Tiny Robot Learning (tinyRL)

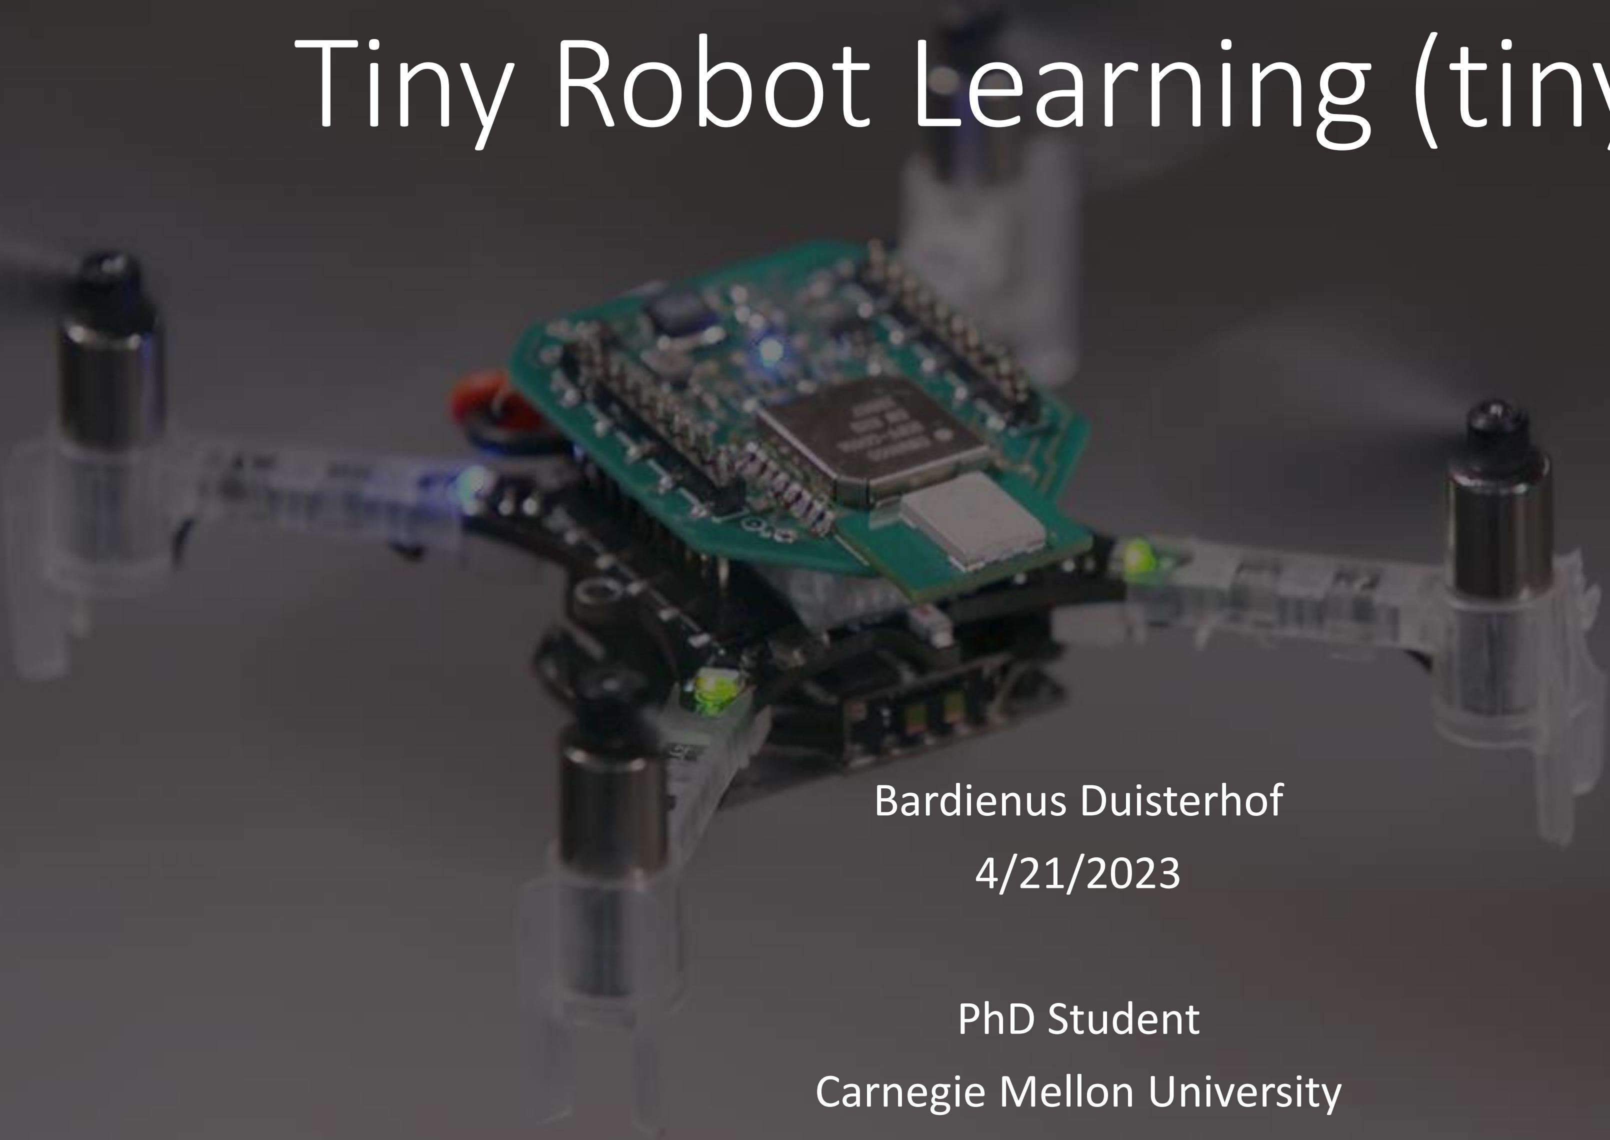

Bardienus Duisterhof

4/21/2023

PhD Student

Carnegie Mellon University

The Robotics Institute

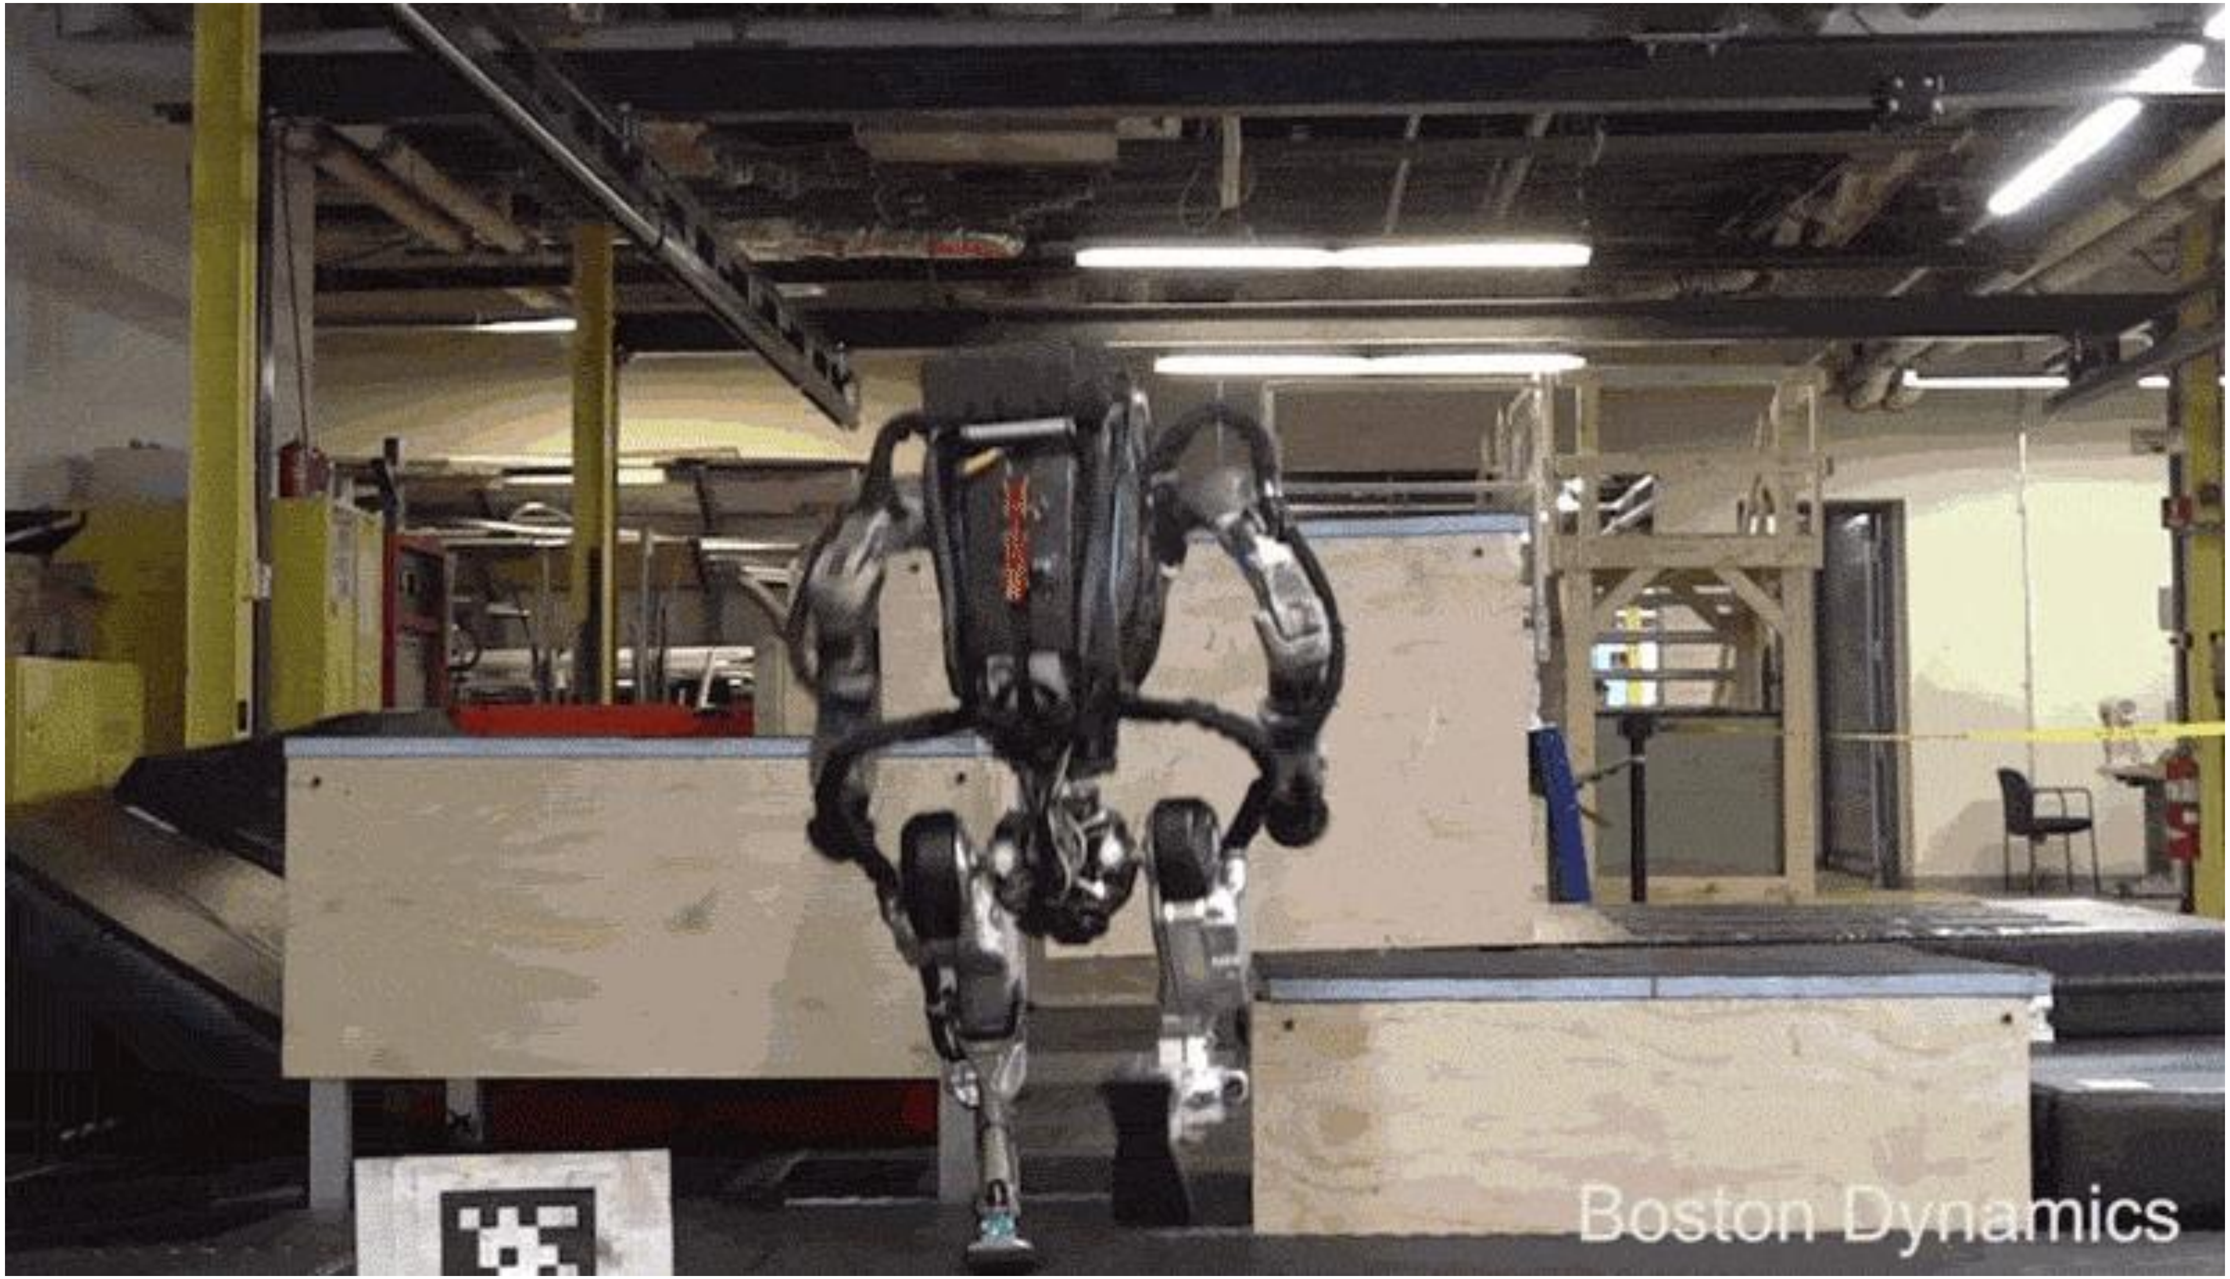
Boston Dynamics


DARPATV

Challenges in Robotics

Autonomous navigation

Resource constrained

- SWaP
  - Size
  - Weight
  - Power

Smaller robots are

- Safer
- Cheaper
- Even more constrained!

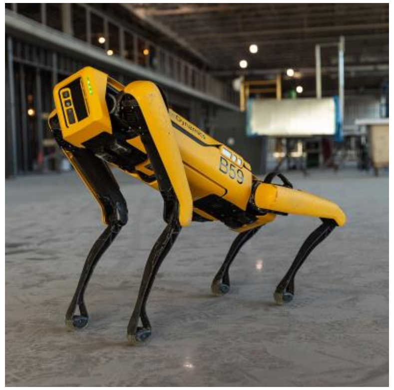

Challenges in Robotics

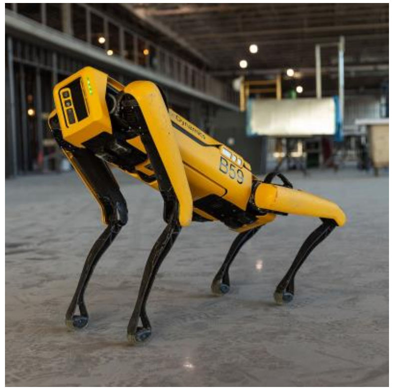

Autonomous navigation

Resource constrained

- SWaP
  - Size
  - Weight
  - Power

Smaller robots are

- Safer
- Cheaper
- Even more constrained!

Harvard John A. Paulson School of Engineering and Applied Sciences

Tiny Robot Learning (tinyRL) for Source Seeking on a Nano Quadcopter

Bardienus P. Duisterhof[1,2], Srivatsan Krishnan[1], Jonathan J. Cruz[1], Colby R. Banbury[1], William Fu[1], Aleksandra Faust[3], Guido C. H. E. de Croon[2], Vijay Janapa Reddi[1]

1 Harvard University 2 Delft University of Technology 3 Google Brain

Motivation

- Insects use simple rules to navigate using little sensory and computational resources.
- Inspired by nature, finite state machines have proven simple rules can solve complex navigation problems on robots.*
- Application-specific FSM’s can be tedious to design and may lack generalization.

