In this Lecture

- Two parts:
  1. Further assignment problems
     - uncertainty in assignment
     - assignment under obfuscation
  2. State-of-the-art and outlook
  3. Final announcements

Mobile Robot Systems

Lecture 10-I: Further Assignment Problems

Dr. Amanda Prorok

asp45@cam.ac.uk

www.proroklab.org

UNIVERSITY OF CAMBRIDGE

The Hungarian Algorithm

Recall Lecture 8

- Assumptions when using an assignment algorithm such as the Hungarian metho
  - Costs (utilities) are known at a **centralized** computation unit.
  - Costs (utilities) are **deterministic** (no noise).
  - Costs (utilities) do not change (**constant**).
  - **1-to-1** assignment (one robot per task, one task per robot).
- Complications:
  - Uncertainty around true utility $U(i,j)$
  - Dynamic environment (changes in utility / agents)
  - Robot / task dependencies (robot heterogeneity / redundancy).

all of these issues are very common in robotics!!

- Consequences:
  - Sub-optimality
  - Problems can become NP-hard (for combinatorial matching problems)
  - Practically infeasible (centralized solutions may not be possible)

Assignment under Uncertainty

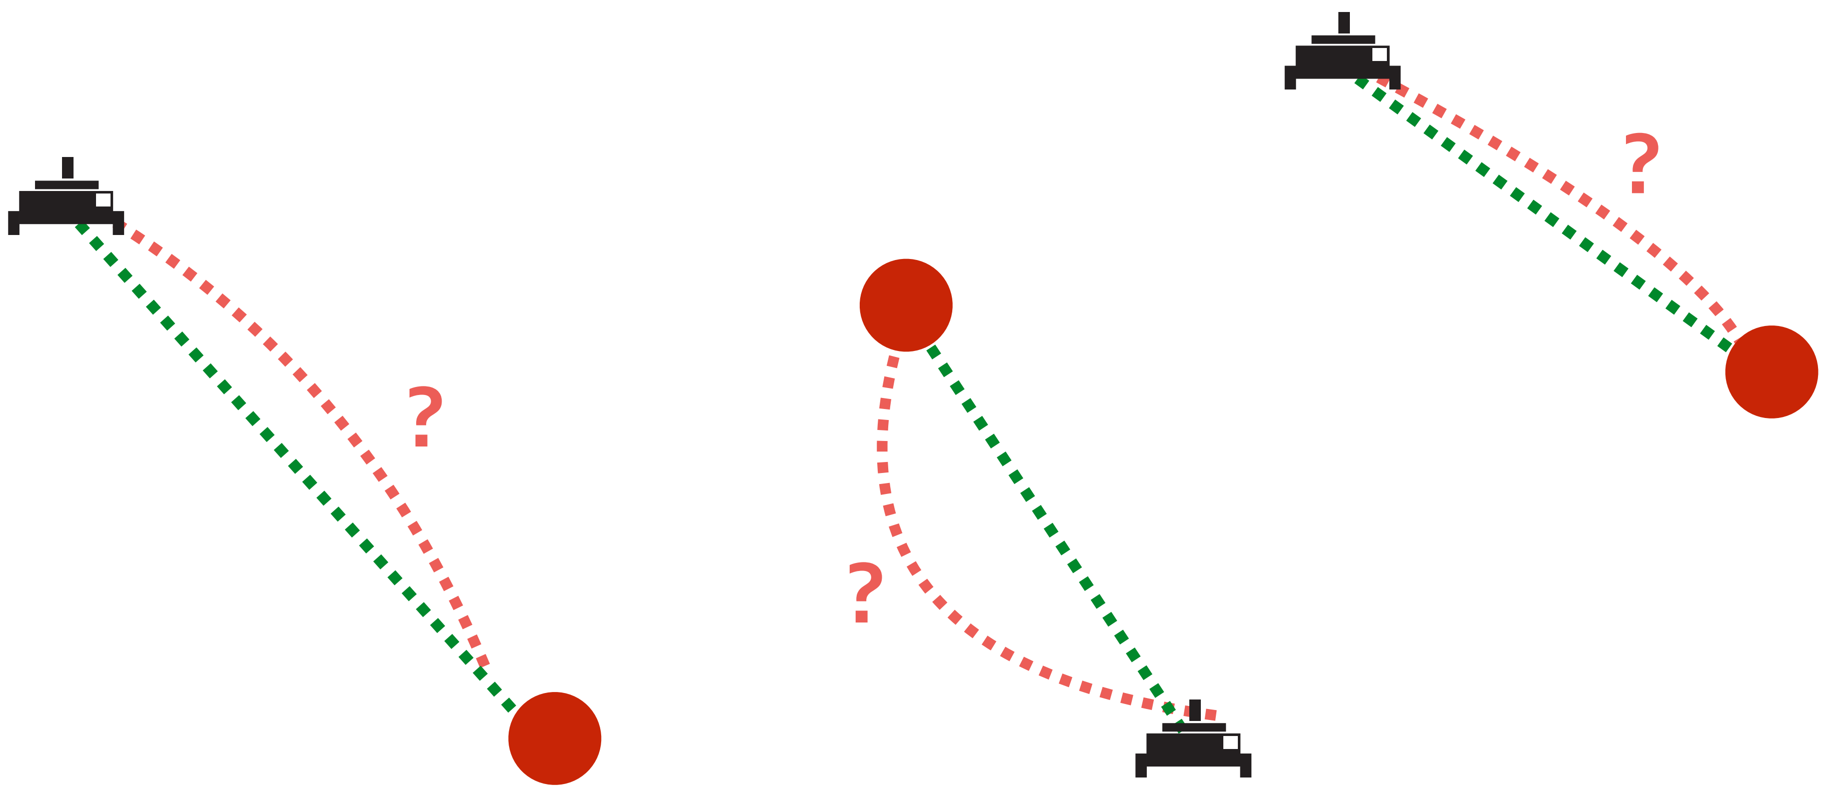

Min. overall assignment cost / **Max.** overall assignment utility

Our focus: **minimize time** to reach destinations

uncertainty along robot travel paths

Assignment under Uncertainty

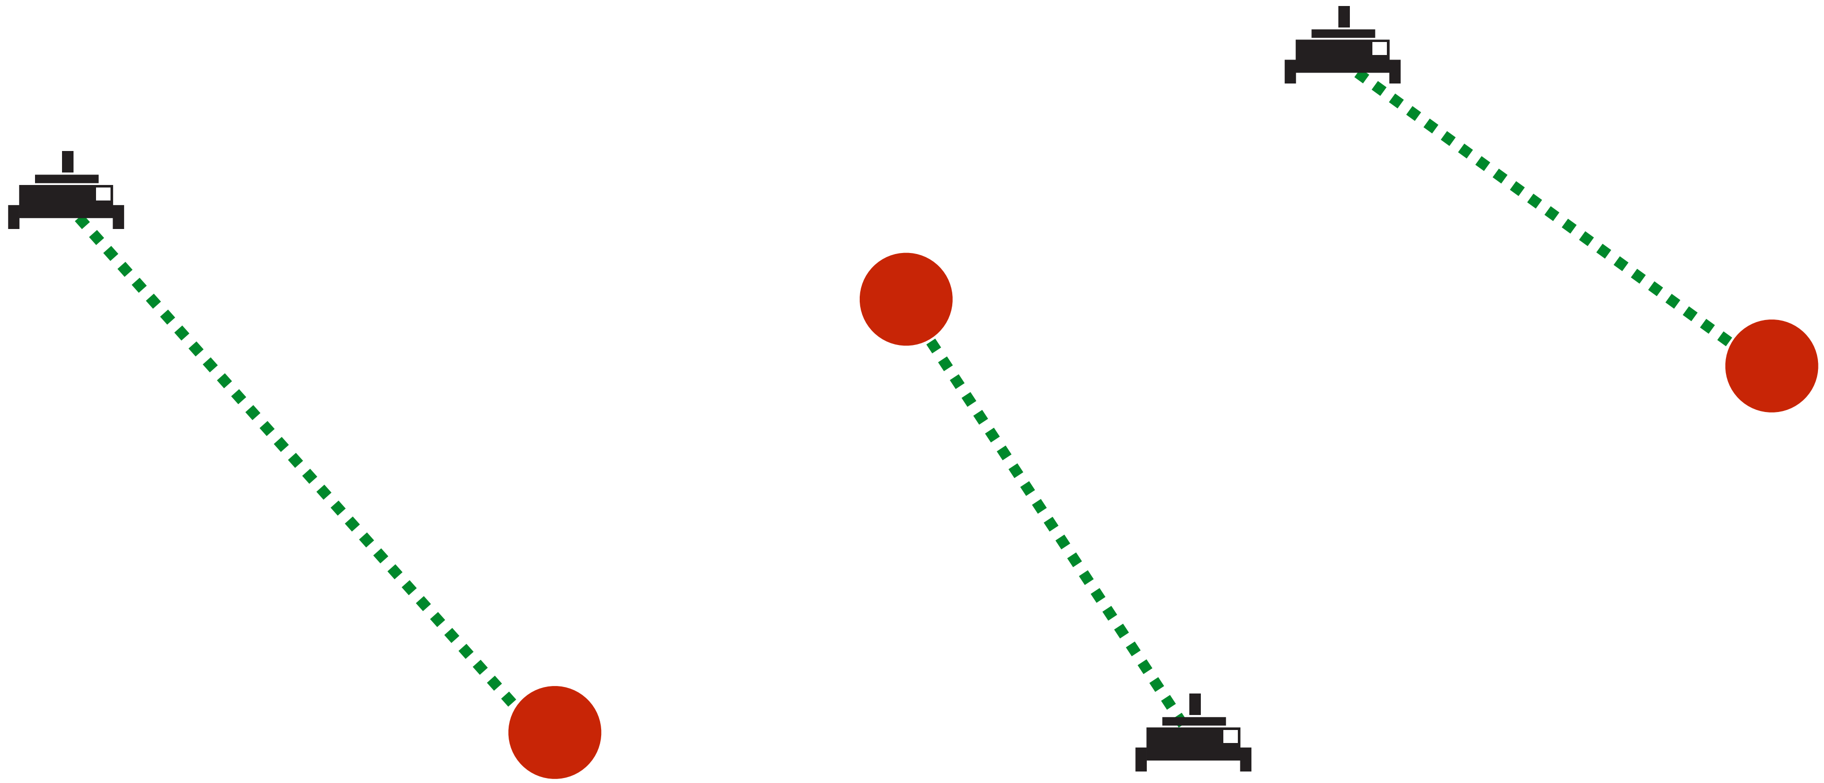

optimal assignment

Assignment under Uncertainty

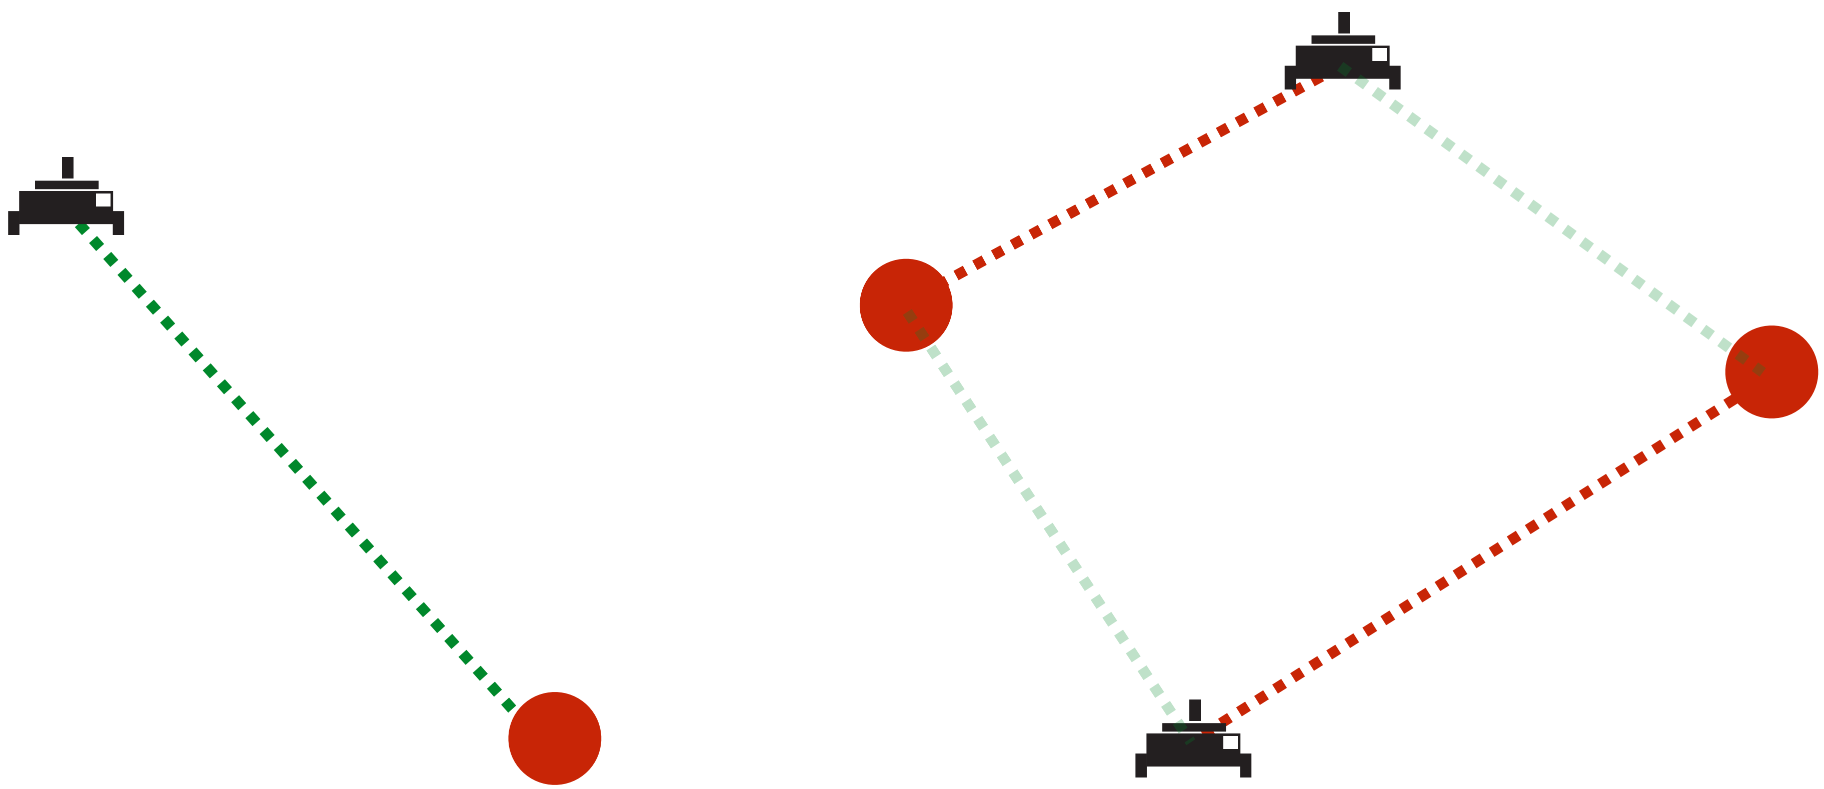

sub-optimal assignment

Assignment under Uncertainty

- [Mills-Tettey, Stentz, Dias; 2007]: dynamically repairs an initial assignment (Hungarian method)
  → dynamic re-assignment is potentially disruptive / expensive

- [Ponda, Johnson, How; 2012]: acceptable risk thresholds; chance-constrained allocation
  → relation of risk threshold to quality of solution as costs change

- [Nam, Shell; 2015], [Nam, Shell; 2017] : sensitivity analysis
  → determines when solutions cross acceptable risk preferences

- [Prorok; 2018]: redundant robots
  → compensate for loss of certainty with system redundancy

The Premise

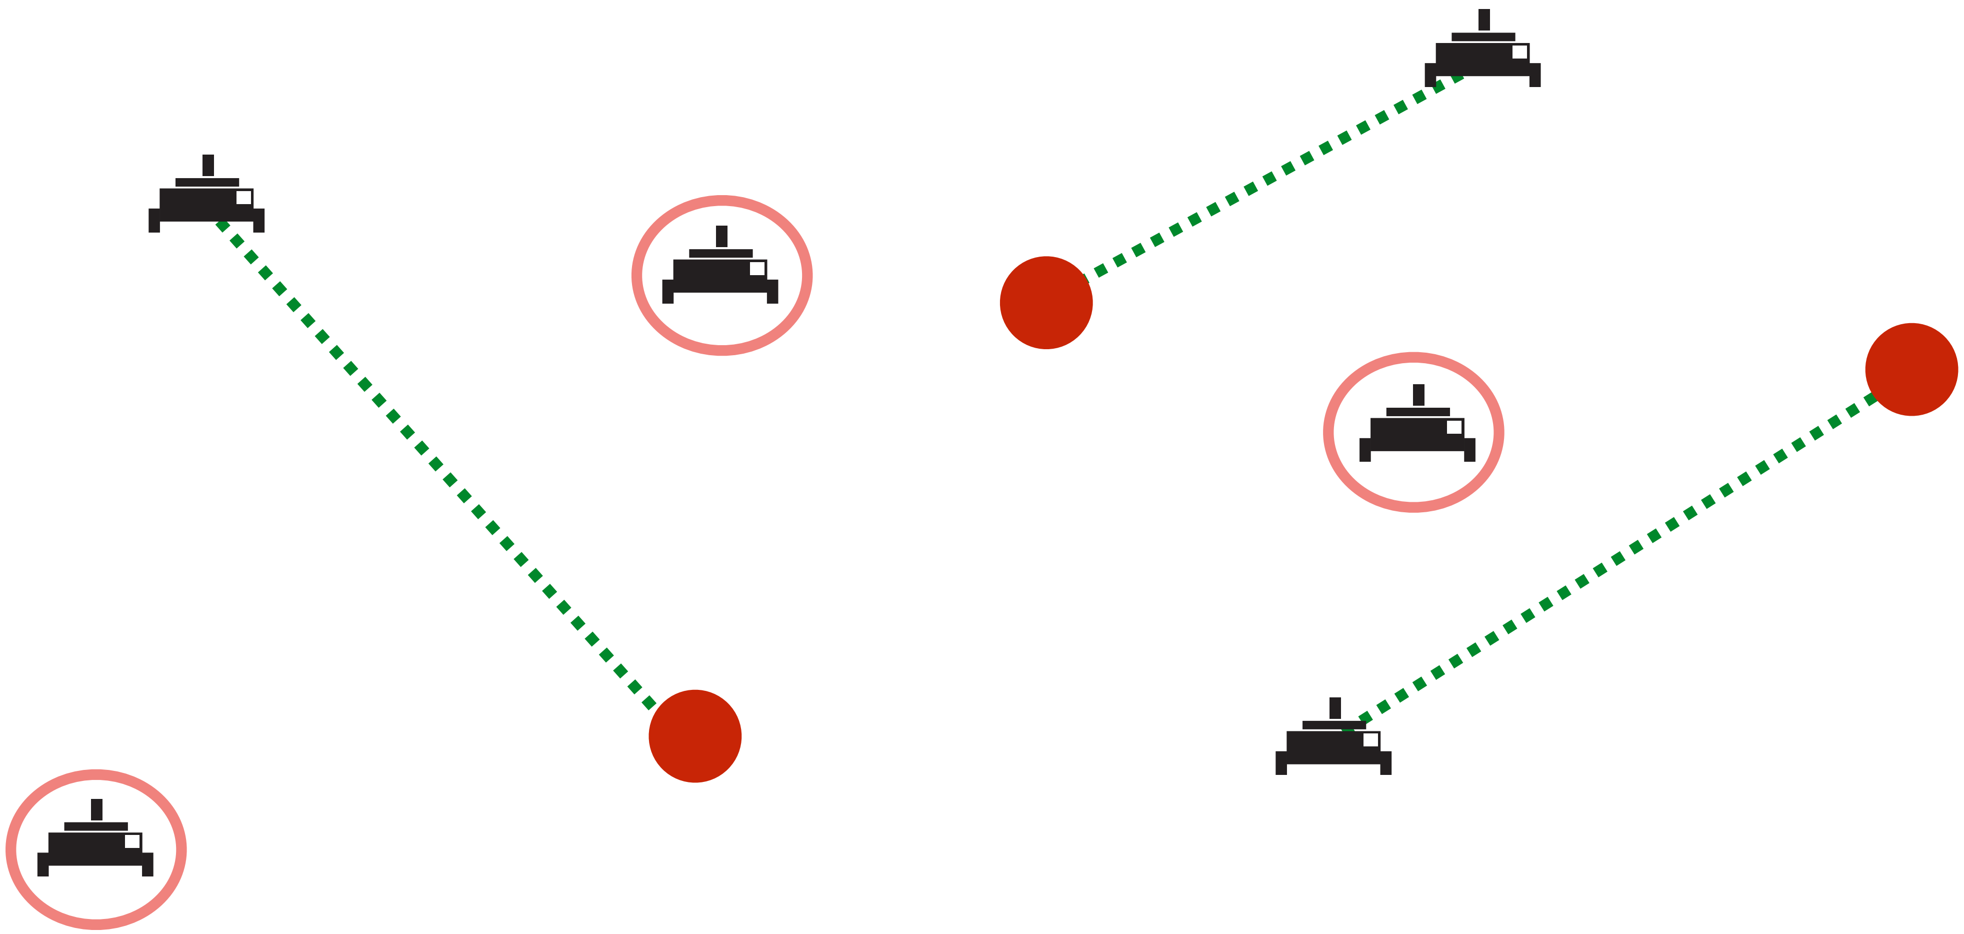

sub-optimal assignment

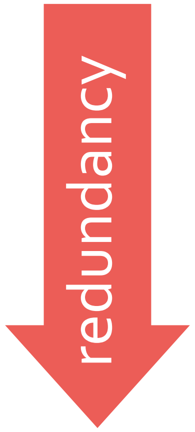

optimal assignment

Problem Statement

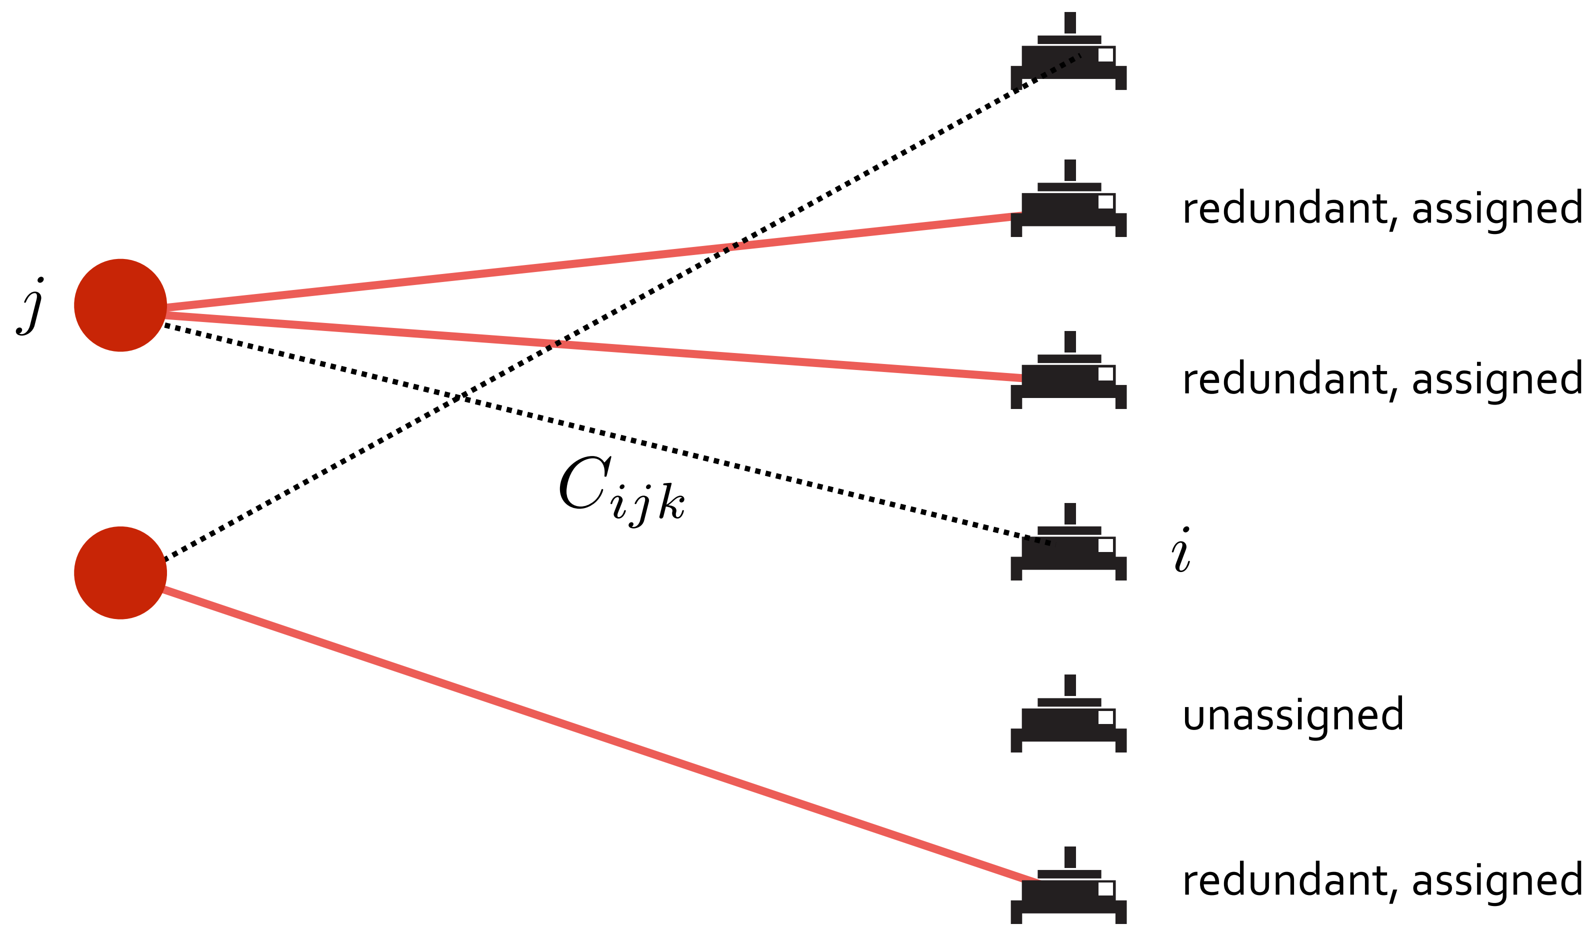

$$\mathcal{C} = \{C_{111}, \ldots, C_{NMK}\} \sim \mathcal{D}$$

N : total robots

M : goals

K : path options

N_{d} : deployable robots

$\mathcal{O}$: initial assignment

$\mathcal{A}$: redundant assign.

