Autonomous Vehicles: A CPS or AI Grand Challenge?

Raj Rajkumar
Carnegie Mellon University
www.ece.cmu.edu/~raj

Keynote at NSF CPS PI Meeting, 2021

1

Motivation for Self-Driving Vehicles

- 1.3 million people die every year in automotive crashes
  - >94% of these crashes due to human error
- Traffic delays are very expensive
- The elderly, legally blind and physically-disabled people depend on others for transportation

2

Grand Challenges for Cyber-Physical Systems

June 25, 2008

Possible Grand Challenges

- Zero automotive traffic fatalities, injuries minimized, and significantly reduced traffic congestion and delays
- Blackout-free electricity generation and distribution
- Reduce testing and integration time and costs of complex CPS systems (e.g. avionics) by one to two orders of magnitude
- Perpetual life assistants for busy, older or disabled people
- Extreme-yield agriculture
- Energy-aware buildings
- Location-independent access to world-class medicine
- Physical critical infrastructure that calls for preventive maintenance
- Self-correcting and self-certifying cyber-physical systems for "one-off" applications

3

CPS Overview from NSF

Dec 8, 2008

Dr Helen Gill

Dr. Kishan Baheti
(In Memoriam)

CPS: Example at Multiple Scales

A BMW is "now actually a network of computers"
[R. Achatz, Seimens, *The Economist*, Oct. 11, 2007]

Autonomous Cars

Credit: PaulStamatiou.com

Smart Infrastructure

Credit: MO Dept. of Transportion.

Credit: Dash Navigation, Inc.

Cars as nodes in a network

Lampson's Grand Challenge:
Reduce traffic deaths to zero
[B. Lampson, Getting Computers to Understand, Microsoft, *J. ACM*, 50:1, pp. 70-72, Jan., 2003]

8

4

CMU's Reference Architecture for Autonomous Vehicles

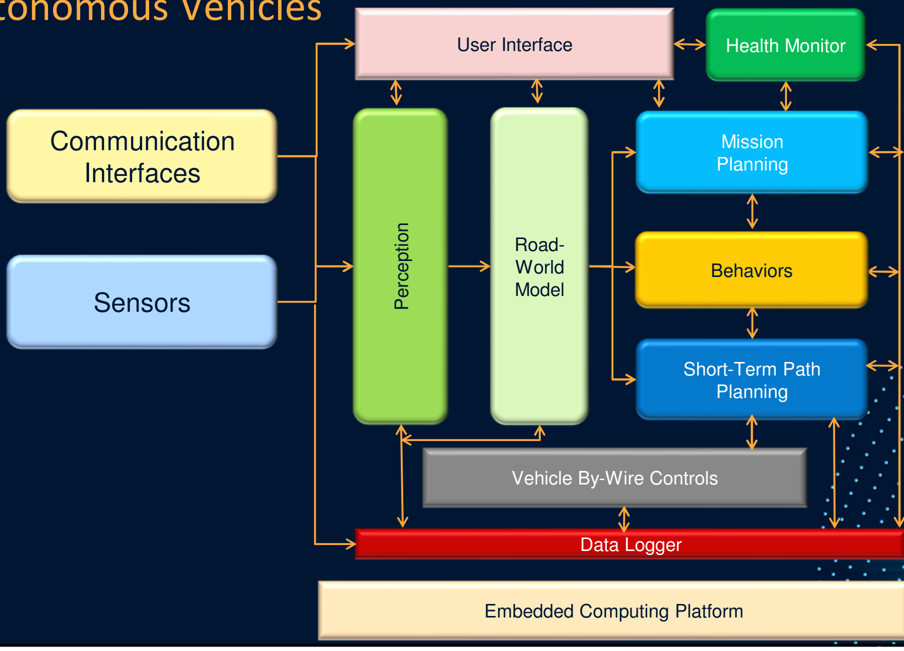

5

AV Challenges 2013

Challenges

- ? Exogenous: The complexities & uncertainties of the real world
  - Weather, lighting, and road conditions; construction; accidents; obsolete information, loss of GPS.
- ✓ Endogenous: Online and safe recovery from failures of sensors, actuators, computing or communications.
  - Mis-calibration, wear and tear, failures.
- ? Verification: How to verify and validate correctness?
- ✓ Interactions: Vehicular Networks
- ✓ Reliability: cost and maintenance, customer acceptance
- ✓ Cyber-Physical Security: thwart connectivity portal attacks
- Human factors
- ✓ Incremental deployment
- ? Legal and regulatory implications

