

# FalconGym: A Photorealistic Simulation Framework for Zero-Shot Sim-to-Real Vision-Based Quadrotor Navigation

Yan Miao<sup>1</sup>, Will Shen<sup>1</sup> and Sayan Mitra<sup>1</sup>

**Abstract**—We present a novel framework demonstrating zero-shot sim-to-real transfer of visual control policies learned in a Neural Radiance Field (NeRF) environment for quadrotors to fly through racing gates. Robust transfer from simulation to real flight poses a major challenge, as standard simulators often lack sufficient visual fidelity. To address this, we construct a photorealistic simulation environment of quadrotor racing tracks, called *FalconGym*, which provides effectively unlimited synthetic images for training. Within *FalconGym*, we develop a pipelined approach for crossing gates that combines (i) a Neural Pose Estimator (NPE) coupled with a Kalman filter to reliably infer quadrotor poses from single-frame RGB images and IMU data, and (ii) a self-attention-based multi-modal controller that adaptively integrates visual features and pose estimation. This multi-modal design compensates for perception noise and intermittent gate visibility. We train this controller purely in *FalconGym* with imitation learning and deploy the resulting policy to real hardware with no additional fine-tuning. Simulation experiments on three distinct tracks (circle, U-turn and figure-8) demonstrate that our controller outperforms a vision-only state-of-the-art baseline in both success rate and gate-crossing accuracy. In 30 live hardware flights spanning three tracks and 120 gates, our controller achieves a 95.8% success rate and an average error of just 10 cm when flying through 38 cm-radius gates.

## I. INTRODUCTION

Autonomous vision-based quadrotor control is crucial for a range of high-impact applications, including racing, search-and-rescue, and exploration. Several challenges exist for visual flight. First, visual control is often fragile: Geles et al. [1] demonstrated impressive quadrotor racing performance with a vision-only actor-critic approach, yet their system degraded when gates were out of view for multiple frames, posing safety risks. Second, sim-to-real transfer remains difficult [2] because visual control policies trained in conventional simulators may not generalize reliably to the real world due to differences in visual features and dynamics.

Recent advances in neural scene representations, such as Neural Radiance Fields (NeRF) [3] and Gaussian Splatting [4], provide an opportunity for creating photorealistic simulation environments, which can be used for training visual control policies. Although prior works have shown promising zero-shot neural-scene to real transfer for humanoid robots [5] and grasping [6], few have focused on quadrotors, which demand higher speeds and control frequencies than ground robots. SOUS VIDE [7] is a notable exception, which uses Gaussian Splat [4] for neural scene representation and achieves impressive zero-shot drone navi-




Fig. 1: Our autonomous quadrotor, equipped with an ArduCam RGB camera and an IMU-enabled Navio2 on a Raspberry Pi 3, can successfully fly through a sequence of small 38 cm-radius gates with an average of only 10 cm error to gates center.

gation by distilling an expert MPC controller to a lightweight end-to-end policy.

In this paper, we address the sim-to-real gap for quadrotor gate crossing by developing a photorealistic simulation environment, *FalconGym*, and designing a new visual control model for quadrotors using it. Within *FalconGym*, we design a Neural Pose Estimator (NPE), which combined with a Kalman filter, to produce pose estimates. We then propose a self-attention-based multi-modal control policy that, unlike conventional fusion, can dynamically reweight the vision features and pose estimation inputs based on gate visibility to enhance robustness against perception errors. Finally, we train our multi-modal controller in *FalconGym* via imitation learning and deploy the resulting policy on a real quadrotor without any additional fine-tuning. In *FalconGym*, our controller completes ten laps on each of the three four-gate tracks (circle, U-turn, and figure-8) with a 100% success rate across all tracks and an average gate-crossing error of 6.3, 10.1, and 5.1 cm for 38 cm-radius gates. In real-world experiments using the same tracks and lap count, the success rate and precision decrease slightly but remain notable: 100%, 87.5%, and 100% success on the circle, U-turn, and figure-8 tracks with average errors of 10.7, 11.1, and 9.4cm. This performance is impressive given the relatively large size of our quadrotor (30 cm × 34 cm × 11 cm), which leaves little margin for error for 38 cm-radius gates.

In summary, our contribution is threefold. First, we

<sup>1</sup>Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, USA

demonstrate a novel framework to achieve zero-shot sim-to-real transfer of NeRF-trained visual control policies for quadrotors navigating through a sequence of small gates. Second, we develop *FalconGym*, a photorealistic simulation environment for agile, vision-based quadrotor flight, providing an open benchmark for visual flight control research. Third, we make two innovations that make our neural controller’s architecture practical: (i) A one-shot *Neural Pose Estimator (NPE)* coupled with a Kalman filter; (ii) A self-attention-based sensor fusion scheme that integrates vision features and pose estimation for quadrotor flight.

