Learning Deep Sensorimotor Policies
for Vision-based Autonomous Drone Racing

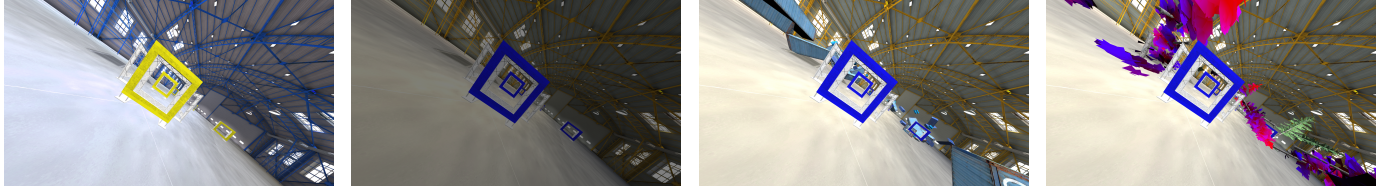Jiawei Fu, Yunlong Song, Yan Wu, Fisher Yu, and Davide Scaramuzza

Fig. 1: We learn deep sensorimotor policies for vision-based autonomous drone racing. The policy maps sensory observations to control commands directly. It achieves results that are robust against various visual disturbances and distractors, which are not experienced during training. From left to right: changing hue value, changing brightness, adding blue boxes, and adding random objects.

Abstract—Autonomous drones can operate in remote and unstructured environments, enabling various real-world applications. However, the lack of effective vision-based algorithms has been a stumbling block to achieving this goal. Existing systems often require hand-engineered components for state estimation, planning, and control. Such a sequential design involves laborious tuning, human heuristics, and compounding delays and errors. This paper tackles the vision-based autonomous-drone-racing problem by learning deep sensorimotor policies. We use contrastive learning to extract robust feature representations from the input images and leverage a two-stage learning-by-cheating framework for training a neural network policy. The resulting policy directly infers control commands with feature representations learned from raw images, forgoing the need for globally-consistent state estimation, trajectory planning, and handcrafted control design. Our experimental results indicate that our vision-based policy can achieve the same level of racing performance as the state-based policy while being robust against different visual disturbances and distractors. We believe this work serves as a stepping-stone toward developing intelligent vision-based autonomous systems that control the drone purely from image inputs, like human pilots.

Video: https://youtu.be/nPlGR83bC0Q

I. INTRODUCTION

Autonomous drones can travel through complex and dynamic environments at very high speed, holding great potential for a wide range of applications, such as industrial inspection, search and rescue, and reconnaissance. Robust vision-based autonomous flight is key to this goal. Autonomous vision-based drones have made significant progress in recent years,

continuously pushing the vehicle to higher speeds and better robustness. Several competitions have been organized to push the limit, such as the IROS 2016-19's Autonomous Drone Racing series [1], NeurIPS 2019's Game of Drones [2], the 2019 AlphaPilot Challenge [3], [4], and the ICRA 2022's DodgeDrone Challenge.

Vision-based autonomous drone racing requires operating the vehicle on the edge of its physical limits, thereby coping with the motion blur and the rapid illumination changes induced by the high speeds and quick rotations of the camera. The tolerance of the system for mistakes is extremely low: any small error can lead to a crash.

While existing works on vision-based autonomous drone racing rely on globally-consistent state-estimation, planning, and control [1]–[5], human pilots race drones by relying solely on a video stream from the drone's onboard camera, that is, by directly mapping visual input to control commands. While human pilots build a mental model of the drone state, they do not perform any explicit state estimation or trajectory planning [6]. In this paper, we make a small step toward emulating human pilots by learning a deep sensorimotor policy for vision-based autonomous drone racing.

Recent progress in the robot learning community demonstrates that learning deep sensorimotor policies for robotic tasks is feasible. Methods of this kind usually predict control commands directly from information extracted from high-dimensional sensory inputs. Deep sensorimotor policies have been heavily investigated in many robotic domains, such as object manipulation using robot arms [7], [8] or benchmark control of simulated robots [9]–[11]. This line of works has the advantage that the policy algorithm relaxes the need for a globally-consistent state information and enlarges the application of the system. However, learning deep sensorimotor policies for vision-based navigation still faces several challenges, including high sample complexity and poor generalization.

arXiv:2210.14985v1 [cs.RO] 26 Oct 2022

An overview of our system is given in Fig. 2. Our main contribution is a deep sensorimotor policy that can jointly solve perception, planning, and control for autonomous drone racing, without relying on an globally-consistent state of the drone nor on trajectory planning. The inputs to our policy are a sequence of images and part of the drone state (orientation, velocity, acceleration) but *no globally-consistent position* information.

