Machine Learning for Intelligent Transportation Systems

Patrick Emami (CISE), Anand Rangarajan (CISE), Sanjay Ranka (CISE), Lily Elefteriadou (CE)

MALT Lab, UFTI

September 6, 2018

ITS - A Broad Perspective

Working definition

Utilizing cutting-edge, synergistic technologies to develop and improve transportation systems of all kinds

ITS - A More Narrow Perspective

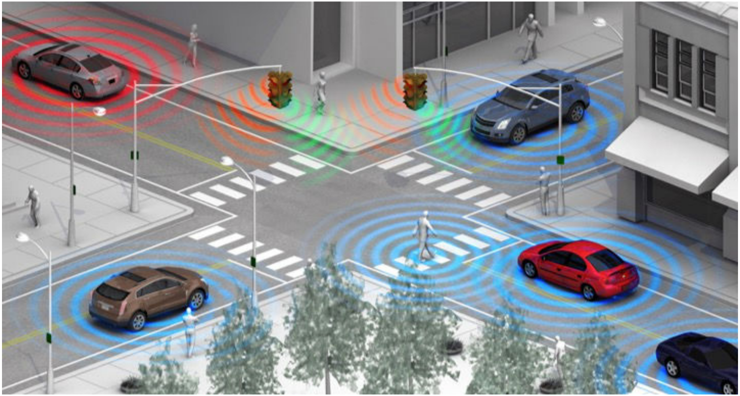

ITS for improved urban mobility

Source: https://www.arch2o.com/future-urban-mobility/

ITS for Urban Mobility - Autonomous Vehicles

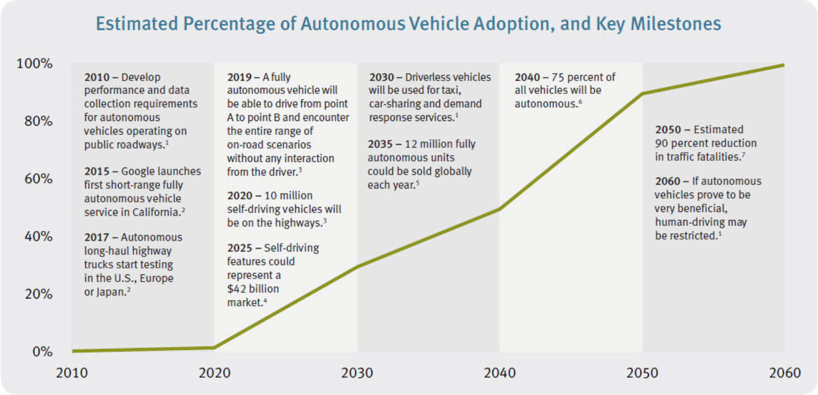

Source: http://www.vtpi.org/avip.pdf

ITS for Urban Mobility - Traffic Surveillance

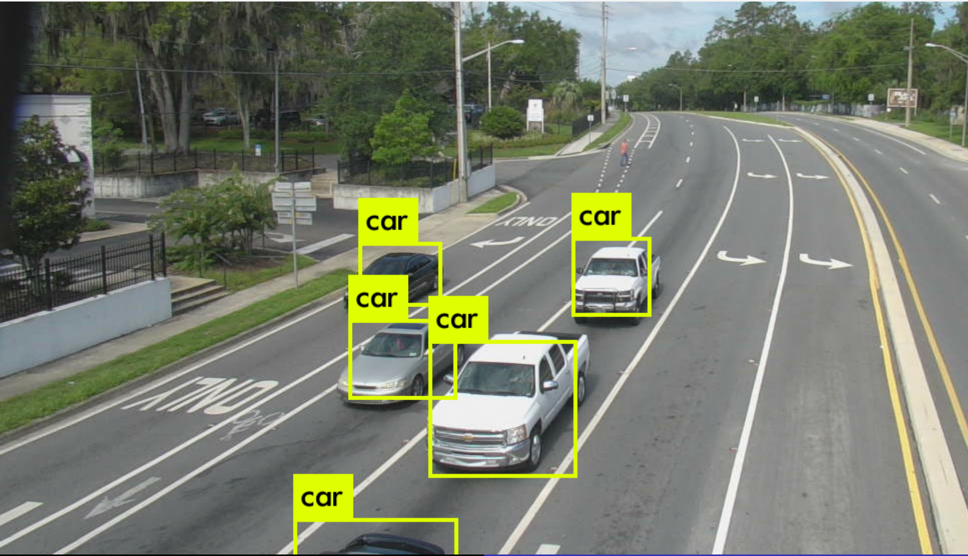

ITS for Urban Mobility - Traffic Optimization

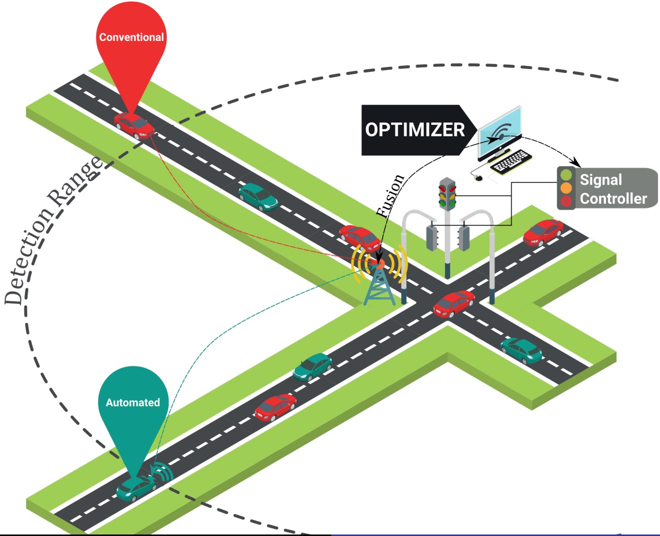

Machine Learning

Machine Learning

Working definition

Extracting patterns and abstractions from datasets to make intelligent decisions on previously unseen data