## II. RELATED WORK

*a) Neural Scene Representations in Robotics:* Neural scene representations like NeRF [3] and Gaussian Splat [4] have recently emerged as powerful tools for photorealistic 3D scene reconstruction. Variants and extensions have been applied to diverse robotics tasks, including 3D scene editing [8], pose estimation [9], grasping [10] and navigation [11]. NeRF2Real [5] and RialTo [6] demonstrate the potential sim-to-real transfer for ground robots and robot arm manipulation. [7] achieves zero-shot drone navigation using policies trained in Gaussian splat scenes. In addition to using NeRF representations for training, our work also uses a completely different solution architecture than [7] and we solve the navigation problem with harder constraints in the shape of 38 cm-radius gates.

*b) Sim-to-Real Transfer in Robotics:* Sim-to-real transfer has long been a cornerstone of robotics research, valued for its efficient and risk-free training. Traditional methods such as domain randomization [12] and system identification [13] are commonly employed to narrow the gap between simulation and reality. These techniques and their variations have enabled successful sim-to-real deployments in quadruped locomotion [14], bipedal walking [15], and autonomous driving [16].

*c) Visual Flight:* Visual flight has attracted considerable interest recently. Kaufmann et al. [17] combined convolutional neural networks with advanced path-planning to achieve drone racing capabilities, while Geles et al. [1] extended agile flight racing to an end-to-end image-based policy, with an impressive speed upto 40 km/h. However, their visual control’s performance deteriorates when racing gates are out of view for extended periods. Other approaches have explored learning aggressive acrobatic flying maneuvers through imitation learning [18] or through carefully simulated sensor noise for robust outdoor flight [19]. Further progress has been made in autonomous flight to match human champion performance with Visual Inertial Odometry (VIO) and deep reinforcement learning with onboard sensors [20].

*d) Sensor Fusion:* Robots frequently integrate multiple sensing modalities to improve accuracy and resilience. In prior work, [7] [18] combine IMU measurements with visual features through a MLP for flight control. Recent sensor fusion approaches have also embraced attention-based architectures [21], including fusing camera and LiDAR signals for 3D perception [22], or employing self-attention

on electromyography (EMG) and IMU data for lower-limb locomotion control [23], fusing goal image with current observation through transformers for better navigation [24]. However, most of these advances have been focused on ground robots rather than aerial systems.




Fig. 2: *FalconGym* can render photorealistic images from different poses. The left column shows four real world images captured by the Arducam in our flight area, while the right column displays images rendered by *FalconGym* at the same coordinates.

## III. METHODOLOGY

In this paper, we address the challenge of enabling a vision-centric quadrotor to autonomously fly through a series of small 38 cm-radius gates while adhering to hard safety constraints—specifically, avoiding collisions and staying on track, as shown in Figure 1 and Figure 4. Our objective is to train the controller entirely in simulation and zero-shot transfer the resulting control policy directly to real-world environments without additional fine-tuning.

In the following subsections, we first describe how we construct our *FalconGym* simulation environments that are aligned with the real-world coordinate system. We then introduce our perception module, which leverages a Neural Pose Estimator (NPE) to infer the quadrotor’s state from onboard camera images, where its estimate is further improved by fusing high-frequency IMU data via a Kalman filter [25]. Lastly, we detail our controller design, which fuses visual features and pose estimation with a self-attention mechanism and is trained via DAGGER-style [26] imitation learning.

### A. *FalconGym* Construction

We define a *track* as a sequence of gates arranged in various poses within a designated flight area, as illustrated in Figure 4. For each track, we want to build a *FalconGym*  $F$  that can render a photorealistic image  $v$  from any virtual camera pose  $p$ , i.e.,  $v = F(p)$ .

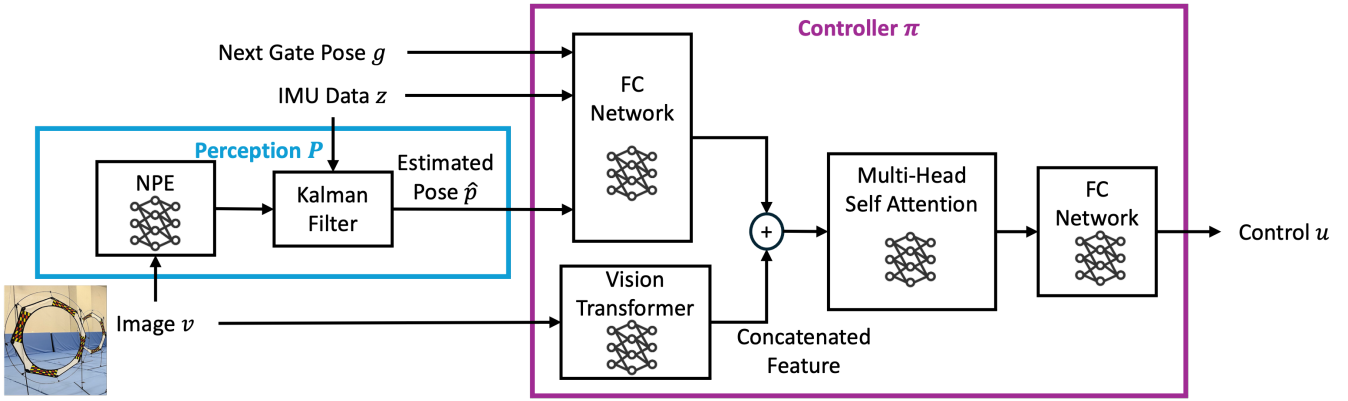


Fig. 3: The architecture of our visual flight system consists of a perception module and a control module. In the blue box, the perception module processes onboard camera images using a ViT-based Neural Pose Estimator (NPE) and improves its estimates with IMU data via a Kalman filter. The purple box details the control module, where a self-attention mechanism fuses state embeddings with vision features.

