SOTIF & Edge Cases

https://bit.ly/2SkTpFl

Carnegie Mellon University

Identifying & Mitigating Hazards

- ISO 26262: Hazard and Risk Analysis (HARA)
  - Identify and mitigate risks in accordance with ASIL requirements

- ISO 21448: Identify and mitigate unsafe scenarios
  - Safety of the Intended Function (SOTIF)
  - Reduce “unknown unsafe” area
    - Restrict ODD if needed
  - Deploy at acceptable residual risk

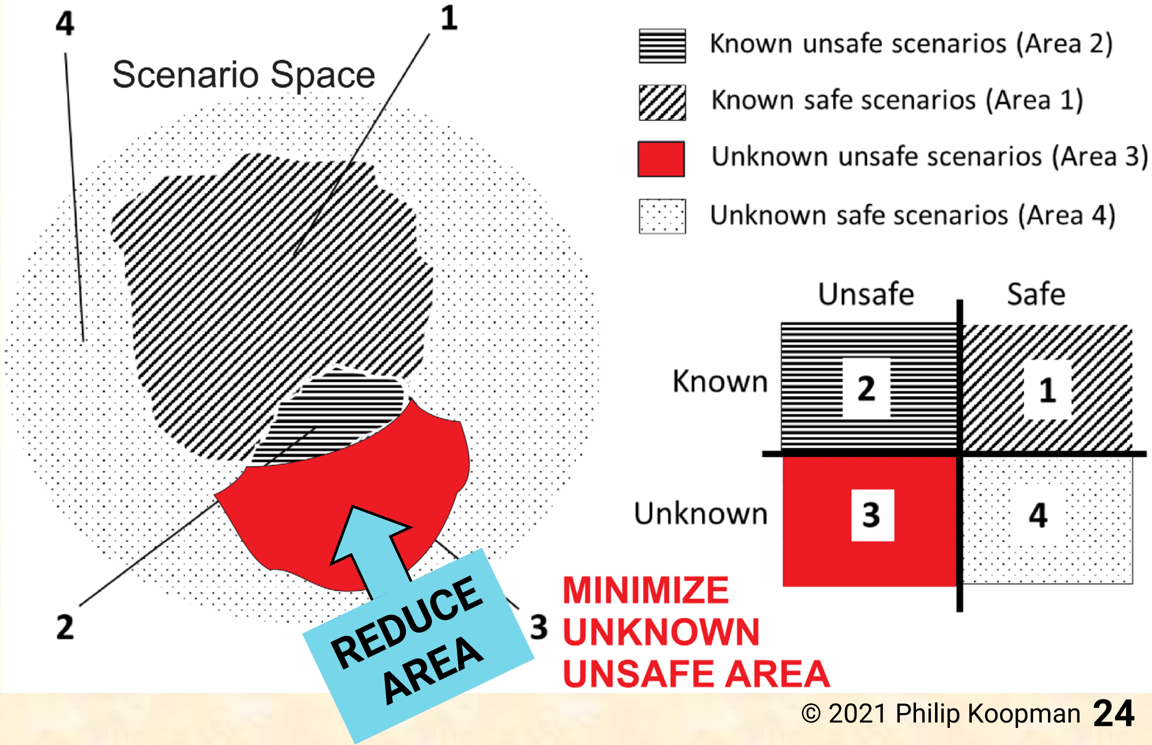

▤ Known unsafe scenarios (Area 2)

▨ Known safe scenarios (Area 1)

🟥 Unknown unsafe scenarios (Area 3)

⬚ Unknown safe scenarios (Area 4)

	Unsafe	Safe
Known	2	1
Unknown	3	4

MINIMIZE UNKNOWN UNSAFE AREA

Six Sigma Isn't Enough for Safety

- Key Performance Indicators (KPIs) help with quality
  - Are all functions working?
  - Is the functionality improving?
  - Is the fault rate decreasing?