Other "Intelligent" Tools

Machine learning is rarely used in isolation, and often overlaps with the following fields:

1. Discrete and continuous optimization
2. Signal processing
3. Distributed systems
4. Control theory
5. And more...!

Machine Learning for ITS

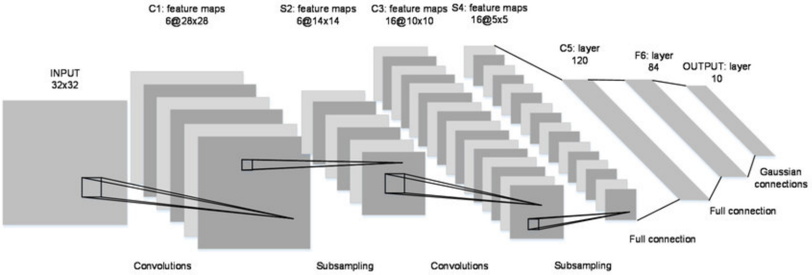

Deep neural networks trained on massive datasets are at the cutting-edge in terms of performance. The theory is lagging behind!

Source: http://yann.lecun.com/exdb/lenet/

Deep Learning

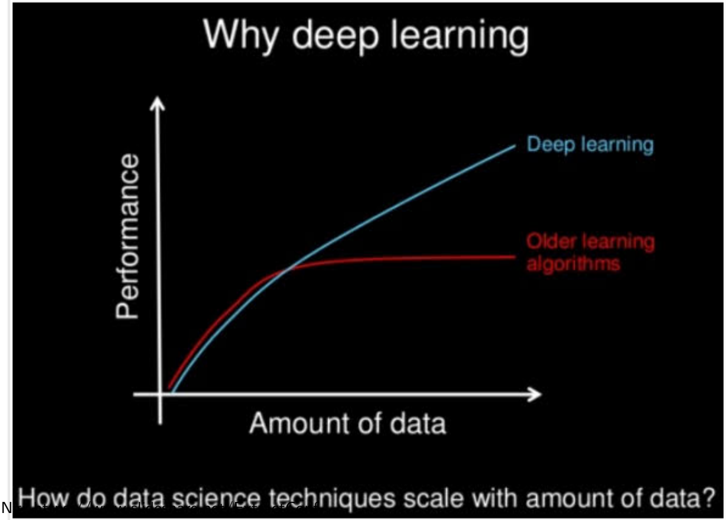

Source: Andrew N

ML ∩ Computer Vision

A primary use of ML in ITS is for intelligent perception

Some key tasks

1. Object detection
2. Multi-object tracking
3. Activity recognition

Autonomous Vehicles

Source: https://www.wired.com/story/waymo-launches-self-driving-minivans-fiat-chrysler/, http://sitn.hms.harvard.edu/flash/2017/self-driving-cars-technology-risks-possibilities/