Our method consists of two key components: privileged policy training and robust feature learning. First, we leverage a two-stage learning-by-cheating framework for policy training. Second, we use contrastive learning and data augmentation to extract robust image embeddings from RGB images.

Furthermore, we compare our vision-based deep sensorimotor policy against a neural control policy that utilizes the full globally-consistent state information [12]. Our experiments, conducted in a realistic simulator [13], show that our vision-based deep sensorimotor policy achieves the same level of racing performance while being robust against different visual disturbances and distractors. Finally, we benchmark the performance of our vision-based policy against the time-optimal trajectory generation algorithm [14], which offers a theoretical minimum time. Our policy achieves lap times close to the time-optimal solution.

II. RELATED WORK

Different approaches have been studied to tackle autonomous drone racing. State-based methods that rely on globally accurate position information have been used extensively. Foehn et al. [14] presented the time-optimal trajectory generation by jointly improving the time allocation and the trajectory. The algorithm enabled them to outperform human experts in drone racing. In [12], [15], [16], authors used reinforcement learning to train a neural network as the policy. For example, Song et al. [12] utilized relative gate positions towards the next gates to achieve near-time-optimal performance. Nagami et al. [16] initialized a network by mimicking a simplified controller and further trained it with reinforcement learning. The hierarchy allowed the policy to outperform a trajectory planning policy. Although promising results can be generated by state-based methods, the assumption of exact position information limits the application of the methods.

Prior work on vision-based drone racing decouples the perception, planning, and control modules. In the work of Foehn et al. [3], visual-inertial odometry (VIO) was fused with a CNN-based gate corner detection for robust state estimation. A receding horizon path planner generates a time-optimal trajectory using motion primitives based on a point-mass model of the drone platform. However, the point-mass assumption cannot represent the true actuation limits of the drone and may lead to dynamically infeasible trajectories. In [17]–[19], authors first use data-driven methods to train the neural networks that can predict the waypoint and the desired speed. Afterward, a minimum jerk trajectory is planned for passing through the waypoint and then tracked by a low-level controller. Muller et al. [20] propose to train a neural network for local trajectory planning, in which a downstream control policy is used to track the trajectory and generate low-level commands for vehicle control. The trajectory labeling requires additional engineering efforts and can result in ambiguity as each image can be labeled with different trajectories. The decoupling of the perception, planning, and control modules inevitably involves simplified assumptions or manual design of parameters, leading to sub-optimality during high-speed flight.

Recent advance in data-driven control [21]–[24] indicates the potential of developing autonomous systems using sensorimotor control, in which a neural network policy can map high-dimensional sensory inputs directly to control commands. However, with naive training, the deep sensorimotor policy might suffer from poor generalization when facing unseen disturbances. Different approaches have been employed to alleviate the overfitting issue, such as data augmentation [11], [25], injecting known biases [26], and extracting invariant information [27]. Most of them are applied to video games [10], [26], robot arm control [25], [28], or autonomous driving [29]. The generalization capability of a neural network policy for minimum-time flight has drawn much less attention due to several challenges, such as low reaction time and a rapid change of the image observations.

III. METHODOLOGY

An overview of our method is visualized in Fig. 2. Our approach consists of two key components: policy training and feature learning. The policy training is done using privileged reinforcement learning and imitation learning, where a student policy mimics the action of a teacher policy. To process high-dimensional image data and allow efficient policy training, we use YOLO [30], [31] to extract low-dimensional image embeddings.

A. Policy Training

Teacher Policy Training: The first step is to obtain a state-based teacher policy that can push the vehicle to its maximum performance. We use reinforcement learning to train a multilayer perceptron (MLP) policy π_{teacher} for passing through a sequence of gates $\{G_i |, i = 1, \cdots, N\}$ in minimum time. At every time step t, the agent is at state s_t and receives information about the gate state g_t. Our goal is to find the optimal policy π^*_{teacher} that maximizes the expected discounted return: $\pi^*_{\text{teacher}} = \arg\max_\pi \mathbb{E}\left[\sum_{t=0}^{T-1} \gamma^t r_t\right]$, where γ is the discount factor and r_t is the reward at time stage t. In privileged learning, the teacher policy has access to all ground truth information, including the vehicle's state s_t and the gate state g_t. Hence, the teacher policy generates an action $\bar{a}_t \sim \pi(s_t, g_t)$ given both states. The policy outputs control commands in the form of mass-normalized collective thrust and angular velocity.