Carnegie Mellon University

6

Recent AV History

7

The AV Hype

2007: DARPA Urban Challenge "No longer science fiction"

2008: Google self-driving vehicles project and Waymo launched in 2016

2013: Ottomatika Inc gets founded

2015:

- Uber Advanced Tech Center in Pittsburgh
- First cross-country autonomous drive in the US with Ottomatika technology
- Ottomatika gets acquired by Delphi (→ Aptiv → Motional)
- TuSimple founded to build driverless trucks

2015-: Tesla "Basic problems have been solved – we'll have full autonomy in 2 years"

- 2018: "Fully autonomous cross-country drive"
- 2019: Robotaxi Teslas will earn $30K per year for the owners
- 2020: Only regulatory considerations stand in the way

2016: Nuro founded to build driverless delivery vehicles

2017: Lyft "Most Lyft vehicles in 2021 will be fully autonomous, car ownership will go away by 2025."

8

… and Ensuing Disillusionment

2019-: Tesla

- 2021: "Full Service Driving Beta" being sold – actual functionality is driver assist
- Multiple fatalities from "AutoPilot" in US and China

2020: Waymo valuation drops from $175B to $30B

2020: Uber ATG sold to a startup company

2021: Lyft Level-5 unit sold to Toyota

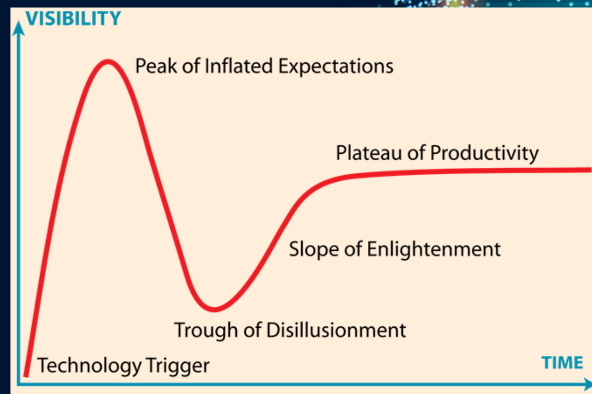

9

What Went Astray?

"An automobile is the most complex consumer device." — *Bill Milam, Ford*

"Driving is the most complex task that adults perform regularly in real-time under dangerous circumstances." — *Raj Rajkumar*

"The $7 trillion promise of self-driving vehicles" — *CNBC*

- Too much money at play
- Too little recognition of the complexity of the problem
  - Imperfect perception
  - Noisy and/or incorrect localization
  - Inaccurate and possibly expensive sensors
  - Very broad range of weather, lighting, road and traffic conditions
  - Significant computational demands
  - Stringent thermal, power and cleaning requirements
- Too much hype

"AI and high-performance hardware will take us to the promised land, pronto."

10

Not Everyone Bought Into The Hype

- "It will take about 10 years, maybe more than 10 years, for a vehicle to be able to drive itself everywhere."

  - Raj Rajkumar, quoted[†] in Voice of America, June 29, 2016
- "The technology just isn't there….There's still a long way to go before we can take the driver away from the driver's seat."

  - Raj Rajkumar, quoted[‡] in the Wall Street Journal, Sept 25, 2016
- "Wouldn't be surprised if it's 2075 before we get [fully autonomous driving]".

  - Steven Schladover, quoted* in Science, December 14, 2017

† https://www.voanews.com/silicon-valley-technology/strides-made-self-driving-cars-much-work-remains

‡ https://www.wsj.com/articles/self-driving-hype-doesnt-reflect-reality-1474821801

* https://www.sciencemag.org/news/2017/12/are-we-going-too-fast-driverless-cars

11

The Non-AI Proposition for AVs

- CPS technologies are sufficient to construct AVs
  - Using a host of sensors, a multi-computer & networked platform, V2X technologies, drive-by-wire actuators and an intelligent software stack

NOTES
An autonomous 33-nile drive from 8mph to 65mph through 11 V2I traffic lights and along two highways from Cranberry Township, PA to the Pittsburgh International Airport in 2013.