Autonomous Vehicles

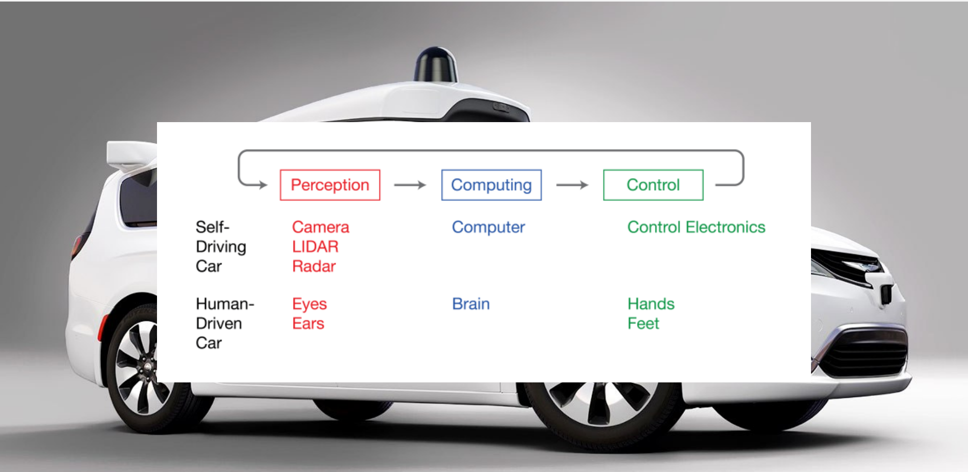

Source: https://www.wired.com/story/waymo-launches-self-driving-minivans-fiat-chrysler/, http://sitn.hms.harvard.edu/flash/2017/self-driving-cars-technology-risks-possibilities/

Traffic Surveillance

Use Computer Vision to try to answer these questions:

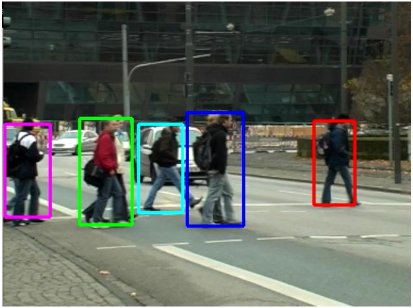

Are pedestrians crossing?

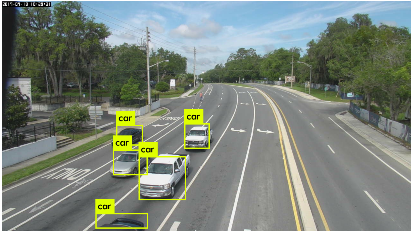

How many vehicles?
Any driving the wrong way?

Object detection

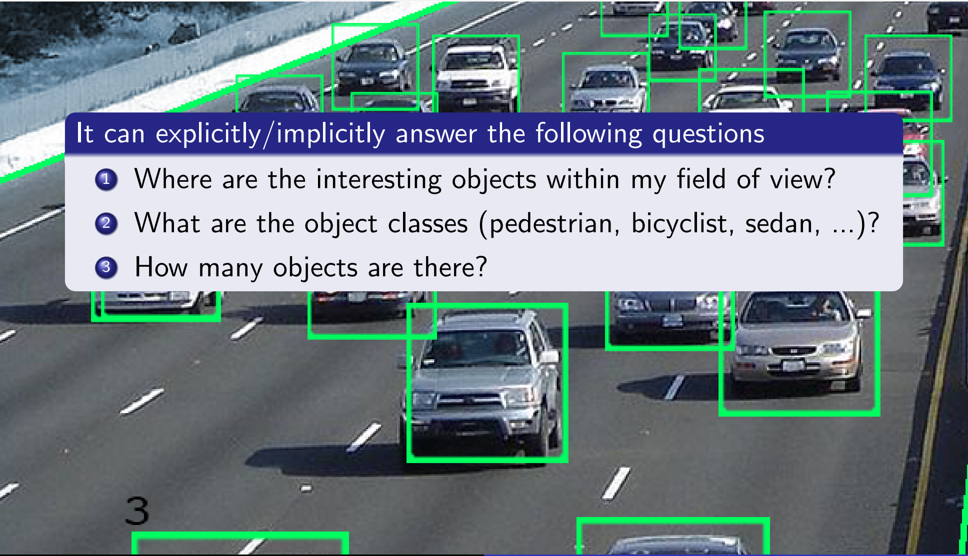

It can explicitly/implicitly answer the following questions

1. Where are the interesting objects within my field of view?
2. What are the object classes (pedestrian, bicyclist, sedan, ...)?
3. How many objects are there?

