

One Net to Rule Them All

Domain Randomization in Quadcopter Racing Across Different Platforms

Ferede, Robin; Blaha, Till; Lucassen, Erin; De Wagter, Christophe; De Croon, Guido C.H.E.

DOI

[10.1109/ICRA55743.2025.11128790](https://doi.org/10.1109/ICRA55743.2025.11128790)

Publication date

2025

Document Version

Final published version

Published in

2025 IEEE International Conference on Robotics and Automation, ICRA 2025

Citation (APA)

Ferede, R., Blaha, T., Lucassen, E., De Wagter, C., & De Croon, G. C. H. E. (2025). One Net to Rule Them All: Domain Randomization in Quadcopter Racing Across Different Platforms. In C. Ott, H. Admoni, S. Behnke, S. Bogdan, A. Boleaon, Y. Choi, F. Ficuciello, N. Gans, C. Gosselin, K. Harada, E. Kayacan, H. J. Kim, S. Leutenegger, Z. Liu, P. Maiolino, L. Marques, T. Matsubara, A. Mavromatti, M. Minor, J. O'Kane, H. W. Park, H.-W. Park, I. Rekleitis, F. Renda, E. Ricci, L. D. Riek, L. Sabattini, S. Shen, Y. Sun, P.-B. Wieber, K. Yamane, ... J. Yu (Eds.), *2025 IEEE International Conference on Robotics and Automation, ICRA 2025* (pp. 6357-6363). (Proceedings - IEEE International Conference on Robotics and Automation). IEEE. <https://doi.org/10.1109/ICRA55743.2025.11128790>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:
the publisher is the copyright holder of this work and the
author uses the Dutch legislation to make this work public.

One Net to Rule Them All: Domain Randomization in Quadcopter Racing Across Different Platforms

Robin Ferede, Till Blaha, Erin Lucassen, Christophe De Wagter, Guido C.H.E. de Croon

Abstract—In high-speed quadcopter racing, finding a single controller that works well across different platforms remains challenging. This work presents the first neural network controller for drone racing that generalizes across physically distinct quadcopters. We demonstrate that a single network, trained with domain randomization, can robustly control various types of quadcopters. The network relies solely on the current state to directly compute motor commands. The effectiveness of this generalized controller is validated through real-world tests on two substantially different crafts (3-inch and 5-inch race quadcopters). We further compare the performance of this generalized controller with controllers specifically trained for the 3-inch and 5-inch drone, using their identified model parameters with varying levels of domain randomization (0%, 10%, 20%, 30%). While the generalized controller shows slightly slower speeds compared to the fine-tuned models, it excels in adaptability across different platforms. Our results show that no randomization fails sim-to-real transfer while increasing randomization improves robustness but reduces speed. Despite this trade-off, our findings highlight the potential of domain randomization for generalizing controllers, paving the way for universal AI controllers that can adapt to any platform.

Index Terms—Drone Racing, Reinforcement Learning, Domain Randomization, Reality Gap, Sim-to-real Transfer

I. INTRODUCTION

The drone market is experiencing rapid growth [1], yet most applications still rely on human pilots. Given that battery life limits flight times, enhancing the speed and autonomy of drones offers significant advantages [2]. As a result, research into autonomous drone racing is a growing field that contributes to advancing technology and enhancing drone capabilities.

One of the significant challenges in drone racing is the development of controllers that generalize well across different quadcopter platforms. While human pilots can easily adapt to various drones with minimal adjustments, current racing AI systems often overfit to their specific platforms, restricting their generalization capabilities. This limitation hampers the transfer of current drone racing technology to other real-world applications [2].

Recent advancements in quadcopter control have predominantly focused on Reinforcement Learning (RL) techniques [3]–[7]. These methods have demonstrated impressive results in drone racing, with neural controllers surpassing human champions [6]. However, these achievements rely on precise

¹The authors are with the Micro Air Vehicle Lab of the Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, The Netherlands robinferede@tudelft.nl, t.m.blaha@tudelft.nl, e.lucassen@student.tudelft.nl, c.deWagter@tudelft.nl, g.c.h.e.deCroon@tudelft.nl

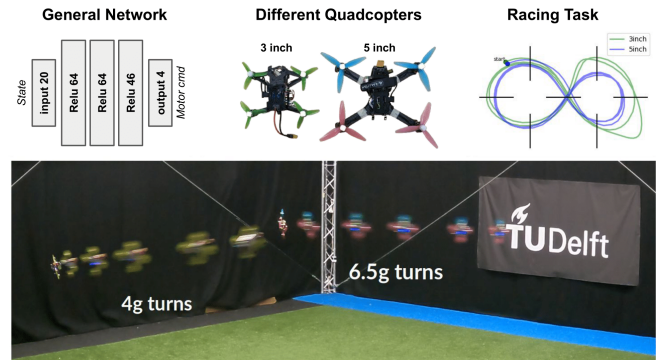


Fig. 1. We train a single neural network to control physically distinct drones. The time-lapse image demonstrates the successful sim-to-real transfer of the reinforcement-learned network, enabling it to perform drone racing with both 3-inch and 5-inch quadcopters.

modeling of the drone platforms and struggle to generalize across different drones.