- With NO GPUs and no neural networks

12

The Non-AI Proposition for AVs

- CPS technologies are sufficient to construct AVs
  - Using a host of sensors (e.g. high-res and FMCW lidars), a multi-computer & networked platform, V2X technologies, drive-by-wire actuators and an intelligent software stack
- The infrastructure, vehicles or people communicate using V2X while on or near roadways.
- Possible but unfortunately *not* entirely practical in the foreseeable future.
- With AI support, fewer and less expensive sensors can be used without pervasive infrastructure support.
  - AI is therefore necessary in practice to construct AVs.
  - **AI is integral for AVs** to "see" the world designed for human drivers
  - But end-to-end AI is NOT the answer either.

13

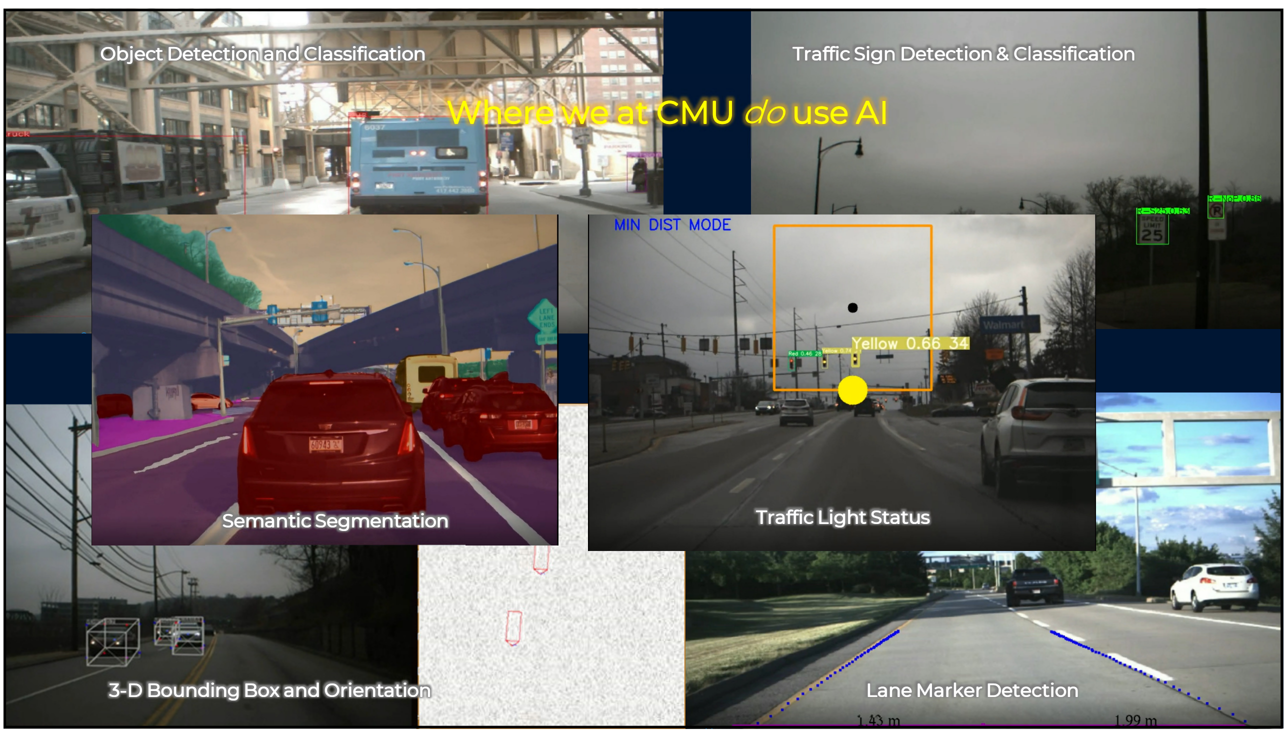

14

Why Not End-To-End AI?

15

The Two Sides of Humans

Super Intelligence

- The Vision system
  - Human eye: very high resolution with depth info, iris/aperture control, focus
- The (Re)cognition system
  - Near-instant detection and identification of objects and organisms
  - Learn from a few examples
- (Unreasonable?) reasoning power
  - Logical reasoning, generalization and specialization
  - Spatial and temporal intelligence
  - Intuition and experience
  - Adaptability (to new vehicles, cultures, rules)
- Coordinated sensing & actuation
  - Interactions: eye contact, social exchanges, game theory

