

Learning to Walk and Fly with Adversarial Motion Priors

Giuseppe L'Erario^{1,3}, Drew Hanover², Ángel Romero², Yunlong Song²,
Gabriele Nava¹, Paolo Maria Viceconte¹, Daniele Pucci^{1,3}, Davide Scaramuzza²

Abstract—Robot multimodal locomotion encompasses the ability to transition between walking and flying, representing a significant challenge in robotics. This work presents an approach that enables automatic smooth transitions between legged and aerial locomotion. Leveraging the concept of Adversarial Motion Priors, our method allows the robot to imitate motion datasets and accomplish the desired task without the need for complex reward functions. The robot learns walking patterns from human-like gaits and aerial locomotion patterns from motions obtained using trajectory optimization. Through this process, the robot adapts the locomotion scheme based on environmental feedback using reinforcement learning, with the spontaneous emergence of mode-switching behavior. The results highlight the potential for achieving multimodal locomotion in aerial humanoid robotics through automatic control of walking and flying modes, paving the way for applications in diverse domains such as search and rescue, surveillance, and exploration missions. This research contributes to advancing the capabilities of aerial humanoid robots in terms of versatile locomotion in various environments. Video: <https://youtu.be/mi6Do-x67CM>

I. INTRODUCTION

The transition between locomotion styles is a phenomenon that many species across the animal kingdom exhibit. Not many robots, however, can combine aerial and terrestrial locomotion, thus leaving the problem of traversing different environments still partially open. The broader research community has begun to study what it takes to enable multimodal locomotion for various bio-inspired robotic systems [1]. Countless examples have demonstrated systems with bespoke forms of locomotion: bat-inspired robots that can fly and walk [2], propeller-driven bipedal robots [3], morphing rovers [4], rolling quadrupeds [5], underwater robots [6], [7] and even flying humanoid robots [8], [9].

However, little work exists on understanding how and when to transition between the locomotion forms. For example, when dealing with an aerial-humanoid robot, what is the proper time to transition between modes such as walking or flying when considering a high-level locomotion task? This goal can be accomplished by breaking out the controller into submodules, where each module handles an individual locomotion task, as proposed by [3]. Nevertheless, the transition between locomotion forms is a significant challenge and remains an open question. This paper moves forward with a learning-based method that enables aerial humanoid robots to exhibit multimodal locomotion capabilities.

¹Artificial and Mechanical Intelligence, Istituto Italiano di Tecnologia, Genoa, Italy ²Robotics and Perception Group, University of Zurich, Switzerland. ³University of Manchester, Manchester, UK. D. Hanover, A. Romero, Y. Song, and D. Scaramuzza were supported by the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT).




Fig. 1: iRonCub, the aerial humanoid robot, on a complex terrain.

A standard approach to address the locomotion problem is to decompose the system into two layers [10]: 1) A trajectory optimization (TO) layer, whose role is to provide a set of feasible trajectories [11], [12]; 2) Whole-body control, which stabilizes the trajectories produced by the TO layer [13].

The TO layer lies at the core of the locomotion architecture, and it is responsible for generating a set of feasible trajectories that the robot can track. It is formulated as an optimization problem that minimizes a cost function – encoding the desired robot behavior – subject to constraints – characterizing the system dynamics, also called *model*. Typical approaches for legged systems as humanoid robots rely on *simplified* models such as the linear inverted pendulum (LIP) [14], or the divergent component of motion (DCM) [15]. While optimization methods using these simplified models are fast, they cannot exploit the robot's complete structure. In contrast, approaches leveraging the robot's *full* dynamics can utilize its inherent system structure, albeit they are computationally expensive. *Reduced* models as the *centroidal momentum dynamics* are a compromise between the two extremes and reduce the problem's complexity [16]–[19]. Nonetheless, TO methods remain computationally intensive and demand prior knowledge of the environment. On the other side of the spectrum, model-free techniques such as Reinforcement Learning (RL) can learn a policy that maps the state of the system to action without the need for a model [20]. These tools enable a new set of control strategies for the robotics realm [21]–[25]. The main challenge in RL for robot control lies in the limited ability of trained agents to exhibit natural behavior in specific tasks, often resulting in inefficient solutions that exploit simulator inaccuracies. To address this limitation, researchers have focused on how nature's ingenuity has shaped the movements of animals, such

as motor skills exhibited by humans, which they embed into simulated through data-driven approaches [26]–[30]. These techniques recently found their way into robotics, where they are used to train robots to imitate animal-like motions [31]–[33]. One such method, called Adversarial motion priors (AMP) [27], takes cues from Generative Adversarial Imitation Learning (GAIL) to train an agent to replicate the “style” embedded in a reference dataset [34].

