Slides: https://dhgo.to/coe-cars

AUTONOMOUS CARS: TECHNIQUES AND CHALLENGES

Prof. Dr. Dominik Herrmann // University of Bamberg (Germany)

“Autonomous car” – a vehicle that drives without human intervention

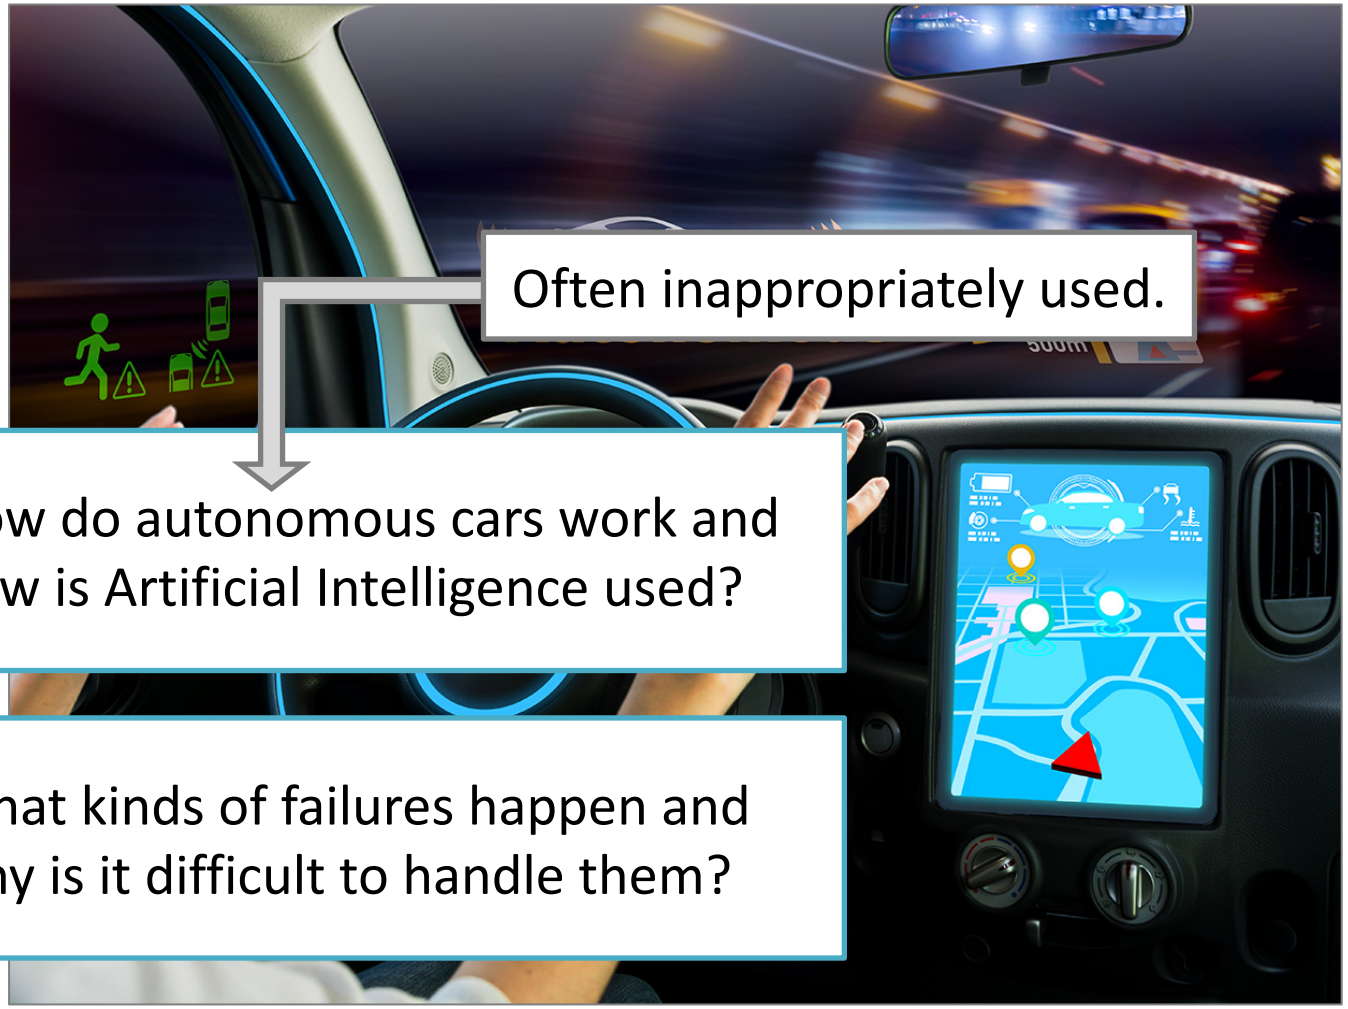

https://spectrum.ieee.org/image/MzAxNTAzMQ.jpeg

What kinds of failures?

