End-to-End Image-Goal Navigation for Agile Quadrotors via Recurrent State Distillation

Simone Nascivera Ismail Geles Davide Scaramuzza
Robotics and Perception Group, University of Zurich, Switzerland

Abstract

Image-goal navigation asks a robot to reach the viewpoint of a target photograph without a map or goal pose. In three dimensions, this requires not only recognizing the goal but also estimating 3D displacement and generating continuous control actions to approach and settle at the target view. This is difficult because the same visual difference between current and goal image can require very different actions depending on the robot's velocity, yet monocular images alone reveal neither metric scale nor body motion. As a result, existing image-goal policies are limited to planar action spaces and rely on a state estimator for flight control. Policies that map directly from images to low-level control commands have been demonstrated for agile flight, but not toward goal images in unseen environments. We present the first image-goal navigation policy that directly commands thrust and angular rates for agile quadrotor flight in unseen 3D environments, without relying on state estimation (e.g., SLAM, VO, or SfM) or a separate velocity controller. We train a recurrent vision-based policy through a four-stage pipeline that combines geometric and motion supervision with privileged-teacher distillation and end-to-end reinforcement learning fine-tuning. On 1000 held-out flights in unseen environments, the policy achieves an 82.3% success rate with a mean terminal position error of 0.10 m across successful flights. Deployed zero-shot on a real quadrotor, the policy reaches an 84% success rate over 25 trials in an unseen indoor environment.

1. Introduction

Human pilots navigate quadrotors through unfamiliar environments using only an onboard camera stream, without explicitly estimating the vehicle's position, velocity, or attitude [10]. We investigate whether autonomous flight can similarly rely on visual feedback alone: given a goal photograph, a quadrotor must reach the viewpoint from which it was taken as quickly as possible (Fig. 1). The photograph specifies a destination without requiring a map or a metric goal pose. We study how to achieve this with a policy

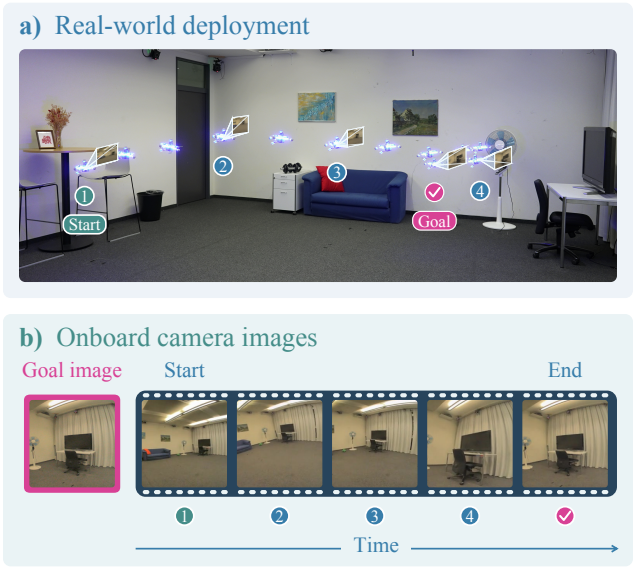

Figure 1. **A goal image specifies where to fly.** Real-world deployment of our method (a). The quadrotor flies from the starting image to the viewpoint goal, only relying on the onboard camera frames (b). No explicit state estimation (i.e., SLAM, VO, or SfM) is used.

that maps visual observations and last applied actions directly to low-level control commands, without explicit state estimation such as Simultaneous Localization and Mapping (SLAM), Visual Odometry (VO), or Structure from Motion (SfM) [2].

Extracting the information needed for flight control from onboard frames and policy's past actions is challenging. Reaching a goal viewpoint requires translating in three dimensions, aligning the camera, and braking to hover at the destination. The appropriate control command depends not only on the displacement to the goal but also on the vehicle's current motion. However, velocity cannot be directly inferred from a single frame, and monocular observations leave metric scale ambiguous.

To address these limitations, we use a recurrent policy that integrates visual information over time, allowing it to implicitly infer the vehicle's motion [32]. The previous con-

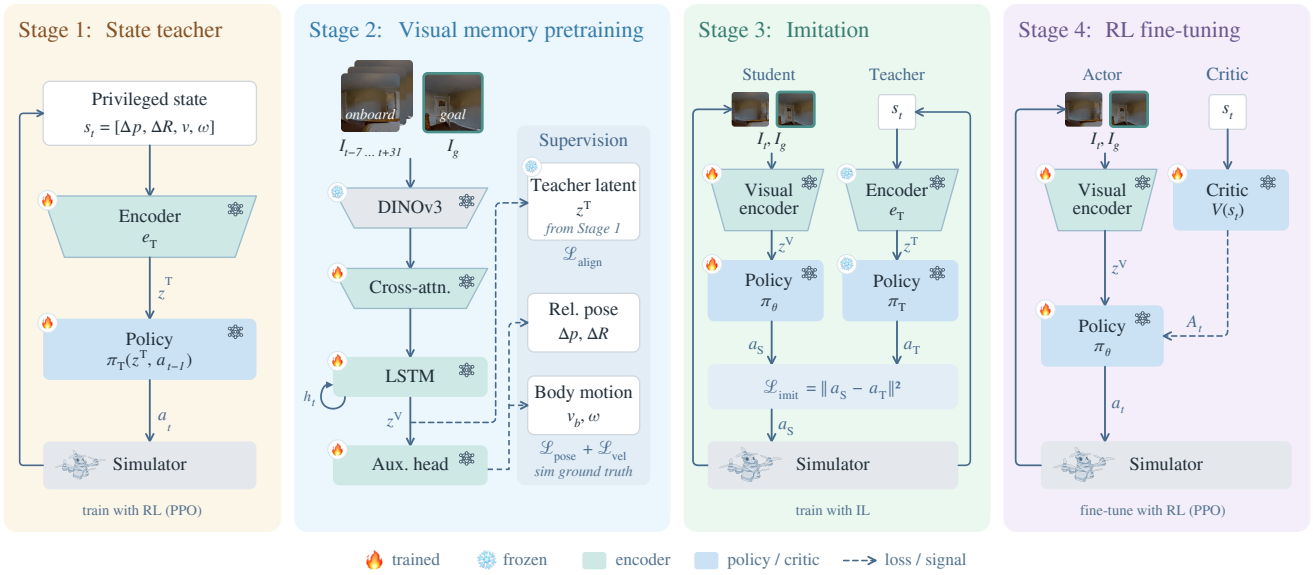

Figure 2. **Four-stage training pipeline.** The policy is trained progressively through state-based teacher learning (Stage 1), visual representation pretraining (Stage 2), imitation learning (Stage 3), and reinforcement learning fine-tuning (Stage 4). Stage 3 weights are initialized with the pretrained visual encoder and the teacher actor's weights, while Stage 4 fine-tunes the resulting policy for improved flight performance.

trol command provides additional information about the vehicle dynamics, helping reduce metric scale ambiguity [21]. Tilting the vehicle also changes its thrust direction and camera view together. A visual policy must account for these coupled changes over time.

Recent work, such as SIGN [36], has explored aerial image-goal navigation. While the authors demonstrate image-goal navigation on quadrotors, including real-world flight, the method outputs planar velocity commands, relying on a separate state-estimation-based controller for flight stabilization. In parallel, vision-based policies have demonstrated agile flight by directly commanding low-level control inputs, without explicit state estimation [10, 11]. Combining image-goal navigation with such direct flight control remains an open challenge, particularly when requiring three-dimensional motion, precise settling at the goal, and generalization to unseen real-world environments.

We address this problem through a four-stage training pipeline (Fig. 2). First, a state-based teacher learns to fly using privileged goal-relative pose and motion information. We then train a visual representation from sequences of onboard images and a goal photograph, using geometric and motion supervision together with alignment to the teacher's latent representation. During the next stage, the student is trained through imitation learning using the teacher from the first stage. Finally, reinforcement learning fine-tunes the student policy in closed loop. At deployment, the policy receives only the current image, the goal photograph, and its previous command, with no access to position, velocity, attitude, or IMU measurements.

Contributions. **(1)** To our knowledge, we present the first image-goal flight policy that directly commands low-level control inputs in unfamiliar environments without explicit state estimation. **(2)** We introduce a four-stage training procedure that combines geometric and motion supervision, latent representation alignment, imitation learning, and reinforcement learning to learn a vision-based flight policy. **(3)** We evaluate the policy on 1000 held-out flights in 100 unseen environments, achieving an 82.3% success rate with a mean terminal position error of 0.10 m across successful flights, and demonstrate zero-shot transfer to a real quadrotor with an 84% success rate over 25 trials in an unseen indoor environment.