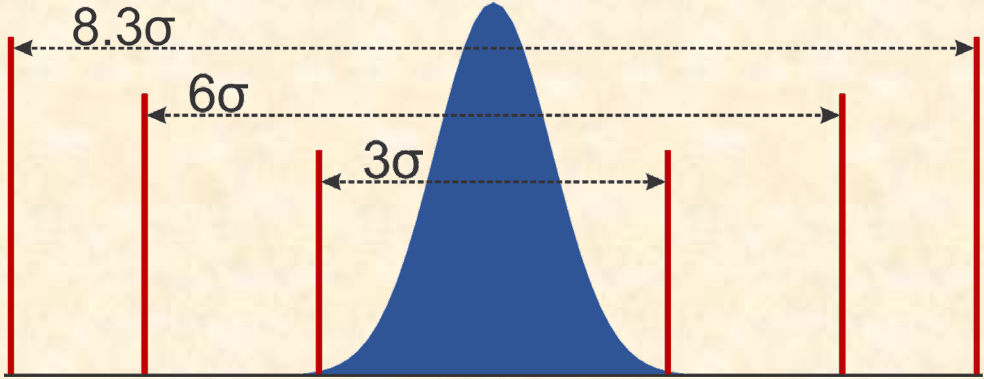

- Good KPIs are just a start
  - Six Sigma Quality: 99.99966% (five nines)
    - Better, but not enough for life critical functions
  - Fatal Crash Avoidance: 99.99999999996% (eleven nines)
    - Safety is 1 million times more demanding! ➔ 8.34 sigma
      - » (example: 1000 opportunities/mile, 250M miles/fatal crash, 1.5σ shift)

It’s All About The Edge Cases

- Gaps in training data can lead to perception failure
  - Safety needs to know: “Is that a person?”
  - Machine learning provides: “Is that thing like the people already in my training data?”

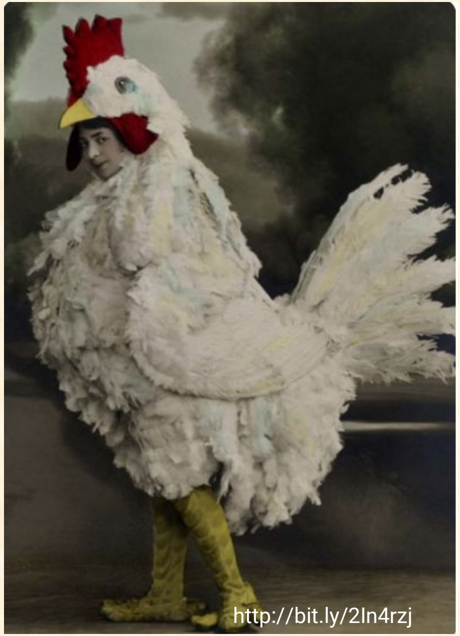

http://bit.ly/2In4rzj

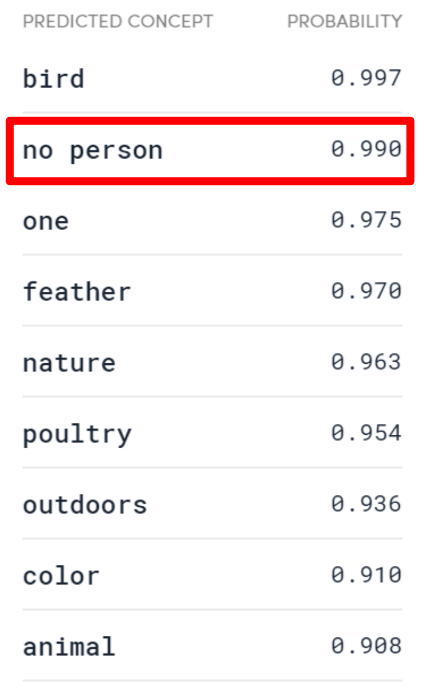

PREDICTED CONCEPT	PROBABILITY
bird	0.997
no person	0.990
one	0.975
feather	0.970
nature	0.963
poultry	0.954
outdoors	0.936
color	0.910
animal	0.908

https://www.clarifai.com/demo

- Edge Case are surprises
  - You won’t see these in training or testing

    ➔ Edge cases are the stuff you didn’t think of!

Why Edge Cases Matter

- Where will you be after 1 Billion miles of drive-fix-drive?