Domain randomization (DR) has emerged as a prominent method for transferring RL policies across different environments. This technique involves training policies in simulated environments with varied parameters to improve their robustness when deployed in real-world scenarios [8]. For instance, DR has facilitated single-shot sim-to-real transfer for robot manipulation [9], visual navigation using synthetic data [10], and object detection [11]. Additionally, in end-to-end quadcopter control, DR, combined with real-time adaptation, has been critical for bridging the reality gap [12]–[14]. However, although DR enhances robustness, it comes with trade-offs, including longer training times and diminished performance optimality [8], [15].

Current research on generalized quadcopter control has only focused on low-level tasks in near-hover or recovery-to-hover conditions, with little exploration into racing scenarios that demand both guidance and control at significantly higher speeds. Recent studies have explored various approaches: one developed a network for stabilizing low-level control (position control and trajectory tracking) that generalizes to various quadcopters using domain randomization [16]. Another method [17] introduced a single-shot real-time technique for controlling an unknown quadcopter through iterative determination of the actuator effectiveness parameters of an INDI controller, expanding on the work of [18]. Additionally, recent work trained a neural network to estimate a latent representation of quadcopter parameters, conditioning another neural network controller for near-hover flight, enabling online adaptation to varying hardware and

disturbances during real-world deployment [19].

In this work, we extend these results to drone racing by using domain randomization to develop a neural network controller for high-speed racing that generalizes across different quadcopters. Our controller maps state information directly to motor commands, enabling end-to-end guidance and control. Our key contributions include:

- 1) Introducing, to the best of our knowledge, the first neural network controller capable of operating across different drones for high-speed racing
- 2) Investigation into the effects of varying degrees of domain randomization on sim-to-real transferability and flight speed performance.

This paper is organized as follows: Sec. II presents the drone model and control formulation; Sec. III describes the experimental setup; Sec. IV analyzes the performance of the general control network and examines the effects of varying degrees of domain randomization; and Sec. V provides the conclusion and outlines directions for future work.

II. METHODOLOGY

A. Quadcopter model

To train our neural policies we simulate our quadcopter dynamics with a parametric model similar to the one outlined in [14]. The quadcopter's state and control are given by:

$$\mathbf{x} = [\mathbf{p}, \mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\Omega}, \boldsymbol{\omega}]^T \quad \mathbf{u} = [u_1, u_2, u_3, u_4]^T$$

Here \mathbf{p} is the position of the drone, \mathbf{v} is the velocity, $\boldsymbol{\lambda}$ are the Euler angles, $\boldsymbol{\Omega}$ are the body rates and $\boldsymbol{\omega}$ are the propeller speed in rad/s. $\mathbf{u} \in [0, 1]^4$ represents the normalized commands sent to the motor controllers. The equations of motion are expressed as follows:

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v} & \dot{\mathbf{v}} &= g\mathbf{e}_3 + R(\boldsymbol{\lambda})\mathbf{F} & (1) \\ \dot{\boldsymbol{\lambda}} &= Q(\boldsymbol{\lambda})\boldsymbol{\Omega} & \dot{\boldsymbol{\Omega}} &= \mathbf{M} \\ \dot{\omega}_i &= (\omega_{ci} - \omega) / \tau & & & (2) \end{aligned}$$

Where R is the rotation matrix and Q transforms body rates to Euler angle derivatives. The steady state motor response is modeled as: $\omega_{ci} = (\omega_{\max} - \omega_{\min}) \sqrt{k_{li} u_i^2 + (1 - k_{li}) u_i} + \omega_{\min}$. The specific force F and the moment M ¹ are modeled as:

$$\mathbf{F} = \left[-\sum_{i=1}^4 k_x v_x^B \omega_i, -\sum_{i=1}^4 k_y v_y^B \omega_i, -\sum_{i=1}^4 k_\omega \omega_i^2 \right]^T$$

$$\mathbf{M} = \begin{bmatrix} -k_{p1} \omega_1^2 - k_{p2} \omega_2^2 + k_{p3} \omega_3^2 + k_{p4} \omega_4^2 \\ -k_{q1} \omega_1^2 + k_{q2} \omega_2^2 - k_{q3} \omega_3^2 + k_{q4} \omega_4^2 \\ -k_{r1} \omega_1 + k_{r2} \omega_2 + k_{r3} \omega_3 - k_{r4} \omega_4 + \\ \dots - k_{r5} \dot{\omega}_1 + k_{r6} \dot{\omega}_2 + k_{r7} \dot{\omega}_3 - k_{r8} \dot{\omega}_4 \end{bmatrix}$$

The parameters $k_{\square}, \omega_{\min}, \omega_{\max}$ are estimated from manual flight data for both the 3-inch and 5-inch quadcopters using linear regression. We normalize the parameters based on the

¹ M represents angular acceleration. The moment of inertia is indirectly estimated through the k_{\square} and gyroscopic effects are ignored.