2. Related Work

Aerial image-goal navigation. SIGN [36] and Omni-IGN [37] demonstrate aerial image-goal navigation with planar velocity commands and dedicated flight-control loops. Their deployed systems use visual-inertial and LiDAR-inertial odometry, respectively, for low-level control. Both address goal discovery and collision avoidance, with Omni-IGN extending perception to four fisheye views. UA-NWM [39] uses a latent world model to score candidate trajectories and plan waypoints, while DreamNav [17] refines relative-motion predictions by generating candidate observations. Our setting instead focuses on reaching and settling at covisible goal viewpoints under quadrotor dynamics, including altitude changes. The policy directly commands collective thrust and body rates from vi-

sual observations and previous commands, without explicit vehicle-state inputs or a separate velocity controller.

Image-goal representations. Image-goal navigation has been extensively studied in simulated embodied environments [16, 35, 40]. FGPrompt [31] emphasizes fine-grained goal conditioning, while Monaci *et al*. [23] analyze visual fusion strategies and relate navigation performance to the relative geometry recoverable from learned representations. DEBiT [1] pretrains a binocular transformer through cross-view completion and relative-pose and visibility supervision, then integrates its features into a recurrent navigation policy. We similarly use current–goal cross-attention, but pretrain the recurrent representation on ordered flight sequences with geometric and motion supervision. Complementary approaches such as AnyImageNav [7] combine exploration with explicit geometric registration of the goal image using depth and pose observations. Our deployed policy instead uses the learned representation directly for flight control.

Visual servoing and relative-pose estimation. Classical image-based visual servoing (IBVS) [3] computes camera-velocity commands from image-feature errors. Learned representations can improve feature matching, as in ViT-VS [28], which combines pretrained DINOv2 correspondences with an IBVS control law. However, avoiding explicit pose reconstruction in the visual front end does not necessarily eliminate state estimation from the complete flight-control pipeline. In the velocity-command IBVS baselines considered here, a separate controller uses vehicle-state estimates to track the commanded motion. Learned matching and relative-pose methods, including MASt3R [15] and Reloc3r [8], provide geometric estimates that can likewise support modular navigation systems.

Our objective is instead to perform image-goal flight without explicit vehicle-state estimation or an intermediate velocity-control loop. The recurrent policy maps visual observations and previous commands directly to collective thrust and body rates, learning to account for the coupling between vehicle attitude, translation, and camera motion.

Privileged learning for visual control. Teacher–student learning transfers privileged training information to deployable visuomotor policies [4, 13, 18]. Xing *et al*. [34] combine state-based teacher training, visual imitation learning, and reinforcement learning fine-tuning for agile flight. We adopt this progression and introduce a separate visual representation-learning stage before action imitation. Our latent alignment is related to RMA [14], which predicts privileged environment and dynamics embeddings from observation–action histories, and perceptive locomotion methods that learn recurrent representations with privileged supervision [22]. Here, the visual representation is supervised by goal-relative pose and body motion and aligned with the teacher's latent space to support reuse of its actor.

Vision-based low-level flight control. Geles *et al*. [10] demonstrate agile flight without explicit state estimation using gate-based visual representations, while Dream to Fly [26] learns flight control from pixels through model-based reinforcement learning. Heeg *et al*. [11] study differentiable training of quadrotor controllers from visual features. Gaussian-splatting environments also support visuomotor policy learning through differentiable dynamics [5] or imitation of an MPC expert [19], with GRaD-Nav++ [6] introducing language conditioning. Our work addresses image-goal-conditioned flight: the target viewpoint is specified by a photograph, and the policy must reach and settle at that viewpoint in unseen environments. Training does not require gradients through the simulator.

3. Method

Figure 2 illustrates the four stages of our training pipeline and their main components.

3.1. Task Definition

The objective is to fly a quadrotor quickly and precisely from an initial stationary pose to a goal viewpoint. The goal is specified by an RGB image I_g, captured at a viewpoint with partial visual overlap with the initial observation and zero roll and pitch. At each time step t, the policy receives the current image I_t, the goal image I_g, and the previous normalized control command a_{t-1}. The policy is recurrent, with hidden state h_{t-1}, and outputs

$$a_t = \pi_\theta(I_t, I_g, a_{t-1}; h_{t-1}), \qquad a_t \in [-1, 1]^4. \quad (1)$$

The normalized action is mapped to collective thrust per unit mass and three body angular rates. Notably, the policy receives no explicit state estimates such as position, velocity, or attitude. A flight is considered successful when the quadrotor reaches the goal viewpoint within specified position, orientation, and velocity tolerances defined in section Sec. 4.

3.2. Quadrotor Simulator

We simulate the quadrotor dynamics using the standalone Warp-based simulator described in [38], derived from the Agilicious architecture [9]. Our custom quadrotor has a mass of $0.21\,\mathrm{kg}$, an arm length of $19.4\,\mathrm{cm}$, a propeller radius of $3.81\,\mathrm{cm}$, and a thrust-to-weight ratio of 6.8. The simulator models calibrated inertia, motor thrust characteristics, battery dynamics, aerodynamic effects, and actuation delays, with a nominal command delay of $40\,\mathrm{ms}$. The policy outputs collective thrust and body-rate (CTBR) commands, which are converted into motor commands by the

Betaflight low-level controller. The dynamics run at 500 Hz, while the policy operates at 50 Hz, holding each command constant for 10 simulation steps.

Visual observations are generated using Isaac Lab's camera interface and NVIDIA RTX renderer, with the rendered pose synchronized to the Warp dynamics. A forward-facing monocular camera mounted at the body-frame origin produces 224×224 images using a calibrated polynomial fisheye model. Depth is additionally rendered for training-data validation and reward computation, but is never provided to the policy.

3.3. Stage 1: State-Based Teacher

We first train a state-based teacher using PPO [29]. Let $\mathbf{p}_t$, R_t, $\mathbf{v}_t$, and $\boldsymbol{\omega}_t$ denote the position, body-to-world orientation, world-frame linear velocity, and body angular velocity, respectively. The teacher observes the 12-dimensional goal-relative state

$$s_t = \big[R_t^\top(\mathbf{p}_g - \mathbf{p}_t),\ \log(R_t^\top R_g)^\vee, \\ R_t^\top \mathbf{v}_t,\ \boldsymbol{\omega}_t\big] \in \mathbb{R}^{12}. \tag{2}$$

Here, $\log$ denotes the matrix logarithm on $\mathrm{SO}(3)$, and $(\cdot)^\vee$ maps a skew-symmetric matrix to its corresponding 3D rotation vector. A linear layer followed by GELU embeds s_t into a 128-dimensional representation $z_t^T = e_T(s_t)$. The actor concatenates z_t^T with the previous action a_{t-1} and passes them through two 512-unit hidden layers to output four normalized commands in $[-1, 1]$. The privileged critic additionally observes the flattened rotation matrix, yielding a 25-dimensional input $[s_t, a_{t-1}, \mathrm{vec}(R_t)]$.

The reward combines position progress, attitude alignment, height tracking, control regularization, and terminal outcomes:

$$r_t = r_t^{\mathrm{pos}} + r_t^{\mathrm{att}} + r_t^{\mathrm{height}} + r_t^{\mathrm{reg}} \\ - 0.005 + 10\mathbb{I}_t^{\mathrm{success}} - 20\mathbb{I}_t^{\mathrm{failure}}. \tag{3}$$

Defining $d_t = \|\mathbf{p}_t - \mathbf{p}_g\|_2$ and $\Delta d_t = d_{t-1} - d_t$, the task rewards are

$$r_t^{\mathrm{pos}} = \begin{cases} 4\Delta d_t, & \Delta d_t \geq 0, \\ 20\Delta d_t, & \Delta d_t < 0, \end{cases} \\ r_t^{\mathrm{att}} = 0.125(\theta_{t-1} - \theta_t), \\ r_t^{\mathrm{height}} = -0.1(p_{t,z} - z_t^{\mathrm{ref}})^2. \tag{4}$$

Here, θ_t is the geodesic orientation error relative to a look-at frame pointing 2 m beyond the goal along its forward axis, with roll determined by the goal's up direction. The reference height z_t^{ref} follows the linear interpolation between the start and goal heights according to horizontal progress. For predominantly vertical flights, it instead penalizes motion beyond the endpoint heights.

