Dream to Fly: Model-Based Reinforcement Learning for Vision-Based Drone Flight

Angel Romero*, Ashwin Shenai*, Ismail Geles, Elie Aljalbout, Davide Scaramuzza

Robotics and Perception Group, University of Zurich, Switzerland

*These authors contributed equally

Abstract—Autonomous drone racing has risen as a challenging robotic benchmark for testing the limits of learning, perception, planning, and control. Expert human pilots are able to fly a drone through a race track by mapping the real-time feed from a single onboard camera directly to control commands. Recent works in autonomous drone racing attempting direct pixel-to-commands control policies, even without explicit state estimation, have relied on either intermediate representations that simplify the observation space or performed extensive bootstrapping using Imitation Learning (IL). This paper presents an approach for training quadrotor navigation policies from scratch — mapping raw onboard camera pixels directly to control commands, much like a human pilot. While model-free methods such as PPO are sample-inefficient and struggle in this setting, we leverage Model-Based Reinforcement Learning (MBRL) to train visuomotor policies capable of agile flight through a racetrack using only raw pixel observations. Moreover, because our policies are trained end-to-end directly from pixels, we no longer require the perception-aware reward term used in previous methods. Instead, we show that this behavior naturally emerges, resulting in policies that guide the camera toward feature-rich areas of the observation space. Our experiments demonstrate in both simulation and real-world flight how the proposed approach can be deployed on agile quadrotors. Our findings emphasize the benefits of training directly from pixel observations and demonstrate that MBRL offers a promising path for real-world robotics research.

I. INTRODUCTION

In recent years, quadrotors have become a central focus of robotics research, emerging as versatile platforms with untapped potential across multiple domains, including search and rescue, inspection, agriculture, cinematography, delivery, passenger air vehicles, and space exploration [1]. Among these applications, the field of autonomous drone racing has attracted particular interest in the research community [2], showcasing the remarkable progress in agile flight control. This domain not only benefits from cutting-edge robotics research [3] but also pushes the limits of what is possible by challenging these flying machines to outperform the most skilled human pilots, as shown by recent successes against the world's best pilots [4, 5]. The ability to navigate complex race tracks at high speeds, making split-second decisions, not only pushes the boundaries of what is possible in autonomous flight, but also stresses the capabilities of machine learning, perception algorithms, planning and control systems, requiring them to perform at peak efficiency under

Fig. 1. Real-world deployment of our DreamerV3 policy in the Figure 8 track. During training, the agent learns a world model from interactions with the environment. At the same time, the actor-critic policy is trained by sampling the predictions of the world model, also called *imagination*. The onboard images consumed by the network are marked in red.

highly dynamic and demanding conditions.

Traditionally, most autonomous drone racing systems have relied on explicit state estimation integrating data from inertial measurement units (IMUs) and other onboard sensors to maintain stability and optimize performance [6, 7, 8, 9, 4]. However, professional human pilots rely solely on visual feedback from an onboard camera, showcasing a remarkable ability to navigate complex environments only from visual input extracted from an unique onboard camera. Emulating this human ability to fly based solely on visual information remains a significant challenge for autonomous systems.

Closing the loop between perception and control – learning directly from pixels to actions without the need for explicit state estimation – remains largely unfulfilled. While reinforcement learning (RL) has shown promise in various robotic applications, applying it to vision-based tasks introduces unique difficulties, especially in real-time robotic systems such as drone racing.

One of the most recent works in this domain [10] manages to train a Proximal Policy Optimization (PPO) [11] policy to fly a drone through a race track from binary image representations. In this case, the visual inputs are first preprocessed and distilled into a simpler intermediate representation in the form

of a binary mask where only the racing gates are visible. This intermediate representation reduces the overall observation space complexity and allows the PPO policy to learn the behavior efficiently. However, this work relies on a simplified observation space: the raw visual input is preprocessed into a binary mask highlighting only the gates. This simplification, while reducing the complexity of the observation space and facilitating learning, also discards valuable information. For instance, excluding background information can hinder the agent's ability to navigate when no gates are immediately visible, or to obtain the gravity direction from the horizon line, for example. In [12] the authors present another recent work that tackles quadrotor flight from pixel observations. However, it relies on IL bootstrapped from a pre-trained expert policy, which had access to the full state and introduces a strong dependency on privileged information. Despite their success, one major challenge in robot learning is the need for large amounts of physical interaction data, which can be very expensive and challenging to obtain. This challenge is compounded in vision-based control tasks, where the high dimensionality of image observations further increases data requirements. In particular, learning directly from pixels has a combined effect with the sample inefficiency of model-free RL methods, which typically demand extensive interactions with the environment.

MBRL presents a promising solution to these challenges because it is generally more sample efficient than its model-free counterpart, reducing the need for extensive environment interaction. MBRL learns the transition dynamics of the system, called the world model, and uses it to predict and optimize future actions. It also generalizes well to different tasks as the transition model does not vary vastly across different tasks in the same system description. However, this approach typically comes at the cost of longer training times, as the system must simultaneously learn a world model and optimize a control policy. Additionally, the inherent complexity of model-based RL architectures, with their multiple interdependent components, often makes them more difficult to tune and optimize.

Recent works in MBRL, such as the DreamerV3 architecture [13], have helped mitigate these issues, offering a model-based RL approach that is more accessible in terms of tuning and optimization. Despite its successes in controlled environments, however, the DreamerV3 framework has so far seen limited application in real-world robotic systems, especially when it comes to real-world robotics tasks involving direct pixel-to-action learning.

This work aims to address these challenges by presenting an approach for vision-based autonomous quadrotor flight, utilizing MBRL to directly map pixel observations to control commands. Using DreamerV3, we train a policy end-to-end from raw pixel observations to collective thrust and body-rate commands, mirroring how a human pilot controls a drone. In contrast to PPO—which, as we show, cannot learn effectively under these conditions—DreamerV3 rapidly acquires agile drone racing skills in simulation and in the real-world.

Contributions

This paper introduces the following contributions:

- To the best of our knowledge, this is the first RL approach that is able to learn to race directly from pixels from scratch, closing the loop between perception and action.
- We show that our agents exhibit emergent behavior that naturally guides the drone camera view towards feature rich areas, such as the next gates, without requiring handcrafted shaped reward terms to incentivize gate-viewing as done in previous work.
- We deploy our system in both simulation and real-world, further validating the applicability of MBRL to real-world mobile robots.