*McGuire, K. N., Wagter, C. D., Tuyls, K., Kappen, H. J. & de Croon, G. C. H. E. Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment. *Sci. Robot.* 4, eaaw9710 (2019).

Motivation

- Tiny machine learning (**tinyML**): ML applications on low-power, cheap, commodity hardware.
- Strong focus on low power consumption for always-on use-cases on battery operated devices.
- **How can tinyML impact robotics?**

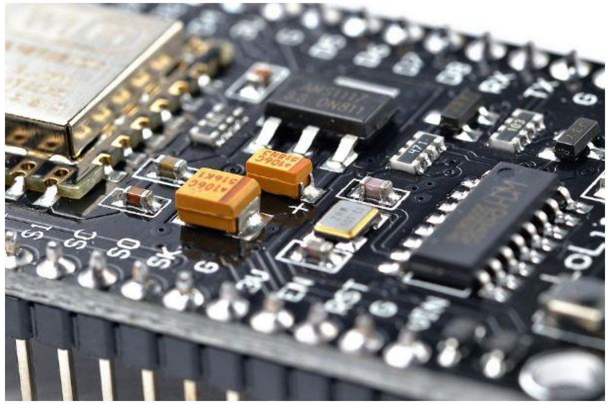

Tiny Robot Learning (tinyRL)

- Inspired by tinyML, in this work, we introduce **tiny robot learning** (**tinyRL**).
- We deploy a tiny ML model onboard a highly constrained nano quadcopter for source seeking.
- Our methodology achieves robust and efficient source seeking, **running a deep-RL model onboard a nano quadcopter.**

Source Seeking

- Autonomous machines locating light, gas, or radiations sources.
- An important task in search and rescue and inspection.
- We imagine small, agile and inexpensive aerial robots for source seeking.
- We use light seeking as an application to show how a tiny deep-RL policy can be a viable alternative to simple finite state machines.

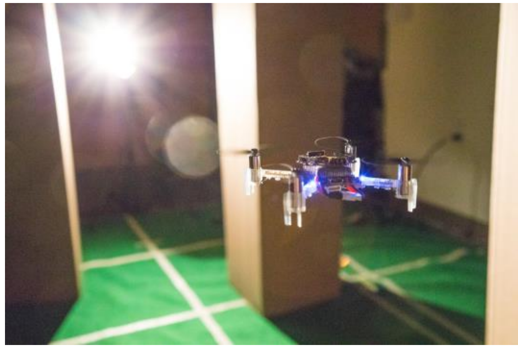

Harvard John A. Paulson School of Engineering and Applied Sciences

Method

- System Design
- Simulation Environment
- POMDP Setup
- Inference Implementation

System design

BitCraze CrazyFlie 2.1

- ARM Cortex-M4
- CPU: 1-core & 168 MHz
- RAM: 196 kB
- Storage: 1MB
- Available RAM: 33 kB
- Weight: 33 grams