Safety

dependability

Security

no malicious interference

Self-driving vehicles consist of two systems.

perception steering

System 1 perceives the environment with various sensors.

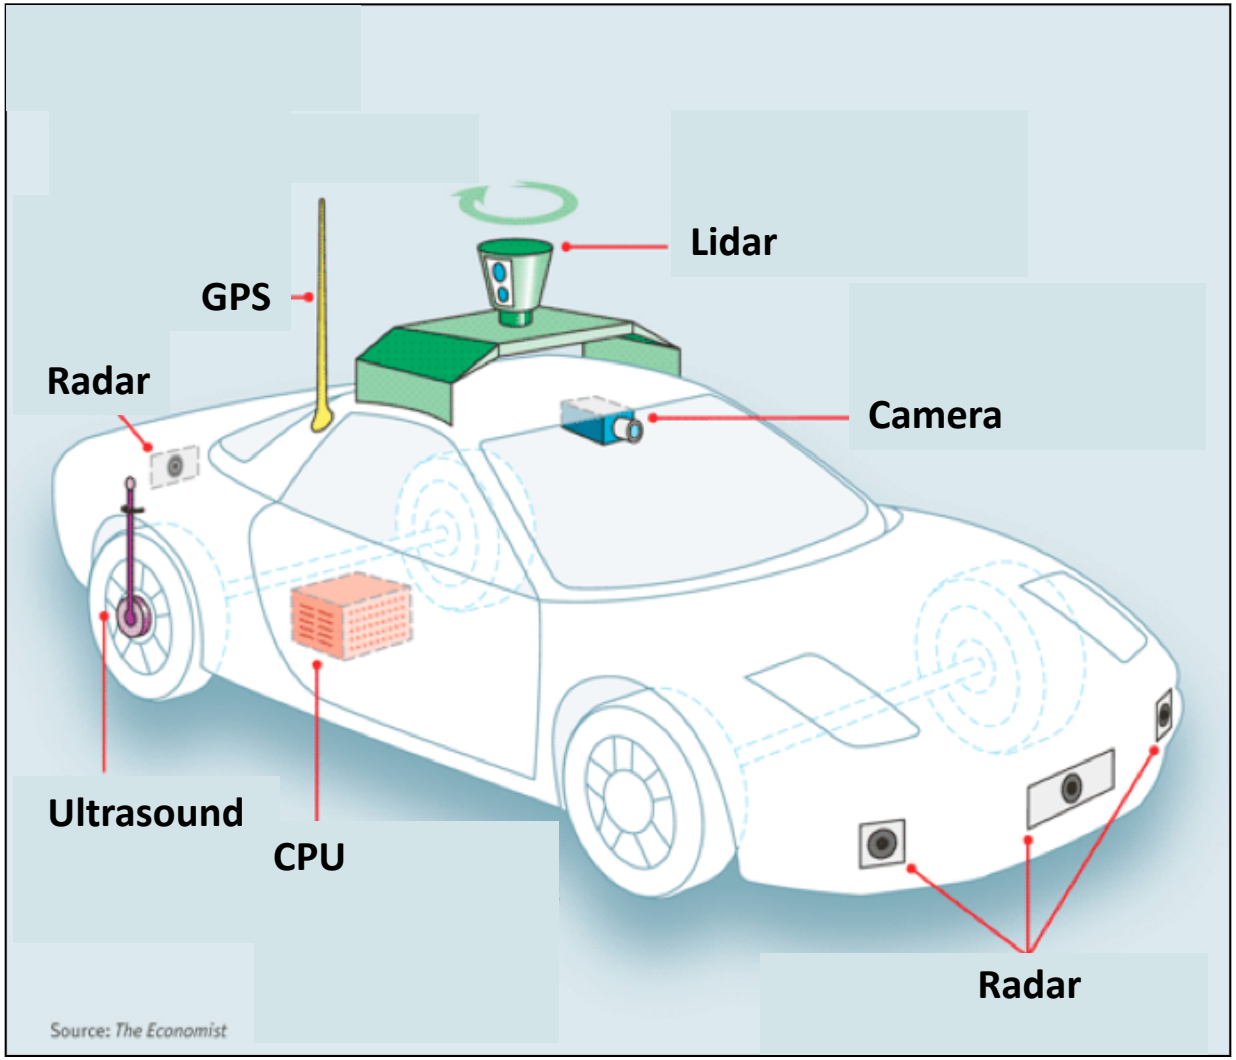

http://econ.st/ZZBLzi

Based on a world model from System 1, System 2 anticipates trajectories of others and makes steering decisions.