Control regularization penalizes changes in thrust and body-rate commands, as well as excessive angular rates:

$$r_t^{\mathrm{reg}} = -\lambda_\ell \big[0.3(\Delta \bar{u}_{T,t})^2 \\ + 0.05\|\Delta \bar{\mathbf{u}}_{\omega,t}\|_2^2 \\ + 0.05\|\bar{\boldsymbol{\omega}}_{xy,t}\|_2^2 \\ + 0.025\|\bar{\mathbf{u}}_{\omega,t}\|_2^2\big]. \tag{5}$$

Here, $\Delta \bar{\mathbf{u}}_t$ denotes command differences normalized by their respective scales, while $\bar{\boldsymbol{\omega}}_t$ and $\bar{\mathbf{u}}_{\omega,t}$ are normalized by the body-rate limits. The curriculum multiplier λ_ℓ increases from 0.25 to 1.25.

We adopt a 20 discrete-steps curriculum that progressively increases goal displacement, yaw differences, and dynamics perturbations. Advancement requires at least 90% success over an assessment window of 1024 completed episodes, aggregated across parallel environments. Success requires position, velocity, and attitude errors to remain within stage-specific tolerances for a prescribed number of consecutive control steps. We additionally randomize mass, inertia, drag, actuation parameters, and delays. The resulting teacher learns goal-directed flight control from privileged state observations, providing the target behavior and latent representation for subsequent stages.

3.4. Stage 2: Visual Encoder

In the second stage, we train a visual encoder to infer the 12-dimensional privileged pose-and-motion state s_t, defined in Eq. (2), from the current image I_t, goal image I_g, and previous action a_{t-1}. A frozen DINOv3 ViT-S/16 backbone [30] extracts 384-dimensional patch descriptors from both images. Similar to [12], we augment each patch with its corresponding 3D camera ray. The visual descriptors are projected to 256 dimensions and combined by elementwise addition with ray embeddings produced by a learned $3 \to 64 \to 256$ MLP:

$$\tilde{\mathbf{f}}_i = \mathrm{Adapter}(\mathbf{f}_i) + \mathrm{MLP}_{\mathrm{ray}}(\mathbf{r}_i). \tag{6}$$

The resulting 256-dimensional tokens are processed by two bidirectional cross-attention blocks, allowing current and goal patches to exchange information. Learned spatial pooling and fusion of the two streams, their difference, and elementwise product produce a 128-dimensional pair feature x_t. The previous normalized command $a_{t-1} \in [-1, 1]^4$ is concatenated with this feature before being passed to a single-layer LSTM:

$$\tilde{x}_t = [x_t; a_{t-1}], \\ (h_t, c_t) = \mathrm{LSTM}(\tilde{x}_t, h_{t-1}, c_{t-1}), \\ z_t^V = P(h_t) \in \mathbb{R}^{128}, \tag{7}$$

where P is a linear projection followed by GELU. The LSTM has a hidden dimension of 256, while the projected representation z_t^V has dimension 128, matching the

teacher's latent space. The recurrent state is reset at episode boundaries and goal changes. An auxiliary head predicts the privileged state $\hat{s}_t = H(z_t^V)$.

We train the encoder using geometric and motion supervision, together with alignment to the frozen teacher representation:

$$\begin{aligned}\mathcal{L}_{\text{pre}} = \lambda_p \Big(& \|\hat{p}_{\text{rel}} - p_{\text{rel}}\|_2^2 \\ & + d_{\text{SO}(3)}(\hat{R}_{\text{rel}}, R_{\text{rel}})^2 \Big) \\ & + \frac{\lambda_v}{3} \|\hat{v}_b - v_b\|_2^2 \\ & + \frac{\lambda_\omega}{3} \|\hat{\omega} - \omega\|_2^2 \\ & + \frac{\lambda_z}{128} \|z_t^V - \text{sg}(z_t^T)\|_2^2.\end{aligned} \tag{8}$$

Here, $z_t^T = e_T(s_t)$ is the frozen teacher embedding, sg stops gradients, and $d_{\text{SO}(3)}$ denotes the geodesic rotation error. The predicted rotation vector is converted into $\hat{R}_{\text{rel}}$ through the exponential map before computing the rotation loss. Position, orientation, linear velocity, and angular velocity errors use reference scales of 1 m, 1 rad, 1 m/s, and 1 rad/s, respectively. All loss weights are set to one. The teacher encoder and DINOv3 backbone remain frozen.

Training uses ordered clips from teacher rollouts, with eight warm-up steps to initialize the recurrent state without gradients, followed by 32 supervised steps. We collect clips from teacher curriculum levels $\{0, 4, 8, 19\}$ and weight their contributions equally in the training objective. Losses are averaged over valid supervised frames after warm-up. The resulting representation is trained to match the teacher's latent features, allowing the pretrained actor to be reused in the subsequent imitation learning stage.

3.5. Stage 3: Imitation Learning

In the third stage, we initialize the student actor with the teacher's weights and replace the teacher's state-based embedding z_t^T with the pretrained recurrent visual embedding z_t^V. The actor receives z_t^V and the previous normalized command $a_{t-1} \in [-1, 1]^4$, without explicit state estimates.

Following the dataset-aggregation principle of DAgger [27], the teacher provides action labels at states visited during data collection. At each control step, we execute the teacher's action with probability β_k and the student's action otherwise. The teacher-selection probability decreases linearly from 0.5 to zero over 1,000 training iterations. Teacher supervision continues after the student assumes full control.

We use the same 20-level competence-based curriculum as in Stage 1. Curriculum advancement is suspended during teacher-assisted collection. Once assistance ends, the curriculum statistics are reset, and advancement requires a student-only success rate of at least 90% over at least 1024 completed episodes. Collected trajectories are stored as ordered replay clips, preserving temporal order and episode boundaries. Each clip contains 8 burn-in steps, processed without gradients to reconstruct recurrent memory, followed by 32 supervised steps. The student minimizes

$$\mathcal{L}_{\text{imit}} = \mathbb{E}_{\mathcal{D}} \left[\frac{1}{4} \left\| \pi_\theta(z_t^V, a_{t-1}) - a_t^T \right\|_2^2 \right], \tag{9}$$

where a_t^T is the teacher's action at the visited state. The loss is averaged over valid supervised frames after burn-in. Gradients jointly update the visual attention encoder, LSTM, projection, and actor, while the DINOv3 backbone remains frozen. No auxiliary state-prediction or embedding-alignment losses are applied during this stage. We disable collision detection during imitation learning, as we empirically found it to hinder training, and re-enable it during the subsequent reinforcement learning stage.

3.6. Stage 4: Reinforcement Learning

In the final stage, we initialize the actor and recurrent visual encoder from Stage 3, while the privileged critic retains the Stage 1 teacher's weights, unchanged during imitation learning. We jointly fine-tune the visual attention encoder, LSTM, projection, and actor using PPO, while keeping DINOv3 frozen. The critic is optimized using privileged simulator state. All trainable components, including the critic, share a fixed learning rate of 10^{-5} through a single Adam optimizer, without a separate critic warm-up.

We retain the Stage 1 reward and introduce an additional image-reprojection penalty near the goal:

$$r_t^{(4)} = r_t^{(1)} + r_t^{\text{repr}}. \tag{10}$$

The penalty measures image alignment by projecting 3D points from the goal image into the current view using ground-truth simulator geometry. It combines reprojection error and geometric image overlap, activates within $0.5\,\text{m}$ of the goal, and is bounded by 0.02 per control step. The penalty smoothly vanishes at the activation boundary. The simulator geometry used to compute this penalty is required only during training. We also re-enable collision detection, which was disabled during imitation learning. The complete reward formulation is provided in the supplementary material.

3.7. Training Details

We precompute approximately 10,000 start-goal pairs per HM3D training building, requiring zero tilt, relative yaw $\leq 90°$, covisibility $\geq 20\%$, valid depth coverage $\geq 70\%$, and 0.6 m clearance along the straight segment, including both endpoints. The Stage 1 curriculum separately allows yaw differences up to $180°$. Valid pairs are cached to avoid repeated geometric checks during training.

