

Learning a Controller Fusion Network by Online Trajectory Filtering for Vision-based UAV Racing

Matthias Müller* Guohao Li* Vincent Casser Neil Smith Dominik L. Michels Bernard Ghanem
Visual Computing Center, King Abdullah University of Science and Technology, Saudi Arabia
{matthias.mueller.2, guohao.li, bernard.ghanem}@kaust.edu.sa

Abstract

Autonomous UAV racing has recently emerged as an interesting research problem. The dream is to beat humans in this new fast-paced sport. A common approach is to learn an end-to-end policy that directly predicts controls from raw images by imitating an expert. However, such a policy is limited by the expert it imitates and scaling to other environments and vehicle dynamics is difficult. One approach to overcome the drawbacks of an end-to-end policy is to train a network only on the perception task and handle control with a PID or MPC controller. However, a single controller must be extensively tuned and cannot usually cover the whole state space. In this paper, we propose learning an optimized controller using a DNN that fuses multiple controllers. The network learns a robust controller with online trajectory filtering, which suppresses noisy trajectories and imperfections of individual controllers. The result is a network that is able to learn a good fusion of filtered trajectories from different controllers leading to significant improvements in overall performance. We compare our trained network to controllers it has learned from, end-to-end baselines and human pilots in a realistic simulation; our network beats all baselines in extensive experiments and approaches the performance of a professional human pilot.

1. Introduction

Recent advances in UAV technology by industry leaders such as DJI, Amazon and Intel make UAV design and point-to-point stabilized flight navigation appear to be a well-solved problem. However, autonomous navigation of UAVs in more complex and real-world scenarios, such as in unknown congested environments, GPS-denied areas, and narrow spaces, is still far from being solved. This is a complex problem, since it requires *both* the sensing of the environment and the execution of appropriate control policies for interaction at low latencies, running on on-board hardware



Figure 1: UAV flying in Sim4CV while being controlled by our controller fusion network.

with typically very limited computational resources. The emerging sport of UAV racing displays a lot of these real-world navigation challenges, and is one of the areas where the performance gap between human pilots and machine-driven navigation approaches is most evident. UAV racing requires human pilots or agents to sequentially control the UAV to fly through a race track based on the feedback (visual information, physical measurements, or both) of previous actions. It requires control over six degrees of freedom (6-DoF) at high speeds while traversing tight spaces, and passing consistently through racing gates. These complex sense-and-understand tasks are conducted at extreme speeds reaching over 100 km/h, and thus can serve as a controlled and challenging benchmark for machine-driven agile navigation approaches.

One of the more difficult tasks in UAV racing is the prediction of the proper trajectory in order to traverse a course while maintaining high speeds. In earlier work, either a PID controller or model predictive controller with Kalman filters has been used. However, in practice a single controller cannot usually cover the whole state space. In certain states the controller may not perform as expected or even fail. In this paper, we propose the fusion of multiple controllers using a DNN to cover a much larger state space and to outperform any single controller.

*equal contribution

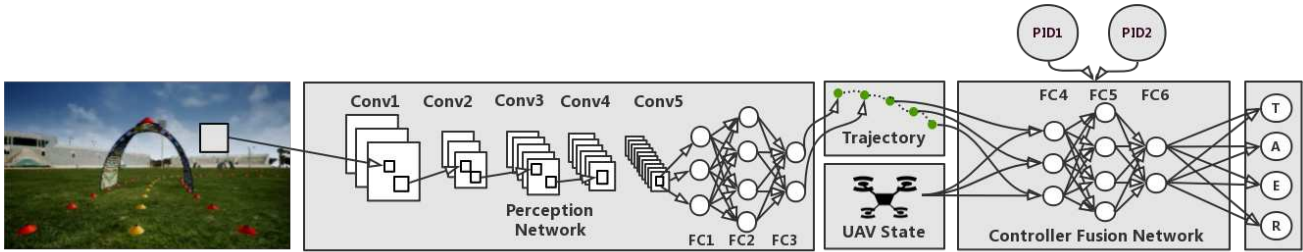


Figure 2: Our complete system consists of a perception network ϕ and a controller fusion network (CFN). The perception network predicts local trajectories from a monocular RGB image. The CFN φ takes the predicted trajectory and UAV state as input and outputs the low-level controls: throttle (T), roll/aileron(A), pitch/elevator (E), yaw/rudder (R). The CFN is trained by fusing filtered trajectories from multiple classical controllers.

However, fusing different controller’s trajectories naively will result in incorporating both their good and bad trajectories due to their high variability and limitations in specific scenarios. We overcome this problem by storing trajectories in an online buffer that can be accessed by our deep neural network (DNN) during learning. By filling this buffer with only good trajectories, we are able to trim out failing trajectories. The network learns an optimized controller by sampling from the buffer containing these pre-filtered trajectories. The result is a network that is able to learn a good fusion of filtered trajectories from different controllers allowing a much larger coverage of the whole state space. We show that this leads to significant improvements in overall performance.

