Mobile Robotics for Computer Architects

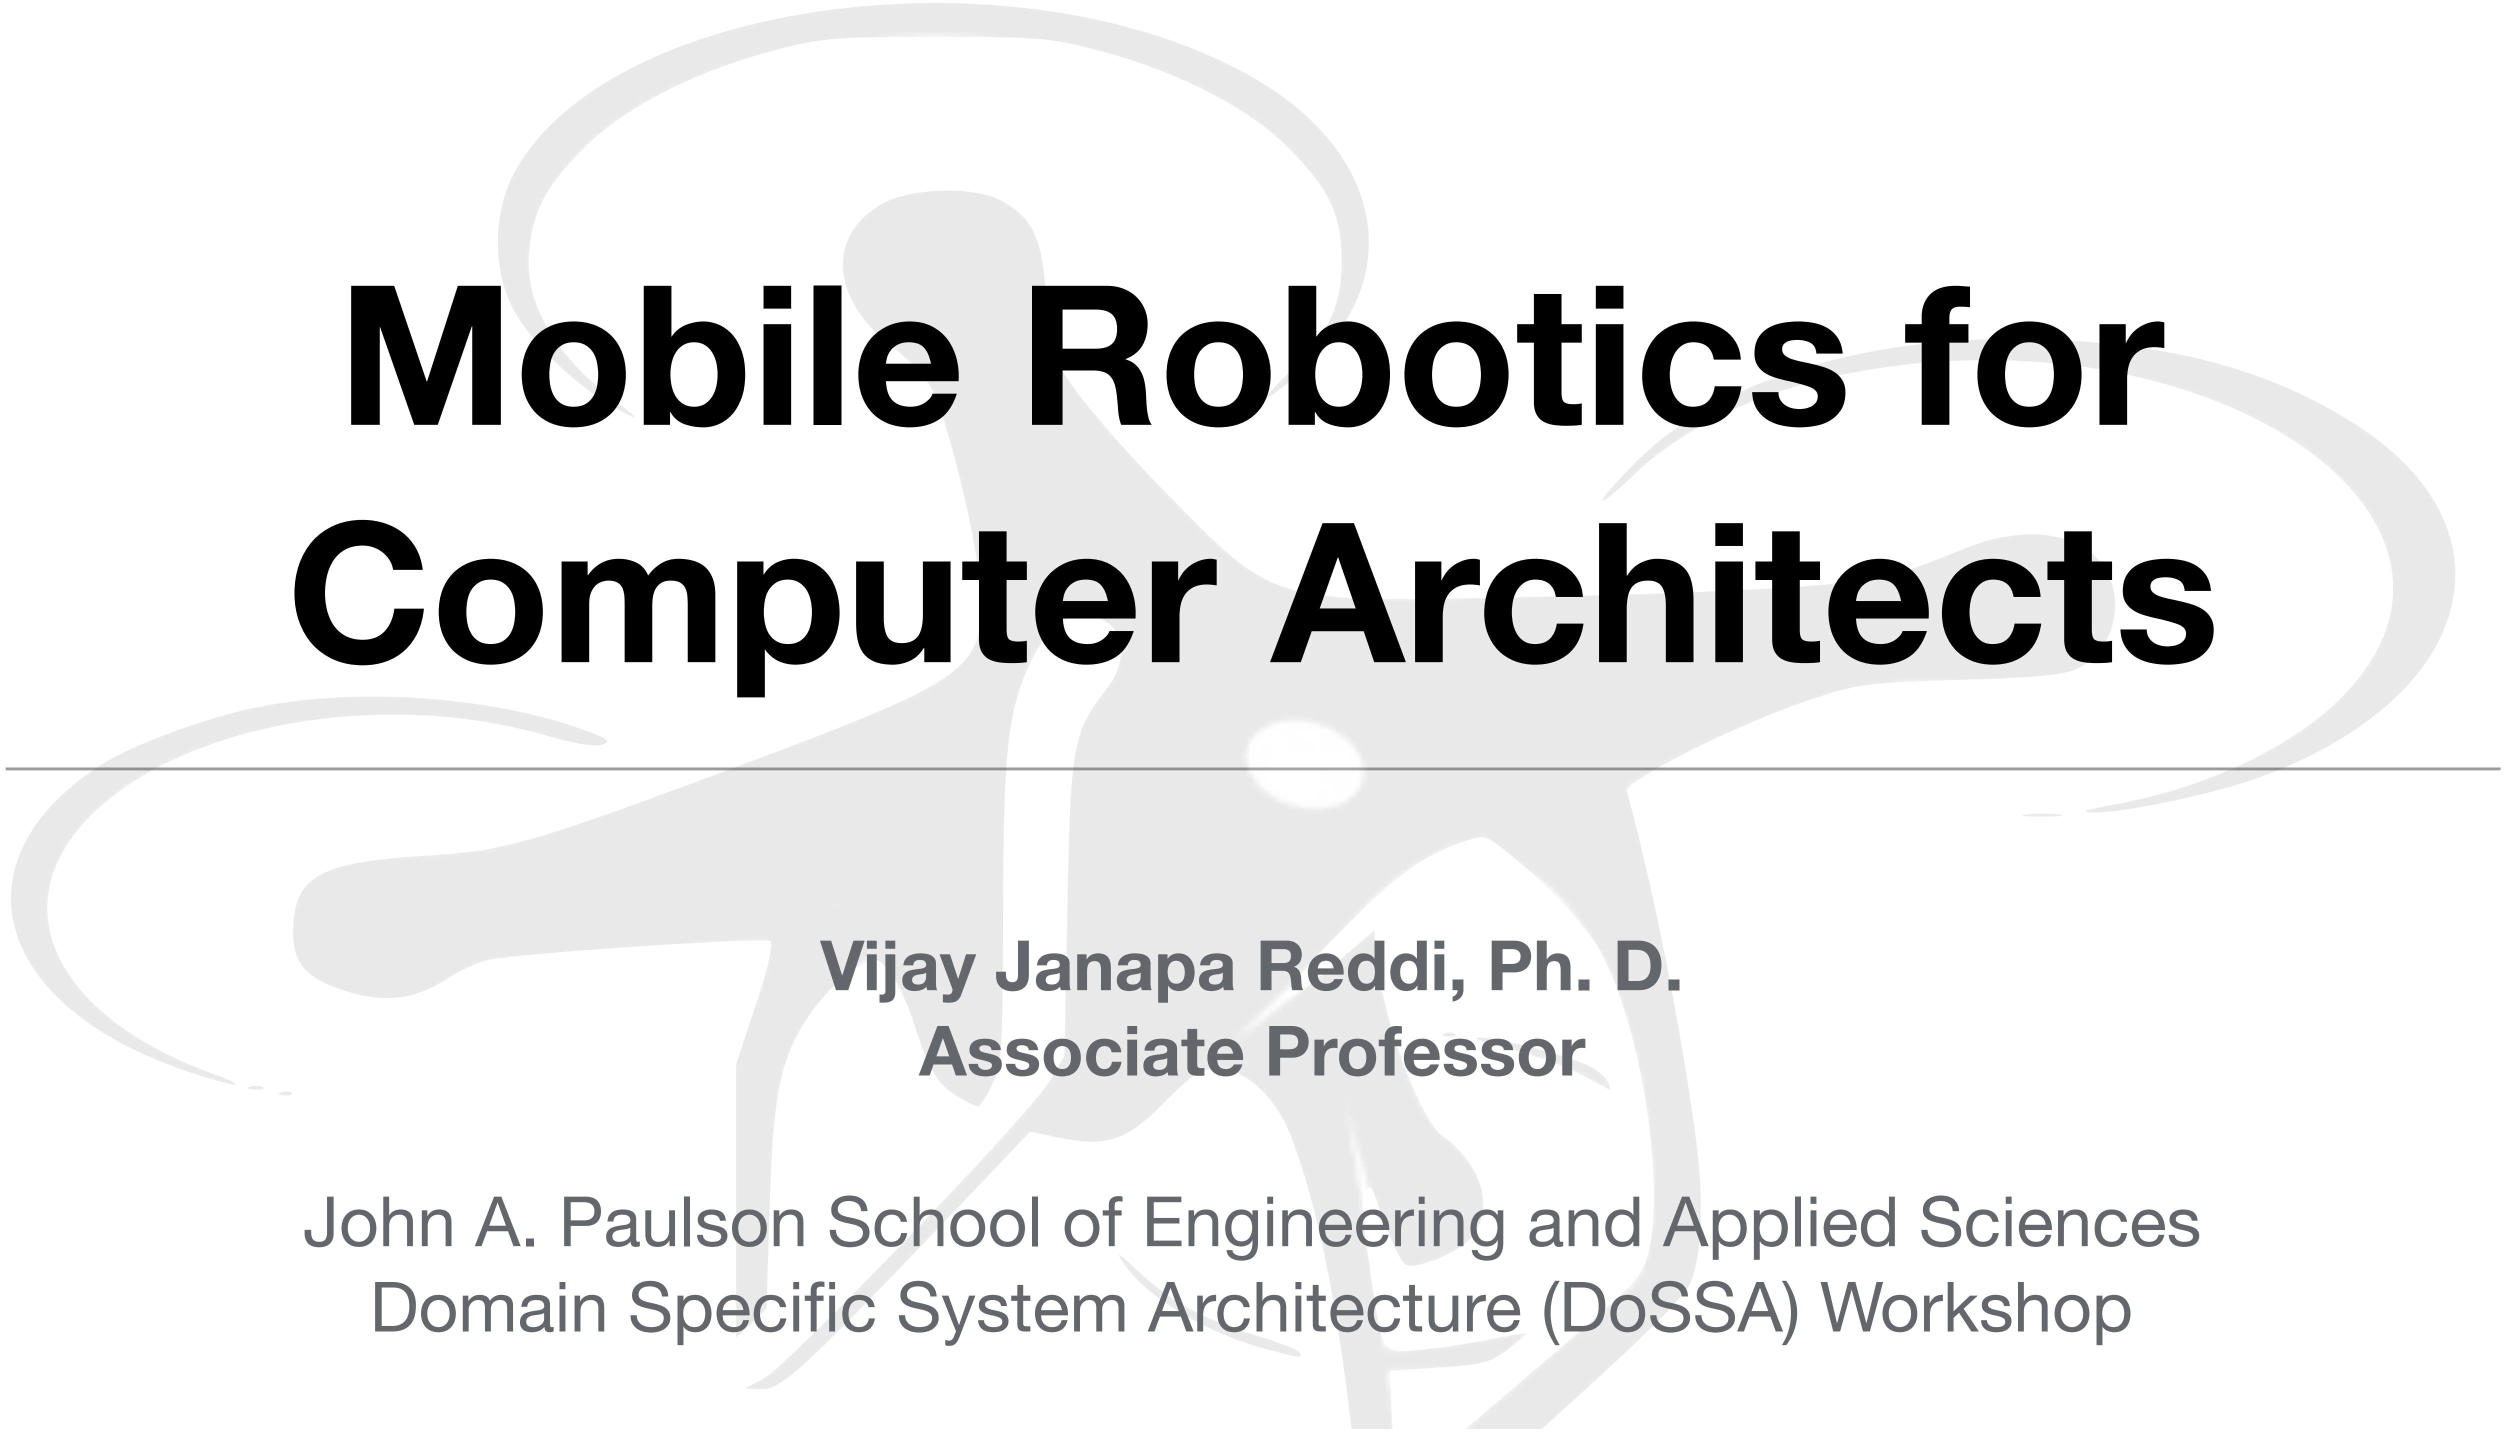

Vijay Janapa Reddi, Ph. D.
Associate Professor

John A. Paulson School of Engineering and Applied Sciences

Domain Specific System Architecture (DoSSA) Workshop

HARVARD
UNIVERSITY

ASIMO

ABB

ASIMO
HONDA

GAP
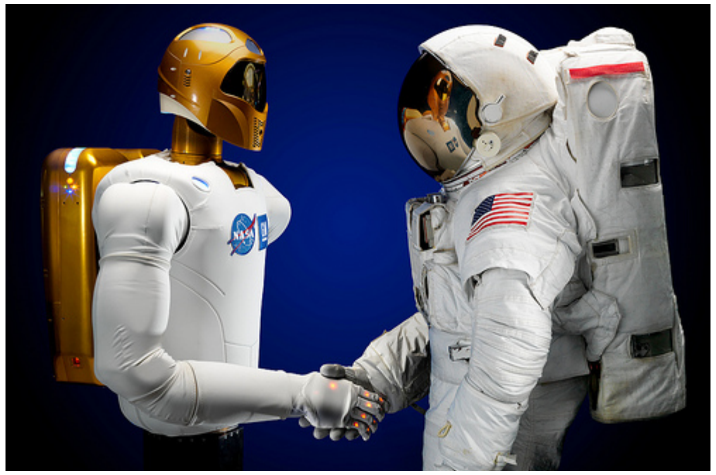
NASA

Amazon

Amazon

Berkley campus (Kiwi)

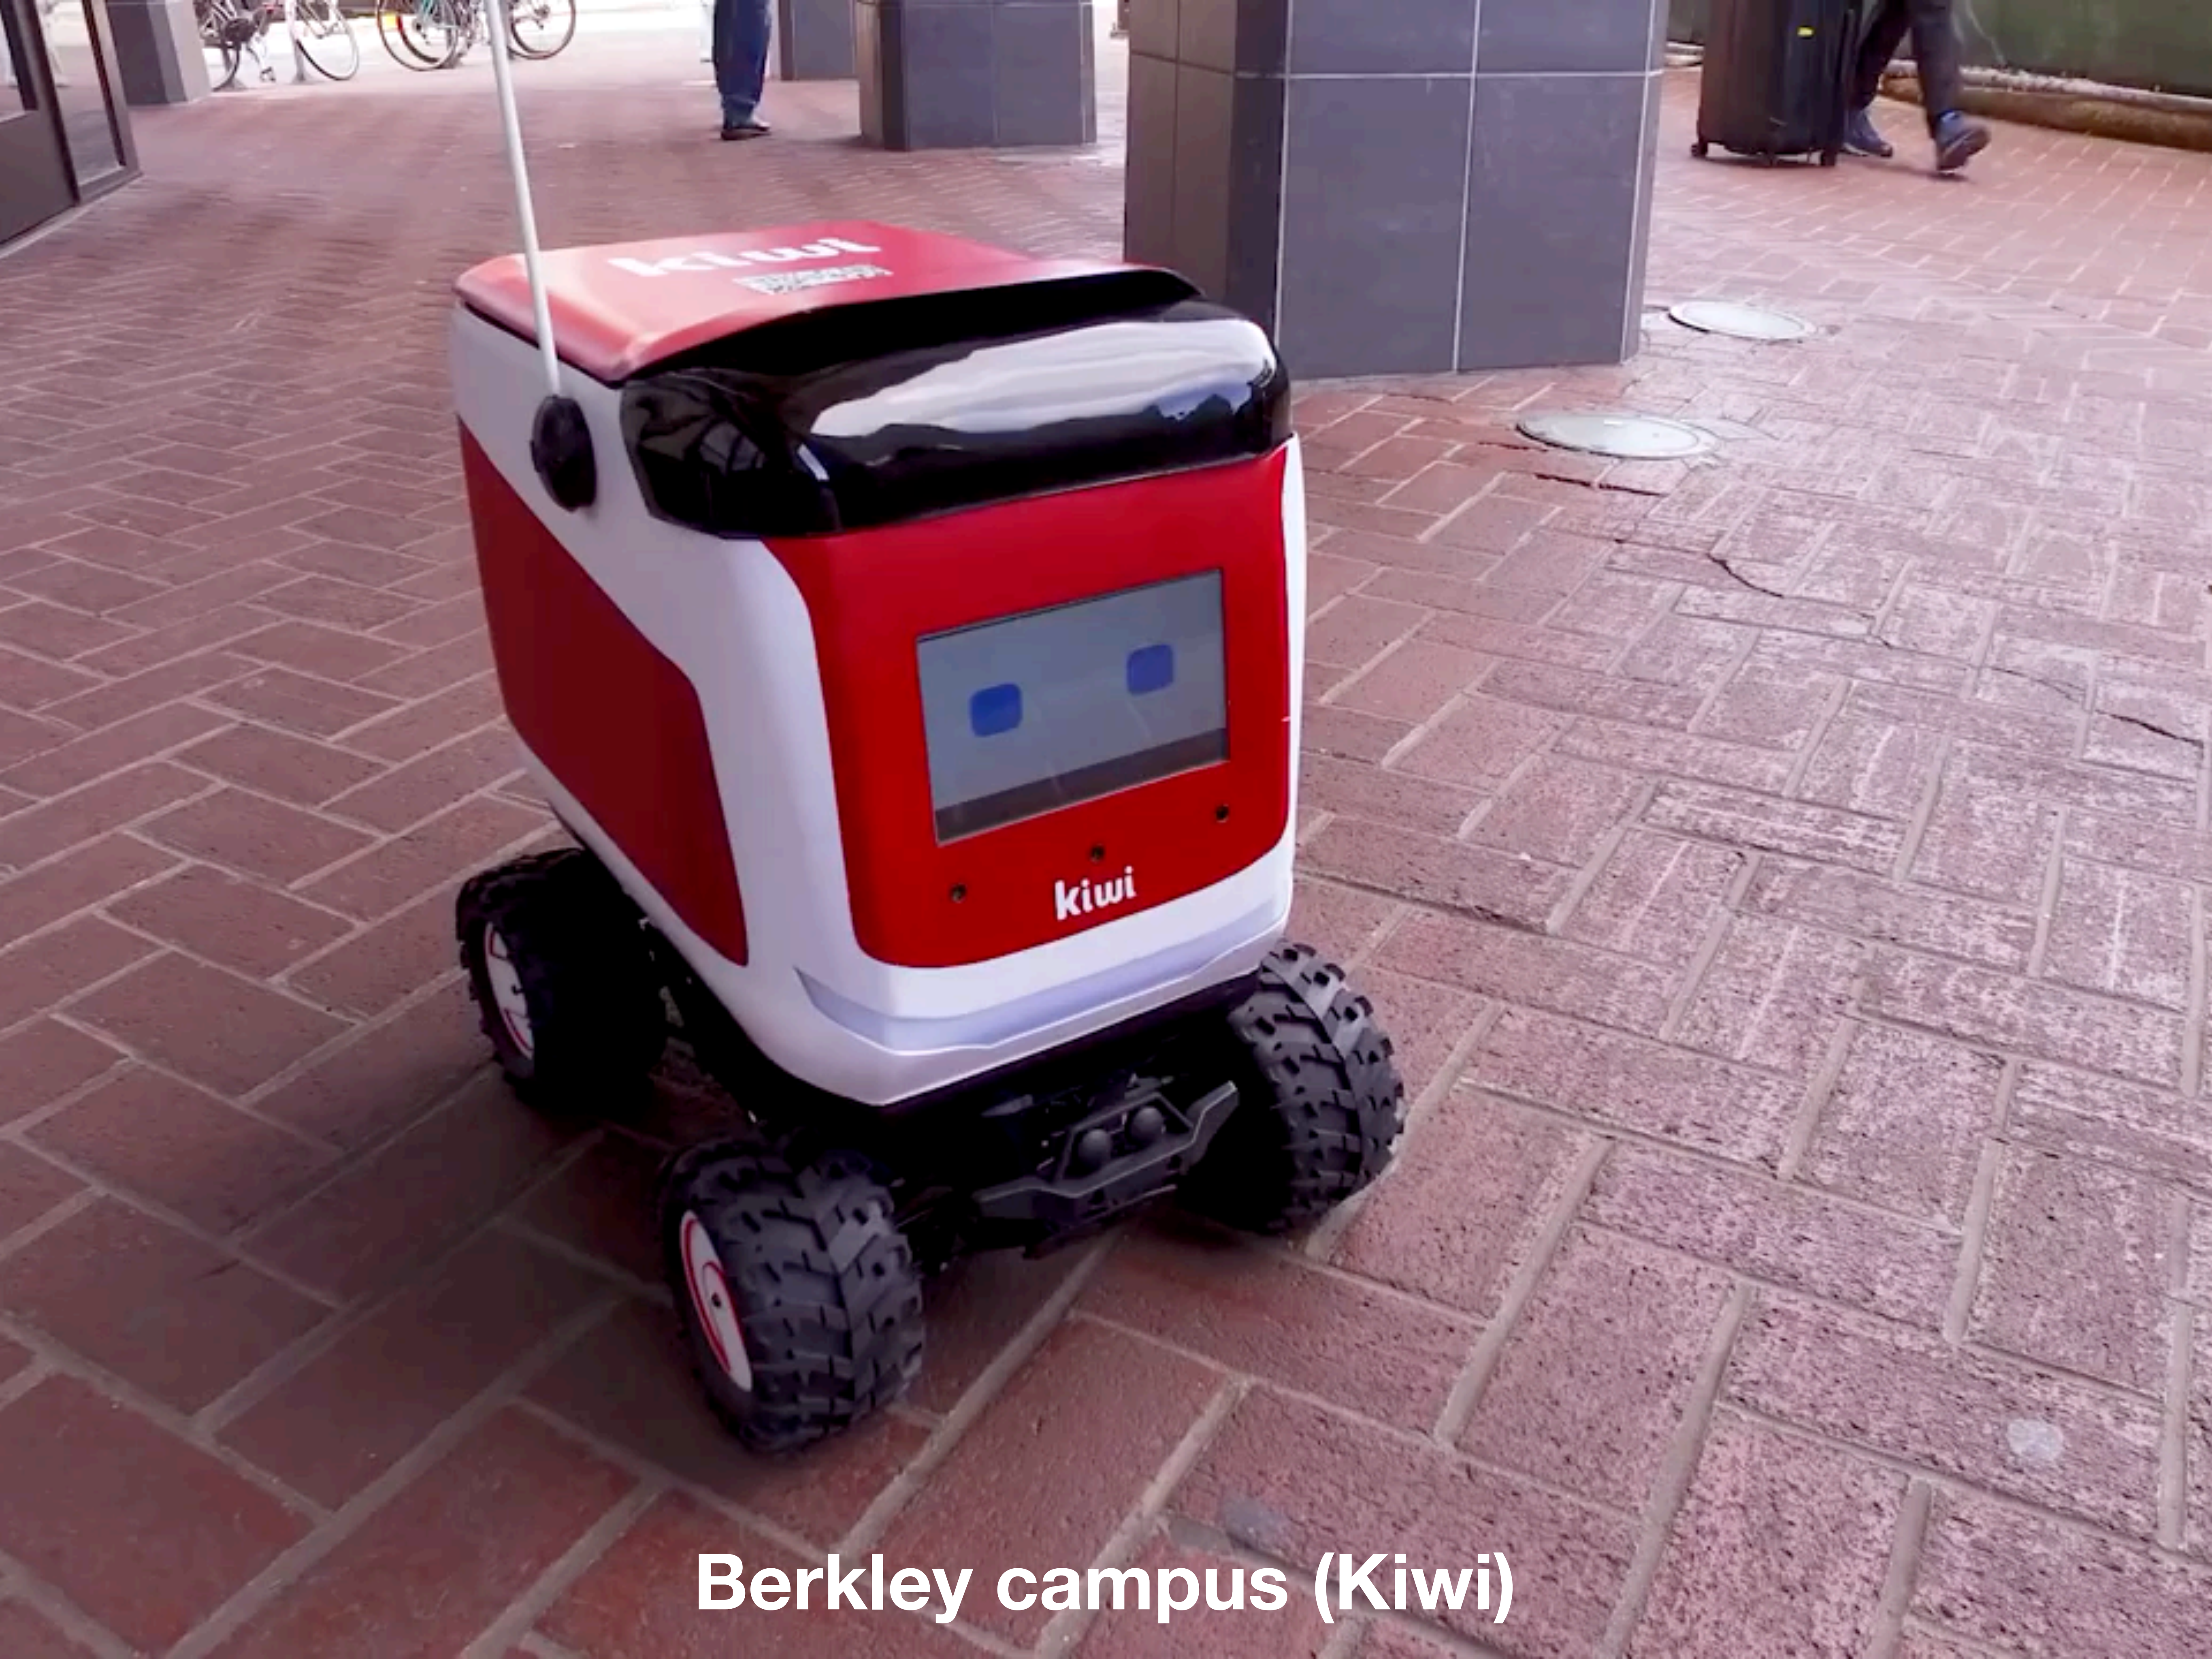

Berkley campus (Kiwi)

Berkley campus (Kiwi)

Berkley campus (Kiwi)

Berkley campus (Kiwi)

Berkley campus (Kiwi)

Security

Wildlife Conservation

Search and Rescue

Cargo Delivery

20 Commercial Drone Use Cases

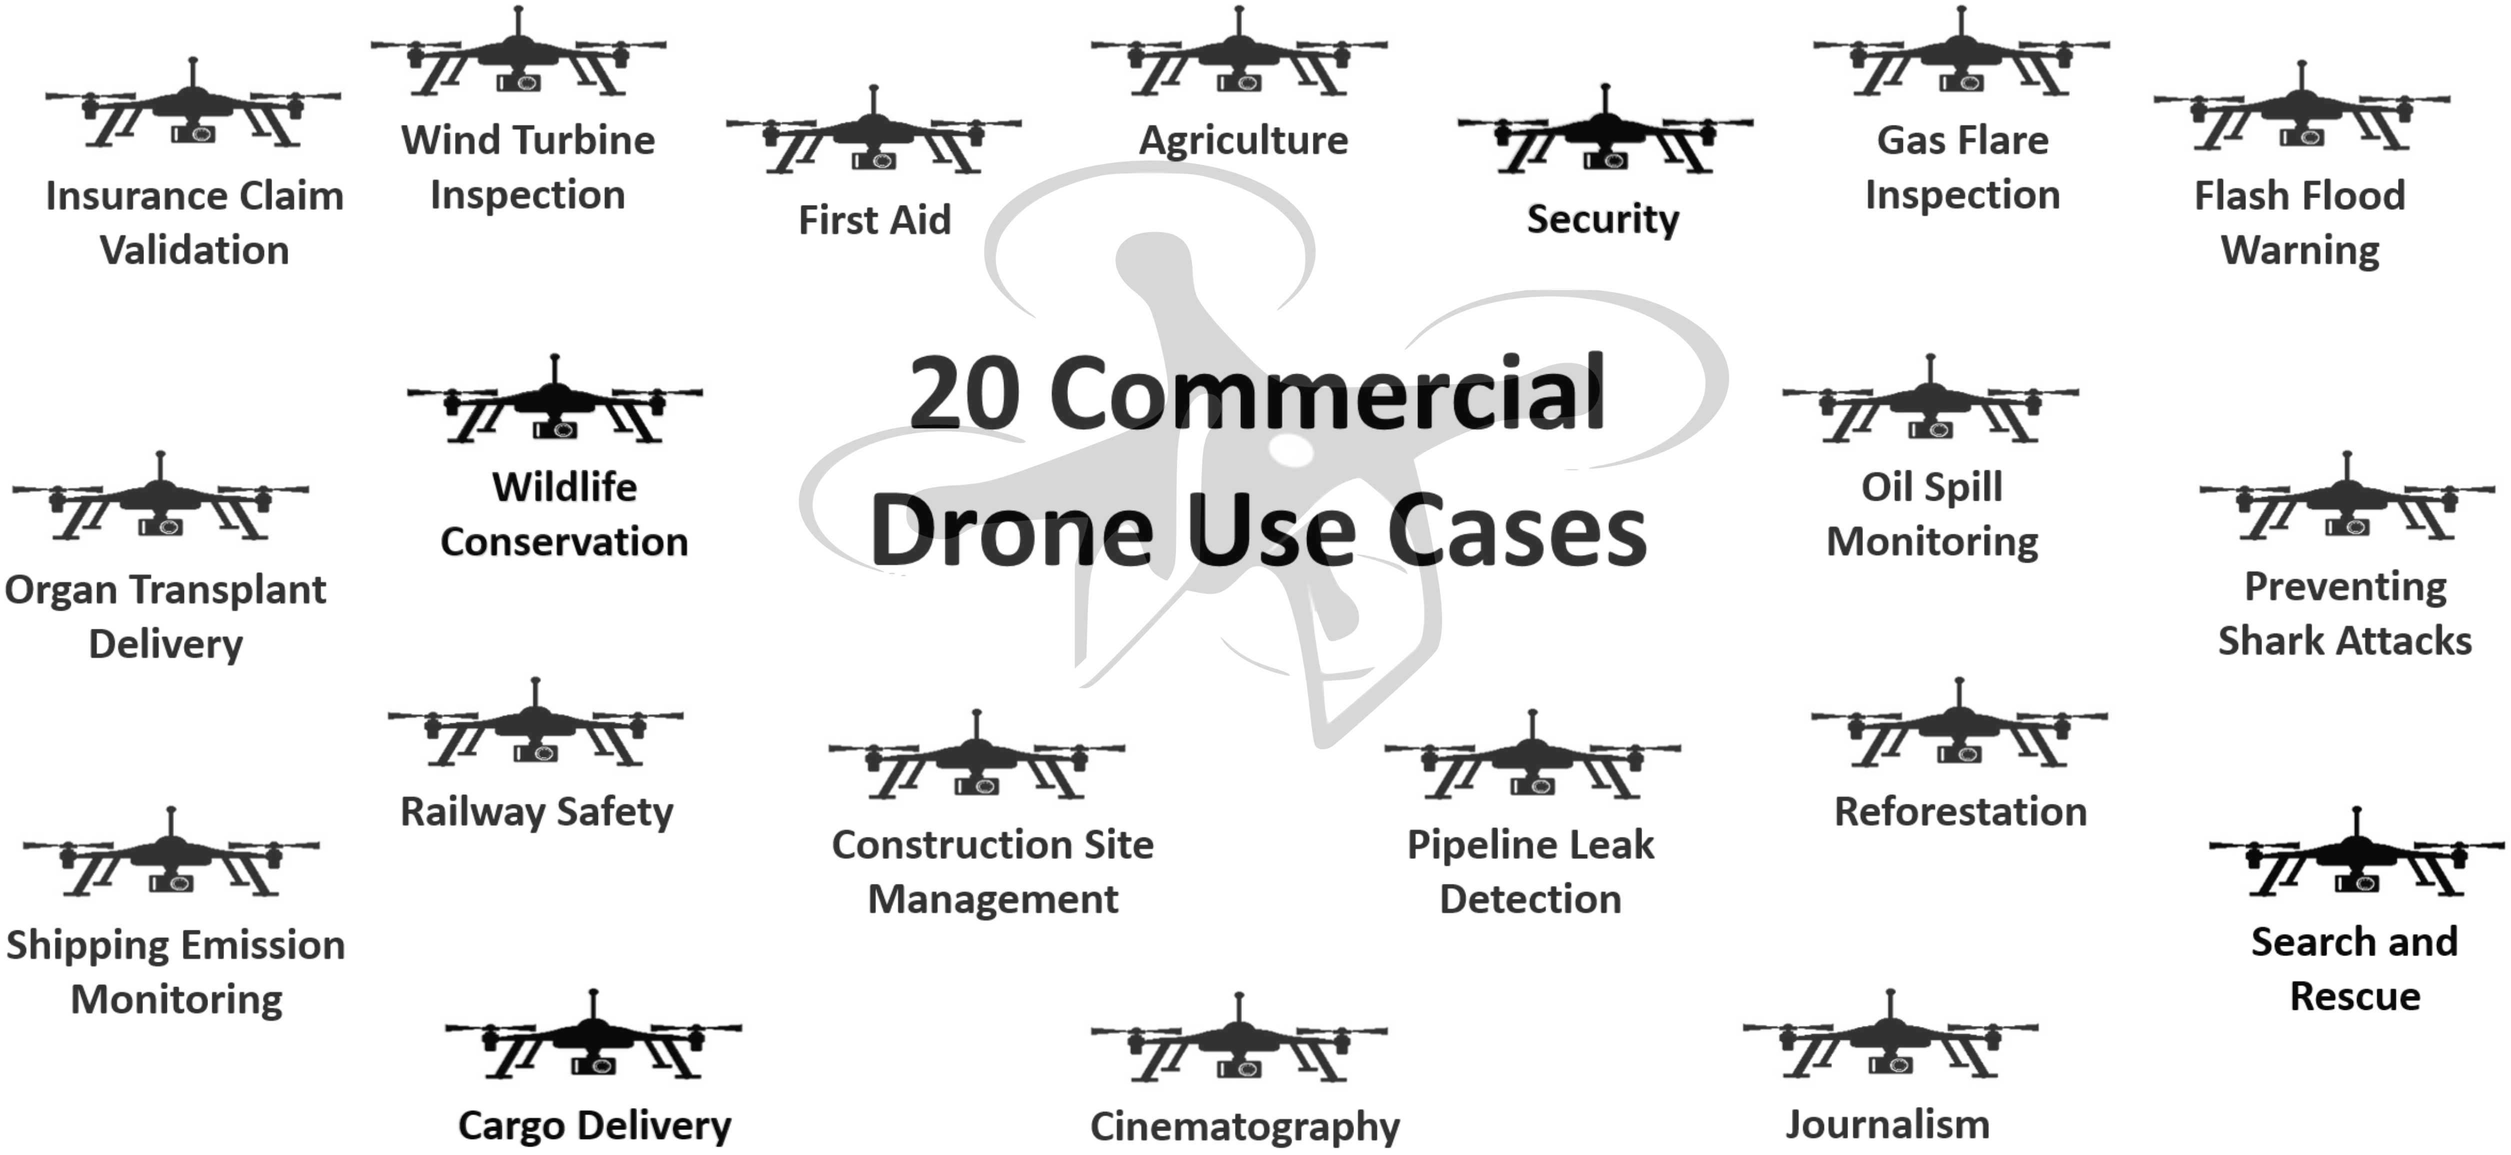

Insurance Claim Validation

Wind Turbine Inspection

First Aid

Agriculture

Security

Gas Flare Inspection

Flash Flood Warning

Organ Transplant Delivery

Wildlife Conservation

Oil Spill Monitoring

Preventing Shark Attacks

Railway Safety

Construction Site Management

Pipeline Leak Detection

Reforestation

Shipping Emission Monitoring

Search and Rescue

Cargo Delivery

Cinematography

Journalism

Global Market Size for Robotics

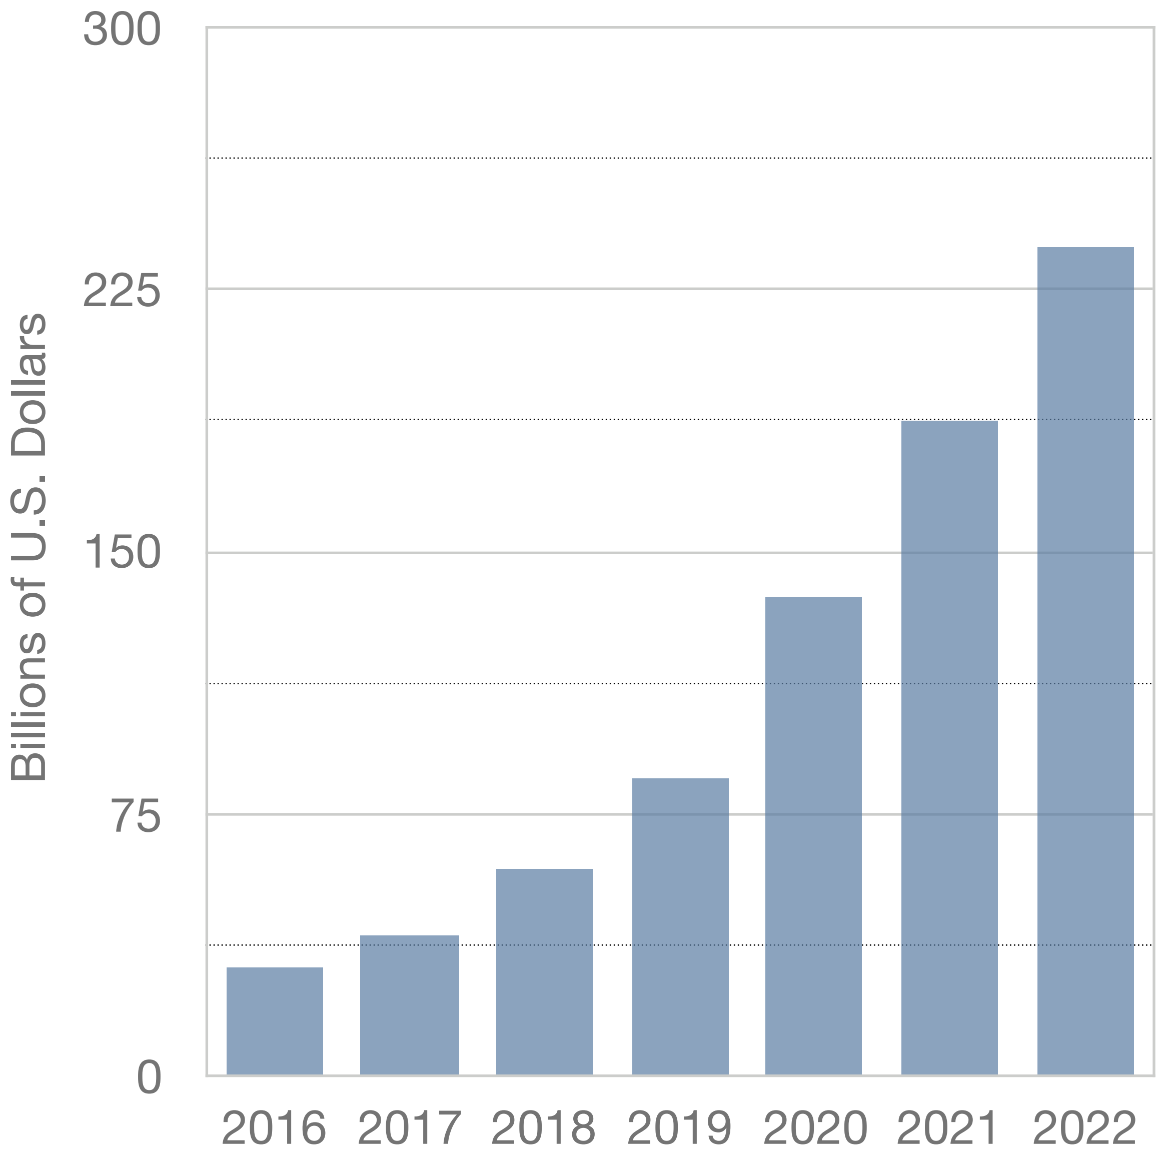

"… revenue generated from the robotics market globally, both industrial and non-industrial, from 2016 to 2022. In 2017, the robotics market is estimated to be worth 40 billion U.S. dollars globally. The industrial robotics market, which has traditionally represented the robotics industry … is giving way to non-industrial robots, such as personal assistant robots, customer service robots, autonomous vehicles, and unmanned aerial vehicles (UAVs)."

Source: [Statista 2018]

Global Market Size for Robotics

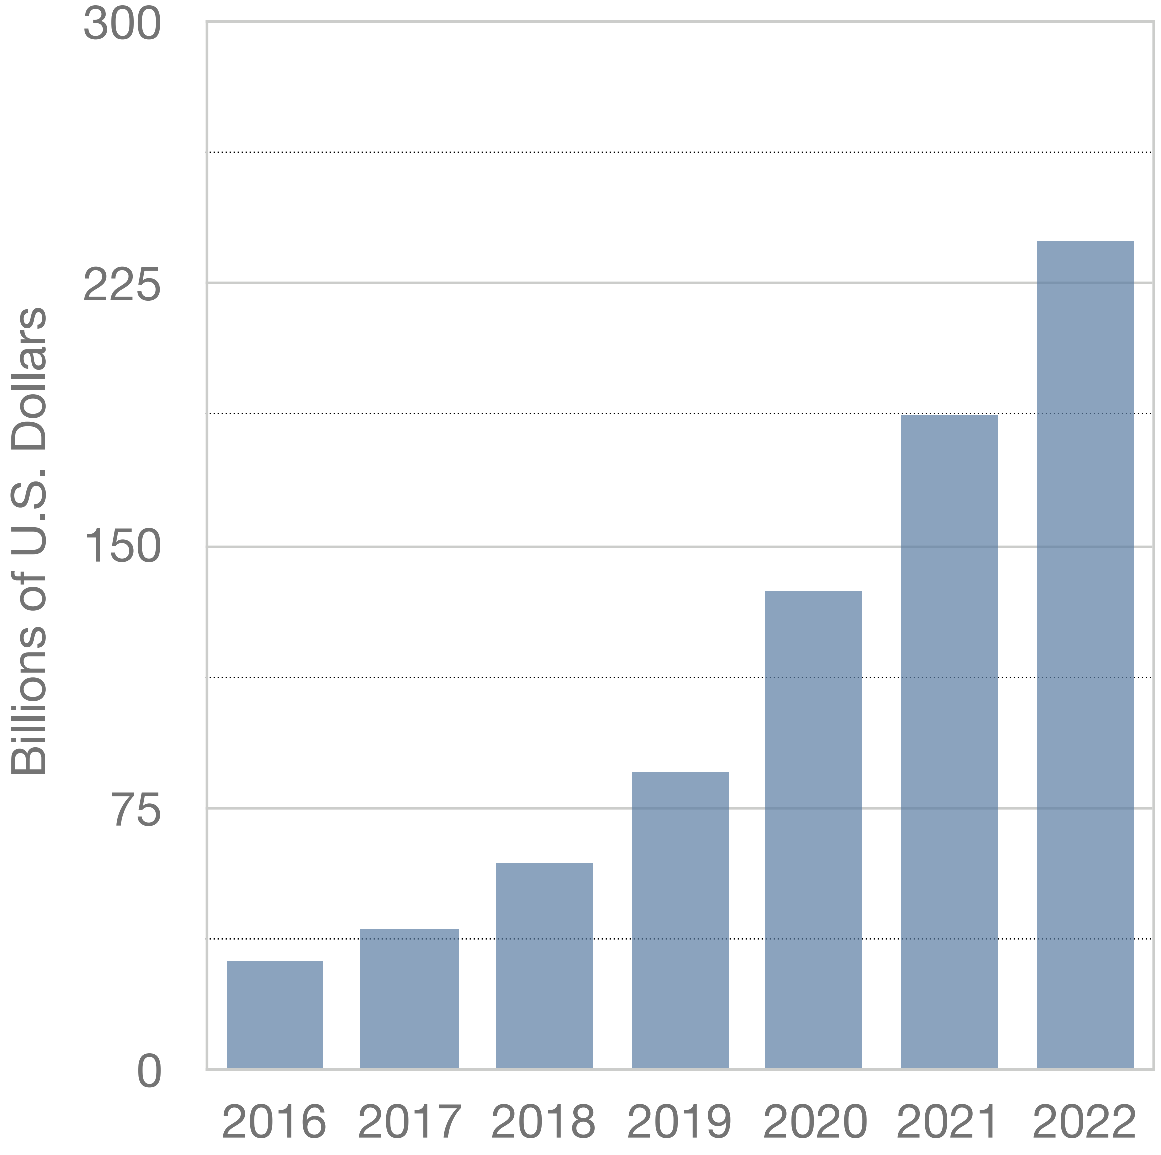

"… revenue generated from the robotics market globally, both industrial and non-industrial, from 2016 to 2022. In 2017, the robotics market is estimated to be worth 40 billion U.S. dollars globally. **The industrial robotics market, which has traditionally represented the robotics industry … is giving way to non-industrial robots, such as personal assistant robots, customer service robots, autonomous vehicles, and unmanned aerial vehicles (UAVs).**"

Source: [Statista 2018]

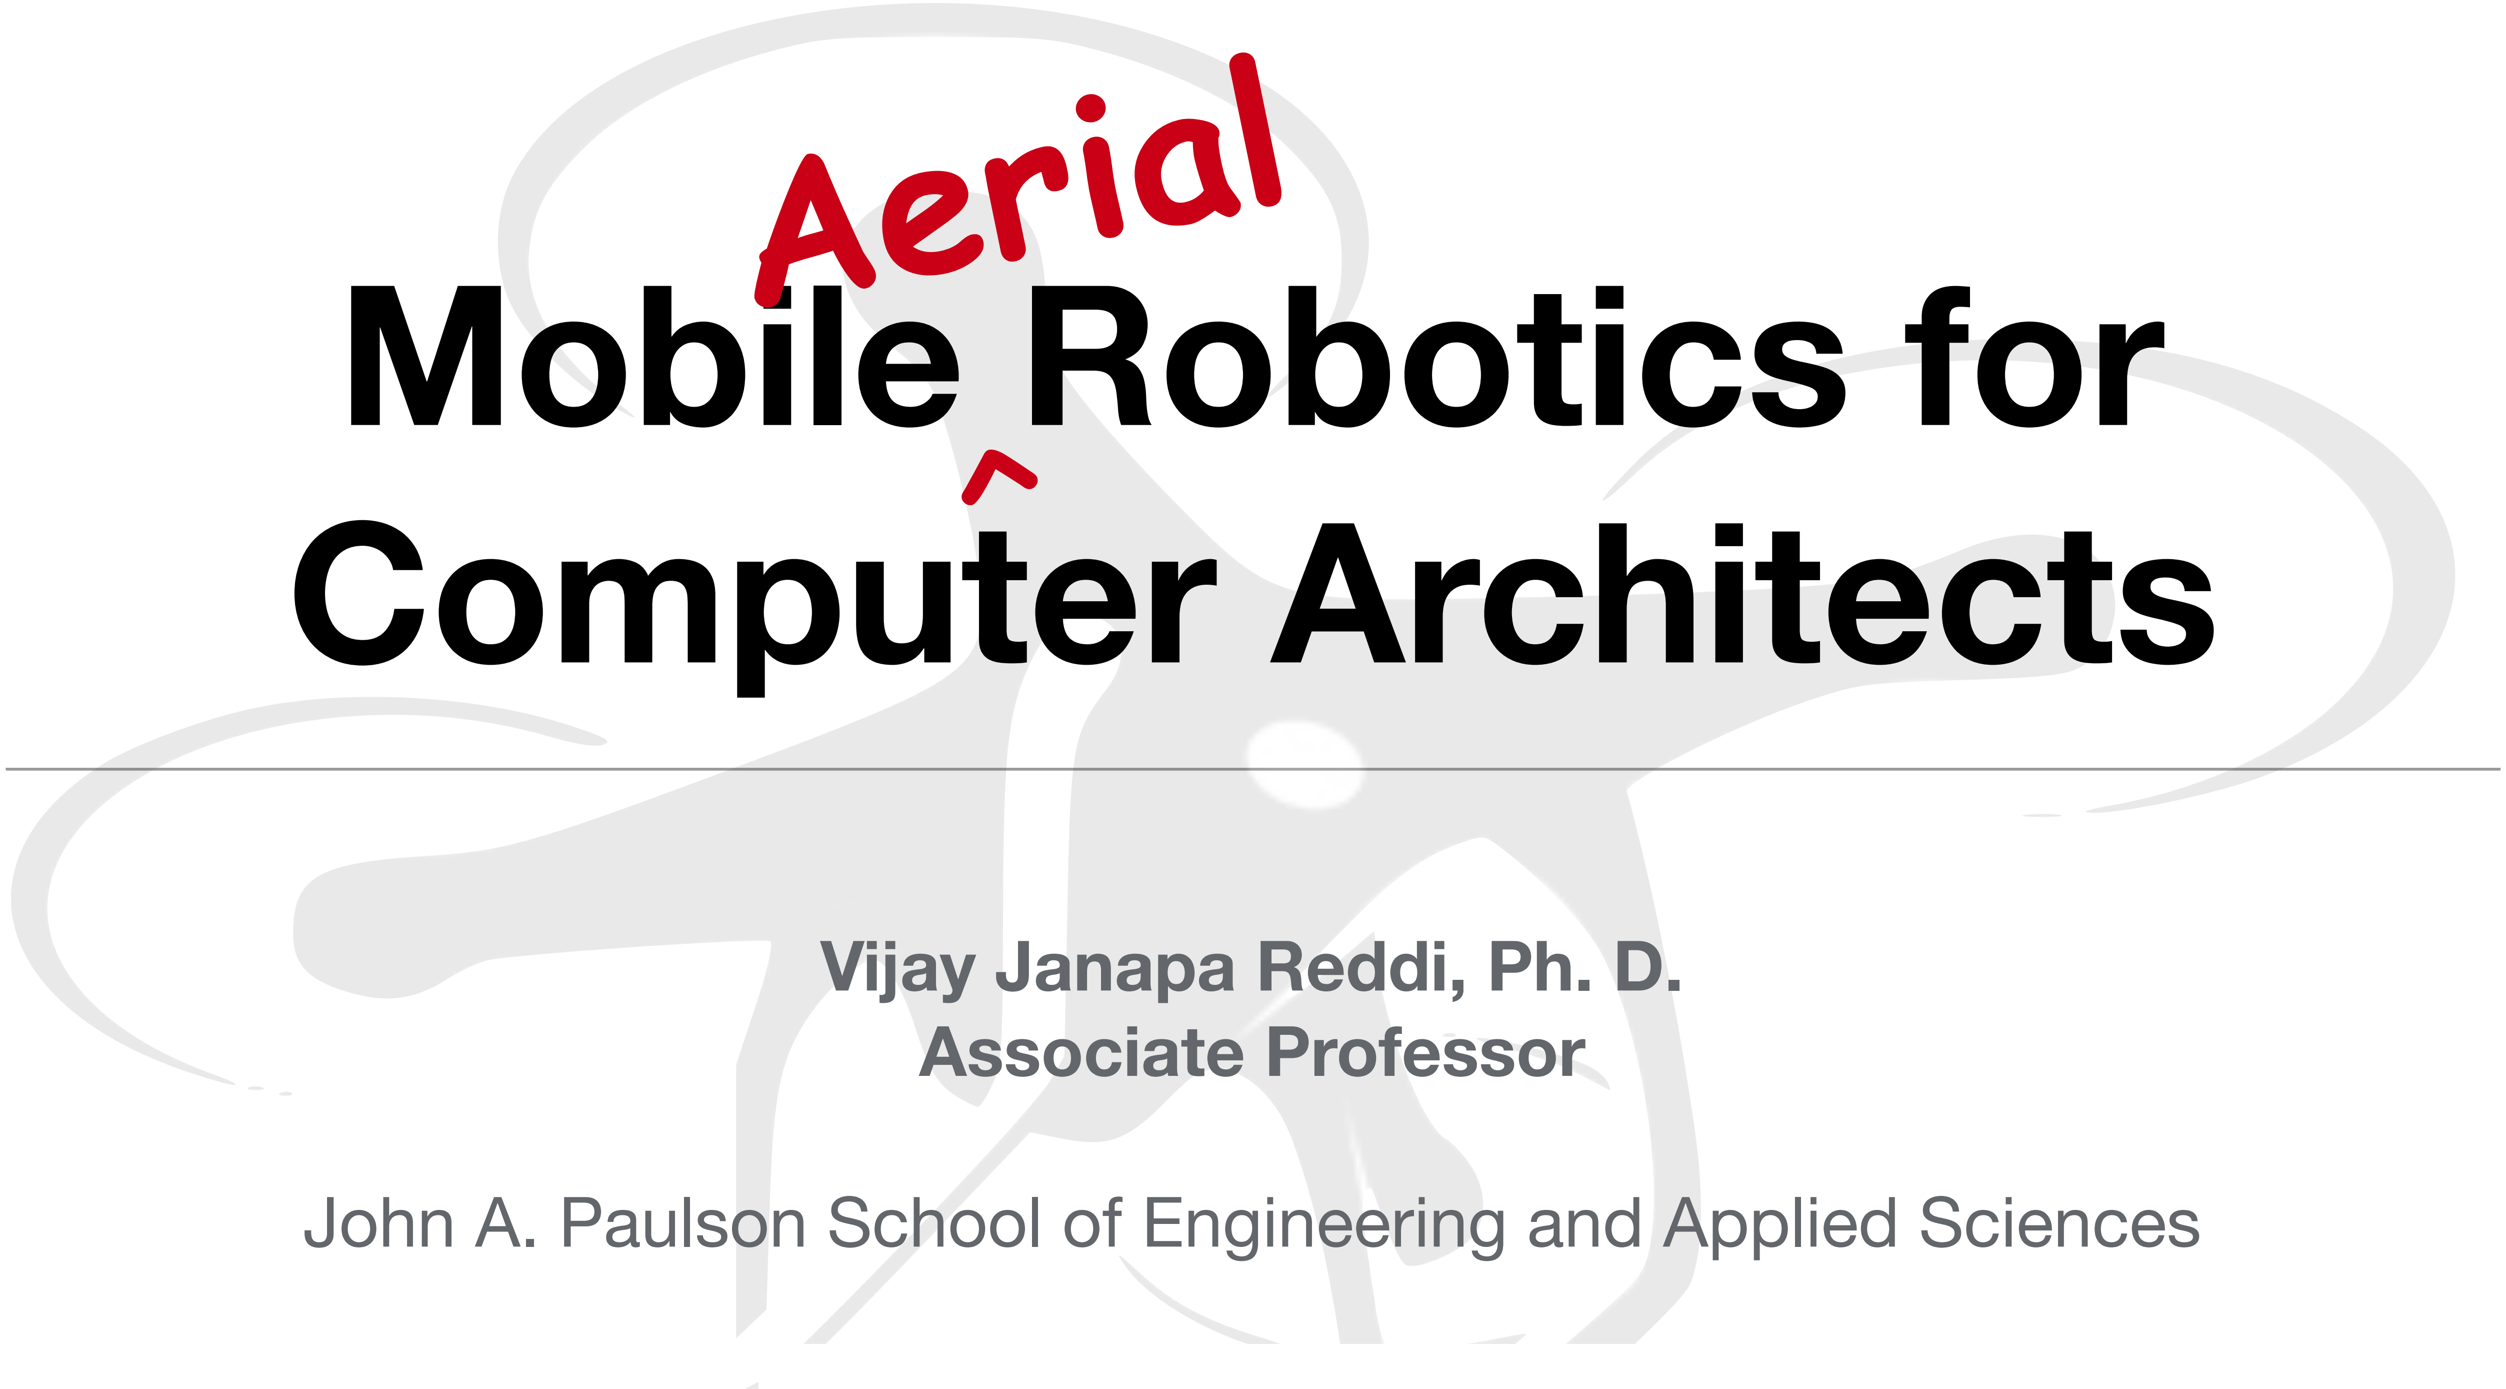

Aerial

Mobile Robotics for Computer Architects

Vijay Janapa Reddi, Ph. D.
Associate Professor

John A. Paulson School of Engineering and Applied Sciences

HARVARD
UNIVERSITY

What are the **challenges** facing aerial robots?

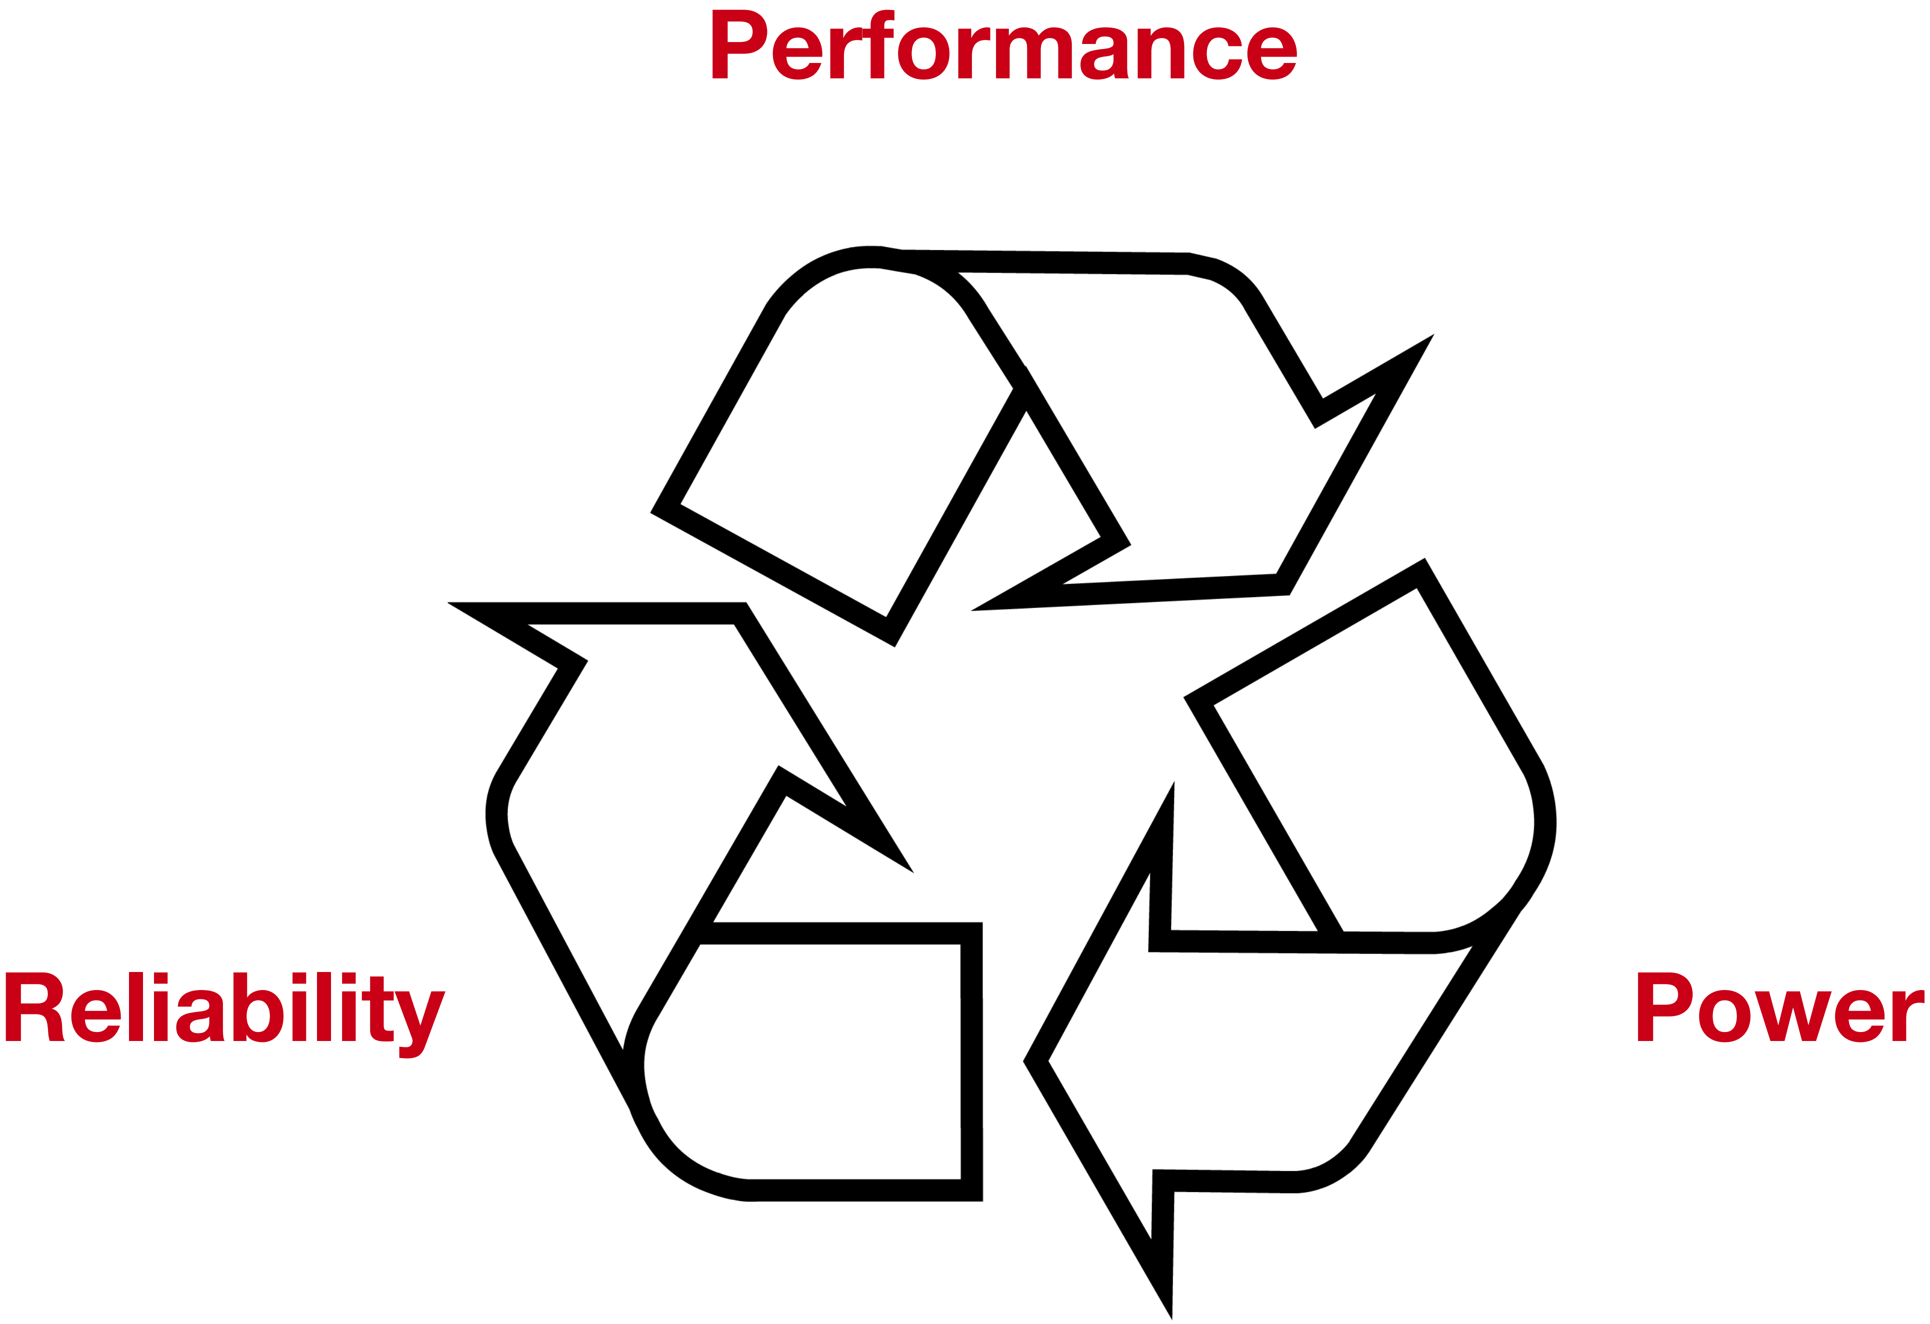
Performance
Reliability
Power

Performance (Real-time)

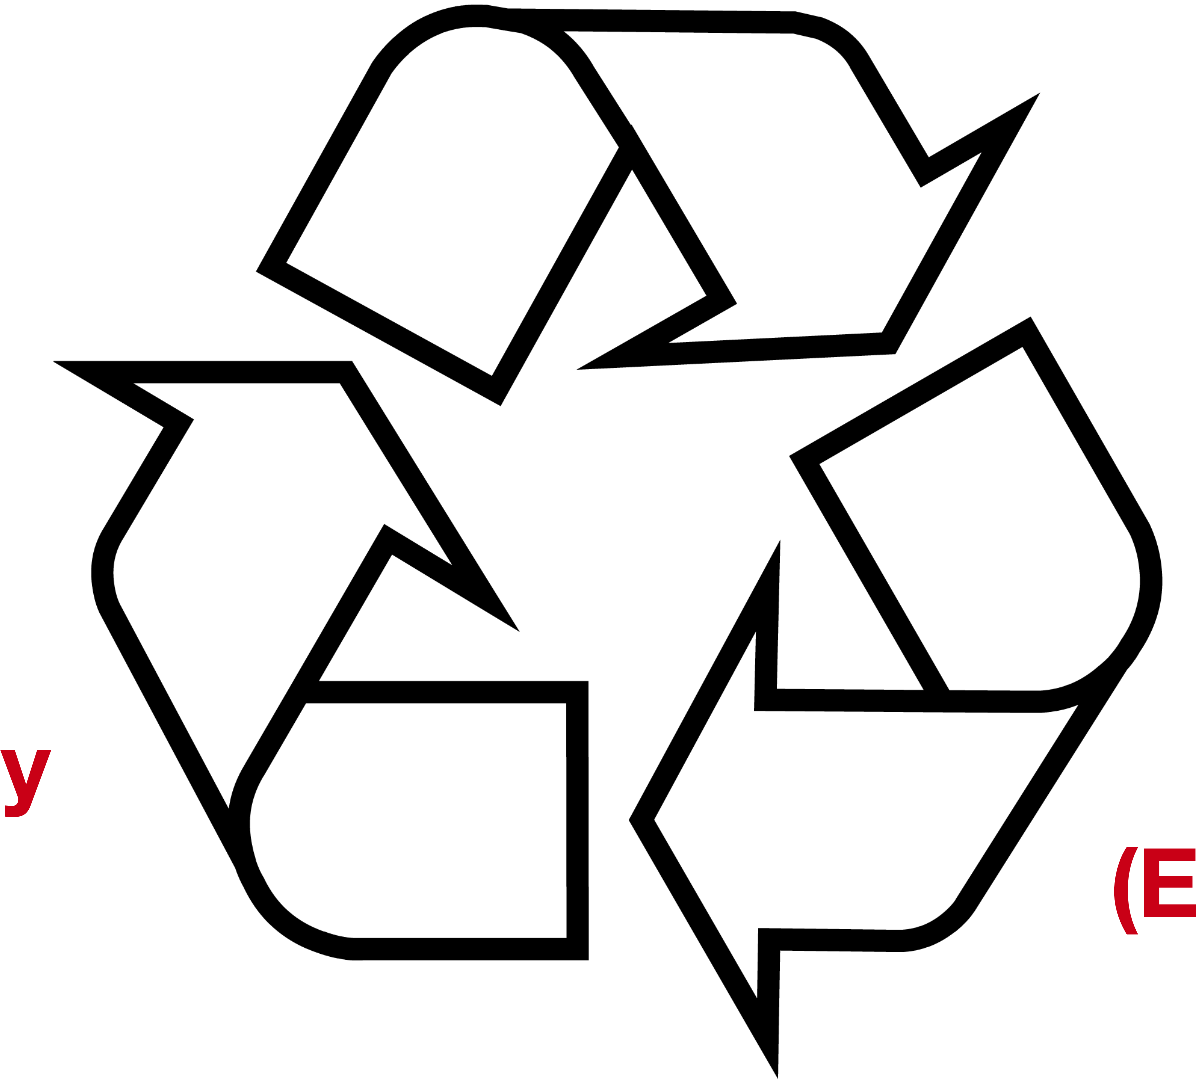

Reliability (Safety)

Power (Endurance)

