The AI Revolution

Doina Precup, PhD, Canada-CIFAR AI chair, McGill University & Mila
Associate Scientific Director of Healthy Brains for Healthy Lives CFREF program, McGill
Research Team Lead, DeepMind Montreal
Senior Member, American Association for Artificial Intelligence
Senior Fellow, CIFAR Learning in Machines & Brains

The Industrial revolution:

Machines extending humans'
MECHANICAL power

The AI revolution:

Machines extending humans'
COGNITIVE power

- Digital economy => AI economy
- All sectors of economy

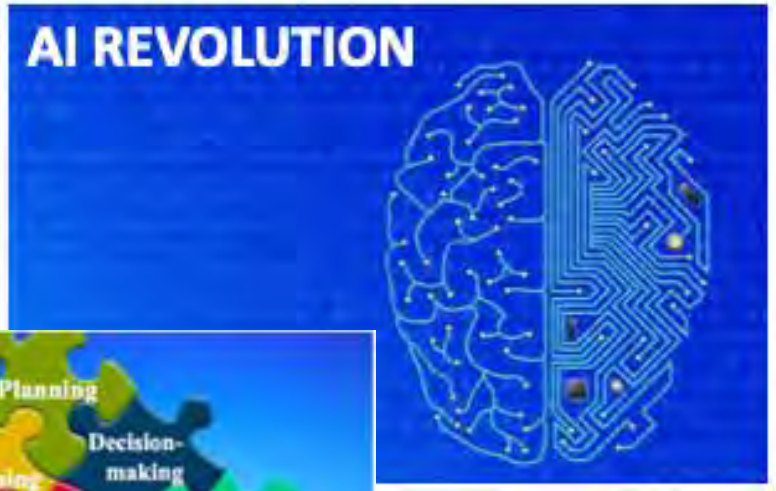

Why AI now?

