16-886 Special Topics

Models & Algorithms for Interactive Robotics

Instructor: Prof. Andrea Bajcsy

Carnegie Mellon University

intent ROBOTICS LAB

Welcome!

What is next?

Course Logistics

Course Contents

Intro Survey

(Intro to Dynamical Systems Models)

Course Logistics

Format: Mix of lectures and paper reading discussions

Typical 80-min class:

~5 min logistics and recap
70 min lecture, invited talk, or paper discussion

Office hours: W, 12:20 - 1:00 pm, NSH 4629 *(or by appointment)*

Resources:

Use ***Canvas*** for uploading assignments
Use ***course website*** for schedule, paper links, etc.

https://abajcsy.github.io/interactive-robotics/

Grading

*See **class syllabus** for detailed info*

Participation	(5%)
Homework (x1)	(10%)
Paper summaries	(10%)
Paper presentations	(15%)
Class project	(60%)

Participation (5%)

Expected to attend class in person—this is how we will all get the most out of the class!

I understand that occasionally you may have challenges attending (e.g., illness, religious observance,..); let me know

Homework (10%)

16-886: Interactive Robotics (Spring 2024) Prof. Andrea Bajcsy

Homework 1: Safety Analysis

In this homework, we will focus on computing backward reachable tubes (BRTs) and safe sets for several dynamical systems. For programming, you are welcome to choose among a variety of

Due Week 6
(Feb. 19)

This is a coding-based homework.
It is *not* meant to be tedious; it is meant to **empower** you! ☺

Paper Summaries + Presentations (25%)

Paper discussion days:

~8 paper reading days
2 papers per reading day

Before class:

write 1-2 paragraphs of paper review / takeaway / questions (must submit on Canvas)

In class:

Split you into small groups, discuss set of questions, I assign a representative from each group to present on the group's takeaways, and the whole class can engage on the answer

On paper reviews

Be **compassionate** *(e.g., invert your position)*

Be **constructive** *(e.g., what would you change to improve it?)*

Be **scholarly** *(e.g., cite sources, justify disagreements with proofs or citations)*

Daniel Dennett
Professor, Philosopher

"You should attempt to re-express your target's position so *clearly*, *vividly*, and *fairly* that your target says,

'Thanks, I wish I'd thought of putting it that way.'"

Slide adapted from Dorsa Sadigh

Class Project (60%)

Two options:

Research project:

Identify a research direction broadly relevant to this class
Propose and take first steps towards an original idea

Literature survey:

Select a topic area and rigorous way in which you will find papers
Characterize this topic area in an insightful way

Example of good literature survey

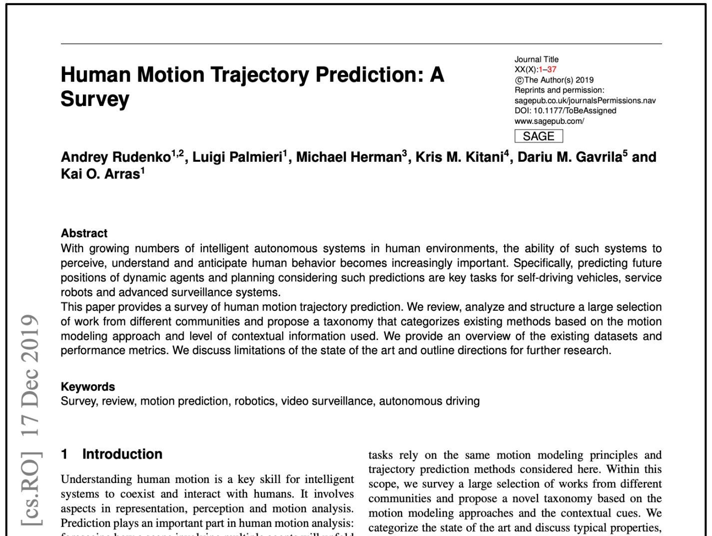
Journal Title
XX(X):1–37
©The Author(s) 2019
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/
SAGE

Human Motion Trajectory Prediction: A Survey

Andrey Rudenko[1,2], Luigi Palmieri[1], Michael Herman[3], Kris M. Kitani[4], Dariu M. Gavrila[5] and Kai O. Arras[1]

Abstract
With growing numbers of intelligent autonomous systems in human environments, the ability of such systems to perceive, understand and anticipate human behavior becomes increasingly important. Specifically, predicting future positions of dynamic agents and planning considering such predictions are key tasks for self-driving vehicles, service robots and advanced surveillance systems.
This paper provides a survey of human motion trajectory prediction. We review, analyze and structure a large selection of work from different communities and propose a taxonomy that categorizes existing methods based on the motion modeling approach and level of contextual information used. We provide an overview of the existing datasets and performance metrics. We discuss limitations of the state of the art and outline directions for further research.

Keywords
Survey, review, motion prediction, robotics, video surveillance, autonomous driving

[cs.RO] 17 Dec 2019

1 Introduction

Understanding human motion is a key skill for intelligent systems to coexist and interact with humans. It involves aspects in representation, perception and motion analysis. Prediction plays an important part in human motion analysis:

tasks rely on the same motion modeling principles and trajectory prediction methods considered here. Within this scope, we survey a large selection of works from different communities and propose a novel taxonomy based on the motion modeling approaches and the contextual cues. We categorize the state of the art and discuss typical properties,

Class Project (60%)

Project proposal (0%) -- due on Feb. 5

~1 page project summary. Identify the problem, background literature, potential solution

Mid-term report (20%) -- due on March. 18

~2 page writeup of progress, updated goals and timeline

Oral project presentation (10%) -- to be scheduled for Apr. 22 & Apr. 24

short presentations (~10 minutes but depends on number of people)

