

PencilNet: Zero-Shot Sim-to-Real Transfer Learning for Robust Gate Perception in Autonomous Drone Racing

Huy Xuan Pham¹, Andriy Sarabakha², Mykola Odnoshykin³ and Erdal Kayacan¹

Abstract—In autonomous and mobile robotics, one of the main challenges is the robust on-the-fly perception of the environment, which is often unknown and dynamic, like in autonomous drone racing. In this work, we propose a novel deep neural network-based perception method for racing gate detection – PencilNet¹ – which relies on a lightweight neural network backbone on top of a pencil filter. This approach unifies predictions of the gates’ 2D position, distance, and orientation in a single pose tuple. We show that our method is effective for zero-shot sim-to-real transfer learning that does not need any real-world training samples. Moreover, our framework is highly robust to illumination changes commonly seen under rapid flight compared to state-of-art methods. A thorough set of experiments demonstrates the effectiveness of this approach in multiple challenging scenarios, where the drone completes various tracks under different lighting conditions.

Index Terms—Aerial Systems: Perception and Autonomy; Aerial Systems: Applications; Aerial Systems: Mechanics and Control

I. INTRODUCTION

VISION-BASED agile navigation for robotics is an emerging research field that has been gaining traction in recent years [1]–[6]. As technical difficulties hinder the extension of batteries’ capacity, there is a strong interest in increasing the speed of robots to expand both the operating range and systems’ capabilities [7]. This is even more crucial for aerial systems, like multicopters [8], that are capable of reaching high speeds in a short time due to their agile nature [9].

Manuscript received: February, 25, 2022; Revised June, 19, 2022; Accepted July, 26, 2022.

This paper was recommended for publication by Editor Pauline Pounds upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported in part by Aarhus University, Dept. of Electrical and Computer Eng. (28173) and in part by the European Union’s Horizon 2020 Research and Innovation Program (OpenDR) under Grant 871449. This publication reflects the authors’ views only. The European Commission is not responsible for any use that may be made of the information it contains. This research was also partially supported by the NTU Presidential Postdoctoral Fellowship (award number 021820-00001).

¹H. X. Pham and E. Kayacan are with the Artificial Intelligence in Robotics Laboratory (Air Lab), Department of Electrical and Computer Engineering, Aarhus University, 8000 Aarhus C, Denmark {huy.pham, erdal} at ece.au.dk

²A. Sarabakha is with the School of Electrical and Electronic Engineering (EEE), Nanyang Technological University (NTU), Singapore, 639798. andriy.sarabakha at ntu.edu.sg

³M. Odnoshykin is with the Munich School of Robotics and Machine Intelligence, Technical University of Munich (TUM), D-80797 Munich, Germany. mykola.odnoshykin at tum.de

Digital Object Identifier (DOI): see top of this page.

¹Source code, trained models, and collected datasets are available at <https://github.com/open-airlab/pencilnet>.




Fig. 1. Composed image of simulated (top left) and real (bottom left) gates and converted counterparts (right-hand side) with pencil filter. The perception network is trained solely on simulation data and can accurately perceive racing gates in real-world environments through a fish-eye camera.

Significant progress for faster quadrotor flights has been made by addressing drone racing [10], previously thought of as merely an entertainment application. Although, autonomous drone racing is an excellent playground for the developed drone technologies, which can be transferred to other domains [11], such as search and rescue. In autonomous drone racing, a drone has to fly autonomously and safely through a track that consists of multiple racing gates [12].

Early works [13]–[15] typically use hand-crafted gate perception on top of traditional control and planning frameworks to achieve safe gate passing. However, they are less robust to real-world conditions [16]. Subsequent work exploits deep neural networks (DNNs) that are more effective in coping with perception uncertainties [17]. A variant ResNet-8 DNN with three residual blocks is proposed in [18] and explicitly estimates a gate distribution. The method in [19] utilizes a variant of U-Net [20] that provides gate segmentation and demonstrates robust performance on top of a well-engineered system. The approach in [16] also segments the gates with U-Net but associates the gate corners differently by combining corner pixel searching and re-projection error rejection. Recently, end-to-end methods [21]–[23] using a single DNN that outputs tracking targets for controllers are considered to reduce the latency introduced by traditional modulation of perception, planning, and control sub-problems. In both lines of work, the quality of a gate perception, either through explicit mapping of gates or implicit feature-based understanding, dictates the

system’s overall performance.

Another point to consider is the ease of training the perception network for drone racing. Unlike many robotics applications, the cost to collect labeled high-speed real-world data for drone racing is often prohibitive, limiting the use of supervised methods. Therefore, learning through simulation data is understandably much desired [24]. For this purpose, [21] uses domain randomization for sim-to-real transfer learning with robust performance. A reinforcement learning method is proposed in [22] to learn in simulation a policy for drones that fly at high speeds but reported high tracking errors in practice due to the abstraction from input data to real physical drones is far more complex in reality than in simulation. Unfortunately, none of these methods has been extensively tested under various real-world conditions, such as illumination alterations, which are strongly relevant for drone racing and widely known to negatively affect the performance of perception networks.

In this work, we propose a hybrid DNN-based approach demonstrating an effective zero-shot sim-to-real transfer learning capability for gate perception in autonomous drone racing navigation. By relying on a smarter input representation using a low-cost morphological operation, called pencil filtering, we can bridge the sim-to-real gap so that our perception network is trained solely in simulation data and works robustly in real-world racing tracks, even under rapid motions and drastic lighting changes. The proposed framework attains a high-speed inference due to the use of our previously proposed efficient backbone network [25] with a modification to allow it to work with the new input representation. Our experiments show that the network’s performance significantly outperforms the original backbone network [25] and other state-of-the-art baseline methods in terms of accuracy and robustness.