RT

RT

Power Depends on the Physical Form Factor

87 kWh

Power Depends on the Physical Form Factor

87 kWh

Power Depends on the Physical Form Factor

87 kWh

0.08 kWh

Power Depends on the Physical Form Factor

87 kWh

1000X

0.08 kWh

Power Depends on the Physical Form Factor

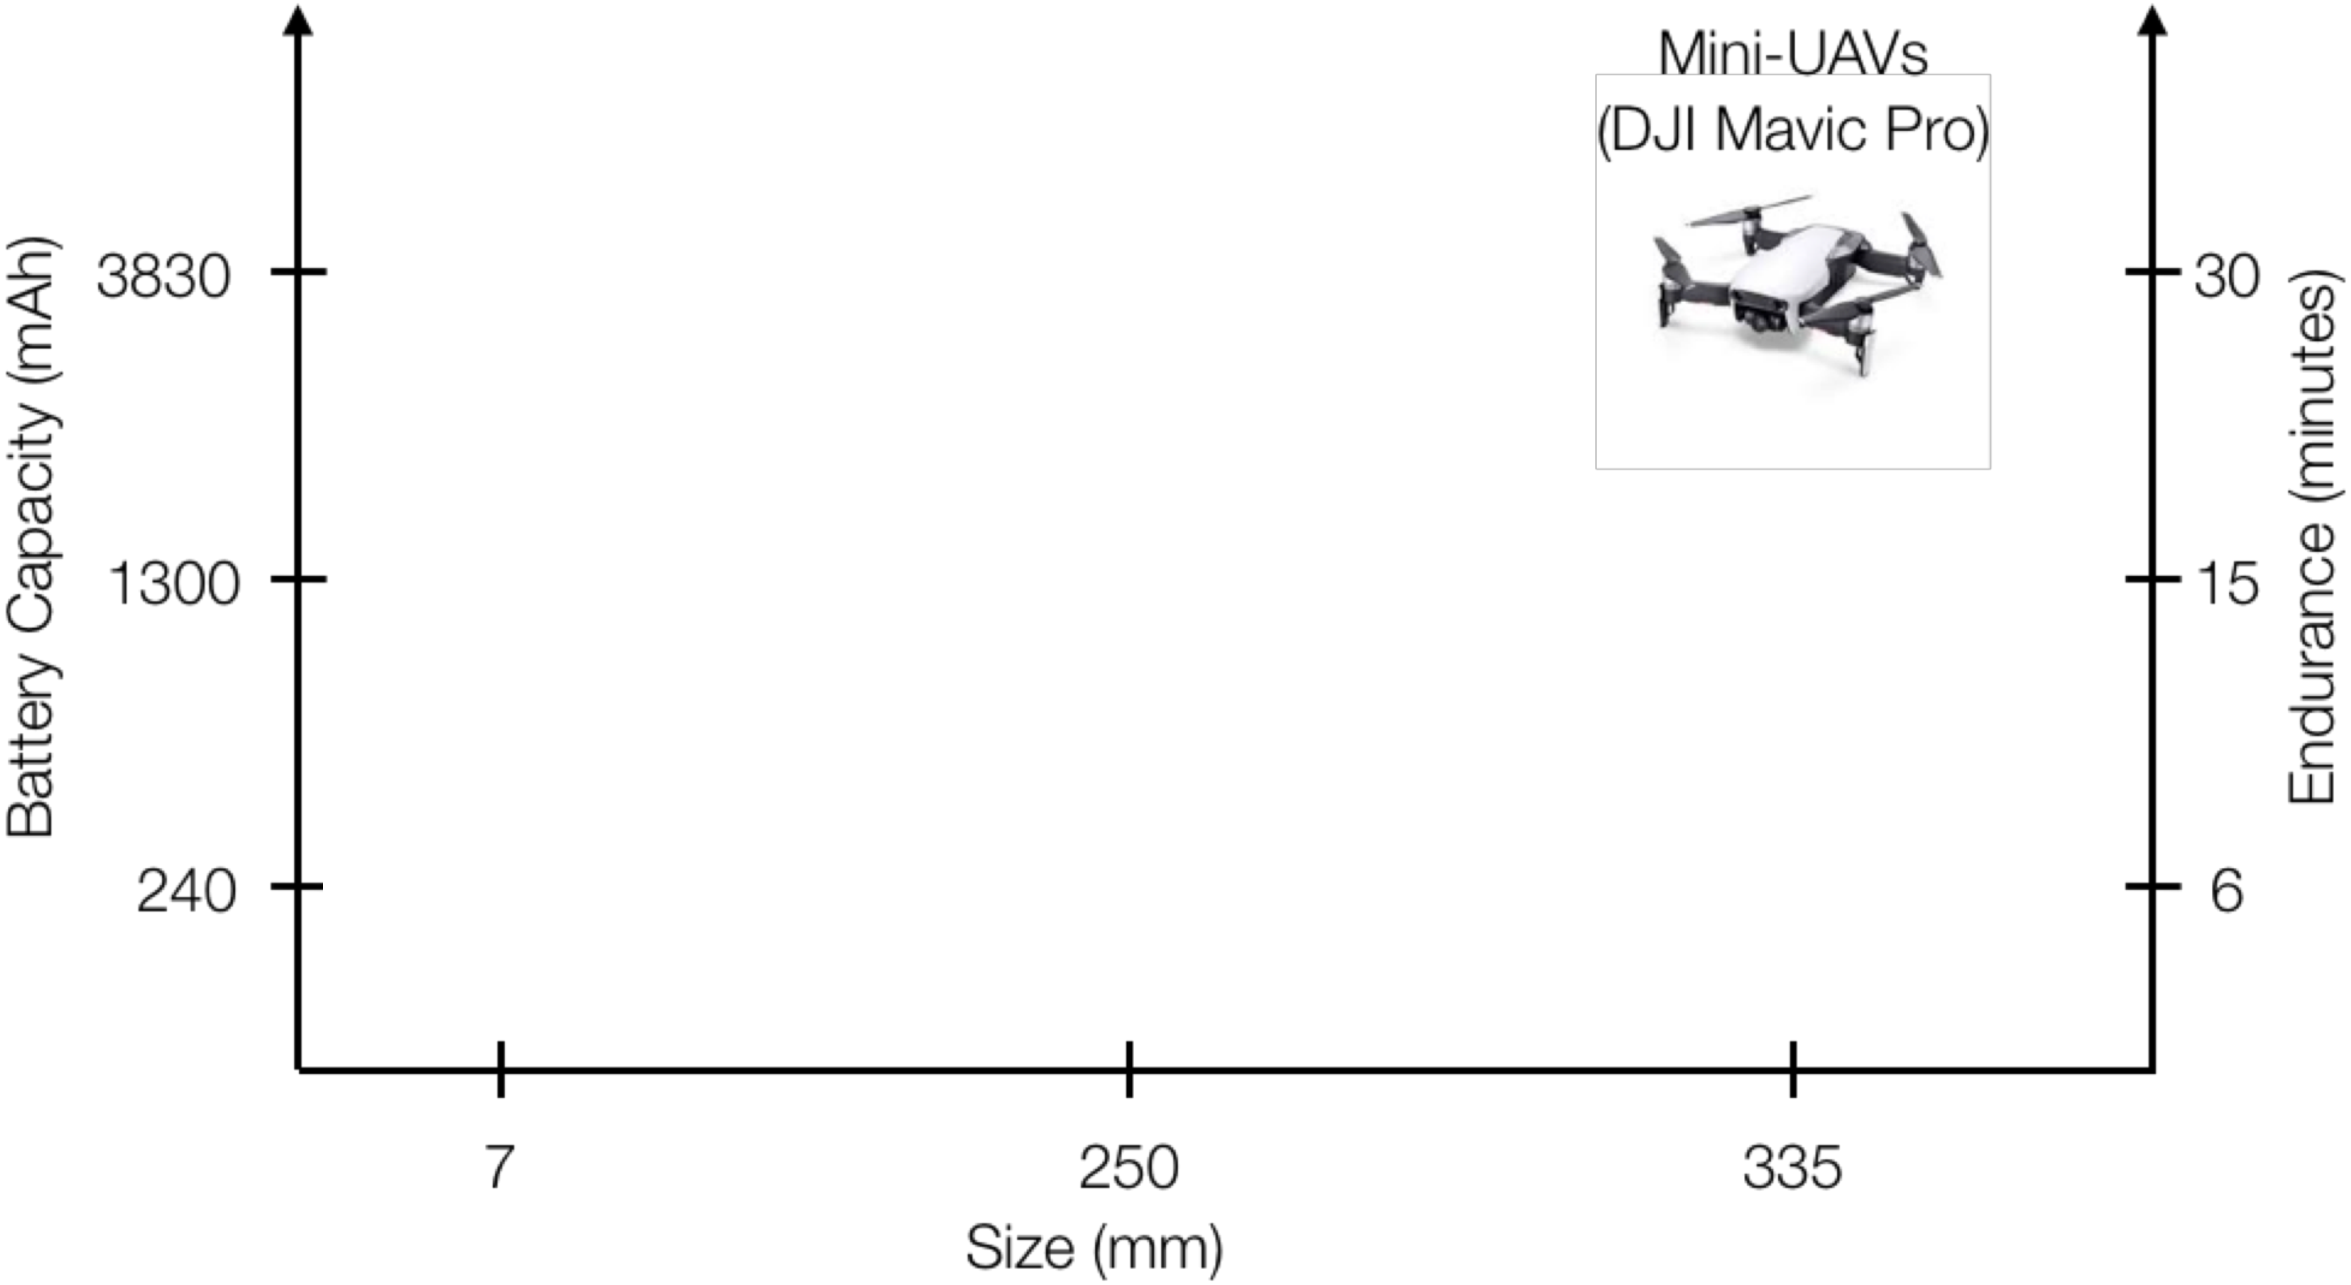

Power Depends on the Physical Form Factor

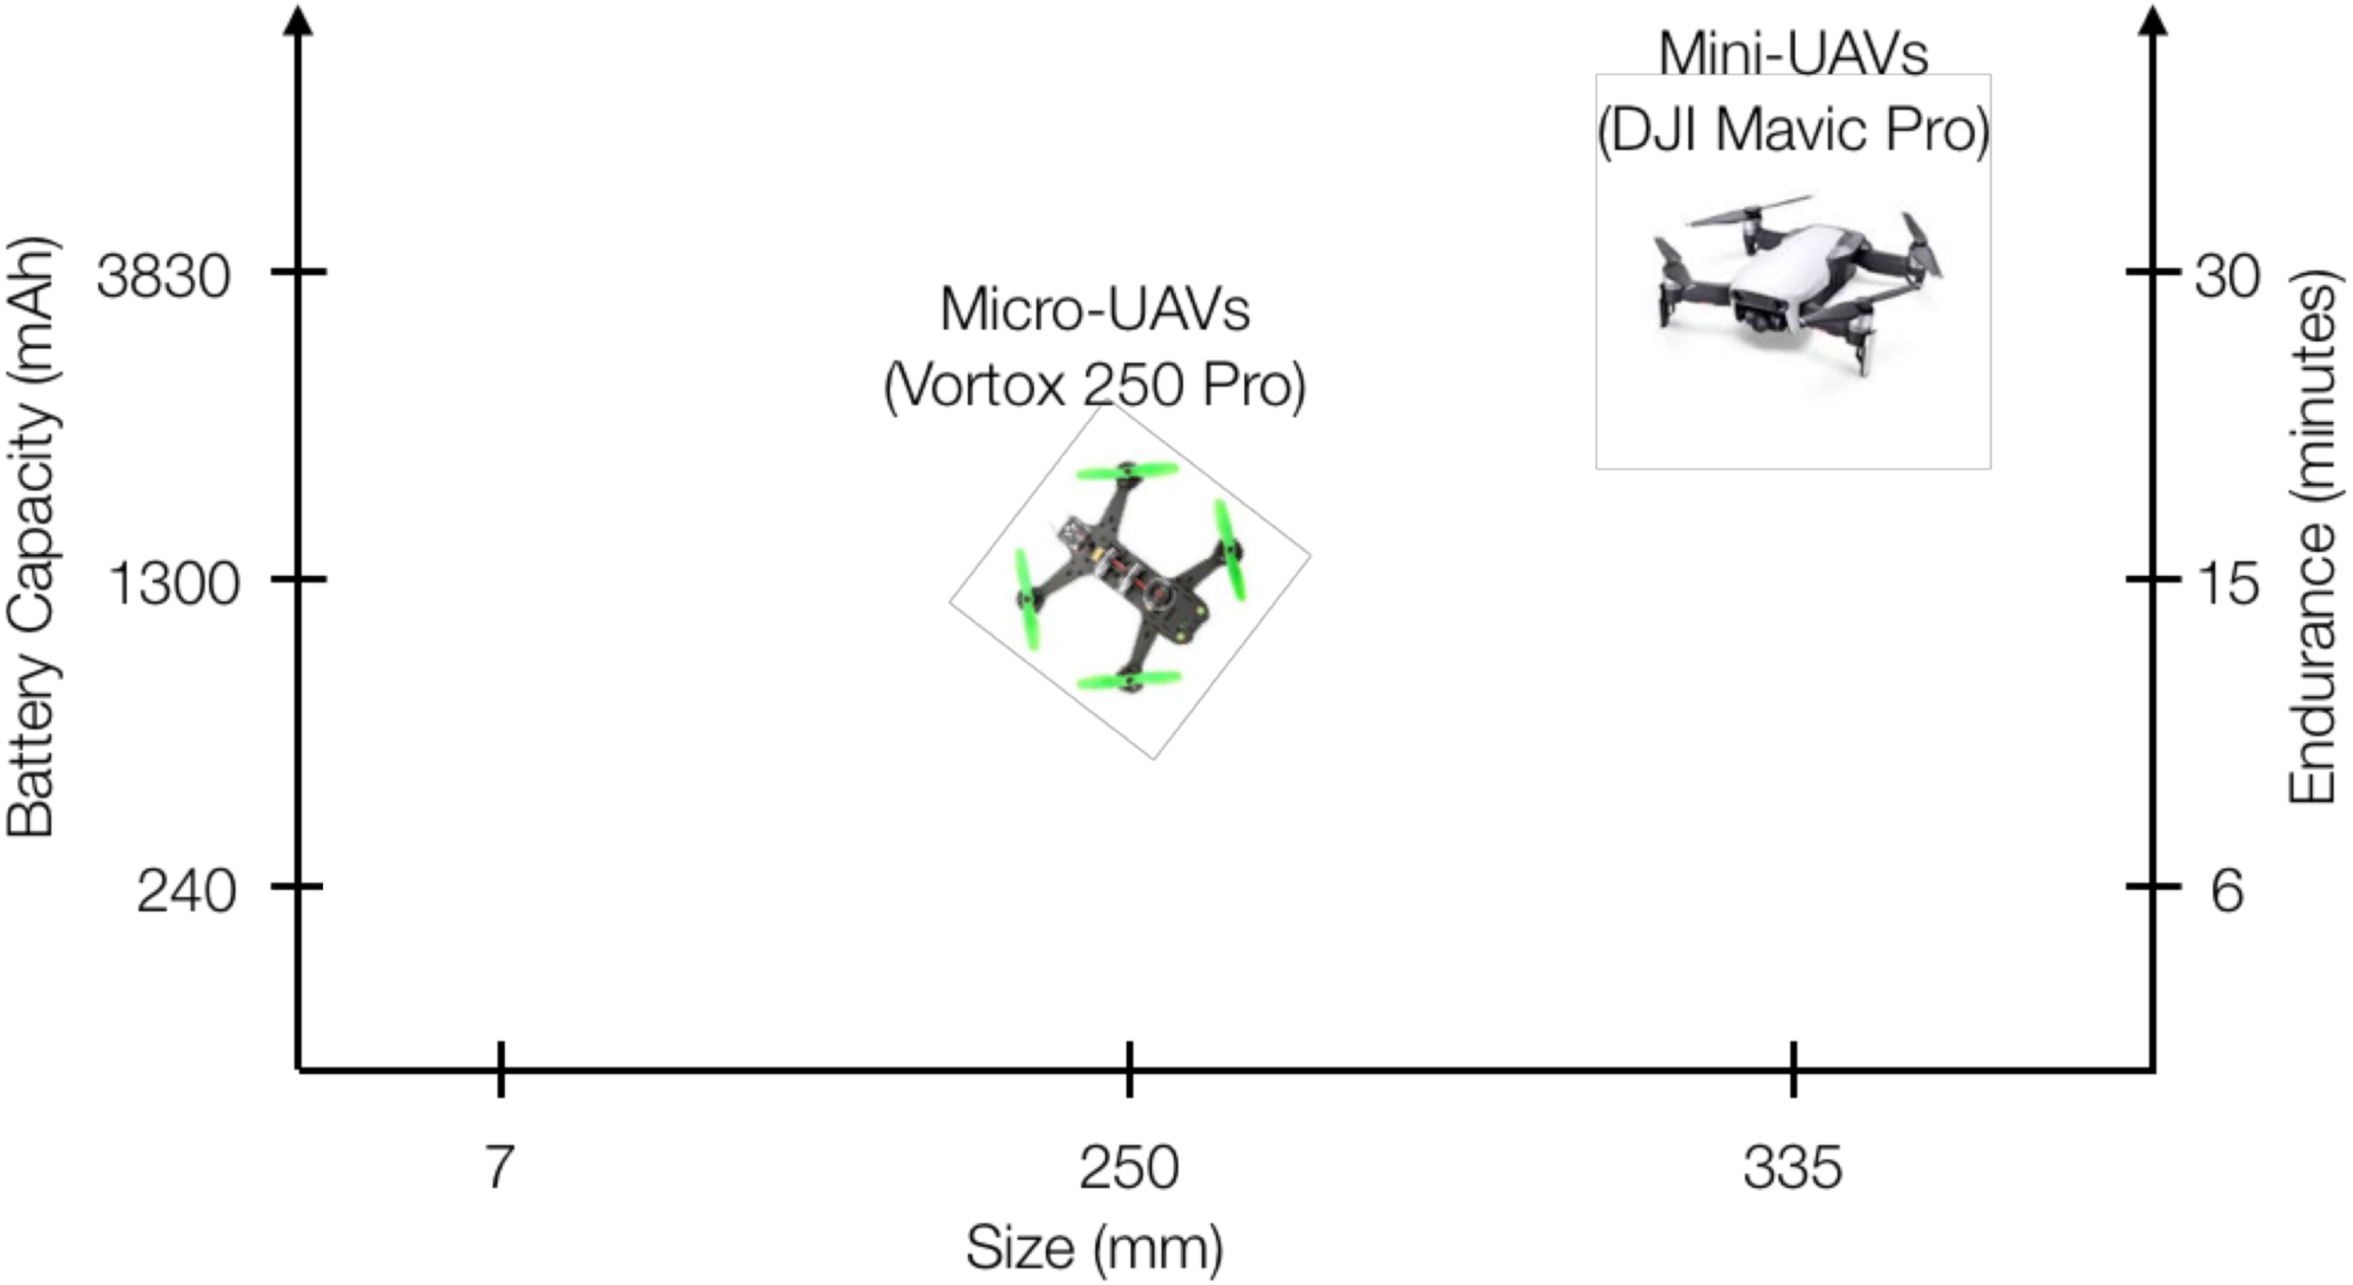

Power Depends on the Physical Form Factor

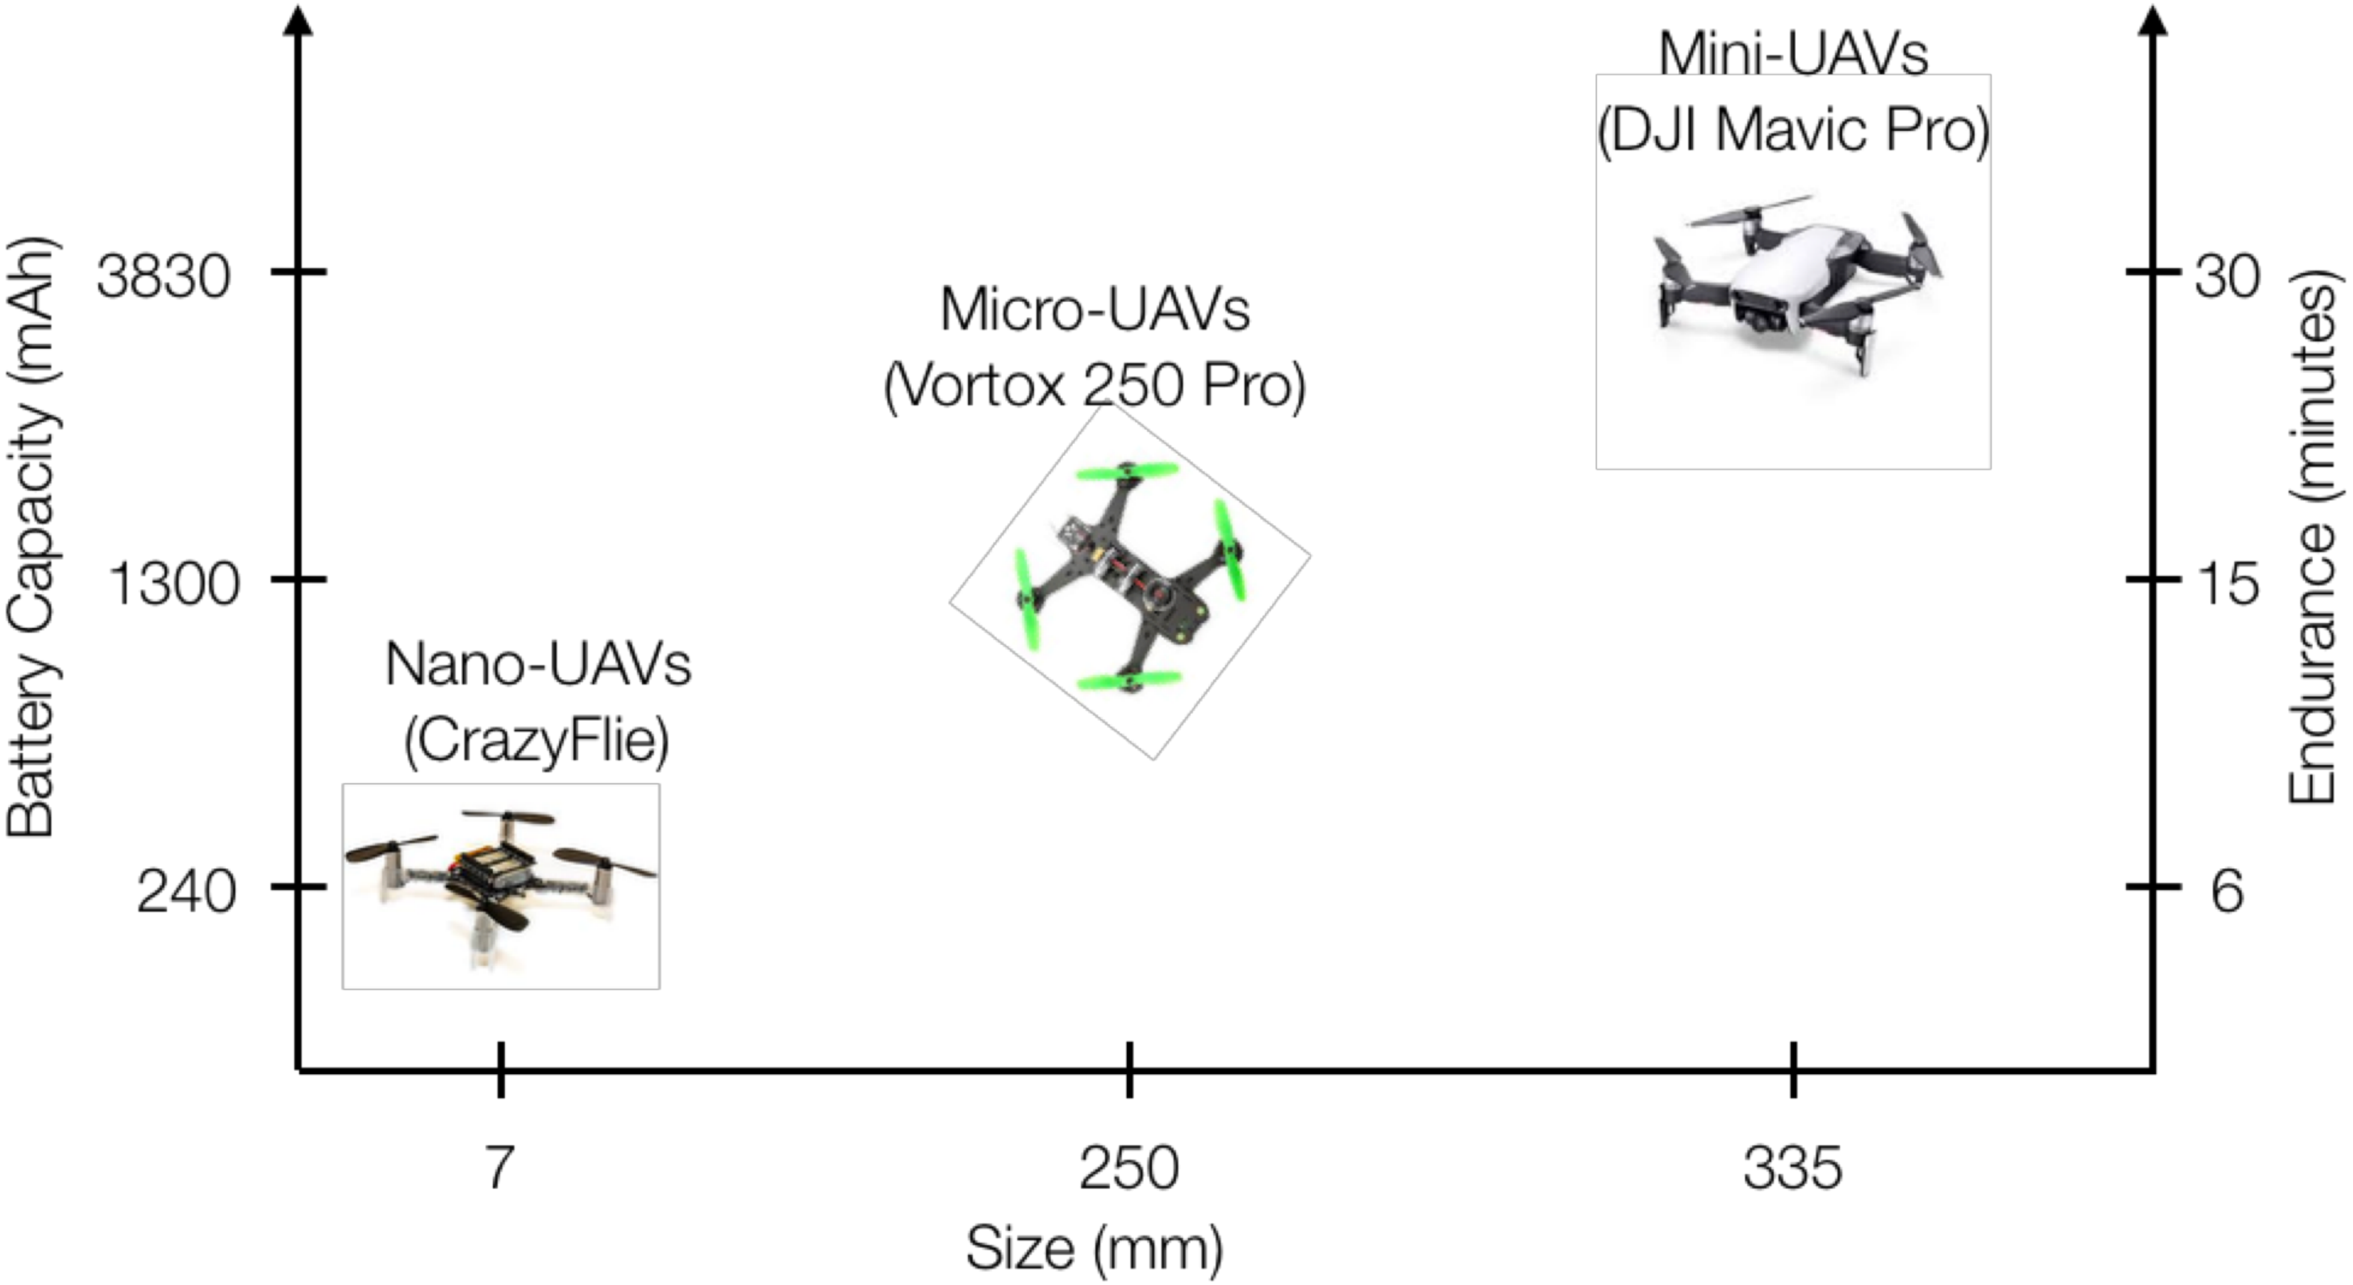

Safety Depends on Robustly Accounting for Failures

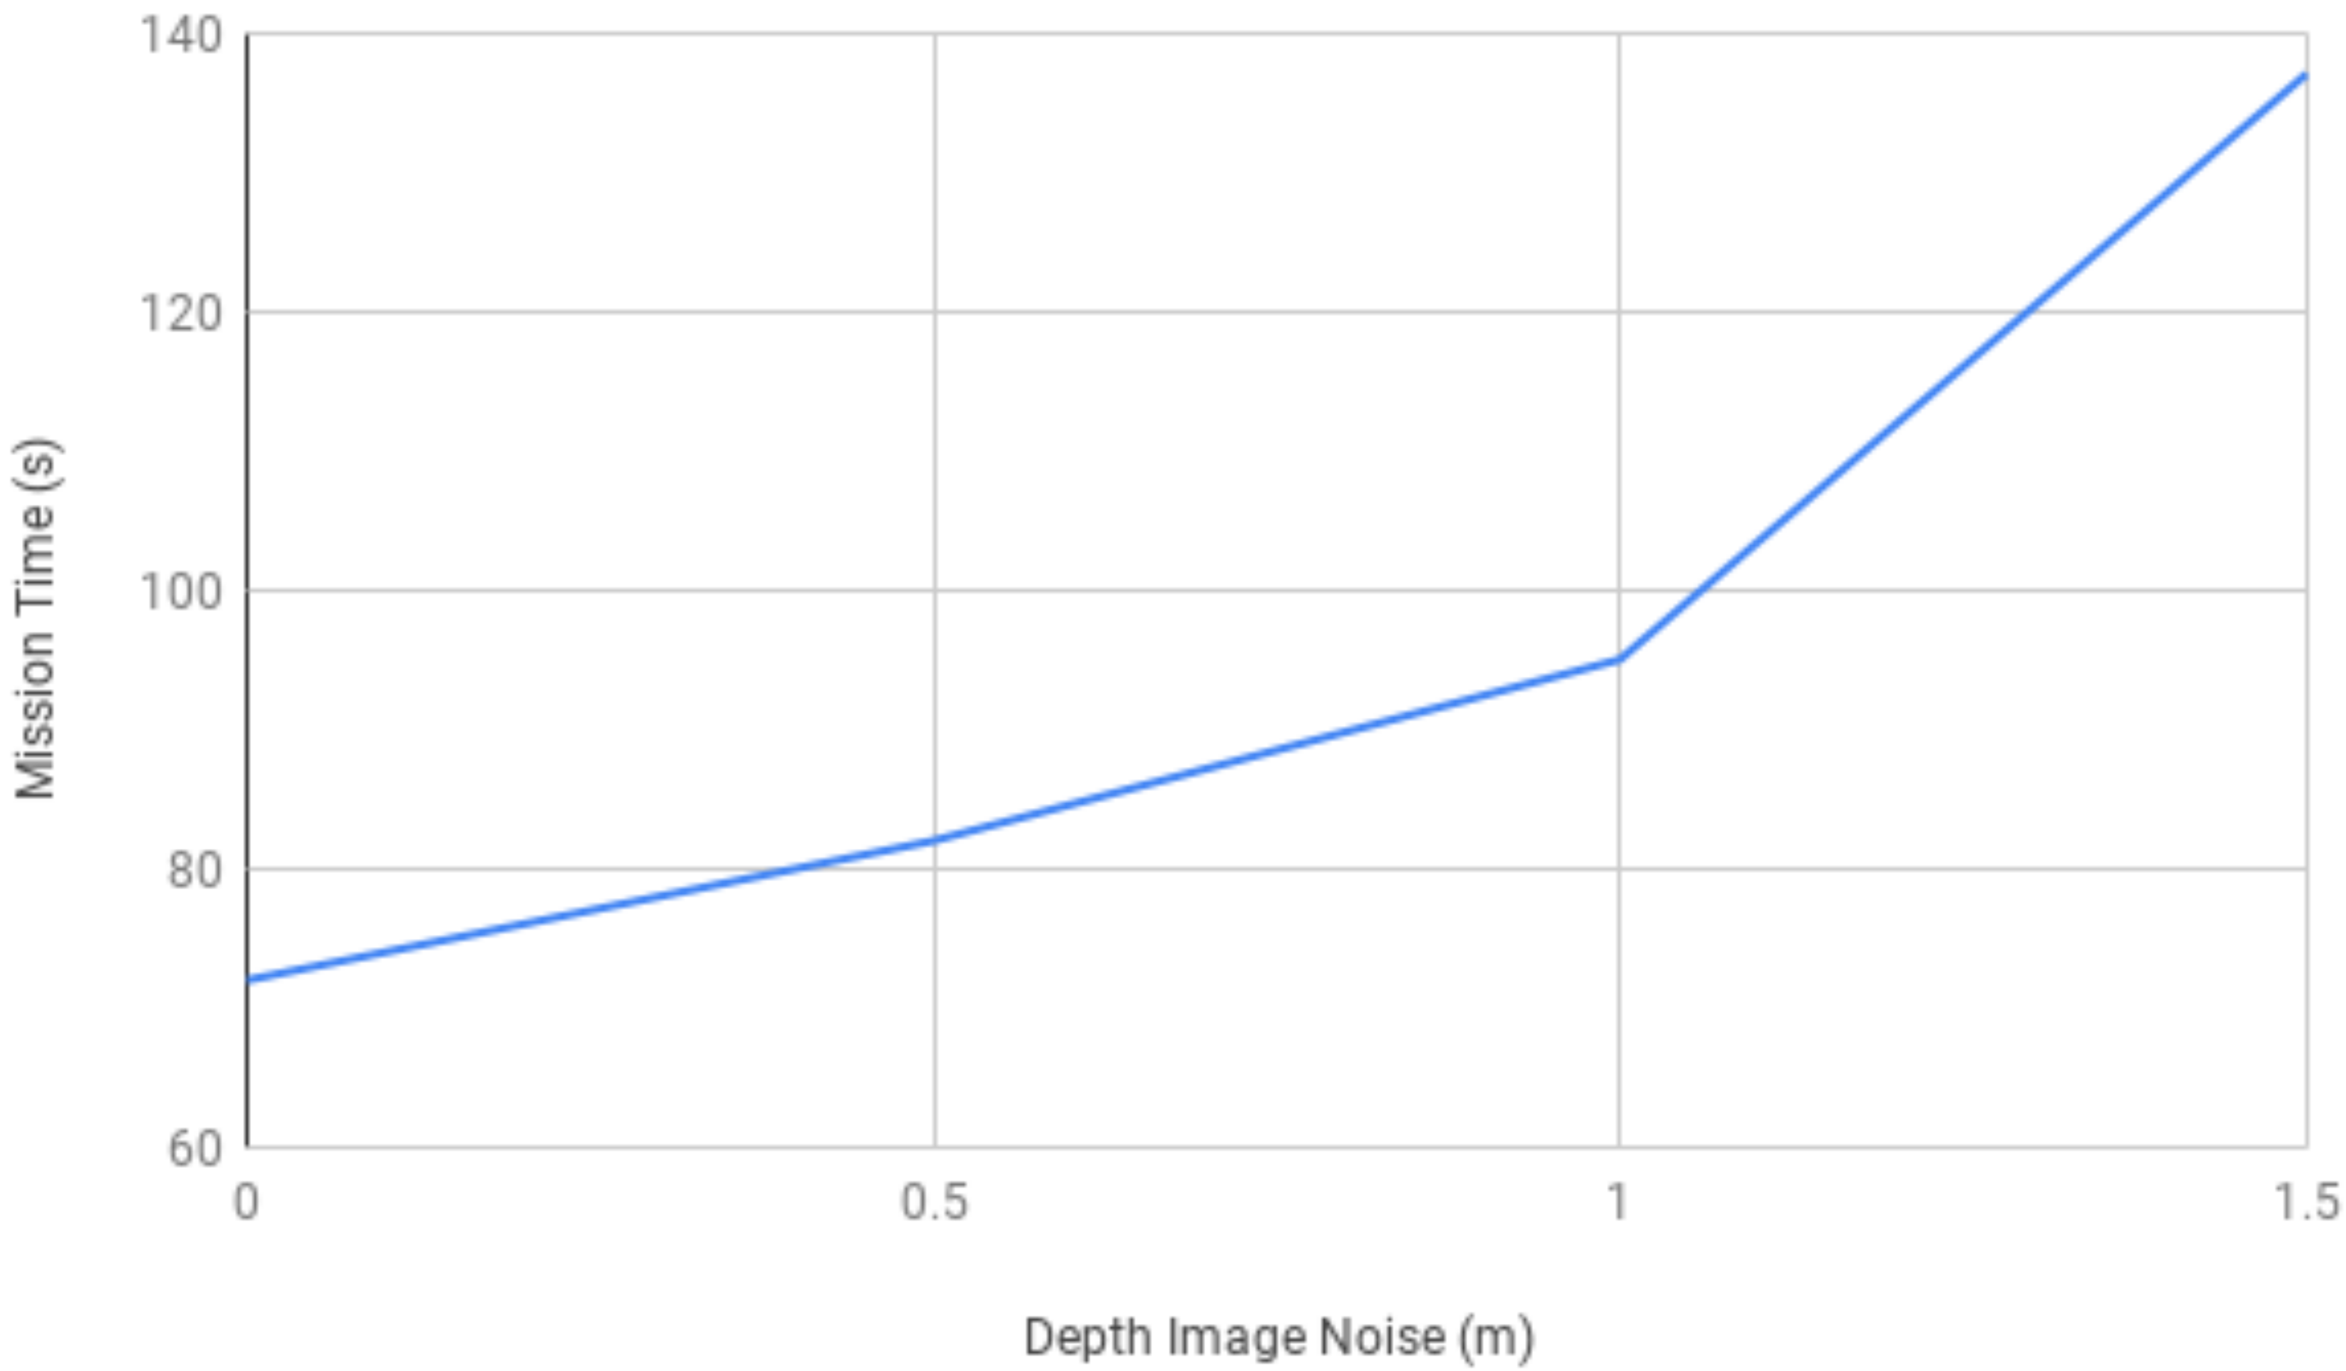

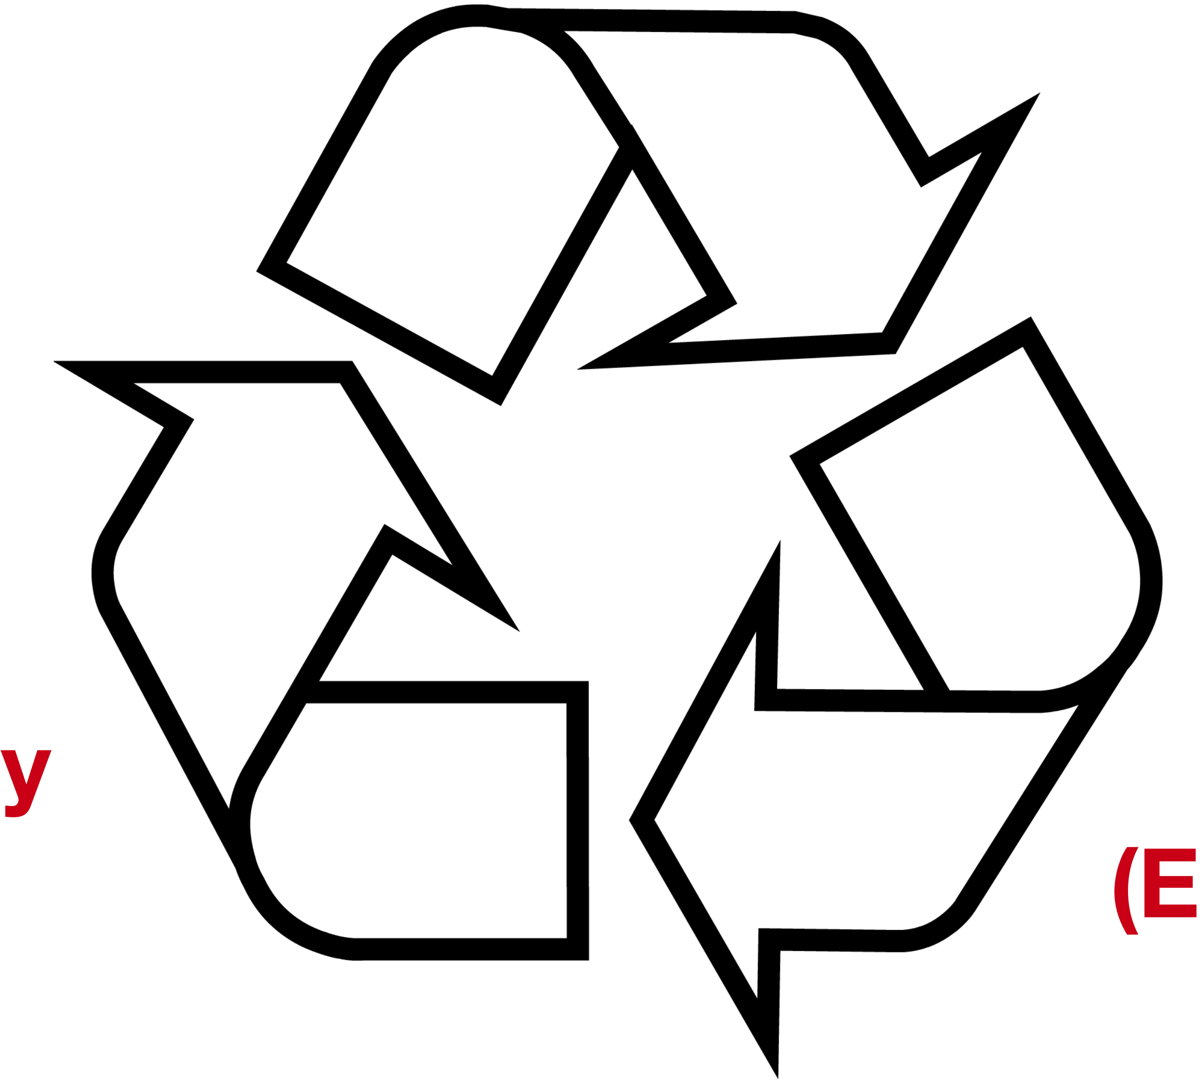

Performance
(Real-time)

Reliability
(Safety)

Power
(Endurance)

What **architectural tools** are needed to enable research on aerial robots?

Traditional Computer Architecture Toolkit

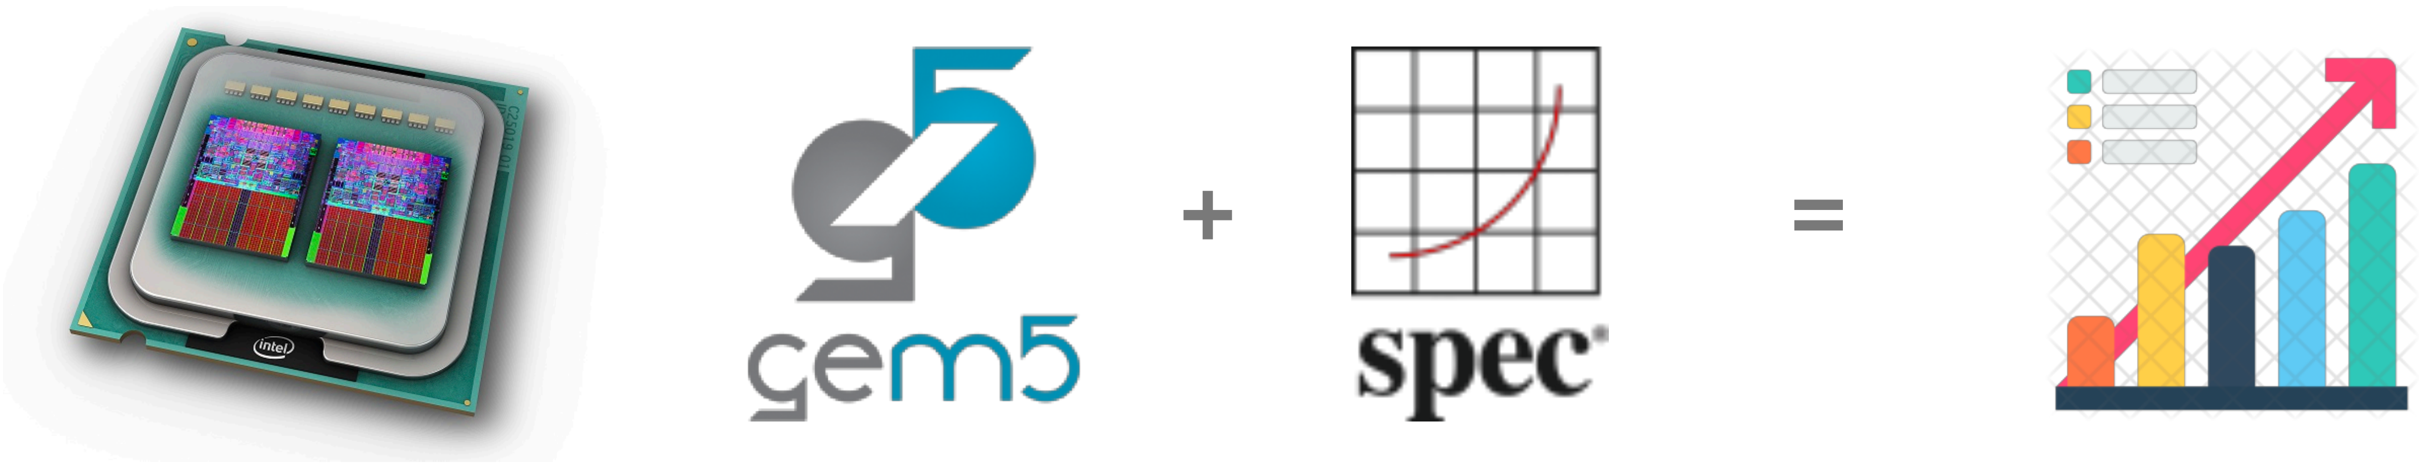

Challenge with the Traditional Toolkit for Robotics

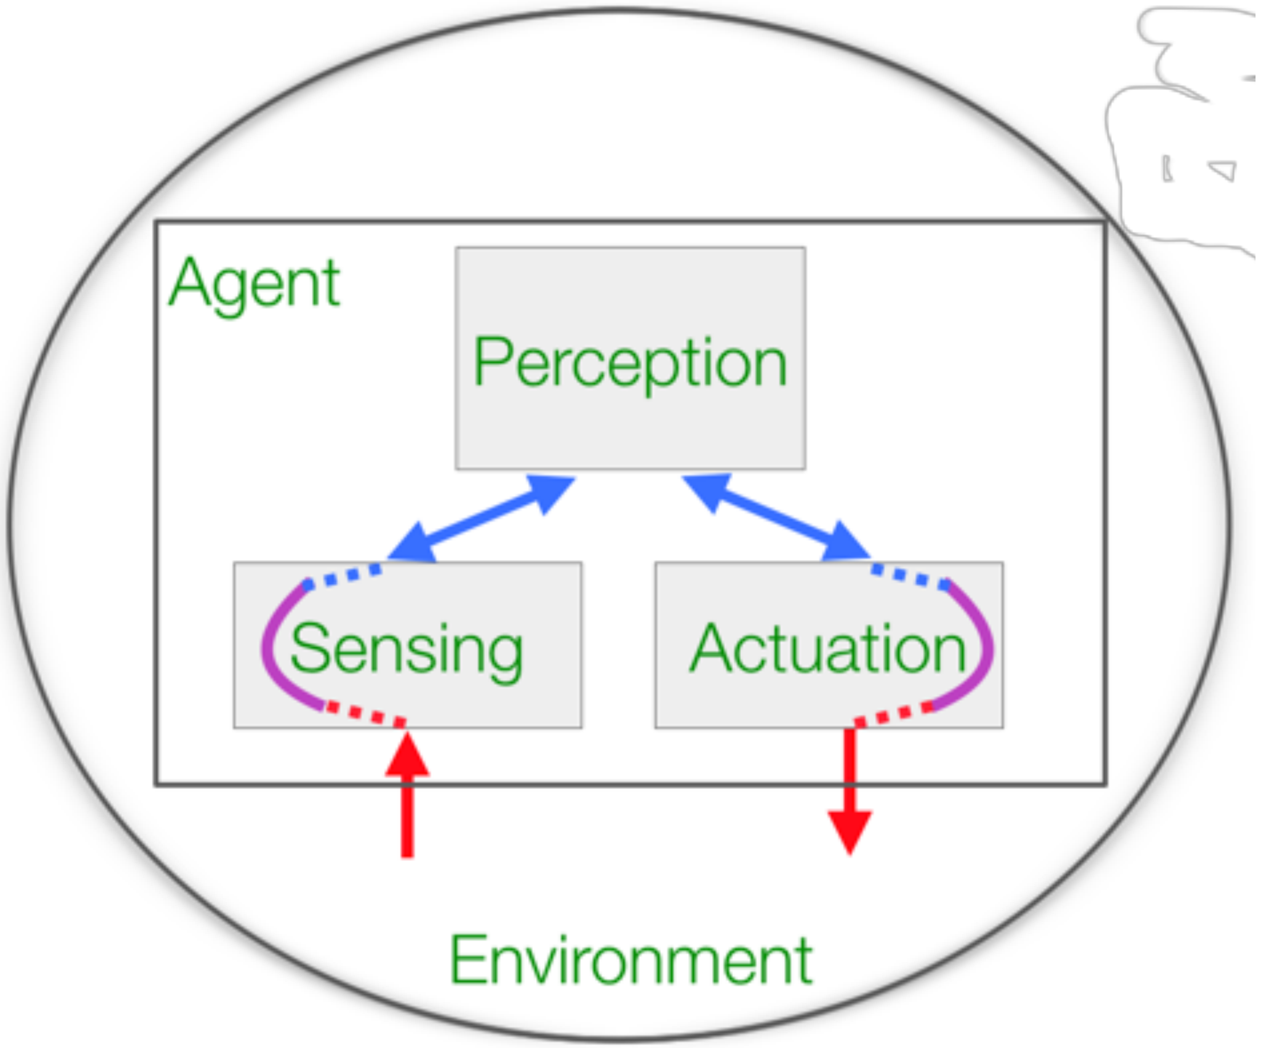

Challenge with the Traditional Toolkit for Robotics

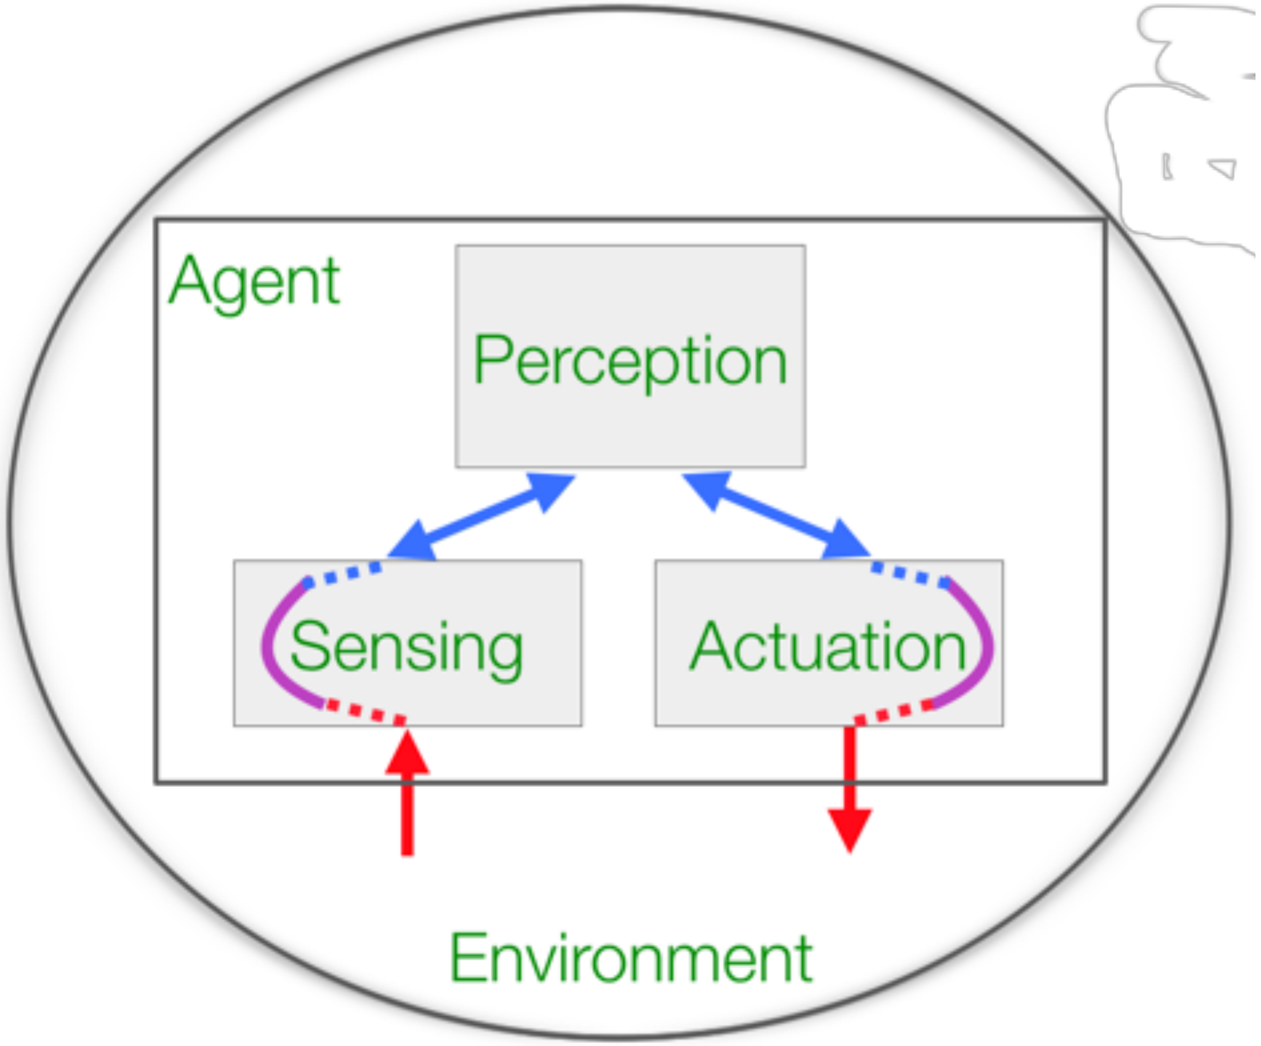

There is a continuous feedback loop.

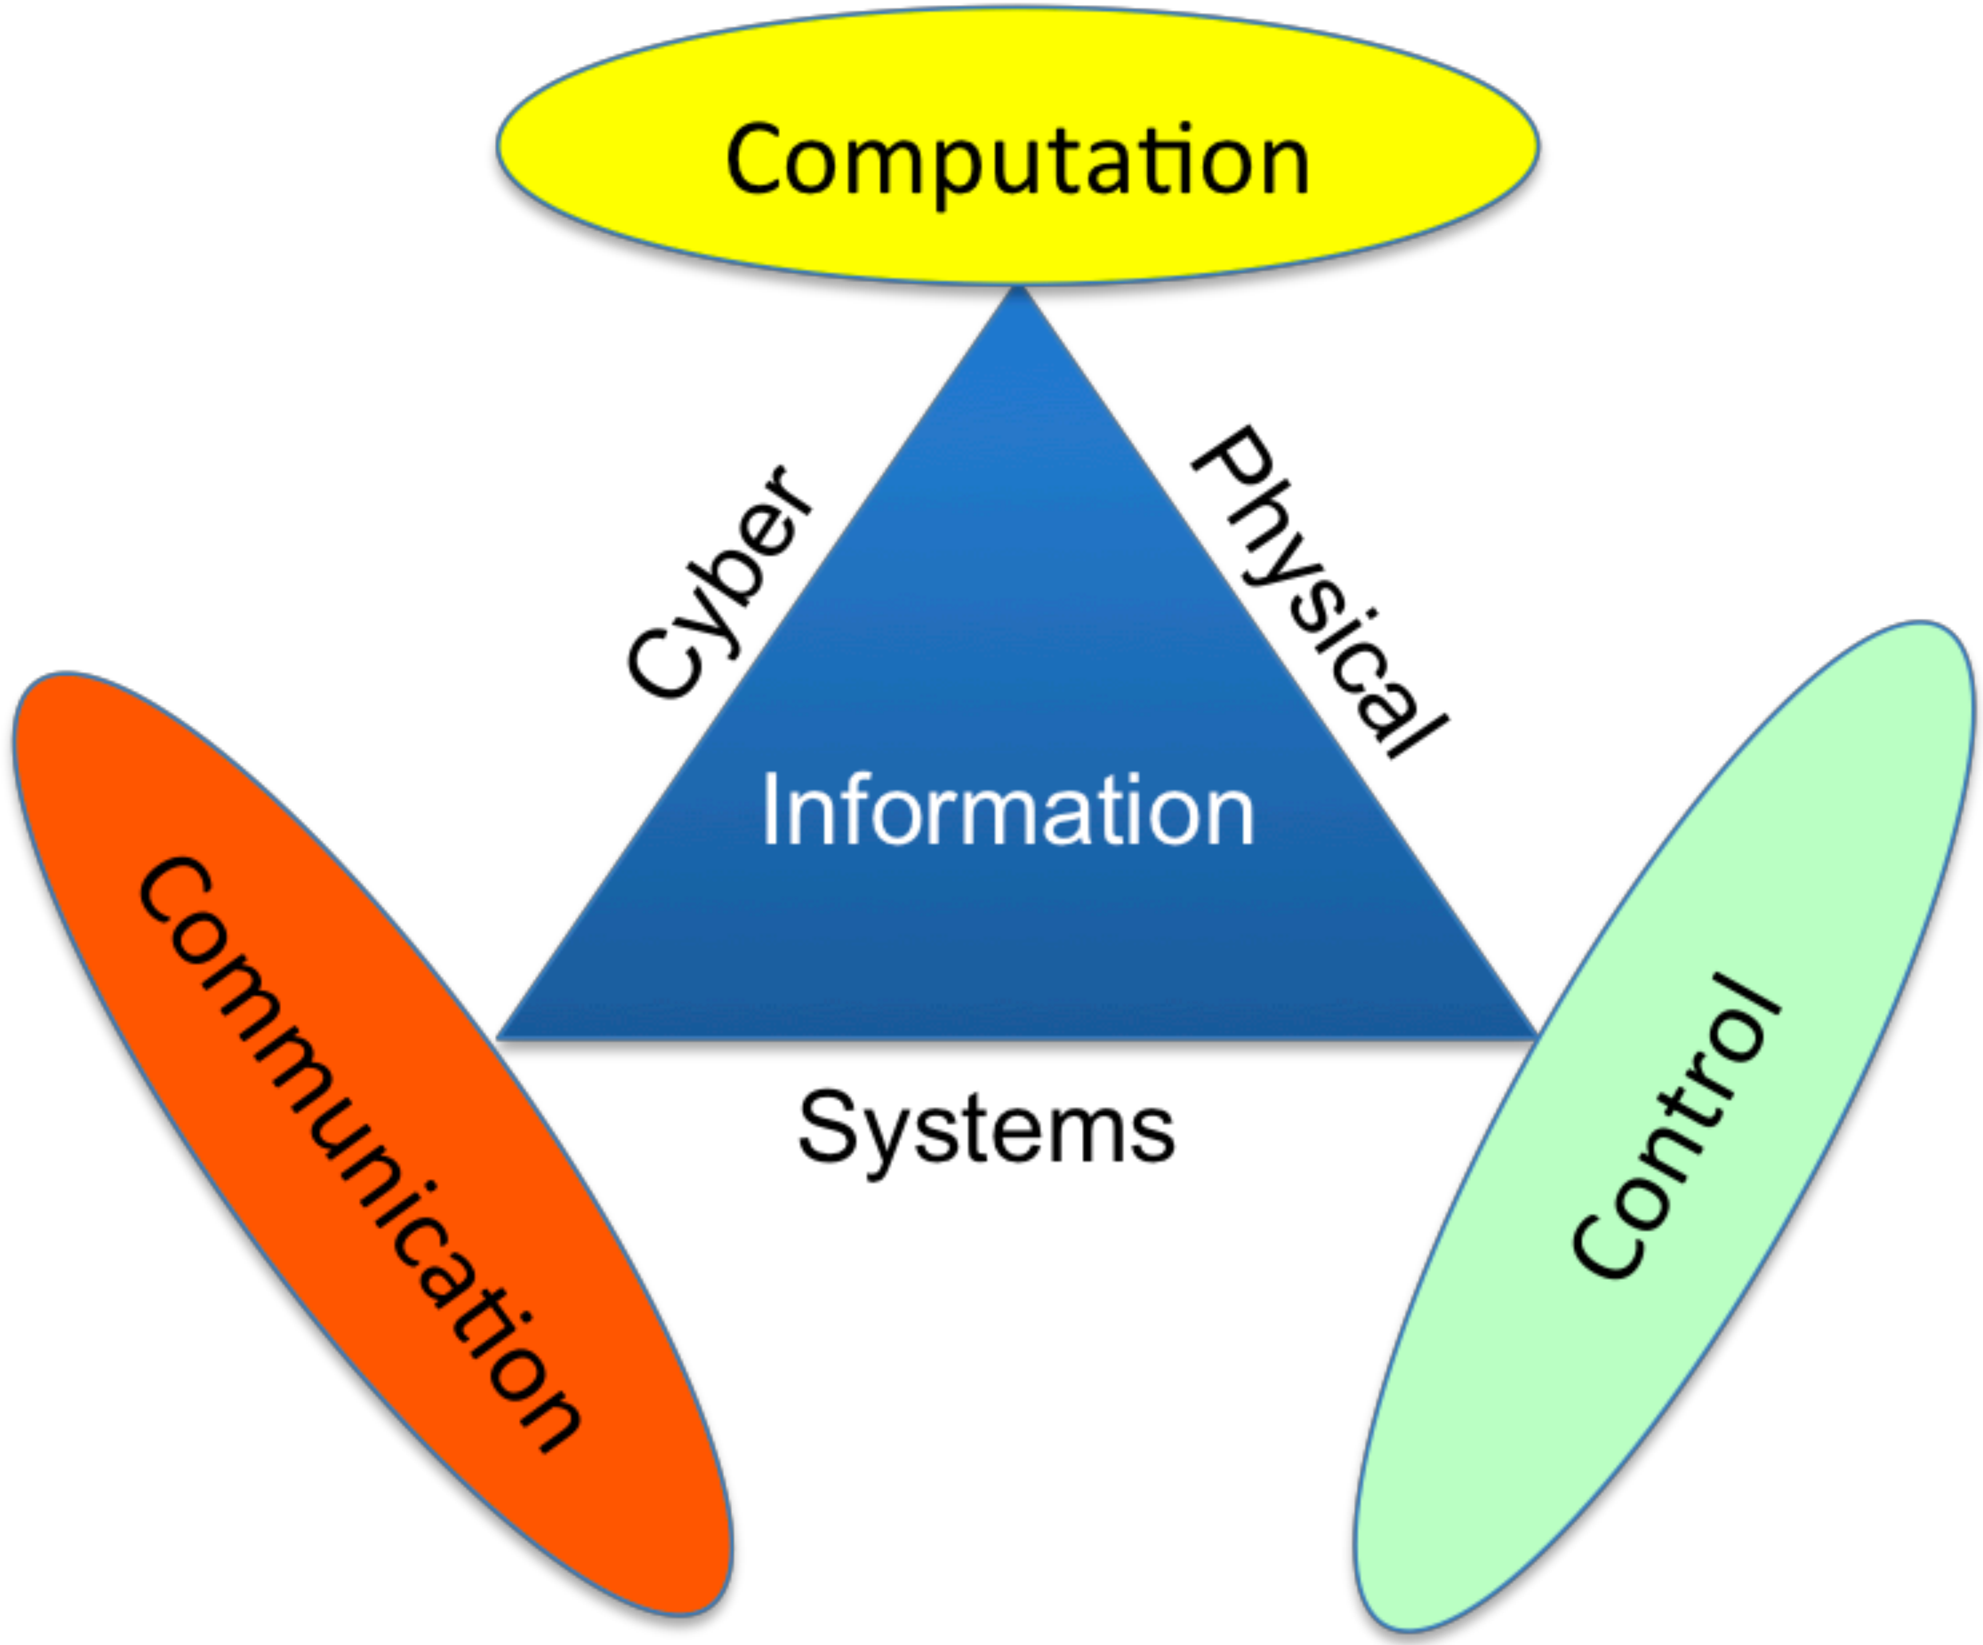
Computation
Cyber
Physical
Information
Communication
Systems
Control

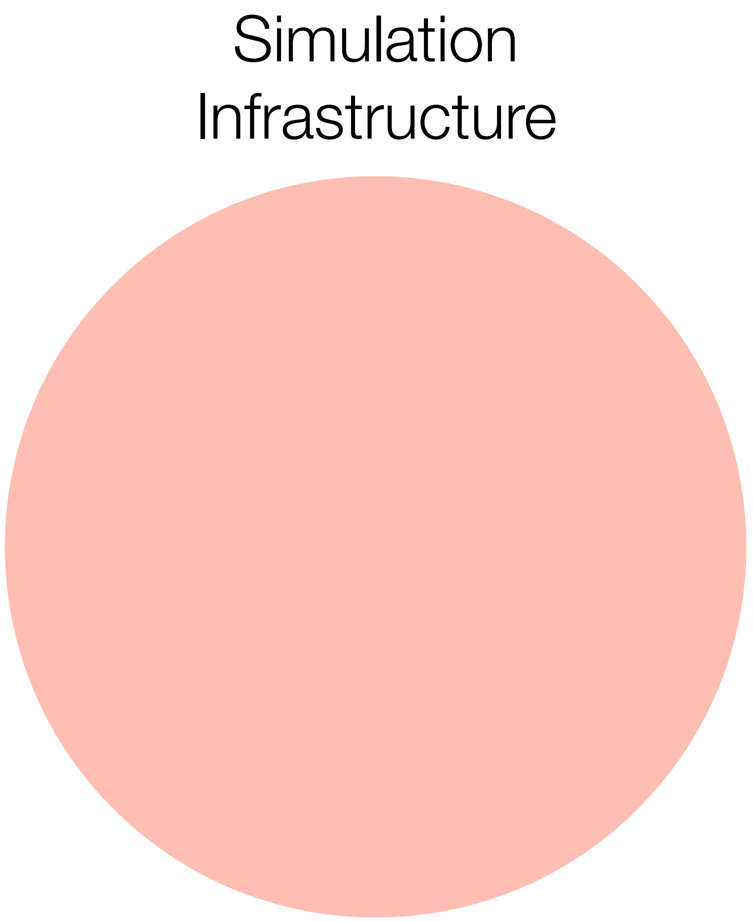

Simulation
Infrastructure

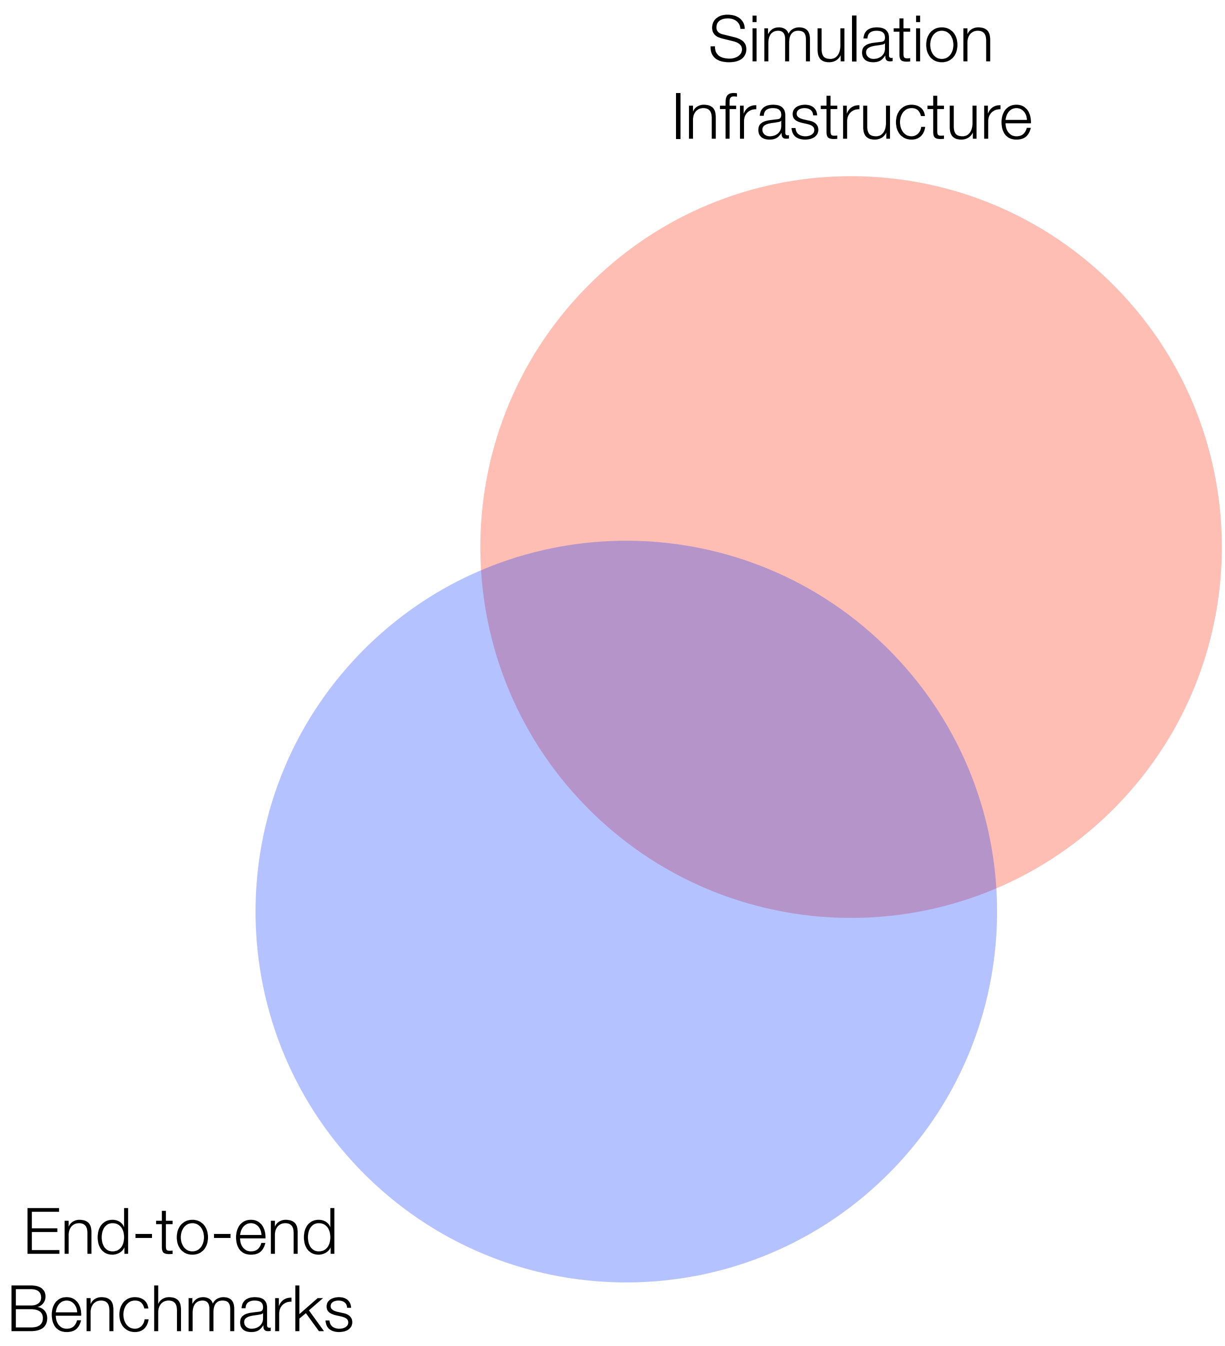
Simulation
Infrastructure
End-to-end
Benchmarks

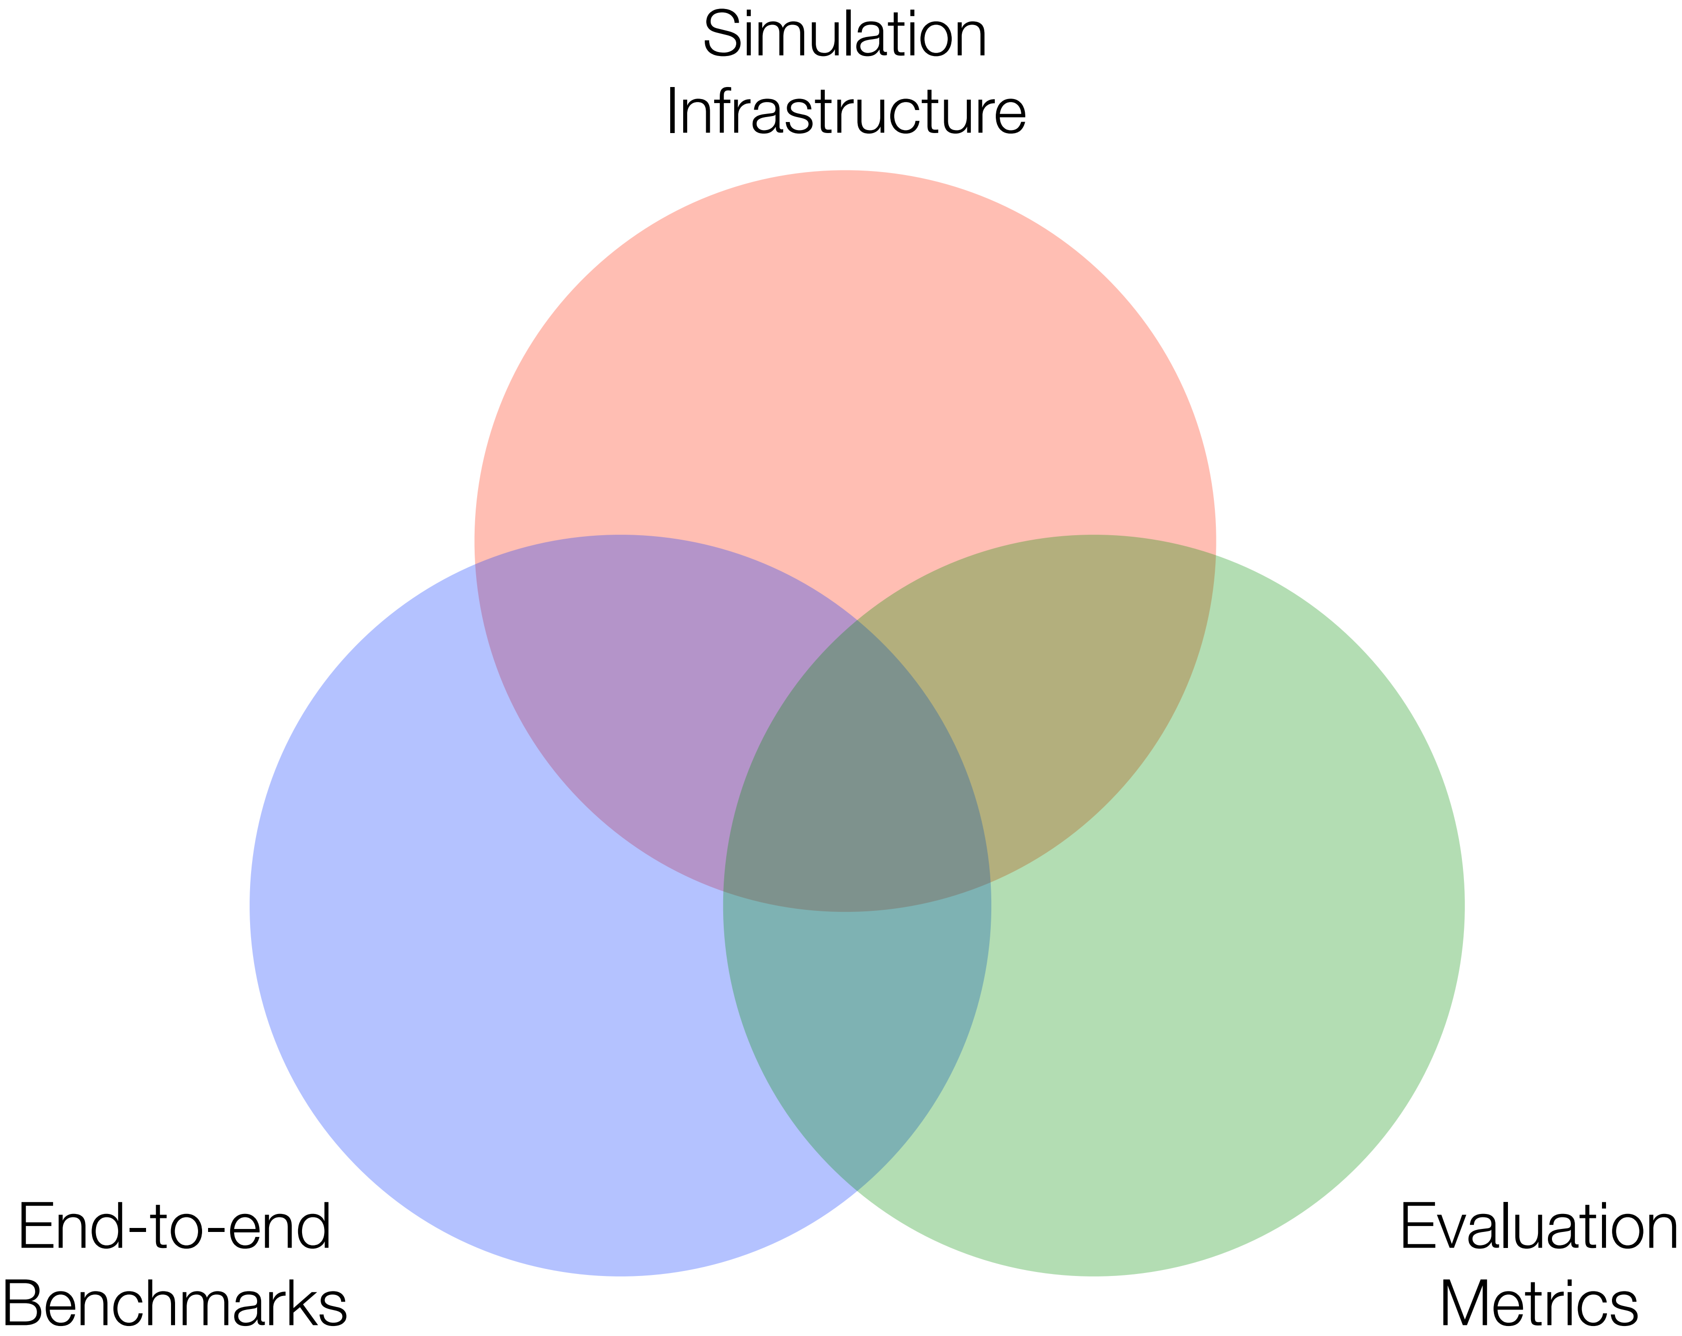

Simulation Infrastructure

End-to-end Benchmarks

Evaluation Metrics

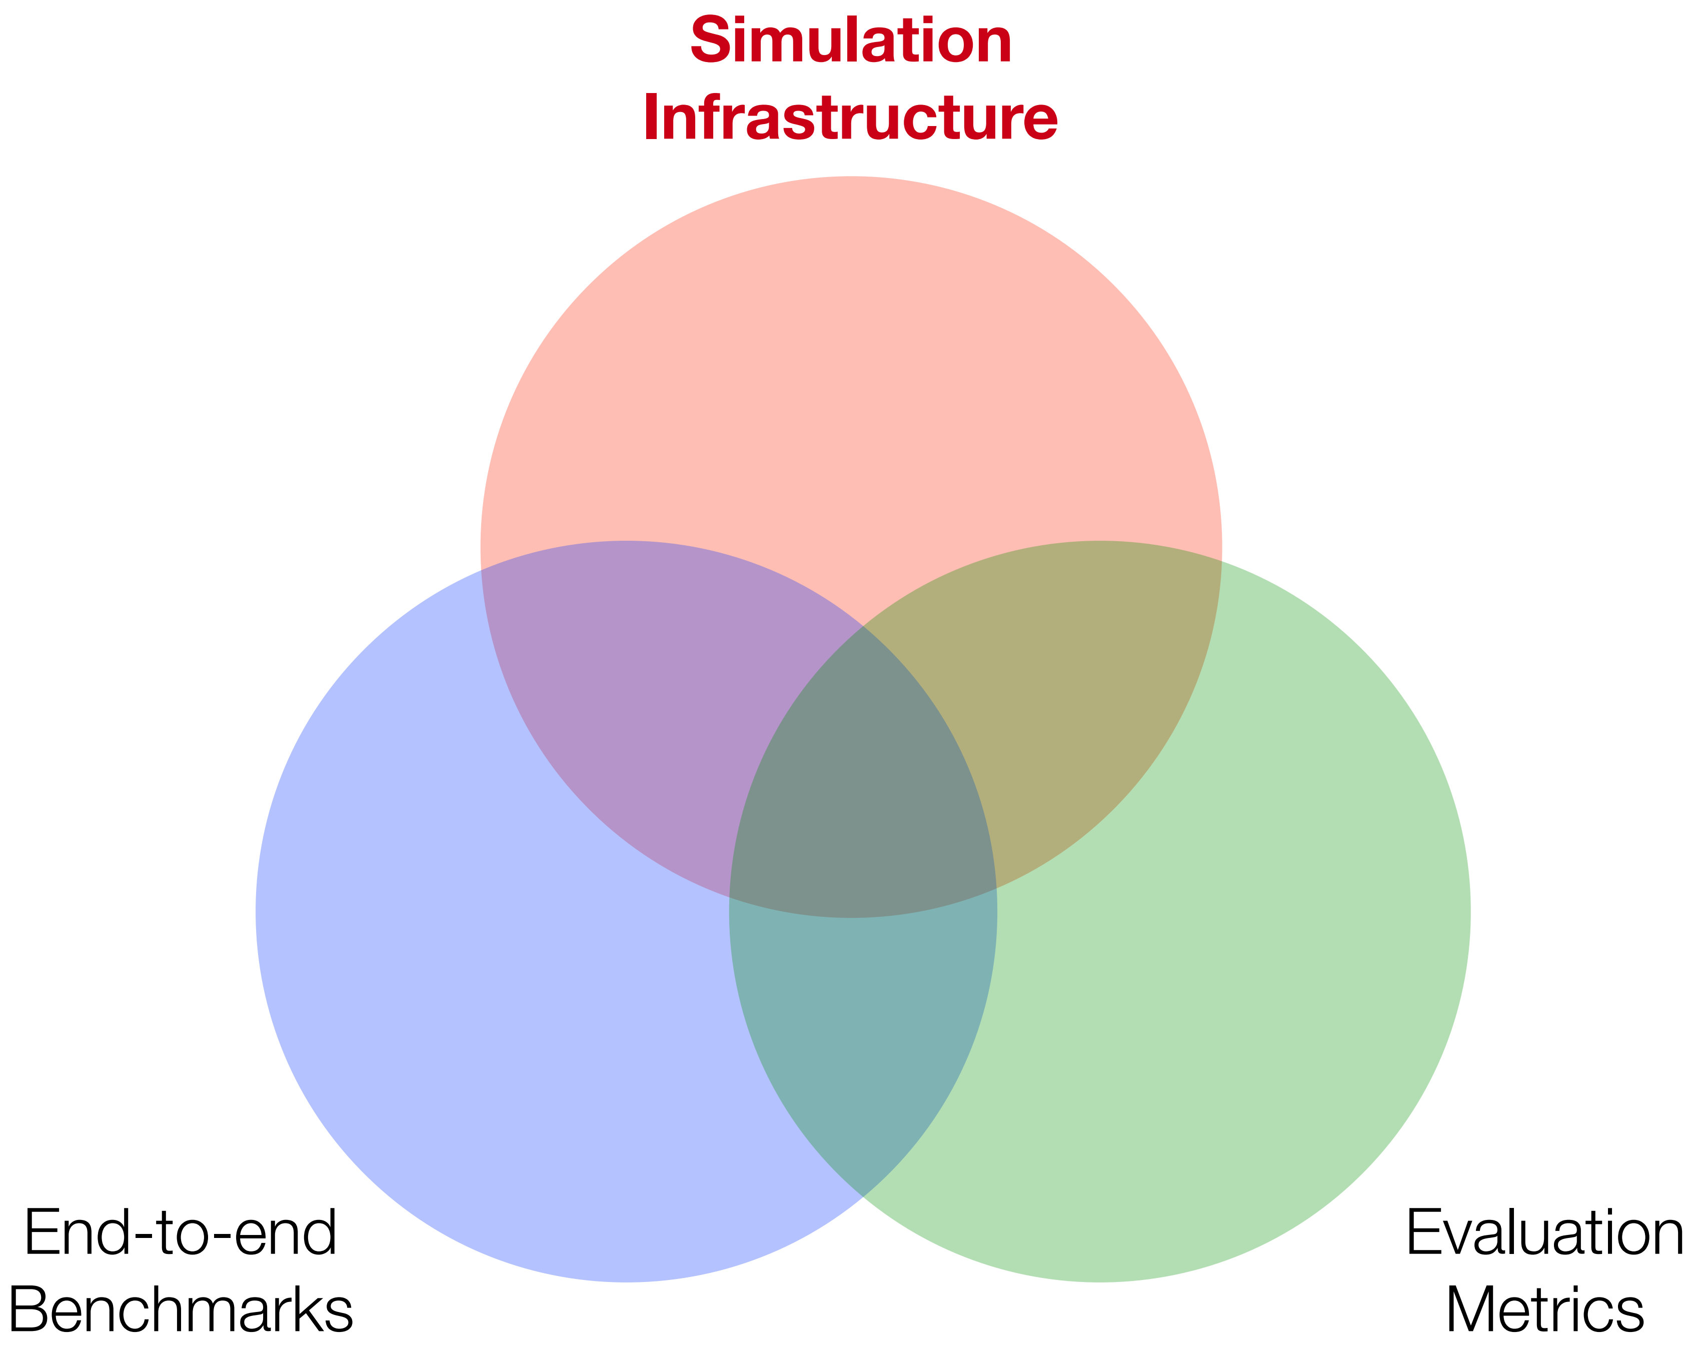

Simulation Infrastructure

End-to-end Benchmarks

Evaluation Metrics

The “World” of an Aerial Agent

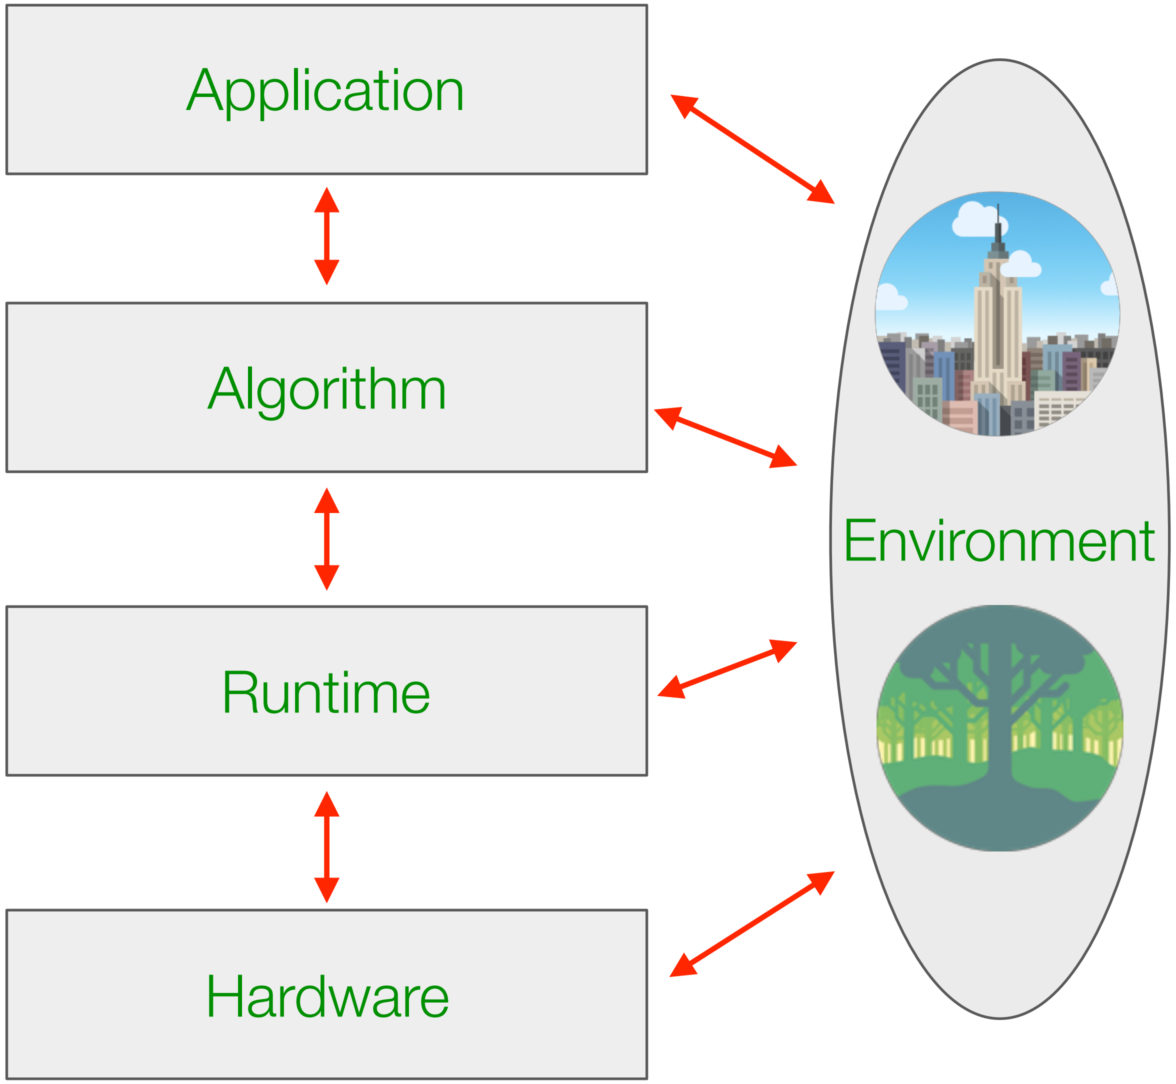

Understanding and Capturing the Interactions

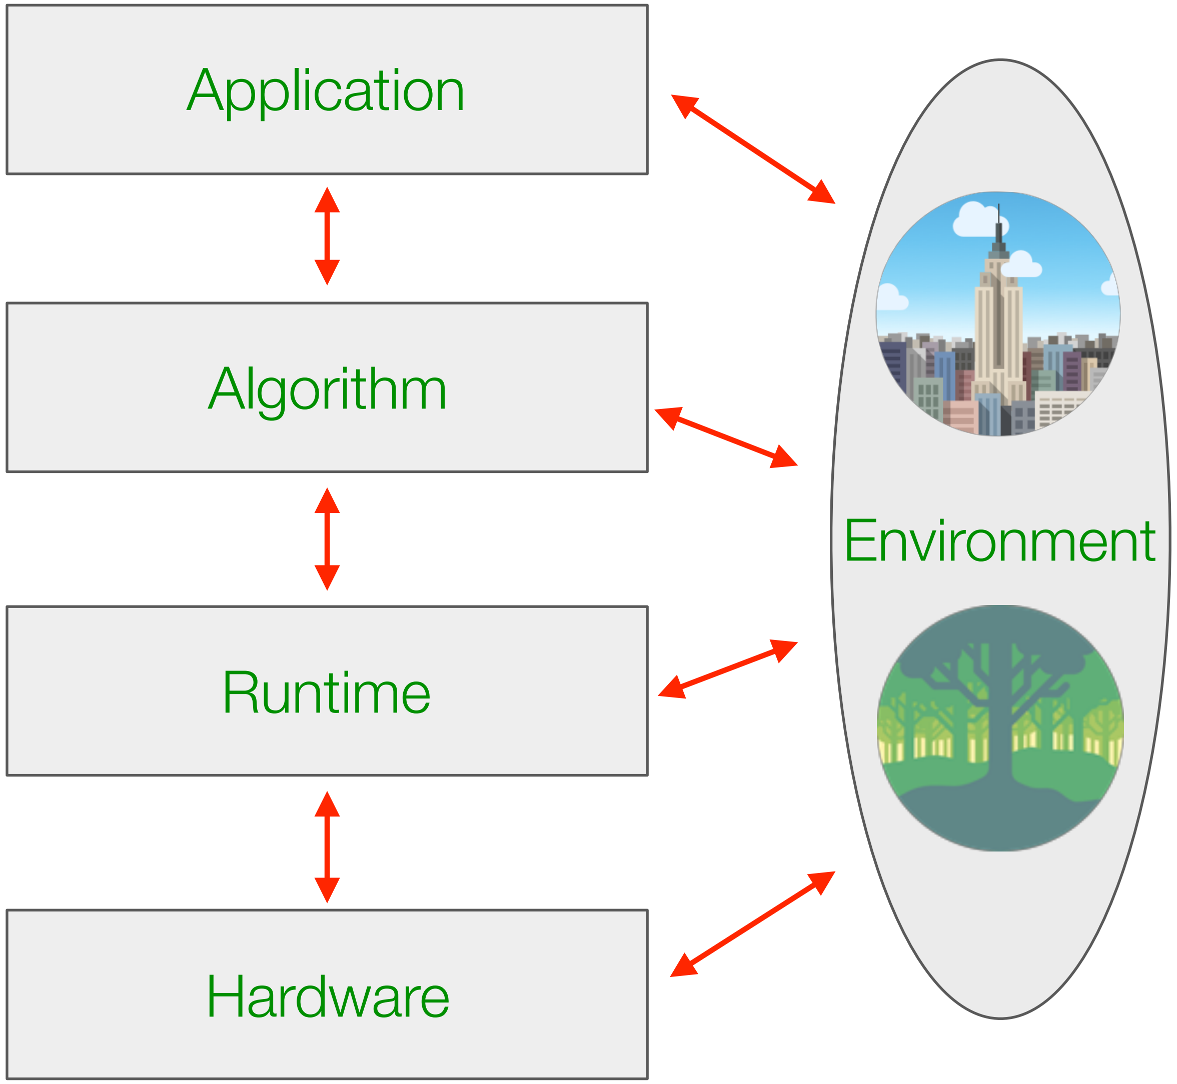

Understanding and Capturing the Interactions

- **Ideally, each level should be readily explorable…**
  - Independently
  - In relation to one another

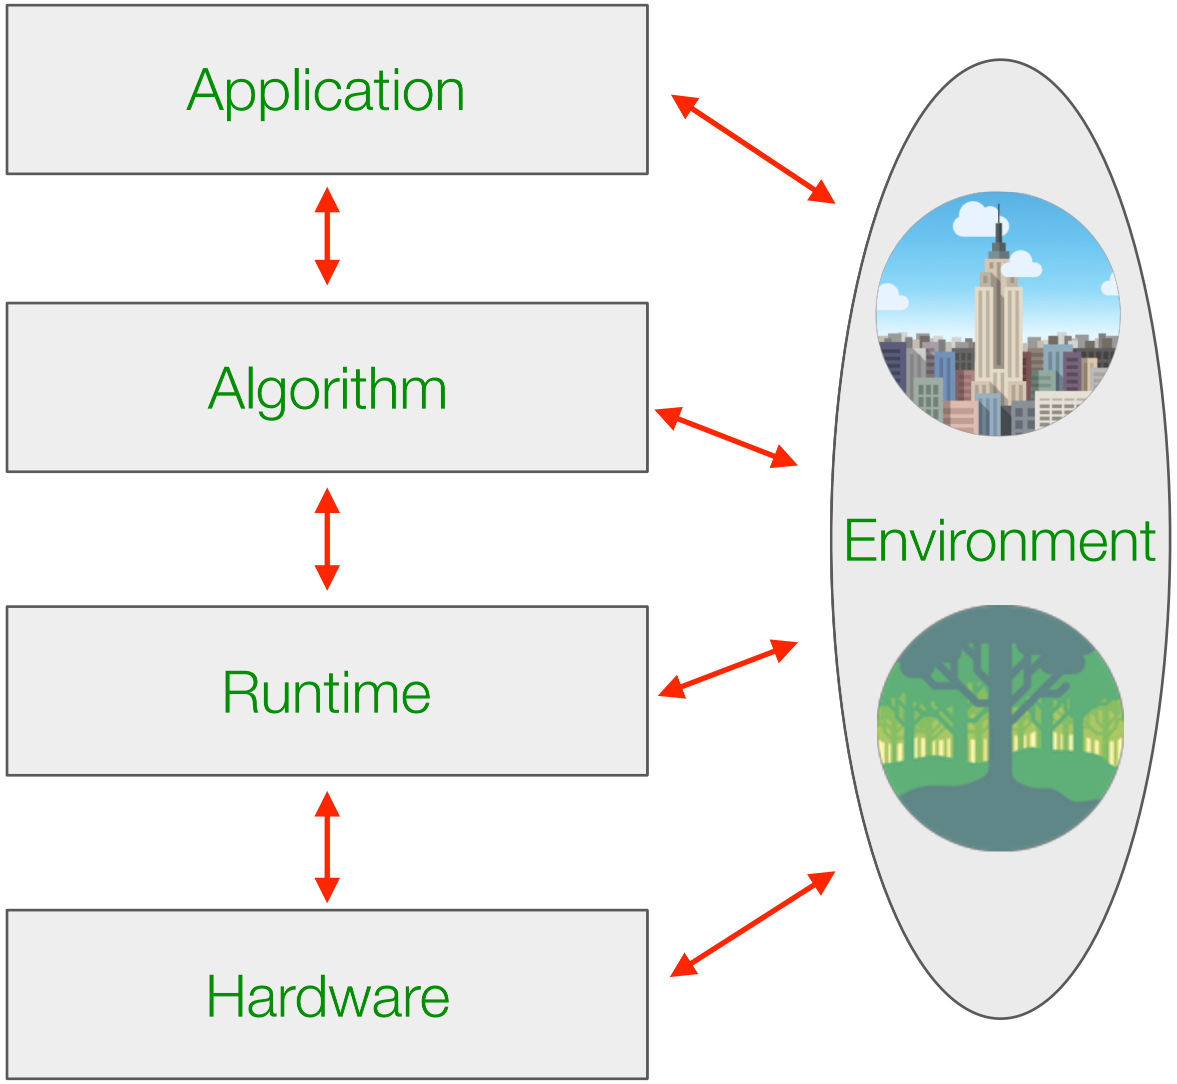

Understanding and Capturing the Interactions

- **Ideally, each level should be readily explorable…**
  - Independently
  - In relation to one another
- **Need development platforms**
  - Enable design, prototyping, development and evaluation
  - Fast, accurate, reproducible, …

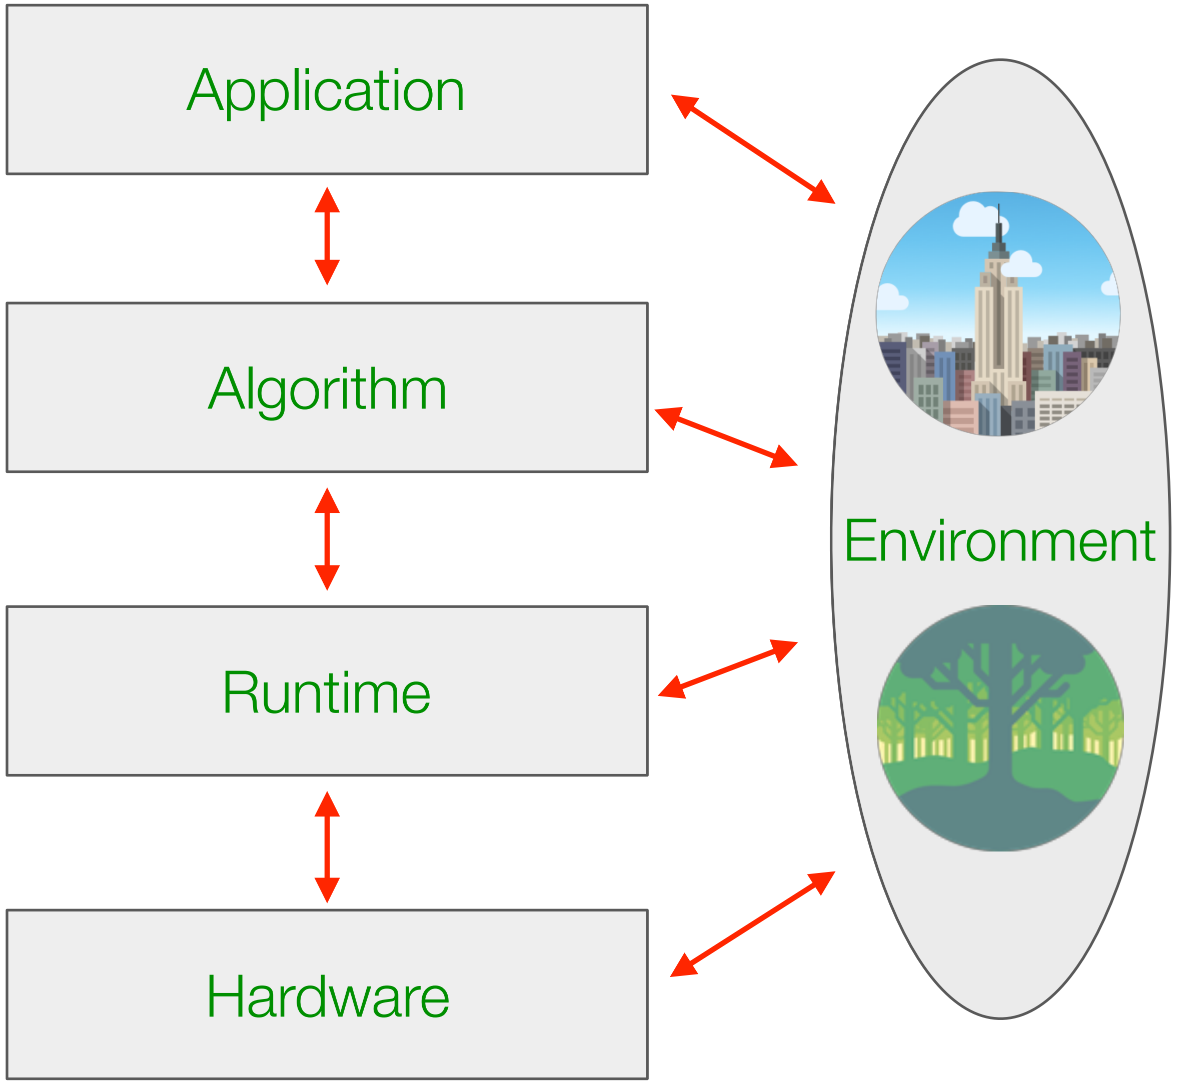

Closed-loop Simulation Setup

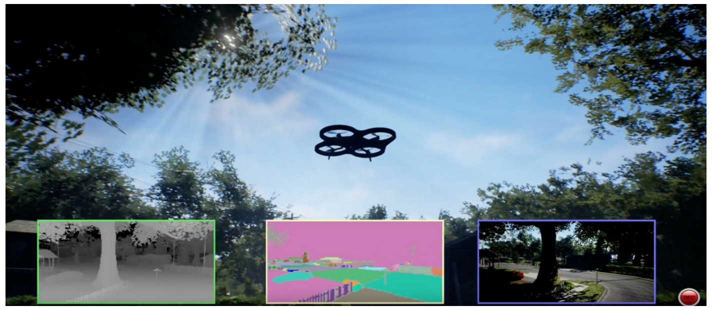

Closed-loop Simulation Setup

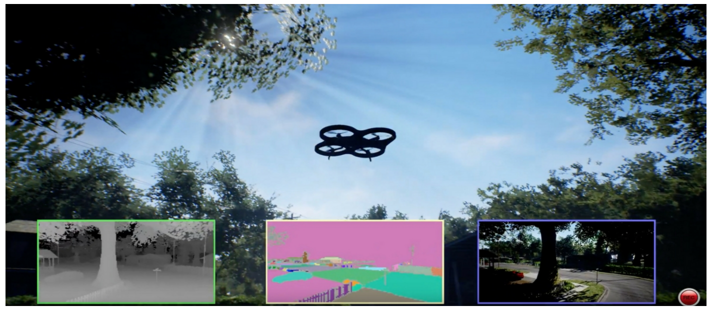

AirSim

Closed-loop Simulation Setup

AirSim

Unreal Engine

Closed-loop Simulation **Setup**

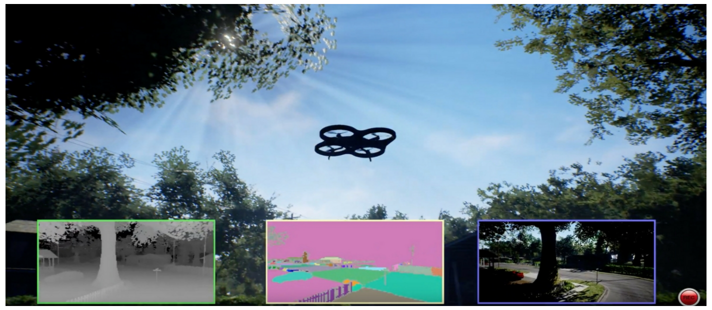

AirSim

Unreal Engine

Sensory Data
(RGBD, GPS)

Closed-loop Simulation Setup

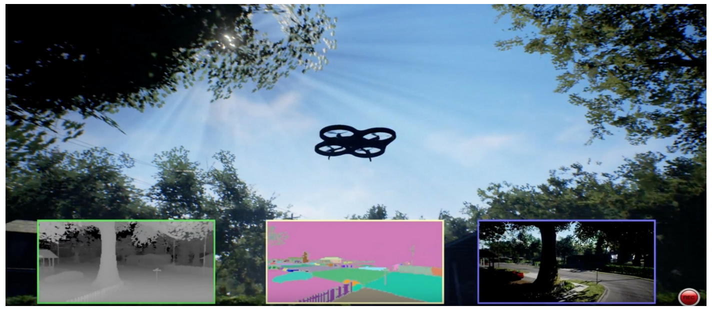

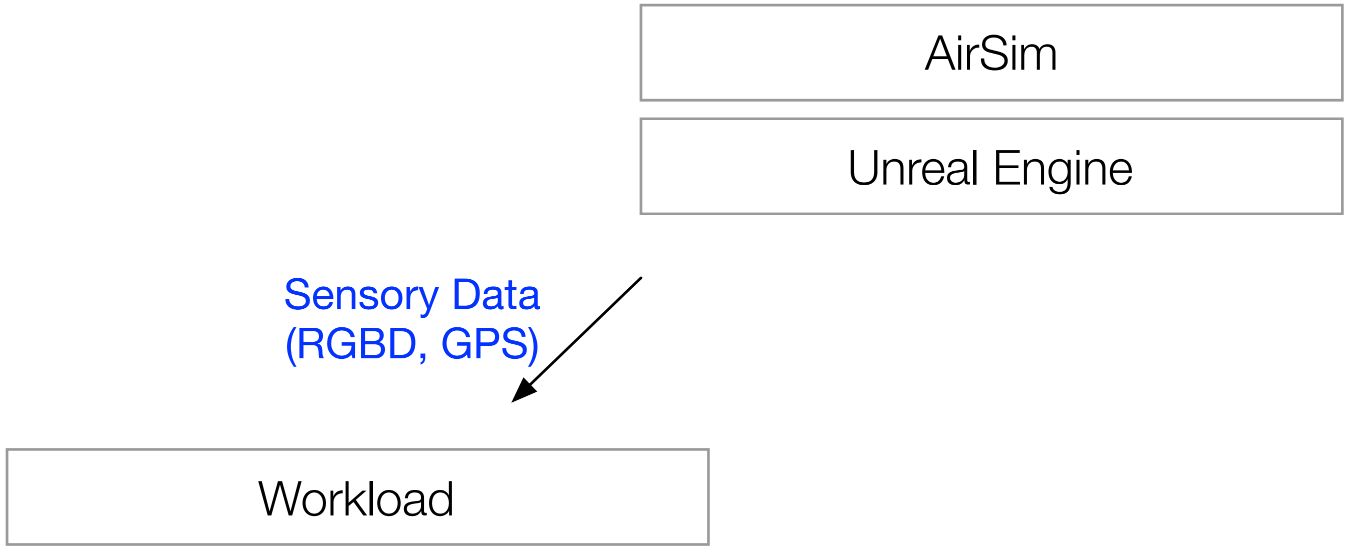

AirSim

Unreal Engine

Sensory Data
(RGBD, GPS)

Workload

Closed-loop Simulation Setup

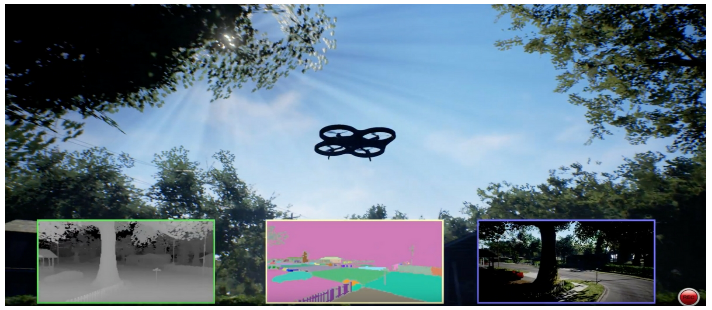

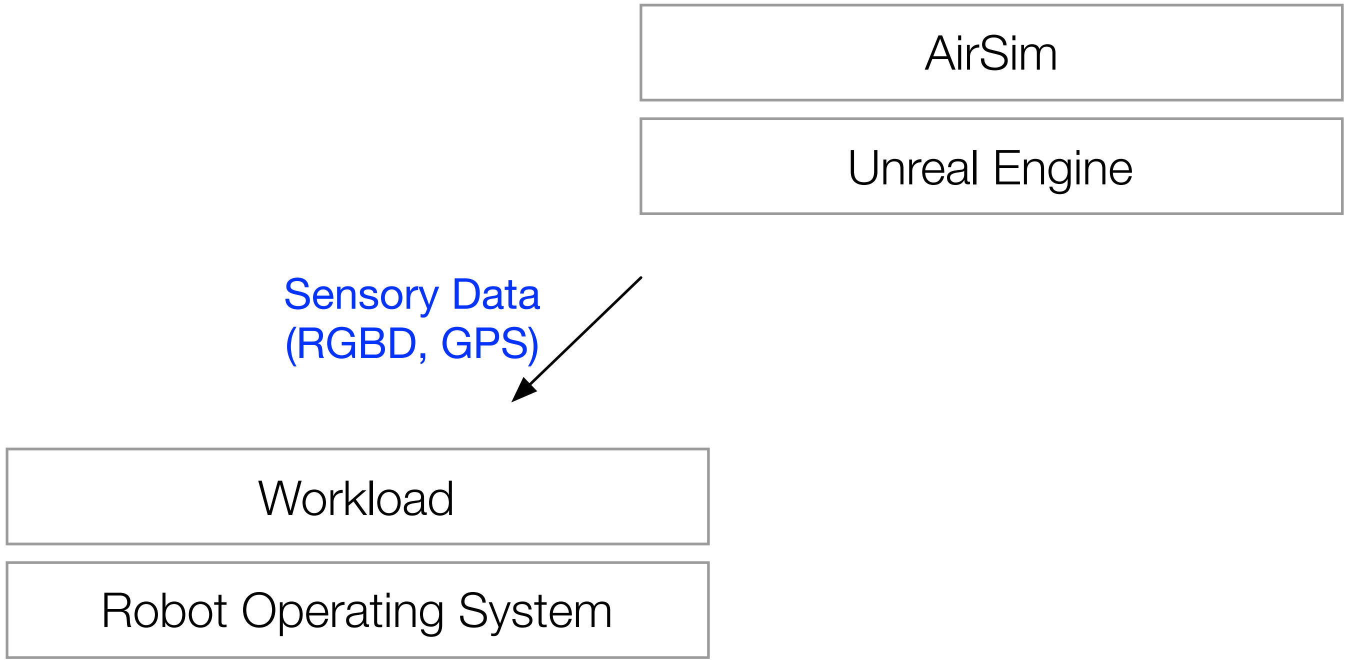

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Workload

Robot Operating System

Closed-loop Simulation **Setup**

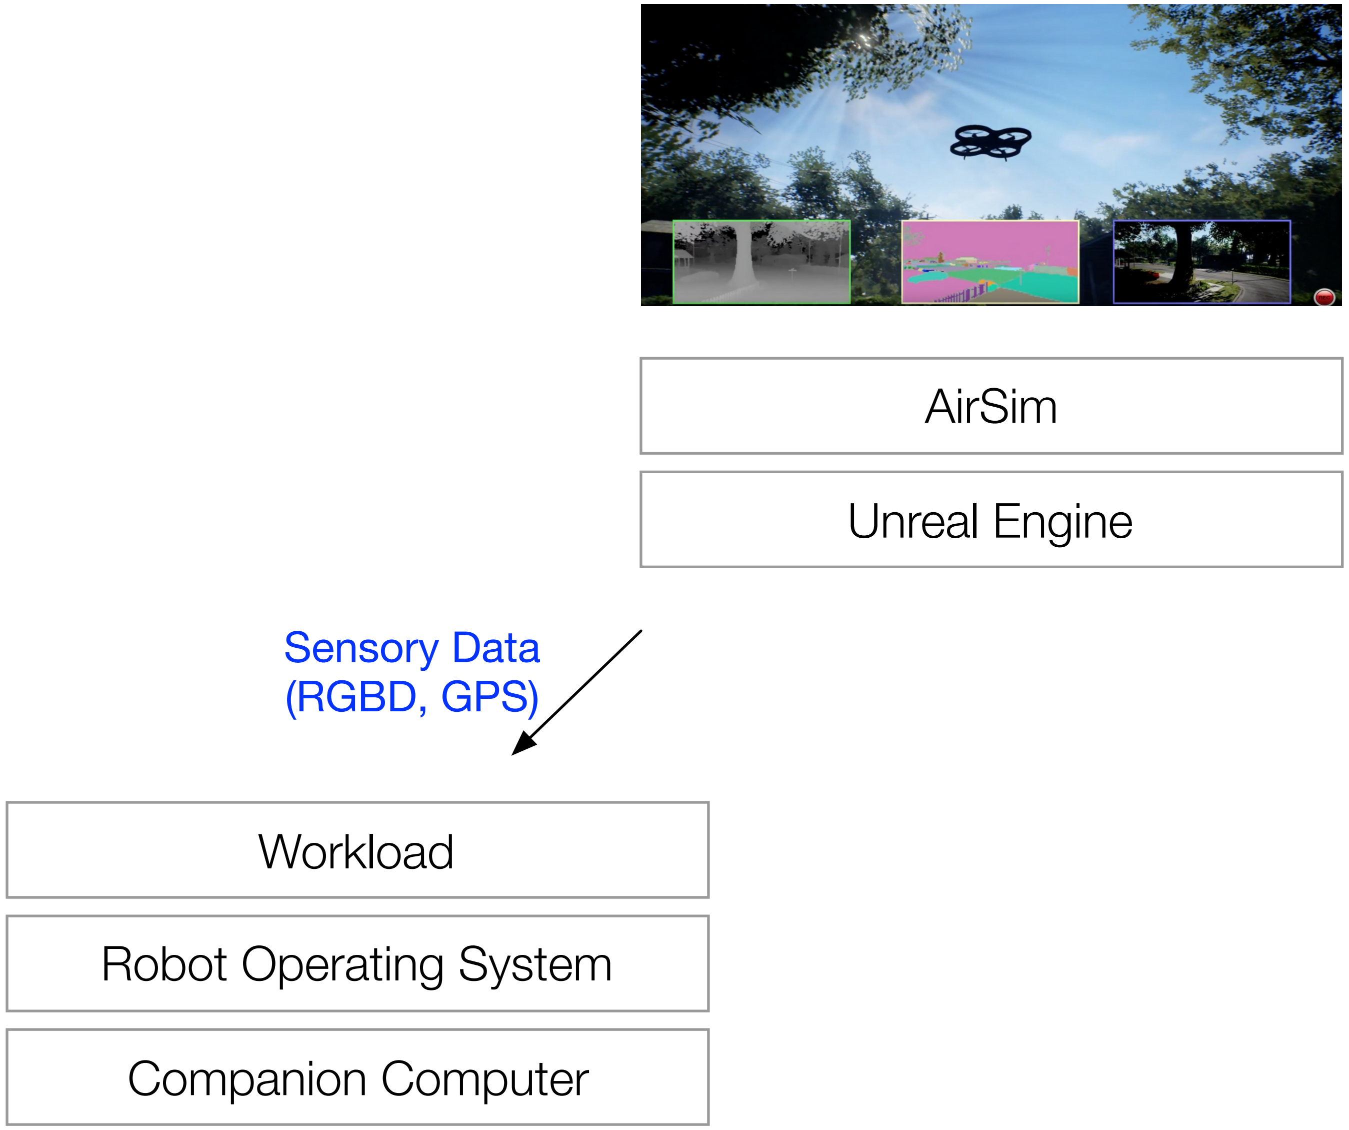

AirSim

Unreal Engine

Sensory Data
(RGBD, GPS)