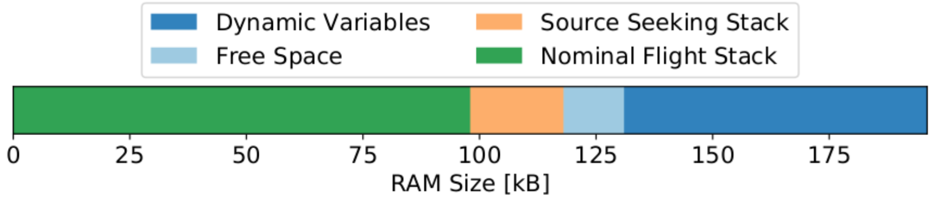

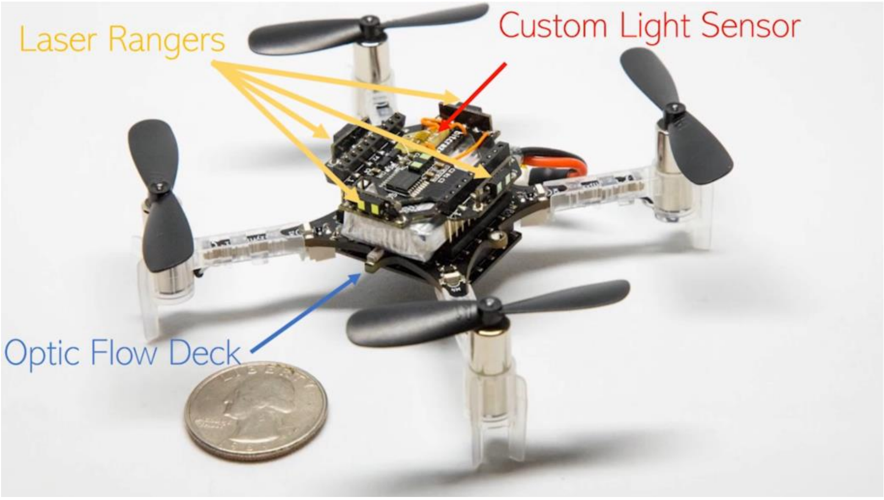

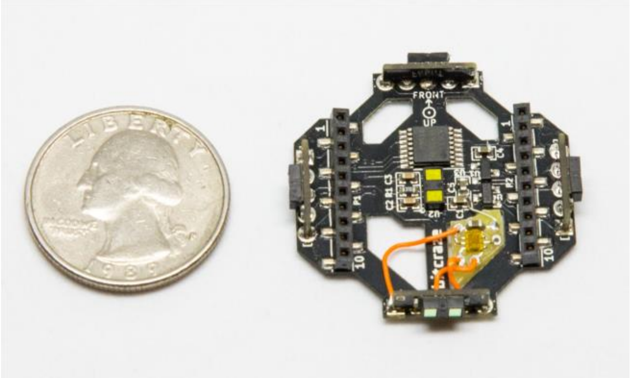

Simulation Environment

- We use the Air Learning platform*, which couples with Microsoft AirSim to provide a deep reinforcement learning back end.
- Source position and obstacle positions are randomized.
- A light source is modeled as a Gaussian distribution, based on data captured in our flight room.

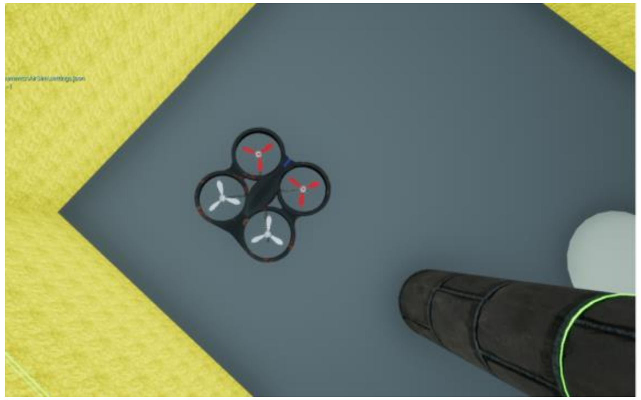

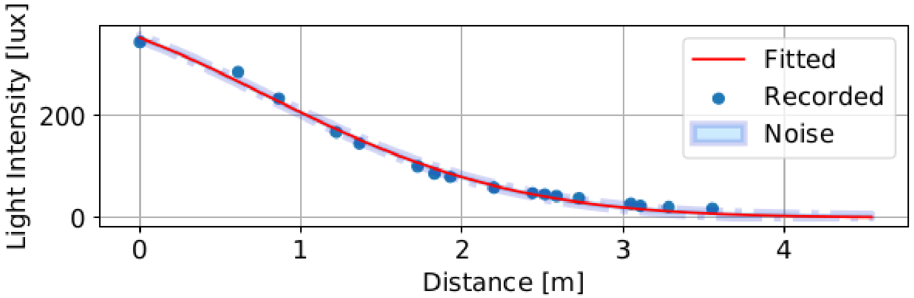

*S. Krishnan, B. Boroujerdian, W. Fu, A. Faust, and V. J. Reddi, "Air learning: An AI research platform for algorithm-hardware benchmarking of autonomous aerial robots", Machine Learning, Special Issue on RL for Real Life, in press, 2021

POMDP Setup

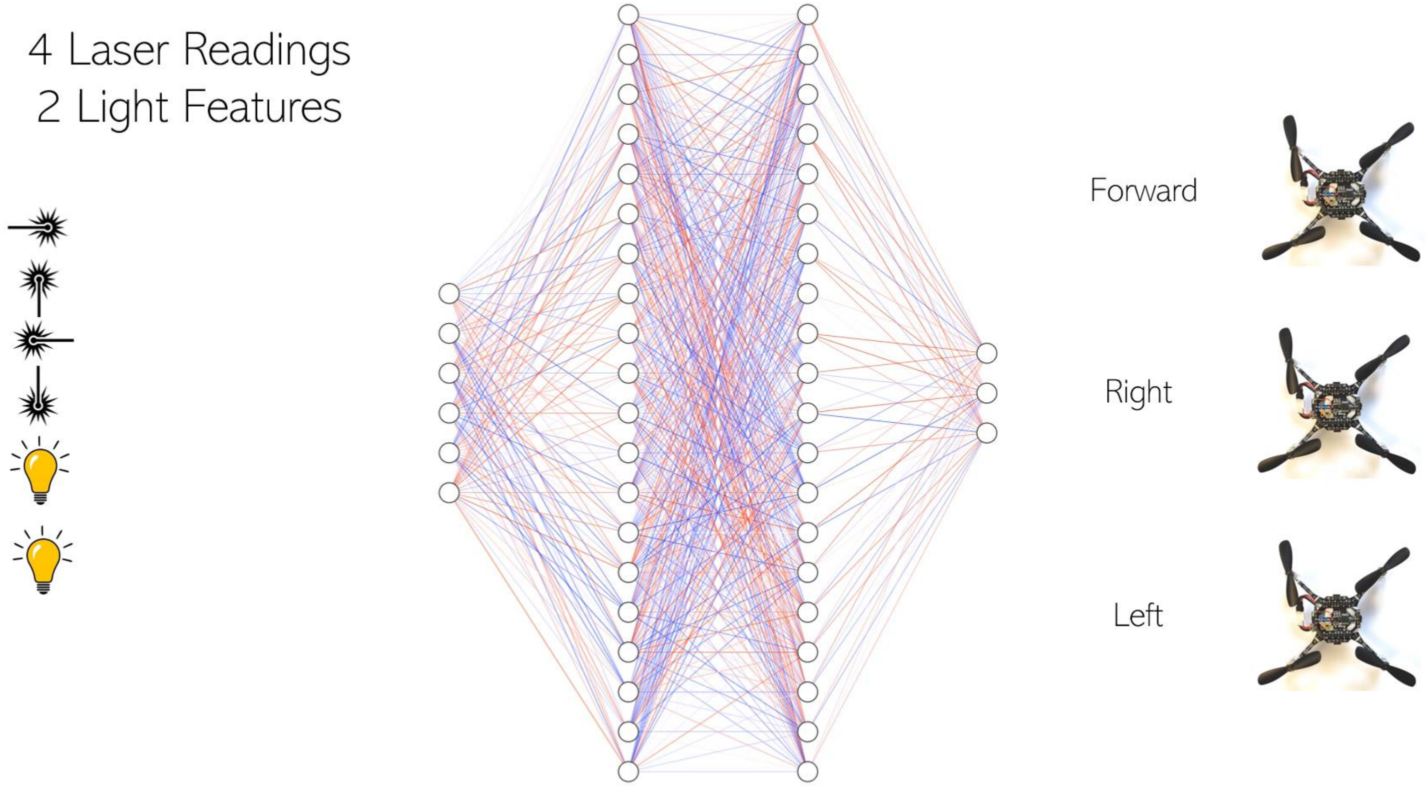

Inference Implementation

- Obstacle avoidance requires low-latency inference.
- Libraries considered:
    - TensorFlow Lite, not fast enough.
    - uTensor, ran out of memory.
- Too much overhead meant we developed a custom lightweight C inference library.
- Result: capable of inference at up to 100Hz, higher than the sensor polling rate.

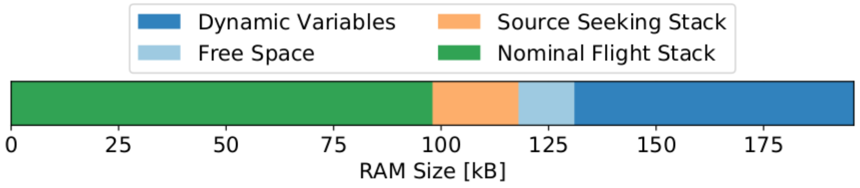

Baseline Comparison in Simulation

Model description	Success
Our deep RL algorithm	96%
FSM baseline	84%
Random actions	30%

- Two baselines in simulation:
  - Random actions
  - FSM baseline geared towards exploration
- To the best of our knowledge, no publicly available algorithm exists that can work with our sensor inputs to avoid obstacles and seek a light source.
- The deep-RL model outperforms the FSM baseline in all metrics.