The main objective is to minimize the lap time, which is equivalent to maximizing the path progress along the center line connection between two consecutive gates [12]. In addition, we maximize a perception-aware reward to maximize the visibility of the next gate. The perception-aware reward incentivizes the policy to face the camera toward the
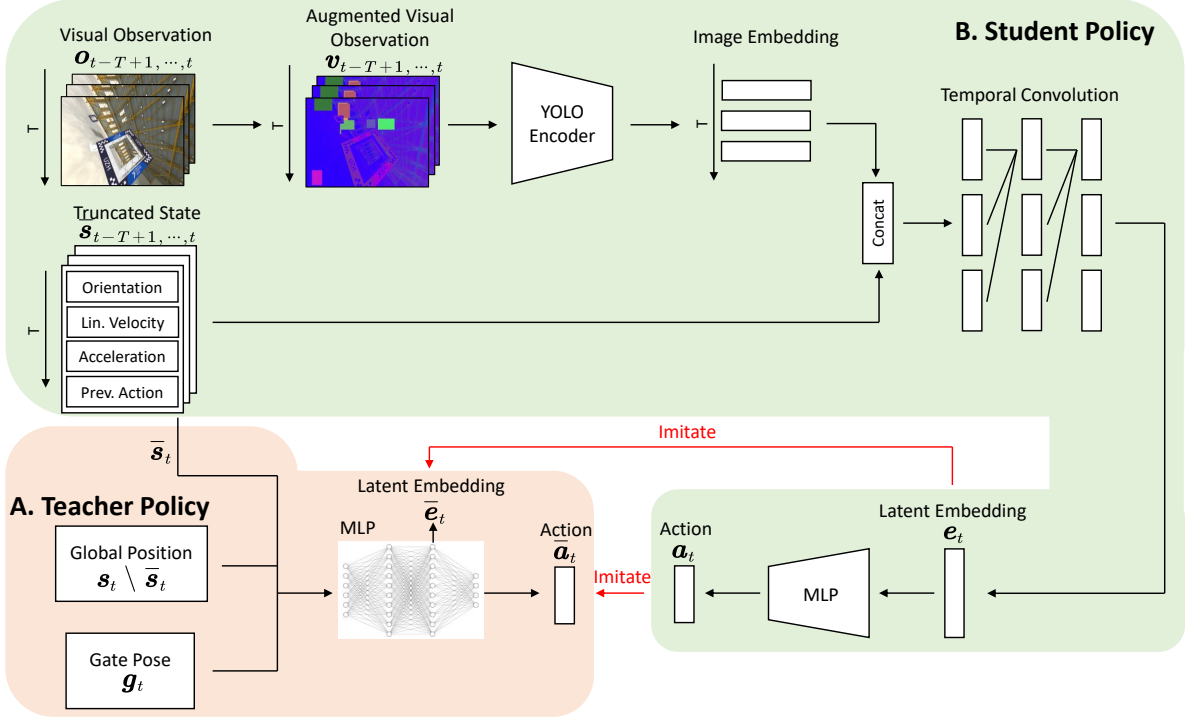
Fig. 2: Overview of our policy training method. We first train a teacher policy with access to privileged state information using model-free reinforcement learning. This teacher policy is then distilled into a student policy, which is trained to do perception, planning, and control jointly.

next passing gate, which is crucial for vision-based flight since our environment is only partially observable when using a camera. We denote the position and velocity of the center of the next gate on the image plane by $\boldsymbol{p}_c$ and $\dot{\boldsymbol{p}}_c$, respectively. The perception award reward is formulated to keep the gate in the image center and reduce the motion blur [32] as $r_p = \exp\left(-||\boldsymbol{p}_c||_2 - ||\dot{\boldsymbol{p}}_c||_2\right)$, where $||\cdot||_2$ refers to 2-norm.

Student Policy Training: After we obtain a teacher policy $\pi^*_{\text{teacher}}(\boldsymbol{s}_t, \boldsymbol{g}_t)$ that can race the drone optimally, we distill the teacher's knowledge to a student policy $\pi_{\text{student}}(\overline{\boldsymbol{s}}_t, \boldsymbol{o}_t)$ that does not have access to the privileged information about the environment. Specifically, the student policy can only observe part of the drone state $\overline{\boldsymbol{s}}_t$, which does not contain the vehicle's global position, and need to infer the gate information from the camera observation $\boldsymbol{o}_t \to \boldsymbol{g}_t$. There are three key components of our student policy: a feature extractor, a memory-based neural network, and a policy network.