Final project report (30%) -- due on Apr. 24

~6 pages final report

Round of Introductions

- Name
- Department
- Year (Masters, PhD…)
- Research interests

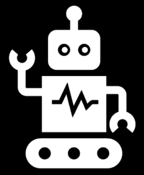

This class: *Interactive Robotics*

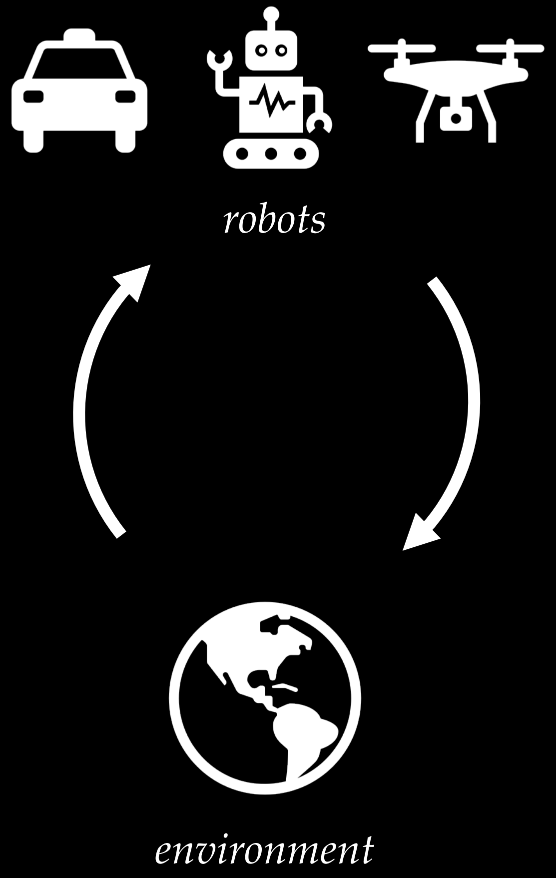
robots
environment

2000 *(ASIMO | Honda)*

Source: https://www.youtube.com/watch?v=82JFCciO3E4

2000 *(ASIMO | Honda)*

Source: https://www.youtube.com/watch?v=82JFCciO3E4

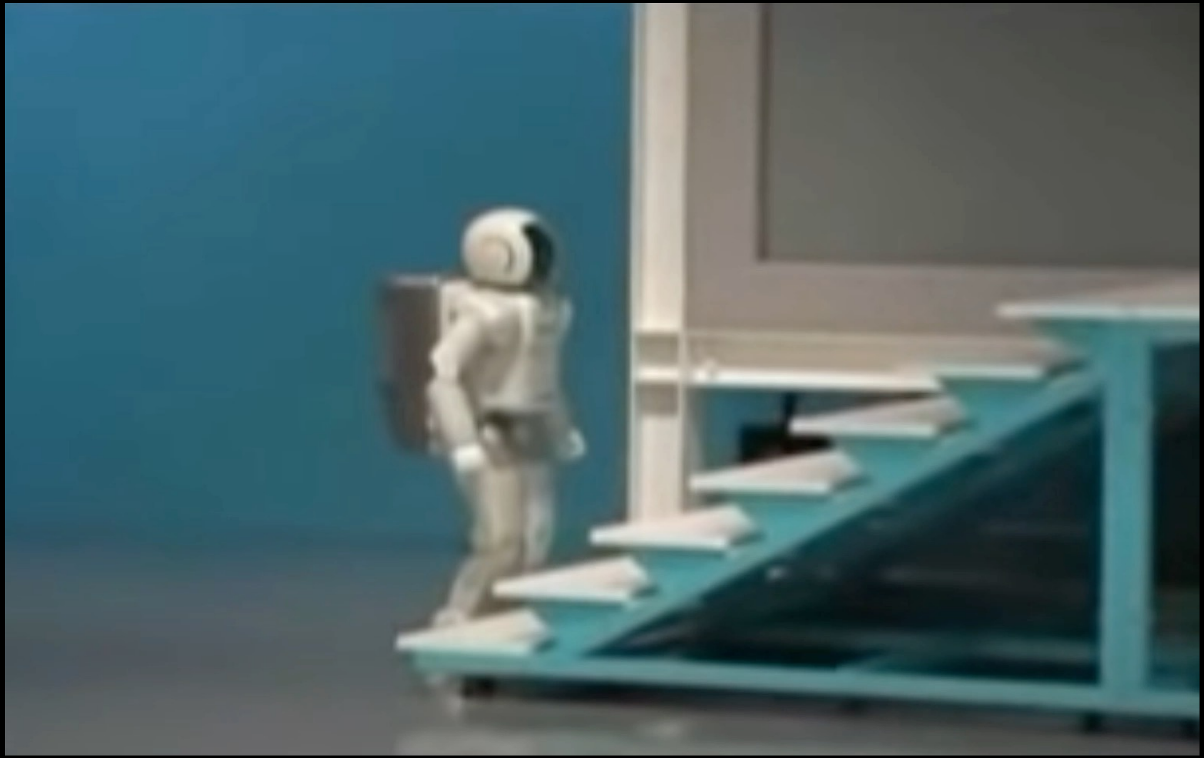

Source: https://www.youtube.com/watch?v=VTlV0Y5yAww

2023 *(Atlas | Boston Dynamics)*

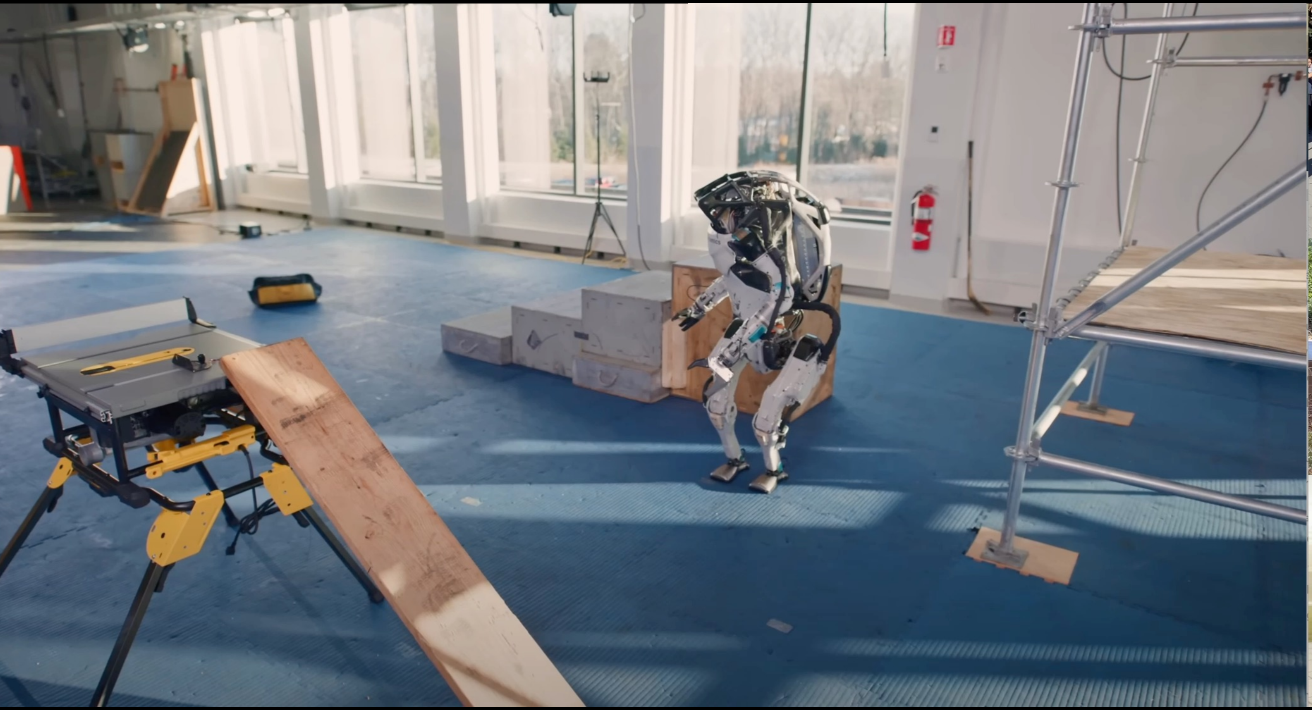

Source: https://www.youtube.com/watch?v=-e1_QhJ1EhQ

2023 *(Digit | Radosavovic & Xiao et. al)*

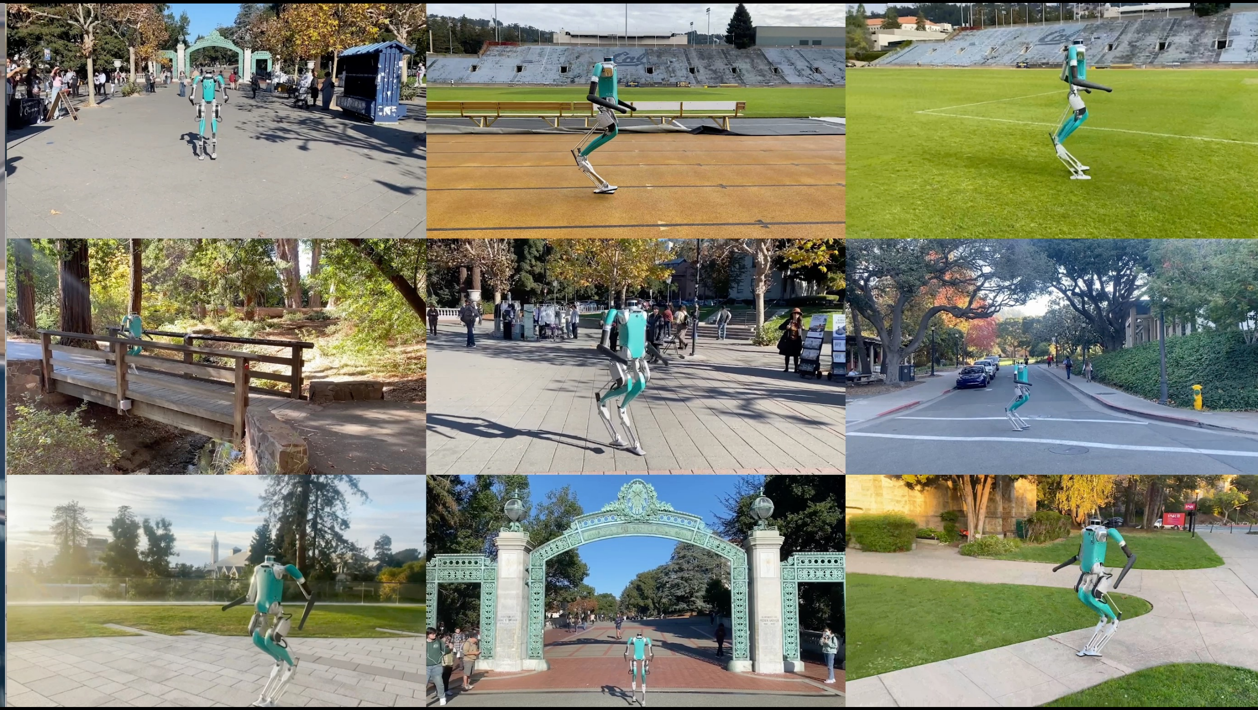

Source: https://learning-humanoid-locomotion.github.io/

1980s *(Ernst Dickmans)* ⟶ **2023** *(Nuro)*

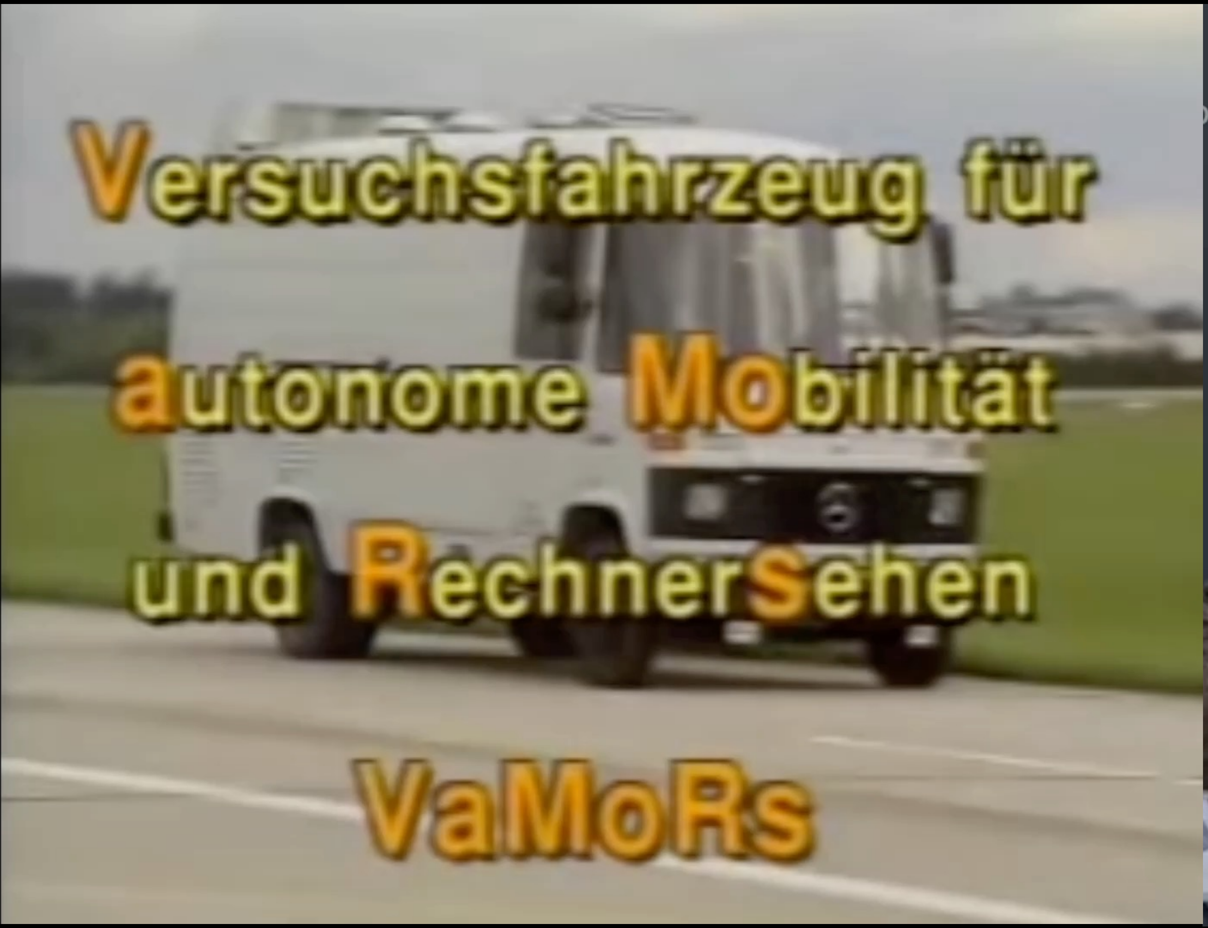

Source: https://www.youtube.com/watch?v=I39sxwYKlEE

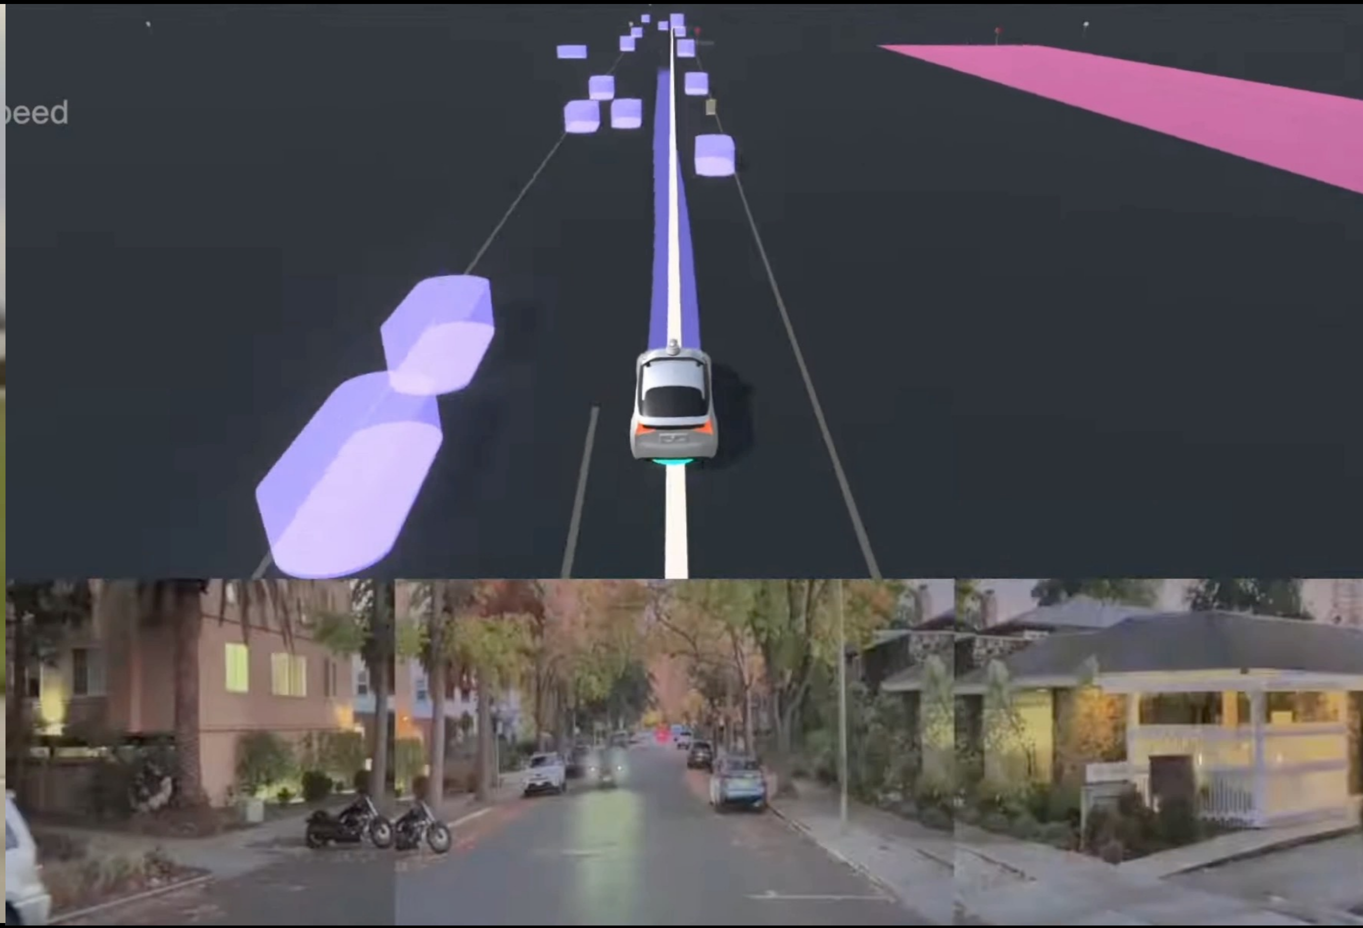

Source: https://www.youtube.com/watch?v=WDeZ3DTyQTI

We can start to consider deploying robots at scale!