Redundant Assignment

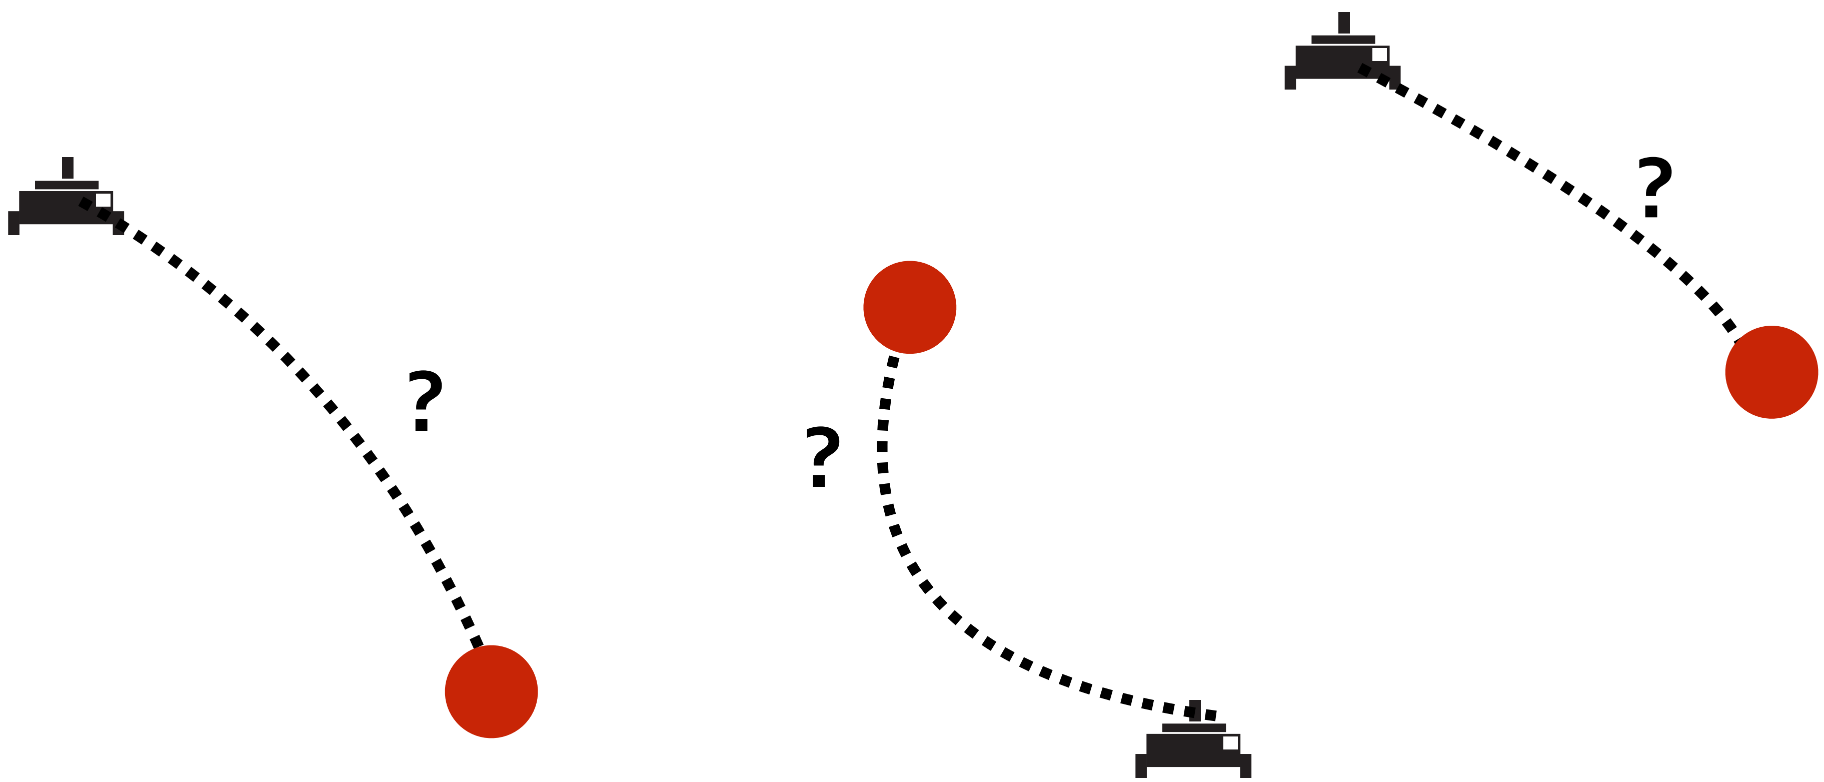

Redundant Assignment

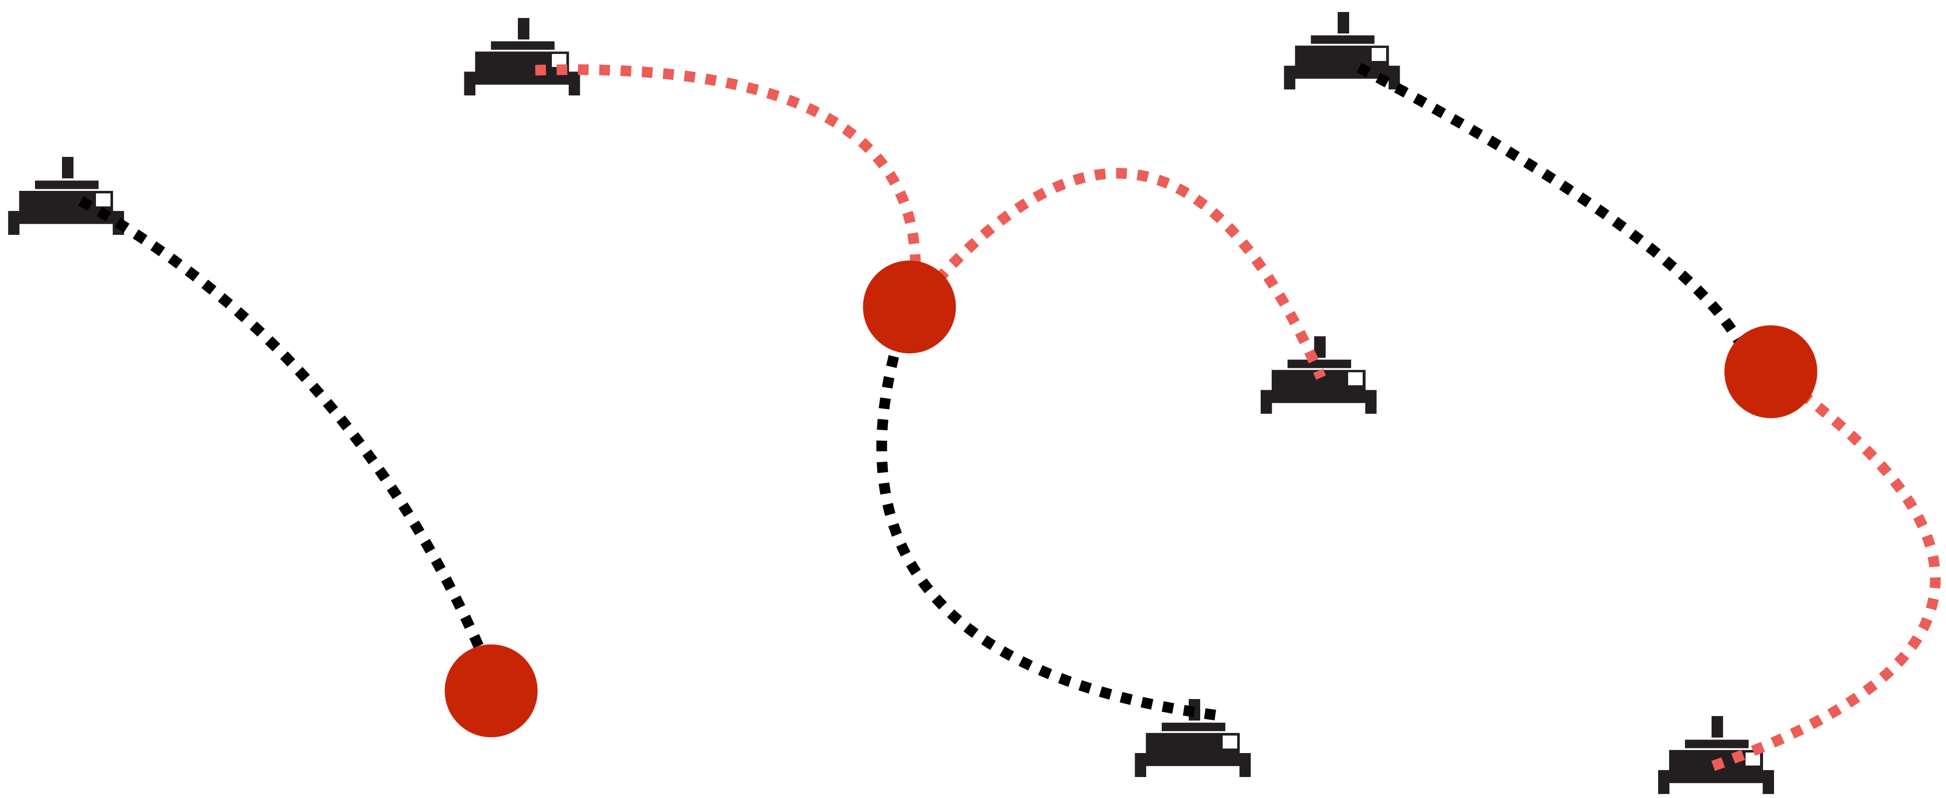

Time is the primary asset

Goal: minimize time to get to goals

Path Uncertainty

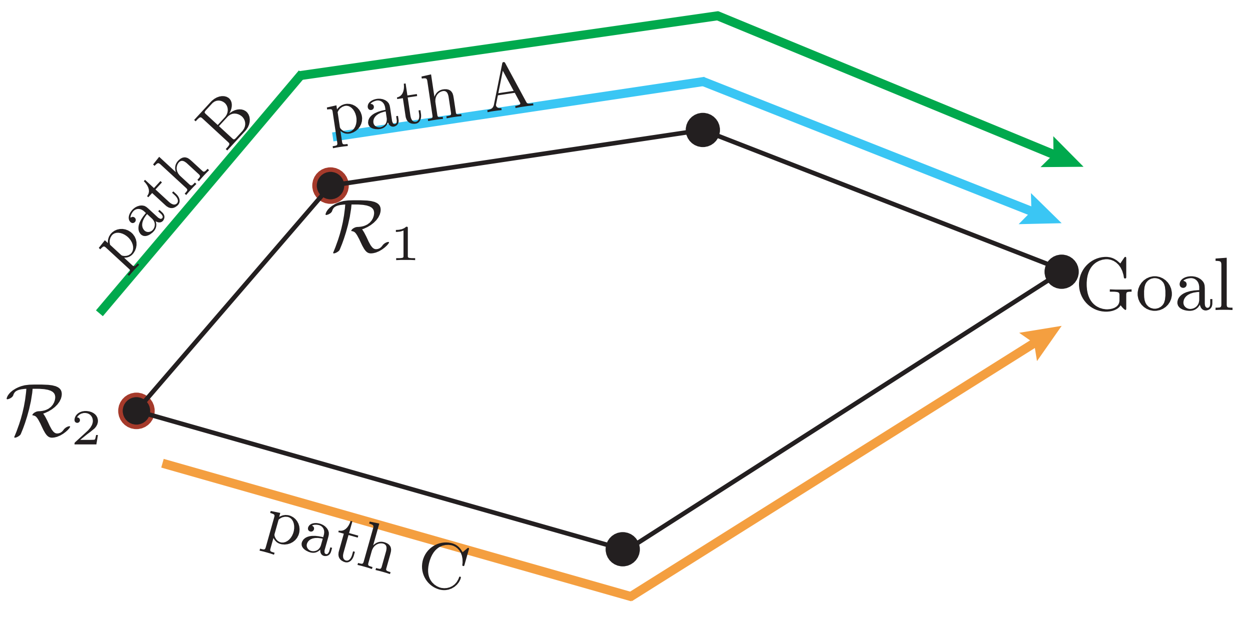

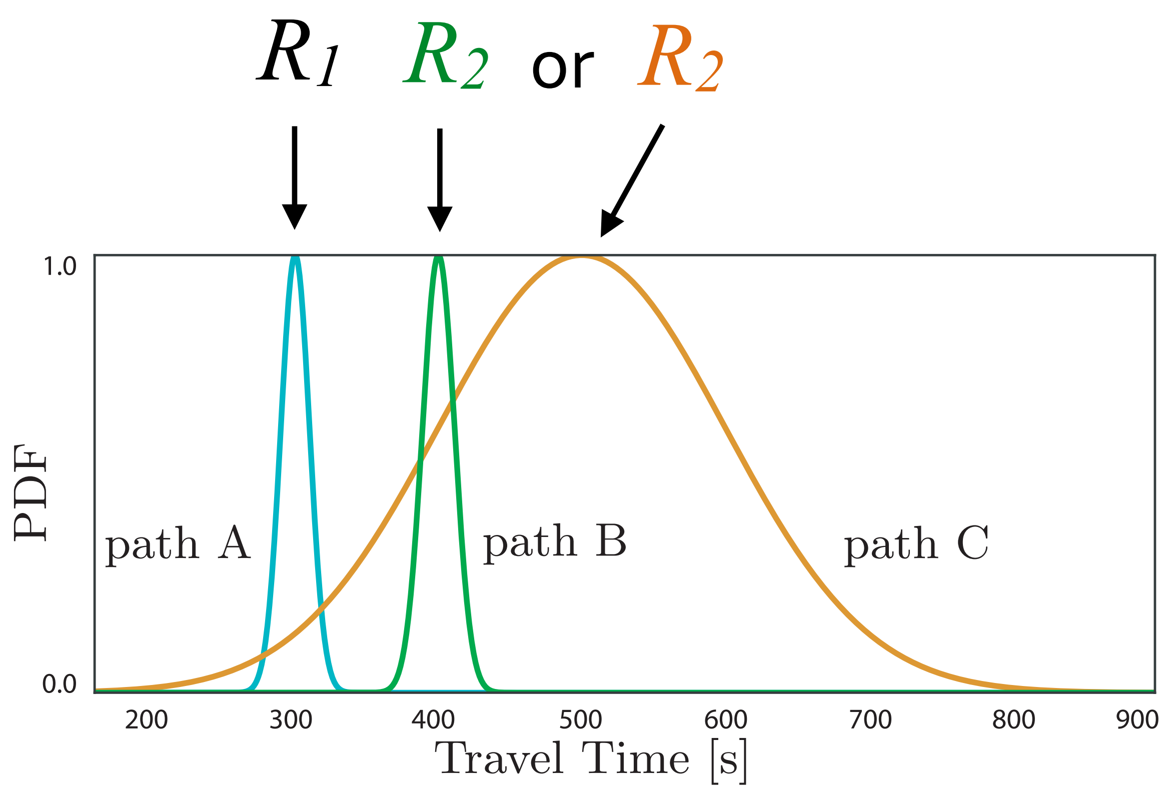

Path Uncertainty

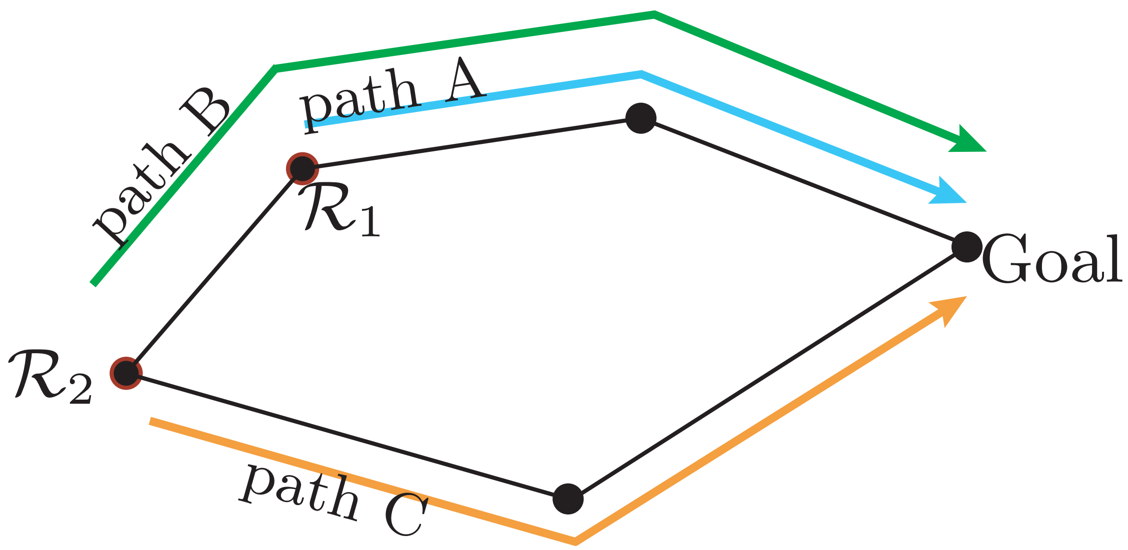

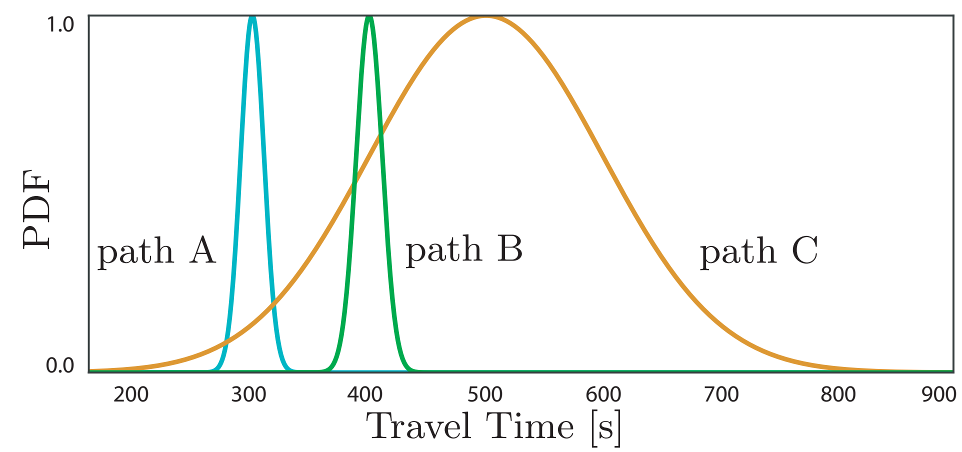

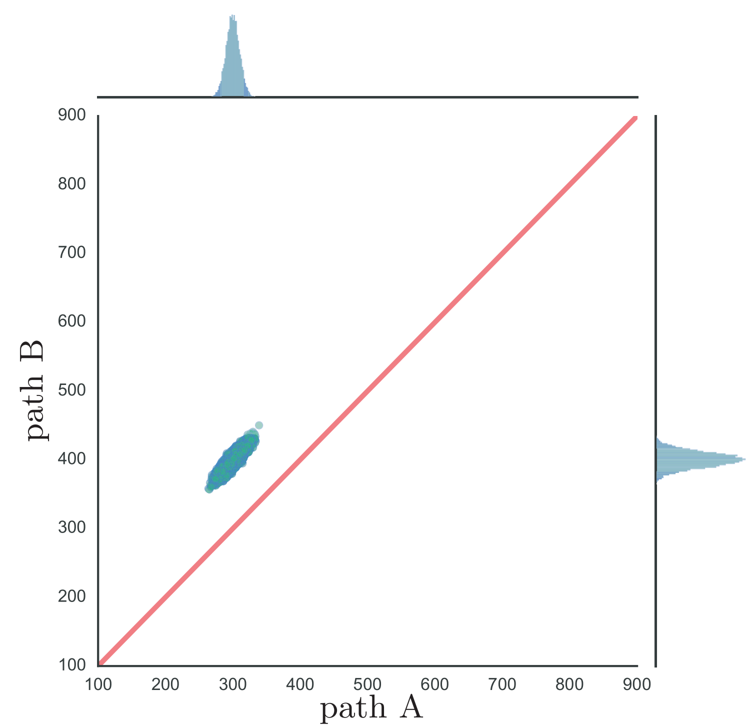

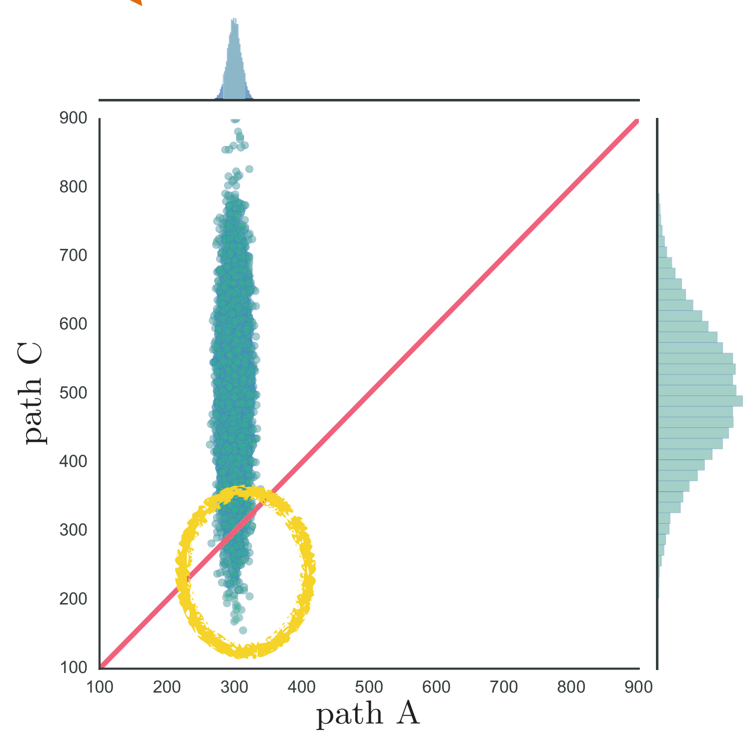

Greedy Redundant Assignment

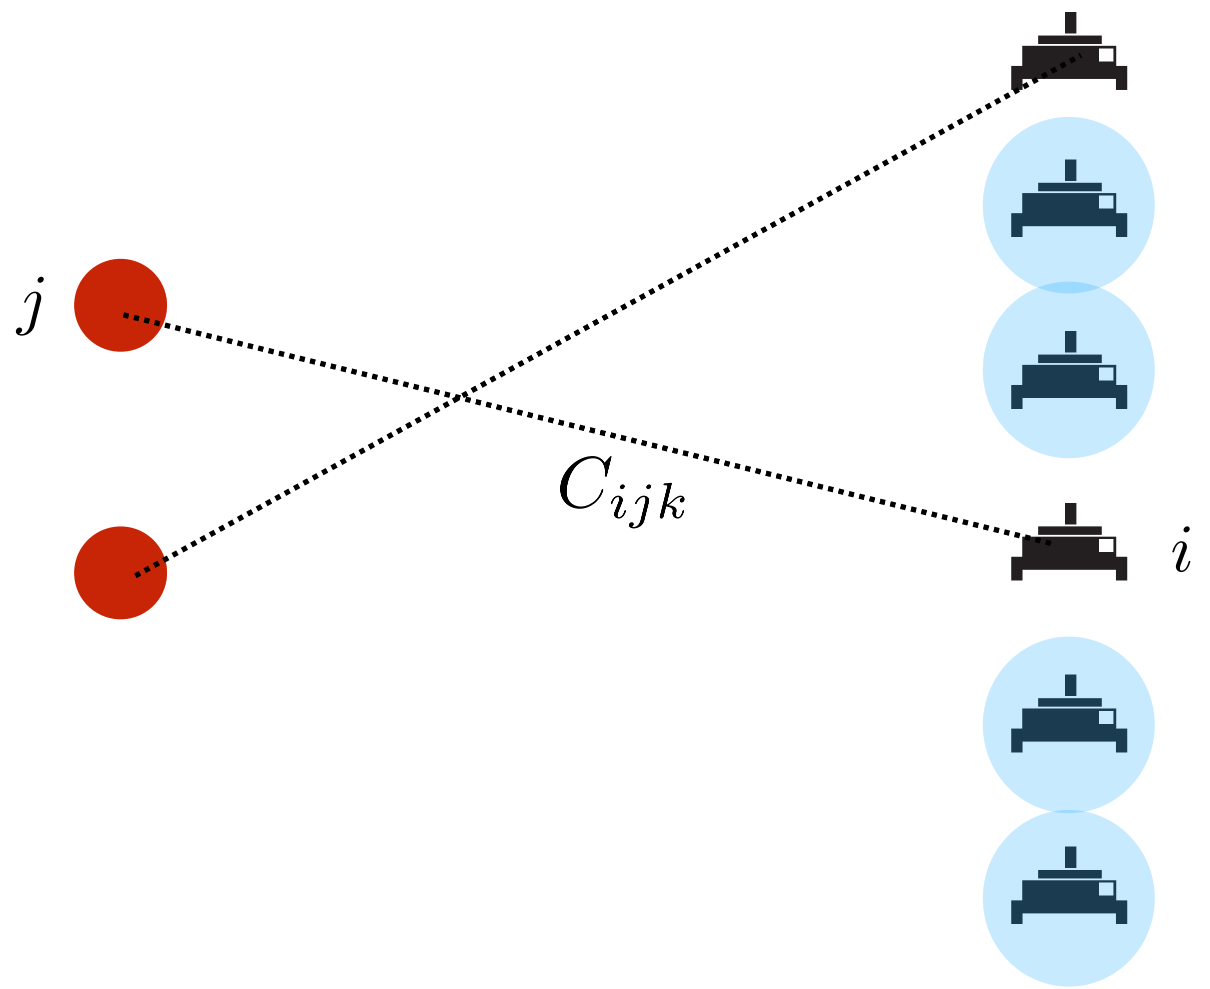

Which redundant robots to assign to which tasks?