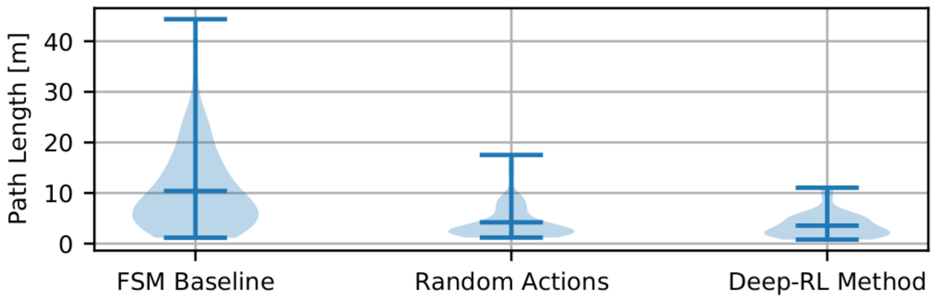

Harvard John A. Paulson School of Engineering and Applied Sciences

Flight Tests

- The deep-RL model reaches **a 94% success rate.**
- The FSM Baseline reaches a 75% success rate.
- Between obstacle densities, our policy found the source 55%-70% faster than the baseline.
- The results show that **our policy generalizes far beyond what was presented in simulation.**

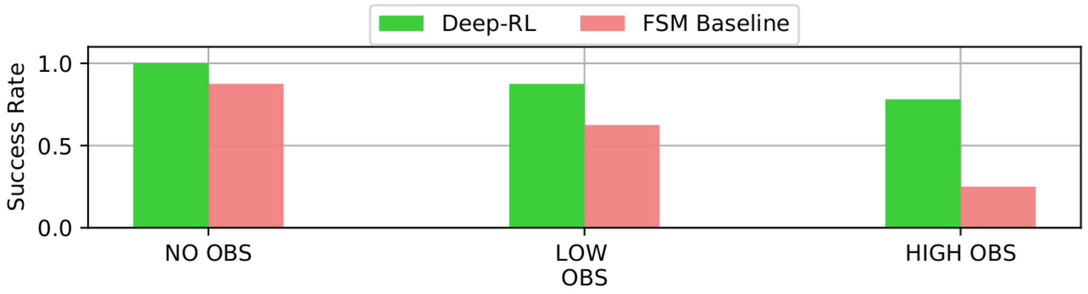

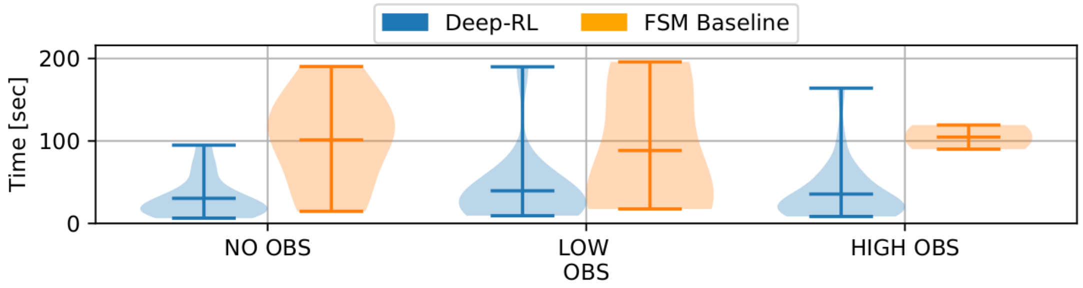

Harvard John A. Paulson School of Engineering and Applied Sciences

Demo

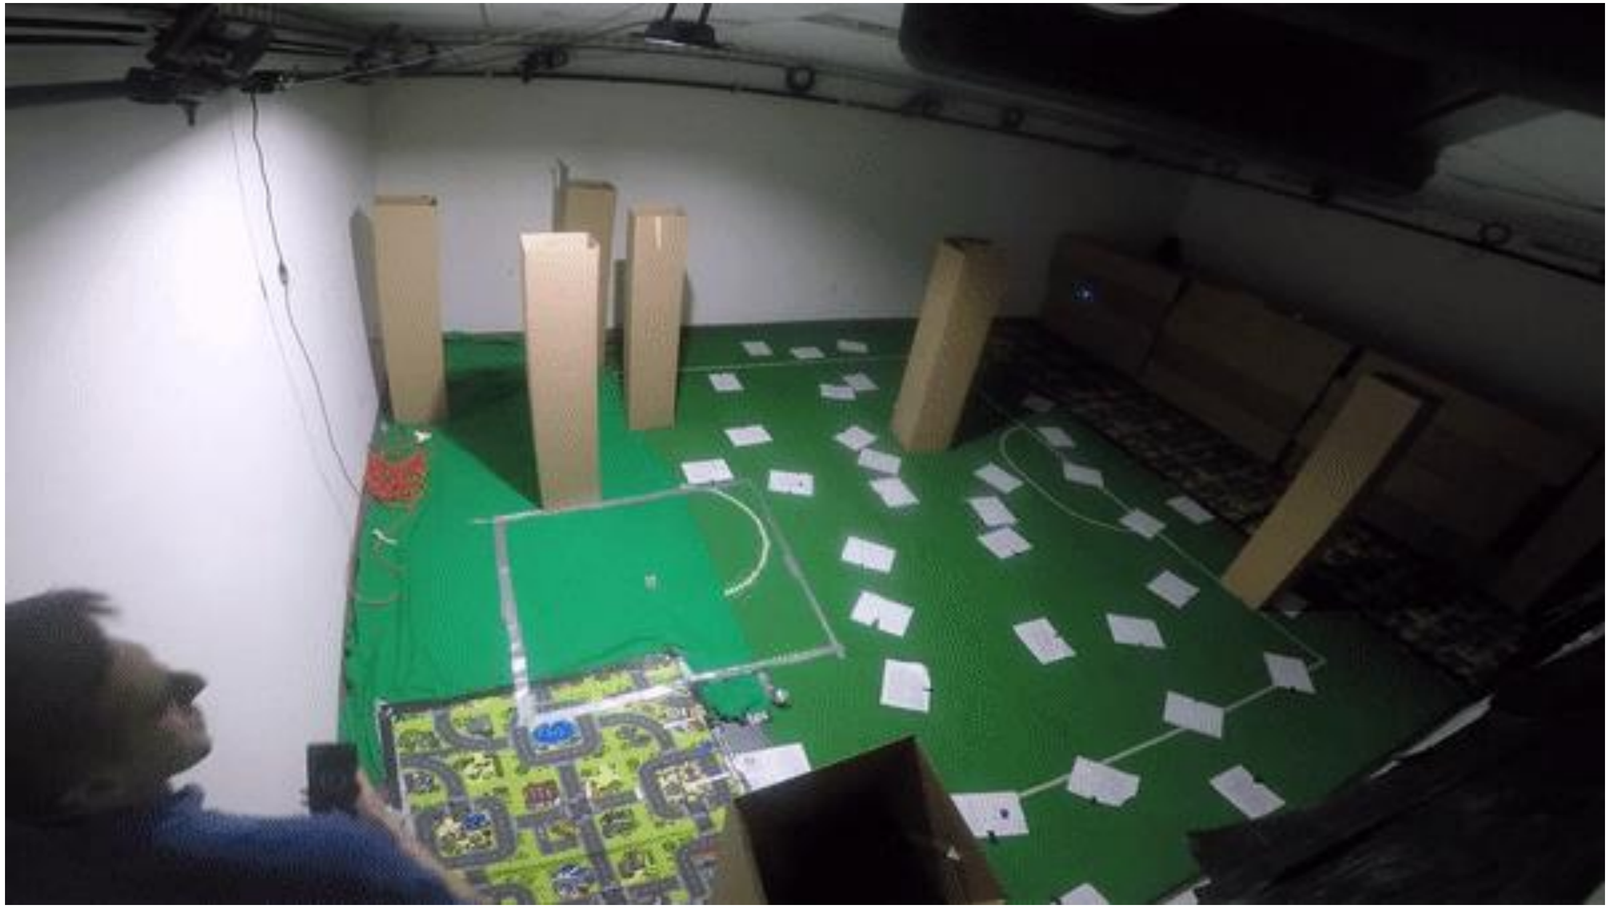

Harvard John A. Paulson School of Engineering and Applied Sciences

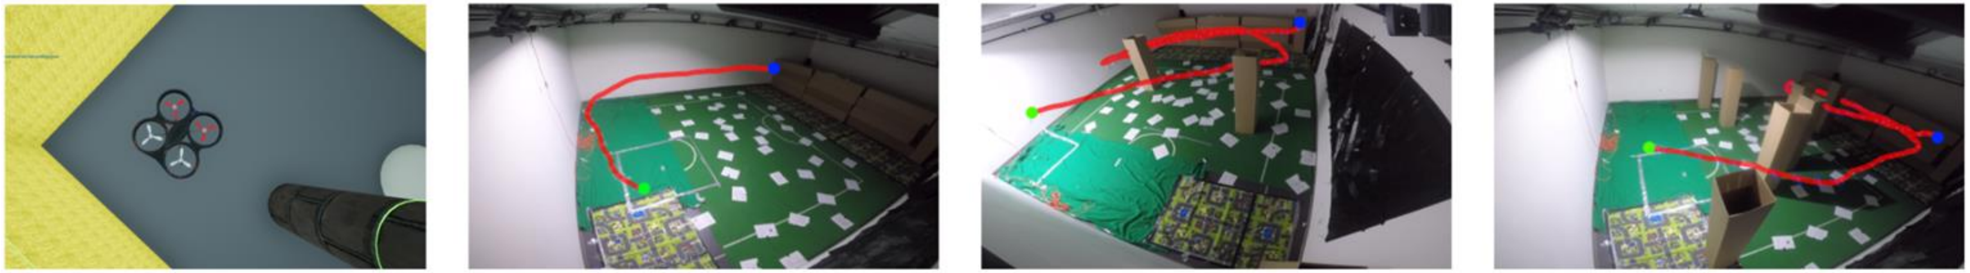

Edge Computing Lab
https://edge.seas.harvard.edu/

ArXiv
https://arxiv.org/abs/1909.11236

Acknowledgements

We thank Robert Wood, Jim MacArthur, Hassan Khawaja and Sharad Chitlangia (from Harvard) for their help in enabling this project

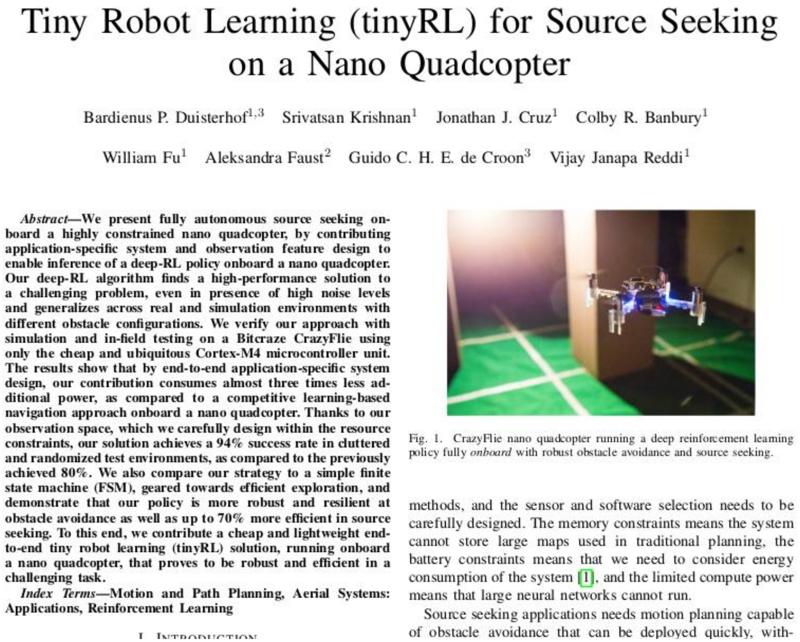

Tiny Robot Learning (tinyRL) for Source Seeking on a Nano Quadcopter

Bardienus P. Duisterhof[1,3] Srivatsan Krishnan[1] Jonathan J. Cruz[1] Colby R. Banbury[1]
William Fu[1] Aleksandra Faust[2] Guido C. H. E. de Croon[3] Vijay Janapa Reddi[1]

Abstract—We present fully autonomous source seeking onboard a highly constrained nano quadcopter, by contributing application-specific system and observation feature design to enable inference of a deep-RL policy onboard a nano quadcopter. Our deep-RL algorithm finds a high-performance solution to a challenging problem, even in presence of high noise levels and generalizes across real and simulation environments with different obstacle configurations. We verify our approach with simulation and in-field testing on a Bitcraze CrazyFlie using only the cheap and ubiquitous Cortex-M4 microcontroller unit. The results show that by end-to-end application-specific system design, our contribution consumes almost three times less additional power, as compared to a competitive learning-based navigation approach onboard a nano quadcopter. Thanks to our observation space, which we carefully design within the resource constraints, our solution achieves a 94% success rate in cluttered and randomized test environments, as compared to the previously achieved 80%. We also compare our strategy to a simple finite state machine (FSM), geared towards efficient exploration, and demonstrate that our policy is more robust and resilient at obstacle avoidance as well as up to 70% more efficient in source seeking. To this end, we contribute a cheap and lightweight end-to-end tiny robot learning (tinyRL) solution, running onboard a nano quadcopter, that proves to be robust and efficient in a challenging task.