We use YOLO [30], [31] as the feature extractor and train it to detect all gates in a given image. We use average pooling to downsample the output of the three convolutional layers in its detection head and concatenate them as the embedding of the image $\boldsymbol{z}_t$. Since the detection head is the rightmost module of YOLO, this embedding contains all the information for detecting the gates, and hence, is a sufficient for representing the image. Note that we additionally normalize the embed-

ding with $l2$-normalization, which is empirically found to be beneficial for the convergence of policy training.

When using a single camera, the environment becomes a partially observable environment. To this end, we use a temporal convolutional network (TCN) [33] for the policy representation. The embedding from the image is concatenated with the truncated vehicle state $\overline{\boldsymbol{s}}_t$. The sequence of concatenated embeddings is then fed into the TCN to extract temporal information from history observations. Finally, we use a MLP to regress the control command. The MLP takes the output of the TCN as input and produces the student policy's action which is of the same format as the teacher policy.

We use imitation learning to train the student policy. We define an action loss $\mathcal{L}_A$ that is the mean square error between the outputs of the teacher policy and the student policy.

$$\mathcal{L}_A(\boldsymbol{\theta}) = ||\pi_{\text{student}}(\overline{\boldsymbol{s}}_t, \boldsymbol{o}_t|\boldsymbol{\theta}) - \pi^*_{\text{teacher}}(\boldsymbol{s}_t, \boldsymbol{g}_t)||_2. \quad (1)$$

In addition, to better enable the knowledge transfer between the teacher policy and the student policy, we also add a latent loss $\mathcal{L}_E$ to supervise the output of the TCN $\boldsymbol{e}_t$ with the intermediate embedding of the teacher $\overline{\boldsymbol{e}}_t$, written as $\mathcal{L}_E(\boldsymbol{\theta}) = ||\overline{\boldsymbol{e}}_t - \boldsymbol{e}_t||_2$. Therefore, we minimize the total loss for the imitation learning

$$\min_{\boldsymbol{\theta}} \quad \mathcal{L} = \mathcal{L}_A(\boldsymbol{\theta}) + \lambda \mathcal{L}_E(\boldsymbol{\theta}), \quad (2)$$

where λ is a coefficient to weight the latent loss.

B. Robust Feature Learning via Data Augmentation

To learn robust image embeddings against disturbance, we train the encoder with contrastive learning. We use the framework introduced in [34]. The framework (shown in Fig. 3) contains an online network and a target network. The online network defined by parameters ϕ contains three components: an encoder f_ϕ, a projection g_ϕ, and a predictor q_ϕ. The target network is an exponential mean average of the online network and defined by parameters $\boldsymbol{\xi}$. It is comprised of two components: a target encoder $f_{\boldsymbol{\xi}}$ and a target projection $g_{\boldsymbol{\xi}}$.
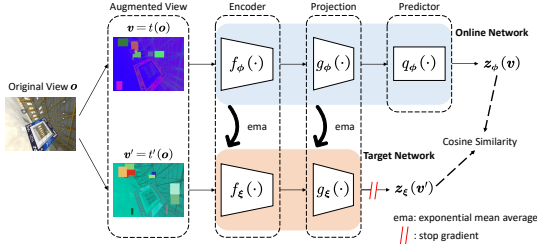

Fig. 3: Contrastive learning framework [34].

An input image $\boldsymbol{o}$ is passed through two augmentations, denoted as t and t', to obtain two augmented views $\boldsymbol{v} = t(\boldsymbol{o})$ and $\boldsymbol{v}' = t'(\boldsymbol{o})$ separately. Then the embedding prediction $\boldsymbol{z}_\phi(\boldsymbol{v}) = q_\phi(g_\phi(f_\phi(\boldsymbol{v})))$ is extracted by the online network while the embedding target $\boldsymbol{z}_{\boldsymbol{\xi}}(\boldsymbol{v}') = g_{\boldsymbol{\xi}}(f_{\boldsymbol{\xi}}(\boldsymbol{v}'))$ is extracted by the target network. A cosine similarity loss is applied to align the embeddings,

$$\mathcal{L}^{cos}(\boldsymbol{v}, \boldsymbol{v}') = 2 - 2 \cdot \frac{< \boldsymbol{z}_\phi(\boldsymbol{v}), \boldsymbol{z}_{\boldsymbol{\xi}}(\boldsymbol{v}') >}{||\boldsymbol{z}_\phi(\boldsymbol{v})||_2 \cdot ||\boldsymbol{z}_{\boldsymbol{\xi}}(\boldsymbol{v}')||_2} \tag{3}$$

where $|| \cdot ||_2$ denotes $l2$-normalization.