Theory shows us that we can use a Greedy assignment algorithm with near-optimal performance: pick the robot that minimizes the overall waiting time the most.

[Prorok; Redundant Robot Assignment on Graphs with Uncertain Edge Costs; 2018]

Greedy Redundant Assignment

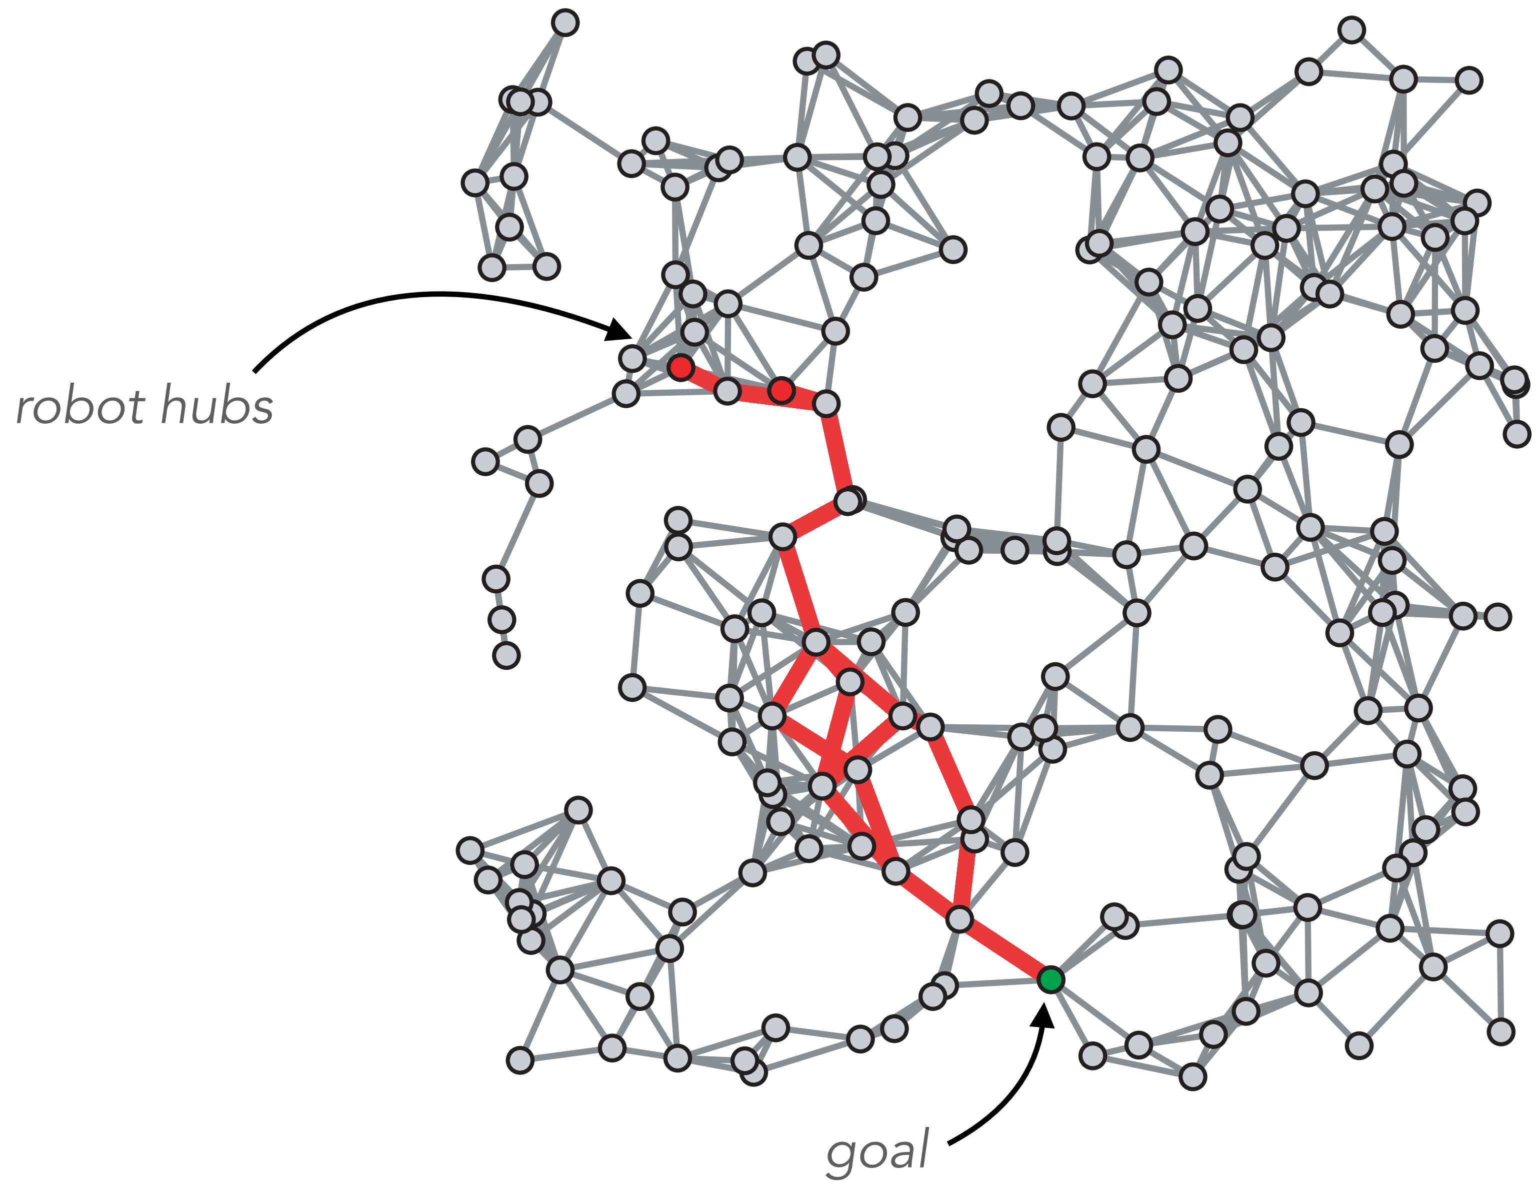

Greedy assignment of robots from 2 hubs to 1 goal

Redundant assignments leads to diversity of paths

[Prorok; Redundant Robot Assignment on Graphs with Uncertain Edge Costs; 2018]

Vehicle-Passenger Assignment

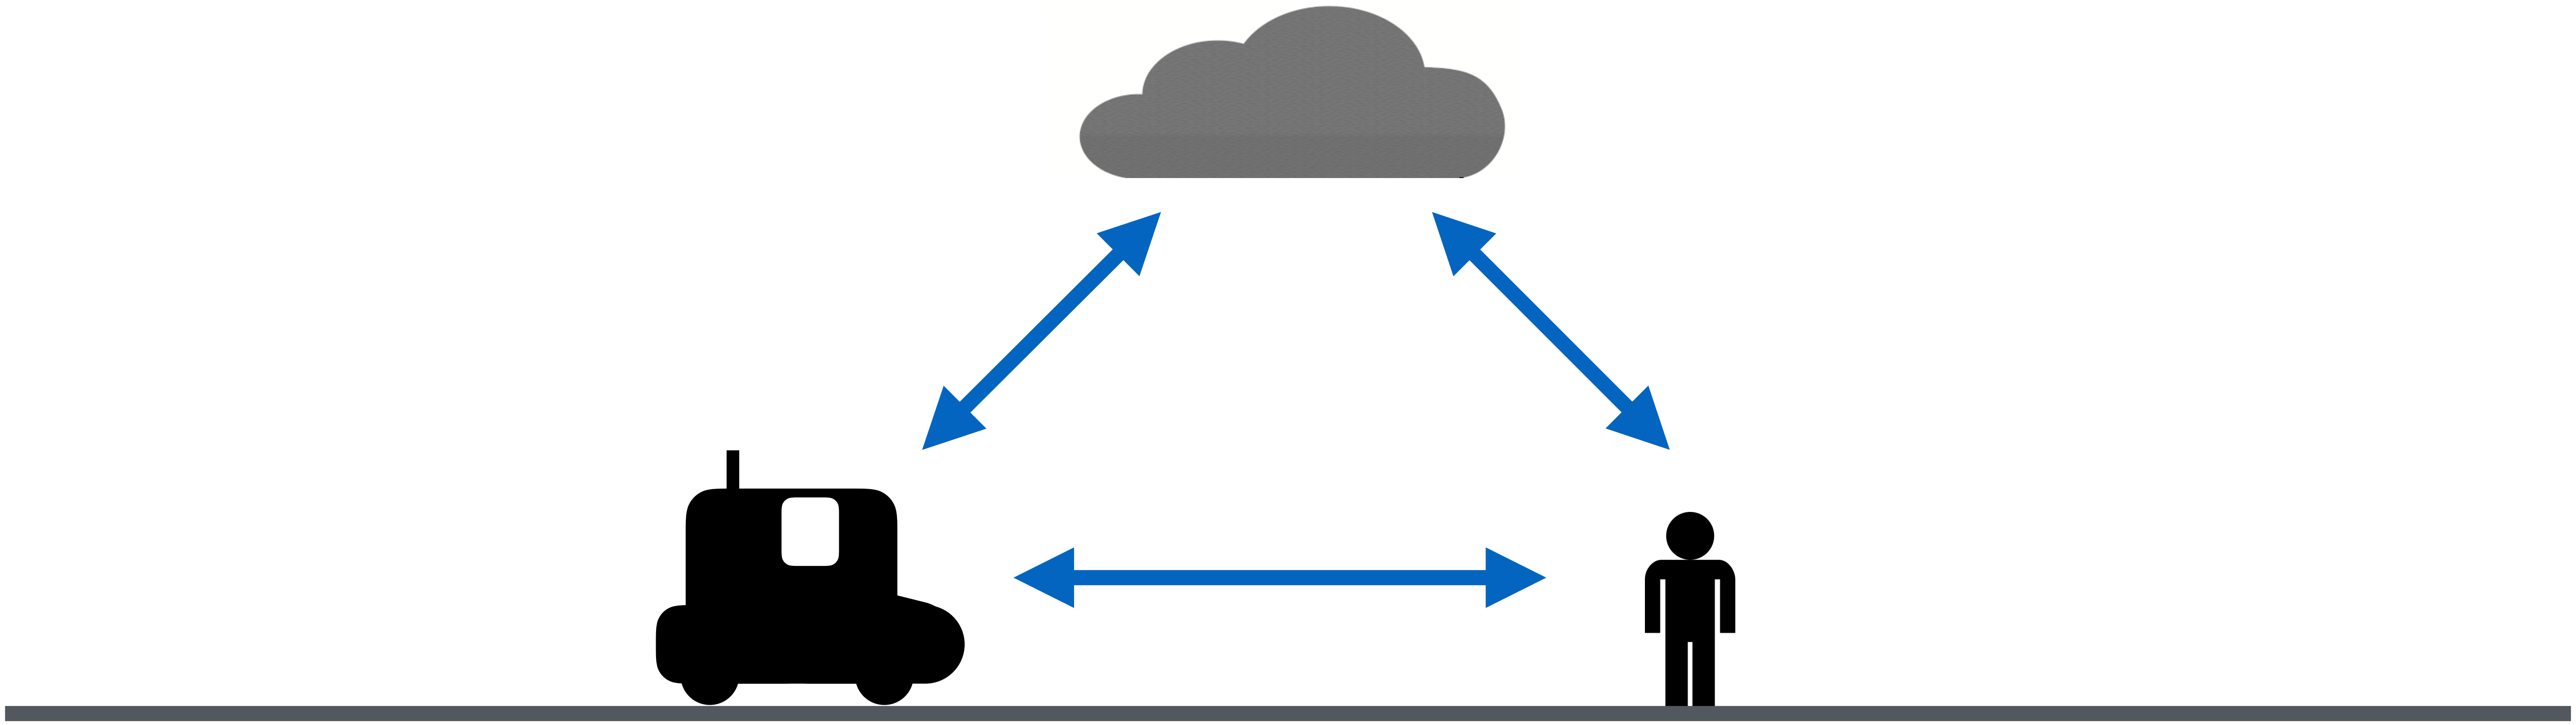

"I want optimized service and short waiting times"

"I want my mobility patterns to remain private "

Vehicle-Passenger Assignment

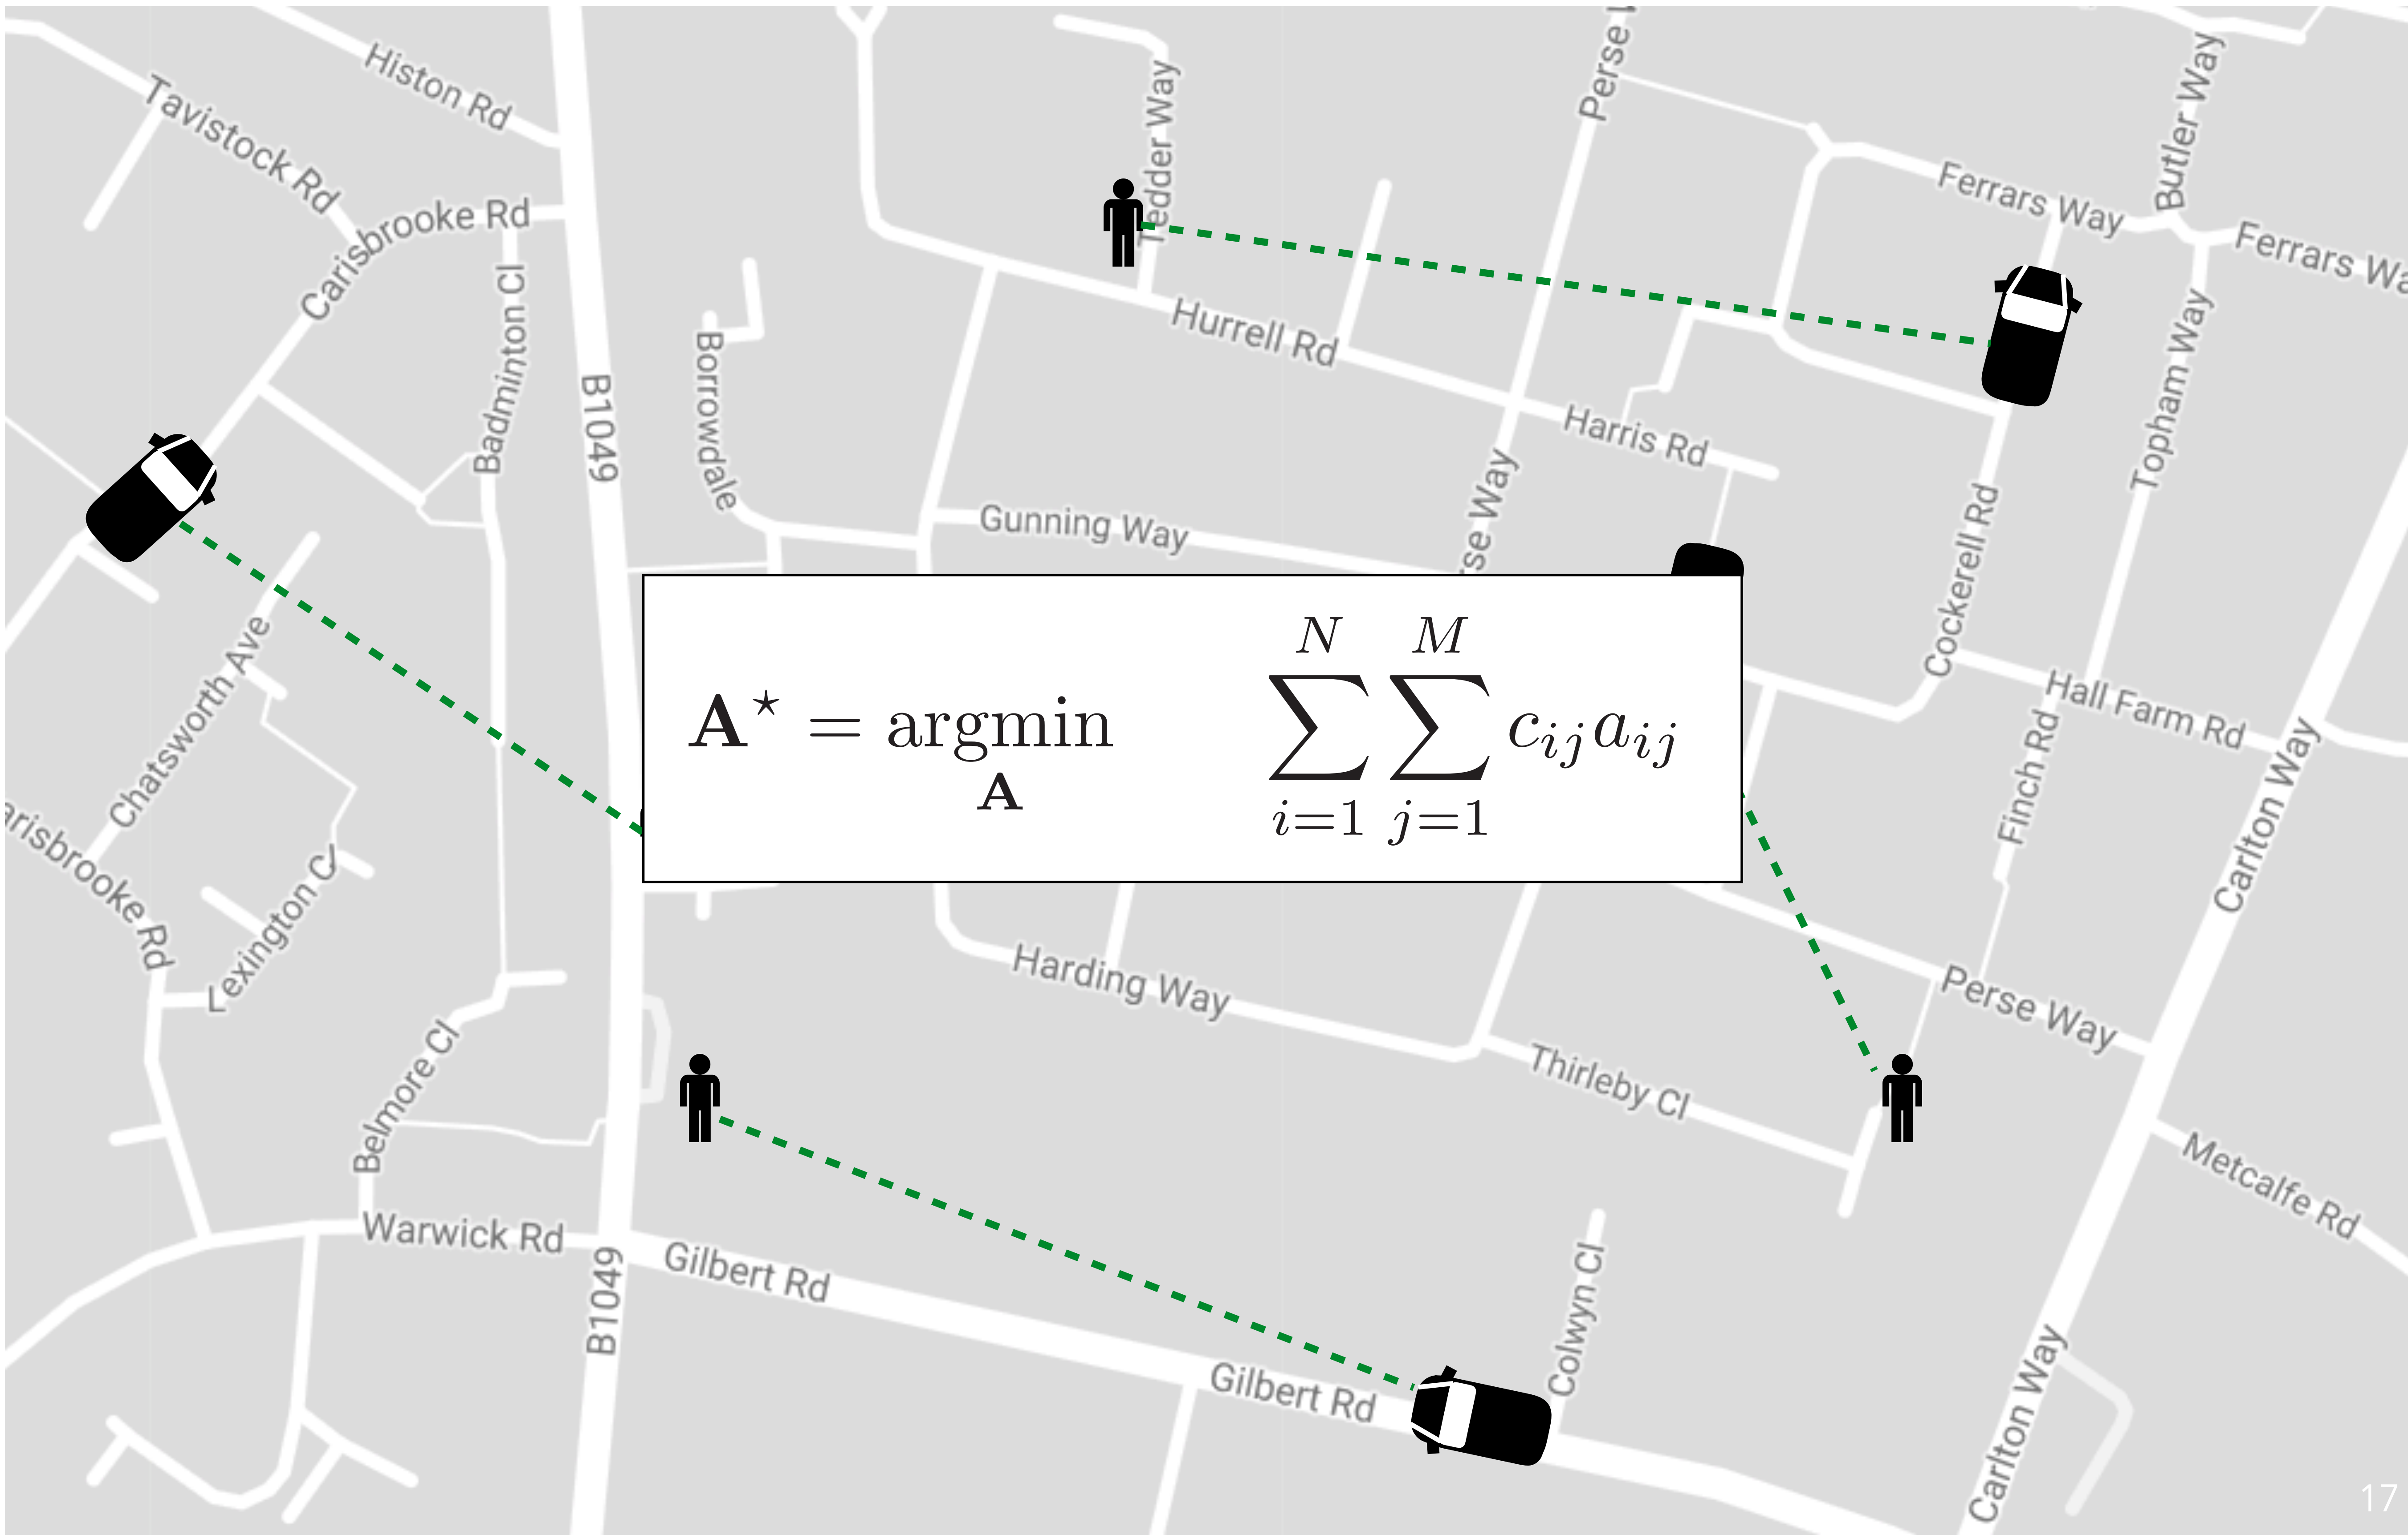

$$\mathbf{A}^{\star} = \underset{\mathbf{A}}{\operatorname{argmin}} \sum_{i=1}^{N} \sum_{j=1}^{M} c_{ij} a_{ij}$$