Index Terms—Motion and Path Planning, Aerial Systems: Applications, Reinforcement Learning

I. INTRODUCTION

Fig. 1. CrazyFlie nano quadcopter running a deep reinforcement learning policy fully *onboard* with robust obstacle avoidance and source seeking.

methods, and the sensor and software selection needs to be carefully designed. The memory constraints means the system cannot store large maps used in traditional planning, the battery constraints means that we need to consider energy consumption of the system [1], and the limited compute power means that large neural networks cannot run.

Source seeking applications needs motion planning capable of obstacle avoidance that can be deployed quickly, with-

Bardienus P. Duisterhof, Srivatsan Krishnan, Jonathan J. Cruz, Colby R. Banbury, William Fu, Aleksandra Faust, Guido C. H. E. de Croon, Vijay Janapa Reddi, Tiny Robot Learning (tinyRL) for Source Seeking on a Nano Quadcopter, ICRA 2021

Harvard John A. Paulson School of Engineering and Applied Sciences

Gas Source Localization (GSL)

Use cases:

- Dangerous leaks with many casualties.
- Gas leaks in chemical plants → $$$
- Chemical attack

In this work

A **swarm** of fully autonomous and collaborative **gas-seeking nano quadcopters.**

A **novel bug algorithm for GSL**, using extremely little resources and **evolved parameters**.

A **pipeline** for end-to-end environment generation and gas dispersion modelling (GDM).

Arena size:
0.8 x 2.0 m
(wind
tunnel)

Related Work

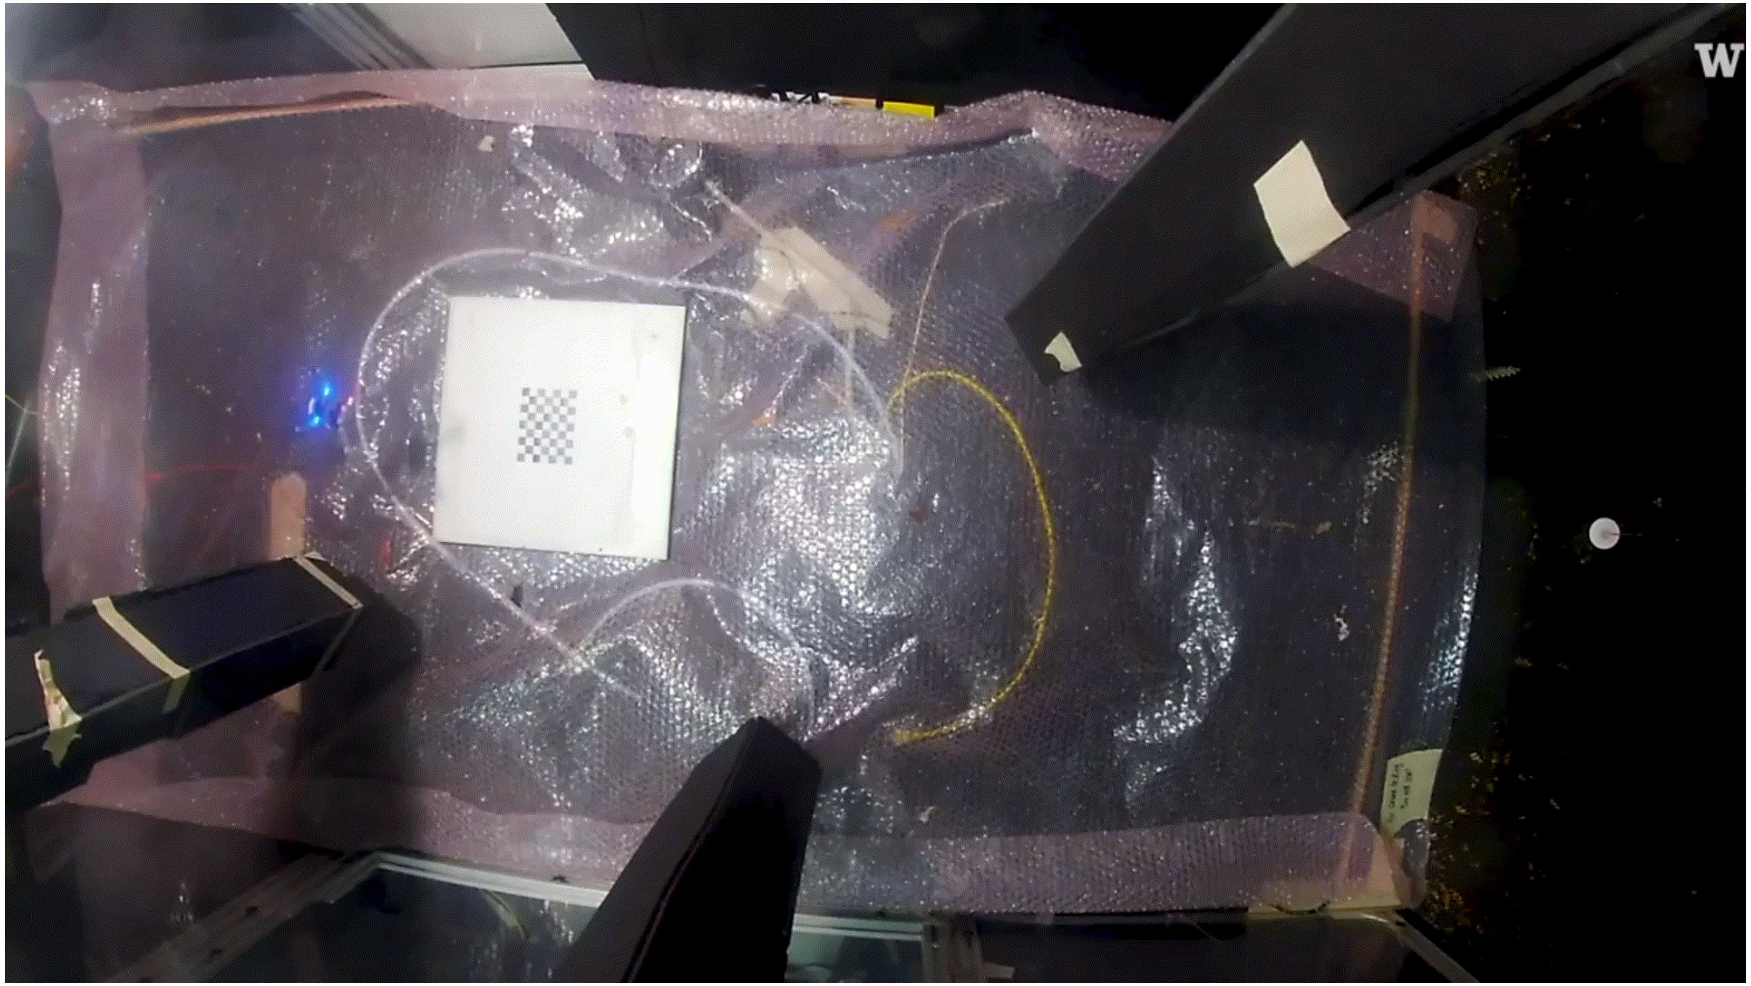

Melanie Joyce Anderson, Joseph Garret Sullivan, Timothy Horiuchi, Sawyer Buckminster Fuller, and Thomas L Daniel. A bio-hybrid odor-guided autonomous palmsized air vehicle. Bioinspiration Biomimetics, 2020.

TU Delft | Harvard John A. Paulson School of Engineering and Applied Sciences | Universitat de Barcelona

Methodology

System Design

Simulation Pipeline

Algorithm Design

System Design

Requirements:

- Obstacle avoidance
- Odometry
- Gas sensing
- Relative ranging
- Communication

Payload:

- Flow deck
- Multiranger deck
- Custom gas/UWB PCB

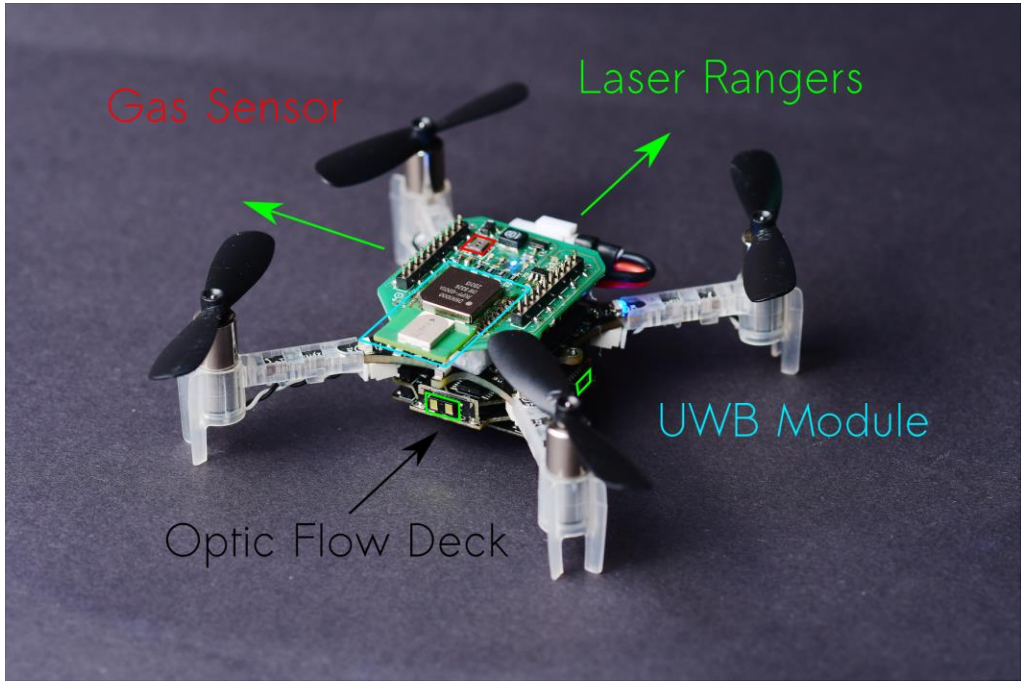

Weight: 37.5g

TU Delft Harvard John A. Paulson School of Engineering and Applied Sciences Universitat de Barcelona