We use Sim4CV [16] to simulate UAV racing with accurate physics using the Unreal Engine 4 (UE4). [17]. We also develop a customizable racing area in the form of a stadium based on a 3D scan of a real-world location. Race tracks are generated within a GUI interface and can be loaded at game-time to allow training on multiple tracks automatically. Inspired by recent work on self-driving cars [4], we are trying to imitate UAV racing at a professional level. The key difference is that we train a perception network that predicts desired trajectories rather than controls. We then train a separate neural network to produce the low-level controls. We call this network Controller Fusion Network (CFN) and train it by fusing filtered trajectories from multiple controllers. Through extensive experiments in simulation we demonstrate that CFN outperforms the controllers it learned from, end-to-end baselines, and even human pilots flying via a remote control used for real UAVs.

Contributions. (1) We propose to learn the fusion of trajectories by a neural network. This allows for combining trajectories from different controllers in a principled way, and separates the control from the perception task. (2) By implementing online trajectory filtering we are able to learn from multiple noisy trajectories without incorporating their imperfections. While the control task is learned online, a

buffer/memory also allows for semi-offline training. Our approach leads to a robust network outperforming several state-of-the-art approaches and human pilots.

2. Related Work

The use of deep neural networks (DNNs) to control UAVs dates back to work on learning acrobatic helicopter flight [1]. More recent work has studied training UAV controller networks with SL, RL or combined methods but with a focus on indoor flight, collision avoidance, and trajectory optimization [9, 13, 20, 22, 2, 11, 21, 10, 17, 3, 23]. An important insight from [13, 9, 12, 20] is that a trajectory optimizer such as a Model Predictive Controller (MPC) can function similar to traditional SL to help regress an agent’s sub-optimal policy towards the optimal one with much fewer iterations. By jointly learning perception and control with the self-supervision of MPC, full end-to-end navigation and collision avoidance can be learned. Recently, [10] implement such an approach where a DNN is used to predict a trajectory and a MPC is used to properly output the motor control of the UAV. Although not functioning at racing speeds, initial experiments demonstrate the advantages of such a setup. A major limitation of this approach is that the DNN is only used for perception, while the MPC requires extensive setup, tuning, and a full knowledge of the UAV dynamics. We propose in this paper that a DNN can also be applied to control allowing the UAV dynamics to be inherently learned.

Our network is able to learn from imperfect proportional-integral-derivative (PID) controllers allowing both self-supervision and extensive exploration. The network training with extensive exploration is most similar to DAGGER [19] and its variants [3, 24, 18, 5]. Our approach differs in that our control network does not learn strictly from the controller. Our network learns the appropriate actions of the PID controllers at each time step and selects the best predictions based on the filtered trajectories. The buffering strategy is similar to DQN’s [15] ”experience replay mecha-

nism” in that it stores a limited set of experiences in a buffer and then selects samples randomly. However, our buffer is dynamic being updated during online training and continually filters the buffer with only good samples. By design, our RL motivated training process enables extensive exploration by only observing the best behavior from various controllers. It differs from [12] in that the controllers never have to deviate from their optimal control to induce exploration. Also, unlike trajectory optimization [7], the trajectory in our setup is without known global 3D position. This enables the prediction of local waypoints without needing precise knowledge of the UAV’s current state and dynamics, which are rarely available in the real-world. Similar to adaptive trajectory optimization, our predicted waypoints are updated every time step allowing for adaptation to environment changes.

3. Methodology

In this section, we introduce the setup of the Controller Fusion Network (CFN) which enables automatic removal of bad trajectory segments from imperfect controllers by filtering bad samples in an online manner before the CFN is updated. We build a modular system for UAV racing (see Figure 2) that separates the task into a high dimensional perception module and a low dimensional control module (CFN). The perception network predicts trajectories which are used by the the CFN along with the UAV state to produce low-level controls.

3.1. Controller Fusion Network (CFN)

We learn a *Controller Fusion Network* agent by a learning strategy that integrates knowledge from multiple controllers and the dynamic environment into the learning process (refer to Algorithm 1 for a detailed description during training). At each time step t , the agent receives a state (or partial state) s_t from the environment and executes an action a_t . Thus, the trajectory of the agent behaviors are denoted as $\tau = (s_1, a_1, s_2, a_2, \dots, s_n)$. The CFN policy is a parameterized function $\pi(s|\theta)$ mapping the state to a deterministic action that can be continuous or discrete. In our case, the action is a 4D continuous control signal for a UAV. The PID controllers’ policy $\mu(s)$ also maps the state to an action. In practice, the controller can be either an automated controller or a human, and demonstration can be performed online or offline from a recording. In this paper, we use two PID (Proportional-Integral-Derivative) controllers, thus, avoiding the need for hours of human recorded control.