Vehicle-Passenger Assignment

Example: sub-area of Manhattan, around Flatiron building

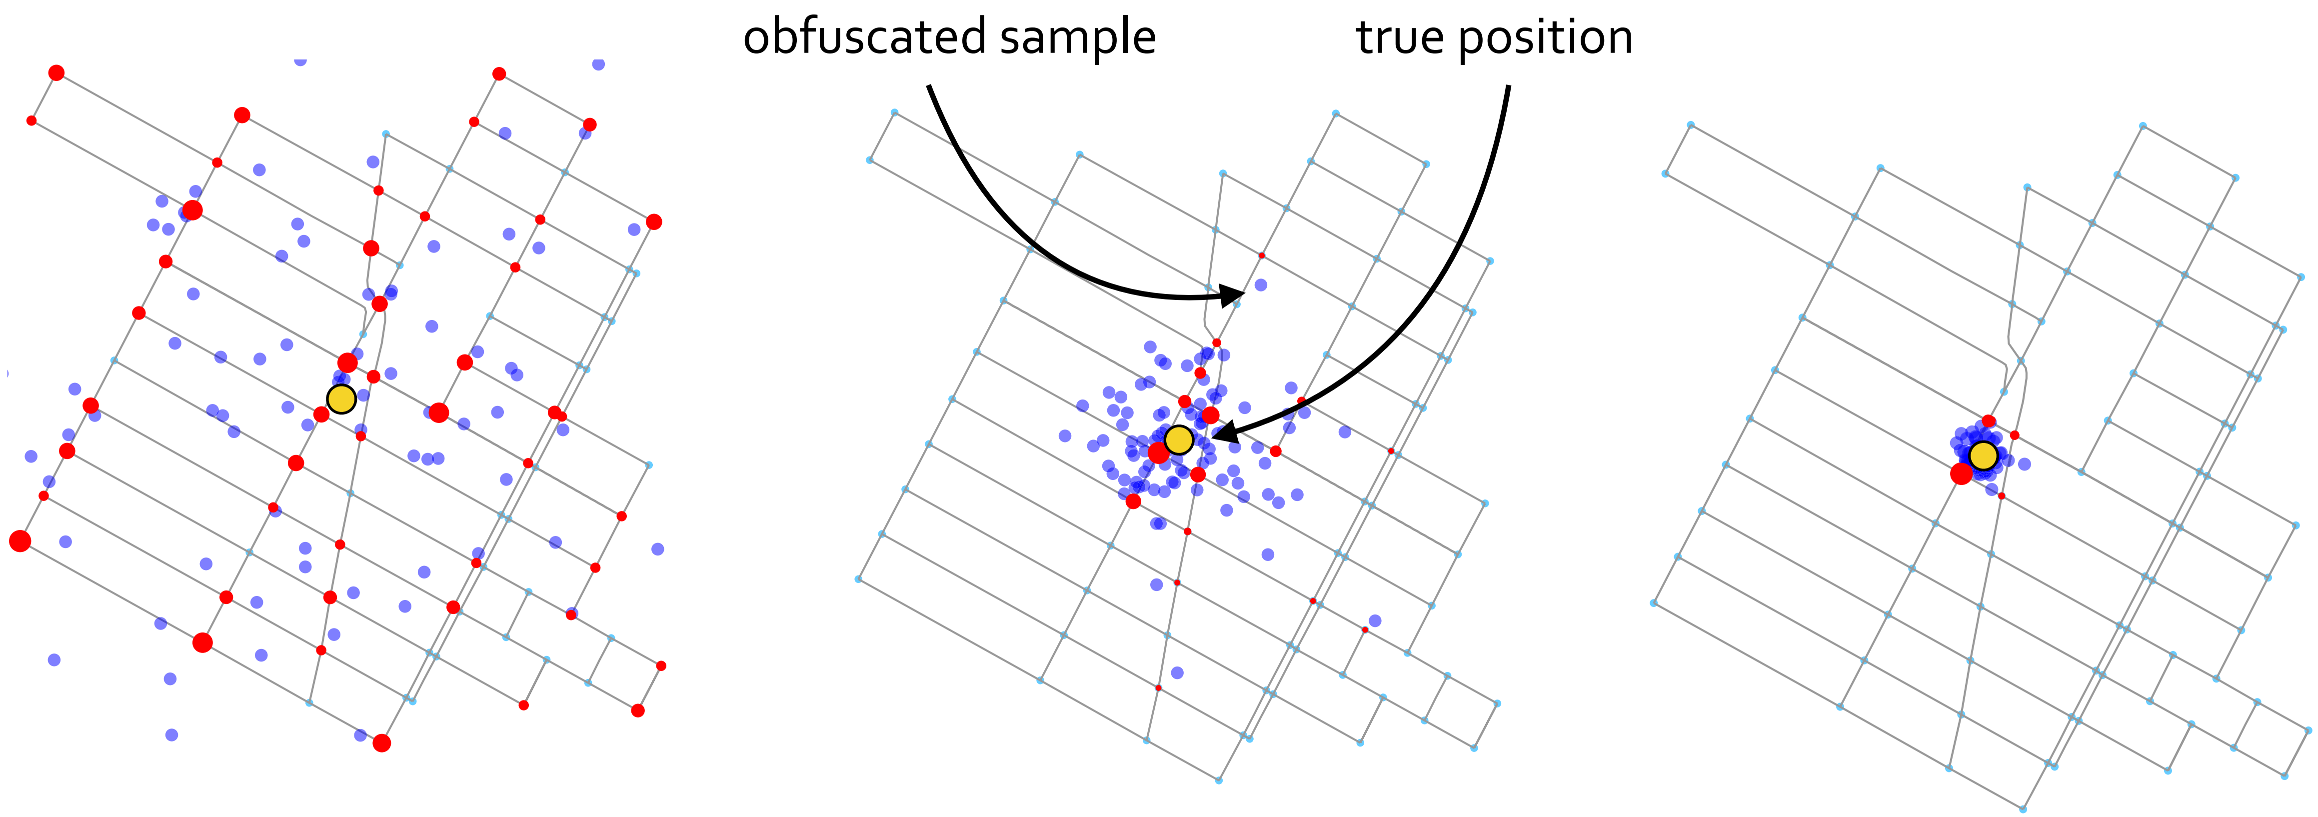

ε=0.005 ε=0.02 ε=0.1

Obfuscate positions with Laplace noise.