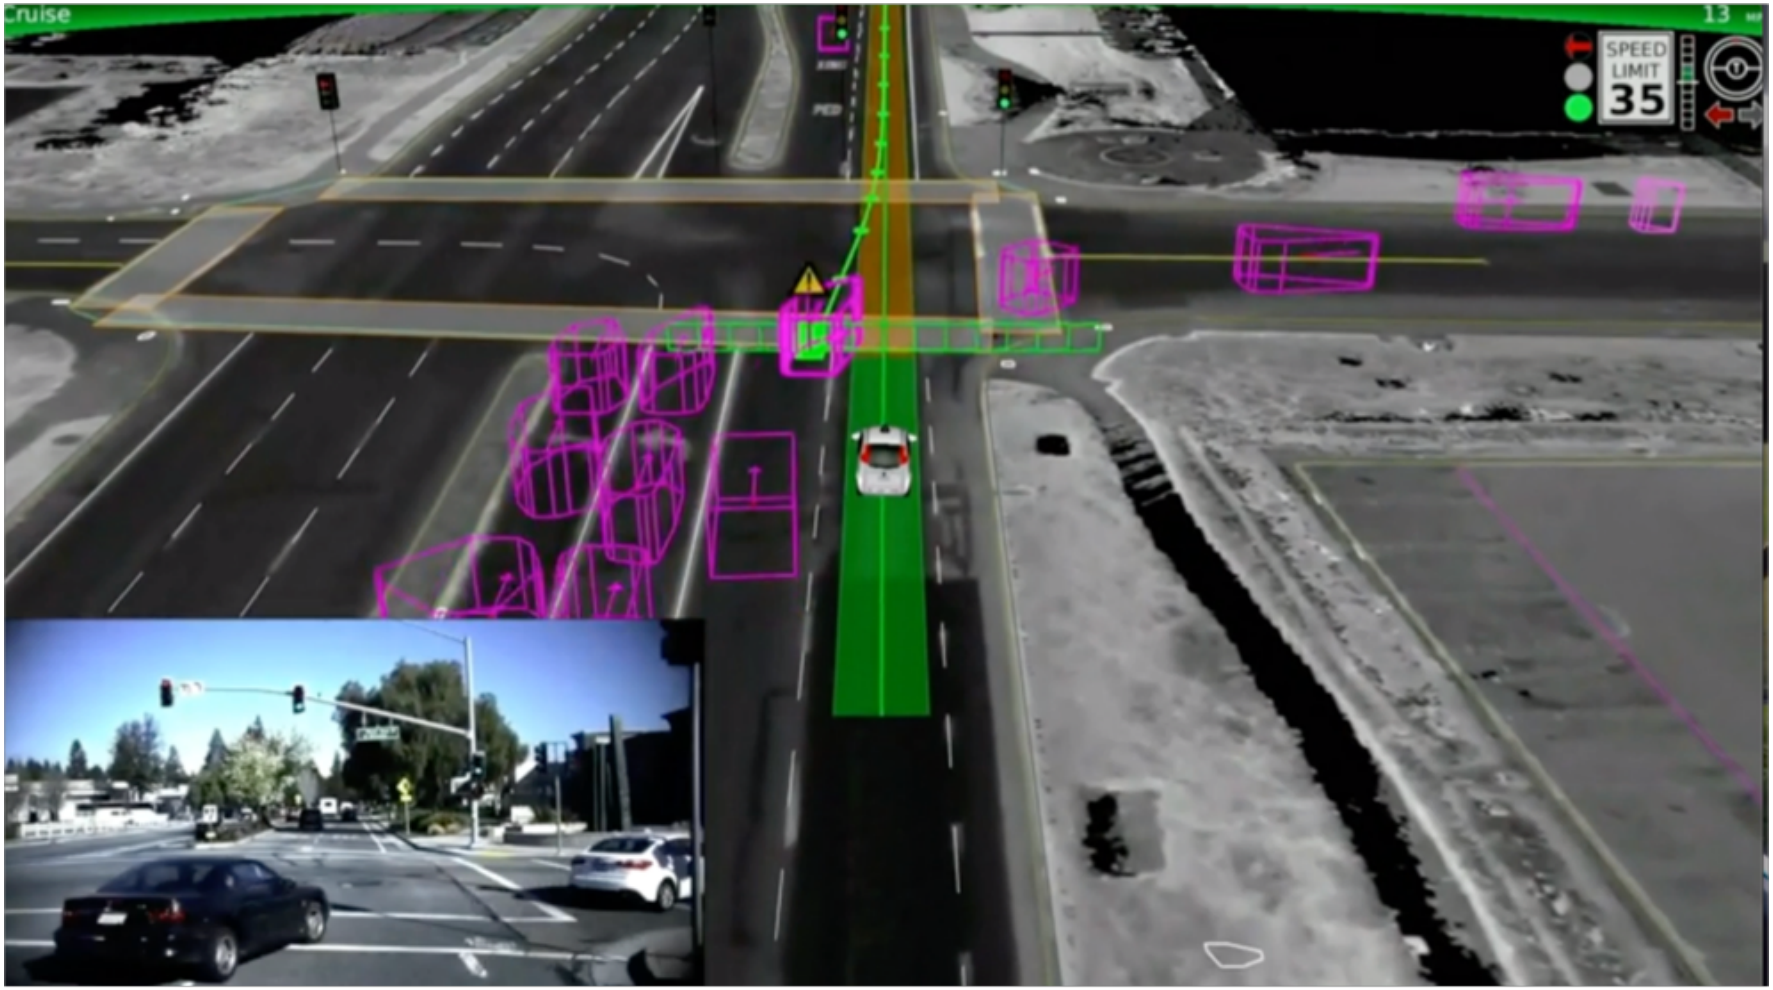

Source: Google

Where is AI used?

Mostly for **perception**,

not so much for steering.