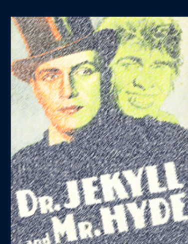

Fundamental Weaknesses

- Distractability
  - Dozing off, conversations, diversion of interest, intoxication
- Imperfect sensing
  - Drivers with poor vision, age impairments,
  - Eyes at the front only
- Non-ideal reasoning
  - Incorrect understanding of vehicle dynamics
  - Poor multitasking
  - Overconfidence
  - Instincts and intuition can be *wrong* at critical moments
- Potentially slow reactions
  - Poor reflexes due to age or health

16

Questions to be Asked of Today's AI

- What does a deepnet really learn?
  - Appeal of deepnets: "no programming"; specify only inputs and outputs
  - Moves the onus of feeding ALL the right inputs to the "programmer"
  - Garbage in → Garbage out
  - Biased inputs → Biased outputs
- What happens when an input not resembling training data is provided to a deepnet?
  - <u>No one</u> knows for sure
- Does a deepnet really learn any generalization or specialization beyond becoming a high-quality pattern matching scheme?
  - With high likelihood, no.

17

The Turing Machine:

Can compute anything that is computable

The link from computability to "human intelligence" however has long been tenuous.

18

A Turing Test for Automated Vehicles

If a self-driving vehicle cannot be distinguished from a human-driven vehicle by humans outside the vehicle, it is said to have passed the *Turing Test for Automated Vehicles*.

19

In Reality, AVs will be *compared to human drivers*

- More conservative much of the time,
- More aggressive some of the time,
- May get stuck occasionally, and
- May take "counter-intuitive" actions that surprise and/or upset humans around the AV

Is this human-like?

Now, let's look at some advances in specific domains…
and their non*-relationship with AI*

20

Counterexample #1: Feedback Control Systems

- Stability, controllability and observability properties
- Ability to deal with disturbances
- Latency implications
- Decidable vs Undecidable problems

Example:
- Inverted Pendulum
  - Yes, one can use machine learning to mimic how a human balances a stick on one's palm – BUT WHY?

21

Counterexample #2: Our Global Computing and Network Fabric

- Human innovations that enable **general-purpose computations**
- Hardware accelerators that perform efficient **special-purpose computations**
- Wired networks and "**network protocol stacks**" that led to the creation of the internet and WWW
- **Wireless networks** that removed the need for a 'leash'

All these are the **products of ingenious human minds** leveraging the power of computing and communications (as serfs).

22

Domain #3: Operating Systems

- Linux, Windows, MacOS, Android, iOS, …
- Brilliant human minds design an architecture comprising of multiple interacting modules
  - "Divide and conquer" for coding, debugging, testing and deployment
  - Yield usability, upgradability and generality.
- AI is very far from being mature enough to design and realize very general and very usable software platforms.

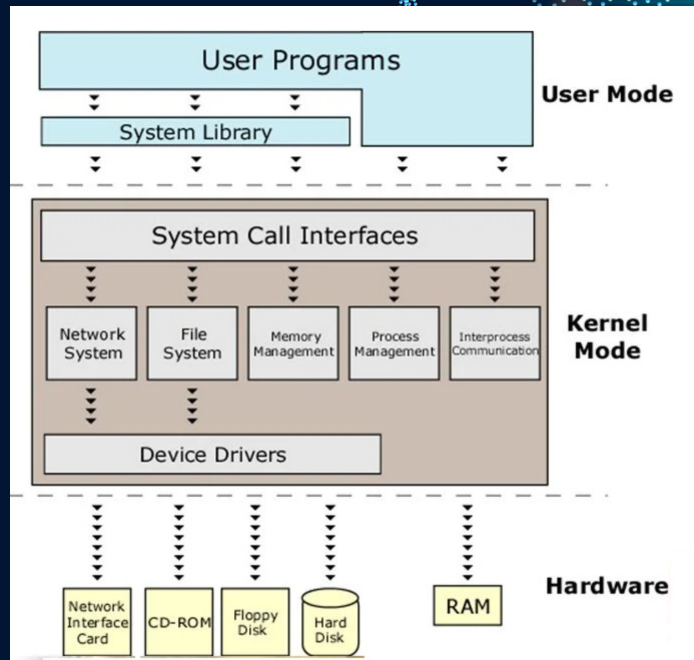

23

Domain #4: Real-Time Systems

- Sound mathematical foundations
  - Worst-case schedulability:
    - Rate-monotonic and earliest-deadline-first scheduling policies for guaranteed timing behavior
  - Provable properties under practical constraints:
    - Synchronization, periodics/aperiodics, mixed criticality, multiprocessors, end-to-end latencies, jitter management, …
  - Has evolved into a science of real-time systems