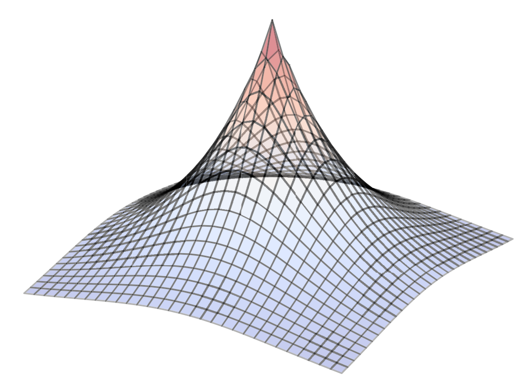

Vehicle-Passenger Assignment

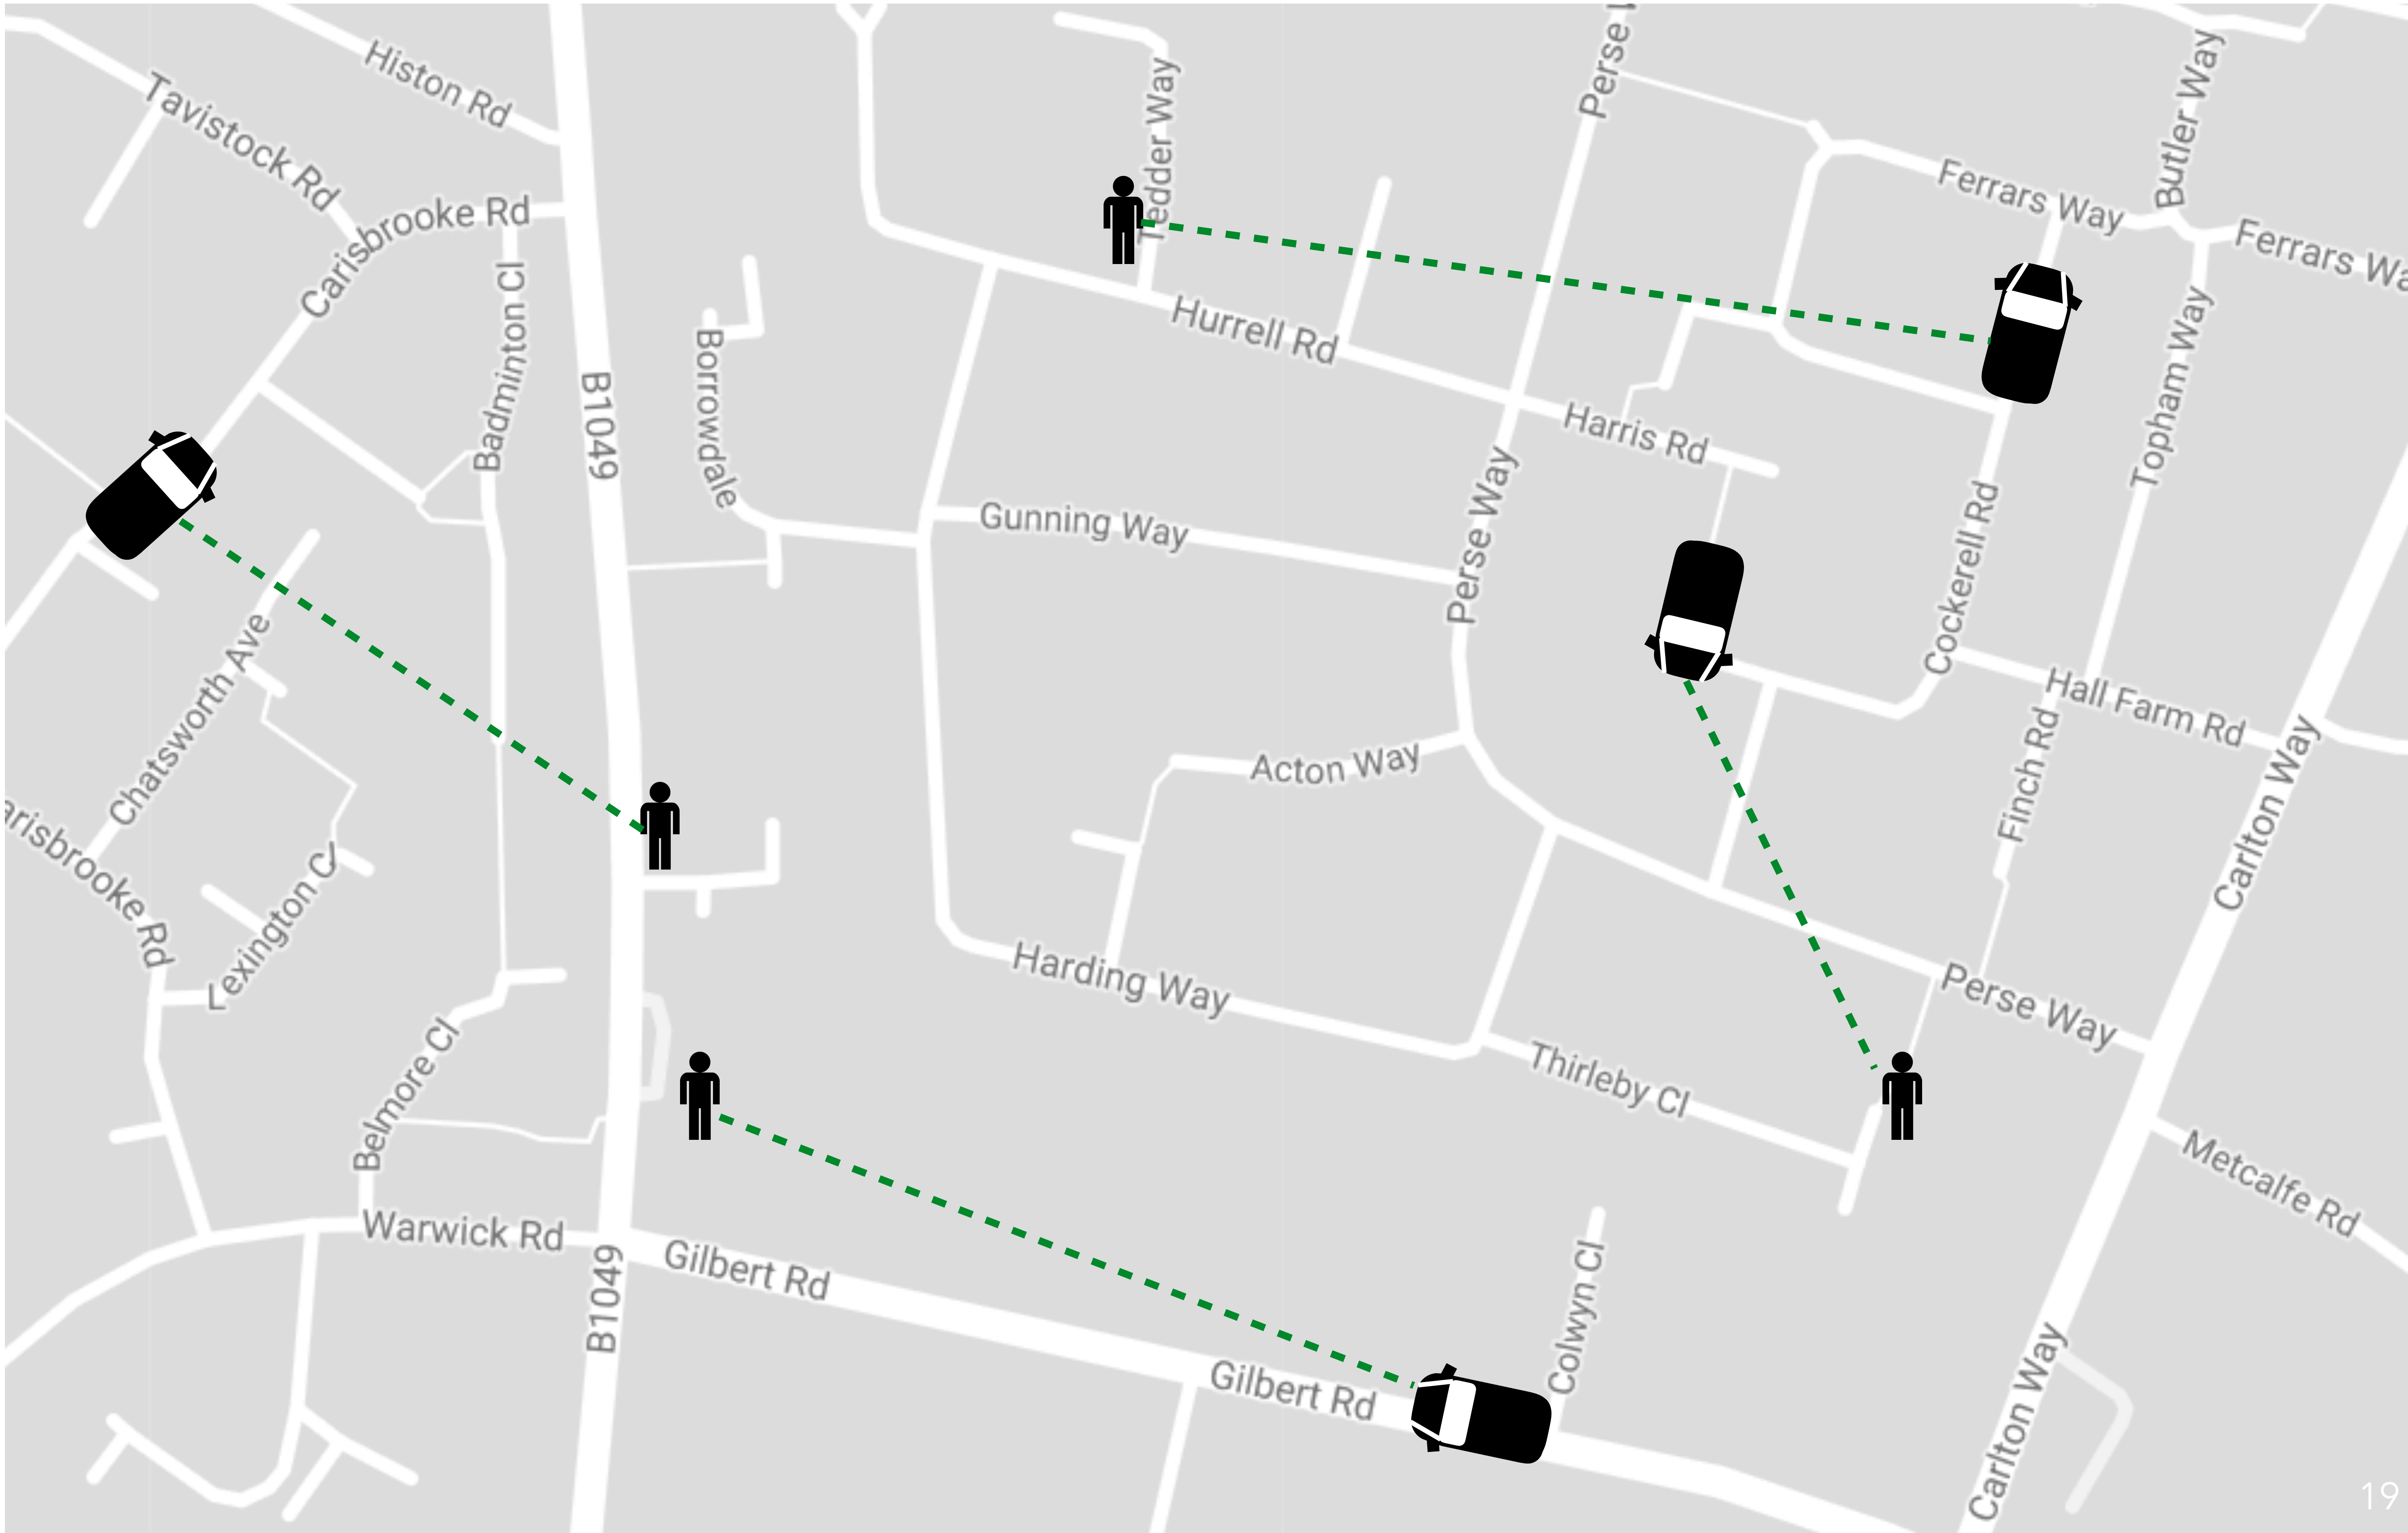

Vehicle-Passenger Assignment

Vehicle-Passenger Assignment

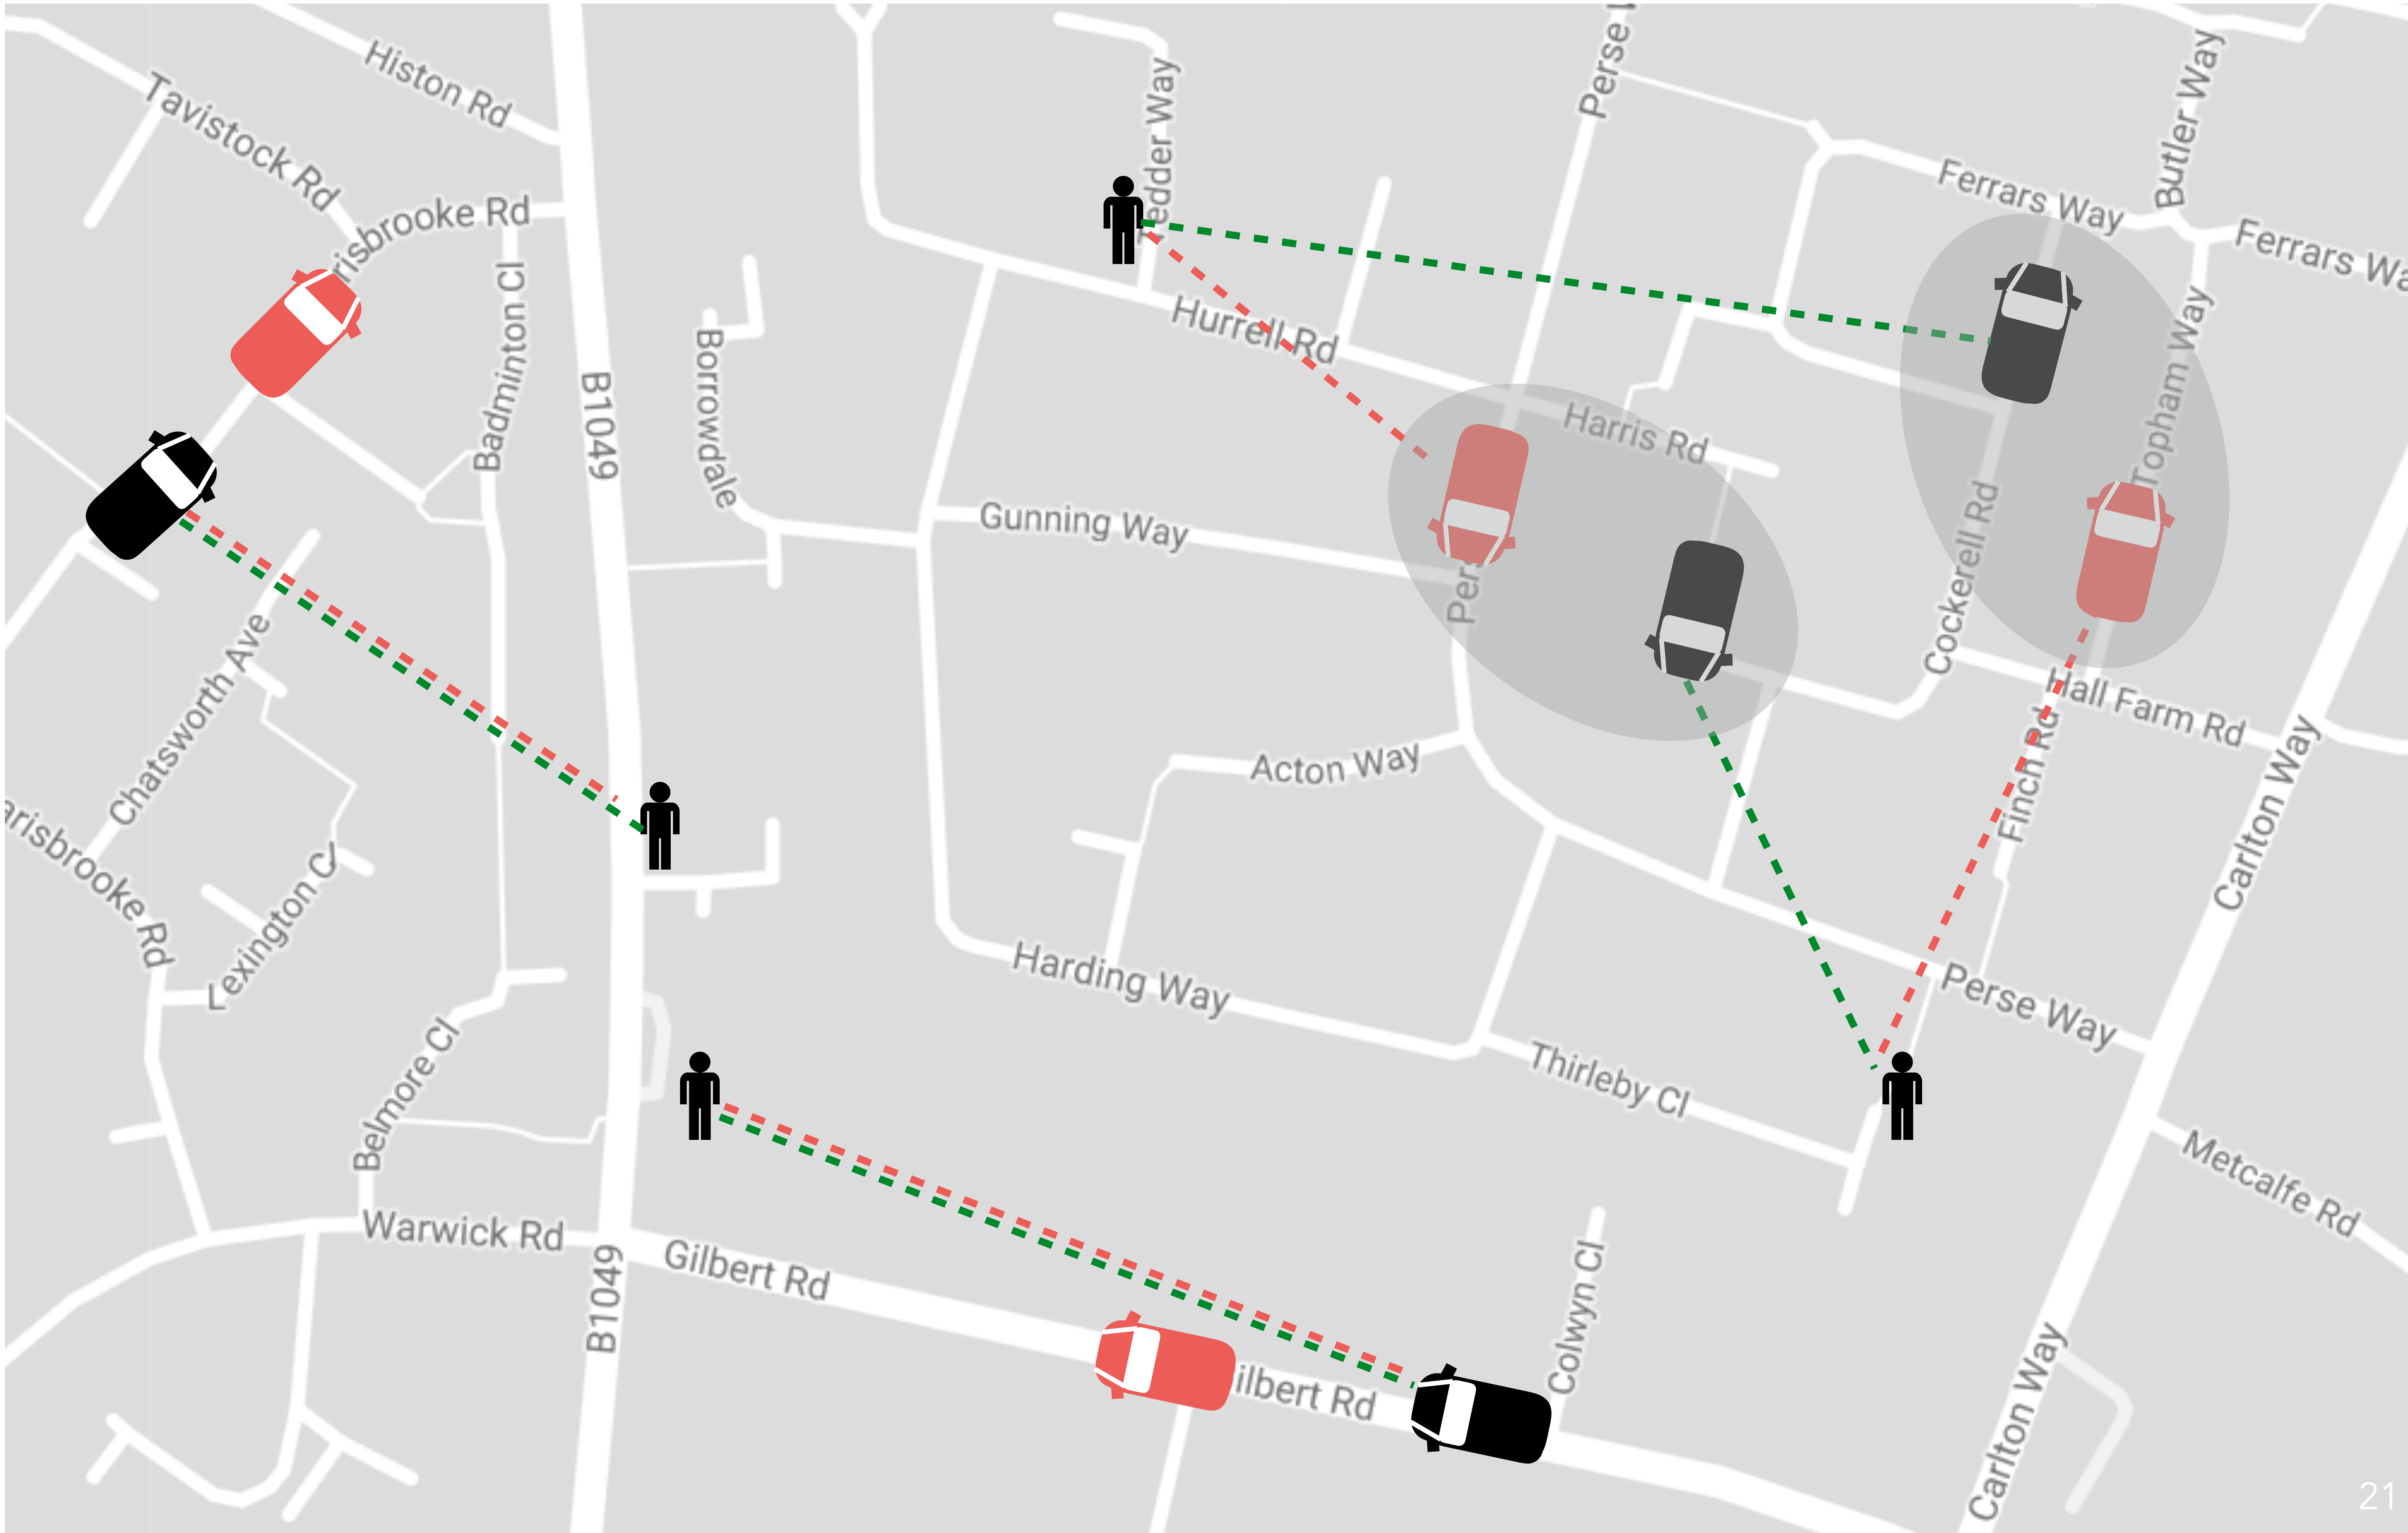

Vehicle-Passenger Assignment

Example: sub-area of Manhattan, around Flatiron building

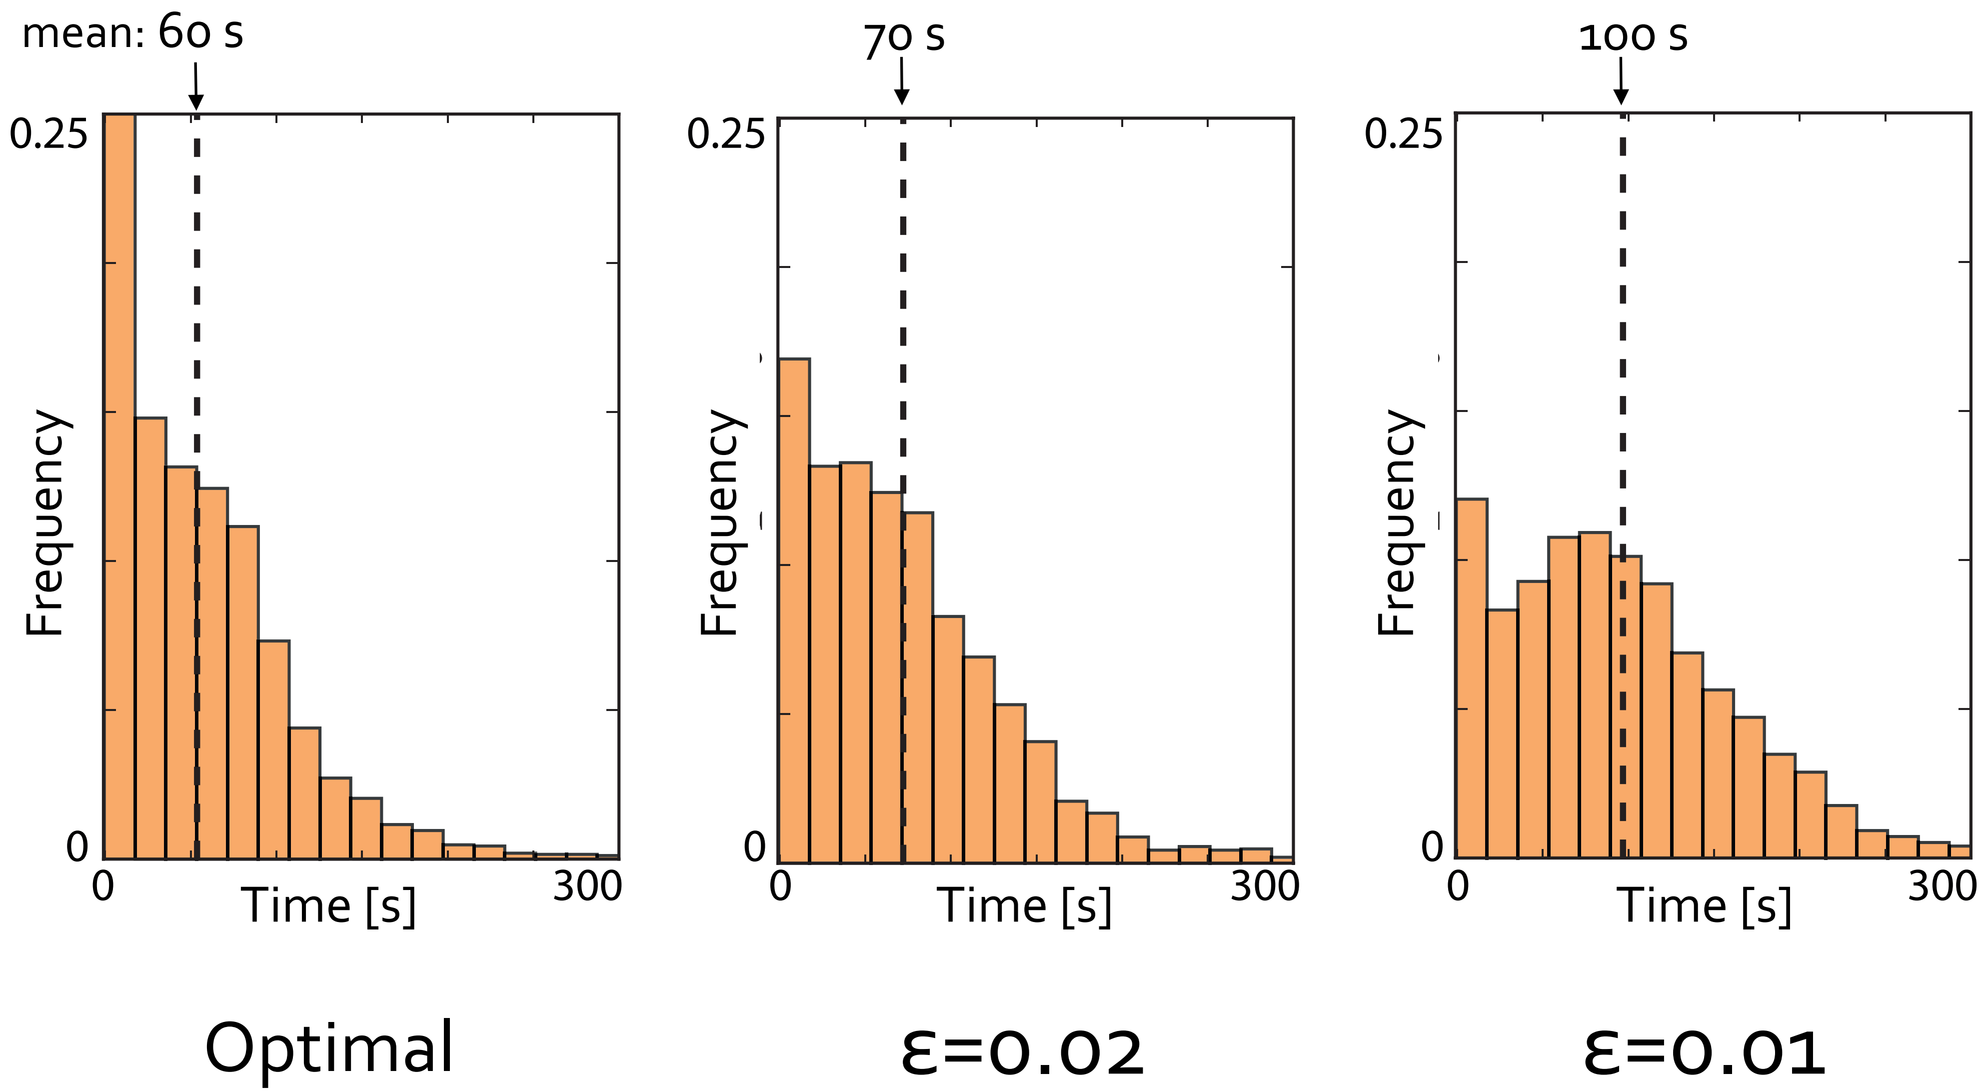

Can we minimize average passenger waiting times?

[Prorok et al.; Privacy-Preserving Vehicle Assignment for Mobility-on-Demand Systems; 2017]

Vehicle-Passenger Assignment

Principle: ***first-come, first-to-serve***!

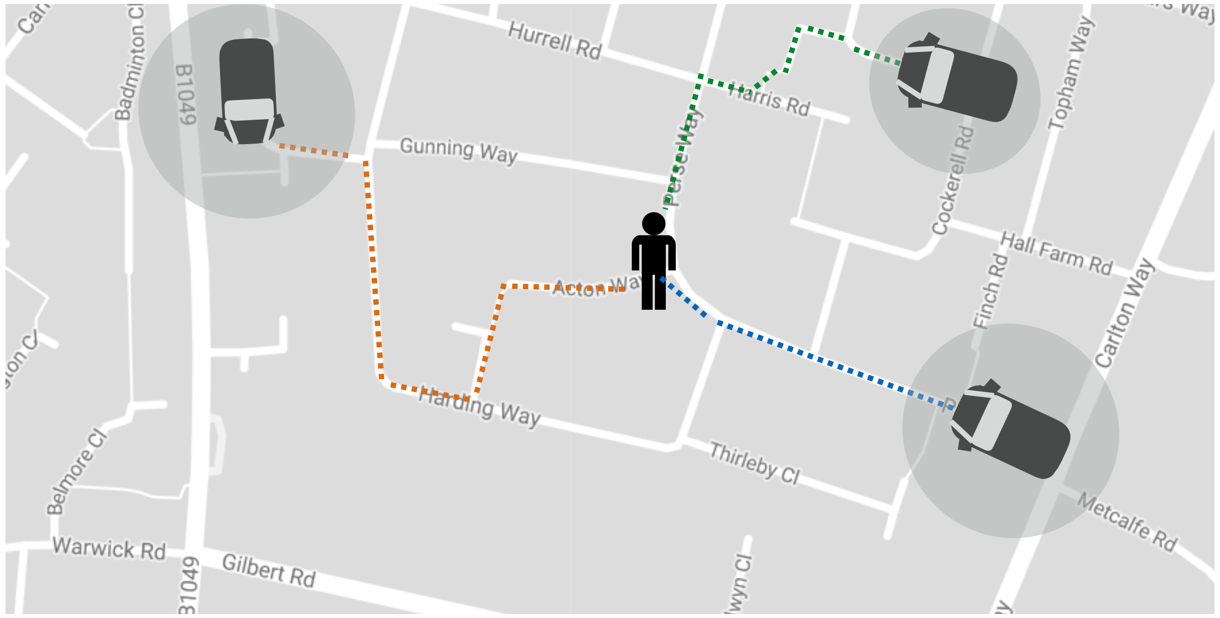

We can use ***Greedy*** to assign redundant vehicles
The algorithm is near-optimal.

Vehicle-Passenger Assignment

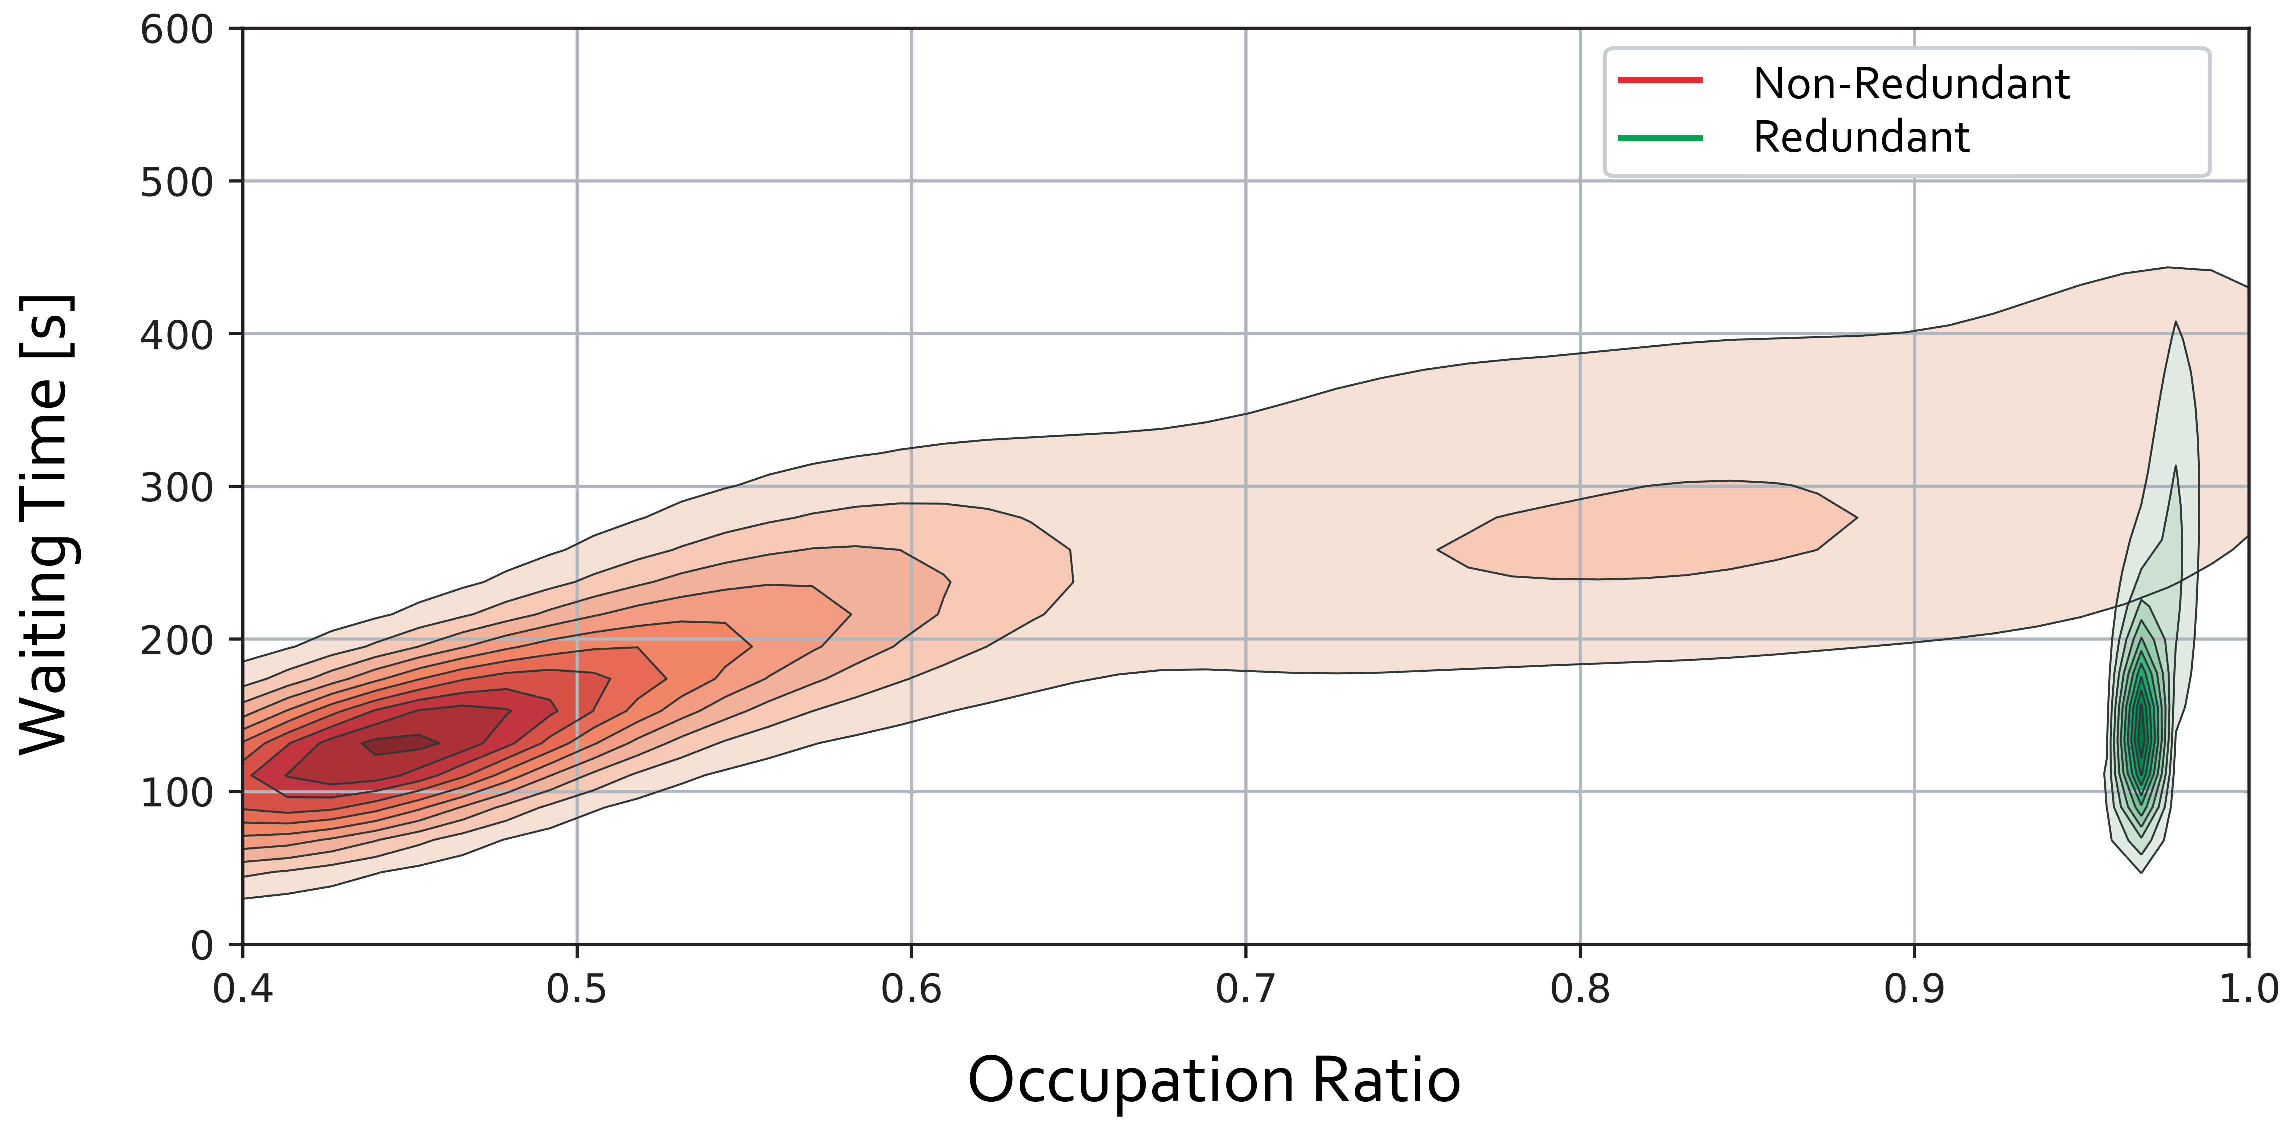

[Prorok et al.; Privacy-Preserving Vehicle Assignment for Mobility-on-Demand Systems; 2017]

Mobile Robot Systems

Lecture 10-II: State-of-the-Art and Outlook

Dr. Amanda Prorok

asp45@cam.ac.uk

www.proroklab.org

UNIVERSITY OF CAMBRIDGE

Grand Challenges of Robotics

1 **New Materials and Fabrication Schemes**

2 **Biohybrid and Bioinspired Robots**

3 **Power and Energy**

4 **Robot Swarms**

5 **Navigation and Exploration**

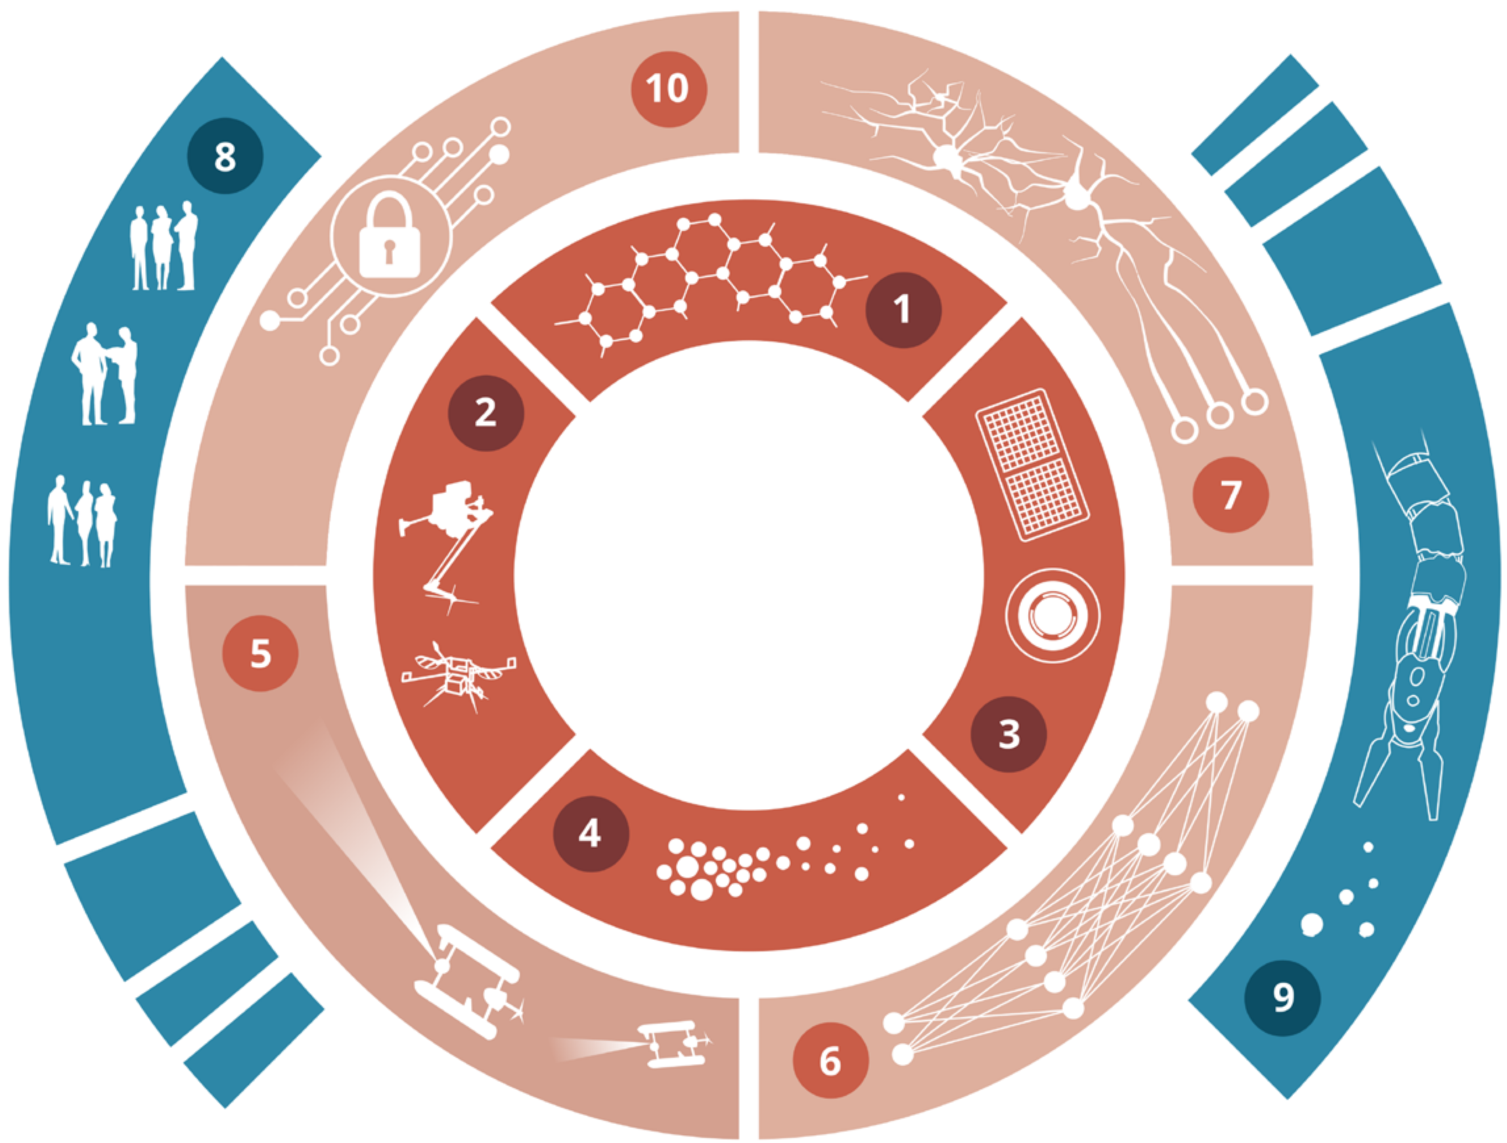

AI for Robotics 6

Brain-Computer Interfaces 7

Social Interaction 8

Medical Robotics 9

Robot Ethics and Security 10

[Yang et al; *The grand challenges of Science Robotics*; 2018]

Robot Swarms

- Technology drivers
  - Falling prices of sensors, processors, storage, communication
  - Convergence of consumer electronics with myriad types of intelligent autonomous systems (drones, robots, self-driving cars, etc.)
  - Mainstream availability of AI and predictive analytics
- ***Hyper-convergence***: Software-centric architecture that tightly couples computation, storage, networking, and virtualization resources



Robot Swarms - The Kilobot

[Rubenstein et al; 2014]

Robot Swarms - Challenges

- Challenges:
  - **Control** models: in large collections of intelligent devices, where the behavioral state-space is a Cartesian product of the individual state spaces: number of interaction modes grows combinatorially
  - **Communication** models: principled models for perception-action-communication loops; current models do not optimize wrt. communication.
  - **Resilience** to faults, non-cooperation, malicious action, intrusion; current paradigms assume perfect cooperation.
  - **Heterogeneity**: re-focusing of coordination methods on systems composed of heterogeneous, complementary robots; current methods model homogeneous systems.