TABLE I
PARAMETERS IDENTIFIED FOR 3INCH AND 5INCH DRONE

param	3 inch	5 inch	param	3 inch	5 inch
\hat{k}_ω	14.3	27.1	ω_{\min} (rad/s)	305.4	238.49
\hat{k}_x	0.16	0.16	ω_{\max} (rad/s)	4887.57	3295.5
\hat{k}_y	0.18	0.24	k_l	0.84	0.95
			τ (sec)	0.04	0.04
\hat{k}_{p1}	615.0	711.0	\hat{k}_{r1}	47.1	35.2
\hat{k}_{p2}	598.0	718.0	\hat{k}_{r2}	47.1	35.2
\hat{k}_{p3}	650.0	691.0	\hat{k}_{r3}	47.1	35.2
\hat{k}_{p4}	479.0	724.0	\hat{k}_{r4}	47.1	35.2
\hat{k}_{q1}	217.0	573.0	\hat{k}_{r5}	5.57	6.49
\hat{k}_{q2}	238.0	637.0	\hat{k}_{r6}	5.57	6.49
\hat{k}_{q3}	280.0	548.0	\hat{k}_{r7}	5.57	6.49
\hat{k}_{q4}	196.0	640.0	\hat{k}_{r8}	5.57	6.49

maximum angular velocity ω_{\max} . The normalized parameters are:

$$\begin{aligned} \hat{k}_\omega, \hat{k}_{pi}, \hat{k}_{qi} &= (k_\omega, k_{pi}, k_{qi}) \times \omega_{\max}^2 \\ \hat{k}_{xi}, \hat{k}_{yi}, \hat{k}_{ri} &= (k_{xi}, k_{yi}, k_{ri}) \times \omega_{\max} \end{aligned}$$

This scaling accounts for variations in ω_{\max} and ensures consistency across the model. The estimated values can be found in Tab.I. These parameters highlight the distinct control characteristics of each platform. For instance, the 3-inch platform has approximately half the thrust-to-weight ratio (proportional to \hat{k}_ω), motors that spin 50% faster, and less than half the pitch effectiveness (\hat{k}_{pi}) compared to the 5-inch platform.

B. RL problem definition

The racing setup consists of seven square gates, each measuring 1.5x1.5 meters, arranged in a figure-eight track, as shown in Fig. 1. The drone starts 1 meter in front of a randomly selected gate, in a similar fashion as [3]. The other initial state variables are defined as follows: $\mathbf{v} \in [-0.5, 0.5]^3$, $\phi, \theta \in [-\frac{\pi}{9}, \frac{\pi}{9}]$, $\psi \in [-\pi, \pi]$, $\boldsymbol{\Omega} \in [-0.1, 0.1]^3$, and $\omega \in [\omega_{\min}, \omega_{\max}]$. Similar to [5] the reward function to be optimised consists of a progress reward, a rate penalty and a collision penalty:

$$r_k = \begin{cases} -10, & \text{if collided} \\ \left(\|\mathbf{p}_{k-1} - \mathbf{p}_{gk}\| - \|\mathbf{p}_k - \mathbf{p}_{gk}\| - c \right) \|\boldsymbol{\Omega}\| & \text{otherwise} \end{cases}$$

Here, \mathbf{p}_{gk} denotes the position of the center of the current target gate, while $\mathbf{p}_k, \mathbf{p}_{k-1}$ are the drone's current and previous positions. The rate penalty coefficient is set to $c = 0.001$. A collision is registered when the drone either makes contact with the ground or when it flies outside of a 10x10x7m bounding box. When a collision is registered, the episode ends. Additionally when a gate is missed (i.e. drone passes gate plane outside of 1.5x1.5m area, the episode also ends.

This reward function, together with the quadcopter model from the previous section, is used to create a gym environment for training with the PPO algorithm [20], utilizing the Python library Stable-Baselines3 [21]. Our implementation simulates 100 drones in parallel, uses a discount factor of $\gamma = 0.999$ and a maximum episode length of 1200 time steps

TABLE II
RANDOMIZATION SCHEME FOR THE GENERAL POLICY

Parameter	Distribution	Parameter	Distribution
ω_{\min}	$\mathcal{U}(0, 500)$	\hat{k}_{ω}	$\mathcal{U}(10, 30)$
ω_{\max}	$\mathcal{U}(3000, 5000)$	\hat{k}_x	$\mathcal{U}(0.1, 0.3)$
k_l	$\mathcal{U}(0, 1)$	\hat{k}_y	$\mathcal{U}(0.1, 0.3)$
τ	$\mathcal{U}(0.01, 0.1)$		
\hat{k}_p	$\mathcal{U}(200, 800)$	$\hat{k}_{p1}, \hat{k}_{p2}, \hat{k}_{p3}, \hat{k}_{p4}$	$\hat{k}_p \pm \mathcal{U}(50)$
\hat{k}_q	$\mathcal{U}(200, 800)$	$\hat{k}_{q1}, \hat{k}_{q2}, \hat{k}_{q3}, \hat{k}_{q4}$	$\hat{k}_q \pm \mathcal{U}(50)$
\hat{k}_r	$\mathcal{U}(20, 80)$	$\hat{k}_{r1}, \hat{k}_{r2}, \hat{k}_{r3}, \hat{k}_{r4}$	\hat{k}_r
\hat{k}_{rd}	$\mathcal{U}(2, 8)$	$\hat{k}_{r5}, \hat{k}_{r6}, \hat{k}_{r7}, \hat{k}_{r8}$	\hat{k}_{rd}

(12 seconds). The training process concludes after a total of 100 million time steps. The code for our implementation is openly available on GitHub².

C. Randomization

1) *General policy*: The randomization outlined in Tab. II trains the general policy to control the 3-inch and 5-inch quadcopters. Uniform distributions were designed to encompass parameters from both sizes.