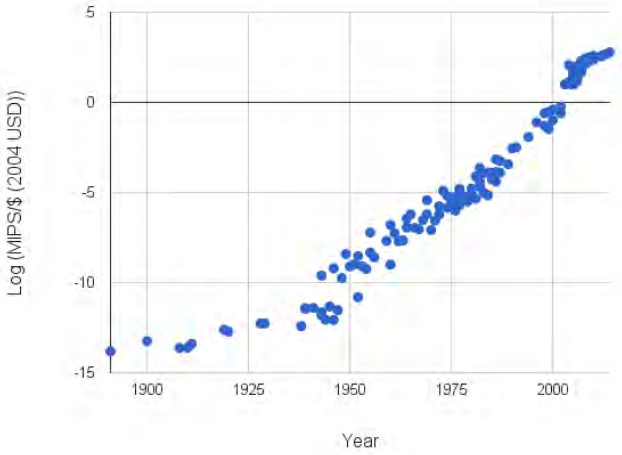

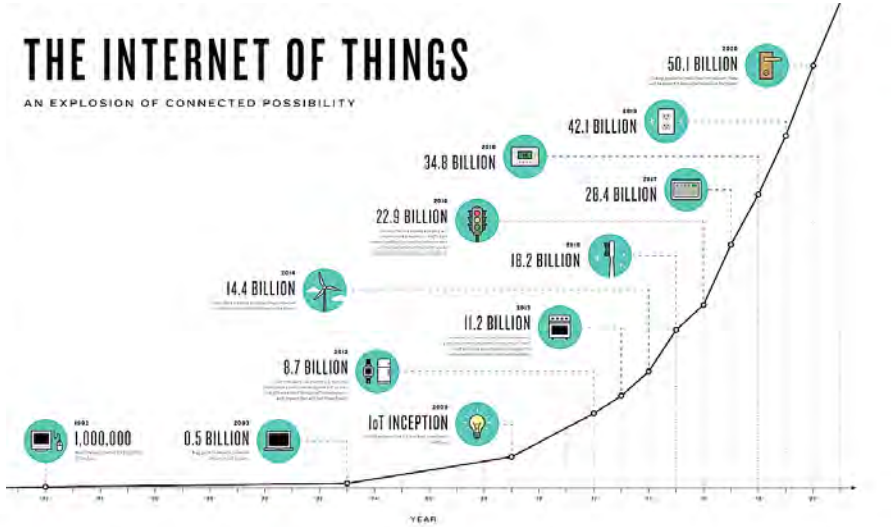

- Cost per computation and memory unit is rapidly decreasing
- Amounts of data generated through new measuring devices (ChipSeq, MRI/fMRI, cameras etc) is exploding
- Eg. IBM estimated that 90% of all data available today was created in the last 2-3 years

McGill

Supervised Learning

Given: Input data and desired output
Eg: Images and parts of interest

Goal: find a function that can be used for new inputs, and that matches the provided examples