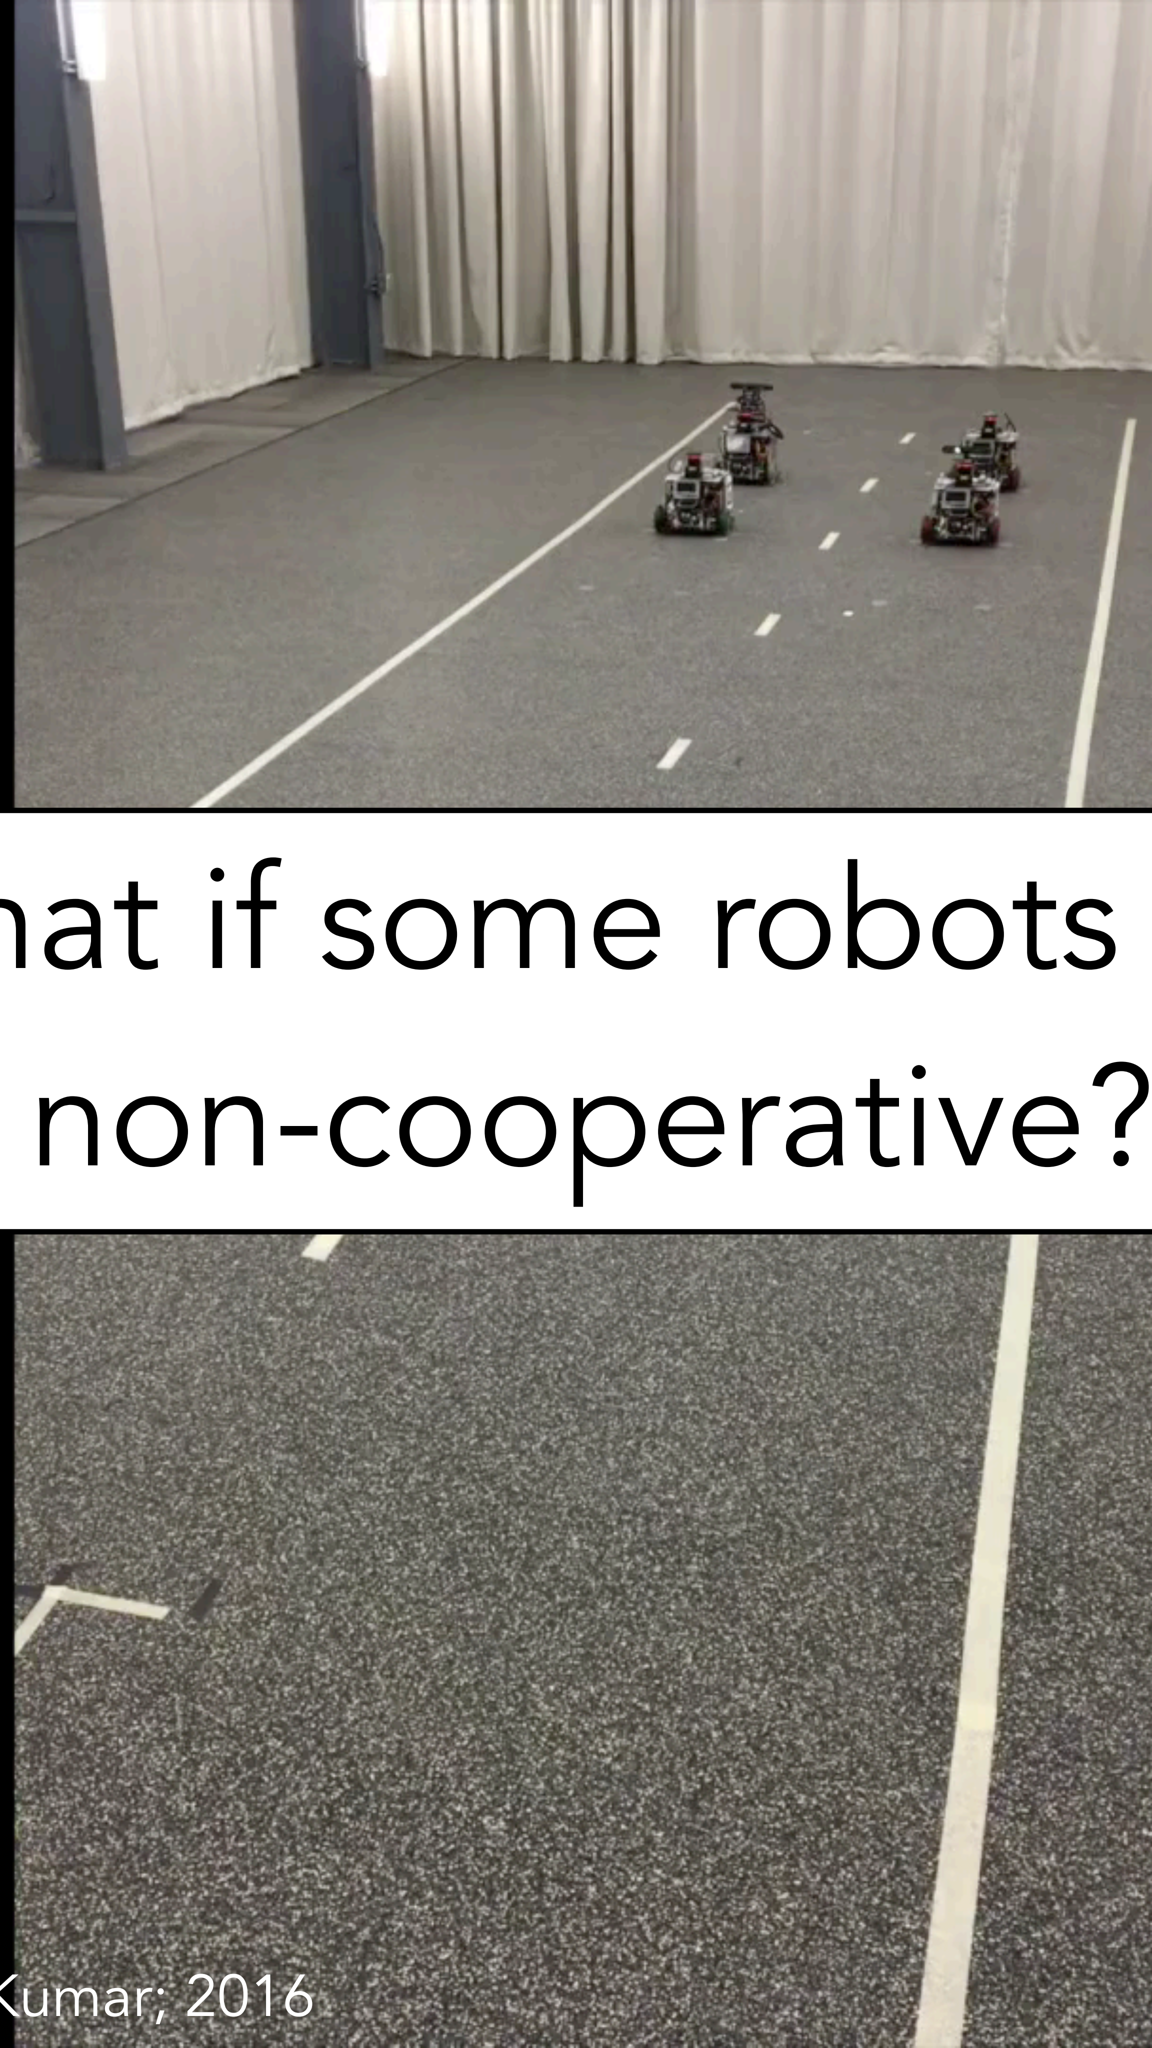

What if some robots are non-cooperative?

Whitzer, Kennedy, **Prorok**, Kumar; 2016

2x

Resilient Swarms

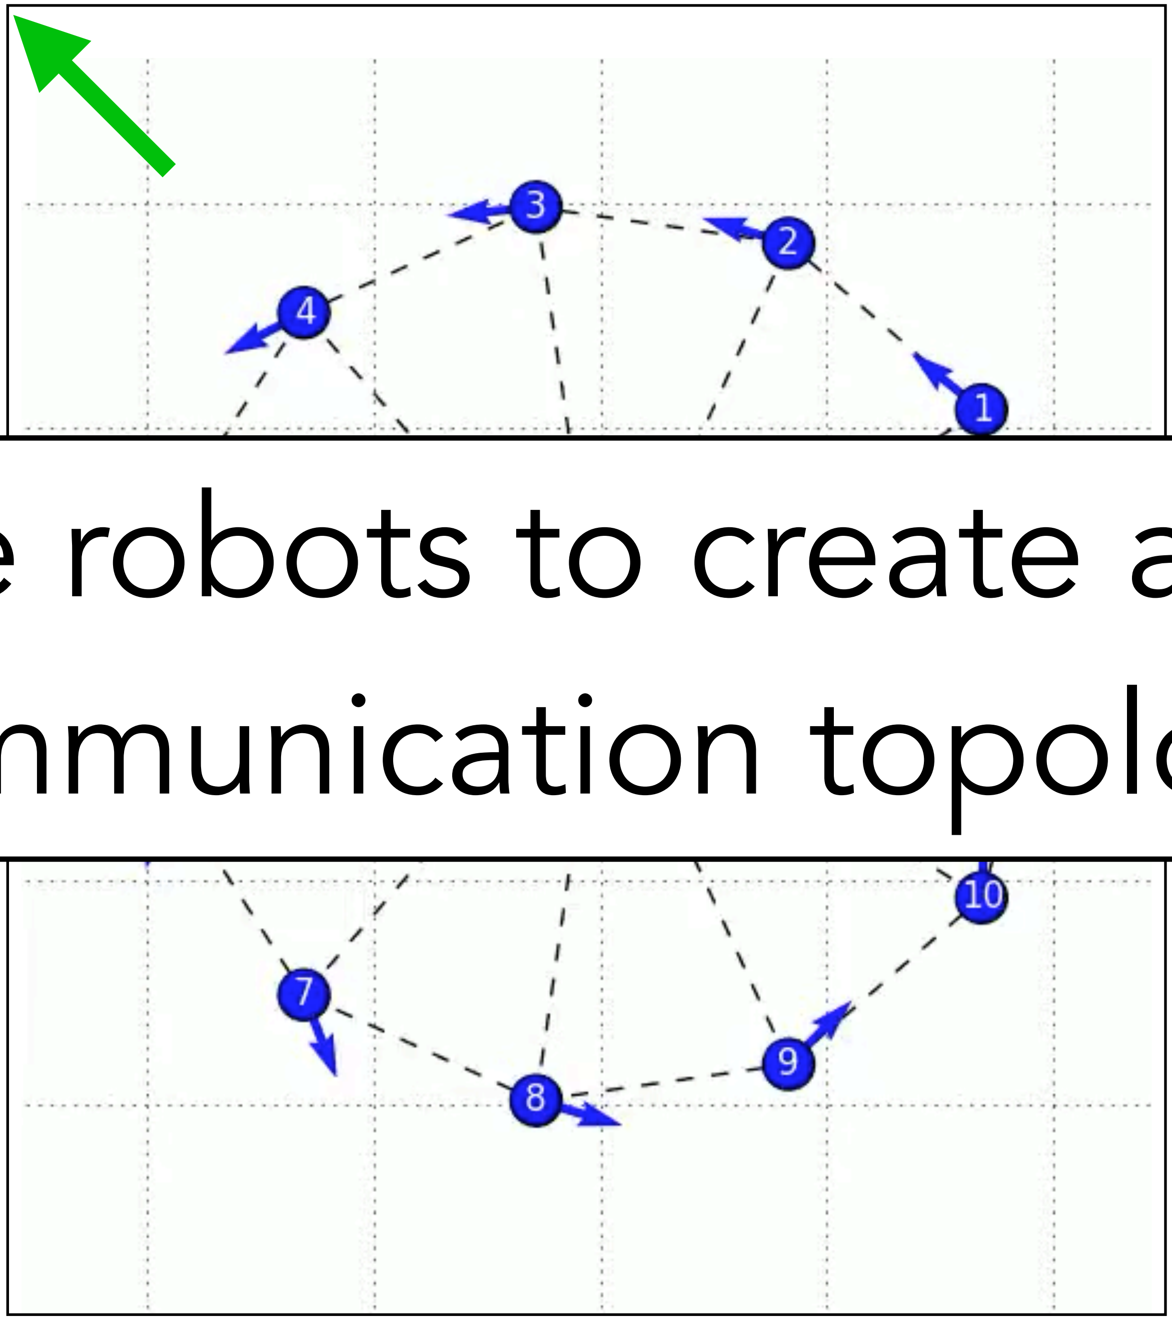

Coordinate robots to create a redundant communication topology.

Resilient Swarms

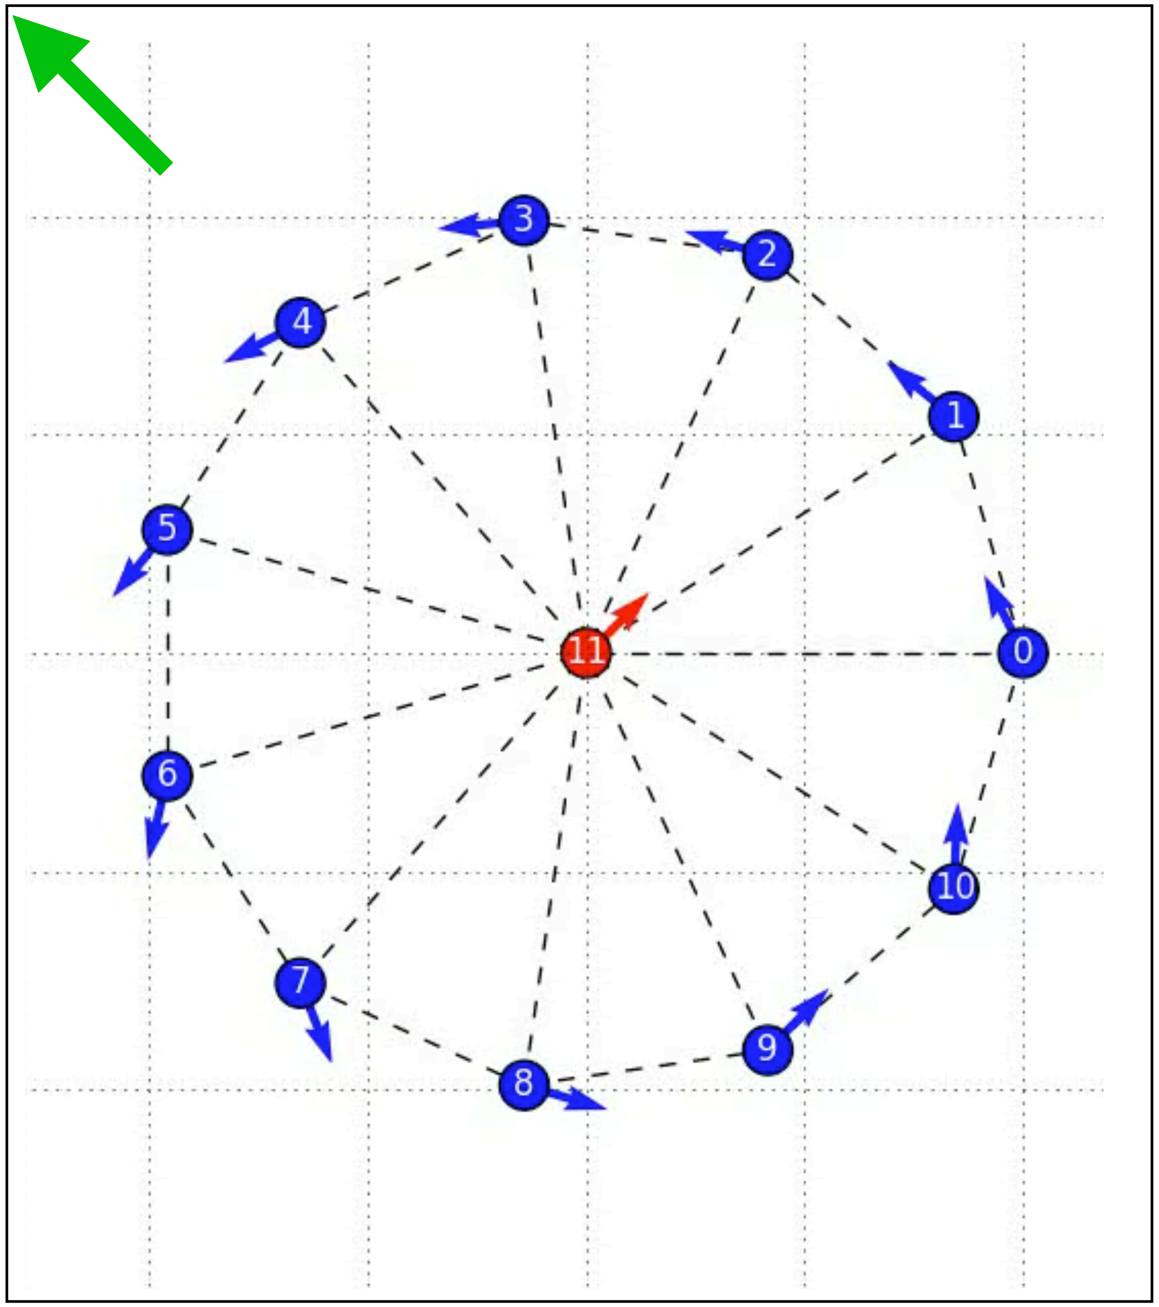

[Saulnier et al; *Resilient Flocking for Mobile Robot Teams*; 2017]

Perception-Action-Communication Loop

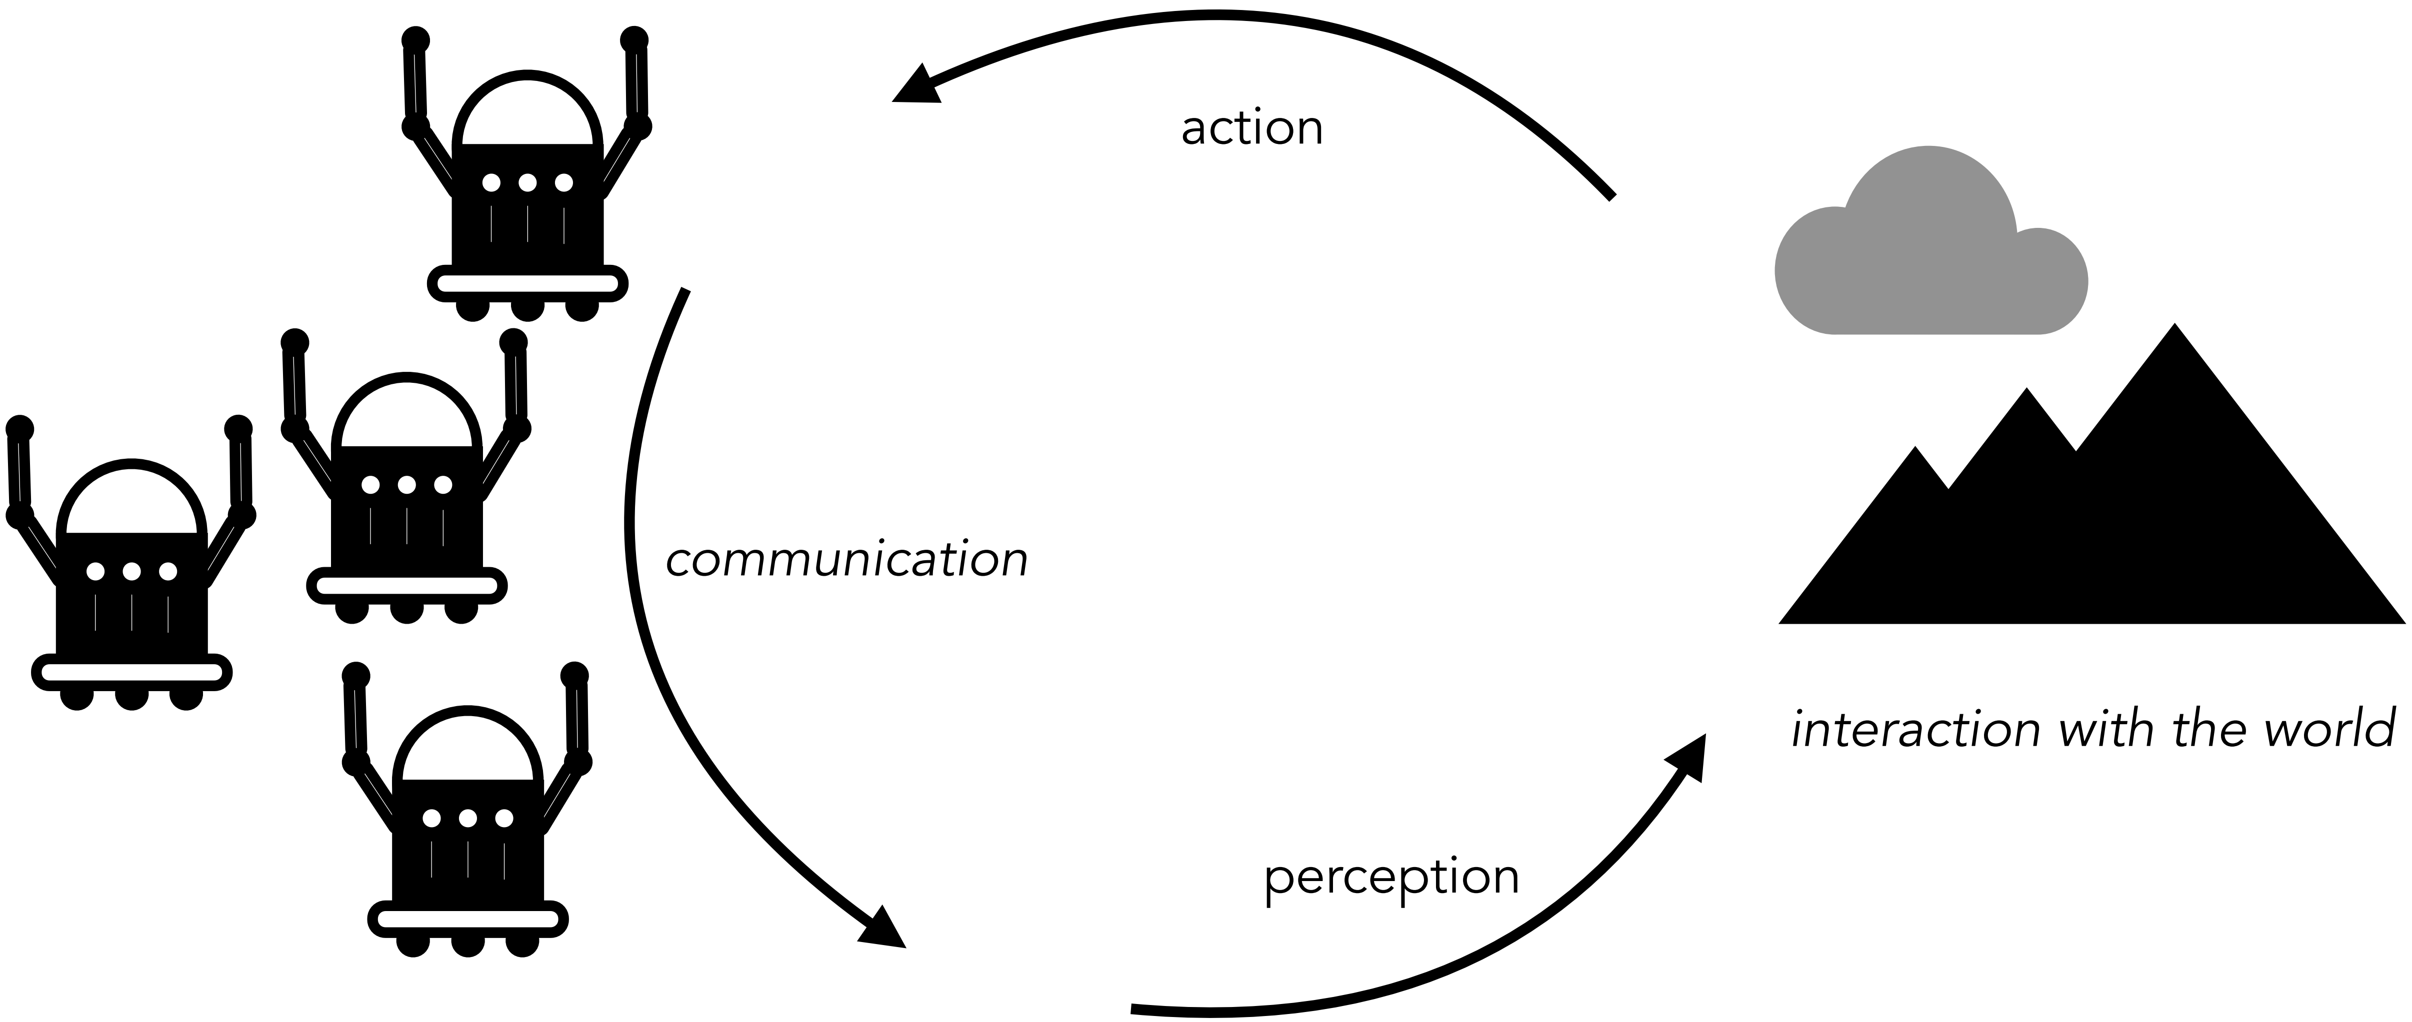

- Communication ability must be embedded in control loop
- Perception-action-communication loop: currently no systematic approaches for multi-dimensional control loops.