Object Recognition / Scene Analysis

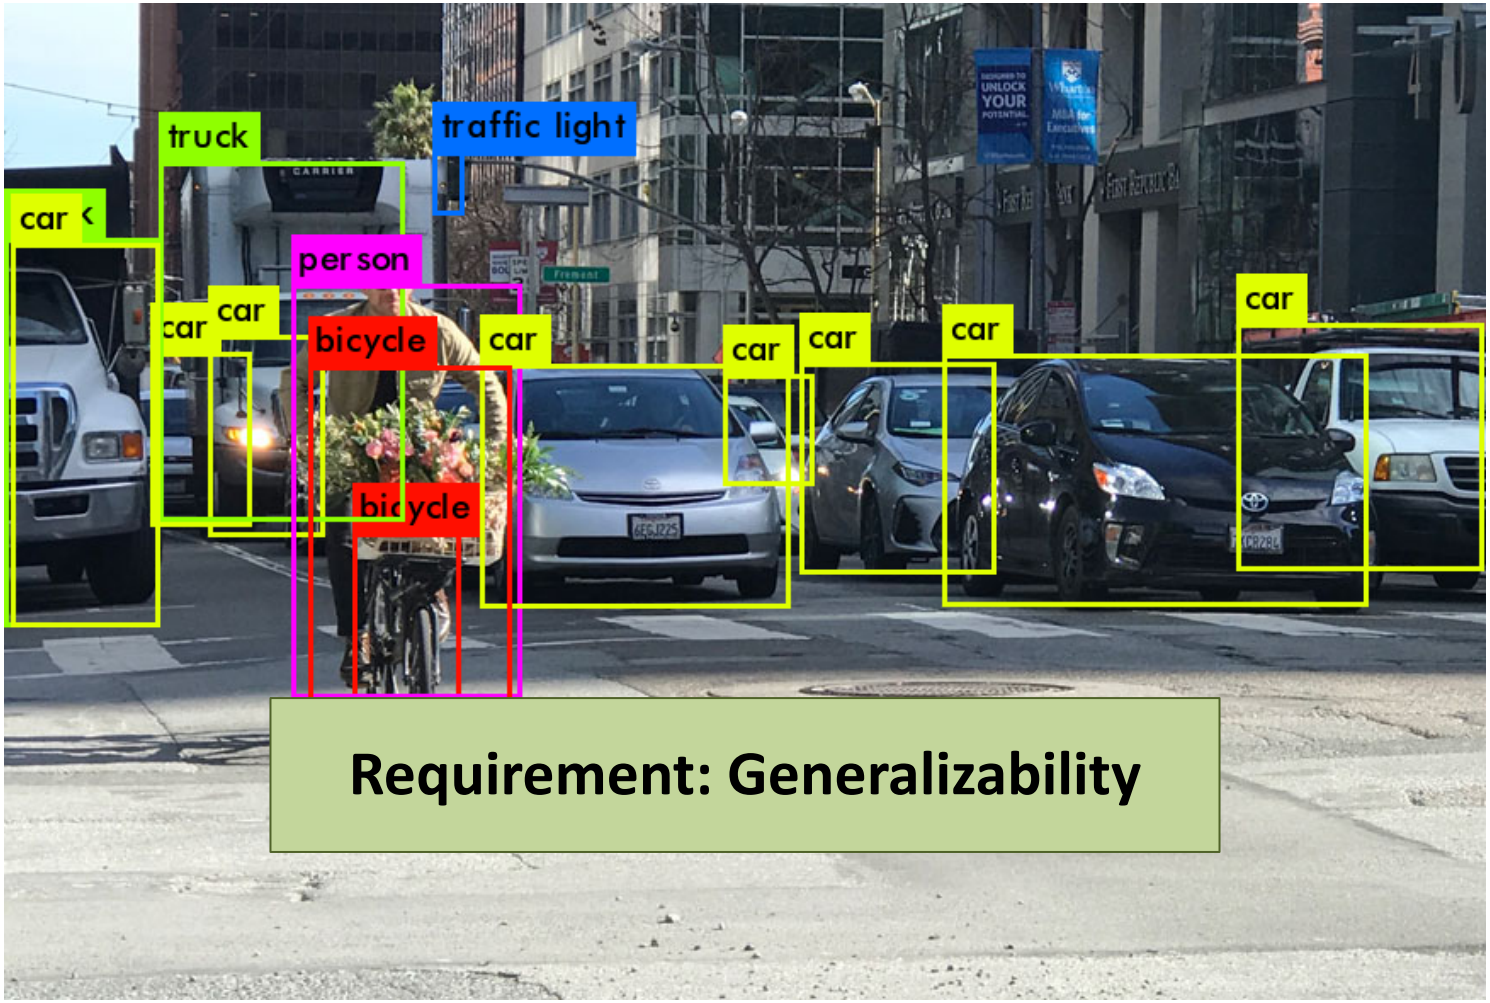

https://medium.com/@jonathan_hui/real-time-object-detection-with-yolo-yolov2-28b1b93e2088