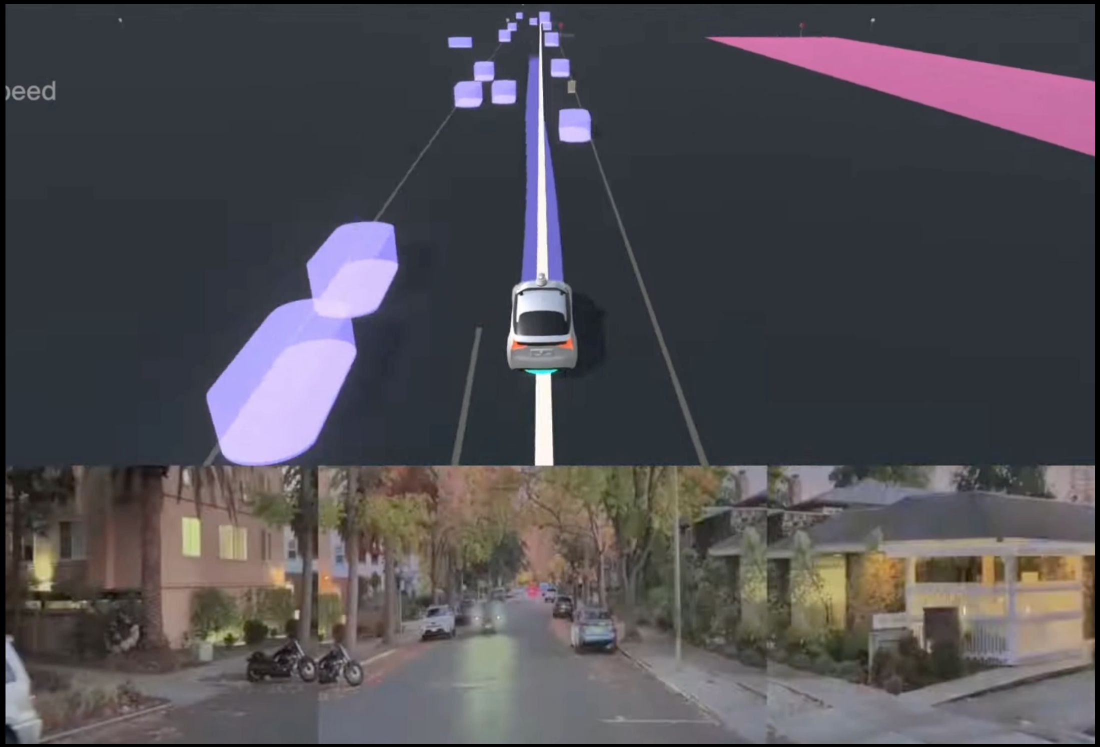

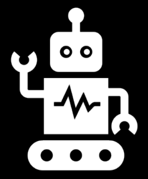

This class: *Interactive Robotics*

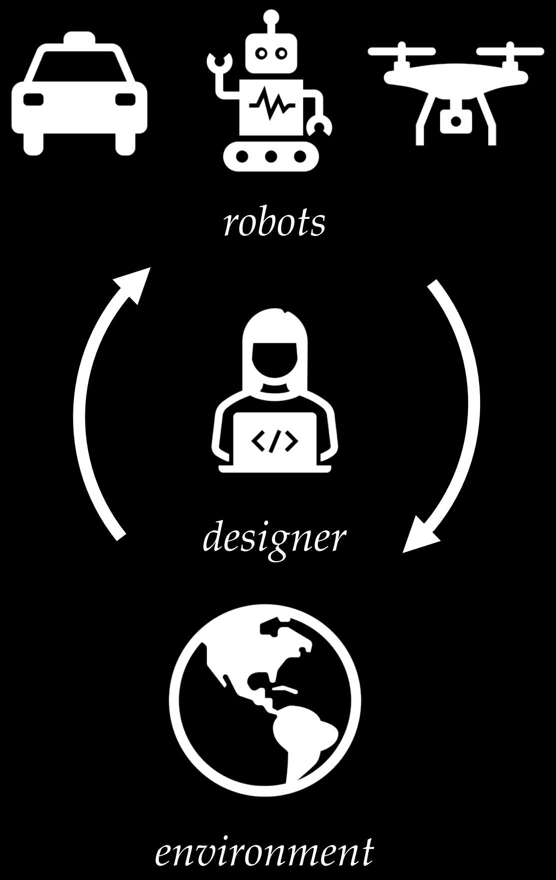
robots
designer
environment

designer

*I want a **safe** autonomous car*

i.e., “don’t collide”


```
             ⎧ brake    if d(you, front_car) < car_len
car_action = ⎨
             ⎩ speed    else
```

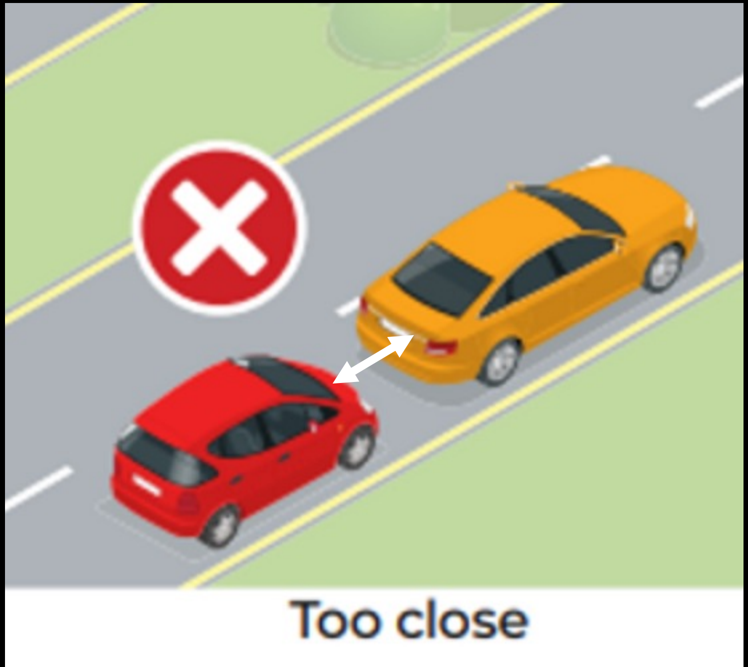

Env. topology

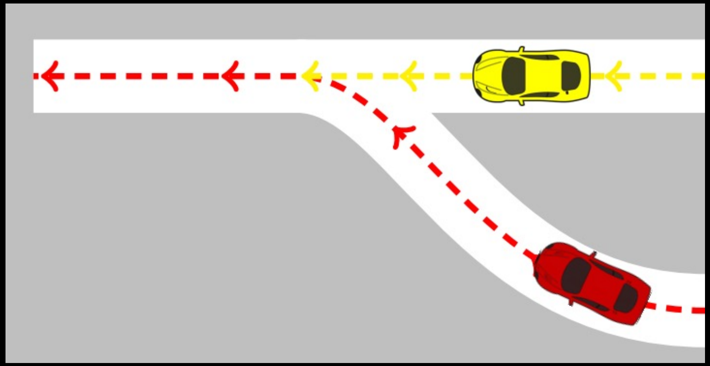

Relative speed

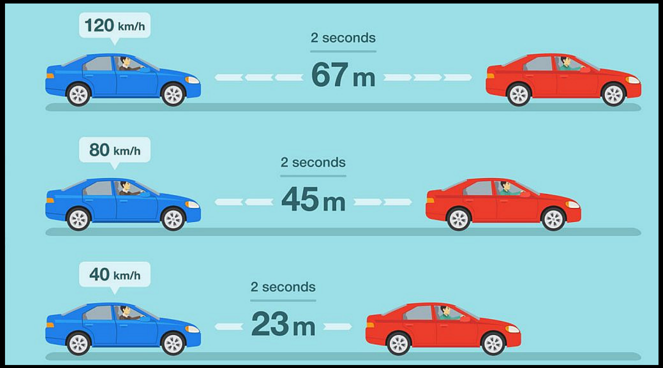

Weather

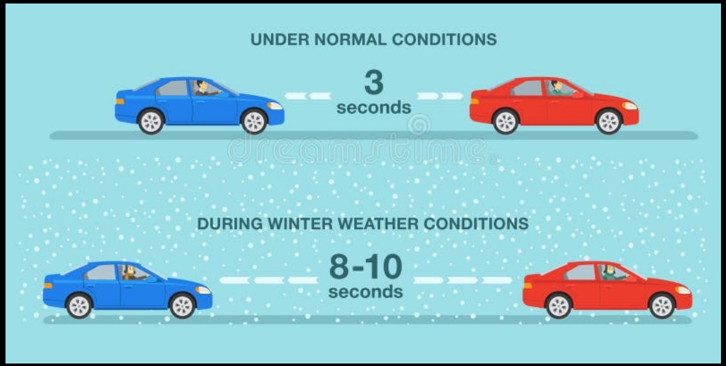

Many drivers

$$\texttt{car\_action} = \begin{cases} \texttt{brake} & \texttt{if d(you, front\_car) < car\_len} \\ \texttt{speed} & \texttt{else} \end{cases}$$

On a Formal Model of Safe and Scalable Self-driving Cars

Shai Shalev-Shwartz, Shaked Shammah, Amnon Shashua

Definition 1 (Safe longitudinal distance — same direction) *A longitudinal distance between a car c_r that drives behind another car c_f, where both cars are driving at the same direction, is safe w.r.t. a response time ρ if for any braking of at most $a_{\max,\text{brake}}$, performed by c_f, if c_r will accelerate by at most $a_{\max,\text{accel}}$ during the response time, and from there on will brake by at least $a_{\min,\text{brake}}$ until a full stop then it won't collide with c_f.*

Lemma 2 below calculates the safe distance as a function of the velocities of c_r, c_f and the parameters in the definition.

Lemma 2 *Let c_r be a vehicle which is behind c_f on the longitudinal axis. Let $\rho, a_{\max,\text{brake}}, a_{\max,\text{accel}}, a_{\min,\text{brake}}$ be as in Definition 1. Let v_r, v_f be the longitudinal velocities of the cars. Then, the minimal safe longitudinal distance between the front-most point of c_r and the rear-most point of c_f is:*

$$d_{\min} = \left[v_r\,\rho + \frac{1}{2} a_{\max,\text{accel}}\,\rho^2 + \frac{(v_r + \rho\, a_{\max,\text{accel}})^2}{2a_{\min,\text{brake}}} - \frac{v_f^2}{2a_{\max,\text{brake}}} \right]_+ ,$$

where we use the notation $[x]_+ := \max\{x, 0\}$.

NVIDIA

The Safety Force Field

David Nistér, Hon-Leung Lee, Julia Ng, Yizhou Wang

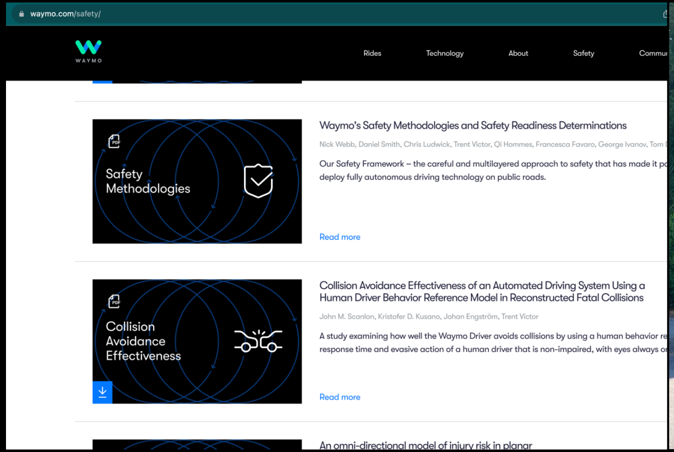

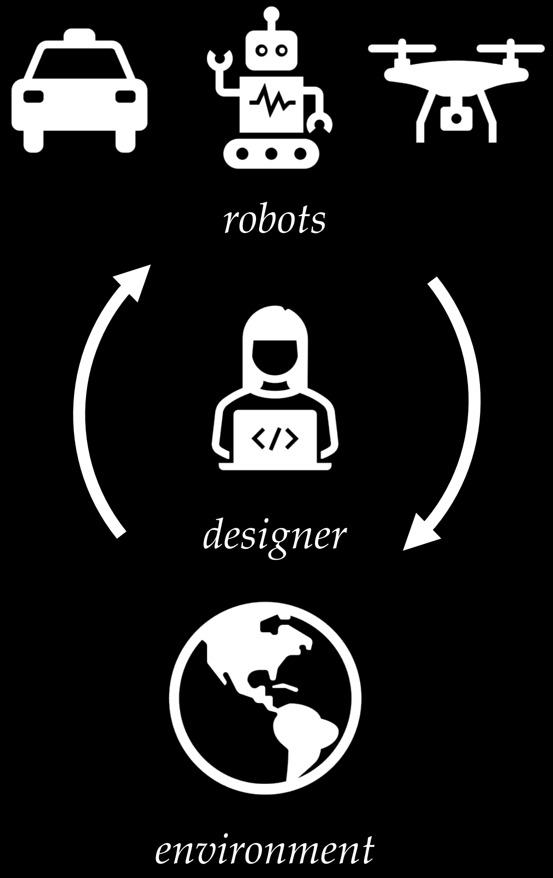

robots
designer
environment

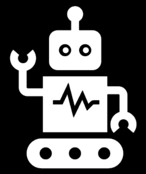

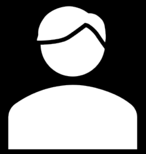

What other kinds of interactions can robots have with people?