To train our FalconGym simulation environment, we first collect a dataset  $\mathcal{D}_F = \{(v_i, p_i)\}_{i=1}^N$  where  $v_i \in \mathbb{R}^{w \times h \times 3}$  is an RGB image taken at camera pose  $p_i \in \mathbb{R}^6$ , which comprises 3D position and orientation. A human operator walks around each track for approximately five minutes, capturing  $N \sim 1500$  images from diverse viewpoints with an ArduCam RGB camera of known intrinsic parameters. Simultaneously, a localization marker is mounted on top of the camera and tracked by a Vicon<sup>®</sup> motion capture system. This system yields sub-millimeter precision ( $\approx 0.15mm$  mean error [27]) and eliminates the need for Structure from Motion (SfM) pipelines such as COLMAP [28], which typically introduce larger alignment and scale uncertainties than a motion capture system. All sensor data, including camera images and pose information, are synchronized via ROS.

After approximately 15 minutes of training  $\mathcal{D}_F$  with NeRF [3] on an NVIDIA RTX 4090 using the open-source nerfstudio library [29], we obtain a FalconGym model capable of generating photorealistic images for any novel camera poses. Figure 2 provides a qualitative comparison between real-world images and those generated by FalconGym.

### B. Perception Module: NPE with Kalman Filter

State estimation is essential for stable quadrotor flight, as it provides crucial pose feedback needed by the controller. In this work, we adopt a standard RGB camera (Arducam), popular in quadrotor setups, rather than more expensive or computationally demanding tracking cameras or LiDARs.

To estimate camera pose directly from an RGB image, we propose a *Neural Pose Estimator (NPE)*, as shown in the blue box in Figure 3. An iterative approach like iNeRF [9] can also estimate camera pose by optimizing the pose to match a target image through gradient-based refinement, but that process is too slow for real-time flight. Instead, we train a single-shot neural network NPE to estimate pose.

For the backbone of NPE, we use a pretrained Vision Transformer (ViT) [30] on ImageNet-21k [31], freezing its early layers to leverage robust, generic feature representations, then we attach a learnable regression head to predict

camera pose. Instead of directly regressing orientation (pitch, roll, and yaw), we predict their sine and cosine values to avoid discontinuities associated with circular data at  $2\pi$ , and later converted to orientation using a trigonometric transformation.

A key advantage of neural scene representation is the ability to generate large synthetic datasets without the overhead of real-world data collection. For each of our *FalconGym* simulation environment, we randomly sample  $M \sim 10k$  camera poses  $p_j$  in the workspace and render corresponding images  $v_j = F(p_j)$  using the previously trained FalconGym to form a dataset  $\mathcal{D}_N = \{(v_j, p_j)\}_{j=1}^M$ . We then train NPE to learn the mapping  $v \mapsto p$ .

Despite this rich synthetic dataset, single-frame predictions may still be noisy under challenging visual conditions (e.g., motion blur or ambiguous viewpoints), leading to abrupt frame-to-frame jumps. To mitigate such transient errors, we fuse the NPE estimates with IMU data  $z \in \mathbb{R}^6$  (representing three-dimensional linear and angular accelerations) in a Kalman filter, assuming a double integrator model for the quadrotor’s dynamics. This fusion strategy not only smooths out pose estimates, but also could accommodate asynchronous sensor rates in hardware experiments: if the camera runs at a lower frame rate, the Kalman filter can still update state estimates from the IMU until the next image arrives.

We model the entire perception pipeline, which consists of NPE followed by the Kalman filter, as a function  $P$ ; that is, given an input image  $v$  and IMU data  $z$ , the model outputs an estimated pose  $\hat{p} = P(v, z)$ . In practice, we observe that it takes about 1 hour to train the NPE model, and the entire perception pipeline has an average of  $7ms$  inference time on a NVIDIA RTX 4090, making the proposed perception pipeline feasible for real-world deployment.

### C. Controller Architecture: Self-Attention

Although one could directly feed the estimated state to a state-based controller (as explored in Section IV-D), large perception errors could affect downstream controller and lead to constraints violations like collision or off-track.




Fig. 4: Trajectory plots for three tracks: circle (left), U-turn (middle), and figure-8 (right). The top row shows the real-world gate setups in our indoor flight area, the middle row visualize 10-lap trajectories in FalconGym using our Multi-Modal (MM) controller, and the bottom row displays trajectories from real hardware experiments, where the same controller was deployed without additional fine-tuning. In FalconGym, our MM controller achieves a 100% gate crossing success rate across all tracks. Although real-world trajectories exhibit greater variability, the MM controller still achieves success rates of 100%, 87.5%, and 100% on the circle, U-turn, and figure-8 tracks, respectively. Red crosses indicate points where the quadrotor collided with a gate.

