



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Bio-inspired Autonomous Exploration Policies with CNN-based Object Detection on Nano-drones

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Lamberti, L., Bompani, L., Kartsch, V.J., Rusci, M., Palossi, D., Benini, L. (2023). Bio-inspired Autonomous Exploration Policies with CNN-based Object Detection on Nano-drones [10.23919/DATE56975.2023.10137154].

Availability:

This version is available at: <https://hdl.handle.net/11585/953244> since: 2024-01-24

Published:

DOI: <http://doi.org/10.23919/DATE56975.2023.10137154>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

L. Lamberti, L. Bompani, V. J. Kartsch, M. Rusci, D. Palossi and L. Benini, "Bio-inspired Autonomous Exploration Policies with CNN-based Object Detection on Nano-drones," *2023 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, Antwerp, Belgium, 2023, pp. 1-6

The final published version is available online at:

<https://doi.org/10.23919/DATE56975.2023.10137154>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Bio-inspired Autonomous Exploration Policies with CNN-based Object Detection on Nano-drones

Lorenzo Lamberti*, Luca Bompani*, Victor Javier Kartsch*, Manuele Rusci[†], Daniele Palossi^{‡§} Luca Benini*[§]

* Department of Electrical, Electronic and Information Engineering, University of Bologna, Italy

[†] Department of Electrical Engineering, KU Leuven, Belgium

[‡] Dalle Molle Institute for Artificial Intelligence, USI-SUPSI, Switzerland

[§] Integrated Systems Laboratory, ETH Zürich, Switzerland

Contact authors: lorenzo.lamberti@unibo.it, luca.bompani5@unibo.it

Abstract—Nano-sized drones, with palm-sized form factor, are gaining relevance in the Internet-of-Things ecosystem. Achieving a high degree of autonomy for complex multi-objective missions (e.g., safe flight, exploration, object detection) is extremely challenging for the onboard chip-set due to tight size, payload (<10 g), and power envelope constraints, which strictly limit both memory and computation. Our work addresses this complex problem by combining bio-inspired navigation policies, which rely on time-of-flight distance sensor data, with a vision-based convolutional neural network (CNN) for object detection. Our field-proven nano-drone is equipped with two microcontroller units (MCUs), a single-core ARM Cortex-M4 (STM32) for safe navigation and exploration policies, and a parallel ultra-low power octa-core RISC-V (GAP8) for onboard CNN inference, with a power envelope of just 134 mW, including image sensors and external memories. The object detection task achieves a mean average precision of 50% (at 1.6 frame/s) on an in-field collected dataset. We compare four bio-inspired exploration policies and identify a pseudo-random policy to achieve the highest coverage area of 83% in a $\sim 36\text{ m}^2$ unknown room in a 3 minutes flight. By combining the detection CNN and the exploration policy, we show an average detection rate of 90% on six target objects in a never-seen-before environment.

SUPPLEMENTARY MATERIAL

In-field experiment video at: <https://youtu.be/BTin8g0nyko>.
Code at: <https://github.com/pulp-platform/pulp-detector>.

I. INTRODUCTION

The rapid evolution of the Internet of Things (IoT) is fueling the advent of flexible edge nodes powered by artificial intelligence (AI). In this context, AI-based nano-sized unmanned aerial vehicles (UAVs), with a form factor of only ~ 10 cm in diameter, can become revolutionary *ubiquitous* smart sensing IoT nodes capable of exploring an environment while interacting with it in full autonomy [1], i.e., fulfilling the mission without the need of external resources. The tiny and lightweight design of these pocket-sized drones provides enhanced safety for close-proximity human-machine interactions [2], as well as for indoor scenarios in narrow spaces, such as rescue missions [3].

Tackling such complex use cases requires the drone to reach a high level of intelligence typical of biological systems, which are autonomous and capable of handling multiple concurrent tasks [4] that span from basic control functionality to high-level perception [5], [6]. While autonomous nano-UAVs can already




Fig. 1. Our fully autonomous prototype, based on a Crazyflie nano-drone.

handle multiple basic functionality blocks at run-time (e.g., state estimation, low-level control), they still struggle to pursue multi-objective missions (e.g., safe navigation, exploration, visual recognition), which require computationally intensive multi-tasking perception [3]. In fact, with their ultra-constrained form factor and payload (a few tens of grams), these robots are prevented from hosting high-density/capacity batteries and bulky sensors or processors. This scenario is even further exacerbated by the fraction of the total power budget allotted for the onboard electronics, i.e., 5-10% of the total, which restricts the onboard processors to low-power microcontroller units (MCUs) [7].