This paper progresses in this direction by proposing a method that utilizes AMP and RL to study the spontaneous emergence of automatic and smooth locomotion transitions in aerial humanoid robots. AMP allows the robot to learn and mimic human-like motions when walking - enhancing the naturalness of the gait - and imitating flying motion obtained via trajectory optimization. The output of the proposed approach is a locomotion pattern that switches between human-like and flying locomotion based on the task and the environment. We train locomotion policies for iRonCub, a jet-powered aerial humanoid robot with terrestrial and aerial locomotion capabilities. We learn to mimic the motion styles provided within the motion datasets while achieving a high-level waypoint tracking task. The transition between flying and walking is managed by including an energy proxy term in the reward function, encouraging the robot to walk when the ground is within reach and fly otherwise. The method is tested in both the cases of *ideal* thrust propulsion and *specific* jet-powered actuation modeled from real-world data.

We performed experiments within the Nvidia Isaac Gym environment [35], [36], forcing the agent to traverse complex terrains and rough courses. To our knowledge, this is the first demonstration of an aerial humanoid robot exhibiting smooth transitions between walking and flying without explicitly tracking a trajectory provided by trajectory optimization or using a state machine to switch between locomotion modes. Instead, our agent learns how to fly, walk, and navigate without explicitly being told how or what locomotion to use.

The paper is organized as follows. Sec II introduces notation and recalls some critical ideas about floating-base systems. Sec III presents the whole approach by recalling the AMP method and introducing the system modeling. Sec IV validates the technique on iRonCub, a flying humanoid robot. Sec V concludes the paper with remarks and future work.

II. BACKGROUND

A. Floating-base modeling

A robot can be modeled as a multi-body system composed of $n + 1$ rigid bodies – called links – connected by n joints with one degree of freedom. Using of the *floating base* formalism [37], the robot’s configuration is defined by the tuple $q = (\mathcal{I}p_B, \mathcal{I}R_B, \theta) \in \mathbb{Q}$, where $\mathcal{I}p_B \in \mathbb{R}^3$ and $\mathcal{I}R_B \in SO(3)$ are the position and orientation of the robot base frame \mathcal{B} w.r.t. the inertial frame \mathcal{I} , and $\theta \in \mathbb{R}^n$ are the joint positions. The element $\nu = (\mathcal{I}\dot{p}_B, \mathcal{I}\omega_B, \dot{\theta})$ is the robot’s velocity, where $\dot{\theta}$ are the joint velocities and $(\mathcal{I}\dot{p}_B, \mathcal{I}\omega_B)$ are the linear and angular velocity of the base frame \mathcal{B} w.r.t. \mathcal{I} such that $\mathcal{I}\dot{R}_B = S(\mathcal{I}\omega_B)\mathcal{I}R_B$ is satisfied. By applying the

Euler-Poincaré formalism, the equation of motion of a system exchanging n_b wrenches with the environment results in

$$M(q)\dot{\nu} + C(q, \nu)\nu + G(q) = \begin{bmatrix} 0_{6 \times 1} \\ \tau \end{bmatrix} + \sum_{i=1}^{n_b} J_i^\top f_i, \quad (1)$$

where $M, C \in \mathbb{R}^{(n+6) \times (n+6)}$ are the mass and Coriolis matrix, $G \in \mathbb{R}^{n+6}$ is the gravity vector, $\tau \in \mathbb{R}^n$ are the internal actuation torques, f_i is the i -th of the n_b external wrench applied on the origin of the frame \mathcal{C}_i and $J_i \in \mathbb{R}^{6 \times (n+6)}$ is the jacobian mapping the system velocity ν to the velocity $(\mathcal{I}\dot{p}_{C_i}, \mathcal{I}\omega_{C_i})$ of the frame \mathcal{C}_i . These quantities are useful to understand the choice of the state and action space of the robot.

B. Reinforcement Learning

We model the problem as a Markov Decision Process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$ [20]. The state space \mathcal{S} is the set of all possible states of the system and the environment, while the action space \mathcal{A} is the set of all possible actions that the agent can execute. The transition probability $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \text{Pr}[\mathcal{S}]$ defines the probability of transitioning from a state s_t to a state s_{t+1} given an action a_t . The discount factor $\gamma \in [0, 1)$ trades off long-term against short-term rewards. In the context of Deep RL, the agent is modeled as a neural network policy $\pi_\sigma(a_t|s_t)$ that maps a state s_t to an action a_t , with σ the parameters of the neural network. The goal of DRL is to find the optimal set of parameters σ that maximize the expected discounted return

$$J(\sigma) = \mathbb{E}_{\pi_\sigma} \left[\sum_{t=0}^{T-1} \gamma^t r_t \right], \quad (2)$$

where $\mathcal{R}(s_t, a_t, s_{t+1})$ represents the reward r_t received when the agent executes an action a_t in s_t and transitions to s_{t+1} .

III. METHOD

Our work aims to train a policy π capable of walking and flying, automatically selecting the best locomotion pattern to accomplish a task without needing a high-level planner. The agent should be able to imitate the locomotion patterns extracted from a set of motion datasets \mathcal{D} , in which each sample represents a snapshot capturing the system’s kinematics at a specific instant.