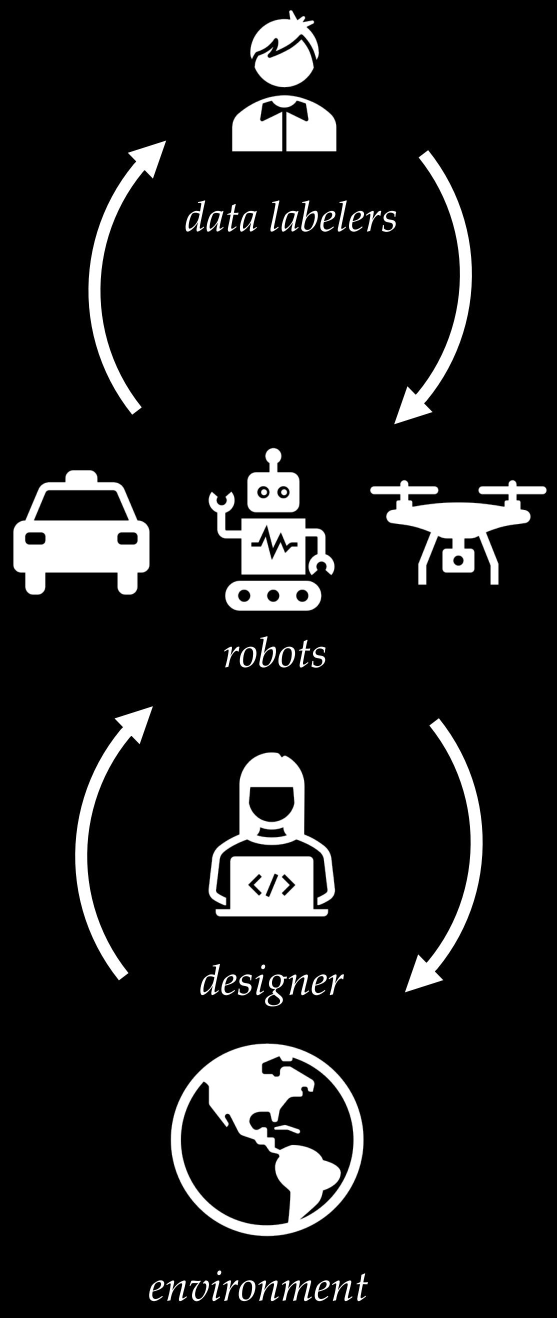
data labelers
robots
designer
environment

1957

Source: https://youtu.be/oNA1_yOq-jw?feature=shared

2024 *(ALOHA | Fu*, Zhao* & Finn)*

Source: https://mobile-aloha.github.io/

CNBC

Search quotes, news & videos | WATCHLIST

MARKETS | BUSINESS | INVESTING | TECH | POLITICS | CNBC TV | INVESTING CLUB | PRO | MAKE IT

Cruise confirms robotaxis rely on human assistance every four to five miles

PUBLISHED MON, NOV 6 2023•7:11 PM EST | UPDATED MON, NOV 6 2023•7:16 PM EST

Lora Kolodny
@IN/LORAKOLODNY/

SHARE

KEY POINTS

- GM-owned Cruise is responding to allegations that its cars are not really self-driving because they require frequent help from humans working as “remote assistants” to get through tricky drives.
- Cruise tells CNBC it worked with roughly one “remote assistant agent,” per every 15 to 20 driverless vehicles in its fleet before grounding operations last month.
- Human advisors generally provide “wayfinding intel” to the robotaxis, and do not drive them remotely, a company spokesperson said.

CNBC TV
Squawk on t
Street
UP NEXT | Money Mo

TRENDING NOW

Source: https://shorturl.at/lnDJ9

Source: https://shorturl.at/cjpEV

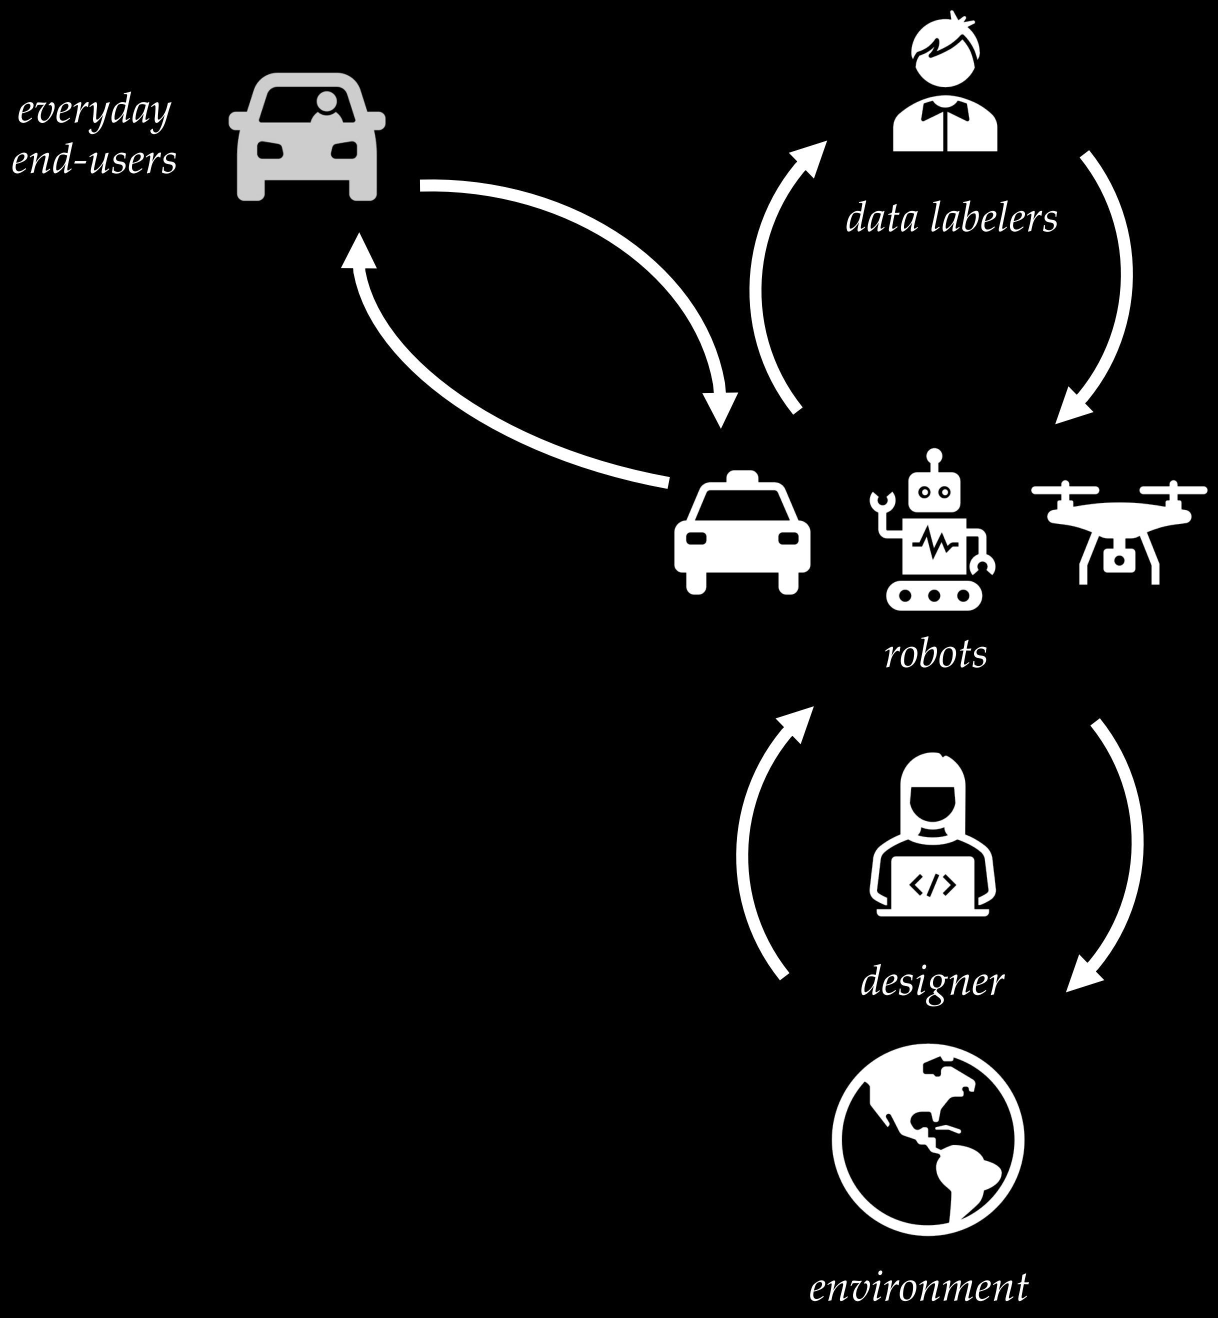
everyday
end-users
data labelers
robots
designer
environment

Source: https://twitter.com/nitguptaa/

[Autoevolution, “Tesla with FSD Beta Avoids Plastic Bag, Some People Mistake It for a Good Thing“]

Human

Place the bowl in the microwave, please.