Inspired by the functional mapping of the different regions of human brains [8], we address this problem by mapping multiple tasks on the two MCUs available aboard our nano-drone. The targeted system design follows the *host-accelerator* template. We leverage *i.*) a single-core STM32F405 MCU, with a peak performance $< 100\text{ M}$ multiply-accumulate (MAC) operations per second, for the execution of lightweight workloads (control, sensor-interfacing tasks), and *ii.*) a parallel ultra-low-power (PULP) multi-core co-processor, GAP8, for coping with convolutional neural network (CNN) workloads ($\sim 1\text{ GMAC/s}$).

In this work, we contribute with a fully autonomous system design that combines bio-inspired exploration policies with a vision-based CNN for object detection aboard a nano-drone – our prototype is shown in Fig. 1. We thoroughly characterize four bio-inspired exploration policies which use the ranging

measures of multiple single-beam Time-of-Flight (ToF) sensors and run on the STM32F405 MCU. Additionally, we deploy the object detection CNN on a GAP8 System-on-Chip (SoC) aboard our nano-drone, and we analyze its precision, computational/power costs, and performance by varying its depth.

Our in-field experiments show a peak performance for the biggest (deepest) object detector up to 1.6 frame/s, which recognizes two classes of objects with a mean average precision (mAP) score of 50% within only 134 mW. On average, the best exploration strategy covers 83% of a $\sim 36\text{m}^2$ room in a time budget of 3 min. Under the same testing environment, our closed-loop system recognizes six objects (two classes) with a mean detection rate of 90% (5 independent runs) with a mean flight speed of 0.5 m/s. Ultimately, to the best of our knowledge, we present the first fully autonomous nano-drone tackling a multi-objective mission, which consists of exploring, preventing collision, and detecting objects in real-time while relying only on onboard sensory and computational resources.

II. RELATED WORKS

AI-powered nano-drones. Recently, several AI-based visual pipelines have been brought to nano-drone systems. *Lamberti et al.* [9] developed a CNN for end-to-end autonomous visual navigation that achieves an inference throughput of 160 FPS on a nano-drone, thanks to a reduced CNN complexity (down to 1.5 MMAC). This network design has two outputs: the steering angle and the collision probability. The latter is trained to detect any obstacle without providing information about its class or position within the image frame. Similarly, *NanoFlowNet* [10] is a CNN for dense optical flow that aims at the obstacle avoidance task. It outputs per-pixel information about the obstacles in the scene, yet it does not give any information about their class. *Palossi et al.* [2] deployed a CNN for human pose estimation aboard a nano-drone, allowing a drone to follow the movement of a human subject. However, this network is not designed to recognize multiple subjects in the same frame or different classes of subjects. In contrast, our work aims to deploy an object detector to a nano-drone to detect multiple instances and classes of objects in its field of view.

Autonomous exploration. Computer vision-based navigation techniques, based on complex feature extraction pipelines [11] or simultaneous localization and mapping (SLAM) techniques [12], are reliable for autonomous robotic navigation. Still, as SLAM-based approaches have large memory requirements and rely on computationally intensive algorithms, they are exclusive to large UAV systems carrying heavy and high-power embedded computing systems [12].

The state-of-the-art autonomous exploration approaches for nano-drones, which suffer from limited computational power and memory, take advantage of bio-inspired (or bug-inspired) algorithms, which rely on lightweight state-machine-based algorithms and low-power sensor readings [13]. For example, [14] compares two bio-inspired exploration policies: the first changes the heading direction randomly after detecting an obstacle, while the second follows the obstacle's boundaries. Similar to the second one, [13] implements a policy that follows the walls of a room, called wall-following, while [15] describes

a spiral motion. Notable examples are described in [14]: an exploration policy either changes the heading direction randomly after detecting an obstacle or it follows the obstacle's boundaries, namely wall-following [13] and spiral [15]. Moreover, there is a subcategory of bug-inspired algorithms that uses ranging measurements for navigating towards a target point [16], [17], even in the presence of obstacles. In our work, we adapt such bio-inspired ranging-based algorithms to our flying nano robotic platform and, for the first time, we integrate them with a visual object detection pipeline, evaluating the effectiveness of the exploration policy within a search mission.

