Prof. Philip Koopman

Autonomous Vehicle Standards & Open Challenges

October 2022

Carnegie Mellon University

www.Koopman.us

Overview

- Autonomous Vehicle safety standards
  - ISO 26262 & ISO 21448
  - ANSI/UL 4600
  - SAE J3018

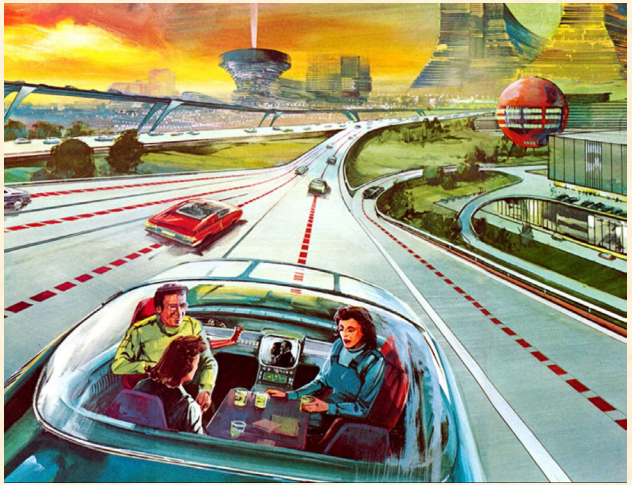

[General Motors]

- The hard bits beyond that are:
  - Fail operational architecture
  - Building an accurate, predictive world model
  - Safety beyond the driving task
  - How safe is safe enough?

© 2022 Philip Koopman

Core AV Design Standards

Carnegie Mellon University

- ISO 26262 – Functional Safety
  - Covers run-time faults & design defects
  - Assume requirements are complete
- ISO 21448 – SOTIF
  - SOTIF: “Safety Of The Intended Function”
  - Iteratively discover & mitigate unknowns
- ANSI/UL 4600: #DidYouThinkofThat?
  - A technically substantive safety argument
  - Evidence of coverage initially + feedback from surprises
  - Aggressive field feedback based on lessons learned

© 2022 Philip Koopman

Standards-Based Engineering Approach

Carnegie Mellon University

SYSTEM SAFETY	ANSI/UL 4600		**Safety Beyond Dynamic Driving**
DYNAMIC DRIVING FUNCTION	ISO 21448	SaFAD/ISO TR 4804	**Environment & Edge Cases**
FUNCTIONAL SAFETY	ISO 26262		**Equipment Faults**
CYBER-SECURITY	SAE J3061	SAE 21434	**Computer Security**
VEHICLE SAFETY	FMVSS	NCAP	**Basic Vehicle Functions**

HIGHLY AUTOMATED VEHICLE SAFETY CASE

ANSI/UL 4600

ROAD TESTING SAFETY

SAE J3018

© 2022 Philip Koopman

AVs Must Fail Operational

Carnegie Mellon University

- "Fail Safe" (fail stop) is not enough
  - Detect failure
  - Switch over to a redundant capability
    - E.g., gracefully terminate mission

https://bit.ly/3VcFzRs

- Safety architecture challenges
  - "Redundancy" is not necessarily enough
  - Safety limited by common mode failures across the redundancy
    - "Diversity" is difficult to measure in all dimensions
  - If two computations disagree, which do you believe?
    - Disagreement is likely for nondeterministic algorithms

