

Self-supervised Domain Adaptation for Visual 3D Pose Estimation of Nano-drone Racing Gates by Enforcing Geometric Consistency

Nicholas Carlotti^{1*}, Michele Antonazzi^{2*}, Elia Cereda¹, Mirko Nava¹,
 Nicola Basilico², Daniele Palossi¹³, and Alessandro Giusti¹

Abstract—We consider the task of visually estimating the relative pose of a drone racing gate in front of a nano-quadrotor, using a convolutional neural network pre-trained on simulated data to regress the gate’s pose. Due to the sim-to-real gap, the pre-trained model underperforms in the real world and must be adapted to the target domain. We propose an unsupervised domain adaptation (UDA) approach using only real image sequences collected by the drone flying an arbitrary trajectory in front of a gate; sequences are annotated in a self-supervised fashion with the drone’s odometry as measured by its onboard sensors. On this dataset, a state consistency loss enforces that two images acquired at different times yield pose predictions that are consistent with the drone’s odometry. Results indicate that our approach outperforms other SoA UDA approaches, has a low mean absolute error in position ($x=26$, $y=28$, $z=10$ cm) and orientation ($\psi=13^\circ$), an improvement of 40% in position and 37% in orientation over a baseline. The approach’s effectiveness is appreciable with as few as 10 minutes of real-world flight data and yields models with an inference time of 30.4ms (33 fps) when deployed aboard the Crazyflie 2.1 Brushless nano-drone.

I. INTRODUCTION

Training deep learning models for perception tasks requires large datasets of annotated data. Obtaining real world data at scale is difficult as the collection process is often expensive and time-consuming. Conversely, simulators allow for easy collection of large amounts of training data; however, differences with the real world in appearance and dynamics pose a challenge, the so-called sim-to-real gap [1]. To bridge the gap between simulated and real-world data, models are adapted with domain transfer approaches [2], [3], with *domain* indicating a specific data distribution, e.g., simulated and real-world environments.

We focus on a drone-racing scenario, in which competing drones have to pass through gates with a characteristic appearance in a precise order to win a race (see Figure 1). In this scenario, we specifically consider the task of monocular visual pose estimation of drone-racing gates in a nano-drone [4], starting from a model pre-trained in simulation, and adapting it to the real world. Given a single image

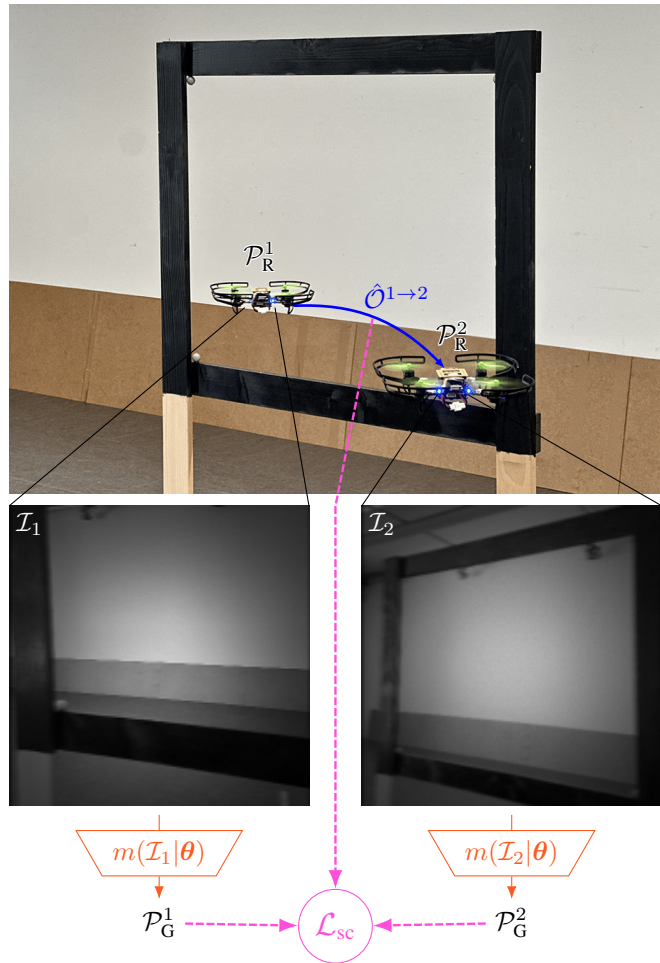


Fig. 1. We estimate the pose (x , y , z , ψ) of drone racing gates by transferring a model trained in simulation to a real-world environment and deploying it aboard the Bitcraze Crazyflie 2.1 Brushless nano-UAV.

taken by a drone’s onboard camera, our model estimates the 3D pose of a gate relative to the drone. We adopt the Bitcraze Crazyflie 2.1 Brushless¹ nano-drone platform, which mounts a grayscale, low-resolution camera with a limited dynamic range: Images are noisy, affected by motion blur and vignetting, and, crucially, significantly differ from those acquired in simulation, which lack these artifacts. A naive adaptation solution is fine-tuning the model on target domain data labeled with ground truth gate poses [5]; however, this approach requires a motion capture system to

* Indicates equal contribution.