Inspired by residual connections in deep networks [32] that can improve model accuracies and attention mechanism [21] in fusing different inputs, we instead design a Multi-Modal (MM) controller that fuses visual feature and state information with self-attention, as illustrated in the purple box of Figure 3.

For the visual branch of the MM controller, we employ another pre-trained, partially frozen ViT [30] to process raw RGB images. The output of its final layer is passed through a trainable head that projects the visual features into a 256-dimensional embedding, capturing essential gate-related vision cues. In parallel, on the state branch of the MM controller, we feed the estimated state (from Section III-B), IMU data, and the next gate’s ground truth pose (obtained when we arranged the track) into a Fully Connected (FC) network, obtaining a second 256-dimensional embedding.

To effectively leverage the visual and state modalities, we apply a multi-head self-attention layer to the concatenated visual and state embeddings. This mechanism enables the network to adaptively focus on whichever branch (visual or state-based) is most relevant in a given context. For instance, when the gate is clearly visible, the attention mechanism can emphasize image features, whereas in poor gate visibility or occlusion scenarios, it can learn to rely more on IMU data. Finally, another trainable FC network outputs the control  $u \in \mathbb{R}^3$  (representing the acceleration in three directions), which

are then fed into closed-loop control in both FalconGym simulation environments and real-world experiments.

Yaw is handled separately by a proportional controller, since the yaw estimates from our perception module is relatively accurate (see Table I and Figure 5).

#### D. Training via Imitation Learning (DAGGER)

Our objective is to train the Multi-Modal (MM) controller with self-attention mechanism to cross gates on various tracks, by learning from a teacher in a DAGGER-style imitation learning framework [26].

1) *Expert Controller*: We begin by implementing an expert state-based controller [33] that uses precise ground-truth states and gate poses to produce expert actions for gate navigation, i.e.  $u_t^* = \pi^*(p_t, z_t, g_t)$ , where  $p_t$  is the ground-truth camera (quadrotor) pose,  $z_t$  is the IMU data,  $g_t \in \mathbb{R}^6$  is the next gate center’s pose, and  $u_t^*$  is the expert acceleration control generated by the expert controller  $\pi^*$  at timestep  $t$ .

During each iteration, we inject carefully-designed mild random noise into expert’s control output, i.e.  $u_t^{*'} = u_t^* + \mathcal{N}(0, \sigma^2)$ , to encourage the quadrotor to visit states beyond its nominal trajectory, yet still remain capable of crossing gates. This practice, akin to domain randomization [19] and DAGGER [26], broadens the data distribution and exposes the policy to less-frequently visited states.

2) *Quantile-Based Perception Error Modeling*: Inspired by the “perception contract” concept [34], [35], we explic-




Fig. 5: Qualitative evaluation of the perception module  $P$  is shown. Blue dots represent a subset of  $\mathcal{D}_Q$  (see Section III-D), which compare the pose estimation from perception module  $P$  against the ground truth pose. The black line indicates perfect estimation, while the orange shaded area denotes the 10–90% quantile range of perception errors, which is used to augment the controller training dataset  $\mathcal{D}_C$ . The close clustering of blue dots around the black line demonstrates that our perception module achieves accurate pose estimation.

itly model the error between ground-truth and estimated states. First, we employ the expert controller (with domain randomized control) in FalconGym, to collect a dataset  $\mathcal{D}_Q = \{(\hat{p}_t, p_t)\}_{t=1}^K$  where  $\hat{p}_t = P(F(p_t), z_t)$  is the pose estimation by NPE and Kalman filter (Section III-B). We gathered  $K \sim 10k$  over 500 runs in FalconGym. Next, we apply quantile regression [36] to  $\mathcal{D}_Q$  to estimate the conditional 10th and 90th percentiles of the estimation error distribution, denoted  $Q_{0.1}(p)$  and  $Q_{0.9}(p)$ . These quantiles define an empirical 80% prediction interval for the error term  $\hat{p} - p$  such that

$$\mathbb{P}\left(Q_{0.1}(p) \leq \hat{p} - p \leq Q_{0.9}(p) \mid p\right) \geq 0.8$$

as demonstrated experimentally in Figure 5.

Finally, rather than assigning a single estimated pose to each ground-truth pose, for each  $p$ , we randomly sample 50 *perturbed* pose  $\tilde{p}$  within this interval:

$$\tilde{p} = p + \text{Unif}\left(Q_{0.1}(p), Q_{0.9}(p)\right).$$

The rationale is that, for a given true pose  $p$ , multiple plausible estimates  $\hat{p}$  might exist, yet all of which should lead to the same expert action. By data augmentation using this quantile range, we prepare our controller for perception uncertainties.

3) *Controller Dataset Collection.*: By above augmentation, we form the final controller dataset with size of  $R \sim 500k$  samples:  $\mathcal{D}_C = \{(v_t, \tilde{p}_t, z_t, g_t), (u_t^*)\}_{t=1}^R$ . where  $v_t = F(p_t)$  is FalconGym-rendered image at ground-truth pose  $p_t$ ,  $\tilde{p}_t$  is the sampled pose estimation from the above quantile range,  $z_t$  is IMU data,  $g_t$  is the next gate’s pose, and  $u_t^* = \pi^*(p_t, z_t, g_t)$  is the expert action.