Object detection on constrained MCUs. Among the CNN-based visual object detection pipelines, single-shot detector (SSD) [18] is a popular approach, consisting of a feature extraction backbone and multiple convolutional detection heads. To reduce the computation and memory costs, shallow backbones have been proposed, e.g., Mobilenet, while preserving the detection scores on widely used dataset [19].

Focusing on the porting of object detectors on embedded systems, *Tran Quang Khoi et al.* [20] deploy an SSD network (SSDLite-MobileNetV2) with a similar architecture as the one used in this paper onto a RaspberryPi B3+ mounted on a standard-sized drone. The model execution reaches 0.71 FPS, lower than our obtained throughput. More importantly, this solution exceeds both the power (up to 3 W) and size constraints of the nano-drone system considered in this work. *Lamberti et al.* [21] presented an SSD algorithm for license plates detection on a static multi-core MCU node (GAP8), achieving up to 1.6 FPS with a power consumption of 117 mW. Starting from this previous work, we design and integrate an SSD-based object detection CNN onto a nano-drone capable of exploring the environment while detecting specific objects.

III. SYSTEM DESIGN

A. Robotic platform

Our robotic platform revolves around Bitcraze Crazyflie 2.1 quadrotor, a commercial off-the-shelf (COTS) nano-drone with a weight of 27 g and a diameter of 10 cm. This quadrotor features an STM32F405 MCU that runs state estimation and actuation controls. We extend this platform with three additional Bitcraze's COTS printed circuit boards (PCBs): the Flow deck, the Multi-ranger deck, and the AI-deck. The Flow deck provides optical flow and height measurements used to increase the state estimation reliability. The Multi-ranger deck features 5 single-beam VL53L1x ToF distance sensors mounted on the drone's top and sides, providing line of sight distance measurements within [0,4] meters range at 20 Hz.

Last, the AI-deck is a visual engine for nano-drones that extends the onboard processing capabilities of the Crazyflie. This PCB embeds the GAP8 SoC, a general-purpose multi-core parallel ultra-low-power (PULP) processor. GAP8 features a Cluster of 8 RISC-V cores that can run up to 175 MHz; it is designed for managing computation-intensive workloads with high energy efficiency, thanks to the parallel programming paradigm. The on-chip memory is organized hierarchically, having 64 kB of L1 scratchpad memory shared among the 8




Fig. 2. Bio-inspired Exploration policies: (A) *Pseudo-random*: inverts the direction by a random angle. (B) *Wall-following*: keeps a fixed distance from the perimeter. (C) *Spiral*: progressively increases the distance from the perimeter. (D) *Rotate-and-measure*: moves along the longest free-space direction.

cores of the CL, and 512 kB of L2 SRAM memory off-the-cluster. Furthermore, the AI-deck includes a low-power QVGA resolution grayscale camera (Himax HM01B0) and additional off-chip memories, including a 8 MB of HyperRAM and a 64 MB HyperFlash.

B. Multi-task integration.

We exploit our platform to explore an environment while searching for specific objects. To this aim, we separate the problem into two tasks: an *exploration task*, targeting navigation inside an unknown mission area while avoiding collisions, and an *object detection task* that runs the visual object detection SSD algorithm, as described in Sec. III-D. The exploration task runs a bio-inspired policy to determine the next set-point to feed the flight controller. As detailed in Sec. III-C, we study four different policies, i.e., the *exploration algorithms*, which take the ranging measurements of the ToF sensors as input to prevent collisions. On the other side, the performance for the *object detection task* is accounted by measuring the successful detections of target objects present in the mission area.