¹Nicholas Carlotti, Elia Cereda, Mirko Nava, Daniele Palossi, and Alessandro Giusti are with the Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Lugano, Switzerland nicholas.carlotti@supsi.ch

²Michele Antonazzi and Nicola Basilico are with University of Milan, Italy [.<surname>@unimi.it](mailto:<surname>@unimi.it)

³Daniele Palossi is also with the Integrated Systems Laboratory (IIS), ETHZ, Zürich, Switzerland.

This work is supported by the Swiss National Science Foundation, grant number 213074.

¹<https://www.bitcraze.io/products/crazyflie-2-1-brushless>

collect pose data. Other solutions consist in zero-shot transfer of the task [4], [6], but are ineffective when the domain shift is too large. Self-supervised approaches from the computer vision literature [7], [8] represent a valid alternative, yet they fail to leverage the full potential of robotic platforms and their sensory apparatus. An unexplored approach for drone racing is robot self-supervision [9], [10], [11], where training data is collected and autonomously labeled by the robots themselves without the need for human supervision.

We propose an approach for sim-to-real domain transfer based on fine-tuning on autonomously-collected target domain data: A nano-drone flying random trajectories in front of a gate acquires images and odometry with its onboard sensors. We derive a supervision signal from the odometry data by enforcing a geometric state consistency loss [12] between pairs of poses predicted by the model: Given two images of the gate taken from different locations, our loss imposes that the two predicted relative gate poses must be compatible with the drone’s measured movement between the two locations (see Figure 1). Compared to general domain adaptation approaches, our work is designed for robotic platforms and uses the robot’s sensors and odometry estimation as the basis of the approach. This renders the approach effective at domain transfer of vision tasks, such as estimating the pose of a drone racing gate. Moreover, the real-world data required by our approach is cheaply obtained using a single drone equipped with a camera and the onboard odometry estimation.

After discussing related works in domain adaptation and self-supervised robot learning in Section II, we present our main **contribution** in Section III: a novel application of the state consistency loss for the domain transfer of a 3D gate pose estimation task; the loss is combined with a self-supervised data collection pipeline making the approach practical and easily applicable in any target domain. In Section IV, we describe the robotic platform, simulated and real-world domains, data collection pipeline, and evaluation metrics. Results, illustrated in Section V, show that our approach effectively transfers perception capabilities from simulation to the real-world, even with as few as 10 minutes of real-world flight data. Specifically, our method achieves a position error of 10 cm and an orientation error of 13° , improving over SoA UDA baselines by 40% and 37%, respectively. Furthermore, we assess the effectiveness of our approach on a small neural network tailored for running at 33 FPS onboard a nano-drone platform. Conclusions are drawn in Section VI along with future work.

II. RELATED WORK

A. Gate Detection for Drone Racing

Accurate gate detection is fundamental to enable fast navigation in drone racing competitions [4], [13]. The straightforward approach to train neural networks for gate detection is to directly use data collected with a motion capture system [5]. However, these systems are expensive and unsuitable for use across multiple environments. As such, alternative approaches learn from simulated data while

mitigating the sim-to-real gap with strong image augmentations [4], or by randomizing the visual aspect of many simulated environments [14]. Another way to align the visual appearance of simulated and real images is to apply a pencil filter that highlights edges and disregards flat image regions [6]. Instead, we propose an approach for sim-to-real transfer using images from the target environment, which utilizes the drone’s odometry as supervision for the alignment process.

B. Domain Generalization

In domain generalization, training data is collected in multiple source domains to cover a large portion of the input space, an effective strategy to generalize models to unseen domains. A prominent approach is to randomize the domain using simulators by altering the number and location of objects, modifying illumination conditions, and randomizing textures to match the complexity of real-world environments [15], [14]. Other approaches leverage data generation to synthesize randomized, realistic training data [16], [17] or alter the input space to render the two domains indistinguishable with autoencoders [18], or a generative adversarial network [19]. A simpler approach for making domains indistinguishable is to preprocess input images with a pencil filter to emphasize edges and remove distracting domain details [6].

Instead of generalizing to any domain, our approach acquires fine-tuning data in the target domain directly and adapts to it by optimizing the state consistency loss.

C. Domain Adaptation

When data can be collected in the target domain, *domain adaptation* approaches fine-tune a model to capture the specific characteristics of the domain; the simple solution is to use high-quality labeled data, often involving laborious hand-labeling [20], [21] or ad-hoc sensing apparatus [5].

