WAYVE

15th Anniversary High Value Manufacturing & 4th New Materials & Graphene Conference 2017

2-3 November 2017

www.cir-strategy.com/events

Alex Kendall

CTO, Wayve Technologies
Research Fellow, Cambridge University

How Can We Build Intelligent Robots?

2017 Robots?

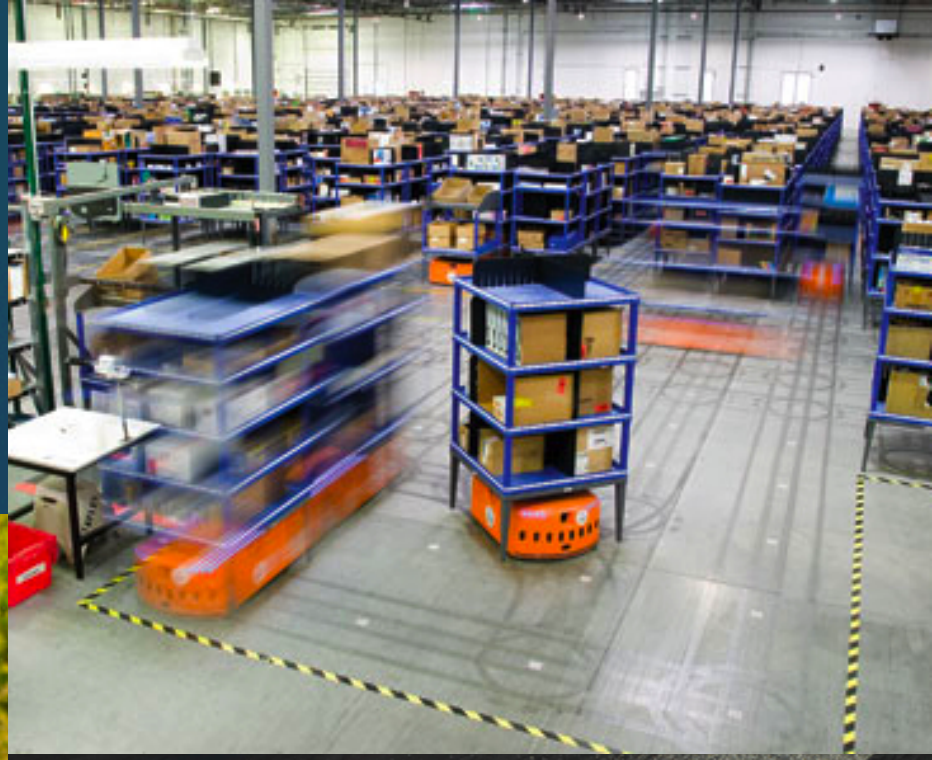


柯洁 KE JIE
00:17:05
ALPHAGO
01:51:38

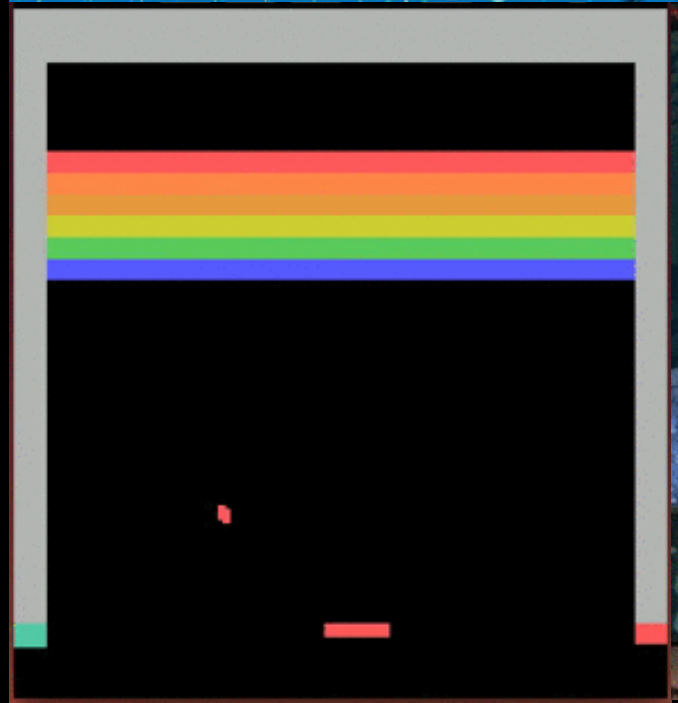

Why are these algorithms not controlling real world robots??

1. Computer Vision
2. Ethics & Safety

Part I: Computer Vision

We aren't born with the ability to see, We need to learn!

- **4 months:** focusing, hand-eye coordination and interest in faces
- **6 months:** depth perception and color vision
- **9 months:** precision grasping and interaction
- **12 months:** object recognition

1. Infant Vision: Birth to 24 Months of Age. https://www.aoa.org/patients-and-public/good-vision-throughout-life/childrens-vision/infant-vision-birth-to-24-months-of-age?

How we learn something as complex as vision?

