Regulating Accuracy-Efficiency Trade-Offs in Distributed Machine Learning Systems

A. Feder Cooper, Karen Levy, Christopher De Sa

July 17, 2020 | Law and Machine Learning Workshop, ICML 2020

Cornell University

Setting the scene

Accuracy-efficiency trade-offs

Policymaking for distributed ML systems

A two-fold call to action

Setting the scene

Accuracy-efficiency trade-offs

- Just *one* of many important engineering trade-offs
- Almost always relevant in deployed systems

Policymaking for distributed ML systems

A two-fold call to action

Setting the scene

Accuracy-efficiency trade-offs

- Just *one* of many important engineering trade-offs
- Almost always relevant in deployed systems

Policymaking for distributed ML systems

- Important to understand for policymaking for emerging tech

A two-fold call to action

Setting the scene

Accuracy-efficiency trade-offs

- Just *one* of many important engineering trade-offs
- Almost always relevant in deployed systems

Policymaking for distributed ML systems

- Important to understand for policymaking for emerging tech

A two-fold call to action

- The implications for both computer scientists and regulators

Accuracy and efficiency

Accuracy: How correct a computation is

Efficiency: How fast a computation is performed

There is an *inherent tension* between the two

Intuition by example: The trade-off in other disciplines

Example 1: Law

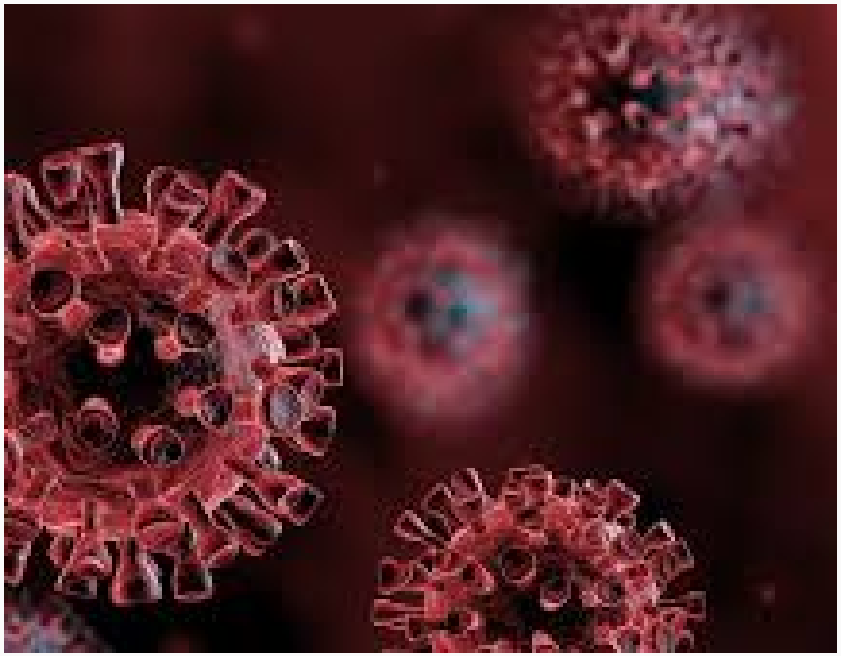

Example 2: Public Health

Intuition by example: The trade-off in other disciplines

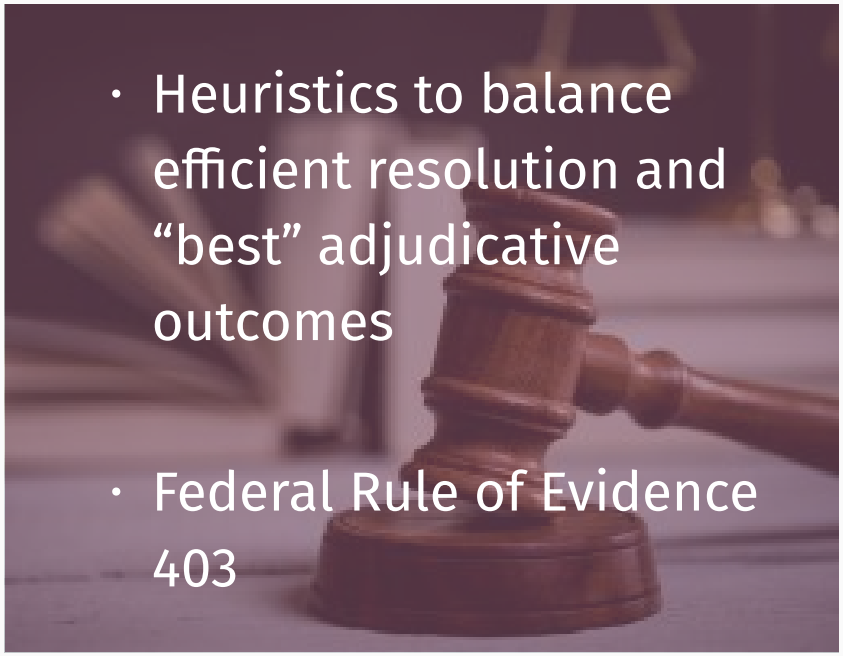

Example 1: Law

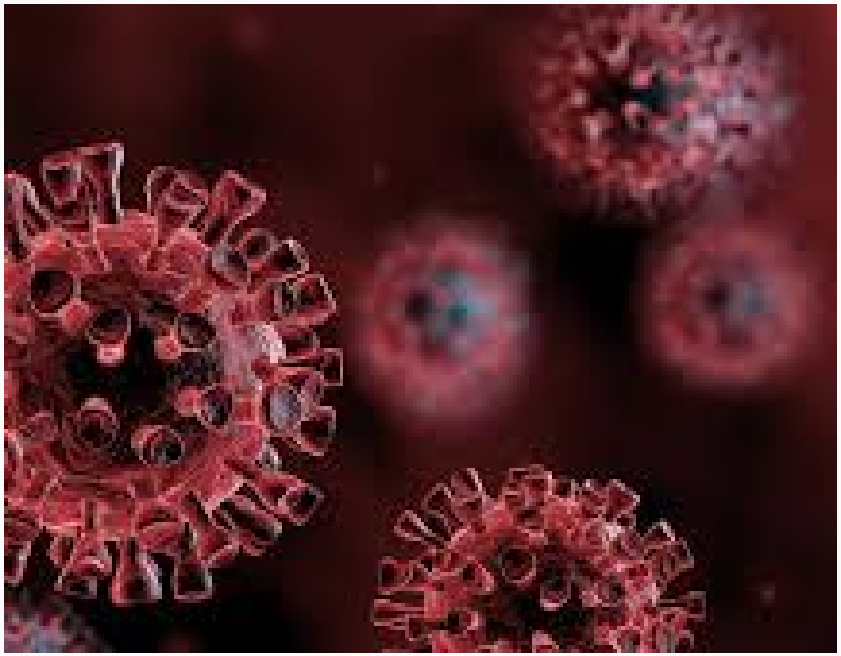

Example 2: Public Health

Intuition by example: The trade-off in other disciplines