Object Recognition / Scene Analysis

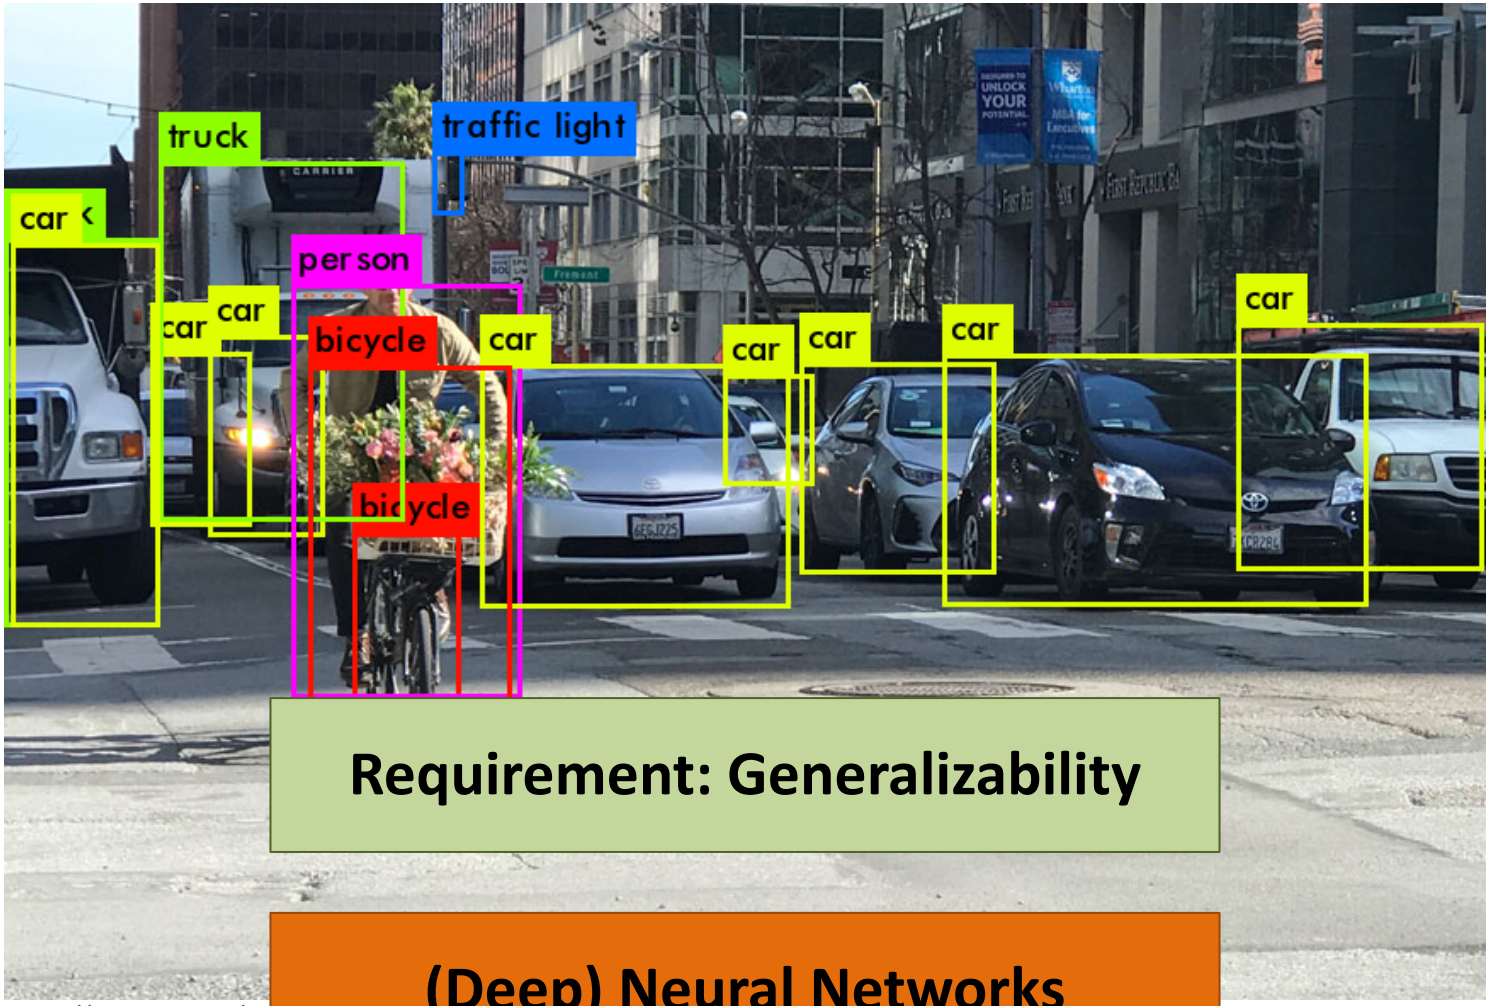

https://medium.com/@

Requirement: Generalizability

(Deep) Neural Networks

Simple neural network for digit detection

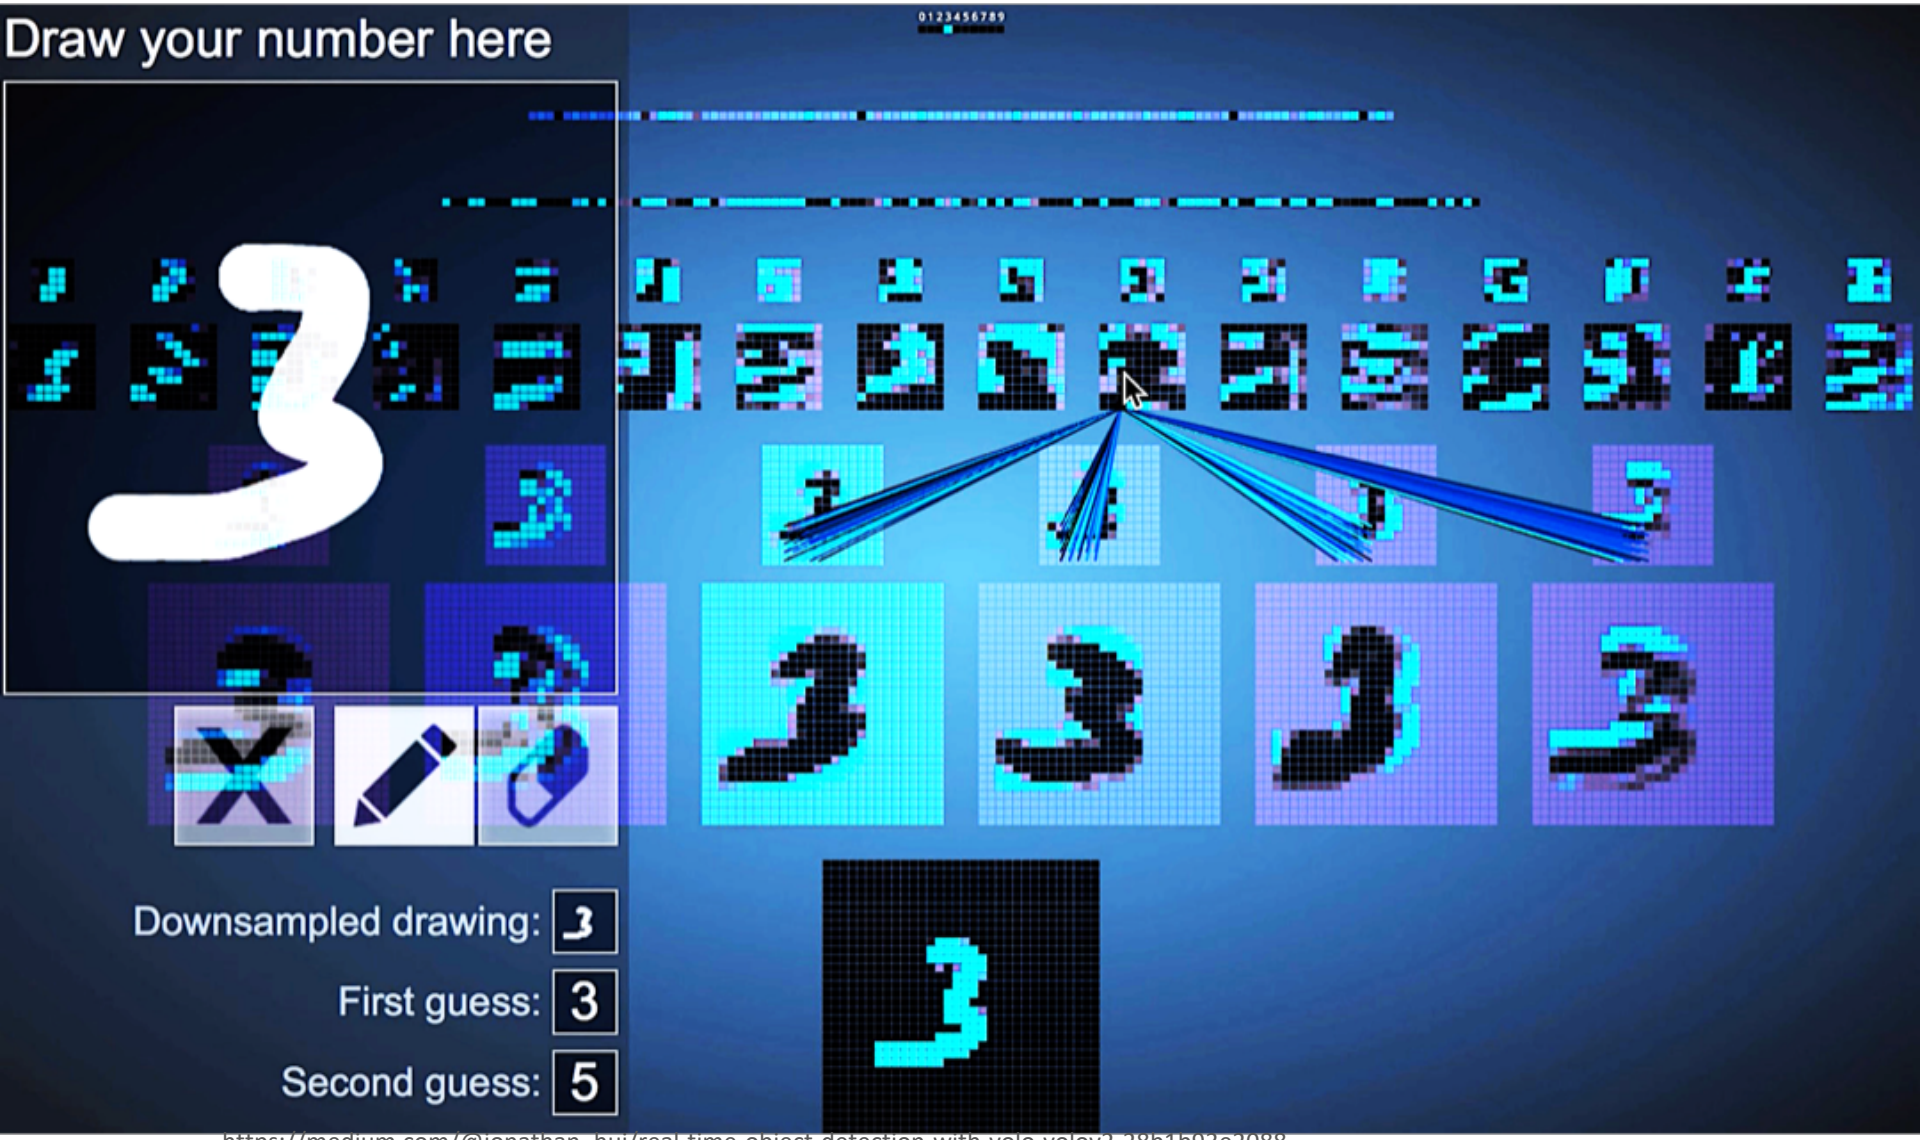

https://medium.com/@jonathan_hui/real-time-object-detection-with-yolo-yolov2-28b1b93e2088

http://scs.ryerson.ca/~aharley/vis/

Traffic sign detection: neural networks outperform humans.

METHOD	TOTAL
Committee of CNNs	99.46%
Human Performance	98.84%
Multi-Scale CNNs	98.31%
Random Forests	96.14%
LDA on HOG 2	95.68%
LDA on HOG 1	93.18%
LDA on HOG 3	92.34%

J. Stallkamp, M. Schlipsing, J. Salmen, C. Igel, Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition, Neural Networks, Volume 32, August 2012, Pages 323-332, ISSN 0893-6080, http://benchmark.ini.rub.de/?section=gtsrb&subsection=results

