



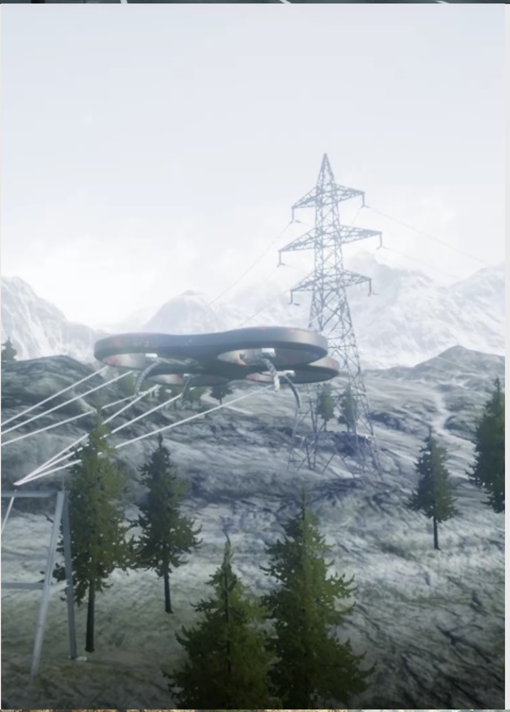

Role of Simulation in Machine Intelligence: Beyond Synthetic Data


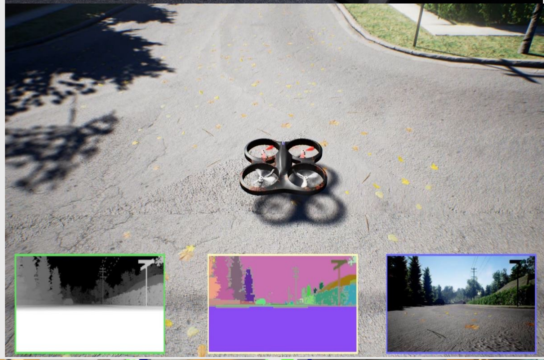

Ashish Kapoor



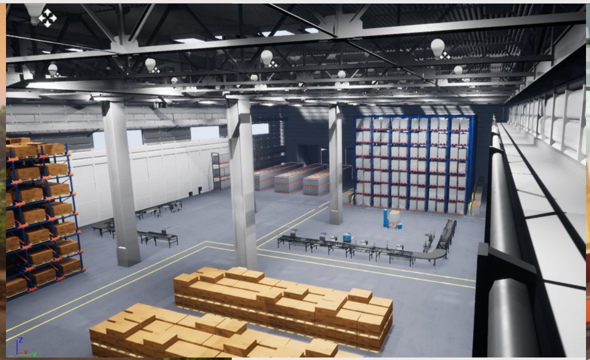


Modeling and Simulation: Foundation of Modern Engineering

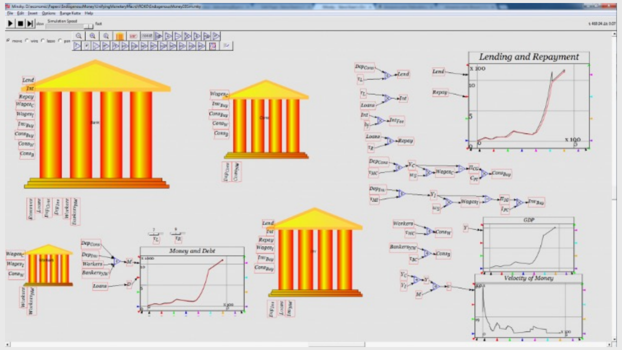

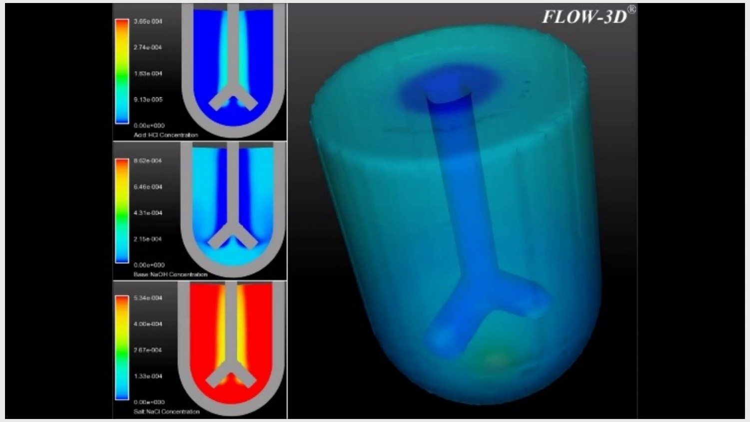

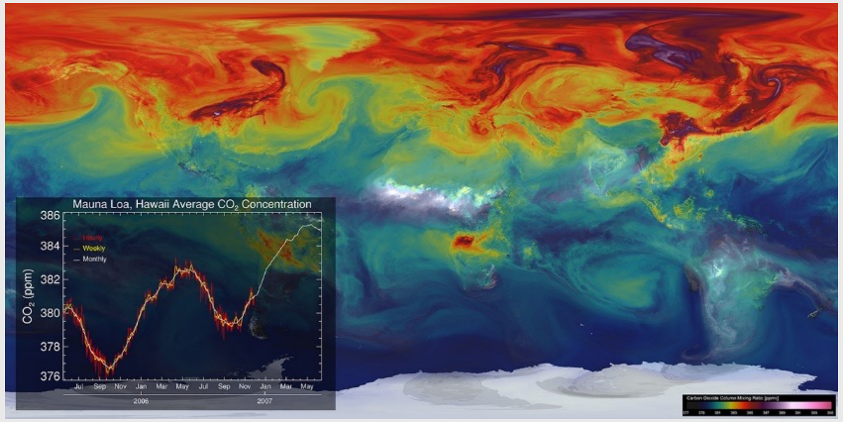

Simulation has enabled several successes of ML

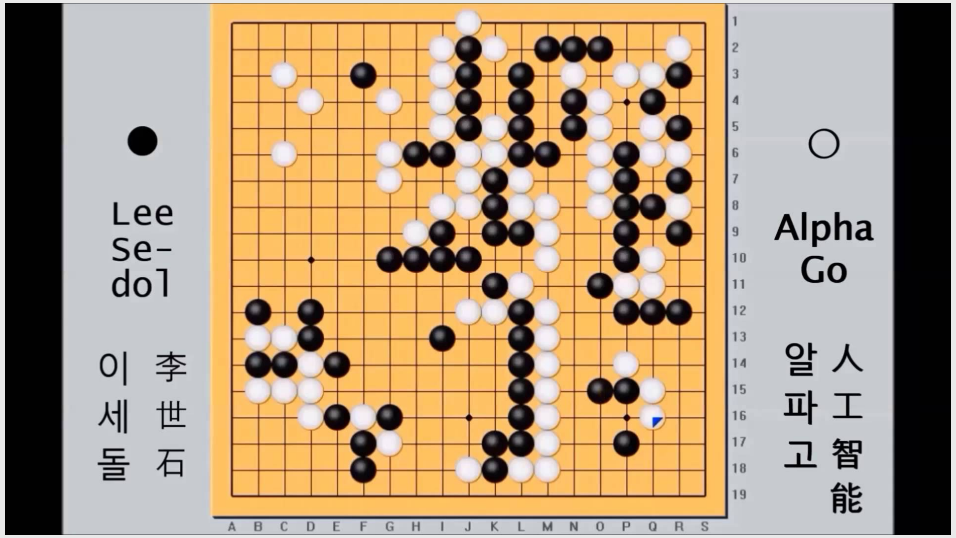

"Closed-world" systems make simulating
millions of trials easy.

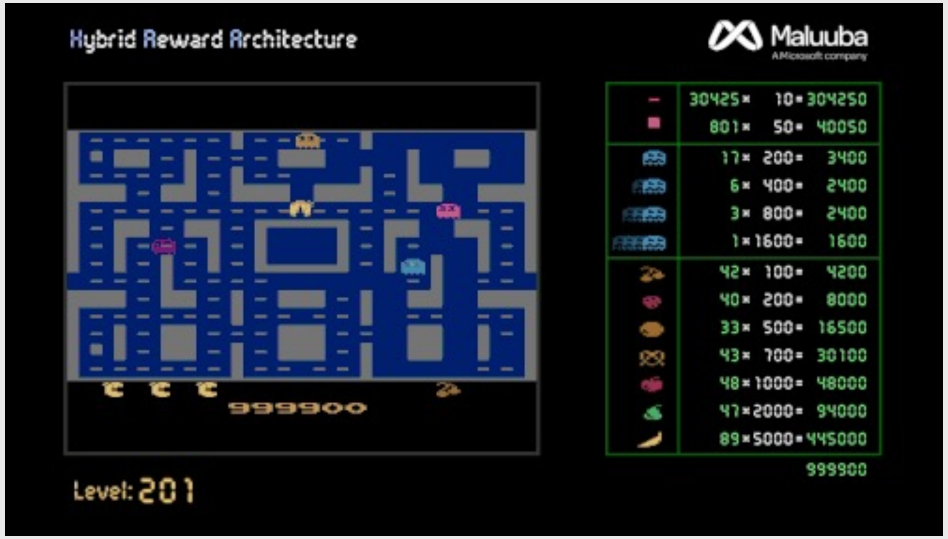

Initial Atari ~ 10M frames
Agent57 ~ 80 billion frames
Alpha Go ~ 130M self-plays
Dota 2 ~ 45000 years of Dota selfplay
(10 months)

Microsoft AirSim

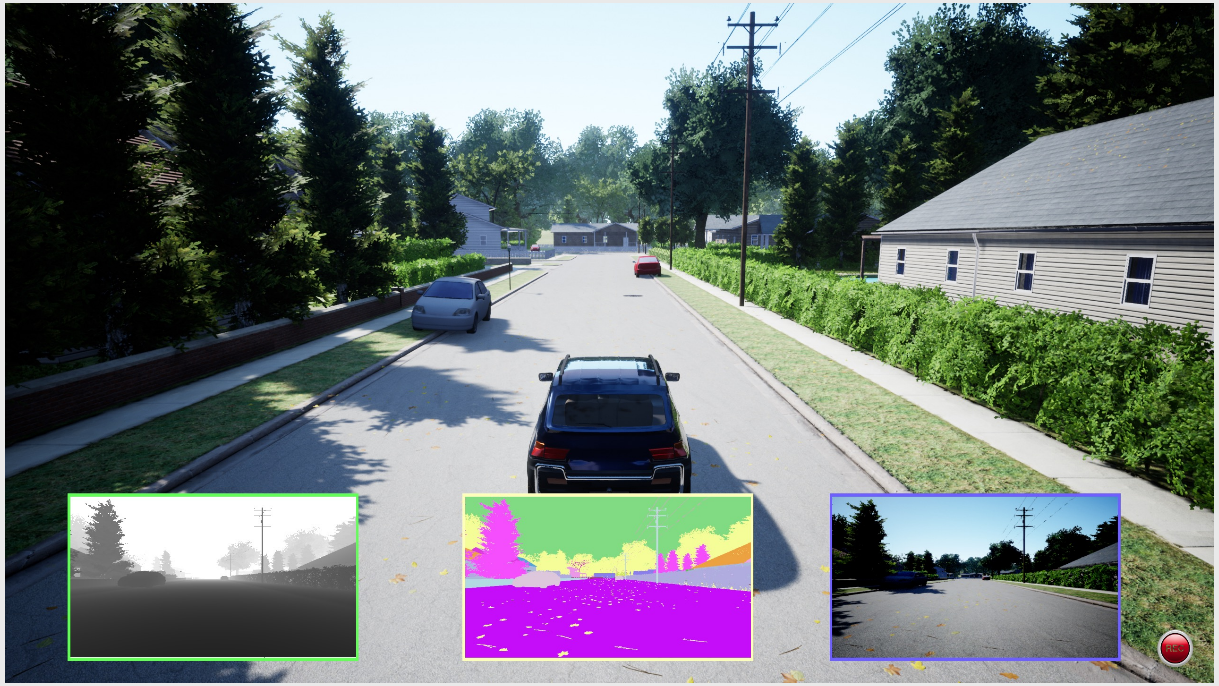

Shah, S., Dey, D., Lovett, C., & Kapoor, A. (2018). Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In Field and service robotics (pp. 621-635). Springer, Cham.