Performing adaptation in *unsupervised* fashion, i.e. without ground-truth annotations for the target domain, is an attractive alternative [3]. Two common solutions are to reduce the domain shift or learn a latent space in the target domain with self-supervised learning. Approaches reduce the domain shift at the input level, rendering the source and target input data indistinguishable by performing style transfer between domains [22] and further aligning image features in a latent space [23], [24]. Reducing the domain shift at feature level is done with adversarial learning, by confusing a discriminator network [25], [26], or by minimizing a notion of domain divergence [27], [28], such as the Maximum Mean Discrepancy (MMD) [29]. Self-supervised domain adaptation without target domain labels is done with teacher-student training of a model pre-trained on source data, to minimize the model entropy on the target domain [30]. Approaches update self-labels at each training iteration to mitigate noisy predictions [31] or update only labels for high-uncertainty samples [32], [33], weight the loss by the model uncertainty under different input augmentations [34], or minimize a triplet loss [35].




Fig. 2. Assuming perfect odometry (left), the state consistency loss forces the relative gate poses $\hat{\mathcal{P}}_G^1$ and $\hat{\mathcal{P}}_G^2$, predicted from the drone’s poses at \mathcal{P}_R^1 and \mathcal{P}_R^2 , to be coherent with the drone’s relative odometry $\mathcal{O}^{1 \rightarrow 2}$ between the two poses. In our experiments, the measured odometry $\hat{\mathcal{O}}^{1 \rightarrow 2}$ is collected by the drone itself and is affected by drift and noise (right).

Compared to other unsupervised approaches, we do not use source domain data for model fine-tuning; instead, we use unlabeled target domain data to optimize a consistency loss specifically formulated for odometry-capable robots.

D. Self-supervised Robot Learning

To lessen the need for labeled training data, self-supervised approaches use trivial learning objectives, whose labels are generated by transforming the input data, leveraging the many sensors equipped on a robot platform. Typical examples include the use of proximity sensors and odometry to self-supervise the training of an obstacle detection model [9], or dense point clouds for self-supervised traversability estimation [11]. Recent self-supervised approaches leverage a pretext task, whose objective is not to directly solve the task of interest but, rather, to learn similar pattern recognition skills as those needed to solve the task of interest: approaches estimate, given an input image, the sound produced by a drone [36], or the state of the LEDs fitted on a robot [37]; others solve masked autoencoding tasks on images [38], or use two sensory streams such as vision and language [39], or vision and haptics [40].

Our approach leverages data collected autonomously by the drone to directly supervise a deep learning model on the transfer from simulation to reality of a drone racing gate detector.

III. METHODOLOGY

We solve the problem of estimating the 6D pose of a drone racing gate relative to an observing drone $\mathcal{P}_G \in SE(3)$, given a single grayscale image \mathcal{I} . The model $m(\mathcal{I}|\theta)$ parametrized by weights θ takes a single image \mathcal{I} as input and produces the estimated gate pose $\hat{\mathcal{P}}_G$. Initially, the model is trained on data coming from a simulator, allowing us to easily collect images labeled with the pose of the simulated gate \mathcal{P}_G . In a second phase, we transfer the model trained in simulation to the real world by fine-tuning on data collected by a real drone



Fig. 3. The Crazyflie 2.1 Brushless nano-drone used in our experiments.

flying in front of a gate, annotated with the drone’s odometry while having no access to information about the gate and its pose. The drone odometry collected at time $t = 1$ w.r.t. the world reference frame is denoted by $\hat{\mathcal{O}}^{1 \rightarrow w}$. The resulting simulated dataset is denoted with \mathcal{D}_{sim} and the real-world one with \mathcal{D}_{real} . The training procedure consists of a first phase where the model is trained on $\mathcal{D}_{sim}^{train}$, while in the second phase, we start from the model parameters minimizing \mathcal{L}_{pose} on the simulated validation set \mathcal{D}_{sim}^{val} and fine-tune them with $\mathcal{D}_{real}^{train}$ using only the \mathcal{L}_{sc} loss.

When training on simulated data, we minimize the mean squared error (MSE) loss between \mathcal{P}_G and $\hat{\mathcal{P}}_G$:

$$\mathcal{L}_{pose} = \text{MSE}(\mathcal{P}_G, \hat{\mathcal{P}}_G) \quad (1)$$

where the position components of the pose are directly passed to the MSE loss, while the orientation components are first converted to the 6D representation proposed in Zhou et al. [41], preserving the continuity of 3D rotations to aid the training process.