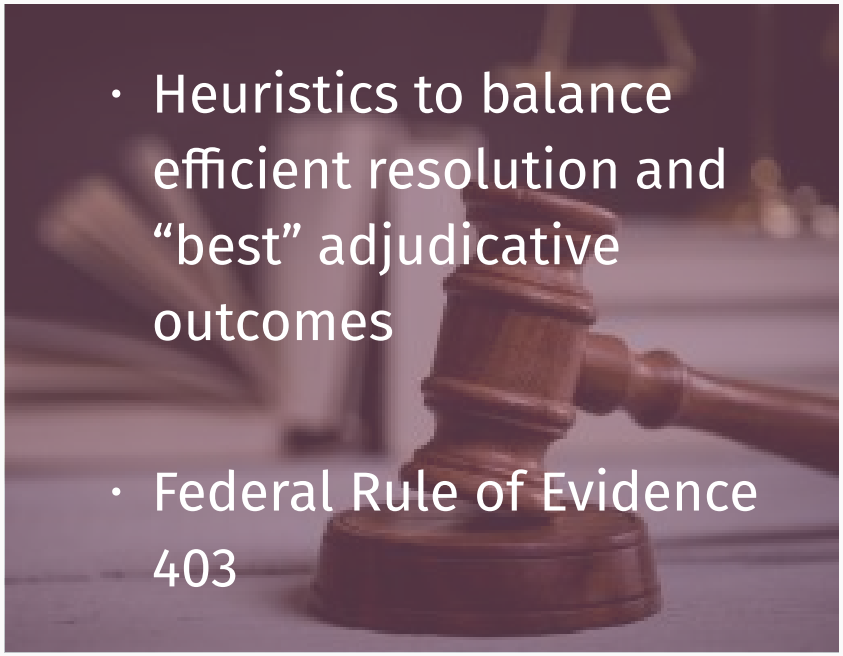

Example 1: Law

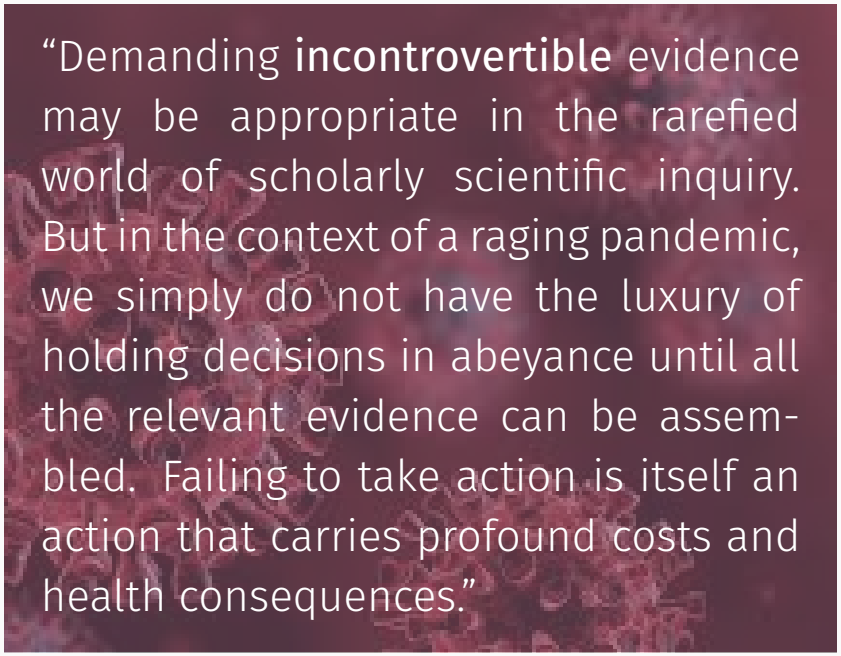

Example 2: Public Health

Accuracy-efficiency trade-offs in computing

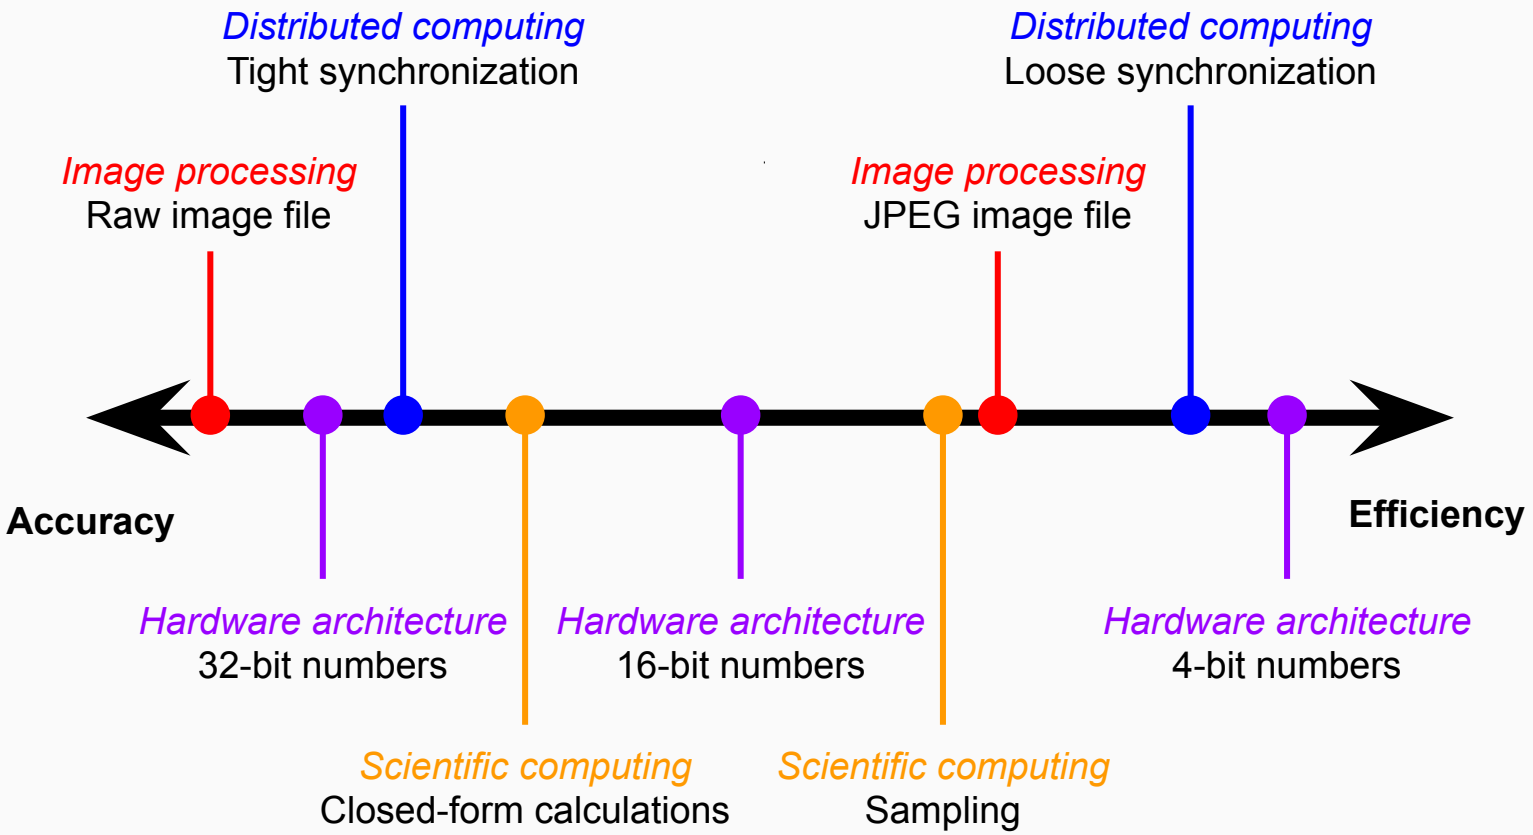

High-stakes computing applications

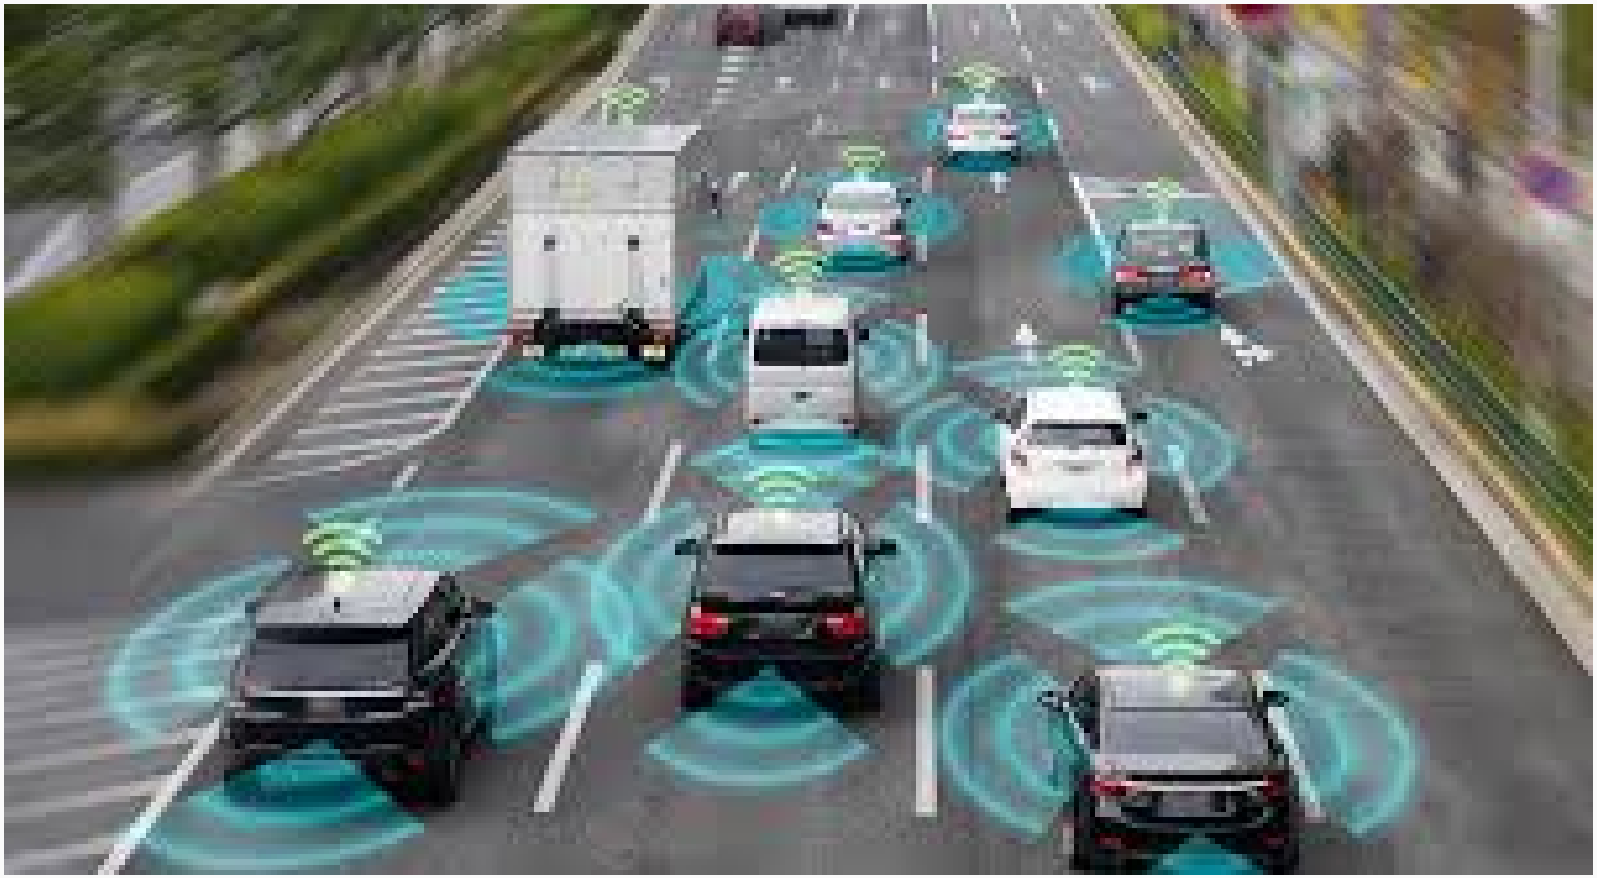

Quick re-cap

Accuracy-efficiency trade-offs are **everywhere** and have **very important**, application-dependent implications

What is a distributed ML system?