The *exploration* and the *object detection* tasks do not require interaction with each other; therefore, we map them on the two MCUs aboard our nano-drone for concurrent execution. More in detail, we use the two MCUs in a *host-accelerator* configuration. The STM32 is the *host* and takes care of the flight controller and the exploration policy based on a lightweight state machine. The exploration policies use the ToF measurements from the Multi-ranger deck to determine the next set-point of the drone in terms of forward speed and yaw rate. On the other hand, GAP8 is the *accelerator* in charge of executing the most computationally demanding CNN algorithm. This task includes camera acquisition and CNN processing and outputs the image-frame coordinates of the objects detected. As a result, the processing is done fully onboard, making the final closed-loop system completely autonomous, i.e., with no external communication or computation.

C. Exploration algorithms

Our four bio-inspired exploration algorithms are shown in Fig. 2. All the policies rely on the ranging measurements from three ToF sensors: the front, left, and right ones. We assess the performance of these policies by tracking the drone’s movements in a $6.5\text{ m} \times 5.5\text{ m}$ room equipped with a motion-capture system, tracking at 50 Hz. By discretizing the room’s area into cells of $0.5 \times 0.5\text{ m}$, we plot the heatmap (Fig. 3) representing the occupancy time of the drone over a 3 min flight. Our four exploration policies are:

A) Pseudo-random. This approach imitates the pseudo-random movement of biological creatures [6], [14]. The drone flies in a straight line as long as the ToF sensor does not identify obstacles within 1 m. When an obstacle is identified, the drone rotates to a random value, which is always greater than $\pm 90^\circ$ from the current heading (Fig. 2-A) to reduce the likelihood of detecting an obstacle previously identified.

B) Wall-following. This algorithm explores the perimeter of the room following its walls at a constant distance of 0.5 m, measured by the side ToF sensors (Fig. 2-B). The navigation stops when a front obstacle is detected and resumes after a $\pm 90^\circ$ turn towards an obstacle-free heading. By construction, this algorithm never explores the inner part of the testing room.

C) Spiral. The drone first explores the environment by performing several concentric perimetral paths (such as wall-following), each with an ever-increasing distance from the external walls (Fig. 2-C). Once the room’s center has been reached, the process is reversed, i.e., the consecutive perimetral explorations are performed with an ever-decreasing distance to the walls. The process starts over once the drone has reached its initial position. This spiral exploration starts with an initial distance from the walls of 0.5 m, and, for each lap, it is increased/decreased by the same amount.

D) Rotate-and-measure. This algorithm features two sequential phases (Fig. 2-D). In the first phase, the drone scans the environment by performing a 360° spin in place while measuring the front distance every 45° . In the second phase, the drone flies towards the most obstacle-free direction for a maximum of 2 m. This policy favors exploring the inner areas of our testing room while frequently neglecting its corners.

D. Object detection algorithm design

Our object detection pipeline is based on an SSD algorithm composed of a MobilenetV2 feature extractor [19], [22], pre-trained on the full OpenImages [23] dataset, and multiple detection heads, which is the ending part of the CNN that predicts

TABLE I
MEAN AVERAGE PRECISION (MAP) OF THE SSD CNNs TRAINED ON GOOGLE OPENIMAGES AND FINETUNED ON THE HIMAX DATASET.

Testing dataset	Fine-tuning	Format	SSD size		
			1×	0.75×	0.5×
OpenImages	no	float32	59%	47%	43%
Himax	no	float32	50%	41%	29%
Himax	yes	float32	55%	46%	43%
Himax	yes	int8	50%	48%	32%