For real-world data, we enforce a State Consistency (SC) loss [12] between two model predictions, exploiting the known drone odometry: Given a pair of images taken by the drone in a static scene from two different locations, the difference between the estimated relative gate’s poses as perceived by the drone in the two locations must be equal to the drone’s movement between the first and second locations (see Fig. 2 for a general overview). Formally, given two input images $\mathcal{I}_1, \mathcal{I}_2$, the gate poses $\mathcal{P}_G^1, \mathcal{P}_G^2$, and the drone odometries $\hat{\mathcal{O}}^{1 \rightarrow w}, \hat{\mathcal{O}}^{2 \rightarrow w}$, we first make the model predict $\hat{\mathcal{P}}_G^1, \hat{\mathcal{P}}_G^2$. Then, we warp $\hat{\mathcal{P}}_G^1$ to the frame of reference of the drone at the position in which \mathcal{I}_2 was taken $\hat{\mathcal{P}}_G^{1 \rightarrow 2} = \hat{\mathcal{P}}_G^1 \cdot \hat{\mathcal{O}}^{2 \rightarrow w} \cdot (\hat{\mathcal{O}}^{1 \rightarrow w})^{-1}$, where \cdot denotes matrix multiplication and \mathcal{M}^{-1} denotes the inverse of matrix \mathcal{M} . The SC loss is then written as:

$$\mathcal{L}_{sc} = \text{MSE}(\hat{\mathcal{P}}_G^{1 \rightarrow 2}, \hat{\mathcal{P}}_G^2) \quad (2)$$

where orientation components of the poses are converted to the 6D representation proposed in Zhou et al. [41].

IV. EXPERIMENTAL EVALUATION

A. Robot Platform

We target the Crazyflie 2.1 Brushless by Bitcraze, a commercial off-the-shelf nano-drone extended with the pluggable GAP9Shield [42] and the Flow-deck expansion modules,



Fig. 4. Examples of (left) simulated images from $\mathcal{D}_{\text{sim}}^{\text{train}}$, (middle) real-world robot images in $\mathcal{D}_{\text{real}}^{\text{train}}$, and (right) after applying the pencil filter [6].

depicted in Figure 3. The GAP9Shield features a GreenWaves Technologies GAP9 ultra-low power System-on-Chip (SoC), composed of an eight-core RISC-V cluster (CL), an 150 MAC/cycle 8-bit hardware accelerator for neural network inference (NE16), a cluster controller (CC) core in charge of orchestrating parallel computation, and a fabric controller core (FC) responsible for data movement with external peripherals. GAP9 boasts a combined throughput up to 220 GOP/s at max frequency 370 MHz, while consuming less than 100 mW. The SoC integrates 1.5 MB L2 memory and 128 kB L1 scratchpad, complemented on the GAP9Shield by off-chip 32 MB OctaSPI RAM and 64 MB Flash memories. A 5MP OmniVision OV5647 camera provides visual perception, and the Espressif ESP32 module offers Wi-Fi connectivity. The Crazyflie 2.1 base platform employs an STM32 microcontroller to handle low-level flight-control tasks and sensor interfacing: a 6-DoF IMU (three-axis accelerometer and gyroscope), an STMicroelectronics VL53L1X laser time-of-flight altitude sensor, and a PMW3901 optical-flow sensor to track horizontal displacement. The sensor data is fused by an Extended Kalman Filter [43] to provide the onboard odometry.

B. Data Collection

We collect training data in simulated and real-world environments. In both domains, our $100\text{cm} \times 80\text{cm}$ gates are composed of four black-painted wooden beams placed at 75 cm height, based on the setup from the first Nanocopter AI challenge at IMAV 2022 [4]. For the simulated data, we use the Webots simulator [44] to generate a dataset of 75K images depicting a gate in random poses. To align the visual aspect of the simulated data in $\mathcal{D}_{\text{sim}}^{\text{train}}$ with the images captured in the real world by the drone, we convert the images into grayscale and we perform data augmentations during training using Gaussian blur, vignetting, multiplicative Gaussian noise, and random exposure (see Figure 4).

To collect real-world images, we use the Crazyflie 2.1 Brushless piloted by a human to follow random trajectories around the gate while keeping it in the camera’s FOV. During flight, we collect images at 25 FPS with a resolution of 160×160 pixels. For testing purposes only, we collect ground truth pose data for both the drone and the gate with a motion capture system installed in our laboratory. Real data is split into the training set $\mathcal{D}_{\text{real}}^{\text{train}}$ (51K samples), the validation set $\mathcal{D}_{\text{real}}^{\text{val}}$ (8K samples), and the testing set $\mathcal{D}_{\text{real}}^{\text{test}}$ (21K samples). Examples of real-world images acquired by our drone are depicted in Figure 4 (middle).