2) *Fine tuned policy*: We train five policies for the 3- and 5 inch quadcopter by applying randomization to the values from Tab. I, introducing variations of 0%, 10%, 20%, and 30%. Specifically, each parameter is multiplied by $\mathcal{U}(1 - p, 1 + p)$, where p represents the randomization percentage. The thrust linearization constant k_l is an exception, as it is capped to ensure it remains below 1.

D. Policy

The selected neural policy is a three-layer fully connected network with ReLU activation functions, with 64 neurons per layer (see Sec. IV-A for motivation). The policy takes in 20 observations, including the quadcopter's state and information about current and future gates, similar to prior work [14]:

$$\mathbf{x}_{obs} = [\mathbf{p}^{g_i}, \mathbf{v}^{g_i}, \boldsymbol{\lambda}^{g_i}, \boldsymbol{\Omega}, \boldsymbol{\omega}, \mathbf{p}_{g_{i+1}}^{g_i}, \boldsymbol{\psi}_{g_{i+1}}^{g_i}]^T \quad (3)$$

Here, g_i denotes the reference frame of the i -th gate, and $\mathbf{p}_{g_{i+1}}^{g_i}$ and $\boldsymbol{\psi}_{g_{i+1}}^{g_i}$ represent the position and orientation of the next gate. The network outputs four motor commands (\mathbf{u}). Because domain randomization is applied to the model parameters, these parameters become part of the state in the Markov decision process (MDP). As a result, the full state is not completely observable with the current observation formulation [15]. Some approaches try to address this by incorporating state or action history into the observation [6], [22], while others use parameter inputs [19]. In Sec. IV-A, we explore these methods and justify our choice of Eq. 3.

III. EXPERIMENTAL SETUP

We test our policies on two quadcopters. Common between them is the compute hardware and experimental infrastructure: we use an STM32H743 microcontroller running INDIflight³, a fork of Betaflight. An onboard EKF fuses

²https://github.com/tudelft/optimal_quad_control_RL/tree/icra2025

³<https://github.com/tudelft/indiflight>

data from a TDK InvenSense ICM-42688-P IMU, and position/attitude observations from an external optical motion capture system (Optitrack). The state is then used by the Neural Network, onboard, at an update rate of 1000Hz.




Fig. 2. This picture shows the drone used for testing. Instead of bolts, four cornerpieces of Moongel are placed on top of the base frame, on which a 3D printed casing with the flight controller is placed. A second layer of Moongel is then placed on the casing, which is then slightly compressed downward by means of two zip-ties.

During initial testing, at high speeds, propeller vibrations caused the accelerometer to saturate, exceeding the $\pm 16g$ limit of the accelerometer. We aimed to address this problem by adding a vibration damping pad as shown in Fig. 2. Sorbothane, adhesive tape and Moongel Damper Pads [23] were compared, but the latter proved to provide the most effective high-frequency resonance control whilst still transmitting the lower frequency true movements of the drone. This result aligns with findings from previous research on damper pad performance, where Moongel was indeed shown to be optimal at eliminating noise [24]. The 3D-printed casing, while not mandatory, aids in distributing the forces more evenly across the FC, whereas the zip-ties connected to the base frame provide preloaded compression for maintaining stability and improving damping performance.

IV. RESULTS

A. Selecting Architecture

We evaluated several architectures using the randomized simulator and selected the 3-layer 64-neuron architecture as the best based on the highest obtained mean episode rewards from three independent runs, each consisting of 100 million time steps. Network size comparisons are illustrated in Fig. 3. Next, we modified the selected architecture to address partial observability by adding various inputs. The modifications included:

- Adding the last 1, 2, or 3 actions to the observation
- Adding the last 1, 2, or 3 actions with stepsize 10 (steps of 0.1s in the past)
- Adding state history (last 1, 2, or 3 states, with or without stepsize 10)
- Adding model parameters (ground truth) or with 10% and 20% noise

Each variation followed the same training setup and was evaluated over 1000 rollouts. See Fig. 4 for results. None of the modified inputs showed a clear improvement, leading us to exclude them from further use. This unexpected result indicates a need for further investigation, which may involve longer training periods, hyperparameter tuning, or alternative

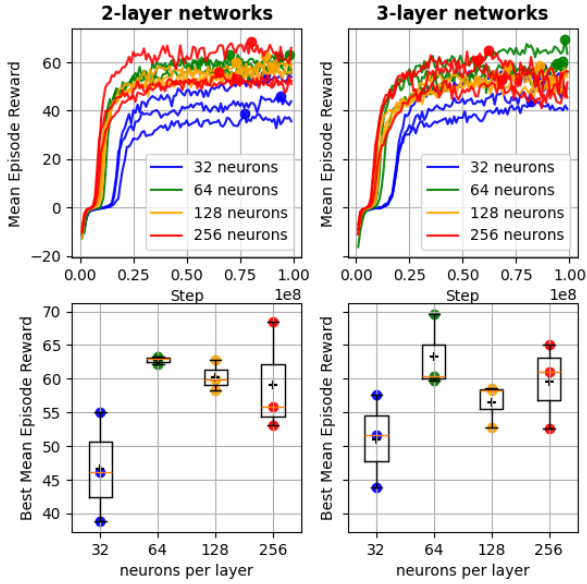


Fig. 3. Comparison of network sizes: The 64,64,64 architecture achieved the highest mean episode reward among the evaluated architectures, based on three independent runs with 100 million time steps each.