Simulation pipeline

AutoGDM

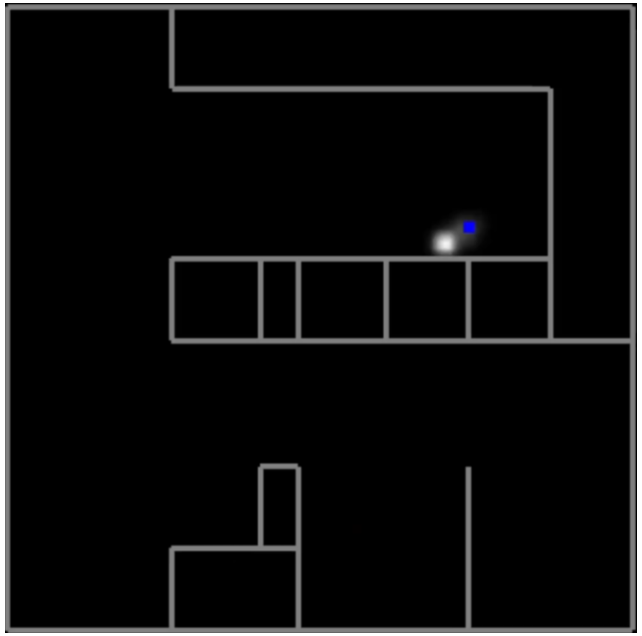

Swarmulator

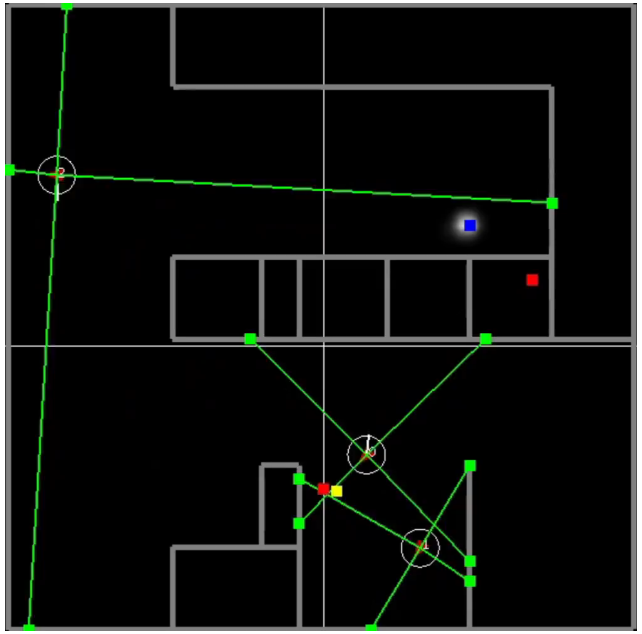

TU Delft

Harvard John A. Paulson School of Engineering and Applied Sciences

UNIVERSITAT DE BARCELONA

AutoGDM – Gas Dispersion Modeling

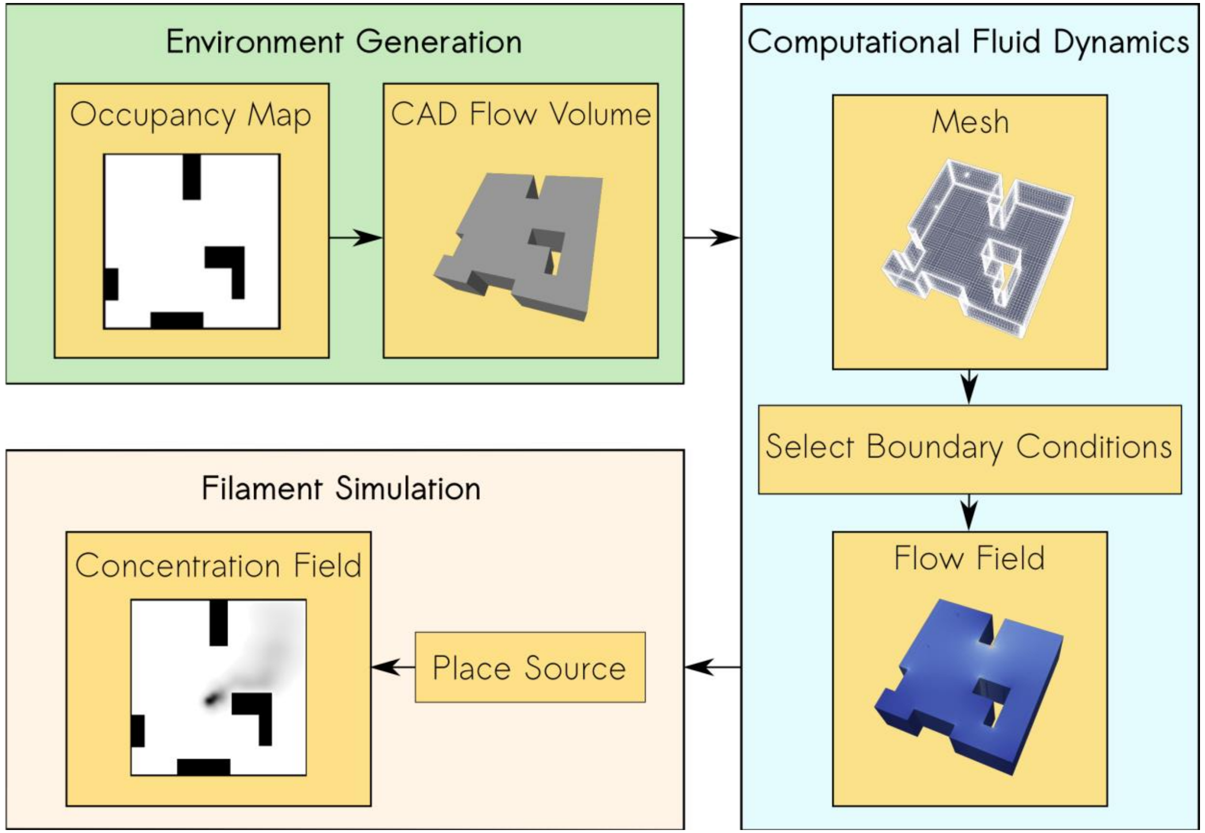

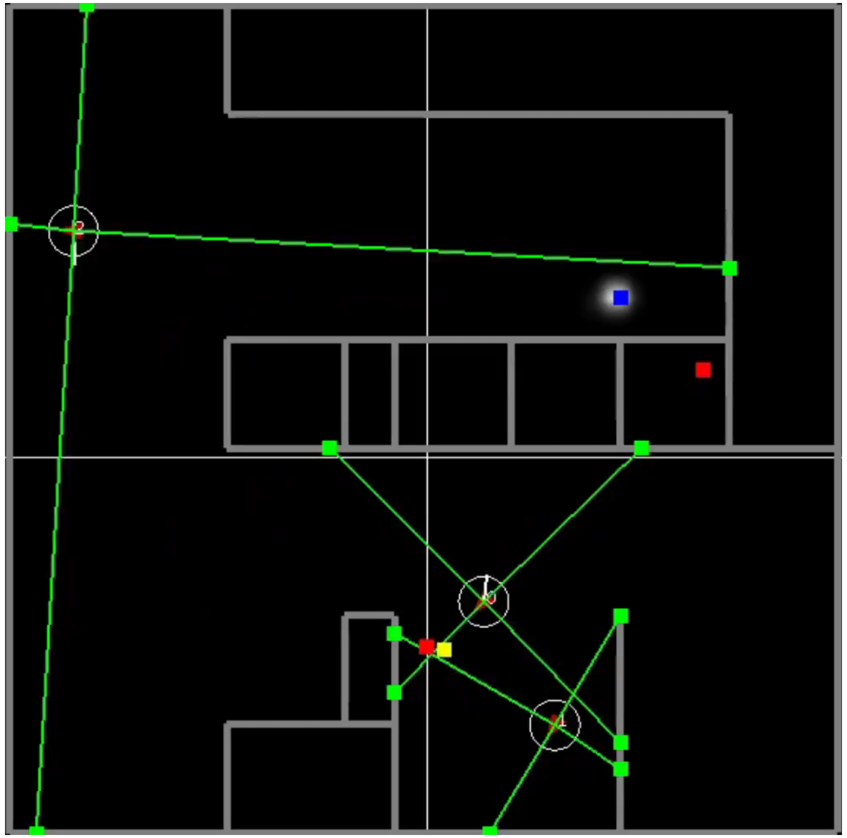

Sniffy Bug

- Our solution: **Sniffy Bug**, a **novel** PSO-powered **bug algorithm for GSL** using **evolved parameters.**
- Generates waypoints in own reference frame using particle swarm optimization (PSO).
- Tracks waypoints using novel bug algorithm.

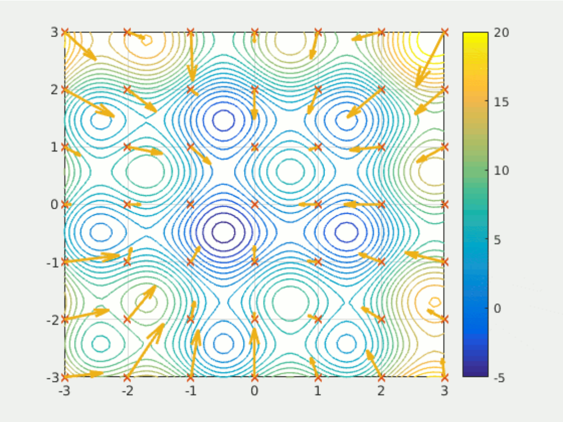

TU Delft
Harvard John A. Paulson School of Engineering and Applied Sciences
Universitat de Barcelona

Evolutionary Optimization

- Genome contains 13 variables:
  - Waypoint generation weights: ‘Exploring’ and ‘Seeking’.
  - All thresholds for Sniffy Bug obstacle avoidance and swarm avoidance.
- Cost of each agent:
  - Average distance to source.
  - Penalty for collision (+ 1.0).