Workload

Robot Operating System

Companion Computer

Closed-loop Simulation **Setup**

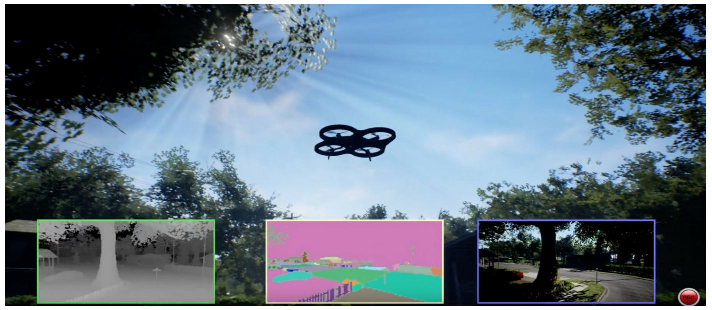

AirSim

Unreal Engine

Sensory Data
(RGBD, GPS)

Workload

Robot Operating System

Companion Computer

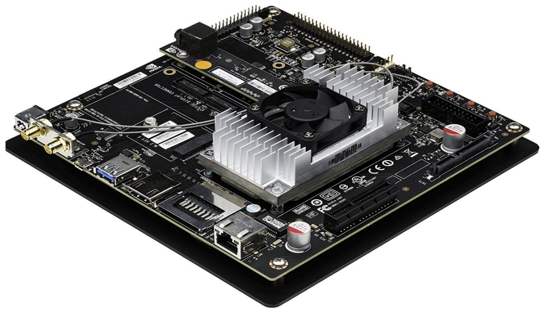

Closed-loop Simulation **Setup**

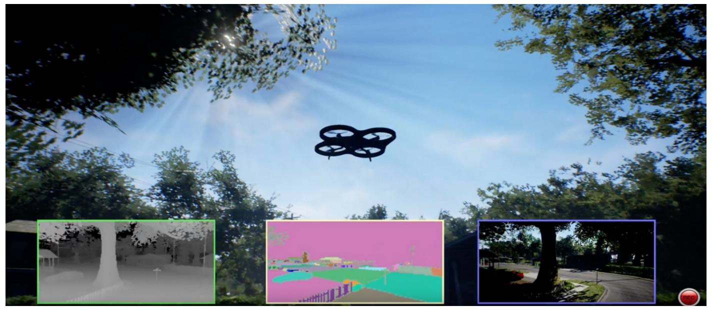

AirSim

Unreal Engine

Sensory Data
(RGBD, GPS)

Workload

Robot Operating System

Companion Computer

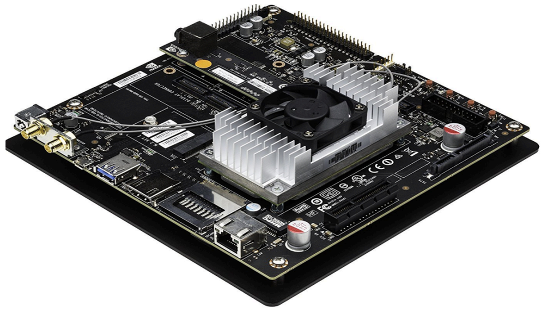

Flight
Commands

Closed-loop Simulation **Setup**

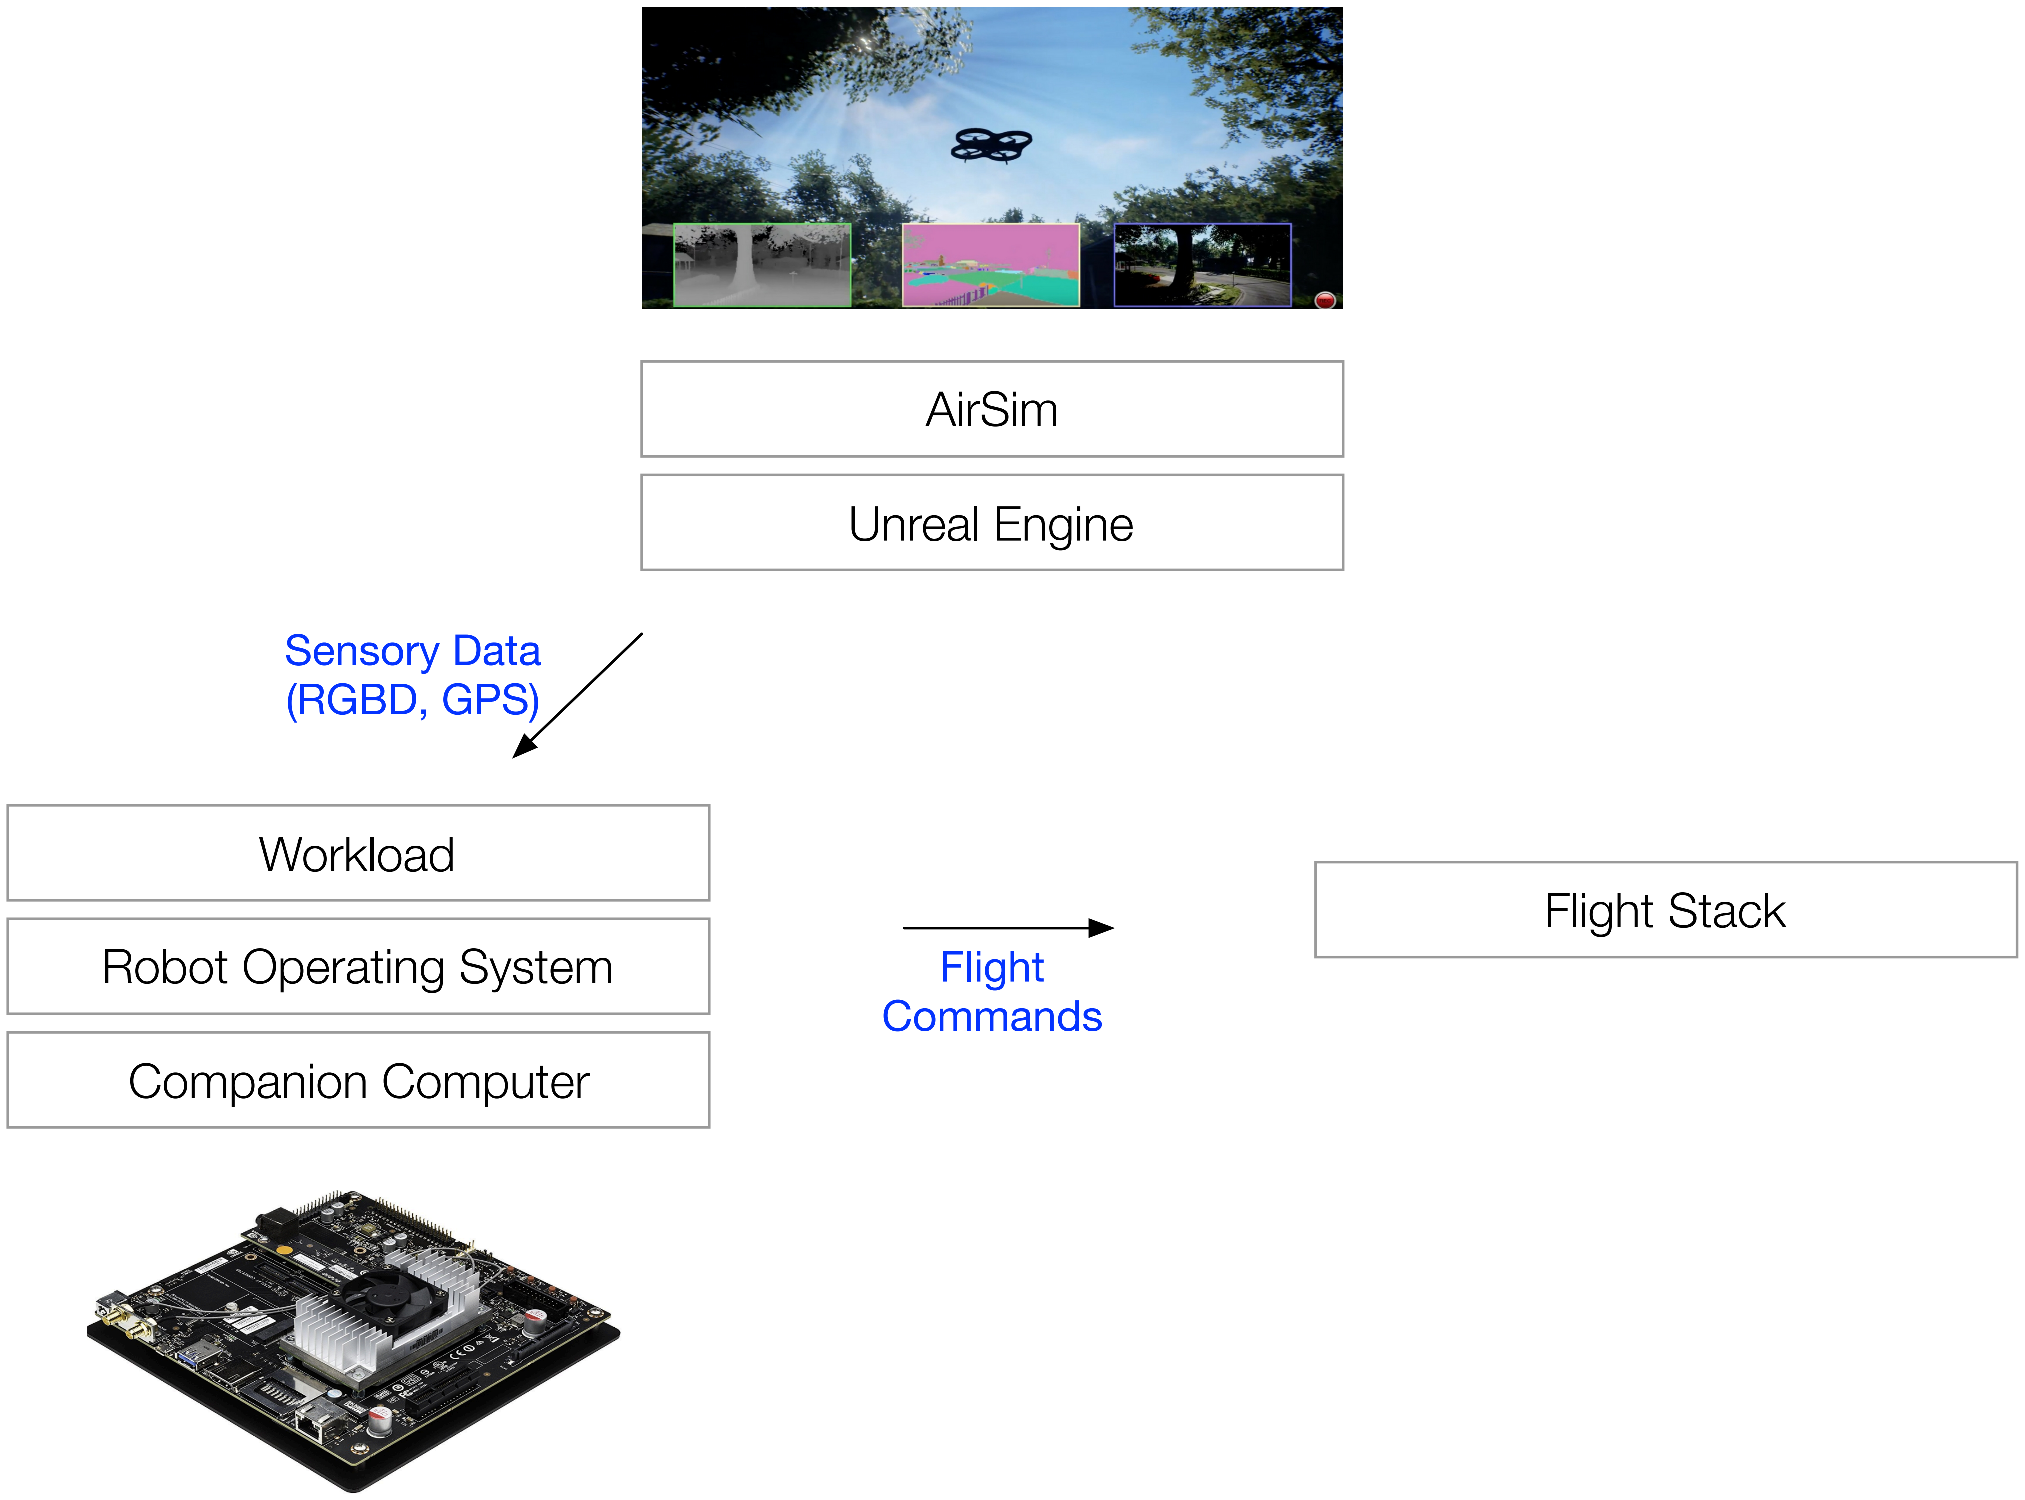

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Closed-loop Simulation Setup

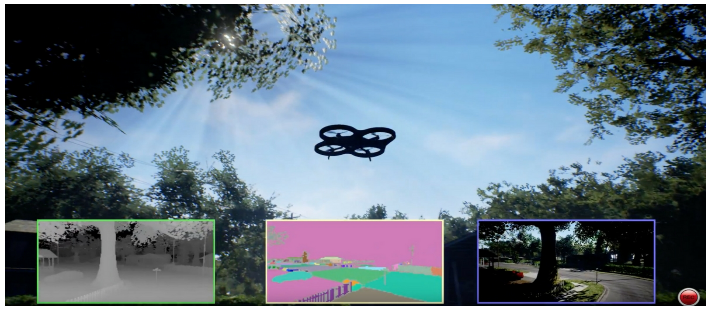

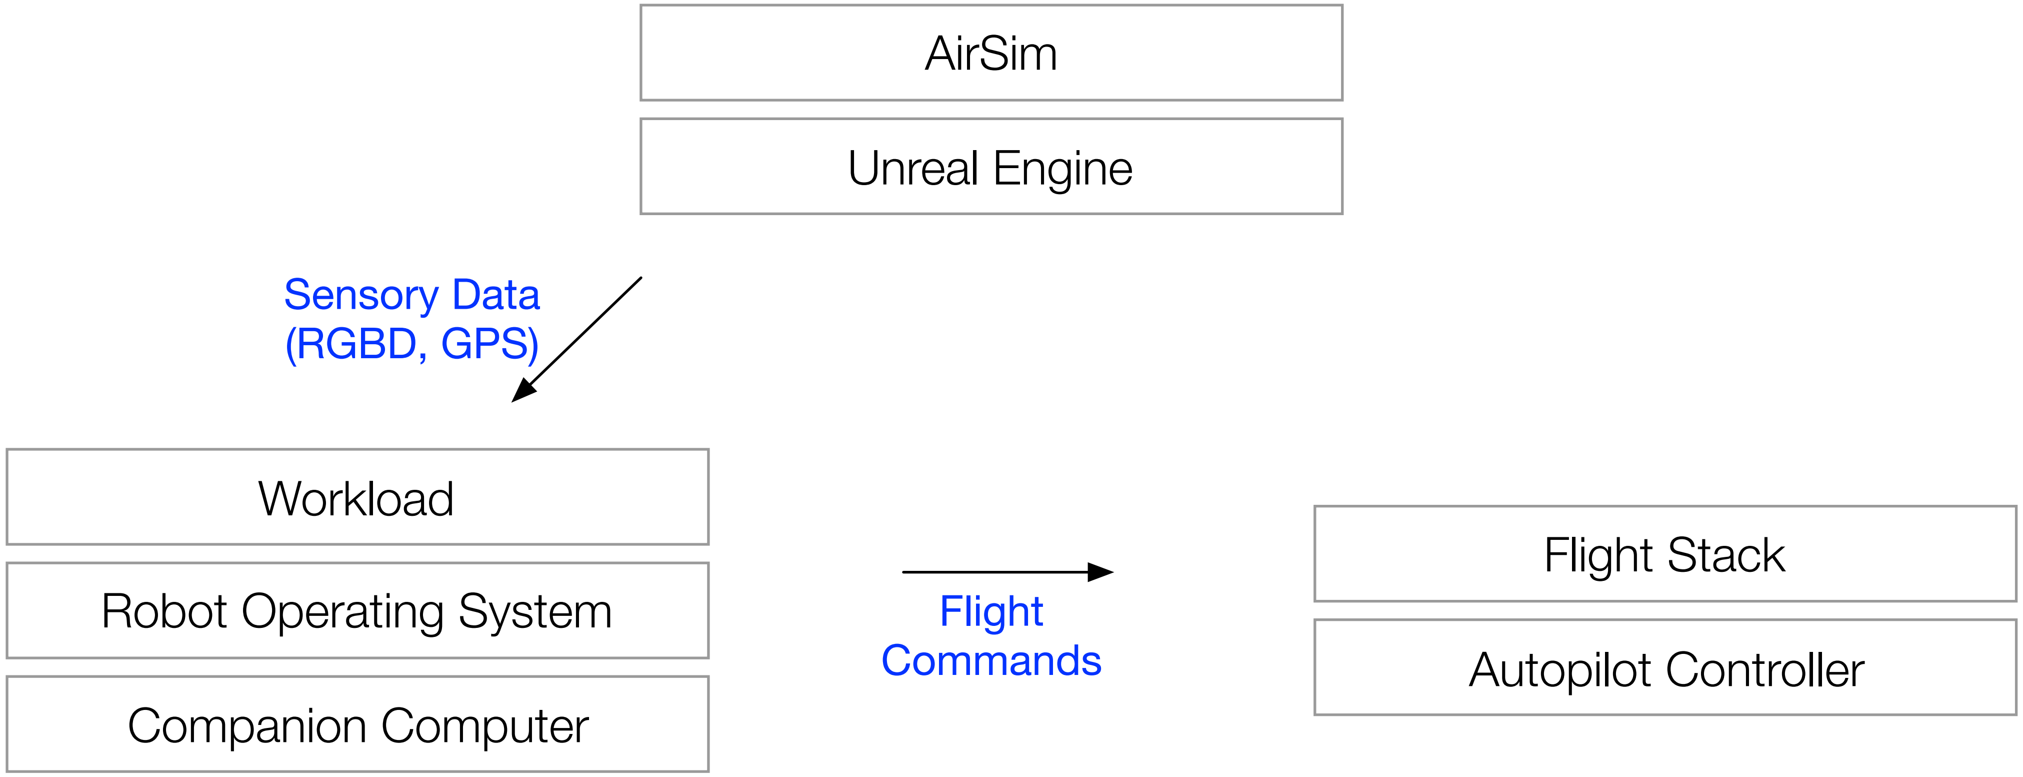

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

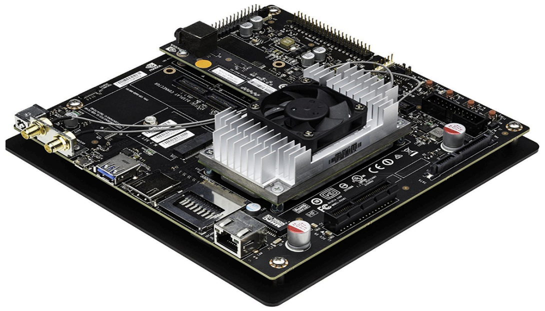

Closed-loop Simulation Setup

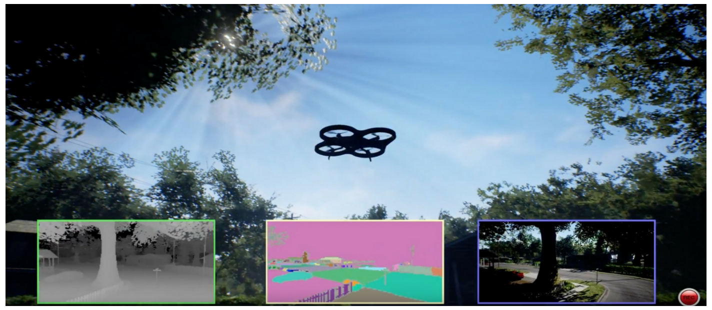

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

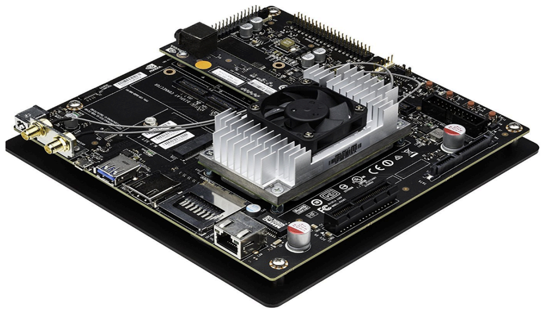

Closed-loop Simulation **Setup**

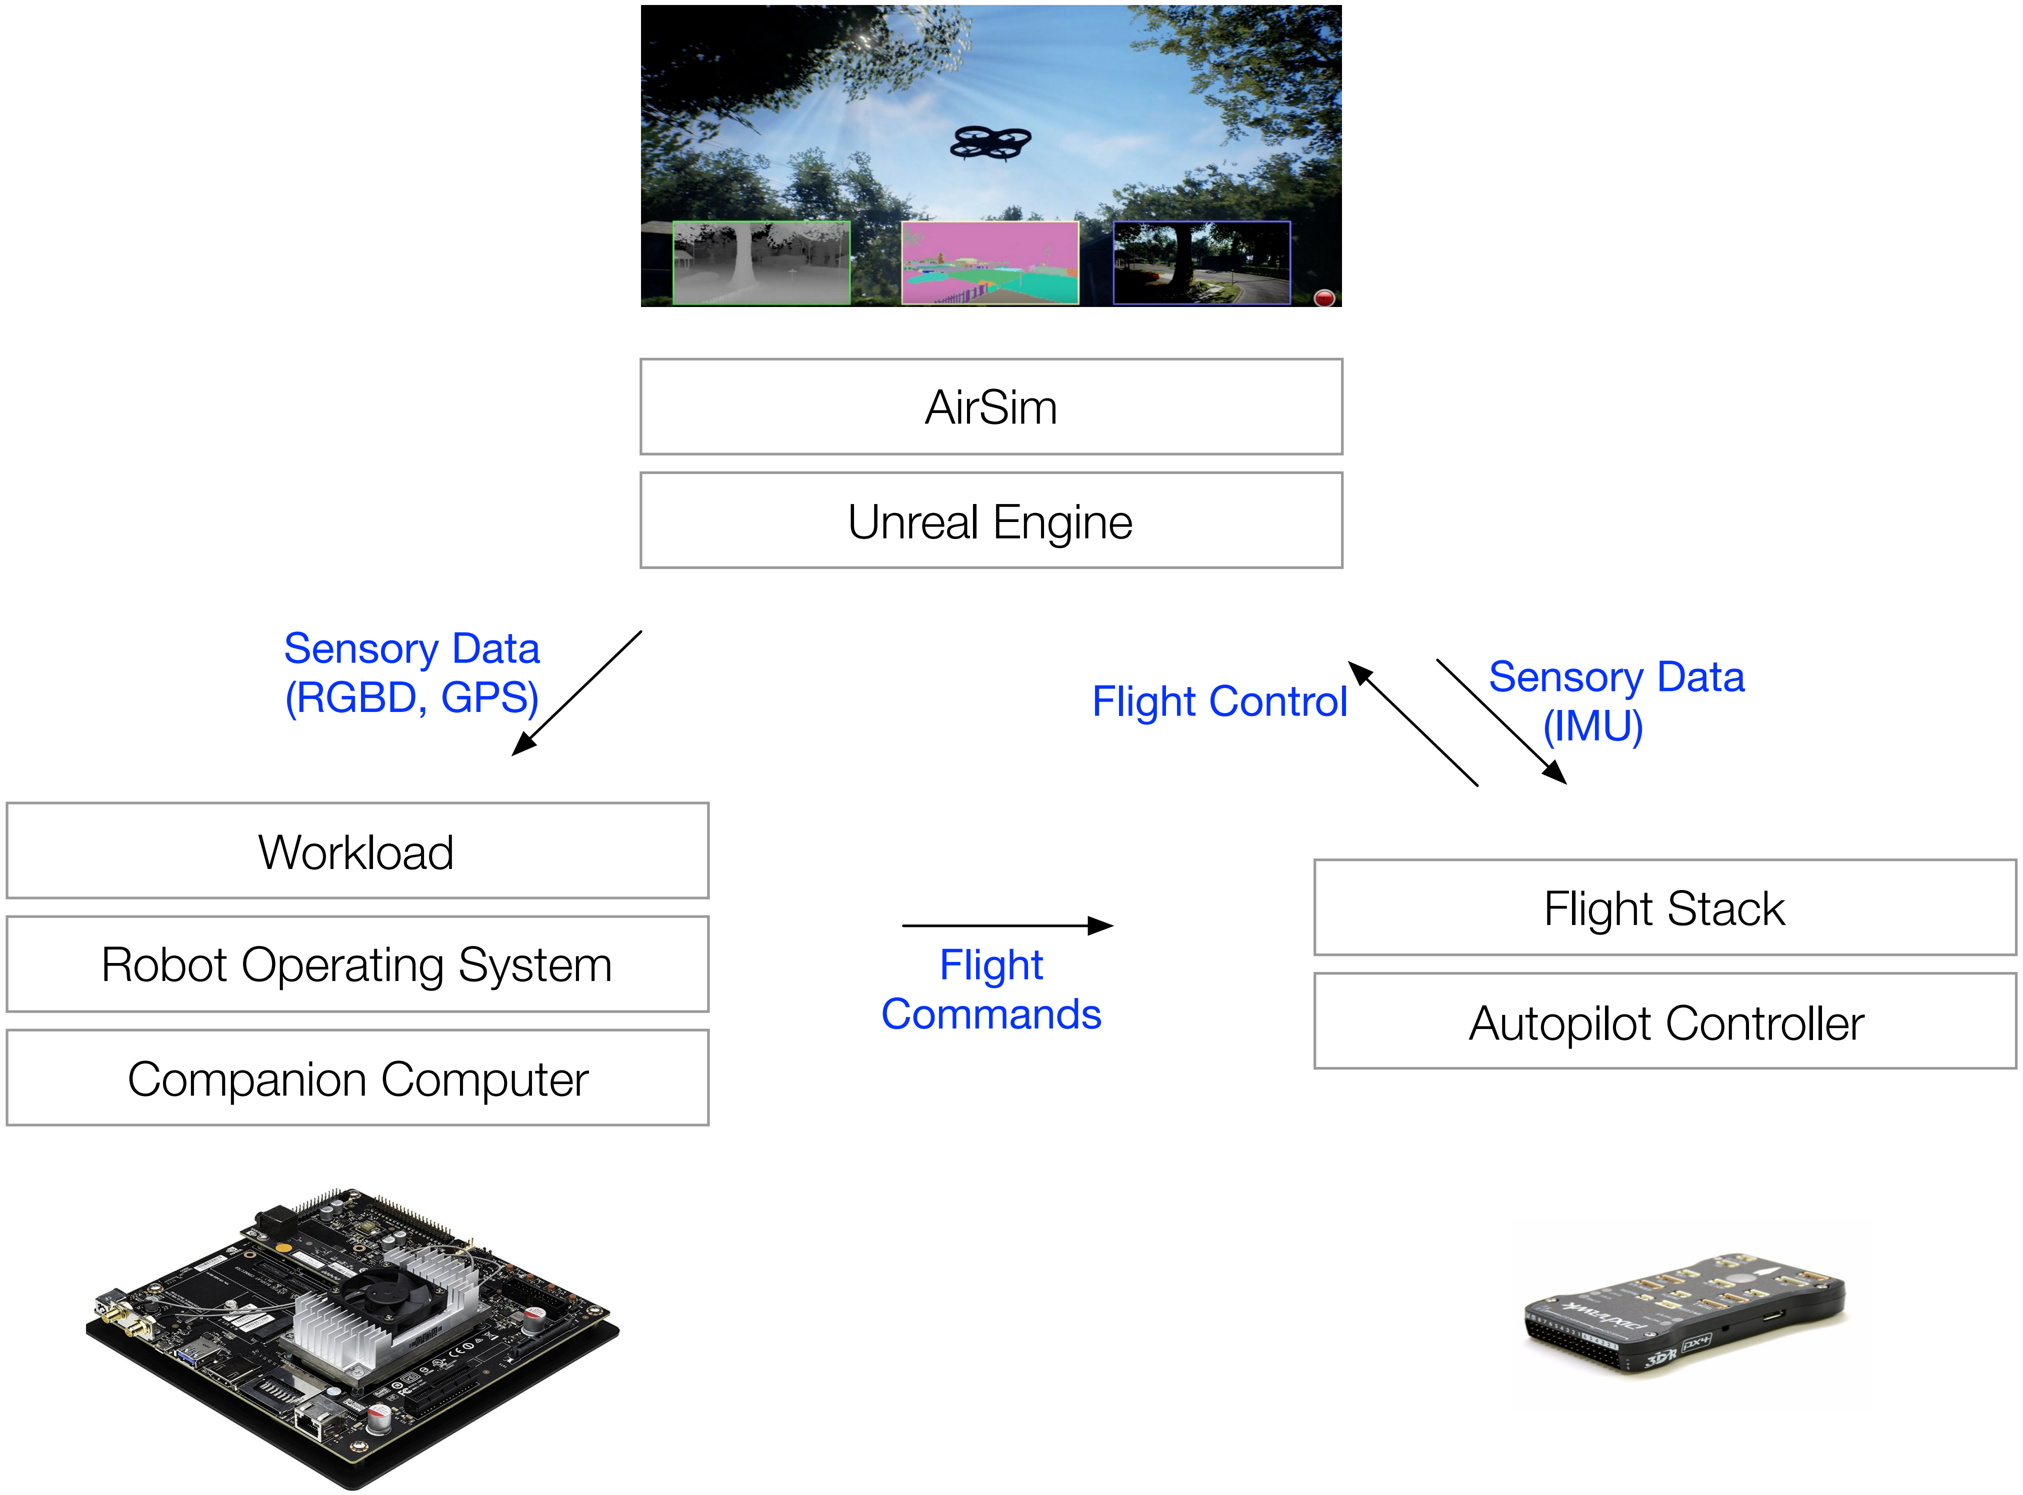

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

Closed-loop Simulation Knobs

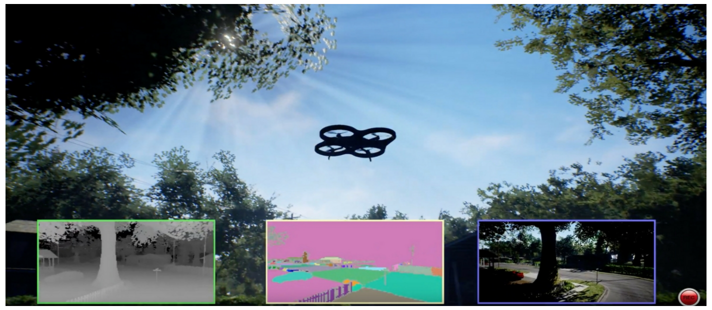

AirSim

Unreal Engine

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

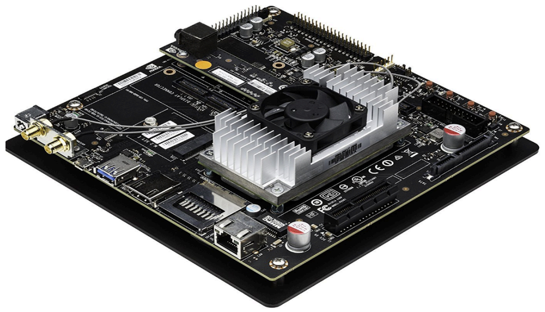

Closed-loop Simulation **Knobs**

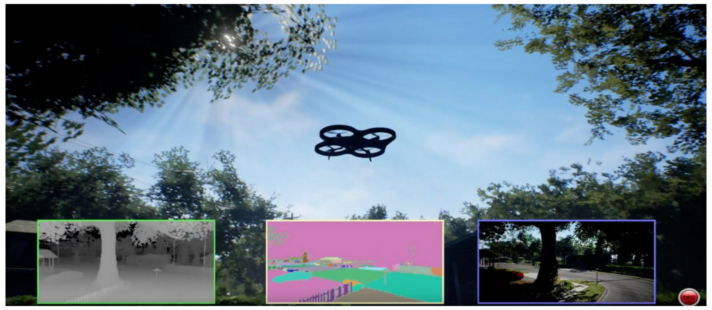

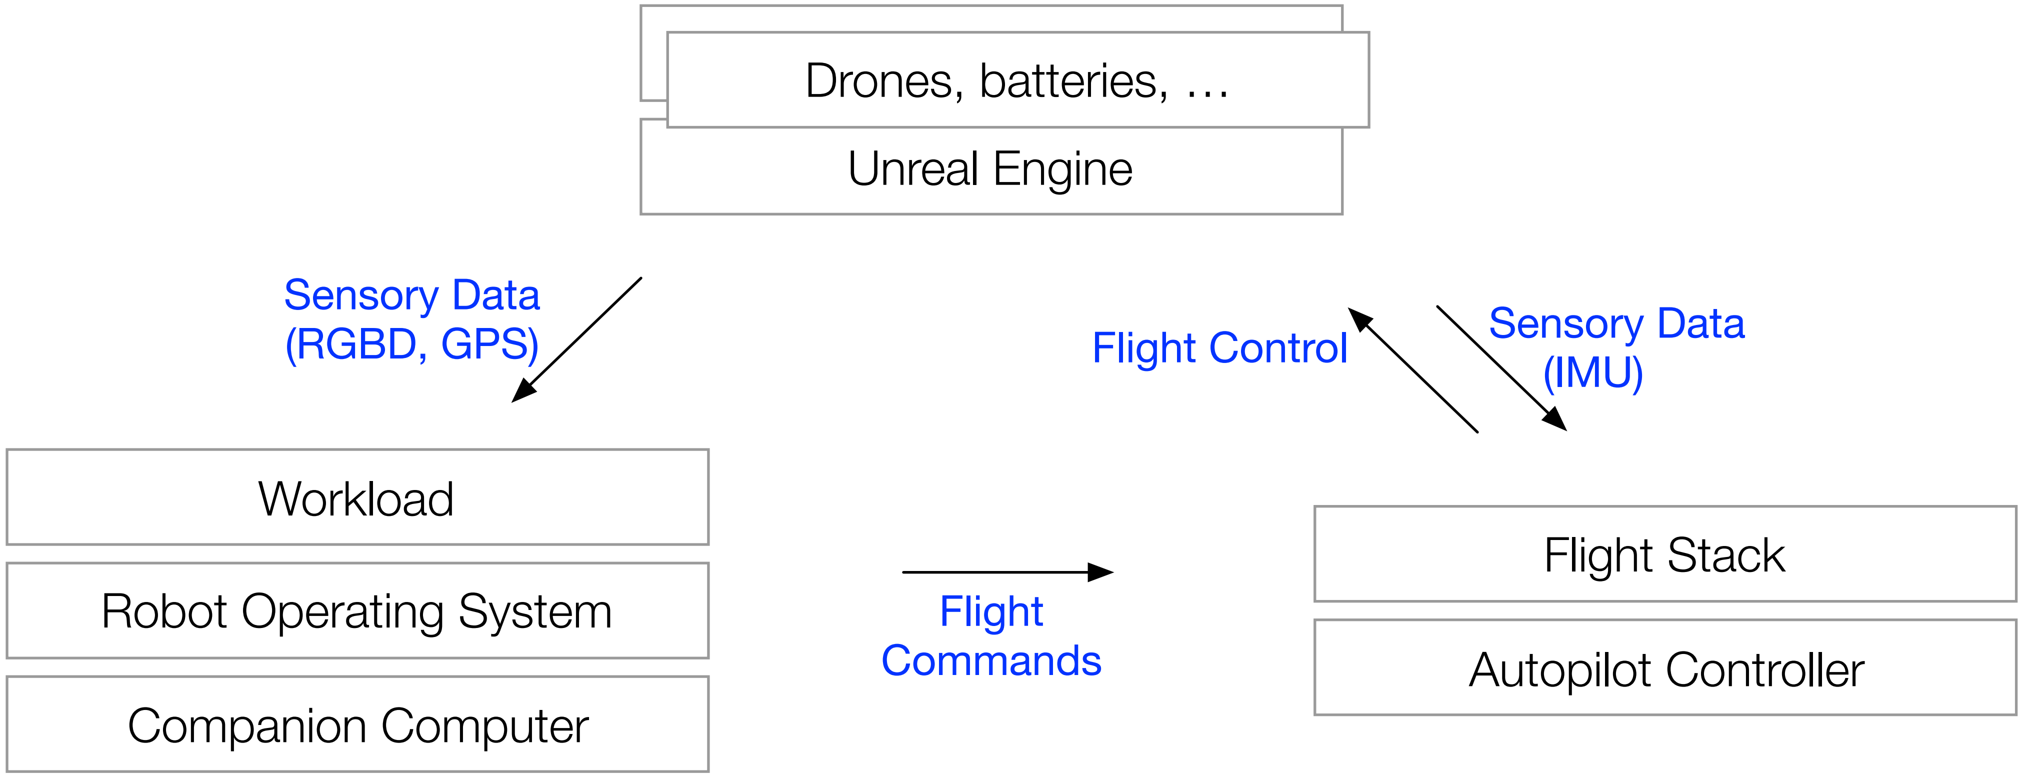

Drones, batteries, …

Unreal Engine

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

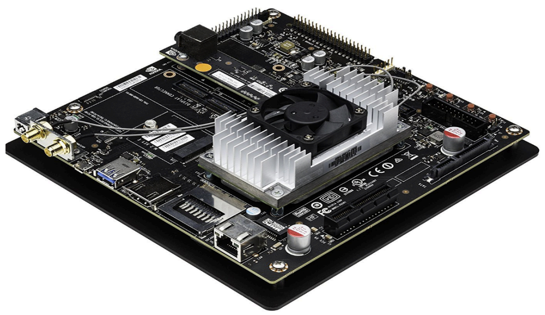

Closed-loop Simulation Knobs

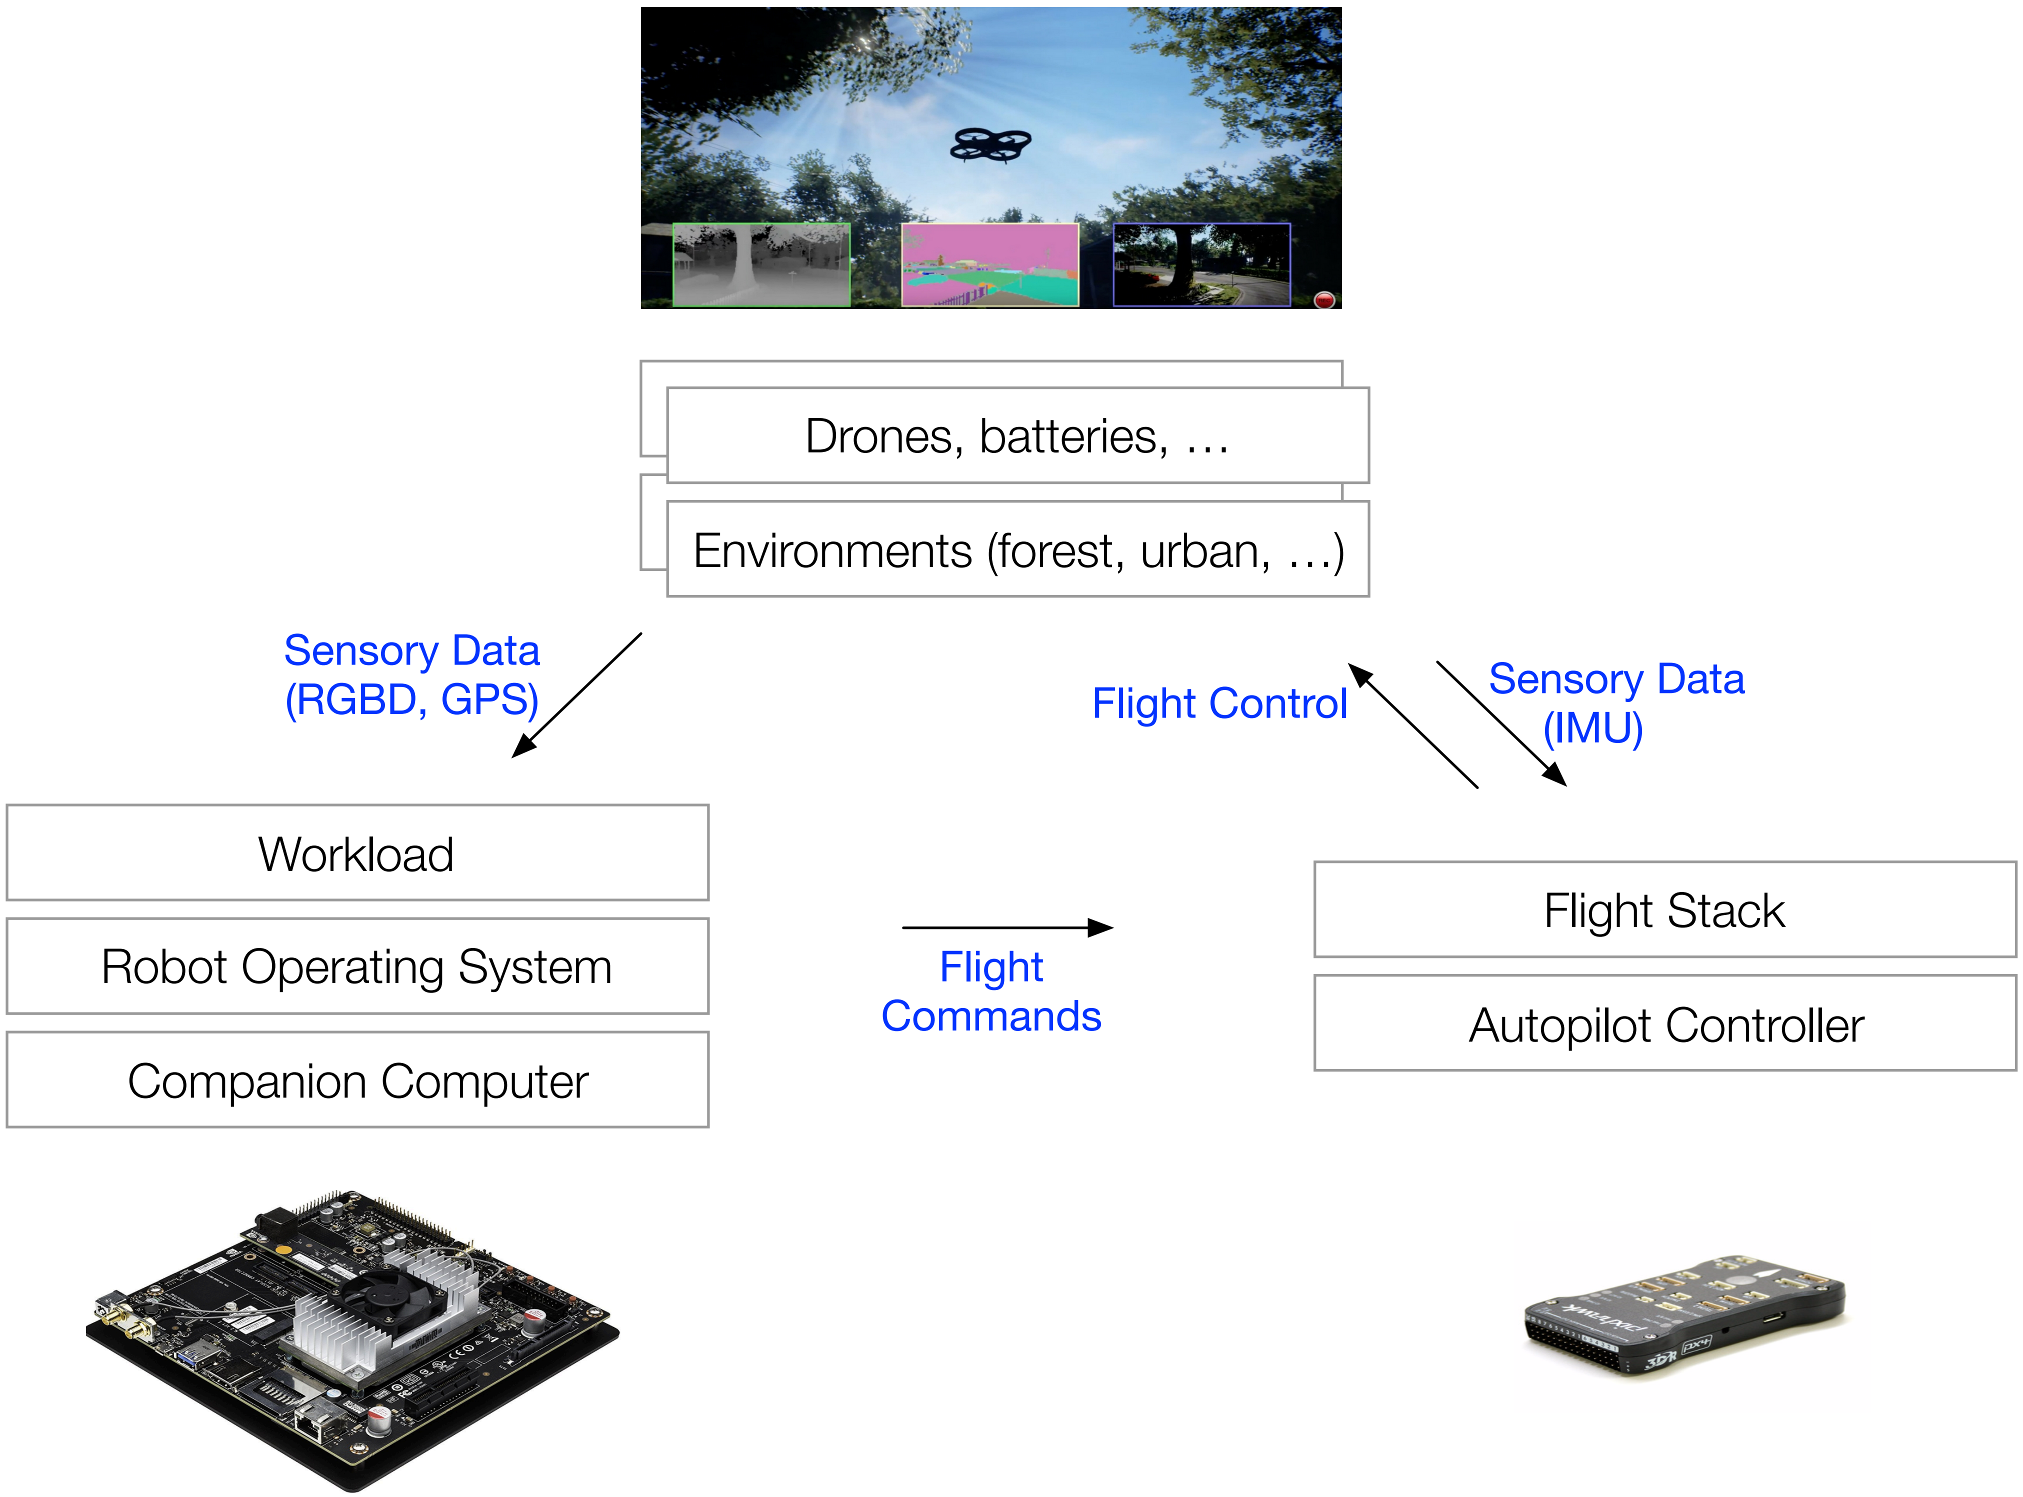

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Workload

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

Closed-loop Simulation Knobs

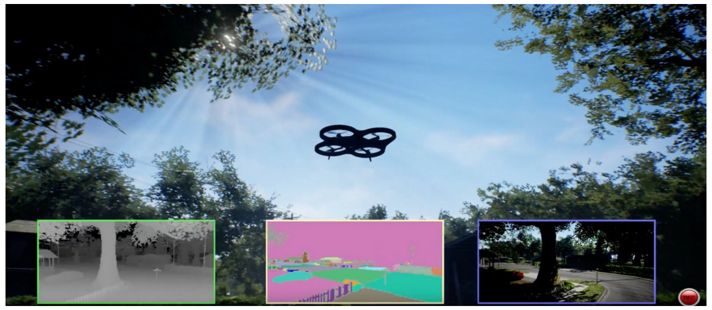

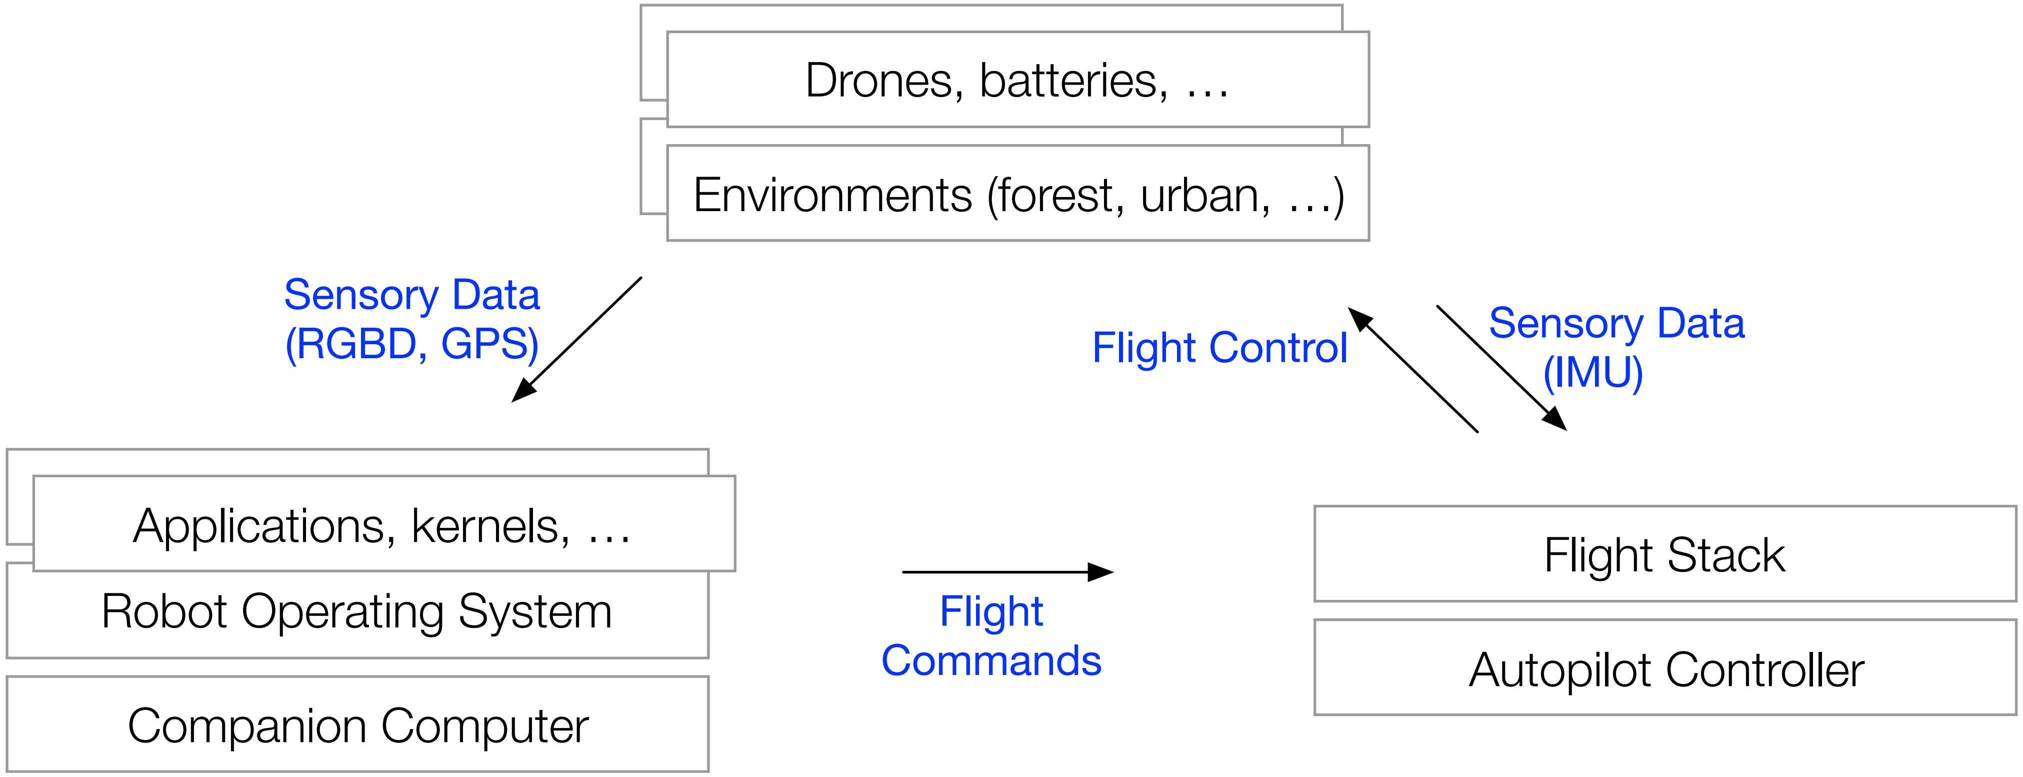

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Robot Operating System

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

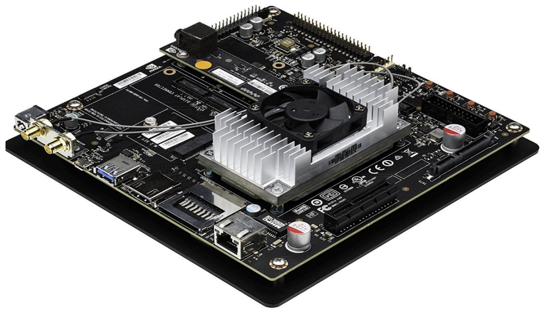

Closed-loop Simulation **Knobs**

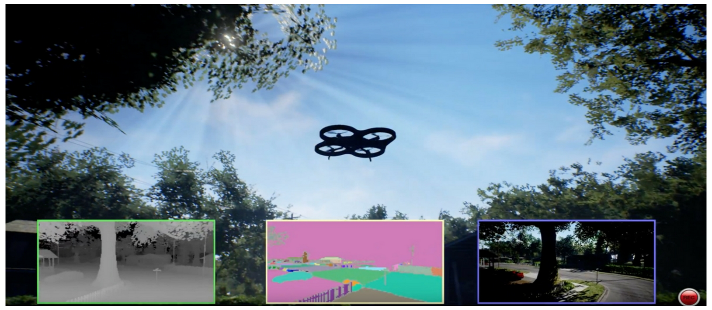

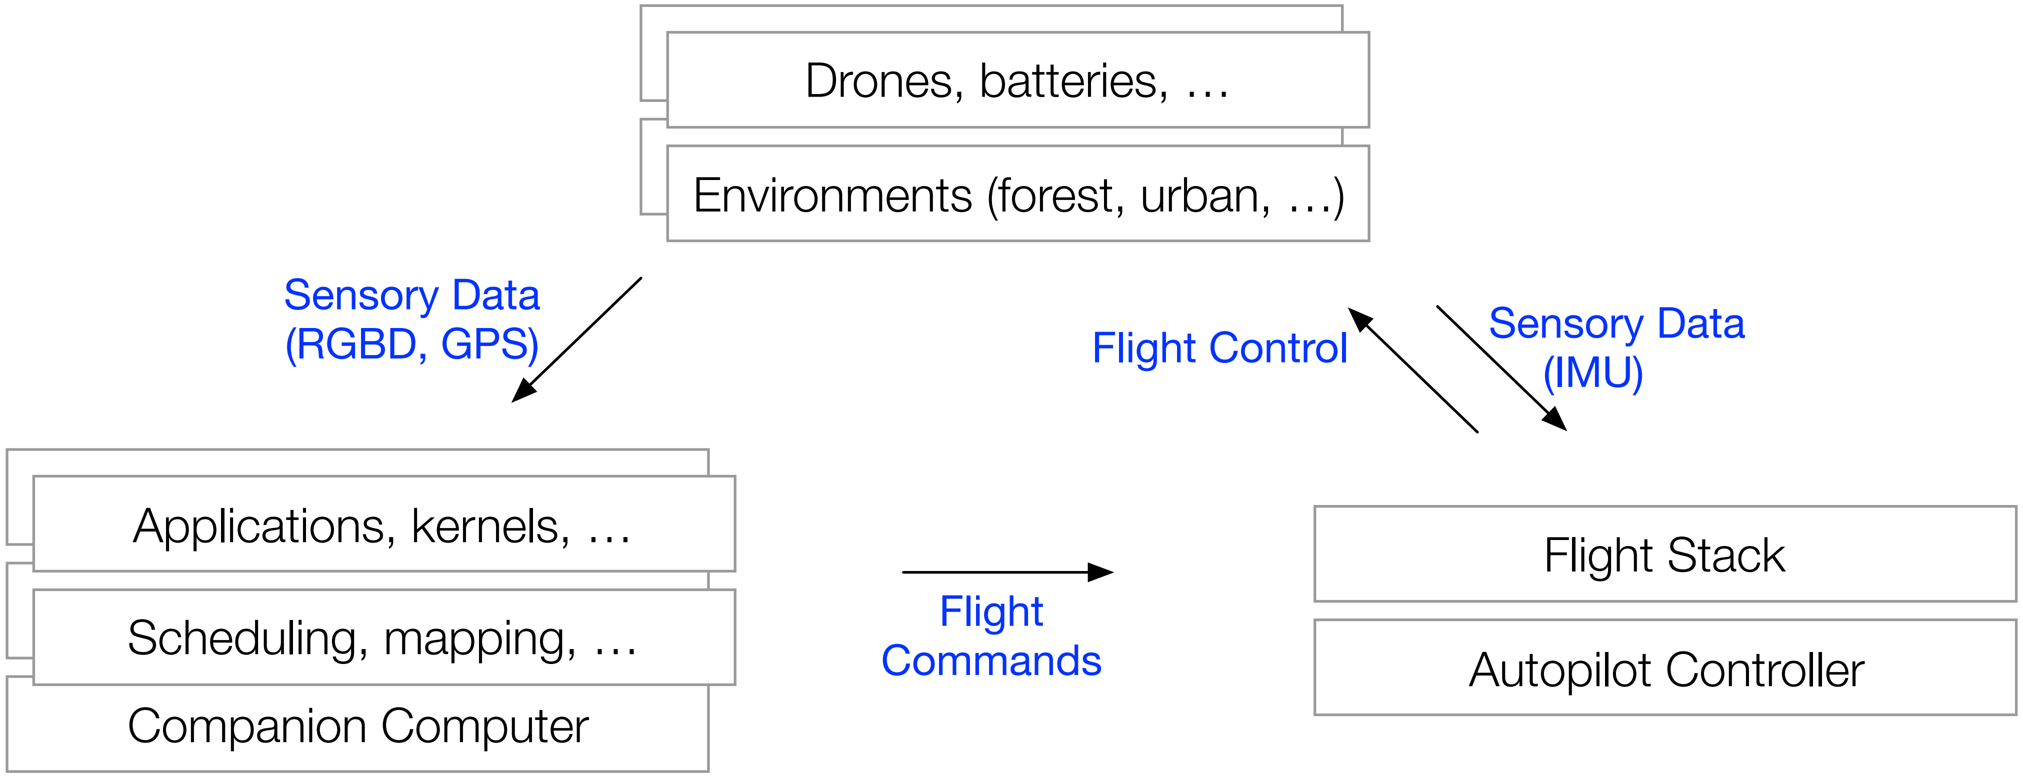

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Companion Computer

Flight Commands

Flight Stack

Autopilot Controller

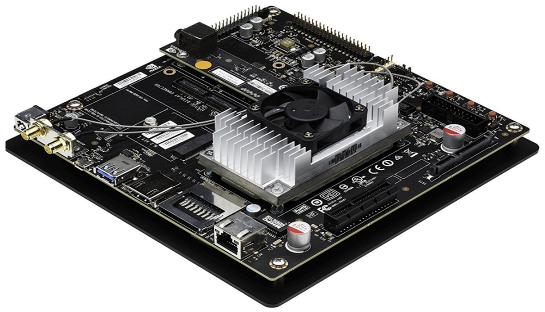

Closed-loop Simulation **Knobs**

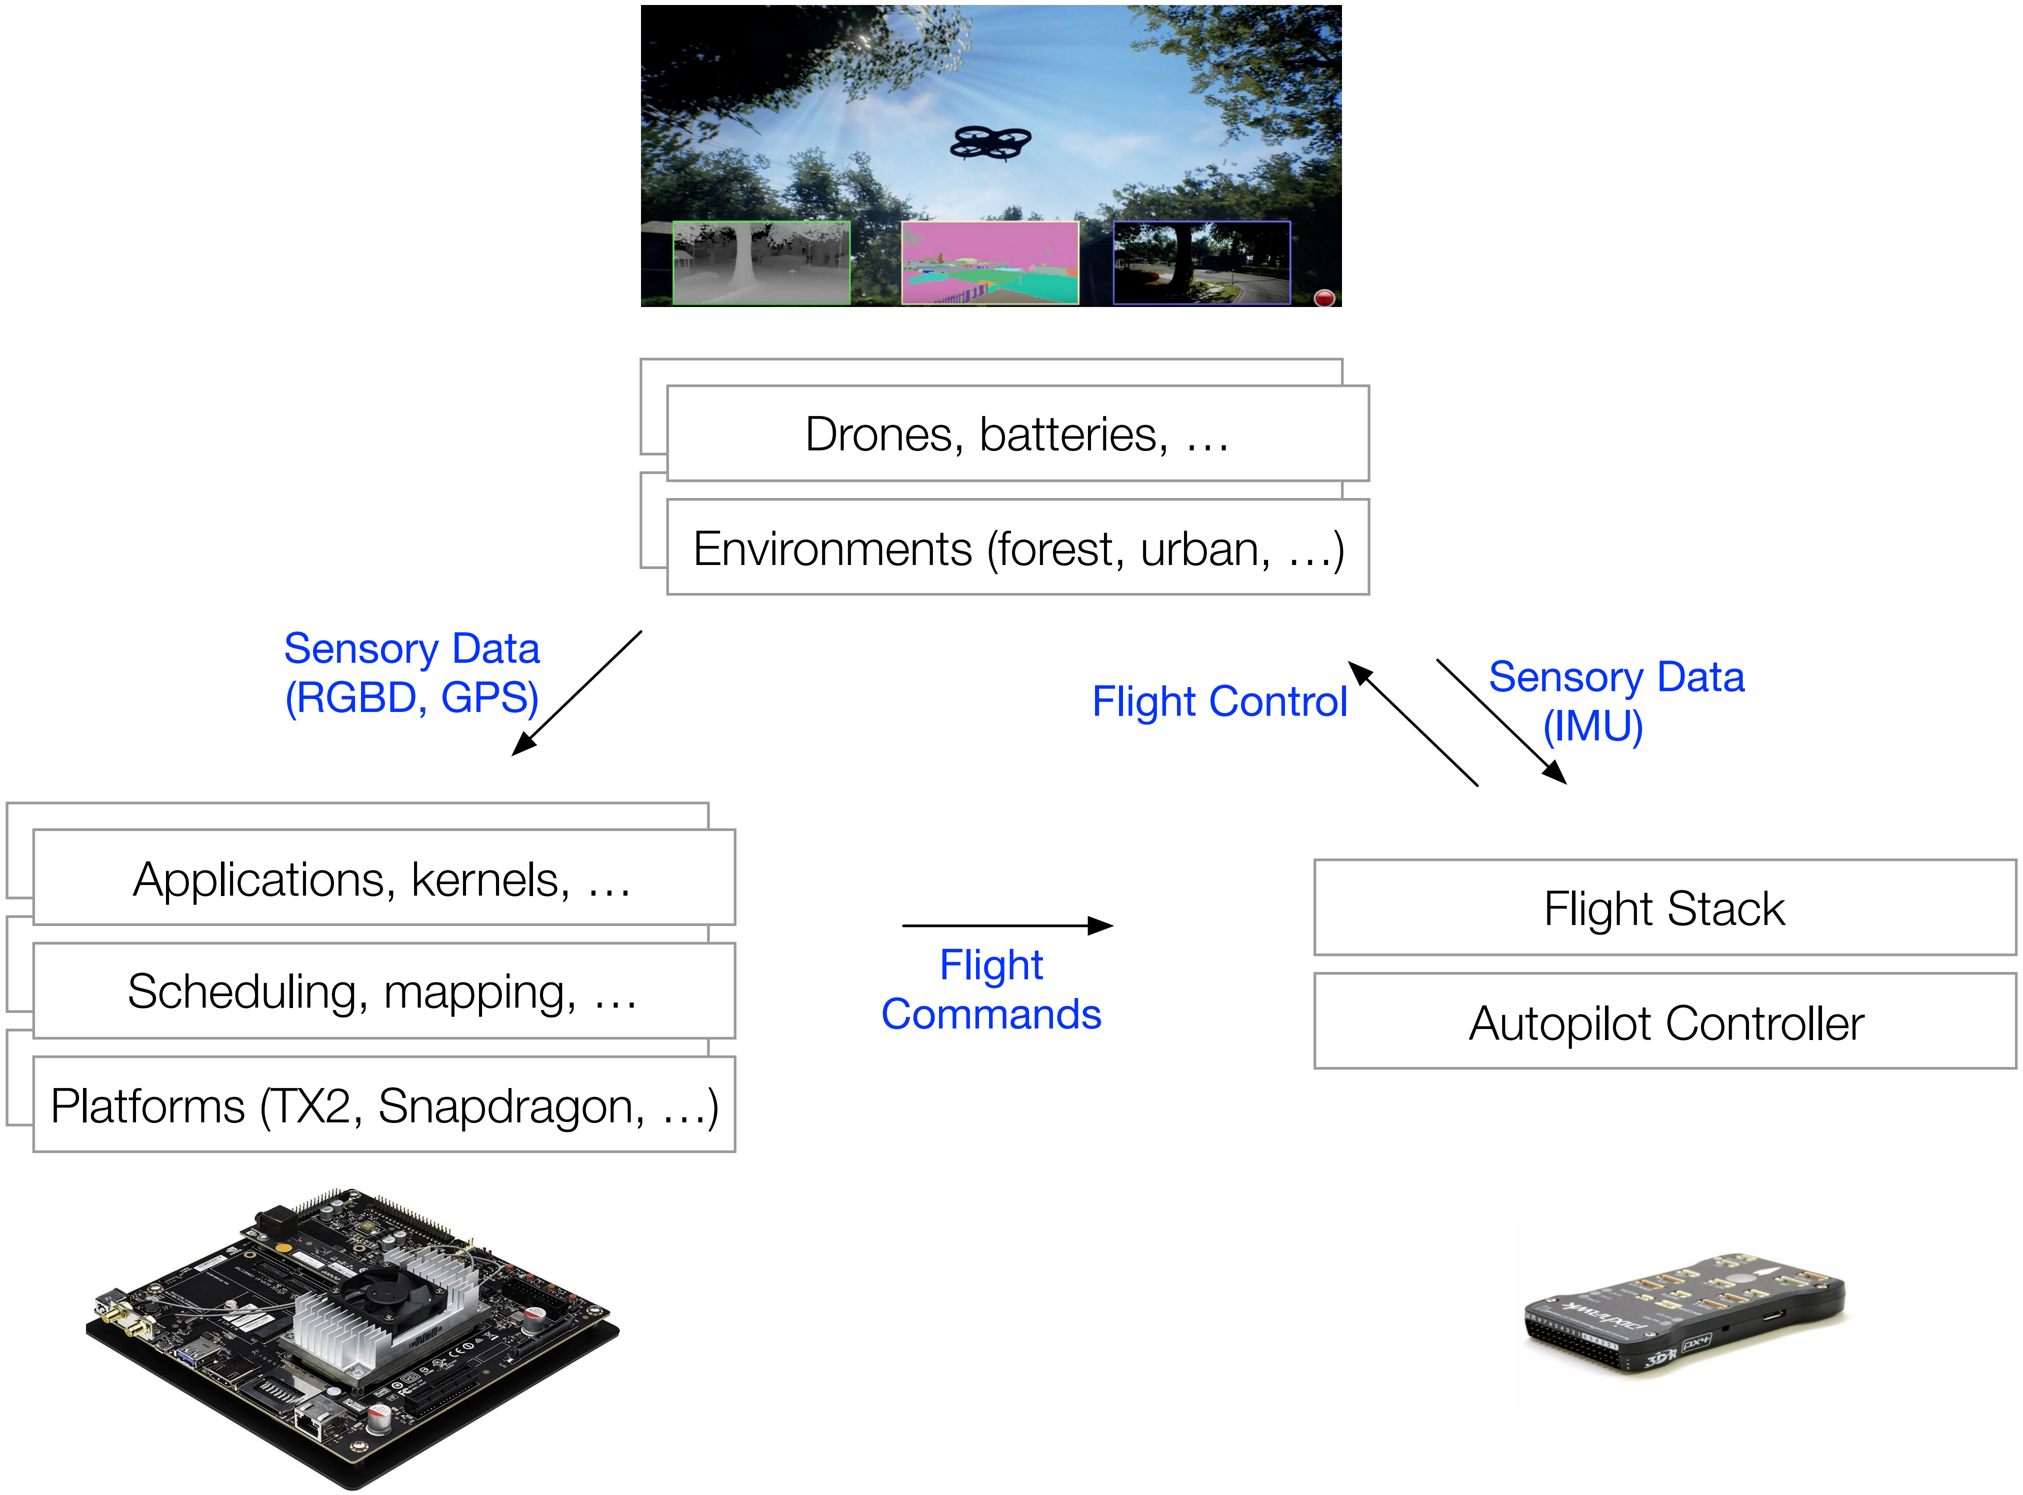

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack

Autopilot Controller

Closed-loop Simulation **Knobs**

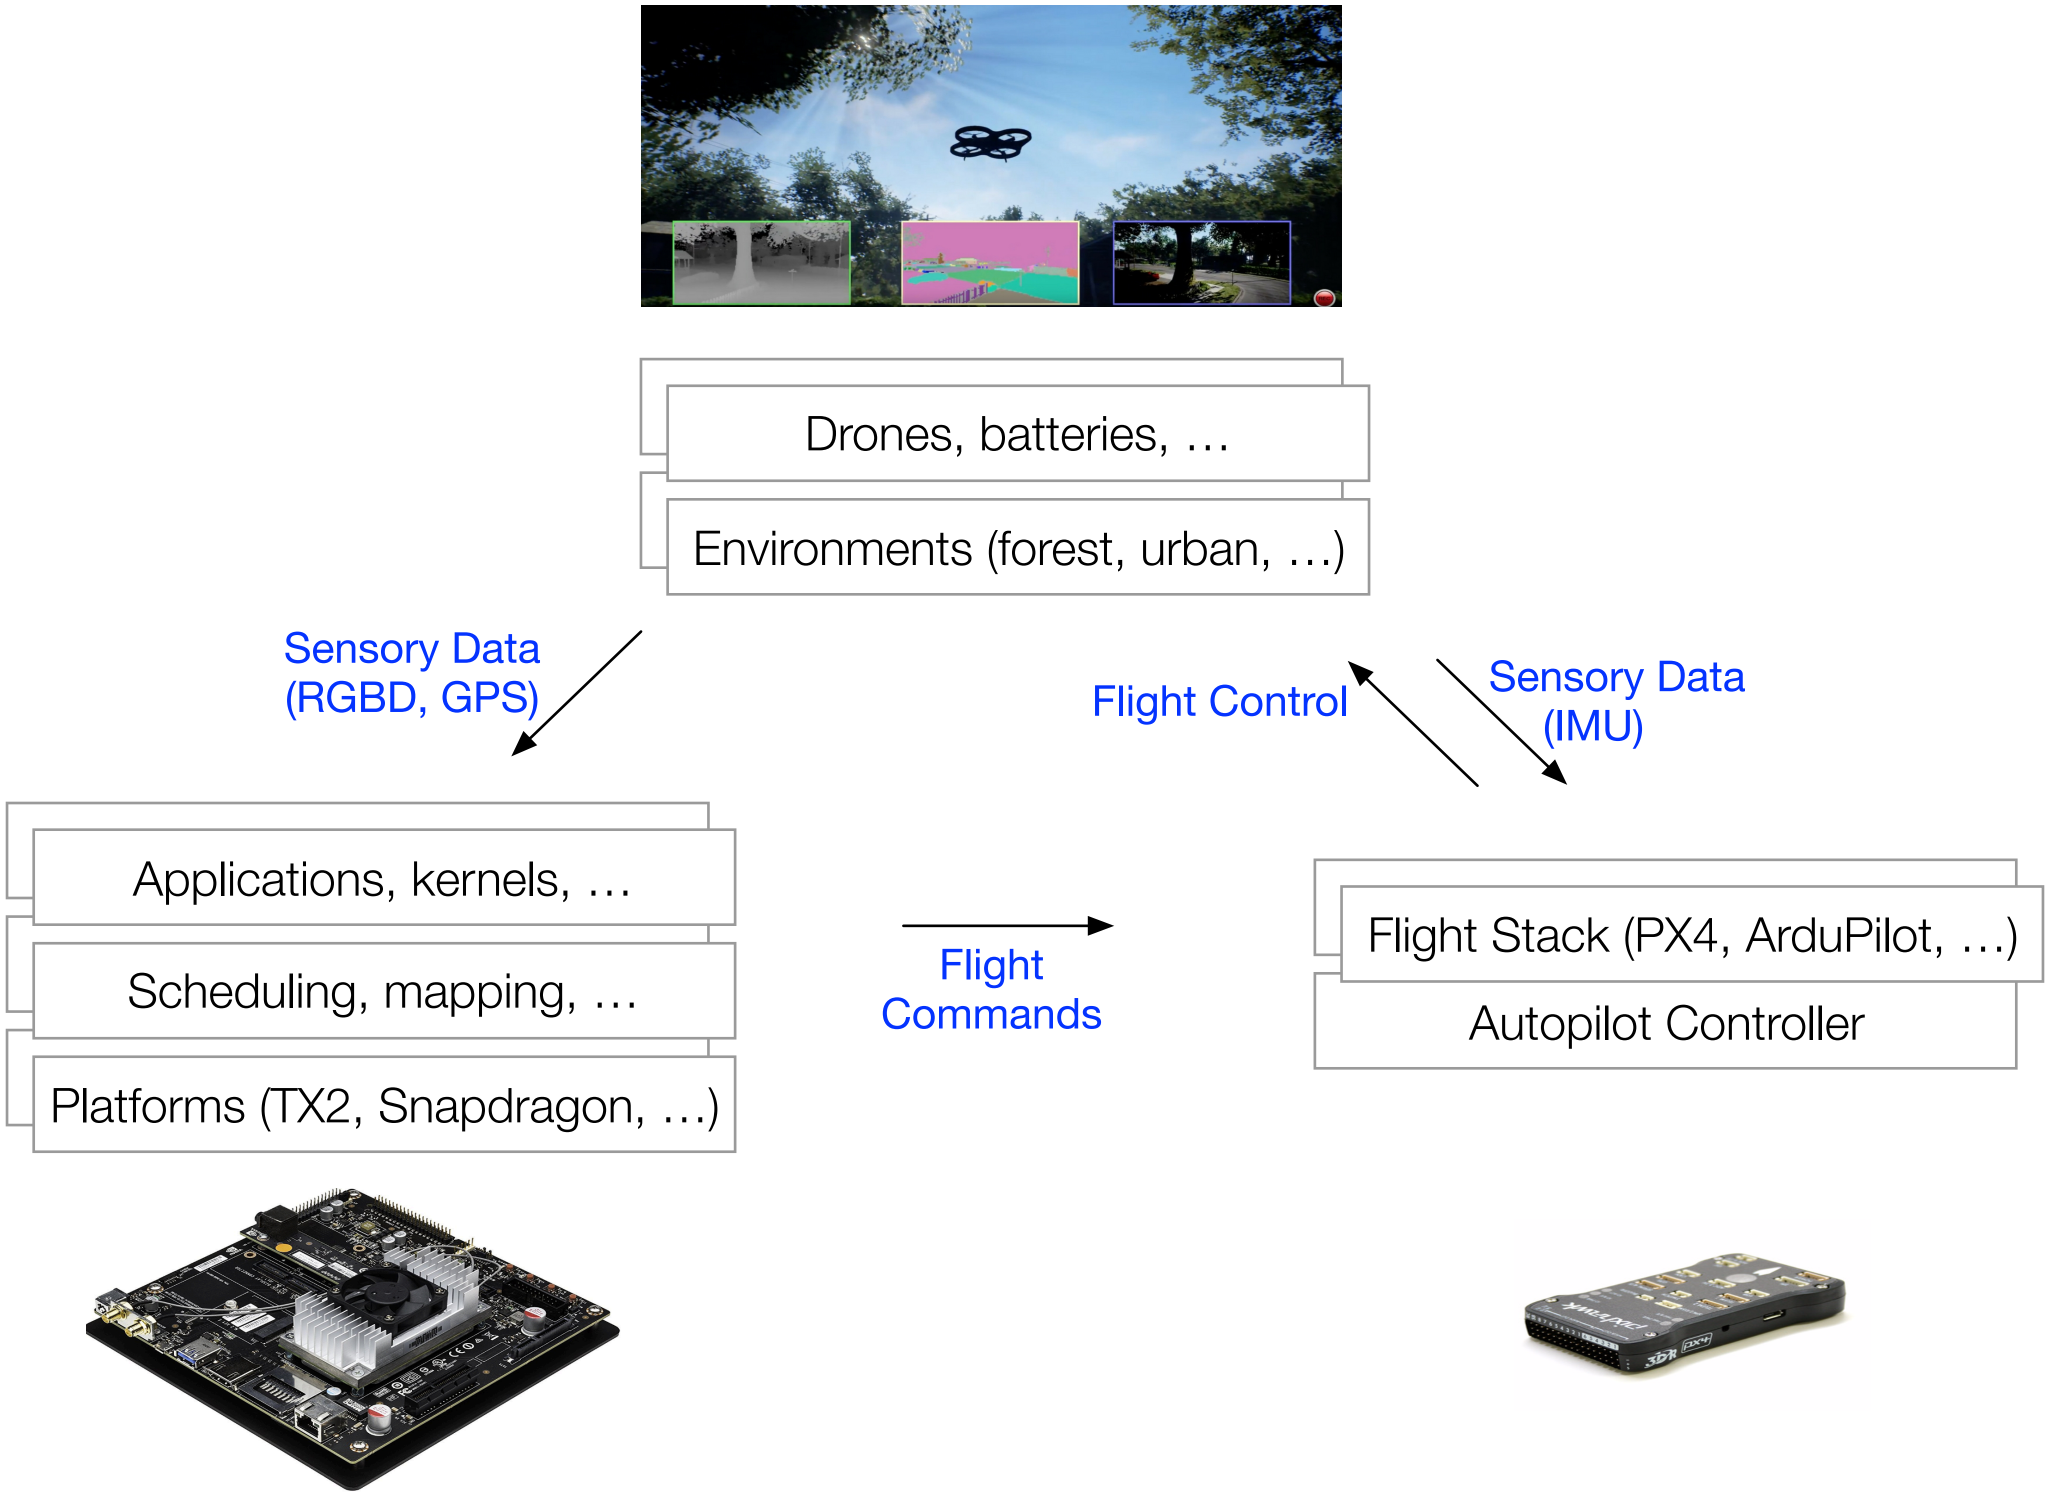

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

Autopilot Controller

Closed-loop Simulation **Knobs**

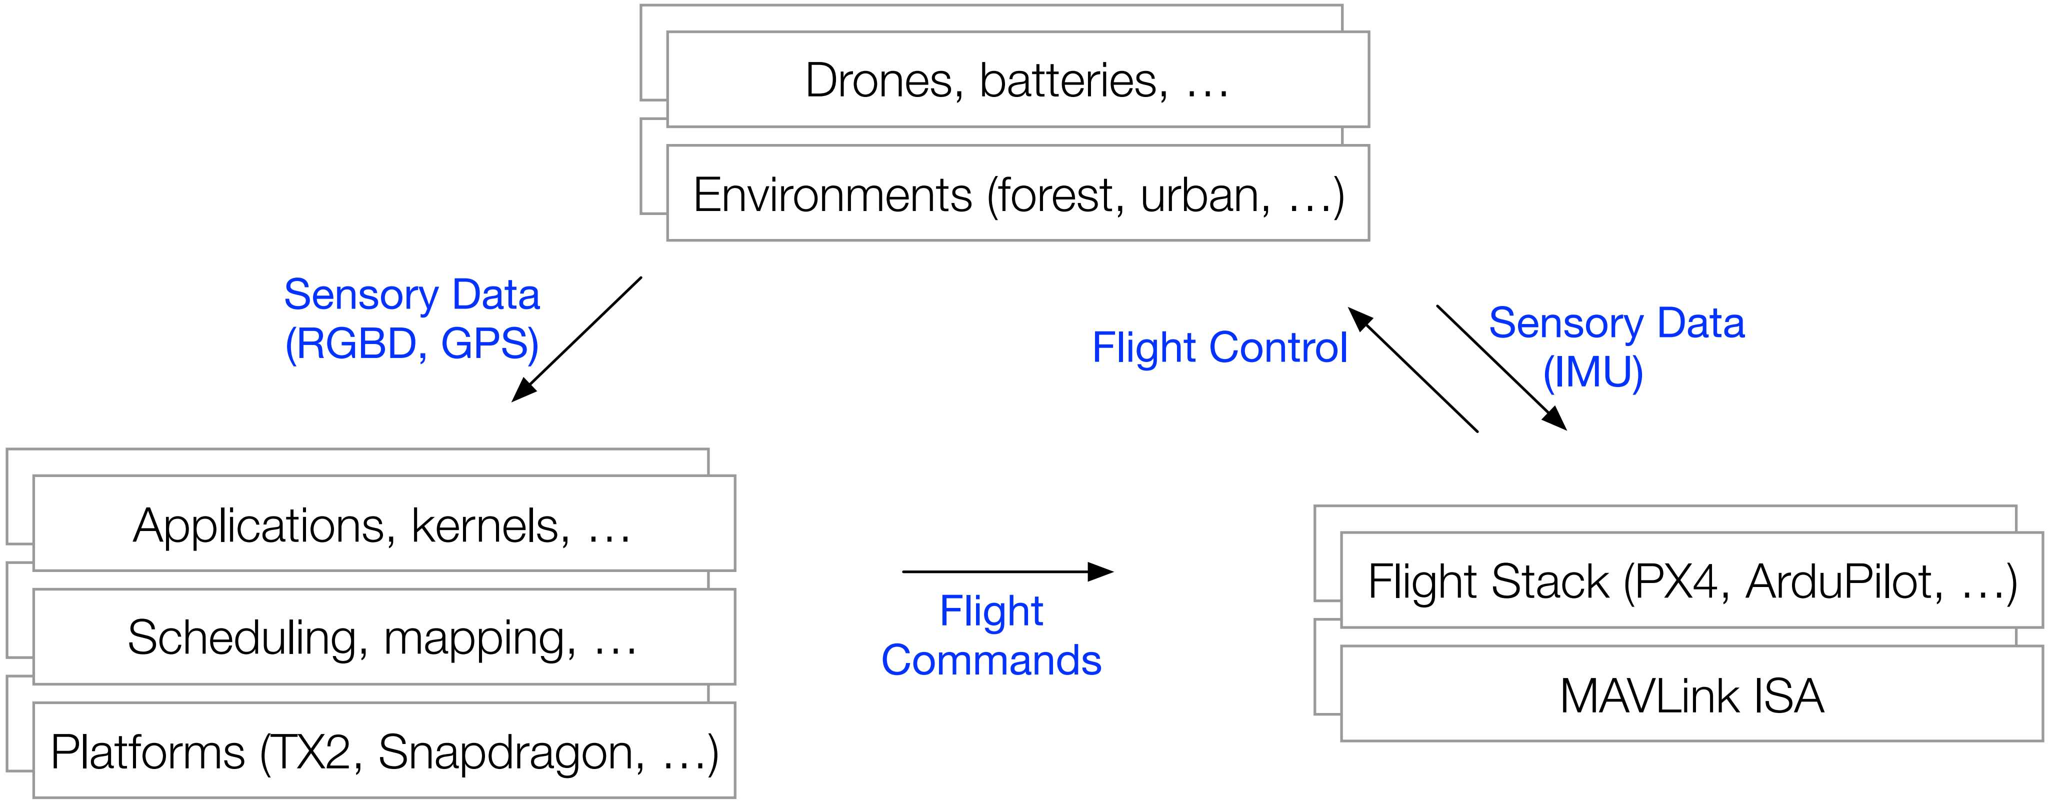

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

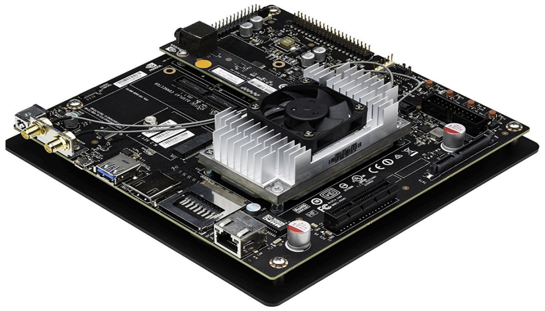

Closed-loop Simulation Profiles

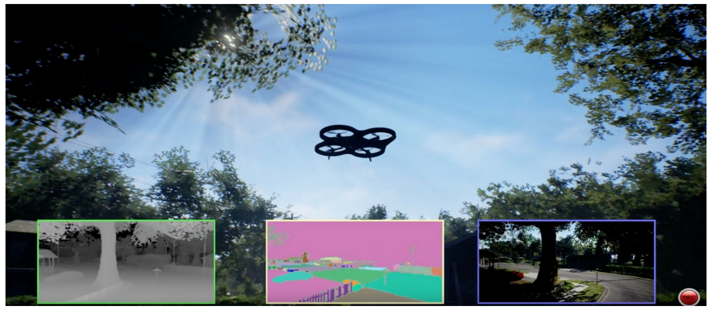

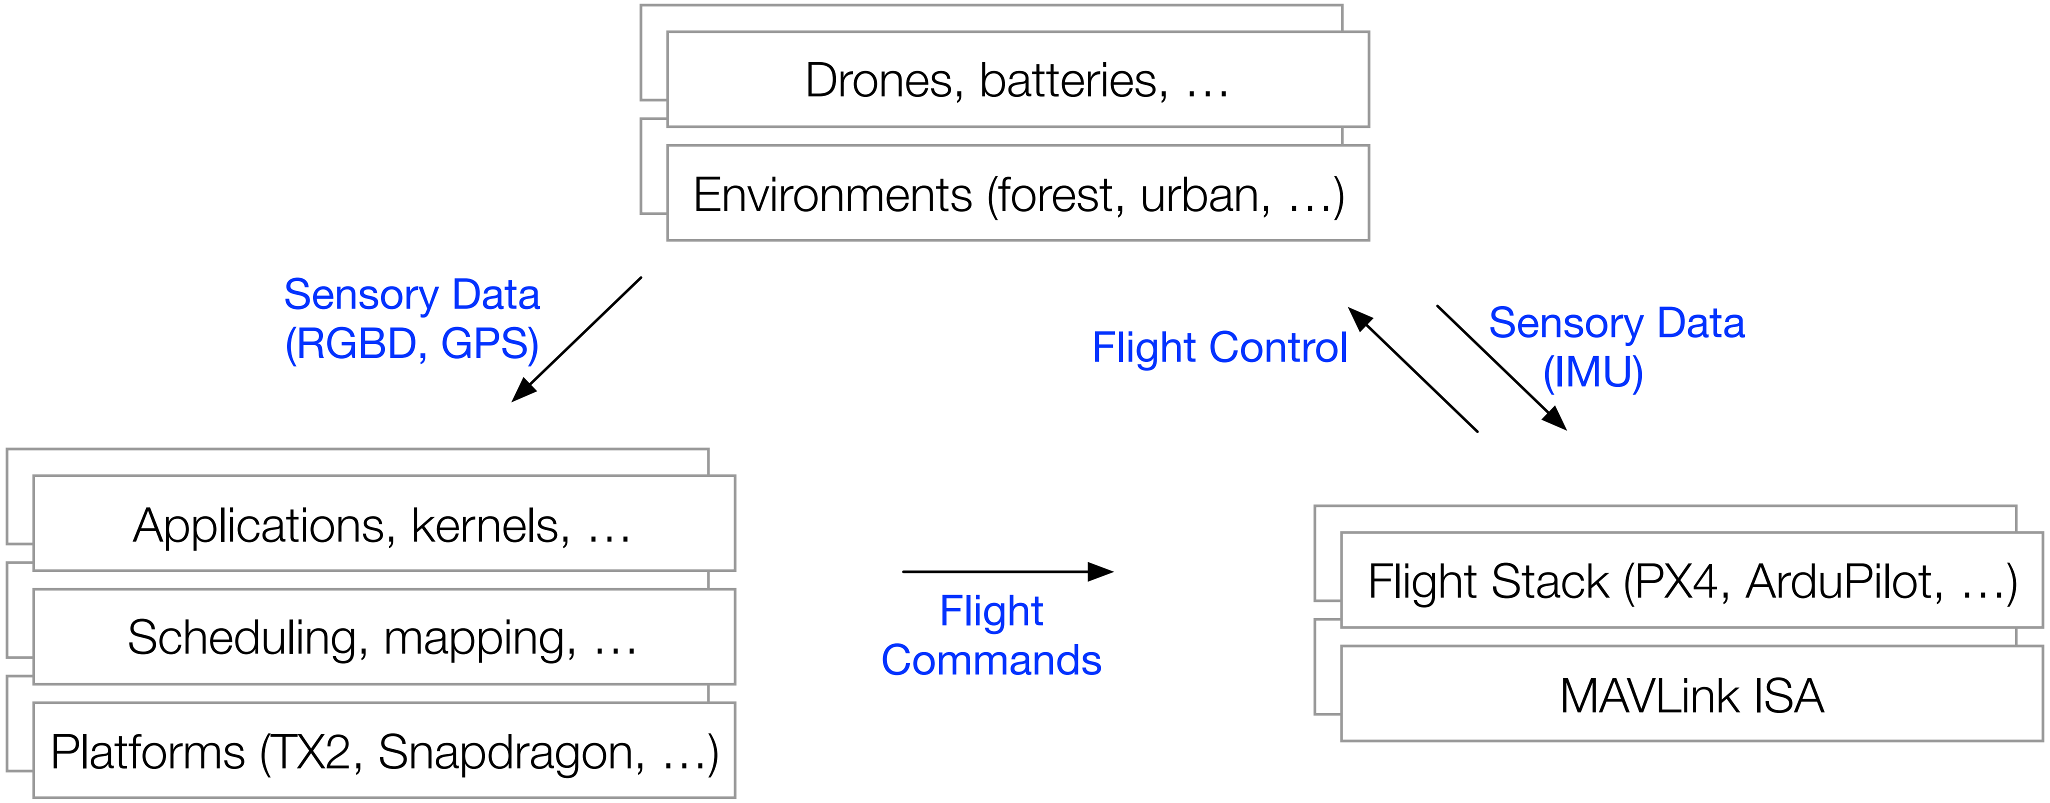

Drones, batteries, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

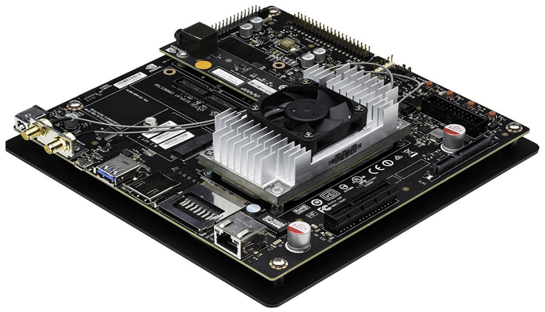

Closed-loop Simulation Profiles

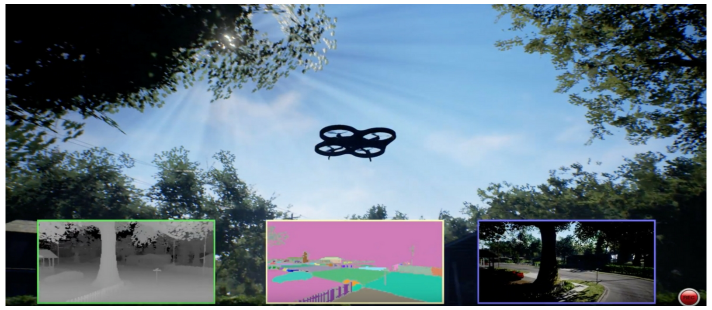

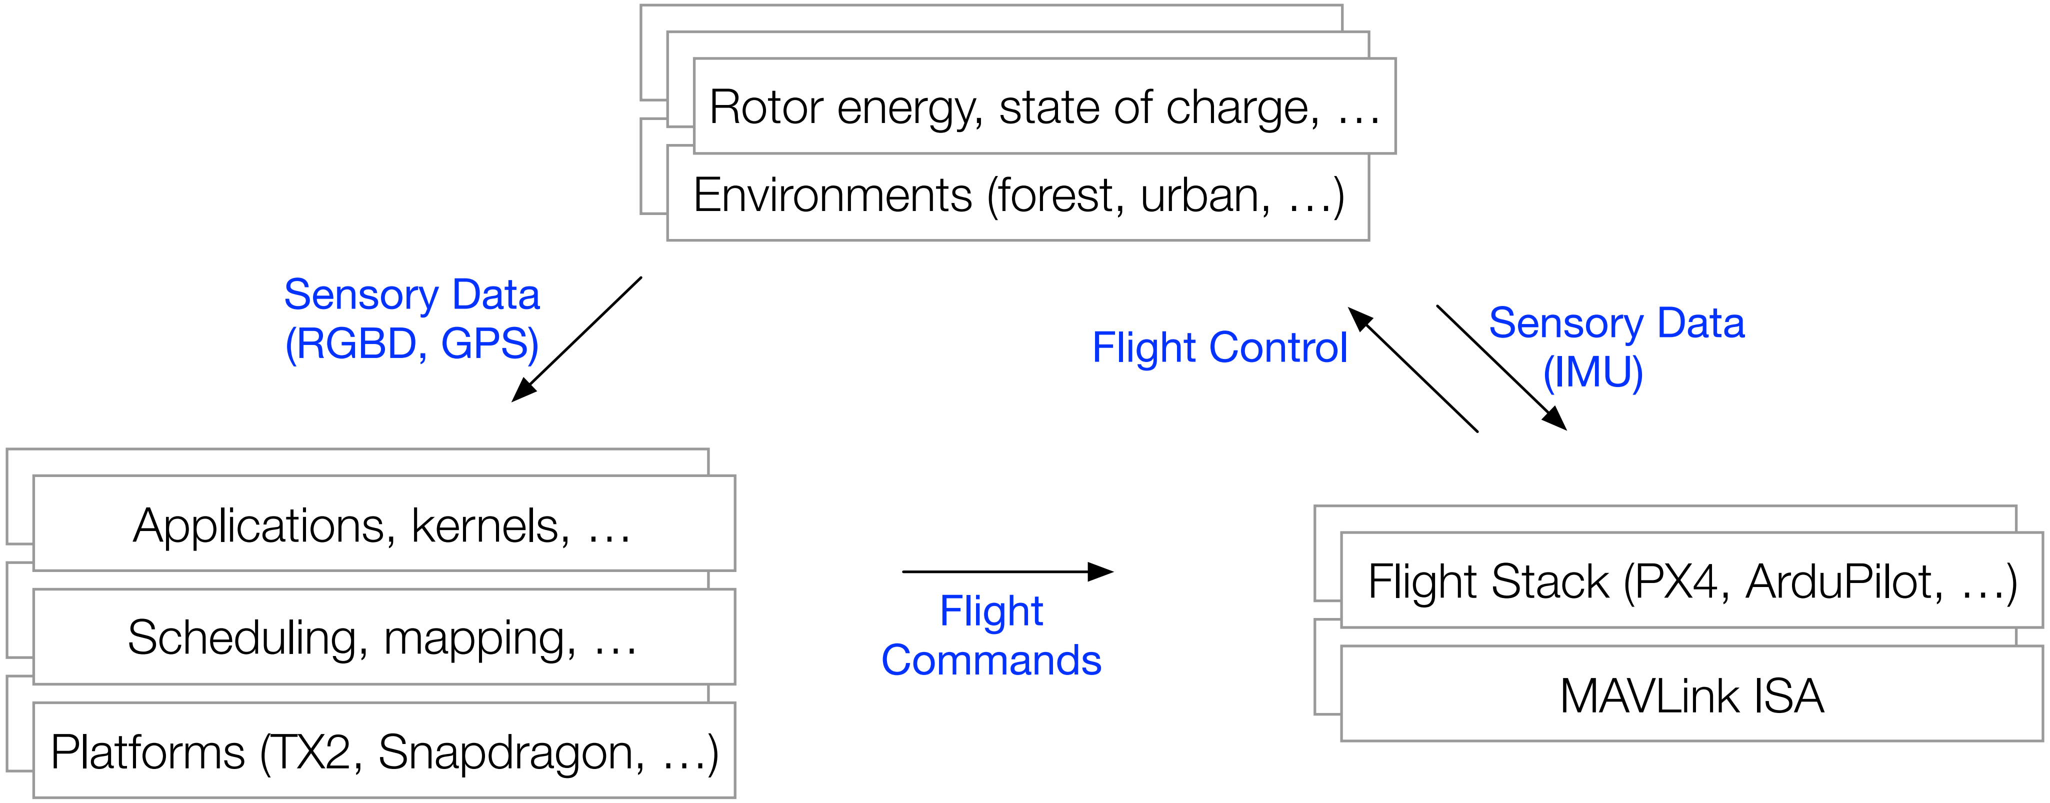

Rotor energy, state of charge, …

Environments (forest, urban, …)