II. RELATED WORK

A. Reinforcement Learning from Pixels

In recent years, RL algorithms have achieved incredible feats in simple environments learning directly from pixels, such as arcade games, used as sample efficient simulators to benchmark the capabilities of AI agents [14]. This success in RL has been repeatedly shown in environments in growing complexity, such as beating human performance in Atari games [15] or in more complex games such as Go [16]. However, these successes have been primarily in environments where the action space is discrete, and not in continous action spaces. This gap is addressed by the DeepMind Control Suite [17], which presents a diverse set of environments with different observation spaces and different action modalities — both continuous and discrete. However, learning directly from pixels in these scenarios has been shown to be less sample efficient, yielding suboptimal performance when compared to state-based learning, even in simulation environments. Recent efforts have tried to improve this by building on top of intermediate visual representations [18, 19, 20, 21, 10]. For example, previous work explored different kinds of state representations via various supervised and self-supervised learning methods, such as auto-encoders [18], future predictions [20], contrastive unsupervised representations [21] or semantic segmentation [10], with the objective of reducing the performance difference between state-based and pixel-based RL.

In order to address the challenge of learning complex behaviors efficiently directly from pixels, some methods have focused on data augmentation [22, 23]. Policies that map camera observations directly to commands — also called sensorimotor or visuomotor policies — are mostly learned through extensive data simulation, making use of domain randomization techniques or incorporating additional privileged information such as joint angles [24, 25, 26] For more complicated tasks, the robotics community has resorted to the usage of expert demonstrations, marking a shift towards imitation learning approaches, contrary to the learning from scratch using RL paradigm. This is showcased by recent works, most notably [27].

B. Vision-based flight

In the last decade, a plethora of works that aim to fly a drone from images to high level setpoints has emerged [7, 28, 29, 30, 31]. These approaches are not using reinforcement learning, and try to tackle the navigation task in the more general sense, therefore far from agile flight. In [32] the authors map features and IMU to control commands to fly aggressive maneuvers, learning via imitation learning. In the realm of state-based policy learning, using RL for the low-level control of quadrotors has been demonstrated to perform better than classical methods [33, 34]. Particularly, in state-based drone racing, RL algorithms have been able to demonstrate agile flight [35, 36], surpassing even state-of-the-art optimal control methods [5]. The adaptability and versatility of RL methods has been showing more and more success in vision-based flight in several recent works. One of the earliest uses of RL for vision-based quadrotor navigation is [37], where a CAD model of the environment was used to train an RL policy that would output discrete velocity commands ('forward', 'right', 'left'). This was enabled thanks to the onboard stabilizing controller and explicit state estimation through a VIO pipeline.

More recently and for the first time, a quadrotor that combines VIO with an RL controller has been able to win a fair drone race against the top human world champions [4]. Even though this represents a landmark achievement for robotics, the drone is still a custom built research drone [38], which runs a state-estimation pipeline on board. On top of this, the policies needed to be fine-tuned by using real-world data captured with a motion-capture system. Therefore, the observation modality is still not the same that the humans have, which is exclusively a single monocular feed from the onboard camera of the platform. In a more recent work [10], the authors were able for the first time to achieve agile flight without explicit state estimation. The authors convert the pixel image into a black and white image that contains only the gates, and from this intermediate representation they train a model-free RL policy from scratch. Even though in that case they were able to fly through the gates, they still were not able to train directly from the camera feed, and the policy is not aware of any of the texture or background information.

In this work, we aim to overcome this limitation by using model-based RL (DreamerV3 [13]) to train a policy that can map directly RGB pixels to control commands, just like humans do.

C. Model-based Reinforcement Learning for Real-World Robots

As mentioned above, there are many approaches that have attempted to train policies using pixels as observations, relying in simplifying assumptions. However, the landscape of model-based RL (MBRL) methods applied to real-world robotics remains relatively limited, and existing research primarily focuses on simulation environments. Within the context of real-world robotic applications, prior MBRL work often relies on state estimations rather than raw pixel observations. Specifically, there are some works tackling the problem of flying a quadcopter using model-based RL, for high-level control [39] or low-level controller learning [34]. However, these works have been using state estimation as observations, and not directly pixels as observations.

When speaking about works that use MBRL from pixels and that deploy in real-world robots, the number of works is even lower. One of the early works in this direction is [40], where a learned world model of the dynamics is built from pixels, and is then used by an MPC planner to get the optimal policy, which outputs torque commands for a real-world manipulator. Another interesting work that applies DreamerV3 to a real-world robotic system is [41], where the table top labyrinth game is solved in the real world by using top-view fixed camera in conjunction with the measurement of the table angle. However, this work does not belong to the mobile robotics category. One of the most salient works that applies DreamerV3 [13] to real-world robotics is [42], where different policies are trained from images, depth and state, for different tasks. These include pick-and-place manipulation utilizing a fusion of RGB images, depth data, and proprioceptive sensing, as well as quadruped locomotion trained solely on proprioceptive sensor data. Notably, the only instance of direct learning from pixels to commands involves a 2D navigation task with a Sphero robot, a relatively simple scenario in terms of dynamics and environmental complexity.

III. METHODOLOGY

A. Problem Statement

Our objective is to navigate a quadrotor through a series of gates in minimum time while avoiding collisions with the environment and the gates. Formally, we seek a policy $\pi(x)$ that maps raw visual observations x directly to control commands a, minimizing the time required to complete the course. We address this challenge using reinforcement learning (RL), framing the problem within the Markov Decision Process (MDP) framework.