Evolutionary Optimization

- Hard instance problem
- Heterogeneity in environments → learn to solve easy environments

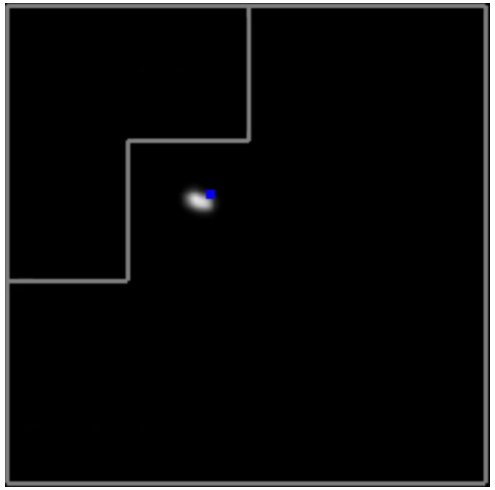

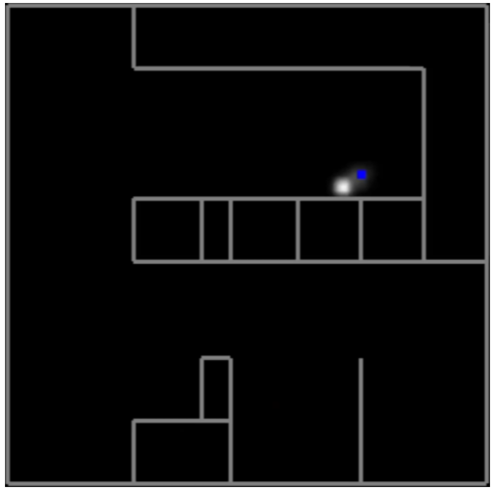

Results - Simulation

	Success Rate	Avg Distance to Source [m]	Avg time to source [s]
Manual Parameters	89 %	3.29	51.1
Evolved without Doping	85 %	2.90	47.1
Evolved with Doping	**93 %**	**2.73**	**39.2**

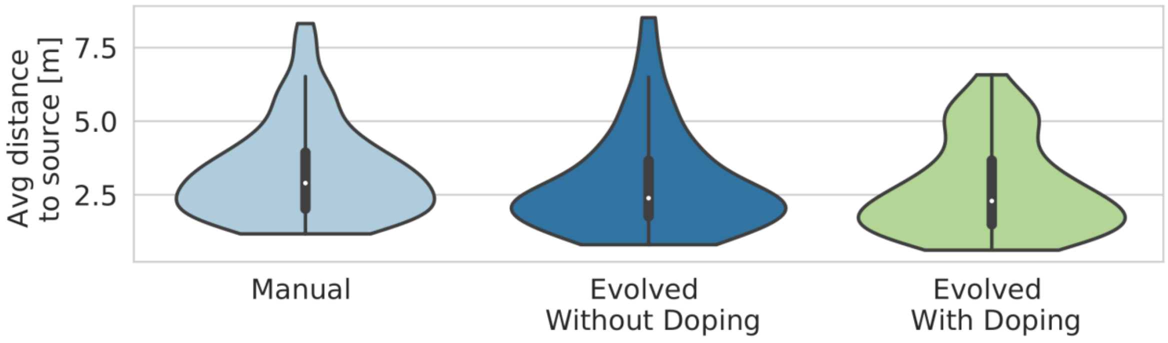

Manual parameters

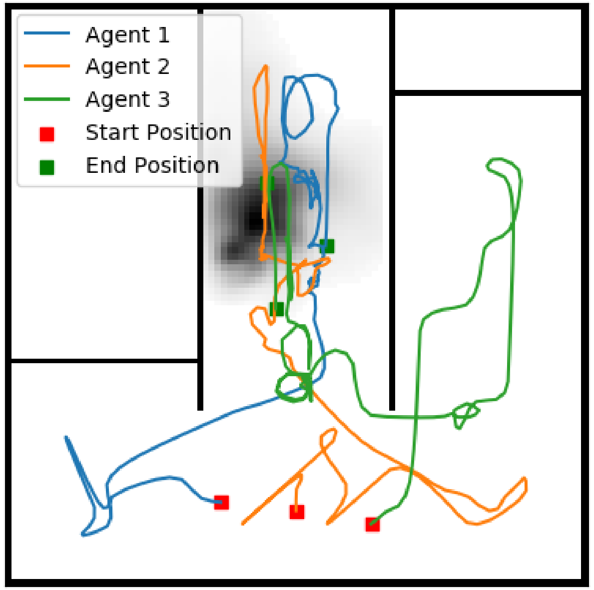

Evolved using doping

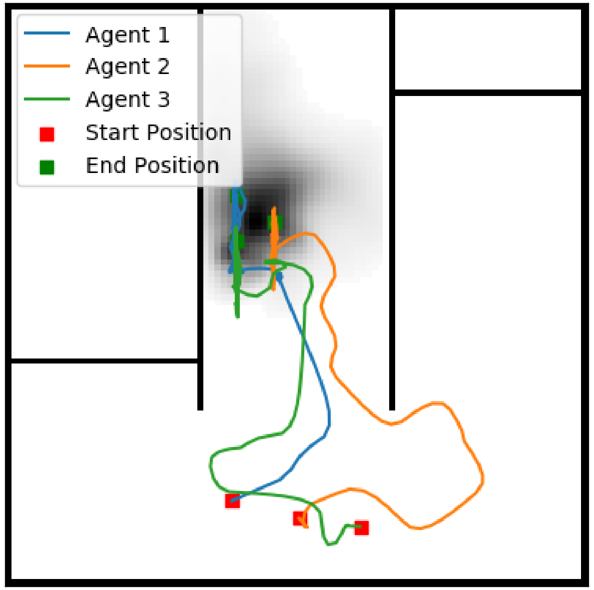

TU Delft | Harvard John A. Paulson School of Engineering and Applied Sciences | Universitat de Barcelona