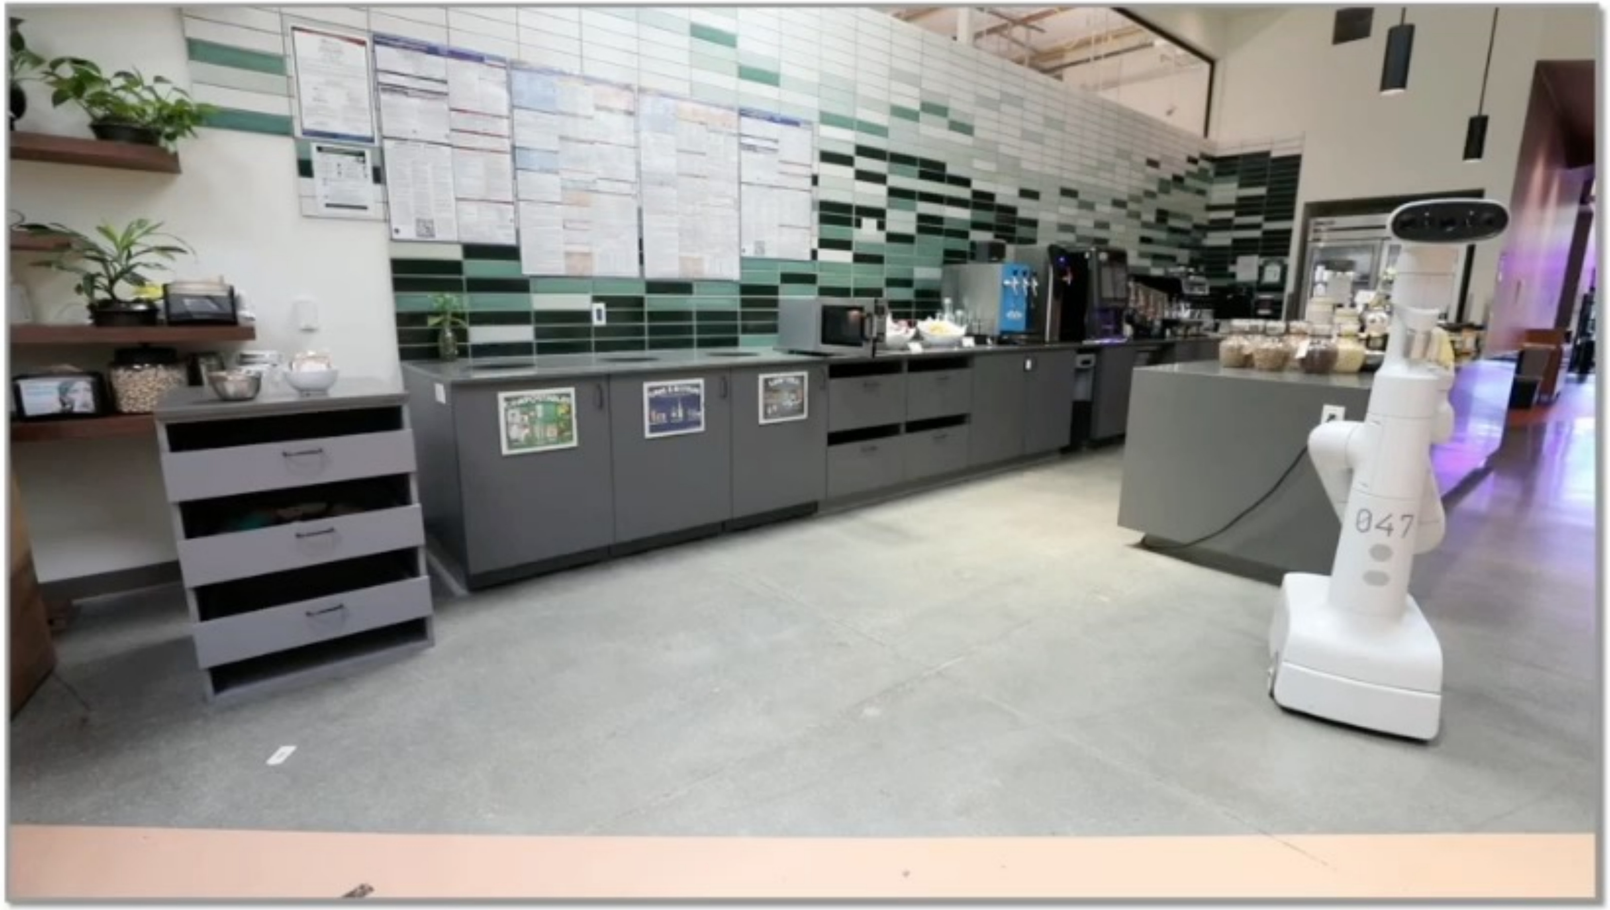

Source: https://robot-help.github.io/

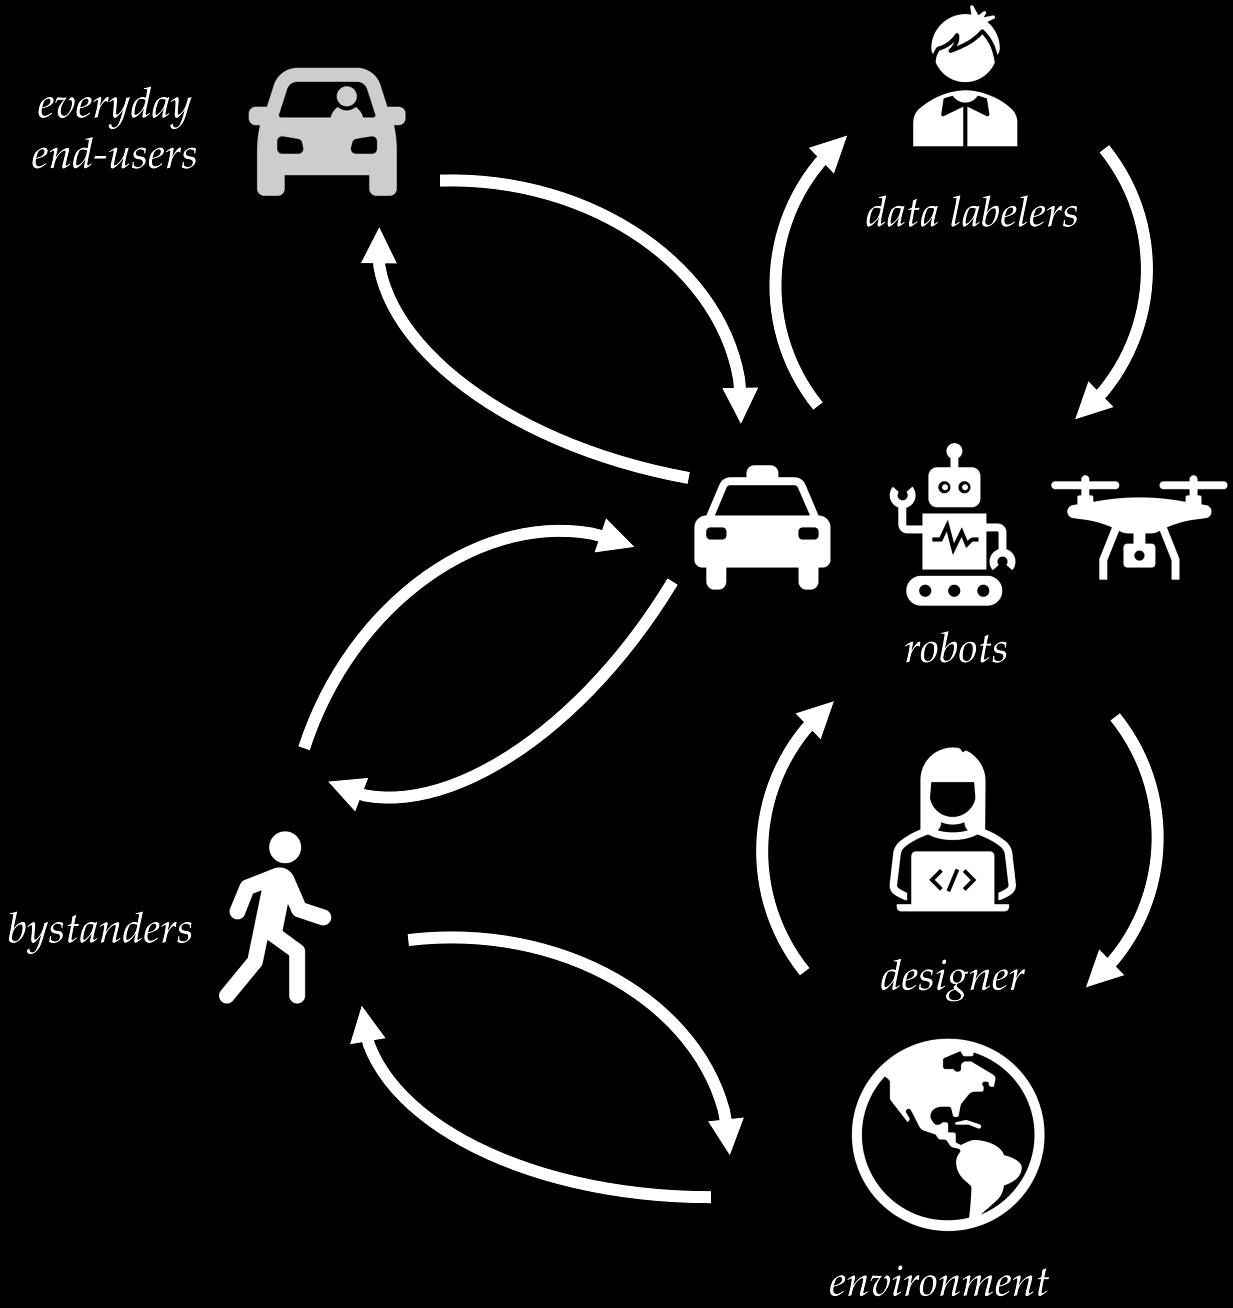
everyday
end-users
data labelers
robots
bystanders
designer
environment

Source: https://abc7news.com/

Source: https://twitter.com/djbaskin

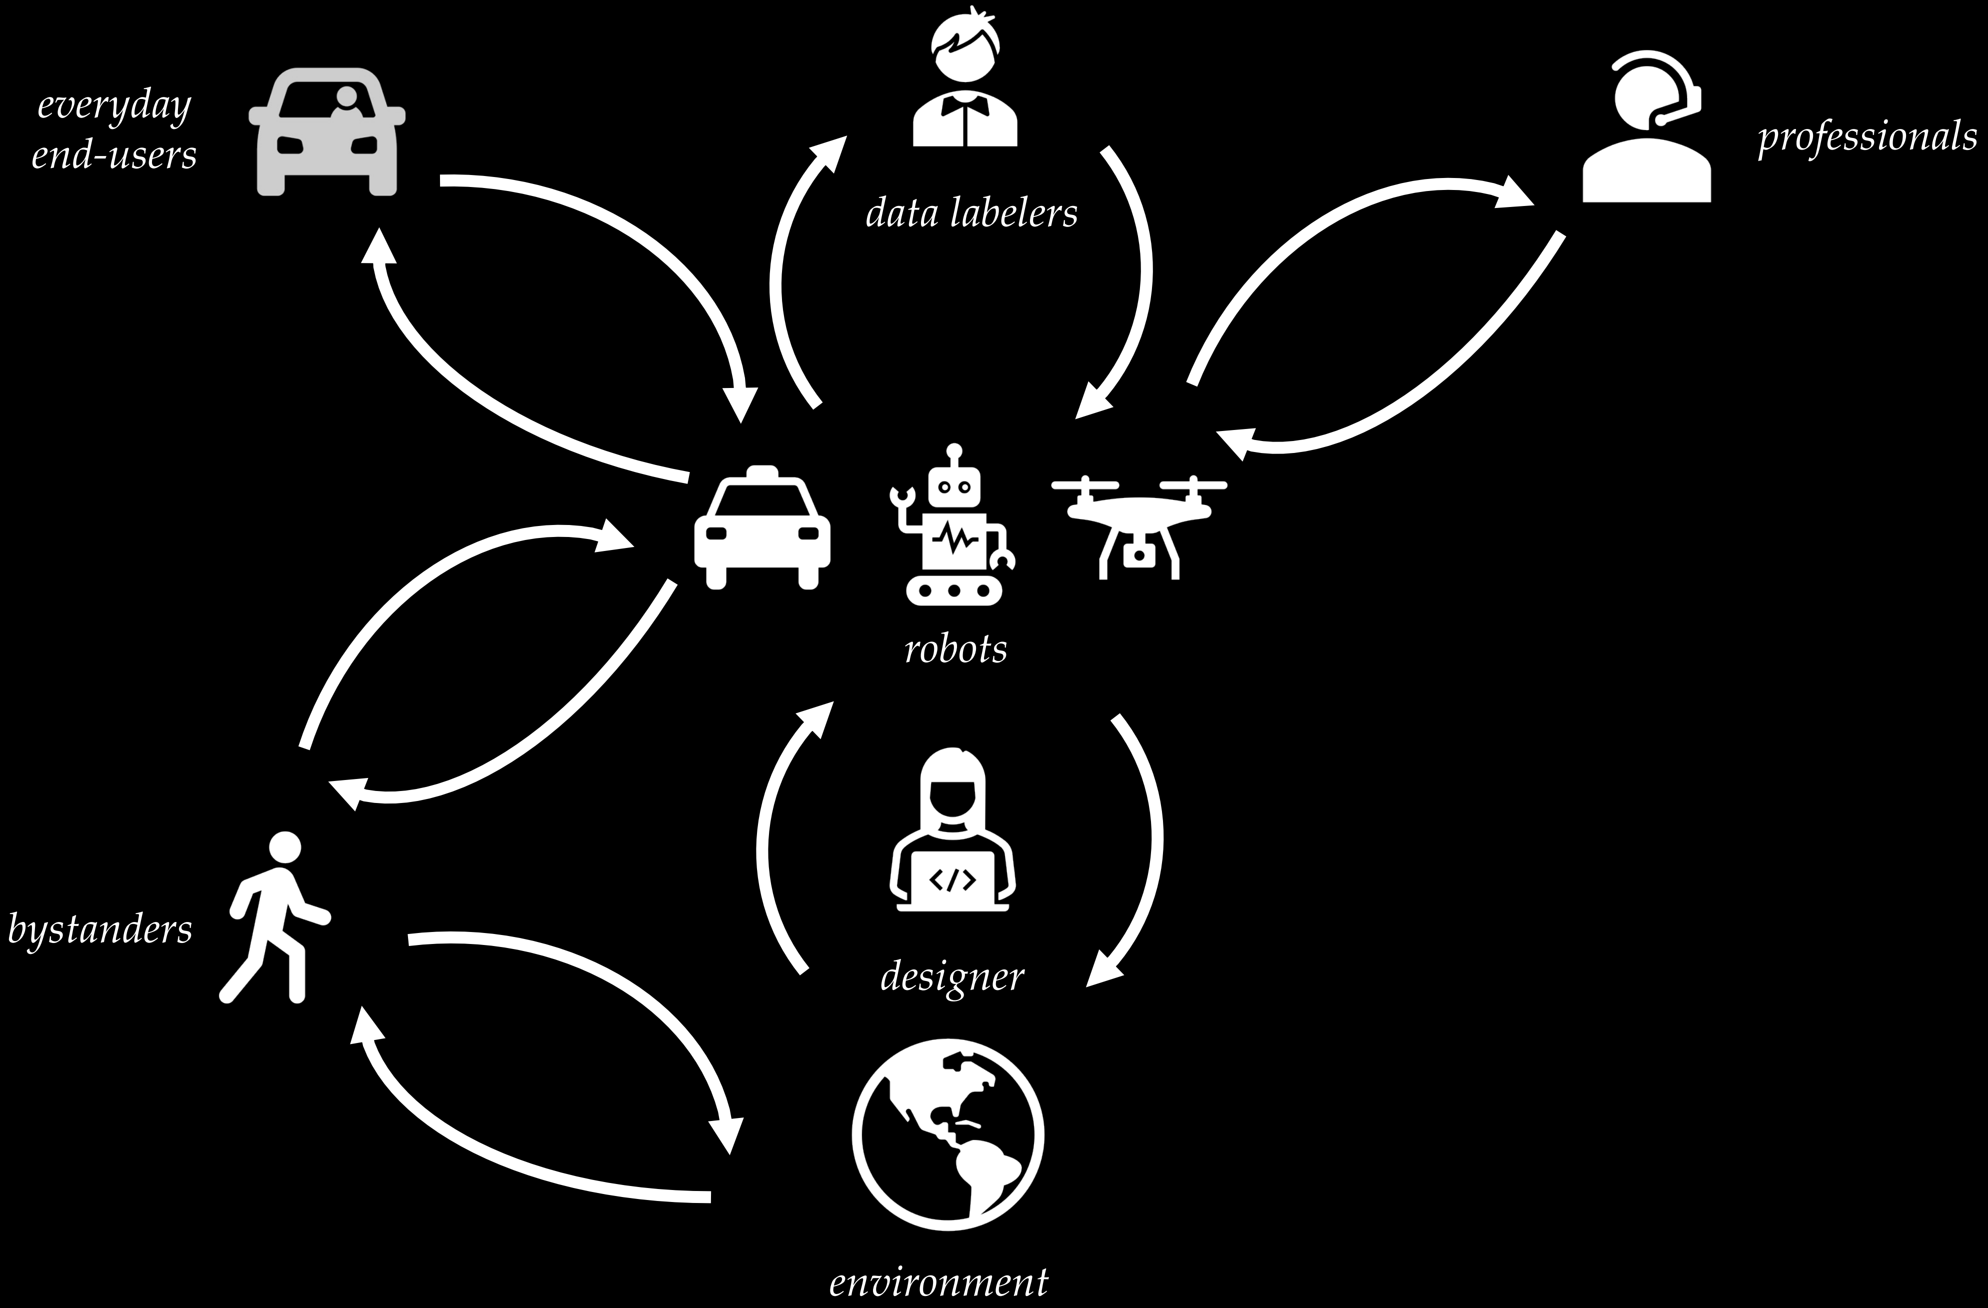
everyday
end-users
data labelers
professionals
robots
bystanders
designer
environment