IV. EXPERIMENTS

We design our experiments to answer the following research questions: 1) our student policy does not depend on the vehicle's global position and can only observe the environment partially; how does such a vision-based control system compare to a state-based system? 2) our policy is trained with data augmentation. Can the data augmentation align image embeddings given different visual disturbances and distractors, and is the policy robust against those disturbances? 3) our policy still relies on some part of the vehicle state; how well can the policy handle estimation errors in the drone state?

A. Experimental Setup

Simulator Environment: We conduct experiments using the Flightmare [13] simulator, a realistic quadrotor simulator with various racing tracks and realistic racing environments. We set up three different race tracks (Circle, Figure8, and SplitS) in a warehouse environment (see the visualization in Fig. 4 left). For training the teacher policy, we use a customized implementation of the proximal policy optimization algorithm (PPO) [35] based on the code from [36]. For

training the student policy, we implement an imitation learning pipeline. For learning robust feature representations from raw images, we use data augmentation with random convolution and random cutout-color (see Fig. 4 middle and right).

Evaluation: To evaluate our policy, we rollout for 10 episodes, with quadrotor starting from different starting positions, which are sampled from a uniform distribution between -0.1m and 0.1m in x, y, z-axis of each. We evaluate the performance of our policy using two different metrics: Lap Time and Success Rate. The lap time indicates the racing performance of our policy, while the success rate indicates the robustness of the policy. We report the lap time by computing the time required by the policy to finish one complete track and calculate the success rate by calculating the ratio that the quadrotor can successfully finish one full lap without crashing among the ten rollouts.

B. Baseline Comparisons

We compare our vision-based policies against two baselines: a state-based learning-based policy [12] and a time-optimal trajectory [14]. The state-based policy controls the drone using ground truth information about the drone state, including position, velocity, orientation, and acceleration, as well as the pose of the next two gates. The time-optimal trajectory serves as the theoretical minimum bound for our platform. The student policy does not have access to ground truth information about the drone's position and the gate poses. Instead, it uses a camera to capture RGB images and controls the drone directly using the image. Therefore, the student policy can only observe the environment partially, similar to how human pilots control the drone using the first-person-view camera. The result is shown in Table I. A visualization of the trajectories is given in Fig. 5. Both policies achieve strong performance on three different race tracks with high success rates. The student policy learns to cut corners, resulting in lower lap time, but more risky behaviors.

In reality, the vehicle states are prone to error due to the drift in state estimation and measurement errors. Despite impressive results in visual-(inertial) odometry in recent years, high-speed flight with six degrees of freedom motion remains challenging for existing estimation algorithms [37]. Hence, the state-based control system is subject to failure since the policy relies heavily on position estimation. We investigate this problem using a simulated VIO pipeline, in which we simulate position drift. Fig. 6 shows how position drifts affect the performance of the state-based policy. Given perfect state information, the policy achieves 100 % success rates on all tracks. However, as we increase the drift in position, the success rates collapse quickly. The VIO drift can be alleviated by relocalizing with respect to the gates but this is challenging because the camera suffers from motion blur and limited field of view. On the other hand, our vision-based policy is not affected by the drift in the position since it does not rely on that information.
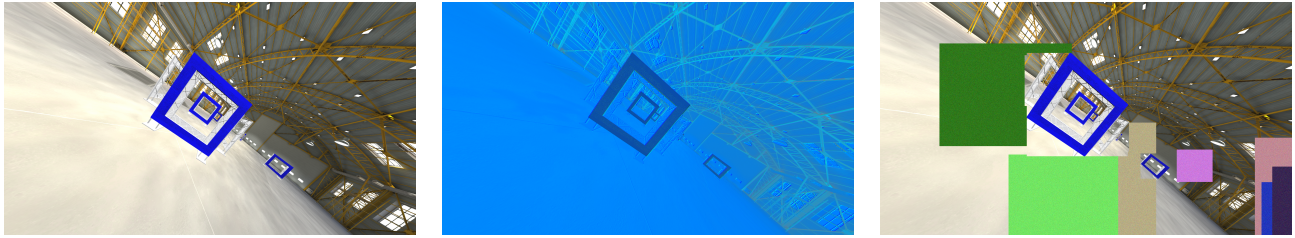
Fig. 4: Visualization of data augmentations used during training. Left: no augmentation. Middle: random convolution. Right: random cutout-color.