C. Model Training and Onboard Deployment

Our model architecture is a convolutional neural network with four convolutional blocks with 32, 32, 64, 128 channels, followed by batch normalization, and a linear layer with 128 neurons, totaling 6M parameters. For simulation training, we optimize the $\mathcal{L}_{\text{pose}}$ loss using AdamW [45] as the optimizer with a learning rate of $1e^{-3}$ for 100 epochs. Then, we pick the model parameters with the lowest loss on $\mathcal{D}_{\text{sim}}^{\text{val}}$ and fine-tune it on real-world data by optimizing the \mathcal{L}_{sc} loss using AdamW with a learning rate of $1e^{-6}$ for 100 epochs. During fine-tuning, we use a small validation set $\mathcal{D}_{\text{real}}^{\text{val}}$, which does not contain pose labels, to pick the best model weights over the training process.

We deploy our model on the GWT GAP9 SoC to run entirely aboard our target nano-drone. For maximum inference throughput, we aim to take advantage of the chip’s NE16 neural accelerator, which only supports mixed-precision integer arithmetic. We employ the GWT NNTool + AutoTiler deployment pipeline to quantize the model to 8-bit integers, with no loss in regression performance, and generate the C inference code. This includes NE16 convolutional and fully-connected kernels, optimized parallel CPU kernels for the remaining layers, and automatically scheduled memory transfers across the chip’s memory hierarchy.

D. Evaluation Metrics

We evaluate the model’s ability to estimate the gate’s relative pose on the real-world testing set $\mathcal{D}_{\text{real}}^{\text{test}}$ by measuring the mean absolute error (MAE) and the Pearson correlation coefficient (ρ) between predicted and ground truth components of the 3D position and yaw of the gate; specifically for the yaw (indicated as ψ) angle, we consider the circular Pearson correlation coefficient [46]. Due to the drone’s flight dynamics, roll and pitch angles of the gate relative to the drone have a low variability in our data. As such, metrics on these variables would be uninformative; we choose not to present them in our evaluation.

Our model predictions correctly capture the pose of the gate but are inherently affected by a constant bias stemming from the nature of the state consistency loss: The loss is minimized as long as the predicted pose pairs are spatially consistent with the drone’s movements, yet they are not forced to be anchored at the gate’s center. This bias can be corrected with a calibration procedure by picking a small batch of images where the position of the gate relative to the camera is known. The model is run on these images,

and the mean offset between the gate’s position and its estimate is measured. To correct the bias, a constant offset is added to predicted positions. In presenting our results, we apply this calibration procedure by considering the whole testing dataset. This correction only affects the MAE metric for x , y , and z , since the linear correlation coefficient is not affected by a constant bias. For fairness, we apply the same bias correction procedure to all the considered baselines.

E. Baselines

We validate our approach considering different baselines. The first one (*Zero-Shot*) is our neural network trained only with the simulated dataset $\mathcal{D}_{\text{sim}}^{\text{train}}$ and tested directly on $\mathcal{D}_{\text{real}}^{\text{test}}$. The second baseline we consider is *PencilNet* [6], a state-of-the-art domain generalization approach for gate pose estimation. It consists of pre-processing the model’s input images using a pencil filter to align the visual aspect of simulated and real-world data, as shown on the right of Figure 4. In this way, the model is trained in simulation and deployed to real scenarios in a zero-shot fashion. Our third baseline (*DA*) is the well-established method for unsupervised domain adaptation described in [28]. This approach uses a two-stream architecture to align the distributions of the source and target domains with Maximum Mean Discrepancy (MMD). Furthermore, an additional loss function forces the parameters of the two streams to be linearly dependent. Similarly to our approach, no ground truth annotations in the real world are required. To ensure a fair evaluation, we re-implement all the baselines using our model architecture (described in Section IV-C) and our training dataset acquired in simulation (see Section IV-B) and report the performance figures obtained by averaging results over three training runs with different parameter initialization. We also report the performance of the *Mean Predictor*, which always predicts the average gate pose in the test set $\mathcal{D}_{\text{real}}^{\text{test}}$.

V. RESULTS

We report the results for our approach and all baselines in Table I. We observe that our method consistently outperforms all the baselines: The *Zero-Shot* baseline, namely the model trained in simulation and tested in the real world, brings only marginal improvements with respect to the lower bound represented by the *Mean Predictor*. This is observed also with *PencilNet* that, while reaching performance comparable to the *Mean Predictor*, remarkably degrades the estimation of the gate’s distance (x) in terms of MAE and ρ . This suggests that the domain generalization approach fails to overcome a large sim-to-real gap, such as the one found in our setup. The results also show that our domain adaptation (*DA*) baseline from [28] is not able to improve the performance of the *Mean Predictor*. Despite being beneficial for y , it degrades the estimation of both x and z , demonstrating that performing distribution alignment is not sufficient to fill the domain gap of our scenario. Compared with *Mean Predictor*, our approach improves the MAE by 54% in x and z (distance and height of the gate) and by 42% in y (horizontal displacement). Moreover, our approach is the




Fig. 5. Our self-supervised domain adaptation approach vs ground truth on $\mathcal{D}_{\text{real}}^{\text{test}}$ for the components of the gate pose (x, y, z, ψ).

only one that manages to reduce the error in the estimation of ψ , with a gain of 37%. Notably, the geometric relation between predictions imposed by the state consistency loss strongly improves the correlation ρ of the model outputs with their respective ground truths (see Figure 5).