4) *Multi-Modal Controller Learning.*: We train our MM controller  $\pi(v_t, \tilde{p}_t, z_t, g_t)$  with supervised learning by minimizing the squared error between its predicted acceleration and the expert action:

$$\min_{\pi} \sum_{t=1}^R \left\| \hat{u}_t - u_t^* \right\|^2, \quad \text{where } \hat{u}_t = \pi(v_t, \tilde{p}_t, z_t, g_t).$$

The training takes about 10 hours on a 4090 GPU to finish due to large amount of training data, but the inference time of the controller is only about 11ms.

## IV. EXPERIMENTS

In this section, we first describe our hardware setup and track configuration. We then evaluate the proposed approach in the previous section using well-established performance metrics and compare it with state-of-the-art baseline.

### A. Closed-Loop System Setup in Hardware

We conduct our hardware experiments on a quadrotor measuring 30cm  $\times$  34cm  $\times$  11cm, equipped with a Raspberry Pi 3<sup>®</sup> and a Navio2<sup>®</sup> board for flight control, as shown in Figure 1. The Navio2<sup>®</sup> incorporates an onboard IMU and streams acceleration data at up to 50hz. An ArduCam global-shutter camera is also mounted on the quadrotor, in order to reduce both computation and communication overhead, we operate the camera at 20Hz and downsample the image size to 320  $\times$  240.

All sensor data, including camera images and IMU measurements, are published to ROS topics on the Pi. An offboard workstation, connected on the same subnet as the Pi and equipped with an NVIDIA RTX 4090 GPU, subscribes to these topics. Then we perform the perception (Section III-B) and controller inference (Section III-C) offboard, and finally sends the computed control actions back to the Pi via another ROS topic. Empirically, we measure an end-to-end latency of about 39ms between when sensor data are sent and when control commands are received. This allows us to maintain a stable 20Hz (50ms delay) control loop in the hardware. To be consistent, we also use a  $\Delta t = 0.05s$  step size (i.e. 20Hz) for training in FalconGym as well.

Our indoor flight area measures 6m  $\times$  6m  $\times$  3m, given this limited operational area, the expert controller is configured to fly at a constant velocity of 1m/s for safety reason. Empirically, we observe that at this speed, the quadrotor’s motion closely follows a double integrator model, justifying its use as the assumed dynamics in FalconGym. We design three distinct tracks within the flight area, each comprising

TABLE I: Evaluation of the perception module. The table compares the Neural Pose Estimator (NPE) used alone versus with Kalman filtering, and evaluates perception module’s performance of zero-shot sim-to-real transfer.

Track	FalconGym				Real World			
	NPE		NPE + Kalman Filter		NPE		NPE + Kalman Filter	
	Pose Err. [cm]↓	Yaw Err. [deg]↓	Pose Err. [cm]↓	Yaw Err. [deg]↓	Pose Err. [cm]↓	Yaw Err. [deg]↓	Pose Err. [cm]↓	Yaw Err. [deg]↓
Circle	33.9	4.7	<b>21.8</b>	<b>4.1</b>	34.9	6.3	<b>22.1</b>	<b>6.0</b>
U-turn	55.5	11.0	<b>32.3</b>	<b>7.6</b>	47.5	19.8	<b>33.6</b>	<b>12.6</b>
Figure-8	42.4	7.7	<b>24.0</b>	<b>6.4</b>	44.3	12.4	<b>26.2</b>	<b>9.0</b>

four gates with 38 cm inner radius, of diverse geometries and sharp turns (namely circle, figure-8, and an acyclic U-turn), as shown in Figure 4. During track setup, we pre-record the precise poses of all gates to ensure accurate gate information for both simulation and hardware experiments.

Although a Vicon<sup>®</sup> motion-capture system is available, as noted in Section III-A, it is used strictly for logging the quadrotor’s trajectories and identifying the next gate center based on the quadrotor’s live position. The next gate’s pose information is relayed to the offboard workstation through ROS as another input to the MM controller, as shown in Figure 3, but our controller does not use any ground truth quadrotor pose information from Vicon<sup>®</sup>.

### B. Validation of FalconGym

In Figure 2, we qualitatively demonstrated the realistic visual fidelity of our FalconGym. We now provide a quantitative analysis using three standard metrics for neural scene representation: peak signal-to-noise ratio (PSNR) which measures the pixel-wise reconstruction quality, structural similarity index (SSIM) [37] which gauges structural similarity, and learned perceptual image patch similarity (LPIPS) [38] that captures perceptual differences by comparing feature embeddings.

For each track, we compare hundreds of the real ArduCam-captured images against the images generated by our FalconGym. As summarized in Table II, the reconstructed scenes achieve an average PSNR of 28.3 dB, SSIM of 0.87, and LPIPS of 0.27 across the three tracks. These values are consistent with those reported by the original NeRF work on a real dataset [3]. Consequently, our FalconGym serves as a solid foundation for the subsequent perception and control modules.