Resilient, Heterogeneous Systems

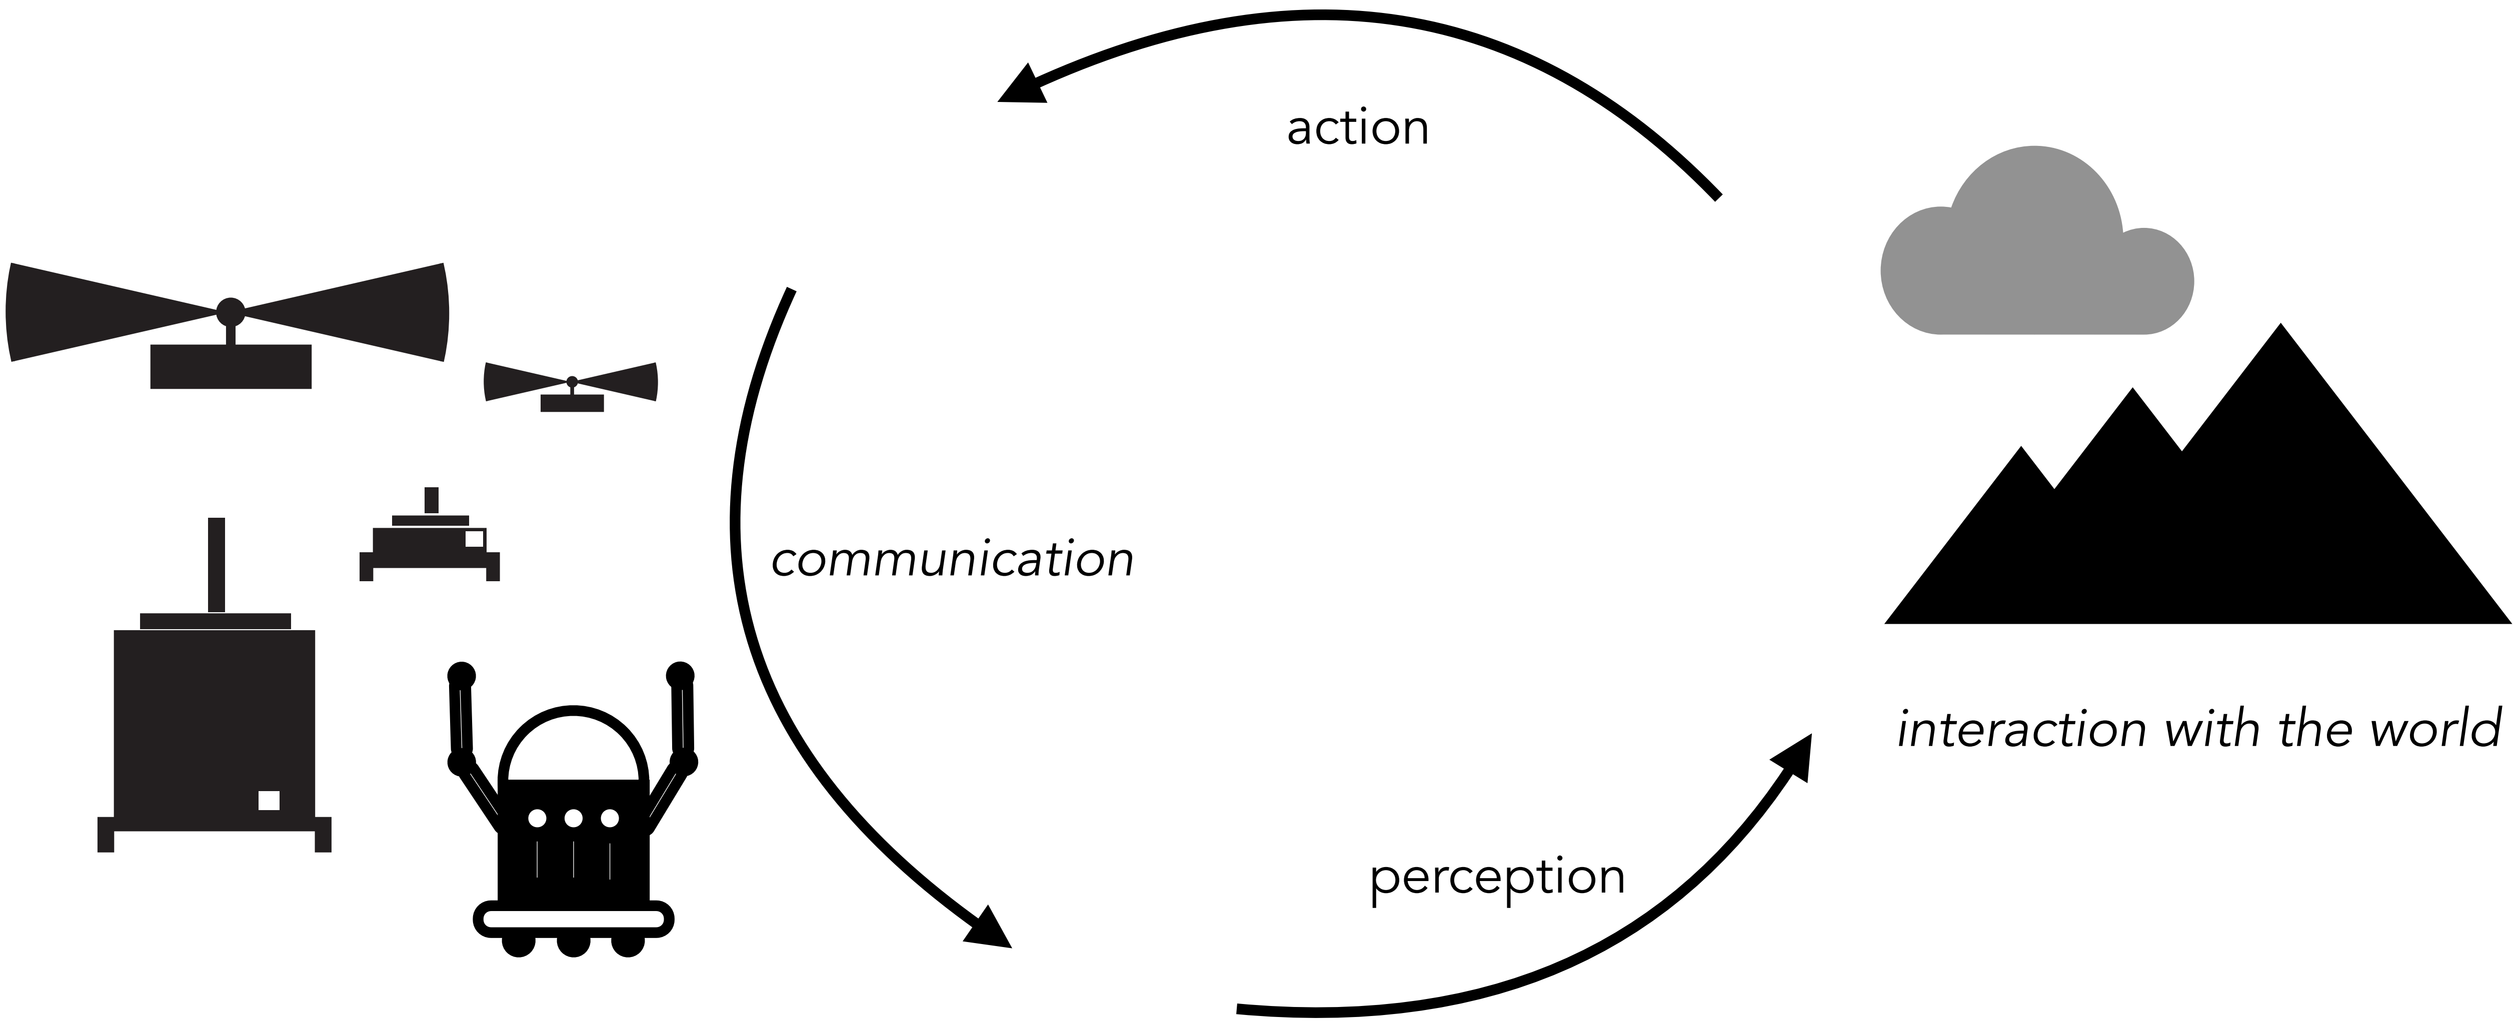

- Robots with **complementary** capabilities can truly collaborate to achieve difficult, complex tasks
- Coordination in heterogeneous systems is poorly addressed: new inter-robot dependencies and combinatorial state-space.

Navigation and Exploration - Challenges

- New terrains: deep sea, tunnels, mines, nuclear sites, outer space:
  - unknown and unmapped areas
  - perception and communication is degraded
- SLAM in non-static, dynamic, deformable environments
- Resource constraints (computational / comms): robots have to learn what is important → semantics

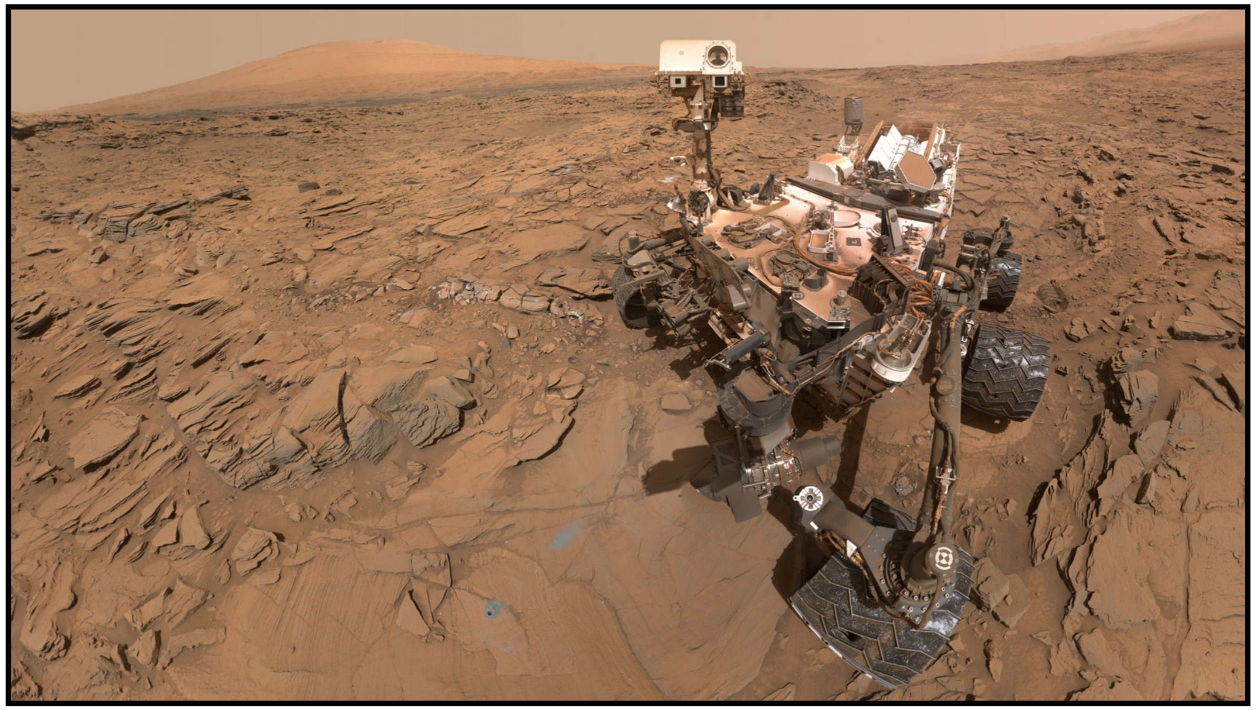

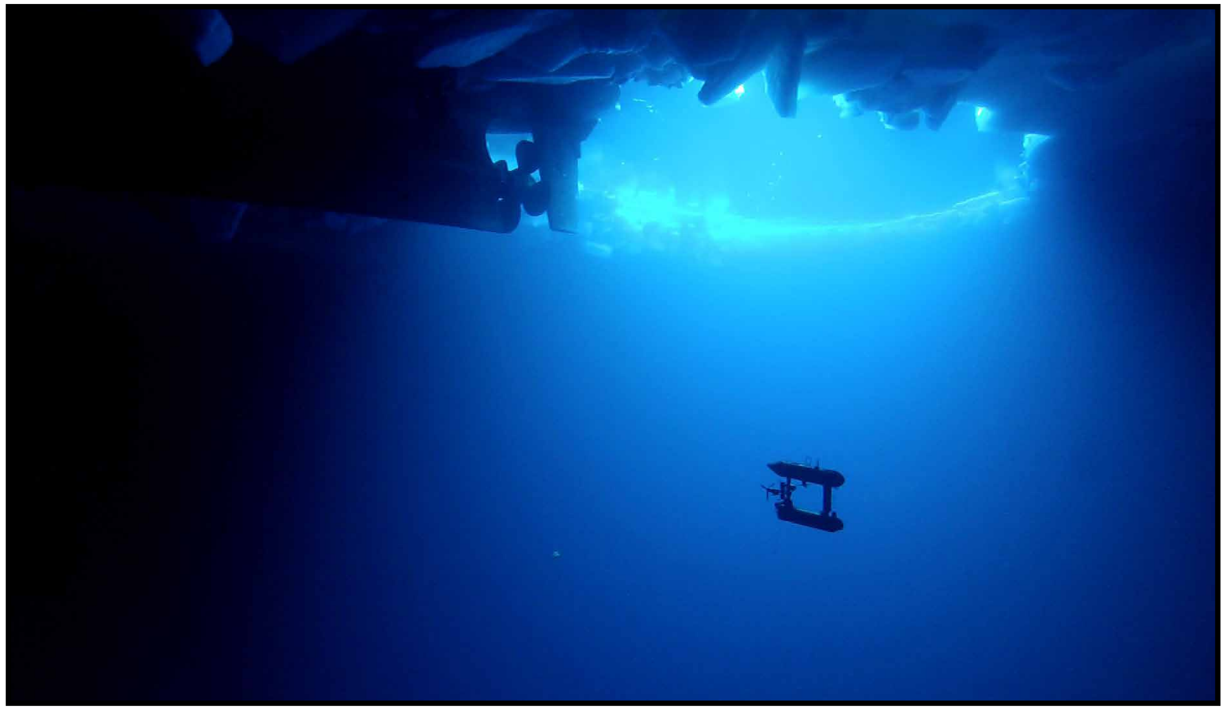

UNIVERSITY OF CAMBRIDGE

Navigation and Exploration - Approaches

- Current trends:
  - SLAM: beyond rigid and static world assumptions [1]
  - Active interaction with the environment to discover key traits [2]
  - End-to-end learning of navigation strategies [3]

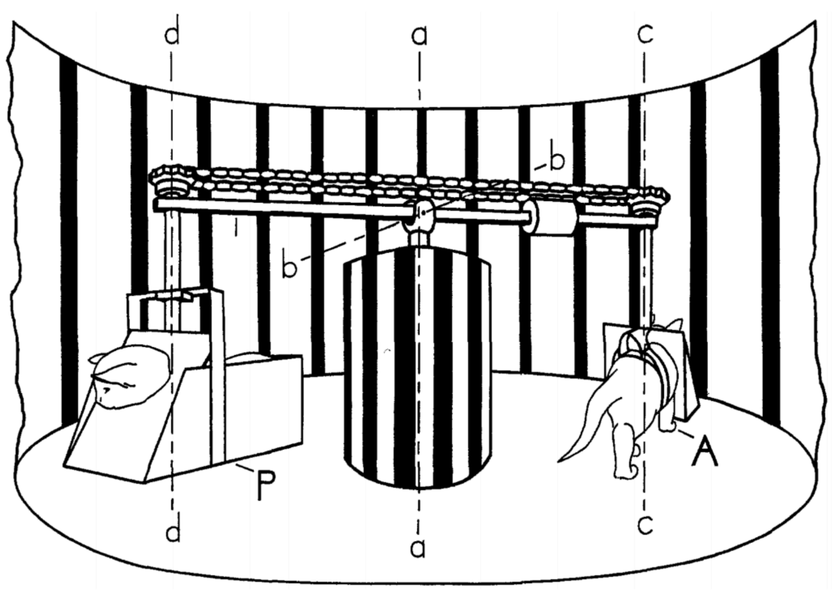

Experiment: One active and one passive kitten explore their world. Only the active kitten developed meaningful visually guided behavior.

[1] [C. Cadena et al.; Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age; 2016]

[2] [J. Bohg et al.; Interactive Perception: Leveraging Action in Perception and Perception in Action; 2017]

[3] [Gupta et al.; Cognitive Mapping and Planning for Visual Navigation; 2017]

Learning and AI

- The great hope:
  - AI that can learn complex tasks on its own and with a minimum of initial training data
  - Artificial General Intelligence (beyond robotics)
- Opportunities:
  - Domain- and task-specific problems: given enough resources (i.e., computing and data), applications can be made to work. General intelligence, however, is far from being solved.
  - Meta-learning: learning how to learn (beyond statistical correlation)
  - Cloud robotics (robotic IoT)

Perception-Action Loop in RL

Note: The agent does not necessarily see the full world state. For simplicity, this slide assumes full world observability.

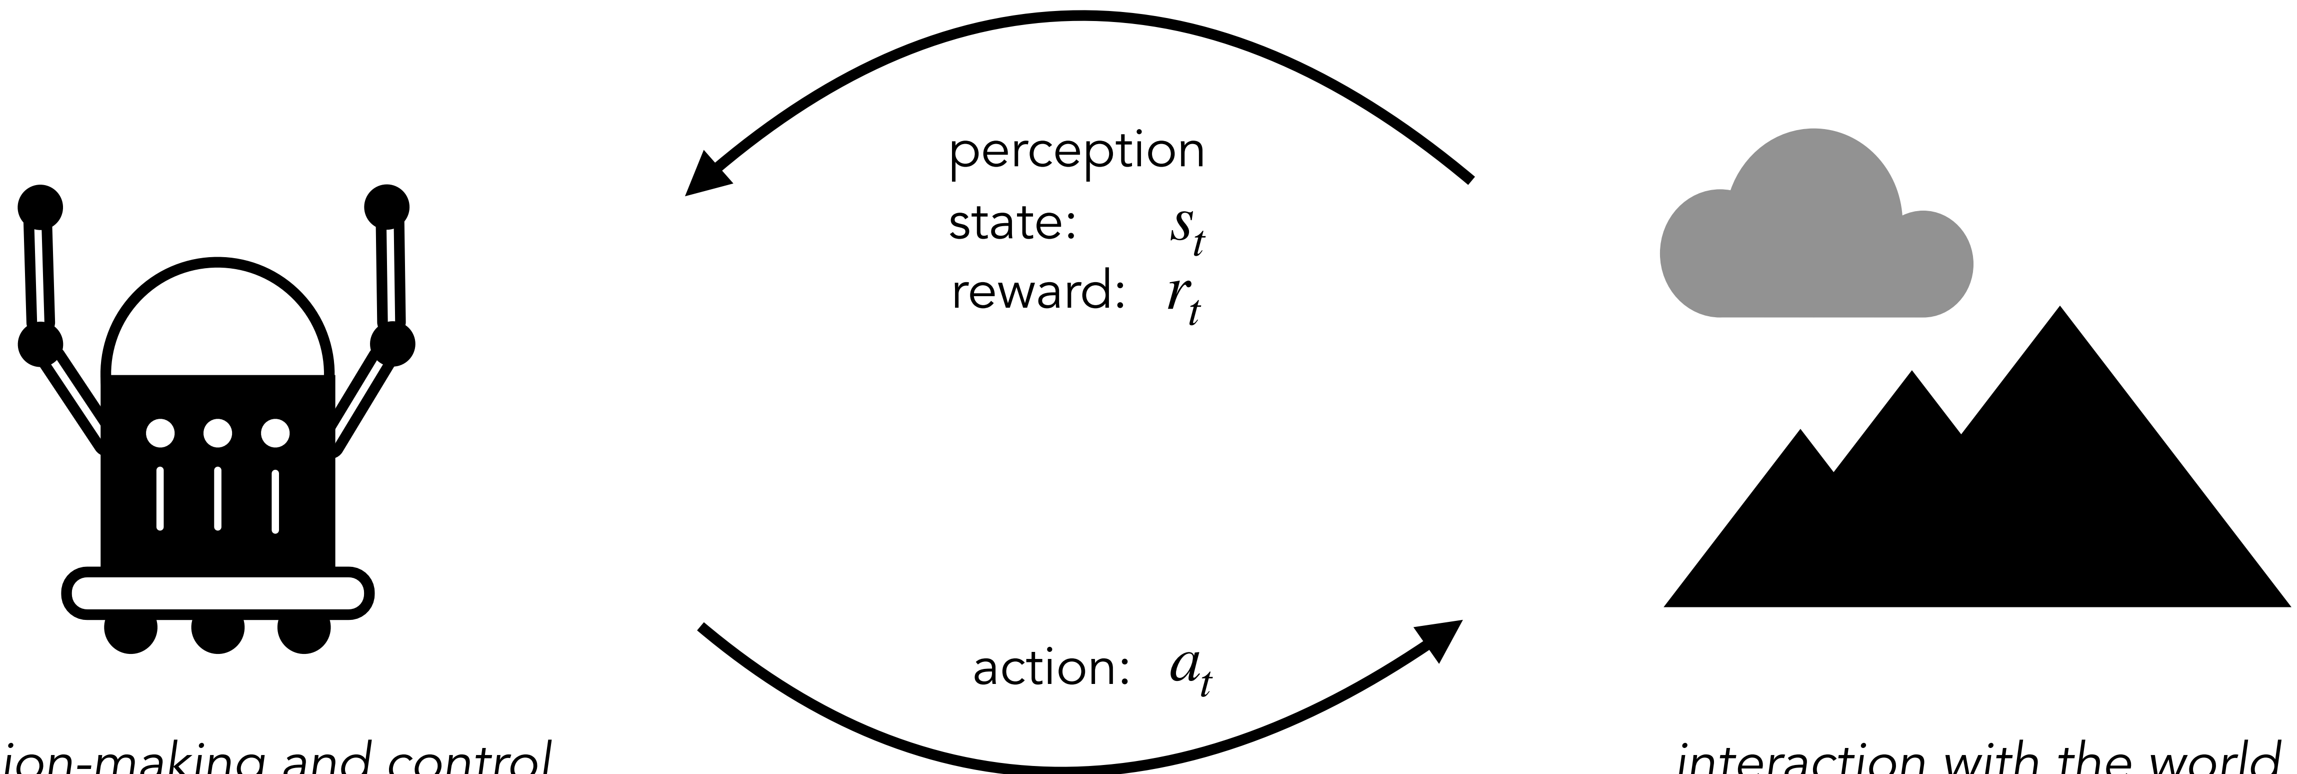

decision-making and control

interaction with the world

policy: $a_t \sim \pi(s_t) = \mathbb{P}(a_t \,|\, s_t)$

learning: $a_t^\star \sim \pi^\star(s_t)$

such that $\mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1}\right]$ is maximized

Markov decision process:

$$s_{t+1} \sim \mathbb{P}(s_{t+1} \,|\, s_t, a_t)$$

$$r_{t+1} \sim \mathbb{P}(r_{t+1} \,|\, s_t, a_t)$$

Successes of RL

DeepMind DQN (Atari) - 2015

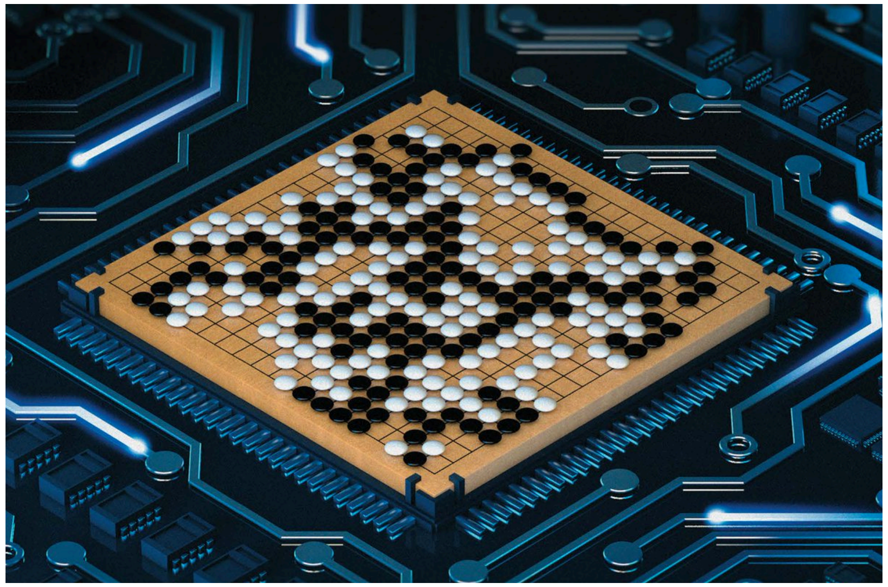

DeepMind AlphaGo (Go) - 2016

OpenAI (Dota 2) - 2018

DeepMind AlphaZero (Chess, Shogi, Go) - 2017

DeepMind AlphaStar (StarCraft II) - 2019

Challenges for Using RL in Robotics

- Remains extremely challenging:
  - **Goal specification** (not limited to robotics): The agent will maximize its sum of rewards. The reward needs to encode what we want the robot to do.
    - Be careful what you wish for! (e.g., paperclip maximizer from Bostrom, 2003).*
    - Reward-shaping is hard
  - **Safety**: There is limited exploration. The robots have to operate under some safety constraints.
  - **Real-world**: The real-world is messy and more noisy than typical games. Also, experiences in the real-world are not repeatable.
  - **Real-time**: We cannot learn faster than real-time. Data efficiency is really important.