- Suppose, a baby experiences 1 saccade per second, for 8 hours a day for 365 days
- $1 \times 60 \times 60 \times 8 \times 365 = 10{,}000{,}000$ training examples to learn to see

Deep Learning for Computer Vision

- Powerful framework for understanding high dimensional data like images, videos, speech, text
- With enough training data, they outperform human baselines for recognition tasks
- Typically computer vision models contain ~10 million parameters, take 3+ days to train on a GPU

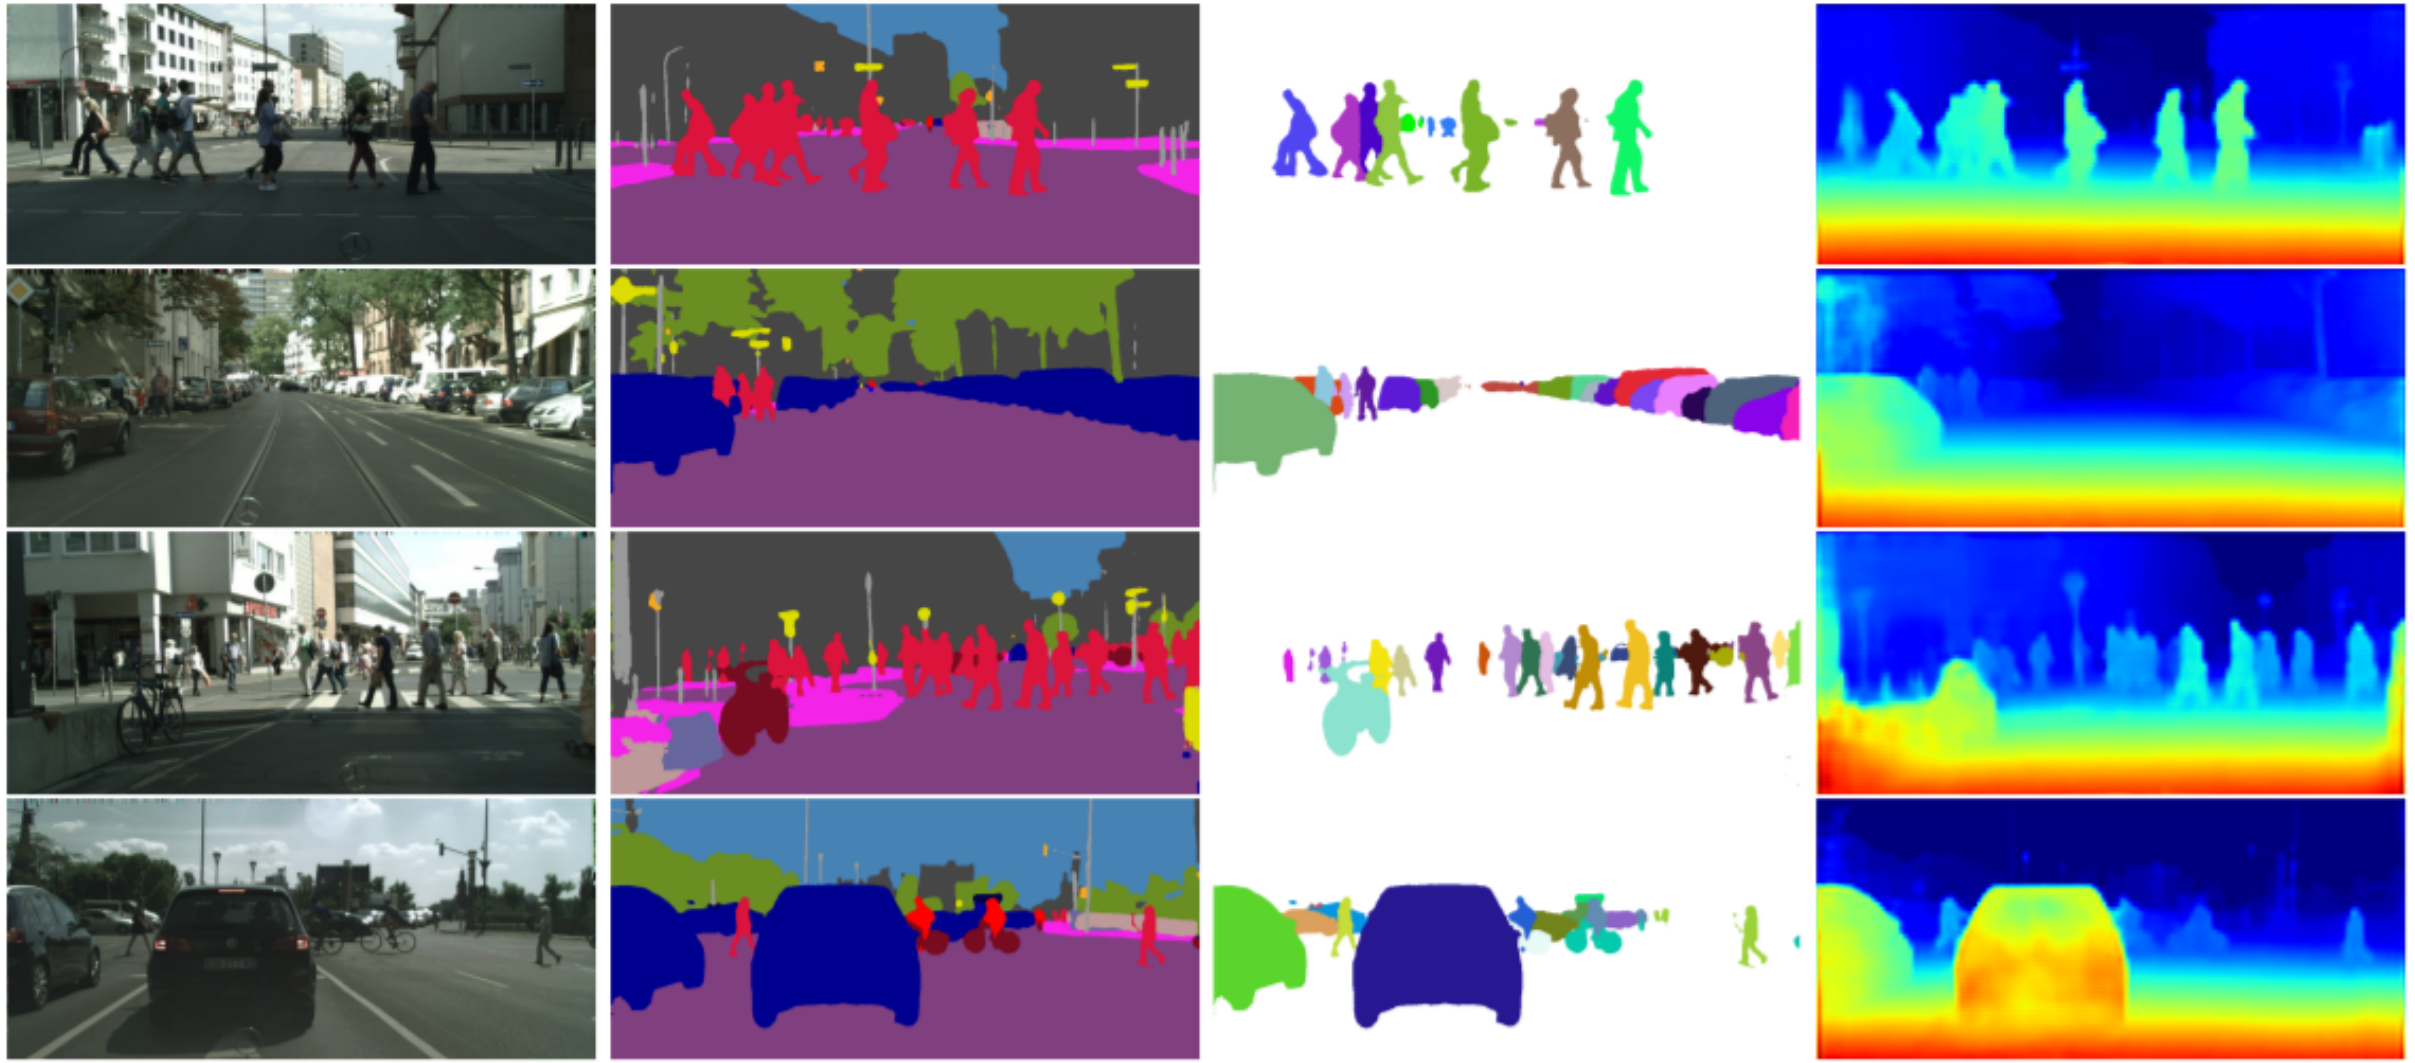

(a) Input image (b) Segmentation output (c) Instance output (d) Depth output

Alex Kendall, Yarin Gal and Roberto Cipolla. **Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics.** *In Submission*, 2017.