TABLE III
SIMULATION RESULTS AVERAGED OVER 1000 POLICY ROLLOUTS

network	3inch sim				5inch sim			
	ep rew	ep len	#gates	%crash	ep rew	ep len	#gates	%crash
general	64.67	1140	40	5.4	76.78	1186	52	1.2
3inch 30%	59.97	1200	41	0.0	3.95	204	4	99.8
3inch 20%	69.02	1200	43	0.0	-6.24	127	3	100
3inch 10%	77.59	1197	46	0.3	-2.24	70	2	100
3inch 0%	78.55	1193	46	0.6	-1.78	89	2	100
5inch 30%	4.23	195	3	100	74.03	1199	49	0.1
5inch 20%	-0.35	125	1	100	88.33	1197	58	0.3
5inch 10%	4.57	218	4	100	92.76	1183	55	1.5
5inch 0%	-0.69	143	1	100	99.35	1179	59	2.0

parameter representations. For the remainder of this paper, we will focus on the network without additional inputs.

B. Simulation

We trained the selected architecture on 9 different environments: the general randomization described in Sec. II-C.1 and the fine-tuned models described in Sec. II-C.2, which include 3-inch and 5-inch drones with 0%, 10%, 20%, and 30% randomization levels. Training followed the same procedure as above. Performance was evaluated with 1000 rollouts on both the 3-inch and 5-inch drone environments (with identical initial conditions sampled from the ranges described in II-B). See Tab. III. As shown, the general model performs effectively across both quadcopter simulations, navigating gates with a crash rate of 5.4% in the 3-inch environment and 1.2% in the 5-inch environment. However, it underperforms slightly compared to the fine-tuned models for each specific drone. Interestingly, the fine-tuned models do not transfer between drone types, resulting in 100% crash




Fig. 4. Comparison of training performance with various input modifications to the 64,64,64 ReLU network.

rates due to non-overlapping parameter ranges, even with 30% randomization. Among the fine-tuned models, lower randomization levels correlate with higher mean episode rewards, though at the cost of a slight increase in crash rates.

C. Real flights

We now test the networks that performed well in simulation on their respective platforms. Specifically, fine-tuned models are evaluated on their designated platforms, while the general model is tested on both the 3-inch and 5-inch quadcopters. All flights begin from a hover state, 1 meter in front of gate 1, and are repeated three times (12 seconds each, matching the maximum episode length used in the RL framework). The first flight uses a fully charged battery, with subsequent flights performed as the battery depletes, leaving it near empty by the third run. Flight logs are subsampled at 100Hz to match the simulator's time step, enabling direct comparison through the episode reward metric. Additionally, each flight is replicated in the simulator using the same initial conditions. The trajectories of both real and simulated flights are shown in Fig. 5. It can be seen here that the general network successfully navigates the race track on both drones, albeit slightly slower. Performance metrics for all flights are presented in Tab. IV-B and IV-B, showing that as randomization decreases from 30% to 10%, performance improves, with the highest rewards at 10%. At 0% randomization, the drone no longer passes through the gates. Fig. 6 provides a box plot of the episode rewards, clearly illustrating how the reality gap widens as domain randomization decreases.

D. Time Optimality

The rewards in the framework are designed to learn time-optimal flight. It is useful to not only compare the achieved episode rewards to the optimal during training