Object detection

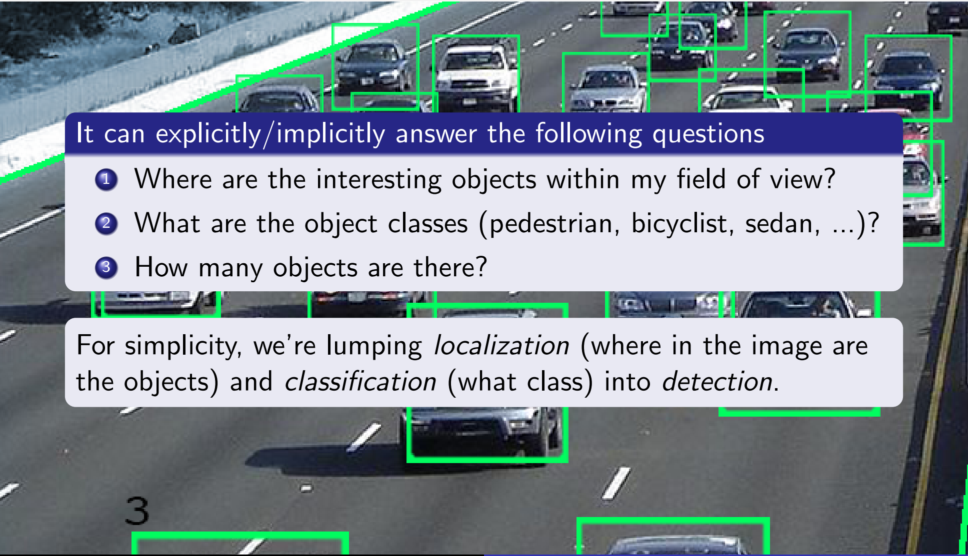

It can explicitly/implicitly answer the following questions

1. Where are the interesting objects within my field of view?
2. What are the object classes (pedestrian, bicyclist, sedan, ...)?
3. How many objects are there?

For simplicity, we're lumping *localization* (where in the image are the objects) and *classification* (what class) into *detection*.

Object Detection with Deep Learning

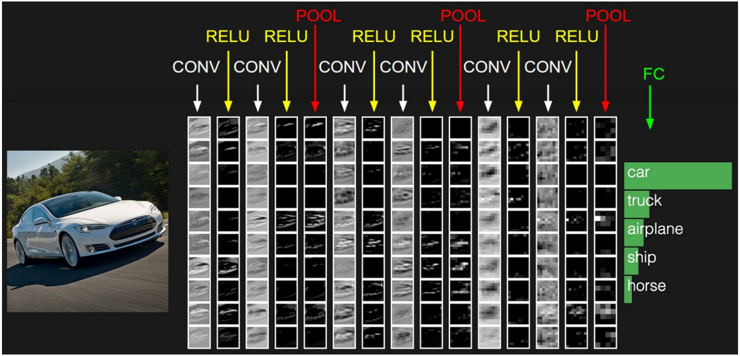

Real world challenges

The current best way to handle variations in lighting, orientation, and scale when deploying is *data augmentation*.

Source: http://cs231n.github.io/convolutional-networks/

Multi-object Tracking

Goal is to estimate the trajectories of all objects in a dynamic scene

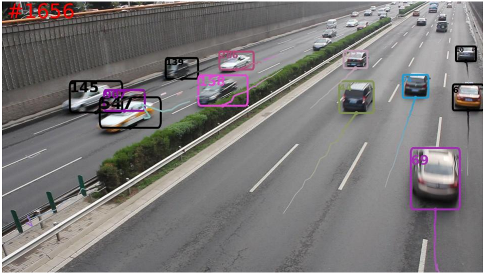

MOT from a stationary traffic cam

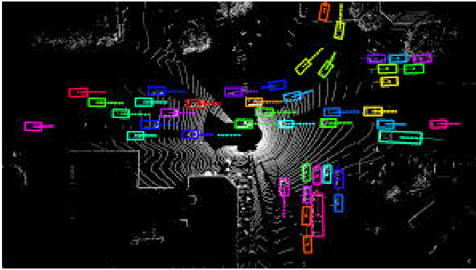

MOT using LiDAR from an AV

Source: Luo, et. al. "Fast and Furious: Real Time End-to-End 3D Detection, Tracking and Motion Forecasting With a Single Convolutional Net." CVPR 2018.

Obstacles to solving MOT

1. Object detectors don't handle partial/full occlusion or drastic variations in lighting, color, orientation very well
2. Stitching detections together over time into tracks is a hard discrete optimization (or inference) problem
3. Sensors are unreliable/noisy
4. MOT systems are typically overly-complex and contain lots of hand-tuned problem-specific parameters

Source: Emami, Patrick, et al. "Machine Learning Methods for Solving Assignment Problems in Multi-Target Tracking." arXiv preprint arXiv:1802.06897 (2018).