To clarify which demonstrations we should learn from, we refer to the demonstrations that lead the CFN agent to good behavior as *desirable* demonstrations, the ones that lead to bad behavior as *unwanted* demonstrations, and the uncertain ones as *unforeseeable* demonstrations. We introduce a *temporary buffer B* to store the demonstrations of

Algorithm 1: Controller Fusion Network (CFN) during training.

```

Initialize controllers  $\{\mu_i(s)\}_{i=1}^n$  and CFN  $\pi(s|\theta)$  with
random weights  $\theta$ ;
Initialize CFN training database  $D \leftarrow \emptyset$  and CFN temporary
buffers  $\{B_i \leftarrow \emptyset\}_{i=1}^n$  corresponding to the controllers;
// for each controller  $\mu_i$ 
for episode  $\leftarrow 1$  to  $M$  do
  Initial state  $s_1$  provided by the environment;
  for  $t \leftarrow 1$  to  $T$  do
    Controller  $\mu_i$  demonstrates  $a_{\mu_i t} = \mu_i(s_t)$ ;
    Execute controller action  $a_{\mu_i t}$ ; observe new state
     $s_{t+1}$  and feedback;
    Update  $B_i \leftarrow B_i \cup \{(s_t, a_{\mu_i t})\}$ ;
    Discard unwanted demonstrations from  $B_i$ 
    (according to buffering strategy);
    Add desirable demonstrations from  $B_i$  to  $D$ ;
    Sample mini-batch  $(s, a)$  from  $D$  and perform SGD
    to minimize  $\mathbb{L}(\pi(s|\theta), a)$ ;
    Break if  $s_{t+1}$  is terminal state;
  end
end

```

the active controller. At each time step t , by observing the feedback of the interactions between controller and environment, the CFN agent can determine whether to discard some *unwanted* demonstrations from B , retain *unforeseeable* demonstrations in B , or augment its training data D by adding *desirable* demonstrations from B . We call this operation *buffer strategy*. Note that the CFN agent maintains a temporary buffer for each controller. See Figure 4 for details.

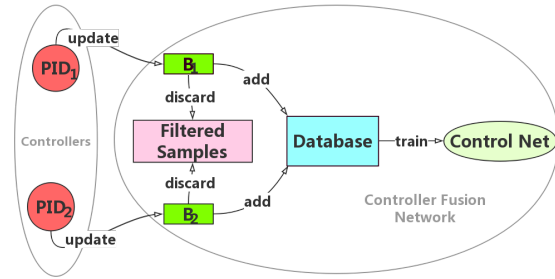


Figure 3: Visualization of the CFN buffer strategy.

CFN buffer strategy. We use two simple PID controllers. The CFN agent maintains two *temporary buffers B* corresponding to each controller. At each time step, the state action pair (s_t, a_t) of each controller will be buffered into B . Then the CFN agent will decide which state action pairs to add to the training database, and discard others based on the *buffer strategy*. Finally, the CFN agent will perform an SGD update on the training database. A schematic diagram

of the buffer used during training is shown in Figure 3.

The training database D can be viewed as a distilled set of demonstrations collected by applying the buffering strategy on B at each time step. Our goal is to train a CFN policy π^* to minimize the loss function $\mathbb{L}(\pi(s|\theta), a^*)$:

$$\pi^*(s|\theta) = \arg \min_{\pi} \mathbb{E}_{(s,a) \sim D} [\mathbb{L}(\pi(s|\theta), a^*)] \quad (1)$$

The total loss consists of the perception loss \mathbb{L}_p and a control loss \mathbb{L}_c :

$$\mathbb{L} = \mathbb{L}_p(\phi(s_I|\theta_p), z^*) + \lambda \mathbb{L}_c(\pi(s|\theta), a^*) \quad (2)$$

$$\mathbb{L}_c(\pi(s|\theta), a^*) = \mathbb{L}_c(\varphi(\phi(s_I|\theta_p), s_M|\theta_c), a^*) \quad (3)$$

Here, θ_p and θ_c are the learnable parameters of the perception and controller fusion network, respectively; λ scales the control loss \mathbb{L}_c ; z^* represents the groundtruth waypoints; s_I and s_M are the input image state and UAV state.

4. Network Architecture and Training Details

Our overall network architecture is depicted in Figure 2. It consists of two networks, one for perception and one for control. The perception network outputs a local trajectory and the control network produces low-level controls given this trajectory and the UAV state as input.

4.1. Perception

Network Architecture. For our perception network, we use a network architecture that is inspired by the one used by Bojarski *et al.* [4] for autonomous driving. However, we make changes to accommodate the complexity of the task at hand and to improve robustness in training. Our network architecture is shown in Figure 2 as the perception module. It consists of eight layers: five convolutional with $\{20, 24, 28, 30, 32\}$ filters and three fully-connected layers with $\{1800, 800, 100\}$ hidden units, respectively. Instead of pooling operations, we introduce convolutional strides of two to downsample the input consecutively, inspired by more recent architecture designs such as MobileNet [8].