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

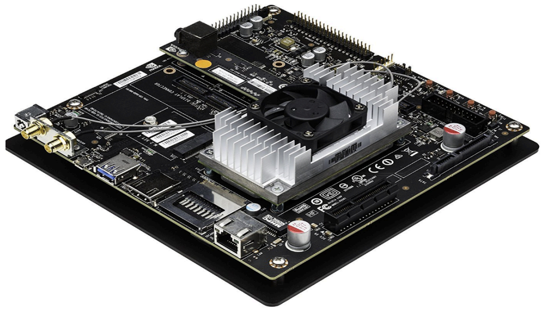

Closed-loop Simulation **Profiles**

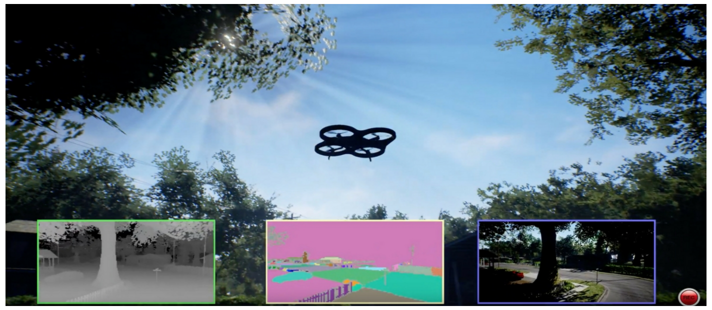

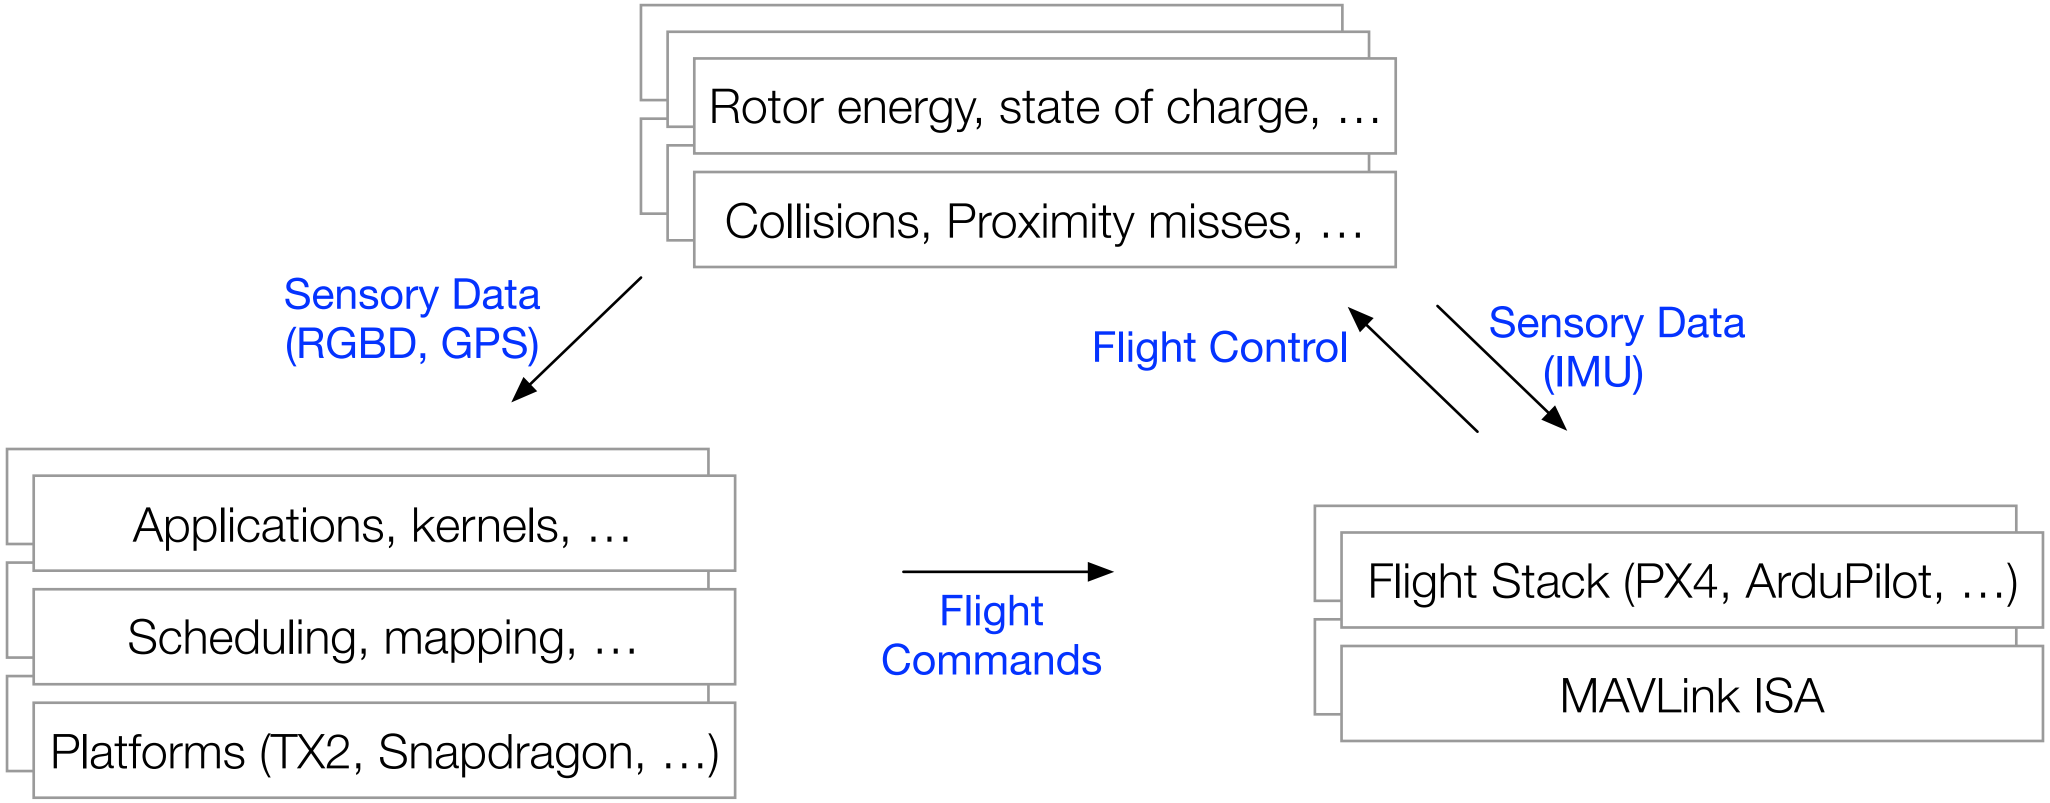

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Applications, kernels, …

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

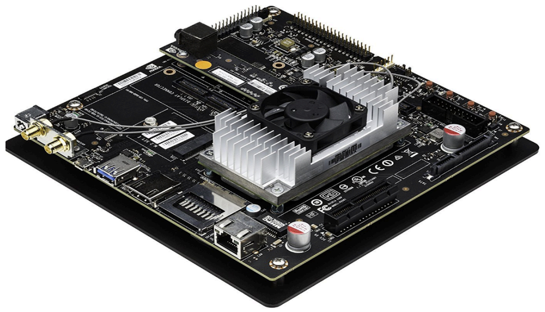

Closed-loop Simulation Profiles

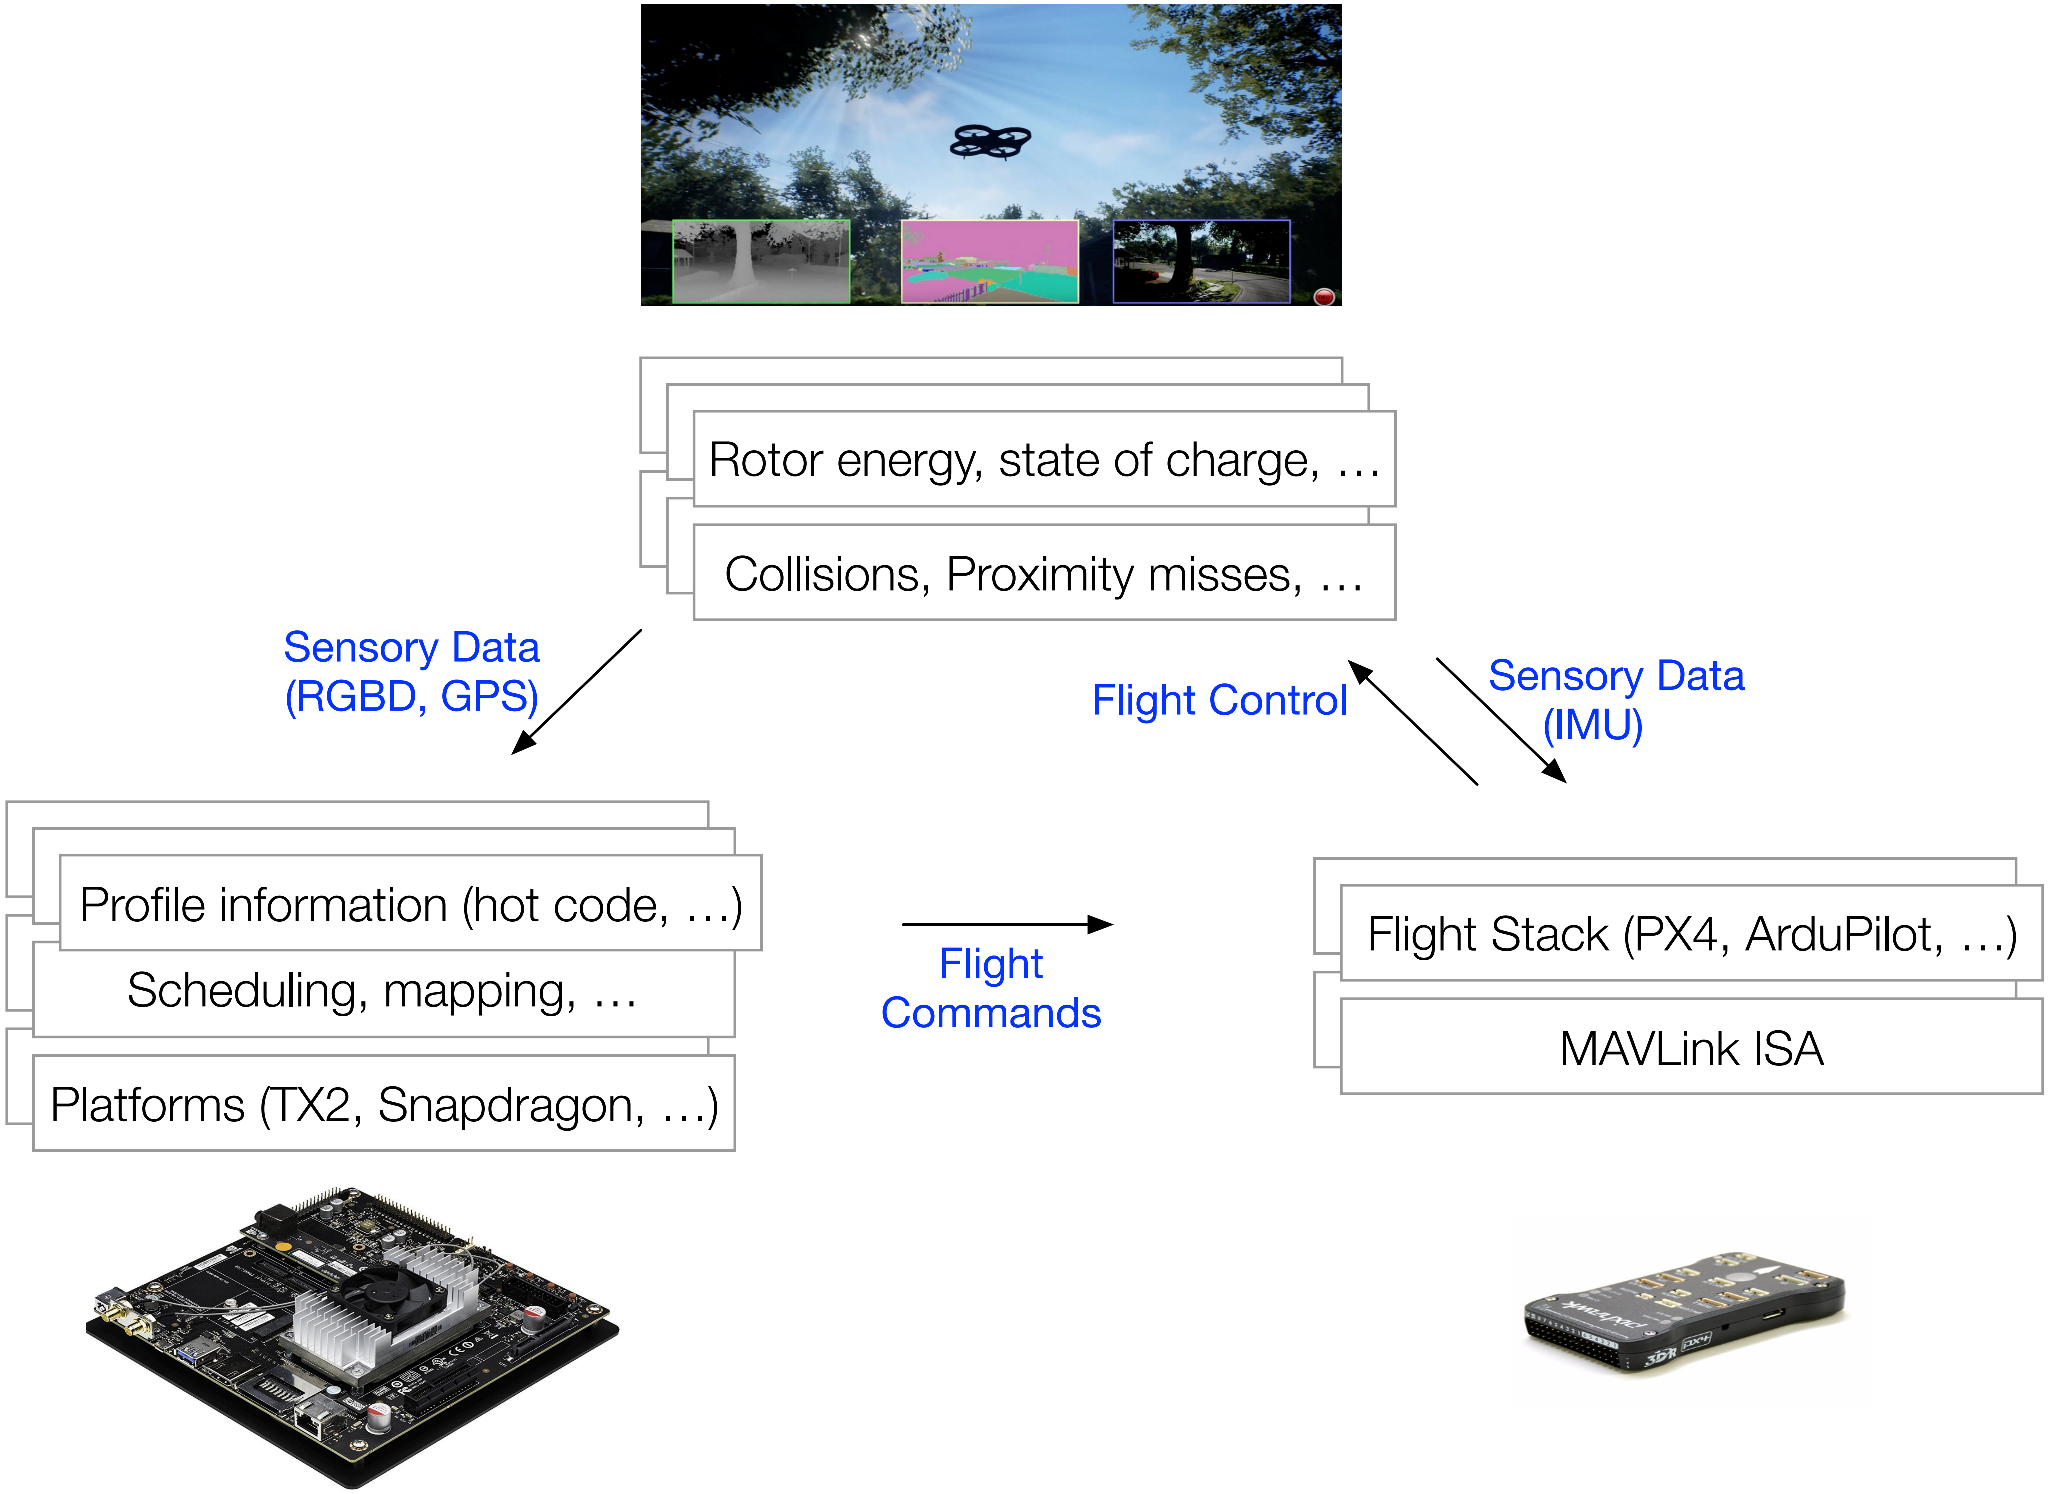

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Profile information (hot code, …)

Scheduling, mapping, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

Closed-loop Simulation **Profiles**

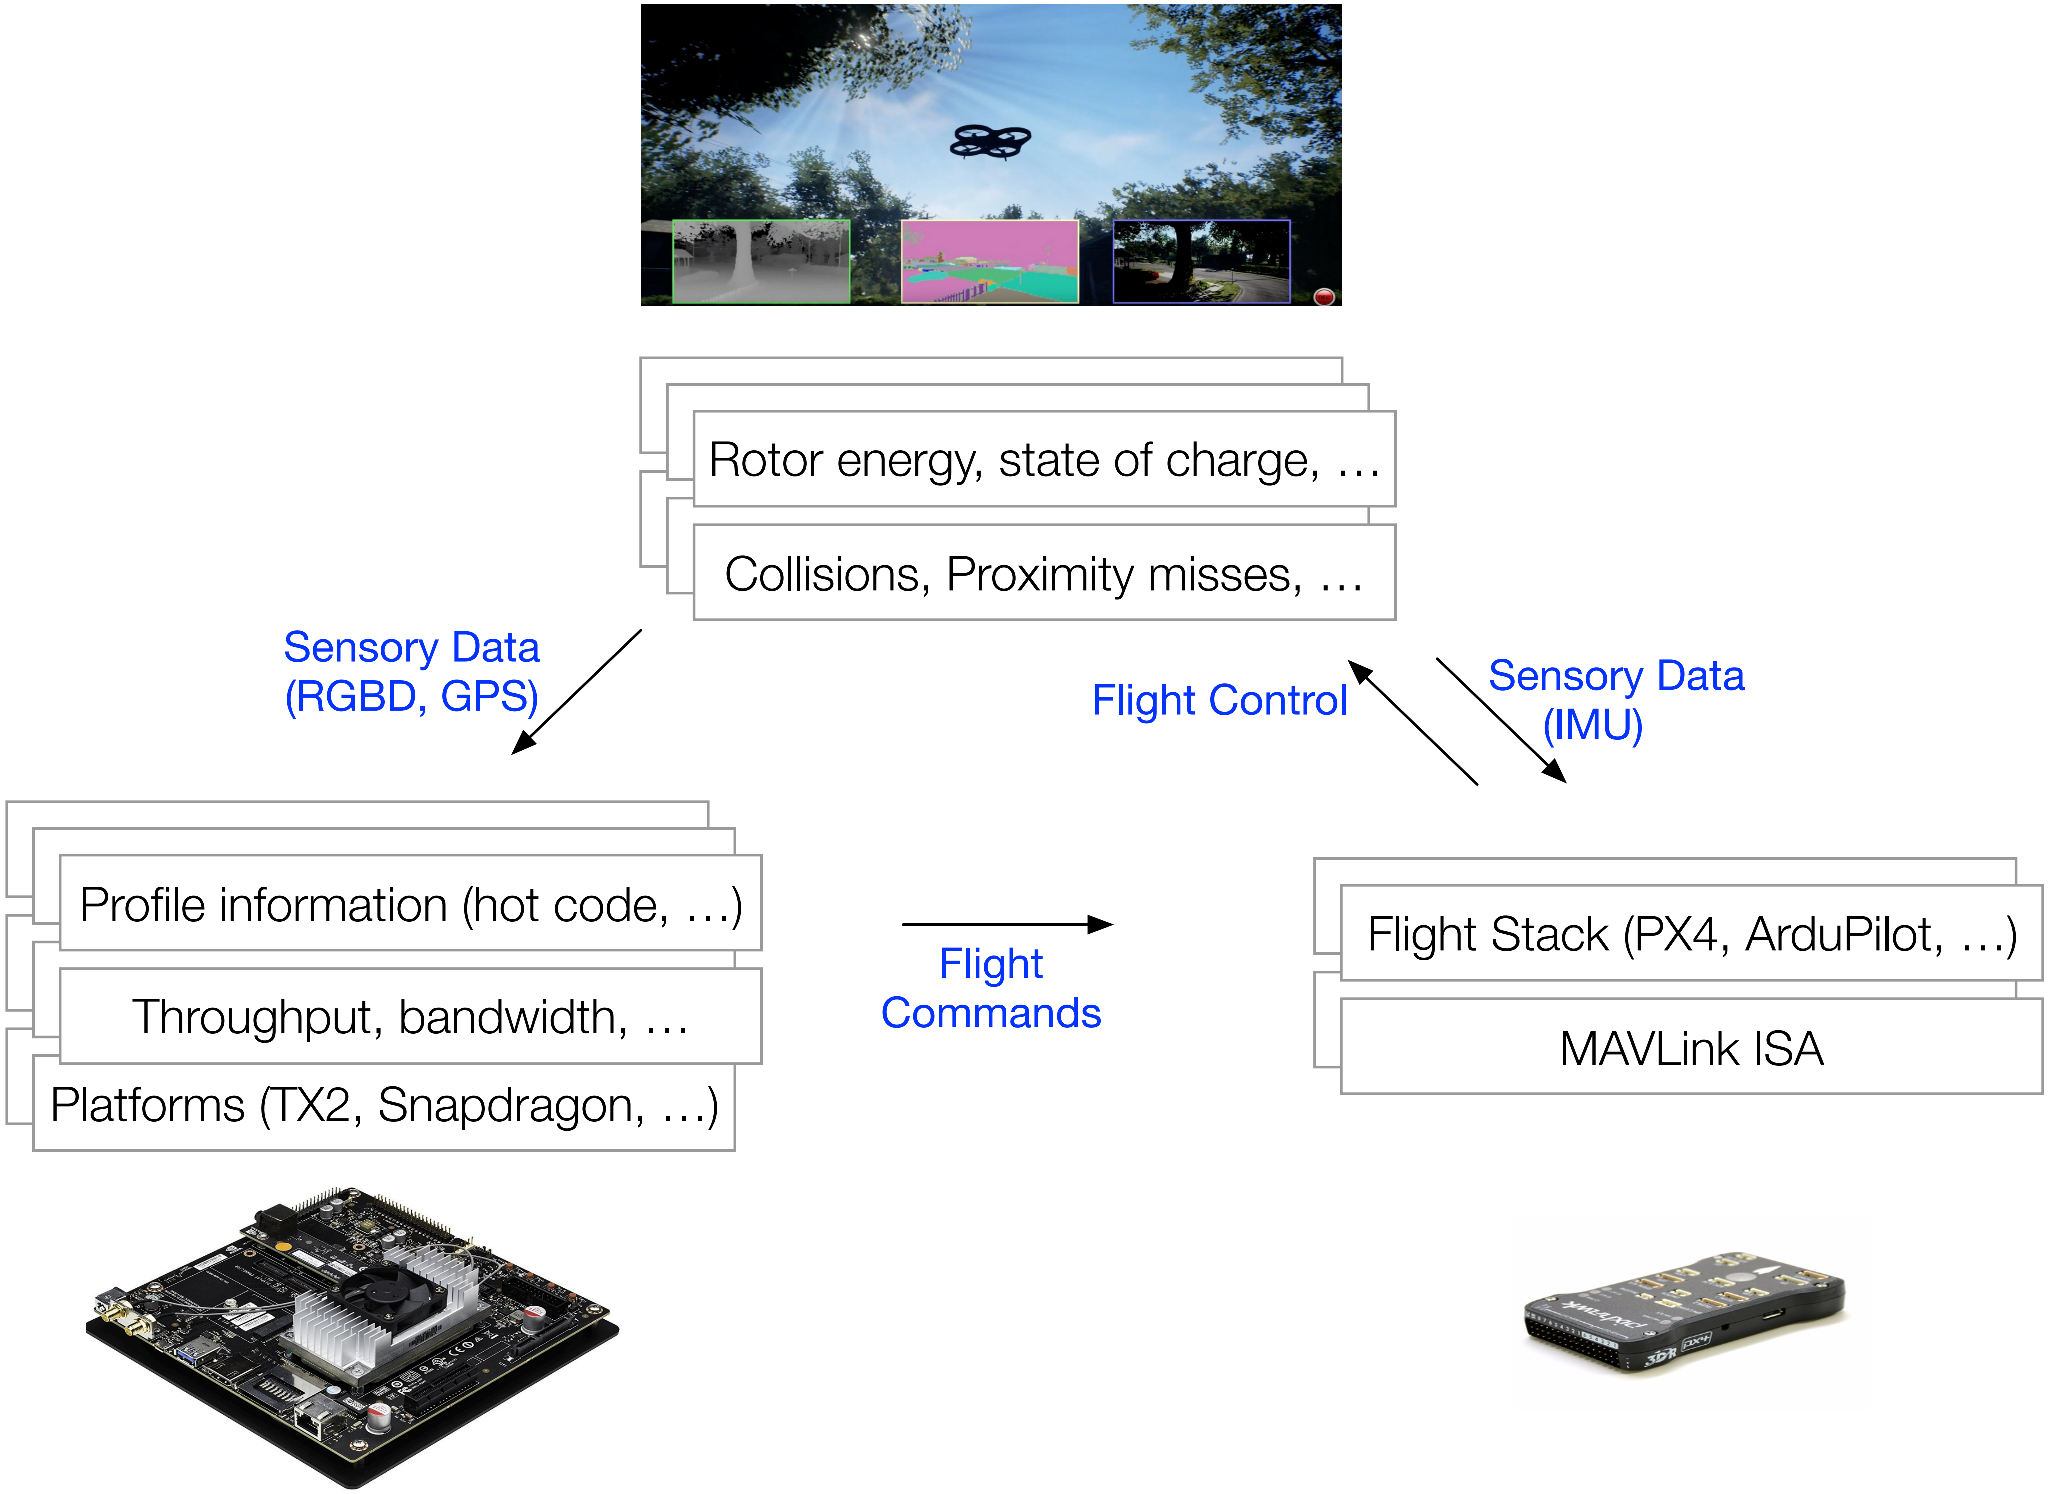

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Profile information (hot code, …)

Throughput, bandwidth, …

Platforms (TX2, Snapdragon, …)

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

Closed-loop Simulation **Profiles**

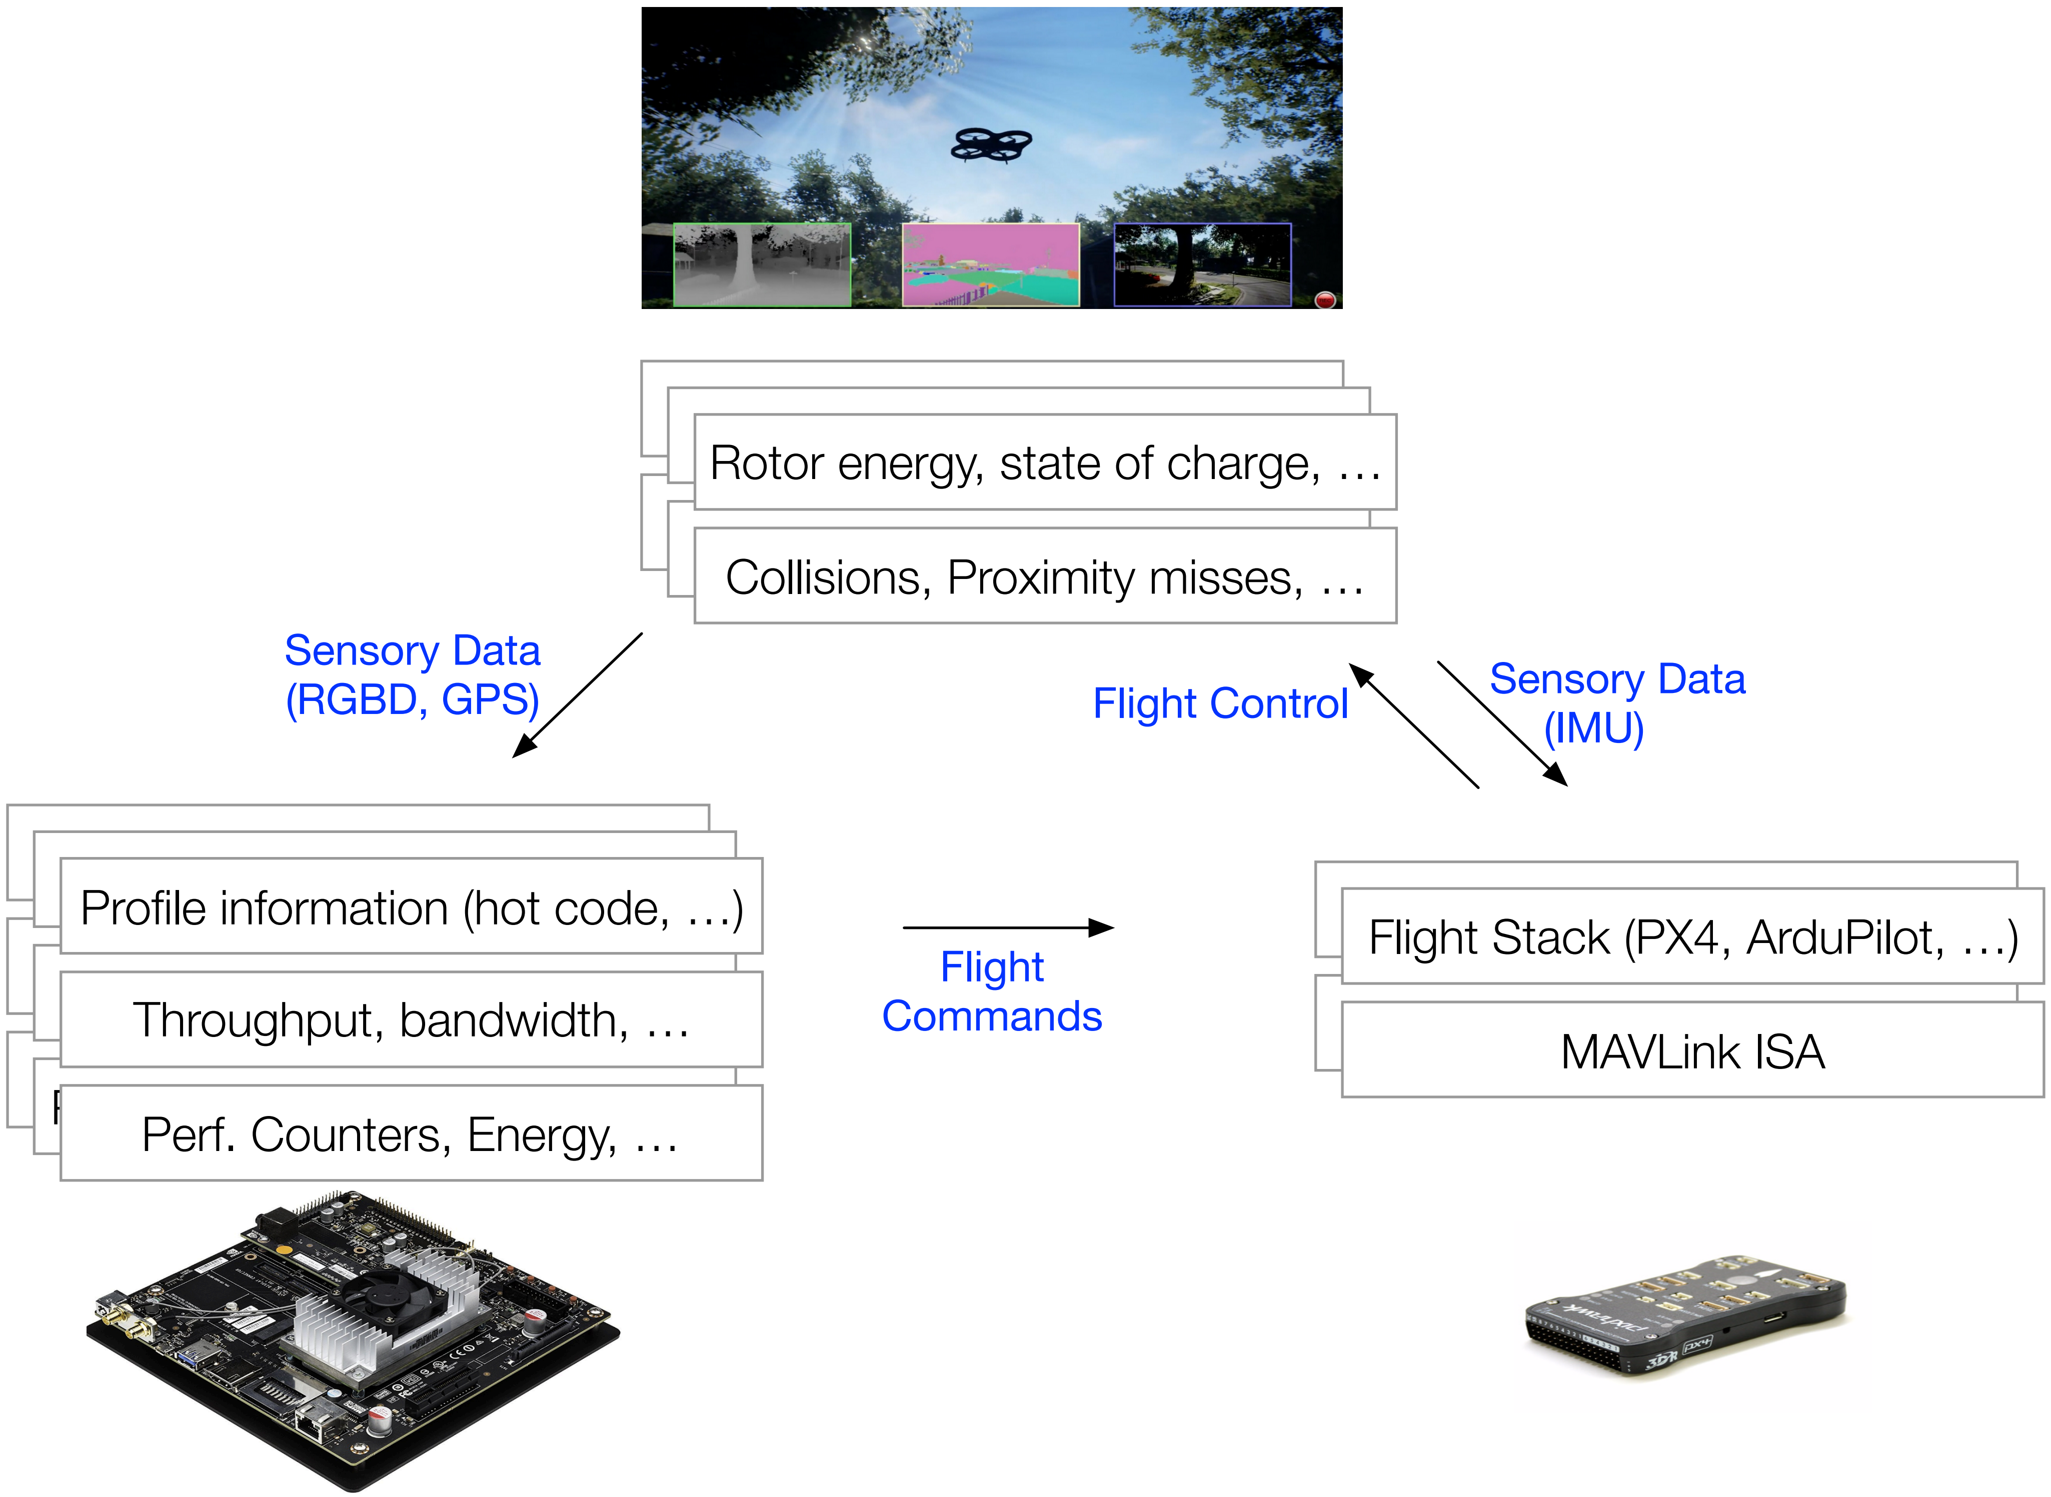

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Profile information (hot code, …)

Throughput, bandwidth, …

Perf. Counters, Energy, …

Flight Commands

Flight Stack (PX4, ArduPilot, …)

MAVLink ISA

Closed-loop Simulation **Profiles**

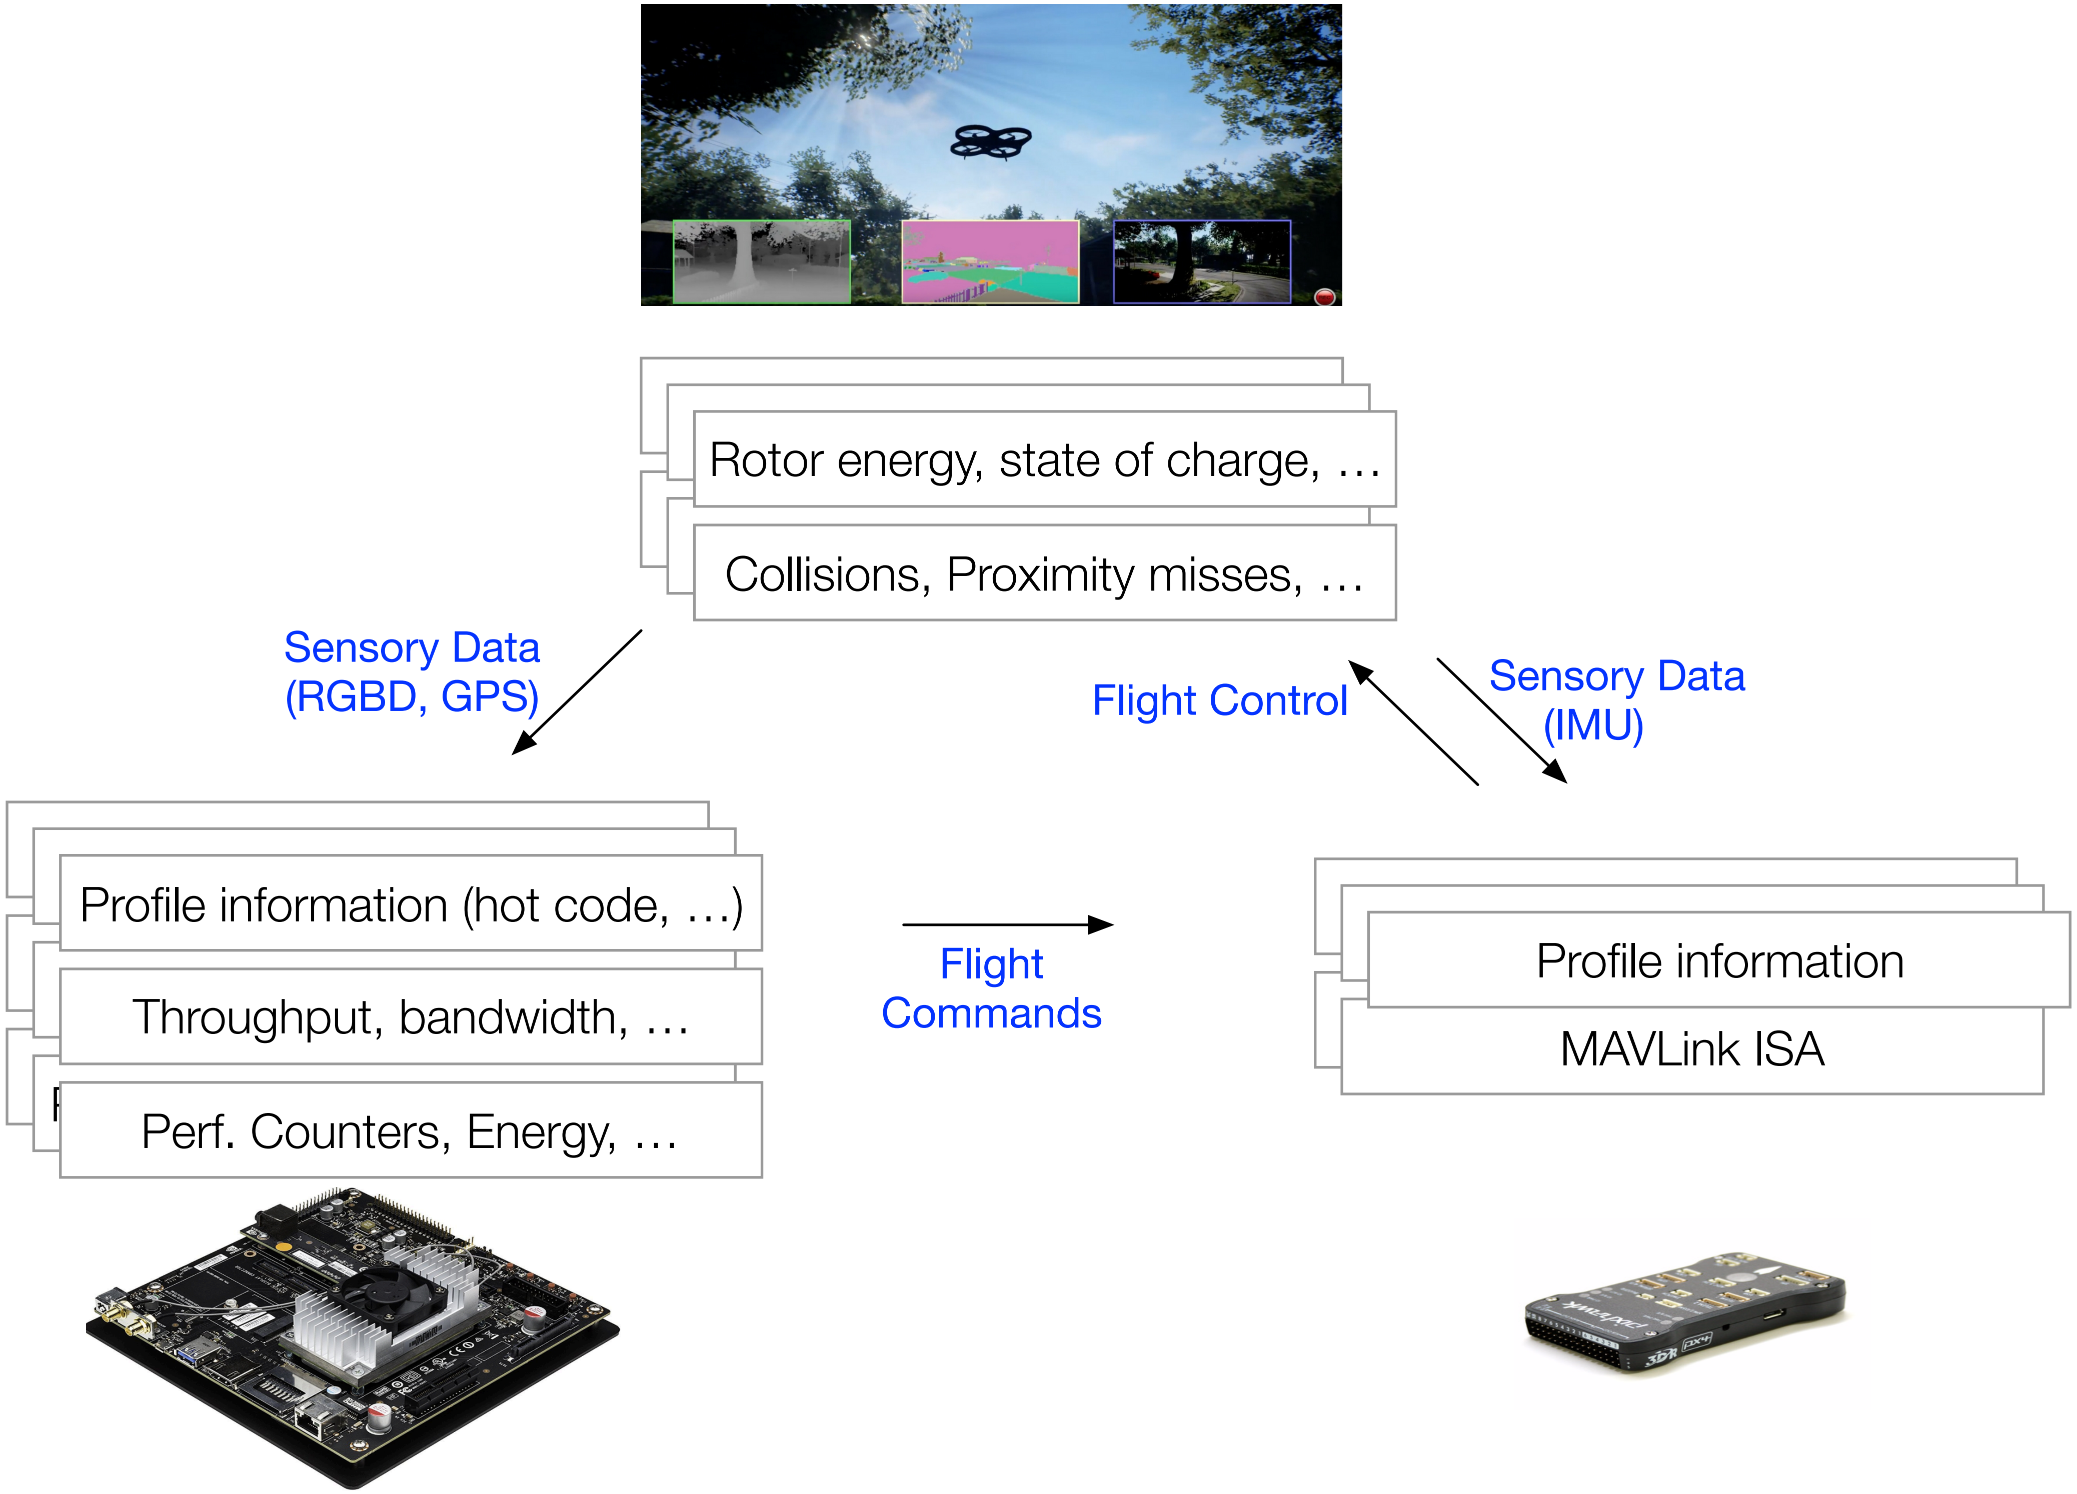

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Profile information (hot code, …)

Throughput, bandwidth, …

Perf. Counters, Energy, …

Flight Commands

Profile information

MAVLink ISA

Closed-loop Simulation **Profiles**

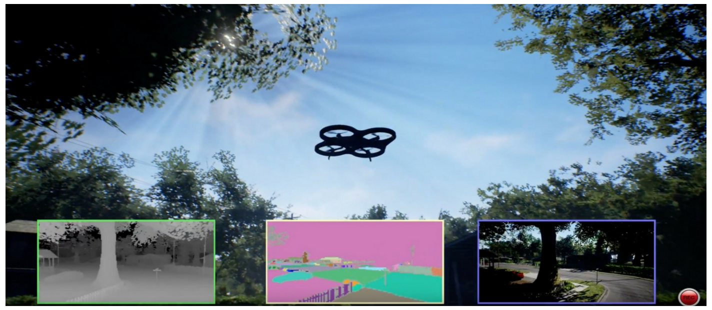

Rotor energy, state of charge, …

Collisions, Proximity misses, …

Sensory Data (RGBD, GPS)

Flight Control

Sensory Data (IMU)

Profile information (hot code, …)

Throughput, bandwidth, …

Perf. Counters, Energy, …

Flight Commands

Profile information

Frequent commands, …

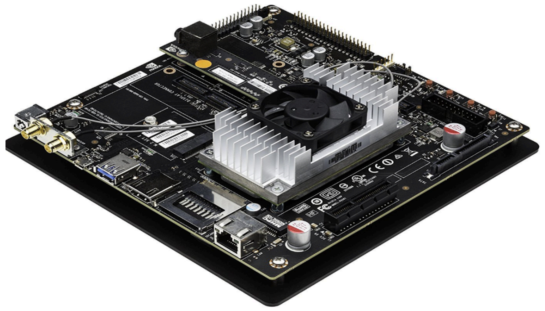

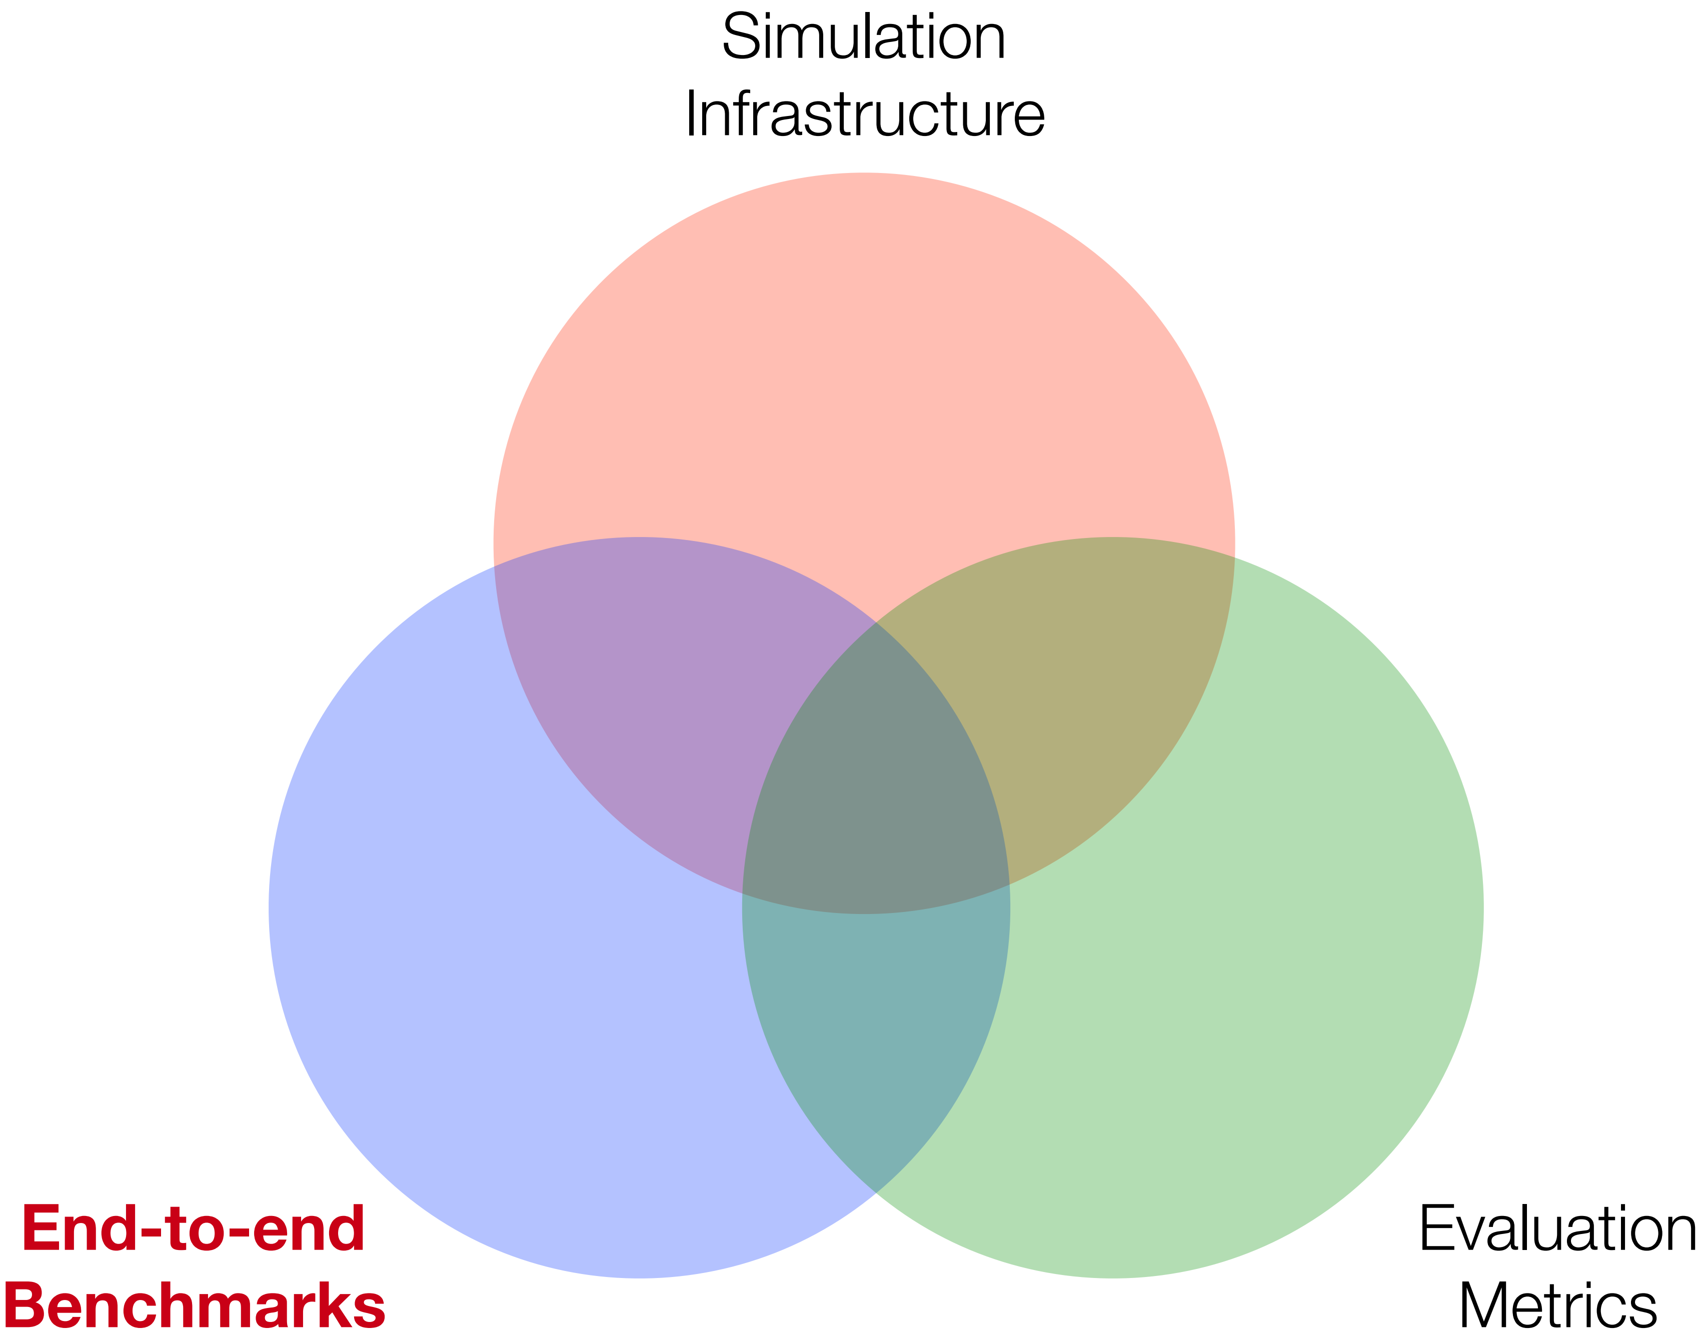

Simulation Infrastructure

End-to-end Benchmarks

Evaluation Metrics

Autonomous Package Delivery

Autonomous Package Delivery

Read depth sensors

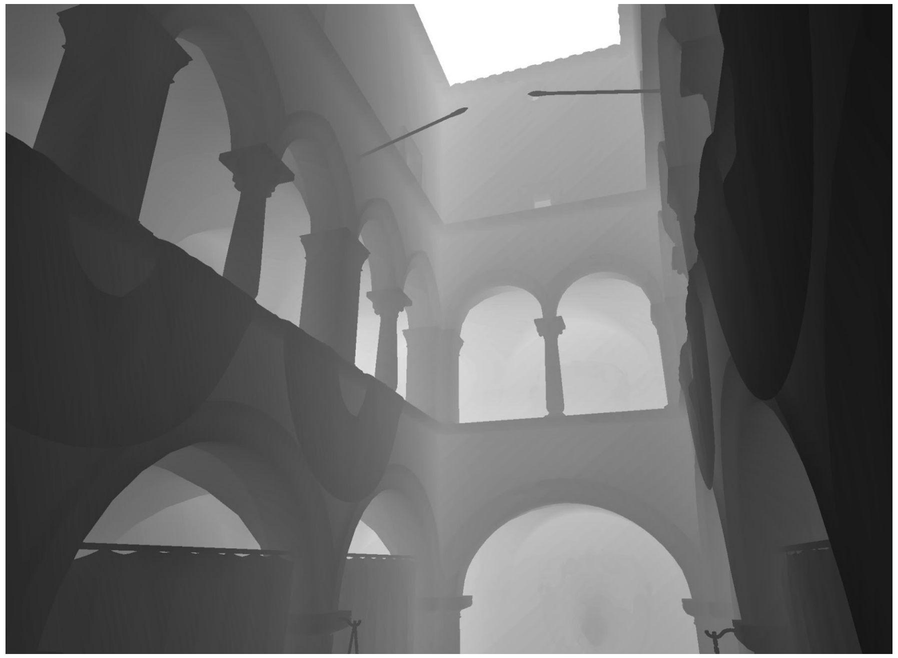

Autonomous Package Delivery

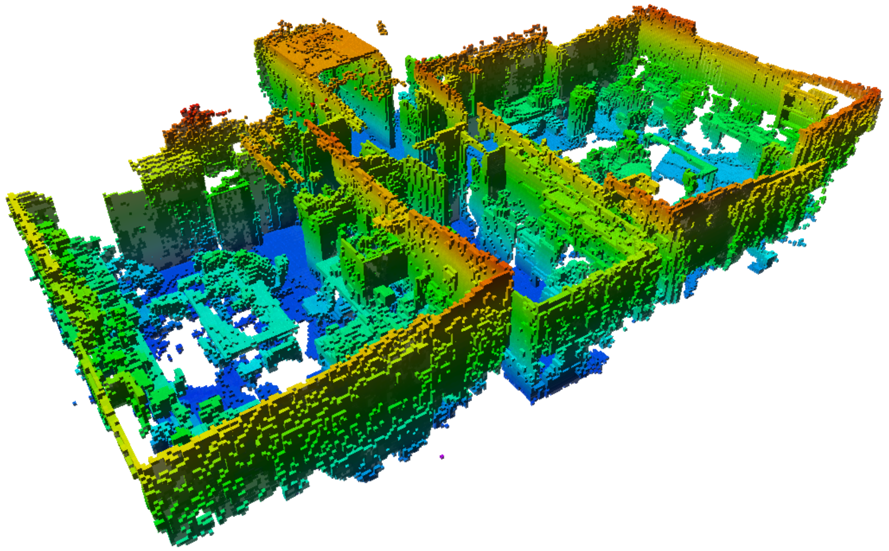

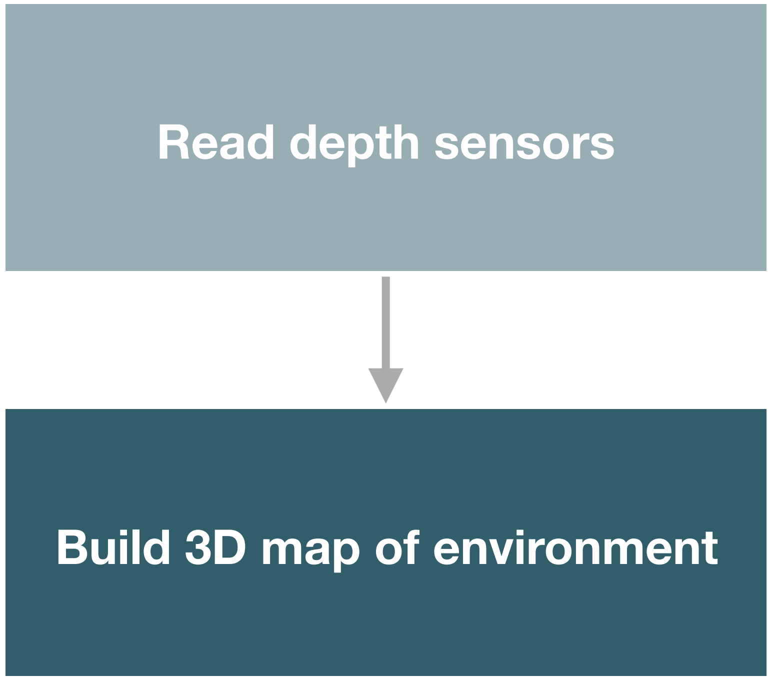

Read depth sensors

Build 3D map of environment

Autonomous Package Delivery

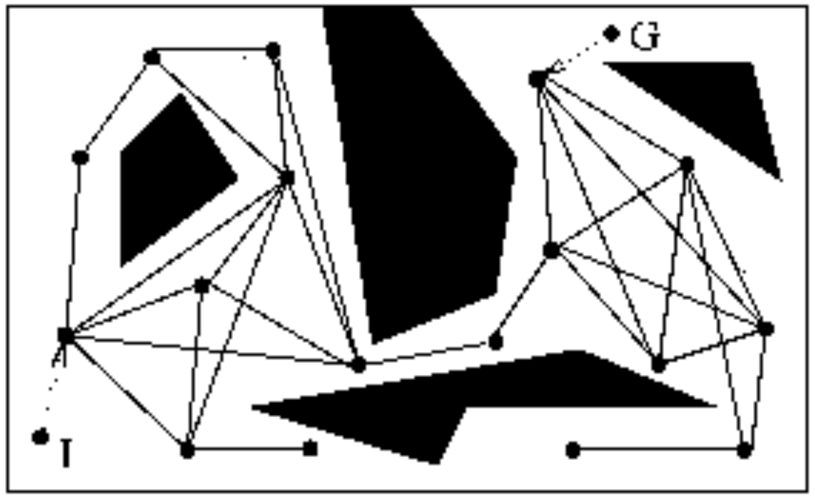

Read depth sensors

Build 3D map of environment

Sample environment

Autonomous Package Delivery

Read depth sensors

Build 3D map of environment

Sample environment

Try to find piecewise path

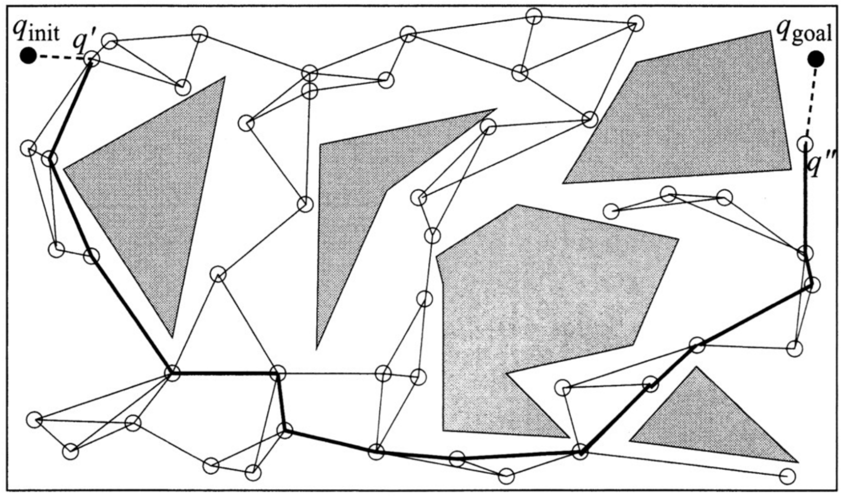

Autonomous Package Delivery

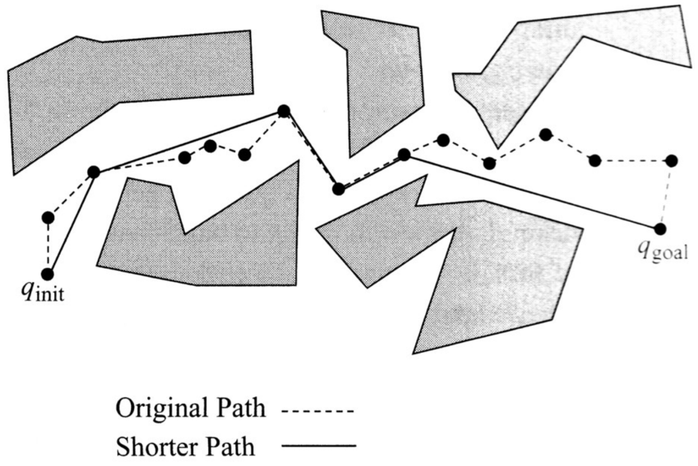

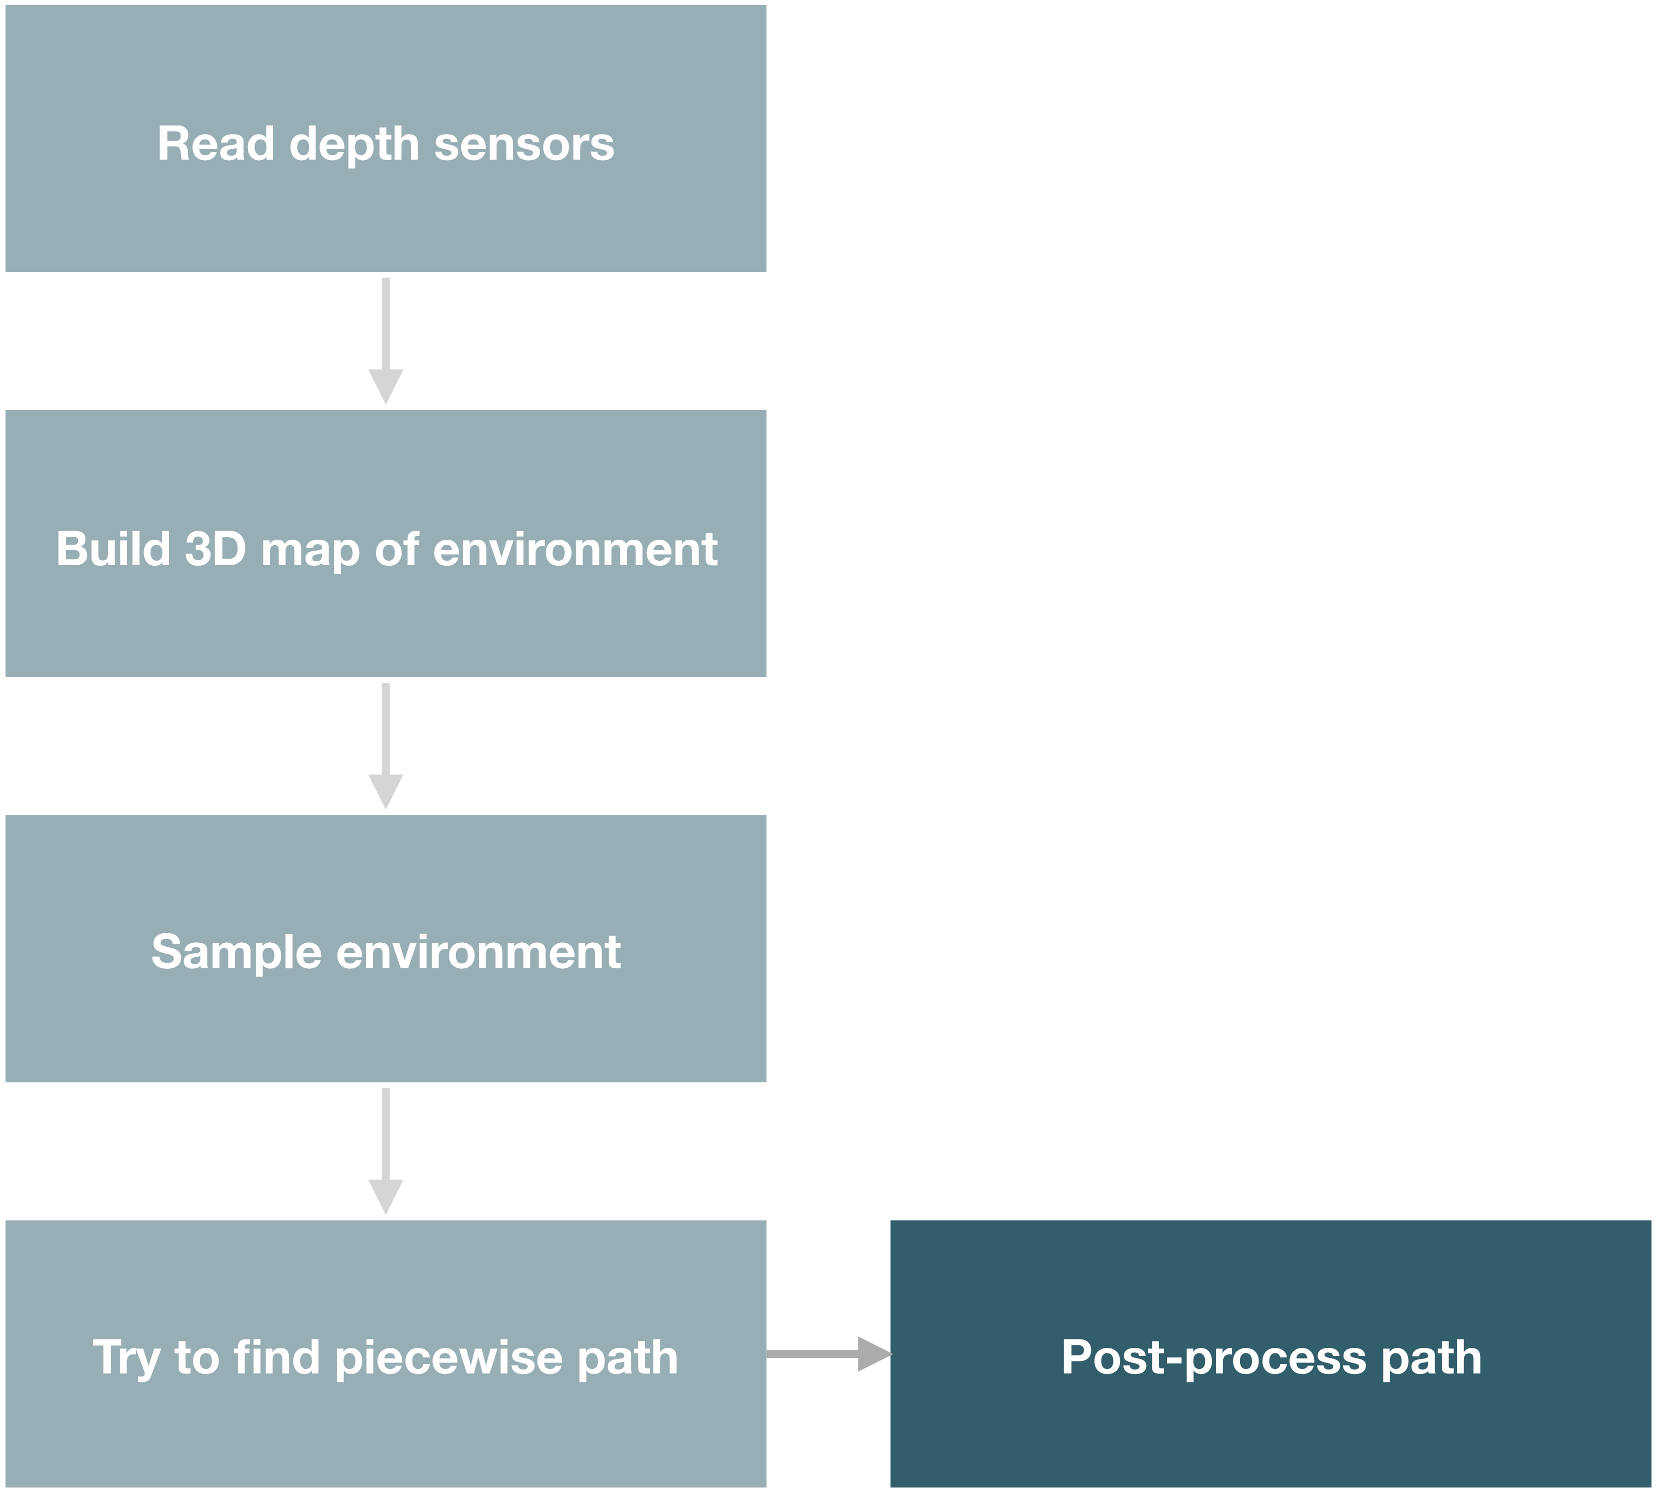
Read depth sensors

Build 3D map of environment

Sample environment

Try to find piecewise path

Post-process path

Autonomous Package Delivery

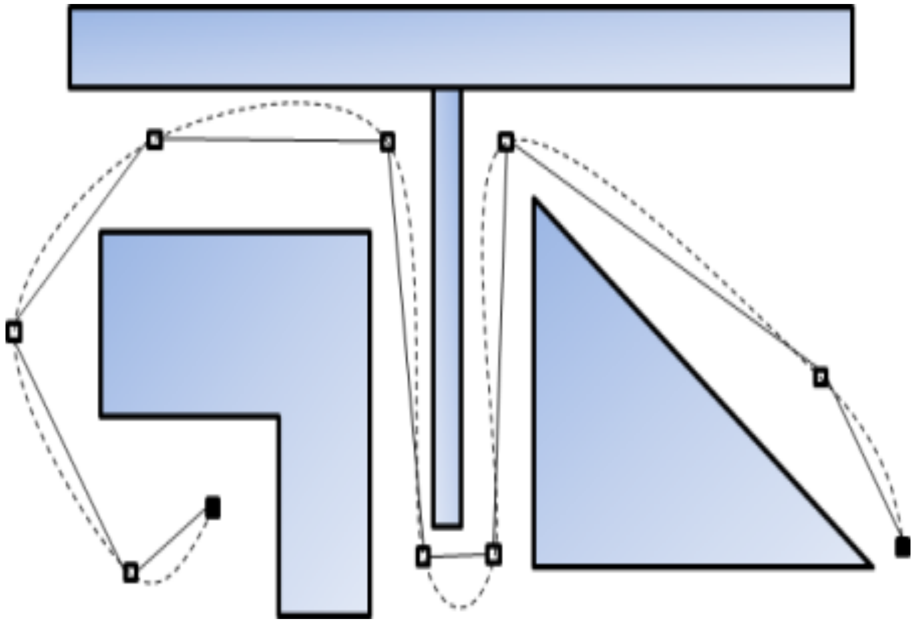

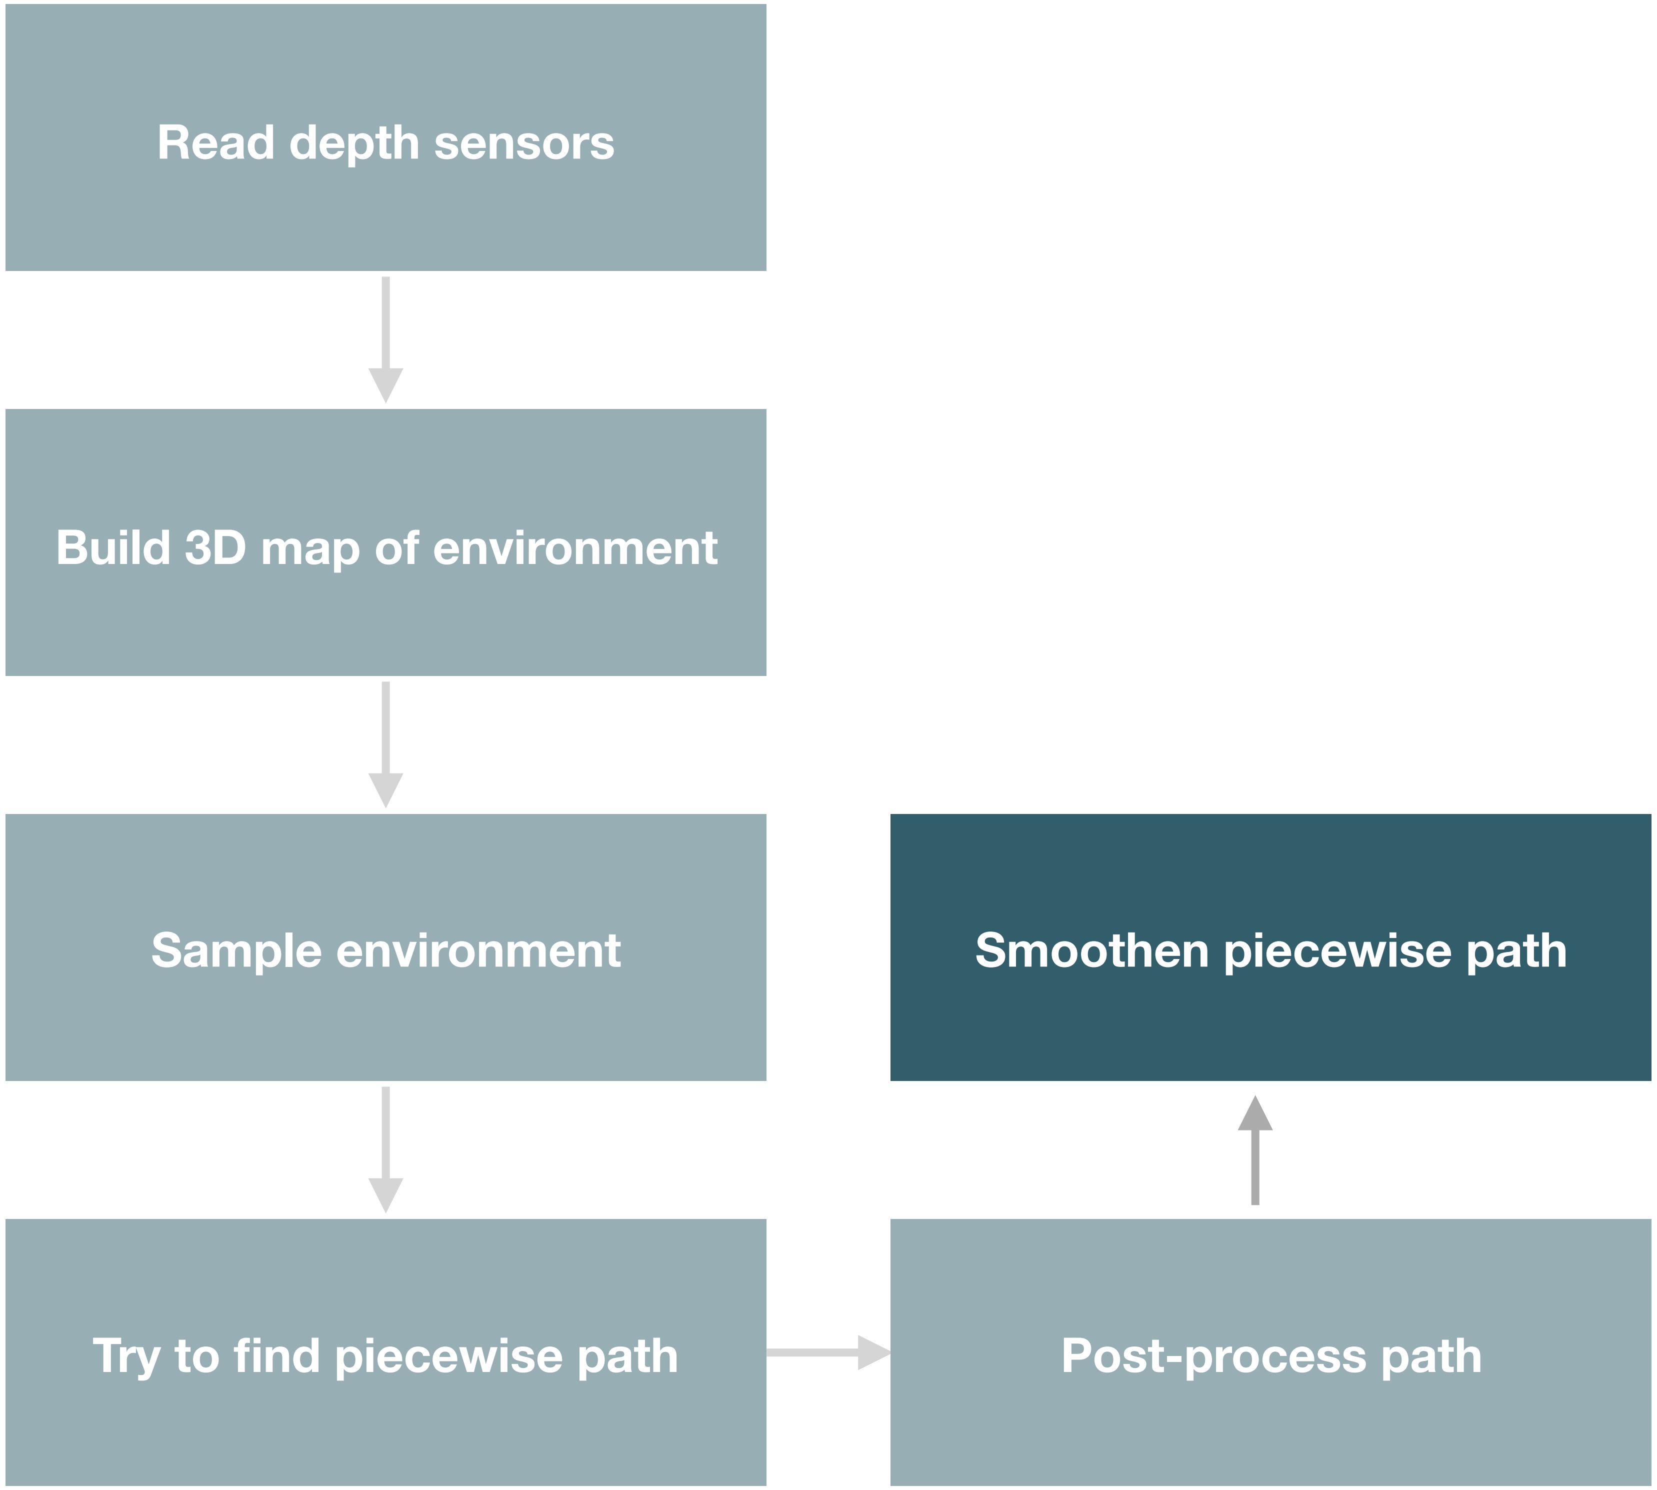

Read depth sensors

Build 3D map of environment

Sample environment

Try to find piecewise path

Post-process path

Smoothen piecewise path

Autonomous Package Delivery

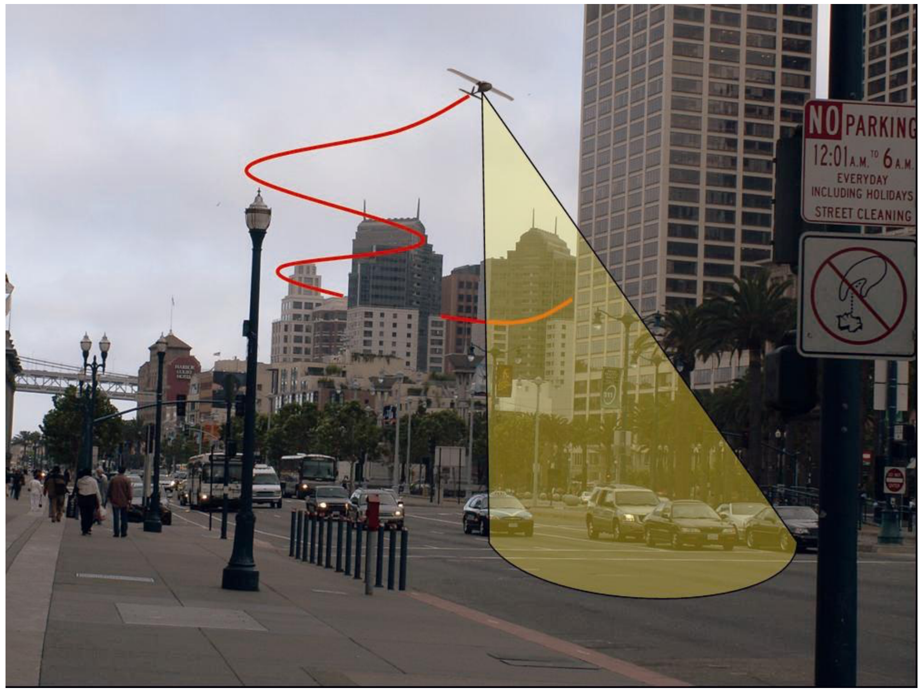

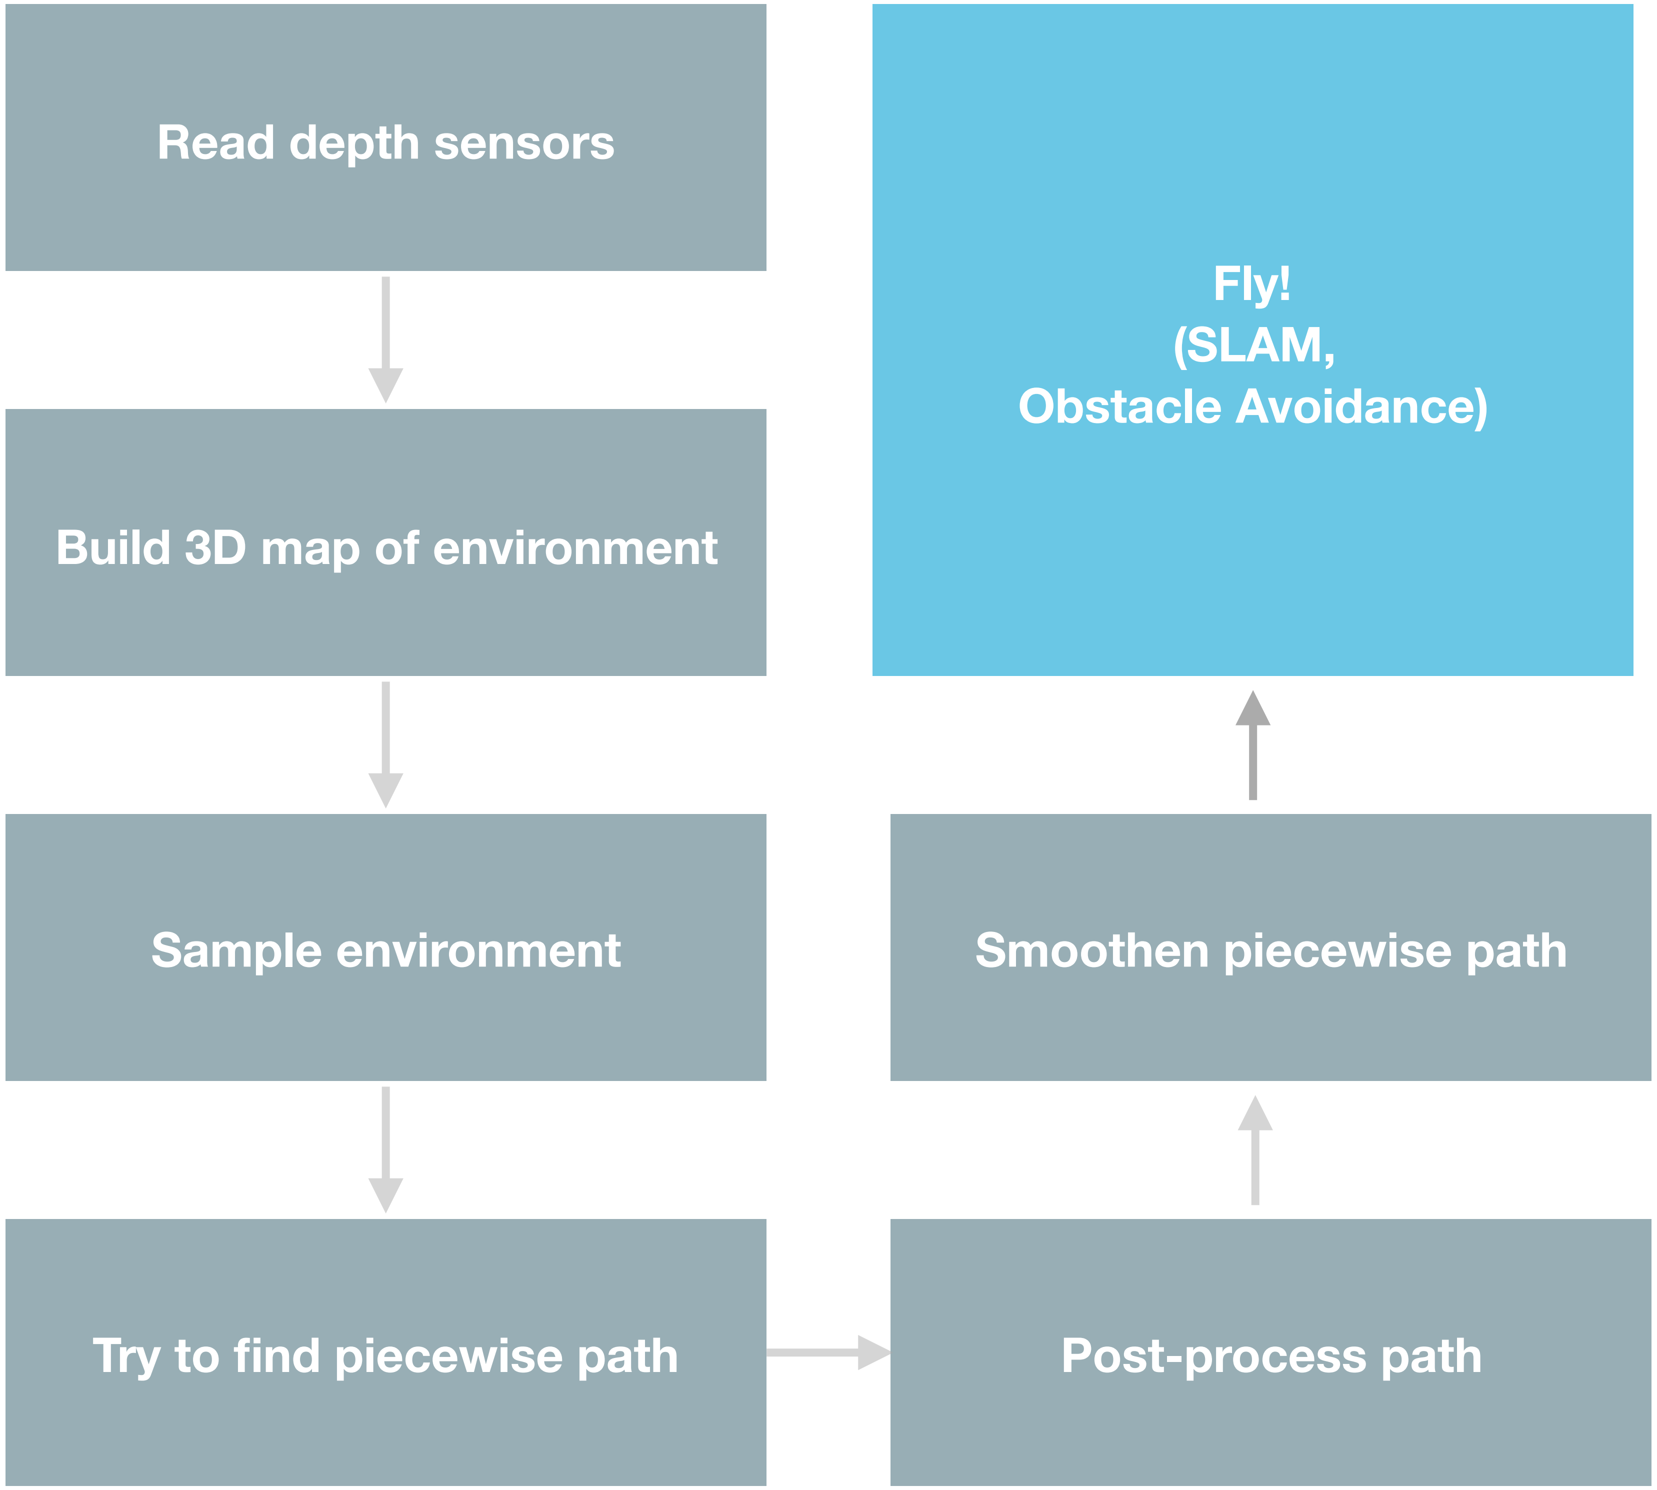

Read depth sensors

Build 3D map of environment

Sample environment

Try to find piecewise path

Post-process path

Smoothen piecewise path

Fly!
(SLAM,
Obstacle Avoidance)