### C. Perception Module Evaluation

Next, we evaluate the perception module from Section III-B to answer two main questions: (i) Does Kalman filtering reduce NPE estimation error? (ii) Can a perception model trained entirely in our FalconGym maintain similar performance when transferred to real world?

We measure two key error metrics: the error between the predicted and ground-truth pose, and the error between the predicted and ground-truth yaw. To isolate perception performance from the controller, we use the expert state-based controller [33] to fly 10 laps on each track from slightly different initial poses, both in FalconGym and in hardware

experiments. Over these trials, we record ground-truth poses, raw NPE estimates and Kalman-filtered estimates.

Table I summarizes these perception results. We see that, in FalconGym, the NPE alone achieves relatively accurate yaw prediction. However, NPE’s position error in FalconGym averages 43.9cm, which poses a risk for gates of only 38cm radius. Once fused with the Kalman filter, we see a significant drop in pose estimation error (26 cm), given that IMU data compensate for some perception error. Comparing FalconGym and real world data, we observe slightly higher perception errors in the real-world setting (an average of 27.3cm) than FalconGym (an average of 26cm), likely reflecting minor domain gaps between rendered FalconGym images and actual camera feeds.

Overall, these findings confirm that Kalman filter enhances state estimation and the overall perception module has a similar sim-to-real performance.

TABLE II: Evaluation of FalconGym Photorealism quality.

Track	# Images	PSNR(dB)↑	SSIM↑	LPIPS↓	Time
Circle	1583	28.3694	0.8720	0.2689	12m 49s
U-turn	1094	28.3205	0.8677	0.2929	12m 33s
Figure-8	839	28.3691	0.8852	0.2665	12m 48s

### D. Controller Evaluation

Following the visualized gate-crossing trajectories shown in Figure 4, we next present a quantitative analysis that addresses two main questions: (i) Does our Multi-Modal (MM) controller outperform a vision-only state-of-the-art (SoTA) baseline? (ii) Can our MM approach transfer effectively from FalconGym to real hardware?

We quantify how well the quadrotor controller upholds the hard constraints (avoiding collisions or going off-track) using two metrics: (i) *Success Rate (SR)* measures the percentage of gates that the quadrotor safely flies through; (ii) *Mean Gate Error (MGE)* evaluates the average distance between the quadrotor and the gate center at the time of crossing.

We implemented four different controllers for our experiments: (i) Expert (State-based): A controller from Sun et al. [33] with full ground-truth state access; (ii) SoTA vision-only Baseline: A method adapted from Geles et al. [1], where a privileged critic observes the full state during training, while the actor relies on images and past actions at inference and do not need IMU; (iii) Direct Perception (DP): our naive method that directly feeds our perception output from NPE and Kalman filter into the state-based

controller; (iv) Multi-Modal(MM): our approach which fuses image embeddings and state embeddings via self-attention

To remain faithful to the baseline’s setup, we apply the same gate-detection method of [1] (HIL experiments) in simulation that projects known 3D gate positions into 2D image coordinates using the camera’s intrinsics and extrinsics. In addition, because the baseline struggles when gates disappear for consecutive frames, we further add a yaw-control heuristic to keep the quadrotor facing the next gate whenever possible.

Each of the four controllers is run for 10 laps on the circle, U-turn, and figure-8 tracks (4 gates per track) with slightly varied initial poses in FalconGym. DP and the baseline are excluded from the real world hardware experiments, as their low simulation success rates raise safety concerns.

TABLE III: Evaluation of controllers by Success Rate (SR) and Mean Gate Error (MGE) in FalconGym and real world.

	Method	FalconGym		Real World	
		SR $\uparrow$	MGE [cm] $\downarrow$	SR $\uparrow$	MGE [cm] $\downarrow$
Circle	State-based	100%	2.49	100%	6.22
	Baseline [1]	100%	6.62	n/a	n/a
	Ours (DP)	90%	16.0	n/a	n/a
	<b>Ours (MM)</b>	<b>100%</b>	<b>6.25</b>	<b>100%</b>	<b>10.7</b>
U-turn	State-based	100%	3.42	100%	9.80
	Baseline [1]	100%	11.9	n/a	n/a
	Ours (DP)	52.5%	11.3	n/a	n/a
	<b>Ours (MM)</b>	<b>100%</b>	<b>10.1</b>	<b>87.5%</b>	<b>11.1</b>
Figure-8	State-based	100%	5.11	100%	9.30
	Baseline [1]	25%	8.44	n/a	n/a
	Ours (DP)	20%	16.2	n/a	n/a
	<b>Ours (MM)</b>	<b>100%</b>	<b>5.13</b>	<b>100%</b>	<b>9.40</b>

Table III summarizes the results. In FalconGym, we see that MM controller consistently achieves higher SR (100%) and lower MGE than DP, primarily due to sensor fusion compensating for inaccurate perception. The baseline also attains 100% SR on circle and U-turn tracks but has slightly worse MGE than our MM controller and experiences a drop in SR on the figure-8 track, where gates remain out of view for extended periods. By contrast, our MM controller maintains 100% SR across all tracks. The expert controller achieves near-perfect MGE thanks to full ground-truth access, which naturally surpasses our MM controller learned by imitation learning.