AirSim Architecture

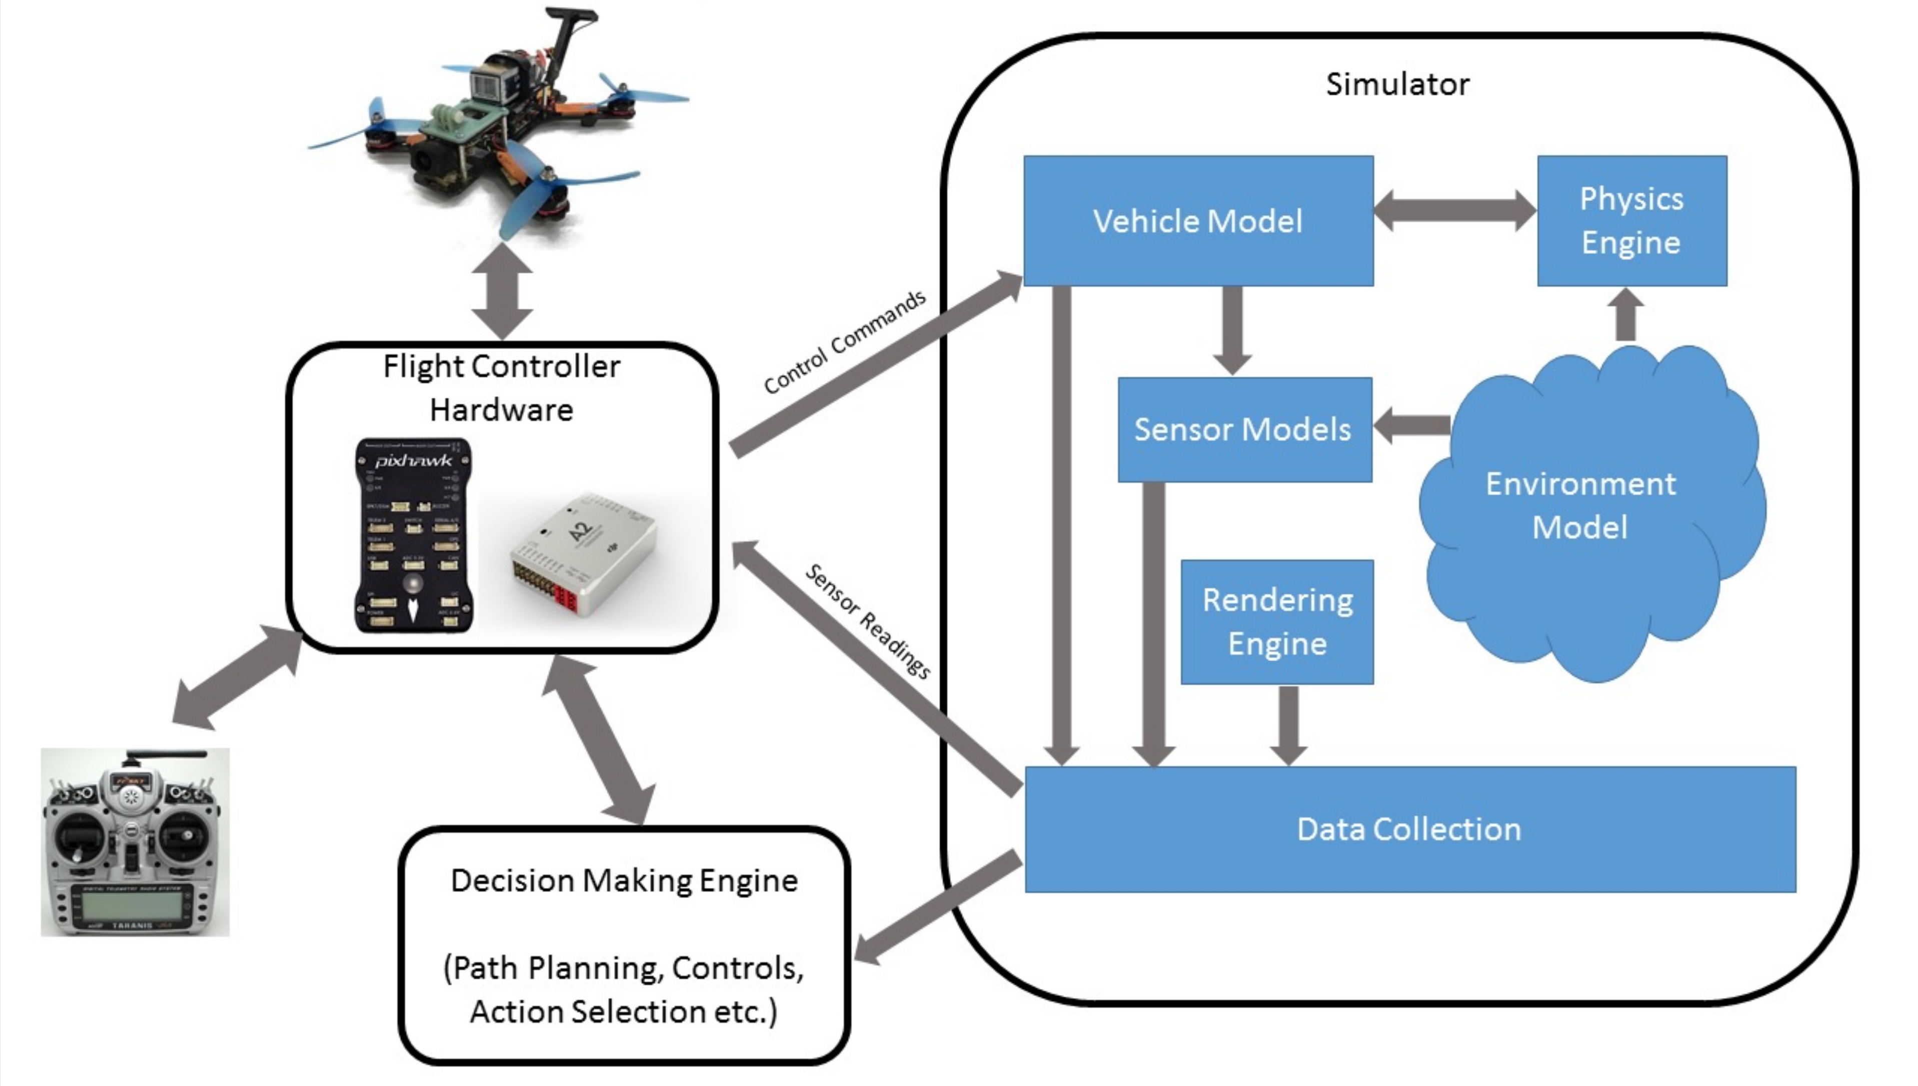

Simulation Centric Data Strategy with AirSim

- 20 Environments
- 500+ Trajectories
- 400,000+ Frames
- Diverse Motion Patterns
- Multiple Sensor Modalities