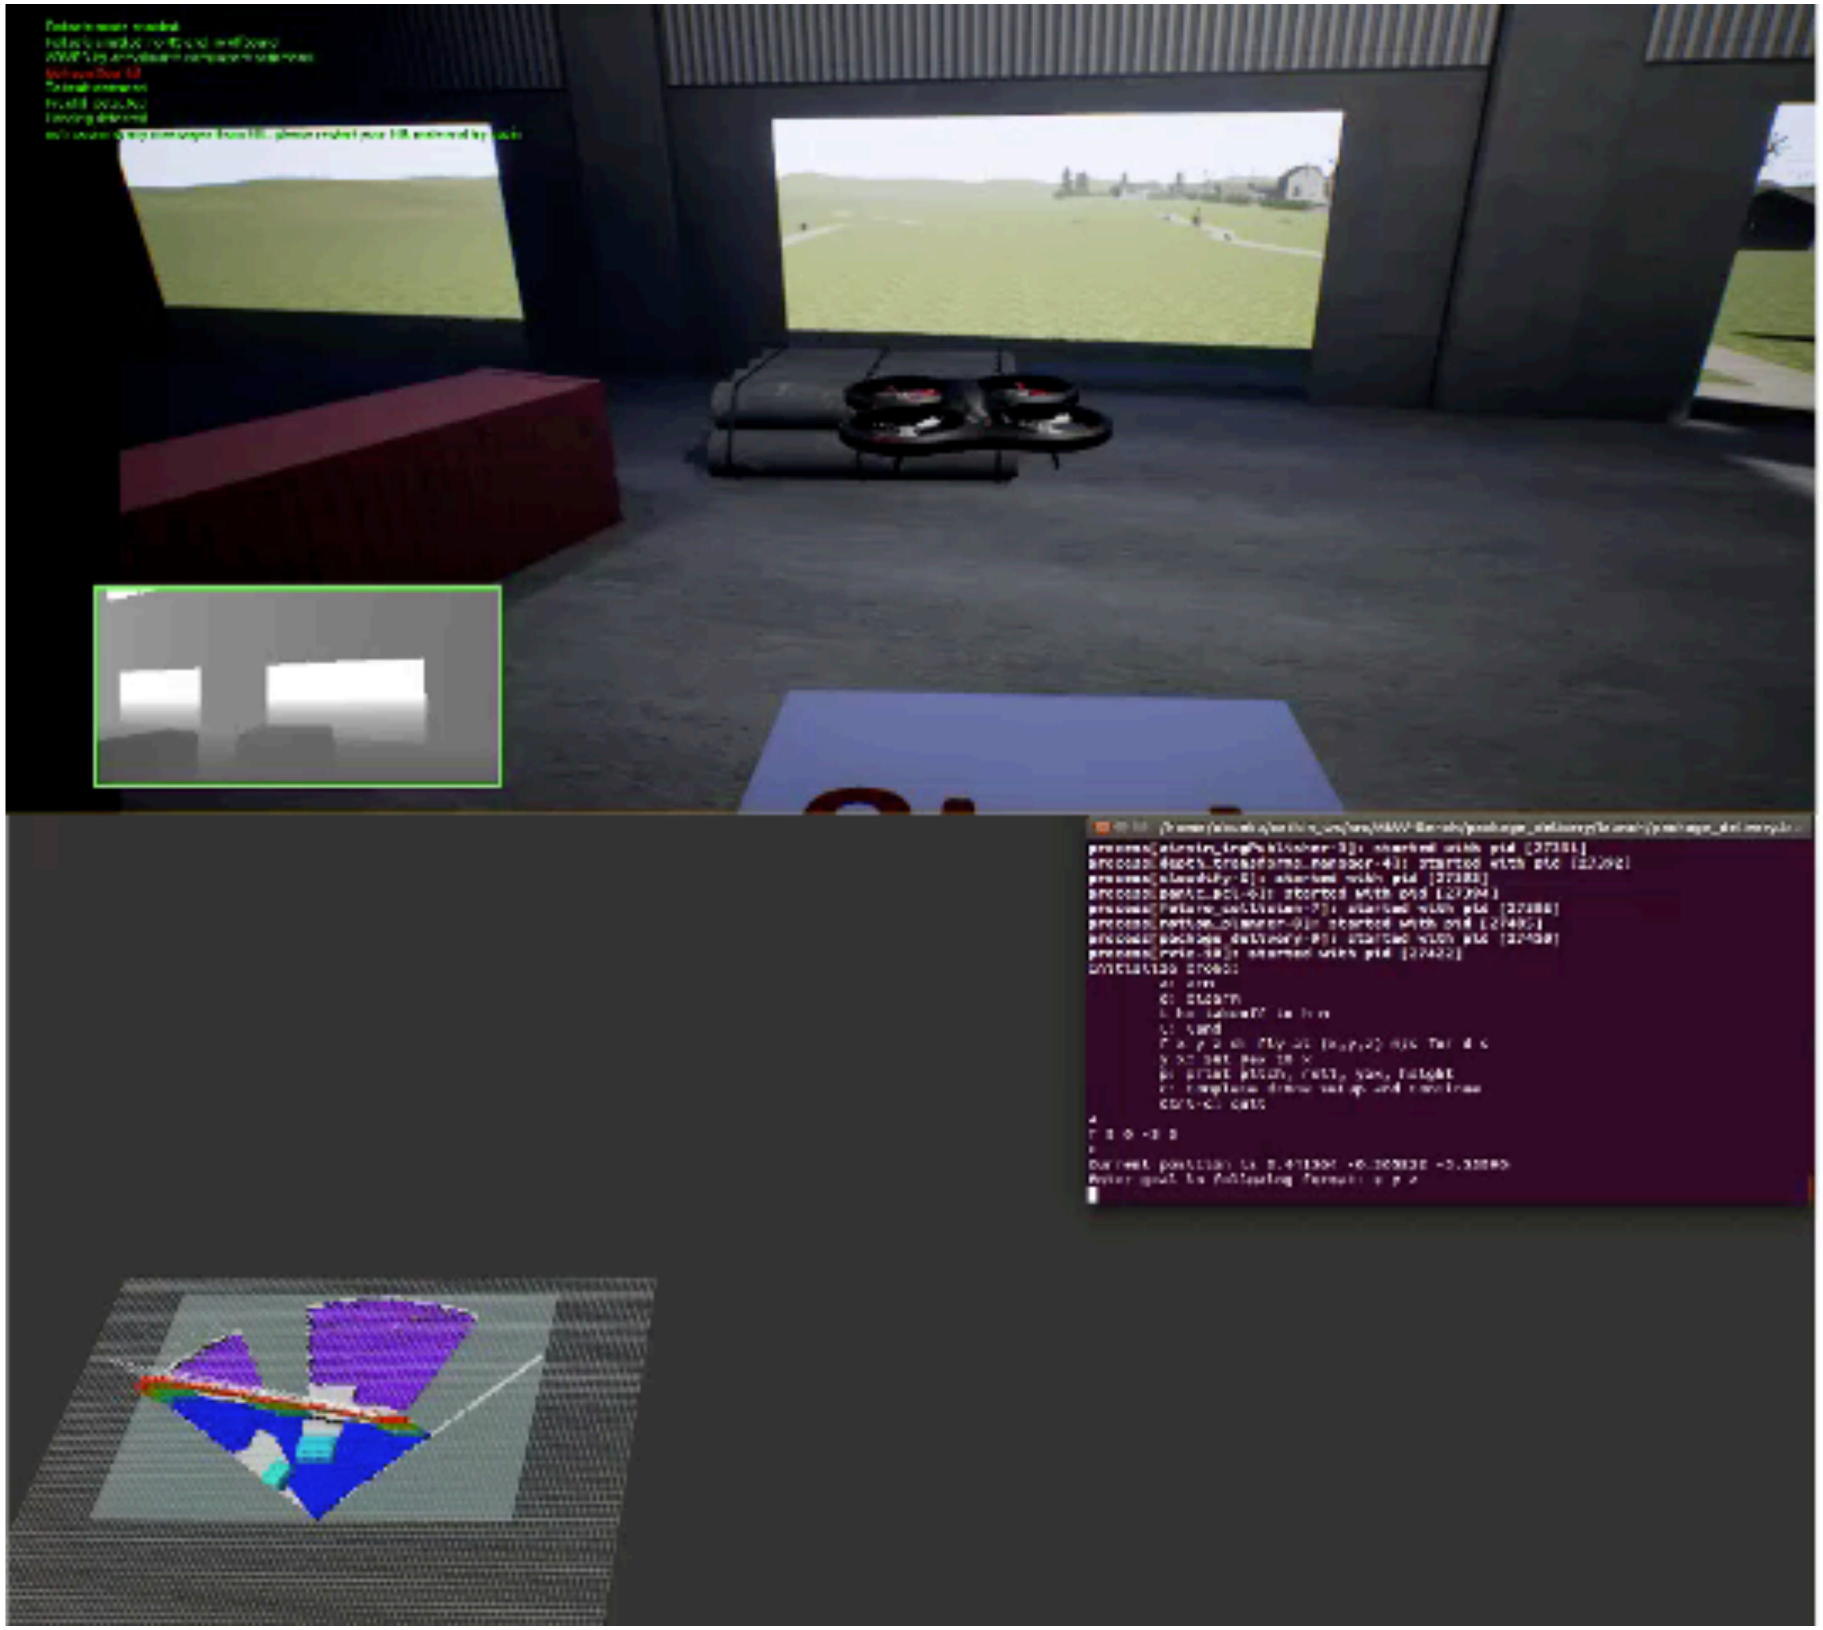

```
process[airsim_imgPublisher-3]: started with pid [27381]
process[depth_transforms_manager-4]: started with pid [27392]
process[cloudify-5]: started with pid [27383]
process[point_pcl-6]: started with pid [27394]
process[future_collision-7]: started with pid [27388]
process[motion_planner-8]: started with pid [27405]
process[package_delivery-9]: started with pid [27430]
process[rviz-10]: started with pid [27422]
Initialize drone:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly to (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        c: complete drone setup and continue
        ctrl-c: quit
a
f 0 0 -3 3
c
Current position is 0.41304 -0.30838 -3.33805
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

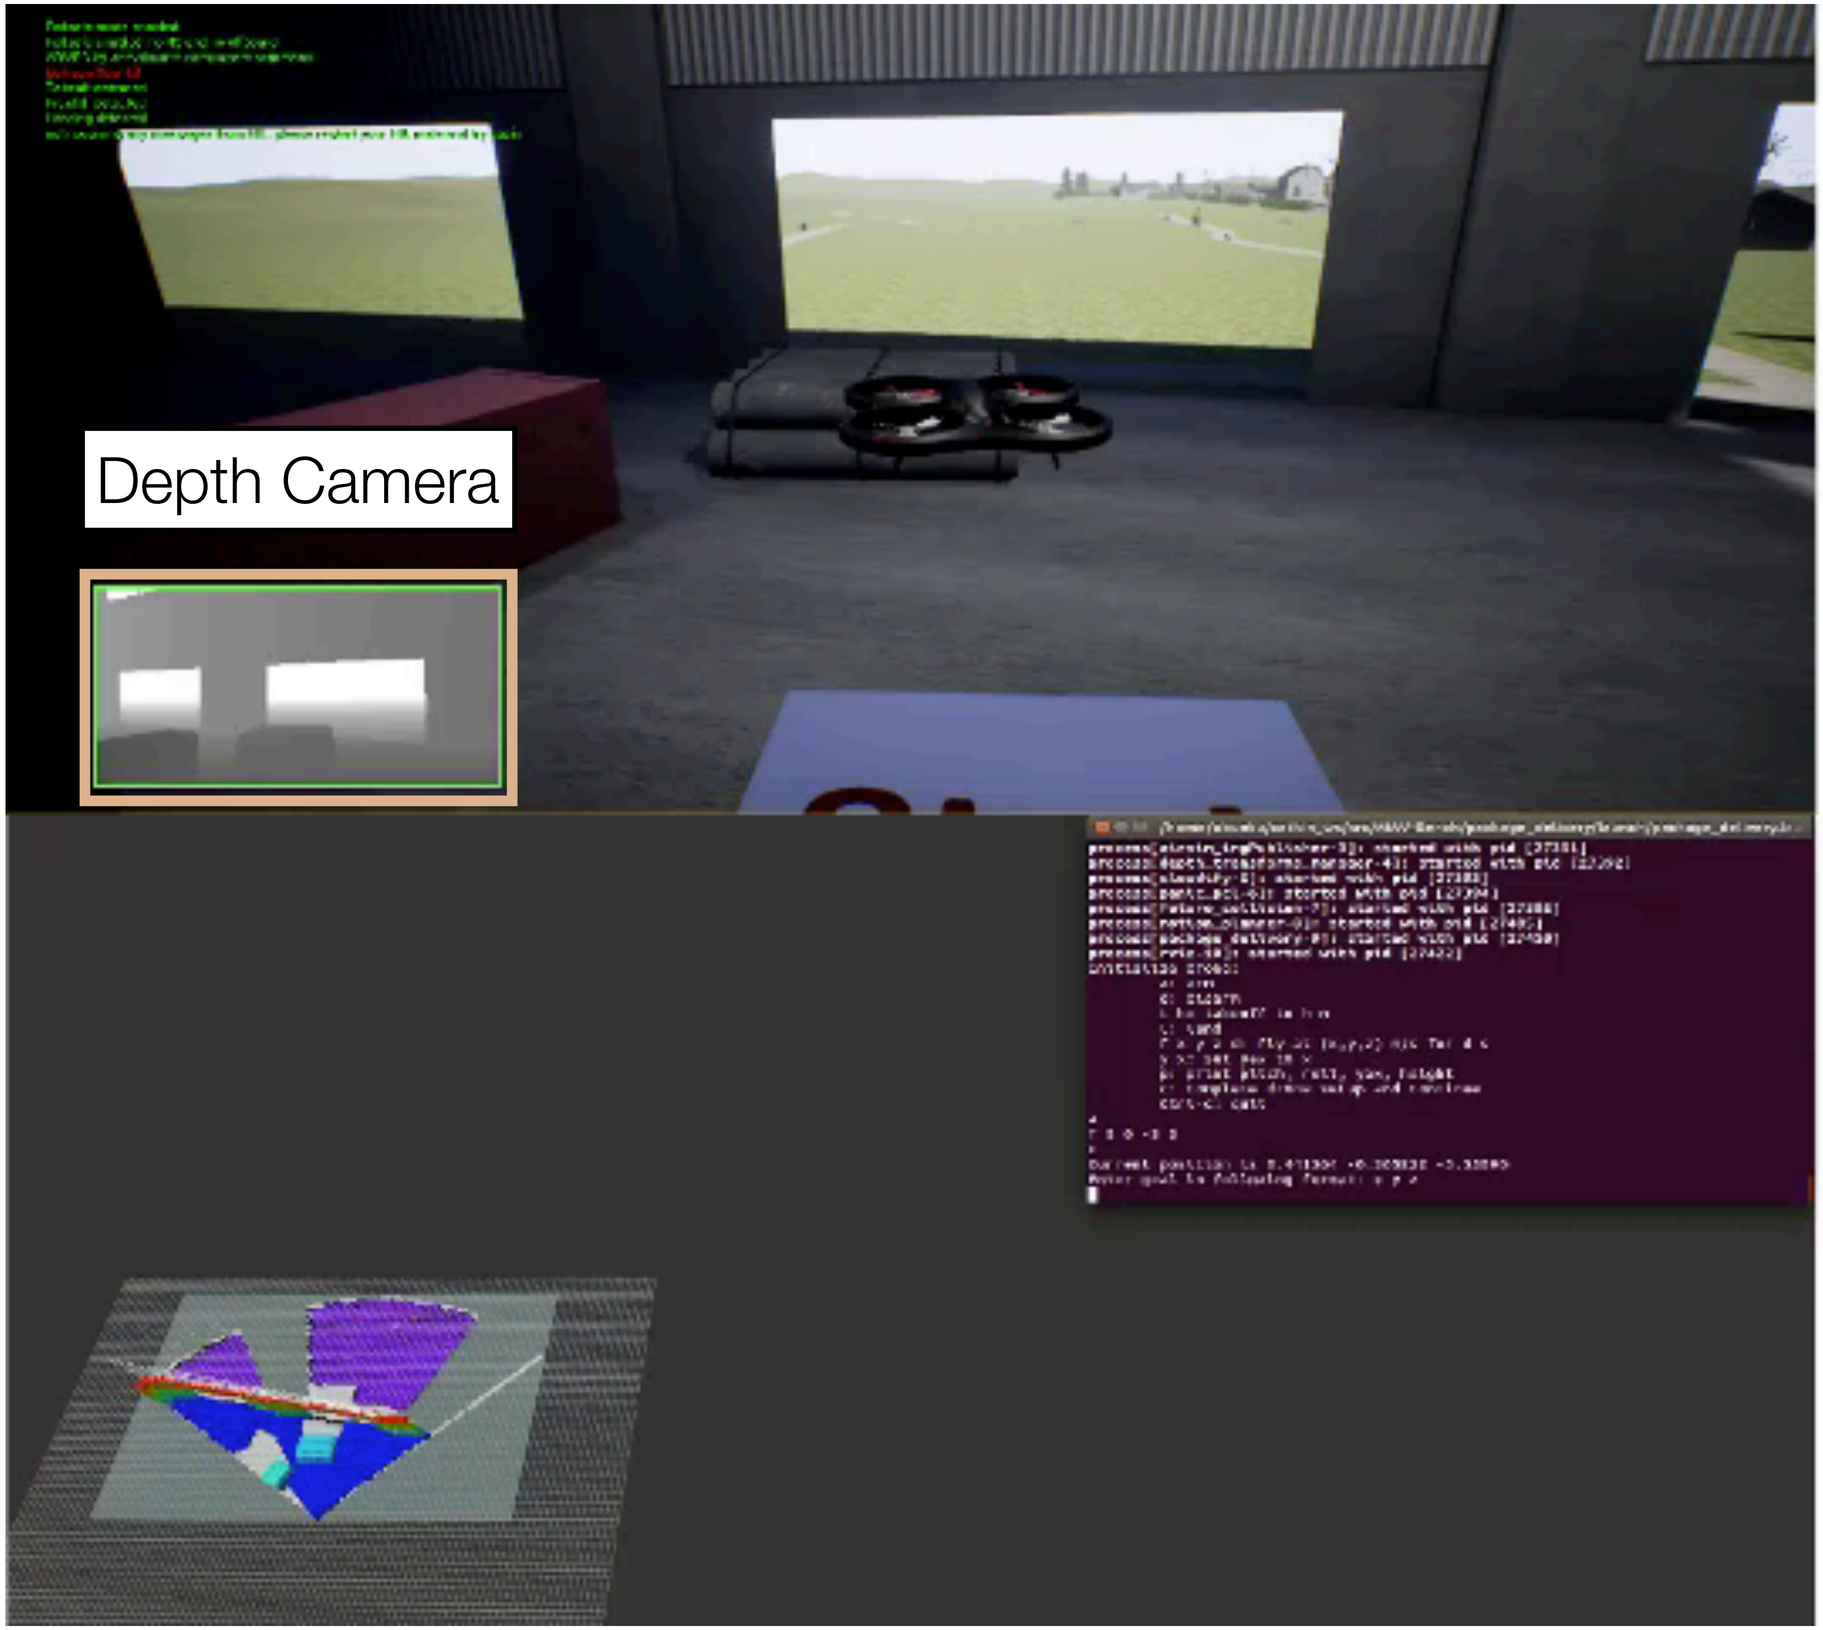

```
/home/.../catkin_ws/src/MAV-Bench/package_delivery/launch/package_delivery.launch
process[airsim_imgPublisher-3]: started with pid [27381]
process[depth_transforms_manager-4]: started with pid [27392]
process[cloudify-5]: started with pid [27383]
process[panic_pcl-6]: started with pid [27394]
process[future_collision-7]: started with pid [27388]
process[motion_planner-8]: started with pid [27405]
process[package_delivery-9]: started with pid [27430]
process[rviz-10]: started with pid [27422]
Initiation steps:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly at (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        c: complete drone set up and continue
        Ctrl-c: quit
a
f 0 0 -3 3
c
Current position is 0.41300 -0.300030 -3.33000
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

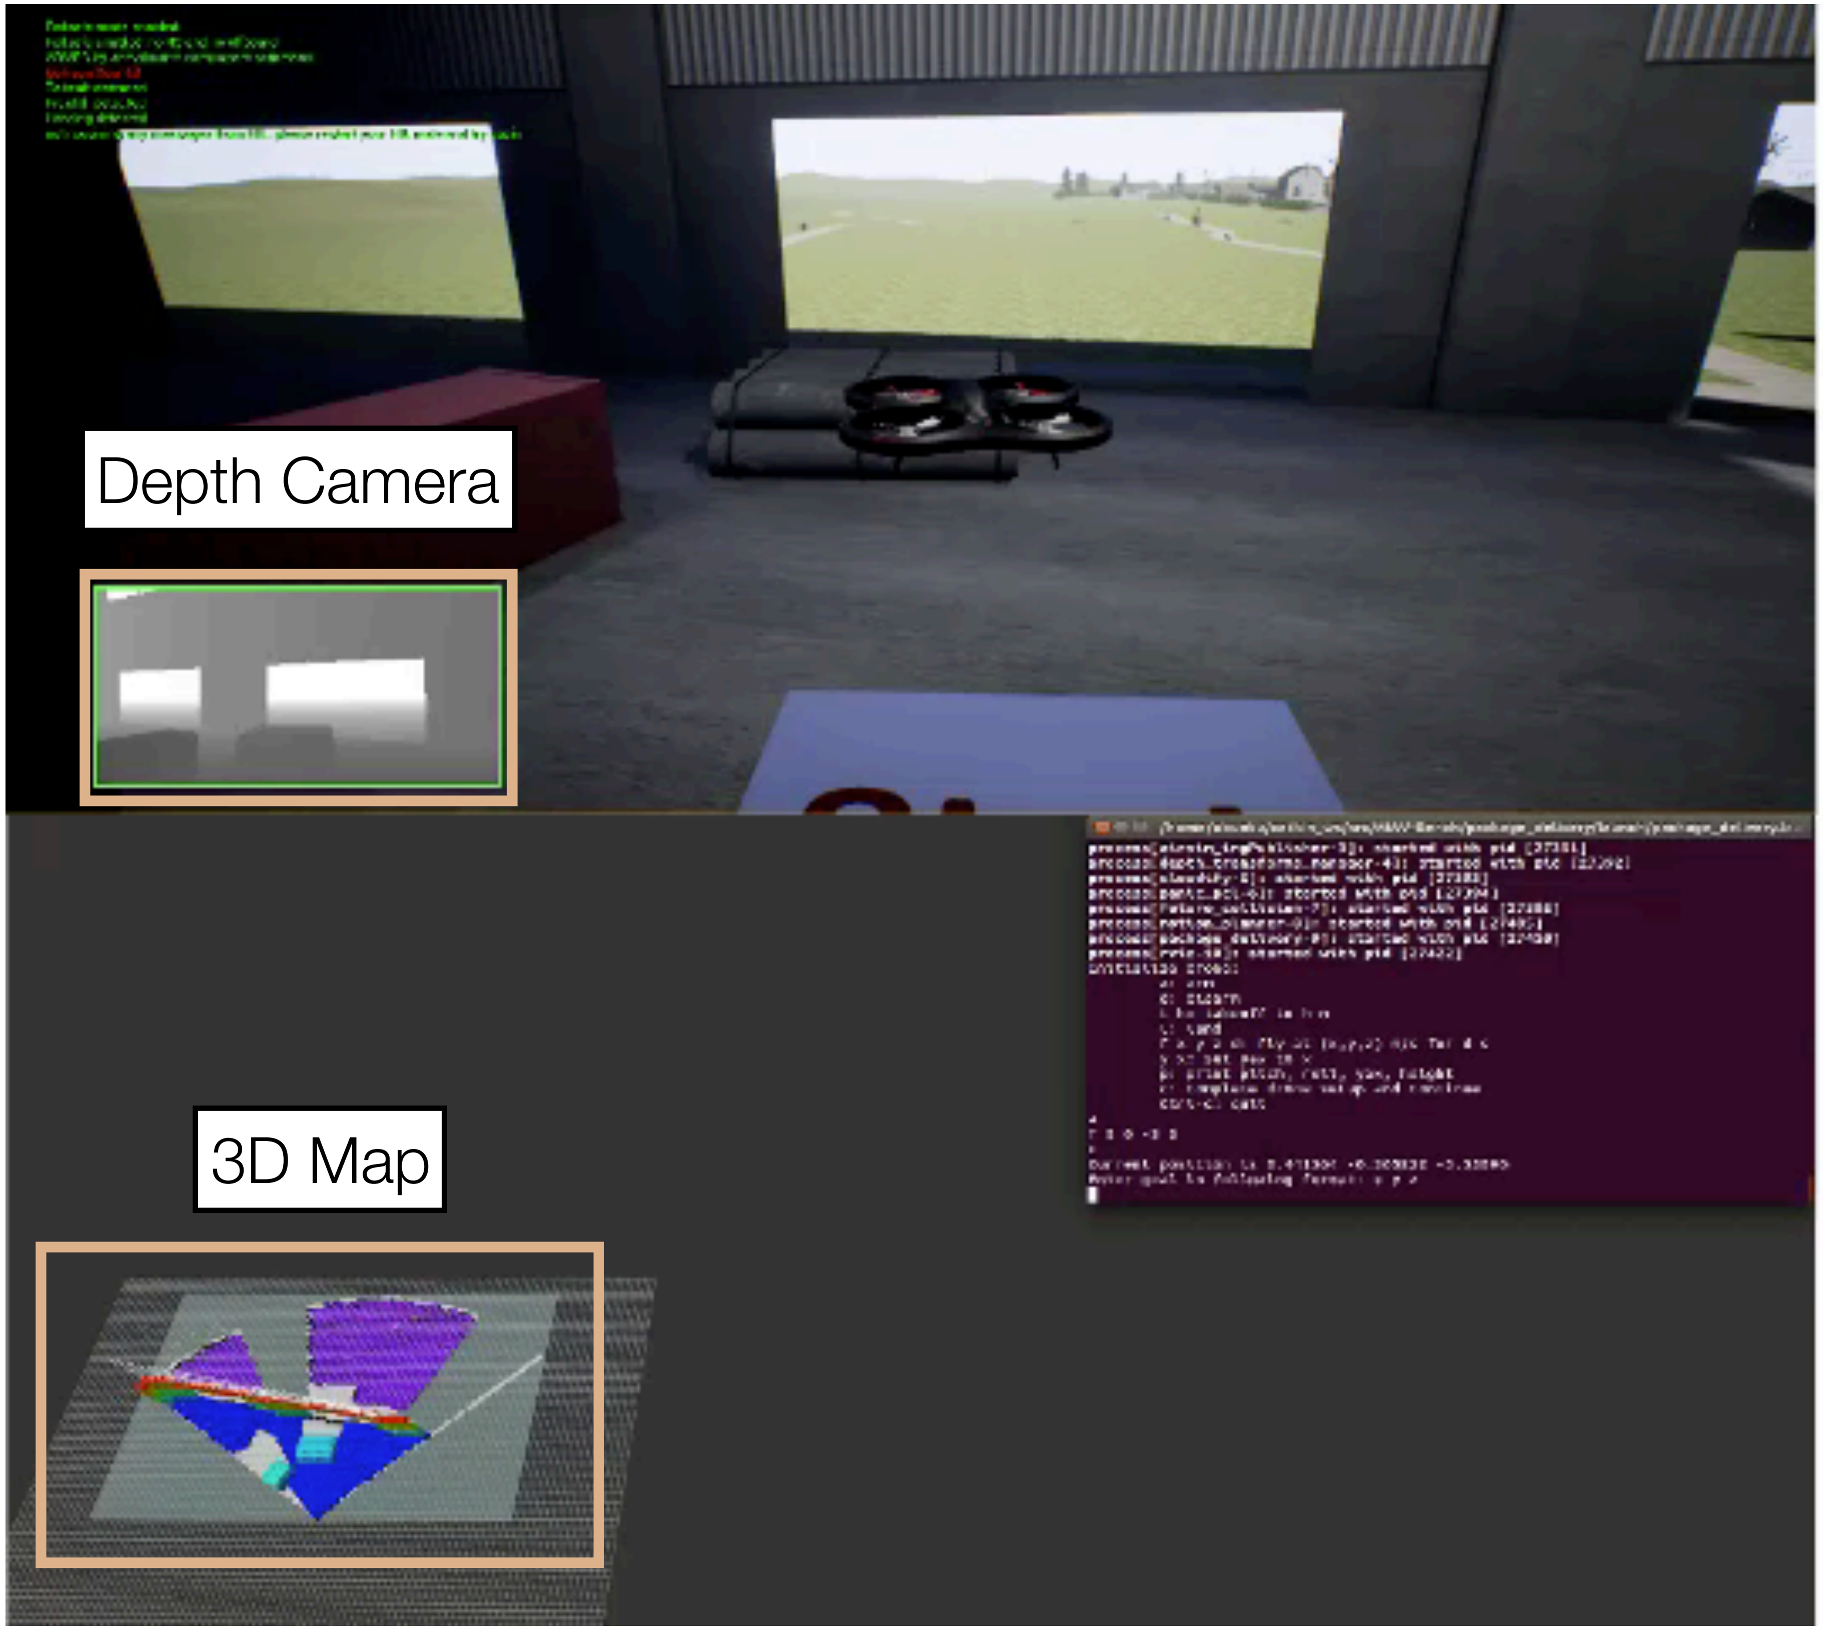

```
process[airsim_imgPublisher-3]: started with pid [27381]
process[depth_transforms_manager-4]: started with pid [27382]
process[cloudify-5]: started with pid [27383]
process[[illegible]-6]: started with pid [27384]
process[future_collision-7]: started with pid [27386]
process[motion_planner-8]: started with pid [27405]
process[package_delivery-9]: started with pid [27410]
process[rviz-10]: started with pid [27422]
Available cmds:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly at (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        r: [illegible] drone set up and continue
        cntrl-c: quit
a
f 0 0 -3 3
t
Current position is 0.41130X -0.369030 -5.33899
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

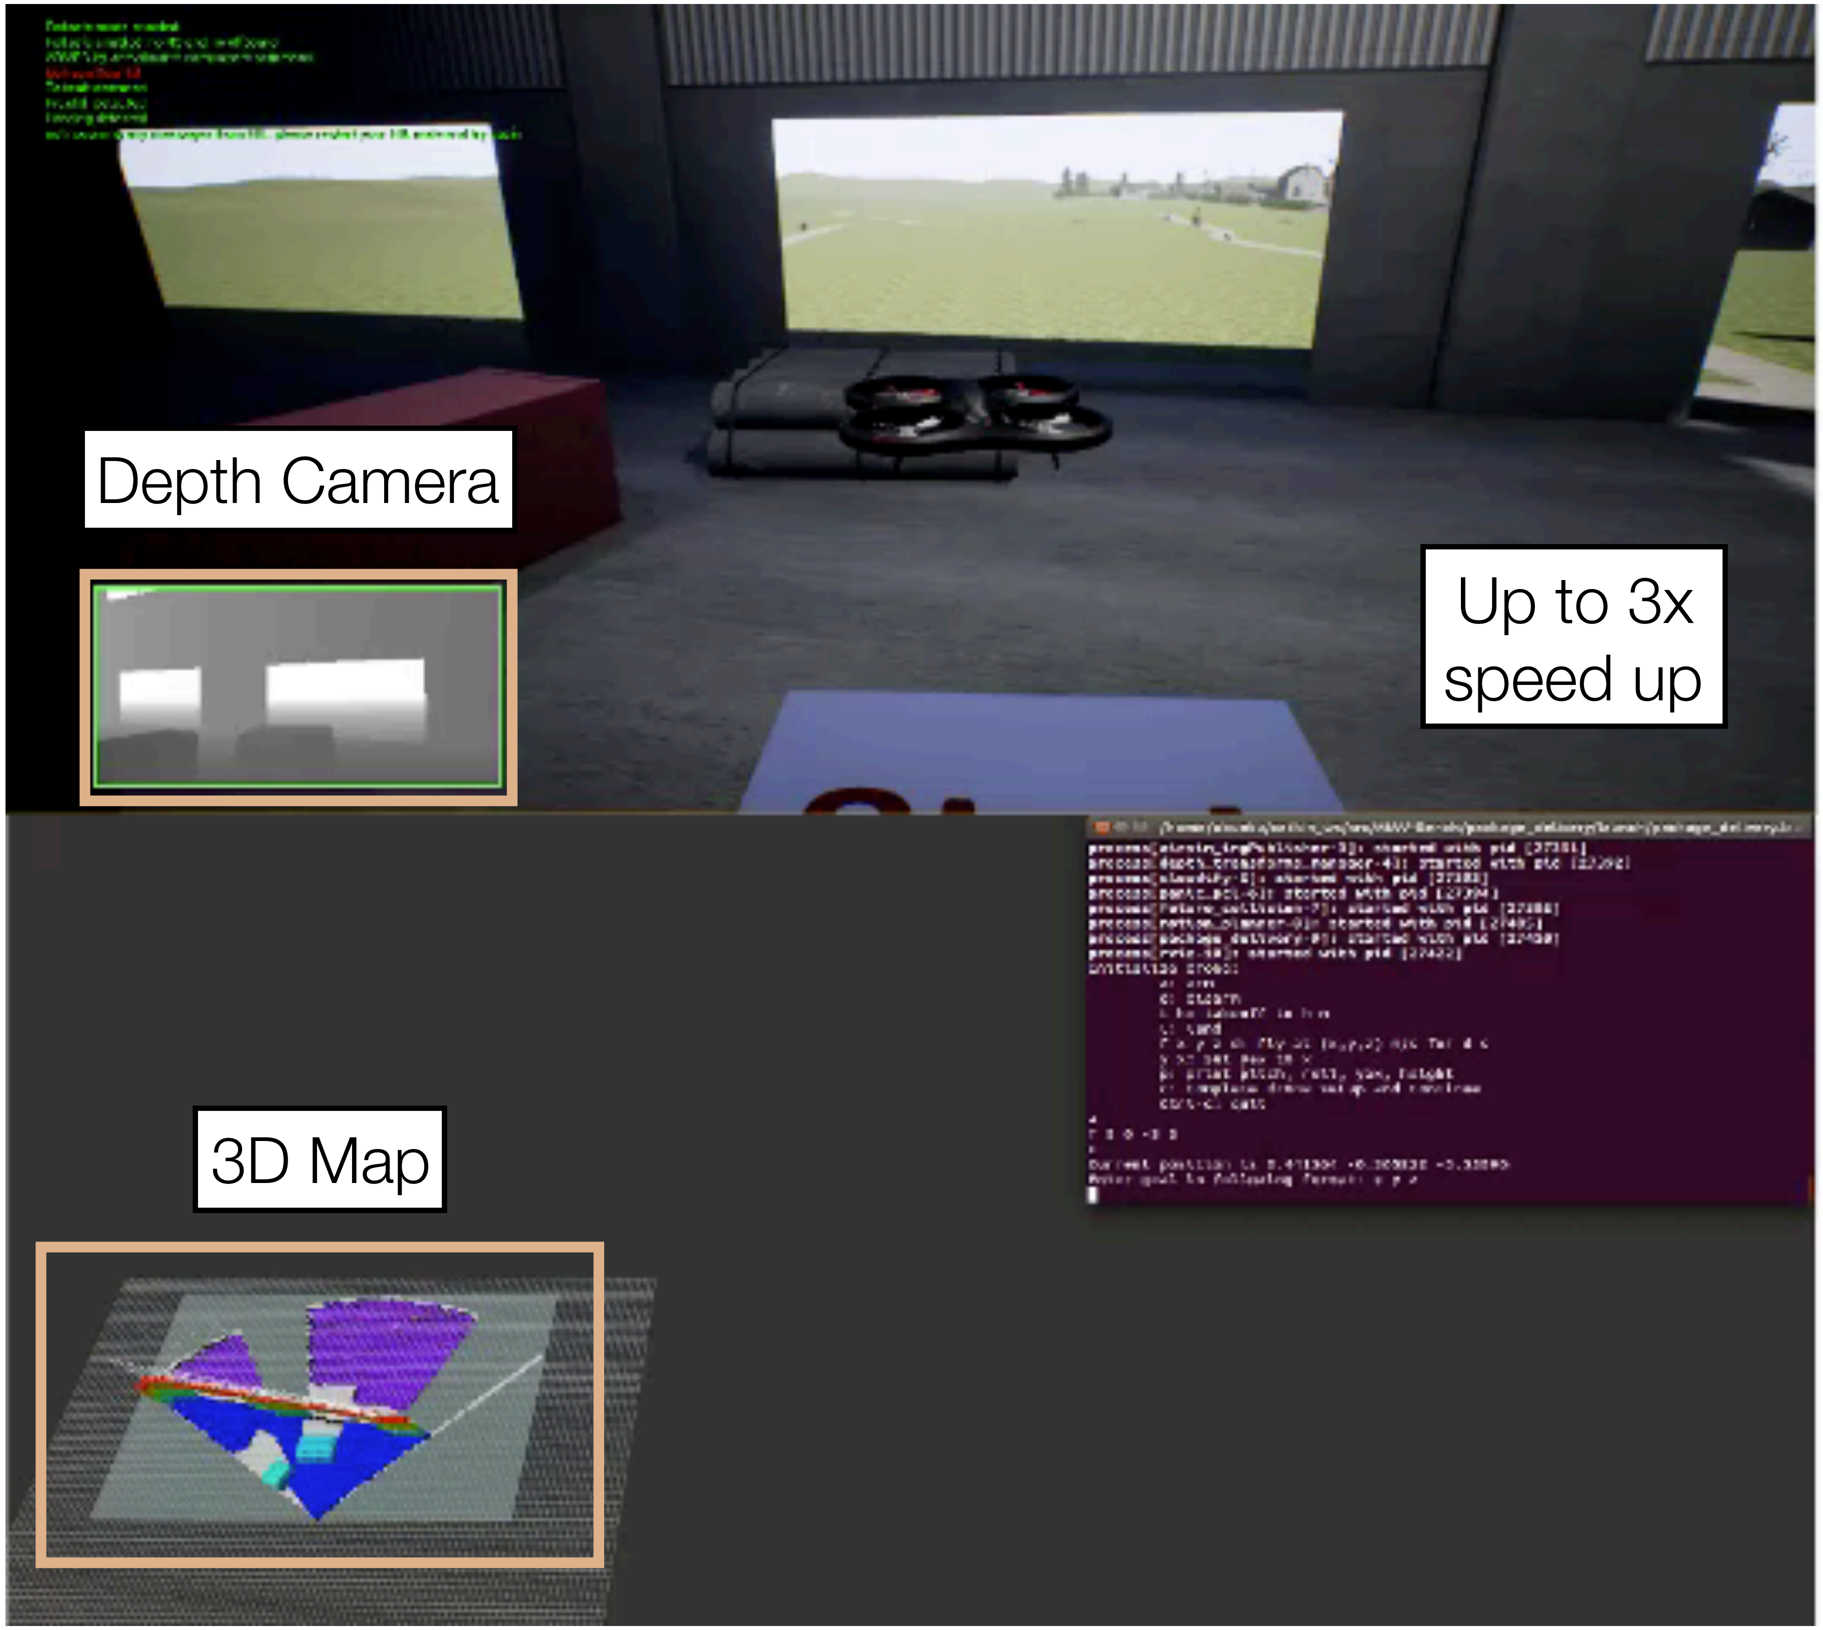

Depth Camera

Up to 3x speed up

3D Map

```
process[airsim_imgPublisher-3]: started with pid [[illegible]]
process[depth_transforms_manager-4]: started with pid [[illegible]]
process[cloudify-5]: started with pid [[illegible]]
process[[illegible]-6]: started with pid [[illegible]]
process[future_collision-7]: started with pid [[illegible]]
process[motion_planner-8]: started with pid [[illegible]]
process[package_delivery-9]: started with pid [[illegible]]
process[rviz-10]: started with pid [[illegible]]
Initialize drone:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly at (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        c: complete drone setup and continue
        ctrl-c: quit
a
f 0 0 -3 3
c
Current position is [illegible] [illegible] [illegible]
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

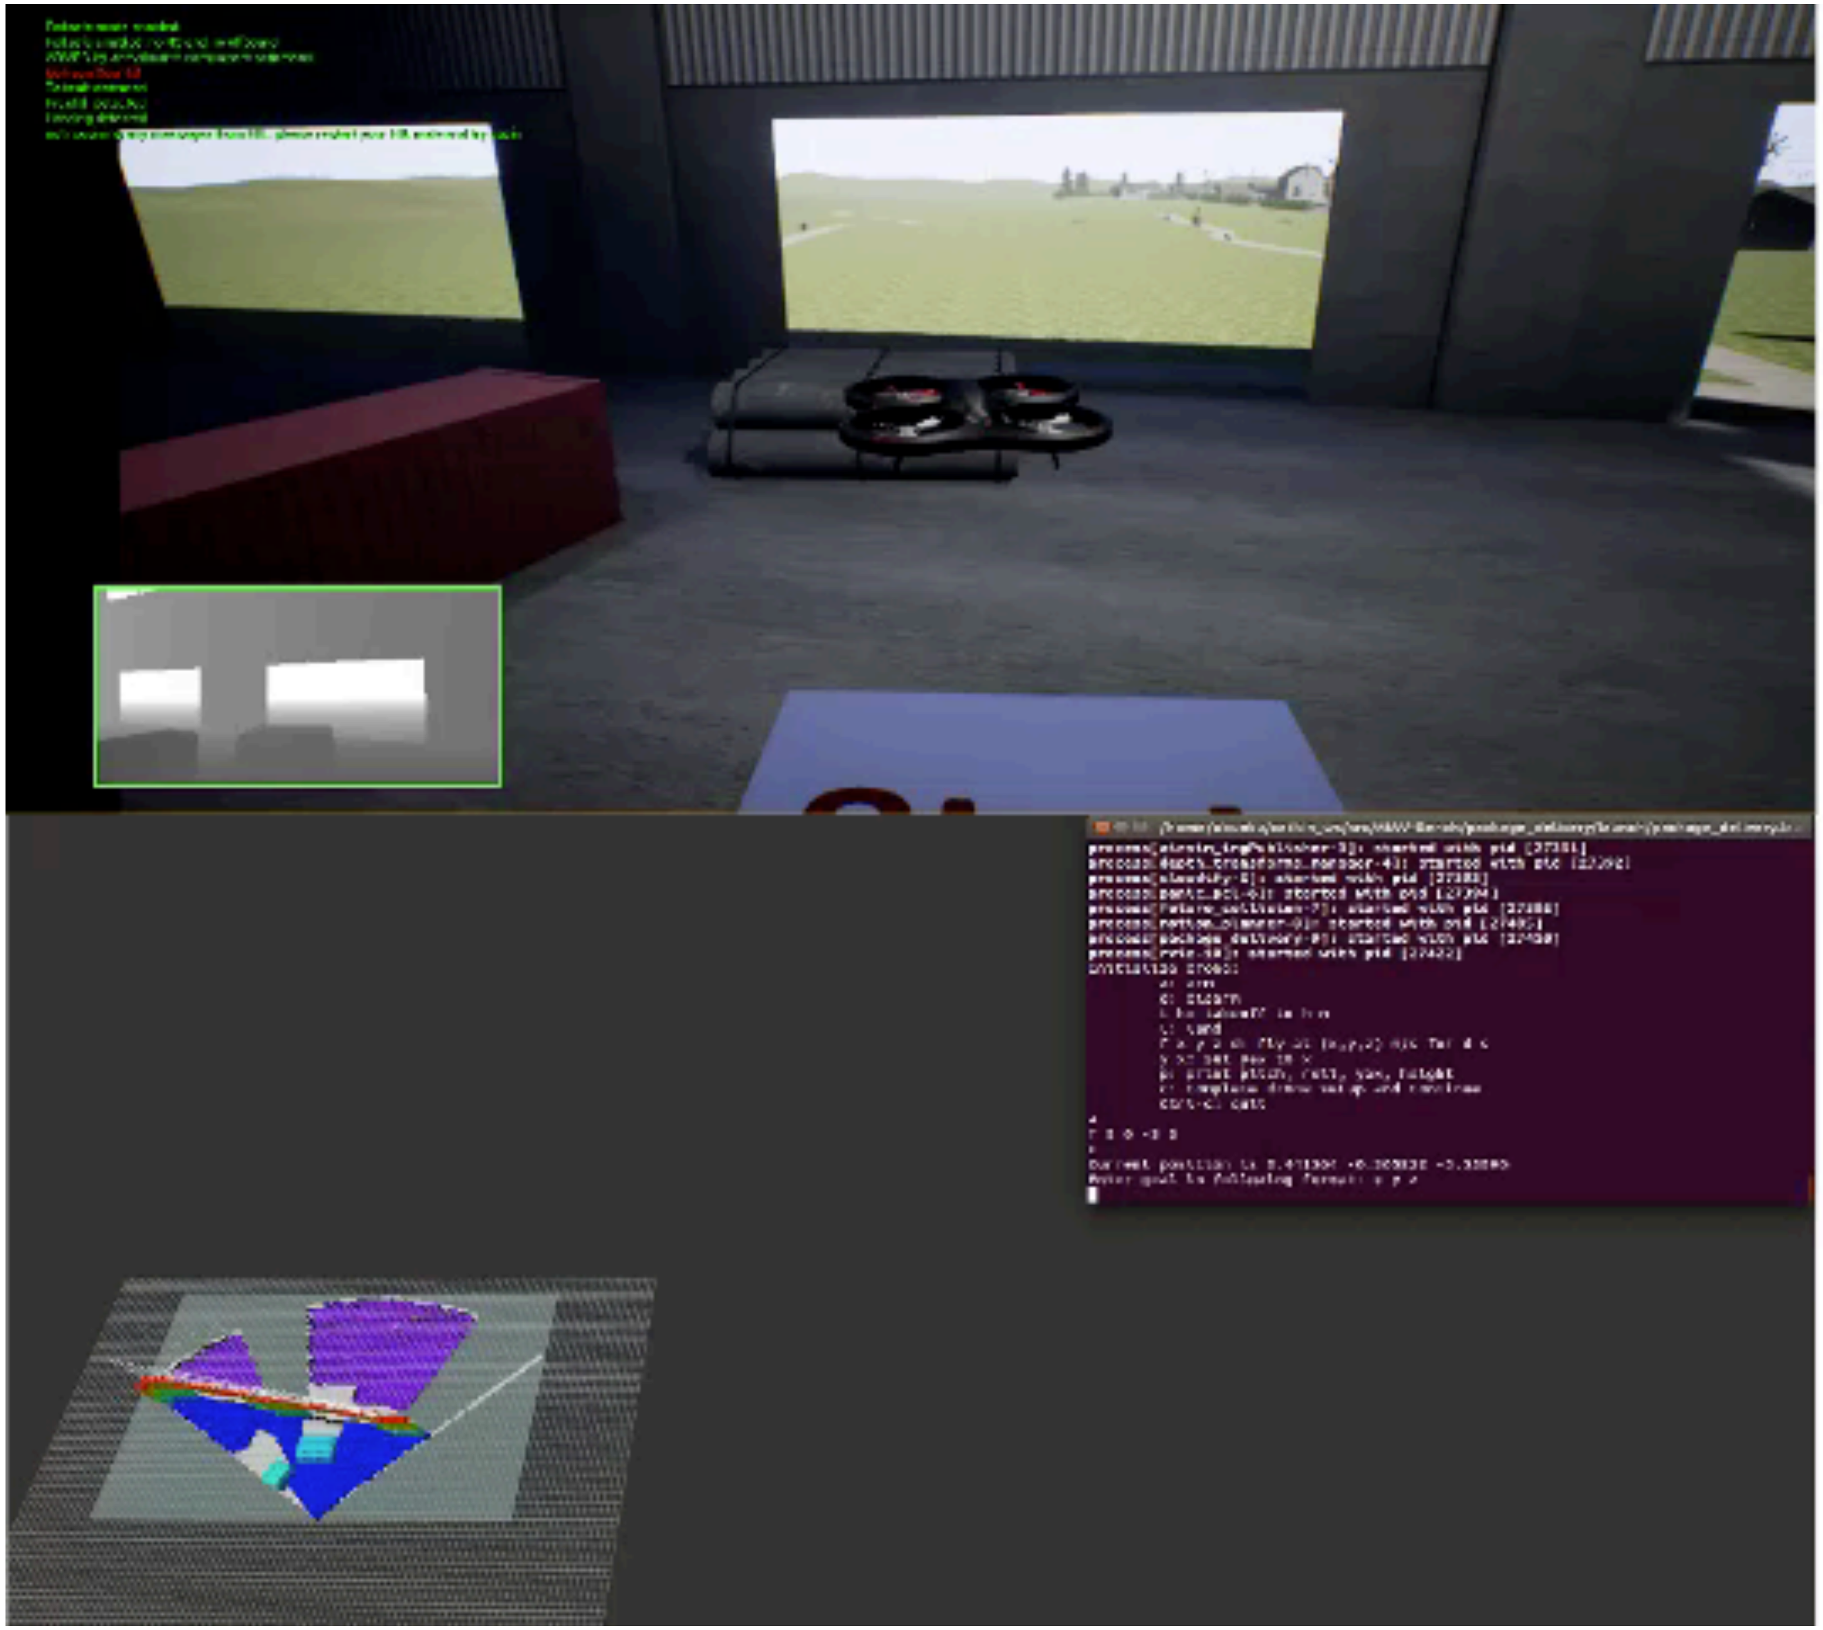

```
process[airsim_imgPublisher-3]: started with pid [27381]
process[depth_transforms_manager-4]: started with pid [27392]
process[[illegible]-5]: started with pid [27383]
process[pcl_[illegible]-6]: started with pid [27394]
process[future_collision-7]: started with pid [27386]
process[motion_planner-8]: started with pid [27405]
process[package_delivery-9]: started with pid [27430]
process[rviz-10]: started with pid [27422]
Initialize drone:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly at (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        c: complete drone set up and continue
        ctrl-c: quit
c
f 0 0 -3 3
c
Current position is 0.41300 -0.300032 -3.33805
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

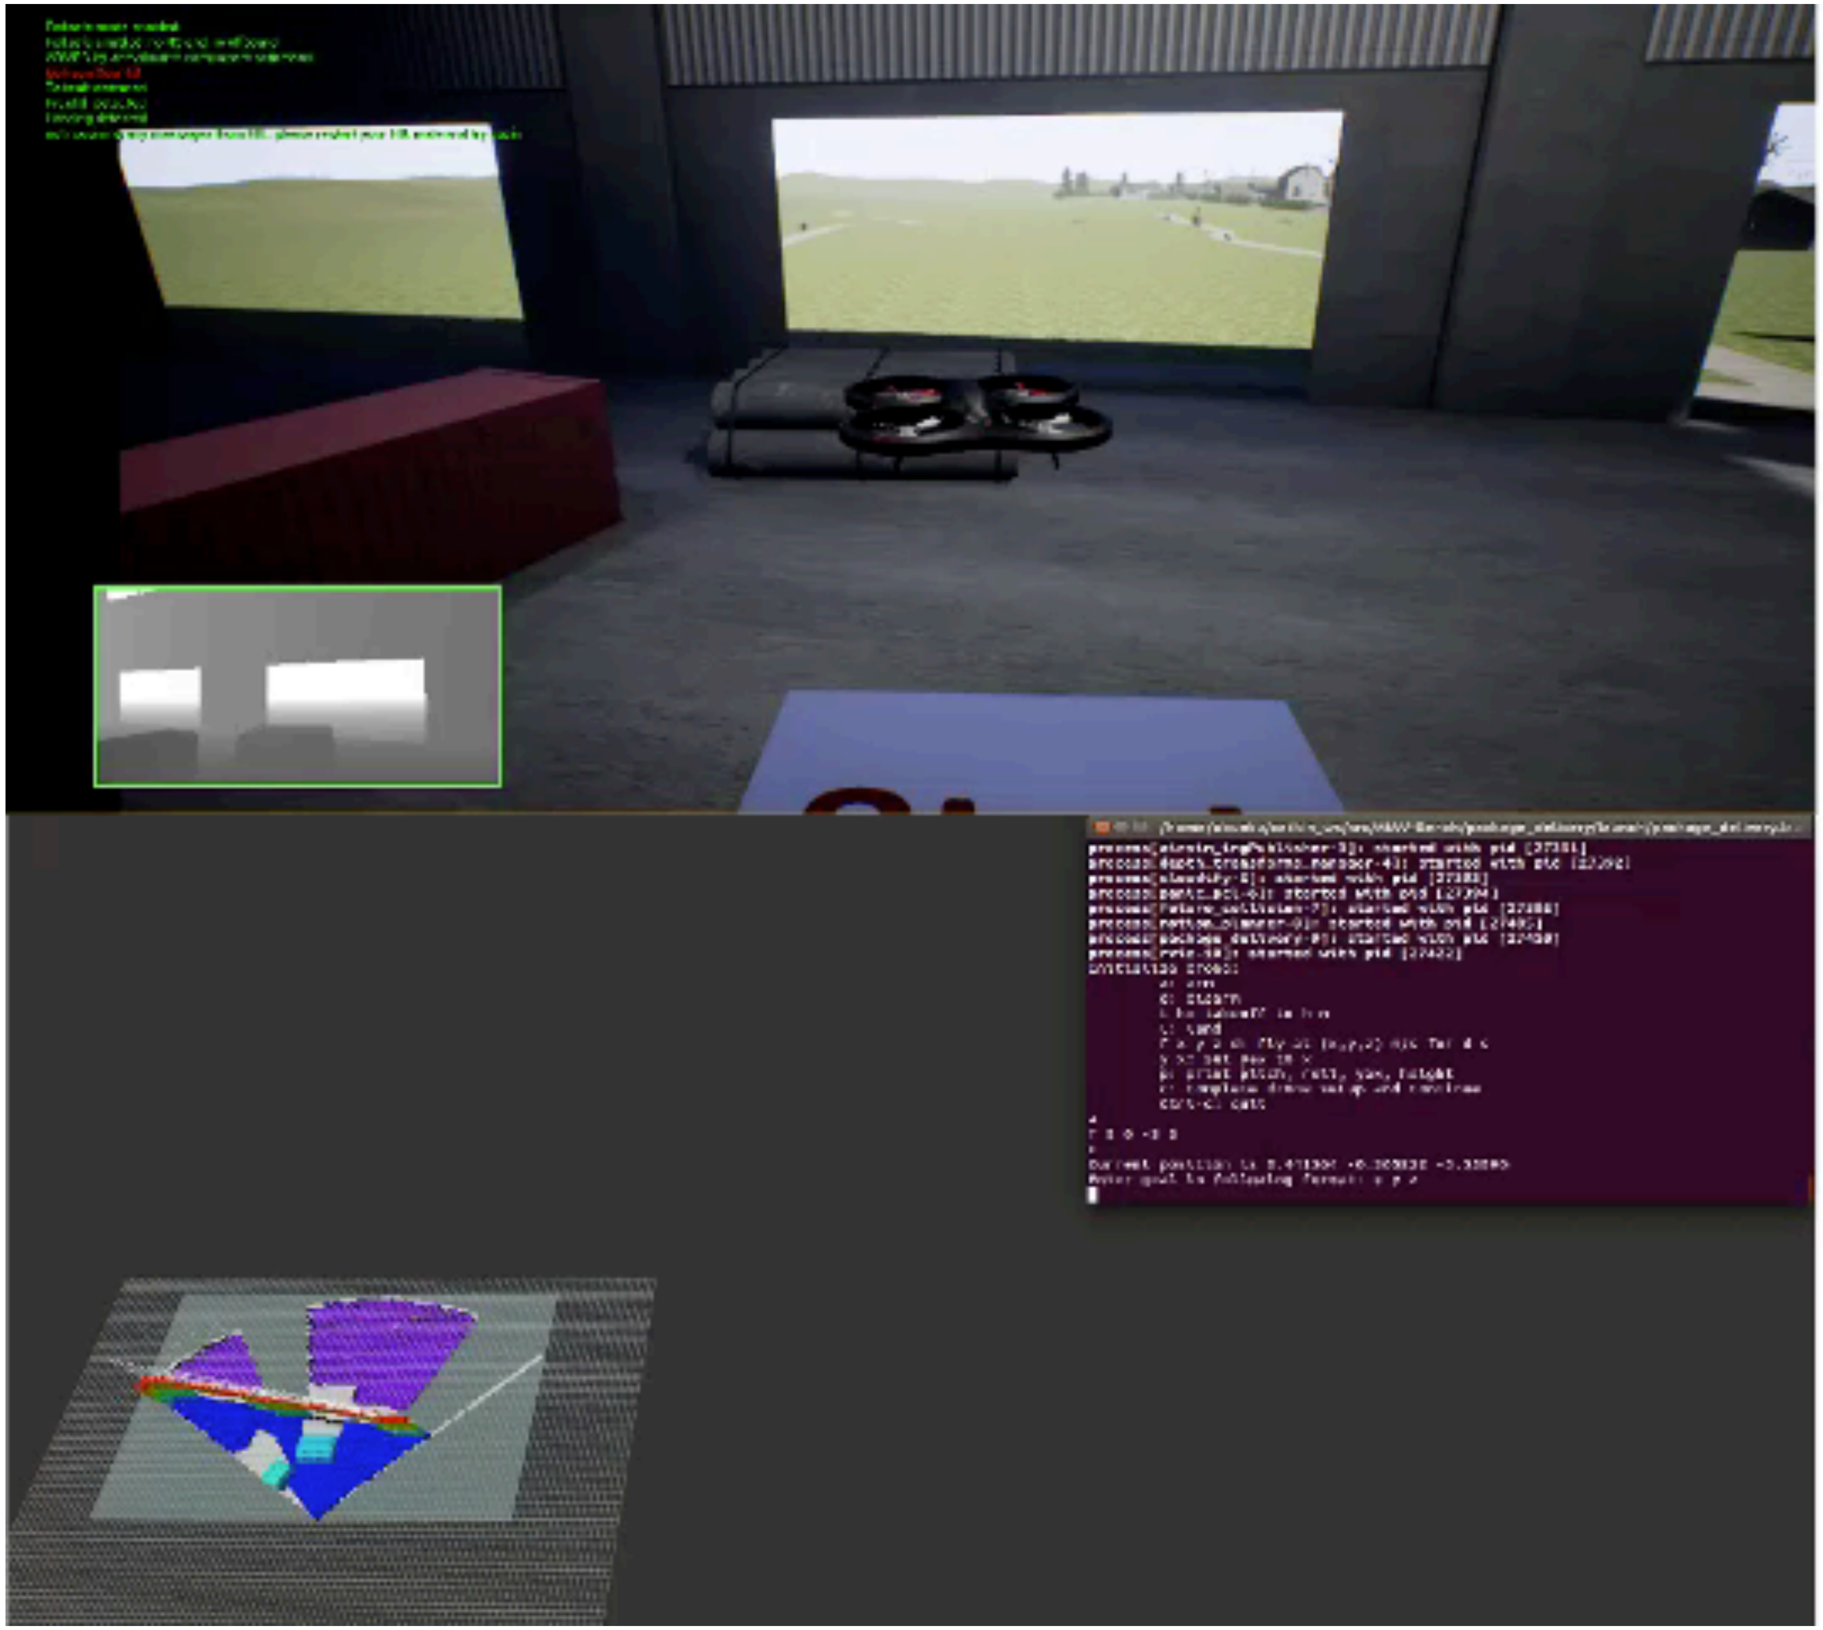

```
process[airsim_imgPublisher-3]: started with pid [27381]
process[depth_transforms_manager-4]: started with pid [27392]
process[cloudify-5]: started with pid [27383]
process[[illegible]-6]: started with pid [27394]
process[future_collision-7]: started with pid [27388]
process[motion_planner-8]: started with pid [27405]
process[package_delivery-9]: started with pid [27430]
process[rviz-10]: started with pid [27422]
Initialize drone:
        a: arm
        d: disarm
        t h: takeoff to h m
        l: land
        f x y z d: fly at (x,y,z) m/s for d s
        y x: set yaw to x
        p: print pitch, roll, yaw, height
        c: complete drone set up and continue
        ctrl-c: quit
a
f [illegible]
c
Current position is [illegible]
Enter goal in following format: x y z
```

What the World Sees

What the Drone Sees

Indoor Navigation

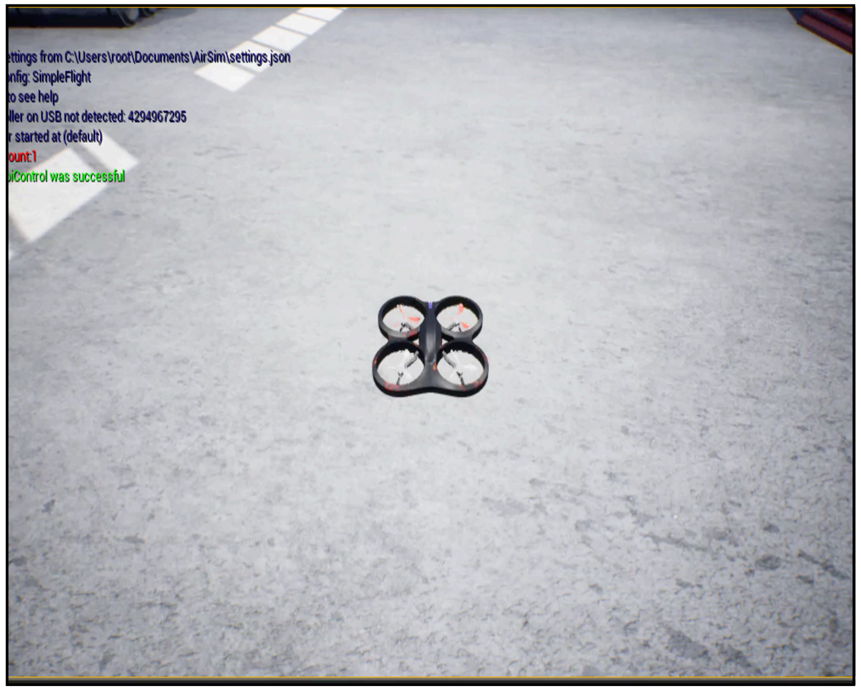

```
ettings from C:\Users\root\Documents\AirSim\settings.json
onfig: SimpleFlight
to see help
ller on USB not detected: 4294967295
r started at (default)
ount:1
iControl was successful
```

Package Delivery

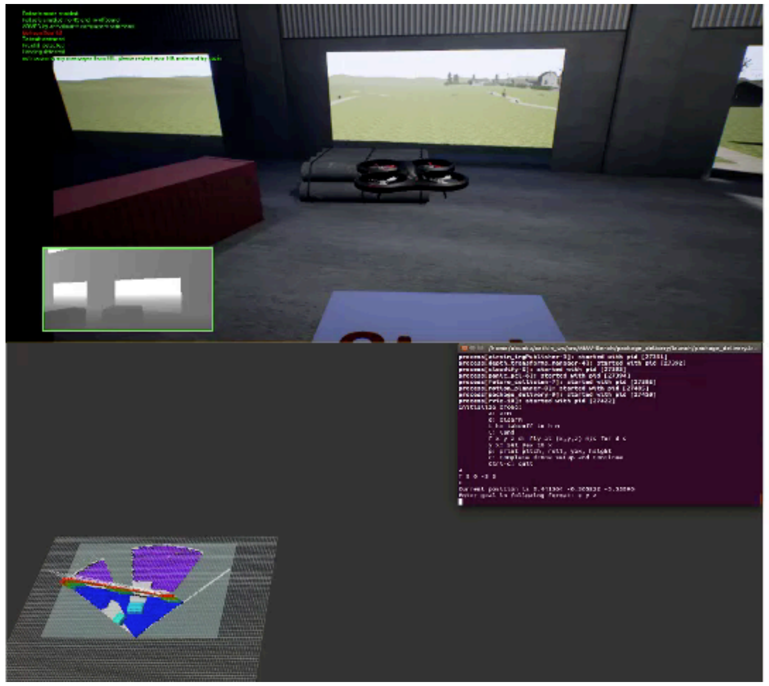

Surveying

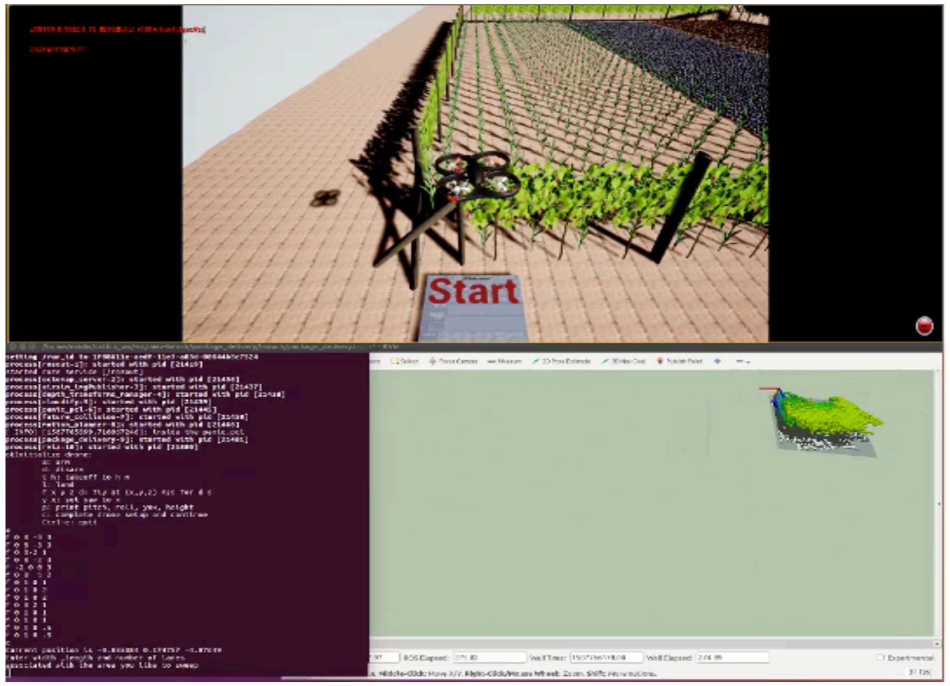

3D Mapping

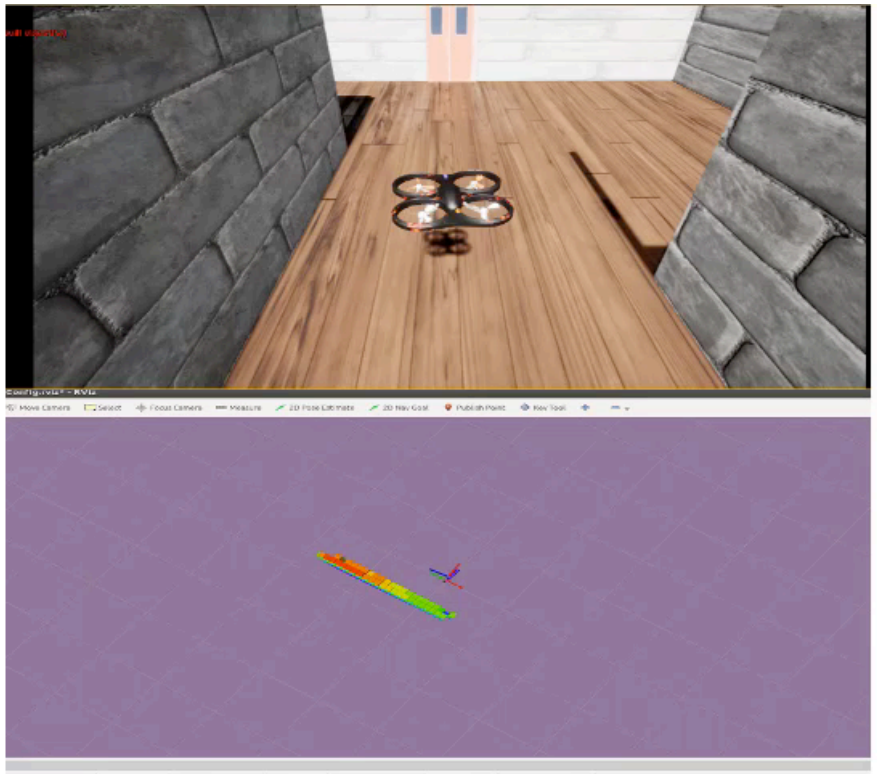

Search and Rescue

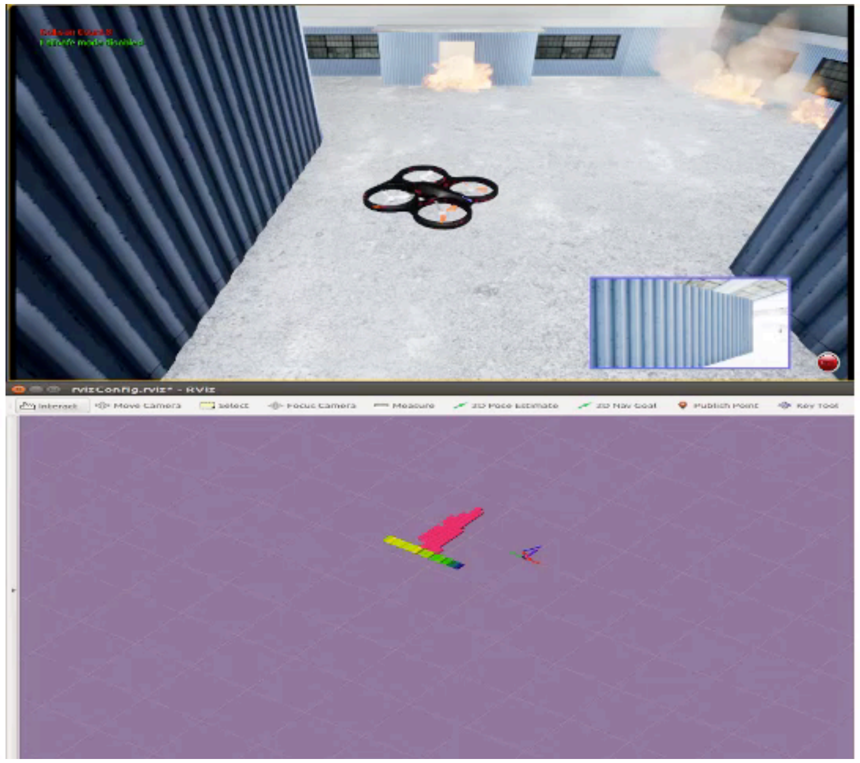

Aerial Photography

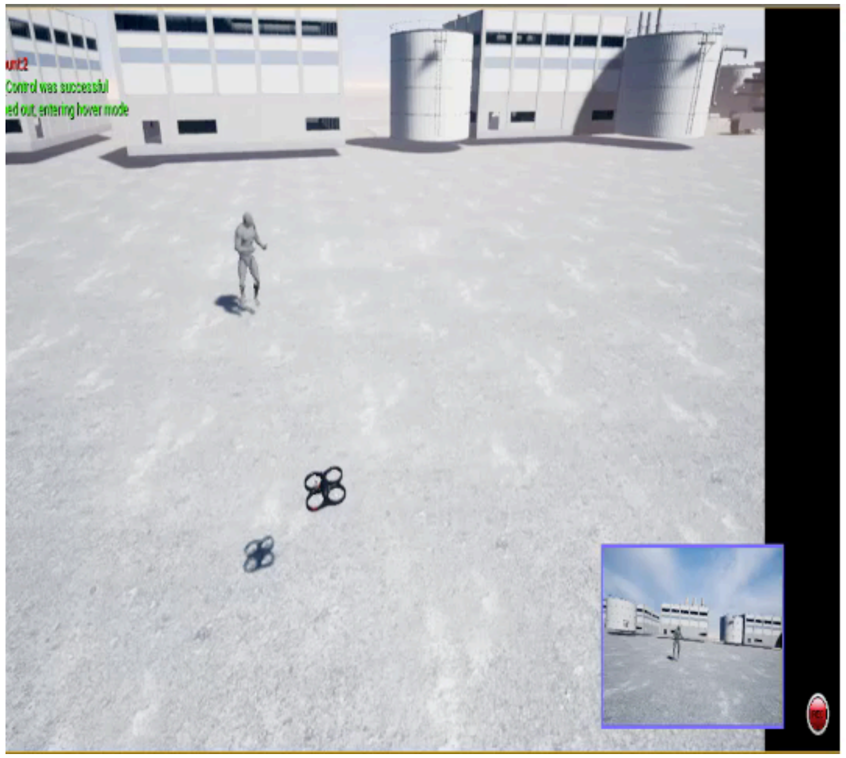

```
unt:2
Control was successful
ed out, entering hover mode
```

Indoor Navigation

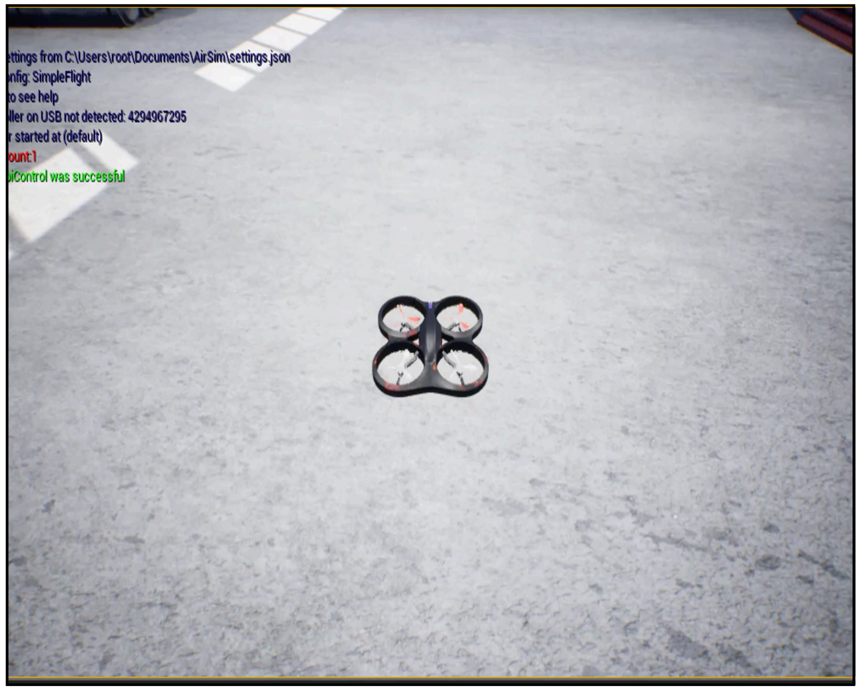

```
ttings from C:\Users\root\Documents\AirSim\settings.json
nfig: SimpleFlight
to see help
ller on USB not detected: 4294967295
r started at (default)
ount:1
lControl was successful
```

Package Delivery

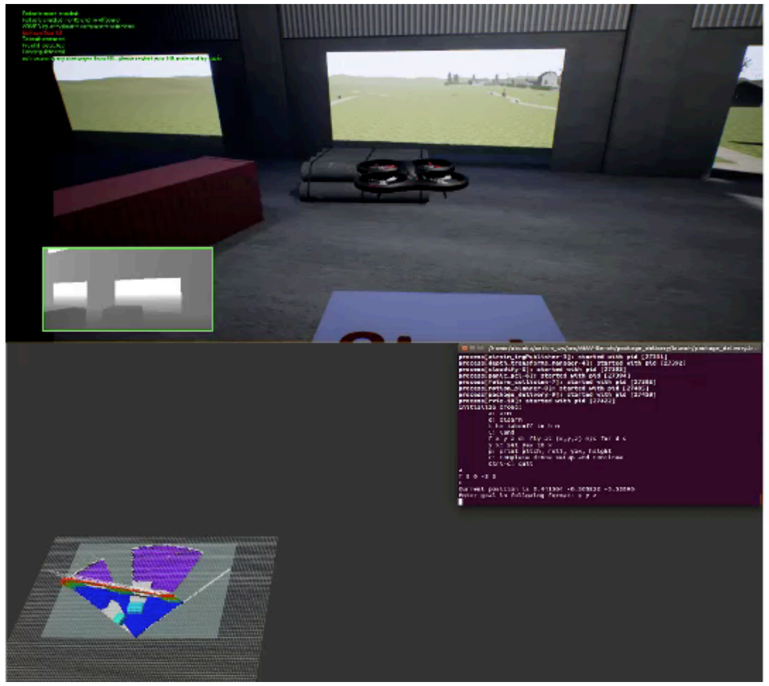

Surveying

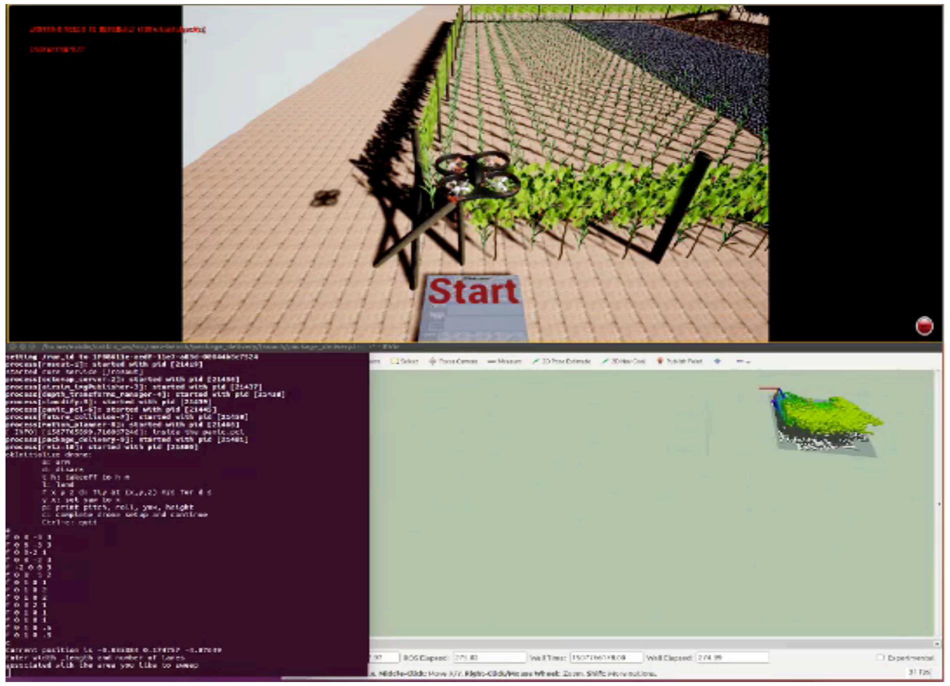

3D Mapping

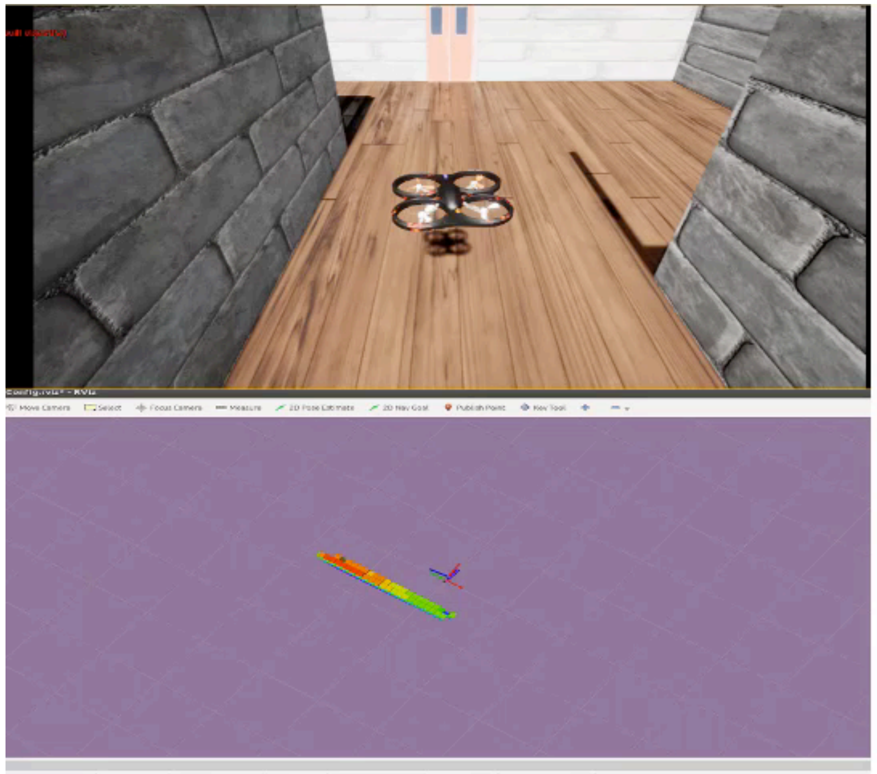

Search and Rescue

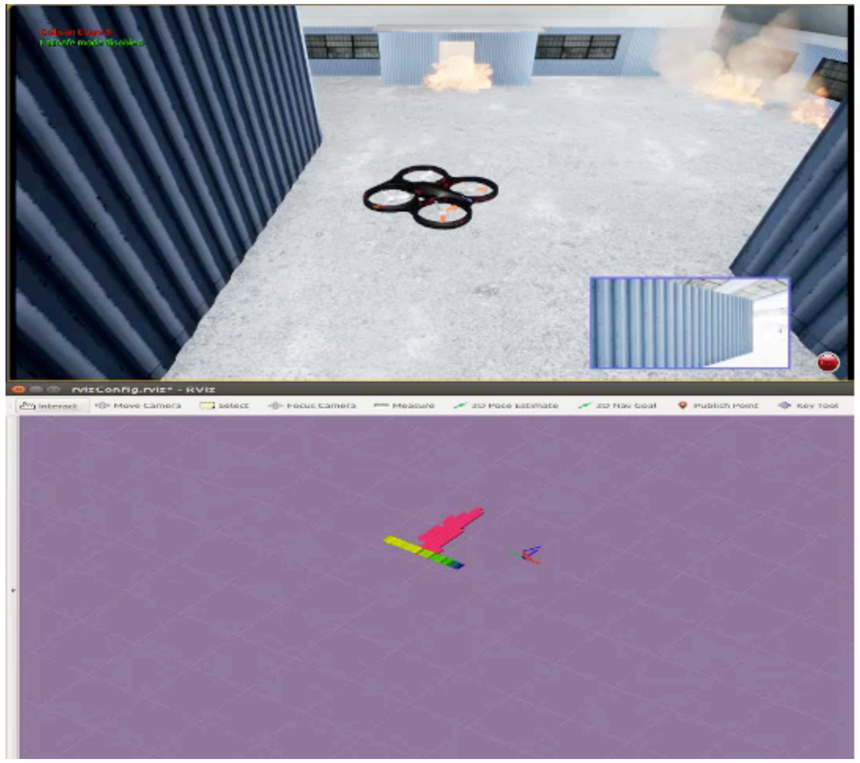

Aerial Photography

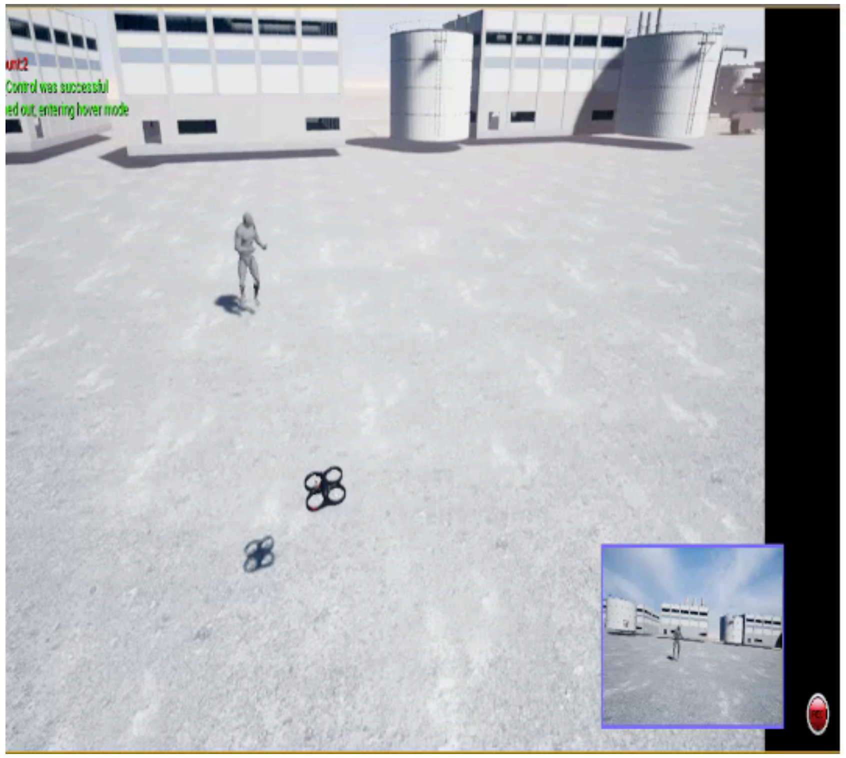