A. Adversarial Motion Priors

The agent is trained to imitate the motion of the dataset and to accomplish a task at the same time. This requirement translates into a reward function composed of two terms, weighted by scalar values w_G, w_S , as proposed in AMP [27]:

$$r_t = w_G r_t^G + w_S r_t^S. \quad (3)$$

The *task reward* r^G specifies the goal the agent should accomplish, e.g., reach a target or have a desired base velocity. The *style reward* r^S enforces the policy to produce locomotion patterns resembling the motion of the dataset. The total reward encourages the agent to imitate a desired set of motion patterns while choosing the one that best accomplishes the desired goal. Given the environment, no high-level planner is required to

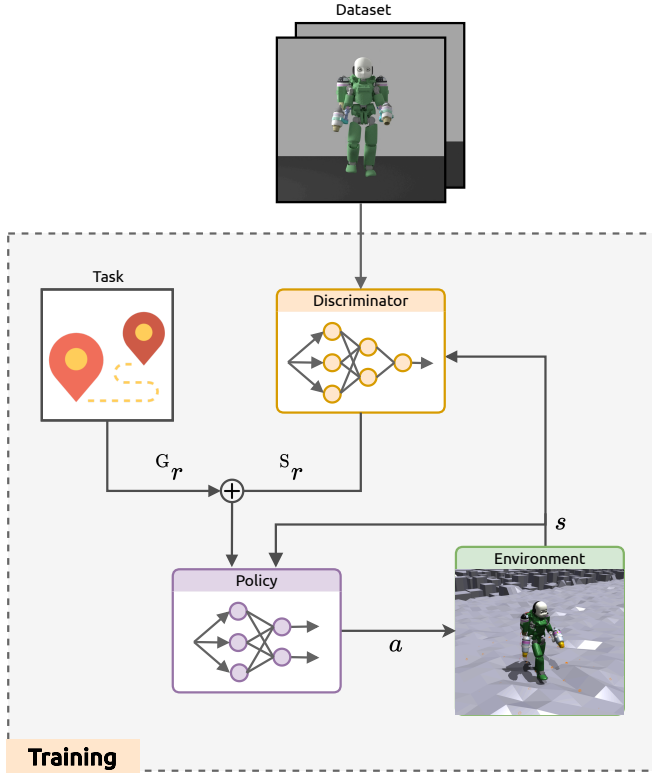


Fig. 2: The discriminator learns to distinguish between samples from the dataset and samples produced by the agent. The policy π_σ is trained to imitate the dataset's motion and accomplish a task simultaneously by maximizing the total reward $r(t)$ that expresses the quality of the motion and the task accomplishment.

select a specific motion, and the switching behavior emerges naturally. An overview of the system is given in Fig. 2.

1) *Style-reward*: Generating data that resemble existing information is a characteristic of Generative Adversarial Networks (GANs) [38]. GANs consist of two competing neural networks: a generator that creates new data that resembles the training set and a discriminator that tries to distinguish between real and generated data, trained in a game-like fashion. Similarly, we can use a discriminator D to distinguish between the samples of the dataset and those produced by the agent π_σ that, in turn, acts as a generator aiming at creating realistic movements.

In practice, the discriminator $D_\Phi : \mathbb{R}^k \times \mathbb{R}^k \rightarrow \mathbb{R}$ is a neural network with parameters Φ that maps a couple of consecutive samples χ_t, χ_{t+1} of dimension k to a scalar value and aims at differentiating motions from the dataset \mathcal{D} from movements produced by the RL agent following the policy π_σ . The objective proposed in [27], [28] to train the discriminator is

$$\begin{aligned} \arg \min_{\Phi} & - \mathbb{E}_{(\chi_t, \chi_{t+1}) \sim \mathcal{D}} [\log(D_\Phi(\chi_t, \chi_{t+1}))] \\ & - \mathbb{E}_{(\chi_t, \chi_{t+1}) \sim \pi_\sigma} [\log(1 - D_\Phi(\chi_t, \chi_{t+1}))] \\ & + w_{gp} \mathbb{E}_{(\chi_t, \chi_{t+1}) \sim \mathcal{D}} [\|\nabla_{\Phi} D_\Phi(\chi_t, \chi_{t+1})\|^2], \end{aligned} \quad (4)$$

where the first two terms force the discriminator to output $D_\Phi(\chi_t, \chi_{t+1}) = 1$ when the discriminator is fed with dataset

transitions and $D_\Phi(\chi_t, \chi_{t+1}) = 0$ when it is fed with samples produced by the agent. The final term is a gradient penalty that penalizes nonzero gradients on dataset transitions, resulting in improved stability and quality of the GAN training [39].

When fed with samples produced by the agent, the discriminator serves as a critic that evaluates the quality of the motion produced by the policy. The output of $D(\chi_t, \chi_{t+1})$ is hence used to compute the style reward [27], [28]

$$s_{r_t} = -\log\left(1 - \frac{1}{1 + e^{-D_\Phi(\chi_t, \chi_{t+1})}}\right), \quad (5)$$

in which the negative logarithm is used to map the output of the discriminator to a reward that is high when the agent produces a transition similar to the dataset.