> We need more samples to train the robots (due to noisy worlds) - but we are hindered by robots being limited by safety requirements and time.

* https://wiki.lesswrong.com/wiki/Paperclip_maximizer

Solution #1: Robot Farms

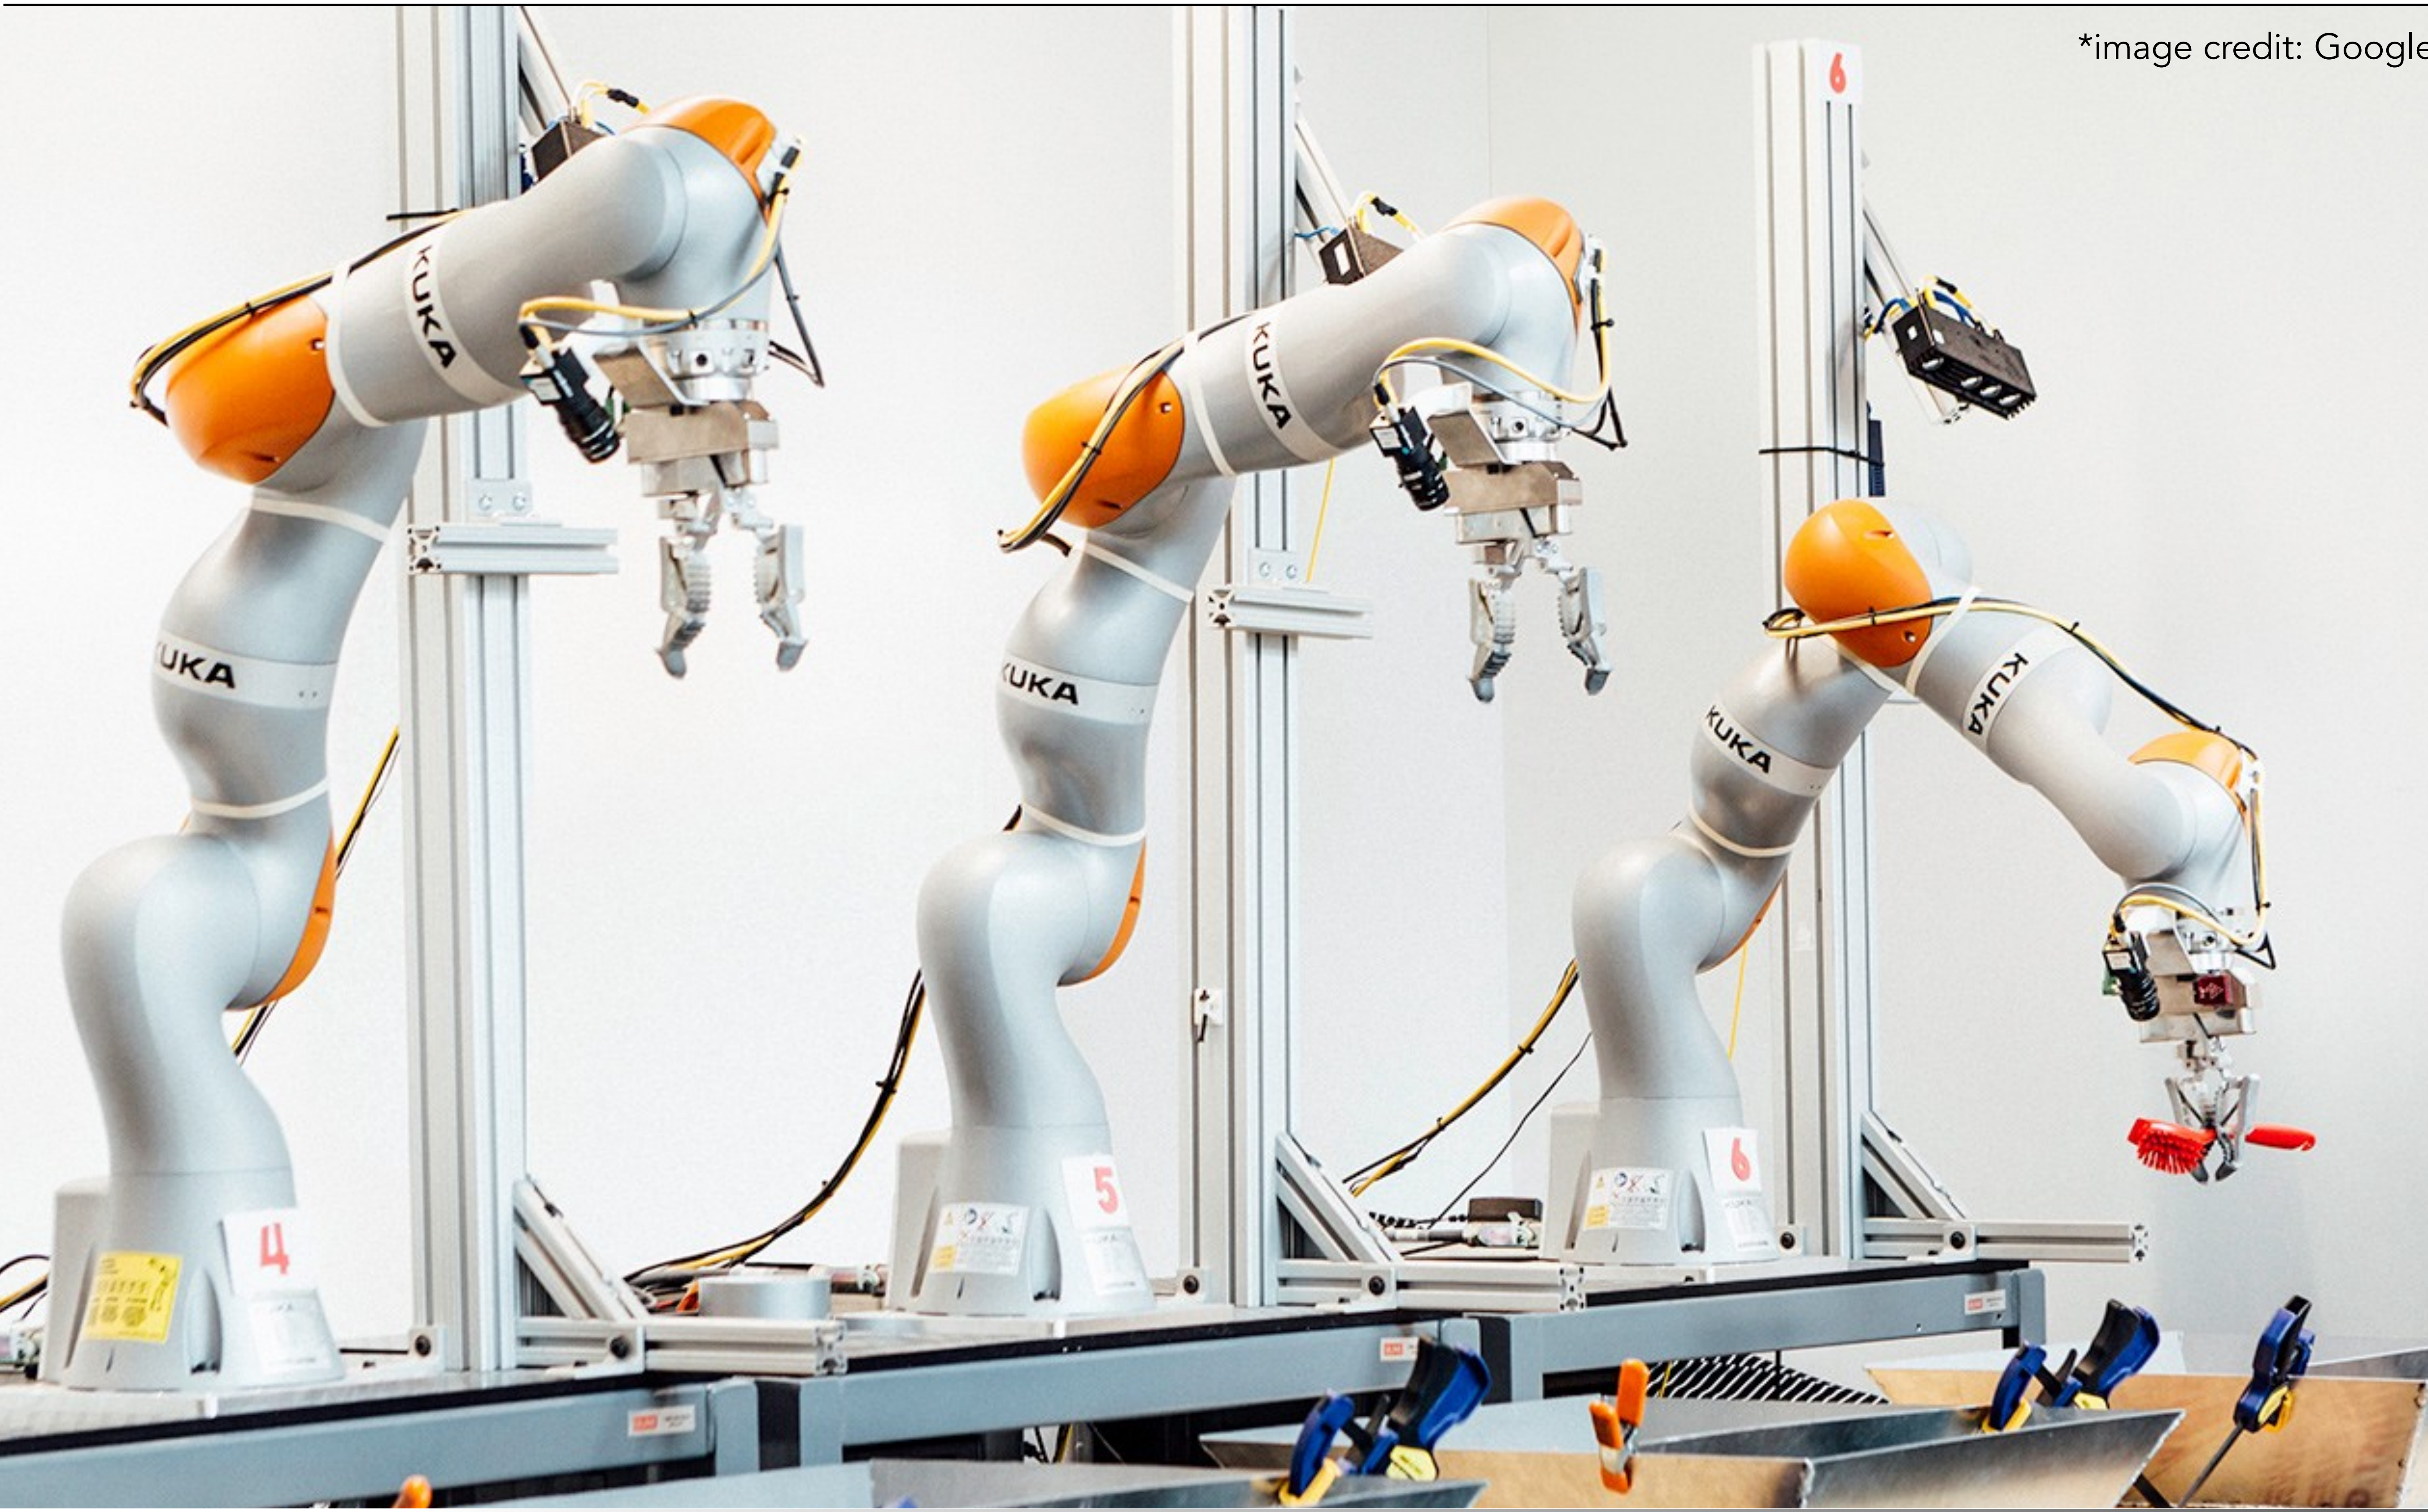

Solution #2: Sim2Real

- In simulation:
  - Create a faithful simulation environment (e.g., using Gazebo or Mujoco).
  - Add randomization to observed state (e.g., different color schemes for cameras, sensor noise), and to physics (e.g., actuation noise).
  - Learn a policy (or ensemble of policies) in simulation.
- In the real-world:
  - Normalize observations (e.g., preprocess real-world images to make them look like the simulation environment).
  - Run a few episodes to fine-tune the policy.

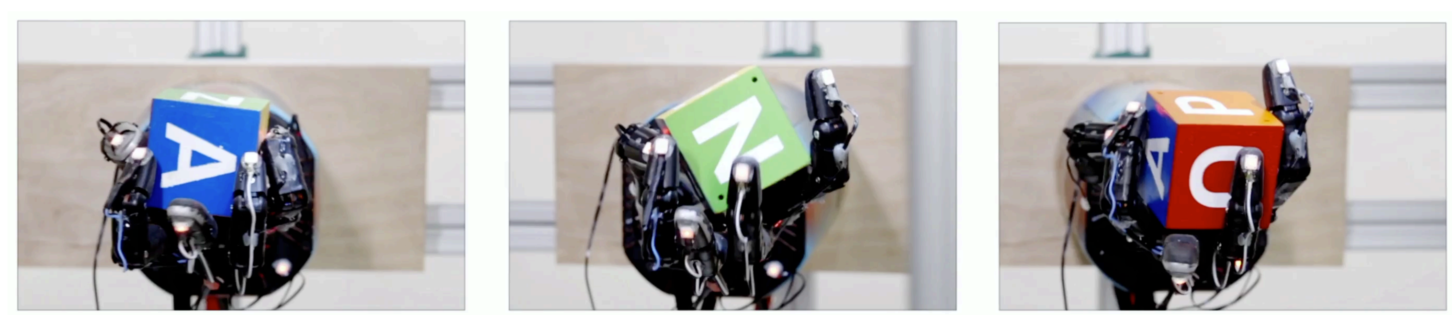

FINGER PIVOTING SLIDING FINGER GAITING

*Credits to OpenAI

Robot Ethics and Security - Challenges

"Robotics and AI may de-responsibilize people whenever an autonomous system could be blamed for a failure."

"Robotics could change the workforce structure and facilitate de-skilling of the work force (even in safety-critical contexts). " What if the AI gets it wrong?

"Excessive reliance on robotics leads to the delegation of tasks that should remain subject to human supervision."

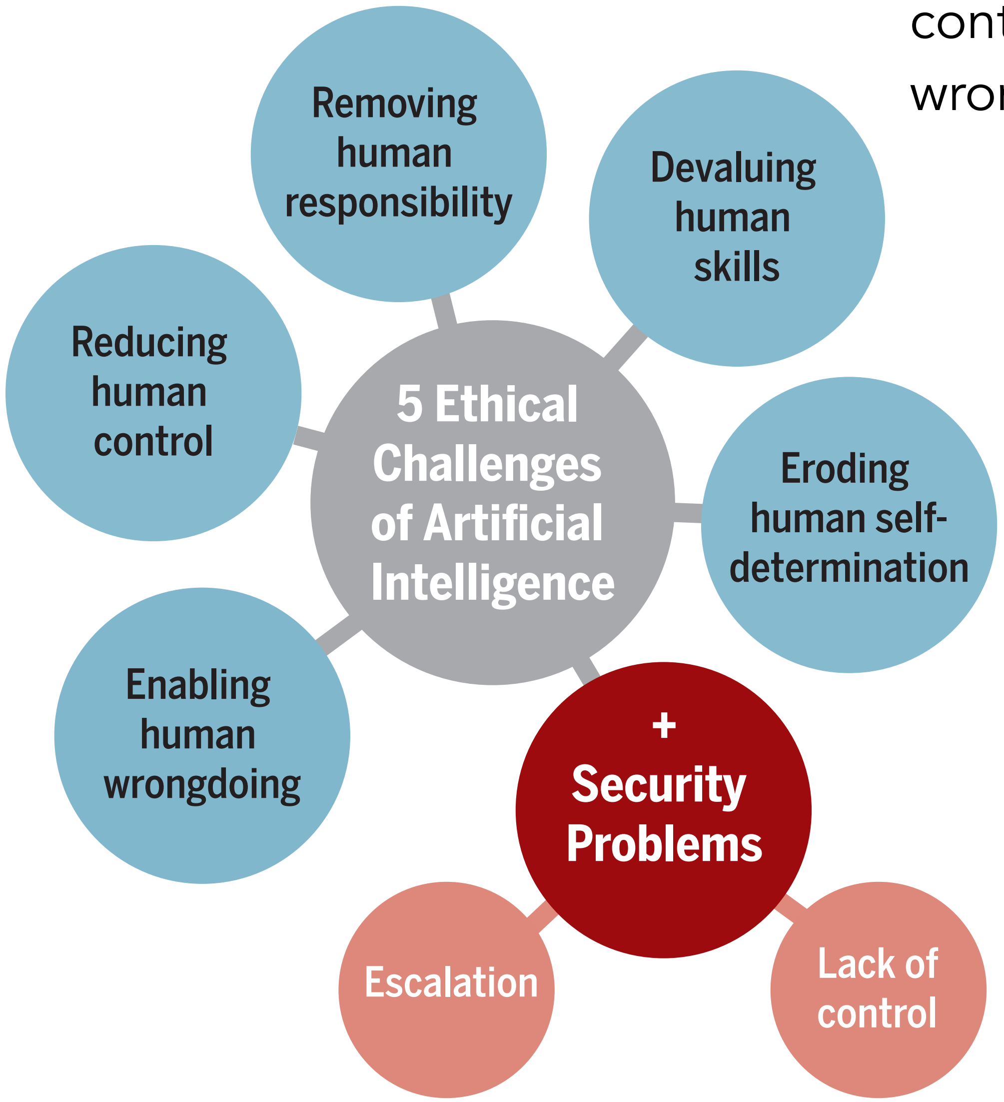

"Unwelcome changes in human behaviors to accommodate the routines that make automation work."

"Unethical application of robotics/AI by those who control it."

[Yang et al; The grand challenges if Science Robotics; 2018]

Robot Ethics and Security - Solutions

Laws need to be changed and adapted.

Education has a big role to play here. Why are we still teaching integrals and derivatives?

Gradual deployment of AI. Ultimately, only systems that are proven to be safe should be able to operate autonomously.

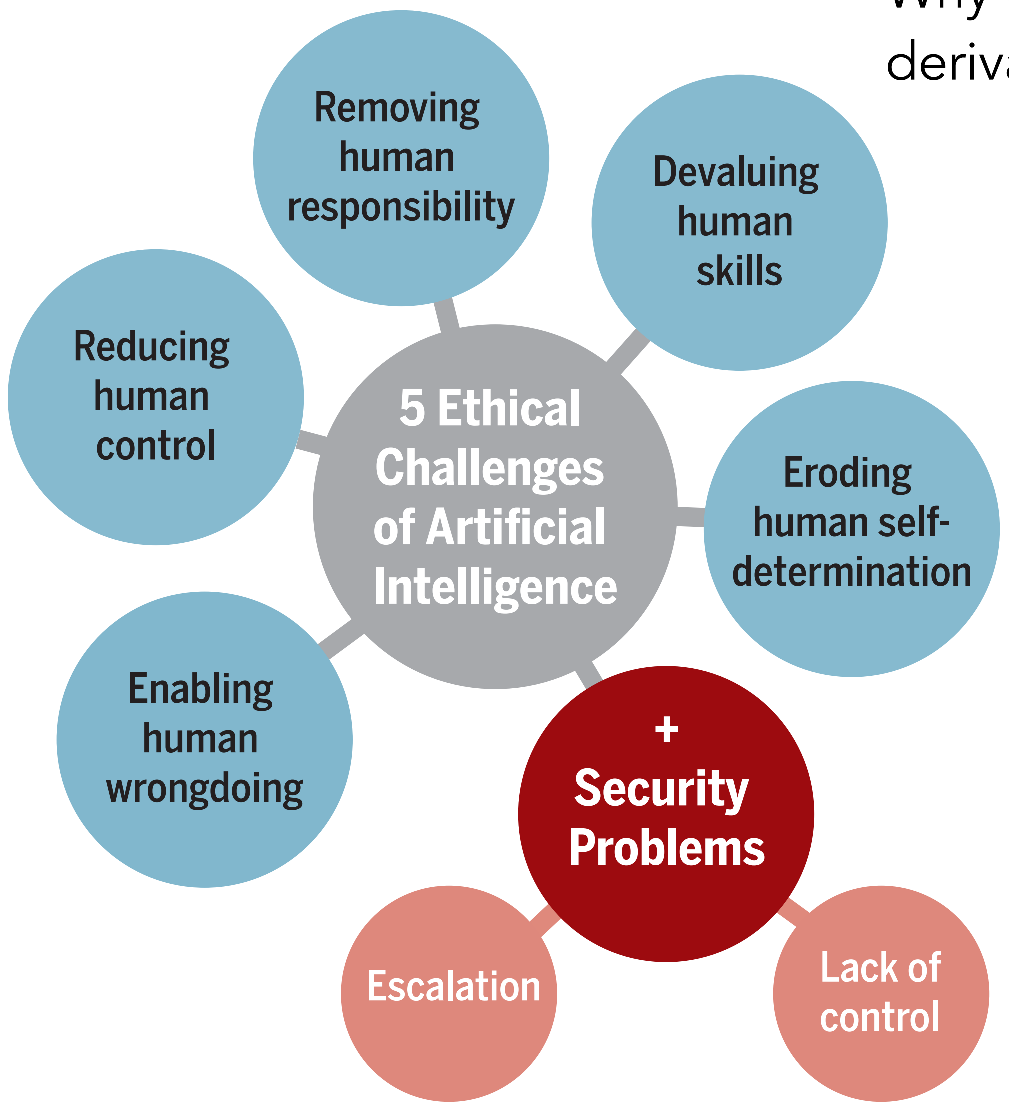

Develop systems that understand human behavior and adapt to it (rather than the opposite).

Enforce clear ethical guidelines at the governmental, academic and company level (e.g., Principles of Google AI, DeepMind Ethics & Society)

[Yang et al; The grand challenges if Science Robotics; 2018]

Further Announcements

- Mini-project support:
  - Tuesdays (up to March 5), 14:00-15:00 in Intel lab
  - Teaching Assistant support
- **Thanks!**
- If you are interested in staying involved… let me know.
  - Research projects
  - PhD in the future?
  - Collaborations