B. MDP formulation

The Markov Decision Process (MDP) formalism provides a structured approach for modeling decision-making problems. An MDP is defined by the tuple $(\mathcal{X}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{X}$ is the set of possible states, $\mathcal{A}$ is the set of possible actions, $P(x_{k+1}|x_k, a_k)$ is the state transition probability, $R(x_k, a_k)$ is the reward function, and $\gamma \in [0, 1)$ is the discount factor. In our vision-based drone racing task, the state x_k is represented by the raw image captured by the onboard camera at time step k. The action a_k consists of the collective thrust and body rates applied to the quadrotor. The reward function is designed to incentivize fast and collision-free navigation through the gates (detailed in Section III-D) The objective in an MDP is to find an optimal policy $\pi_\theta^* : \mathcal{X} \to \mathcal{A}$ that maximizes the expected cumulative discounted reward by optimizing the parameters θ of a neural network,

$$\pi_\theta^* = \arg\max_\pi \mathbf{E} \left[\sum_{k=0}^{\infty} \gamma^k r_k \right].$$
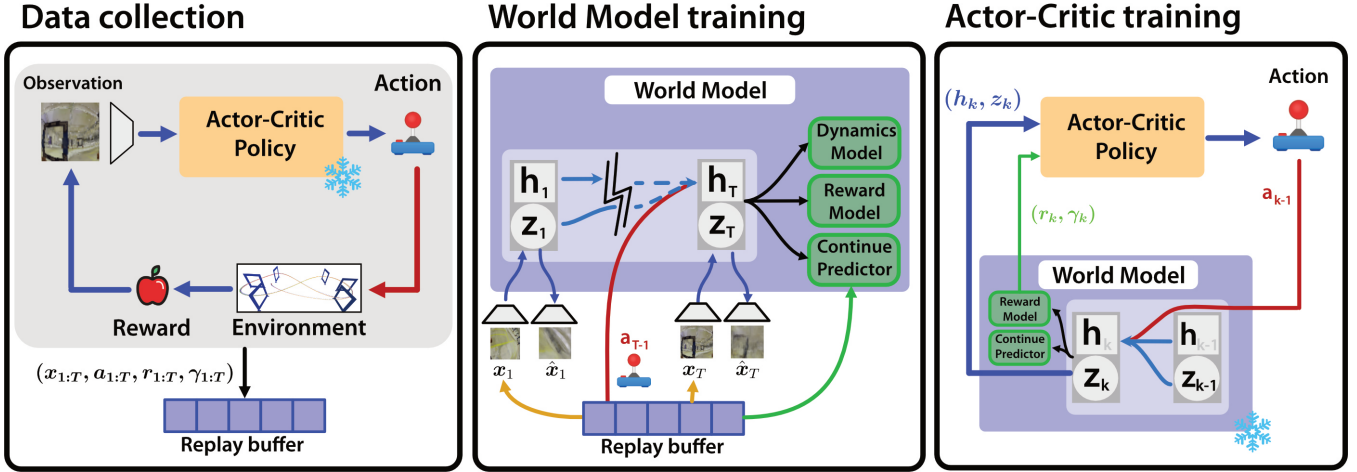
Fig. 2. The process begins with data collection in the simulation environment using the current policy, storing experiences in a replay buffer. This buffer is used to train the world model components: the encoder, decoder, RSSM, dynamics model, reward model, and continue predictor (Section III-E). Subsequently, an actor-critic policy is trained within the learned world model to maximize expected (imagined) returns. This updated policy is then used to collect new data, restarting the loop.

This optimization considers both immediate and future rewards, with future rewards discounted by γ.

C. Observation and Action Spaces

The RL framework learns a visuomotor, end-to-end policy that directly maps raw RGB images as observations to control inputs. These images provide a rich and high-dimensional representation of the environment, enabling the agent to infer its state and surroundings without relying on explicitly estimated states.

1) Observation Space: The observation space consists of RGB images captured at each time step, denoted as $\mathbf{x}_k \in \mathbb{R}^{H \times W \times 3}$, where H and W represent the image height and width, respectively, and 3 corresponds to the RGB color channels. These images are normalized to the range $[0, 1]$ by dividing pixel values by 255. Unlike approaches that rely on explicit state estimation (e.g., position, velocity, or attitude), this image-based approach directly leverages rich visual information to guide the policy.

2) Action Space: At each time step k, the policy outputs a four-dimensional action vector $\mathbf{a}_k = [c, \omega_x, \omega_y, \omega_z]$, where c represents the mass-normalized collective thrust, and $\omega_x, \omega_y, \omega_z$ are the body rate setpoints in the drone's body frame. These actions are expressed in the Collective Thrust and Body Rates (CTBR) format, a control interface commonly used by professional drone pilots [4, 43]. This format directly commands the low-level actuation of the drone, bypassing the need for intermediate, high-level abstractions such as position or velocity commands. To ensure bounded and physically realizable control actions, the action space is constrained to $\mathcal{A} = [-1, 1]^4$, which are then mapped to the actual collective thrust and body rates limits. This constraint is enforced by applying a hyperbolic tangent (tanh) activation function to the output of the actor policy network.

D. Reward Function

The racing track is defined as a sequence of linearly connected waypoints placed at the center of the gates. The reward function is designed to incentivize progress along this track while penalizing undesirable behaviors. The reward at time step k, $r(k)$, is defined as

$$
r_k = \begin{cases} r_{collision}, & \text{if collision;} \\ r_{passed}, & \text{if gate passed;} \\ b_1(\|\mathbf{g}_k - \mathbf{p}_{k-1}\| - \|\mathbf{g}_k - \mathbf{p}_k\|) \\ \quad - b_2\|\boldsymbol{\omega}_k\|, & \text{otherwise,} \end{cases} \tag{1}
$$

where $\mathbf{g}_k$ is the center of the target gate, $\mathbf{p}_k$ and $\mathbf{p}_{k-1}$ are the drone's positions at the current and previous time steps, respectively, and $\boldsymbol{\omega}_k$ is the body rate vector. The primary component of the reward, $b_1(\|\mathbf{g}_k - \mathbf{p}_{k-1}\| - \|\mathbf{g}_k - \mathbf{p}_k\|)$, is called the progress term, as it encourages the drone to move closer to the target gate. The term $b_2\|\boldsymbol{\omega}_k\|$ penalizes excessive body rates. We use the coefficients $b_1 = 1.0, b_2 = 0.01$. This formulation directly rewards progress toward the gate center. Importantly, deviations from the exact path defined by the waypoints are not penalized, allowing the agent to discover potentially more efficient trajectories.

In addition to the continuous reward component, discrete rewards are provided for specific events. A large negative reward $r_{collision} = -4.0$ is given upon collision with the environment or a gate, terminating the episode. A positive reward $r_{passed} = +10.0$ is awarded when the drone successfully passes through a gate, triggering the transition to the next target gate.