Figure 6 reports qualitative examples of our model on gate pose estimation. In the first three images, the model precisely localizes the gate, even when it is far away from the drone. In the last two examples, challenging situations where the gate partially exits the image boundaries or is perceived from a narrow point of view cause the model to fail.

Following the pipeline described in Section IV-C, we deploy our network in the Crazyflie 2.1 quadrotor. Counting 6M parameters, the model requires 318 MMAC per inference and 30.4ms for processing a single frame on the NE16. With an inference throughput of 32.8 frames/s, our architecture is fully suitable for closed-loop autonomous control aboard the drone during real-world flights.

A. Fine-tuning Data Ablation

As a final experiment, we analyze the performance of our approach with varying amounts of target domain data: 1, 10, 20, 40, and 60 recording sequences, each one lasting 30 seconds and amounting to 850 samples. Results reported in Figure 7, show that training our approach with 10 sequences is enough to surpass all of the baselines’ performance. By training with 40 sequences, we closely match the performance of the full dataset, showing diminishing returns, especially in estimating the gate’s position. These results prove that the approach is effective, and even small amounts of target domain data suffice to outperform state-of-the-art baselines.

VI. CONCLUSIONS

This work presents a self-supervised approach for gate pose estimation specifically tailored to mobile robots. Using

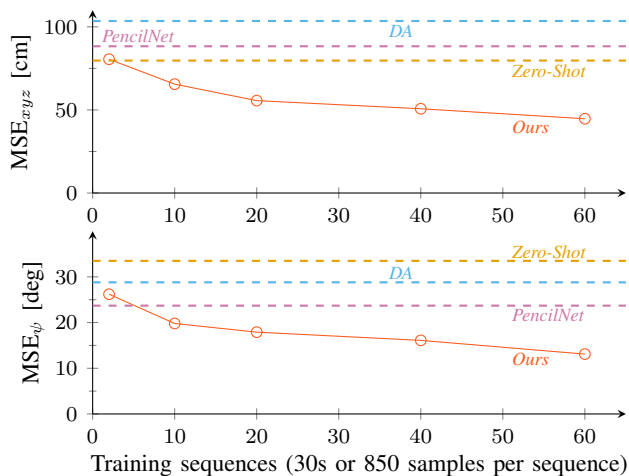
TABLE I

MEAN ABSOLUTE ERROR MAE AND PEARSON CORRELATION COEFFICIENT ρ ON THE REAL-WORLD TESTING SET $\mathcal{D}_{\text{real}}^{\text{test}}$, 3 RUNS PER ROW.

Model	Supervision	MSE \downarrow xyz [cm]	MAE \downarrow			ρ \uparrow				Plot for MSE $_{xyz}$ [cm] \leftarrow Error bars mark 95% CI	
			x [cm]	y [cm]	z [cm]	ψ [deg]	x [%]	y [%]	z [%]		ψ [%]
<i>Mean Predictor</i>	$\mathcal{D}_{\text{real}}^{\text{test}}$	90.7	57.0	48.7	23.1	20.9	—	—	—	—	●
<i>Zero-Shot</i>	$\mathcal{D}_{\text{sim}}^{\text{train}}$	79.7	51.6	41.0	24.5	33.5	38	53	40	5	●
<i>PencilNet [6]</i>	$\mathcal{D}_{\text{sim}}^{\text{train}}$	88.3	56.7	46.8	22.5	32.7	4	40	44	2	●
<i>DA [28]</i>	$\mathcal{D}_{\text{sim}}^{\text{train}}, \mathcal{D}_{\text{real}}^{\text{train}}$	103.5	66.0	34.7	33.0	28.8	30	69	52	5	●
<i>Ours</i>	$\mathcal{D}_{\text{real}}^{\text{train}}$ (Odometry only)	44.7	25.6	28.2	10.5	13.1	88	79	91	82	●



Fig. 6. Input frames overlaid with gate skeleton (in orange) using the pose predicted by our approach. The model correctly locates the gate (1-3), while failures occur when a large portion of the gate is occluded and for shallow angles (4-5).