In real world hardware experiments (right portion of Table III), the expert again demonstrates flawless gate crossings but shows a slight increase in MGE, likely due to discrepancies between the real quadrotor and the approximate double-integrator dynamics used in FalconGym. Our MM controller achieves 100% SR on the circle and figure-8 tracks and 87.5% SR on the U-turn track (5 failures out of 40 gates over 10 runs). Further investigation shows that these failures consistently occur at the second gate of the U-turn track. We identified a few minor artifacts in the corresponding FalconGym simulation environment that appear to introduce a non-trivial sim-to-real visual gap: when we provide the MM controller with an artifact-affected image instead of a real-world image—while keeping the camera pose and

other inputs identical—we observe a marked difference in its control outputs. Nonetheless, the MGE for all tracks remains sufficiently low (average of 10.4cm) relative to the 38 cm gate radius, and only slightly higher than its FalconGym counterpart (average of 7.2cm). These results indicate that our controller effectively conforms to the hard constraints of avoiding collisions and staying within gate boundaries in real-world scenarios.

Overall, these experiments show that (i) our MM controller surpasses both naive DP approach and vision-only SoTA baseline in FalconGym by effectively fusing sensor inputs through self-attention, and (ii) our MM controller generalizes well from FalconGym to hardware, validating our approach’s ability for zero-shot sim-to-real quadrotor crossing gates.

## V. CONCLUSION

In this work, we present a novel visual control framework for quadrotor gate navigation that achieves zero-shot sim-to-real transfer. Leveraging our photorealistic FalconGym simulation environment, we train a visual controller that integrates visual features and pose estimation through a self-attention-based sensor fusion mechanism, making it robust to perception errors. Experiments in both FalconGym and hardware demonstrate high performance of our approach.

Nonetheless, performance can degrade due to visual domain gaps, mismatches between FalconGym’s simulated dynamics and the true quadrotor, and offboard computation or communication delays. Future directions include replacing our NeRF backend with Gaussian Splatting [39] for faster, higher-fidelity rendering; training policies that generalize to unseen tracks; incorporating physics-based dynamics [40] beyond a double-integrator model for aggressive maneuvers; distilling perception and control modules for fully onboard execution; and exploring high-speed flights to handle motion blur and tight perception–action coupling.

## ACKNOWLEDGMENT

The authors conducting this research were supported in part by the Air Force Research Laboratories (Award number FA8651-24-1-0007) and the Boeing Company. The authors would like to thank Yangge Li and Ege Yuceel for their valuable feedback.

## REFERENCES

- [1] I. Geles, L. Bauersfeld, A. Romero, J. Xing, and D. Scaramuzza, “Demonstrating agile flight from pixels without state estimation,” in *Robotics: Science and Systems XX, Delft, The Netherlands, July 15–19, 2024*, D. Kulis, G. Venture, K. E. Bekris, and E. Coronado, Eds., 2024. [Online]. Available: <https://doi.org/10.15607/RSS.2024.XX.082>
- [2] W. Zhao, J. P. Queralta, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: a survey,” *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 737–744, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:221971078>
- [3] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: representing scenes as neural radiance fields for view synthesis,” *Commun. ACM*, vol. 65, no. 1, p. 99–106, Dec. 2021. [Online]. Available: <https://doi.org/10.1145/3503250>