	Lap Time [s]			Success Rate		
	Circle	Figure8	SplitS	Circle	Figure8	SplitS
Time-optimal Trajectory [14]	4.68	6.26	7.93	-	-	-
State-based Policy [12]	4.97±0.01	6.84±0.05	8.74±0.01	1.0	1.0	1.0
Vision-based Policy (ours)	4.95±0.01	6.76±0.01	8.58±0.01	1.0	1.0	1.0

TABLE I: Performance of the time-optimal trajectory, state-based teacher policy, and vision-based student policy on three different race tracks.
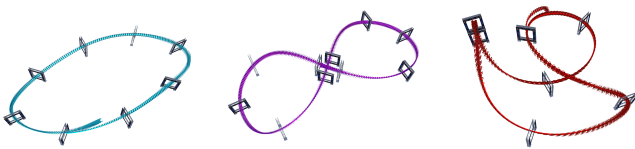

Fig. 5: Visualization of trajectories. Left: Circle. Middle: Figure8. Right: SplitS.
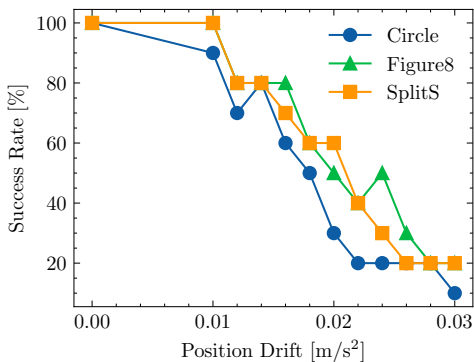

Fig. 6: Success rates of the state-based policy over position drift.

C. Handling Visual Disturbances and Unseen Distractors

We deploy our vision-based system in various unseen contexts to investigate how it performs against unseen visual disturbances, such as environments with color changes, brightness changes, and environments with many randomly arranged unseen objects. We darken the environment by lowing the brightness value from 1 to 0.5 and 0.8, and we also change environment colors by tuning the image hue value from 0 to both 0.5 and -0.5. Fig. 1 left and middle-left provide examples of environment with brightness values of 0.5 and hue values of 0.5, respectively. In addition, we also place some visual distractors randomly around the environment, including blue boxes that are similar in color and shape to the racing gates (see visualization in Fig. 1 middle-right), and some random objects with irregular shapes (see visualization in Fig. 1 right).

As presented in Table II, our system is robust against various types of visual disturbances while still maintaining a high success rate and comparable lap time on all three racing tracks, which demonstrates the effectiveness of our image feature learning mechanism. In the following section, we further investigate how the image encoder trained using contrastive learning can generalize to these visual disturbances.

D. Aligning Image Embeddings

To ensure robust feature extraction, we use contrastive learning (Setion III-B). In the contrastive learning framework, the similarity loss ensures that the encoder learns the invariance between the two augmented views. As a result, the image embeddings between augmented views are aligned in the embedding space. We choose random convolution as the augmentation for hue changes and brightness changes; and use random cutout-color against distractors, such as blue boxes and random objects. In Fig. 7, we present the qualitative results of aligning image embeddings between augmentations and disturbances. For each of the three tracks, we collect the images along the flight trajectory of the teacher policy with either augmentations or disturbances. For each of the trajectories, we extract the image embeddings with the YOLO encoder and reduce the dimension to 2 with t-distributed stochastic neighbor embedding (t-SNE). The embeddings of each test-time disturbance are then evaluated with those of the corresponding augmentation during training. We can observe that the image embeddings of all the disturbances are well aligned with those of the corresponding augmentations. It ensures that our policy receives matching image embeddings in test time and behaves robustly. Thus, our policy still maintains a high success rate under all the disturbances (Table II).

		Lap Time [s]			Success Rate		
		Circle	Figure8	SplitS	Circle	Figure8	SplitS
Brightness Change	0.5	4.88±0.01	6.71±0.01	8.60±0.01	1.0	0.8	1.0
	0.8	4.91±0.02	6.68±0.01	8.65±0.01	1.0	1.0	1.0
Hue Change	-0.5	4.91±0.01	6.72±0.03	8.66±0.01	1.0	1.0	1.0
	0.5	4.92±0.02	6.71±0.01	8.65±0.01	1.0	1.0	1.0
Blue Boxes	10	4.94±0.01	6.87±0.01	8.64±0.01	1.0	1.0	1.0
	60	4.90±0.02	6.77±0.01	8.73±0.01	1.0	1.0	0.9
Random Objects	10	4.94±0.02	6.75±0.02	8.67±0.01	1.0	1.0	1.0
	60	4.91±0.02	6.81±0.03	8.67±0.01	1.0	0.6	1.0

TABLE II: Success rate and lap time of the vision-based sensorimotor control policy. Our policy is robust against different visual disturbances, including brightness change (darker environment with brightness value changing from 1 to 0.5 and 0.8, respectively) and hue change (environment color changes with hue value tuning from 0 to 0.5 and -0.5, respectively), and the presence of various distractors, including blue boxes and random objects, of various densities (10 and 60 distractors are added, respectively).