All stages are trained on a single NVIDIA RTX 5090 GPU for 10,000, 10,000, 25,000, and 35,000 iterations, respectively. Stage 1 uses 2,048 parallel environments with 24 rollout steps (49,152 transitions per iteration), five PPO epochs, four mini-batches, and an adaptive learning rate initialized at 10^{-3} (target KL divergence of 0.02). Stages 2 to 4 use 256 parallel environments. Stages 2 and 3 optimize mini-batches of 128 and 1,024 supervised frames, respectively, using 32-step clips with eight additional burn-in steps and masking invalid frames. Both use a fixed learning rate of 10^{-4}. Stage 4 collects 6,144 transitions per iteration, with three PPO epochs, 16 mini-batches, and a fixed learning rate of 10^{-5}.

4. Experiments

We evaluate our policy's ability to reach goal viewpoints in buildings unseen during training. We compare against classical and learning-based methods in simulated environments and evaluate zero-shot transfer to a real-world setting. Finally, we conduct ablation studies to assess the contribution of the main components of our approach.

4.1. Experimental Setup

Datasets We train on the 800 HM3D training buildings [25] and evaluate on 1,000 fixed flights across the 100 validation buildings, with 10 flights per building. The flights are divided into three goal-distance ranges: 334 at 0.3-2 m, 333 at 2-5 m, and 333 at 5-10 m. Rather than using the trajectories provided by the dataset, we sample our own start-goal pairs to satisfy the covisibility requirements of our task. Specifically, we require at least 20% visual overlap between the initial and goal views, a maximum relative yaw of 90 degrees, and zero roll and pitch at both poses. We require 0.6 m clearance along the straight start-goal segment, endpoints included. This pair filter does not constrain the executed flight path. All validation buildings are unseen during training. To further assess generalization, we evaluate the policy zero-shot on 1,000 fixed flights across 14 Gibson [33] buildings, none of which are used for training.

Finally, we evaluate the policy zero-shot on a physical quadrotor using its onboard camera in an indoor flight arena. We use 25 start-goal image pairs to assess real-world flight performance.

Baselines. To our knowledge, no existing image-goal navigation method directly commands collective thrust and body rates without explicit state estimation. We therefore compare against methods that rely on state estimates for low-level flight control, despite the additional information available to them.

We compare against two variants of Image-Based Visual Servoing (IBVS) [3]. *Oracle IBVS* uses ground-truth correspondences between the current and goal images, while *SIFT IBVS* obtains correspondences using SIFT features [20] and brute-force descriptor matching. Both variants use ground-truth depth for the current image and noise-free state estimates, with Oracle IBVS additionally using the goal-image depth. We evaluate both variants with translational velocity command limits of 1, 3, and 5 $\mathrm{m/s}$, noting that the achieved flight speed may exceed these limits. The resulting camera velocity and yaw-rate commands are passed to a geometric velocity and attitude controller, which converts them into collective thrust and body-rate commands. Both baselines use the same quadrotor, camera configuration, and low-level controller as our method.

We also evaluate SIGN [36] using its native camera configuration and a 1 s decision interval. SIGN outputs forward velocity and yaw-rate commands, without vertical control. As with the IBVS baselines, we use a geometric controller with access to the ground-truth simulator state to convert these commands into collective thrust and body rates, ensuring the same low-level control interface and as our method.

Evaluation Metrics. A flight succeeds when position error is below 0.15 m, attitude error below 10 °, and speed below 0.08 $\mathrm{m/s}$ for one control step (0.02 s). Flights that neither succeed nor crash within the time limit are counted as timeouts. To detect collisions, we sweep a body-centered sphere of radius 0.2 m along the executed trajectory, sampling at intervals of at most 0.05 m, and count contact with the building mesh as a crash. We report the success, crash, and timeout rates across the evaluation set. For successful flights, we additionally report the mean completion time [s], final attitude error [°], and final position error [m].

4.2. Synthetic Evaluation

HM3D. Table 1 reports the performance of our method and the baselines on the HM3D validation set. Our method achieves the highest success rate among all baselines, despite operating without explicit state estimation or a velocity controller. However, it exhibits longer completion times and larger final position errors than the IBVS baselines. Oracle IBVS achieves a significantly higher success rate than SIFT IBVS, but also suffers from more collisions. In contrast, SIFT IBVS predominantly fails through timeouts. Over 99% of these timed-out flights have a terminal speed below 0.003 $\mathrm{m/s}$; their median remaining distance to the goal is 5 m. This suggests that unreliable feature correspondences cause the controller to stall before reaching the goal. For SIGN, we report both the success rate under our criteria and the one presented in the paper itself, which is within 1 m and 25 ° from the target pose. In addition, in this secondary success rate we consider a 2D projected distance, instead of 3D like all the other baselines, since the original paper maintained constant altitude.

Table 1. HM3D validation results on 1,000 trajectories across 100 held-out buildings. Our method and IBVS use a 20 s time limit, while SIGN uses 500 s to account for its lower maximum velocity. Completion time and final errors are reported only for successful flights. For SIGN, we report two success rates, with the value in parentheses computed using the success criterion from the original paper [36]. Bold indicates the best results among non-oracle methods.

Method	Velocity [m/s] Limit	Success [%] ↑ Rate	Failures ↓ Crashes / Timeouts	Time [s] ↓ Mean / Median	Position error [m] ↓ Mean / Median	Attitude error [°] ↓ Mean / Median
Oracle IBVS	1	74.6	252 / 2	3.7 / 3.3	0.04 / 0.04	3.31 / 2.03
	3	70.1	299 / 0	2.4 / 2.3	0.04 / 0.03	2.76 / 2.71
	5	66.7	333 / 0	2.3 / 2.3	0.04 / 0.03	2.77 / 2.45
SIFT IBVS	1	40.8	**26** / 566	2.9 / 2.5	**0.04** / 0.04	3.43 / 3.52
	3	38.8	52 / 560	**2.4** / **2.1**	**0.04** / **0.03**	**2.56** / 2.29
	5	38.2	61 / 557	2.5 / **2.1**	**0.04** / **0.03**	2.64 / **2.28**
SIGN	0.25	1.1 / (53.0)	71 / 918	66.5 / 25.9	0.12 / 0.14	7.96 / 9.74
Ours	–	**82.3**	89 / **88**	5.1 / 4.7	0.10 / 0.11	2.77 / 2.52

Table 2. Gibson evaluation on 1,000 matched validation flights, with a time limit of 500 s for all methods. Time and error statistics summarize successful flights only. For SIGN, we report two success rates, with the value in parentheses computed using the success criterion from the original paper [36]. Bold indicates the best results among non-oracle methods.

Method	Velocity [m/s] Limit	Success [%] ↑ Rate	Failures ↓ Crashes / Timeouts	Time [s] ↓ Mean / Median	Position error [m] ↓ Mean / Median	Attitude error [°] ↓ Mean / Median
Oracle IBVS	1	95.7	42 / 1	2.6 / 2.3	0.03 / 0.02	3.18 / 2.13
	3	92.0	79 / 1	2.1 / 2.0	0.03 / 0.03	2.94 / 2.67
	5	91.5	84 / 1	2.1 / 2.0	0.03 / 0.03	2.95 / 2.61
SIFT IBVS	1	80.1	67 / 132	6.4 / 2.2	**0.03** / **0.02**	3.27 / 2.48
	3	74.5	136 / 119	**3.4** / **2.0**	**0.03** / 0.03	2.89 / 2.69
	5	74.6	136 / 118	**3.4** / **2.0**	**0.03** / 0.03	2.94 / 2.65
SIGN	0.25	11.7 / (79.0)	**0** / 883	24.3 / 23.2	0.10 / 0.11	7.89 / 9.80
Ours	–	**93.0**	53 / **17**	5.4 / 4.0	0.10 / 0.10	**2.53** / **2.35**

Gibson. Our policy demonstrates strong generalization to the unseen Gibson dataset, achieving a success rate of 93%. As on HM3D, Table 2 shows that collisions are the dominant failure mode for the Oracle IBVS baseline, becoming more frequent as the velocity command limit increases. SIFT IBVS instead fails due to both collisions and timeouts. SIGN achieves a higher success rate than on HM3D. However, the trajectories sampled for both synthetic evaluations do not enforce the constant-altitude constraint used during SIGN's training. Since SIGN does not command vertical motion, we include it as a reference for existing image-goal navigation methods rather than a directly comparable solution to our three-dimensional navigation task. Finally, our method remains the only evaluated policy that does not rely on explicit state estimation or a separate velocity controller.