E. Model-based Reinforcement Learning: DreamerV3

In this section, we briefly describe the key aspects of the DreamerV3 algorithm. For more details, we refer the reader to the original paper [13]. DreamerV3 is an off-policy,

model-based reinforcement learning algorithm. The algorithm is structured in two main blocks: the **world model**, and the **actor-critic policy**. These two blocks are trained in an alternating fashion using an experience replay buffer while the agent interacts with the environment. A high level depiction of the training process is described in Fig. 2.

1) World Model training: Our approach employs a world model that captures the state transition dynamics in a compact latent space. Specifically, we consider a latent state $s_k = (h_k, z_k)$ and model the state transition probability $P(s_{k+1} \mid s_k, a_k)$. By encoding high-dimensional sensory inputs into low-dimensional representations, the world model enables predicting future latent states and rewards based on the agent's actions. The world model is implemented as a Recurrent State Space Model (RSSM) [44], and consists of the following components:

- **Encoder:** An encoder network maps raw observations x_k into a stochastic latent representation z_k. This provides a compressed representation of the sensory observations.

$$z_k \sim q_\phi(z_k \mid h_k, x_k). \tag{2}$$

- **Recurrent Sequence Model:** A recurrent sequence model, parameterized by the recurrent state h_k, predicts the evolution of the latent representation. Given the previous latent state $[h_{k-1}, z_{k-1}]$ and the action a_{k-1}, it predicts the current recurrent state h_k.

$$h_k = f_\phi(h_{k-1}, z_{k-1}, a_{k-1}). \tag{3}$$

- **Dynamics, Reward and Continue Prediction:** The dynamics predictor predicts the stochastic state $\hat{z}_k$ given the recurrent state h_k. The reward and continue predictors are conditioned on the latent state $s_k = (h_k, z_k)$, and predict the immediate reward r_k and the episode continuation flag $c_k \in \{0, 1\}$, indicating whether the episode terminates or continues.

$$\text{Dynamics predictor:} \quad \hat{z}_k \quad \sim \quad p_\phi(\hat{z}_k \mid h_k) \tag{4}$$
$$\text{Reward predictor:} \quad \hat{r}_k \quad \sim \quad p_\phi(\hat{r}_k \mid h_k, z_k) \tag{5}$$
$$\text{Continue predictor:} \quad \hat{c}_k \quad \sim \quad p_\phi(\hat{c}_k \mid h_k, z_k). \tag{6}$$

- **Decoder:** A decoder reconstructs the original observations from the latent representations. This reconstruction loss ensures that the latent variables z_k retain essential information from the environment.

$$\hat{x}_k \sim p_\phi(\hat{x}_k \mid h_k, z_k). \tag{7}$$

The encoder and decoder use convolutional neural networks (CNNs) for image inputs and multi-layer perceptrons (MLPs) for vector inputs. These models are trained by minimizing different losses: the prediction loss $\mathcal{L}_{pred}(\phi)$, which trains the decoder, the reward and the continue flags; the dynamics loss $\mathcal{L}_{dyn}(\phi)$, which trains the sequence model; and the representation loss $\mathcal{L}_{rep}(\phi)$, aims to make the representations more predictable. For more details about these losses, we refer the reader to [13].

The overall world model objective is a linear combination of the above defined losses:

$$\mathcal{L}(\phi) = \mathbf{E}_{q_\phi} \left[\sum_{k=1}^{T} \beta_{pred} \mathcal{L}_{pred}(\phi) + \beta_{dyn} \mathcal{L}_{dyn}(\phi) + \beta_{rep} \mathcal{L}_{rep}(\phi) \right],$$

where $\beta_{pred} = 1, \beta_{dyn} = 1, \beta_{rep} = 0.1$ and T is the batch sequence length. The world model is trained by randomly sampling T-length snippets of inputs $x_{1:T}$, actions $a_{1:T}$, rewards $r_{1:T}$ and continuation flags $c_{1:T}$ from episodes in the replay buffer.

By jointly training these components, the RSSM-based world model learns a compact, predictive representation of the environment, and a latent dynamics model, supporting more efficient decision-making and planning in latent space.

2) Actor-Critic Training: Our approach employs an actor-critic architecture trained using imagined trajectories generated by a learned world model (see Fig. 2). This allows the agent to learn complex behaviors without requiring extensive real-world interactions. Given a starting state representation, the actor generates an imagined trajectory consisting of model states $(h_{1:T}, z_{1:T})$, actions $a_{1:T}$, rewards $r_{1:T}$, and continuation flags $c_{1:T}$. The critic then learns to evaluate the quality of these imagined trajectories by predicting the distribution of bootstrapped λ-returns. This bootstrapping allows the critic to estimate long-term returns even within the limited prediction horizon of the world model.

The actor learns to maximize these λ-returns while exploring through an entropy regularizer. To ensure robust exploration across diverse reward scales and frequencies, the returns are normalized to approximately lie in $[0, 1]$. To optimize the policy, DreamerV3 uses the REINFORCE estimator combining unbiased but high-variance policy gradients with a stop-gradient operation on the value targets.

For further details on the actor-critic training and loss functions, we refer the reader to [13]. In our case, the prediction horizon for the imagined trajectories is set to $T = 16$. This horizon balances computational cost and the ability to capture longer-term dependencies.

F. Quadrotor Dynamics

In this section, we describe the nominal dynamics $\vec{f}(\vec{s}, \vec{u})$, where $\vec{s} \in \mathbb{R}^{10}$ is the state of the quadrotor and $\vec{u} \in \mathbb{R}^4$ is the input to the system. The state of the quadrotor is given by $\vec{s} = [\vec{p}_I, \vec{q}_{IB}, \vec{v}_I]^T$, where $\vec{p}_I \in \mathbb{R}^3$ is the position, $\vec{q}_{IB} \in \mathbb{SO}(3)$ is the unit quaternion that describes the rotation from the body to the inertial frame, and $\vec{v}_I \in \mathbb{R}^3$ is the linear velocity vector. $\vec{\omega}_B \in \mathbb{R}^3$ are the bodyrates in the body frame B. For ease of readability, we drop the frame indices, as they remain consistent throughout the description. The nominal dynamic equations are given by:

$$\dot{\vec{p}} = \vec{v} \qquad \dot{\vec{v}} = \vec{g} + \frac{\mathbf{R}(\vec{q})\vec{f}_T}{m}$$
$$\dot{\vec{q}} = \frac{\vec{q}}{2} \odot [0 \quad \vec{\omega}]^T, \tag{8}$$