24

Domain #5: Software Apps

- Complex software applications:
  - Web browsers (https, TCP/IP, GUI, client-server)
  - Web servers (cloud, client-server, cybersecurity)
  - MS Office (easy-to-use WYSIWYG for non-technical people, tight integration, inexpensive hardware)
  - Compilers (parsing, lexical analysis, code optimization, code generation, platform independence)
  - "Zoom" (high-bandwidth low-latency communications, audio/video sync)
- Yes, the zoom "virtual background" uses AI
  - But… it also chops off the shoulder, parts of the head, etc. at times
  - Would you rely on it for a safety-critical system? Should you?

25

Domain #6: Aerospace Applications

- Passenger aircraft
- Fighter aircraft
- Planetary rovers
  - Rocket launch
  - Fuel tank/launch module separation
  - Cruise to destination in the openness of space
  - Orbit around planet
  - Lander separates for journey to planetary surface
  - Rover emerges
  - Baby rovers (or helicopters) get deployed
  - Communicate information to/from Earth

26

AV Challenges for CPS

27

Handling Uncertainties of the Road World Continuously and Safely

- Dealing with the uncertainties of an unconstrained road-world operating environment
- Highways, urban, suburban, ex-urban and rural roads
- Crashes, Road debris, Emergency Vehicles
- School zones, kids, the elderly, strollers, wheelchairs, jaywalkers and absent-minded crossing
- Weather conditions
- Road conditions
- Lighting conditions
- Construction zones
- Wild life

28

The CPS Verification Challenge

Can we prove the following proposition?

The strong version:

1. Given a vehicle with its AV hardware † and software ‡ stack, the vehicle will be **safe under all possible conditions**.

The weak versions:

2. Given a vehicle with its AV hardware and software stack, the vehicle will be **unsafe only under a known finite set of conditions.**
3. Given a vehicle with its AV hardware and software stack, the vehicle will be **safe under a known non-trivial set of conditions**.

†Compute h/w, sensors, actuators and communications interfaces
‡The entire software tack from sensor processing to actuation decisions

29

What Does Being "Safe" Mean?

Strong safety:

Zero crashes, zero injuries and zero fatalities with the AV.

Weak (practical) safety:

The AV is **not "at fault"** when a crash occurs, per accepted legal precedents and insurance practices.

30

Other CPS AV Challenges

- Sensor Fusion
  - minimizing false positives and false negatives
- Generalization and Adaptation
  - Knowledge Encoding
  - Customizability to Different Countries & Driving Cultures
- Behavioral
  - Co-Existence with Other Vehicles and Humans
  - Human Expectation and Comfort
  - Safe behaviors in the presence of faults
  - “Ethical” decision-making
  - Dealing with Pranksters
- Cyber-Physical Security

31

The CPS Toolbox

- Modeling and analysis :
  - The sciences (mathematics, statistics, physics, chemistry)
  - The engineering disciplines (electrical, computer, civil, materials, mechanical, robotics)
  - Data analytics
  - Structured testing
- System architectures
- CPS software engineering
- Verification and Validation
  - Theorem-proving, Model-checking
  - Simulation, HIL, SIL
- …

32

AI or CPS in the Physical World?

The Myth
Modern AI is the be-all, end-all approach for many challenges like self-driving vehicles.

The Reality
In the physical world, CPS techniques and technologies must and will continue to play the over-arching role:

- Explainable
- Analyzable
- Provable
- Fixable
- Understandable

The Practice
Yes, AI can and should be used as a tool in the CPS toolbox
The goal of mimicking the creativity and genius of human intelligence will be a fool's errand for quite some time to come.

Human ingenuity, collective human intelligence, scholarship and **cyber-physical thinking** will enable practical and useful benefits now and in the future.

33

Where Do We Stand vis-à-vis AVs?

- Multi-sensor fusion with AI as *one* key enabler
- Aggregation of well-designed subsystems
- Composability of behaviors
- Safe path planning
- Controllability and observability

34

LEVEL 2 HIGHWAY DRIVING
25 MILES @ 65MPH

Level-4 Driving from 15 mph to 65mph

Level-2 Highway Driving @ 65mph

Notes:
Videos playing in
background during
Q&A

LEVEL 4 DRIVING
FROM 15MPH TO 65MPH

35