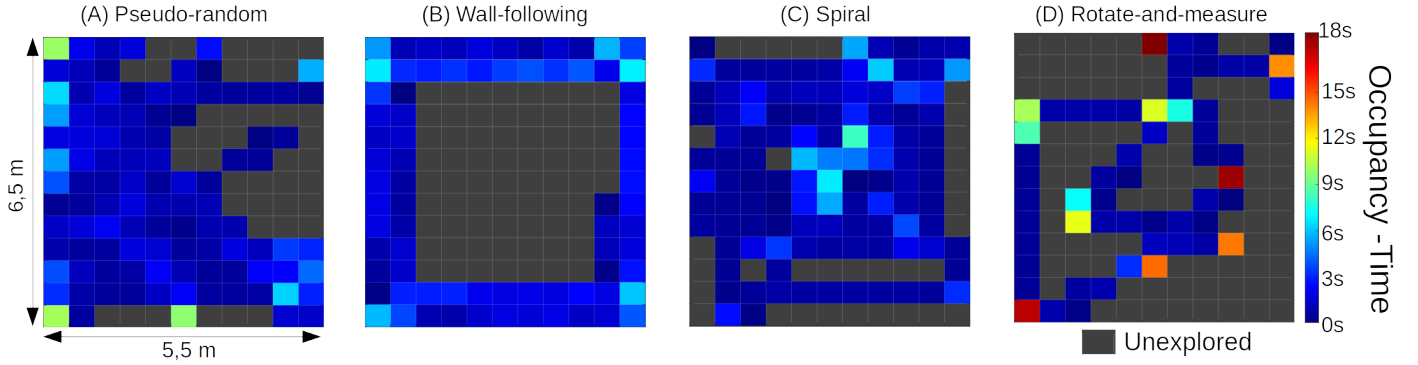


Fig. 3. Occupancy maps measured using the four exploration policies. During the experiments, the nano-drone flies at 0.5 m/s. The color of any cell, which represents a 0.5×0.5 m area, indicates the time spent by the nano-agent in that area (up to 18 sec). Black is used if the cell has not been explored.

the locations, categories, and confidence scores of objects at different resolution scales. To trade off latency/memory and detection accuracy, we deploy three different CNNs by varying the width multiplier α of the MobilenetV2 backbone. We refer to the different object detection algorithms as *SSD-MbV2- α* , where $\alpha = \{0.5, 0.75, 1.0\}$. The largest *SSD-MbV2-1.0* model features 4.67 M parameters and requires 534 MMAC operations while the *SSD-MbV2-0.75* and *SSD-MbV2-0.5* have, respectively, 2.68 M and 1.34 M parameters and require 358 MMAC and 193 MMAC operations.

We train our SSD models to detect two object categories, *bottles* and *tin cans*, using a subset of images from the OpenImages collection. Given the unbalanced training set, i.e., 1306 and 11306 images for tin cans and bottles, respectively, we balance the dataset by generating additional tin can images through horizontal translation (up to 10% of the image's width). The final split consists of 19142 images for training, 208 for validation, and 663 for testing, where the validation and testing portions match the original dataset splitting. Finally, to overcome the domain shift between the training and real-world data (Fig. 4), we add a fine-tuning phase on an additional dataset we collected and called *Himax Dataset*, which includes 321 training images and 279 testing images.

To deploy the object detector on GAP8, we quantize the CNN to 8-bit. We perform an additional fine-tuning step of quantization aware training (QAT) to minimize the mAP loss due to the 8-bit conversion. Then, we exploit the QAT routine included within the Tensorflow Object Detection API framework [22] using symmetric quantization ranges because the software primitives of the target hardware rely on symmetric integer ranges. Lastly, Greenwaves's *GAPflow* toolset¹ is used to produce the C code of the object detection task, constraining the L2 buffer size to 250 kB.

IV. RESULTS

A. Object detection evaluation

The SSD models are trained on OpenImages for 1200 epochs using the RMSProp optimizer. We set a learning rate of $8 \cdot 10^{-4}$ with an exponential decay of 0.95 every 24 epochs and a batch

size of 24. The OpenImages frames are resized to 320×240 pixels to match the resolution of the drone onboard camera. During training, the images are extended with photometric augmentations, such as flipping, brightness adjustment, random cropping, and grayscale conversion, individually applied with a probability of 0.5. Before deployment, the SSD models are fine-tuned, also applying QAT, on the Himax dataset for 100 epochs, with a learning rate of 10^{-4} retaining the same exponential decay of 0.95 every 10 epochs.

Tab. I reports the mAP scores (as defined for the COCO dataset [24]) of the multi-sized SSD models on both the OpenImages and our Himax testing datasets. The largest model achieves an mAP of 59% when trained and tested on images from the OpenImages dataset, up to 16% vs. the smallest model. These scores reduce by 9%, 6%, and 14% for the three models when tested on the Himax dataset. We argue this effect to be due to the onboard camera image quality being lower than the web-retrieved images. After fine-tuning on the Himax training set, the mAP score improves up to 55% for the best model, bridging the accuracy gap for the mAP score on the OpenImages test set. Given the best scores of 50% and 48% after 8-bit quantization, we consider both *SSD-MbV2-1* and *SSD-MbV2-0.75* models for the in-field evaluation.