Obstacles to solving MOT

1. Object detectors don't handle partial/full occlusion or drastic variations in lighting, color, orientation very well
2. Stitching detections together over time into tracks is a hard discrete optimization (or inference) problem
3. Sensors are unreliable/noisy
4. MOT systems are typically overly-complex and contain lots of hand-tuned problem-specific parameters

Interesting research question keeping me up at night

Is there a principled way to learn the concept of object permanence within an MOT system?

Source: Emami, Patrick, et al. "Machine Learning Methods for Solving Assignment Problems in Multi-Target Tracking." arXiv preprint arXiv:1802.06897 (2018).

Activity Recognition

Using object detections and trajectories, can we then extract patterns at the level of behaviors?

1. Pedestrian safety; ID'ing whether a person is walking/about to walk into the street
2. Vehicle collision prediction
3. Multi-agent modeling at traffic intersections and merging zones for AVs

Collision Prediction

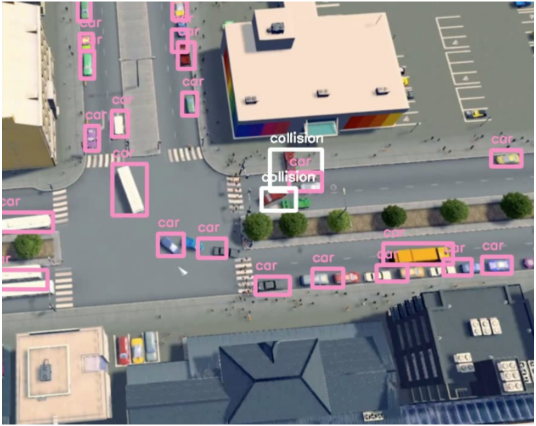

Source: Xiaohui Huang, Sanjay Ranka and Anand Rangarajan. Real-time Multi-Object Tracking and Road Traffic Safety Measurement. In preparation.

Traffic Optimization

Guiding question

Using sensors and edge computing, can we maximize the efficiency of traffic flow through a road network in real-time?

Traffic Sensors

Short-term Traffic Flow Prediction

Accurate forecasting of congestion levels enables real-time traffic planning

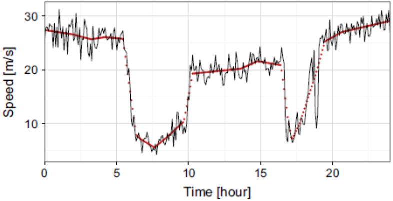

Train a model (e.g., deep network or Random Forest) to predict next 15-30 minutes of traffic flow.

Source: Polson, Nicholas G., and Vadim O. Sokolov. "Deep learning for short-term traffic flow prediction." Transportation Research Part C: Emerging Technologies 79 (2017): 1-17.

Traffic Intersection Optimization

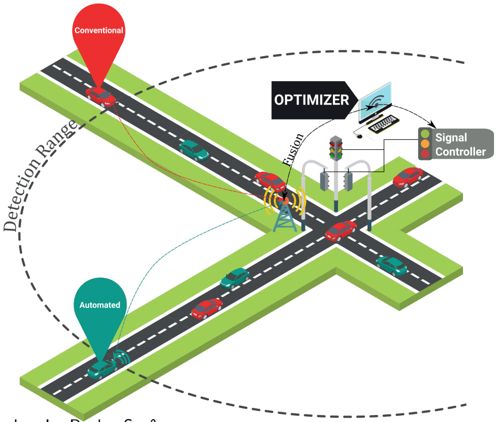

Source: Pourmehrab, M., Elefteriadou, L., Ranka, S., & Martin-Gasulla, M. "Optimizing Signalized Intersections Performance under Conventional and Automated Vehicles Traffic." arXiv:1707.01748 (2017)

Conclusion

Plenty of challenges when applying ML to ITS

1. Collecting, cleaning, and labeling large-scale datasets
2. Law-makers and policy has to keep up with the tech
3. Brittle models that break when applied to new domains
4. Security and privacy

Conclusion

Plenty of challenges when applying ML to ITS

1. Collecting, cleaning, and labeling large-scale datasets
2. Law-makers and policy has to keep up with the tech
3. Brittle models that break when applied to new domains
4. Security and privacy

But we've made great progress!

Thank you!

Questions?

Twitter: @patrickomid, email: pemami@ufl.edu

Slides available at: https://pemami4911.github.io