McGill

Deep learning

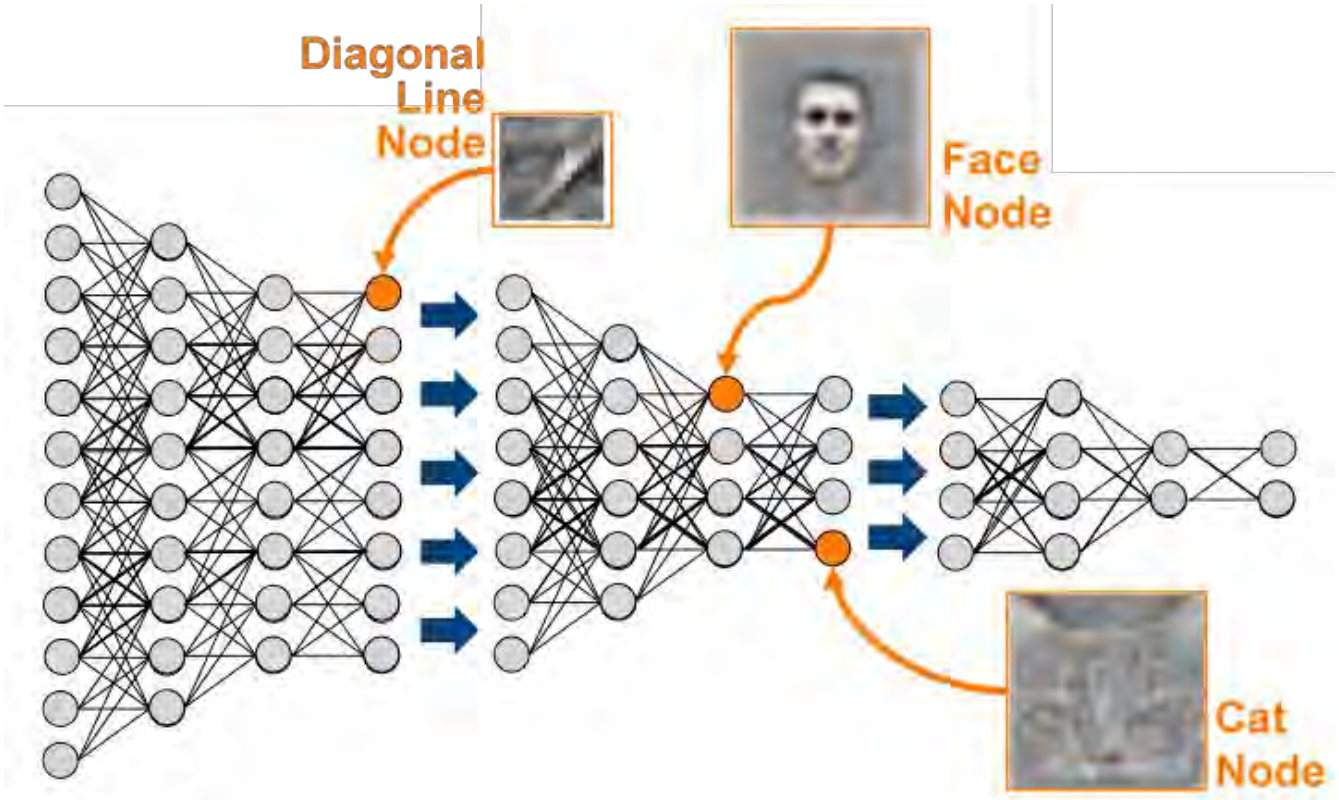

Deep Learning Revolution: Image Recognition

ImageNet: 1M images

1000+ object categories

2015: Human-level performance

What's the secret?