Source: https://shorturl.at/nqyRW

Source: https://spectrum.ieee.org/tag/davinci-robot

Boeing Built Deadly Assumptions Into 737 Max, Blind to a Late Design Change

Share full article 1.5K

After Boeing removed one of the sensors from an automated flight system on its 737 Max, the jet's designers and regulators still proceeded as if there would be two. Ruth Fremson/The New York Times

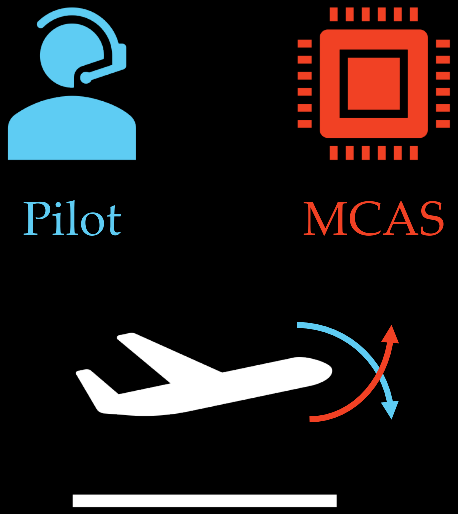

the 737 Max to market. And The Times's investigation details how an essential software system known as MCAS was implemented with insufficient oversight and inadequate pilot training.

Source: https://www.nytimes.com/2019/06/01/business/boeing-737-max-crash.html

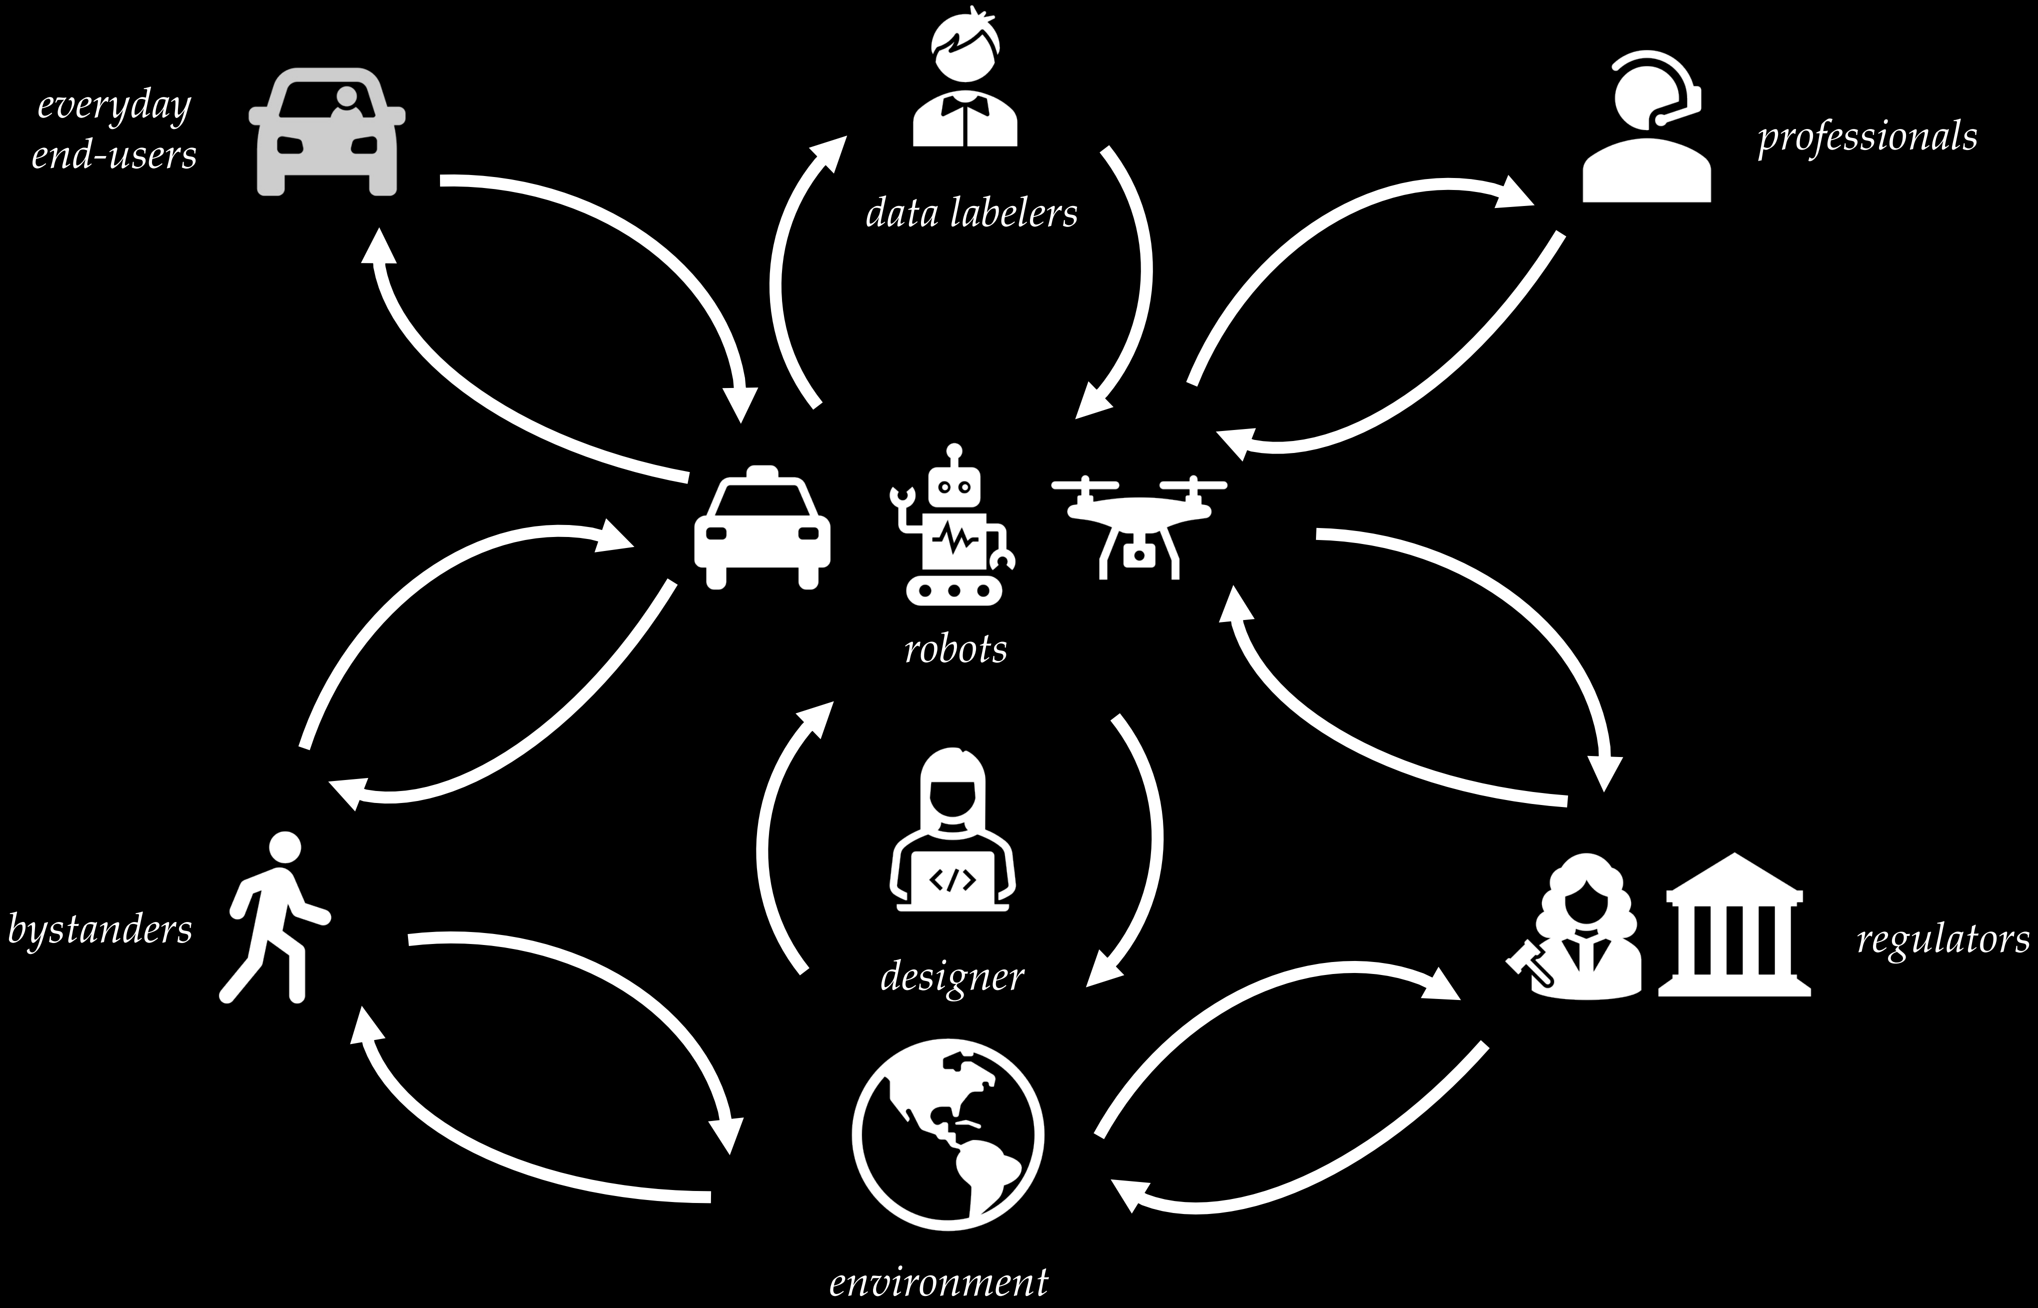
everyday
end-users
data labelers
professionals
robots
bystanders
designer
regulators
environment

Standards

2846-2022 - IEEE Standard for Assumptions in Safety-Related Models for Automated Driving Systems

Publisher: IEEE Cite This PDF

Additional content is available

Status: Active - Approved

5 Cites in Papers

1057 Full Text Views

Abstract

Figures

References

Citations

Keywords

Definitions

Abstract:

This standard applies to road vehicles. It defines a minimum set of reasonable assumptions and foreseeable scenarios that shall be considered in the development of safety related models that are part of an automated driving system (ADS).