4.3. Real World Deployment

We deploy our policy zero-shot in an indoor flight arena, following a setup similar to [10, 38]. We use a modified version of *Agilicious* [9], streaming control commands from a workstation to the quadrotor via an RF transmitter and receiving low-latency video from its onboard camera. The policy is optimized using NVIDIA TensorRT [24] with FP32 precision, achieving a median inference time of 1.06 ms on an NVIDIA RTX 4090 GPU.

Figure 3 shows a subset of the 25 start-goal image pairs used in our real-world experiments. We apply the same success criteria as in simulation, while using a motion capture system only for recording the resulting performance.

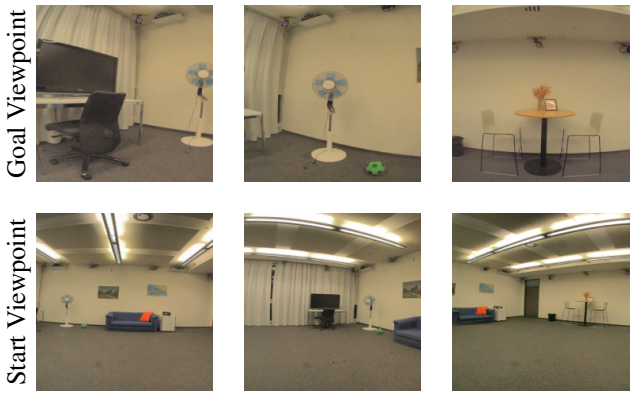

(a) Sequence 0 (b) Sequence 10 (c) Sequence 11

Figure 3. Three examples of the 25 start and goal viewpoints used during the real world experiments in the indoor flying arena.

Our policy achieves a success rate of **84%**, with a mean final position error of **0.118** m and orientation error of **3.07** ° across successful flights. The initial position and orientation differences average 3.02 m and 18.95 °, respectively. These results demonstrate that the policy transfers to a previously unseen real-world environment without fine-tuning, despite

variable video transmission delays and substantial frame drops. Across the 25 experimental recordings, 30.51% of received frames were dropped.

4.4. Ablations

We evaluate the contribution of the recurrent visual encoder and the importance of each training stage. For the visual encoder, we study the effect of previous control commands, latent dimension, and recurrence on geometric and motion prediction. For the training pipeline, we compare different initialization strategies and assess the contribution of each training stage.

Recurrent Visual Encoder. We ablate previous control commands, recurrence, and LSTM hidden dimension. Figure 4 shows training losses smoothed with an exponential moving average (decay 0.995), along with the validation loss evaluated at iteration 10000. Validation uses 400 held-out trajectories from curriculum levels 0, 4, 8, and 19 across 100 unseen HM3D buildings.

Removing previous commands increases training and validation losses by factors of 3.7 and 4, respectively. Shuffling command sequences across trajectories, while preserving teacher control and supervision, produces losses comparable to removing commands entirely. This suggests that the benefit comes from the temporal correspondence between actions and visual observations, rather than command values alone.

Increasing the LSTM hidden dimension from 128 to 256 and 512 reduces validation loss from 0.078 to 0.066 and 0.056, respectively. We select 256 to limit computational cost, as 512 increases training time by 11%. Finally, removing recurrence increases validation loss by nearly 12×, with angular velocity prediction loss increasing by 23.6×, highlighting the importance of visual memory for motion inference.

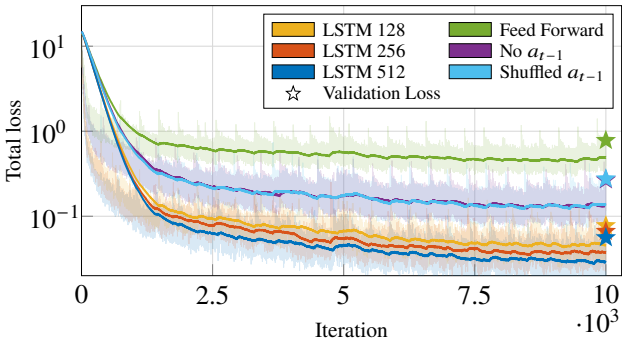

Figure 4. Stage 2 visual encoder ablations showing the evolution of training loss, along with the final validation loss.

Staged Training. We evaluate the contribution of each training stage on the 1,000 held-out HM3D trajectories used in our main evaluation. The student after imitation learning (Stage 3) achieves a success rate of 64.8%, which increases to 82.3% after reinforcement learning fine-tuning (Stage 4).

To assess the importance of initialization, we train PPO directly from scratch, starting at curriculum level 0 with an adaptive learning rate initialized at 10^{-4} and a target KL divergence of 0.01. This variant fails to progress beyond the initial curriculum level. Initializing PPO from the pretrained visual encoder (Stage 2), without imitation learning, allows limited progress but fails to advance beyond the first curriculum level. Similarly, training Stage 3 without Stage 2 visual pretraining fails to progress beyond the initial curriculum level.

These results highlight the importance of staged initialization: visual pretraining enables imitation learning to progress through the curriculum, while imitation provides an effective initialization for reinforcement learning, which further improves flight success.

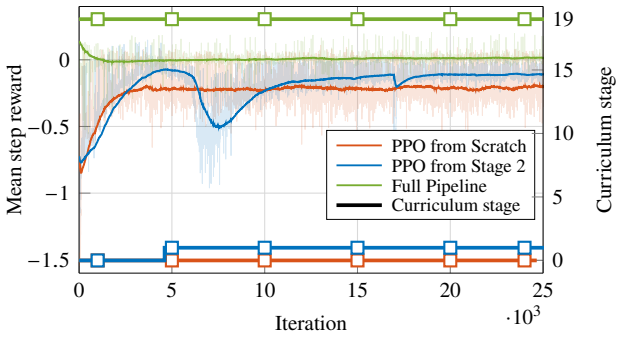

Figure 5. Stage 4 initialization ablation. The full pipeline progresses through the curriculum, while scratch and Stage 2 initializations stall at early levels.

5. Conclusion

We present the first image-goal navigation policy that enables quadrotors to reach a target viewpoint using onboard images and previous control commands, directly outputting collective thrust and body rates without explicit state estimation. Our approach uses a four-stage training pipeline, progressing from a state-based teacher to a recurrent visual representation learned from foundation model features, followed by imitation learning and reinforcement learning fine-tuning. Our ablations demonstrate the importance of staged training for efficient convergence and successful flight. We evaluated the policy on held-out HM3D buildings and demonstrated zero-shot generalization to the Gibson dataset. Finally, we deployed the policy on a physical quadrotor, achieving an 84% success rate in a previously unseen real-world environment.

Our current setting assumes partial visual overlap between the initial and goal images and does not address exploration. Future work will investigate how to combine goal-directed exploration with direct visual flight control to relax this assumption.

LLM Usage Statement

OpenAI's ChatGPT and Anthropic's Claude were used to assist with language editing, improving manuscript clarity, and identifying relevant references. OpenAI Codex was used to refine figures and assist with code development, debugging, and refactoring. All AI-assisted content and code modifications were reviewed and validated by the authors, who retain full responsibility for the scientific content and experimental results.

Acknowledgments

This work was supported by the European Union's Horizon Europe Research and Innovation Programme under grant agreement No. 101120732 (AUTOASSESS), the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT), and the National University of Singapore.