15-yr old algorithms

+ GPUs + 10x more data

Application: Sensor processing on autonomous cars

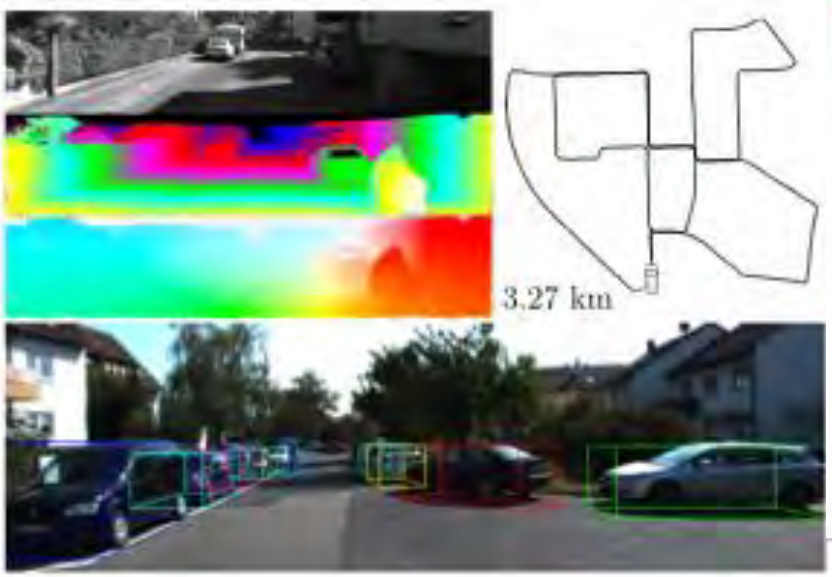

Cf. Urtasun et al, Univ. Toronto & Uber

McGill

Unsupervised Learning

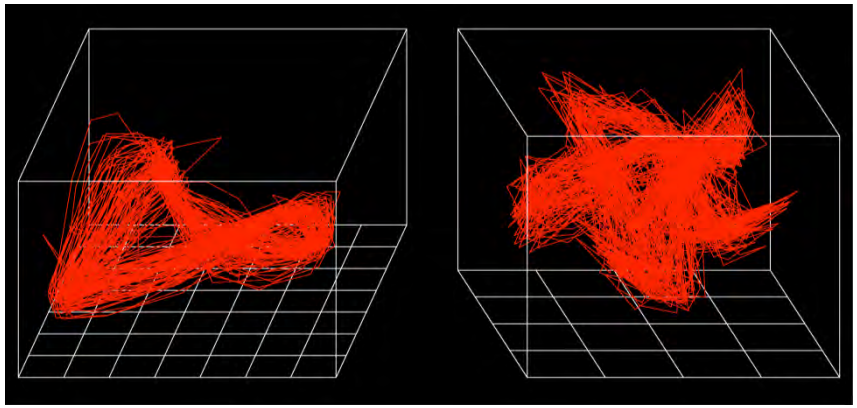

Given: Just data!
Eg: accelerometer information from a mobile phone

Goal: find “interesting” patterns
Often there is no single correct answer

Example: Mode of transportation

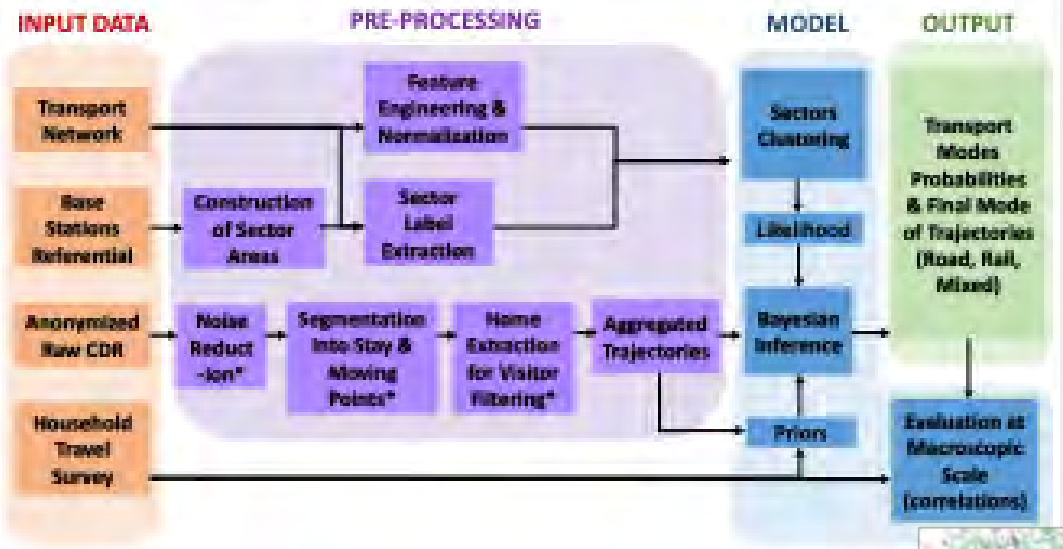

Cf Bachir et al, 2018

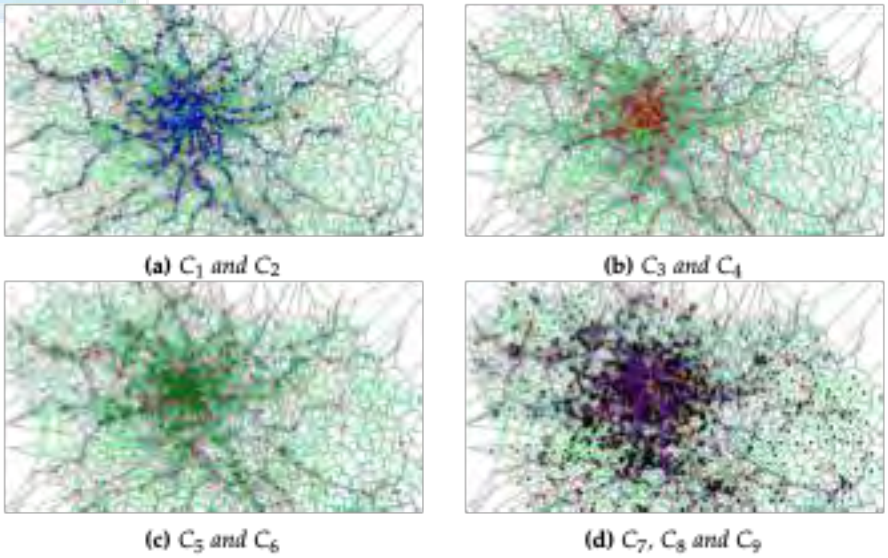

(a) C_1 and C_2

(b) C_3 and C_4

(c) C_5 and C_6

(d) C_7, C_8 and C_9

Reinforcement Learning

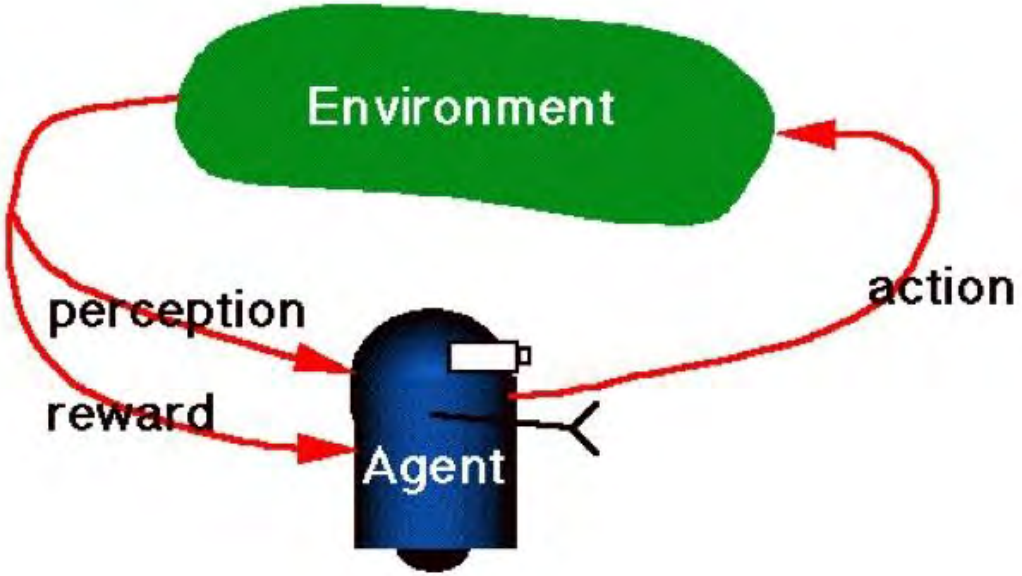

Reward: Food or shock

Reward: Positive and negative numbers

- Learning by trial-and-error
- Reward is often delayed