- [4] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Transactions on Graphics*, vol. 42, no. 4, July 2023. [Online]. Available: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- [5] A. Byravan, J. Humplik, L. Hasenclever, A. Brussee, F. Nori, T. Haarnoja, B. Moran, S. Bohez, F. Sadeghi, B. Vujatovic, and N. M. O. Heess, “Nerf2real: Sim2real transfer of vision-guided bipedal motion skills using neural radiance fields,” *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9362–9369, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:252815541>
- [6] M. Torne, A. Simeonov, Z. Li, A. Chan, T. Chen, A. Gupta, and P. Agrawal, “Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation,” *Arxiv*, 2024.
- [7] J. Low, M. Adang, J. Yu, K. Nagami, and M. Schwager, “Sous vide: Cooking visual drone navigation policies in a gaussian splatting vacuum,” *IEEE Robotics and Automation Letters (under review)*, 2024, available on arXiv: <https://arxiv.org/abs/2412.16346>.
- [8] J. Ni, W. Zhao, D. Wang, Z. Zeng, C. You, A. Wong, and K. Huang, “Efficient interactive 3d multi-object removal,” 2025. [Online]. Available: <https://arxiv.org/abs/2501.17636>
- [9] L. Yen-Chen, P. Florence, J. T. Barron, A. Rodriguez, P. Isola, and T.-Y. Lin, “iNeRF: Inverting neural radiance fields for pose estimation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [10] Q. Dai, Y. Zhu, Y. Geng, C. Ruan, J. Zhang, and H. Wang, “Graspnerf: Multiview-based 6-dof grasp detection for transparent and specular objects using generalizable nerf,” 2023.
- [11] M. Adamkiewicz, T. Chen, A. Caccavale, R. Gardner, P. Culbertson, J. Bohg, and M. Schwager, “Vision-only robot navigation in a neural radiance world,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4606–4613, 2022.
- [12] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. 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)*. IEEE Press, 2017, p. 23–30. [Online]. Available: <https://doi.org/10.1109/IROS.2017.8202133>
- [13] T. Söderström and P. Stoica, *System identification*. USA: Prentice-Hall, Inc., 1988.
- [14] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science Robotics*, vol. 5, no. 47, p. eabc5986, 2020. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abc5986>
- [15] W. Yu, V. C. Kumar, G. Turk, and C. K. Liu, “Sim-to-real transfer for biped locomotion,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Press, 2019, p. 3503–3510. [Online]. Available: <https://doi.org/10.1109/IROS40897.2019.8968053>
- [16] X. Hu, S. Li, T. Huang, B. Tang, R. Huai, and L. Chen, “How simulation helps autonomous driving: a survey of sim2real, digital twins, and parallel intelligence,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.01263>
- [17] E. Kaufmann, A. Loquercio, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, “Deep drone racing: Learning agile flight in dynamic environments,” in *Conference on Robot Learning (CoRL)*, 2018.
- [18] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, “Deep drone acrobatics,” *RSS: Robotics, Science, and Systems*, 2020.
- [19] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, “Learning high-speed flight in the wild,” *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abg5810>
- [20] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, “Champion-level drone racing using deep reinforcement learning,” *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS’17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 6000–6010.
- [22] J. Lee, H. Lee, and J. Oh, “Fusionloc: Camera-2d lidar fusion using multi-head self-attention for end-to-end serving robot relocalization,” *IEEE Access*, vol. 11, pp. 75 121–75 133, 2023.
- [23] C. Zhao, K. Liu, H. Zheng, W. Song, Z. Pei, and W. Chen, “Cross-modality self-attention and fusion-based neural network for lower limb locomotion mode recognition,” *IEEE Transactions on Automation Science and Engineering*, pp. 1–14, 2024.
- [24] D. Shah, A. Sridhar, N. Dashora, K. Stachowicz, K. Black, N. Hirose, and S. Levine, “ViNT: A foundation model for visual navigation,” in *7th Annual Conference on Robot Learning*, 2023. [Online]. Available: <https://arxiv.org/abs/2306.14846>
- [25] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME - Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [26] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, G. Gordon, D. Dunson, and M. Dudík, Eds., vol. 15. Fort Lauderdale, FL, USA: PMLR, 11–13 Apr 2011, pp. 627–635. [Online]. Available: <https://proceedings.mlr.press/v15/ross11a.html>
- [27] P. Meriaux, Y. Dupuis, R. Bouteau, P. Vasseur, and X. Savatier, “A Study of Vicon System Positioning Performance,” *Sensors*, vol. 17, no. 7, July 2017. [Online]. Available: <https://hal.science/hal-01710399>
- [28] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 4104–4113.
- [29] M. Tancik, E. Weber, E. Ng, R. Li, B. Yi, J. Kerr, T. Wang, A. Kristoffersen, J. Austin, K. Salahi, A. Ahuja, D. McAllister, and A. Kanazawa, “Nerfstudio: A modular framework for neural radiance field development,” in *ACM SIGGRAPH 2023 Conference Proceedings*, ser. SIGGRAPH ’23, 2023.
- [30] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=YicbFdNTTy>
- [31] T. Ridnik, E. Ben-Baruch, A. Noy, and L. Zelnik-Manor, “Imagenet-21k pretraining for the masses,” in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition*, ser. CVPR ’16. IEEE, June 2016, pp. 770–778. [Online]. Available: <http://ieeexplore.ieee.org/document/7780459>
- [33] D. Sun, S. Jha, and C. Fan, “Learning Certified Control using Contraction Metric,” in *Conference on Robot Learning*, 2020.
- [34] Y. Li, B. C. Yang, Y. Jia, D. Zhuang, and S. Mitra, “Refining perception contracts: Case studies in vision-based safe auto-landing,” 2023. [Online]. Available: <https://arxiv.org/abs/2311.08652>
- [35] C. Hsieh, Y. Li, D. Sun, K. Joshi, S. Misailovic, and S. Mitra, “Verifying controllers with vision-based perception using safe approximate abstractions,” *Trans. Comp.-Aided Des. Integ. Cir. Sys.*, vol. 41, no. 11, p. 4205–4216, Nov. 2022. [Online]. Available: <https://doi.org/10.1109/TCAD.2022.3197508>
- [36] R. Koenker and K. F. Hallock, “Quantile regression,” *The Journal of Economic Perspectives*, vol. 15, no. 4, pp. 143–156, 2001. [Online]. Available: <http://www.jstor.org/stable/2696522>
- [37] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [38] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [39] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, and X. Wang, “4d gaussian splatting for real-time dynamic scene rendering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 20 310–20 320.
- [40] L. Bauersfeld\*, E. Kaufmann\*, P. Foehn, S. Sun, and D. Scaramuzza, “Neurobem: Hybrid aerodynamic quadrotor model,” in *Robotics: Science and Systems XVII*, ser. RSS2021. Robotics: Science and Systems Foundation, July 2021. [Online]. Available: <http://dx.doi.org/10.15607/RSS.2021.XVII.042>