© 2022 Philip Koopman

Perception Limits To Safety

Carnegie Mellon University

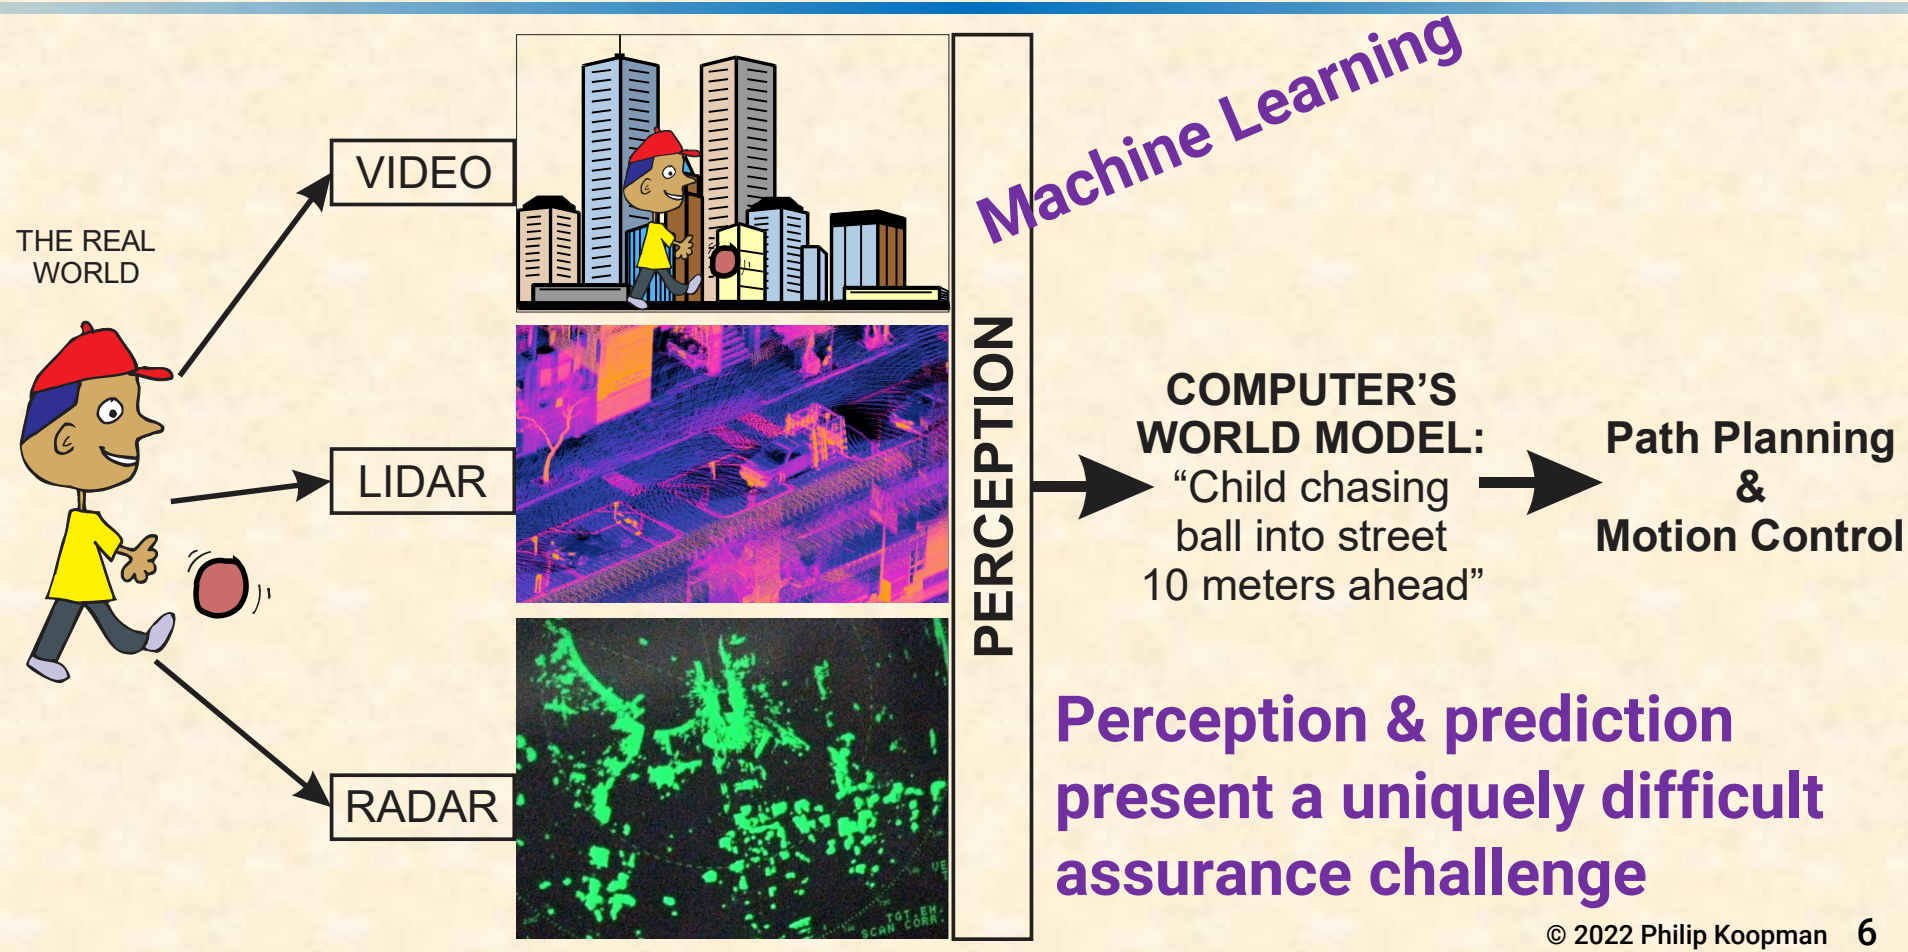

Perception & prediction present a uniquely difficult assurance challenge

© 2022 Philip Koopman

Safety Requires an Accurate World Model

Carnegie Mellon University

- Good prediction based on the world model
  - Classification accuracy affects prediction
  - Probability cloud for object motion
- Safety limited by heavy tail scenarios (rare, important)
  - Probabilities might be context dependent
  - Rare cases tend to dominate safety

?