System 2: Steering

mostly **rule-based**

complex and error-prone

https://techthelead.com/look-easily-autonomous-cars-tricked/

Rule-based systems increase pressure on developers to make ethical decisions.

no helmet

higher chance to die (and a VIP)

with helmet

higher chance to survive

http://www.theguardian.com/environment/green-living-blog/2010/apr/07/david-cameron-cycles-without-helmet

System 2: Steering

Future directions:

Train cars to drive with machine learning **(reinforcement learning)**

pro: no need for hand-written rules and detailed maps

con: difficult to *learn* "common sense"

Reinforcement learning demonstration (June 2018)

https://wayve.ai/blog/learning-to-drive-in-a-day-with-reinforcement-learning

Self-driving cars as discussed are not autonomous.

Training only in the lab, model read-only on the road.

Behavior is entirely deterministic, yet **unpredictable** (complexity).

Research problems:

- Improve explainability of models (but for whom?)
- Additional safeguards ("artificial common sense")

Truck crossed highway, reflecting the sun – never happened during training.

Was it "only" a bug or is Tesla liable because of insufficient training? Or is it the truck driver's fault?

Extension: Retraining on the road.

"The whole Tesla fleet operates as a network. **When one car learns something, they all learn it** … each driver using the autopilot system essentially becomes an expert trainer for how the autopilot should work"

– Elon Musk

"True" autonomy is undesirable.
Manufacturers will want to be in the loop.

Security issue: Risk of malicious injection of faulty training data.

Computer vision is still very brittle and can be attacked cleverly.

https://arxiv.org/abs/1808.03305

Cars will not solely rely on their own perception, but communicate with others. This makes it more difficult to understand the reason of failures.

https://www.youtube.com/watch?v=5vkQJljZ2Qo

Foreseeable consequences

Self-driving cars *are* a black box.

Their behavior is complex and **difficult to predict** – even without retraining on the road.

Manufacturers will **collect a lot of data**. They might be inclined to provide only favorable evidence.

Citizens might be at a disadvantage to prove their case.

Managing this asymmetry is an important policy issue.