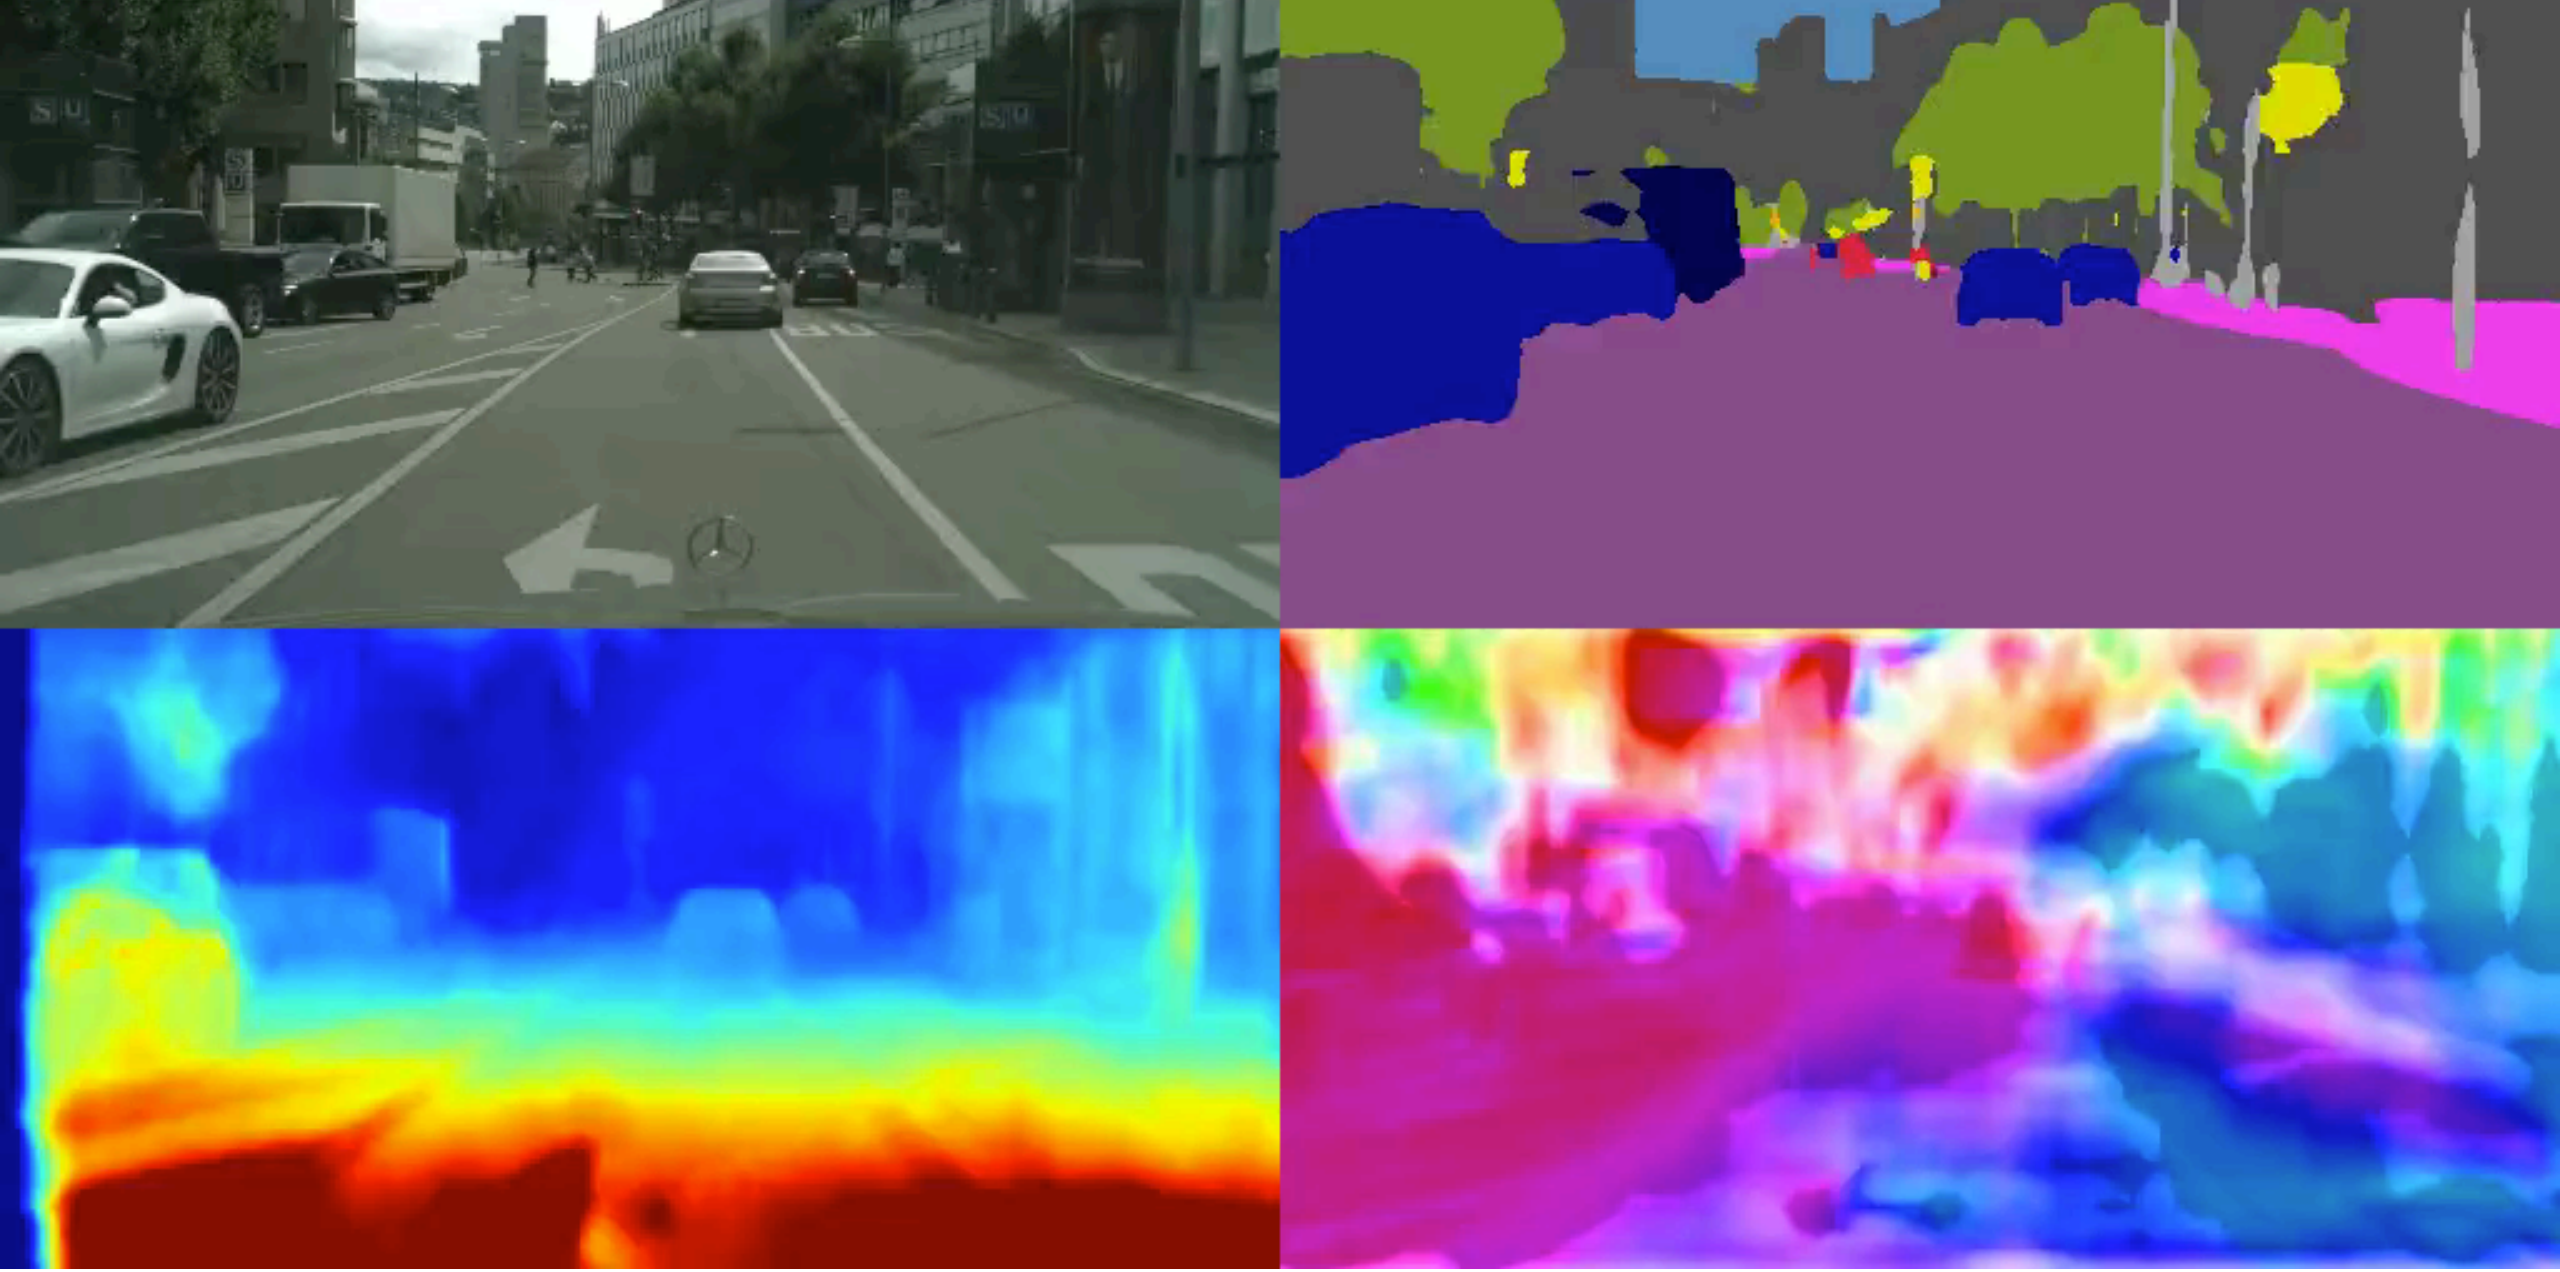

Alex Kendall and Roberto Cipolla. **VideoSegNet: Self-Supervised Motion and Depth for Video Semantic Segmentation.** *In Submission*, 2017.