where $\odot$ represents the Hamilton quaternion multiplication, $\mathbf{R}(\vec{q})$ the quaternion rotation, m the quadrotor's mass, $\mathbf{J}$ the
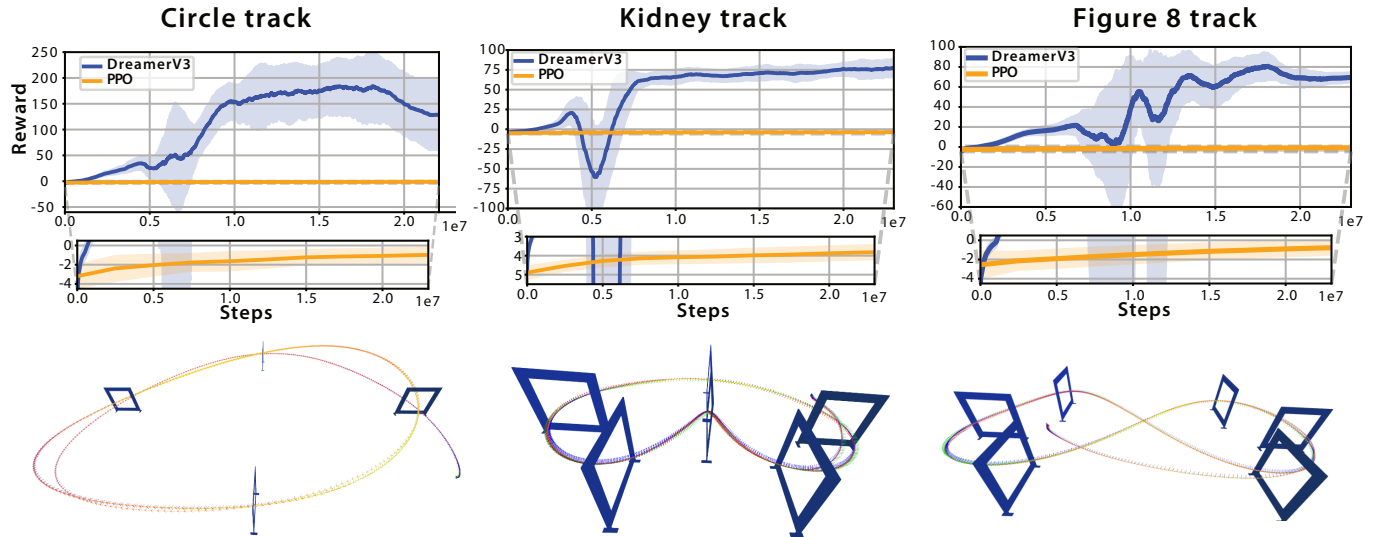
Fig. 3. Reward evolution by number of steps for three different tracks: Circle track, Kidney Track and Figure 8 track. The training performance of DreamerV3 is shown in blue, and for PPO in orange. In the zoomed, lower figure, we show that PPO is not able to achieve any considerable training in 10 million environment interactions, while DreamerV3 is able to train to convergence for the three tracks.

quadrotor's inertia, and $\vec{f}_T$ the collective thrust. The input space is given in the collective thrust and body rates modality, hence $\vec{u} = [\vec{f}_T, \vec{\omega}]$.
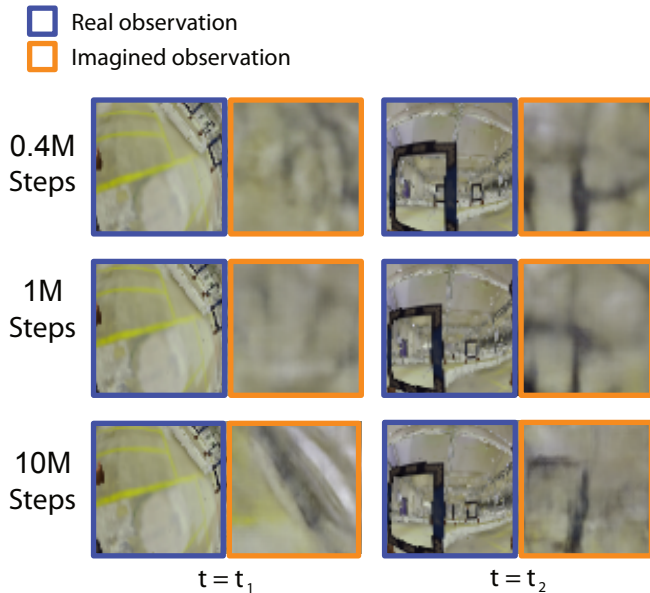

Fig. 4. Comparison of real observations and imagined observations for the Figure 8 track. Imagined observations are observations that are reconstructed by the world model. The figure shows the reconstructed observations after 0.4M steps (early stage training), 1M timesteps (mid stage training), and after 10M steps (training convergence). One can observe how the reconstruction gets better and better as the training evolves.

IV. SIMULATION EXPERIMENTS

We implemented DreamerV3 in PyTorch, leveraging the dreamerv3-torch open-source implementation and building on

top of the stable-baselines3 library [45]. Our experiments were conducted within a high-fidelity simulation environment combining Flightmare [46] and Agilicious [38] for realistic quadrotor dynamics and track generation. To enable fast, real-time rendering and direct access to the image feed during training, we integrated the Habitat simulator [47, 48], enabling training with the renderer in the loop at several thousand frames per second. This same training environment has been successfully employed in prior research [12].

Our method's performance is benchmarked against a model-free Proximal Policy Optimization (PPO) baseline [49, 5, 4, 10]. The PPO baseline used a CNN followed by four-layer Multi-Layer Perceptrons (MLPs) with 768 neurons per layer for both the actor and critic networks. The simulated quadrotor dynamics are described in Section III-F, and its physical parameters were consistent across all experiments: mass $m = 0.6$ kg, diagonal inertia matrix $J = \text{diag}([0.002410, 0.001800, 0.003759])$ kg m^2, rotor torque constant $\kappa = 0.022$, and arm length 0.14 m. The maximum rotor thrust was limited to 4.0 N, resulting in a thrust-to-weight ratio of 2.7. The dynamics, platform and parameters are the same one as in [10].