- Assume 1 Million miles between unsafe "surprises"
  - Example #1:
    100 "surprises" @ 100M miles / surprise
    - All surprises seen about 10 times during testing
    - With luck, all bugs are fixed
  - Example #2:
    100,000 "surprises" @ 100B miles / surprise
    - Only 1% of surprises seen during 1B mile testing
    - Bug fixes give no real improvement (1.01M miles / surprise)

https://goo.gl/3dzguf

Real World: Heavy Tail Distribution

Carnegie Mellon University

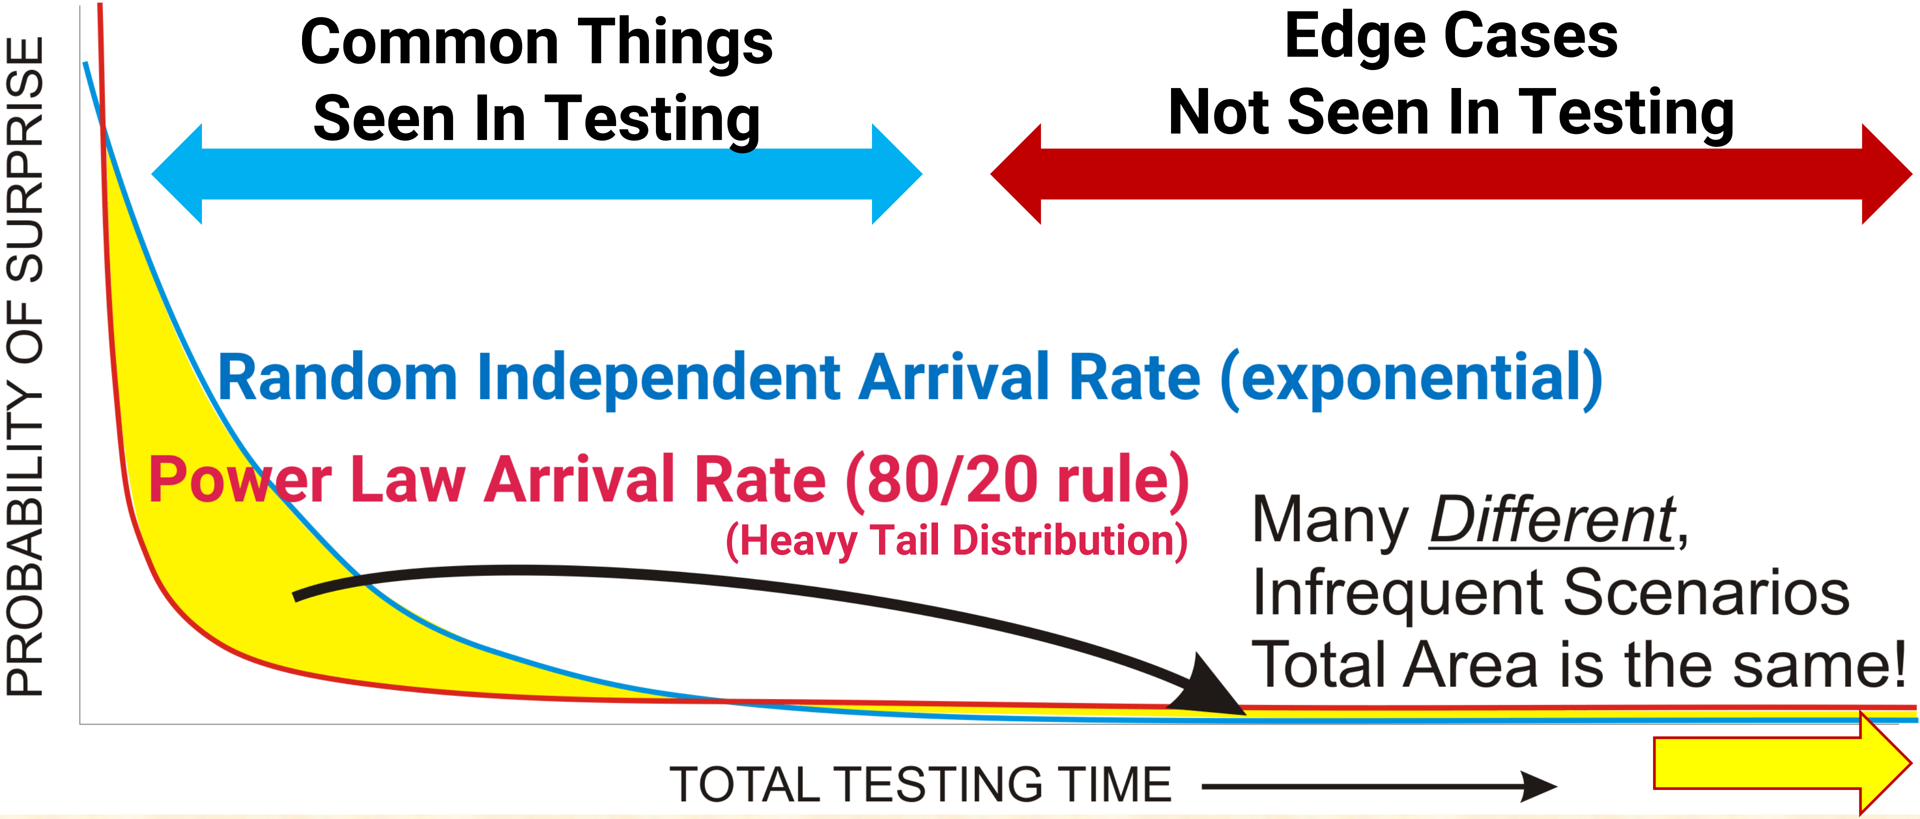

Humans are good at heavy tail

From Driver Assist to Automation

- Driver Assistance (Advanced Driver Assistance System/ADAS)
  - Effective driver monitoring
  - Safety credit if low false positives
    - Every activation can be a life saved
    - Non-activation was driver's fault anyway
- Automated Vehicle (AV)
  - Scenario completeness & coverage
  - Sensor fusion, perception, prediction
  - Blamed for false negatives in heavy tail
    - Every mistake can be a life lost