In the works: 100+ Environments, 2000+ Trajectories, 2M+ Frames, Novel Sensors

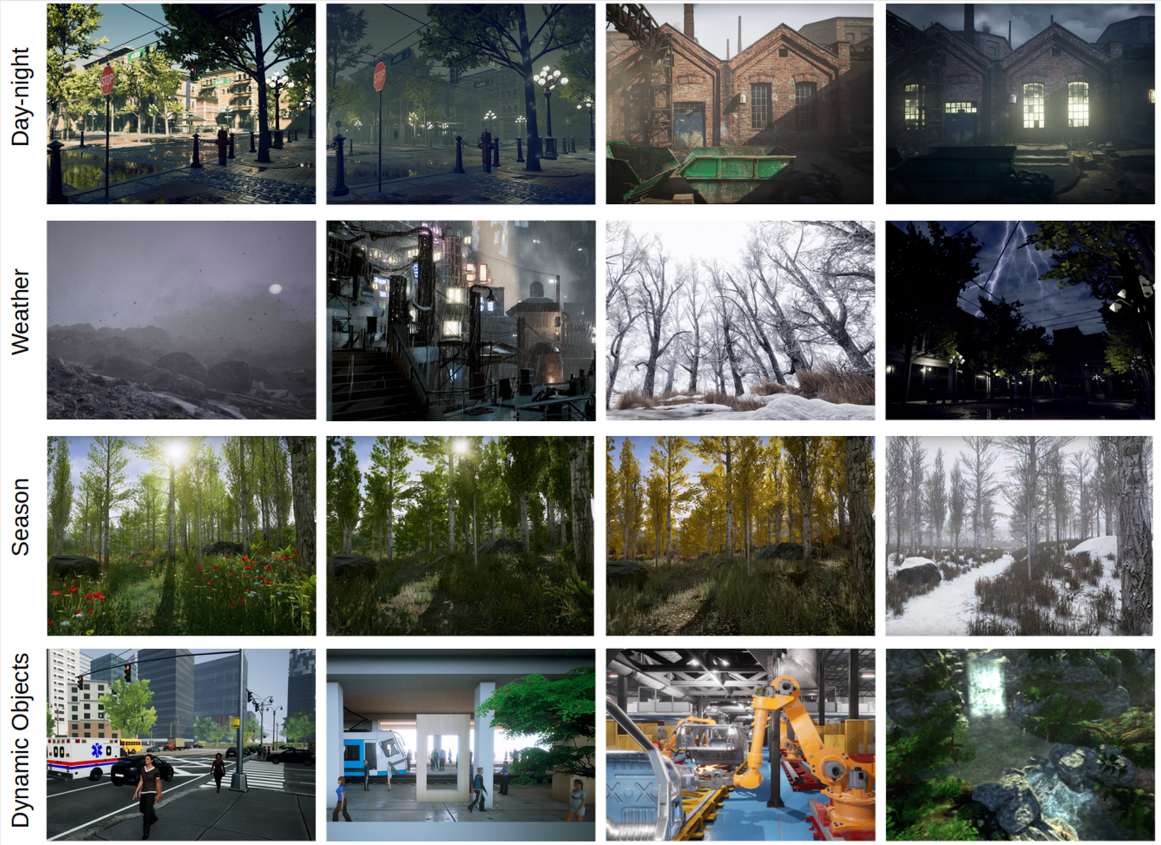

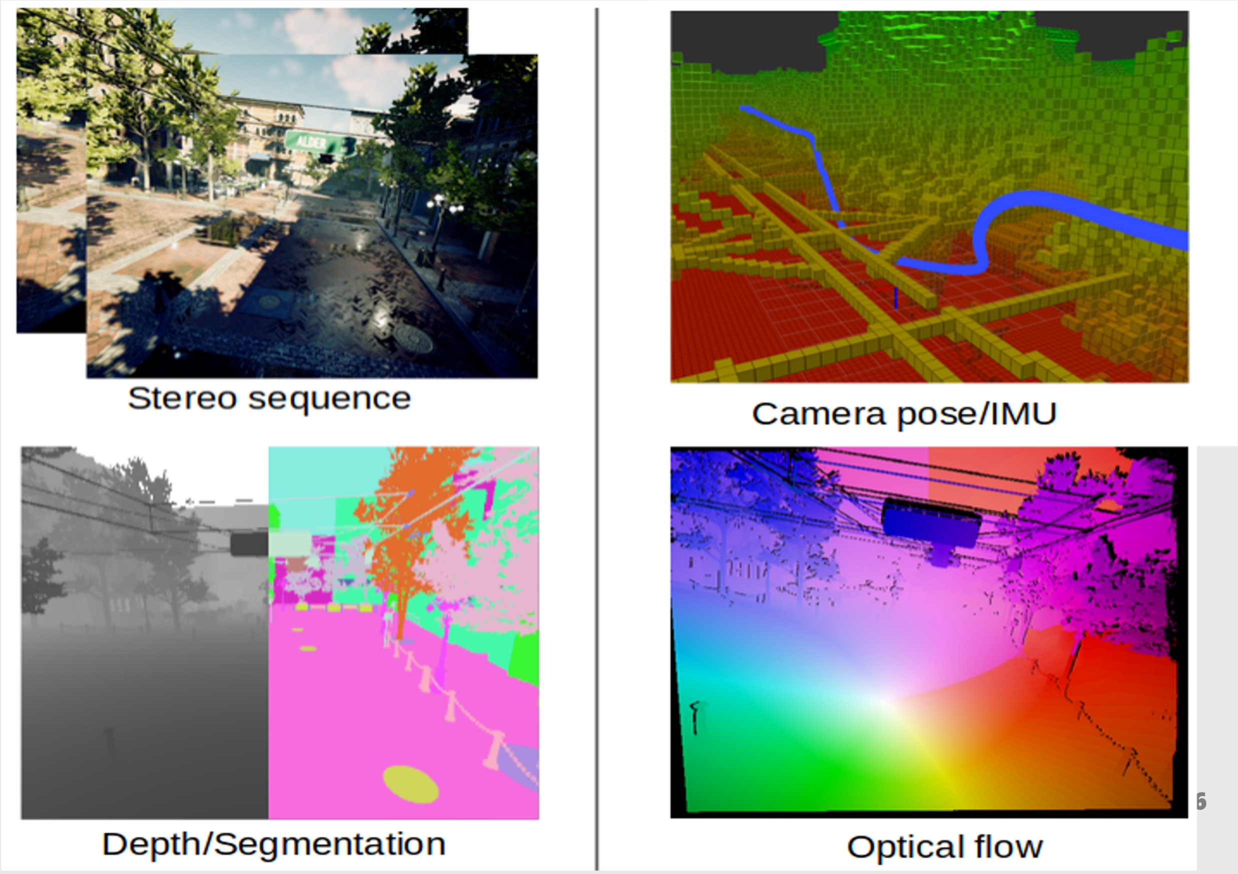

Wang et al. 2020; https://arxiv.org/abs/2003.14338

NEAR INFINITE DATA

Supervised Learning to Detect Poachers

Faster RCNN

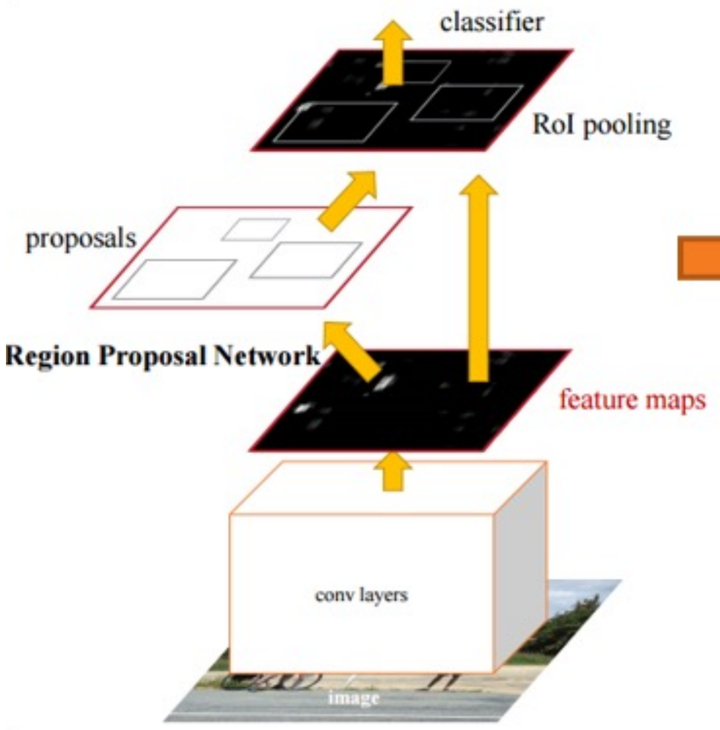

Simulated and real imagery for training

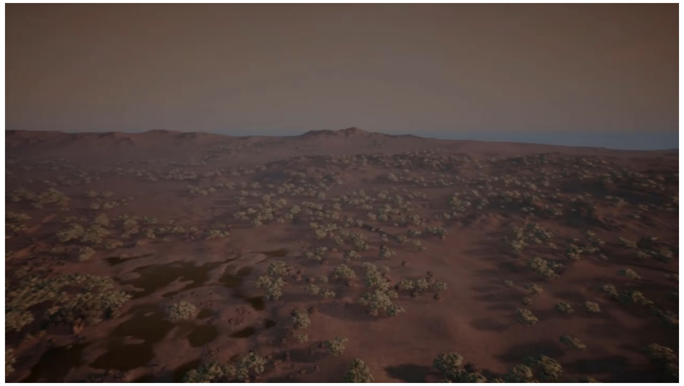

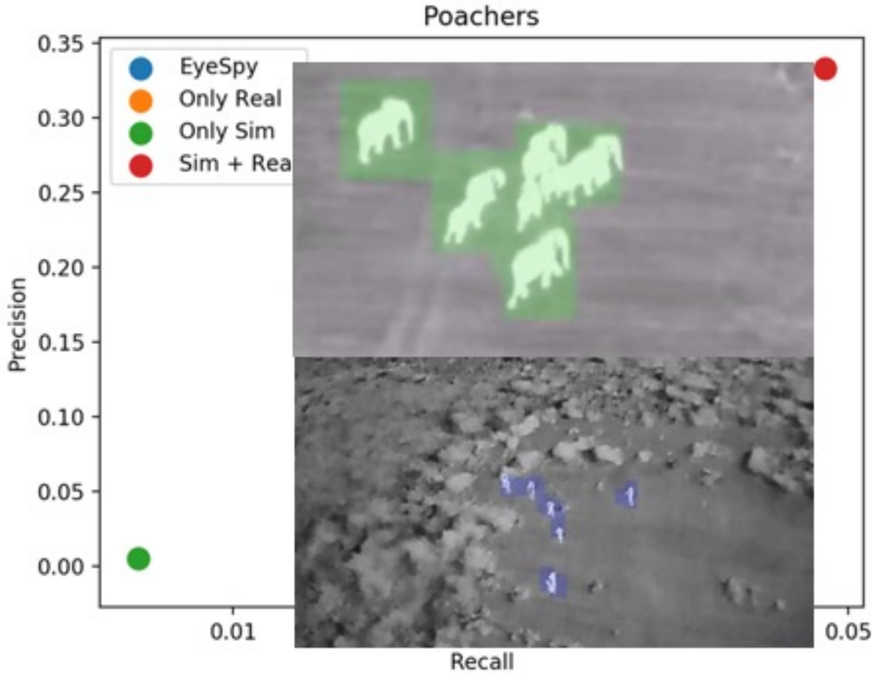

Example: Poacher Detection.
Bondi et al. COMPASS 2018

New Sensor Simulations – GANs etc

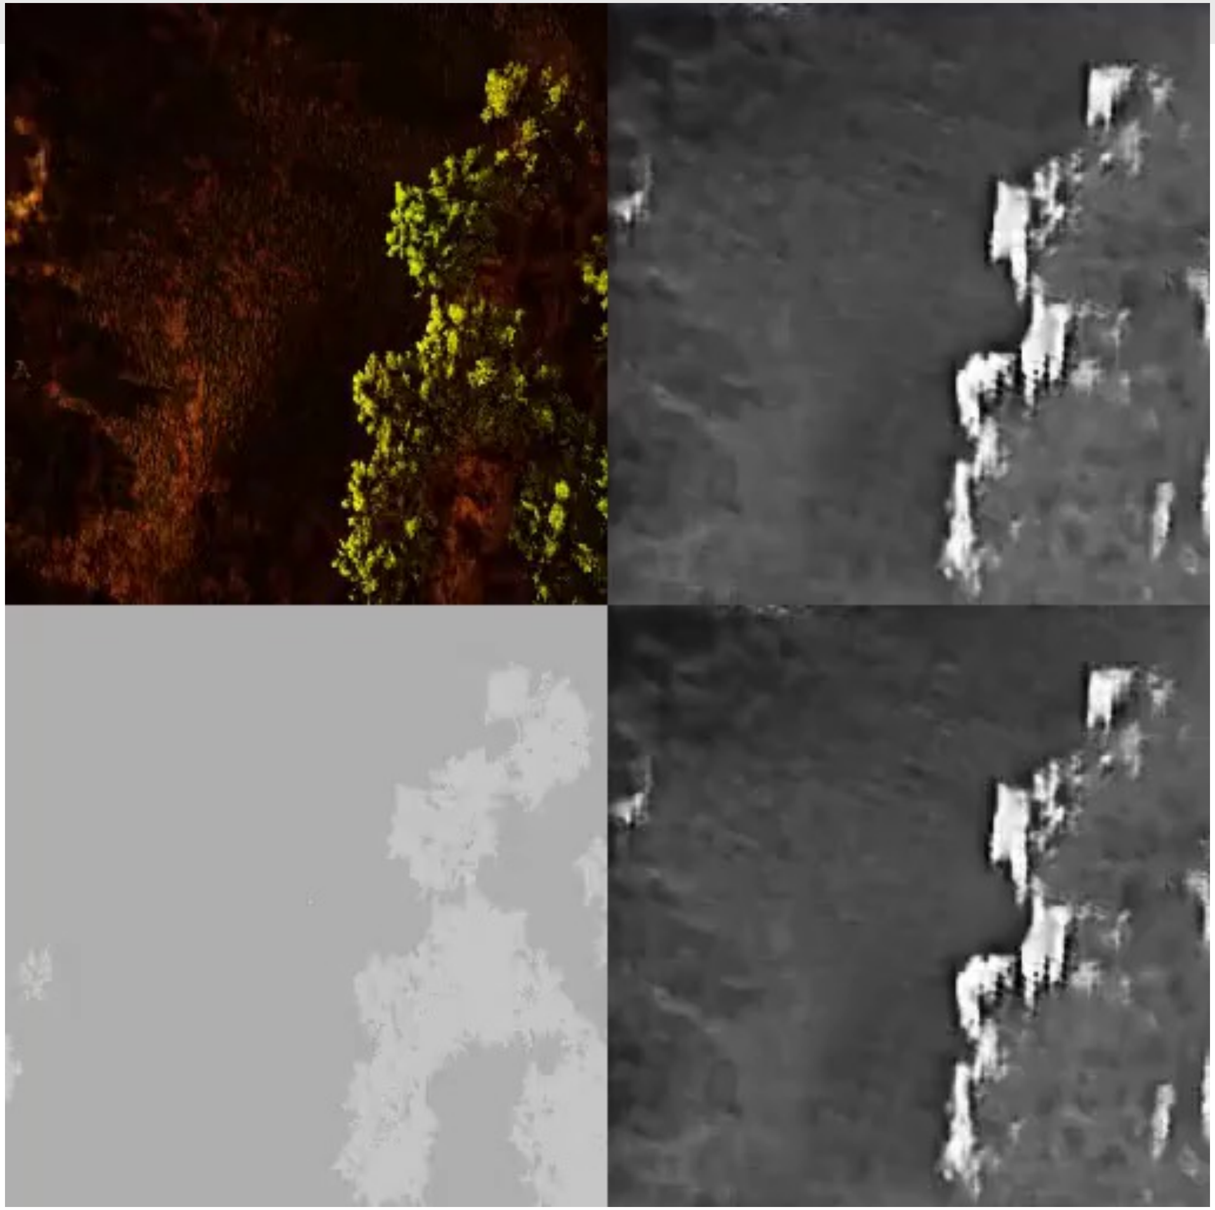

Berry et al. 2019

Berry et al. 2019

Data-driven composition of simulators

Data-driven
Differentiable
Neural

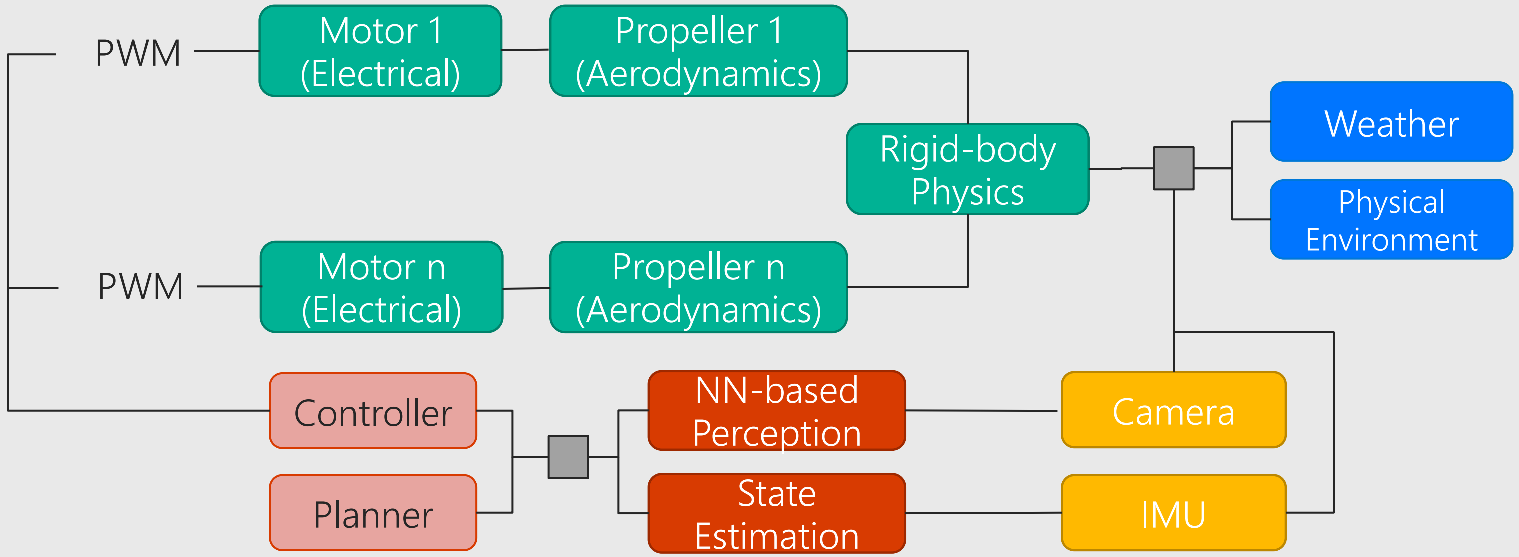

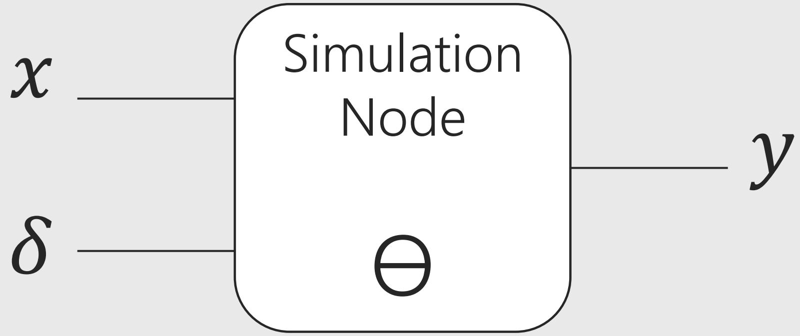

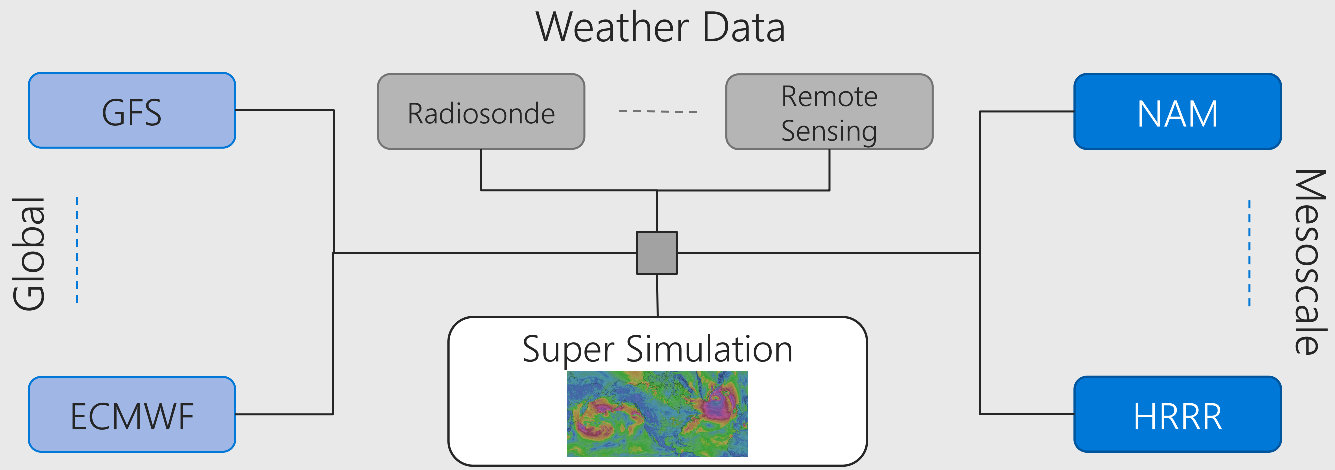

COMPASS.V0 Current Achievements – Drone Racing

Input

Onboard camera image
of gate from drone

Output

3D velocity command
for drone to reach gate

CM-VAE

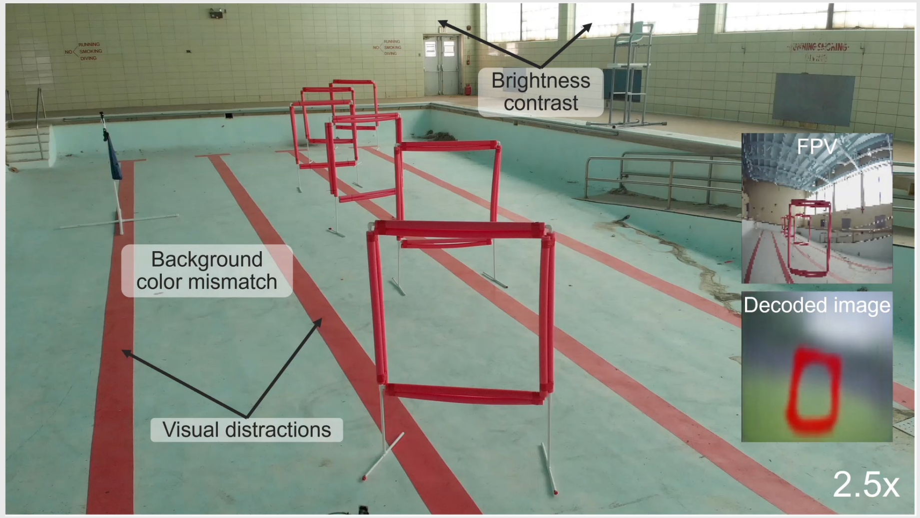

COMPASS

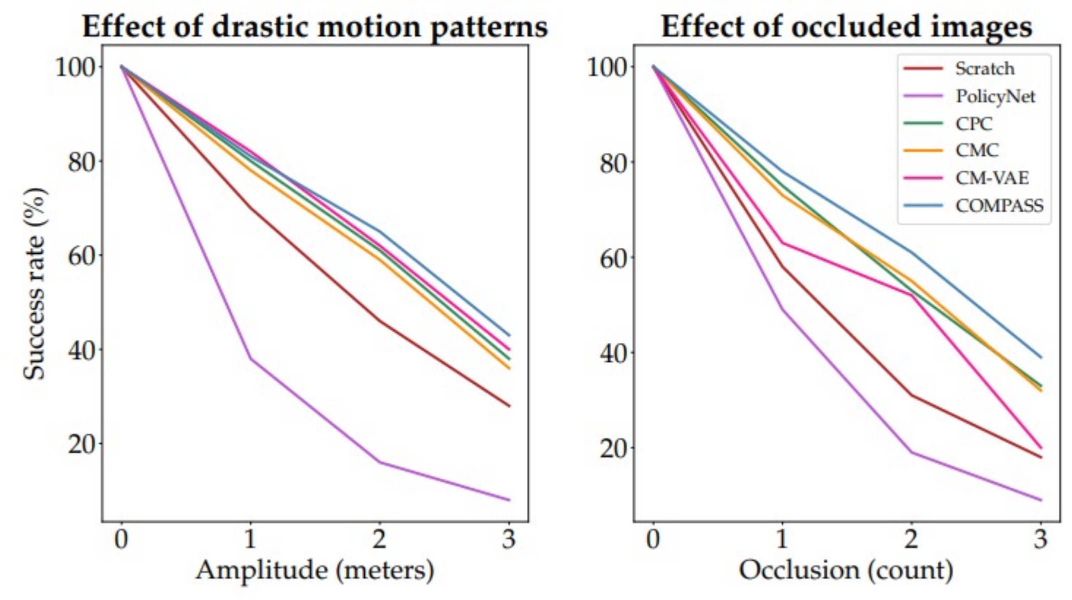

Perception:

>> SCRATCH

> SOTA pre-training approaches

>= SOTA task-specific approach

Policy:

Robust to drastic motion patterns

Robust to poor quality inputs

Model	Radius r[m]	Azimuth θ[○]	Polar ψ[○]	Yaw ϕ[○]
SCRATCH	0.41±0.013	2.40±0.14	2.50±0.14	11.0±0.67
CM-VAE	0.39±0.023	2.30±0.23	2.10±0.23	**9.70±0.75**
CPC	0.36±0.020	2.35±0.21	2.00±0.24	11.1±0.75
CMC	0.38±0.018	2.26±0.25	2.12±0.25	11.0±0.72
COMPASS	**0.31±0.016**	**2.09±0.17**	**1.98±0.17**	10.03±0.82

Bonatti et al. IROS 2020; https://arxiv.org/abs/1909.06993

Enabling Autonomy – Imitation and RL for Racing

Zadok et al. 2020

Example: Optimal Scanning. What Trajectory to Fly?