Training Details. In contrast to the network by Bojarski *et al.* [4], our network regresses local trajectories rather than raw controls. Predicting raw controls has several shortcomings; raw controls are specific to a vehicle and controller. In addition, different sequences of controls can lead to the same trajectory which makes data augmentation very difficult. As a result, we are able to train a very robust perception network and can validate it with ground-truth trajectories which are well defined. Our perception network takes images from a monocular RGB camera as input and outputs a trajectory relative to the current position of the UAV. We represent the trajectory with five uniformly sampled points. We regress these points to the groundtruth (waypoints along the center line of the track) applying a L2-loss and dropout

of 0.5 in the fully-connected layers. We implement our model in TensorFlow and train it with a learning rate of $5e-4$ using the Adam optimizer. Given the dynamic nature of the racing task we use the maximum frame rate of 60 fps unlike other works that downsample [6, 4, 22]. Note that when predicting controls directly from the input image, a high frame rate can be problematic, since fast transitions in control can lead to similar images with different labels, causing a regression towards averages. Since the image to trajectory correspondence is well defined and more stable, our approach is not affected by sampling rate variations.

4.2. Control

Here, we present the details of our controller fusion network, including network architecture and training strategy. The network takes the predicted trajectory from the perception network and the UAV state (orientation and velocity) as input and predicts the four UAV controls: throttle (T), roll/aileron (A), pitch/elevator (E), yaw/rudder (R).

Learning the Controller Fusion Network. In our experiments, we use two naive PID controllers and denote them as μ_1 and μ_2 . We briefly tune two PID controllers on the training tracks. It is not necessary to tune controllers to be optimal on all the training tracks since CFN is robust to learn from sub-optimal controllers. The first PID controller μ_1 is conservative, as it accurately follows the center of the track and flies through gates precisely but at relatively low speeds. Its output control values are a function of the first predicted waypoint wp_1 and UAV state. The second PID controller μ_2 is more dynamic, as it flies at maximum speed but can often overshoot gates on sharp turns due to inertia and limitations of the UAV. Its output control values are a function of the fourth predicted waypoint wp_4 and the UAV state. We use a three-layer fully connected network to approximate the policy φ of the CFN agent. The state of the CFN agent is a vector concatenation of the predicted waypoints and the UAV state (physical measurements): $s = [s_{wp_{1-n}}, s_M] = [\phi(s_I|\theta_p), s_M]$.

As such, the states of the CFN agent φ , the conservative PID μ_1 , and the aggressive PID μ_2 , are $[s_{wp_{1-n}}, s_M]$, $[s_{wp_1}, s_M]$ and $[s_{wp_4}, s_M]$ respectively. At each time step, the state-action pairs (s, a) in the *temporary buffers* B_1 and B_2 are $([s_{wp_{1-n}}, s_M], \mu_1(s_{wp_1}))$ and $([s_{wp_{1-n}}, s_M], \mu_2(s_{wp_4}))$.

An illustration of the *buffering strategy* is shown in Figure 4 with buffer size $k = 3$. At time step t , there are k unforeseeable samples $\{(s_t, a_t)\}_{t-k}^{t-1}$ in B . The next sample is (s_t, a_t) . The samples ahead of (s_{t-k}, a_{t-k}) are stored in the ground truth training database D . Figure 4 illustrates two cases of buffering operations.

Network Architecture and Training Details. The goal of the control network is to find a control policy φ that mini-

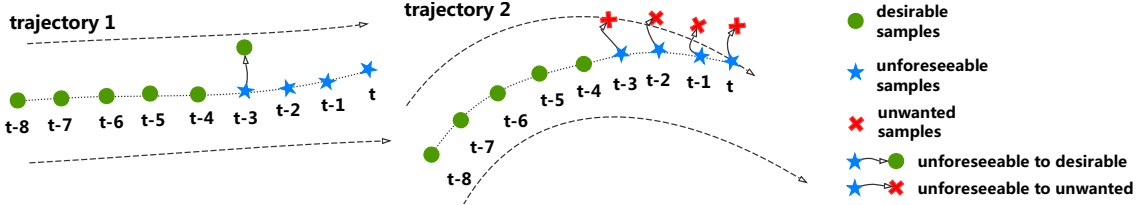


Figure 4: Illustration of the Buffering Strategy. In Trajectory 1 the PID controller remains on track: the sample (s_{t-k}, a_{t-k}) will become a desirable sample and will be added to the ground truth database D . In Trajectory 2 the PID controller leaves the track: the samples $\{(s_t, a_t)\}_{t-k}^t$ will be discarded.

minimizes the control loss \mathbb{L}_c :

$$\varphi^*(s|\theta_c) = \arg \min_{\varphi} \mathbb{E}_{(s,a) \sim D} [\mathbb{L}_c(\varphi(s|\theta_c), a^*)] \quad (4)$$