The rest of this work is organized as follows. The proposed method for gate perception is detailed in Section II. Section III presents experiments to evaluate the efficiency of the approach against the state-of-the-art baselines. A real-world validation study is presented in Section IV. Finally, Section V concludes this work.

II. PROPOSED METHOD

A. Pencil Filter

Drastic illumination changes and blurriness caused by rapid flights have been extensively documented in previous work as the main factors behind the performance degradation of visual-based navigation frameworks. In this work, we tackle these challenges by employing an edge-enhancing technique known as the pencil filter [26], but unlike [26], we study its performance under various light intensities and motion blur. The pencil filter applied to an image emphasizes geometrical features, such as edges, corners, and outlines but preserves the overall image structure. This reduces the real-to-sim gap since images from simulation lack accuracy in objects’ sharpness and illumination compared to real images, allowing learning on synthetic data and applying in the real-world environment. In addition, the pencil filter is an illumination agnostic method that keeps the light intensity constant in the images. To extract geometrical features pencil filter first converts an RGB

Algorithm 1 Pencil filter.

```

1: Input: RGB image  $\mathbf{I}$ 
2: Output: pencil image  $\mathbf{P}$ 
3: function PENCIL_FILTER( $\mathbf{I}$ )
4:    $\mathbf{G} \leftarrow \text{grayscale}(\mathbf{I})$ 
5:    $\mathbf{P} \leftarrow \text{dilate}(\mathbf{G}, \text{ellipse})$ 
6:   for each  $y \in [1, \text{rows}(\mathbf{P})]$  do
7:     for each  $x \in [1, \text{columns}(\mathbf{P})]$  do
8:       if  $\mathbf{P}(y, x) = 0$  then
9:          $\mathbf{P}(y, x) \leftarrow 255$ 
10:      else
11:         $\mathbf{P}(y, x) \leftarrow \text{int} \left( 255 \frac{\mathbf{G}(y, x)}{\mathbf{P}(y, x)} \right)$ 
12:      end if
13:    end for
14:  end for
15:  return  $\mathbf{P}$ 
16: end function

```

image to a gray-scale format. Then, it applies dilation by convolving the gray-scale image with an ellipse kernel and thresholding to obtain local maxima for each pixel in its neighborhood. This morphological operation generally exaggerates the bright regions of the image. To avoid information loss in the dark regions, the difference in pixel values of the gray-scale image and the dilated one is calculated by normalizing their respective values. The resulting image retains sharp and bolder edges, accentuating object boundaries. Thus, sim-to-real transfer performance is strongly enhanced by converting training RGB images into pencil images. The resulting pencil image is shown on the right-hand side of Fig. 1; while the pseudo-code of the pencil filter is provided in Algorithm 1.

B. Training Data Generation

The key feature to reduce the gap between simulation and real-world appearance is the use of photo-realistic images. Graphical gaming engines allow simulating intricate ambient effects, mimicking the visual appearance of real environments and the surrounding noise and perturbations. In this work, photo-realistic images are generated by Unreal Engine containing racing scenes with vertical gates in multiple environments with different conditions. Similar to previous drone racing studies in [19], [25], we use square-shaped gates with checkerboard-like patterns around each corner. To generate a large number of distinct images, a wide-angle RGB camera and the gates are spawned randomly in the environment. They have to be distorted and interpolated to obtain simulated images with the same camera parameters as those from a real camera with a fish-eye lens. Finally, the pencil filter is applied to the fish-eye images. The pipeline for generating training images is depicted in Fig. 2.

Using ground-truth information from the simulator, the RGB fish-eye images are labeled with five variables as a tuple $\{x, y, d, \theta, c\}$. x and y are normalized pixel differences between the top-left corner of the grid and the gate’s center, d is the relative distance to the gate (in meters), and θ is the relative orientation of the gate (in radians), calculated as the

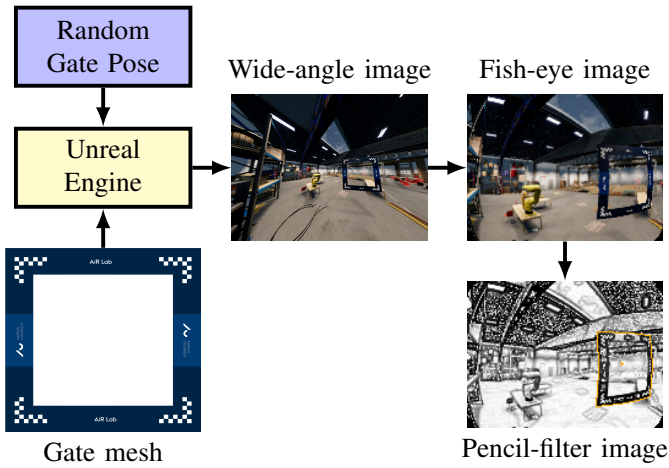



Fig. 2. Synthetic images generation pipeline: first, gates are randomly spawned in the simulation environment; then, wide-angle images are distorted with a fish-eye model; finally, the pencil filter is applied.

difference between gate’s yaw (heading) and camera’s yaw projected onto the horizontal plane. c is the confidence value of a corresponding grid, which has a value of 1, if a gate center locates inside, and 0, otherwise. The training dataset includes more than 30K samples drawn from 13 different environments with 18 distinct backgrounds and various lighting conditions. Each data sample consists of annotations of the gates and a 160×120 RGB image (illustrated in Fig. 3a). We do not apply any augmentation to the gates and images. In the training phase, the RGB images are converted into corresponding pencil-filtered images (illustrated in Fig. 3b), which retain the same annotations. The advantage of using simulation for data collection is that the data distributions are fully controlled, as shown in Fig. 4a. Note that the training data only contain labels for the front side of a gate, as the back side does not provide relevant information.