2) *Task-reward*: We want the robot to reach a target and follow a route made of checkpoints. Ideally, we could reward the agent when it hits the target. This approach makes associating a reward with an individual action complex since it introduces reward sparsity. A common practice is to use proxy rewards that guide the agent to the true objective and provide continuous feedback to the agent. We use three rewards to accomplish the task.

At each time instant, the agent can be associated with a checkpoint. First, we want to minimize the distance of the robot base to the target at each time instant t . This requirement translates into the reward:

$$c_{r_t} = \exp(-c_1 \|x_d - p_t\|^2), \quad (6)$$

where $x_d \in \mathbb{R}^3$ is the target position, $p_t \in \mathbb{R}^3$ is the position of the robot base, and c_1 a hyperparameter.

Second, we want the robot to approach the checkpoint with a desired velocity $v_d \in \mathbb{R}$. This reward is written as:

$$v_{r_t} = \exp\left(-c_2 \left\|v_d - \frac{g(p_t) - g(p_{t-1})}{\Delta t}\right\|^2\right), \quad (7)$$

where $g(p) = \|p - x_d\|$ is a function that returns the scalar distance between the robot base and the target, Δt is the time step, and c_2 is a hyperparameter. This term encourages the agent to travel at a desired velocity v_d on a line connecting the robot base to the target.

The last reward demands the agent to face the checkpoint, i.e., the robot base x -axis should point to the projection of the target in the horizontal plane:

$$f_{r_t} = \min(0, f(p_t)_{xy} \cdot \mathcal{I}_{i_{xy}}), \quad (8)$$

where $f(p) = \frac{x_d - p}{\|x_d - p\|}$ is a function that returns the unit vector pointing from the robot base to the target, the subscript xy extracts the horizontal components, and the unit vector \mathcal{I}_i is the robot x -axis in the inertial frame.

The reward components (6), (7), and (8) fully describe the task and become equal to 1 when the robot reaches the target.

Additionally, we want to minimize thrust usage. The thrust penalty can be considered a proxy term that minimizes propulsion expenditure and encourages the agent to use legged locomotion when possible. The thrust penalty is written as

$$T_{r_t} = -\left\|\frac{T_t}{T_{\max}}\right\|^2, \quad (9)$$

where T_{\max} is the maximum thrust each jet can exert while $T_t \in \mathbb{R}^m$ is the thrust vector acting on the system.

Finally, the total task reward is written as

$$G r_t = w_c c r_t + w_v v r_t + w_f f r_t + w_T T r_t, \quad (10)$$

where w_c, w_v, w_f , and w_T are hyperparameters that weight the contribution of each reward.

B. Model representation

Sec II-A recalls how the quantities describing a multi-body system affect the robot's evolution. We build the action and the observation spaces accordingly, considering also the environment and task information.

1) *Observation Space*: The observation space consists of two components: one relates to the robot, and one distillates information about the task and the environment. We define the robot observation vector as

$$o_{\text{robot}} = \left[\rho^\top \quad \mathcal{B}v_B^\top \quad \mathcal{B}\omega_B^\top \quad \theta^\top \quad \dot{\theta}^\top \quad p_{\text{EE}}^\top \quad T^\top \right]^\top, \quad (11)$$

where ρ is the quaternion representation of the rotation of the base ${}^{\mathcal{I}}R_B$, $\mathcal{B}v_B$ and $\mathcal{B}\omega_B$ are the base velocity in the base frame, θ and $\dot{\theta}$ are the joints position and velocity, p_{EE} is the vector containing the relative position of the end effectors, namely hands and feet, and $T = [T_1, \dots, T_m]$ is the thrust acting on the system as specified in Eq. (16).

The environment observation gives information about the robot's interaction with the environment. We use an elevation map simulating a scan of the terrain around the robot

$$o_{\text{env}} = \begin{bmatrix} h_{11} & h_{12} & \dots & h_{0L} \\ h_{21} & h_{22} & \dots & h_{2L} \\ \vdots & \vdots & \ddots & \vdots \\ h_{V1} & h_{V2} & \dots & h_{VL} \end{bmatrix}. \quad (12)$$

The elevation map is centered at the robot base and moves with the robot. The grid is discretized into $V \times L$ cells. $h_{ij} \in \mathbb{R}$ is the height of the i -th row and j -th column of the grid map with respect to the robot base.

The task observation vector gives the agent information about the goal. We use the position of the target x_d with respect to the robot base p

$$o_{\text{task}} = \mathcal{B}p_d = {}^{\mathcal{I}}R_B^\top (x_d - p). \quad (13)$$

The observation space is the concatenation of the robot, environment, and task observations

$$o = \left[o_{\text{robot}}^\top \quad \bar{o}_{\text{env}}^\top \quad o_{\text{task}}^\top \right]^\top, \quad (14)$$

where \bar{o}_{env} is the flattened environment observation matrix.