Example: AlphaGo & AlphaZero

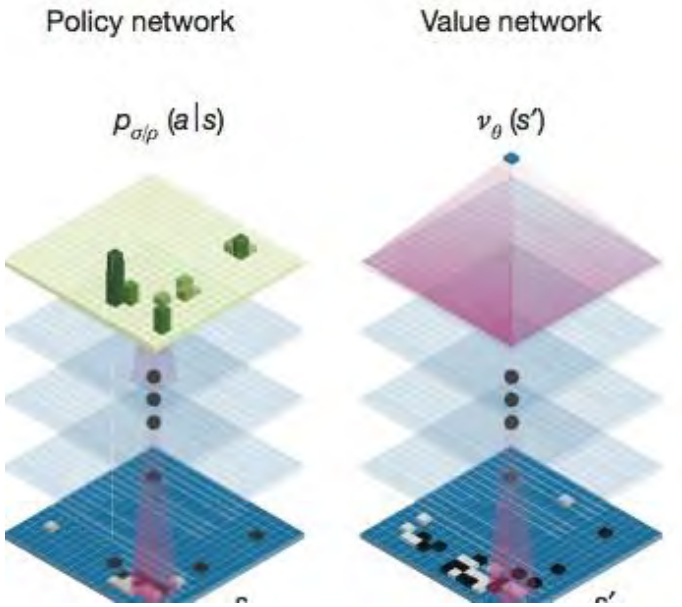

- Perceptions: state of the board
- Actions: legal moves
- Reward: +1 or -1 at the end of the game
- Trained by playing games against itself
- Invented new ways of playing which seem superior

McGill

Example: AlphaGo (DeepMind)

ARTICLE

doi:10.1038/nature16961

Mastering the game of Go with deep neural networks and tree search

David Silver[1]*, Aja Huang[1]*, Chris J. Maddison[1], Arthur Guez[1], Laurent Sifre[1], George van den Driessche[1], Julian Schrittwieser[1], Ioannis Antonoglou[1], Veda Panneershelvam[1], Marc Lanctot[1], Sander Dieleman[1], Dominik Grewe[1], John Nham[2], Nal Kalchbrenner[1], Ilya Sutskever[2], Timothy Lillicrap[1], Madeleine Leach[1], Koray Kavukcuoglu[1], Thore Graepel[1] & Demis Hassabis[1]