C. Backbone Detection Network

The detection network takes a 160×120 image converted by the pencil filter as input before normalization. For feature detection, PencilNet, illustrated in Fig. 5, utilizes as the backbone a convolutional neural network (CNN), from our previous work [25], comprising of six 2D convolutional




(a) Synthetic images after applying fish-eye distortion.




(b) Synthetic images after applying pencil filter.


Fig. 3. Samples of the simulation data before (3a) and after (3b) the pencil filter are applied. The images are generated from different domains with varying angles and illumination settings. Note that the training images do not contain bounding boxes of gates, which are illustrated here merely for better visualization.



(a) Synthetic training dataset.



(b) Synthetic testing dataset.



(c) Real-world testing dataset.

Fig. 4. Distributions of x and y -coordinates, distance and orientation of the gate on images in training and testing datasets.

(conv2D) layers with a small number of filters in each layer and one fully-connected (dense) layer to regress detection outputs. The tensors of PencilNet are modified to process one-channel images leading to a smaller number of parameters. Each conv2D layer is followed by a batch normalization and a rectified linear unit (ReLU) activation. Besides, the first five layers use max-pooling with 2×2 kernels. In the last conv2D layer, a small $3 \times 5 \times 16$ tensor is flattened so that each hidden neuron in the following dense layer can be connected with the extracted features. As in [25], the network output is reshaped to $R \times C \times F$, where each grid cell from $R = 4$ rows and $C = 3$ columns contains regressed $F = 5$ values of $\{x, y, d, \theta, c\}$.

D. Training

Similarly to [25], PencilNet uses four losses \mathcal{L}_{xy} , \mathcal{L}_d , \mathcal{L}_θ and \mathcal{L}_c , which are center coordinate, distance, orientation and confidence deviations from true values, respectively:

$$\begin{cases} \mathcal{L}_{xy} &= \sum_{i=1}^R \sum_{j=1}^C \mathbb{1}_{ij}^{\text{obj}} \left[(x_{ij} - \hat{x}_{ij})^2 + (y_{ij} - \hat{y}_{ij})^2 \right], \\ \mathcal{L}_d &= \sum_{i=1}^R \sum_{j=1}^C \mathbb{1}_{ij}^{\text{obj}} \left(d_{ij} - \hat{d}_{ij} \right)^2, \\ \mathcal{L}_\theta &= \sum_{i=1}^R \sum_{j=1}^C \mathbb{1}_{ij}^{\text{obj}} \left(\theta_{ij} - \hat{\theta}_{ij} \right)^2, \\ \mathcal{L}_c &= \sum_{i=1}^R \sum_{j=1}^C \left[\mathbb{1}_{ij}^{\text{obj}} + \alpha(1 - \mathbb{1}_{ij}^{\text{obj}}) \right] (c_{ij} - \hat{c}_{ij})^2. \end{cases} \quad (1)$$

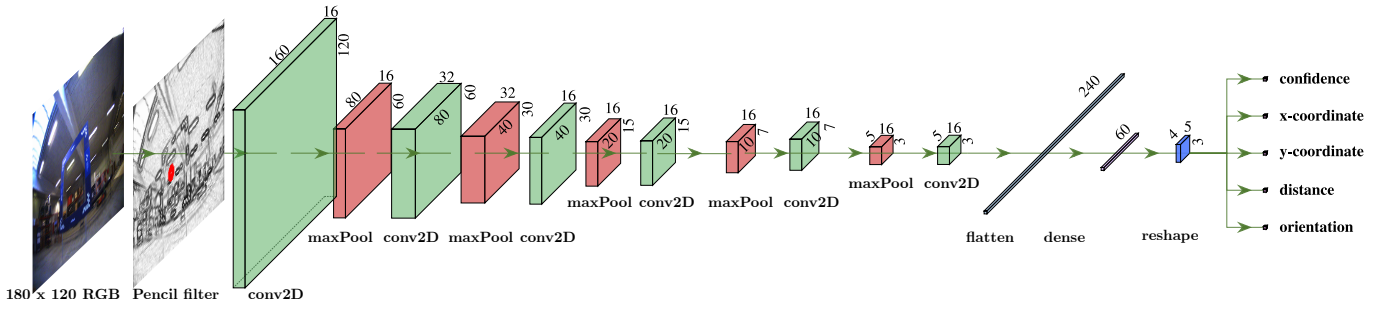


Fig. 5. Architecture of the proposed method: PencilNet. The pencil filter is applied to RGB images before feeding them to the first conv2D layers. The backbone network is similar to [25], with the modification in reducing the number of channels for each filter to work with one-channel input. Each of the first five conv2D layers is followed by a max-pooling with a 2×2 kernel. The tensor in the last conv2D layer is flattened to increase connectivity with the following dense layer. The output tensor is reshaped as a grid of 4×3 cells, where each cell contains five prediction values for pixel offsets to the gate’s center, relative distance and orientation of the gate, and a confidence value.