Tab. II reports the performance of the SSD CNN running on the GAP8 SoC. We set the operating voltage at 1.2 V and a clock frequency of 160 MHz for the multicore cluster, while the peripheral clock is set to 250 MHz. The largest models (*SSD-MbV2-1.0*) can process up to 1.6 frame/s, with a computational efficiency of 5.3 *MAC/clock cycles*. Instead, the smaller *SSD-MbV2-0.75* and *SSD-MbV2-0.5* result, respectively, $1.6 \times$ and



Fig. 4. Google OpenImages sample (left) vs. onboard sensor image (right).

¹<https://greenwaves-technologies.com/gapflow/>

TABLE II
SSD CNNs' ONBOARD PERFORMANCE.

SSD	Parameters	Operations	Efficiency	Throughput
1×	4.7M	534 MMAC	5.3 MAC/cycles	1.6 FPS
0.75×	2.7M	358 MMAC	5.9 MAC/cycles	2.3 FPS
0.5×	1.2M	193 MMAC	5.3 MAC/cycles	4.3 FPS

2.7× faster than the larger model. The power consumption of the AI-deck reaches a peak of 143.5 mW when running the SSD-MbV2-0.75 model, where the inference task shows the highest compute efficiency, maximizing the memories' bandwidth and processing logic utilization. Conversely, the power consumption decreases to 134.5 mW if running the most accurate SSD-MbV2-1.0 model.

B. Exploration policies evaluation

To assess the performance of our four exploration policies, we calculate the *coverage area* (%) in the testing room described in III-C by calculating the ratio between the visited cells w.r.t. their total (i.e., 143 cells). We mark a cell "visited" when the drone's center of mass falls into it. We evaluate each exploration policy with three average flight speeds (i.e., 0.1 m/s, 0.5 m/s, and 1 m/s), obtaining 12 test configurations. For every configuration, we perform five runs of 3 min each, accounting for a total of 60 runs (3 h flight time).

Fig. 5 reports the coverage area of each configuration (policy and speed) averaged on the five runs. The *pseudo-random* and *spiral* policies are those that mainly benefit from higher speeds, passing from 35% coverage area (at 0.1 m/s) to 74% and 82% (at 0.5 m/s), and finally to 80% and 83% (at 1 m/s), respectively. Instead, the *wall-following* and *rotate-and-measure* exploration policies show a slight improvement passing from a mean flight speed of 0.1 m/s to 0.5 m/s, as much as +17% and +14%, respectively. Contrary to the previous policies, the highest flight speed does not improve the coverage area, i.e., +2% and -1%, respectively. The *wall-following* policy has a limited performance since the drone only explores the room's perimeter (see Fig. 3-B). Similarly, the *rotate-and-measure* policy spends the vast majority of the 3 min flight spinning in place and focusing on the center of the room (see Fig. 3-D).




Fig. 5. Average coverage area (in %) for each exploration policy, varying its mean flight speed (i.e., 0.1 m/s, 0.5 m/s, and 1 m/s).

TABLE III
AVERAGE DETECTION RATE: 6 OBJECTS, 5 RUNS OF 3 MINUTES EACH.

SSD	Flight speed [m/s]	Detection rate			
		Pseudo random	Wall following	Spiral	Rotate and measure
1.0×	0.1	27%	63%	67%	53%
	0.5	90%	50%	73%	53%
	1	83%	53%	70%	47%
0.75×	0.1	27%	50%	33%	47%
	0.5	80%	37%	43%	50%
	1	37%	27%	43%	33%

C. In-field closed-loop system evaluation

To evaluate the closed-loop in-field performance, we measure the detection rate of each policy/CNN by placing three bottles and three tin cans in the room, relying on the same testing configuration of Sec. IV-B: 6.5×5.5 m testing room, 3 min flights. One bottle and one tin can are close to the center, while the other four are near the corners. The final detection rate depends on *i.*) detector's precision and throughput and *ii.*) covered area (partially depending on the flight speed). In Tab. III, we evaluate the best two SSD CNNs (see Sec. IV-A) with all four exploration policies and three flight speeds. The bigger *SSD-MbV2-1.0* consistently achieves, for any configuration, an equal or greater detection rate than the medium 0.75× model, suggesting that, in our setup, the detector's accuracy is more important than its throughput (*SSD-MbV2-0.75* has higher throughput than *SSD-MbV2-1.0*, but lower mAP).