2) *Action Space*: The action space \mathcal{A} is composed of the desired joint positions of the robot $\theta_d \in \mathbb{R}^n$ and the desired thrust dynamics input $u = [u_1, \dots, u_m]^\top \in \mathbb{R}^m$

$$a = \left[\theta_d^\top \quad u^\top \right]^\top, \quad (15)$$

where n is the number of joints of the robot and m is the number of links on which the thrust acts. The desired joint

positions θ_d are fed to a PD controller. The resulting thrust intensities T are computed as the increment of the thrust intensity with respect to the previous time step $k-1$ as $T = T[k-1] + g(T[k-1], u)\Delta t$ and act on the specified links, generating a pure force

$$f_i = {}^{\mathcal{I}}R_i(q) \begin{bmatrix} 0 \\ 0 \\ T_i \end{bmatrix}, \forall i \in [1, m], \quad (16)$$

where ${}^{\mathcal{I}}R_i(q)$ represents the orientation of the frame i on which the force f_i is applied and Δt is the time step between two consecutive actions. $g(T[k-1], u)$ is a function that computes the thrust dynamics given the previous thrust intensities $T[k-1]$ and the input u , which in the *ideal* case is considered to be $g(T[k-1], u) = \dot{T}$, i.e., thrust intensity rate-of-change, while in the *specific* use case of the jet-powered actuation, it is a function identified from real-world data, see Sec. IV-G for more details.

C. Discriminator Observation Space

The discriminator D gives feedback about the quality of the motion produced by the agent. Selecting a meaningful set of features to feed into the discriminator is crucial since it should be able to capture the robot's motion. We choose the discriminator observation space to be the same as the robot observation vector $\chi = o_{\text{robot}}$. The discriminator is fed with couples of consecutive observations χ_t, χ_{t+1} and trained with batches of samples from the motion dataset \mathcal{D} and samples produced by the policy π_σ .

IV. RESULTS

In this section, we present the simulated results of the proposed approach. We designed our experiments to answer the following questions: (i) Can smooth multimodal locomotion be achieved with an AMP-based method? (ii) What type of information does the policy need to guide the emergence of the desired locomotion pattern? (iii) Can the same result be obtained using a classic RL approach?

Our testbench is iRonCub, a flying humanoid robot that expresses a degree of terrestrial and aerial locomotion. iRonCub has 23 joints, weights 44 kg, and is equipped with four jet engines, two fixed on the chest and two moved by the arms, which in the specific use case are the commercial JetCat P250 engines [40], capable of exerting $T_{\max} = 250$ N of thrust each.

The training consists of two main tasks. The first task is a multimodal locomotion scenario on flat ground, in which the robot has to catch several waypoints located on the ground and in the air. The second scenario consists of a world containing different terrains, in which the robot has to understand which locomotion modality is more suited. While the scenarios above are tested in the *ideal* thrust case, we also validate the approach in the *specific* use case of jet-powered actuation, in which the model of the jet engines is identified from real-world data with a neural network and embedded in the simulation environment, see Sec. IV-G.



Fig. 3: Snapshots of the flying motion obtained using TO.

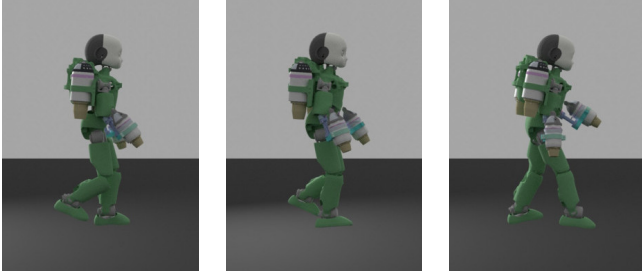


Fig. 4: Snapshots of the walking motion obtained using inverse kinematics from the CMU human walking dataset.

TABLE I: PPO parameters.

Parameter	Value
Discount rate γ	0.99
Learning rate	5e-5
GAE parameter λ	0.95
Entropy coefficient	0.0
Clip parameter	0.2
Mini-batch size	32768
Critic loss coefficient	5
KL threshold	0.008
Number of actors	4096

A. Motion Priors

The used motion data comes from two sources. (i) Walking motions are produced from recorded mocap clip [26], [41] and retargeted to the iRonCub model using inverse kinematics implemented using *iDynTree* [42] (Fig. 4). (ii) Trajectory optimization for aerial humanoid robots produces flying motions containing the full state [43]. These datasets consist of motions described in terms of the system’s kinematics, i.e., joint positions, base poses, and thrust intensities (Fig. 3) In the case of walking motions, the thrust is set to zero. If the original demonstrator is different from the agent, the motion might not be feasible. In this case, the agent will produce a motion similar to the original one but feasible for the agent.

B. Experimental Setup

The training environment is developed using the IsaacGym simulator, which allows massive parallel training. We trained 4096 agents using PPO, controlled at 60 hz. The training requires ~ 2 hours on an NVIDIA Quadro RTX 6000. The training parameters are shown in Tab. I. Fig. 5 shows the learning curves of average reward and episode duration over ten training runs for the test presented in Sec. IV-D. The




Fig. 5: Training curves over 10 runs on the scenario from Sec. IV-D.

TABLE II: Reward coefficients.