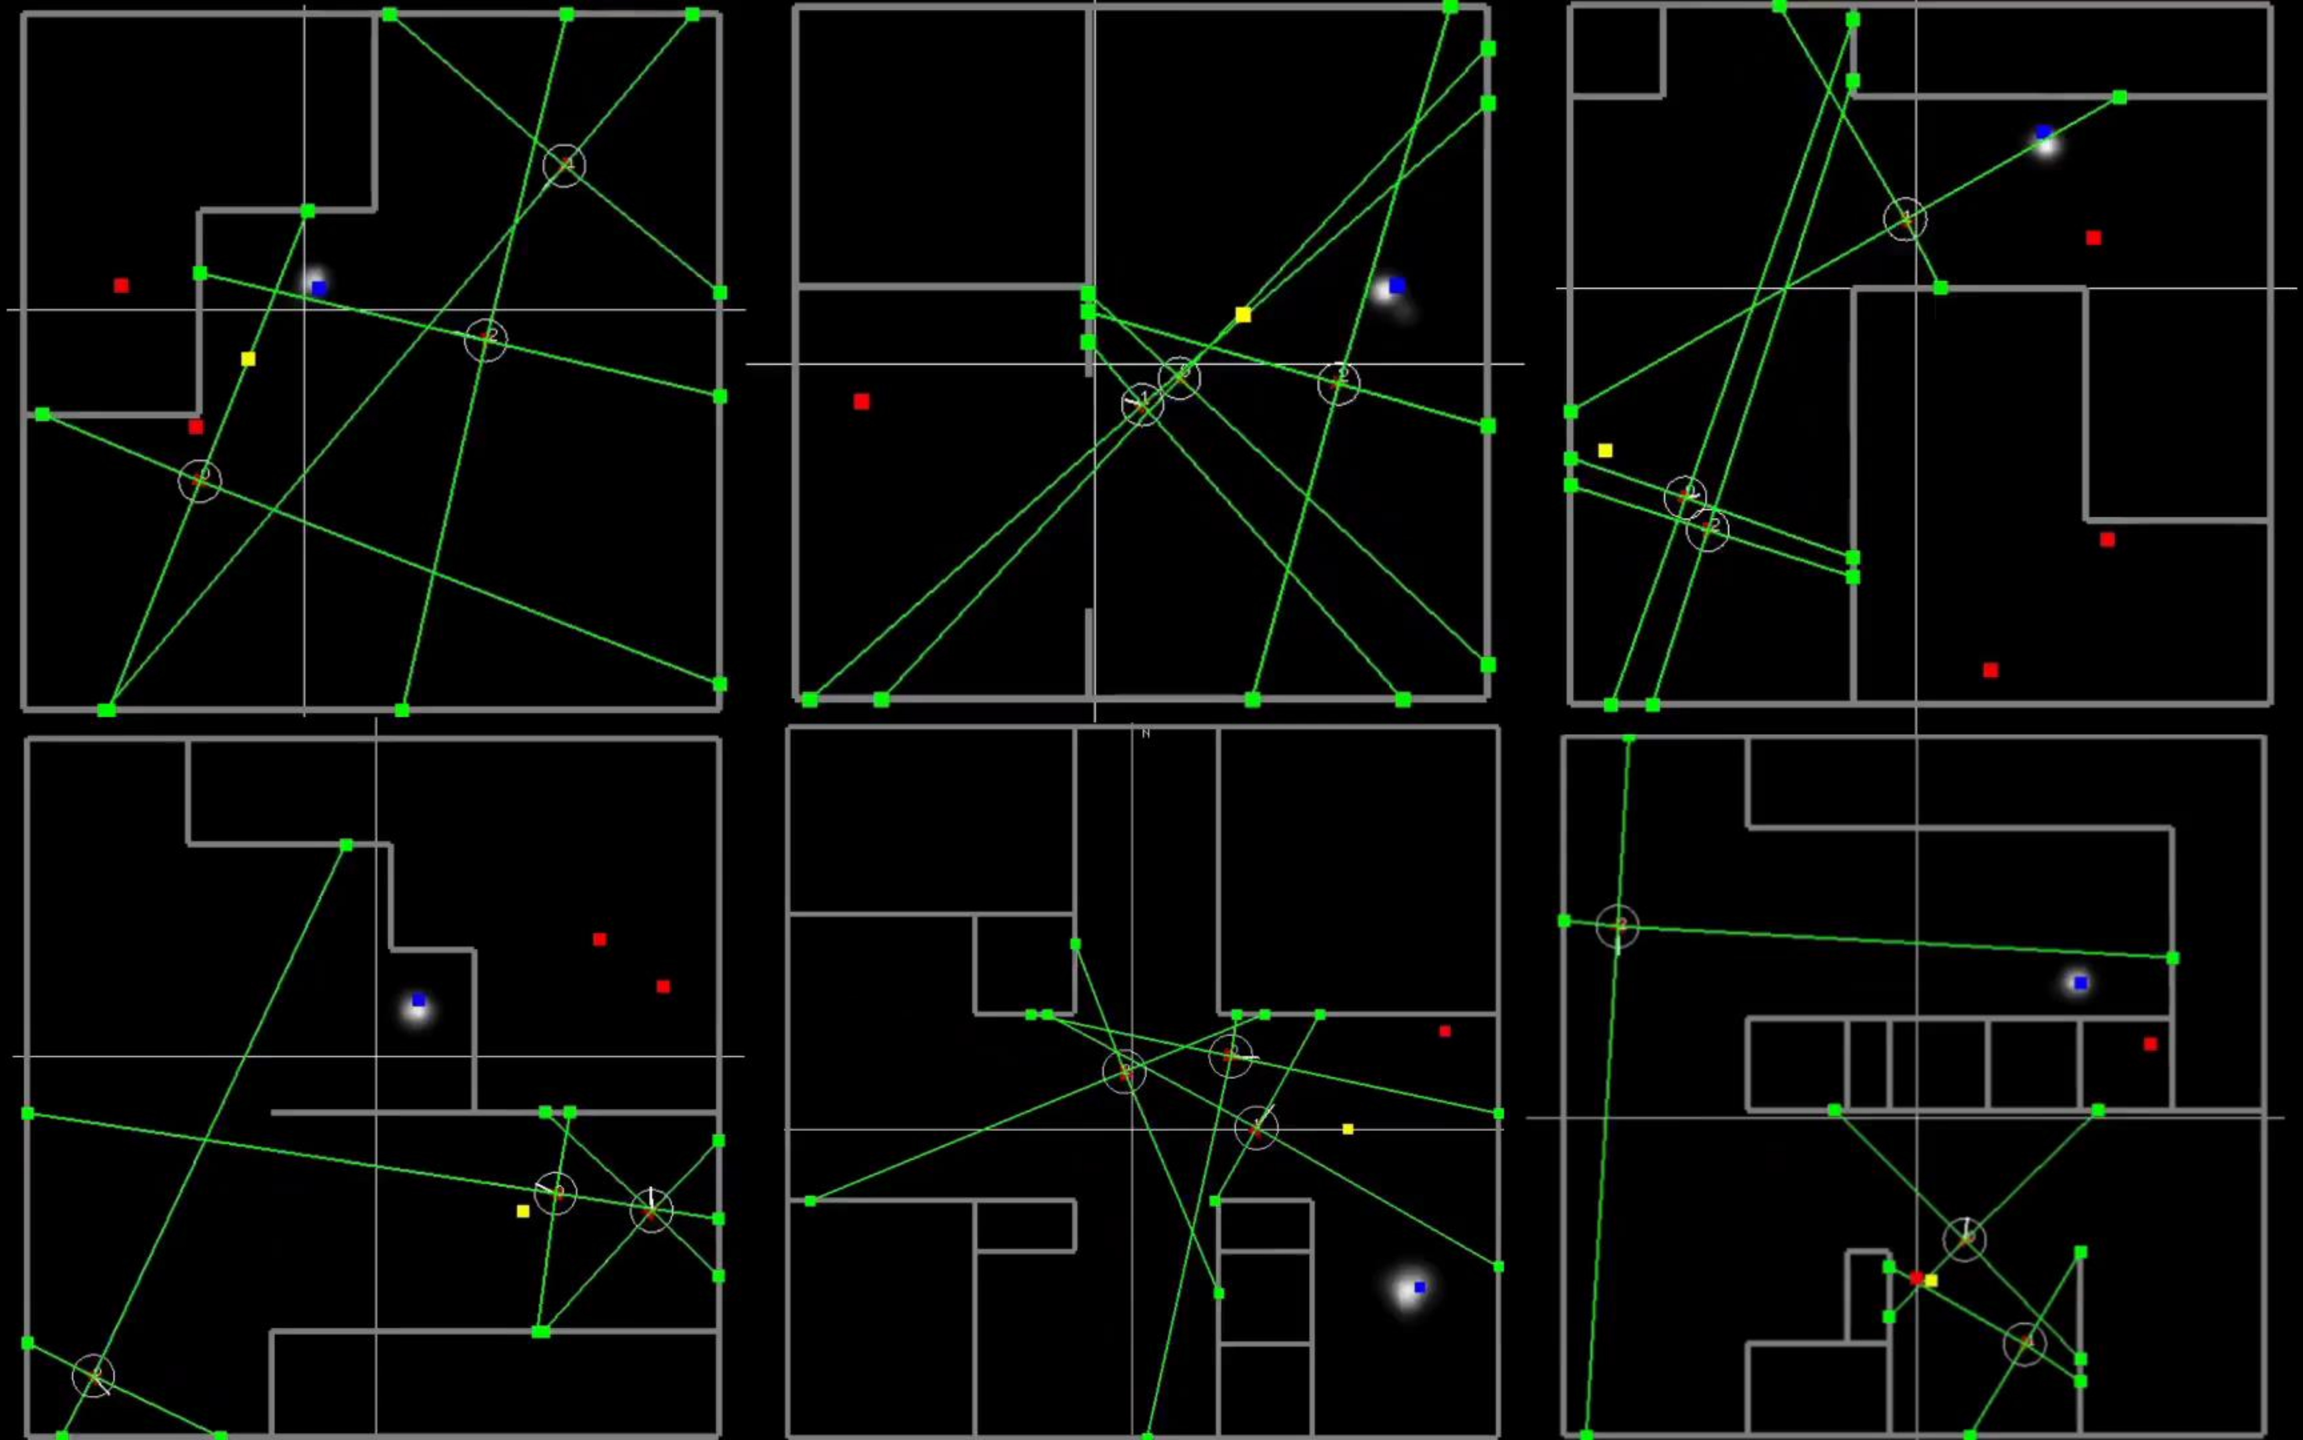


2020-11-17 Tue 06:08:44
CyberZoo TU Delft
iso view

Conclusion

We contribute:

- The **first fully autonomous swarm of gas-seeking nano-quadcopters** in cluttered GPS-denied environments.

https://arxiv.org/abs/2107.05490

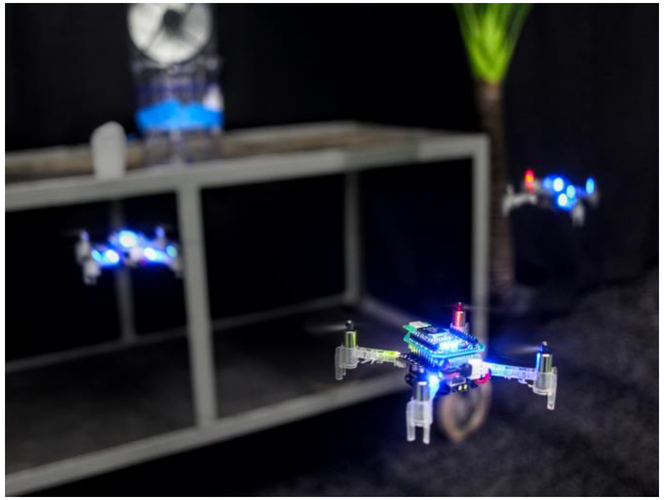

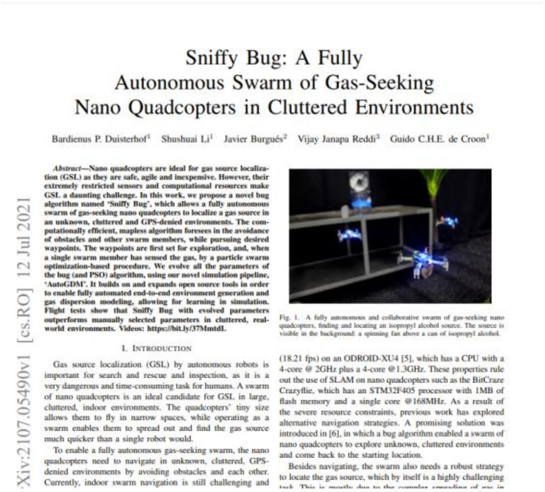

Sniffy Bug: A Fully Autonomous Swarm of Gas-Seeking Nano Quadcopters in Cluttered Environments

Bardienus P. Duisterhof[1] Shushuai Li[1] Javier Burgués[2] Vijay Janapa Reddi[3] Guido C.H.E. de Croon[1]

TU Delft | Harvard John A. Paulson School of Engineering and Applied Sciences | Universitat de Barcelona

Thanks for attending!

Collaborators and Mentors




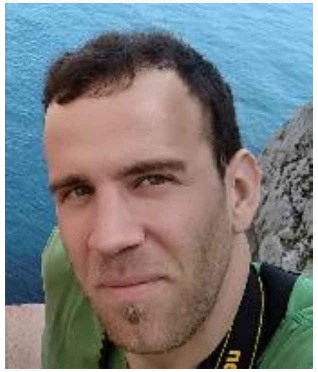

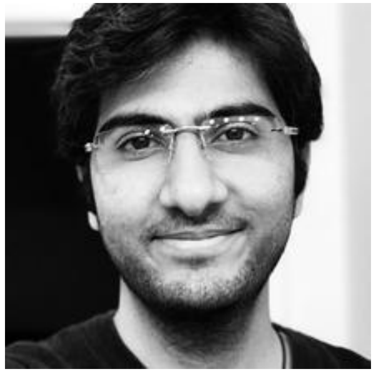

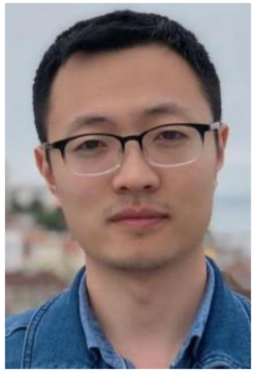

TU Delft

Harvard John A. Paulson School of Engineering and Applied Sciences

Universitat de Barcelona

Google Brain