Distributed system

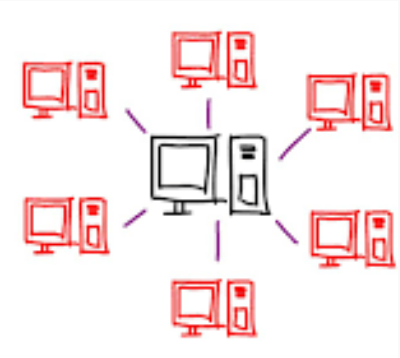

Machine learning

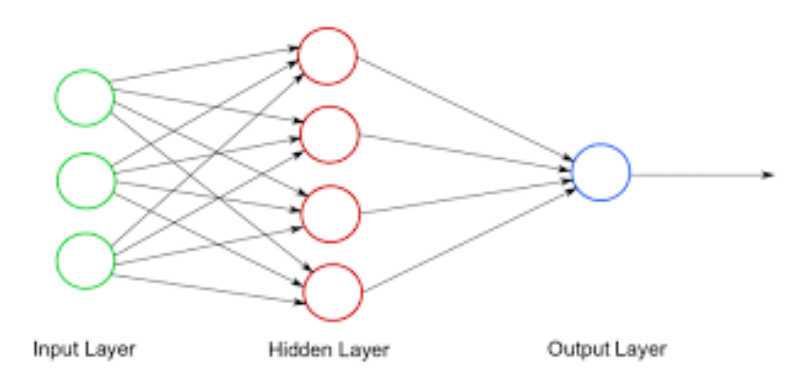

Distributed ML system

Accuracy-efficiency trade-offs in machine learning

- Data subsampling
- Low-precision computation
- Resource-constrained ML
- Markov chain Monte Carlo (MCMC) sampling
- Asynchronous ML