For DreamerV3, we adopted the hyperparameters detailed in [13] and used the *Large (L)* network configuration described in Appendix B of [50]. This configuration consists of four-layer MLPs with 768 units for the decoder, predictors, actor, and critic networks, and 2048 recurrent units for the world model's recurrent component.

All experiments were performed on a single Quadro RTX 8000. Input images from the simulated camera were resized to 64×64 RGB pixels before being fed into the DreamerV3 agent. These input images are visualized in Figure 4.

A. Results

We conducted training experiments on our agent using two distinct reinforcement learning algorithms: DreamerV3 and PPO. These experiments were carried out across three challenging drone racing tracks: Circle, Kidney, and Figure 8. The results of these experiments are visualized in Figure 3. The top row of the figure presents the overall reward evolution for each track, comparing the performance of both DreamerV3 and PPO. The bottom row provides a zoomed-in view of the PPO reward curves, for a more detailed examination of their evolution. Each run was conducted with 5 different random seeds. The resulting plots depict the average reward across these 5 runs, represented by a solid center line, and the shaded area surrounding the line indicates the standard deviation. Our findings indicate that DreamerV3, by leveraging a learned world model, effectively learns directly from pixel observations. This enables the agent to acquire policies that are capable of racing through each track. In contrast, PPO, which relies solely on direct interaction with the environment, struggles to converge to high-reward solutions. Its policies often lead to policies that fail to execute any meaningful flight maneuvers, resulting in consistently low reward values. An important observation is that our system is trained using curriculum learning, where the parameter b_2 in Eq. (1) is initially set to 0.0 and gradually increases once the total reward surpasses 50.0. This explains the oscillatory reward values observed in Fig. 3, as the system adjusts to the changing rewards.

This aligns with the findings presented in [12], where the authors encountered similar challenges in training end-to-end PPO policies directly from pixel observations. To overcome these difficulties, they adopted a hybrid approach, combining imitation learning with subsequent reinforcement learning fine-tuning.

As illustrated in Figure 2, the DreamerV3 architecture incorporates an observation reconstruction module. This component ensures that the latent state captures sufficient information to accurately represent the observed sensory input. Figure 4 presents a comparative analysis of real observations and their corresponding reconstructions for the Figure 8 track at two distinct time steps, t_1 and t_2. The reconstructions are sampled at three different training stages: early training (0.4 million steps), mid-training (1 million steps), and late-stage training (10 million steps). A visual inspection of the figure reveals a progressive improvement in reconstruction quality as training progresses. Notably, the fine details of the environment, such as the yellow floor lines at time step t_1, become identifiable only in the late-stage reconstructions. Similarly, the shape of the gate at time step t_2 is accurately reconstructed by the model only at late-stage training.

B. Perception-aware emergent behaviour

An additional interesting observation is that after training, the agent exhibits a consistent tendency to orient the camera towards the gates, a behavior that naturally emerged during training rather than being explicitly incentivized by the reward
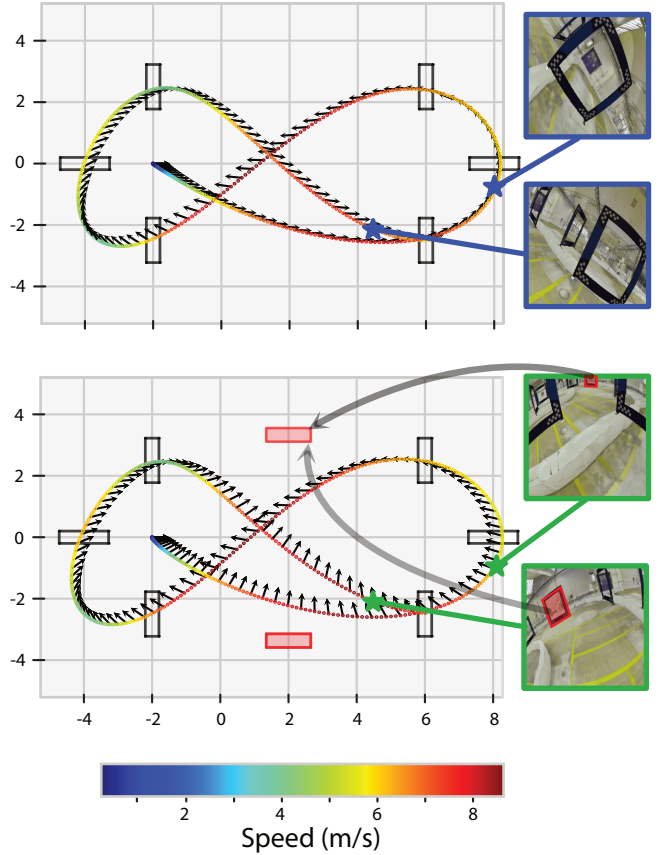

Fig. 5. **Ablation study on the perception aware behaviour of our policies.** **Top:** DreamerV3 policy trained on pixel observations in an environment where the only rendered gates are the actual gates. As indicated by the black arrows (representing camera direction), the platform predominantly focuses on the next gate. **Bottom:** We introduce two additional gates to the rendering engine (marked in red color). These gates are not required to be traversed but serve as a valuable source of information for platform localization. Consequently, the policy's behavior shifts, and the platform now distributes its camera view more evenly across both the actual and the extra gates.

function, as it is generally done in previous works [10, 12]. This emergent behavior can be attributed to the fact that we are optimizing end-to-end from pixels to commands, therefore allowing for closing the action-perception loop. As depicted in Figure 4, the gates remain visually clear throughout the entire track, while the background details become increasingly blurred due to the downsampling of the input image to 64×64 pixels. This suggests that the policy strategically prioritizes focusing on the information-rich gates, which are essential for successful navigation. To reinforce this hypothesis, we have conducted an additional ablation experiment where we place two additional gates in the periphery of the Figure 8 track. These gates are there only in the rendering, but do not need to be passed through. Fig. 5 shows the same policy trained with and without these additional gates. In the top part of Fig. 5, we show the policy trained without the extra gates. As indicated by the black arrows (representing camera view direction), the platform predominantly focuses on the following gates. In the

bottom part of the figure, we show the policy trained with the two additional gates to the rendering engine (marked in red color). As it can be seen, the policy's behavior changes, and the platform now distributes its camera view more evenly across both the actual and the extra gates. In the supplementary video we show the first person view of the two policies depicted in Fig. 5 deployed in a simulated environment.