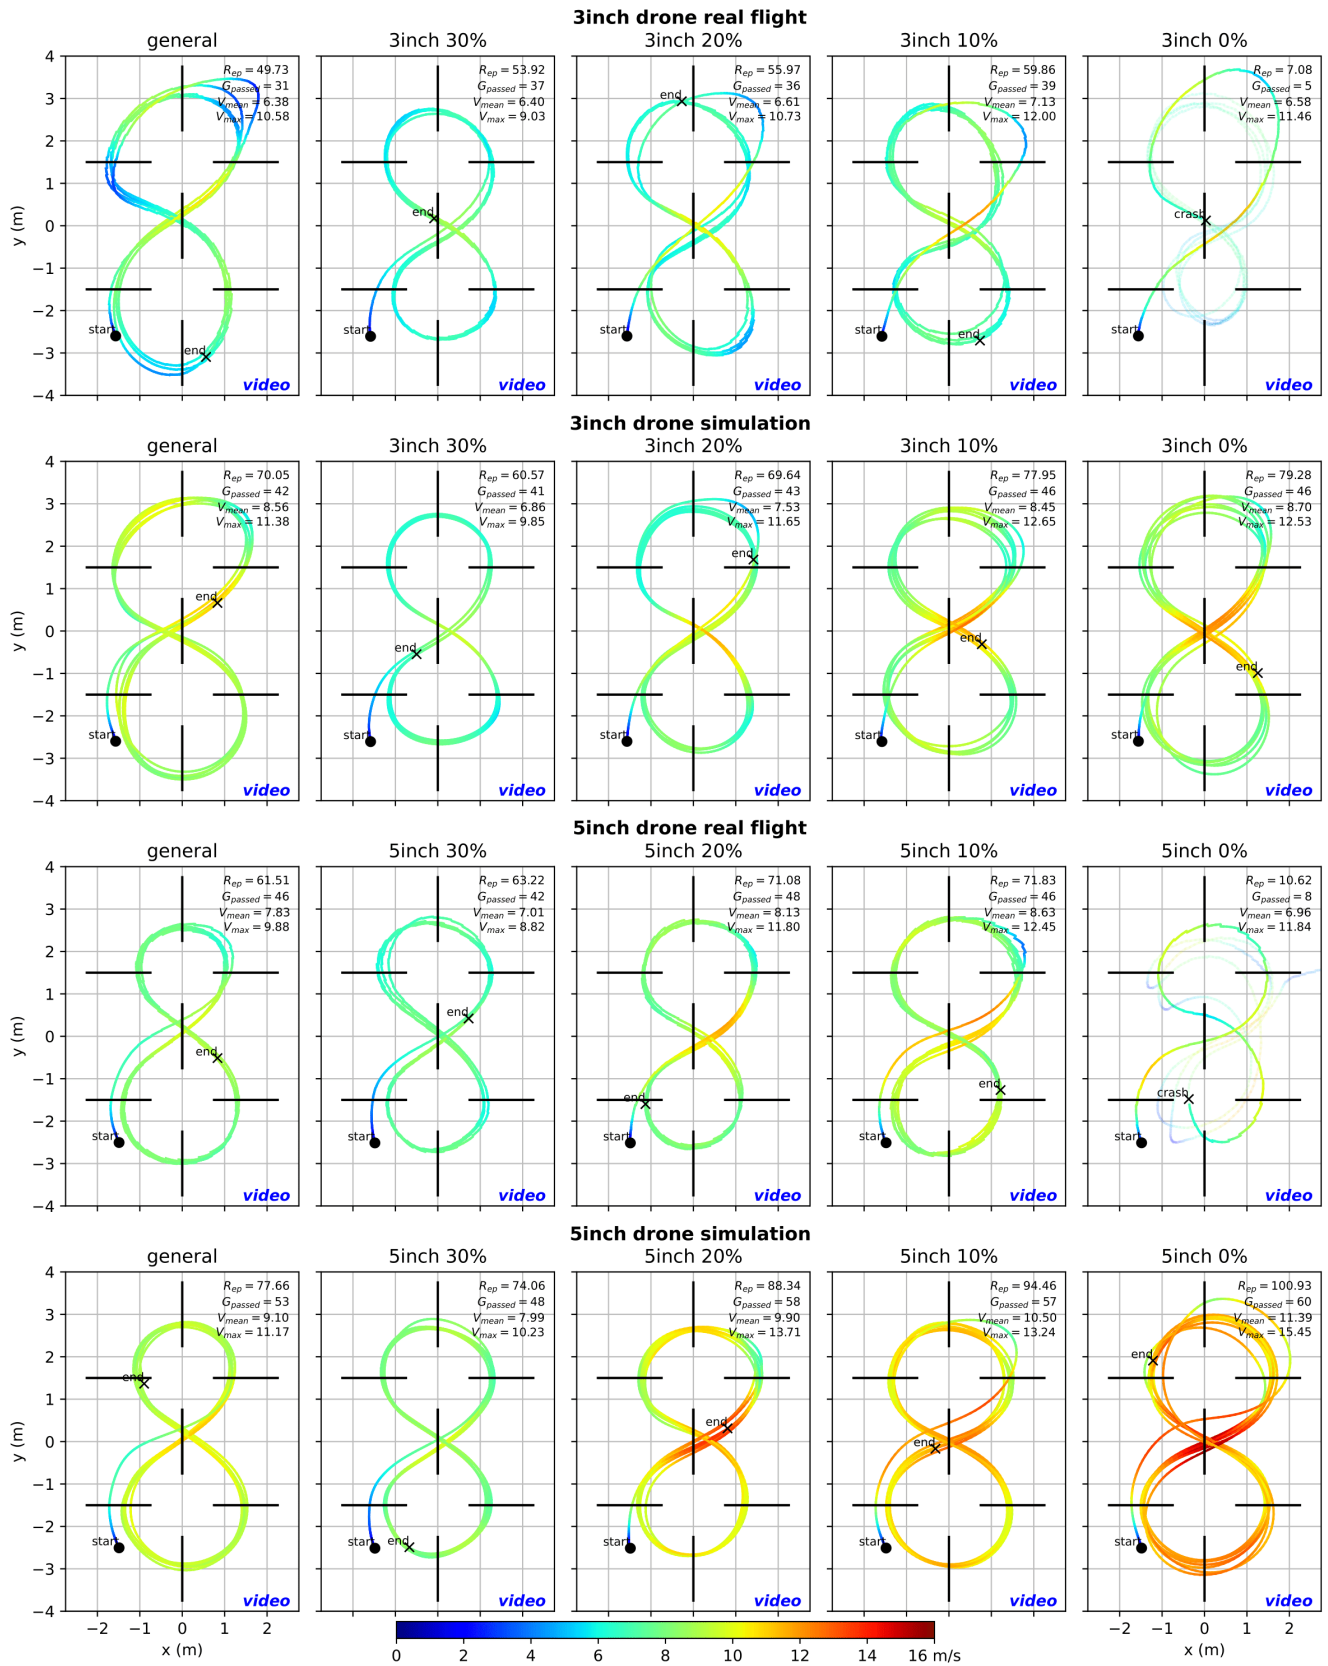


Fig. 5. Comparison of real-world and simulated flight trajectories for 3-inch and 5-inch quadcopters, using the general policy (for both platforms) and fine-tuned policies (trained on their respective parametric models) with 30%, 20%, 10%, and 0% randomization