Reward term	Value
Target reward weight w_c	0.1
Velocity reward weight w_v	0.7
Facing reward weight w_f	0.2
Thrust penalty weight w_T	0.11
Target reward hyperparameter c_1	0.5
Velocity reward hyperparameter c_2	0.5

curves show a small variance, demonstrating the stability of the training. The reward coefficients are shown in Tab. II.

C. Traversing flat ground

The first scenario involves the robot walking, take-off, flying, and landing on flat ground at a desired velocity v_d of 0.8 m/s. The robot must reach several waypoints set in an interval between 0.7 m and 2.0 m. Once a waypoint is hit, a new one appears forward, and the robot moves toward it, as shown in Fig. 6. In this case, the height map is reduced to a single value, i.e., the height of the robot base w.r.t. the terrain. The test shows that the robot effectively learns terrestrial and aerial locomotion and can choose when to switch between them. Each episode is terminated when the distance of the robot base from the ground is less than 0.4 m or the maximum number of steps is reached.

D. Terrain-aware locomotion

In the second scenario, the robot deals with diverse terrains, i.e., flat and rough terrain, stepping stones, and pits at a desired velocity v_d of 0.8 m/s. As in the previous scenario, the route comprises waypoints the robot needs to catch. The policy is fed with a height map of 7×9 cells, where each cell is 0.3 m wide. Using the height map as observation, the robot learns to use small portions of reachable terrain to step on, walk when possible, and fly when the ground is not reachable, as shown in Fig. 7. The termination strategy is the same as the previous case: when the base height is smaller than 0.4 m, the episode is over.

Tuning the thrust penalization is not trivial. If the penalty is too high, the robot prefers to stay at the edge of the terrain rather than fly and reach the following reachable terrain since using the thrust leads to a higher penalty than the reward. If the penalty is too low, the robot prefers to fly rather than walk when the terrain is sparser. The penalty should be high enough to make the robot prefer walking over flying but not too high to make the robot choose to walk when flying is needed.




Fig. 6: Walk-to-fly maneuver. The robot lands and walks to catch a waypoint on the ground level, takes off, and flies to hit the aerial target.




Fig. 7: The robot walks over any reachable terrain. The height map enables *terrain-aware* multimodal locomotion.




Fig. 8: The radar chart shows the average quantities in percentage. The policy trained with only walk motion prior shows less usage of the thrust but collects less reward. The other three policies collect higher rewards, with the policy trained with both motion priors – in yellow – being efficient in thrust usage and covering a larger area.

E. Ablation Studies

We perform ablation studies to show the influence of the motion priors. We trained the policies using the same parameters but without the motion priors or part of them. We keep the thrust penalization constant and use the scenario described in Sec. IV-C as a testbench. Fig. 8 shows episode duration over the maximum duration, reward r collected over a maximum observed reward \bar{r} , and thrust usage as $1 - T/T_{max}$.

1) *No Motion Prior*: The policy trained without any motion prior does not show any particular behavior. The reward

collected is similar to the policy trained with both motion priors, but the thrust usage is higher.

2) *Only Walking Motion Prior*: The policy trained without aerial motion priors learns how to walk, but it does not show any flying behavior, preventing the robot from reaching the waypoints in the air and collecting rewards. The usage of the thrust is lower than the other policies.

3) *Only Flying Motion Prior*: The policy trained without terrestrial motion priors does not learn to walk and fly to reach all the waypoints, collecting a reward similar to the other policies. Despite the thrust usage being higher than the policy trained with both motion priors, it is lower than the policy trained without any motion prior IV-E.1, suggesting that the aerial motion prior embeds an optimized thrust usage coming from the trajectory optimization solution.

4) *Walking + Flying Motion Prior*: The policy trained with both motion priors learns how to walk and fly. The collected reward is similar to other policies, but the thrust usage is lower, indicating that switching between walking and flying is also more efficient in propulsion expenditure.

F. Comparison with Trajectory Optimization

The trajectory optimization method proposed in [43] is formulated as a multiple-shooting optimal control problem using the centroidal dynamics over an optimized horizon of 100 nodes that requires minimization of thrust usage. The contact sequence is computed using a complementarity condition formulation, and the problem is solved using the IPOPT solver [44]. The TO method is able to produce a transition from legged to aerial locomotion but requires a complex cost function and a long computation time preventing replanning capabilities. A trajectory optimization approach




Fig. 9: The plot shows the identification results obtained on the JetCat P250. The orange line is the thrust measured from a load cell in a custom test bench for jet data collection. The red line is the thrust obtained by integrating the identified LSTM model.

TABLE III: Identification errors.

Error	Value
Mean Absolute Error	4.995 N
Root Mean Squared Error	8.061 N

would need the concatenation of several offline trajectories computed for each possible scenario, stabilized afterward by an online controller, while the proposed method can adapt to the environment online. Conversely, the policy can learn how to switch between walking and flying without the need for a complex reward function while inferring the terrain information from the elevation map, which is infeasible with the compared trajectory optimization method or classical control method, needing a separate module that deals with the terrain information. Table IV compares the two approaches.