V. REAL-WORLD EXPERIMENTS

Fig. 6. Real-world experimental setup. Our drone is equipped with a RF receiver similar to that used by professional human pilots. The observations are onboard RGB images, similar to those seen by a human pilot.

A. Setup

The software setup is identical to the one used in the simulation experiments, explained in Section IV. Regarding the hardware, we use a modification of the *Agilicious* platform [38] for the real-world deployment. We have replaced the onboard computer with an RF receiver, which is connected directly to the flight controller[1] and takes care of parsing the collective thrust and bodyrate commands from the offboard computer. Our hardware setup is the same setup as in [10]. This configuration is similar to the one used by professional drone racing pilots. This quadrotor and racing setup is shown in Fig. 6. For the deployment in the real world, we use a hardware-in-the-loop (HIL) setup, where the images are obtained from the habitat simulator and fed into the network, and the commands produced are sent directly to the real drone platform. This way, we have the real world dynamics in the loop, which allows us to asses the sim-to-real gap and our policy performance when tested in the real system.

B. Results

We demonstrate the efficacy of our policies by deploying them in the real world on the Figure 8 track. Figure 7 presents a comparative analysis of the simulated and real-world trajectories and speed profiles for this track. A strong similarity is observed between the two, indicating a minimal

[1] https://www.betaflight.com

a) Simulation

b) Real-world

Fig. 7. Real-world deployment of the trained policy for the Figure 8 track. We show the deployment in simulation (top) and in the real-world (bottom). By looking at the 3D trajectories and the speed profile, we note that our policies transfer and result in a small sim-to-real gap.

sim-to-real gap in our dynamic model. Moreover, the camera axis, visualized by black arrows in Figure 7, aligns closely between the simulated and real-world scenarios. As already mentioned in Section IV-B, both the simulated and deployed policies show the same perception-aware behavior: they keep the camera view aligned with the next gate, even if there are no reward terms guiding or incentivizing it. To the best of our knowledge, our real-world demonstration marks the first RL approach to learn drone racing directly from pixel inputs to control commands, effectively closing the loop between perception and action. The supplementary video shows the deployment of our system in the real world.

VI. LIMITATIONS

The proposed model-based RL approach has shown great success at learning complex behaviours from pixels to commands at agile flight, also demonstrating successful real-world deployment. However, there are some limitations, mainly related to the fact that our approach needs a large amount of computation for both training and deployment. To be able to train a policy that performs to convergence (for 20 million steps, as shown in Fig. 3, our architecture takes an average time of 240 hours. This time comes primarily from the forward

and backward propagation through the RSSM world model, and not from the rendering. This means that more efficient implementations of DreamerV3 hold the potential to reduce the training times. For example, using the open source JAX implementation [13] instead of the PyTorch one can potentially reduce the training times by a factor of 3-4 times.

VII. Conclusion

This paper introduces a novel application of MBRL using the DreamerV3 architecture to address the challenge of vision-based agile quadrotor flight. Our approach eliminates the need for intermediate representations, explicit state estimation, or extensive bootstrapping and is able of mapping raw RGB pixel inputs directly to control commands.

Our results demonstrate the efficacy of DreamerV3 by executing a drone racing task, both in simulation and the real world. Compared to model-free RL methods (in particular PPO), our approach achieves significantly higher sample efficiency and produces policies capable of navigating challenging race tracks at high speeds with policies trained from scratch and completely end-to-end, from camera feed to control commands. Moreover, the emergent perception-aware behaviors that arise in trained policies illustrate how end-to-end learning can leverage perceptual information to optimize actions and simultaneously use those very actions to refine perception. In general, our work provides a step forward applying pixel-to-command model-based RL to real-world mobile robotics.

References

[1] Syed Agha Hassnain Mohsan, Nawaf Qasem Hamood Othman, Yanlong Li, Mohammed H. Alsharif, and Muhammad Asghar Khan. Unmanned aerial vehicles (uavs): practical aspects, applications, open challenges, security issues, and future trends. *Intelligent Service Robotics*, 16:109–137, 3 2023. ISSN 18612784. doi: 10.1007/s11370-022-00452-4.

[2] Drew Hanover, Antonio Loquercio, Leonard Bauersfeld, Angel Romero, Robert Penicka, Yunlong Song, Giovanni Cioffi, Elia Kaufmann, and Davide Scaramuzza. Autonomous drone racing: A survey. *IEEE Transactions on Robotics*, 2024.

[3] Hyungpil Moon, Jose Martinez-Carranza, Titus Cieslewski, Matthias Faessler, Davide Falanga, Alessandro Simovic, Davide Scaramuzza, Shuo Li, Michael Ozo, Christophe De Wagter, et al. Challenges and implemented technologies used in autonomous drone racing. *Intelligent Service Robotics*, 2019.

[4] Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620(7976):982–987, 2023.

[5] Yunlong Song, Angel Romero, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Reaching the limit in autonomous racing: Optimal control versus reinforcement learning. *Science Robotics*, 8(82):eadg1462, 2023.

[6] Sunggoo Jung, Sungwook Cho, Dasol Lee, Hanseob Lee, and David Hyunchul Shim. A direct visual servoing-based framework for the 2016 iros autonomous drone racing challenge. *Journal of Field Robotics*, 35(1):146–166, 2018.

[7] Elia Kaufmann, Antonio Loquercio, Rene Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. Deep drone racing: Learning agile flight in dynamic environments. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 133–145. PMLR, 29–31 Oct 2018.

[8] Christophe De Wagter, Federico Paredes-Vallés, Nilay Sheth, and Guido de Croon. The artificial intelligence behind the winning entry to the 2019 ai robotic racing competition. *arXiv preprint arXiv:2109.14985*, 2021.

[9] Chao Qin, Maxime SJ Michet, Jingxiang Chen, and Hugh H-T Liu. Time-optimal gate-traversing planner for autonomous drone racing. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8693–8699. IEEE, 2024.

[10] Ismail Geles, Leonard Bauersfeld, Angel Romero, Jiaxu Xing, and Davide Scaramuzza. Demonstrating agile flight from pixels without state estimation. *Robotics: Science and Systems*, 2024.