SKYDIO

Part II: Ethics & Safety

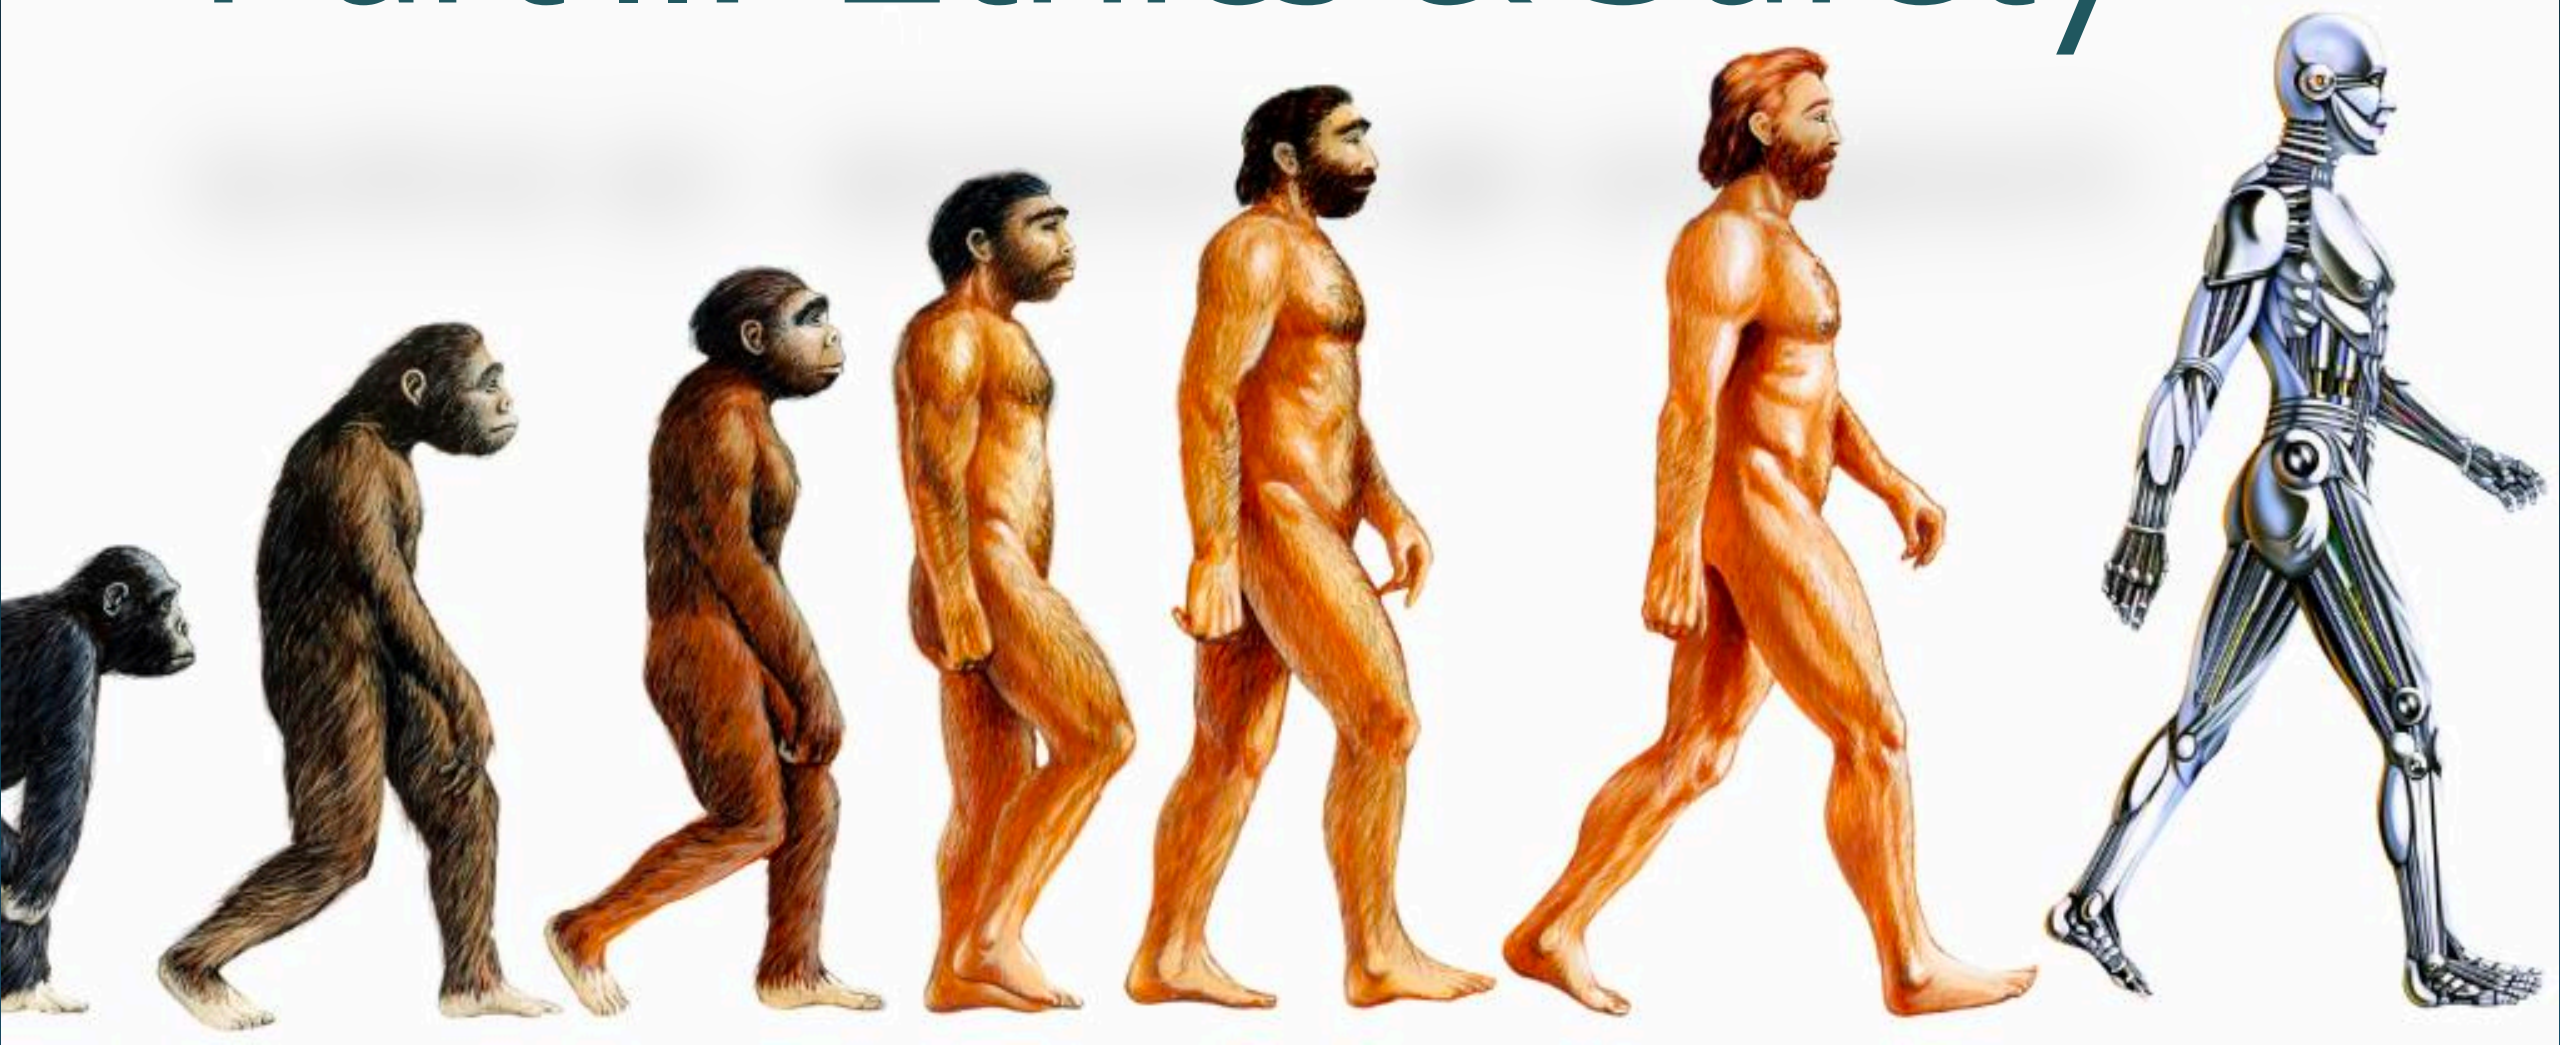

The Trolley Problem

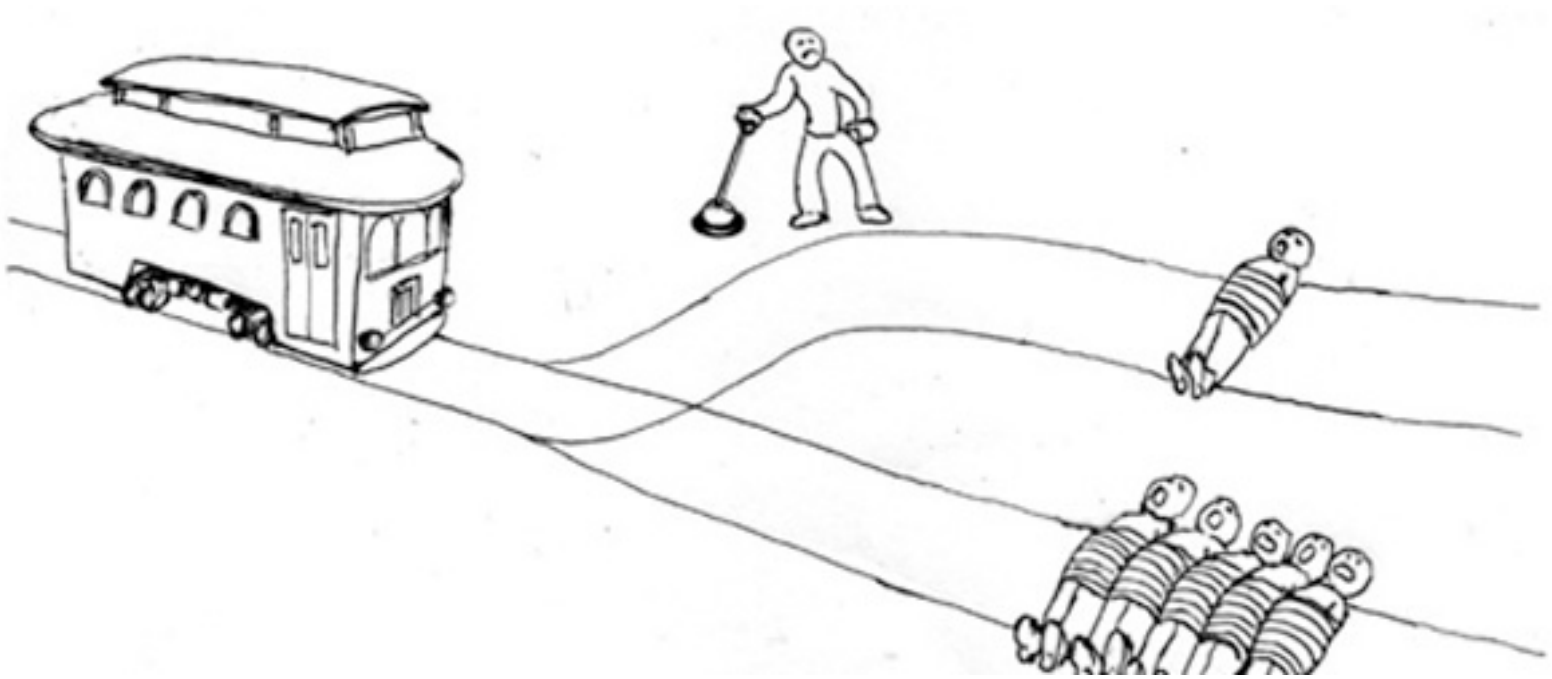

Do you…

1. Do nothing, and the trolley kills the five people on the main track.
2. Pull the lever, diverting the trolley onto the side track where it will kill one person.

Which is the most ethical choice?

1. *The Trolley Problem, Wikipedia.* https://en.wikipedia.org/wiki/Trolley_problem
2. *Rowan McAllister, Yarin Gal, Alex Kendall, Mark van der Wilk, Amar Shah, Roberto Cipolla, and Adrian Weller. Concrete Problems for Autonomous Vehicle Safety: Advantages of Bayesian Deep Learning. IJCAI, 2017.*

Examples of Un-Ethical AI, Today

- US Justice System Re-Offending Rate algorithm.. Biased against minorities
  https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
- Google's speech recognition system is better at male voices
  https://makingnoiseandhearingthings.com/2016/07/12/googles-speech-recognition-has-a-gender-bias/
- Many self-driving cars only work in California
  https://www.ft.com/content/4377b4c0-0479-11e7-aa5b-6bb07f5c8e12

Concrete Problems for AI Safety

- Trust
  - Improve model performance and accuracy
  - Models understanding uncertainty in decisions
- Fairness
  - Improve data efficiency and reduce bias
  - Algorithms which require less training data and generalise better
- Honesty
  - Interpretability of results and model saliency
  - Causal reasoning
  - Avoid reward hacking

---

1. *Rowan McAllister, Yarin Gal, Alex Kendall, Mark van der Wilk, Amar Shah, Roberto Cipolla, and Adrian Weller. Concrete Problems for Autonomous Vehicle Safety: Advantages of Bayesian Deep Learning. IJCAI, 2017.*