G. Use case: Jet-powered actuation

This section presents the results and considerations in the use case of jet-powered actuation. The jet engines are modeled using a Long Short-Term Memory (LSTM) neural network [45], which is trained to predict the thrust given the input and the state of the jet engine. The model is trained using the data collected from the JetCat P250, a commercial jet engine capable of exerting a thrust of $T_{max} = 250$ N [40]. The data comprises thrust measured with a load cell at different throttle inputs in an ad-hoc designed test bench [46]. The resulting discrete jet dynamics model is

$$T = T[k-1] + g(T[k-1], u)\Delta t, \quad (17)$$

where T and $T[k-1]$ are the actual and the previous thrust, u is the throttle input, g is the LSTM model and Δt is the time step. The model is validated by comparing the thrust obtained by integrating the model against the thrust measured from the load cell, as shown in Fig. 9, with performance in Table III.

The model is then embedded in the simulation environment, and the policy is trained as in scenarios described in Sec. IV-D. The policy outputs the throttle input u , which is then passed to the LSTM model (17) along with the previous thrust $T[k-1]$ to obtain the thrust T . The response dynamics

TABLE IV: Comparison between the policy and the TO method.

Feature	Method	
	Our	TO
Computation time	≈ 2 hours	≈ 30 min
Length	Episode length	7 seconds
Cost function	4 terms	12 terms
Terrain-aware	✓	✗
Online	✓	✗
Automatic Transition	✓	✓

of the jet propulsion system exhibit a comparatively slower response rate when compared to the ideal thrust case: attaining the take-off thrust requires a longer time. For this reason, the policy is trained by reducing the thrust penalization to $1e-8$. Furthermore, the minimum throttle input is set to 15%, which leads to a minimum thrust of ~ 40 N. The policy can learn how to transition between walking and flying, although the transition is slower than the ideal thrust case, e.g., when traversing the stepping stones, the robot employs the aerial locomotion mode.

V. CONCLUSIONS

Our paper presents a method that enables aerial humanoid robots to seamlessly transition between walking and flying modes. The proposed strategy leverages the concept of Adversarial Motion Priors to learn a natural gait pattern from human-like gaits and an efficient aerial locomotion pattern from motions obtained using trajectory optimization.

The robot can traverse complex terrains, switching automatically between both locomotion forms, without an explicit constraint on the form of navigation. Although our method has been tested in simulation environments and necessitates further investigations of its reliability in real-world domains, this result marks progress towards the application of aerial humanoid robots in various fields, such as search and rescue and monitoring missions, where diverse and demanding scenarios are encountered.

Looking forward, the integration of more motion priors might broaden the potential of aerial humanoid robotics in future challenges.