The losses are calculated for each grid cell of the output layer with row and column indices $i \in [1, R]$ and $j \in [1, C]$. Terms with the hat operator ($\hat{\cdot}$) denote predicted values. The term $\mathbb{1}_{ij}^{\text{obj}}$ is a binary variable to penalize the network only when a gate center locates in a particular grid. The confidence values for grid cells where a gate center does not present is minimized with a weight α so that in run time we can threshold the low confidence predictions, thus reducing false positives. Finally, the loss function is a weighted term:

$$\mathcal{L} = \lambda_{xy}\mathcal{L}_{xy} + \lambda_d\mathcal{L}_d + \lambda_\theta\mathcal{L}_\theta + \lambda_c\mathcal{L}_c, \quad (2)$$

where λ_{xy} , λ_d , λ_θ , λ_c are weights reflecting importance for each of the losses.

III. GATE DETECTION EVALUATION

A. Baseline Methods

In the first part of our study, we evaluate PencilNet performance for the task of racing gate perception with a few notable milestone methods that include DroNet variants (used in [18] and [27]), and ADRNet variants [28]. We also include the performance of our previously published GateNet [25], to evaluate the contribution of the pencil filter.

Remark 1: The above baselines are chosen because of their similarities to our method. Some other state-of-the-art baselines such as [19] and [16] are not included in this study due to the differences in training data and target outputs.

1) *DroNet variants* [18], [21]: A variant of ResNet-8 DNN with three residual blocks is proposed in [21] to originally learn a steering policy for a quadrotor to avoid obstacles. However, due to its efficient structure, many subsequent papers utilize it as a backbone feature extractor with different output layers. Kaufmann et al. [18] replace the DroNet’s original output layers with multi-layer perceptron to estimate the mean and variance of the next gate’s pose expressed in spherical coordinates and use the predicted pose for a model predictive control framework to compute a trajectory for the drone. Loquercio et al. [21] train DroNet with simulation data using domain randomization and also change the output layers to produce a velocity vector that drives the robot through the gate. Pfeiffer et al. [23] combine DroNet with an attention map as an end-to-end network to capture visual-spatial information allowing the system to track high-speed trajectories. In this

work, we compare PencilNet with two variants of DroNet: DroNet-1.0 with the full set of parameters as in [18], and DroNet-0.5 with only half of the filters as in [21]. The original output layers are also modified to produce similar outputs as our method for compatibility reasons.

2) *ADRNet variant* [28]: A single-shot object detector based on AlexNet [29] is proposed in [28] to detect the bounding box of a gate in an image frame, which is used to calculate offset velocities needed to guide a drone past the gate. Since the original ADRNet only detects the gate center and lacks distance and orientation estimation, this work also replaces its output layer with our output layer to predict additional attributes (named ADRNet-mod).

3) *GateNet variants* [25]: To comprehensively evaluate the effect of the pencil filter, PencilNet is compared against the same backbone network, i.e., GateNet, tailored with other edge detection filters. Three different classical filters are considered, namely the Sobel filter [30] and Canny filter [31]. Similar to the pencil filter, these filters transform the input images before feeding them to the backbone detection network, as illustrated in Fig. 6.

B. Baselines Training

As mentioned in Section II-B, PencilNet and the other baseline methods are trained with the same simulation dataset containing 160×120 RGB images. Note that filter models, such as the proposed method, apply the respective filter on these images before feeding them into its network. All methods




Fig. 6. Visualization of different filters that are considered and evaluated under bright (the first row) and dark (the second row) lighting: (6a) original RGB images, and images converted by (6b) pencil filter, (6c) Sobel filter, (6d) Canny filter, and (6e) Canny filter with a more conservative threshold.

use the same training settings, that is set with a batch size of 32, the initial learning rate of 0.01, and linearly decayed by 0.1 at epochs 5 and 8. The models are optimized by Adam [32] with default parameters of $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$.

C. Evaluation Metrics

In the evaluation experiments, mean absolute errors (MAE) are calculated for the aforementioned methods as a measure of accuracy in predicting gate's center on the image plane (E_c), distance to the gate (E_d), and orientation of the gate (E_θ) relative to the drone's body frame with the ground-truth information as follows:

$$\begin{cases} E_c = \frac{1}{N} \sum_{i=1}^R \sum_{j=1}^C (|\hat{x}_{ij} - x_{ij}| + |\hat{y}_{ij} - y_{ij}|), \\ E_d = \frac{1}{N} \sum_{i=1}^R \sum_{j=1}^C |\hat{d}_{ij} - d_{ij}|, \\ E_\theta = \frac{1}{N} \sum_{i=1}^R \sum_{j=1}^C |\hat{\theta}_{ij} - \theta_{ij}|, \end{cases} \quad (3)$$

where N is the number of test samples. The metrics in (3) provide the detection accuracy but do not take into account cases where the model fails to make a prediction and, consequently, produces no errors. Therefore, in addition to the metrics in (3), we consider another metric that calculates the percentage of false negative (FN) predictions, i.e., the number of times a network does not detect a gate present in the image over the total number of samples. If a method produces a high FN, it performs poorly even though the other metrics may appear decent.

D. Sim-to-Sim and Sim-to-Real Evaluation

Remark 2: By assumption, all gates are upright in all simulation and real-world environments.