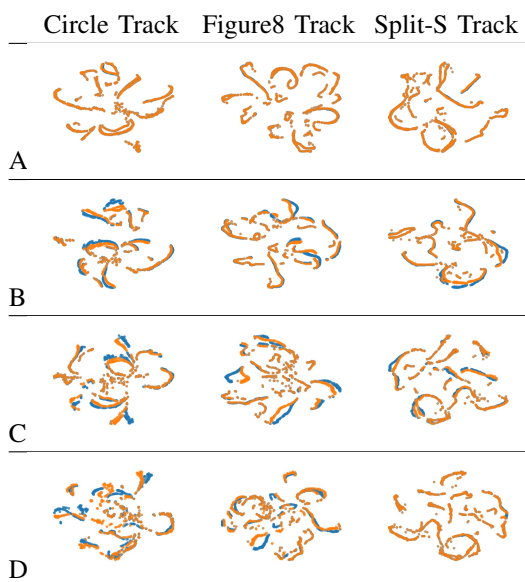|	Circle Track	Figure8 Track	Split-S Track

Fig. 7: A time-lapse t-SNE visualization of image embeddings used by our policy. We collect images along the flying trajectory on three different tracks. The blue dots represent image embeddings from augmentations during training time and the orange dots represent image embeddings from test-time disturbances. A: Hue change. B: Brightness change. C: Blue boxes. D: Random Objects.

E. Handling Noisy State

Our policy relies on part of the drone state, including the orientation, linear velocity, and acceleration. The state information can be estimated using measurements from onboard sensors, such as IMUs, which are usually noisy. We further investigate the robustness of our policy against disturbances on the truncated sates by adding Gaussian noise $\mathcal{N}(0, std)$ individually to each component of the states, where std is the standard deviation. Table III shows the result. We can observe that the success rate decreases mildly when the standard deviation increases, which proves that our policy is robust

	std	Circle	Figure8	SplitS
State Disturbance	0.04	1.0	1.0	1.0
	0.12	1.0	1.0	0.7
	0.20	1.0	1.0	0.4
	0.28	0.3	0.5	0.0

TABLE III: Success rates of the student policy when adding Gaussian noises to the drone states.

against noises from the sensor measurements.

V. DISCUSSION AND CONCLUSION

This work presented a method to learn deep sensorimotor policies for vision-based autonomous drone racing. We showed that a vision-based control policy allows predicting control commands with information extracted from images without explicitly estimating position information, trajectory planning, and tracking. The vision-based policy can achieve the same level of racing performance as the state-based policy while being robust against different visual disturbances and distractors. On the other hand, a state-based control policy is sensitive to position errors in state estimation. The key to achieving robust sensorimotor control is to learn well-aligned image embeddings using contrastive learning and data augmentation. These findings suggest that deep sensorimotor control has the potential for vision-based agile drone flight and merits further investigation.

A major limitation of the presented work is a lack of real-world experiments to demonstrate the effectiveness and robustness of our vision-based policy. The deployment of the student policy on a real drone still requires further research on transfer learning or adaptive learning. Although relaxing the need for globally-consistent position information about the drone and the gate, the student policy still relies on part of the vehicle's state to predict the control commands. We plan to tackle this in the near future by using memory-based policy representations, such as RNNs to learn hidden state representations from a history of images alone. We believe our study is a stepping stone towards this goal.

REFERENCES

[1] H. Moon, J. Martinez-Carranza, T. Cieslewski, M. Faessler, D. Falanga, A. Simovic, D. Scaramuzza, S. Li, M. Ozo, C. De Wagter *et al.*, "Challenges and implemented technologies used in autonomous drone racing," *Intelligent Service Robotics*, 2019.

[2] R. Madaan, N. Gyde, S. Vemprala, M. Brown, K. Nagami, T. Taubner, E. Cristofalo, D. Scaramuzza, M. Schwager, and A. Kapoor, "Airsim drone racing lab," in *NeurIPS 2019 Competition and Demonstration Track*. PMLR, 2020.