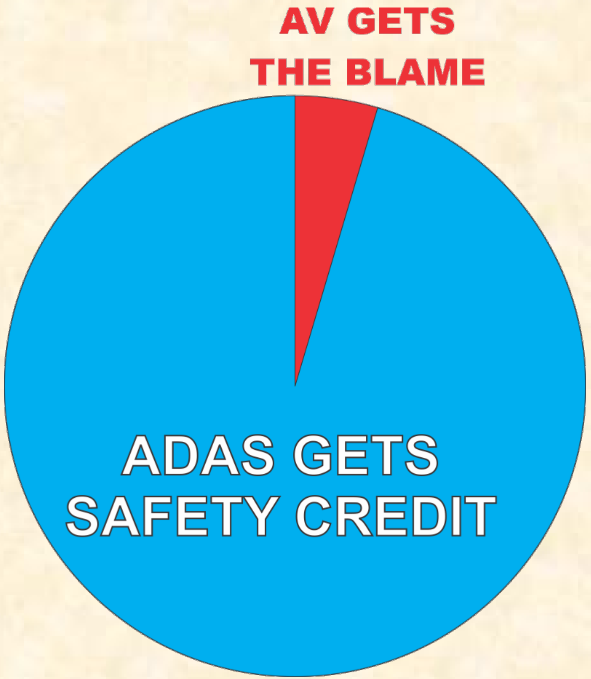

ADS Must Handle Unusual Situations

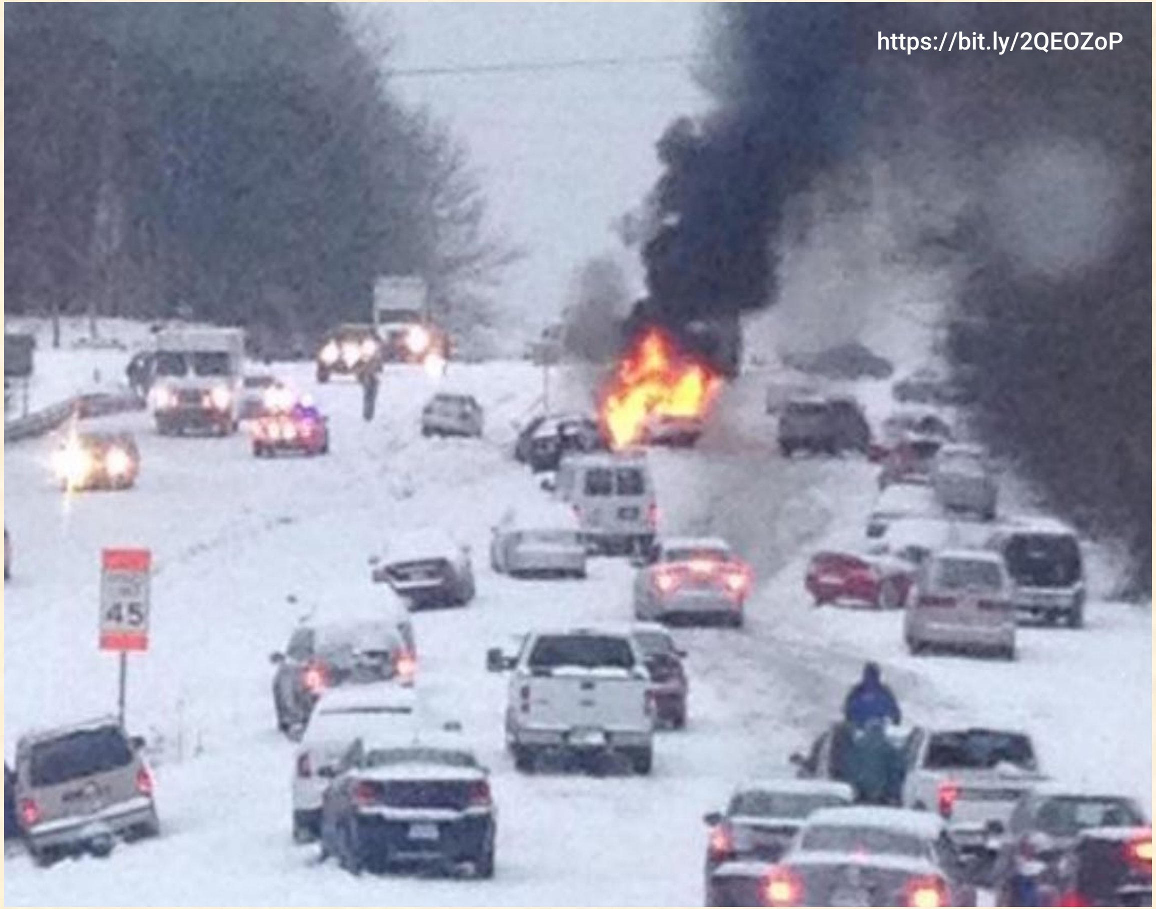
https://bit.ly/2QEOZoP

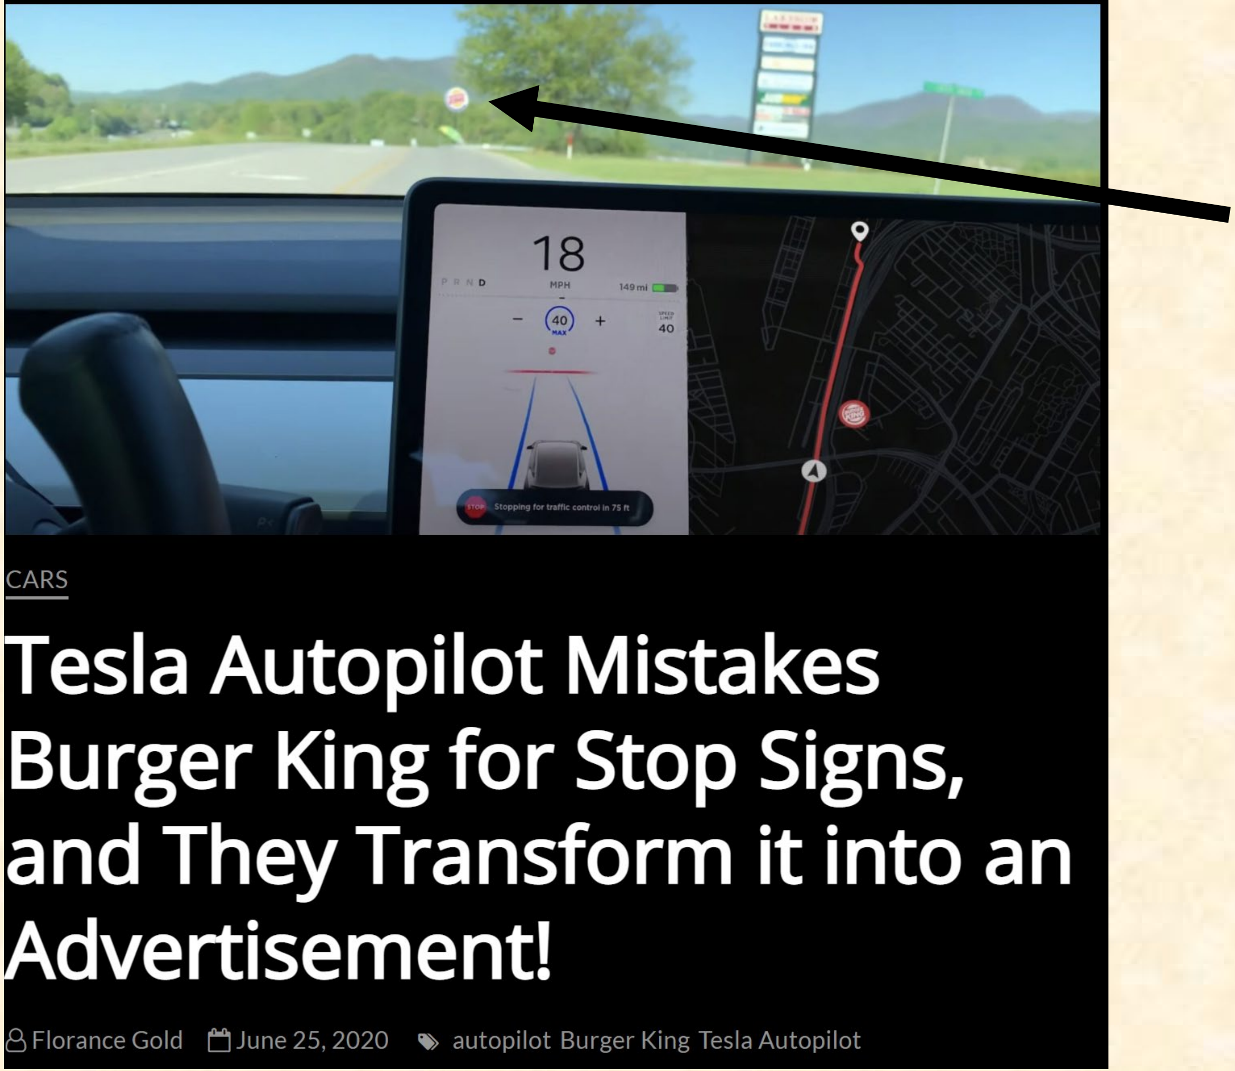

Stopping for traffic control in 75 ft

CARS

Tesla Autopilot Mistakes Burger King for Stop Signs, and They Transform it into an Advertisement!