REFERENCES

- [1] R. Lock, S. Burgess, and R. Vaidyanathan, "Multi-modal locomotion: from animal to application," *Bioinspiration & biomimetics*, vol. 9, no. 1, 2013.
- [2] L. Daler, S. Mintchev, C. Stefanini, and D. Floreano, "A bioinspired multi-modal flying and walking robot," *Bioinspiration & biomimetics*, vol. 10, no. 1, 2015.
- [3] K. Kim, P. Spieler, E.-S. Lupu, A. Ramezani, and S.-J. Chung, "A bipedal walking robot that can fly, slackline, and skateboard," *Science Robotics*, vol. 6, no. 59, 2021.
- [4] E. Sihite, A. Kalantari, R. Nemovi, A. Ramezani, and M. Gharib, "Multi-modal mobility morphobot (m4) with appendage repurposing for locomotion plasticity enhancement," *Nature communications*, vol. 14, no. 1, 2023.
- [5] M. Bjelonic, C. D. Bellicoso, Y. de Viragh, D. Sako, F. D. Tresoldi, F. Jenelten, and M. Hutter, "Keep rollin'—whole-body motion control and planning for wheeled quadrupedal robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, 2019.
- [6] Q. Yu and N. Gravish, "Multimodal locomotion in a soft robot through hierarchical actuation," *Soft Robotics*, 2023.
- [7] J. Yu, M. Wang, W. Wang, M. Tan, and J. Zhang, "Design and control of a fish-inspired multimodal swimming robot," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011.
- [8] D. Pucci, S. Traversaro, and F. Nori, "Momentum control of an underactuated flying humanoid robot," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, 2017.
- [9] G. Nava, L. Fiorio, S. Traversaro, and D. Pucci, "Position and Attitude Control of an Underactuated Flying Humanoid Robot," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, 2018.
- [10] G. Romualdi, S. Dafarra, Y. Hu, P. Ramadoss, F. J. A. Chavez, S. Traversaro, and D. Pucci, "A benchmarking of dcm-based architectures for position, velocity and torque-controlled humanoid robots," *International Journal of Humanoid Robotics*, vol. 17, no. 01, 2020.
- [11] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, vol. 59, no. 4, 2017. [Online]. Available: <https://epubs.siam.org/page/terms>
- [12] S. Dafarra, G. Romualdi, and D. Pucci, "Dynamic complementarity conditions and whole-body trajectory optimization for humanoid robot locomotion," *IEEE Transactions on Robotics*, vol. 38, no. 6, 2022.
- [13] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, "Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot," *Autonomous robots*, vol. 40, 2016.
- [14] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," in *2003 IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 2003.
- [15] J. Engelsberger, C. Ott, and A. Albu-Schäffer, "Three-dimensional bipedal walking control using divergent component of motion," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013.
- [16] D. E. Orin, A. Goswami, and S.-H. Lee, "Centroidal dynamics of a humanoid robot," *Autonomous robots*, vol. 35, no. 2.
- [17] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014.
- [18] B. Ponton, A. Herzog, A. Del Prete, S. Schaal, and L. Righetti, "On Time Optimization of Centroidal Momentum Dynamics," in *Proceedings - IEEE International Conference on Robotics and Automation*. Institute of Electrical and Electronics Engineers Inc., 9 2018. [Online]. Available: <https://git-amd.tuebingen.mpg.de/bponton/timeoptimization.https://arxiv.org/abs/1709.09265v3>
- [19] G. Romualdi, S. Dafarra, G. L'Erario, I. Sorrentino, S. Traversaro, and D. Pucci, "Online non-linear centroidal mpc for humanoid robot locomotion with step adjustment," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022.
- [20] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [21] V. Tsounis, M. Alge, J. Lee, F. Farshidian, and M. Hutter, "Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, 2020.
- [22] Y. Song, A. Romero, M. Mueller, V. Koltun, and D. Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, 2023.
- [23] D. Rodriguez and S. Behnke, "Deepwalk: Omnidirectional bipedal gait by deep reinforcement learning," in *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2021.
- [24] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for robust parameterized locomotion control of bipedal robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [25] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone-racing with deep reinforcement learning," *Nature*, 2023.
- [26] X. B. Peng, P. Abbeel, S. Levine, and M. Van de Panne, "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Transactions On Graphics (TOG)*, vol. 37, no. 4, 2018.
- [27] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa, "Amp: Adversarial motion priors for stylized physics-based character control," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, 2021.
- [28] X. B. Peng, Y. Guo, L. Halper, S. Levine, and S. Fidler, "Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters," *ACM Transactions On Graphics (TOG)*, vol. 41, no. 4, 2022.
- [29] S. Starke, Y. Zhao, T. Komura, and K. Zaman, "Local motion phases for learning multi-contact character movements," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 4, 2020.
- [30] S. Starke, I. Mason, and T. Komura, "Deepphase: Periodic autoencoders for learning motion phase manifolds," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, 2022.
- [31] A. Escontrela, X. B. Peng, W. Yu, T. Zhang, A. Iscen, K. Goldberg, and P. Abbeel, "Adversarial motion priors make good substitutes for complex reward functions," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022.
- [32] X. B. Peng, E. Coumans, T. Zhang, T.-W. E. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," in *Robotics: Science and Systems*, 07 2020.
- [33] P. M. Viceconte, R. Camoriano, G. Romualdi, D. Ferigo, S. Dafarra, S. Traversaro, G. Oriolo, L. Rosasco, and D. Pucci, "Adherent: Learning human-like trajectory generators for whole-body control of humanoid robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, 2022.
- [34] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, 2016.
- [35] V. Makoviychuk, et al., "Isaac gym: High performance gpu-based physics simulation for robot learning," 2021.
- [36] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022.
- [37] S. Traversaro, "Modelling, estimation and identification of humanoid robots dynamics," Ph.D. dissertation, Italian Institute of Technology Genoa, Italy, 2017.
- [38] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems*, vol. 27, 2014.
- [39] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for gans do actually converge?" in *International conference on machine learning*. PMLR, 2018.
- [40] "JetCatP250." [Online]. Available: www.jetcat.de/en/productdetails/produkte/jetcat/produkte/Professionell/p250%20pro%20s
- [41] C. M. University, "Cmu graphics lab motion capture database," <http://mocap.cs.cmu.edu/>, accessed: 2023-06-07.
- [42] F. Nori, S. Traversaro, J. Eljaik, F. Romano, A. Del Prete, and D. Pucci, "icub whole-body control through force regulation on rigid noncoplanar contacts," *Frontiers in Robotics and AI*, vol. 2, no. 6, 2015. [Online]. Available: <http://www.frontiersin.org/humanoid-robotics/10.3389/frobt.2015.00006/abstract>
- [43] G. L'Erario, G. Nava, G. Romualdi, F. Bergonti, V. Razza, S. Dafarra, and D. Pucci, "Whole-body trajectory optimization for robot multimodal locomotion," in *2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, 2022.
- [44] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, 2006.
- [45] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, 1997.
- [46] G. L'Erario, L. Fiorio, G. Nava, F. Bergonti, H. A. O. Mohamed, E. Benenati, S. Traversaro, and D. Pucci, "Modeling, identification and control of model jet engines for jet powered robotics," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, 2020.