- Safety: How risky are the flight maneuvers?
- Efficiency: How long does the process take?
- Performance: How good is the 3D model?

Roberts et al. ECCV 2018

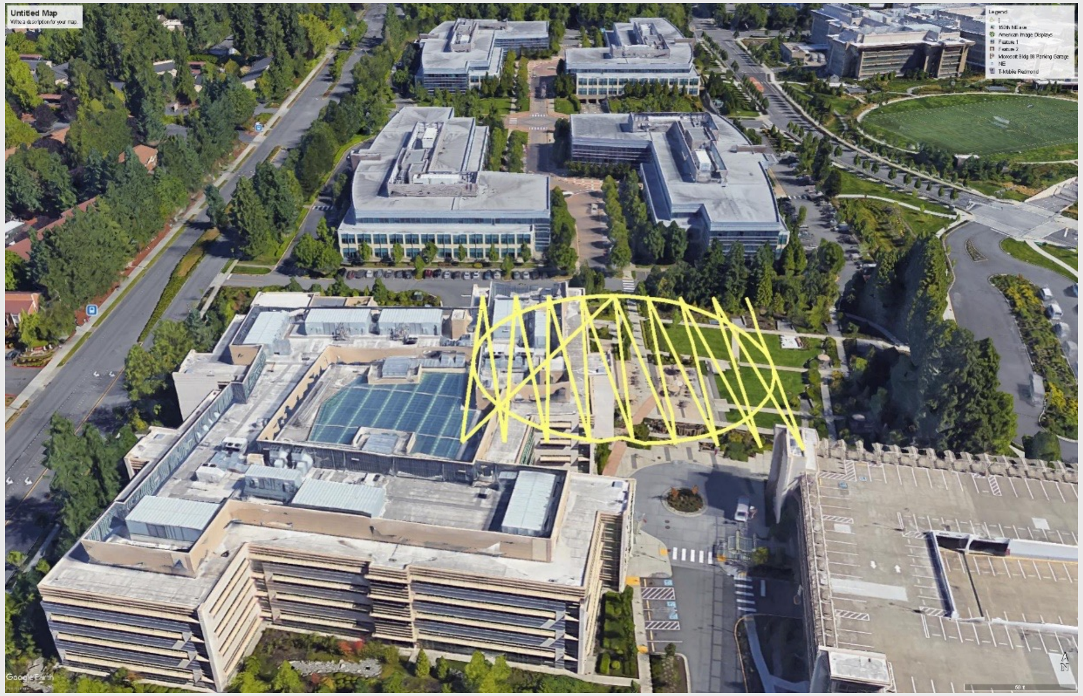
Untitled Map
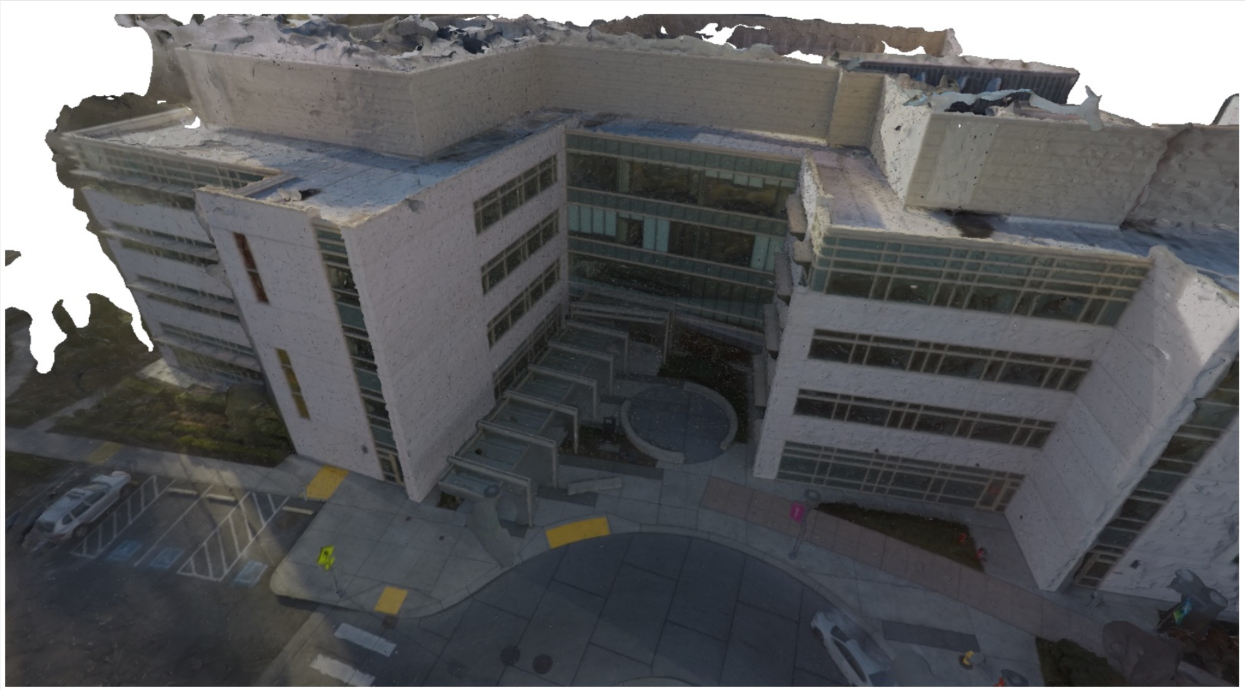
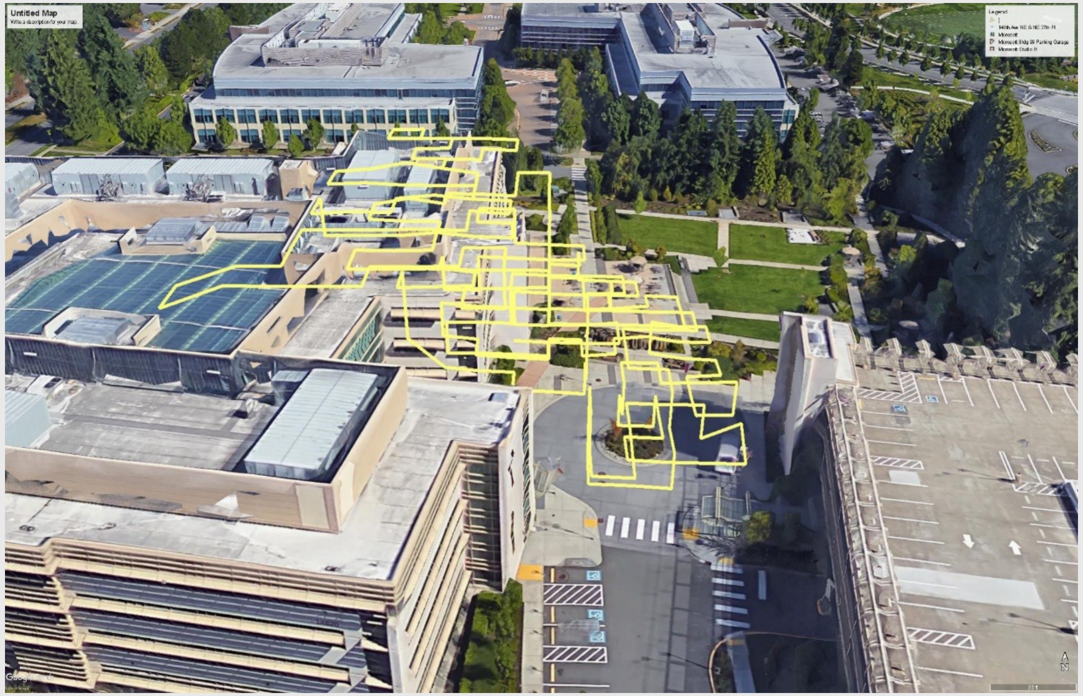
Untitled Map
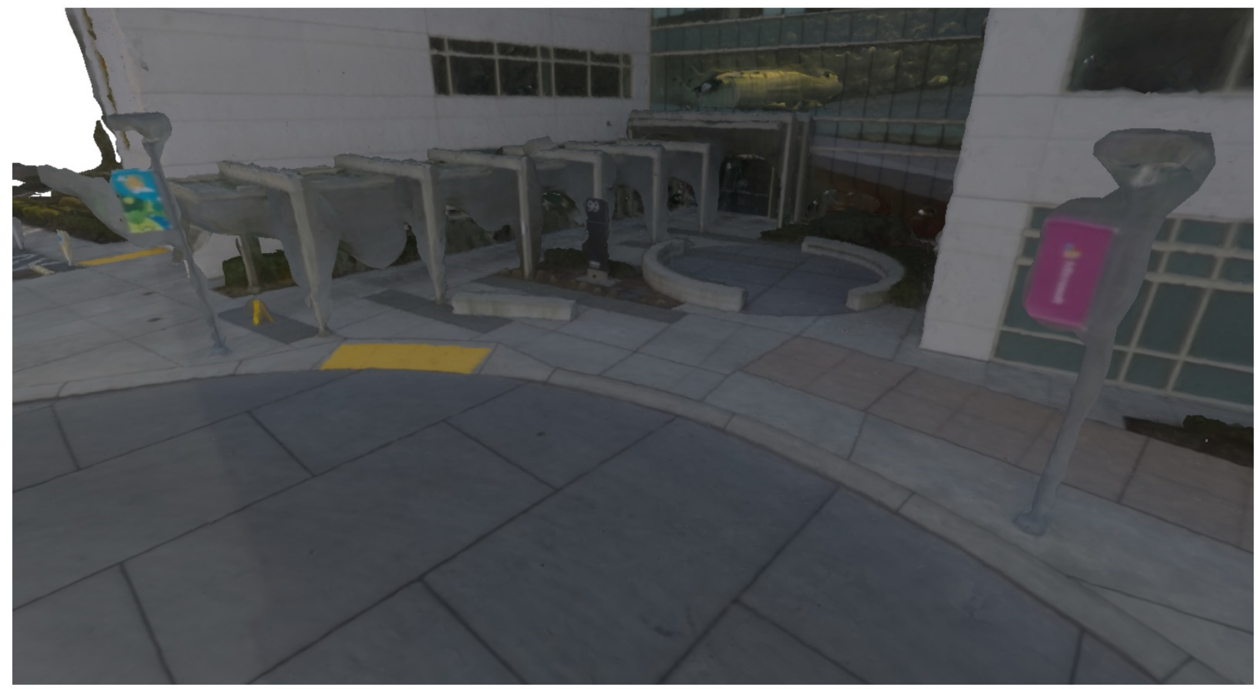

ExpLOre: Imitating an oracle

ExpLOre trains a policy to imitate the cost-to-go provided by an oracle

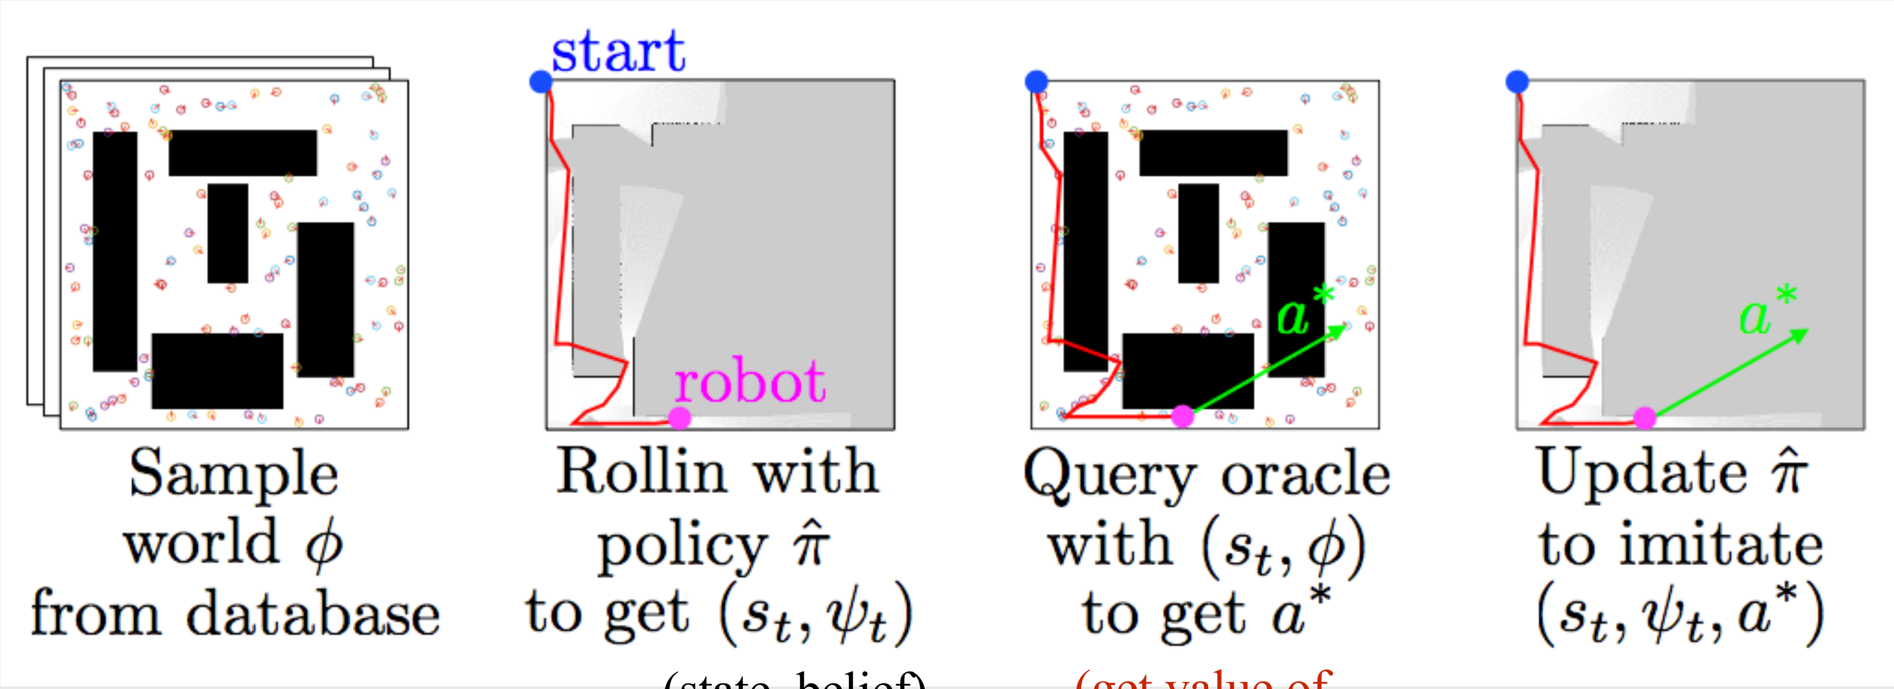

(state, belief)