Dynamic programming

$$V(S_t) \leftarrow E_\pi\left[R_{t+1} + \gamma V(S_{t+1})\right] = \sum_a \pi(a|S_t) \sum_{s',r} p(s', r|S_t, a)[r + \gamma V(s')]$$

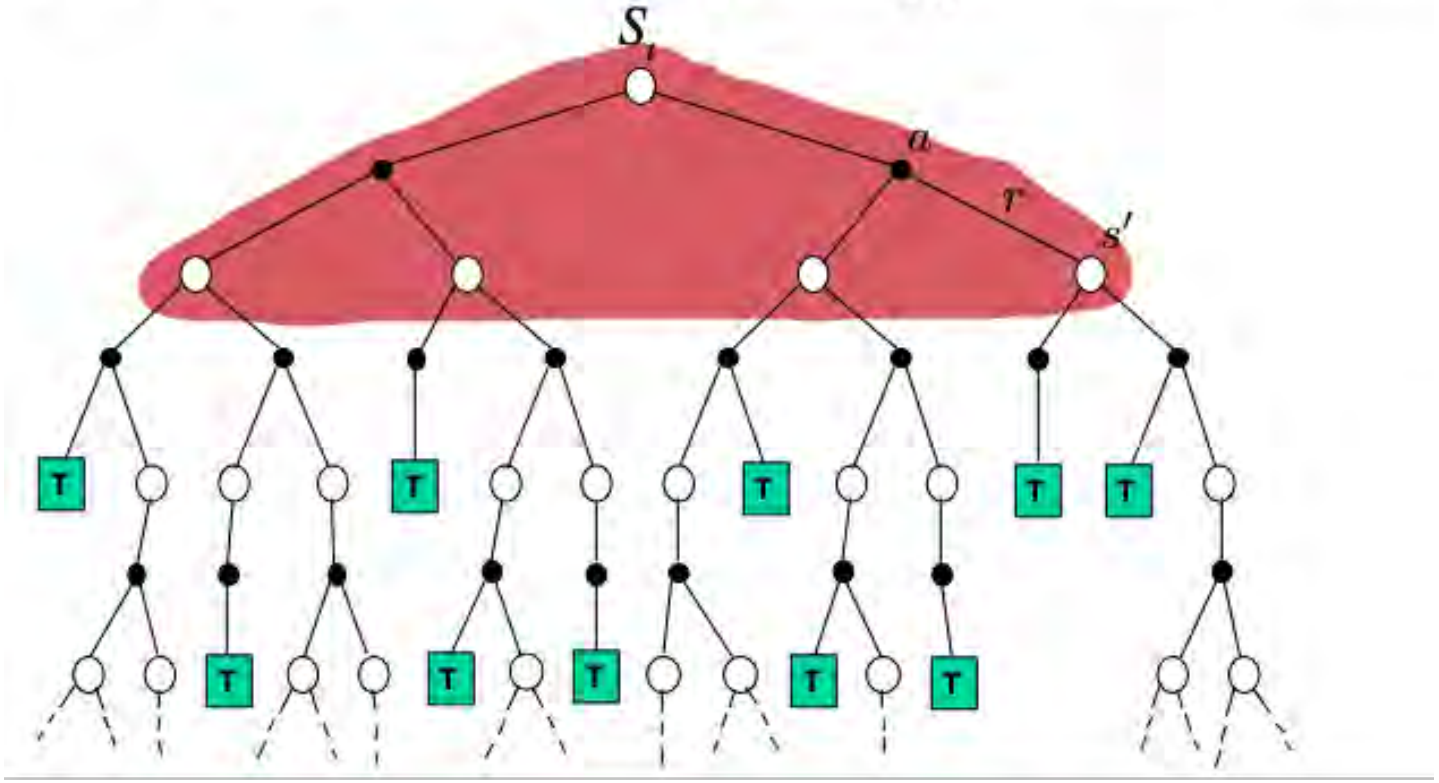

Monte Carlo