Fig. 7. Position (top) and orientation (bottom) performance measure with MSE of our approach with increasing amounts of fine-tune data. DA is trained on the full real-world dataset (60 sequences), while PencilNet and Zero-Shot are trained only on $\mathcal{D}_{\text{sim}}^{\text{train}}$.

solely the estimated odometry as the supervision signal for adapting a source-domain model, our method forces the model’s predictions to be consistent with the drone’s movements with the state consistency loss. The extensive experimental evaluation in the real world proves that our method outperforms both zero-shot and unsupervised domain adaptation baselines. In addition, we demonstrate that our approach is effective, even when a limited amount of target domain data is available. Lastly, we assess the effectiveness of our approach on a lightweight neural network that runs at 33 fps onboard the Crazyflie 2.1 nano-drone. In future work, we plan to extend our approach with few-shot adaptation techniques to further reduce the amount of target domain

data, and introduce uncertainty estimation into the training setup, aiming to minimize the entropy in a similar fashion to source-free domain adaptation [30]. Furthermore, we plan to validate our approach in a closed-loop drone racing setup.

REFERENCES

- [1] E. Salvato, G. Fenu, E. Medvet, and F. A. Pellegrino, “Crossing the reality gap: A survey on sim-to-real transferability of robot controllers in reinforcement learning,” *IEEE Access*, vol. 9, pp. 153 171–153 187, 2021.
- [2] M. Schwonberg, J. Niemeijer, J.-A. Termöhlen, N. M. Schmidt, H. Gottschalk, T. Fingscheidt *et al.*, “Survey on unsupervised domain adaptation for semantic segmentation for visual perception in automated driving,” *IEEE Access*, vol. 11, pp. 54 296–54 336, 2023.
- [3] J. Li, Z. Yu, Z. Du, L. Zhu, and H. T. Shen, “A comprehensive survey on source-free domain adaptation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 8, pp. 5743–5762, 2024.
- [4] L. Lamberti, E. Cereda, G. Abbate, L. Bellone, V. J. K. Morinigo, M. Barciś, A. Barciś, A. Giusti, F. Conti, and D. Palossi, “A sim-to-real deep learning-based framework for autonomous nano-drone racing,” *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1899–1906, 2024.
- [5] L. Scarciglia, A. Paolillo, and D. Palossi, “A map-free deep learning-based framework for gate-to-gate monocular visual navigation aboard miniaturized aerial vehicles,” 2025.
- [6] H. X. Pham, A. Sarabakha, M. Odnoshynkin, and E. Kayacan, “Pencilnet: Zero-shot sim-to-real transfer learning for robust gate perception in autonomous drone racing,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 847–11 854, 2022.
- [7] P. Nguyen, T.-D. Truong, M. Huang, Y. Liang, N. Le, and K. Luu, “Self-supervised domain adaptation in crowd counting,” in *2022 IEEE International Conference on Image Processing (ICIP)*, 2022, pp. 2786–2790.
- [8] H.-Y. Yoon, J.-H. Kim, and J.-W. Jeong, “Classification of the sidewalk condition using self-supervised transfer learning for wheelchair safety driving,” *Sensors*, vol. 22, no. 1, 2022.
- [9] M. Nava, J. Guzzi, R. O. Chavez-Garcia, L. M. Gambardella, and A. Giusti, “Learning long-range perception using self-supervision from short-range sensors and odometry,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1279–1286, 2019.