Scope:

This standard applies to road vehicles. For a set of scenarios, a minimum set of assumptions regarding reasonably foreseeable behaviors of other road users are defined that shall be considered in the development of safety-related models for automated driving systems (ADS).

Source: https://ieeexplore.ieee.org/document/9761121

Biden Issues Executive Order to Create A.I. Safeguards

The sweeping order is a first step as the Biden administration seeks to put guardrails on a global technology that offers great promise but also carries significant dangers.

Share full article

The order is an effort by President Biden to show that the United States, considered the leading power in fast-moving artificial intelligence technology, will also take the lead in its regulation. Doug Mills/The New York Times

Source: https://shorturl.at/kuFP2

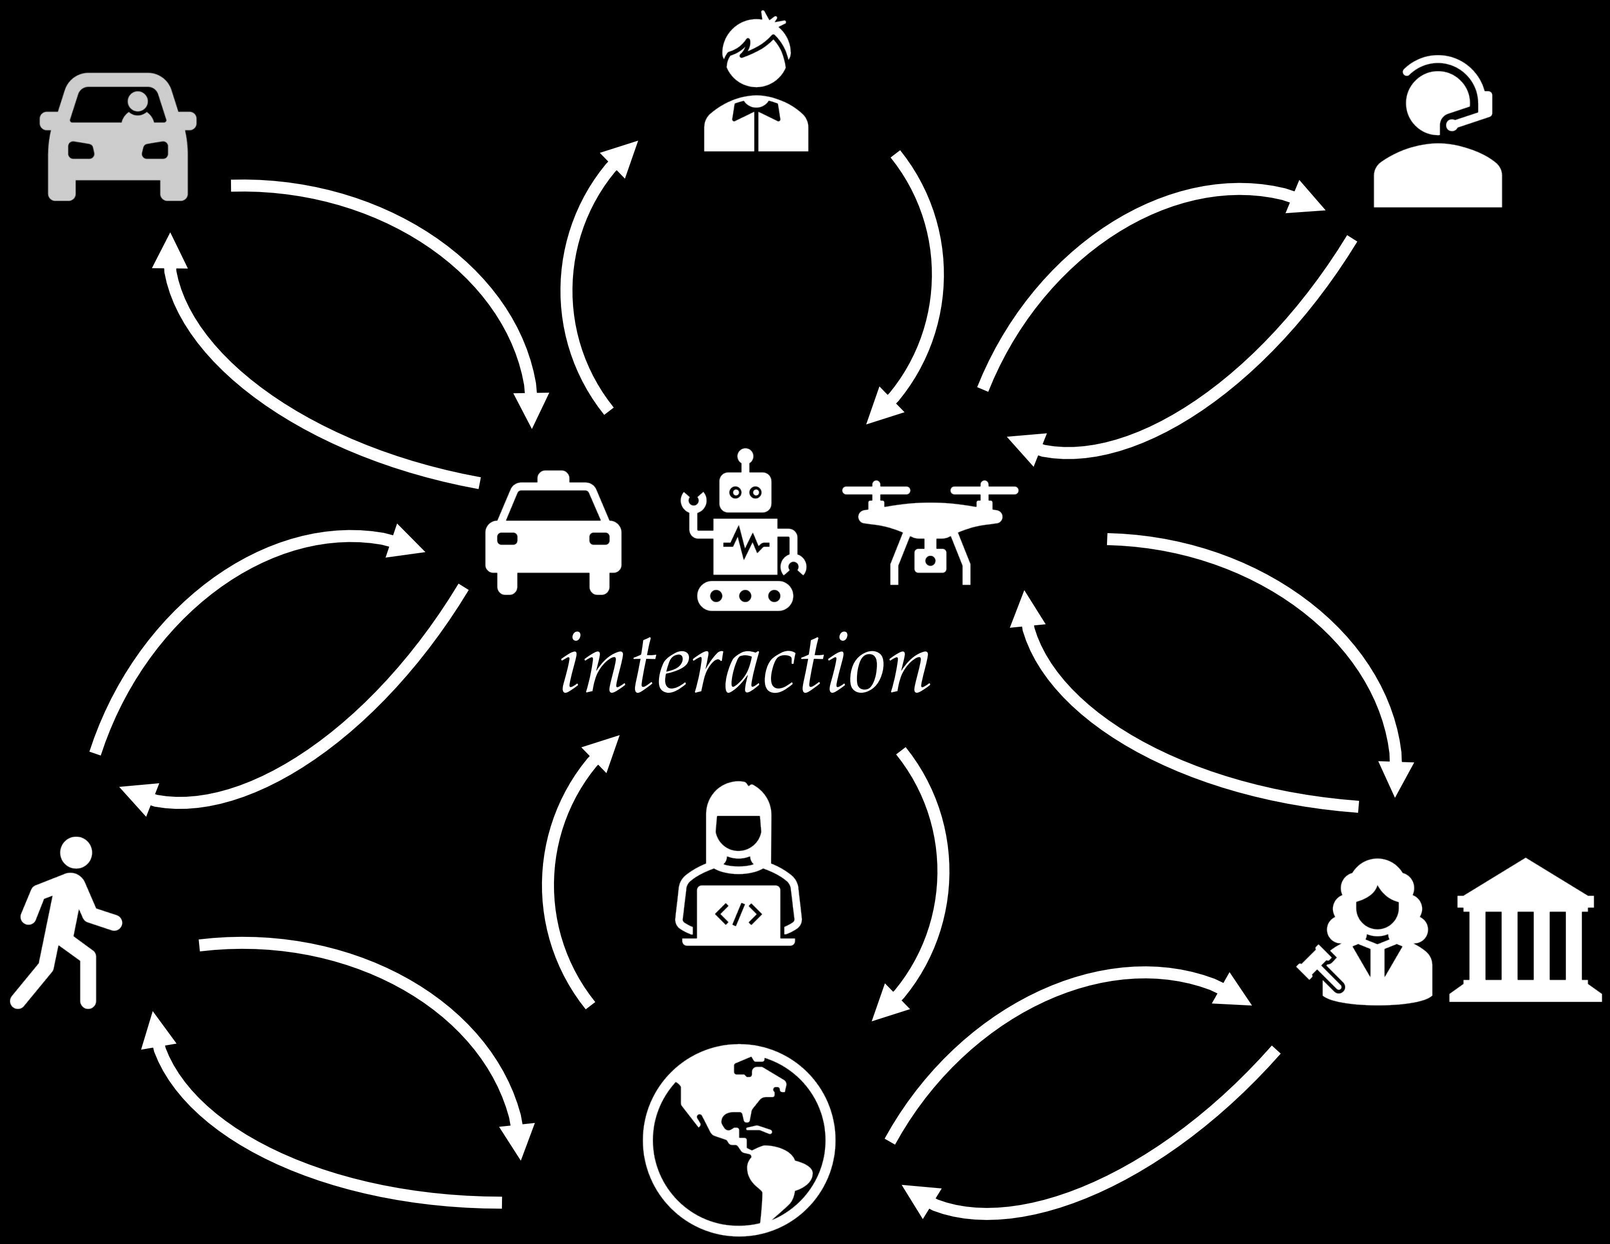
interaction

interaction means there exists a **feedback loop** between **human stakeholders** and **autonomous robots**