Focusing on the *SSD-MbV2-1.0*, the peak performances are achieved by the *pseudo-random* and *spiral* policies, which show the highest detection rate at the intermediate flight speed of 0.5 m/s. In contrast, in Sec. IV-B, the best coverage area was obtained by the highest flight speed, which suggests 1 m/s being too high for the limited inference rate of the bigger 1× SSD (i.e., 1.6 frame/s). Then, the *wall-following* and *rotate-and-measure* policies show the lower detection rates, 63% and 53% at most, respectively. This behavior is the consequence of the limited coverage area of the policies, where the former misses all the objects placed in the center of the room while the latter marginally explores the perimeter.

In Fig. 6, we focus our analysis on the best performing model: *pseudo-random* exploration policy with the nano-drone




Fig. 6. Coverage area of the *random* policy over 5 runs (mean and variance). The detection time of the 6 target objects (blue dots) refers to one single run.

TABLE IV
POWER BREAKDOWN OF THE ROBOTIC PLATFORM.

	Motors	CF elect.	AI-deck	Multi-ranger	Total
Power [W]	7.32	0.277	0.134	0.286	8.02
Percentage	91.31%	3.45%	1.67%	3.57%	100%

flying at 0.5 m/s and running the *SSD-MbV2-1.0*. The yellow line shows the coverage area over five 3 min flights, which scores a maximum average of 72% coverage, with a variance of 21%. The blue dots indicate the time each object is recognized for a run achieving 100% detection rate in 154s.

Lastly, Tab. IV shows the power breakdown of the nano-drone platform, measured by profiling each component individually with the Power Profiler Kit 2 by Nordic Semiconductor. The power cost of the AI-deck, running the biggest *SSD-MbV2-1.0*, accounts for 1.67% of the total power, which is dominated by the motors (91.31%). The other electronics components, i.e., Crazyflie's MCU and the ToF sensors, require the remaining 7.03% of the total power.

V. CONCLUSIONS

This work presents the first multi-objective autonomous nano-drone, pursuing the exploration of an unknown environment while avoiding collision and searching objects with a vision-based CNN detector. Our work compares four bio-inspired exploration policies and three versions of object detectors running independently on two resource-limited MCUs aboard the nano-drone. The best configuration reaches a final detection rate of 90%, exploiting *i.)* a pseudo-random policy for exploration, *ii.)* the largest object detection model, and *iii.)* a mean flight speed of 0.5 m/s. This result shows how the higher detection rate can be reached by trading off the detection capabilities of the CNN, its throughput, and the mean flight speed of our nano-drone.

ACKNOWLEDGMENTS

We thank the Center for Research on Complex Automated Systems and Aurora Di Giampietro for their support. This work has been partially funded by the Autonomous Robotics Research Center of the UAE Technology Innovation Institute.