Uncertainty in Deep Learning

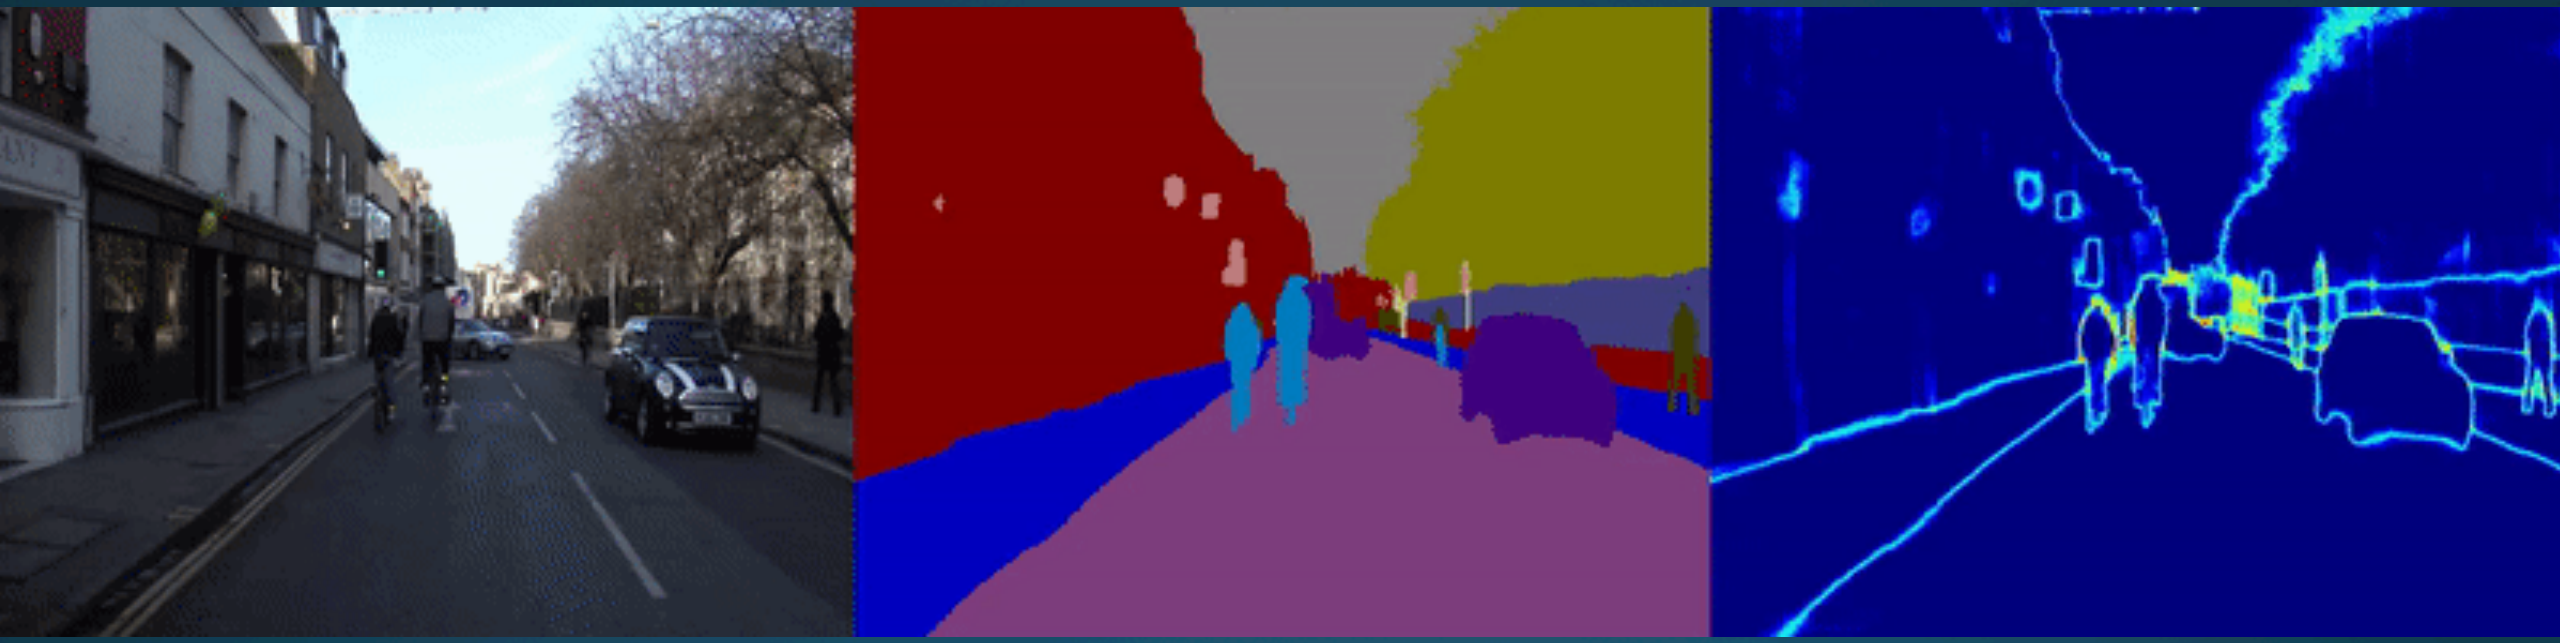

Input Video | Semantic Segmentation | Uncertainty

1. *Alex Kendall and Yarin Gal. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? Advances in Neural Information Processing Systems (NIPS), 2017.*

Conclusions

- We **need machine learning to scale** to hard problems with intelligent robots – hand engineering isn't good enough
- We can **learning to perceive and act from data** with deep learning
- We cannot explicitly reason about ethical situations on a case-by-case basis – our **models need to understand ethics** themselves
- **Computer vision is holding back robotics** from real-world applications

Questions?


alexgkendall.com/publications/


@alexgkendall


agk34@cam.ac.uk

Thank you to the amazing people who made this work possible:

Roberto Cipolla, Vijay Badrinarayanan, Yarin Gal, Tom Roddick, Matthew Grimes, Adrian Weller, Amar Shah, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abe Bachrach, Adam Bry