The first objective of the evaluation is to understand how well the models above generalize their understanding of the gate perception task in different environments. For this purpose, all models are evaluated on collected datasets. For the synthetic test dataset (Sim), samples are drawn from three simulation environments that have not been used during the training of PencilNet and other baseline methods. The real-world images, illustrated in Fig. 7, are obtained from a quadrotor flying fast random trajectories through gates positioned in different layouts at various lighting intensities. The quadrotor is equipped with a high-frequency camera with a fish-eye lens. A motion capture system is used to provide ground-truth information for the poses of the drone and gates. Four real-world datasets (*Real N-100*, *Real N-40*, *Real N-20* and *Real N-10*) are collected at night time with 100%, 40%,

20% and 10% of artificial light intensity, respectively. Additionally, four datasets (*Blur Real N-100*, *Blur Real N-40*, *Blur Real N-20* and *Blur Real N-10*) are collected containing blur images by increasing both the exposure time of the camera and flight speeds of the drone at different illumination levels. The statistics of the synthetic test dataset are similar to those of the training dataset; while the real-world dataset covers a wide range of situations with great variations in gate poses, as shown in Figs. 4b and 4c.

Table I shows the comparison of the gate detection performance in various test environments with different backgrounds and illumination conditions. The number of parameters in each network and their inference speed measured on an NVIDIA Jetson TX2 onboard computer in frames per second (fps) are compared. The proposed method performs reliably in all environments and illumination conditions. It also outperforms the baseline methods in real-world datasets, particularly in lower illumination. Sobel filter appears to have outstanding performance in the *Real-10* dataset, but it produces a high number of FN predictions and, consequently, performs poorly in practice. Since PencilNet uses GateNet as its backbone network and performs a low-cost operation of pencil filter on small-size images ($\approx 0.3\text{ms}$ on 160×120 images), it attains a high inference rate in real-time.

When illumination decreases, all baseline methods predict with lower accuracy and produce a higher number of FN detections. At the same time, the performance of PencilNet retains high accuracy and a low number of FN predictions. PencilNet works well in darker environments is due to the dilation operation in the pencil filter that makes the edges sharper by extracting local maxima within the neighborhood of each pixel. Thus, PencilNet can enhance dim edges and enable the detection of the objects' geometry. Although other edge detection filters used for the GateNet variants could theoretically extract similar features, their performance does not translate well in reality. The main distinction between PencilNet and other GateNet variants is the robustness to noise. By its nature, the Sobel filter is based on gradient calculation which is more sensitive to the noise in the image. On the other hand, the Canny filter is more robust to noise thanks to the application of non-maximum suppression. However, the hysteresis thresholding requires selecting thresholds that have to be tuned for various light intensities, as illustrated in Figs. 6d and 6e. The local maxima extraction through dilation in the pencil filter does not require parameter tuning. Therefore, PencilNet is more suitable for drone racing applications where sudden illumination changes occur frequently.

On the other hand, PencilNet also performs notably better in sim-to-real schemes compared to other RGB-based models (ADRNNet, DronNet variants, and GateNet). Even though being trained on photo-realistic images and domain randomization and having additional color information, it is difficult for these baseline models to understand the real-world environment represented in RGB images. Nevertheless, the loss of color information can negatively affect PencilNet's performance in some situations where the color of the gate can be clearly observable. It is possible to observe from Table I, that PencilNet produces a slightly higher number of FN predictions when




Fig. 7. Real-time experiment scenes with different conditions such as (a) high, (b) medium and (c) low illumination, and (d) blurriness.

TABLE I

COMPARISON OF VARIOUS METHODS ON THE TEST DATASETS FROM DIFFERENT ENVIRONMENTS IN TERMS OF THE NUMBER OF PARAMETERS (#P), INFERENCE SPEED IN FRAMES PER SECOND (FPS), ACCURACY METRIC (E_c IN PIXELS, E_d IN METERS, E_θ IN RADIANS) AND FN DETECTIONS IN PERCENTS.

Method	#p	fps	Sim				Real N-100				Real N-40				Real N-20				Real N-10			
			E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN
ADRNNet-mod [28]	2.5M	14	0.013	0.284	0.033	24.26	0.067	1.028	0.252	19.64	0.059	0.713	0.143	86.56	0.06	0.687	0.14	98.88	0.058	0.685	0.136	97.73
DroNet-1.0 [18]	478K	22	0.048	0.119	0.023	62.44	0.145	0.797	0.217	61.46	0.116	0.657	0.129	99.91	0.115	0.637	0.125	99.85	0.115	0.629	0.123	99.84
DroNet-0.5 [27]	250K	40	0.045	0.106	0.022	51.58	0.144	0.778	0.216	47.01	0.115	0.648	0.131	98.32	0.115	0.626	0.129	99.55	0.114	0.622	0.126	99.62
GateNet [25]	32K	57	0.018	0.036	0.02	24.96	0.071	0.323	0.228	15.74	0.062	0.222	0.151	66.01	0.066	0.234	0.145	65.23	0.064	0.238	0.139	72.07
GateNet + sobel	32K	57	0.032	0.068	0.022	27.96	0.077	0.385	0.217	14.32	0.054	0.219	0.13	79.08	0.054	0.219	0.131	88.36	0.053	0.223	0.128	90.79
GateNet + canny	32K	57	0.039	0.112	0.026	78.67	0.072	0.299	0.219	95.30	0.061	0.264	0.145	90.56	0.067	0.278	0.144	93.10	0.065	0.268	0.142	89.88
PencilNet	32K	57	0.017	0.031	0.018	20.5	0.065	0.287	0.210	18.78	0.051	0.2	0.129	35.10	0.052	0.209	0.135	36	0.053	0.235	0.14	50.18

TABLE II

COMPARISON OF VARIOUS METHODS ON THE BLUR IMAGES DATASETS UNDER DIFFERENT ILLUMINATION CONDITIONS IN TERMS OF THE NUMBER OF PARAMETERS (#P), INFERENCE SPEED IN FRAMES PER SECOND (FPS), ACCURACY METRIC (E_c IN PIXELS, E_d IN METERS, E_θ IN RADIANS) AND FN DETECTIONS IN PERCENTS.