```
ut2
Control was successful
ed out, entering hover mode
```

MAVBench: Kernel Decomposition

MAVBench: Kernel Decomposition

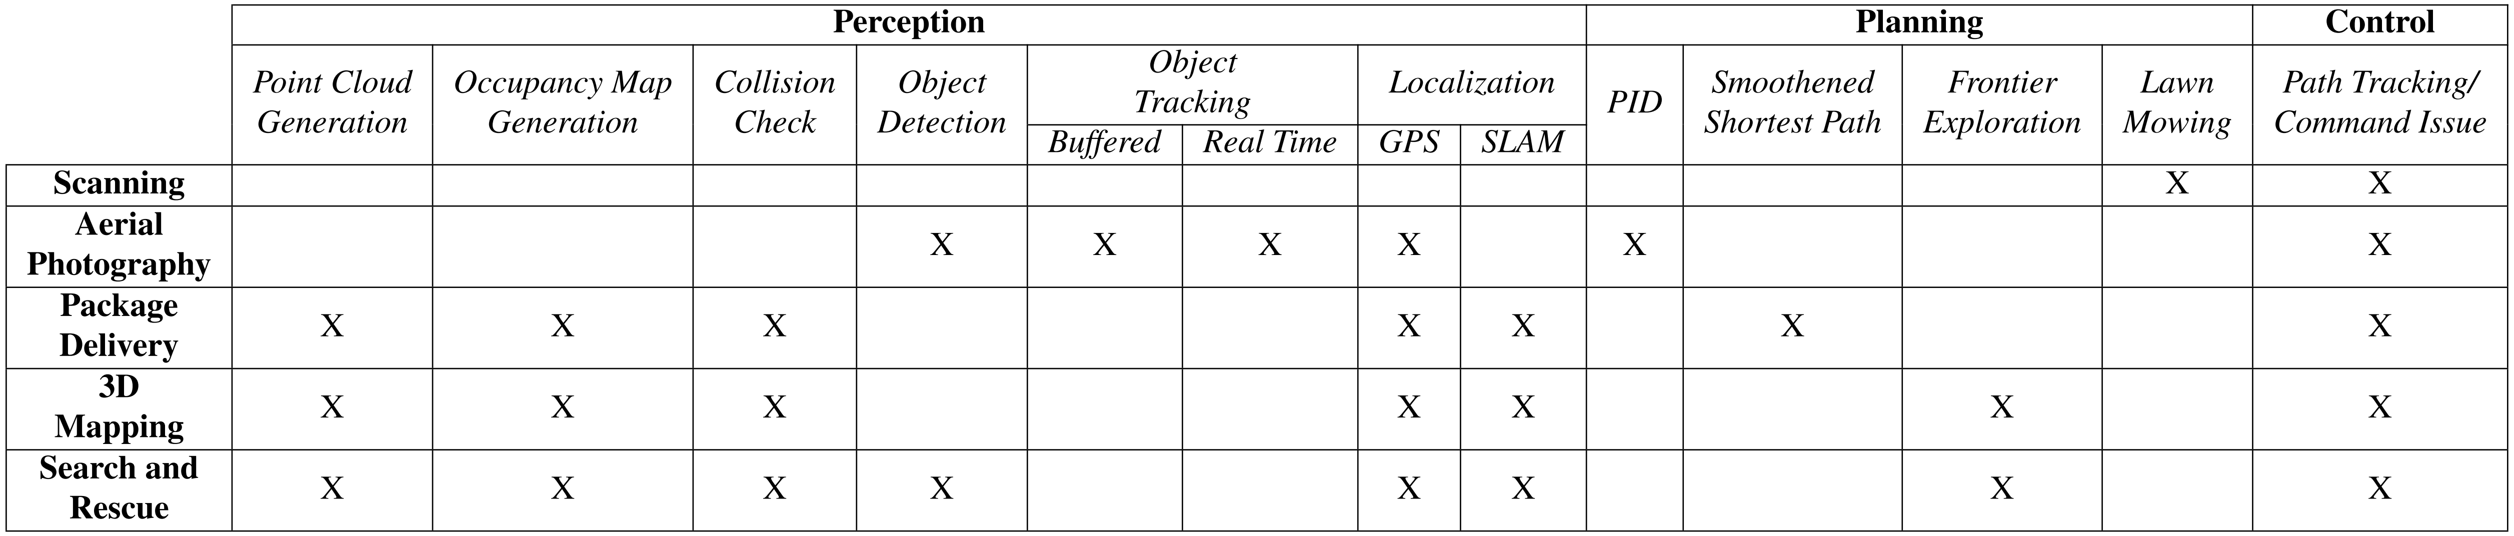

	Perception								Planning				Control
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												X	X
Aerial Photography				X	X	X	X		X				X
Package Delivery	X	X	X				X	X		X			X
3D Mapping	X	X	X				X	X			X		X
Search and Rescue	X	X	X	X			X	X			X		X

MAVBench: Kernel Decomposition

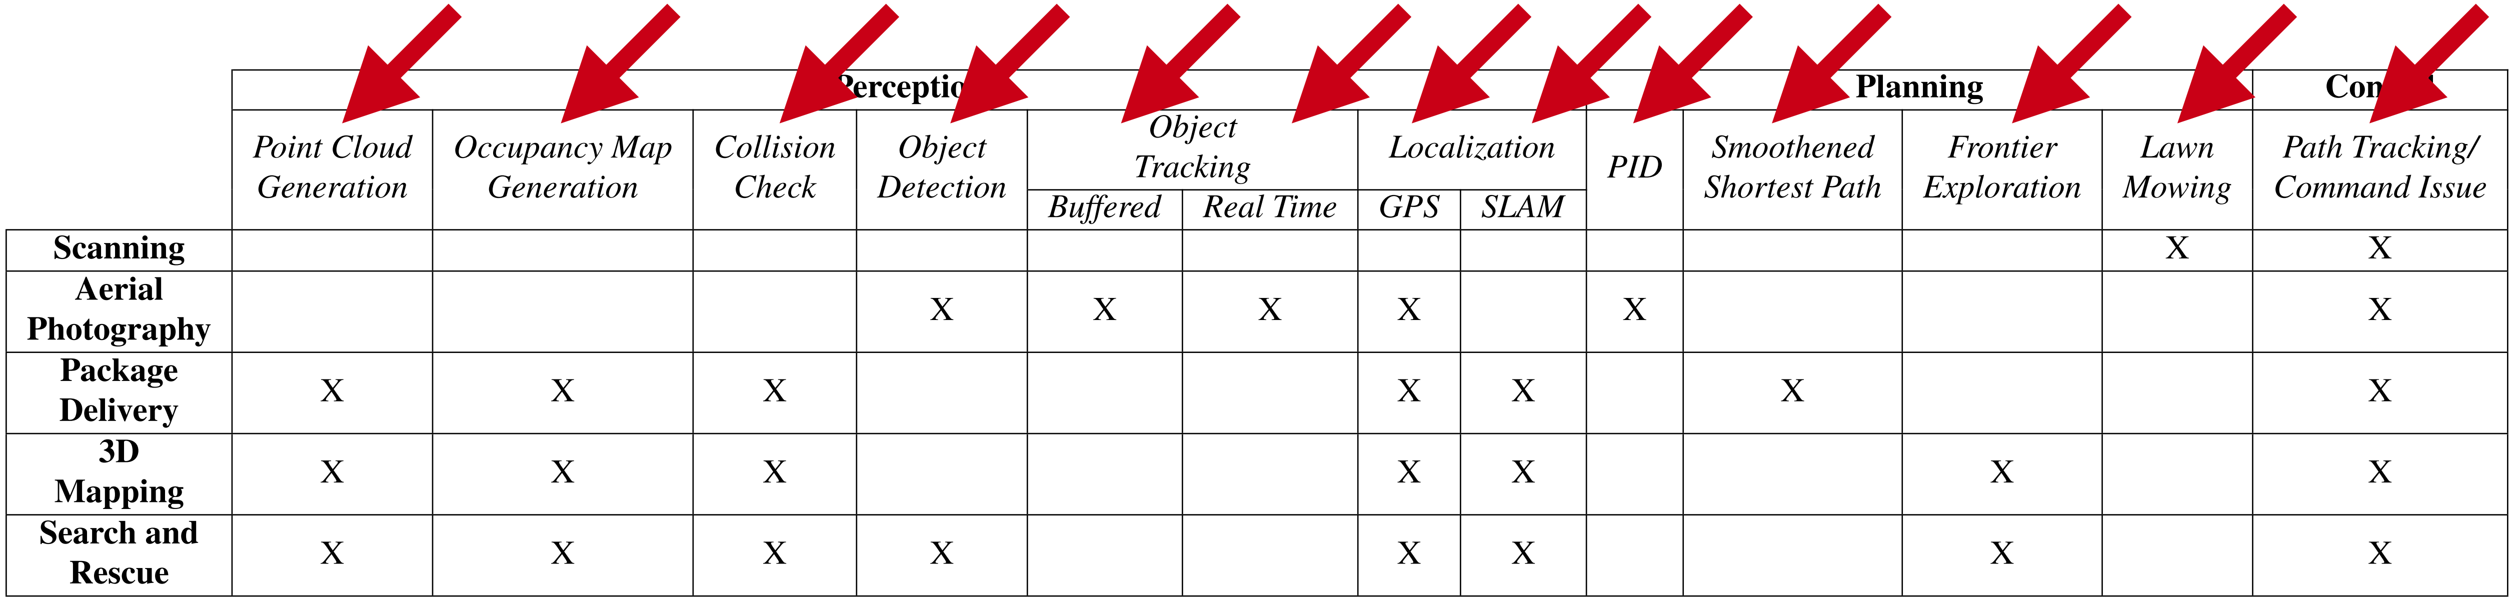

	Perception								Planning				Control
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												X	X
Aerial Photography				X	X	X	X		X				X
Package Delivery	X	X	X				X	X		X			X
3D Mapping	X	X	X				X	X			X		X
Search and Rescue	X	X	X	X			X	X			X		X

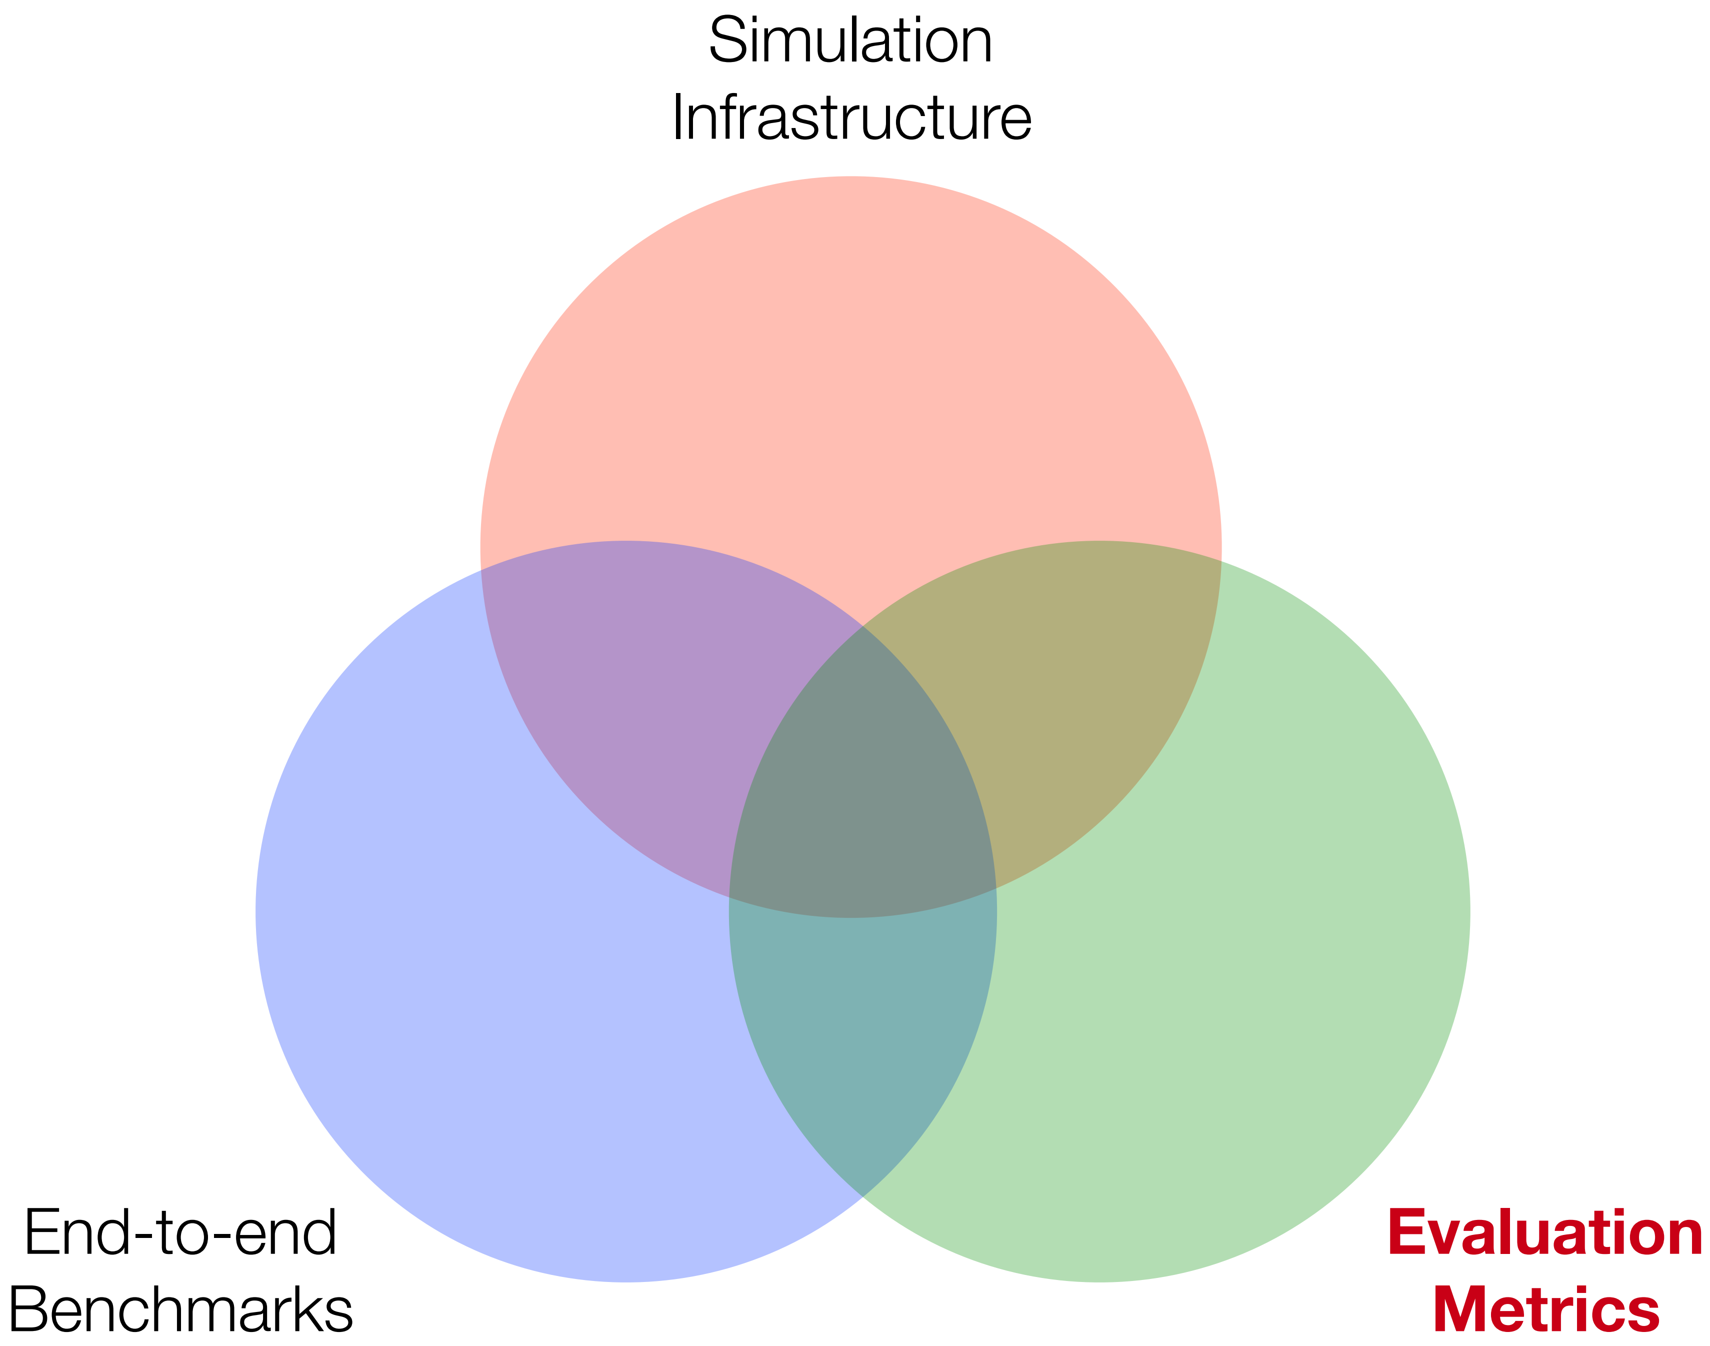
Simulation
Infrastructure
End-to-end
Benchmarks
**Evaluation
Metrics**

“Quality of Flight” Metrics

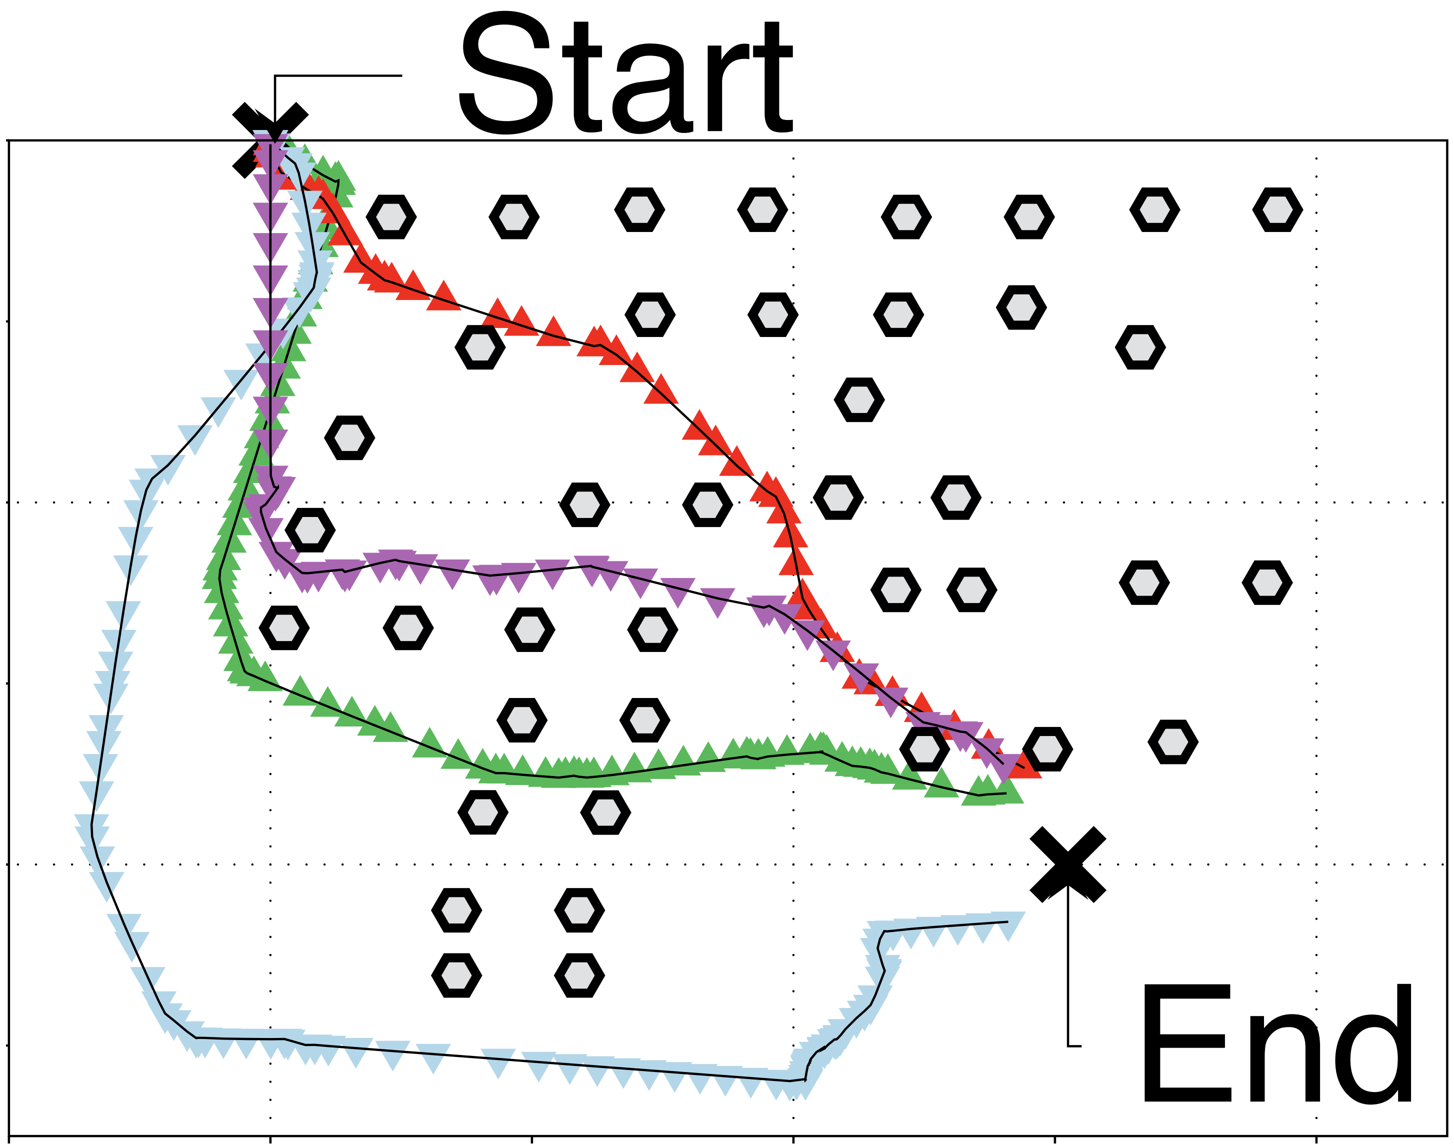

“Quality of Flight” Metrics

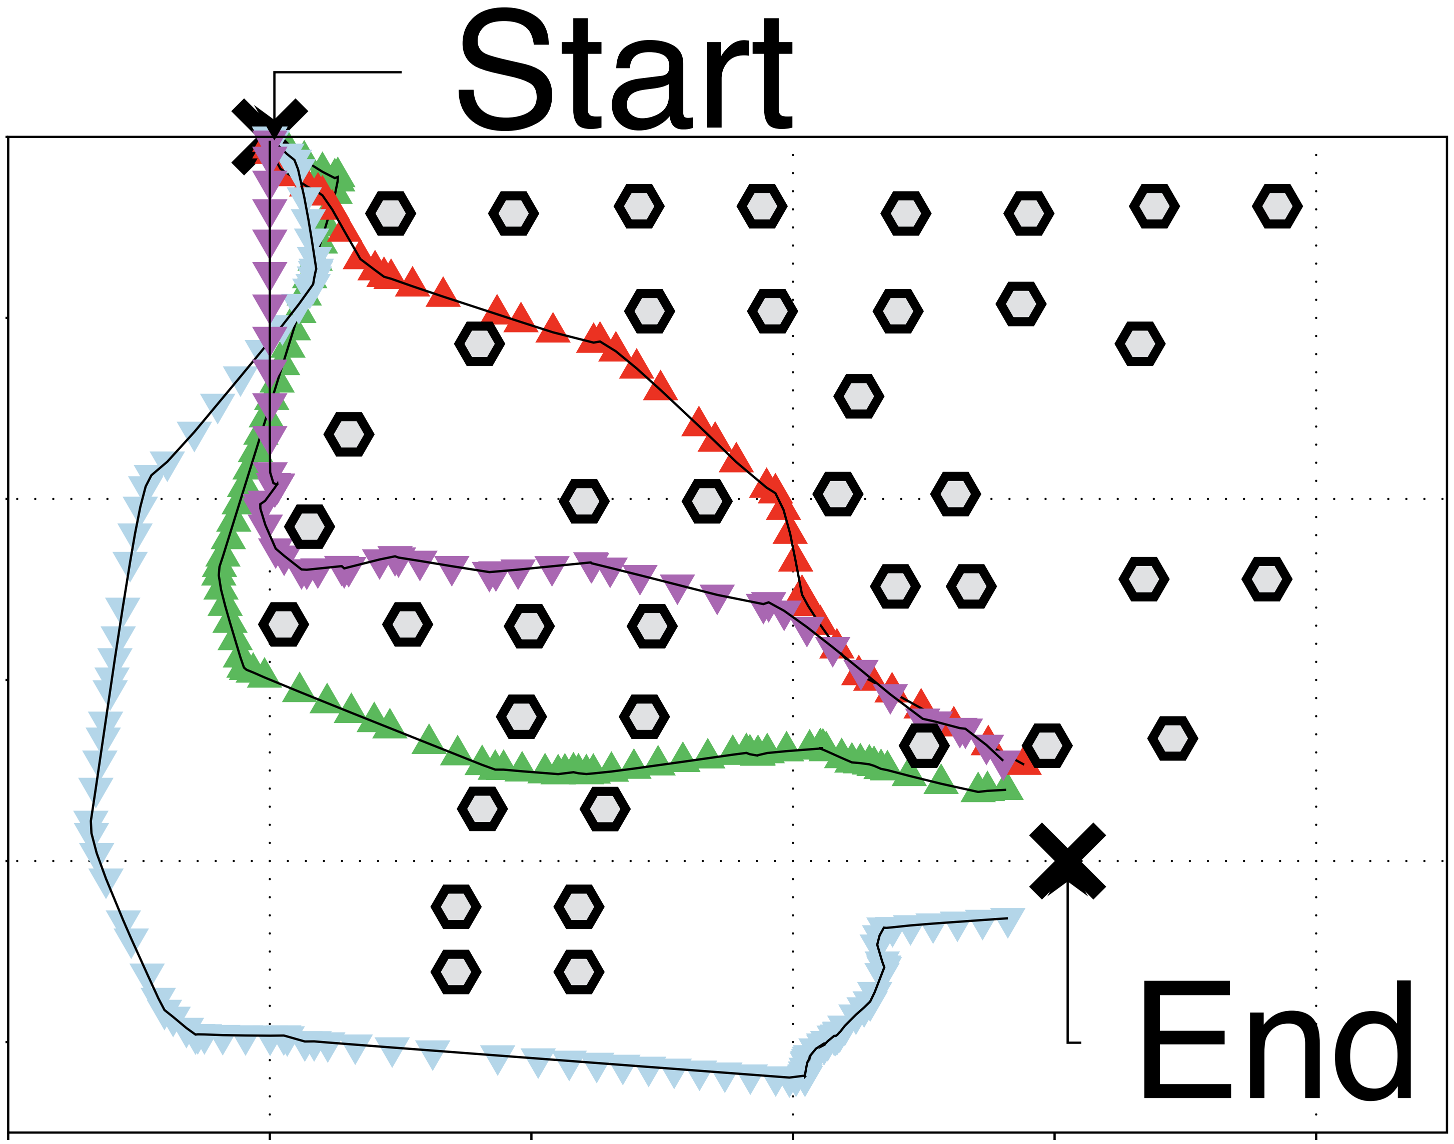

“Quality of Flight” Metrics

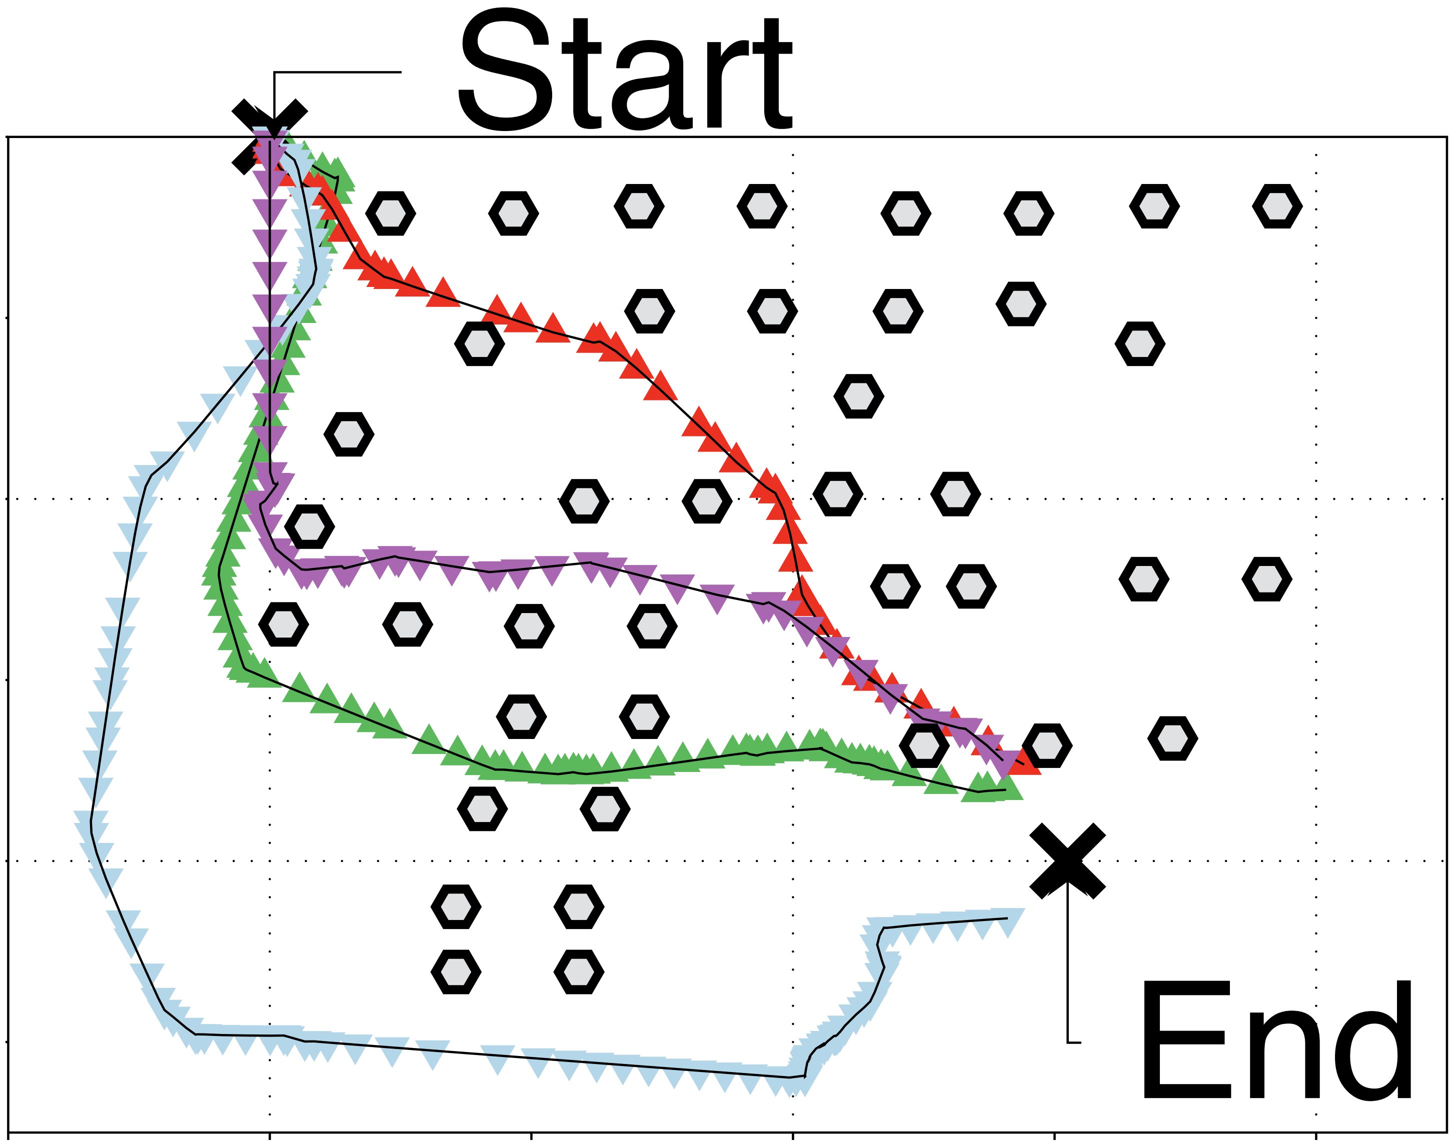

Trajectory

“Quality of Flight” Metrics

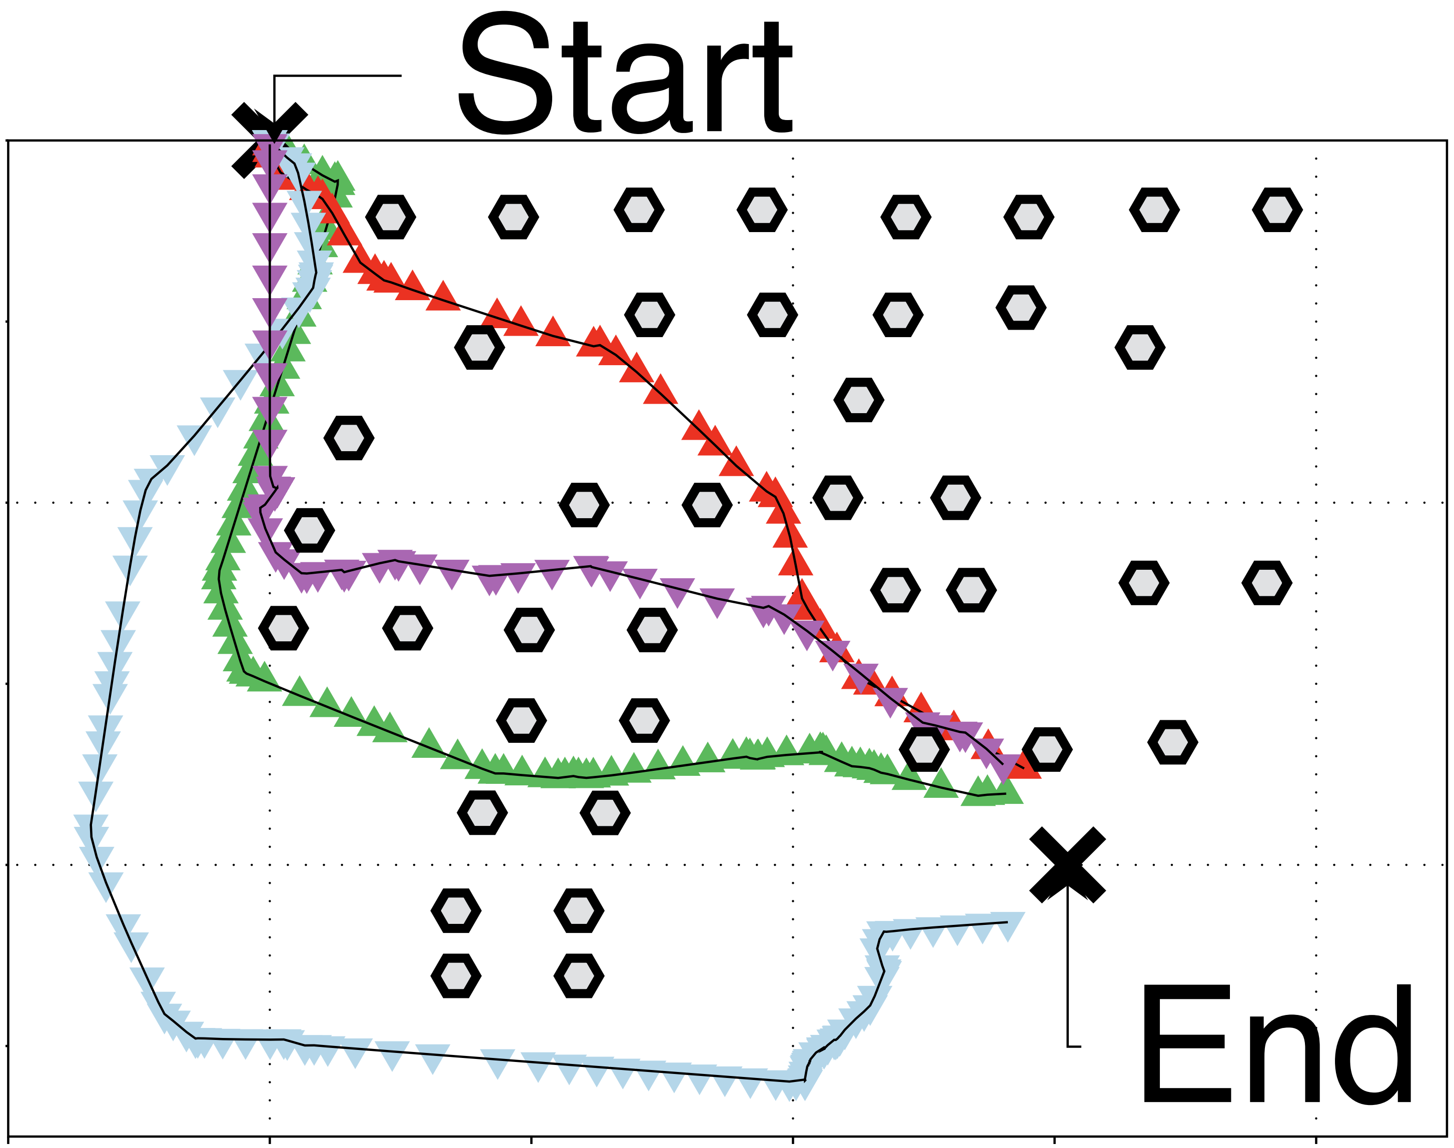

Trajectory

Energy

“Quality of Flight” Metrics

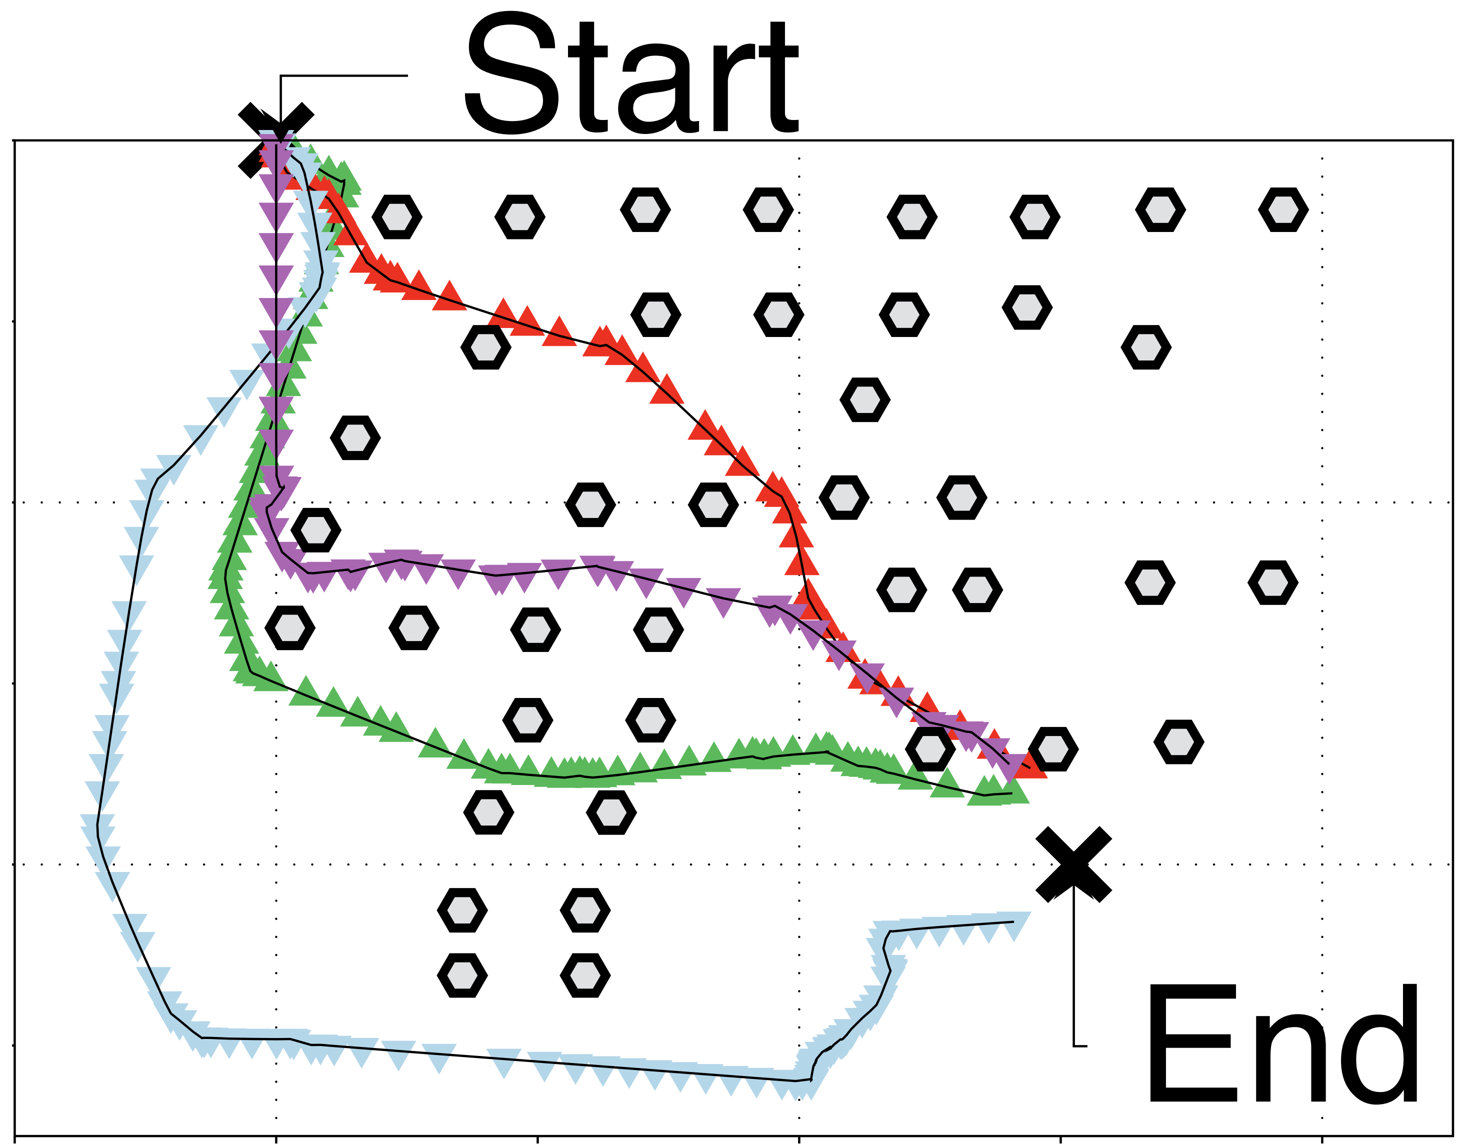

Trajectory

Energy

Mission time

“Quality of Flight” Metrics

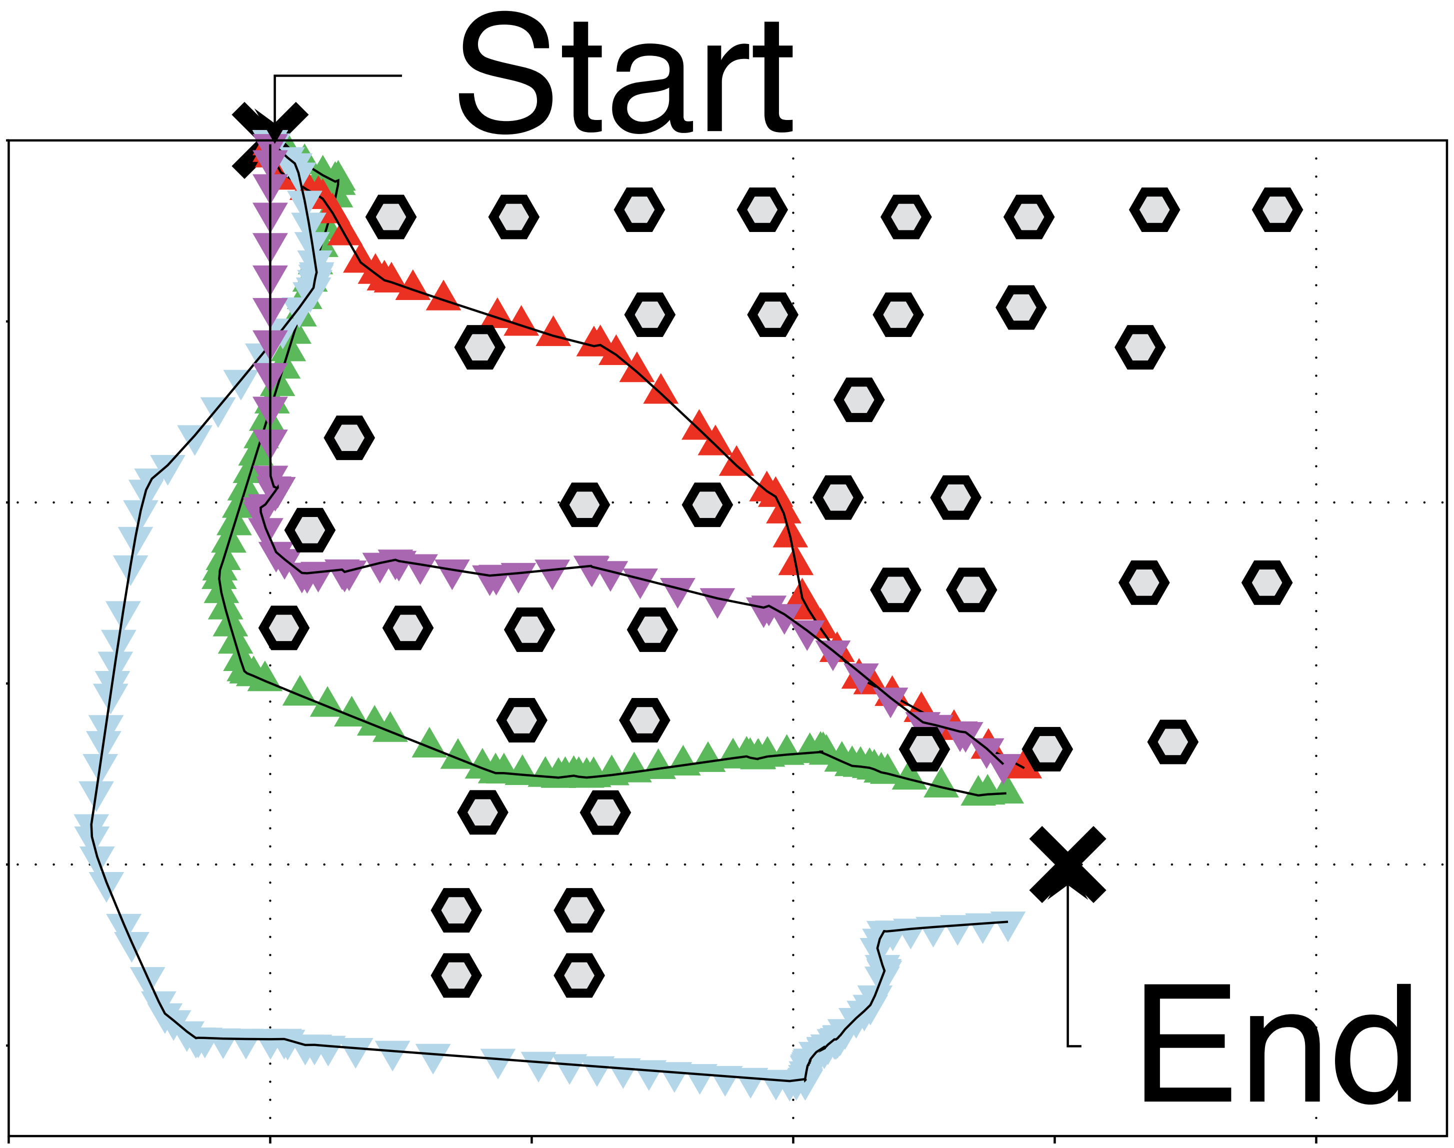

Trajectory

Energy

Mission time

Velocity

“Quality of Flight” Metrics

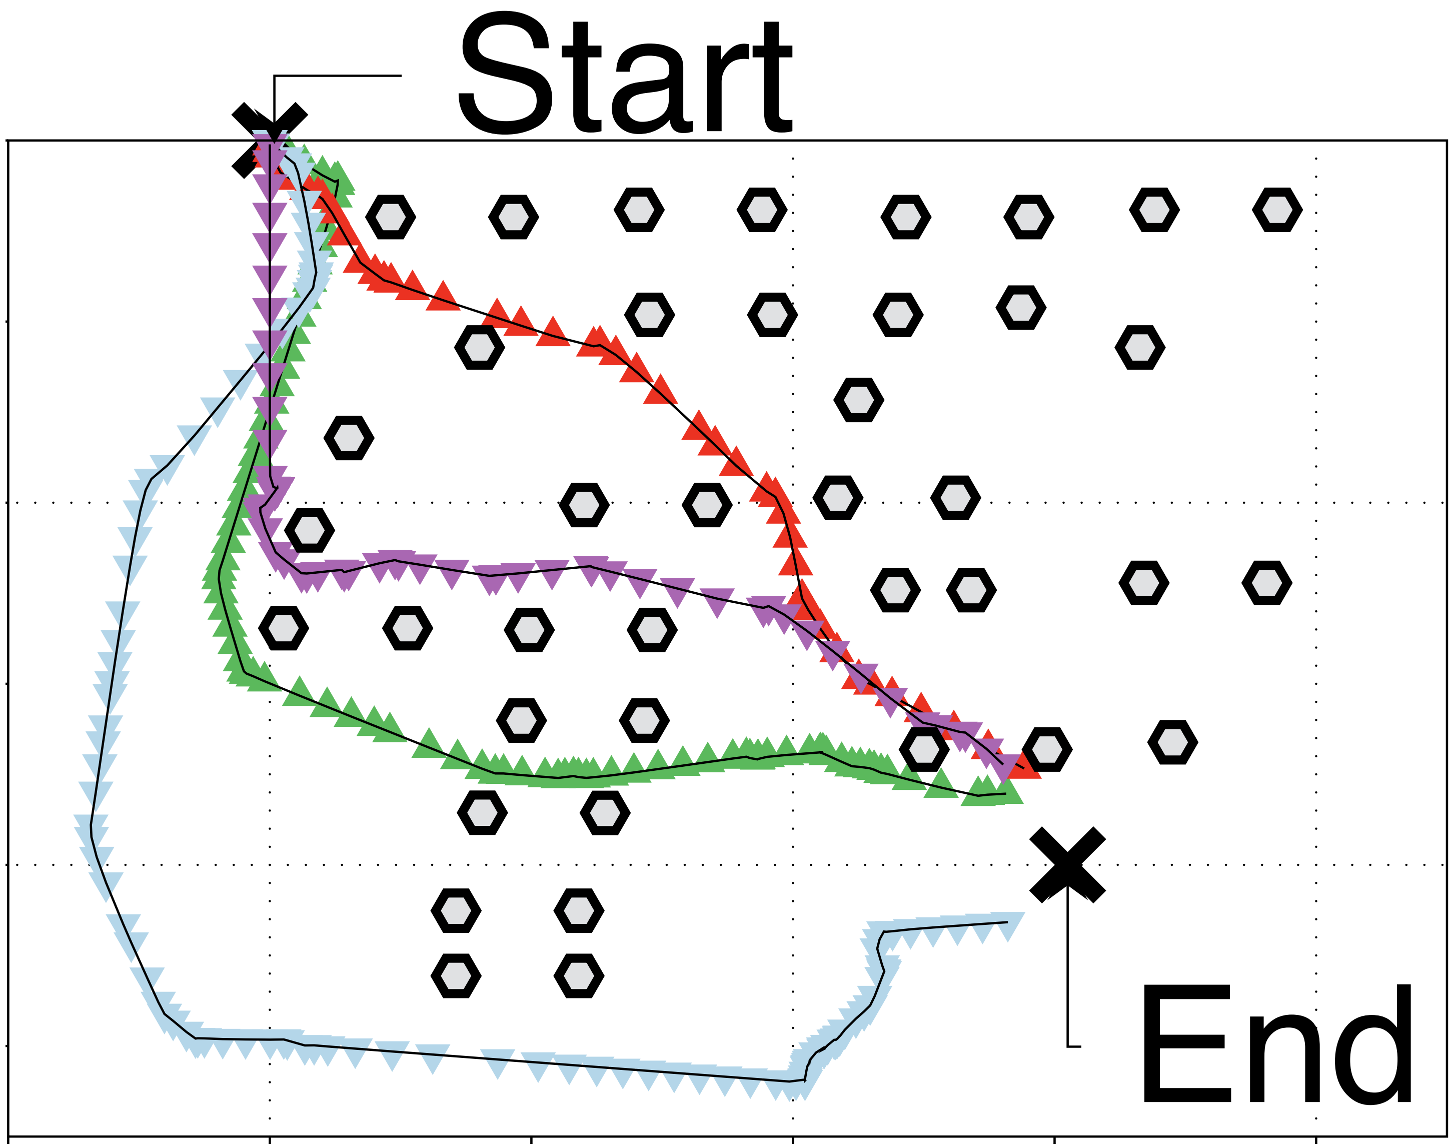

Trajectory

Energy

Mission time

Velocity

Jerk

“Quality of Flight” Metrics

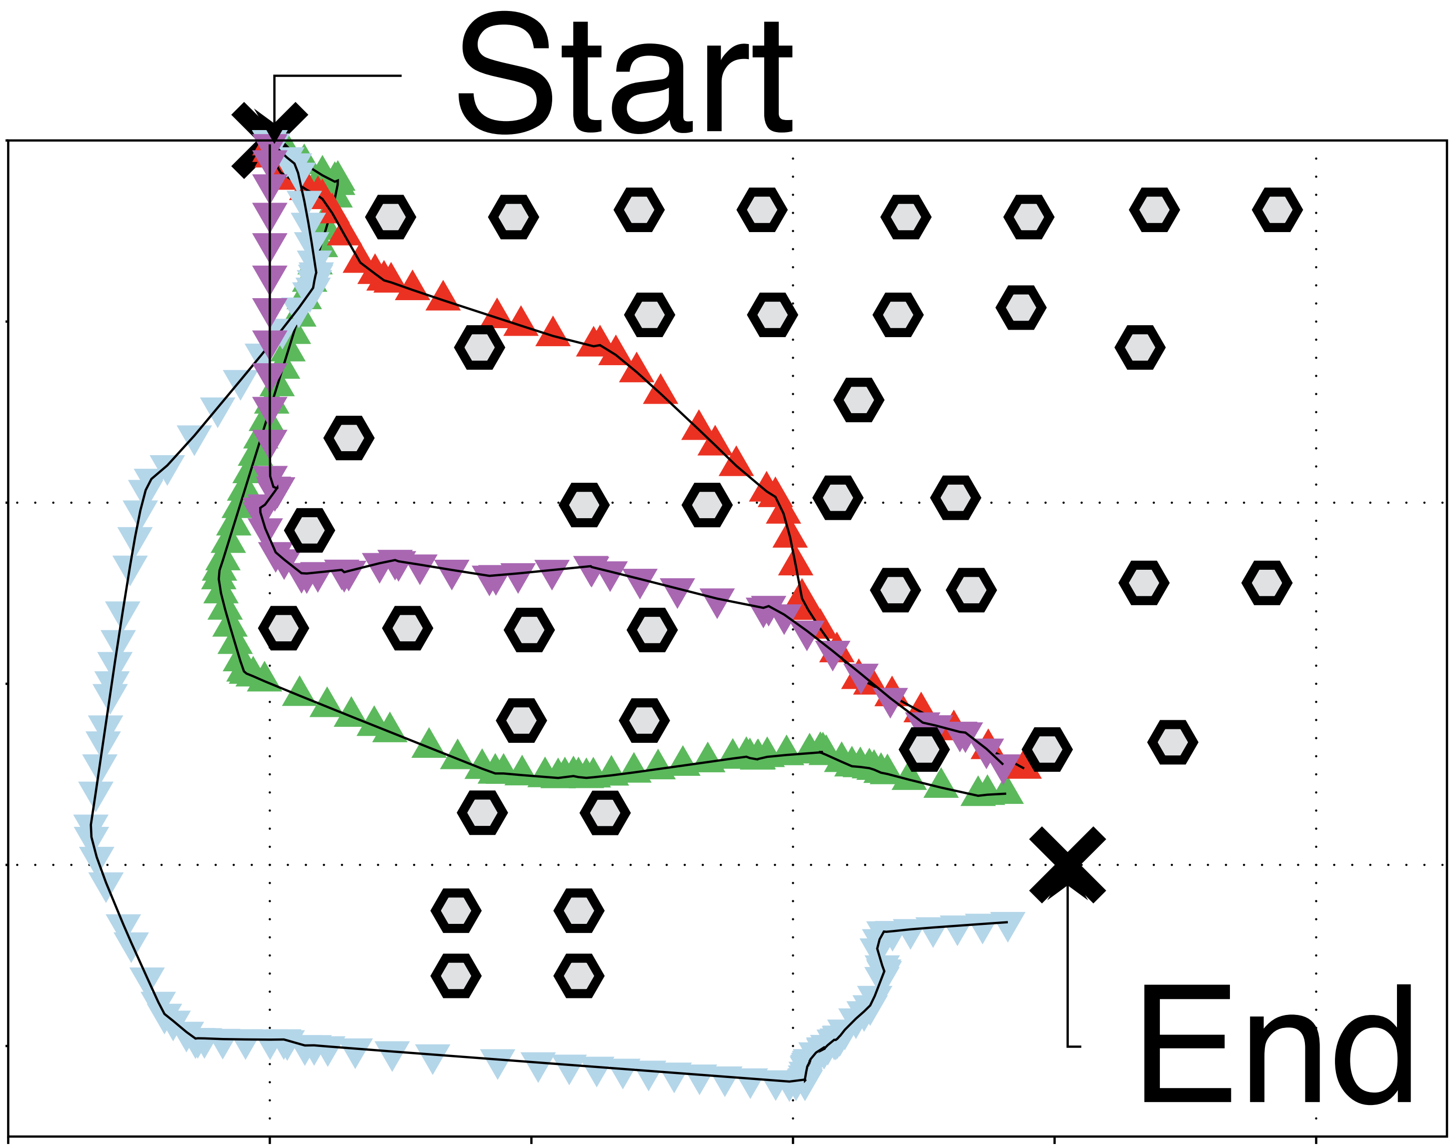

Trajectory

Energy

Mission time

Velocity

Jerk

…

A Case Study of Power Efficiency

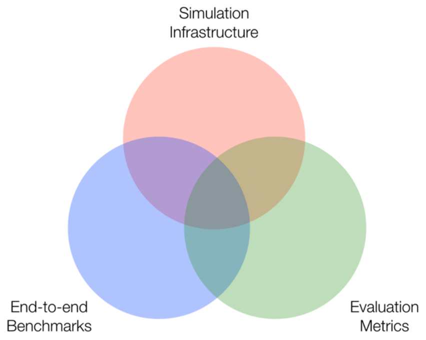

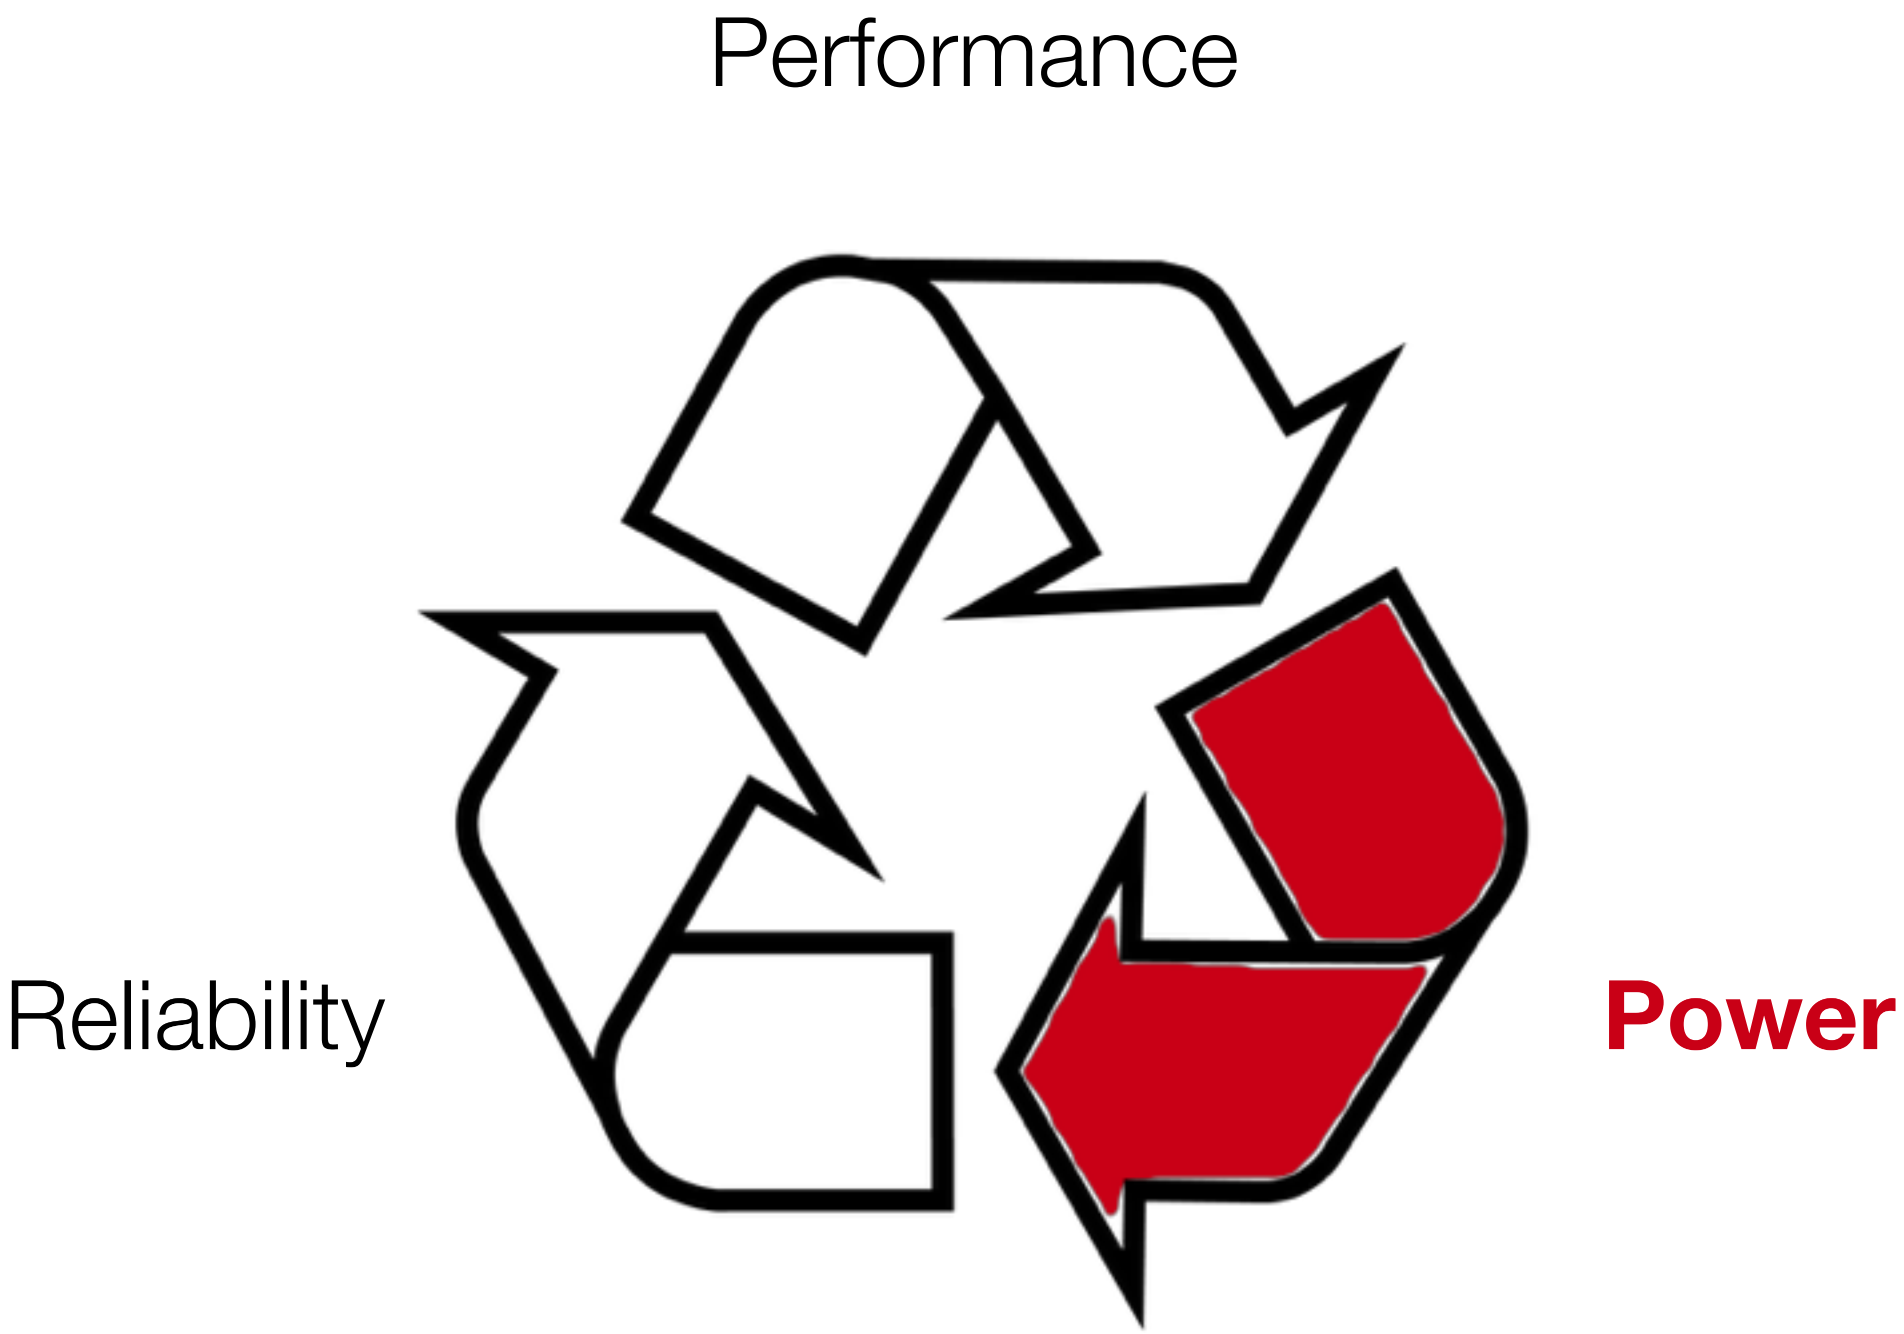

Where have my Joules gone?

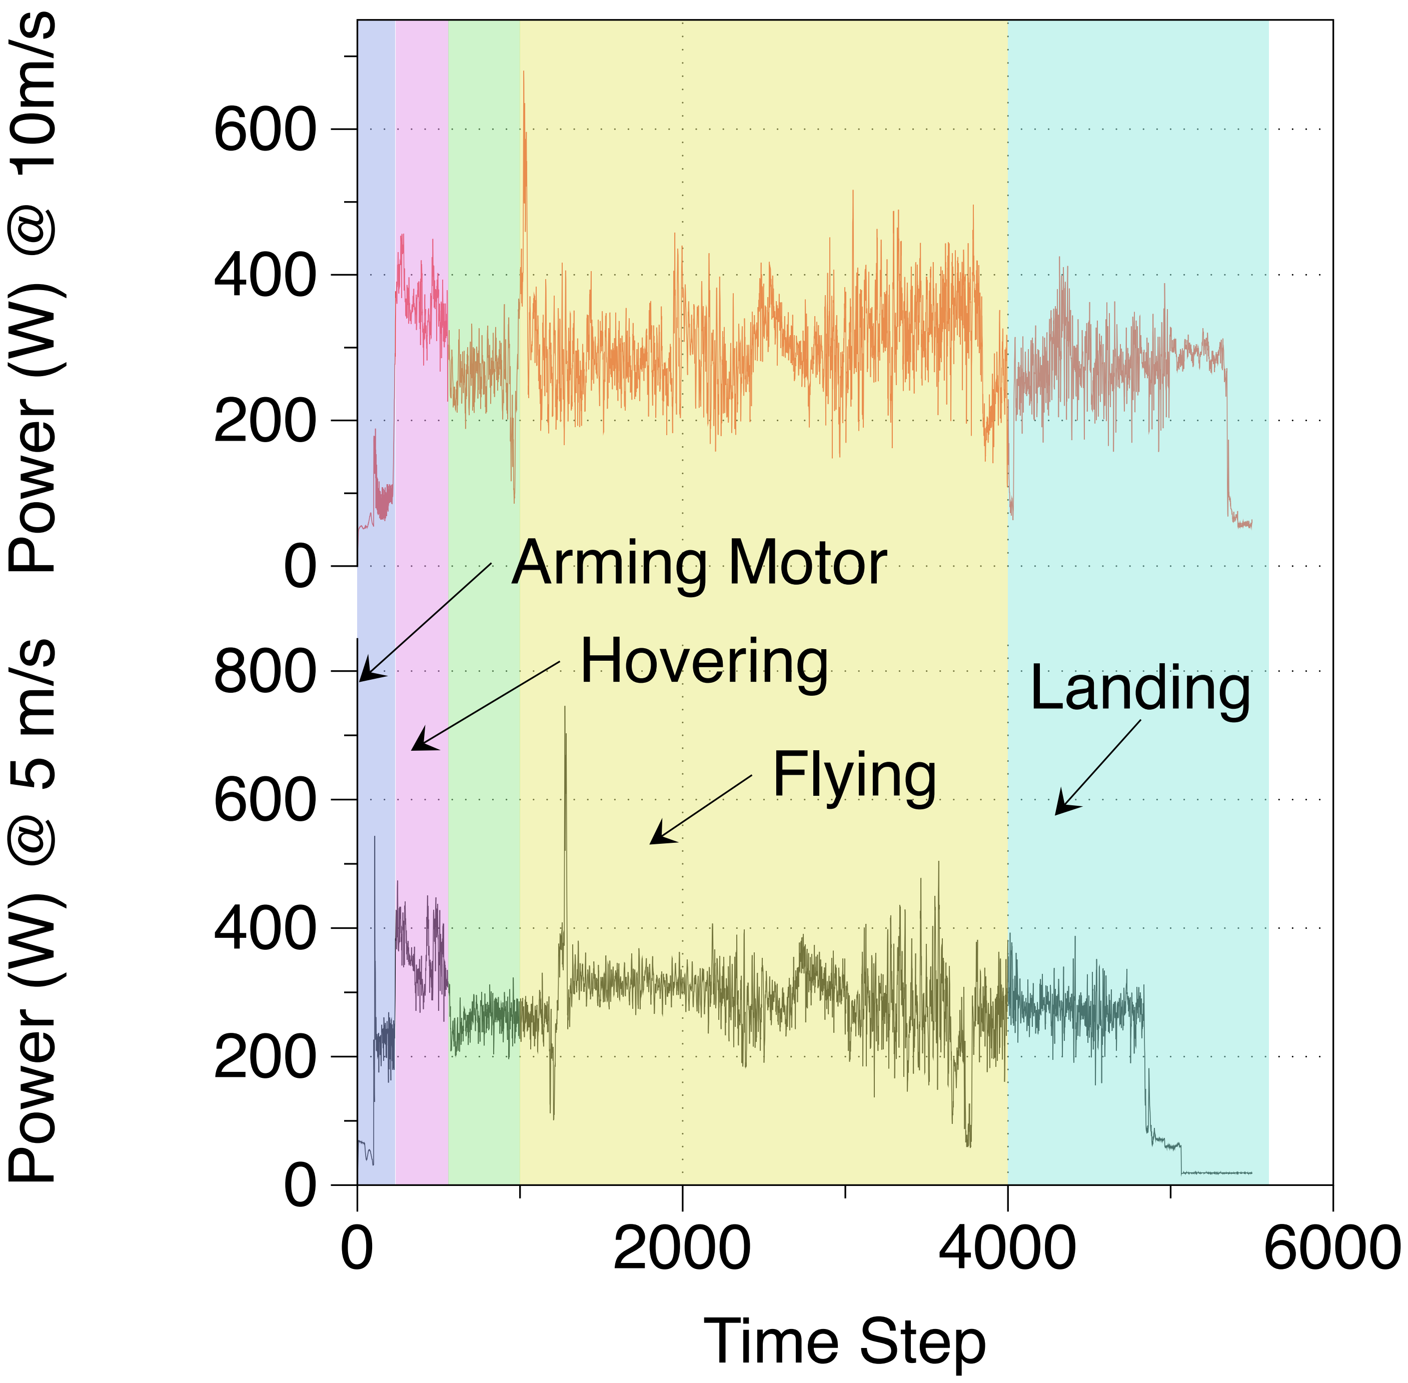

Where have my Joules gone?

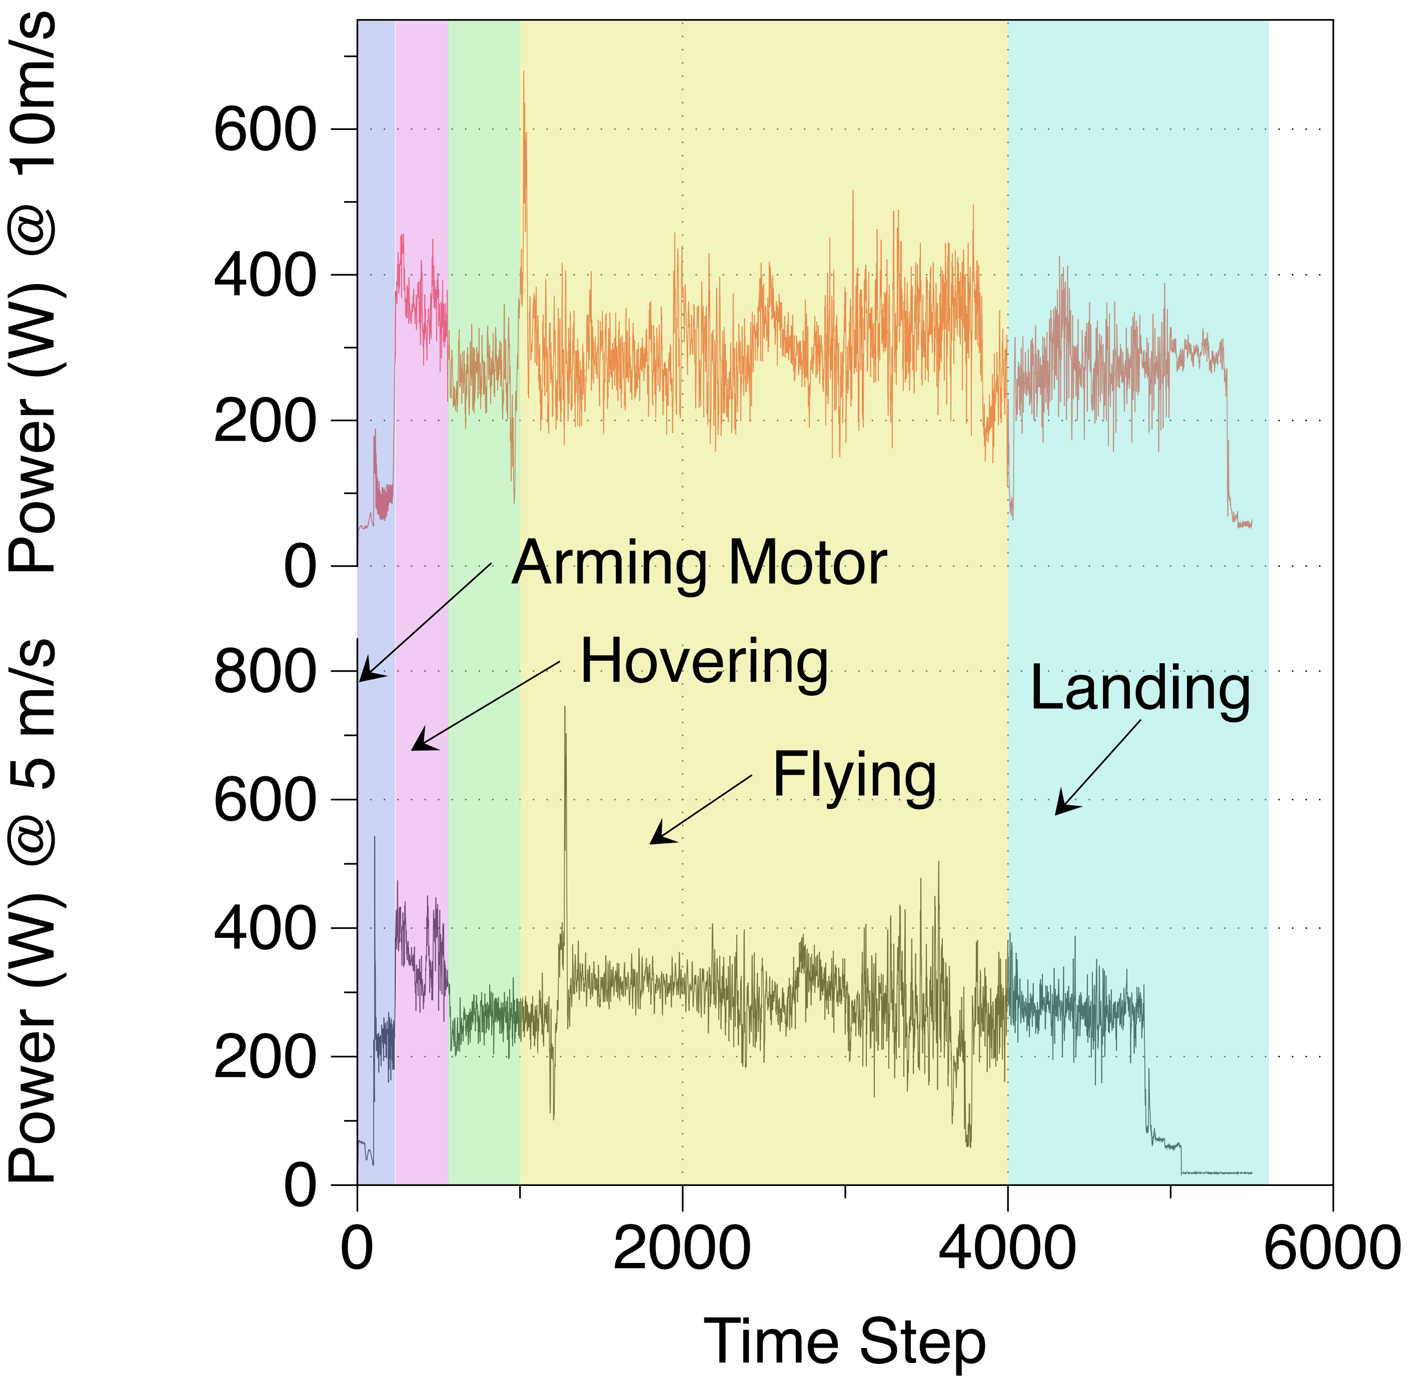

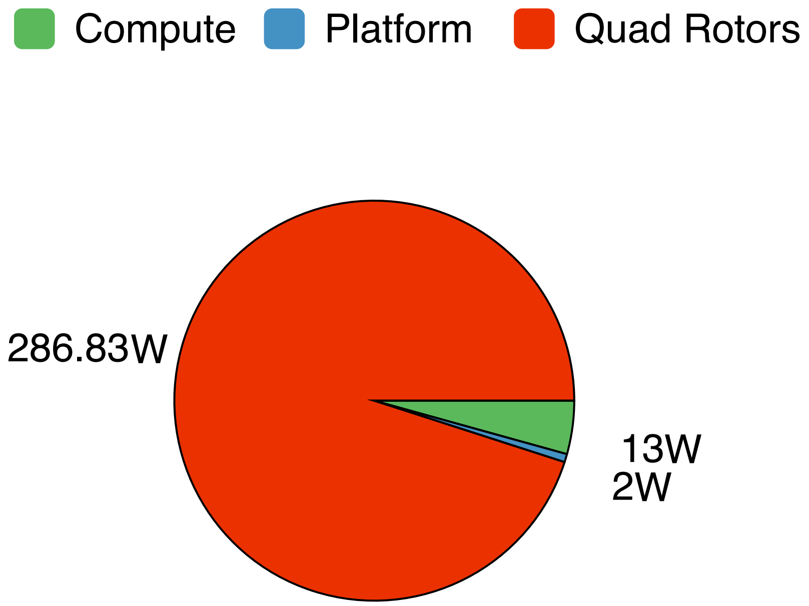

Most of the energy is burned in the mechanical subsystem.

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

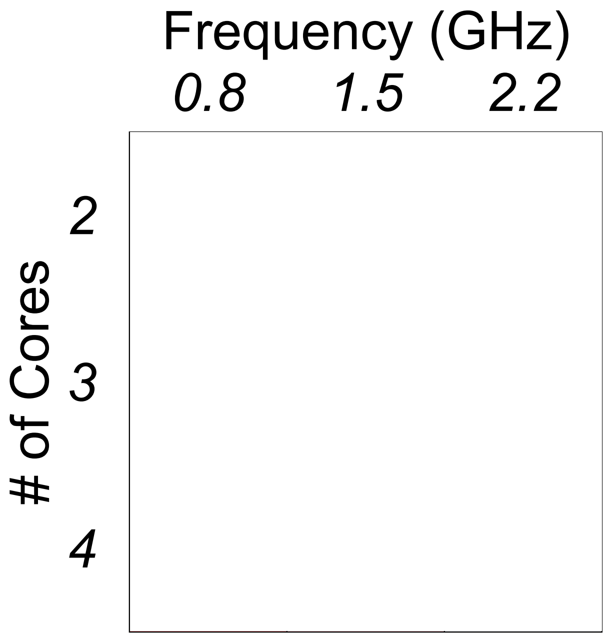

Mission Time (s)

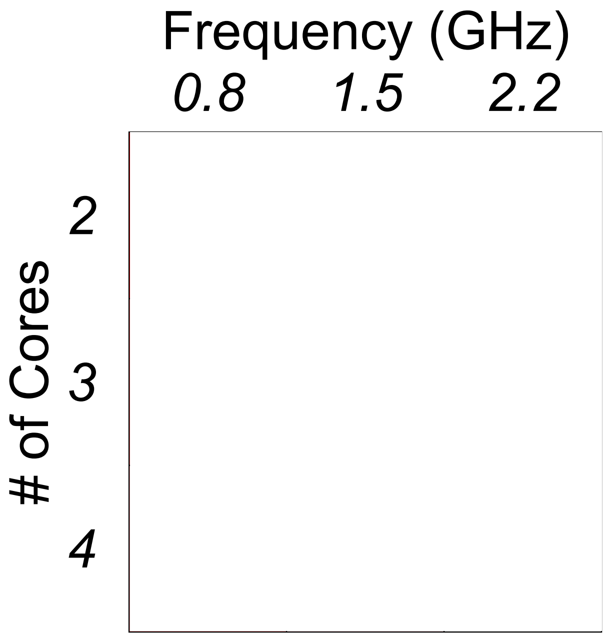

Energy (kJ)

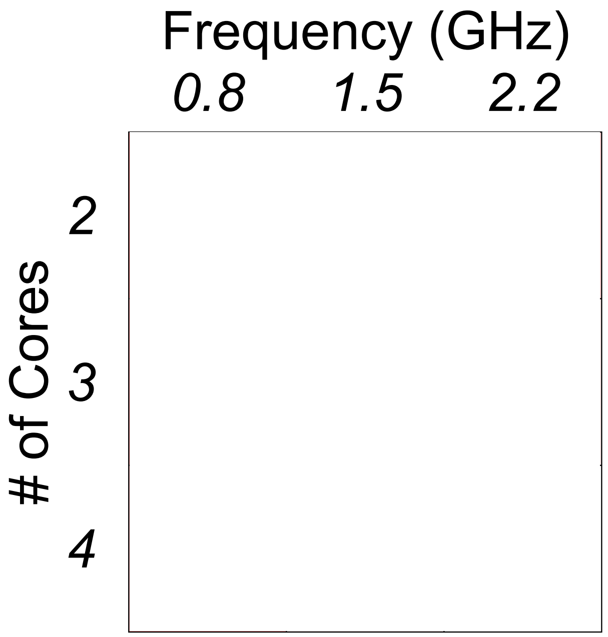

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2	0.34	0.65	0.75
3			
4			

Mission Time (s)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2			
3			
4			

Energy (kJ)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2			
3			
4			

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

	Frequency (GHz)		
# of Cores	*0.8*	*1.5*	*2.2*
2	0.34	0.65	0.75
3			
4			

Mission Time (s)

	Frequency (GHz)		
# of Cores	*0.8*	*1.5*	*2.2*
2	1388.1	676.8	530.5
3			
4			

Energy (kJ)

	Frequency (GHz)		
# of Cores	*0.8*	*1.5*	*2.2*
2			
3			
4			

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

# of Cores \ Frequency (GHz)	*0.8*	*1.5*	*2.2*
2	0.34	0.65	0.75
3			
4			

Mission Time (s)

# of Cores \ Frequency (GHz)	*0.8*	*1.5*	*2.2*
2	1388.1	676.8	530.5
3			
4			

Energy (kJ)

# of Cores \ Frequency (GHz)	*0.8*	*1.5*	*2.2*
2	625.4	370	308.1
3			
4			

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

# of Cores	Frequency (GHz)		
	0.8	1.5	2.2
2	0.34	0.65	0.75
3	0.49	0.86	1.07
4	0.72	0.92	1.12

Mission Time (s)

# of Cores	Frequency (GHz)		
	0.8	1.5	2.2
2	1388.1	676.8	530.5
3	980.1	477.1	393.1
4	716	440.7	326.7

Energy (kJ)

# of Cores	Frequency (GHz)		
	0.8	1.5	2.2
2	625.4	370	308.1
3	554.3	297.5	260.3
4	438.1	269.4	223.9

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	0.34	0.65	0.75
3	0.49	0.[illegible]6	1.07
4	0.72	0.92	1.12

Mission Time (s)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	1388.1	676.8	530.5
3	980.1	477.1	393.1
4	716	440.7	326.7

Energy (kJ)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	625.4	370	308.1
3	554.3	297.5	260.3
4	438.1	269.4	223.9

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

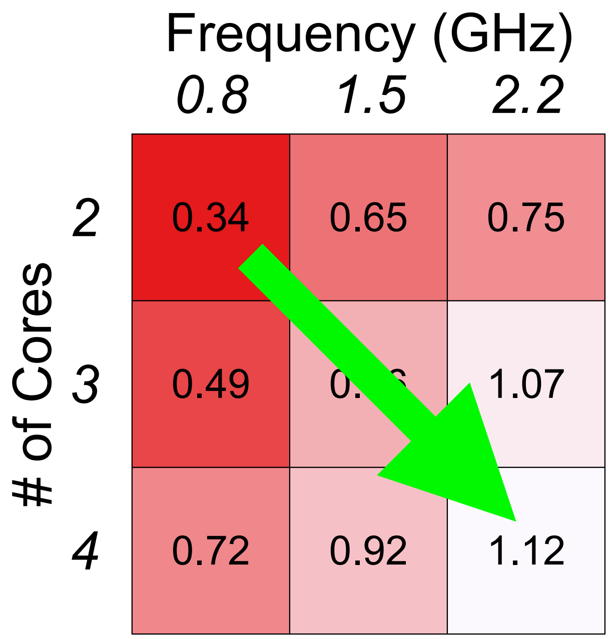

# of Cores \ Frequency (GHz)	0.8	1.5	2.2
2	0.34	0.65	0.75
3	0.49	0.[illegible]6	1.07
4	0.72	0.92	1.12

Mission Time (s)

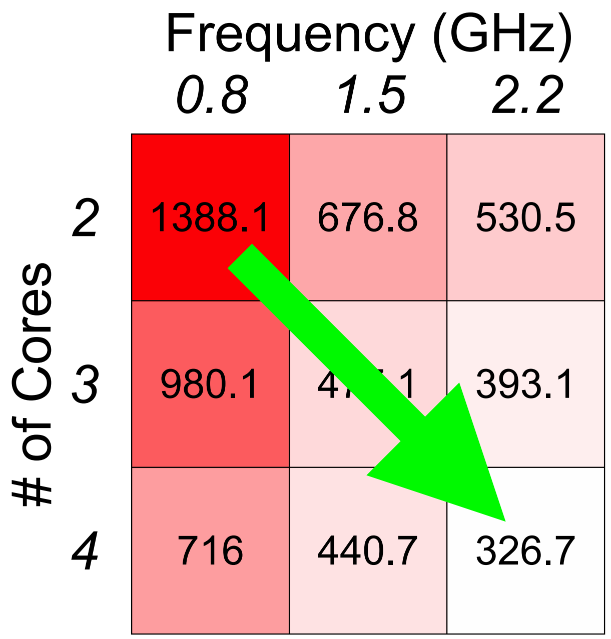

# of Cores \ Frequency (GHz)	0.8	1.5	2.2
2	1388.1	676.8	530.5
3	980.1	47[illegible].1	393.1
4	716	440.7	326.7

Energy (kJ)

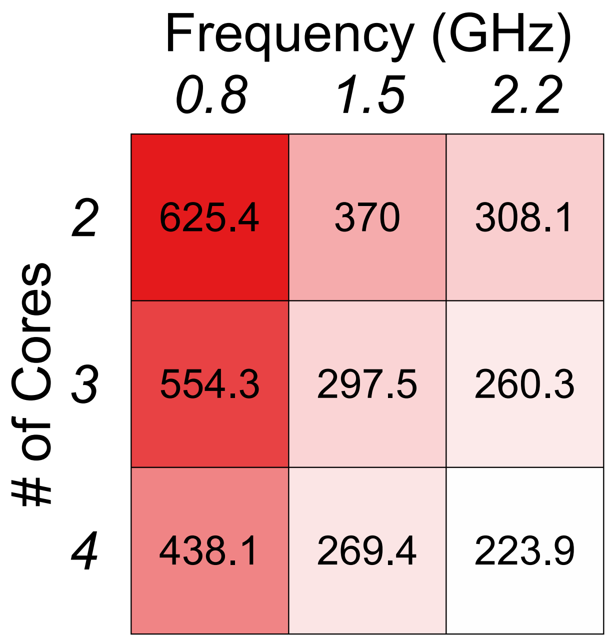

# of Cores \ Frequency (GHz)	0.8	1.5	2.2
2	625.4	370	308.1
3	554.3	297.5	260.3
4	438.1	269.4	223.9

3D Mapping

Core/Frequency Sensitivity Analysis

Average Velocity (m/s)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2	0.34	0.65	0.75
3	0.49	0.[illegible]6	1.07
4	0.72	0.92	1.12

Mission Time (s)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2	1388.1	676.8	530.5
3	980.1	47[illegible]1	393.1
4	716	440.7	326.7

Energy (kJ)

Frequency (GHz)

# of Cores	0.8	1.5	2.2
2	625.4	370	308.1
3	554.3	29[illegible]5	260.3
4	438.1	269.4	223.9

3D Mapping

Core/Frequency Sensitivity Analysis

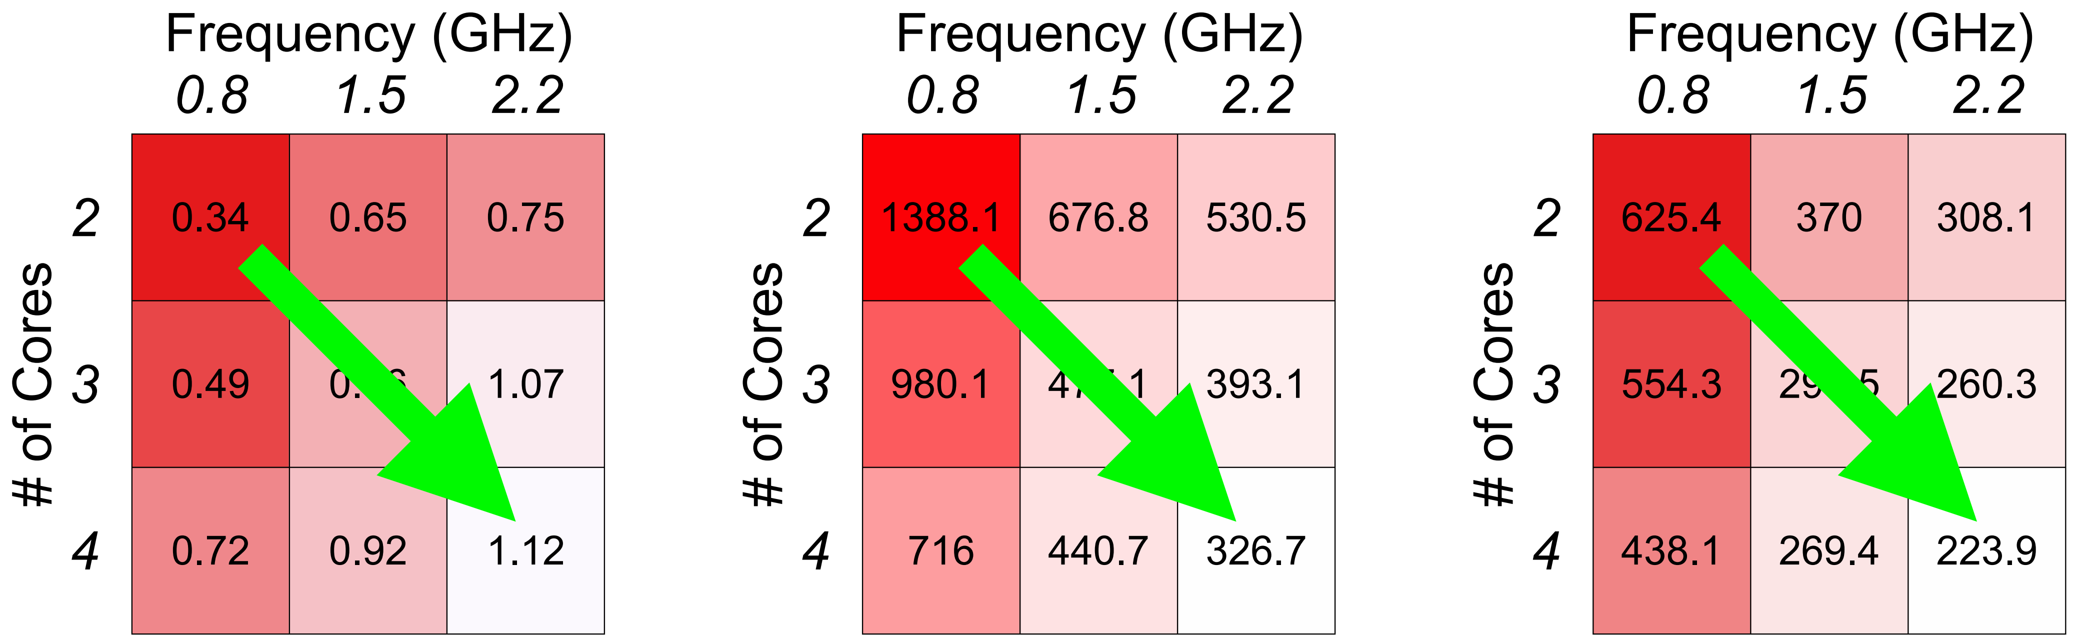

Average Velocity (m/s)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	0.34	0.65	0.75
3	0.49	[illegible]	1.07
4	0.72	0.92	1.12

Mission Time (s)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	1388.1	676.8	530.5
3	980.1	[illegible]	393.1
4	716	440.7	326.7

Energy (kJ)

# of Cores	Frequency (GHz) 0.8	1.5	2.2
2	625.4	370	308.1
3	554.3	[illegible]	260.3
4	438.1	269.4	223.9

3D Mapping

3X improvement in speed, flight time and energy as compute capability increases

Recall: Power Depends on Physical Form Factor

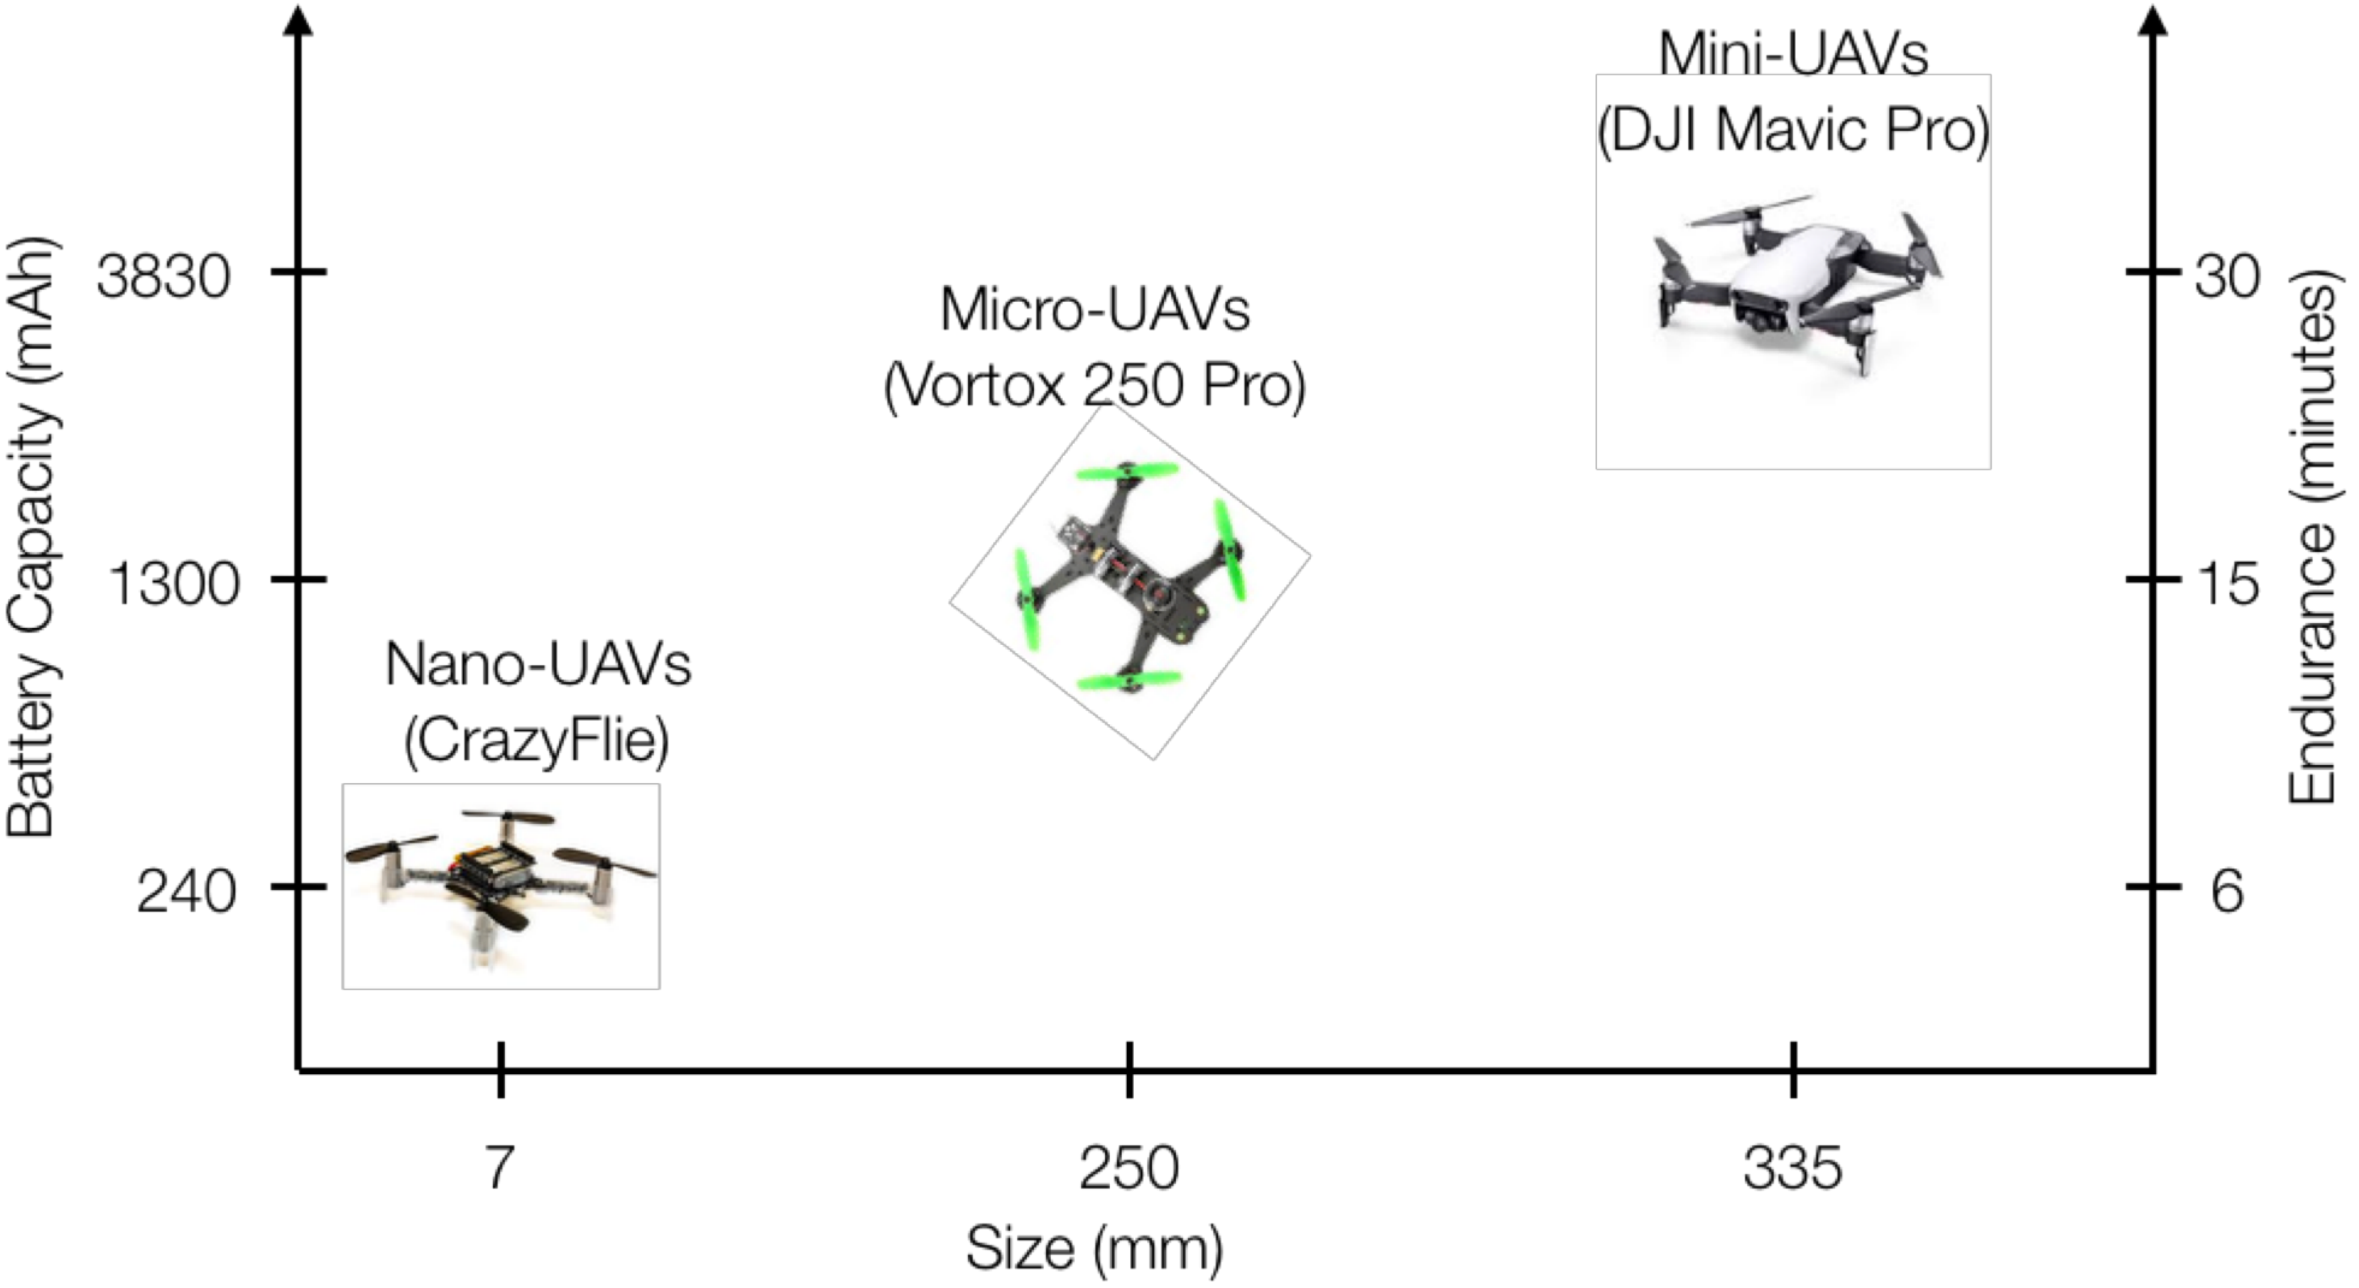

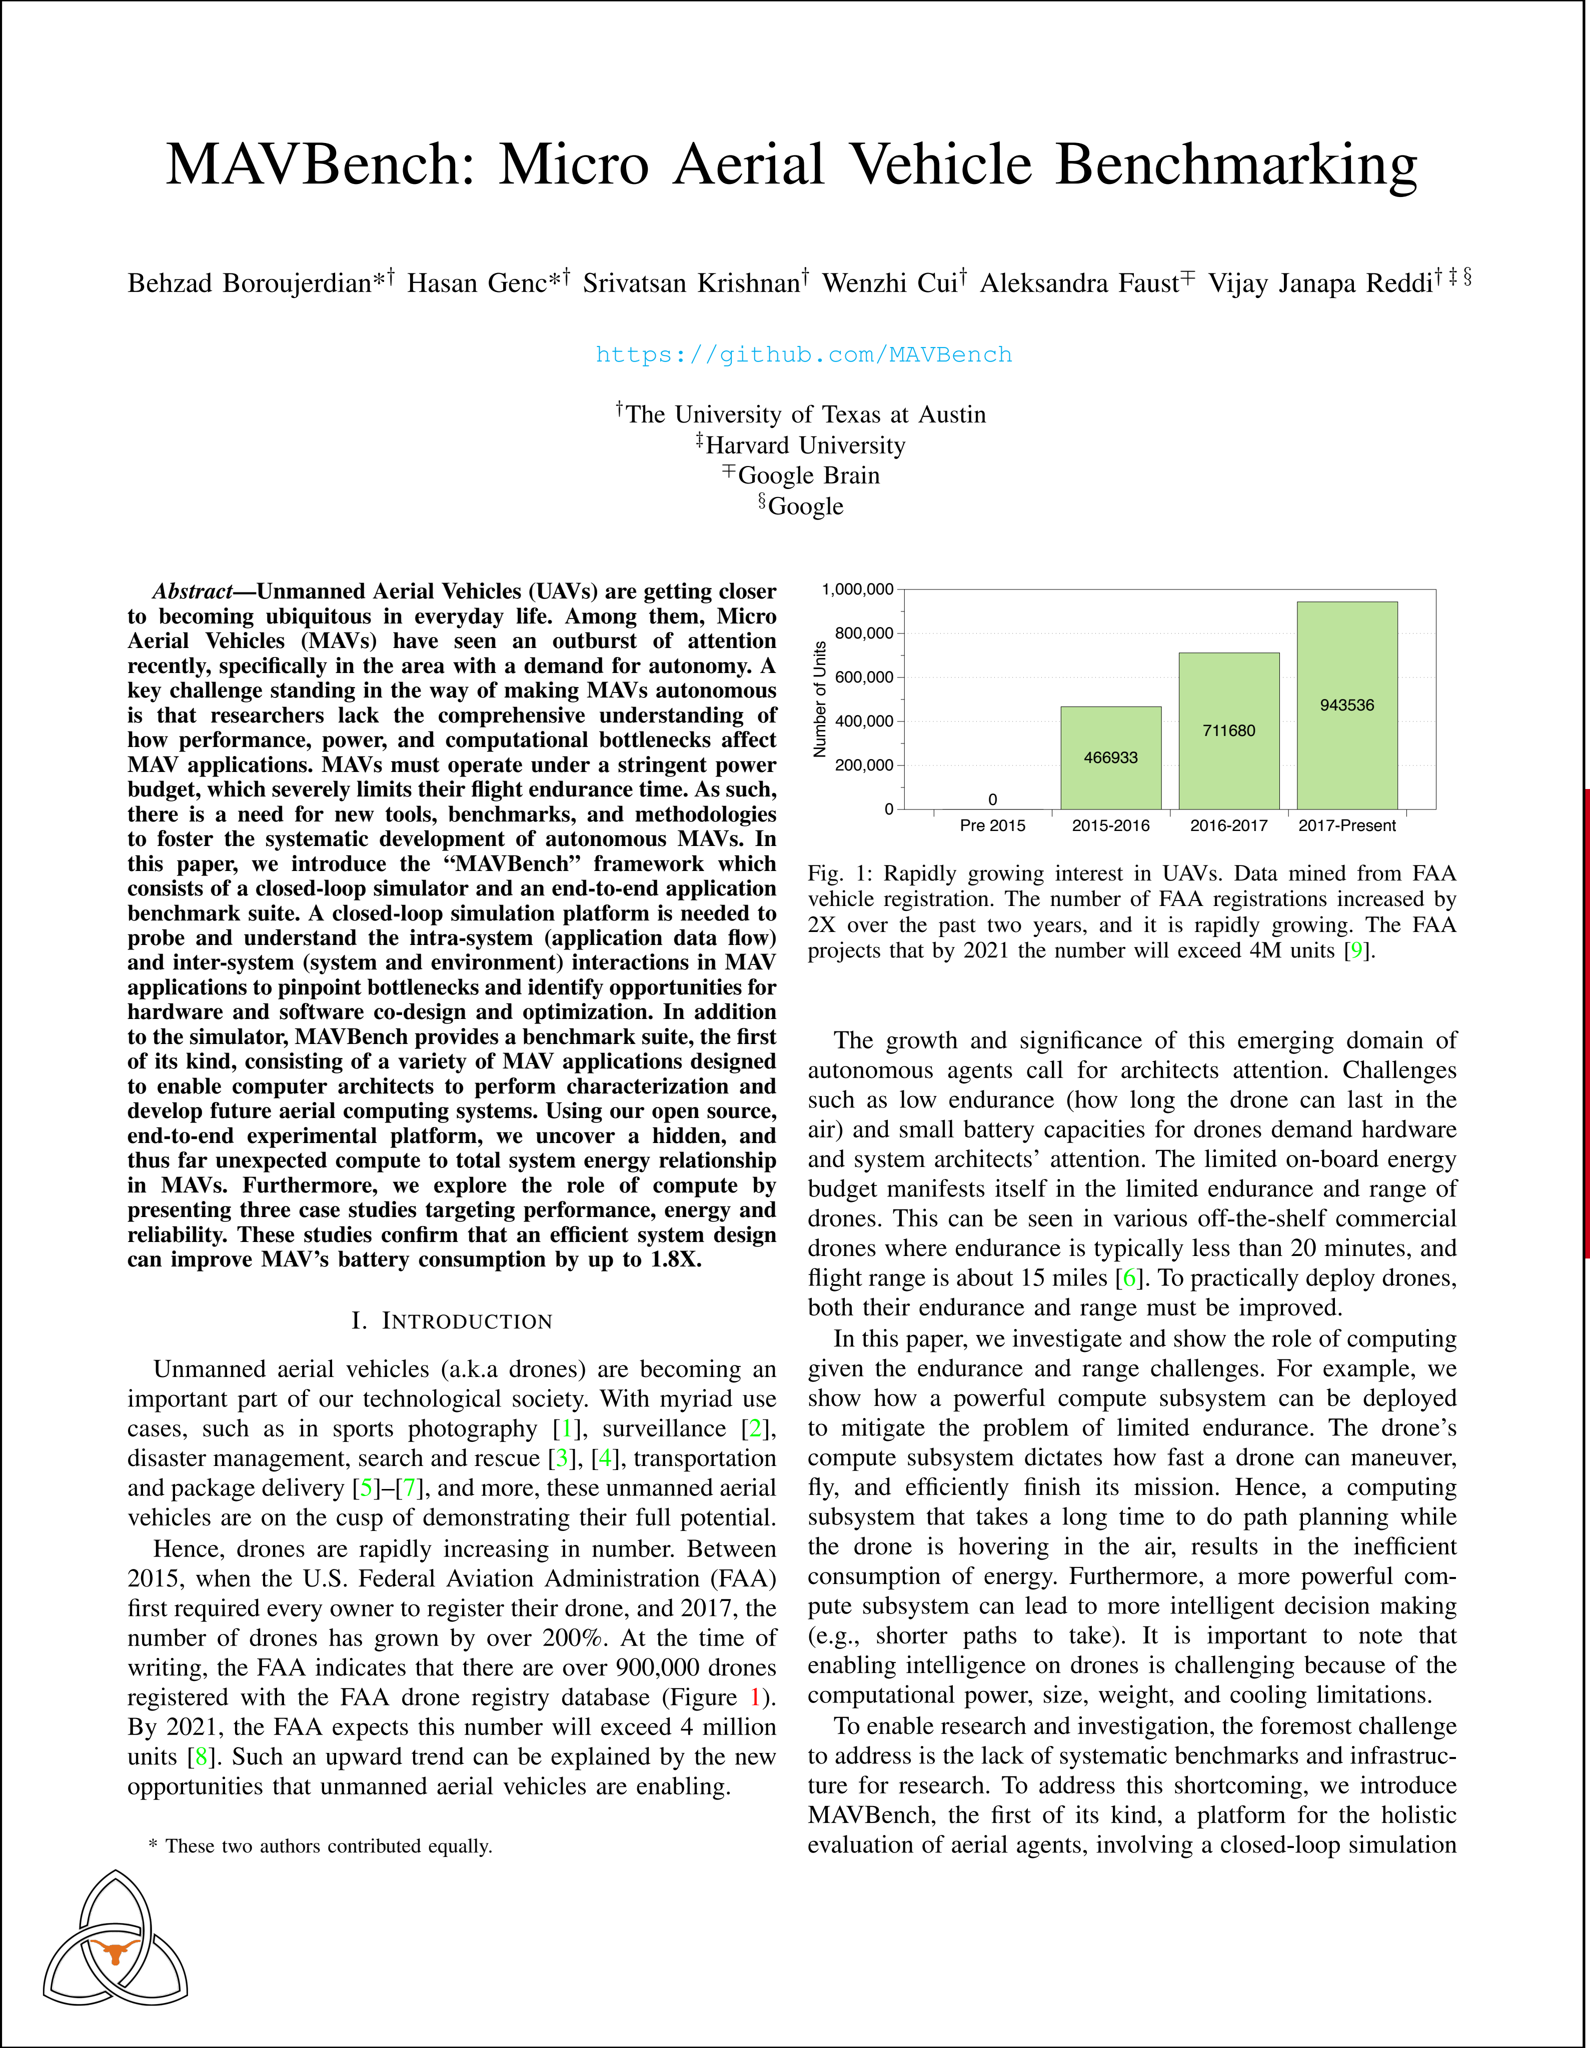

MAVBench: Micro Aerial Vehicle Benchmarking

Behzad Boroujerdian[*†] Hasan Genc[*†] Srivatsan Krishnan[†] Wenzhi Cui[†] Aleksandra Faust[∓] Vijay Janapa Reddi[† ‡ §]

https://github.com/MAVBench

[†]The University of Texas at Austin
[‡]Harvard University
[∓]Google Brain
[§]Google

Abstract—**Unmanned Aerial Vehicles (UAVs) are getting closer to becoming ubiquitous in everyday life. Among them, Micro Aerial Vehicles (MAVs) have seen an outburst of attention recently, specifically in the area with a demand for autonomy. A key challenge standing in the way of making MAVs autonomous is that researchers lack the comprehensive understanding of how performance, power, and computational bottlenecks affect MAV applications. MAVs must operate under a stringent power budget, which severely limits their flight endurance time. As such, there is a need for new tools, benchmarks, and methodologies to foster the systematic development of autonomous MAVs. In this paper, we introduce the "MAVBench" framework which consists of a closed-loop simulator and an end-to-end application benchmark suite. A closed-loop simulation platform is needed to probe and understand the intra-system (application data flow) and inter-system (system and environment) interactions in MAV applications to pinpoint bottlenecks and identify opportunities for hardware and software co-design and optimization. In addition to the simulator, MAVBench provides a benchmark suite, the first of its kind, consisting of a variety of MAV applications designed to enable computer architects to perform characterization and develop future aerial computing systems. Using our open source, end-to-end experimental platform, we uncover a hidden, and thus far unexpected compute to total system energy relationship in MAVs. Furthermore, we explore the role of compute by presenting three case studies targeting performance, energy and reliability. These studies confirm that an efficient system design can improve MAV's battery consumption by up to 1.8X.**

I. Introduction

Unmanned aerial vehicles (a.k.a drones) are becoming an important part of our technological society. With myriad use cases, such as in sports photography [1], surveillance [2], disaster management, search and rescue [3], [4], transportation and package delivery [5]–[7], and more, these unmanned aerial vehicles are on the cusp of demonstrating their full potential.

Hence, drones are rapidly increasing in number. Between 2015, when the U.S. Federal Aviation Administration (FAA) first required every owner to register their drone, and 2017, the number of drones has grown by over 200%. At the time of writing, the FAA indicates that there are over 900,000 drones registered with the FAA drone registry database (Figure 1). By 2021, the FAA expects this number will exceed 4 million units [8]. Such an upward trend can be explained by the new opportunities that unmanned aerial vehicles are enabling.

Fig. 1: Rapidly growing interest in UAVs. Data mined from FAA vehicle registration. The number of FAA registrations increased by 2X over the past two years, and it is rapidly growing. The FAA projects that by 2021 the number will exceed 4M units [9].

The growth and significance of this emerging domain of autonomous agents call for architects attention. Challenges such as low endurance (how long the drone can last in the air) and small battery capacities for drones demand hardware and system architects' attention. The limited on-board energy budget manifests itself in the limited endurance and range of drones. This can be seen in various off-the-shelf commercial drones where endurance is typically less than 20 minutes, and flight range is about 15 miles [6]. To practically deploy drones, both their endurance and range must be improved.

In this paper, we investigate and show the role of computing given the endurance and range challenges. For example, we show how a powerful compute subsystem can be deployed to mitigate the problem of limited endurance. The drone's compute subsystem dictates how fast a drone can maneuver, fly, and efficiently finish its mission. Hence, a computing subsystem that takes a long time to do path planning while the drone is hovering in the air, results in the inefficient consumption of energy. Furthermore, a more powerful compute subsystem can lead to more intelligent decision making (e.g., shorter paths to take). It is important to note that enabling intelligence on drones is challenging because of the computational power, size, weight, and cooling limitations.

To enable research and investigation, the foremost challenge to address is the lack of systematic benchmarks and infrastructure for research. To address this shortcoming, we introduce MAVBench, the first of its kind, a platform for the holistic evaluation of aerial agents, involving a closed-loop simulation

[*] These two authors contributed equally.

International Symposium on Microarchitecture (MICRO-51)

https://github.com/MAVBench

The AI Revolution

The Traditional “PPC” Pipeline

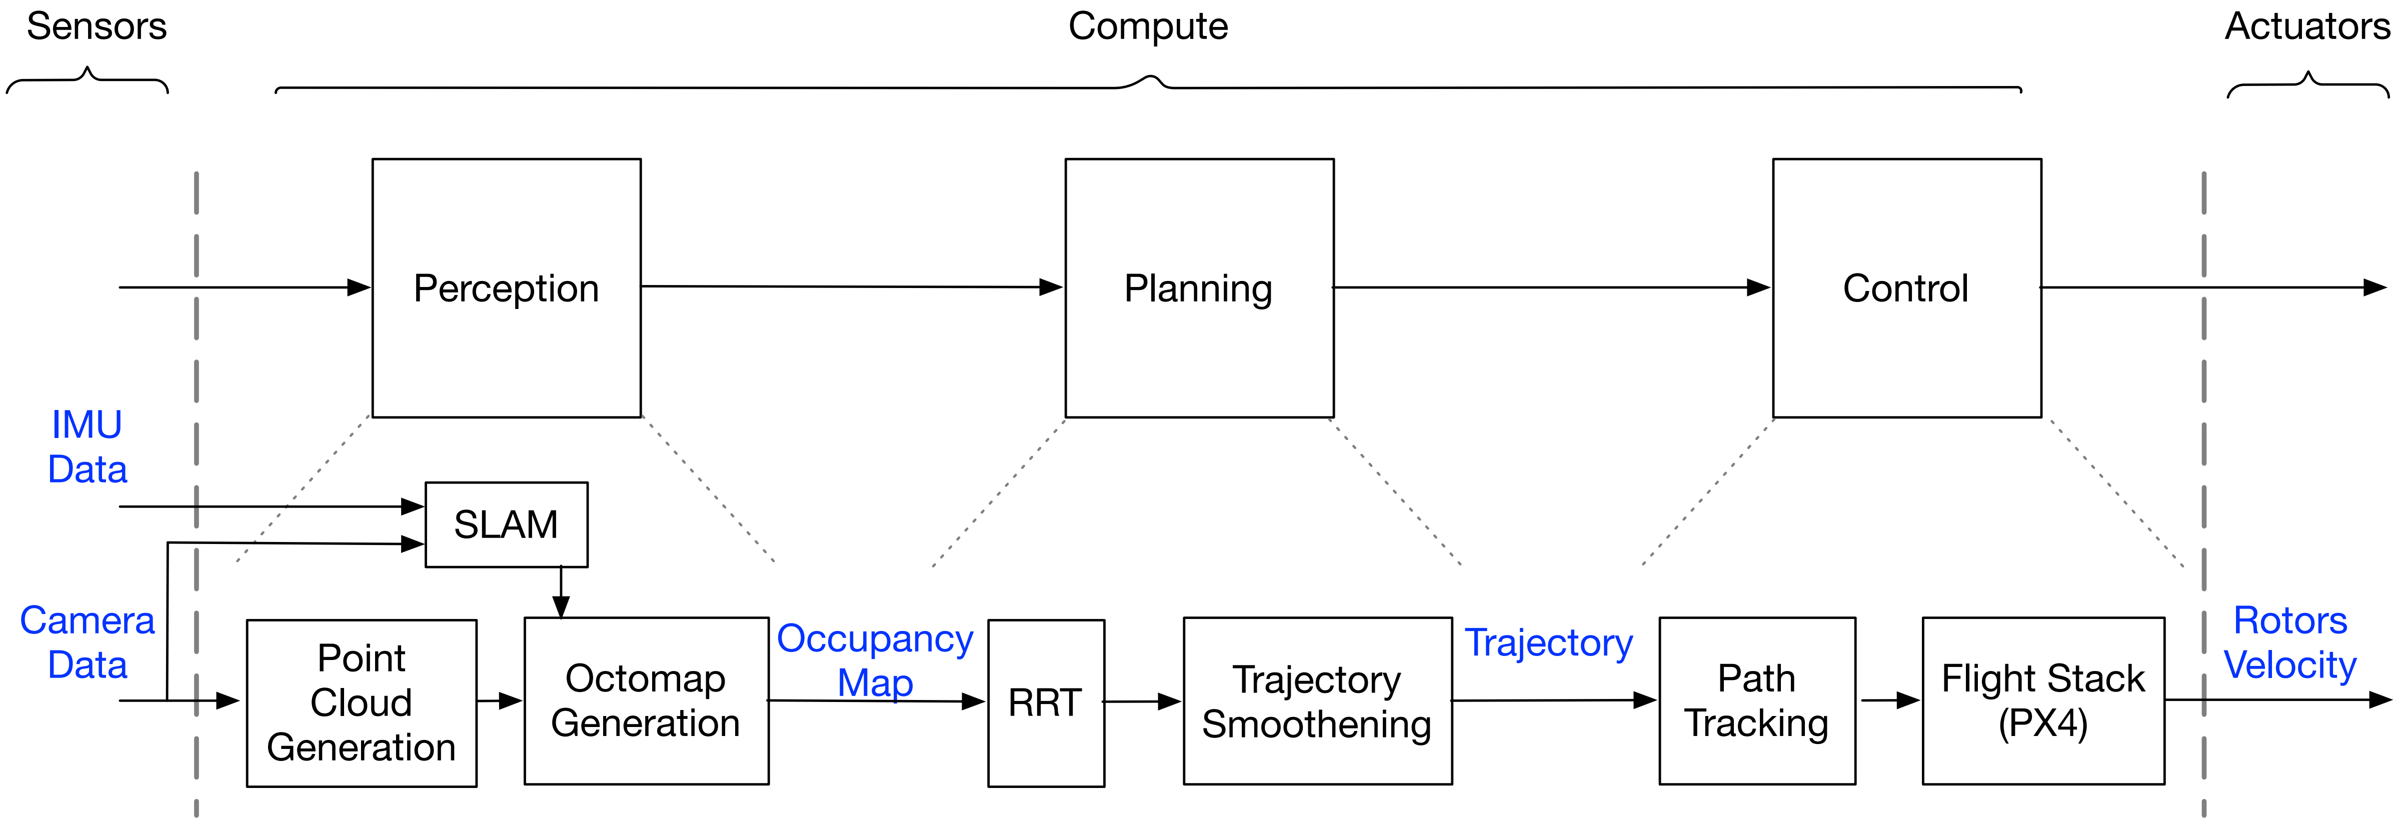

Deep Neural Nets

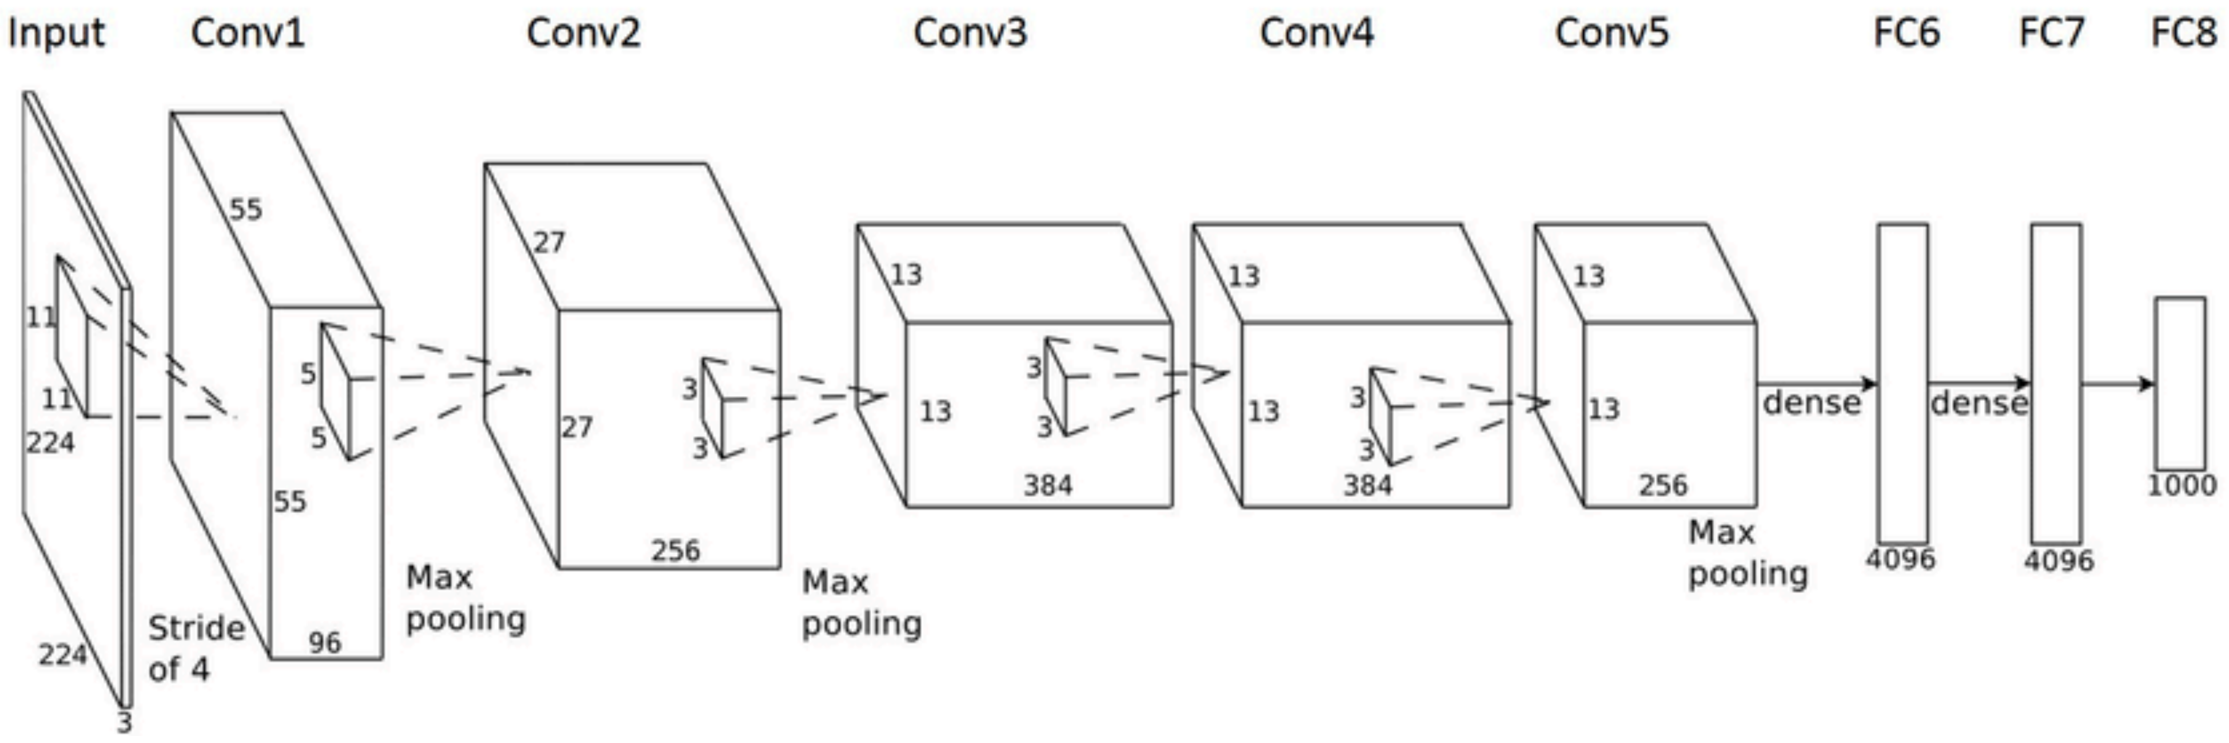

Autonomous Navigation

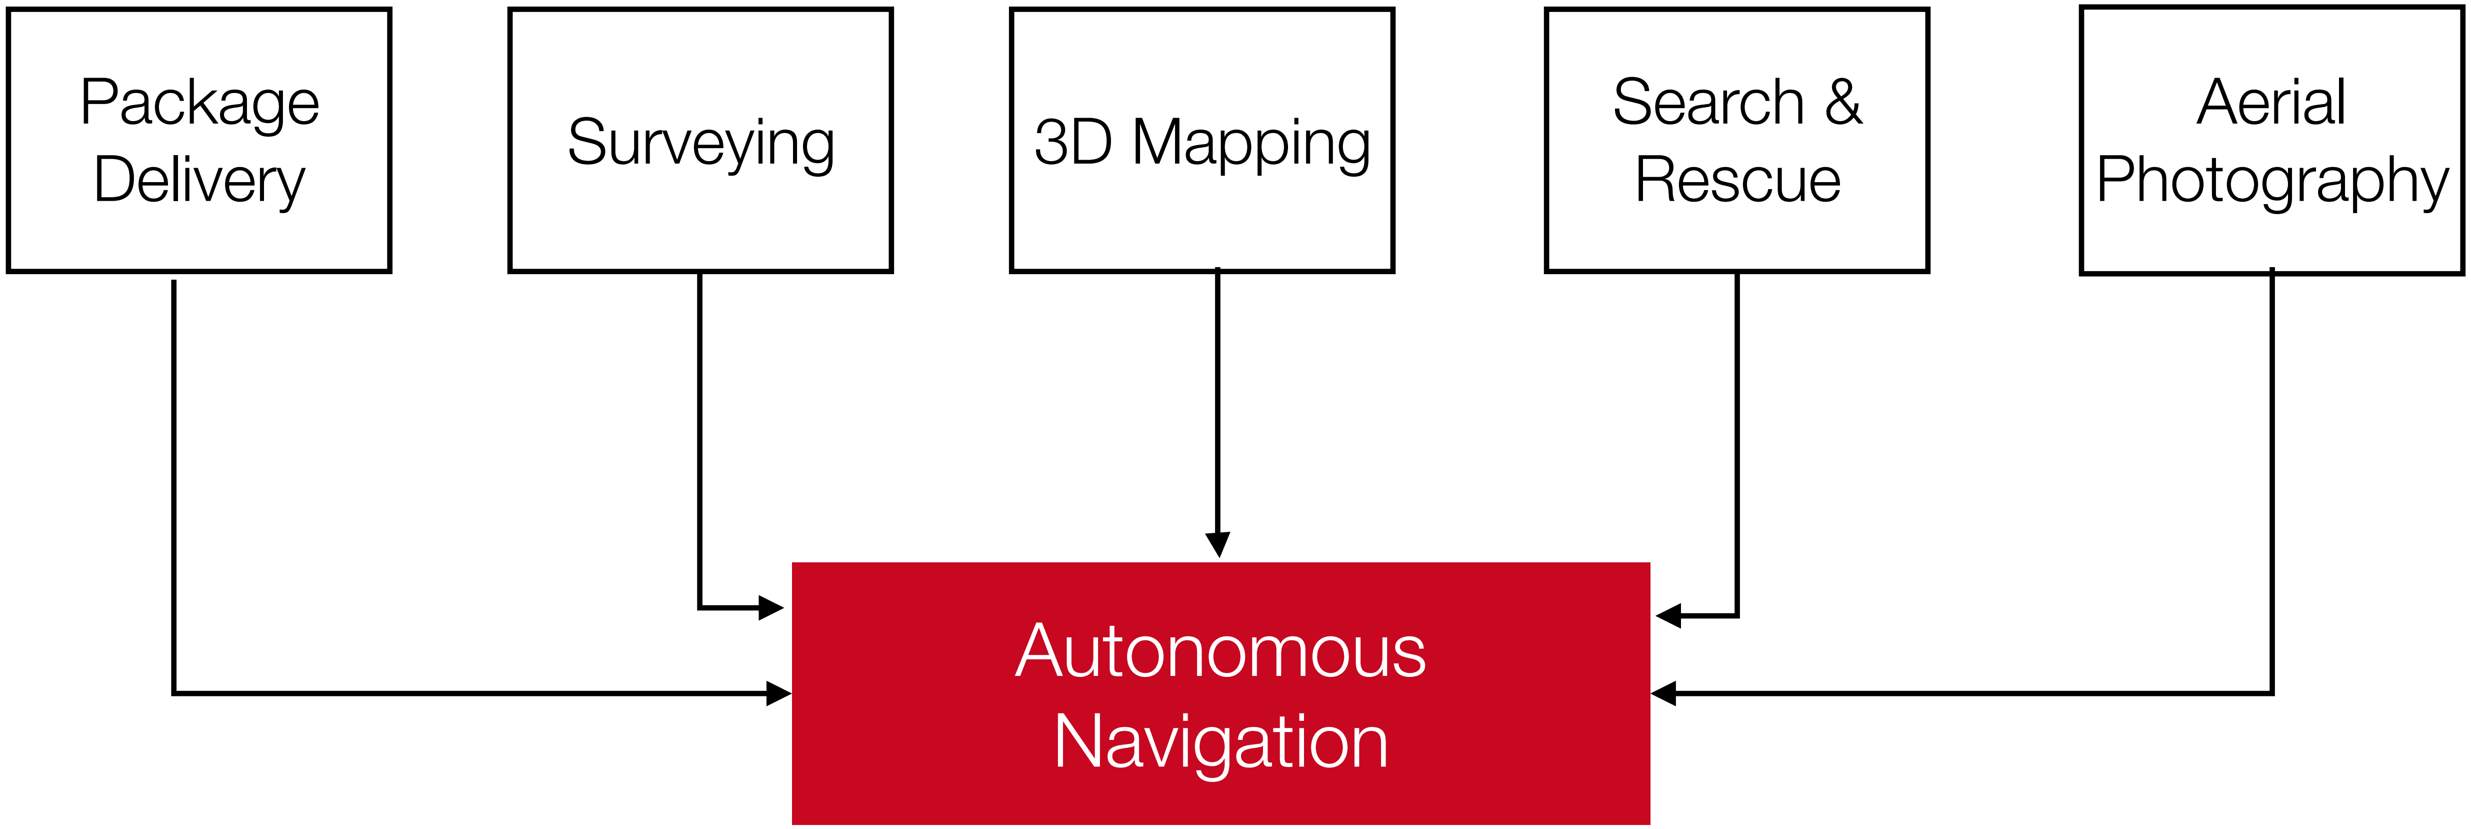

Autonomous Navigation

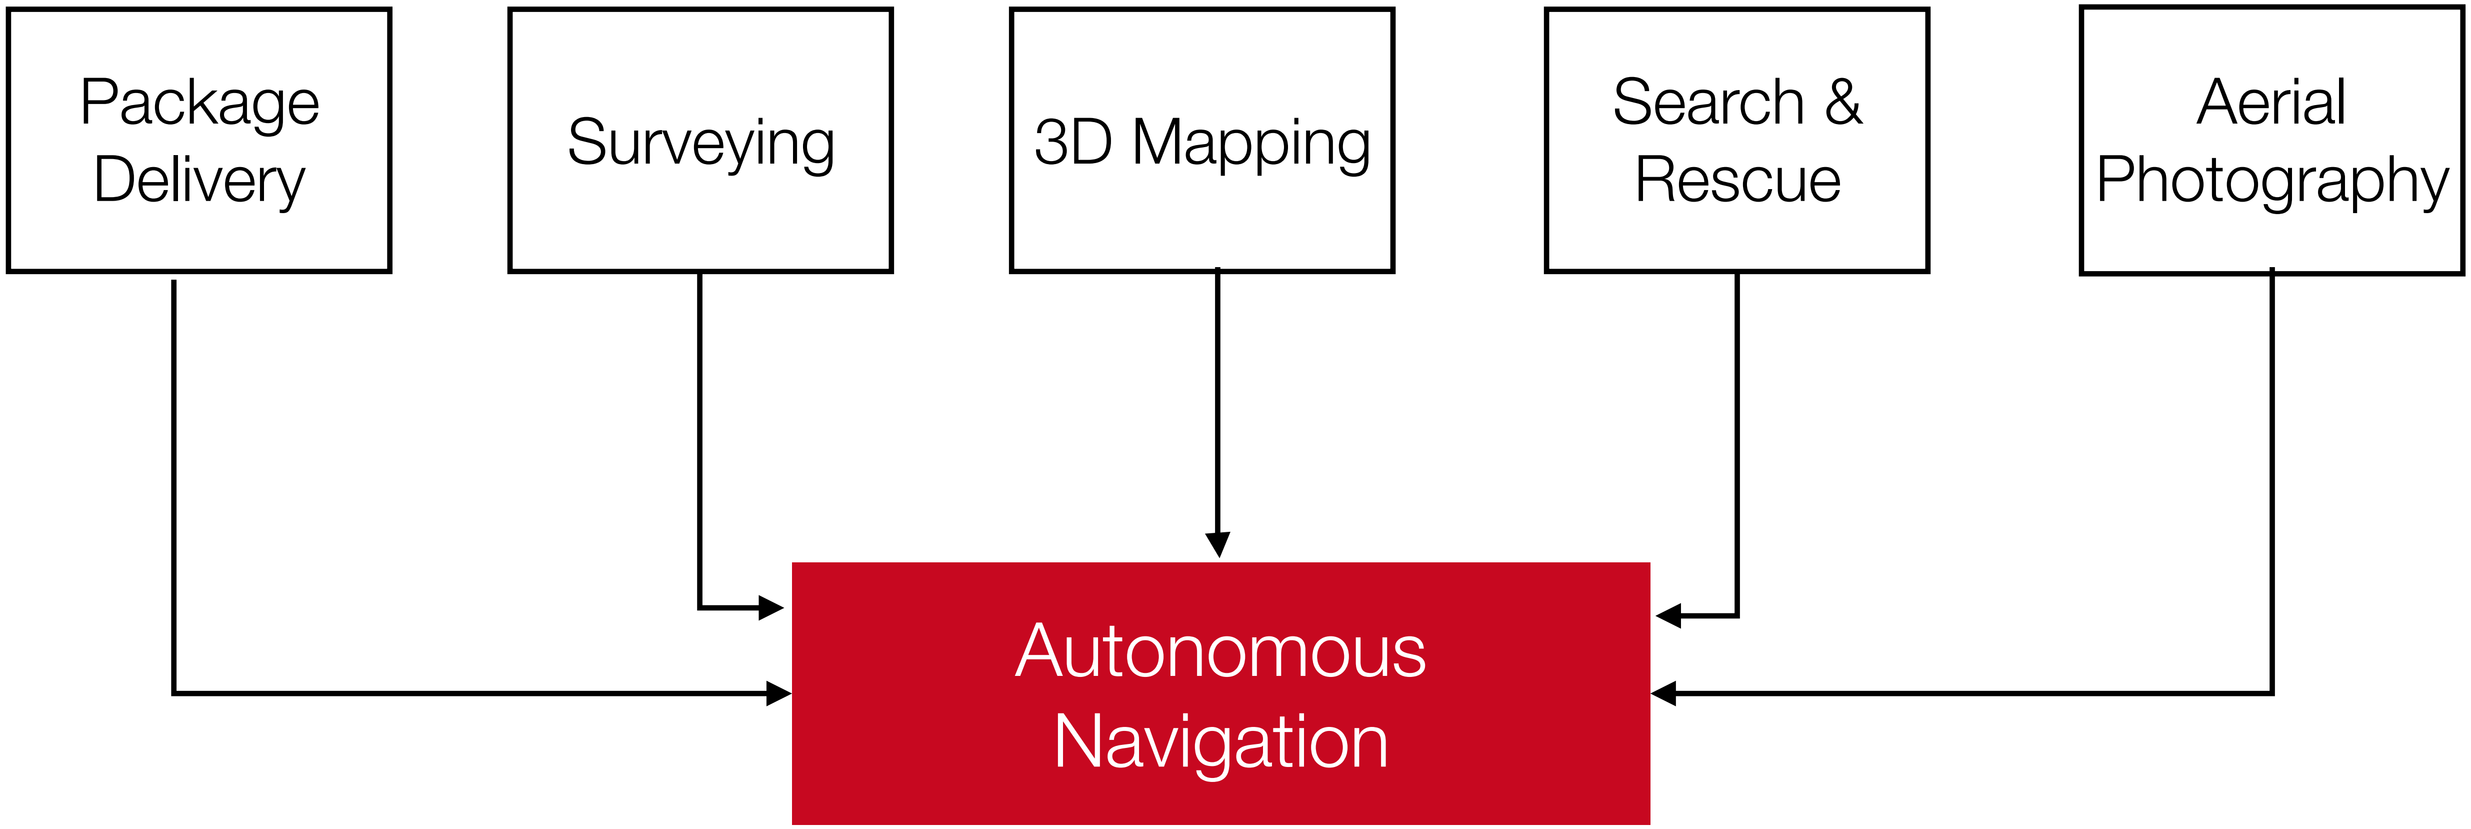

Autonomous navigation is a common kernel across multiple application

Autonomous Navigation - Two Paradigms

Autonomous Navigation

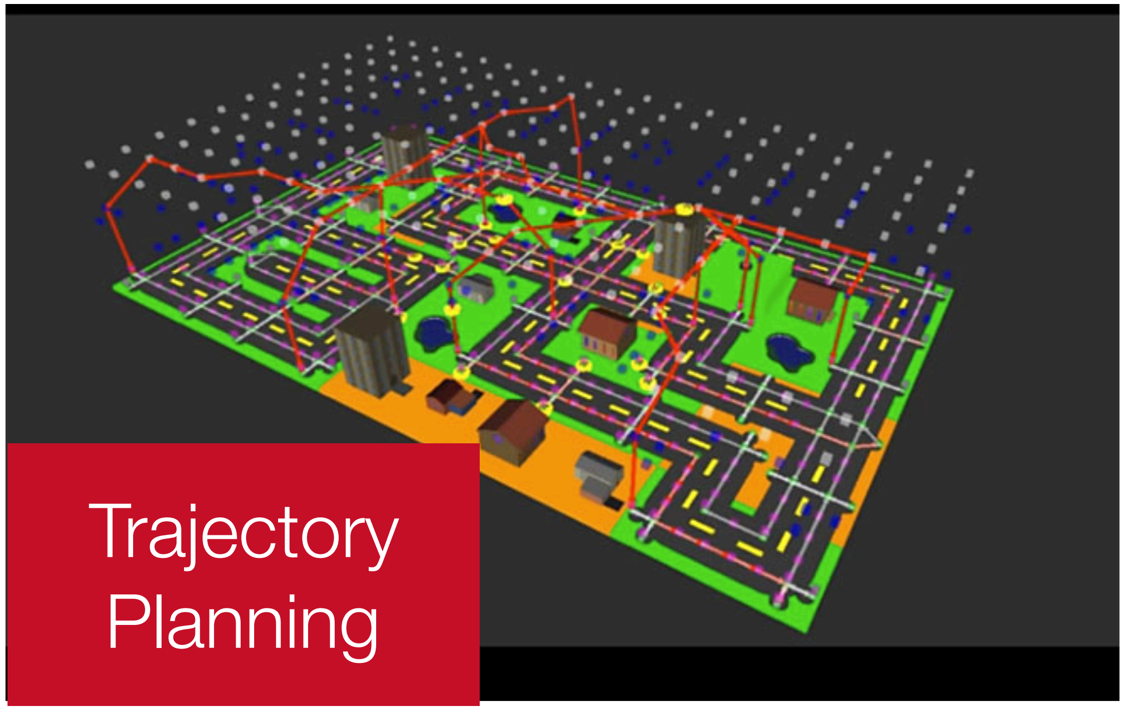

Trajectory Planning

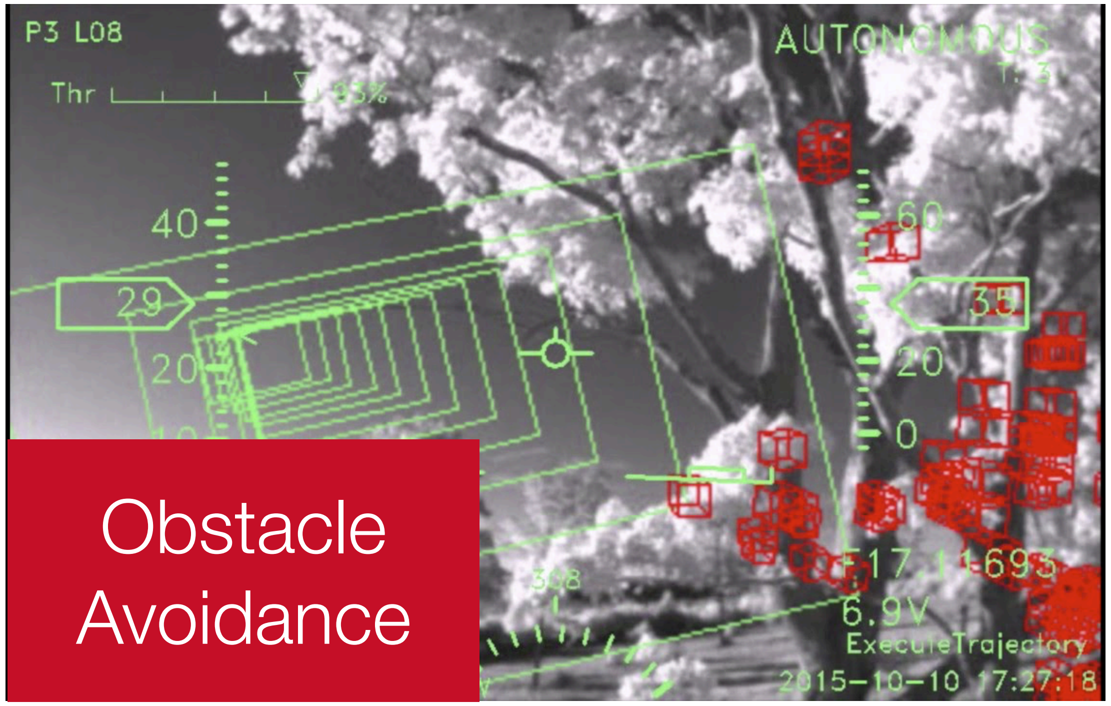

Obstacle Avoidance

PPC

Pixel2Control

AirLearning Infrastructure

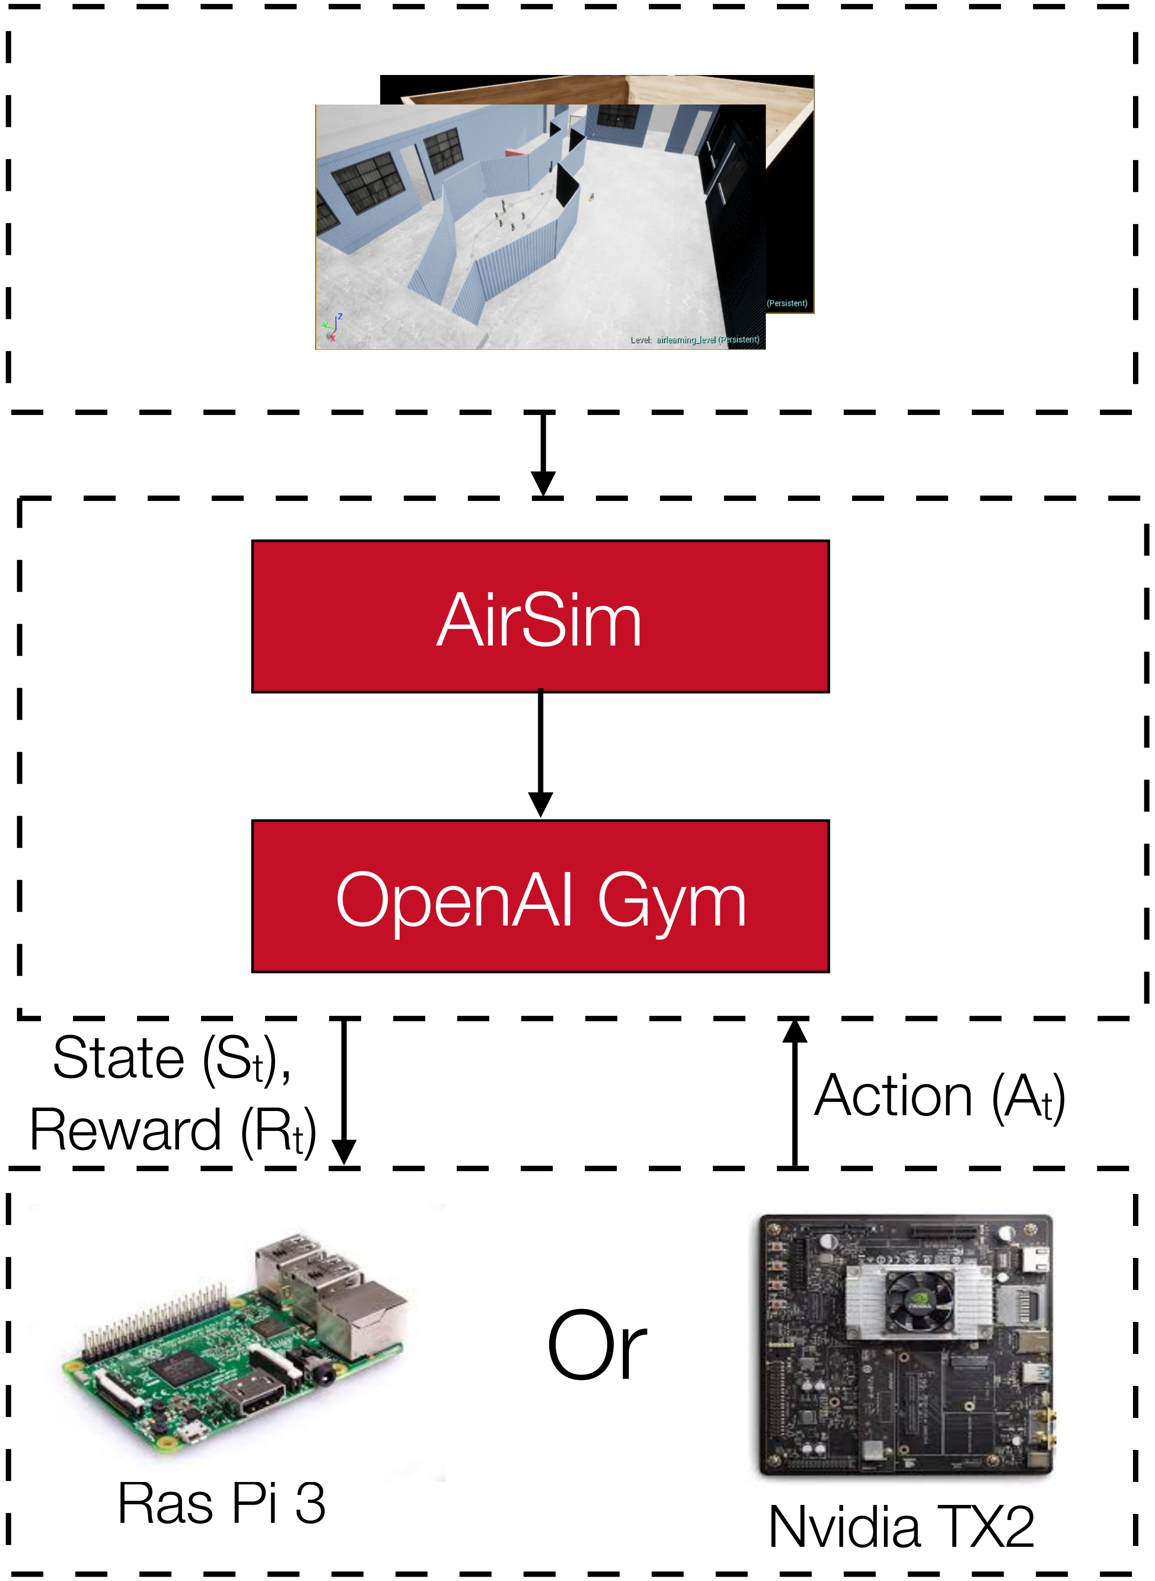

AirSim

OpenAI Gym

State (S_t),
Reward (R_t)

Action (A_t)

Ras Pi 3

Or

Nvidia TX2

AirLearning Infrastructure

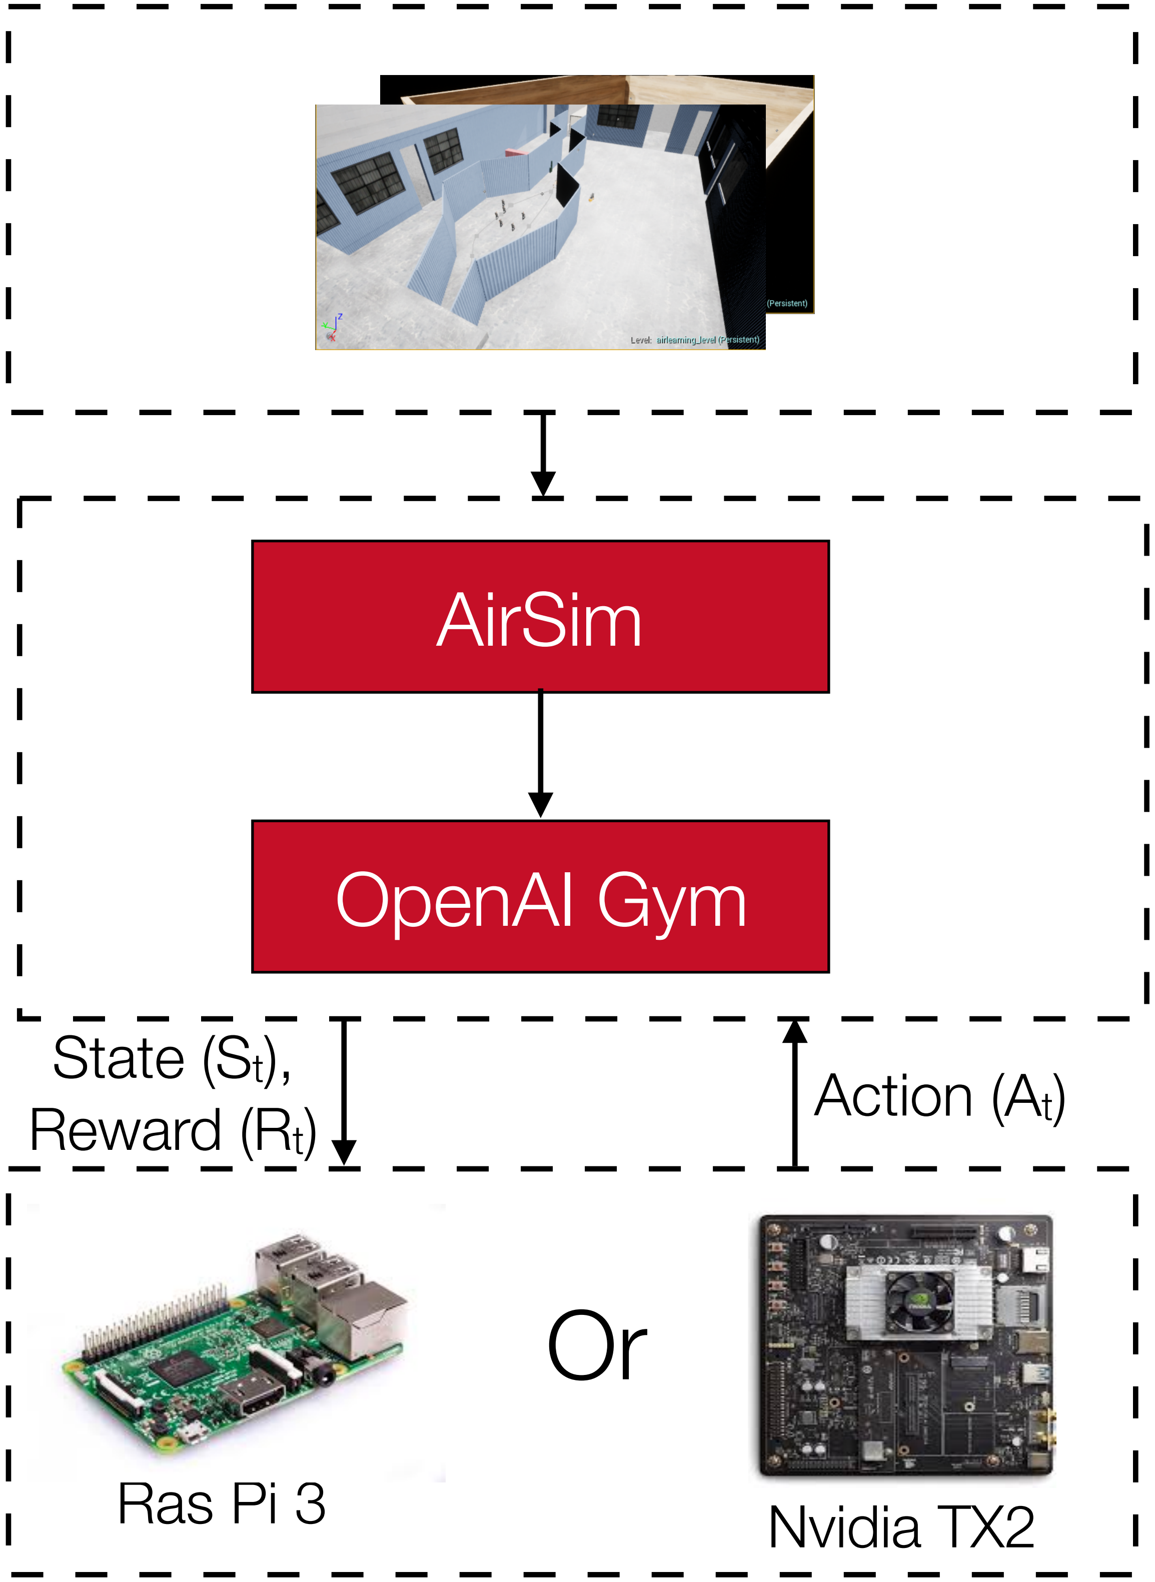

AirSim

OpenAI Gym

State (S_t),
Reward (R_t)

Action (A_t)

Ras Pi 3

Or

Nvidia TX2

Simulation Environment

DRL Framework

Compute System

AirLearning Infrastructure

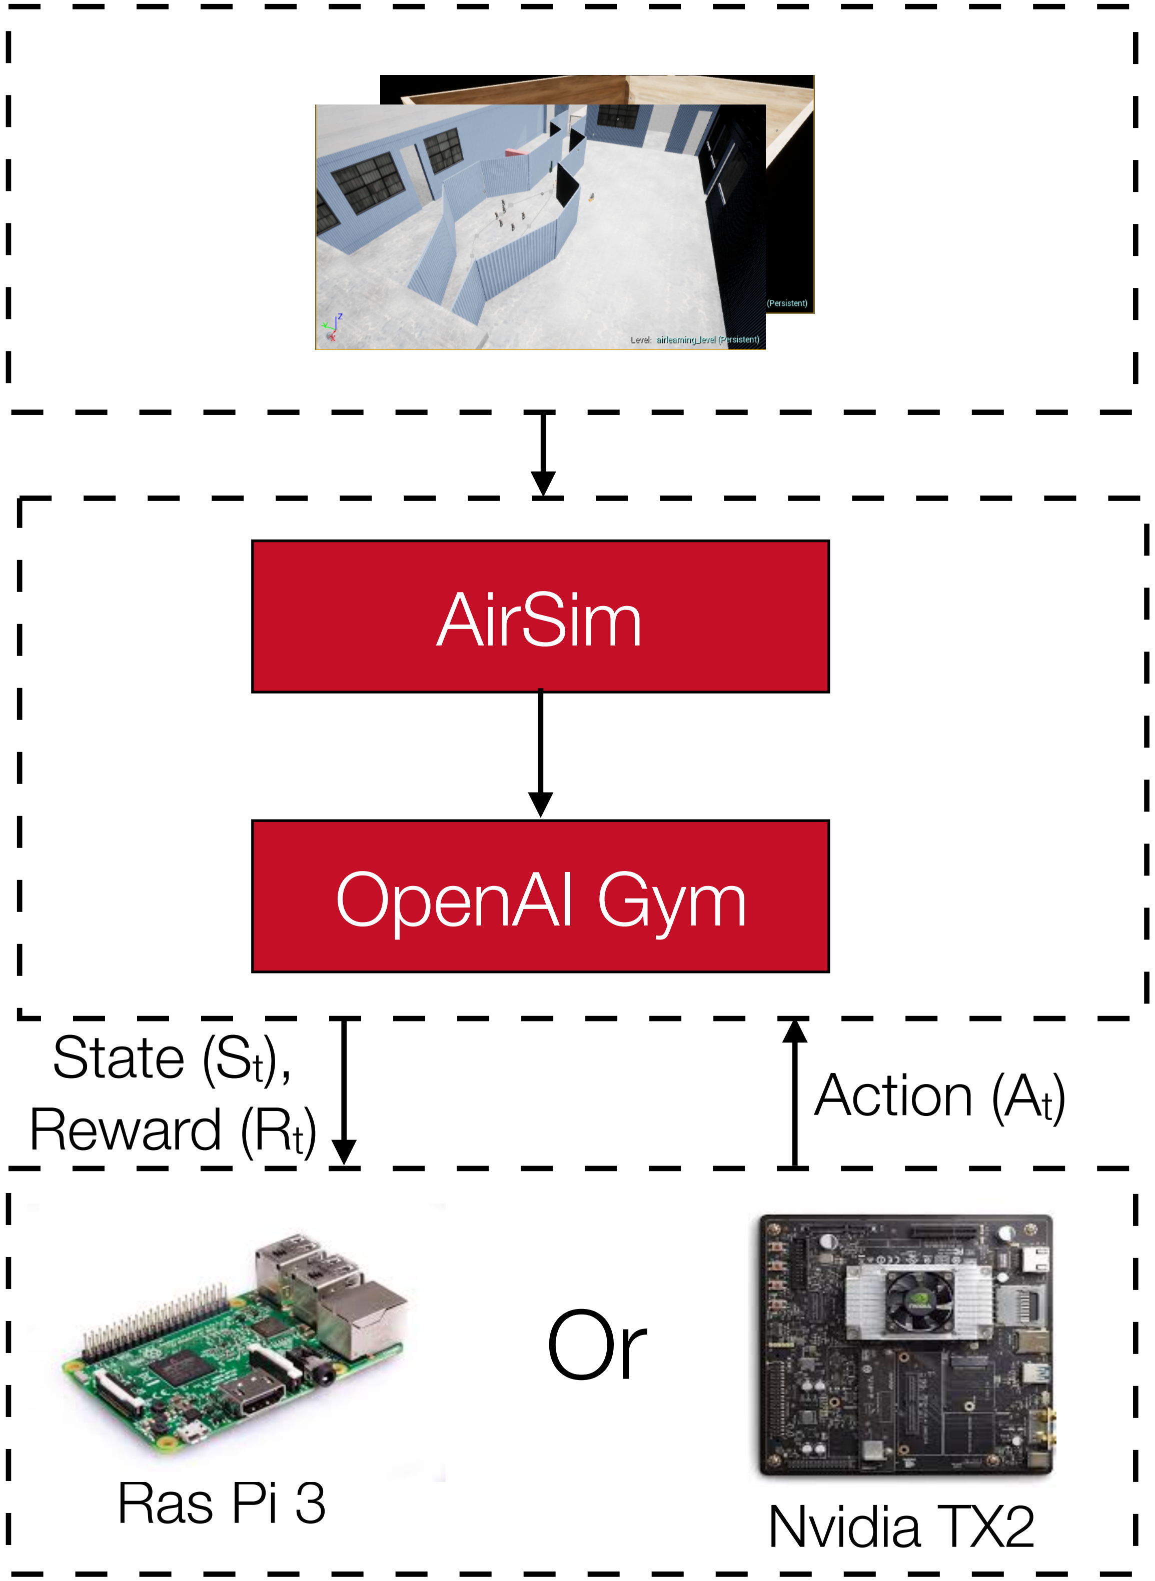

AirSim

OpenAI Gym

State (S_t),
Reward (R_t)

Action (A_t)

Ras Pi 3

Or

Nvidia TX2

Simulation Environment

DRL Framework

Compute System

Simulation Environment

Constrained Space

LIGHTING NEEDS TO BE REBUILT (24 unbuilt object(s))
'DisableAllScreenMessages' to suppress

Sample 1

https://youtu.be/DMCj0bwHmds

LIGHTING NEEDS TO BE REBUILT (10 unbuilt object(s))
'DisableAllScreenMessages' to suppress

Sample 1*

https://youtu.be/JRZ1IxHMWnc

LIGHTING NEEDS TO BE REBUILT

Sample 2

https://youtu.be/Yi2aaSSqZ0Y

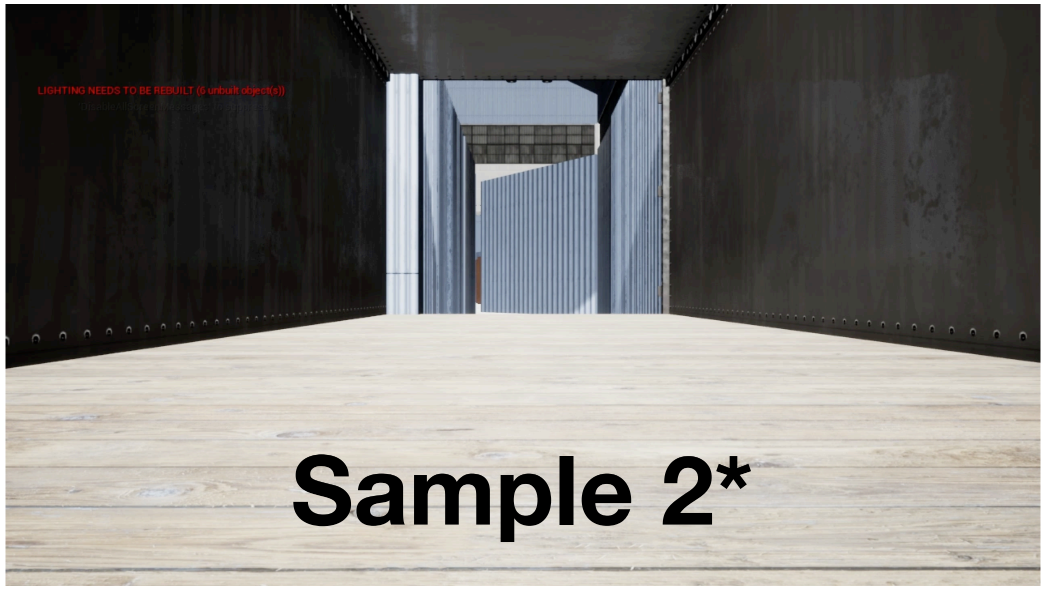

Sample 2*

https://youtu.be/Oh3Hj0HsG50

Increased Obstacle Density

Simulation Environment

Constrained Space

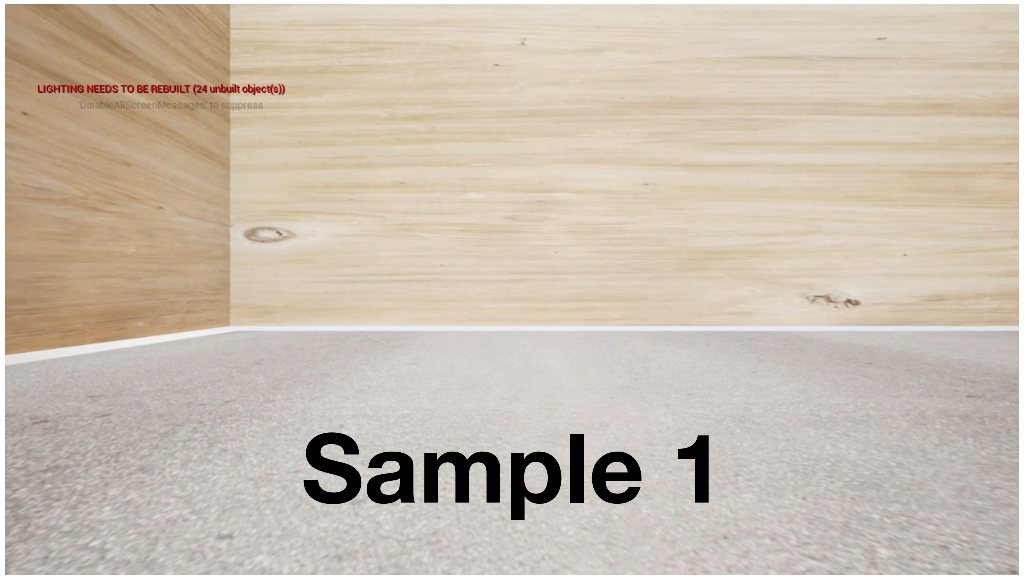
LIGHTING NEEDS TO BE REBUILT (24 unbuilt object(s))
'DisableAllScreenMessages' to suppress
Sample 1

https://youtu.be/DMCj0bwHmds

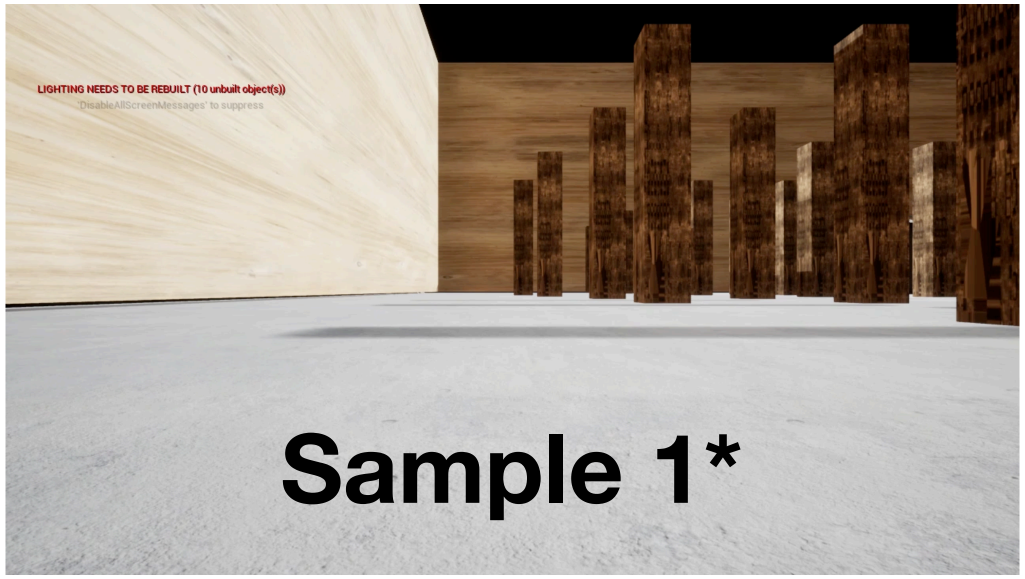
LIGHTING NEEDS TO BE REBUILT (10 unbuilt object(s))
'DisableAllScreenMessages' to suppress
Sample 1*

https://youtu.be/JRZ1IxHMWnc

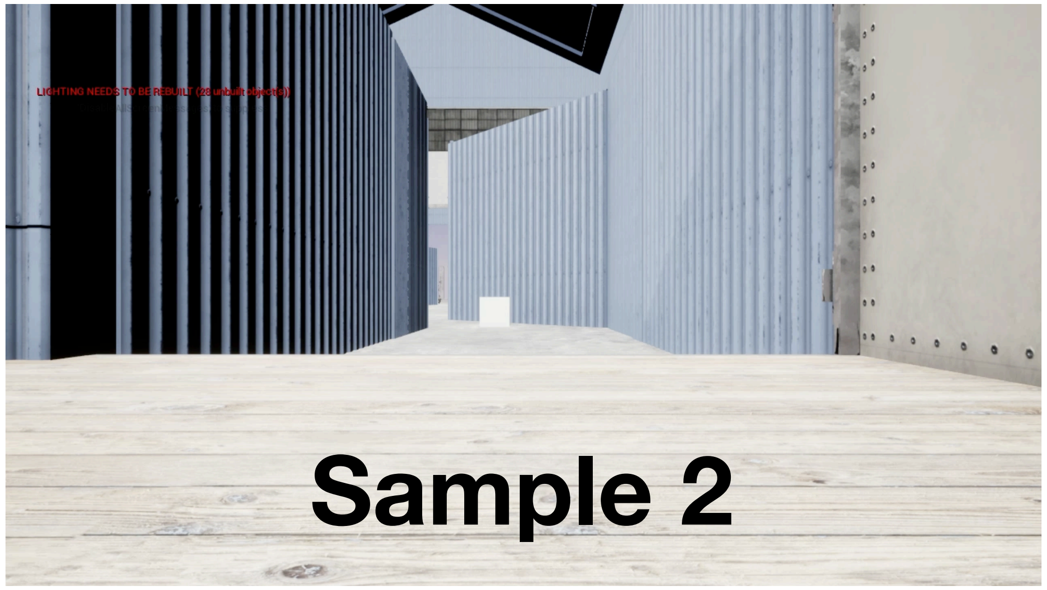
LIGHTING NEEDS TO BE REBUILT (28 unbuilt object(s))
'DisableAllScreenMessages' to suppress
Sample 2

https://youtu.be/Yi2aaSSqZ0Y

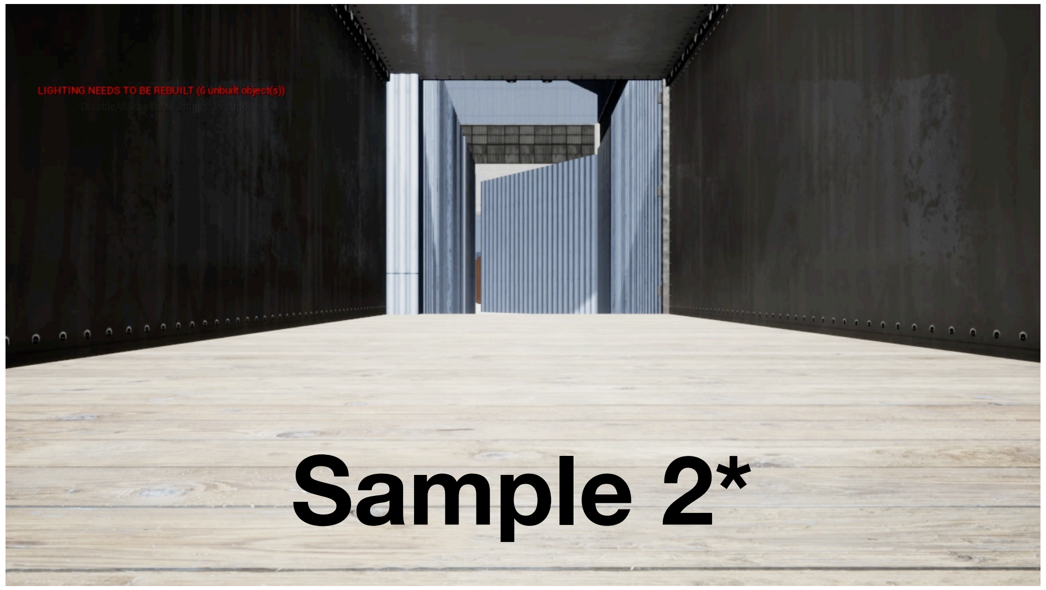
LIGHTING NEEDS TO BE REBUILT (6 unbuilt object(s))
'DisableAllScreenMessages' to suppress
Sample 2*

https://youtu.be/Oh3Hj0HsG50

Increased Obstacle Density

AirLearning Infrastructure

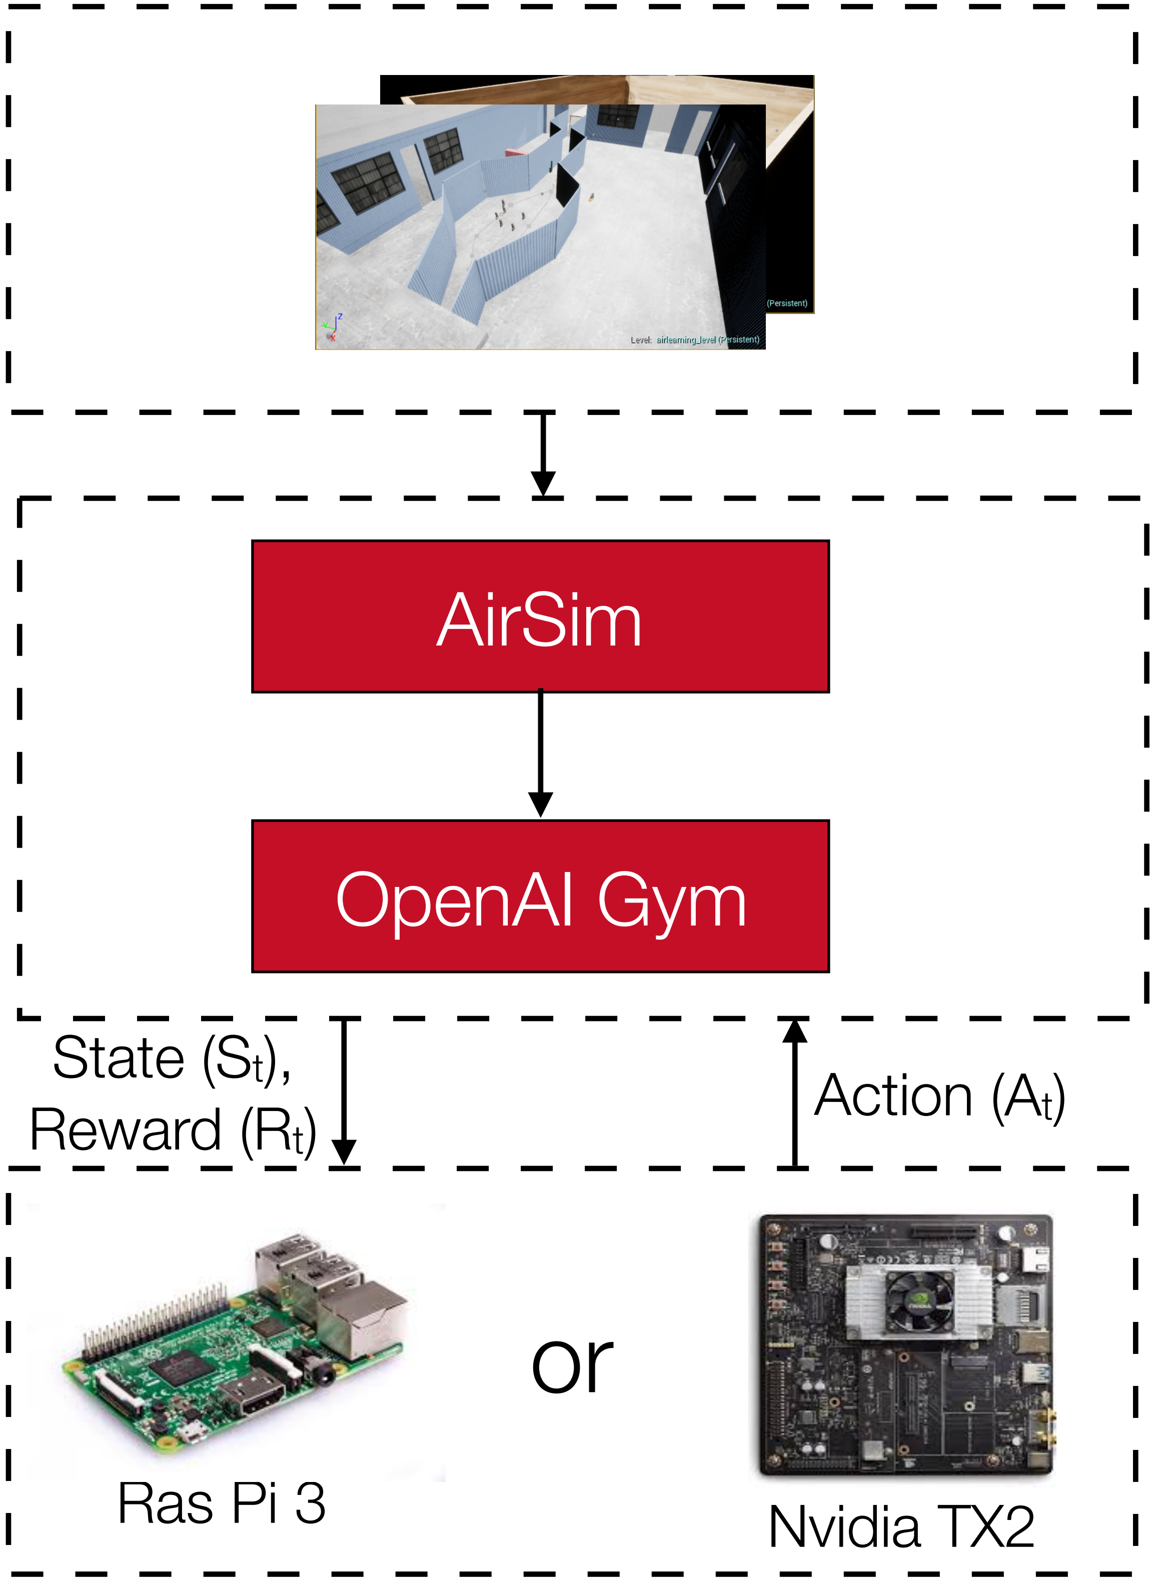

Simulation Environment

DRL Framework

Compute System

Deep RL Framework

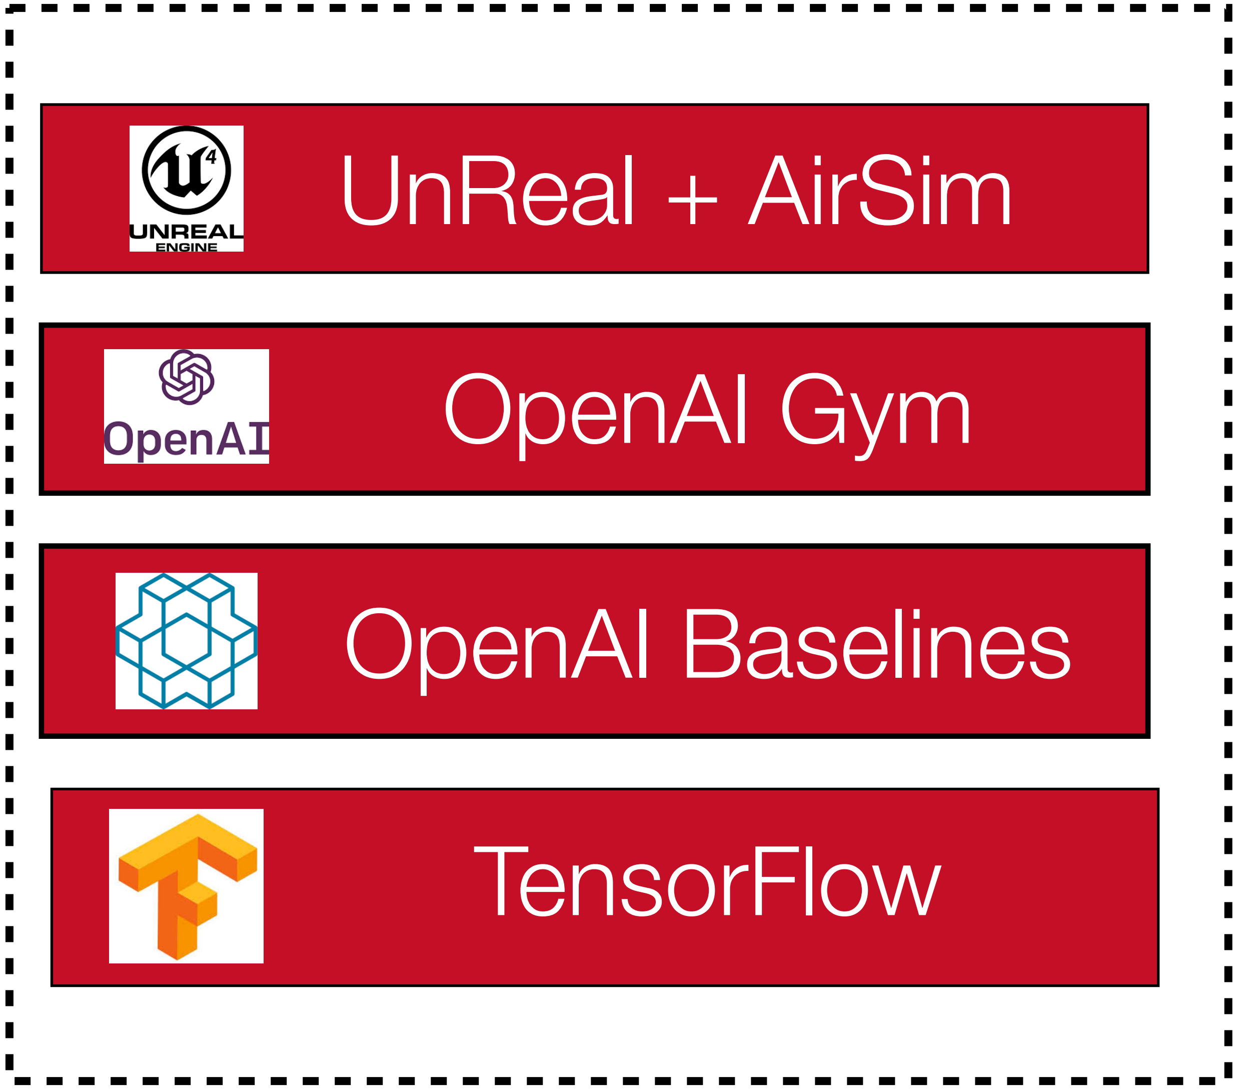

UnReal + AirSim	Depth Image + RGB + IMU
OpenAI Gym	Reinforcement Learning
OpenAI Baselines	Algorithms (DQN, DDQN, A3C)
TensorFlow	Machine Learning Backend

Reinforcement learning based Drone Control

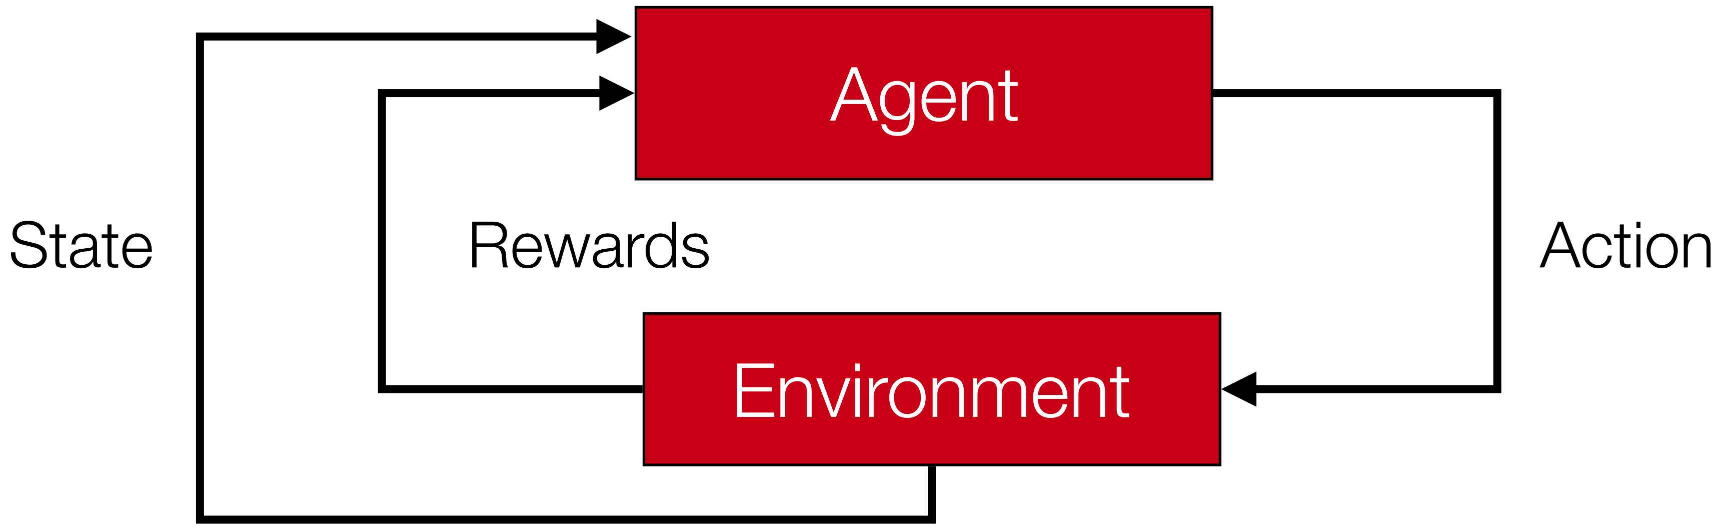

Agent

Environment

State

Rewards

Action

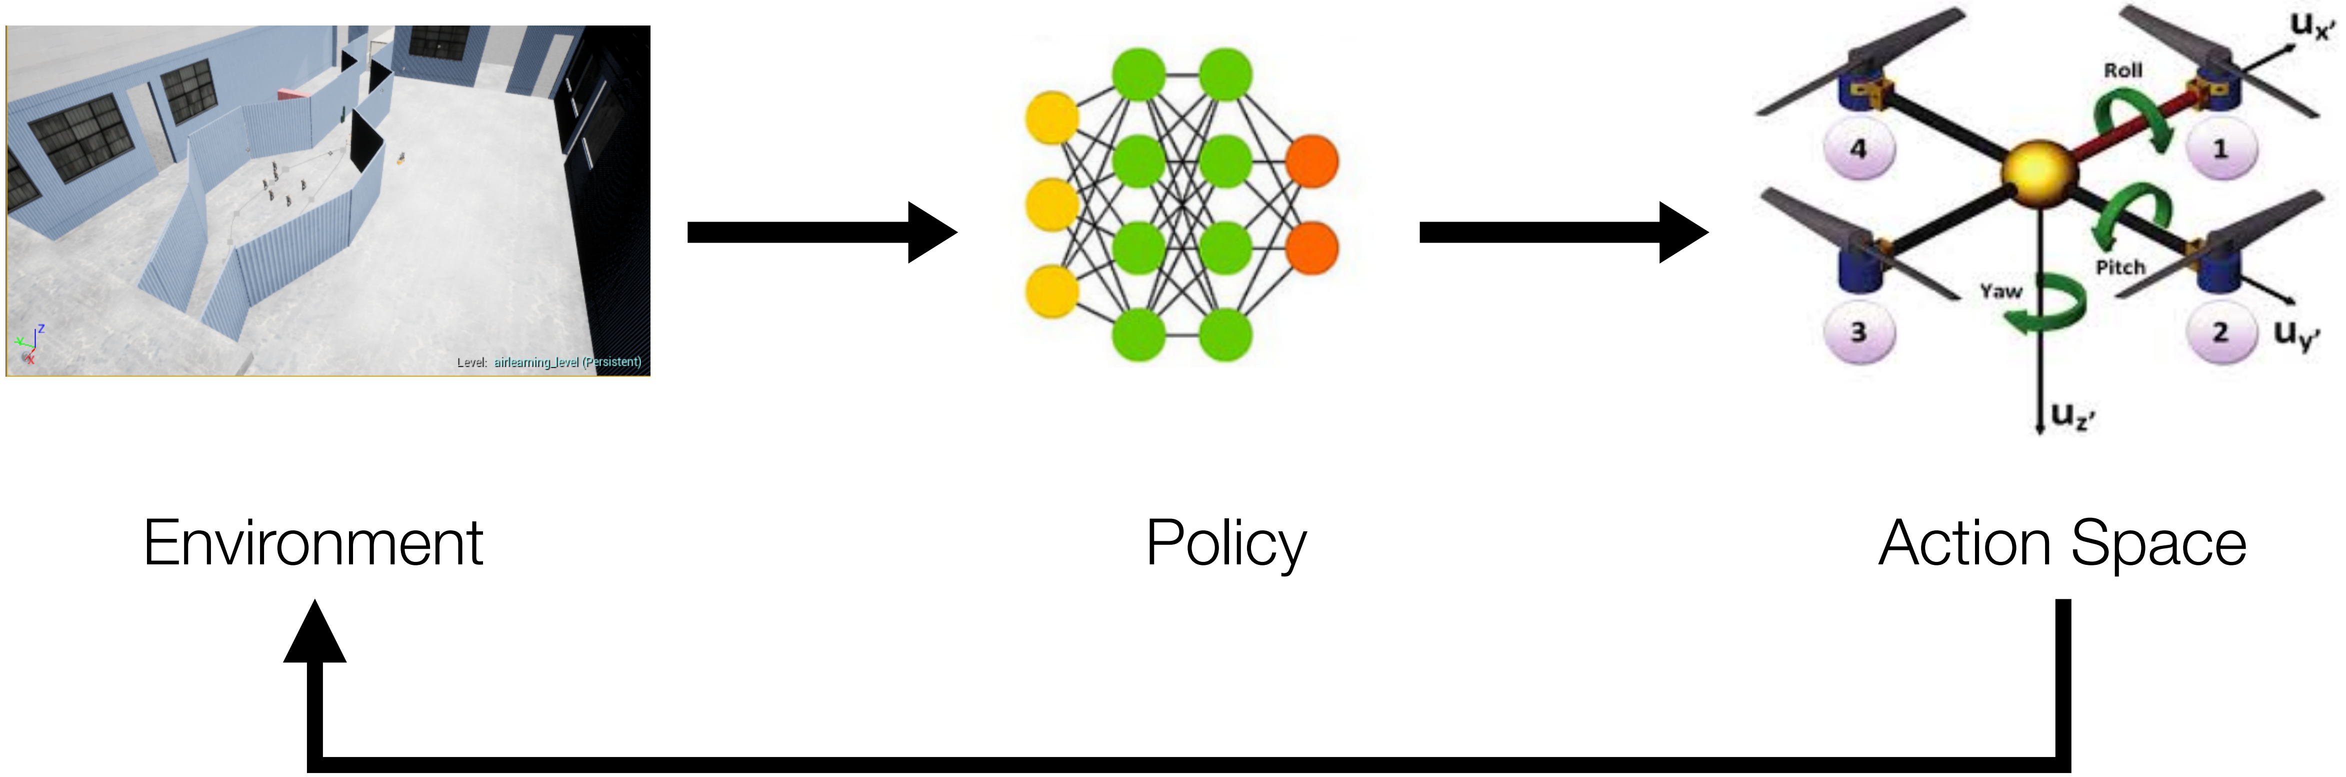

Environment

Policy

Action Space

AirLearning Demonstration

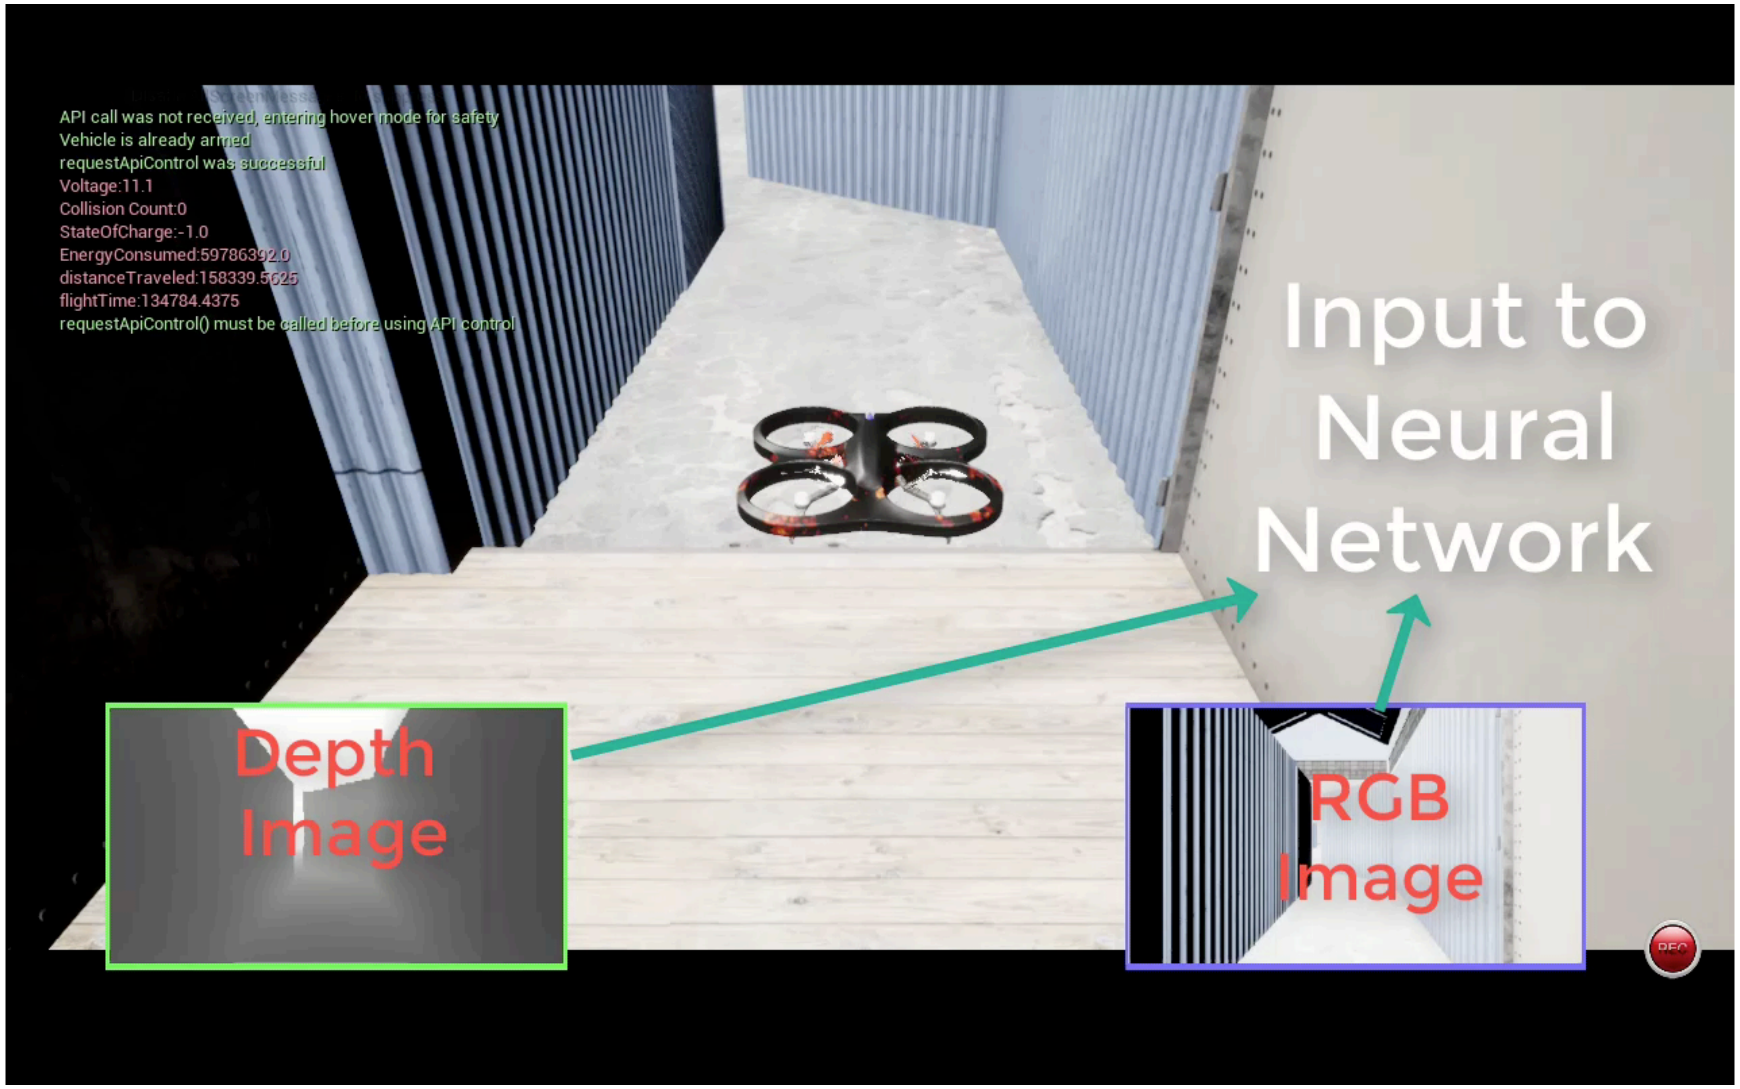