References

[1] Guillaume Bono, Leonid Antsfeld, Boris Chidlovskii, Philippe Weinzaepfel, and Christian Wolf. End-to-end (instance)-image goal navigation through correspondence as an emergent phenomenon. In *International Conference on Learning Representations (ICLR)*, 2024. 3

[2] Luca Carlone, Ayoung Kim, Timothy Barfoot, Daniel Cremers, and Frank Dellaert, editors. *SLAM Handbook. From Localization and Mapping to Spatial Intelligence*. Cambridge University Press, 2026. 1

[3] Francois Chaumette and Seth Hutchinson. Visual servo control, part i: Basic approaches. *IEEE Robotics & Automation Magazine*, 13(4):82–90, 2006. 3, 6

[4] Dian Chen, Brady Zhou, Vladlen Koltun, and Philipp Krähenbühl. Learning by cheating. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2019. 3

[5] Qianzhong Chen, Jiankai Sun, Naixiang Gao, JunEn Low, Timothy Chen, and Mac Schwager. GRaD-Nav: Efficiently learning visual drone navigation with gaussian radiance fields and differentiable dynamics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7941–7948, 2025. 3

[6] Qianzhong Chen, Naixiang Gao, Suning Huang, JunEn Low, Timothy Chen, Jiankai Sun, and Mac Schwager. GRaD-Nav++: Vision-language model enabled visual drone navigation with gaussian radiance fields and differentiable dynamics. *IEEE Robotics and Automation Letters*, 11(2):1418–1425, 2026. 3

[7] Yijie Deng, Shuaihang Yuan, and Yi Fang. AnyImageNav: Any-view geometry for precise last-meter image-goal navigation. *arXiv preprint*, 2026. 3

[8] Siyan Dong, Shuzhe Wang, Shaohui Liu, Lulu Cai, Qingnan Fan, Juho Kannala, and Yanchao Yang. Reloc3r: Large-scale training of relative camera pose regression for generalizable, fast, and accurate visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3

[9] 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. 3, 7

[10] Ismail Geles, Leonard Bauersfeld, Angel Romero, Jiaxu Xing, and Davide Scaramuzza. Demonstrating agile flight from pixels without state estimation. In *Robotics: Science and Systems (RSS)*, 2024. 1, 2, 3, 7

[11] Johannes Heeg, Yunlong Song, and Davide Scaramuzza. Learning quadrotor control from visual features using differentiable simulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 2, 3

[12] Wonbong Jang, Philippe Weinzaepfel, Vincent Leroy, Lourdes Agapito, and Jerome Revaud. Pow3r: Empowering unconstrained 3d reconstruction with camera and scene priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1071–1081, 2025. 4

[13] Elia Kaufmann, Antonio Loquercio, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Deep drone acrobatics. In *Robotics: Science and Systems (RSS)*, 2020. 3

[14] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. RMA: Rapid motor adaptation for legged robots. In *Robotics: Science and Systems (RSS)*, 2021. 3

[15] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3D with MASt3R. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 3

[16] Pengna Li, Kangyi Wu, Jingwen Fu, and Sanping Zhou. REGNav: Room expert guided image-goal navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025. 3

[17] Yaxuan Li, Jiarui Zeng, Shaofei Huang, and Zhedong Zheng. Last-meter precision navigation for UAVs: A diffusion-refined aerial visual servoing approach. *arXiv preprint*, 2026. 2

[18] Antonio Loquercio, Elia Kaufmann, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Learning high-speed flight in the wild. *Science Robotics*, 6(59): eabg5810, 2021. 3

[19] JunEn Low, Maximilian Adang, Javier Yu, Keiko Nagami, and Mac Schwager. SOUS VIDE: Cooking visual drone navigation policies in a gaussian splatting vacuum. *IEEE Robotics and Automation Letters*, 10(5):5122–5129, 2025. 3

[20] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 6

[21] Mohit Ludhiyani, Vishvendra Rustagi, Ranjan Dasgupta, and Arnab Sinha. Multirotor dynamics based online scale estimation for monocular slam. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 6475–6481. IEEE, 2019. 2

[22] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust

perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7(62), 2022. 3

[23] Gianluca Monaci, Philippe Weinzaepfel, and Christian Wolf. What does really matter in image goal navigation? In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 640–651, 2026. 3

[24] NVIDIA. NVIDIA TensorRT. https://developer.nvidia.com/tensorrt, 2026. Accessed: 2026-11-14. 7

[25] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M. Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X. Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3D dataset (HM3D): 1000 large-scale 3D environments for embodied AI. In *NeurIPS Datasets and Benchmarks Track*, 2021. 6

[26] Angel Romero, Ashwin Shenai, Ismail Geles, Elie Aljalbout, and Davide Scaramuzza. Dream to fly: Model-based reinforcement learning for vision-based drone flight. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2026. 3

[27] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 627–635, 2011. 5

[28] Alessandro Scherl, Stefan Thalhammer, Bernhard Neuberger, Wilfried Wöber, and José García-Rodríguez. ViT-VS: On the applicability of pretrained vision transformer features for generalizable visual servoing. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 17769–17776, 2025. 3

[29] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint*, 2017. 4

[30] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3. *Transactions on Machine Learning Research*, 2026. 4

[31] Xinyu Sun, Peihao Chen, Jugang Fan, Thomas H. Li, Jian Chen, and Mingkui Tan. FGPrompt: Fine-grained goal prompting for image-goal navigation. In *Advances in Neural Information Processing Systems*, 2023. 3

[32] Rui Wang, Stephen M. Pizer, and Jan-Michael Frahm. Recurrent neural network for (un-)supervised learning of monocular video visual odometry and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5555–5564, 2019. 1

[33] Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9068–9079, 2018. 6

[34] Jiaxu Xing, Angel Romero, Leonard Bauersfeld, and Davide Scaramuzza. Bootstrapping reinforcement learning with imitation for vision-based agile flight. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2024. 3

[35] Karmesh Yadav, Arjun Majumdar, Ram Ramrakhya, Naoki Yokoyama, Alexei Baevski, Zsolt Kira, Oleksandr Maksymets, and Dhruv Batra. OVRL-V2: A simple state-of-art baseline for ImageNav and ObjectNav. *arXiv preprint*, 2023. 3

[36] Zichen Yan, Rui Huang, Lei He, Shao Guo, and Lin Zhao. SIGN: Safety-aware image-goal navigation for autonomous drones via reinforcement learning. *IEEE Robotics and Automation Letters*, 11(2):1962–1969, 2026. 2, 6, 7

[37] Zichen Yan, Rui Huang, Shenao Wang, Yucong Wang, Nuthasith Gerdpratoom, and Lin Zhao. Omni-IGN: Onboard image-goal navigation for drones with omnidirectional vision. *arXiv preprint*, 2026. 2

[38] Yifan Zhai, Elia Raimondi, Yunfan Ren, Ismail Geles, Yannick Armati, Jiaxu Xing, and Davide Scaramuzza. Learning to throw: Agile and accurate cable-suspended payload delivery with a quadrotor. *arXiv preprint arXiv:2606.27603*, 2026. 3, 7

[39] Deyi Zhu, Haoyu Fan, Yinan Zhu, Weichen Zhang, Shilin Ma, Xinlei Chen, and Yansong Tang. Uncertainty-aware world model for aerial image-goal navigation. *arXiv preprint*, 2026. 2

[40] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017. 3

End-to-End Image-Goal Navigation for Agile Quadrotors via Recurrent State Distillation

Supplementary Material

A. Stage 4 Reprojection Penalty Term

The additional bounded image-reprojection penalty near the goal introduced in Stage 4 is defined as:

$$r_t^{(4)} = r_t^{(1)} + r_t^{\text{repr}}. \tag{11}$$

Let $\mathcal{G}$ denote the set of goal-image pixels with valid simulator depth. Each pixel $\mathbf{u} \in \mathcal{G}$ is back-projected into 3D, transformed into the current camera frame, and reprojected as

$$\mathbf{u}'t = \Pi\big(Rt \leftarrow g[D_g(\mathbf{u})\mathbf{r}_g(\mathbf{u})] \quad + \mathbf{t}_{t\leftarrow g}\big), \tag{12}$$

where D_g is the distance from the goal camera center, $\mathbf{r}_g$ is the unit viewing ray, and Π denotes calibrated camera projection. The relative camera transform $(R_{t\leftarrow g}, \mathbf{t}_{t\leftarrow g})$ is obtained from simulator poses.

Let $\mathcal{V}_t \subseteq \mathcal{G}$ contain points whose projections fall inside the current image. For nonempty sets, the mean reprojection error and geometric overlap are

$$e_t = \frac{1}{|\mathcal{V}_t|} \sum_{\mathbf{u}\in\mathcal{V}_t} |\mathbf{u}'_t - \mathbf{u}|2, \; o_t \quad = \frac{|\mathcal{V}_t|}{|\mathcal{G}|}. \tag{13}$$

We combine them into an alignment score:

$$A_t = \begin{cases} o_t \exp\left[-(e_t/\sigma)^2\right], & |\mathcal{V}_t| > 0, \\ 0, & \text{otherwise}, \end{cases} \tag{14}$$

where $\sigma = 5$ pixels. The reprojection penalty is then

$$r_t^{\text{repr}} = \begin{cases} -\lambda_{\text{repr}} \left(1 - \frac{d_t^2}{\rho^2}\right)^2 (1 - A_t), & \substack{d_t<\rho,\\ |\mathcal{G}|>0,} \\ 0, & \text{otherwise}, \end{cases} \tag{15}$$

where $d_t = |\mathbf{p}_t - \mathbf{p}_g|2$, $\rho = 0.5, \text{m}$, and $\lambda\text{repr} = 0.02$. The penalty is bounded in $[-0.02, 0]$, vanishes at perfect alignment, and smoothly decreases to zero at the activation boundary. When no goal points project into the current image, the maximum distance-gated penalty is applied, while missing goal depth produces no penalty.