Method	Blur Real N-100				Blur Real N-40				Blur Real N-20				Blur Real N-10			
	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN	E_c	E_d	E_θ	FN
GateNet	0.056	0.3	0.117	12.45	0.057	0.234	0.134	19.21	0.073	0.236	0.138	59.24	0.077	0.248	0.145	48.52
GateNet + sobel	0.067	0.26	0.125	38.58	0.068	0.25	0.123	39.61	0.068	0.253	0.119	82.23	0.069	0.277	0.129	88.41
GateNet + canny	0.069	0.253	0.133	85.81	0.08	0.234	0.135	80	0.099	0.275	0.135	85.99	0.101	0.281	0.147	85.17
PencilNet	0.048	0.248	0.128	20.07	0.049	0.225	0.122	14.71	0.055	0.207	0.121	23.69	0.061	0.215	0.137	26.14

evaluated on the real-world dataset *Real N-100* compared to an RGB-based baseline, i.e., GateNet. In this data, the light intensity is maximum. Therefore, the color of the gate is clearly visible, providing an advantage for the RGB-based baseline. However, the rest of the evaluation and experiments illustrate that, in general, learning geometry through enhanced edge detection has better performance.

To further highlight the robustness of the pencil filter over other edge enhancement filter types, Table II shows the comparison of different filters for gate detection tasks in real-world blur images that are also commonly observed in drone racing scenarios. PencilNet significantly outperforms other filters in gate center and distance prediction metrics, while having strong performance on orientation prediction with a much lower number of FNs.

A pivotal idea to ensure a successful sim-to-real transfer learning is to narrow the gap between the quality of images rendered by the simulator and the real images obtained from the drone's camera. Thanks to the pencil filter's intelligent abstraction of input data, the network performance is not affected much by ambience noise and intricate lighting in real-world conditions, as shown in Fig. 8. This finding is aligned

with previous methods that are trained only on simulation data but achieve good real-time performances by adopting a simpler input representation, such as depth map [33], [34]. The pencil filter should also be considered as another smart input representation for sim-to-real applications.

IV. EXPERIMENTAL VALIDATION

To demonstrate the usefulness of our proposed method, real-world experiments are set up with typical conditions as in recent autonomous drone racing applications, in which a small quadrotor traverses through various sequences of square gates. The drone is equipped with a single RGB camera running at 50Hz for gate perception, and an Intel Realsense Tracking camera T265 provides pose and velocities estimation. In the following experiments, PencilNet and the baseline methods are run on an NVIDIA Jetson TX2 computer, together with other perception, planning, and control pipelines whose details are provided below.

A. Gate Mapping

As in similar studies [18], [19], [25], we construct a map (as shown in Fig. 9) of the gates in the racing tracks from




Fig. 8. Samples of the gate perception by PencilNet. The network can distinguish a gate from objects with similar geometry such as the back of another gate, or in extreme dark lighting.




Fig. 9. Example of the constructed global map. The pink blocks represent the last raw estimation of the gate's pose; while the green blocks represent the estimations by an extended Kalman filter. The generated trajectory is shown with red spheres.

the output of a perception network, which is a prediction vector $\hat{\mathbf{v}} = [\hat{c}_x \ \hat{c}_y \ \hat{d} \ \hat{\theta}]^T$ of a gate center pixel coordinates (\hat{c}_x, \hat{c}_y) , relative distance from the robot to the gate \hat{d} , and relative heading angle $\hat{\theta}$. Using the method of fish-eye back-projection proposed in [25], one can obtain the pose of the gate G expressed in world frame W (${}_W\hat{\mathbf{p}}_G$) as follows:

$$\begin{cases} {}_C\hat{\mathbf{p}}_G = \mathbf{D}^{-1}\mathbf{K}^{-1}\hat{\mathbf{d}} [\hat{c}_x, \hat{c}_y, 1]^T, \\ {}_W\hat{\mathbf{p}}_G = {}_W\mathbf{p}_D \\ + (\mathbf{R}_D^W)^{-1} \left((\mathbf{R}_C^D)^{-1} \begin{bmatrix} 0 & 0 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix} {}_C\hat{\mathbf{p}}_G + {}_D\mathbf{p}_C \right). \end{cases} \quad (4)$$

where W , D , and C are the world, drone’s body, and camera coordinate frames, respectively. \mathbf{R}_B^A expresses the transformation from frame A to frame B . \mathbf{D} and \mathbf{K} are the camera distortion matrix and intrinsic matrix. In real-time flights, an extended Kalman filter (EKF) is used for each detected gate with prior knowledge of coarse gate locations within a radius of six meters on the track, to correctly associate the prediction values with a gate.

B. Motion Planning and Control

After the global pose of the next gate is estimated, a minimum-snap trajectory is generated [35], [36] from the drone’s current position to a perceived gate’s center with a receding-horizon local planner that can replan if the estimation of the next gate pose exceeds a threshold. The trajectory is tracked by a geometric controller [37].