TABLE IV
PERFORMANCE MEASURES: 3-INCH FLIGHTS

net	Real Flight				Simulation			
	ep rew	#gates	V_{mean}	V_{max}	ep rew	#gates	V_{mean}	V_{max}
Gen	49.73	31	6.38	10.58	69.76	42	8.59	10.99
	49.38	31	6.38	10.33	69.77	42	8.59	10.96
	48.78	31	6.31	10.4	69.97	42	8.6	10.94
30%	53.92	37	6.4	9.03	60.44	41	6.86	9.76
	53.48	37	6.34	8.73	60.61	41	6.86	9.68
	53.25	37	6.33	8.85	60.6	41	6.86	9.67
20%	55.97	36	6.61	10.73	69.67	43	7.56	11.58
	55.4	35	6.54	10.71	69.67	43	7.55	11.58
	55.32	35	6.49	10.48	69.72	43	7.56	11.58
10%	59.86	39	7.13	12	77.74	46	8.45	13.15
	59.99	39	7.13	11.56	78.17	46	8.46	13.24
	59.34	38	7.05	11.39	78.11	46	8.46	13.25
0%	7.08	5	6.58	11.46	80.82	46	8.79	12.8
	7.11	5	6.49	11.25	81.09	46	8.78	12.49
	10.18	7	6.51	10.87	80.93	46	8.77	12.38

TABLE V
PERFORMANCE MEASURES: 5-INCH FLIGHTS

net	Real Flight				Simulation			
	ep rew	#gates	V_{mean}	V_{max}	ep rew	#gates	V_{mean}	V_{max}
Gen	61.51	46	7.83	9.88	77.14	53	9.09	11.03
	61.22	46	7.84	9.83	77.5	53	9.12	11.37
	60.65	46	7.8	9.84	77.74	53	9.14	11.26
30%	63.22	42	7.01	8.82	74.38	48	8.01	10.22
	62.87	42	7.01	8.73	74.1	48	8.01	10.15
	62.59	41	6.98	8.7	73.7	48	8	10.1
20%	71.08	48	8.13	11.8	89.44	58	9.99	13.8
	71.44	49	8.16	11.74	89.56	58	9.99	13.9
	71.3	48	8.13	11.52	89.84	58	10	14.09
10%	71.83	46	8.63	12.45	95.68	58	10.53	12.81
	72.67	47	8.73	12.18	94.93	58	10.52	12.81
	71.98	46	8.63	12.07	95.67	58	10.54	13.05
0%	10.62	8	6.96	11.84	99.9	60	11.34	15.4
	17.16	13	7.46	13.13	100.55	59	11.34	15.53
	11.34	8	7.43	12.39	97.48	59	11.27	15.37

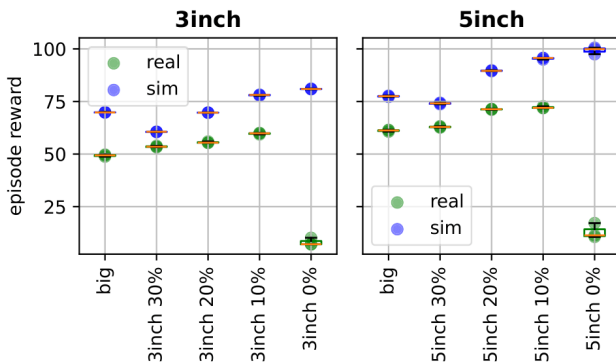


Fig. 6. Domain randomization enables the same network to fly different quadcopters; however with performance penalties. No domain randomization results in failed sim2real transfer: our DR approach stands as an alternative to higher modeling effort, at a small performance cost.




Fig. 7. The optimal control solution through the race track for the 5inch quadcopter

but also compare actual lap times to time-optimal flight as a means of validation and context. The optimal figure 8 path (see Fig.7) was formulated as a boundary value problem and solved with [25] for the 5-inch model, yielding a theoretical lap-time of 1.07s, compared to 1.6s for our fastest RL policies in simulation, and 1.9s in real-life. Even though the time-optimal program did not include

drag terms, we attribute the bulk of the difference to the sub-optimal path taken by the RL policies; relatively far away from the edges of the gates. However, the optimal control approach returns 1.53s when forcing passing the gate-centers. This suggests that our reward function does not perfectly align with optimizing lap times. As shown in Tab. IV-B, the 10% fine-tuned network achieves the highest real-life performance in terms of episode reward, yet it flies through fewer gates than the 20% variant. We therefore conclude that, to achieve more time-optimal performance on a tight track, the reward function should include terms that allow better tuning of the risk-reward balance, encouraging the agent to fly closer to the gate edges.