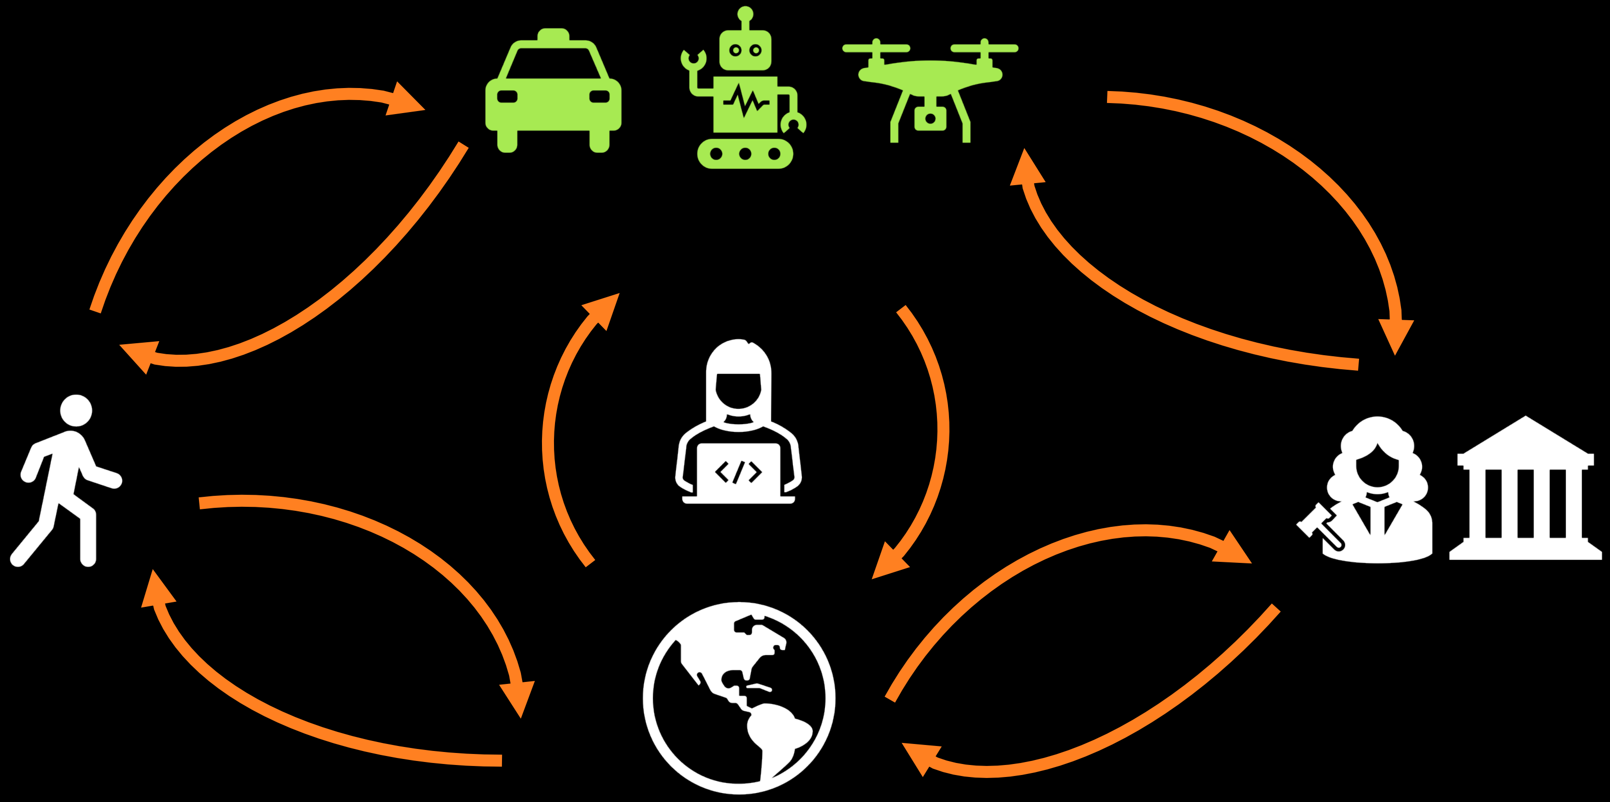

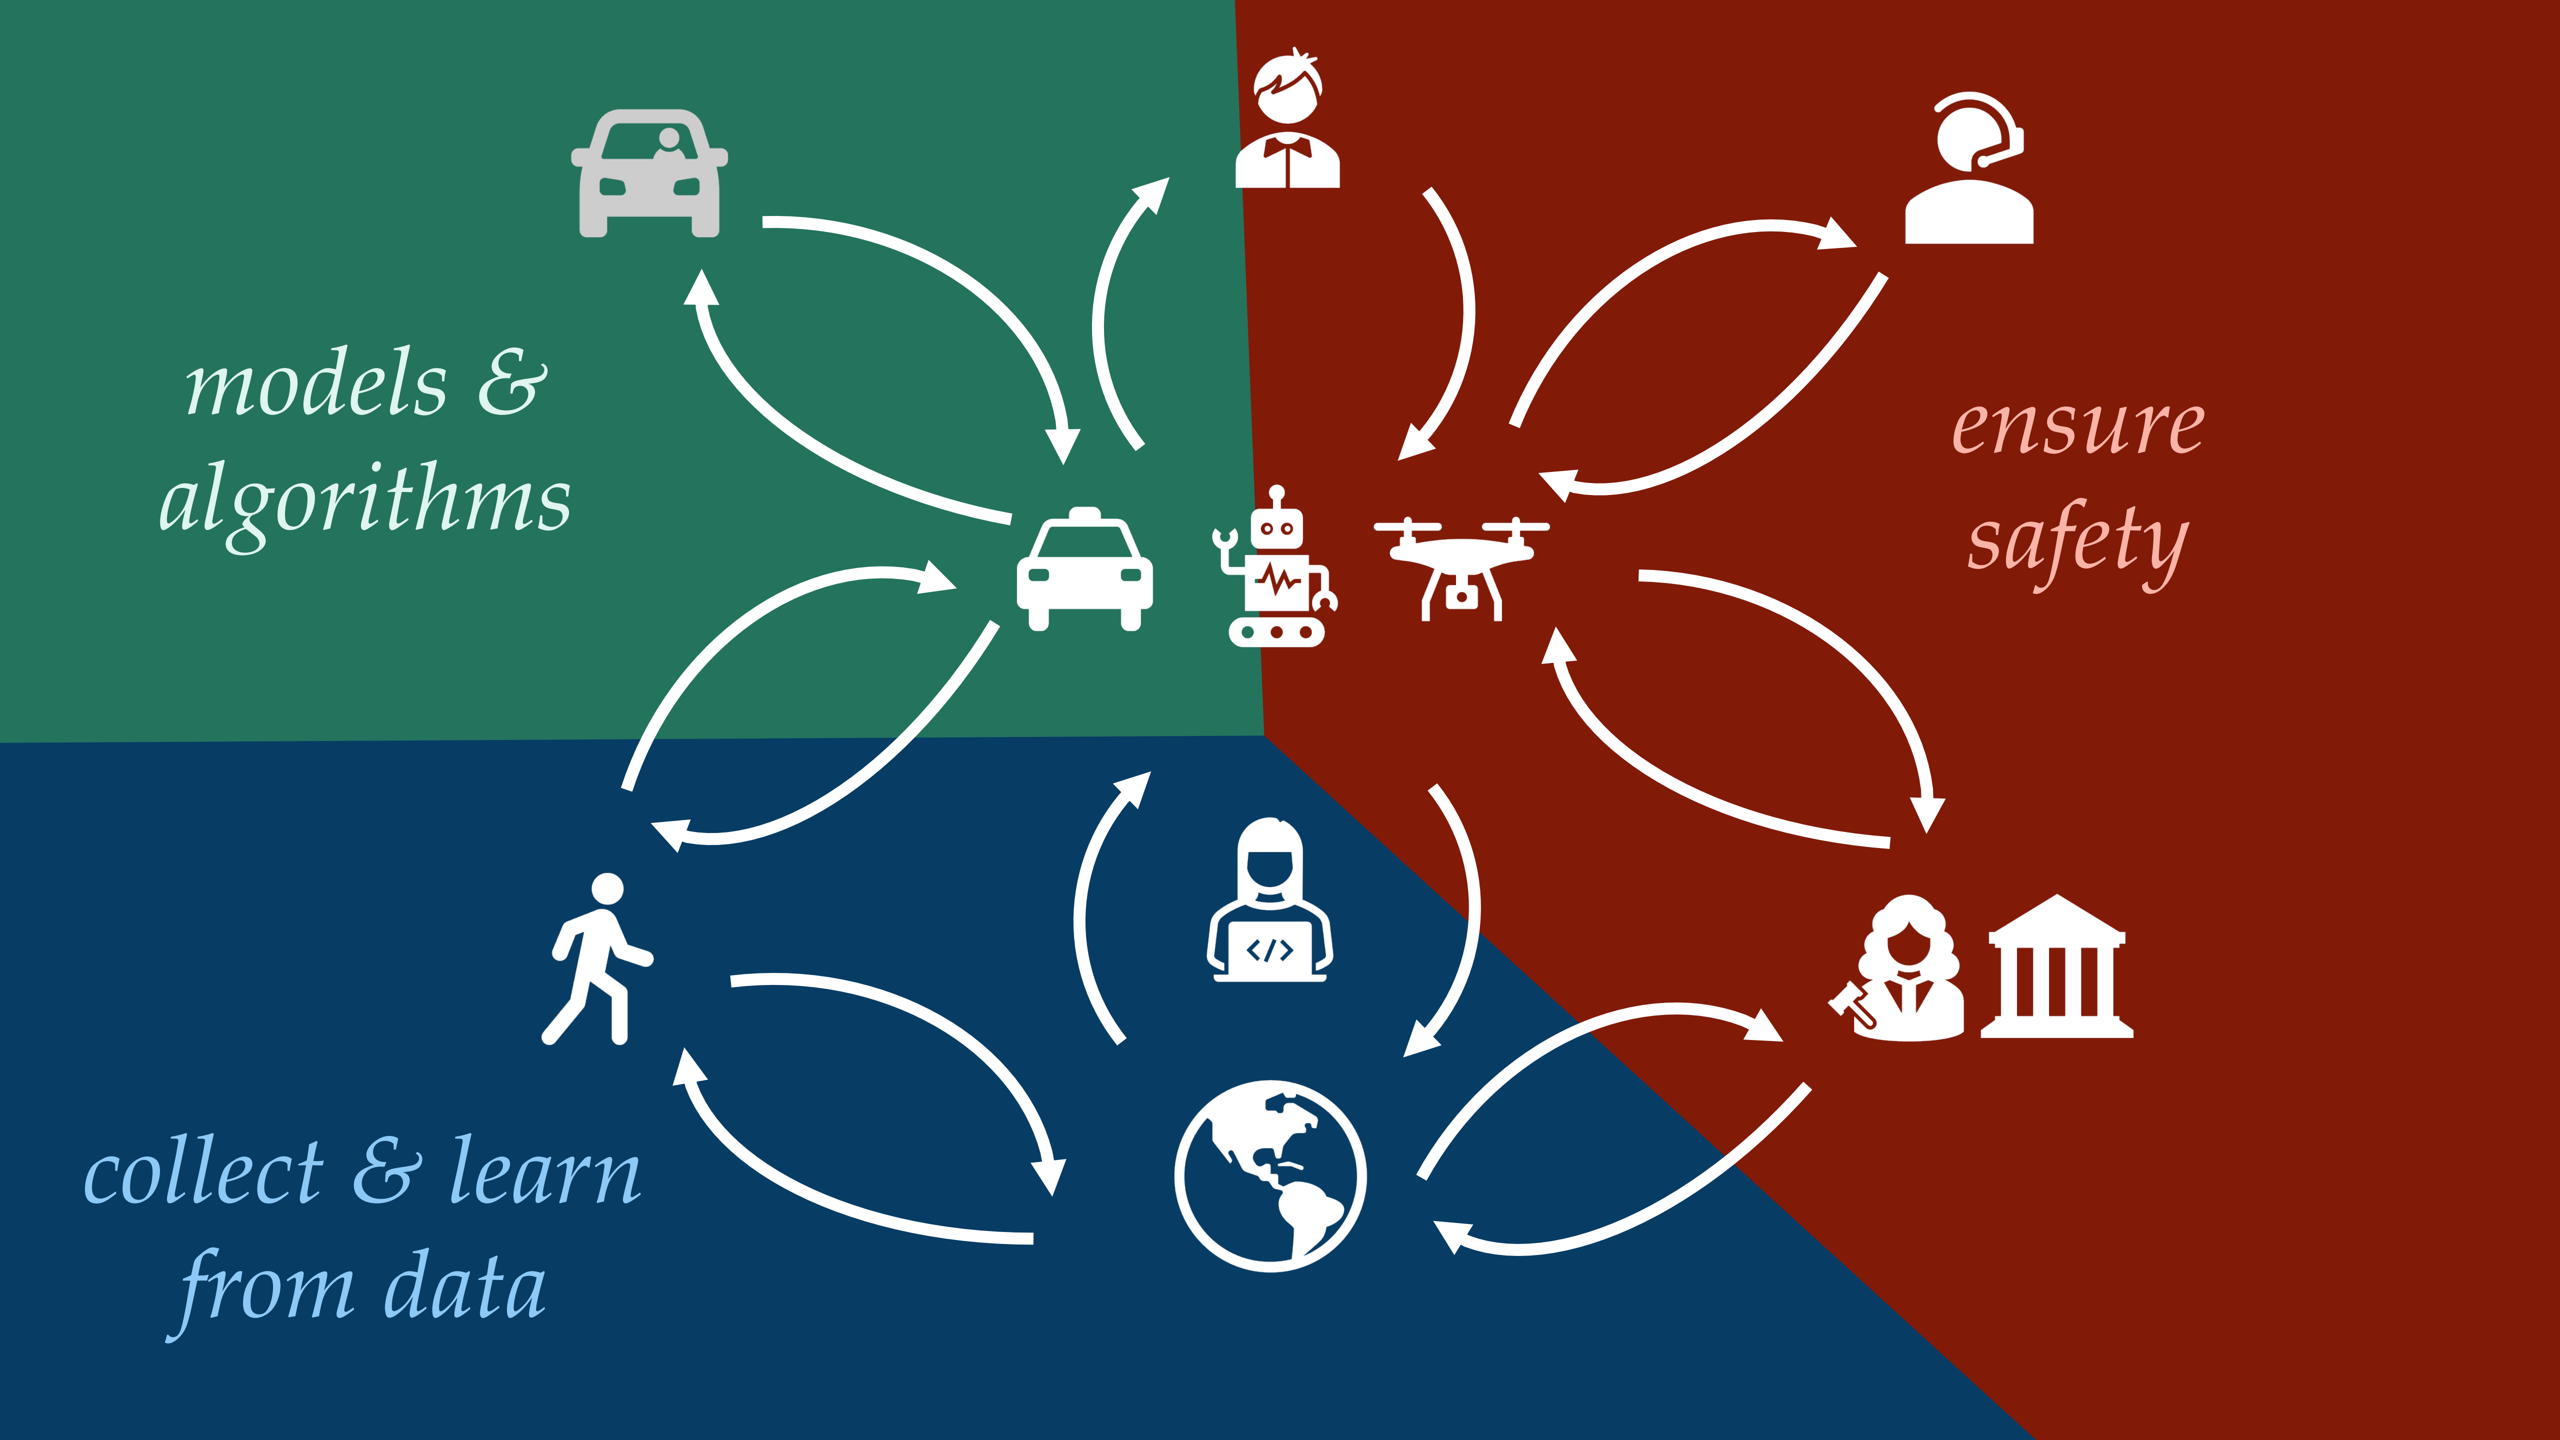

models &
algorithms
ensure
safety
collect & learn
from data

What you will learn in this course

Part I: Safe Interaction

Safety analysis (single & multi-agent)
Scalable computational safety tools
Safety filtering for robot interaction with humans

Part II: Robot Learning From Human Data

Human behavior prediction (game-theoretic & data-driven)
Embedding human models into safety
Sources of human data

Part III: Emerging Research Frontiers

Reliable / robust learning from human data
Alignment and AI safety
Latent-space safety

Guest Lectures

Safety Filtering

Jason Choi
PhD Candidate @ UC Berkeley

Game-theoretic Interaction

Lasse Peters
PhD Candidate @ TU Delft

Data-driven Behavior Prediction

Dr. Boris Ivanovic
Manager @ NVIDIA

Reward Learning in Multi-Agent Games

David Fridovich-Keil
Prof @ UT Austin

"Imitation Learning: It's Only a Game!"

Sanjiban Choudhury
Prof @ Cornell

Safety via ML Robustness

Aditi Raghunathan
Prof @ CMU

Survey (5 min)

https://forms.gle/CjsyUS2nDRD4PiVE8

16-886 Special Topics

Models & Algorithms for Interactive Robotics

Instructor: Prof. Andrea Bajcsy

Carnegie Mellon University

intent ROBOTICS LAB