C. Real-Time Experiments

The vision-based autonomous navigation system is tested in multiple scenarios with different track layouts and backgrounds, as shown in Table III. The level of difficulties for each track is reflected by the time window (T), measured by the average gate-to-gate distance and max drone speeds, in which a perception network must predict a gate correctly, and the controller reacts timely. The narrow tracks (Circular and S-shape) have a shorter time window ($T = 1.6\text{s}$ and $T = 1.16\text{s}$) so it is more difficult to navigate successfully. The 8-shape track has a longer time window ($T = 3.25\text{s}$). The illumination

settings of the test environment can be controlled using a led lighting system at night and roof windows during the daytime.

The performance of two baseline methods, namely GateNet [25] and DroNet-0.5 [18], are also reported due to their similarities with our method and their decent speed of inferences. In Table III, we report the success rate for gate crossing for each scenario; while, in Fig. 10, the Euclidean errors of each method are depicted. As the inner gate dimension is around 1m, a deviation above 0.5m normally results in a crash. Since the gates are predicted and then mapped globally, the accuracy of the perception pipeline in the navigation task is directly affected also by the accuracy of the state estimation, which could drift significantly during high-speed flights under degraded visual conditions. This partly explains the poor success rates in the extreme lighting condition (N-20, N-10) for all tested tracks using SLAM. However, the results still show that compared to the similar perception pipeline in [25], the pencil filter has significantly improved the accuracy of gate detection in practice, especially in darker illumination settings, hence consequently the overall success rate of the mission. Such improvements are valuable in drone racing contests where gates are placed in visually degraded conditions. Readers are encouraged to watch the accompanying video for a detailed review of the performance of the system.




Fig. 10. Euclidean position and orientation errors for the gate detection with various illumination settings.

TABLE III

SUCCESS RATES OF GATE CROSSING WITH VARIOUS TRACK LAYOUT AND ILLUMINATION SETTINGS.

Method	Track	Lighting				
		Day	N-100	N-40	N-20	N-10
PencilNet (ours)	Circular, narrow	90%	100%	90%	60%	10%
	8-shape, wide	90%	100%	75%	40%	10%
	S-shape, narrow	95%	90%	60%	10%	10%
GateNet [25]	Circular, narrow	60%	100%	50%	30%	0%
	8-shape, wide	25%	80%	12.5%	0%	0%
	S-shape, narrow	12.5%	90%	0%	0%	0%
DroNet-1.0 [18]	Circular, narrow	60%	60%	0%	30%	0%
	8-shape, wide	75%	70%	0%	0%	0%
	S-shape, narrow	50%	60%	0%	0%	0%
DroNet-0.5 [27]	Circular, narrow	80%	60%	0%	0%	0%
	8-shape, wide	75%	80%	12.5%	0%	0%
	S-shape, narrow	40%	60%	0%	0%	0%

V. CONCLUSION

This work proposes PencilNet – a DNN method utilizing morphological operation to improve both the sim-to-real transfer learning capability and the robustness of the gate perception system even in extreme illumination settings. Through an extensive evaluation study, PencilNet’s performance is shown to significantly outperform state-of-the-art methods for both sim-to-real transfer learning and robustness in detection. We demonstrate the effectiveness of PencilNet in multiple real-time autonomous racing scenarios under different lighting settings that reflect realistic conditions. The system’s performance can be improved by accounting for the drifts and degraded

performance in state estimation in darker lighting, which negatively affects the accuracy of the perception system.

ACKNOWLEDGMENT

We would like to thank Micha Hei, Jonas le Fevre, Abdel Hakim, Jakob Grimm, Ilker Bozcan, and Michal Kozlowski for fruitful scientific discussions.