(get value of action)

“Learning to Gather Information”, Choudhury et al. ICRA 2017

Example: 3D reconstruction of office scene

Train Data: Office desks created in Simulation

Test Data: RGBD data collected by [3]

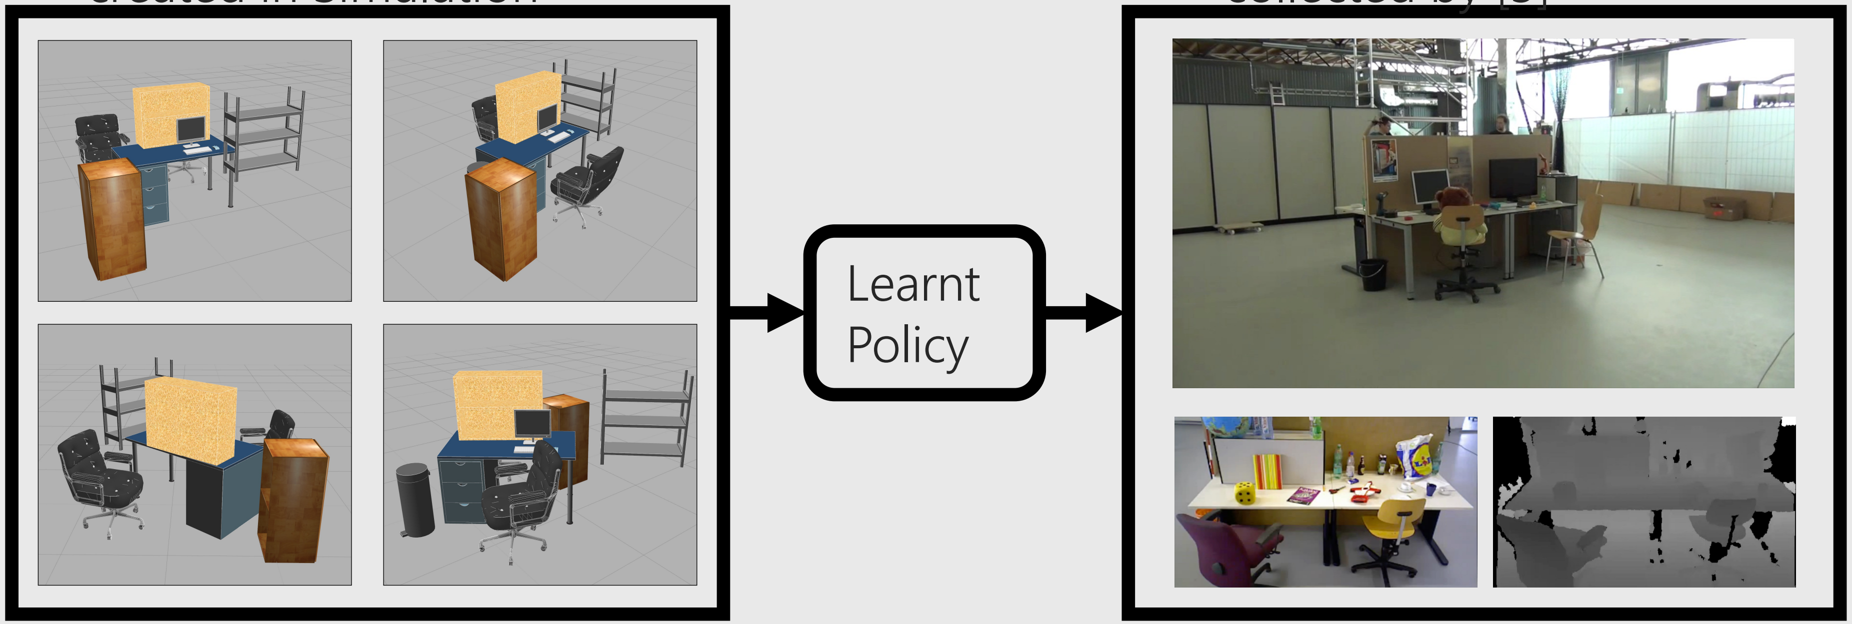

We train ExpLORE on synthetic data, and test on real data.

3. "A benchmark for the evaluation of rgb-d slam systems"
J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers. *IROS*, 2012

ExpLOre learns “desk exploring” policy

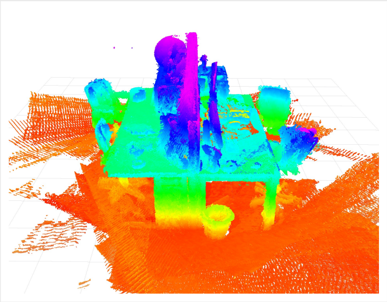

Ground
Truth
3D Model

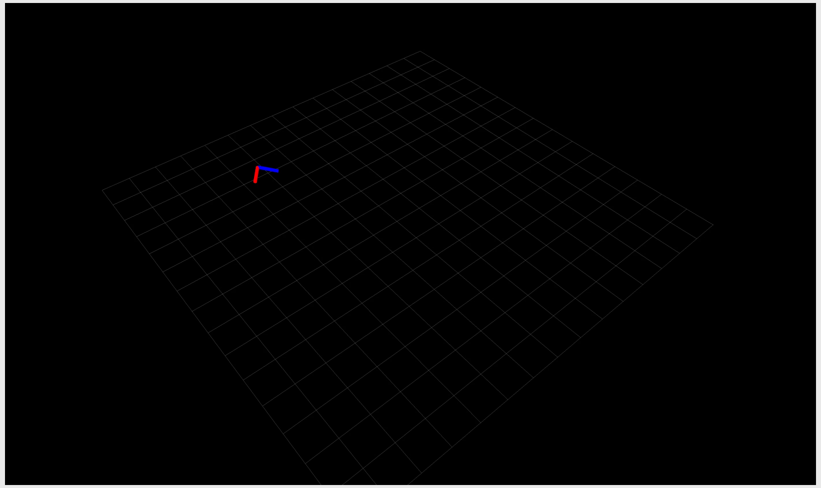

Occlusion Aware Heuristic

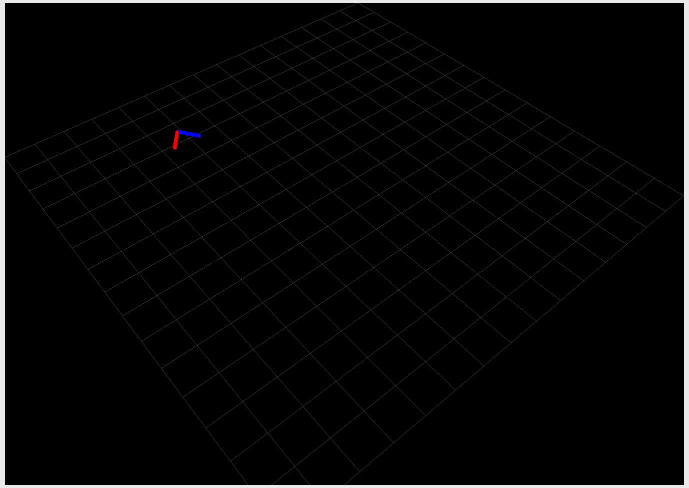

ExpLOre

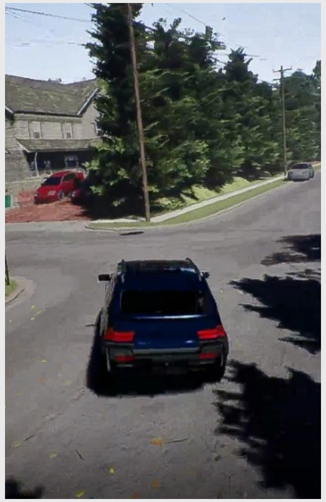

Retrain

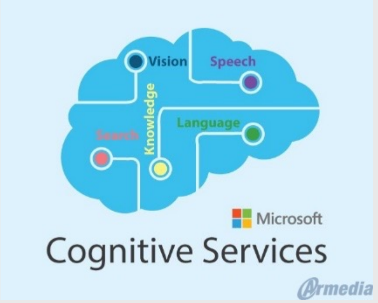

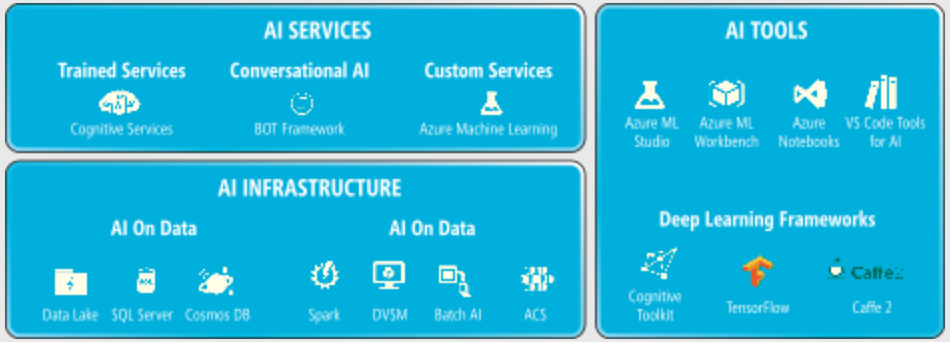

Identify Biases

Simulation

ML / AI Models

AI systems are often biased because the data used to train them is biased.

We use high-fidelity **simulations** to **diagnose biases** within ML classifiers.