```
API call was not received, entering hover mode for safety
Vehicle is already armed
requestApiControl was successful
Voltage:11.1
Collision Count:0
StateOfCharge:-1.0
EnergyConsumed:59786392.0
distanceTraveled:158339.5625
flightTime:134784.4375
requestApiControl() must be called before using API control
```

Demonstrates that the UAV has learned to navigate in a narrow passageway

https://www.youtube.com/watch?v=RlrkjnmXEAs&feature=youtu.be

AirLearning Demonstration

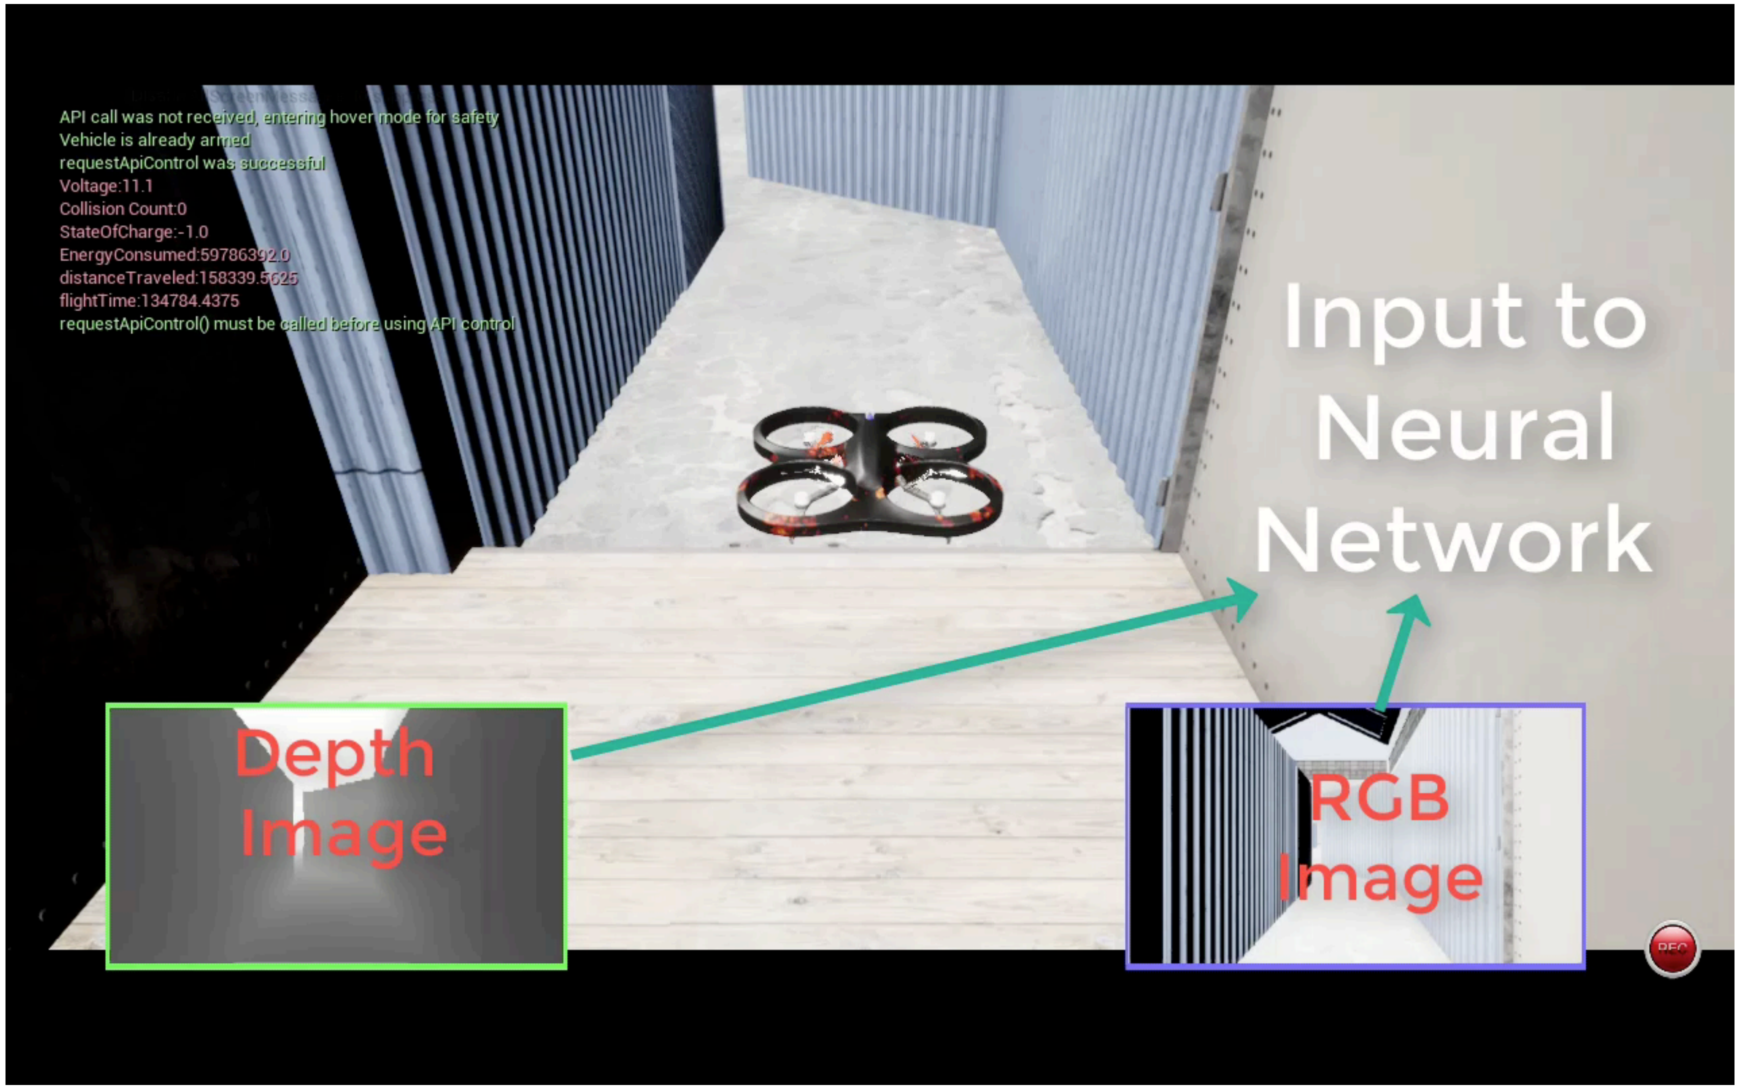