REFERENCES

- [1] K. F. Andersen, H. X. Pham, H. I. Ugurlu, and E. Kayacan, "Event-based navigation for autonomous drone racing with sparse gated recurrent network," *2022 European Control Conference (ECC)*, 2022.
- [2] E. Camci and E. Kayacan, "Learning motion primitives for planning swift maneuvers of quadrotor," *Autonomous Robots*, vol. 43, no. 7, pp. 1733–1745, 2019.
- [3] M. Blsch, S. Weiss, D. Scaramuzza, and R. Siegwart, "Vision based MAV navigation in unknown and unstructured environments," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 21–28.
- [4] M. Mehdiratta and E. Kayacan, "Gaussian process-based learning control of aerial robots for precise visualization of geological outcrops," in *2020 European Control Conference (ECC)*, 2020, pp. 10–16.
- [5] M. Dharmadhikari, T. Dang, L. Solanka, J. Loje, H. Nguyen, N. Khedekar, and K. Alexis, "Motion primitives-based path planning for fast and agile exploration using aerial robots," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 179–185.
- [6] E. Camci and E. Kayacan, "Planning swift maneuvers of quadcopter using motion primitives explored by reinforcement learning," in *2019 American Control Conference (ACC)*, 2019, pp. 279–285.
- [7] Z. Wang, X. Xiao, A. J. Nettekoven, K. Umasankar, A. Singh, S. Bommakanti, U. Topcu, and P. Stone, "From Agile Ground to Aerial Navigation: Learning from Learned Hallucination," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 148–153.
- [8] S. Patel, A. Sarabakha, D. Kircali, and E. Kayacan, "An Intelligent Hybrid Artificial Neural Network-Based Approach for Control of Aerial Robots," *Journal of Intelligent & Robotic Systems*, May 2019.
- [9] S. Patel, A. Sarabakha, D. Kircali, G. Loianno, and E. Kayacan, "Artificial Neural Network-Assisted Controller for Fast and Agile UAV Flight: Onboard Implementation and Experimental Results," in *2019 Workshop on Research, Education and Development of Unmanned Aerial Systems (RED-UAS)*, 2019, pp. 37–43.
- [10] H. Moon, J. Martinez-Carranza, T. Cieslewski, M. Faessler, D. Falanga, A. Simovic, D. Scaramuzza, S. Li, M. Ozo, C. De Wagter, G. de Croon, S. Hwang, S. Jung, H. Shim, H. Kim, M. Park, T.-C. Au, and S. J. Kim, "Challenges and implemented technologies used in autonomous drone racing," *Intelligent Service Robotics*, vol. 12, no. 2, pp. 137–148, Apr 2019.
- [11] H. X. Pham, H. I. Ugurlu, J. Le Fevre, D. Bardakci, and E. Kayacan, "Chapter 15 - Deep learning for vision-based navigation in autonomous drone racing," in *Deep Learning for Robot Perception and Cognition*. Academic Press, 2022, pp. 371–406.
- [12] C. De Wagter, F. Paredes-Valls, N. Sheth, and G. de Croon, "Learning fast in autonomous drone racing," *Nature Machine Intelligence*, vol. 3, no. 10, pp. 923–923, Oct 2021.
- [13] S. Jung, S. Cho, D. Lee, H. Lee, and D. H. Shim, "A direct visual servoing-based framework for the 2016 IROS Autonomous Drone Racing Challenge," *Journal of Field Robotics*, vol. 35, no. 1, pp. 146–166, 2018.
- [14] S. Li, M. M. Ozo, C. De Wagter, and G. C. de Croon, "Autonomous drone race: A computationally efficient vision-based navigation and control strategy," *Robotics and Autonomous Systems*, vol. 133, p. 103621, 2020.
- [15] L. O. Rojas-Perez and J. Martinez-Carranza, "Metric monocular SLAM and colour segmentation for multiple obstacle avoidance in autonomous flight," in *2017 Workshop on Research, Education and Development of Unmanned Aerial Systems (RED-UAS)*, 2017, pp. 234–239.
- [16] C. De Wagter, F. Paredes-Valles, N. Sheth, and G. de Croon, "The Artificial Intelligence behind the winning entry to the 2019 AI Robotic Racing Competition," *arXiv preprint arXiv:2109.14985*, 2021.
- [17] T. Morales, A. Sarabakha, and E. Kayacan, "Image Generation for Efficient Neural Network Training in Autonomous Drone Racing," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.
- [18] E. Kaufmann, M. Gehrig, P. Foehn, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Beauty and the beast: Optimal methods meet learning for drone racing," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 690–696.
- [19] P. Foehn, D. Brescianini, E. Kaufmann, T. Cieslewski, M. Gehrig, M. Muglikar, and D. Scaramuzza, "Alphapilot: Autonomous drone racing," *Autonomous Robots*, vol. 46, no. 1, pp. 307–320, 2022.
- [20] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [21] A. Loquercio, A. I. Maqueda, C. R. Del-Blanco, and D. Scaramuzza, "Dronet: Learning to fly by driving," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1088–1095, 2018.
- [22] 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.
- [23] C. Pfeiffer, S. Wengeler, A. Loquercio, and D. Scaramuzza, "Visual attention prediction improves performance of autonomous drone racing agents," *PLOS ONE*, vol. 17, pp. 1–16, 03 2022.
- [24] S. Zakhharov, I. Shugurov, and S. Ilic, "Dpod: 6D pose object detector and refiner," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 1941–1950.
- [25] H. X. Pham, I. Bozcan, A. Sarabakha, S. Haddadin, and E. Kayacan, "GateNet: An Efficient Deep Neural Network Architecture for Gate Perception Using Fish-Eye Camera in Autonomous Drone Racing," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 4176–4183.
- [26] J. Rambach, C. Deng, A. Pagani, and D. Stricker, "Learning 6DOF object poses from synthetic single channel images," in *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, 2018, pp. 164–169.
- [27] 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.
- [28] S. Jung, S. Hwang, H. Shin, and D. H. Shim, "Perception, guidance, and navigation for indoor autonomous drone racing using deep learning," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2539–2544, 2018.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [30] N. Kanopoulos, N. Vasanthavada, and R. Baker, "Design of an image edge detection filter using the sobel operator," *IEEE Journal of Solid-State Circuits*, vol. 23, no. 2, pp. 358–367, 1988.
- [31] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, 1986.
- [32] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *CoRR*, vol. abs/1412.6980, 2015.
- [33] E. Camci, D. Campolo, and E. Kayacan, "Deep reinforcement learning for motion planning of quadrotors using raw depth images," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–7.
- [34] A. Loquercio, E. Kaufmann, R. Ranftl, M. Mller, V. Koltun, and D. Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.
- [35] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Robotics Research*. Springer, 2016, pp. 649–666.
- [36] M. Burri, H. Oleynikova, M. W. Achtelik, and R. Siegwart, "Real-Time Visual-Inertial Mapping, Re-localization and Planning Onboard MAVs in Unknown Environments," in *Intelligent Robots and Systems (IROS 2015)*, *2015 IEEE/RSJ International Conference on*, Sept 2015.
- [37] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor UAV on SE(3)," in *49th IEEE conference on decision and control (CDC)*. IEEE, 2010, pp. 5420–5425.