Here, θ_c represents the learnable parameters of the controller fusion network φ . We use Tensorflow to implement our CFN. A three-layer fully-connected network with $\{64, 32, 16\}$ hidden units is used to represent CFN. To regularize the network, we apply dropout to the second layer with a 0.5 ratio. A weighted $L1$ -loss is used for our loss function \mathbb{L}_c . We use an Adam Optimizer with a learning rate of $1e-3$ to train CFN in an online fashion, while the UAV flies through a training track. A *temporary buffer* with size k is used to temporally store the last k state-action pairs from each controller ($k_1 = 1$ for B_1 and $k_2 = 50$ for B_2). The buffer size mainly depends on the controller’s trajectory quality. For example, if an aggressive PID controller overshoots or crashes after 100 steps at a straight and 50 steps at a bend with a high probability, 40–60 would be considered a good buffer size; a buffer size of 0–20 would lead the CFN agent to learn dangerous behaviours, while a buffer size of over 100 would make the CFN agent unable to benefit from the speed advantages of the aggressive PID controller on straightaways. In our case, we chose $k_1 = 1$ for the conservative PID controller to benefit from its accuracy and $k_2 = 50$ for the aggressive PID controller to filter out its dangerous behaviors in curves. To improve learning, we add an Ornstein Uhlenbeck process [14] noise to the output of the controller fusion network to allow for exploration at the beginning and move the UAV by the agent’s control predictions, but the action is labeled by the conservative PID. After the initial exploration, the controllers take over control. At each time step, if a controller leaves the track, its temporary buffer is flushed and no state-action pairs are added to D for training. If the controller keeps the UAV within the track boundaries, the oldest state-action pair in its temporary buffer is added to D , as shown in Figure 4. At each time step, the network parameters θ_c are updated by back propagation to minimize the difference between the controller’s UAV policy φ and the demonstrations (state-action pairs) in D , which are considered to be believable after discarding unwanted behaviors.

5. Experiments

Creation of the UAV Racing Simulation. Many professional pilots compete in time trials on well-known tracks such as those posted by the MultiGP Drone Racing League. Following this paradigm, our simulator race course is modeled after a football stadium, where local professional pilots regularly setup MultiGP tracks. Using a combination of LiDAR scanning and aerial photogrammetry, we captured the stadium with an accuracy of 0.5 cm (see Figure 5). A team of architects used the dense point cloud and textured mesh to create an accurate solid model with physics based rendering (PBR) textures in 3DS Max for export to Unreal. This resulted in a geometrically accurate and photo-realistic race course that remains low in poly count, so as to run within Sim4CV in real-time, in which all training and testing experiments are conducted. We refer to Figure 6 for a side-by-side comparison of the real and virtual stadiums.

Experimental Setup. We use our UAV racing environment in Sim4CV [16] (see Figure 1); we design seven racing tracks for training and seven tracks for testing. To avoid user bias, we collect online images and trace their contours to create uniquely stylized tracks. We select the tracks with two aspects in mind. (1) The tracks should be similar to what racing professionals are accustomed to, and (2) they should offer enough diversity for network generalization on unseen tracks (see Figure 7). For all of the following evaluations, both the trained networks and human pilots are tasked

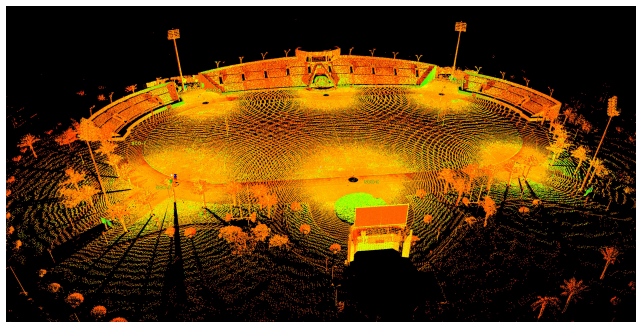


Figure 5: The point cloud from the LiDAR scan of the stadium collected from six different locations.



Figure 6: *Left*: Aerial image captured from an UAV hovering above the stadium racing track. *Right*: Rendering of the reconstructed stadium generated at a similar altitude and viewing angle within the simulator.

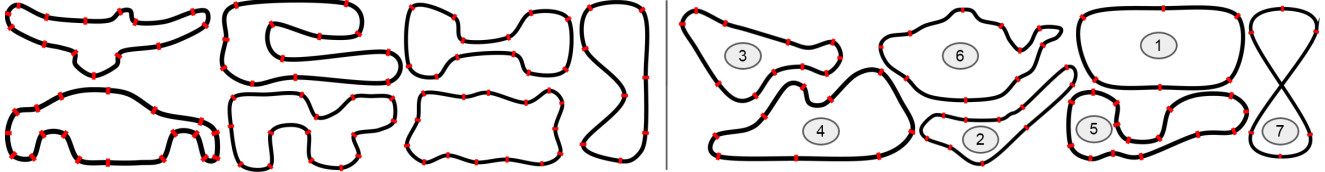


Figure 7: The seven training tracks (left) and the seven testing tracks (right). Gates are marked in red.

with flying two laps on each of the test tracks. The score comprises three components: the percentage of successfully passed gates, the time to complete both laps, and the number of required resets. The UAV is reset at the next gate, if it does not reach it within 10 seconds after passing through the previous gate. This occurs if the UAV crashes beyond recovery or drifts off the track. Visualizations of the UAV’s trajectory for all models on each track are provided in the supplementary material.