© 2022 Philip Koopman

Field Engineering Feedback

- Architectures will need to support lifecycle field feedback
  - Safety Performance Indicators (SPI) data linked to safety case
    - Transition from safety recall model to continuous improvement

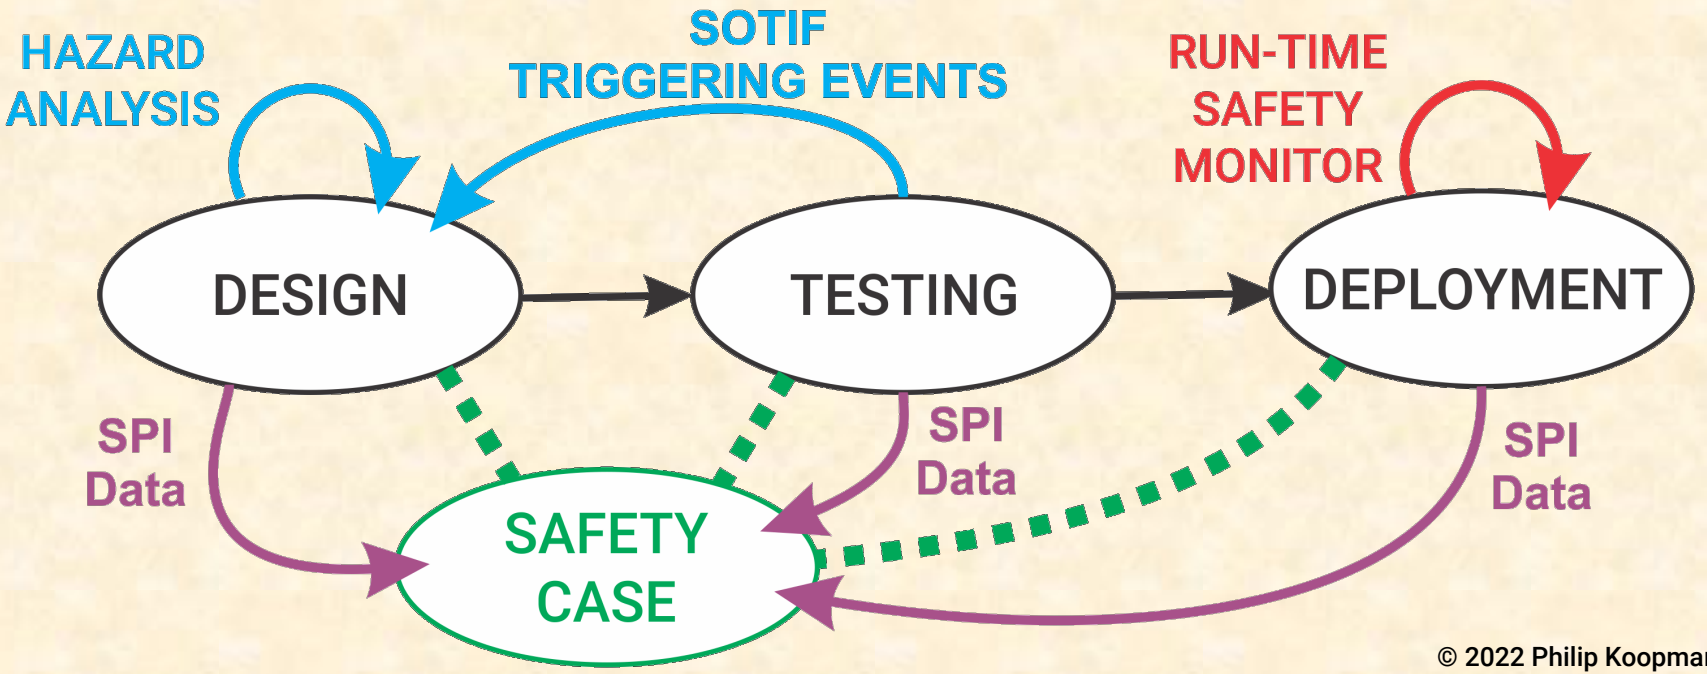

© 2022 Philip Koopman

SPIs and Lifecycle Feedback

Carnegie Mellon University

- SPI: direct measurement of safety case claim failure
  - Independent of reasoning ("claim is X … yet here is ~X")
- A falsified safety case claim:
  - Safety case has some defect
- Root cause analysis might reveal:
  - Product or process defect
  - Invalid safety argument
  - Issue with supporting evidence
  - Assumption error
- Continual Safety case improvement