Florance Gold June 25, 2020 autopilot Burger King Tesla Autopilot

(Burger King owns this trademark.
They have not endorsed this slide)

Human Intuition Isn’t Enough

- Some (perhaps most?) surprises are not obvious to humans
  - Characteristics human test designers think shouldn’t matter
  - Rare events humans know are important but are under-represented
    - High visibility clothing

How good is your ADS at knowing it doesn’t know?

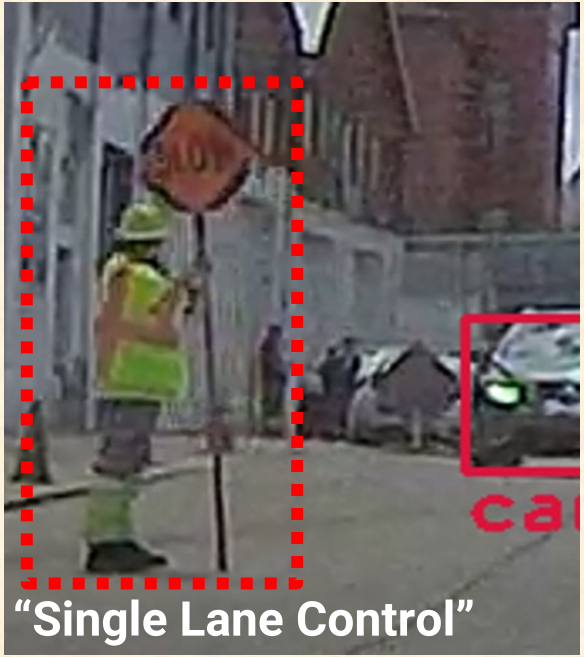

[Edge Case Research] “Single Lane Control”

Changing Relevance of Perception Defects

- Functional safety ➔ SOTIF & system safety
- Heavy tail/edge cases determine safety
- Need to do something safe for unknown unknowns