Data subsampling

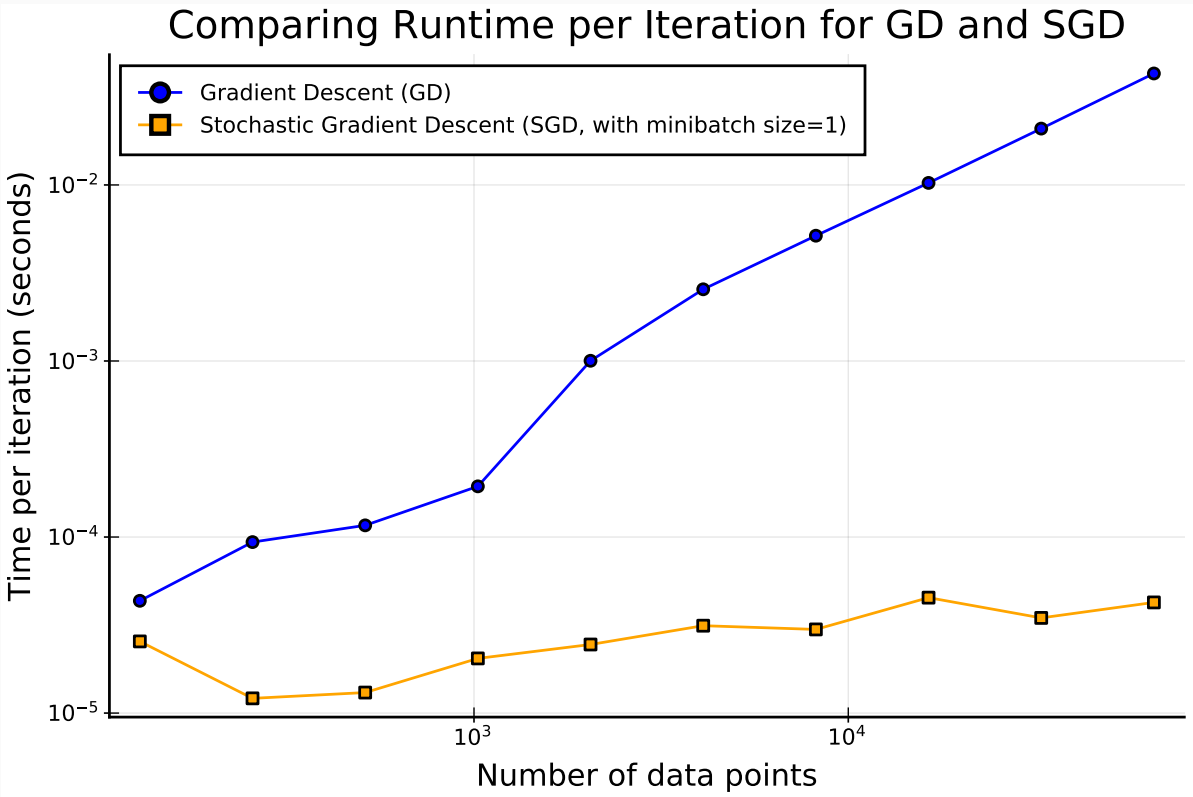

The trade-off in distributed systems

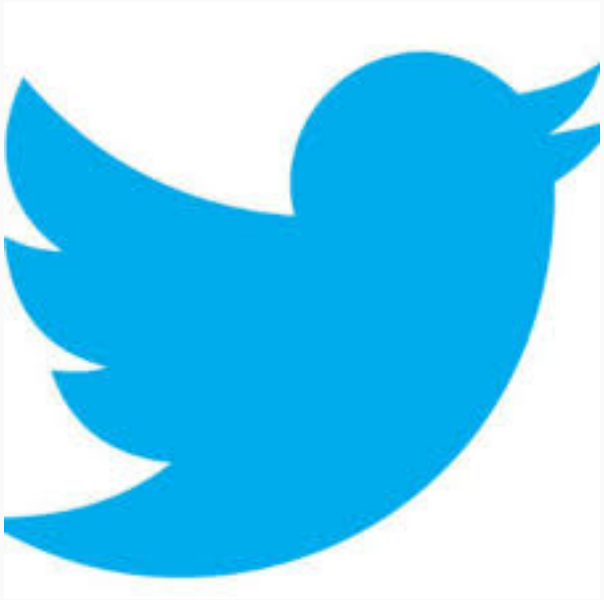

Social media

Blockchain

Distributed ML systems

Collision Between Vehicle Controlled by Developmental
Automated Driving System and Pedestrian
Tempe, Arizona
March 18, 2018

Accident Report

NTSB/HAR-19/03
PB2019-101402

National
Transportation
Safety Board

What’s different here?

Both correctness- and speed-critical

Unclear how much inaccuracy or slowness will be tolerable

Issues with existing systems

- Opt for efficiency
- Do **not** necessarily guarantee correctness
- Are empirically-tuned, rather than theoretically-backed

Our two-fold call to action

Computer scientists

Policymakers

Our two-fold call to action

Computer scientists

- Build distributed ML systems that can be assessed directly in terms of accuracy-efficiency trade-offs

Policymakers

Our two-fold call to action

Computer scientists

- Build distributed ML systems that can be assessed directly in terms of accuracy-efficiency trade-offs

Policymakers

- Understand the trade-off and that it relates directly to low-level engineering decisions
- Use these technical tools to properly assess systems when things go wrong

Concluding thoughts

Computer scientists should build tools that expose decisional **uncertainty** that depends on time

Policymakers should identify normative, domain-specific values at play and frame potential **risks** that follow from this uncertainty

Thank you

A. Feder Cooper (”Cooper”)

Email: afc78@cornell.edu

Website: *https://CacioEPe.pe*