- [10] A. J. Sathiamoorthy, K. Weerakoon, T. Guan, J. Liang, and D. Manocha, "Terrapin: Unstructured terrain navigation using online self-supervised learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 7197–7204.
- [11] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, "Scate: A scalable framework for self-supervised traversability estimation in unstructured environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 888–895, 2023.
- [12] M. Nava, L. M. Gambardella, and A. Giusti, "State-consistency loss for learning spatial perception tasks from partial labels," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1112–1119, 2021.
- [13] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, "Alphapilot: Autonomous drone racing," *Springer Autonomous Robots*, vol. 46, no. 1, pp. 307–320, 2022.
- [14] 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, 2020.
- [15] J. Tobin, L. Biewald, R. Duan, M. Andrychowicz, A. Handa, V. Kumar, B. McGrew, A. Ray, J. Schneider, P. Welinder, W. Zaremba, and P. Abbeel, "Domain randomization and generative models for robotic grasping," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3482–3489.
- [16] M. M. Rahman, C. Fookes, M. Baktashmotlagh, and S. Sridharan, "Multi-component image translation for deep domain generalization," in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019, pp. 579–588.
- [17] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [18] T. Inoue, S. Choudhury, G. De Magistris, and S. Dasgupta, "Transfer learning from synthetic to real images using variational autoencoders for precise position detection," in *IEEE International Conference on Image Processing (ICIP)*, 2018, pp. 2725–2729.
- [19] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 627–12 637.
- [20] M. Antonazzi, M. Luperto, N. A. Borghese, and N. Basilico, "Development and adaptation of robotic vision in the real world: The challenge of door detection," *Journal of Field Robotics*, 2025.
- [21] N. Zimmerman, T. Guadagnino, X. Chen, J. Behley, and C. Stachniss, "Long-term localization using semantic cues in floor plan maps," *IEEE Robotics and Automation Letters*, vol. 8, no. 1, pp. 176–183, 2023.
- [22] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE/CVF International Conference on Computer Vision*, Oct 2017.
- [23] J. Hoffman, E. Tzeng, T. Park, J.-Y. Zhu, P. Isola, K. Saenko, A. Efros, and T. Darrell, "CyCADA: Cycle-consistent adversarial domain adaptation," in *International Conference on Machine Learning*, vol. 80, 10–15 Jul 2018, pp. 1989–1998.
- [24] Y.-C. Chen, Y.-Y. Lin, M.-H. Yang, and J.-B. Huang, "Crdoco: Pixel-level domain transfer with cross-domain consistency," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.
- [25] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, July 2017.
- [26] Y. Zhang and Z. Wang, "Joint adversarial learning for domain adaptation in semantic segmentation," *AAAI conference on artificial intelligence*, vol. 34, pp. 6877–6884, Apr. 2020.
- [27] J. Wang, W. Feng, Y. Chen, H. Yu, M. Huang, and P. S. Yu, "Visual domain adaptation with manifold embedded distribution alignment," in *ACM International Conference on Multimedia*. Association for Computing Machinery, 2018, p. 402–410.
- [28] A. Rozantsev, M. Salzmann, and P. Fua, "Beyond sharing weights for deep domain adaptation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 4, pp. 801–814, 2019.
- [29] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola, "A kernel method for the two-sample-problem," *Advances in neural information processing systems*, vol. 19, 2006.
- [30] Y. Kim, D. Cho, K. Han, P. Panda, and S. Hong, "Domain adaptation without source data," *IEEE Transactions on Artificial Intelligence*, vol. 2, no. 6, pp. 508–518, 2021.
- [31] D. Kothandaraman, R. Chandra, and D. Manocha, "Ss-sfda: Self-supervised source-free domain adaptation for road segmentation in hazardous environments," in *IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3049–3059.
- [32] W. Chen, L. Lin, S. Yang, D. Xie, S. Pu, and Y. Zhuang, "Self-supervised noisy label learning for source-free unsupervised domain adaptation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 10 185–10 192.
- [33] T. Chu, Y. Liu, J. Deng, W. Li, and L. Duan, "Dennoised maximum classifier discrepancy for source-free unsupervised domain adaptation," in *AAAI conference on artificial intelligence*, vol. 36, no. 1, 2022, pp. 472–480.
- [34] H. Yan, Y. Guo, and C. Yang, "Augmented self-labeling for source-free unsupervised domain adaptation," in *Advances in Neural Information Processing Systems, Workshop*, 2021.
- [35] R. Wang, Z. Wu, Z. Weng, J. Chen, G.-J. Qi, and Y.-G. Jiang, "Cross-domain contrastive learning for unsupervised domain adaptation," *IEEE Transactions on Multimedia*, vol. 25, pp. 1665–1673, 2022.
- [36] M. Nava, A. Paolillo, J. Guzzi, L. M. Gambardella, and A. Giusti, "Learning visual localization of a quadrotor using its noise as self-supervision," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2218–2225, 2022.
- [37] N. Carlotti, M. Nava, and A. Giusti, "Learning to estimate the pose of a peer robot in a camera image by predicting the states of its leds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024, pp. 2763–2769.
- [38] I. Radosavovic, T. Xiao, S. James, P. Abbeel, J. Malik, and T. Darrell, "Real-world robot learning with masked visual pre-training," in *PMLR Conference on Robot Learning*, 2023, pp. 416–426.
- [39] S. Karamcheti, S. Nair, A. S. Chen, T. Kollar, C. Finn, D. Sadigh, and P. Liang, "Language-driven representation learning for robotics," in *Robotics: Science and Systems*, 2023.
- [40] Q. Liu, Q. Ye, Z. Sun, Y. Cui, G. Li, and J. Chen, "Masked visual-tactile pre-training for robot manipulation," in *IEEE International Conference on Robotics and Automation*, 2024, pp. 13 859–13 875.
- [41] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [42] H. Müller, V. Kartsch, and L. Benini, "GAP9Shield: A 150GOPS AI-capable ultra-low power module for vision and ranging applications on nano-drones," in *European Robotics Forum 2024*, C. Secchi and L. Marconi, Eds. Cham: Springer Nature Switzerland, 2024, pp. 292–297.
- [43] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 1730–1736.
- [44] Webots, "http://www.cyberbotics.com," open-source Mobile Robot Simulation Software. [Online]. Available: <http://www.cyberbotics.com>
- [45] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2017.
- [46] K. V. Mardia and P. E. Jupp, *Directional statistics*. John Wiley & Sons, 2009.