Comparison to State-of-the-Art Baselines. We compare our system for UAV racing to the two most related and recent network architectures, the first denoted as Nvidia (for self-driving cars [4]) and the second as MAV (for forest path navigating UAVs [22]). Both the Nvidia and MAV networks use data augmentation from an additional left and right camera. For the Nvidia network, the exact offset choices for training are not publicly known, so we use a rotational offset of $\{-30^\circ, 30^\circ\}$. For the MAV network, we use the same augmentation parameters in the paper, *i.e.* a rotational offset of $\{-30^\circ, 30^\circ\}$. We modify the MAV network to allow for a regression output instead of its original classification (left, center and right controls). This is necessary, since our task requires fine-grained control, and discrete controls would be insufficient. We assign corrective controls to the augmentation views using a fairly simple but effective strategy. Depending on the camera view, we apply the following offset parameters: one that acts as a horizontal offset (roll-offset) and one that acts as a rotational offset (yaw-offset). For rotational offsets, we couple the yaw correction with a proportional roll correction because the UAV is in motion while rotating, causing it to drift outwards due to its inertia.

While the domains of these methods are similar, it should

be noted that flying a high-speed racing UAV is a particularly challenging task, since the effect of inertia is more significant and there are more degrees of freedom than for ground vehicles. To ensure a strong end-to-end baseline, we build an end-to-end network that takes the state of the UAV (exactly like our control module) as input along with the image. We also augment the data with 18 additional camera views (exactly like our perception module) and assign the best corrective controls after cross-validation search [17].

Table 1 compares the performance of these baselines against our method. The MAV reference network needs more than 7 resets and only completes about 60% of gates on average, while taking more than twice the time. The Nvidia-inspired architecture performs slightly better, but still needs about 4 resets and only completes 80% of gates on average. While the end-to-end trained version of our network achieves better performance than MAV and Nvidia, our modular network with CFN clearly outperforms it without the need for supervision or approximate corrective controls. In fact, CFN outperforms all baselines by a considerable margin in all three evaluation metrics; it completes all seven race tracks with 100% accuracy compared to 62.69%, 79.85% and 95.52% in about half the time on average.

Comparison to PID Controllers and Human Performance. We compare our CFN trained control network to the PID controllers it learned from. The perception network stays the same. The summary of this experiment is given in Table 2, where we find that our learned policy is able to outperform both PID controllers (PID1 and PID2). Further, as Figure 8 shows, both PID controllers are imperfect. PID1 completes most gates while flying very slowly and PID2 misses many gates while flying at very high speeds.

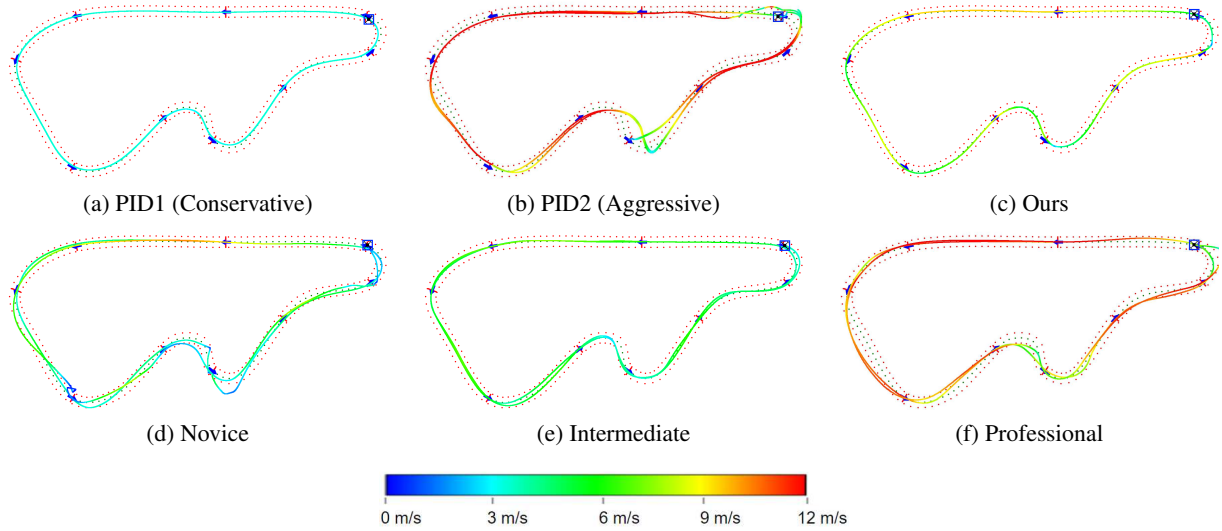


Figure 8: Comparison between our learned CFN policy and PID controllers (*row 1*) and human pilots (*row 2*), on a test track. Color encodes speed as a heatmap, where blue is the minimum speed and red is the maximum speed.