Geometric overlap is computed without checking occlusions in the current view. Simulator geometry and camera poses are used only for reward computation during training and are not required by the deployed policy.

B. Real World Deployment

Table 3 reports detailed information on each individual real world sequence.

Table 3. Results for the 25 real-world trials. Errors are measured relative to the saved goal using a motion capture system. Orientation error is the shortest full 3D rotation angle. Time runs from policy activation to termination, including warmup and excluding recovery.

			Initial error		Final error	
Run	Outcome	Time (s)	Position (m)	Orientation (°)	Position (cm)	Orientation (°)
0	✓	6.22	3.09	39.99	11.95	2.55
1	✓	13.65	3.38	3.37	13.00	4.87
2	✓	10.97	3.71	69.41	7.23	1.57
3	× (crash)	1.05	4.24	23.48	350.36	42.08
4	✓	20.20	3.17	17.19	7.47	5.85
5	✓	6.07	2.86	45.06	10.36	2.38
6	× (crash)	2.17	3.53	4.05	211.95	20.33
7	✓	21.27	2.17	10.16	14.90	3.68
8	✓	7.07	3.27	15.84	8.10	2.02
9	✓	19.59	2.45	5.19	12.10	1.90
10	✓	8.21	2.97	36.33	14.88	4.70
11	✓	6.94	3.50	34.69	13.29	6.31
12	✓	9.27	2.96	28.14	14.97	4.10
13	✓	7.01	2.86	12.82	8.48	1.37
14	× (crash)	1.10	2.97	0.94	228.41	17.28
15	✓	8.52	1.88	3.96	12.43	2.50
16	× (timeout)	86.52	2.71	20.83	38.56	8.32
17	✓	26.66	2.78	5.32	14.66	3.74
18	✓	6.66	3.00	8.82	10.56	1.19
19	✓	7.71	3.05	2.98	12.60	2.66
20	✓	7.33	2.53	22.74	14.92	1.48
21	✓	11.06	2.69	20.04	12.53	2.99
22	✓	7.63	3.04	15.91	13.49	2.42
23	✓	7.56	3.40	11.37	11.74	3.33
24	✓	7.28	3.23	15.15	8.03	2.92

C. Ablations: Stage 3

We report here the plots comparing Stage 3 trained from scratch, without the visual encoder and actor pretraining, and the full pipeline.

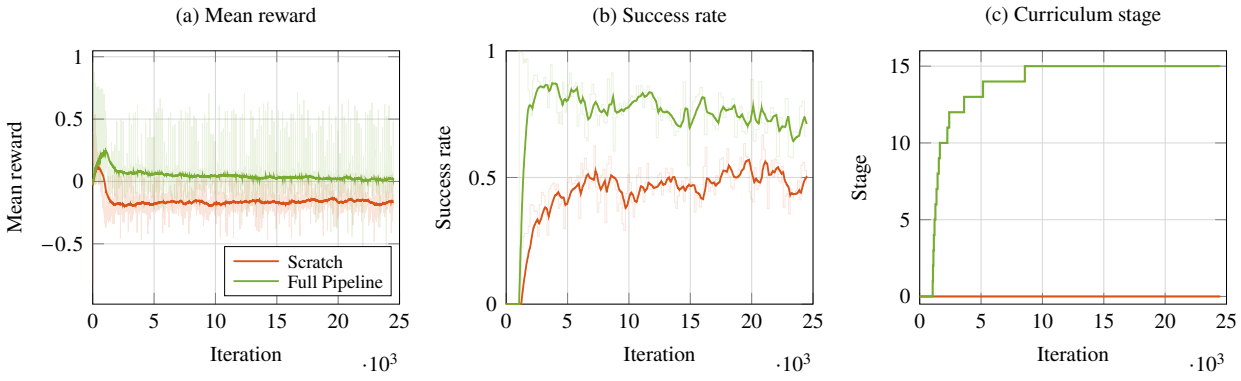

Figure 6. Stage 3 initialization ablation. The full pipeline progresses through the curriculum, while freshly initialized run does not pass over the first curriculum level.

D. Cross-Attention Visualization

To better understand the learned visual representation, we visualize attention weights between current and goal image patches. Among the inspected heads, the third head of the first cross-attention layer exhibits spatially localized responses compatible

with patch correspondences. Figure 7 compares these attention maps with cosine similarities computed from frozen DINOv3 features across three HM3D validation sequences.

For each selected query patch, we visualize its attention weights over the opposite image, using Layer 1, Head 3. The attention operates on 14×14 patch grids with a head dimension of 64. For comparison, we compute cosine similarities between the frozen DINOv3 patch descriptors, before the learned adapter and attention projections. Both maps are independently min-max normalized to $[0, 1]$ for visualization. Attention weights represent learned feature aggregation rather than explicit correspondence probabilities.



(a)

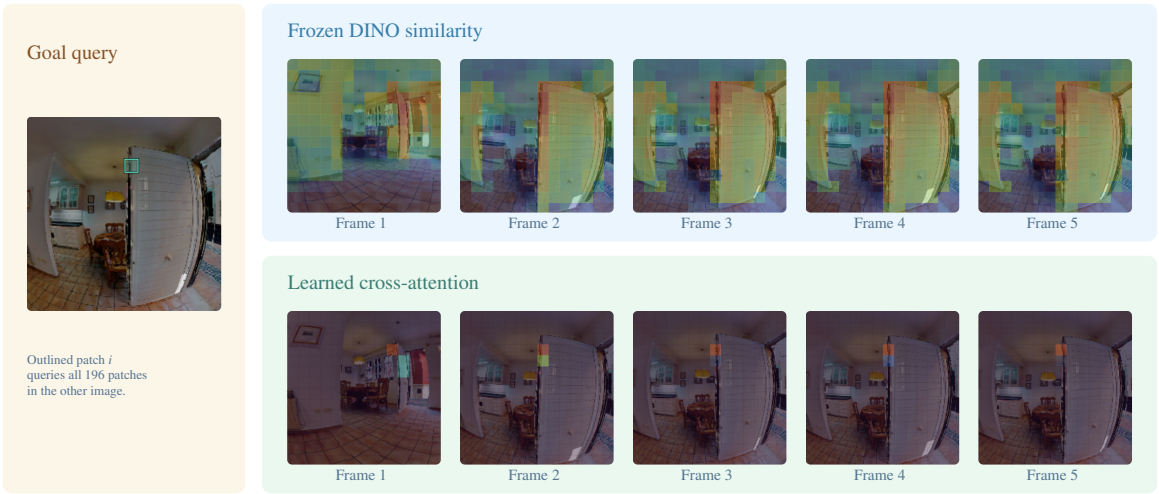

(b)

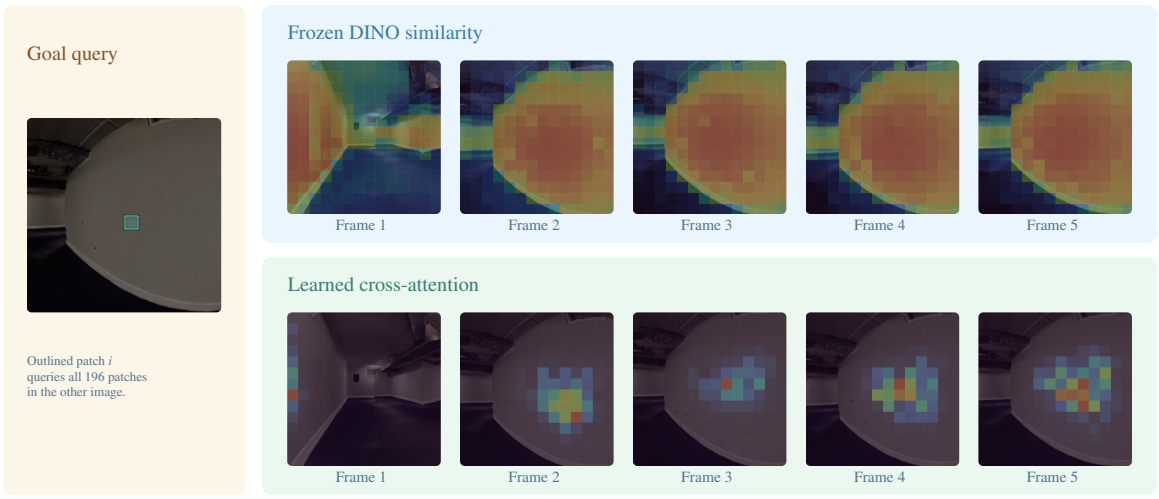

(c)

Figure 7. Cross-attention visualization on three HM3D validation sequences. (a)–(c) compare frozen DINOv3 cosine similarity with learned cross-attention from Layer 1, Head 3. Heatmaps are independently normalized for visualization.

