jingkai Sun, Yiqun Duan, Arthur Zhang, and Renjing Xu. Whole-body humanoid robot locomotion with human reference. arXiv preprint arXiv:2402.18294, 2024.
- [6] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. arXiv preprint arXiv:2406.06005, 2024.
- [7] Yashuai Yan, Esteve Valls Mascaro, Tobias Egle, and Dongheui Lee. I-ctrl: Imitation to control humanoid robots through constrained reinforcement learning. arXiv preprint arXiv:2405.08726, 2024.
- [8] Matthias Plappert, Marcin Andrychowicz, Alex Ray, Bob McGrew, Bowen Baker, Glenn Powell, Jonas Schneider, Josh Tobin, Maciek Chociej, Peter Welinder, et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. arXiv preprint arXiv:1802.09464, 2018.
- [9] Thomas Muender and Thomas Röfer. Model-based fall detection and fall prevention for humanoid robots. Springer-Verlag, 2017.
- [10] Min Li, Guanghua Xu, Bo He, Xiaolong Ma, and Jun Xie. Pre-impact fall detection based on a modified zero moment point criterion using data from kinect sensors. *IEEE Sensors Journal*, 2018.
- [11] Shivaram Kalyanakrishnan and Ambarish Goswami. Learning to predict humanoid fall. *I. J. Humanoid Robotics*, 2011.
- [12] Rajesh Subburaman, Jinoh Lee, Darwin G. Caldwell, and Nikos G. Tsagarakis. Multi-sensor based fall prediction method for humanoid robots. In 2016 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), 2016.
- [13] Dongdong Liu, Hoon Jeong, Aoxue Wei, and Vikram Kapila. Bidirectional lstm-based network for fall prediction in a humanoid. In 2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pages 129–135, 2020.
- [14] M Eva Mungai, Gokul Prabhakaran, and Jessy W Grizzle. Fall prediction for bipedal robots: The standing phase. 2024 IEEE International Conference on Robotics and Automation (ICRA, 2024.
- [15] Wenxun Dai, Ling-Hao Chen, Jingbo Wang, Jinpeng Liu, Bo Dai, and Yansong Tang. Motionlcm: Real-time controllable motion generation via latent consistency model. arXiv preprint arXiv:2404.19759, 2024.
- [16] Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, Jan Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 2018.
- [17] Harish Ravichandar, Athanasios S Polydoros, Sonia Chernova, and Aude Billard. Recent advances in robot learning from demonstration. *Annual review of control, robotics, and autonomous systems*, 2020.
- [18] Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. Machine Learning Research. PMLR, 2011.
- [19] Stuart Russell. Learning agents for uncertain environments. In *The* eleventh annual conference on Computational learning theory, 1998.
- [20] Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. Advances in neural information processing systems, 1988.
- [21] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

- [22] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. {SQIL}: Imitation learning via reinforcement learning with sparse rewards. In International Conference on Learning Representations, 2020.
- [23] Josiah W, Albert T, Andrey K, Ajay M, Li F-F, Silvio S, and Roberto M-M. Error-aware imitation learning from teleoperation data for mobile manipulation. In *Conference on Robot Learning*. PMLR, 2022.
- [24] Arnaud Fickinger, Samuel Cohen, Stuart Russell, and Brandon Amos. Cross-domain imitation learning via optimal transport. In *International Conference on Learning Representations*, 2022.
- [25] Christian Ott, Dongheui Lee, and Yoshihiko Nakamura. Motion capture based human motion recognition and imitation by direct marker control. In *IEEE-RAS International Conference on Humanoid Robots*, 2008.
- [26] Toru Takenaka, Takashi Matsumoto, and Takahide Yoshiike. Real time motion generation and control for biped robot. In 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems.
- [27] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. ACM Trans. Graph., (4), 2018.
- [28] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. ACM Trans. Graph., 40(4), July 2021.
- [29] Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated avatars. In *IEEE/CVF International Conference on Computer Vision*, pages 10895–10904, 2023.
- [30] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning*, 2023.
- [31] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: Learning a unified policy for manipulation and locomotion. In *Conference on Robot Learning*. PMLR, 2023.
- [32] Shelley Duval and Robert A Wicklund. A theory of objective self awareness. 1972.
- [33] Elisabeth Pacherie. The phenomenology of action: A conceptual framework. *Cognition*, 2008.
- [34] Stefano Dafarra, Gabriele Nava, Marie Charbonneau, Nuno Guedelha, Francisco Andradel, Silvio Traversaro, Luca Fiorio, Francesco Romano, Francesco Nori, Giorgio Metta, and Daniele Pucci. A control architecture with online predictive planning for position and torque controlled walking of humanoid robots. 2018.
- [35] Zhibin Li, Chengxu Zhou, Qiuguo Zhu, Rong Xiong, Nikos Tsagarakis, and Darwin Caldwell. Active control of under-actuated foot tilting for humanoid push recovery. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.
- [36] Joonho Lee, Jemin Hwangbo, and Marco Hutter. Robust recovery controller for a quadrupedal robot using deep reinforcement learning. *arXiv preprint arXiv:1901.07517*, 2019.
- [37] Tuomas H, Ben M., Guy L., Sandy H H., Dhruva T., Jan H., Markus W., Saran T., Noah Y S., Roland H., et al. Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *Science Robotics*, 2024.
- [38] Yashuai Yan, Esteve Valls Mascaro, and Dongheui Lee. Imitationnet: Unsupervised human-to-robot motion retargeting via shared latent space. In 2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids), pages 1–8.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. Advances in neural information processing systems (NeurIPS), 30, 2017.
- [40] Alexey D., Lucas B., Alexander K., Dirk W., Xiaohua Z., Thomas U., Mostafa D., Matthias M., Georg H., Sylvain G., et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.
- [41] M. Okugawa, K. Oogane, M. Shimizu, Y. Ohtsubo, T. Kimura, T. Takahashi, and S. Tadokoro. Proposal of inspection and rescue tasks for tunnel disasters — task development of japan virtual robotics challenge. In 2015 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), pages 1–2, 2015.