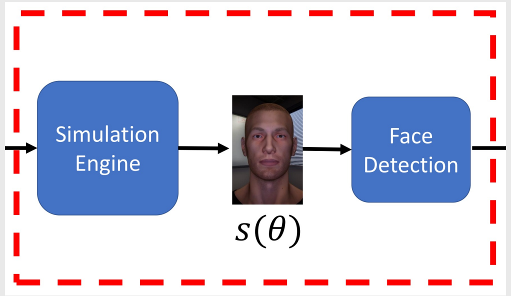

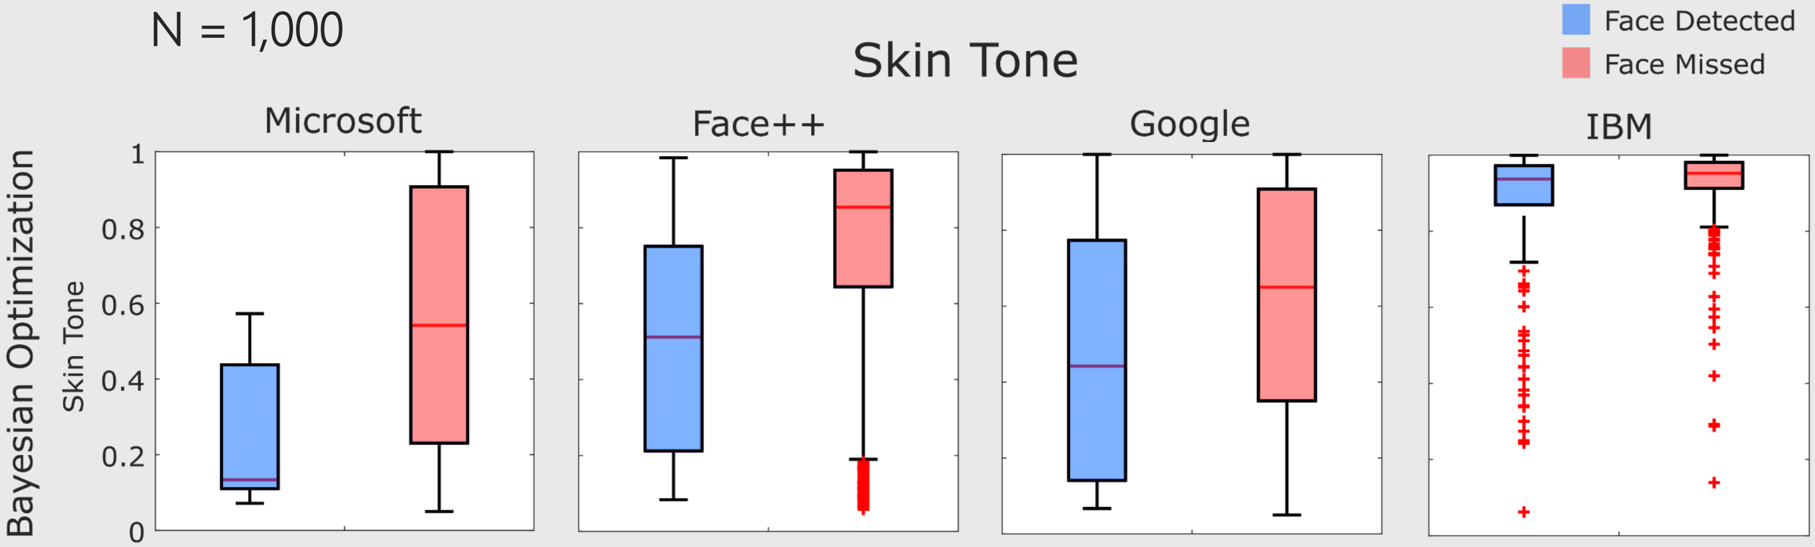

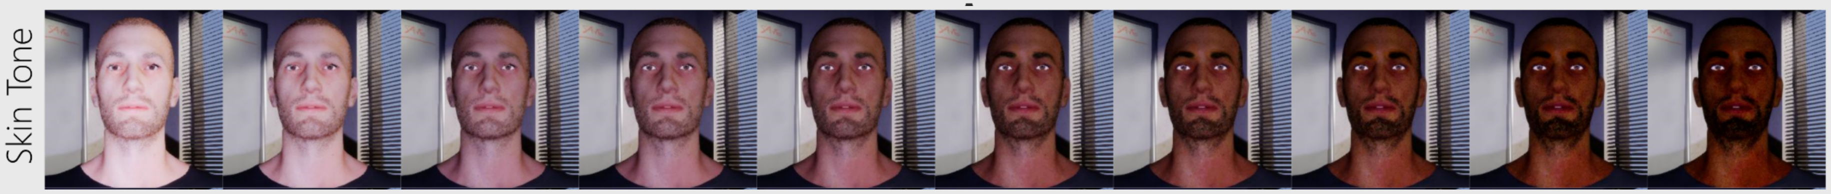

Testing, Verification and Robustness

Microsoft

Testing, Verification and Robustness

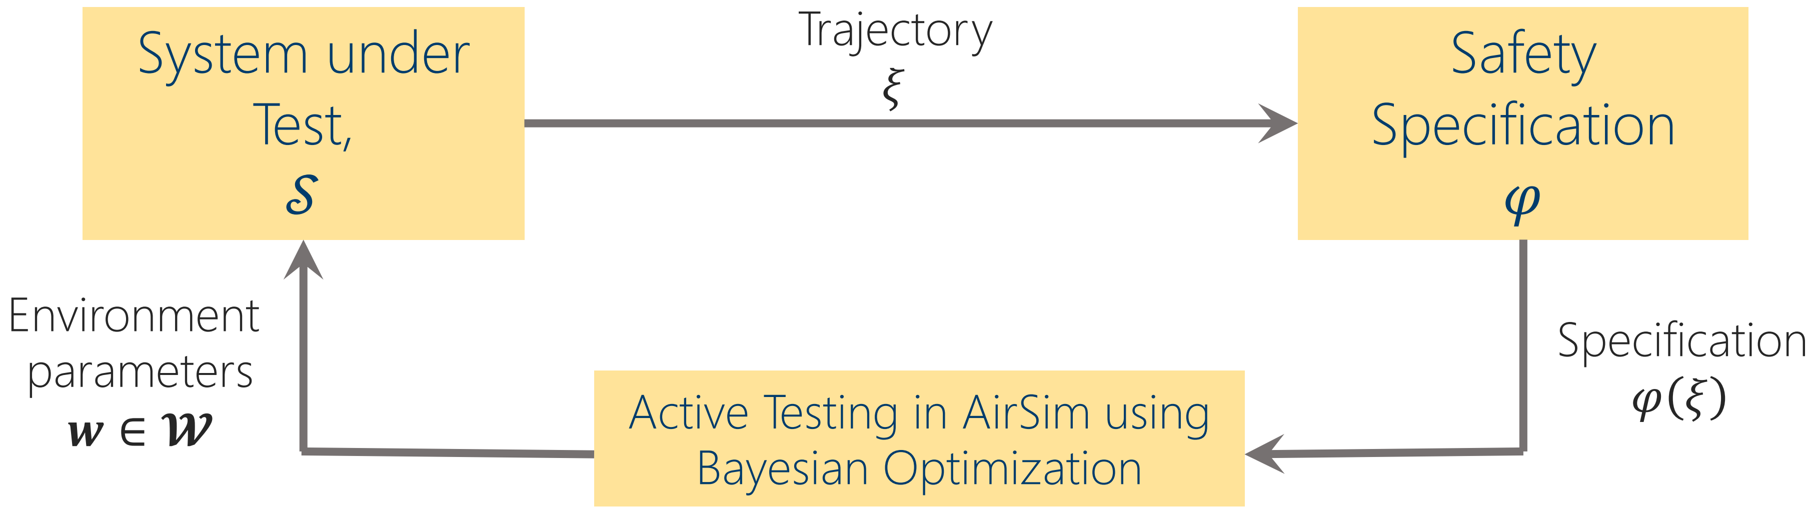

Ghosh et al. ICRA 2018

Adversarial Attacks on Optimization based Planners

$$f(x) = w_d J_d(x) + w_c J_c(x, a)$$

Path derivative cost

Collision cost

Adversary

This adversary (✖) can make the optimization problem harder.

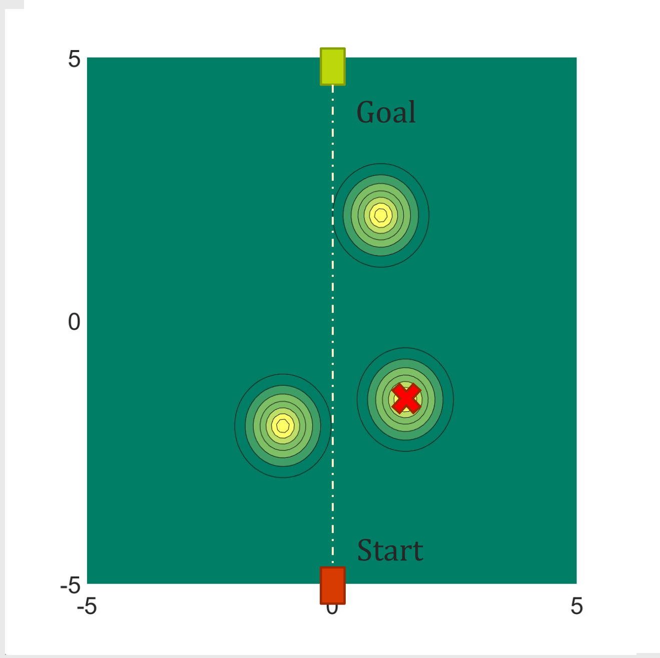

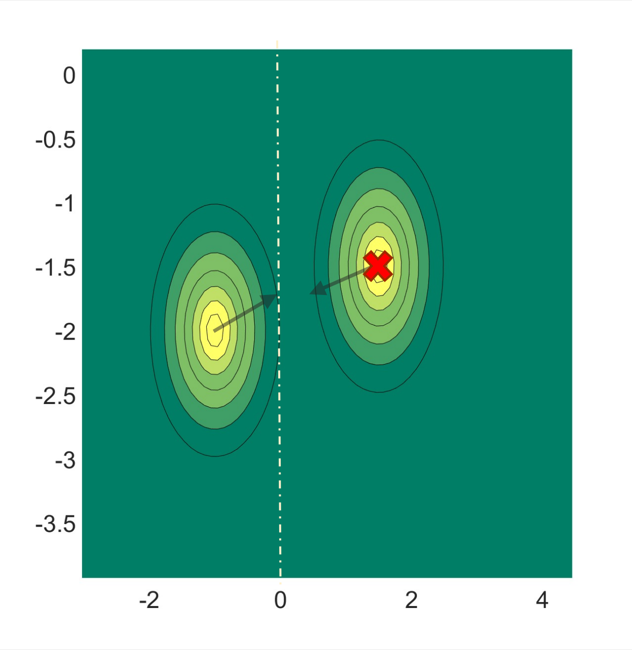

Poor condition number!

Vemprala et al. ICRA 2021

https://arxiv.org/abs/2011.00095

Adversarial Attacks on Optimization based Planners

$$f(x) = w_d J_d(x) + w_c J_c(x, a)$$

Path derivative cost — Collision cost

This adversary (✖) can make the optimization problem harder.

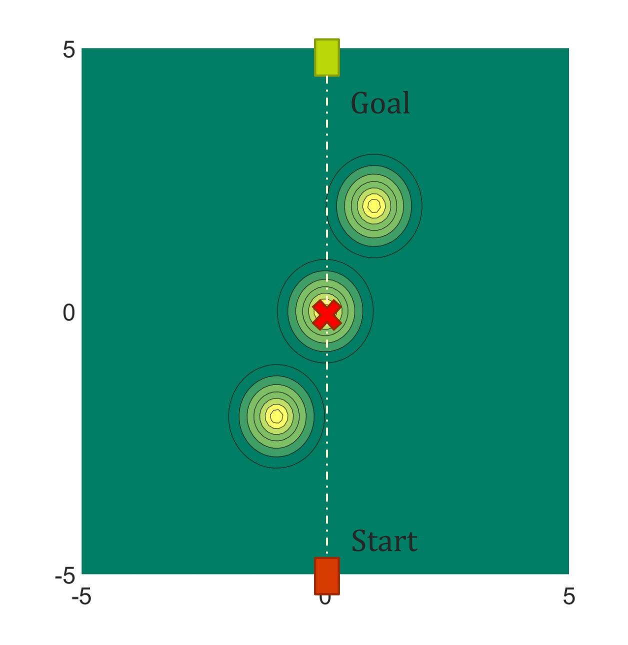

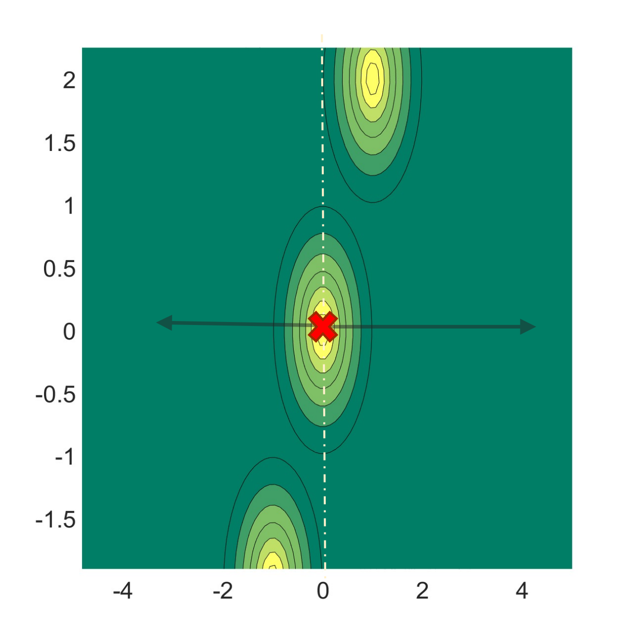

Poor condition number!

Vemprala et al. 2020
https://arxiv.org/abs/2011.00095

Black box attacks:

1. Model planner behavior by observing extent of deviation of trajectory for a given obstacle location.
   Fit a Gaussian Process between position of adversary and expected deviation

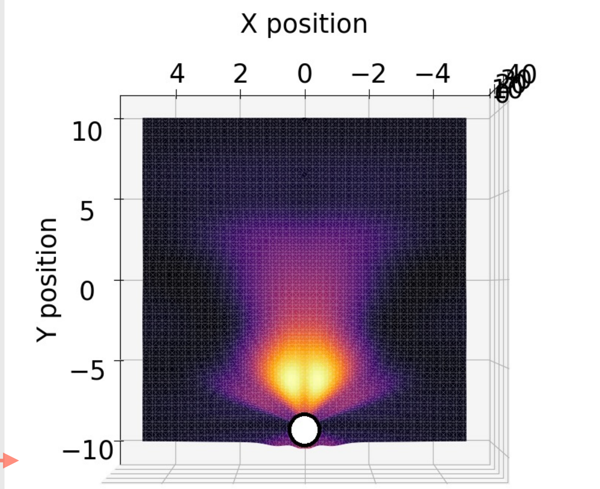

Adversary position vs. expected deviation

2. Find maximum of the above function to identify locations that can cause the most deviation.
   Use Bayesian Optimization to compute optimal adversary position given target state

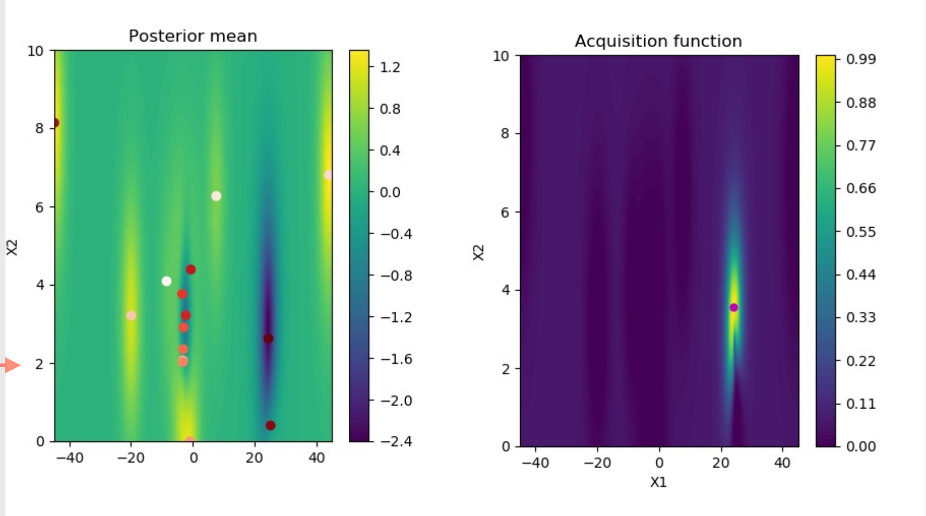

Optimal adversary position via Bayesian opt.