V. CONCLUSION

In this work, we introduced a neural network controller for high-speed drone racing capable of operating across multiple quadcopter platforms. We demonstrated its effectiveness with real flights on both 3-inch and 5-inch quadcopters at speeds up to 10 m/s. Our benchmarking involved comparing this generalized controller against neural policies specifically trained for the 3-inch and 5-inch quadcopters with varying levels of domain randomization. While our generalized model showed slightly reduced performance compared to the fine-tuned controllers, it exhibited superior transferability across platforms. Our experiments revealed that increased domain randomization initially enhances robustness and helps bridge the reality gap. However, excessive randomization ultimately compromises flight speed. We explored incorporating online adaptation into our network, as discussed in [17] and [19], but did not achieve the desired improvements. Potential factors include training duration, hyperparameter tuning, or the choice of parameter representations. Nevertheless, the significant success of our generalized approach, even without additional inputs, highlights its promise and paves the way for future research. Addressing the identified challenges could lead to substantial performance enhancements, making this an exciting avenue for further exploration.

REFERENCES

- [1] M. Hassanalian and A. Abdelkefi, "Classifications, applications, and design challenges of drones: A review," *Progress in Aerospace Sciences*, vol. 91, pp. 99–131, 2017.
- [2] D. Hanover, A. Loquercio, L. Bauersfeld, A. Romero, R. Penicka, Y. Song, G. Cioffi, E. Kaufmann, and D. Scaramuzza, "Autonomous Drone Racing: A Survey," *IEEE Transactions on Robotics*, vol. 40, pp. 3044–3067, Jan. 2024.
- [3] Y. Song, M. Steinweg, E. Kaufmann, and D. Scaramuzza, "Autonomous drone racing with deep reinforcement learning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 1205–1212.
- [4] R. Penicka, Y. Song, E. Kaufmann, and D. Scaramuzza, "Learning minimum-time flight in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7209–7216, 2022.
- [5] Y. Song, A. Romero, M. Müller, V. Koltun, and D. Scaramuzza, "Reaching the limit in autonomous racing: Optimal control versus reinforcement learning," *Science Robotics*, vol. 8, no. 82, p. eadg1462, 2023. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adg1462>
- [6] 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, Aug. 2023. [Online]. Available: <https://doi.org/10.1038/s41586-023-06419-4>
- [7] J. Eschmann, D. Albani, and G. Loianno, "Learning to Fly in Seconds," Nov. 2023, arXiv:2311.13081.
- [8] J. Josifovski, M. Malmir, N. Klarmann, B. L. Žagar, N. Navarro-Guerrero, and A. Knoll, "Analysis of randomization effects on sim2real transfer in reinforcement learning for robotic manipulation tasks," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 10 193–10 200.
- [9] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [10] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep drone racing: From simulation to reality with domain randomization," *IEEE Transactions on Robotics*, vol. 36, no. 1, pp. 1–14, 2019.
- [11] 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, 2017, pp. 23–30.
- [12] R. Ferede, G. de Croon, C. De Wagter, and D. Izzo, "End-to-end neural network based optimal quadcopter control," *Robotics and Autonomous Systems*, vol. 172, p. 104588, 2024.
- [13] S. Origer, C. De Wagter, R. Ferede, G. C. de Croon, and D. Izzo, "Guidance & control networks for time-optimal quadcopter flight," *arXiv preprint arXiv:2305.02705*, 2023.
- [14] R. Ferede, C. De Wagter, D. Izzo, and G. C. De Croon, "End-to-end reinforcement learning for time-optimal quadcopter flight," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 6172–6177.
- [15] G. Tiboni, P. Klink, J. Peters, T. Tommasi, C. D'Eramo, and G. Chaltatzaki, "Domain randomization via entropy maximization," *arXiv preprint arXiv:2311.01885*, 2023.
- [16] A. Molchanov, T. Chen, W. Hönig, J. A. Preiss, N. Ayanian, and G. S. Sukhatme, "Sim-to-(multi)-real: Transfer of low-level robust control policies to multiple quadrotors," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 59–66.
- [17] T. Blaha, E. Smeur, and B. Remes, "Control of Unknown Quadrotors from a Single Throw," *arXiv*, vol. arXiv:2406.11723, 2024.
- [18] E. J. J. Smeur, Q. Chu, and G. C. H. E. de Croon, "Adaptive Incremental Nonlinear Dynamic Inversion for Attitude Control of Micro Air Vehicles," *Journal of Guidance, Control, and Dynamics*, vol. 39, no. 3, pp. 450–461, Mar. 2016.
- [19] D. Zhang, A. Loquercio, X. Wu, A. Kumar, J. Malik, and M. W. Mueller, "Learning a Single Near-hover Position Controller for Vastly Different Quadcopters," May 2023.
- [20] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [21] 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>
- [22] I. Geles, L. Bauersfeld, A. Romero, J. Xing, and D. Scaramuzza, "Demonstrating agile flight from pixels without state estimation," *arXiv preprint arXiv:2406.12505*, 2024.
- [23] RTOM Corporation, "Moongel damper pads," <https://rtom.com/moongel-damper-pad/>, accessed: August 20, 2024.
- [24] S. F. ul Haq Gilani, M. H. bin Mohd Khir, R. Ibrahim, E. ul Hassan Kirmani, and S. I. ul Haq Gilani, "Modelling and development of a vibration-based electromagnetic energy harvester for industrial centrifugal pump application," *Microelectronics Journal*, vol. 66, pp. 103–111, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0026269216305948>
- [25] J.-B. Caillaud, O. Cots, J. Gergaud, P. Martinon, and S. Sed, "OptimalControl.jl: a Julia package to model and solve optimal control problems with ODE's." [Online]. Available: <https://control-toolbox.org/OptimalControl.jl>