[3] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, "Alphapilot: Autonomous drone racing," *Autonomous Robots*, 2021.

[4] C. De Wagter, F. Paredes-Valles, N. Sheth, and G. de Croon, "The sensing, state-estimation, and control behind the winning entry to the 2019 artificial intelligence robotic racing competition," *Field Robotics*, 2022.

[5] L. O. Rojas-Perez and J. Martínez-Carranza, "On-board processing for autonomous drone racing: an overview," *Integration*, vol. 80, pp. 46–59, 2021.

[6] C. Pfeiffer and D. Scaramuzza, "Human-piloted drone racing: Visual processing and control," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3467–3474, 2021.

[7] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.

[8] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.

[9] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq *et al.*, "Deepmind control suite," *arXiv preprint arXiv:1801.00690*, 2018.

[10] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, "Reinforcement learning with augmented data," *Advances in Neural Information Processing Systems*, vol. 33, pp. 19 884–19 895, 2020.

[11] N. Hansen and X. Wang, "Generalization in reinforcement learning by soft data augmentation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 611–13 617.

[12] Y. Song, M. Steinweg, E. Kaufmann, and D. Scaramuzza, "Autonomous drone racing with deep reinforcement learning," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2021.

[13] Y. Song, S. Naji, E. Kaufmann, A. Loquercio, and D. Scaramuzza, "Flightmare: A flexible quadrotor simulator," in *Conference on Robot Learning*, 2020.

[14] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, 2021.

[15] J. Hwangbo, I. Sa, R. Siegwart, and M. Hutter, "Control of a quadrotor with reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2096–2103, 2017.

[16] K. Nagami and M. Schwager, "Hjb-rl: Initializing reinforcement learning with optimal control policies applied to autonomous drone racing." in *Robotics: science and systems*, 2021.

[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*. PMLR, 2018, pp. 133–145.

[18] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: From simulation to reality with domain randomization," *IEEE Trans. Robotics*, vol. 36, no. 1, pp. 1–14, 2019.

[19] T. Wang and D. E. Chang, "Robust navigation for racing drones based on imitation learning and modularization," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 724–13 730.

[20] M. Muller, G. Li, V. Casser, N. Smith, D. L. Michels, and B. Ghanem, "Learning a controller fusion network by online trajectory filtering for vision-based uav racing," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.

[21] E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Deep drone acrobatics," *RSS: Robotics, Science, and Systems*, 2020.

[22] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.

[23] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, 2020.

[24] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 222–15 232.

[25] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 627–12 637.

[26] X. Wang, L. Lian, and S. X. Yu, "Unsupervised visual attention and invariance for reinforcement learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6677–6687.

[27] B. Li, V. François-Lavet, T. Doan, and J. Pineau, "Domain adversarial reinforcement learning," *arXiv preprint arXiv:2102.07097*, 2021.

[28] K. Lee, K. Lee, J. Shin, and H. Lee, "Network randomization: A simple technique for generalization in deep reinforcement learning," *arXiv preprint arXiv:1910.05396*, 2019.

[29] L. Fan, G. Wang, D.-A. Huang, Z. Yu, L. Fei-Fei, Y. Zhu, and A. Anandkumar, "Secant: Self-expert cloning for zero-shot generalization of visual policies," *arXiv preprint arXiv:2106.09678*, 2021.

[30] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.

[31] G. Jocher, A. Chaurasia, A. Stoken, J. Borovec, NanoCode012, Y. Kwon, TaoXie, J. Fang, imyhxy, K. Michael, Lorna, A. V, D. Montes, J. Nadar, Laughing, tkianai, yxNONG, P. Skalski, Z. Wang, A. Hogan, C. Fati, L. Mammana, AlexWang1900, D. Patel, D. Yiwei, F. You, J. Hajek, L. Diaconu, and M. T. Minh, "ultralytics/yolov5: v6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference," Feb. 2022. [Online]. Available: https://doi.org/10.5281/zenodo.6222936

[32] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, "Pampc: Perception-aware model predictive control for quadrotors," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. IEEE, 2018, pp. 1–8.

[33] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.

[34] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent-a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21 271–21 284, 2020.

[35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[36] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: http://jmlr.org/papers/v22/20-1364.html

[37] J. Delmerico, T. Cieslewski, H. Rebecq, M. Faessler, and D. Scaramuzza, "Are we ready for autonomous drone racing? the uzh-fpv drone racing dataset," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2019.