Table 1: Comparison of CFN to baselines

	End2End (MAV)			End2End (Nvidia)			End2End (Ours)			Ours (WP + CFN)		
	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets
Track1	6/12	98.60	6	7/12	101.15	5	12/12	80.11	0	12/12	52.20	0
Track2	16/20	113.55	4	15/20	140.92	5	20/20	91.88	0	20/20	64.75	0
Track3	11/22	161.85	11	19/22	110.94	3	22/22	81.26	0	22/22	62.00	0
Track4	10/18	152.27	8	15/18	121.07	3	18/18	97.10	0	18/18	71.93	0
Track5	18/30	207.07	12	21/30	197.11	9	30/30	100.47	0	30/30	71.16	0
Track6	15/20	136.69	5	20/20	108.42	0	16/20	137.14	4	20/20	81.66	0
Track7	8/12	115.05	4	10/12	105.37	2	10/12	104.97	2	12/12	64.86	0
Avg.	62.69%	140.72	7.14	79.85%	126.43	3.86	95.52%	98.99	0.86	100%	66.94	0

Table 2: Ablation study

	PID1 (Conservative)			PID2 (Aggressive)			Ours (No Buffer)			Ours (WP + CFN)		
	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets
Track1	12/12	130.76	0	12/12	40.04	0	10/12	70.35	2	12/12	52.20	0
Track2	20/20	136.19	0	17/20	77.41	3	18/20	75.58	2	20/20	64.75	0
Track3	22/22	121.54	0	11/22	149.45	11	17/22	102.72	5	22/22	62.00	0
Track4	18/18	139.09	0	15/18	81.08	3	14/18	102.27	4	18/18	71.93	0
Track5	30/30	144.49	0	12/30	212.79	18	28/30	89.93	2	30/30	71.16	0
Track6	20/20	151.95	0	12/20	118.69	8	13/20	126.77	7	20/20	81.66	0
Track7	10/12	139.28	2	9/12	72.90	3	7/12	86.53	5	12/12	64.86	0
Avg.	98.51%	137.61	0.29	65.67%	107.48	6.57	79.85%	93.45	3.86	100%	66.94	0

Table 3: Comparison of CFN to humans

	Human (Novice)			Human (Intermediate)			Human (Professional)			Ours (WP + CFN)		
	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets	Score	Time	Resets
Track1	12/12	87.44	0	12/12	62.80	0	12/12	40.50	0	12/12	52.20	0
Track2	20/20	166.11	0	20/20	88.21	0	20/20	49.23	0	20/20	64.75	0
Track3	21/22	118.41	1	22/22	82.17	0	22/22	47.67	0	22/22	62.00	0
Track4	17/18	126.47	1	18/18	91.53	0	18/18	50.10	0	18/18	71.93	0
Track5	30/30	129.49	0	30/30	87.62	0	30/30	55.72	0	30/30	71.16	0
Track6	20/20	196.16	0	20/20	95.99	0	20/20	57.90	0	20/20	81.66	0
Track7	12/12	113.14	0	12/12	74.91	0	12/12	46.98	0	12/12	64.86	0
Avg.	98.51%	133.89	0.29	100%	83.32	0	100%	49.73	0	100%	66.94	0



Figure 9: Simulated UAV racing stadium with low and high quality grass textures (*a,b*), different ground materials (*c,d*), different weather conditions (*e,f*) and different lighting conditions (*g,h*).

Table 4: Adaptation through modularity. The reported results reflect average performance across all seven testing tracks. Please refer to the supplementary material for detailed results per track.

	Baseline		Weather		Lighting		Texture	
	Grass	GrassHD	Fog	Rain	Sunrise	Night	Snow	Mud
Score	100%	100%	95.5%	100%	99.2%	96.2%	100%	100%
Time	66.94	69.75	85.48	68.68	73.77	86.01	68.26	67.54
Reset	0	0	0.86	0	0.14	0.71	0	0

However, since our control module is designed to learn only from the best behaviour of both PID controllers, it completes all gates at a high speed. We also want to highlight the importance of the *temporary buffer*. Removing it results in a significant drop in performance (score: 79.85%, time: 93.45, resets: 3.86) since the CFN agent is forced to learn from all the controller demonstrations including the undesirable ones.

We also compare our system to three pilots with different skill levels: novice (has never flown before), interme-

diated (a moderately experienced pilot), and a professional (a competitive racing pilot with many years of experience). The pilots are given the opportunity to fly the seven training tracks as many times as needed until they successfully complete the tracks at their best time while passing through all gates. The pilots are then scored on the test tracks in the same fashion as the trained networks. The results are summarized in Table 3. CFN achieves at least the same accuracy as human pilots but is about two times faster than a novice, 20% faster than an intermediate pilot, and within 25% of a professional pilot. Although slower, our network flies more consistently than even the professional racing pilot while remaining reliably on the track (see Figure 8).

Adaptation through Modularity. We replace the low-quality grass textures with high-quality grass and show that our perception network generalizes without any modification (see Figure 9 and Table 4). When changing the weather and lighting conditions, the performance of the perception network starts degrading. The network is not affected by rain, slightly degrades with sunrise lighting and degrades noticeably with fog and night-time lighting. If our perception network was trained on diverse environments/textures, it would learn even more invariance to the background, as demonstrated in [20]. However, generalization only works up to some extent and usually requires heavy data augmentation. In some applications, it might not even be desirable to generalize too broadly as the performance in the target domain often suffers as a result. For such cases, our modular approach allows to simply swap out the perception module to adapt to any environment. To demonstrate this, we train the perception network on different textured environments while keeping the controller fusion network fixed and show successful transfer of the control policy.