[11] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[12] Jiaxu Xing, Angel Romero, Leonard Bauersfeld, and Davide Scaramuzza. Bootstrapping reinforcement learning with imitation for vision-based agile flight. *8th Conference on Robot Learning (CoRL)*, 2024.

[13] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.

[14] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, June 2013. ISSN 1076-9757.

[15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015. ISSN

0028-0836, 1476-4687.

[16] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016. ISSN 1476-4687. Number: 7587 Publisher: Nature Publishing Group.

[17] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. DeepMind Control Suite, January 2018. URL http://arxiv.org/abs/1801.00690. arXiv:1801.00690 [cs].

[18] Denis Yarats, Amy Zhang, Ilya Kostrikov, Brandon Amos, Joelle Pineau, and Rob Fergus. Improving sample efficiency in model-free reinforcement learning from images. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(12):10674–10681, May 2021. ISSN 2374-3468, 2159-5399.

[19] Elie Aljalbout, Ji Chen, Konstantin Ritt, Maximilian Ulmer, and Sami Haddadin. Learning vision-based reactive policies for obstacle avoidance. In *Conference on Robot Learning*, pages 2040–2054. PMLR, 2021.

[20] Danijar Hafner, Timothy P. Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2555–2565. PMLR, 2019.

[21] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: Contrastive unsupervised representations for reinforcement learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5639–5650. PMLR, 13–18 Jul 2020.

[22] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

[23] Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Mastering visual continuous control: Improved data-augmented reinforcement learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.

[24] Andrei A. Rusu, Matej Vecerík, Thomas Rothörl, Nicolas Heess, Razvan Pascanu, and Raia Hadsell. Sim-to-real robot learning from pixels with progressive nets. In *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings*, volume 78 of *Proceedings of Machine Learning Research*, pages 262–270. PMLR, 2017.

[25] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 23–30. IEEE, 2017.

[26] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.*, 17:39:1–39:40, 2016.

[27] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *arXiv*, 2024.

[28] Antonio Loquercio, Ana I Maqueda, Carlos R Del-Blanco, and Davide Scaramuzza. Dronet: Learning to fly by driving. *IEEE Robotics and Automation Letters*, 3(2):1088–1095, 2018.

[29] Dhruv Shah, Ajay Sridhar, Arjun Bhorkar, Noriaki Hirose, and Sergey Levine. GNM: A General Navigation Model to Drive Any Robot. In *International Conference on Robotics and Automation (ICRA)*, 2023. URL https://arxiv.org/abs/2210.03370.

[30] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. ViNT: A foundation model for visual navigation. In *7th Annual Conference on Robot Learning*, 2023. URL https://arxiv.org/abs/2306.14846.

[31] Ajay Sridhar, Dhruv Shah, Catherine Glossop, and Sergey Levine. NoMaD: Goal Masked Diffusion Policies for Navigation and Exploration. *arXiv pre-print*, 2023. URL https://arxiv.org/abs/2310.07896.

[32] Elia Kaufmann, Antonio Loquercio, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Deep drone acrobatics. In *Proceedings of Robotics: Science and Systems*, Corvalis, Oregon, USA, July 2020.

[33] William Koch, Renato Mancuso, Richard West, and Azer Bestavros. Reinforcement learning for uav attitude control. *ACM Transactions on Cyber-Physical Systems*, 3(2):1–21, 2019.

[34] Nathan O Lambert, Daniel S Drew, Joseph Yaconelli, Sergey Levine, Roberto Calandra, and Kristofer SJ Pister. Low-level control of a quadrotor with deep model-based reinforcement learning. *IEEE Robotics and Automation Letters*, 4(4):4224–4230, 2019.

[35] Robin Ferede, Christophe De Wagter, Dario Izzo, and Guido CHE de Croon. End-to-end reinforcement learning for time-optimal quadcopter flight. *arXiv preprint*

arXiv:2311.16948, 2023.

[36] Jonas Eschmann, Dario Albani, and Giuseppe Loianno. Learning to fly in seconds. *arXiv e-prints*, pages arXiv–2311, 2023.

[37] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.

[38] Philipp Foehn, Elia Kaufmann, Angel Romero, Robert Penicka, Sihao Sun, Leonard Bauersfeld, Thomas Laengle, Giovanni Cioffi, Yunlong Song, Antonio Loquercio, et al. Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight. *Science Robotics*, 7 (67):eabl6259, 2022.

[39] Philip Becker-Ehmck, Maximilian Karl, Jan Peters, and Patrick van der Smagt. Learning to fly via deep model-based reinforcement learning. *arXiv preprint arXiv:2003.08876*, 2020.

[40] Niklas Wahlström, Thomas B Schön, and Marc Peter Deisenroth. From pixels to torques: Policy learning with deep dynamical models. *arXiv preprint arXiv:1502.02251*, 2015.

[41] Thomas Bi and Raffaello D'Andrea. Sample-efficient learning to solve a real-world labyrinth game using data-augmented model-based reinforcement learning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7455–7460. IEEE, 2024.

[42] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Daydreamer: World models for physical robot learning. In *Conference on Robot Learning (CoRL)*. PMLR, 2022.

[43] Christian Pfeiffer and Davide Scaramuzza. Human-piloted drone racing: Visual processing and control. *IEEE Robotics and Automation Letters*, 6(2):3467–3474, 2021.

[44] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2555–2565. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/hafner19a.html.

[45] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL http://jmlr.org/papers/v22/20-1364.html.

[46] Yunlong Song, Selim Naji, Elia Kaufmann, Antonio Loquercio, and Davide Scaramuzza. Flightmare: A flexible quadrotor simulator. In *Conference on Robot Learning*, 2020.

[47] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9339–9347, 2019.

[48] Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander Clegg, Michal Hlavac, So Yeon Min, Vladimír Vondruš, Theophile Gervet, Vincent-Pierre Berges, John M Turner, Oleksandr Maksymets, Zsolt Kira, Mrinal Kalakrishnan, Jitendra Malik, Devendra Singh Chaplot, Unnat Jain, Dhruv Batra, Akshara Rai, and Roozbeh Mottaghi. Habitat 3.0: A co-habitat for humans, avatars, and robots. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=4znwzG92CE.

[49] Yunlong Song, Mats Steinweg, Elia Kaufmann, and Davide Scaramuzza. Autonomous drone racing with deep reinforcement learning. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2021.

[50] Danijar Hafner, Timothy P. Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.