3DB Workflow

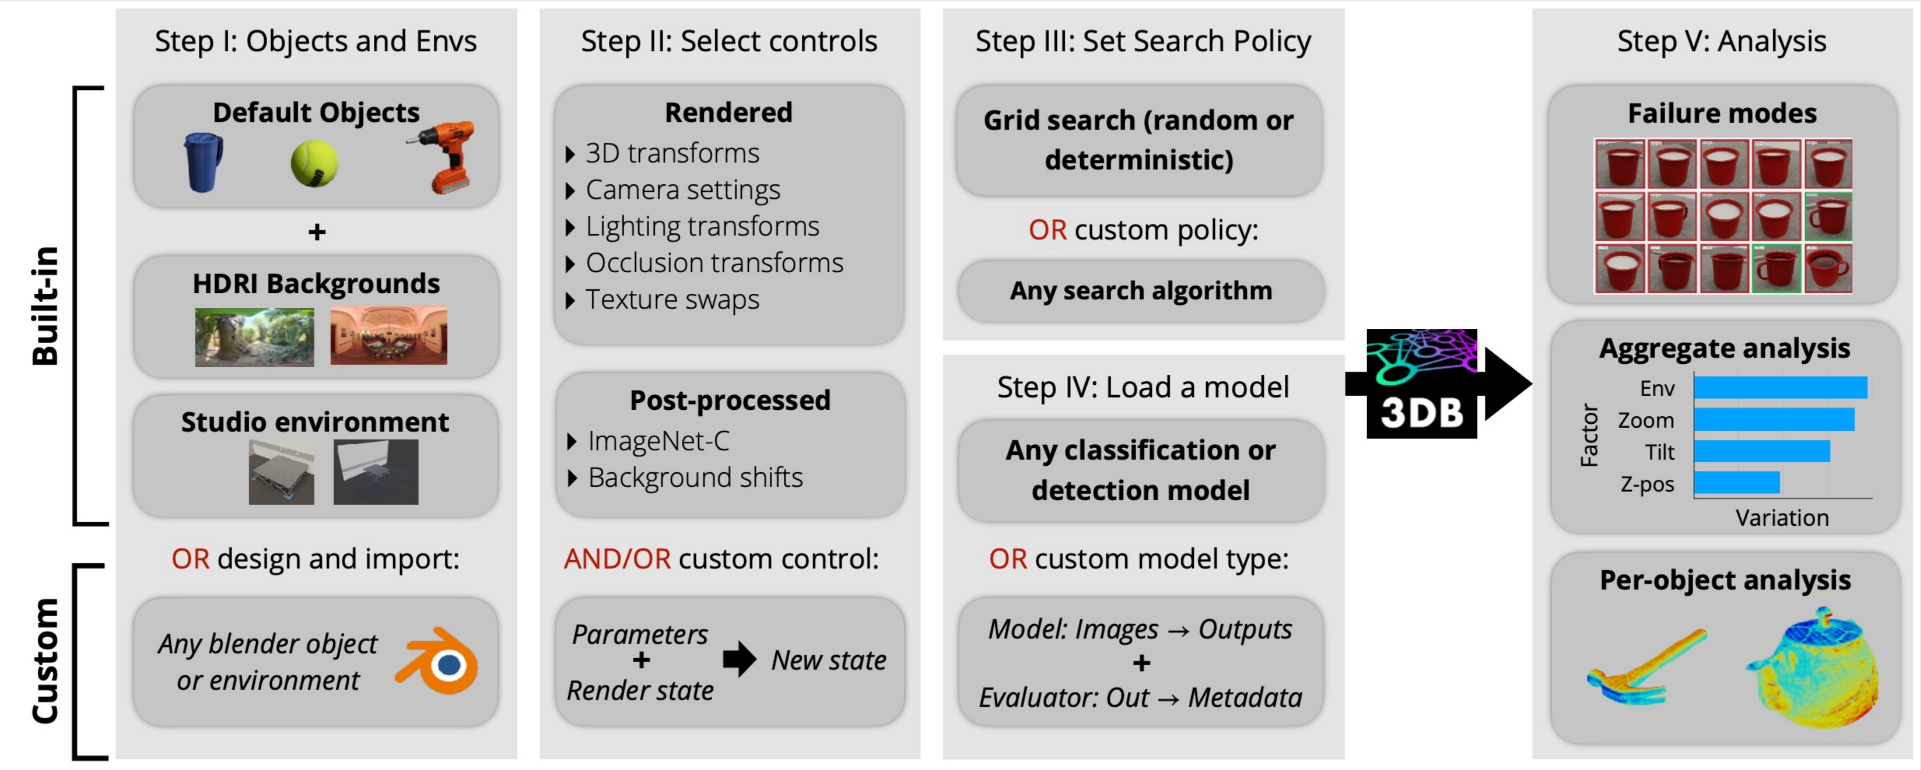

Example Images

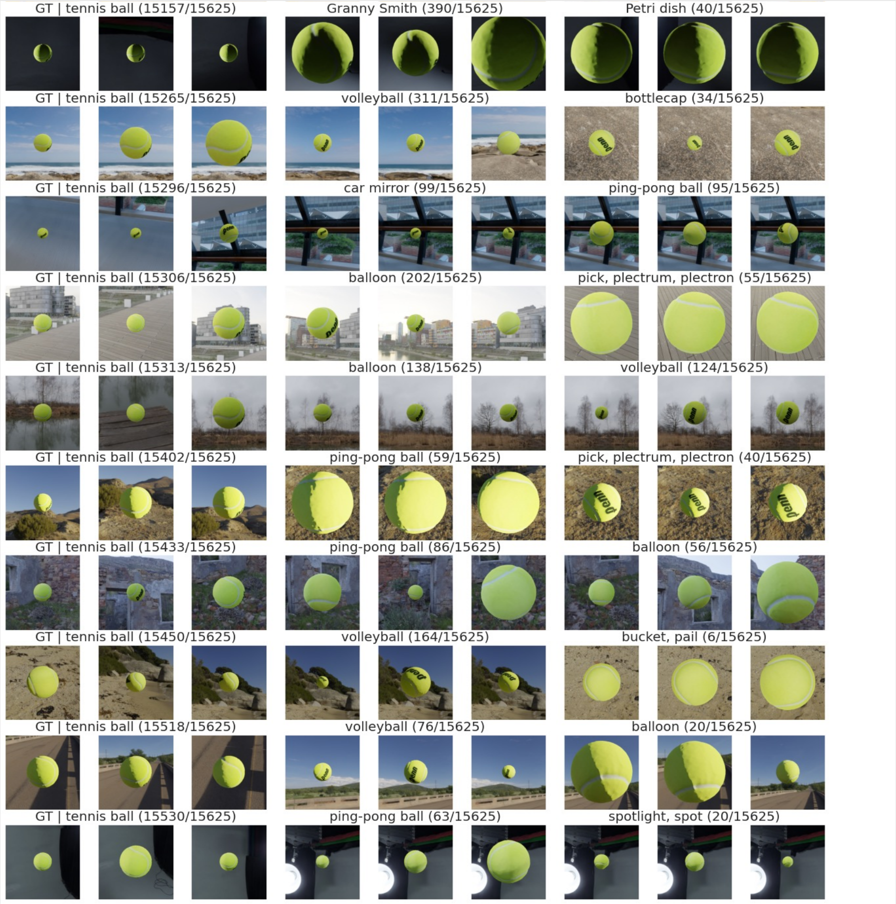

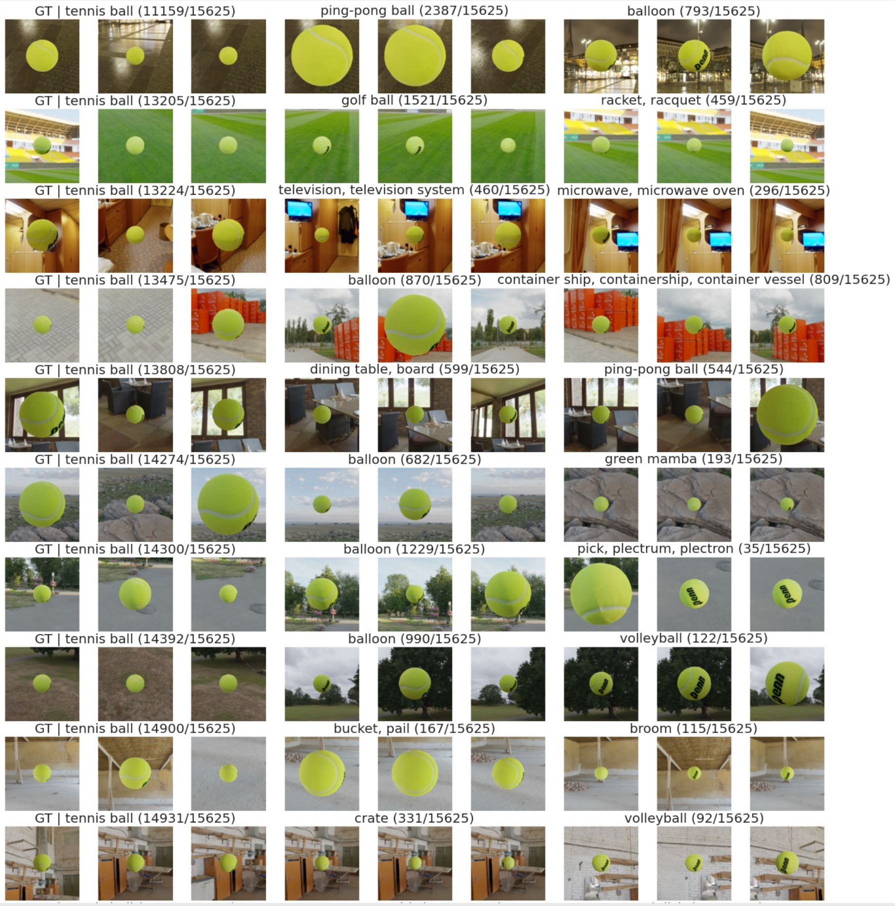

Case-Study: Coffee Mug

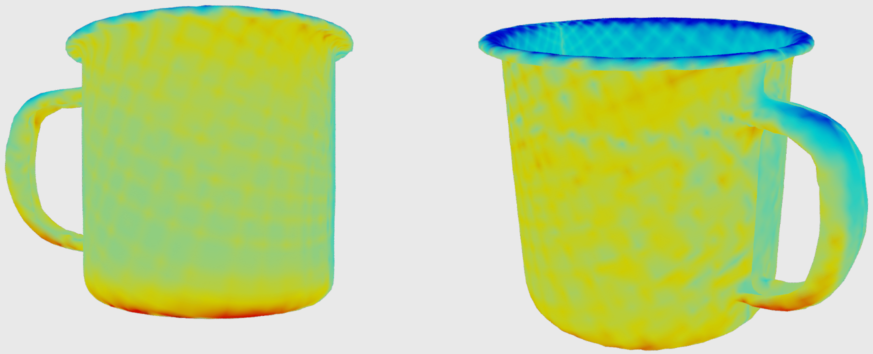

Class: Coffee Mug

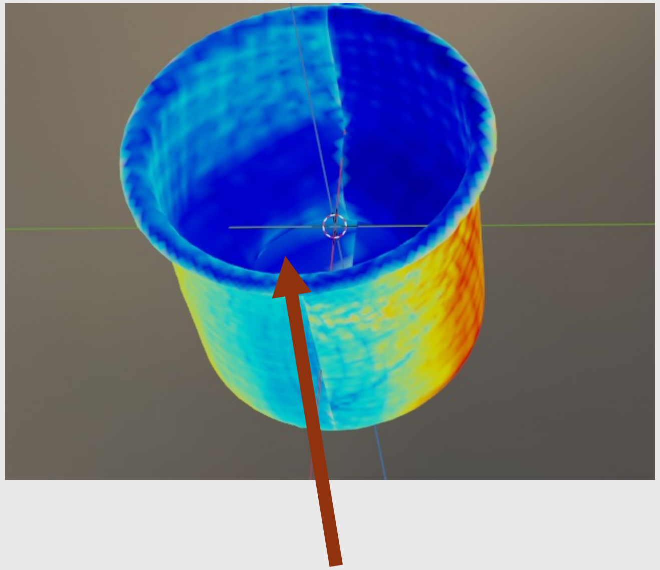

Wrong prediction of the cup if the inside is appearing

Hypothesis: classifier relies on the contents of the mug to correctly classify it

Can we leverage 3DB to **confirm or refute** this hypothesis?

Case-Study: Coffee Mug

Experiment: Fill the mug with various fluids and check predictions

NN prediction

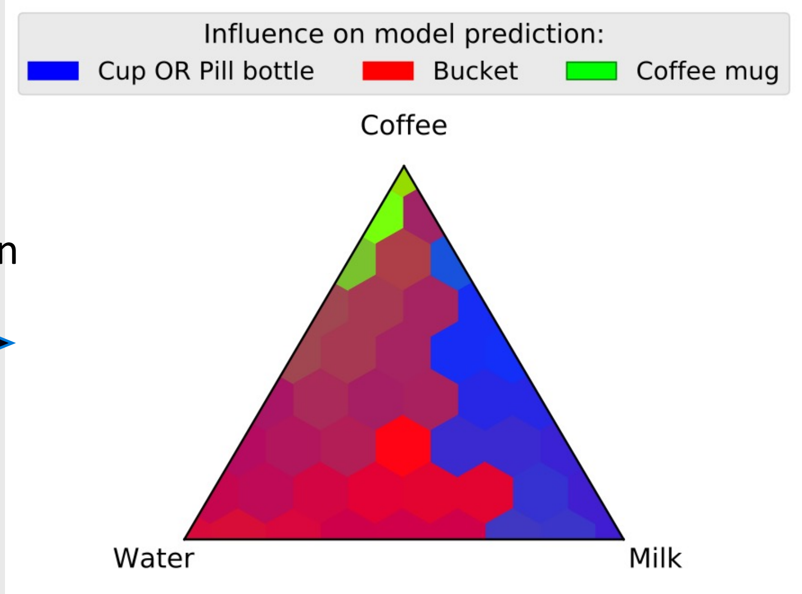

Indeed the mug's prediction **relies on the fluid** inside!