```
API call was not received, entering hover mode for safety
Vehicle is already armed
requestApiControl was successful
Voltage:11.1
Collision Count:0
StateOfCharge:-1.0
EnergyConsumed:59786392.0
distanceTraveled:158339.5625
flightTime:134784.4375
requestApiControl() must be called before using API control
```

Demonstrates that the UAV has learned to navigate in a narrow passageway

https://www.youtube.com/watch?v=RlrkjnmXEAs&feature=youtu.be

AirLearning Infrastructure

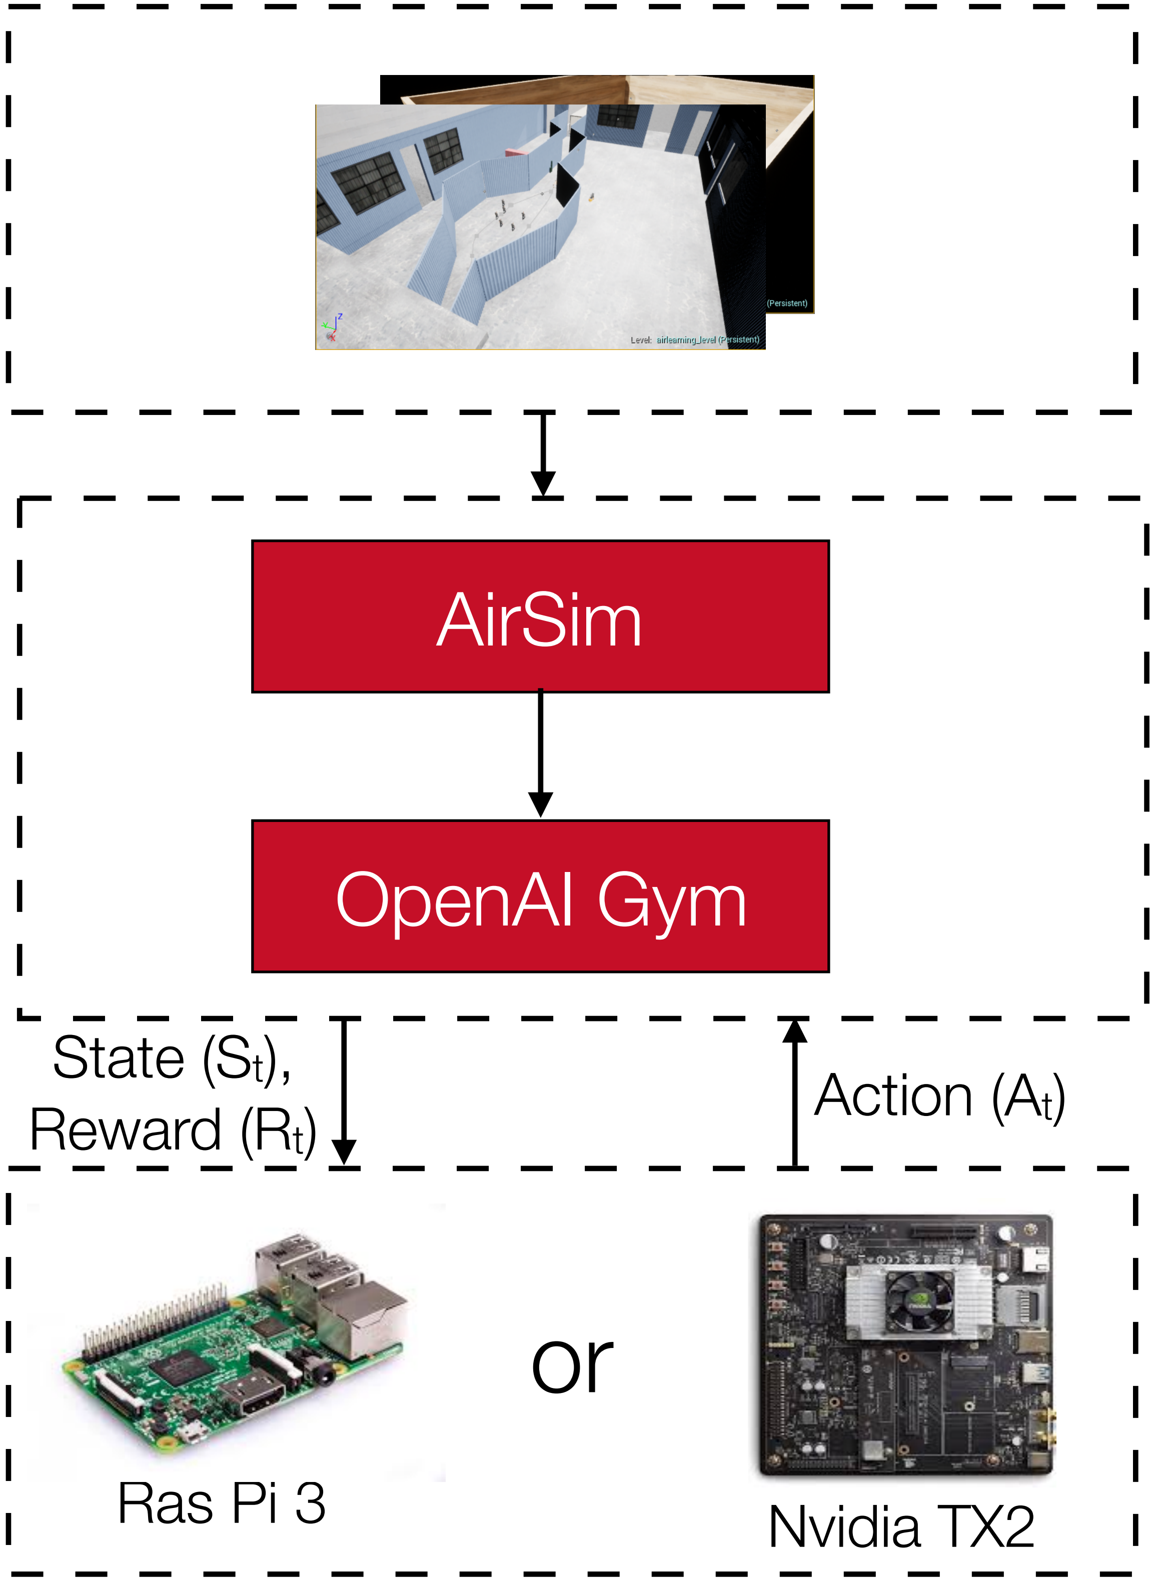

AirSim

OpenAI Gym

State (S_t),
Reward (R_t)

Action (A_t)

Ras Pi 3

or

Nvidia TX2

Simulation Environment

DRL Framework

Compute System

Accelerator Design

SiFive

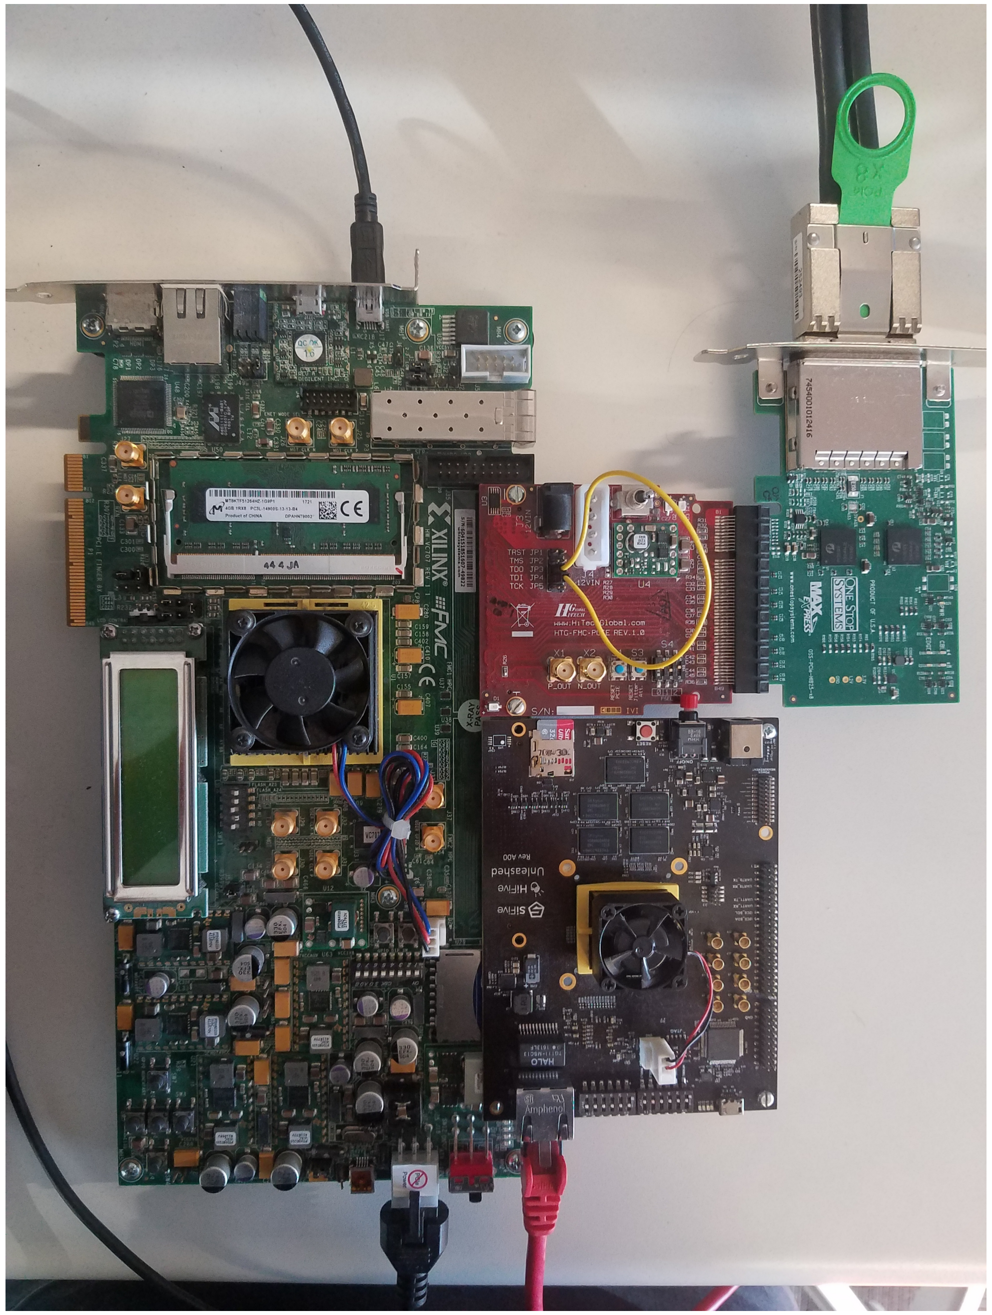

+

ROS

Open Source Robotics Foundation

+

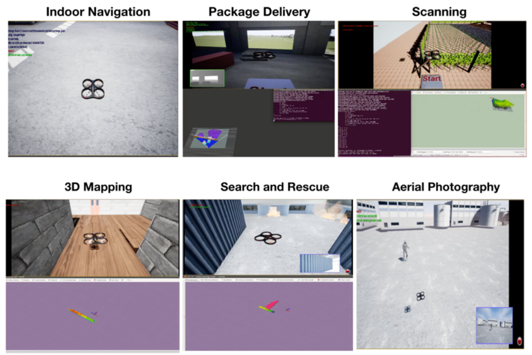

= XPU

Accelerator Design

SiFive

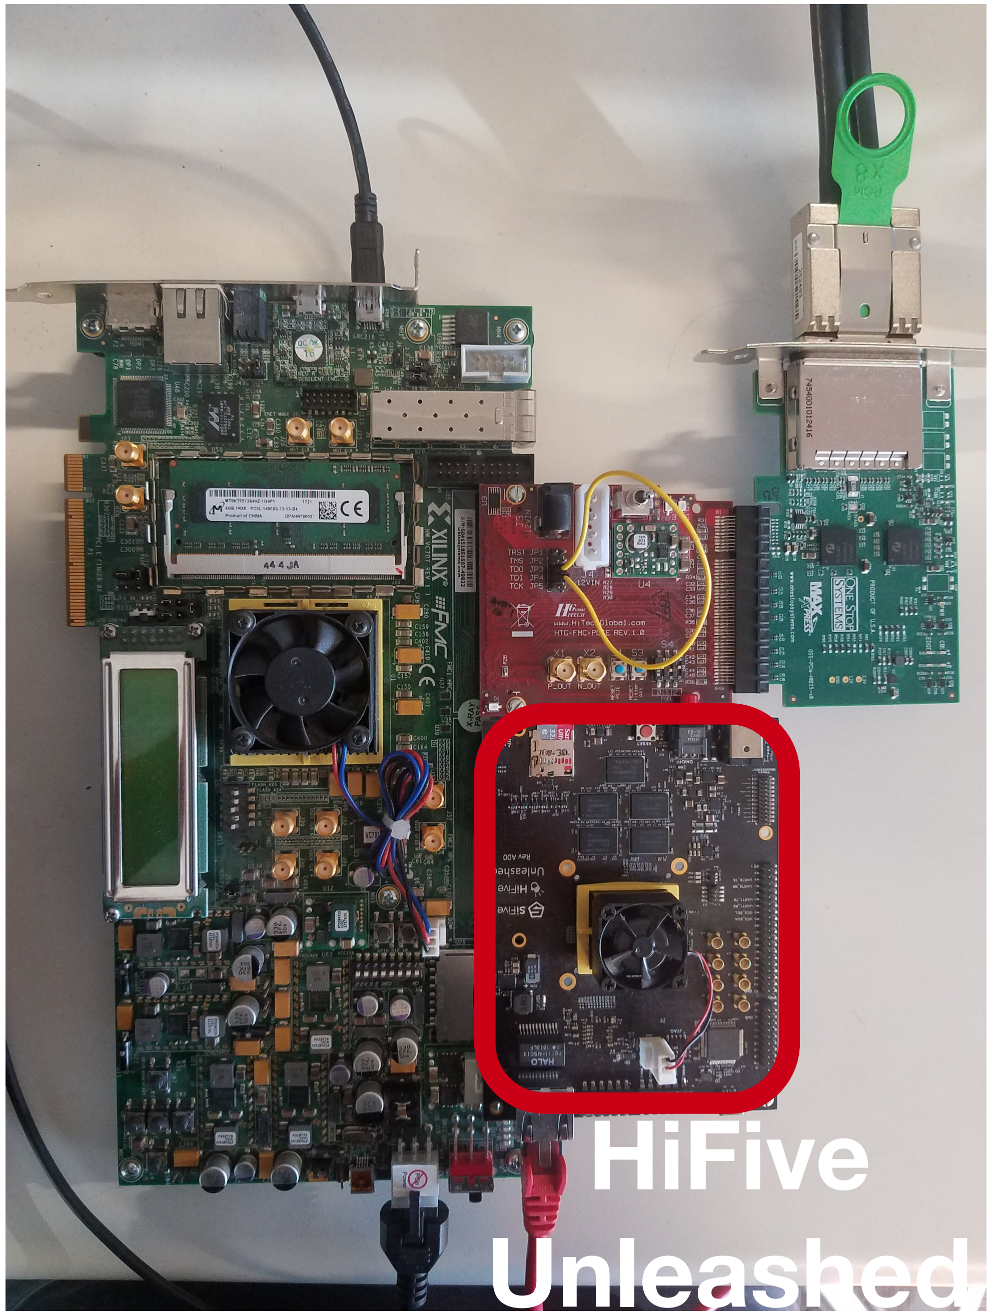

+

ROS

Open Source Robotics Foundation

+

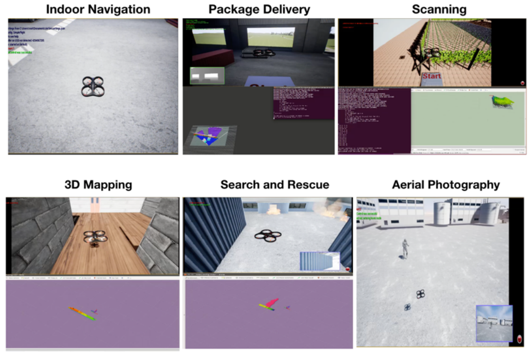

= XPU

Accelerator Design

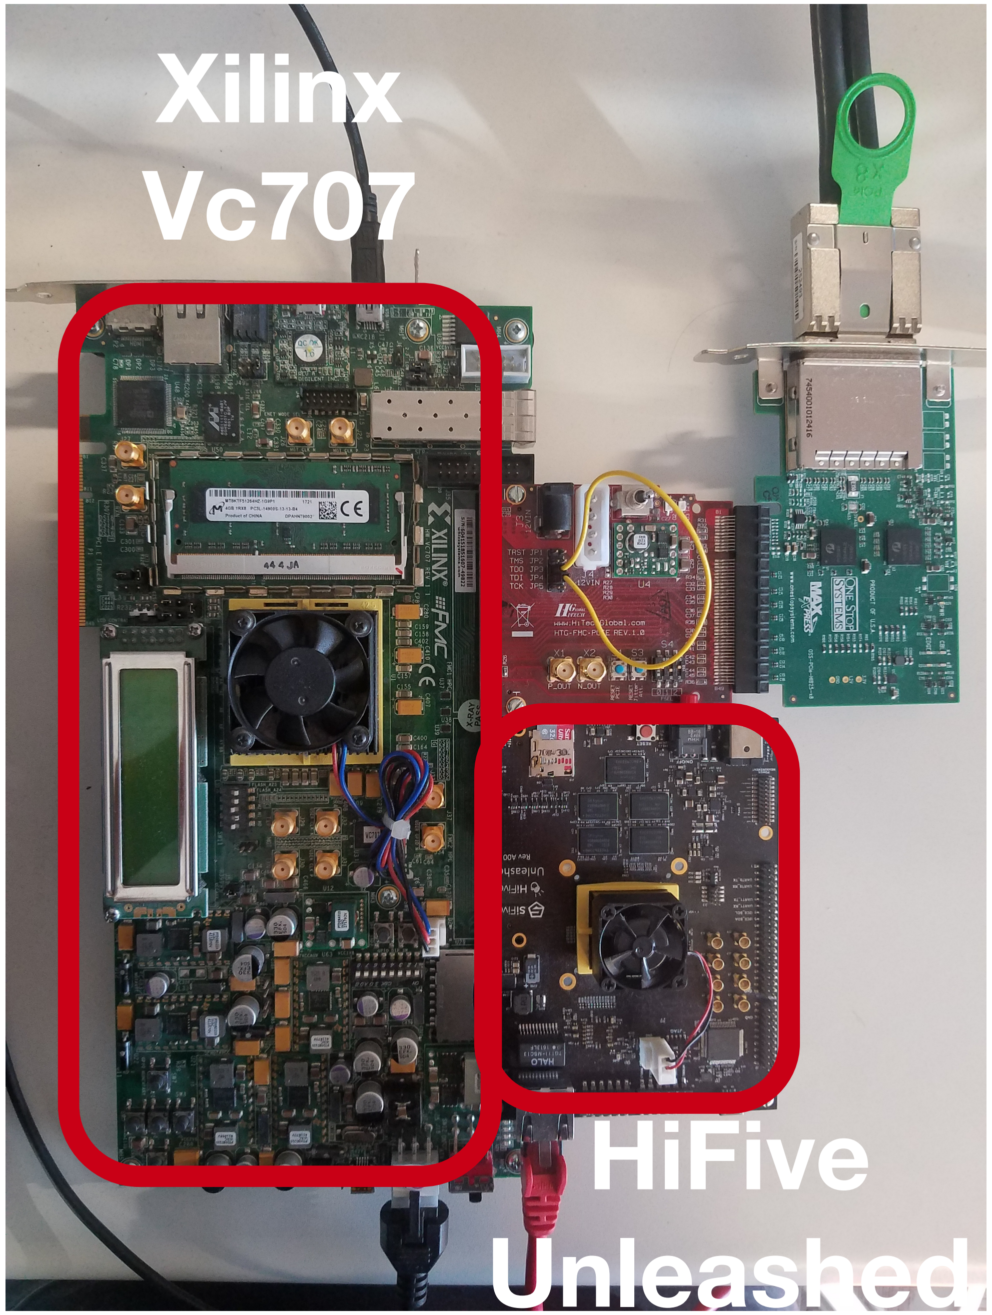

+

ROS

Open Source Robotics Foundation

+

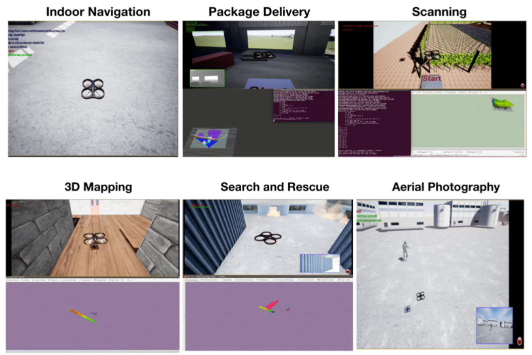

= XPU

Accelerator Design

SiFive

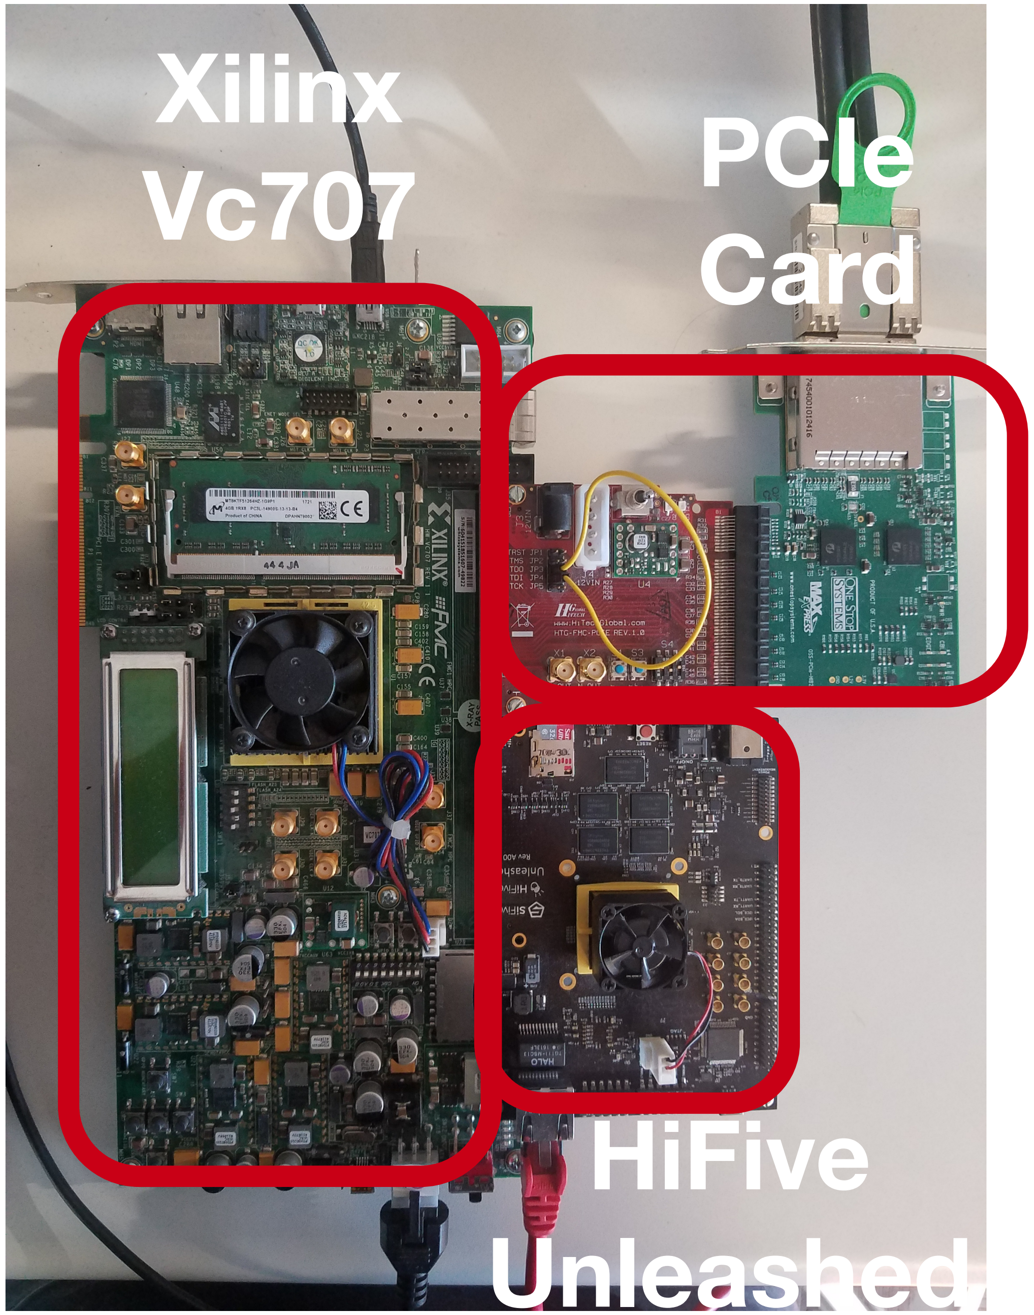

+

ROS

Open Source Robotics Foundation

+

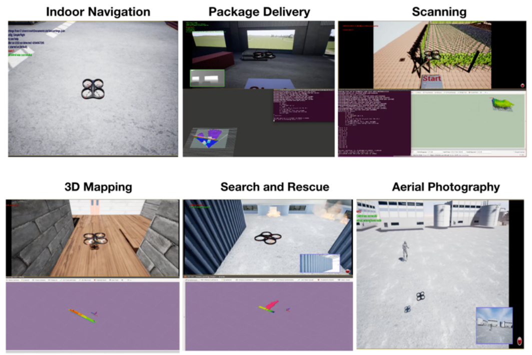

= XPU

Accelerator Design

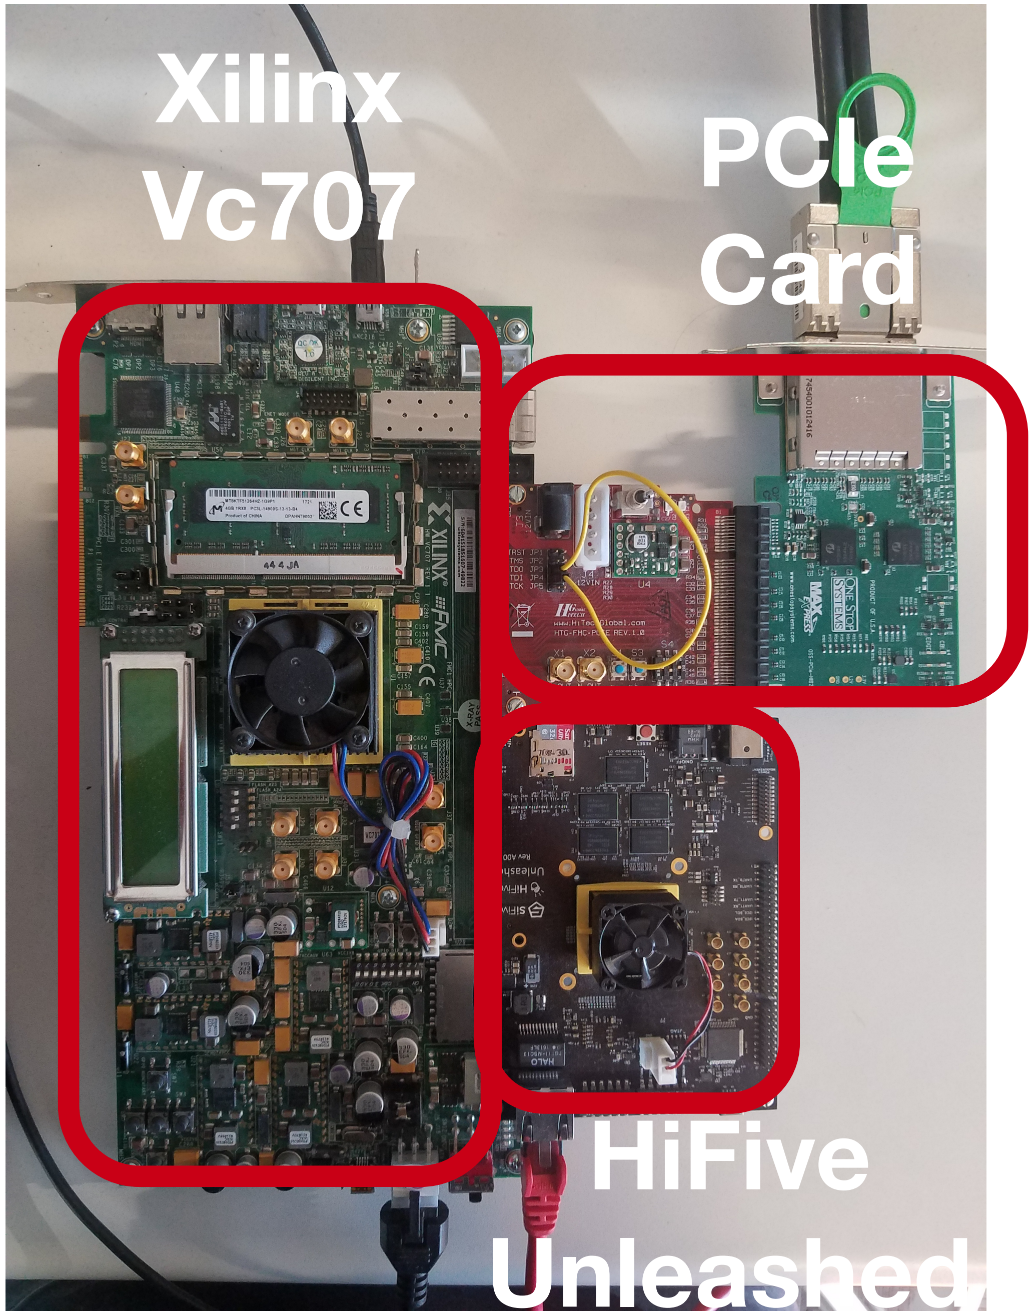

+

ROS

Open Source Robotics Foundation

+

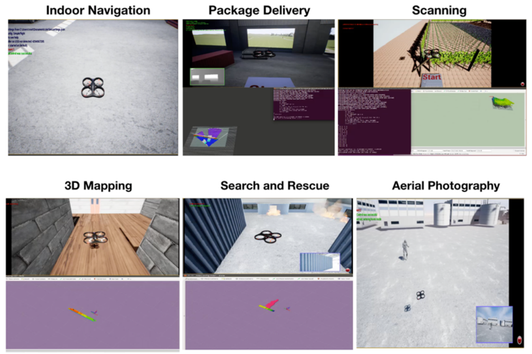

=

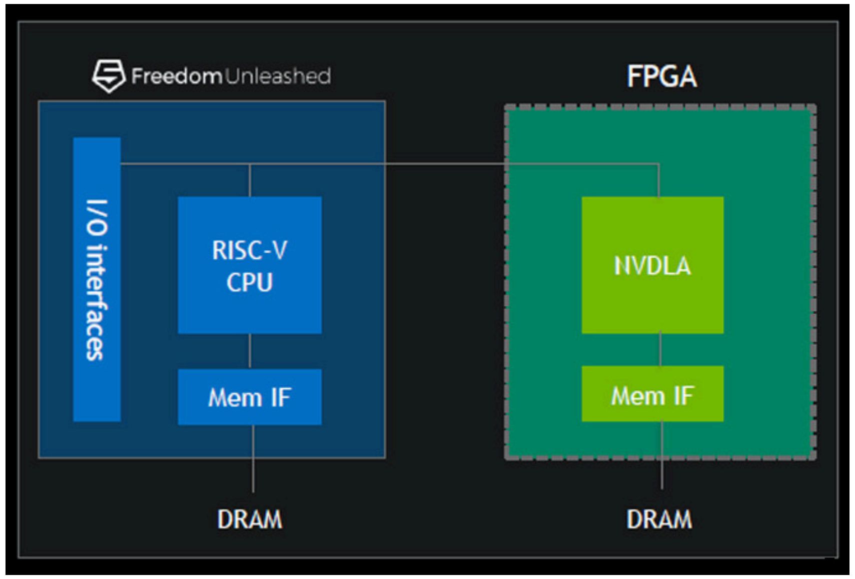

Algorithm Performance: Small vs. Large Policy

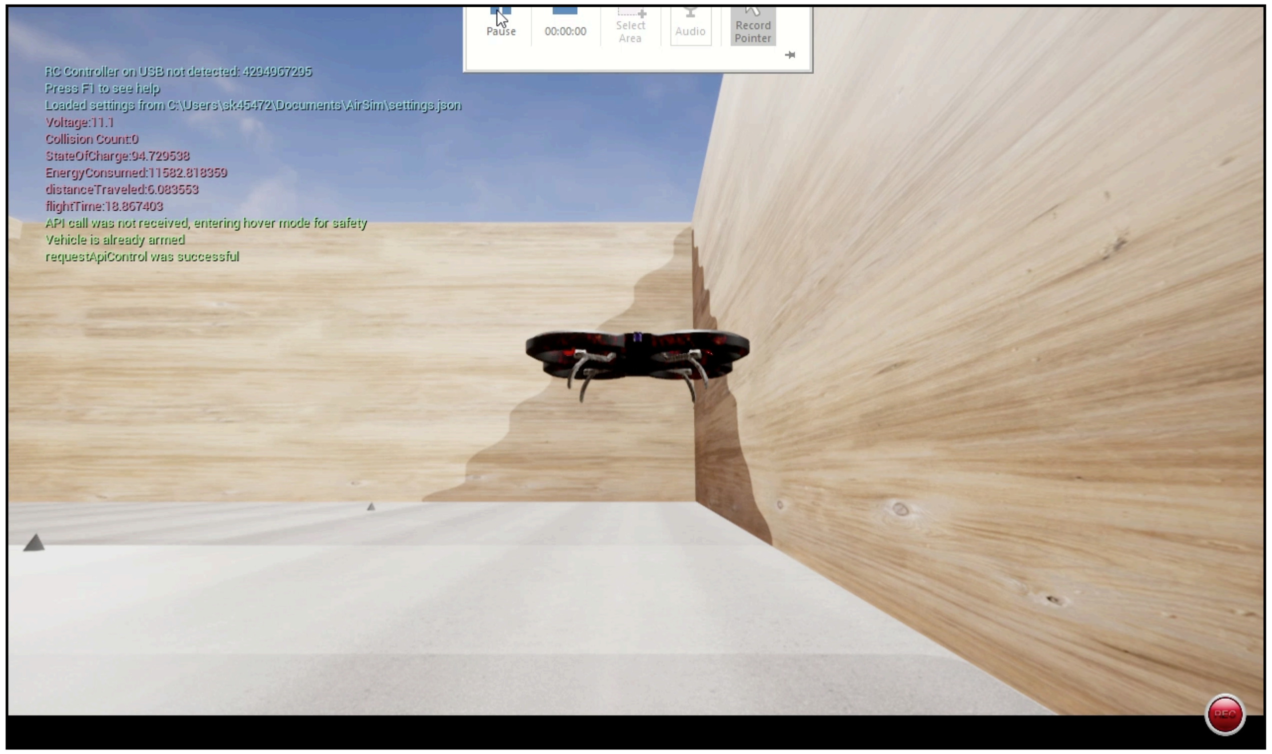

```
RC Controller on USB not detected: 4294967295
Press F1 to see help
Loaded settings from C:\Users\sk45472\Documents\AirSim\settings.json
Voltage:11.1
Collision Count:0
StateOfCharge:94.729538
EnergyConsumed:11582.818359
distanceTraveled:6.083553
flightTime:18.867403
API call was not received, entering hover mode for safety
Vehicle is already armed
requestApiControl was successful
```

Small Policy

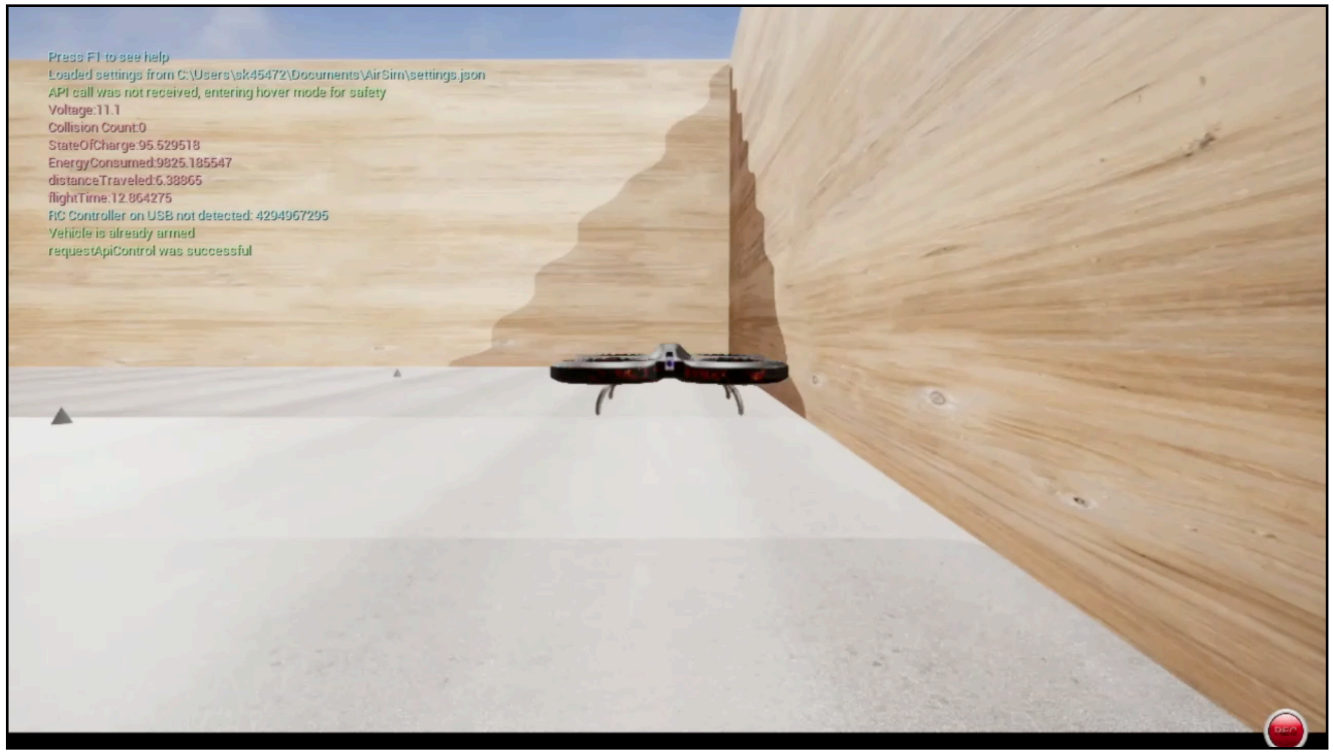

```
Press F1 to see help
Loaded settings from C:\Users\sk45472\Documents\AirSim\settings.json
API call was not received, entering hover mode for safety
Voltage:11.1
Collision Count:0
StateOfCharge:95.529518
EnergyConsumed:9825.185547
distanceTraveled:6.38865
flightTime:12.864275
RC Controller on USB not detected: 4294967295
Vehicle is already armed
requestApiControl was successful
```

Large Policy

Algorithm Performance: Small vs. Large Policy

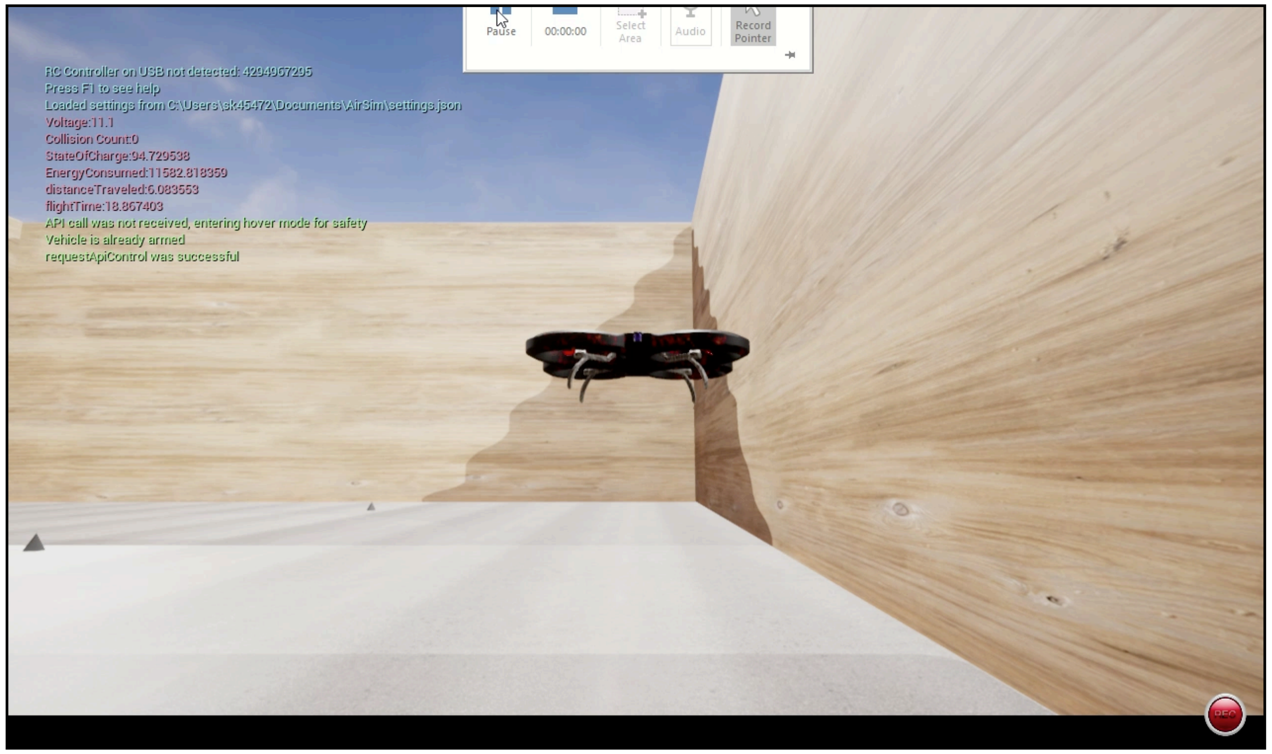

```
RC Controller on USB not detected: 4294967295
Press F1 to see help
Loaded settings from C:\Users\sk45472\Documents\AirSim\settings.json
Voltage:11.1
Collision Count:0
StateOfCharge:94.729538
EnergyConsumed:11582.818359
distanceTraveled:6.083553
flightTime:18.867403
API call was not received, entering hover mode for safety
Vehicle is already armed
requestApiControl was successful
```

Small Policy

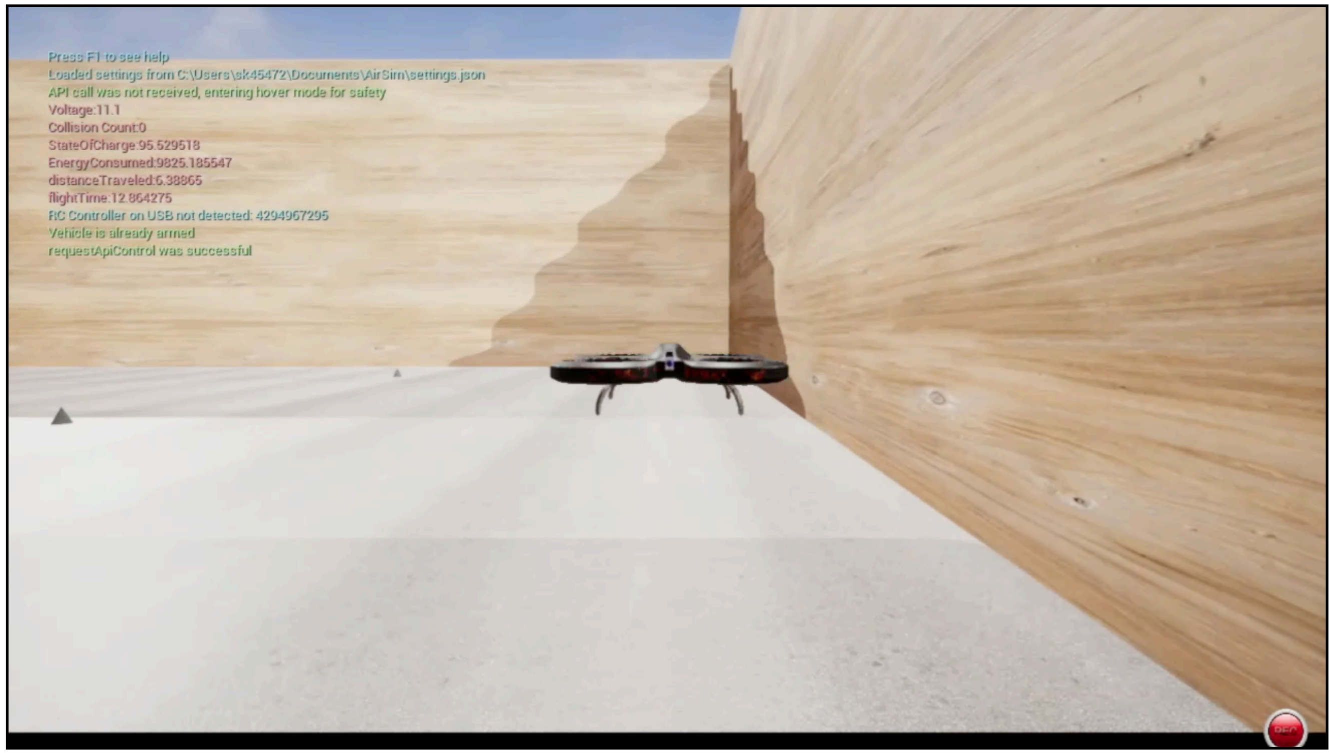

```
Press F1 to see help
Loaded settings from C:\Users\sk45472\Documents\AirSim\settings.json
API call was not received, entering hover mode for safety
Voltage:11.1
Collision Count:0
StateOfCharge:95.529518
EnergyConsumed:9825.185547
distanceTraveled:6.38865
flightTime:12.864275
RC Controller on USB not detected: 4294967295
Vehicle is already armed
requestApiControl was successful
```

Large Policy

MAVBench: Kernel Decomposition

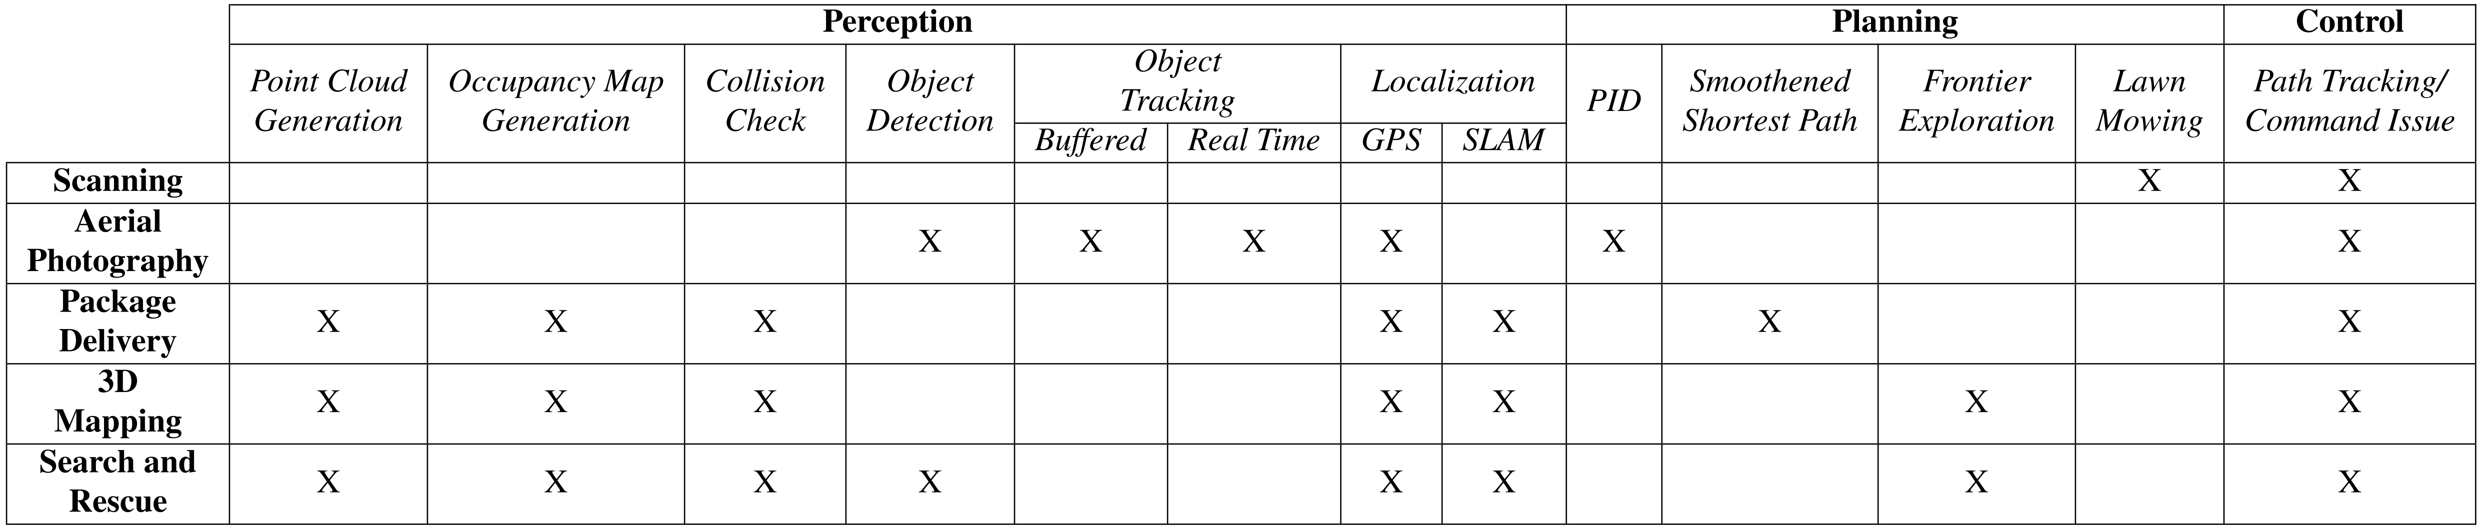

	Perception								**Planning**				**Control**
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												X	X
Aerial Photography				X	X	X	X		X				X
Package Delivery	X	X	X				X	X		X			X
3D Mapping	X	X	X				X	X			X		X
Search and Rescue	X	X	X	X			X	X			X		X

MAVBench: Kernel Decomposition

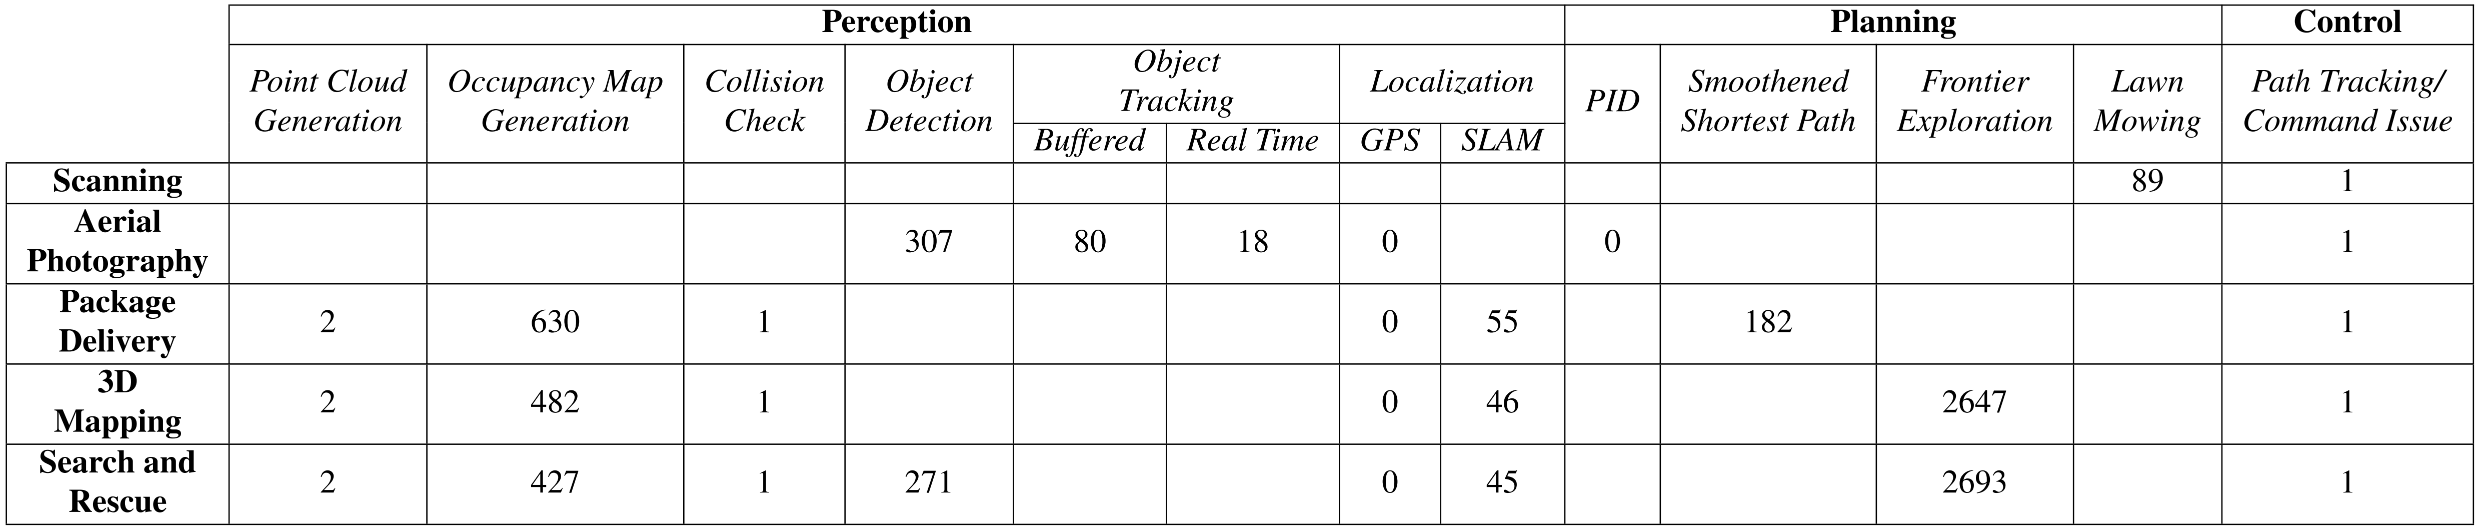

	Perception								**Planning**				**Control**
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												89	1
Aerial Photography				307	80	18	0		0				1
Package Delivery	2	630	1				0	55		182			1
3D Mapping	2	482	1				0	46			2647		1
Search and Rescue	2	427	1	271			0	45			2693		1

MAVBench: Kernel Decomposition

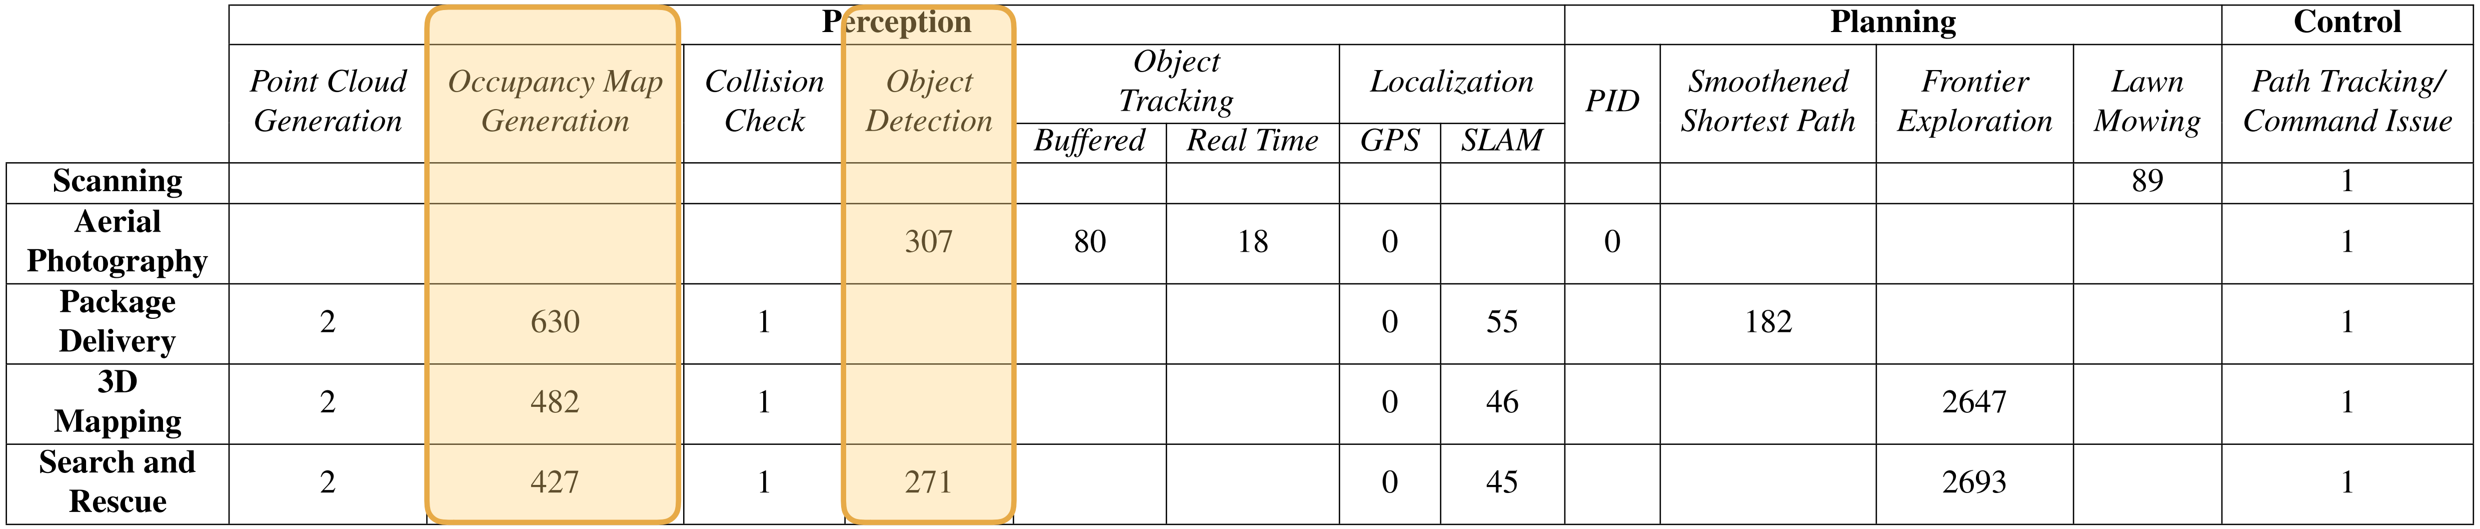

	Perception								Planning				Control
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												89	1
Aerial Photography				307	80	18	0		0				1
Package Delivery	2	630	1				0	55		182			1
3D Mapping	2	482	1				0	46			2647		1
Search and Rescue	2	427	1	271			0	45			2693		1

MAVBench: Kernel Decomposition

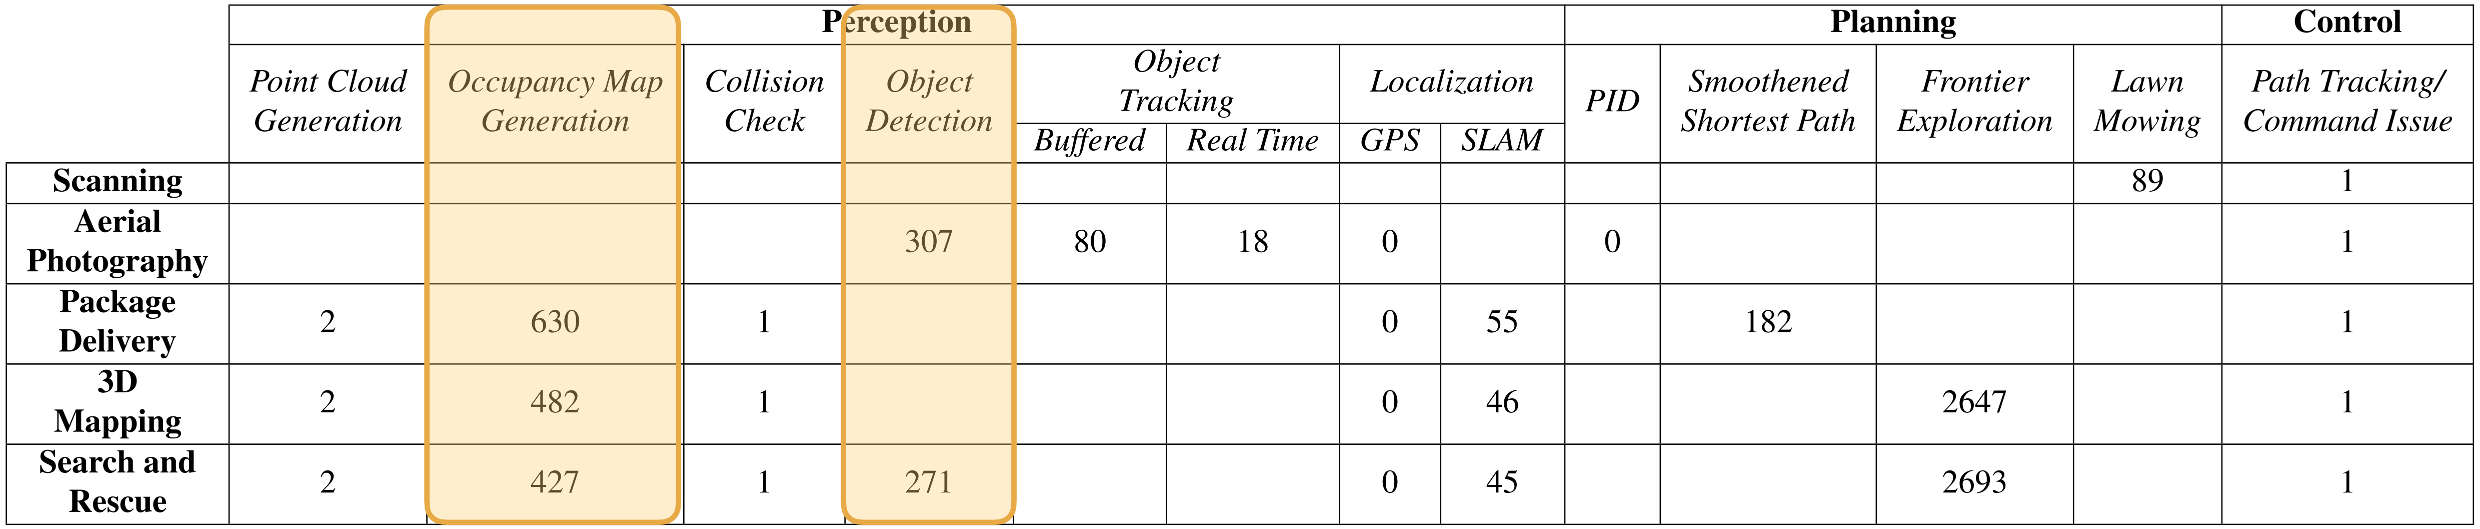

	Perception								Planning				Control
	Point Cloud Generation	*Occupancy Map Generation*	*Collision Check*	*Object Detection*	*Object Tracking*		*Localization*		*PID*	*Smoothened Shortest Path*	*Frontier Exploration*	*Lawn Mowing*	*Path Tracking/ Command Issue*
					Buffered	*Real Time*	*GPS*	*SLAM*					
Scanning												89	1
Aerial Photography				307	80	18	0		0				1
Package Delivery	2	630	1				0	55		182			1
3D Mapping	2	482	1				0	46			2647		1
Search and Rescue	2	427	1	271			0	45			2693		1

Identify bottlenecks and accelerate them through domain specific logic.

Summary

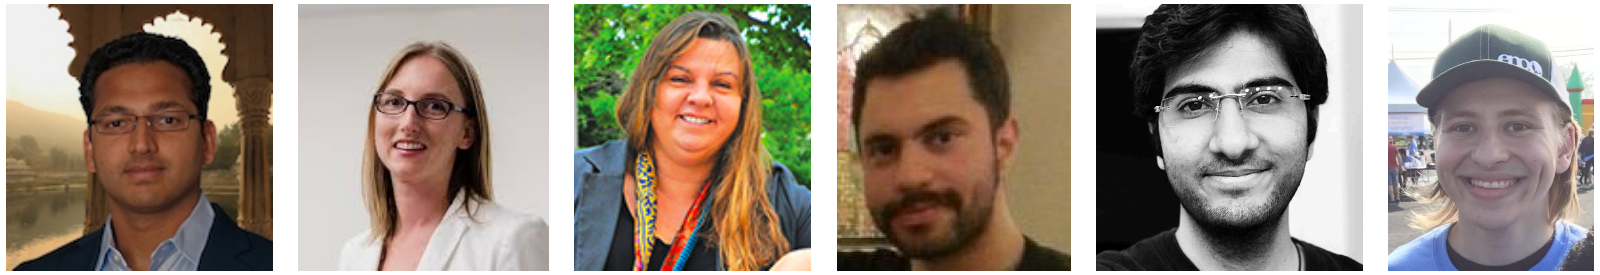

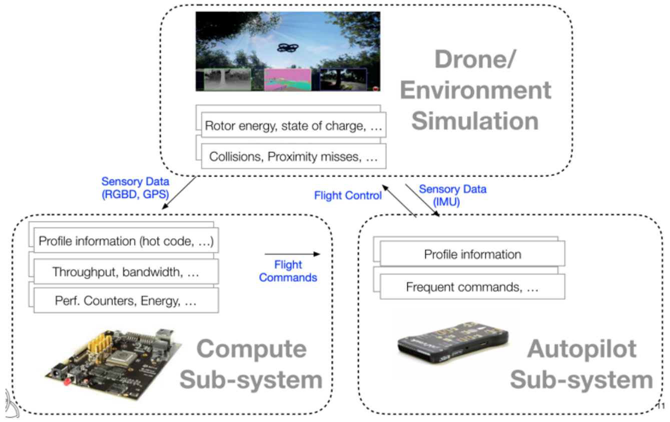

Closed-loop simulation is important for the development of architectural solutions for autonomous machines.

End-to-end applications are needed to do hardware and software co-design to design domain specific accelerators

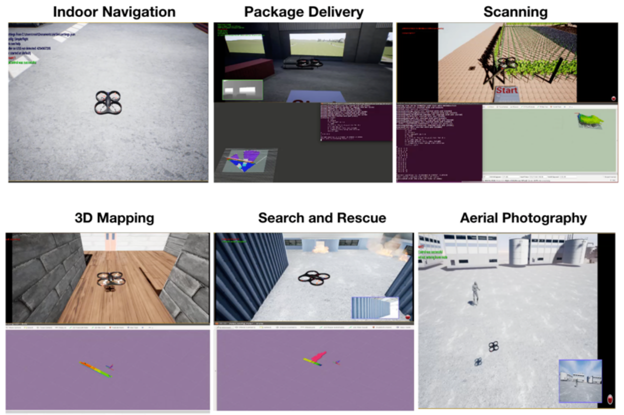

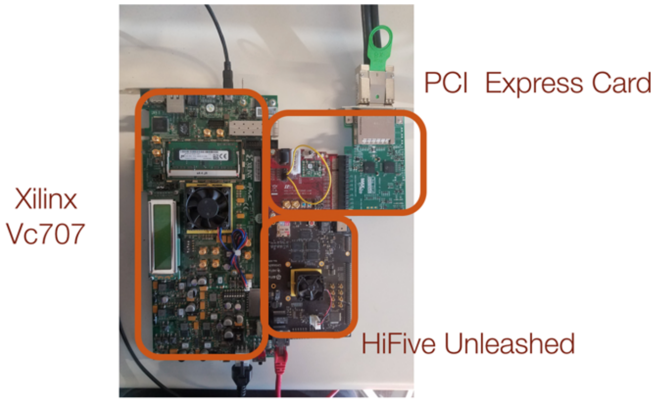

RISC-V with ROS unlocks the potential of designing custom hardware accelerators that are open source to drive innovation