6. Conclusions and Future Work

In this work, we present a controller fusion network (CFN) that allows fusing multiple classical controllers. Extensive experiments demonstrate that a CFN based network outperforms state-of-the-art methods and flies more consistently than human pilots. This is a product of both the ability for the network to fuse multiple controller’s trajectories and at the same time filter out controller actions leading to poor performance. We expect the framework can be adapted for other robotic and controller based dynamic tasks such as visual grasping tasks or visual placing tasks by making minor changes in buffer strategy. Instead of relying on extensive fine-tuning of a controller or defining an explicit model of a system, a CFN is able to produce an optimized predictive control of dynamic systems.

Acknowledgments This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research.

References

- [1] Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y. Ng. An application of reinforcement learning to aerobatic helicopter flight. In P. B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1–8. MIT Press, 2007.
- [2] Olov Andersson, Mariusz Wzorek, and Patrick Doherty. Deep learning quadcopter control via risk-aware active learning. In *Thirty-First AAAI Conference on Artificial Intelligence (AAAI), 2017, San Francisco, February 4-9. .*, 2017. Accepted.
- [3] Valts Blukis, Nataly Brukhim, Andrew Bennett, Ross A. Knepper, and Yoav Artzi. Following high-level navigation instructions on a simulated quadcopter with imitation learning. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018.
- [4] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. End to end learning for self-driving cars. *CoRR*, abs/1604.07316, 2016.
- [5] Kai-Wei Chang, Akshay Krishnamurthy, Alekh Agarwal, Hal Daume III, and John Langford. Learning to search better than your teacher. 2015.
- [6] Chenyi Chen, Ari Seff, Alain Kornhauser, and Jianxiong Xiao. Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), ICCV '15*, pages 2722–2730, Washington, DC, USA, 2015. IEEE Computer Society.
- [7] M. Hehn and R. DAndrea. Real-time trajectory generation for quadcopters. *IEEE Transactions on Robotics*, 31(4):877–892, Aug 2015.
- [8] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [9] Gregory Kahn, Tianhao Zhang, Sergey Levine, and Pieter Abbeel. PLATO: policy learning using adaptive trajectory optimization. *CoRR*, abs/1603.00622, 2016.
- [10] Elia Kaufmann, Antonio Loquercio, Rene Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. Deep drone racing: Learning agile flight in dynamic environments. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 133–145. PMLR, 29–31 Oct 2018.
- [11] Dong Ki Kim and Tsuhan Chen. Deep neural network for real-time autonomous indoor navigation. *CoRR*, abs/1511.04668, 2015.
- [12] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.*, 17(1):1334–1373, Jan. 2016.
- [13] Sergey Levine and Vladlen Koltun. Guided policy search. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1–9, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [14] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *ICLR*, abs/1509.02971, 2016.
- [15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmasharan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, Feb. 2015.
- [16] Matthias Müller, Vincent Casser, Jean Lahoud, Neil Smith, and Bernard Ghanem. Sim4cv: A photo-realistic simulator for computer vision applications. *Int. J. Comput. Vision*, 126(9):902–919, Sept. 2018.
- [17] M. Müller, V. Casser, N. Smith, D. L. Michels, and B. Ghanem. Teaching UAVs to Race: End-to-End Regression of Agile Controls in Simulation. In *European Conference on Computer Vision Workshop (ECCVW)*, Sept. 2018.
- [18] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos Theodorou, and Byron Boots. Agile autonomous driving using end-to-end deep imitation learning. 06 2018.
- [19] Stéphane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. No-regret reductions for imitation learning and structured prediction. *CoRR*, abs/1011.0686, 2010.
- [20] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *CoRR*, abs/1611.04201, 2016.
- [21] Utsav Shah, Rishabh Khawad, and K Madhava Krishna. Deepfly: Towards complete autonomous navigation of mavs with monocular camera. In *Proceedings of the Tenth Indian Conference on Computer Vision, Graphics and Image Processing, ICVGIP '16*, pages 59:1–59:8, New York, NY, USA, 2016. ACM.
- [22] N. Smolyanskiy, A. Kamenev, J. Smith, and S. Birchfield. Toward Low-Flying Autonomous MAV Trail Navigation using Deep Neural Networks for Environmental Awareness. *ArXiv e-prints*, May 2017.
- [23] Riccardo Spica, Davide Falanga, Eric Cristofalo, Eduardo Montijano, Davide Scaramuzza, and Mac Schwager. A real-time game theoretic planner for autonomous two-player drone racing. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018.
- [24] Jiakai Zhang and Kyunghyun Cho. Query-efficient imitation learning for end-to-end autonomous driving. *arXiv preprint arXiv:1605.06450*, 2016.