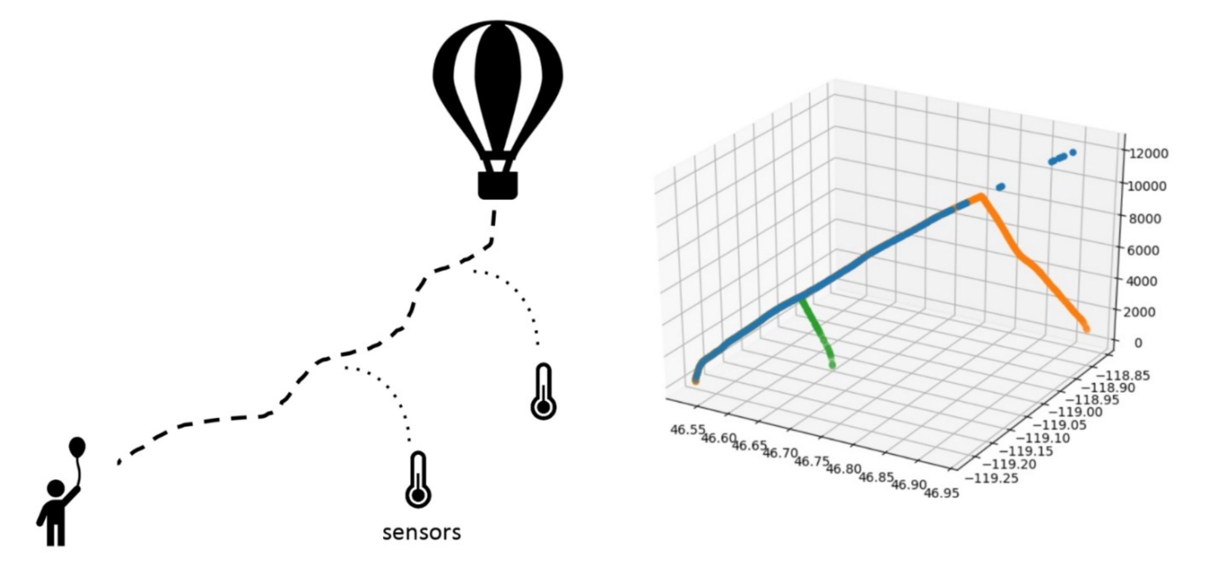

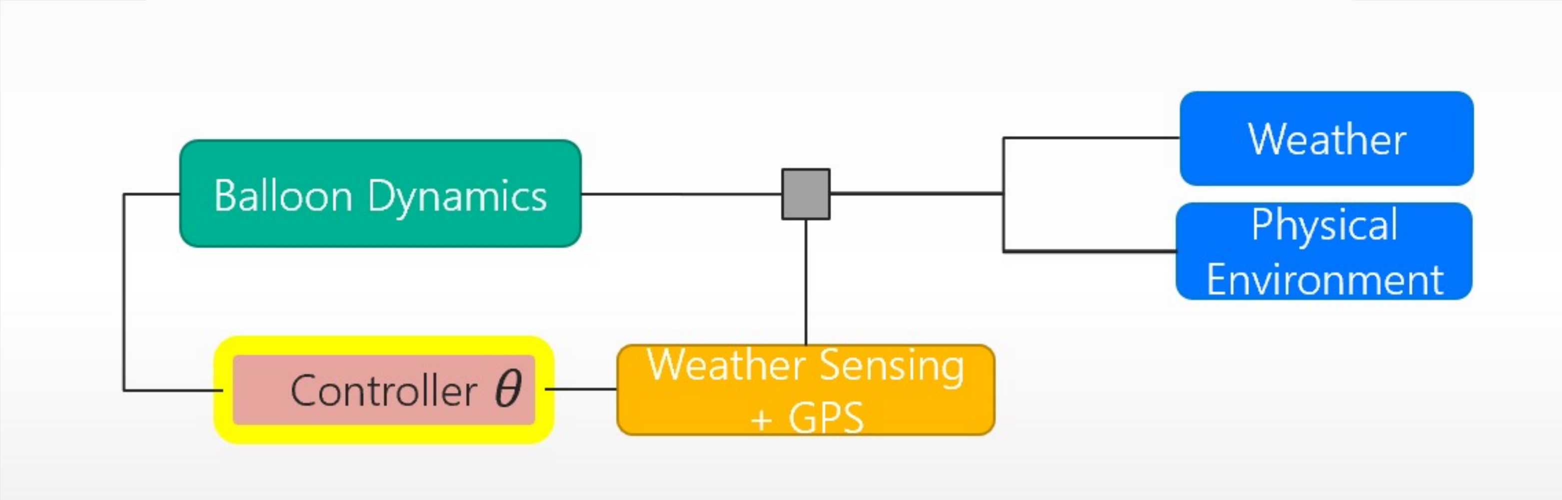

Stratospheric Weather Modeling with Active Sensing ^

Instrument Landing System

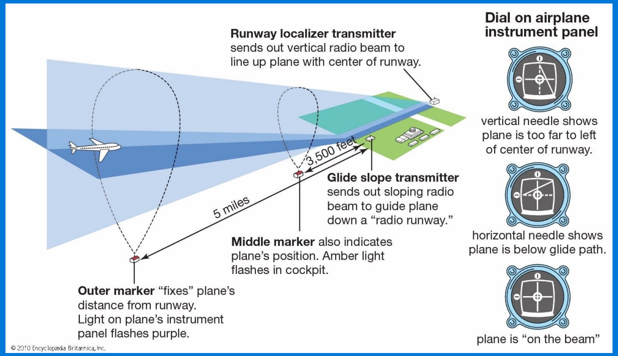

Fewer than 20% Airports have any kind of ILS, Expensive, Requires trained human-in-loop

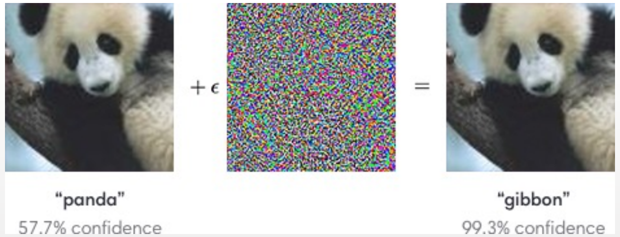

Attack

$$\hat{x} = argmax_{x' \in N(x,\epsilon)} L(F(x'), y)$$

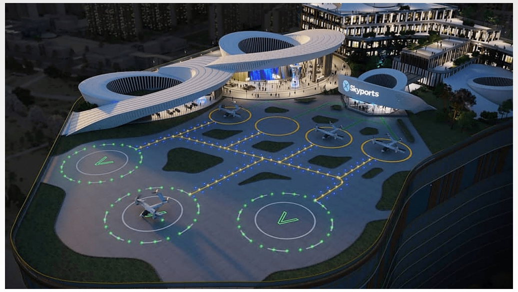

Boosting

$$\hat{x} = argmin_{x' \in Env(x,\theta)} L(F(x'), y)$$

Landing Pad Texture

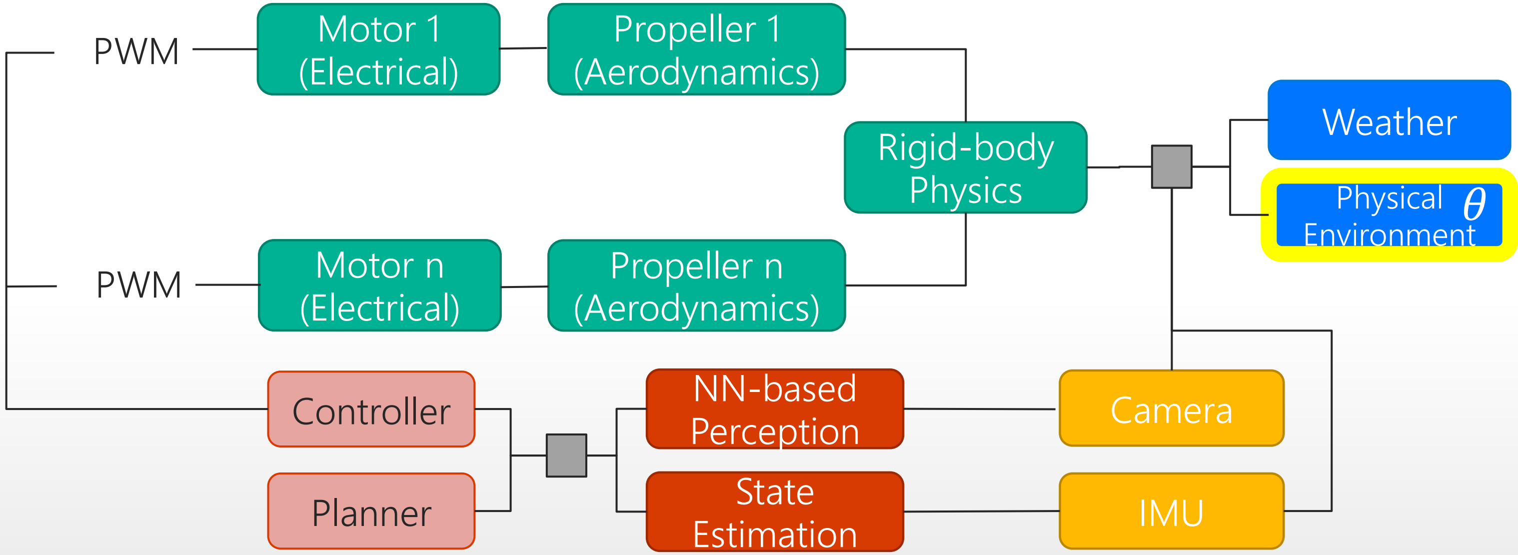

Unadversarial Examples: Designing Real-World Objects to Aid Neural Networks

Intuition: ML models love “non-robust features"-> let’s use those features!

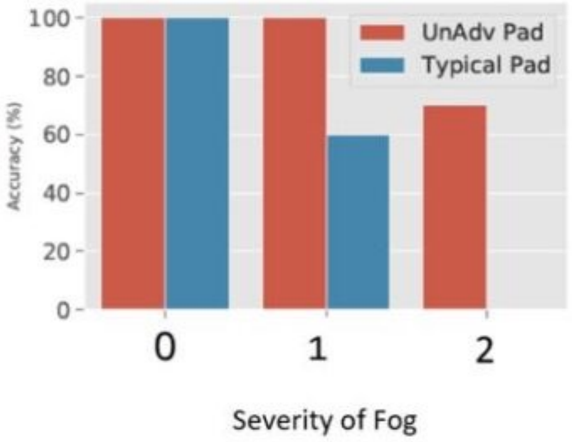

Plane with patch (78% accuracy)

Car with patch (96% accuracy)

Unadversarial

Normal

Plane without patch (67% accuracy)

Car without patch (46% accuracy)

See Salman et al. [21] for details

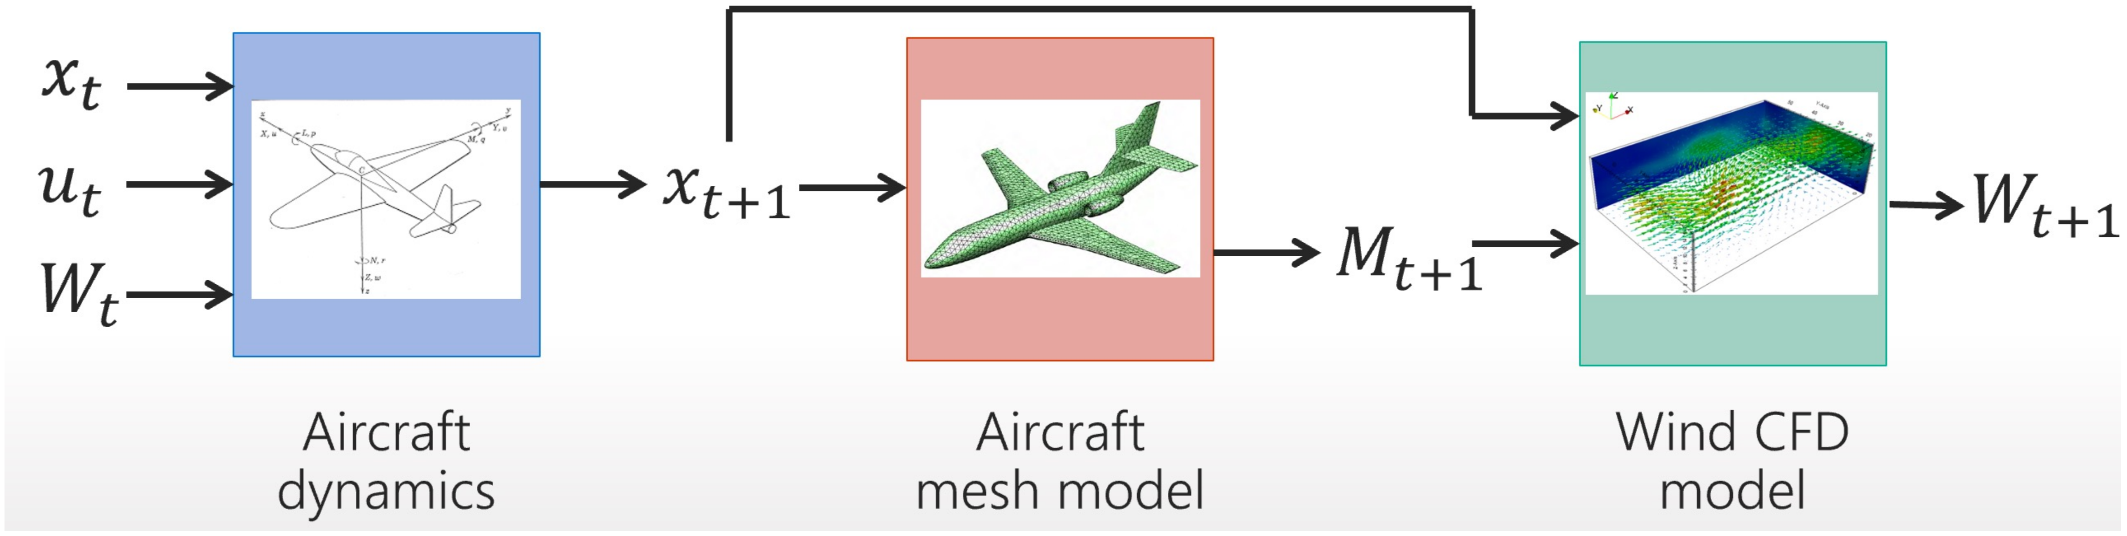

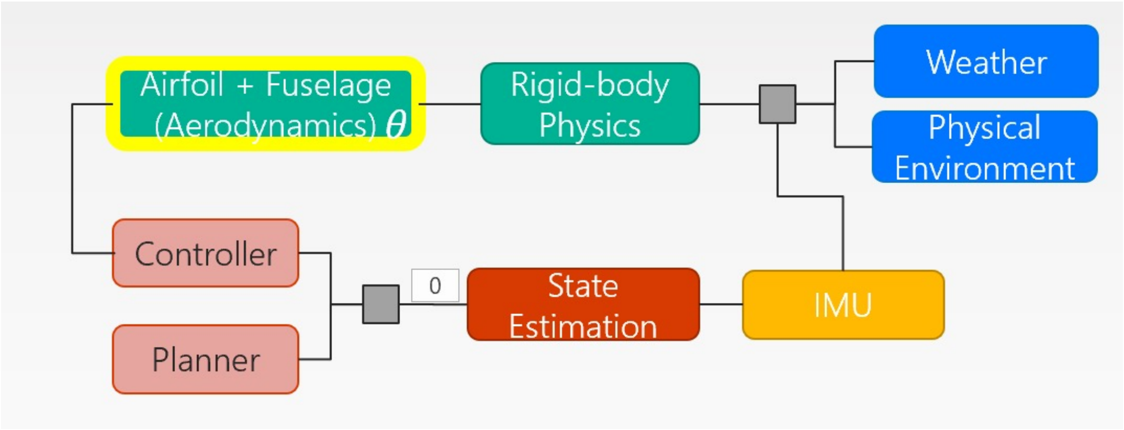

Smart Shape Aerodynamic Models*

Example: AgIoT + ML

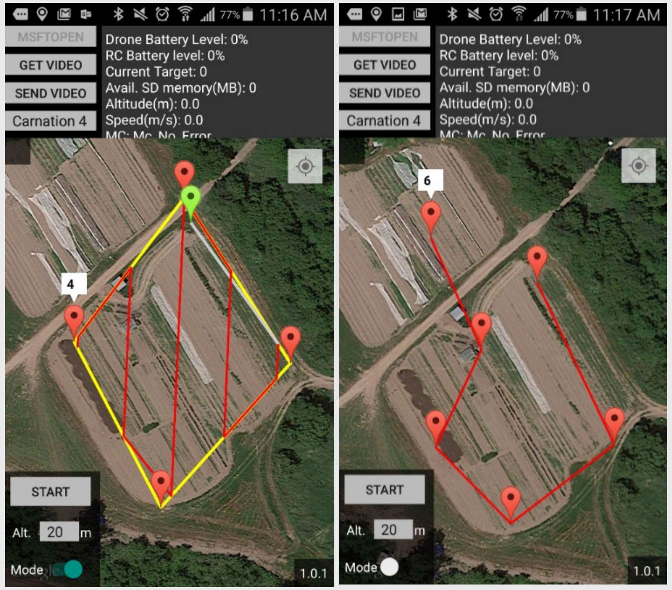

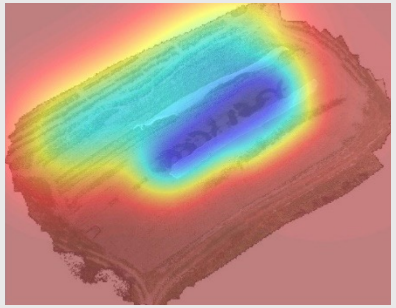

Farmbeats, Chandra et al.

Microsoft

CES Display

AirSim visualization/simulation with real-time data.

Microsoft

CES Display

Bell’s CES live display