E. Trajectory visualizations

Figures 8 and 9 show the first real-world flight and four HM3D validation trajectories, respectively. Each plot shows the 3D trajectory above the goal, initial, and final images and three timestamped intermediate observations.

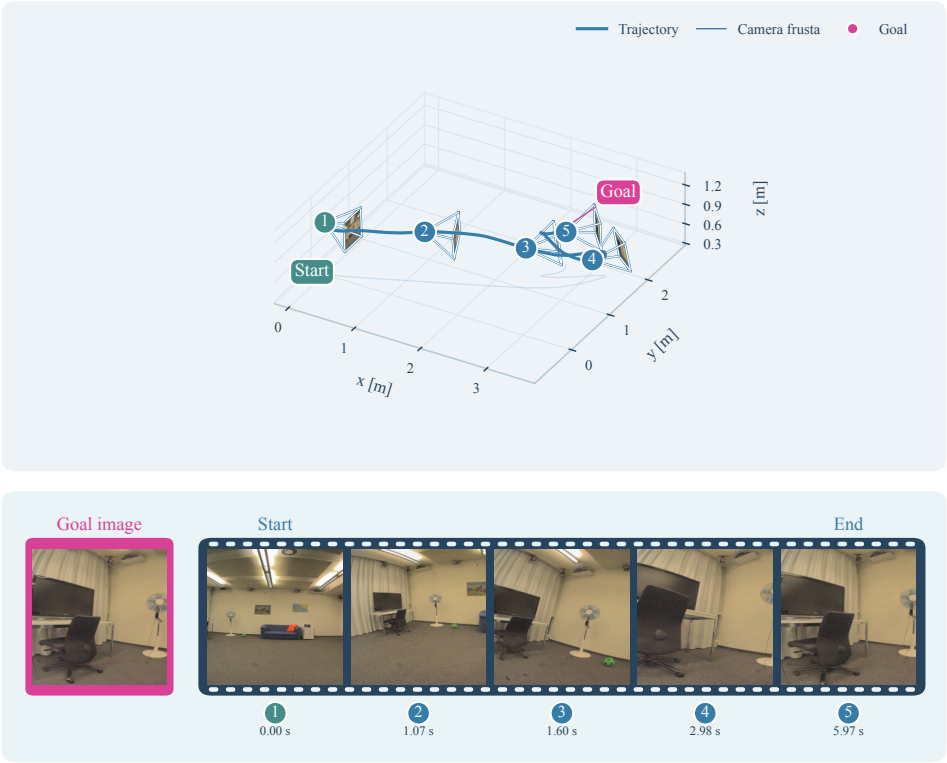

Figure 8. Qualitative result on a real-world flight sequence. The 3D trajectory is shown above, with the goal, initial, final, and three timestamped intermediate frames below.



(a)

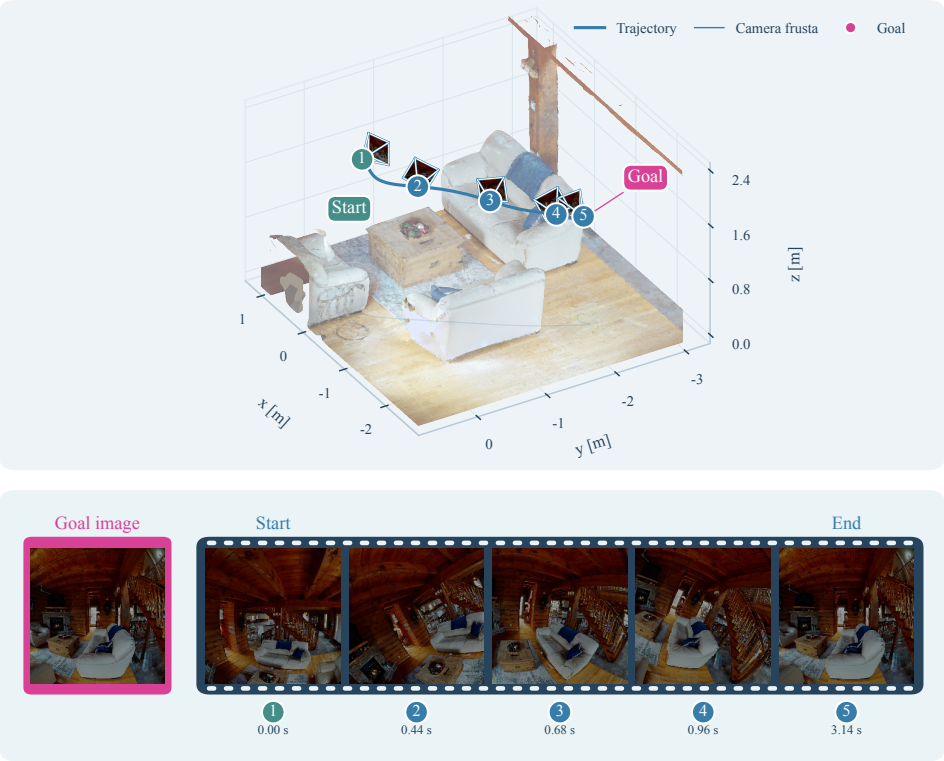

(b)

Figure 9. Qualitative results on four HM3D validation sequences. Each example shows the 3D trajectory above and the goal, initial, final, and three timestamped intermediate frames below.

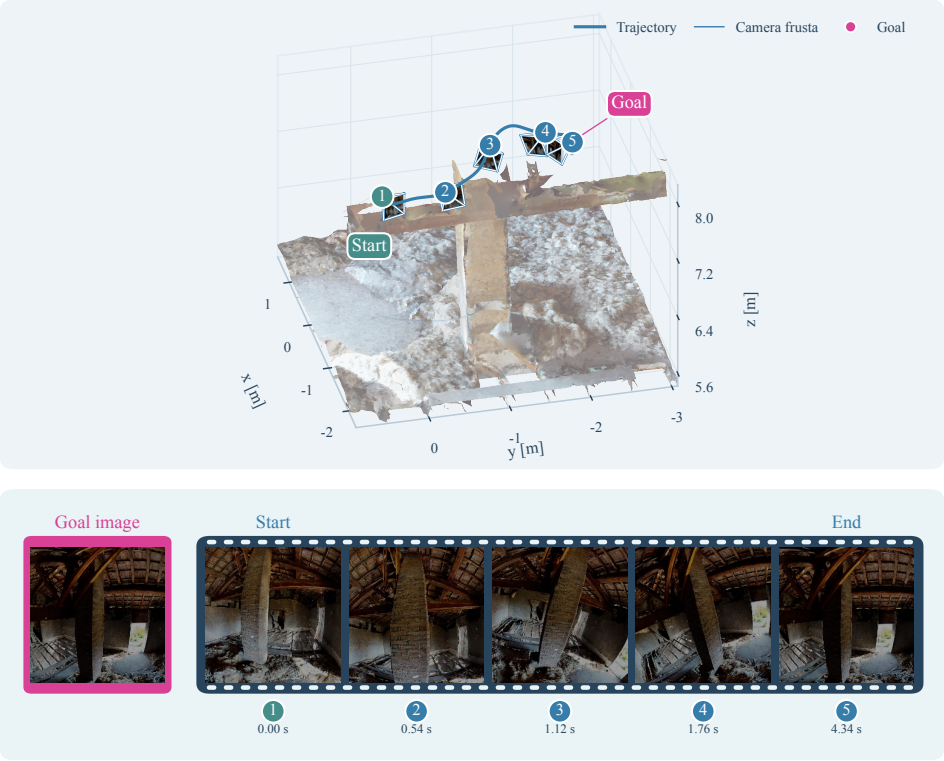

(c)



(d)

Figure 9. HM3D validation sequences (continued).