REFERENCES

- [1] D. Palossi, A. Gomez, S. Draskovic, K. Keller, L. Benini, and L. Thiele, "Self-sustainability in nano unmanned aerial vehicles: A blimp case study," in *Proceedings of the computing frontiers conference*, 2017.
- [2] D. Palossi, N. Zimmerman, A. Burrello, F. Conti, H. Müller, L. M. Gambardella, L. Benini, A. Giusti, and J. Guzzi, "Fully onboard ai-powered human-drone pose estimation on ultra-low power autonomous flying nano-uavs," *IEEE Internet of Things Journal*, pp. 1–1, 2021.
- [3] A. Birk, B. Wiggerich, H. Bülow, M. Pflingstorn, and S. Schwertfeger, "Safety, Security, and Rescue Missions with an Unmanned Aerial Vehicle (UAV)," *Journal of Intelligent & Robotic Systems*, vol. 64, no. 1, pp. 57–76, Oct. 2011.
- [4] M. V. Srinivasan, "Honeybees as a model for the study of visually guided flight, navigation, and biologically inspired robotics," *Physiological Reviews*, vol. 91, no. 2, pp. 413–460, 2011, PMID: 21527730.
- [5] M. Giurfa and R. Menzel, "Insect visual perception: complex abilities of simple nervous systems," *Current Opinion in Neurobiology*, vol. 7, no. 4, pp. 505–513, 1997.
- [6] B. Webb, "The internal maps of insects," *Journal of Experimental Biology*, vol. 222, 2019.
- [7] R. J. Wood, B. Finio, M. Karpelson, K. Ma, N. O. Pérez-Arancibia, P. S. Sreetharan, H. Tanaka, and J. P. Whitney, "Progress on "Pico" air vehicles," in *Robotics Research : The 15th International Symposium ISRR*, H. I. Christensen and O. Khatib, Eds. Cham: Springer International Publishing, 2017, pp. 3–19.
- [8] I. Dinstein, U. Hasson, N. Rubin, and D. J. Heeger, "Brain areas selective for both observed and executed movements," *Journal of neurophysiology*, vol. 98, no. 3, pp. 1415–1427, 2007.
- [9] L. Lamberti, V. Niculescu, M. Barciš, L. Bellone, E. Natalizio, L. Benini, and D. Palossi, "Tiny-pulp-dronets: Squeezing neural networks for faster and lighter inference on multi-tasking autonomous nano-drones," in *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 2022, pp. 287–290.
- [10] R. J. Bouwmeester, F. Paredes-Vallés, and G. C. H. E. de Croon, "Nanoflownet: Real-time dense optical flow on a nano quadcopter," *arXiv preprint arXiv:2209.06918*, 2022.
- [11] D. Palossi, F. Tombari, S. Salti, M. Ruggiero, L. Stefano, and L. Benini, "Gpu-shot: Parallel optimization for real-time 3d local description," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2013, pp. 584–591.
- [12] B. Bodin, H. Wagstaff, S. Saecdi, L. Nardi, E. Vespa, J. Mawer, A. Nisbet, M. Lujan, S. Furber, A. J. Davison, P. H. J. Kelly, and M. F. P. O'Boyle, "Slambench2: Multi-objective head-to-head benchmarking for visual slam," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 3637–3644.
- [13] K. McGuire, G. de Croon, and K. Tuyls, "A comparative study of bug algorithms for robot navigation," *Robotics and Autonomous Systems*, vol. 121, p. 103261, 2019.
- [14] Q.-l. Xu, "Randombug: Novel path planning algorithm in unknown environment," *The Open Electrical & Electronic Engineering Journal*, vol. 8, no. 1, 2014.
- [15] E. Magid and E. Rivlin, "Cautiousbug: a competitive algorithm for sensory-based robot navigation," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 3, 2004, pp. 2757–2762 vol.3.
- [16] Q.-L. Xu and G.-Y. Tang, "Vectorization path planning for autonomous mobile agent in unknown environment," *Neural Computing and Applications*, vol. 23, no. 7, pp. 2129–2135, 2013.
- [17] I. Kamon, E. Rivlin, and E. Rimon, "A new range-sensor based globally convergent navigation algorithm for mobile robots," in *Proceedings of IEEE International Conference on Robotics and Automation*, vol. 1. IEEE, 1996, pp. 429–435.
- [18] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 21–37.
- [19] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *The IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [20] T. Q. Khoi, N. A. Quang, and N. K. Hieu, "Object detection for drones on raspberry pi potentials and challenges," *IOP Conference Series: Materials Science and Engineering*, vol. 1109, no. 1, p. 012033, mar 2021.
- [21] L. Lamberti, M. Rusci, M. Fariselli, F. Paci, and L. Benini, "Low-power license plate detection and recognition on a risc-v multi-core mcu-based vision system," in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2021, pp. 1–5.
- [22] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, "Speed/accuracy trade-offs for modern convolutional object detectors," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [23] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov *et al.*, "The open images dataset v4," *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1956–1981, 2020.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer International Publishing, 2014, pp. 740–755.