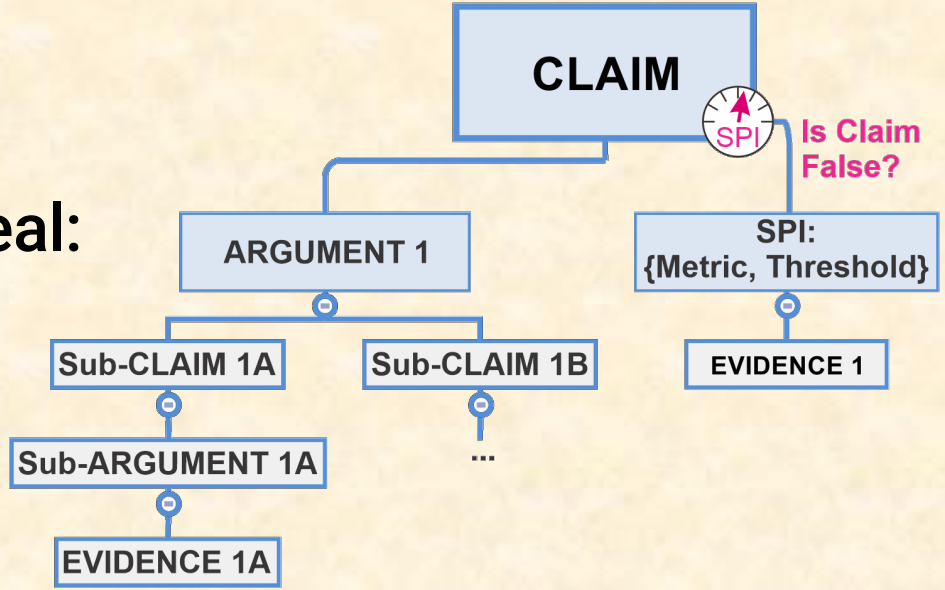

© 2022 Philip Koopman

Role of Humans

Carnegie Mellon University

- There is no “captain of the ship”
  - Autonomy must assume responsibility
- Interacting with people
  - Occupants, cargo loading
  - Pedestrians & mobility device users
  - Potential abuse, misuse
- Role of humans as drivers?
  - Remote operators and wireless data have their limits
  - Avoid “Moral Crumple Zone” operational concept
- Safety culture for all stakeholders

Is it safe to drive now?

© 2022 Philip Koopman

Safe Behavior & Safe Enough

Carnegie Mellon University

- Contextual safety for safe vehicle shutdown
  - Is in-lane stop in fast moving highway “safe”?
  - What if stopped AV blocks an emergency vehicle?
- Where is the “safe enough” bar set?
  - Better than human, but...
    - Prediction uncertainty
    - Equity & risk redistribution issues
  - Safety engineering reduces uncertainty
  - Field feedback of SPIs manages uncertainty
- Governance model: who decides to deploy?
  - What basis is used for decision?

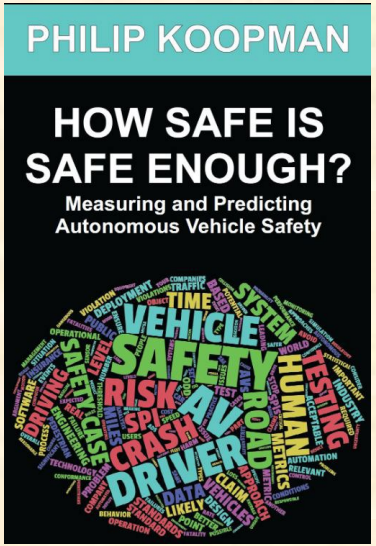

© 2022 Philip Koopman

Summary

Carnegie Mellon University

- Follow safety standards for a foundation
  - Identify & mitigate hazards
    - Within vehicle
    - Presented by operational environment
    - At system level, beyond driving task
  - Safety engineering beyond just road testing
- Be prepared to wrestle with these parts:
  - Fail operational architecture
  - Accuracy of building a world model
  - Safety beyond the driving task
  - How safe is safe enough?

[Shutterstock]

© 2022 Philip Koopman