$$V(S_t) \leftarrow V(S_t) + \alpha\left[G_t - V(S_t)\right]$$

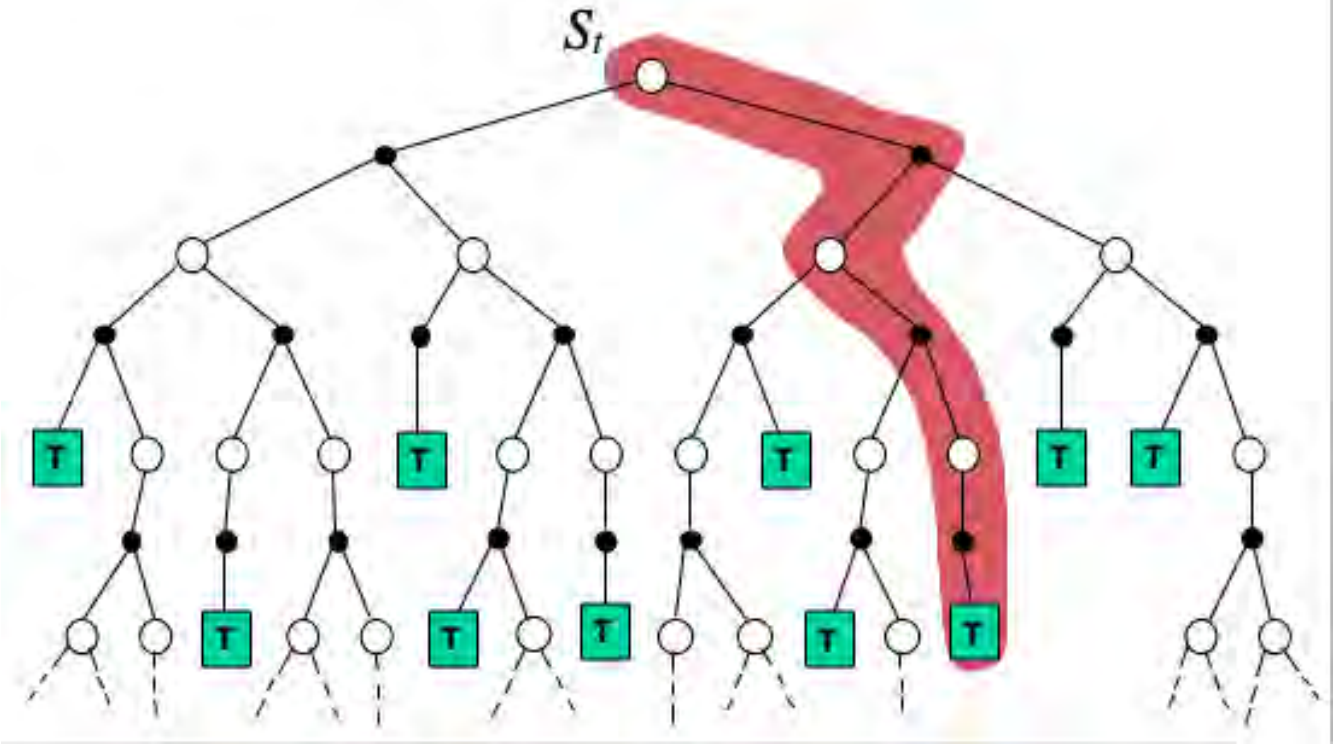

Temporal-difference learning

$$V(S_t) \leftarrow V(S_t) + \alpha\left[R_{t+1} + \gamma V(S_{t+1}) - V(S_t)\right]$$

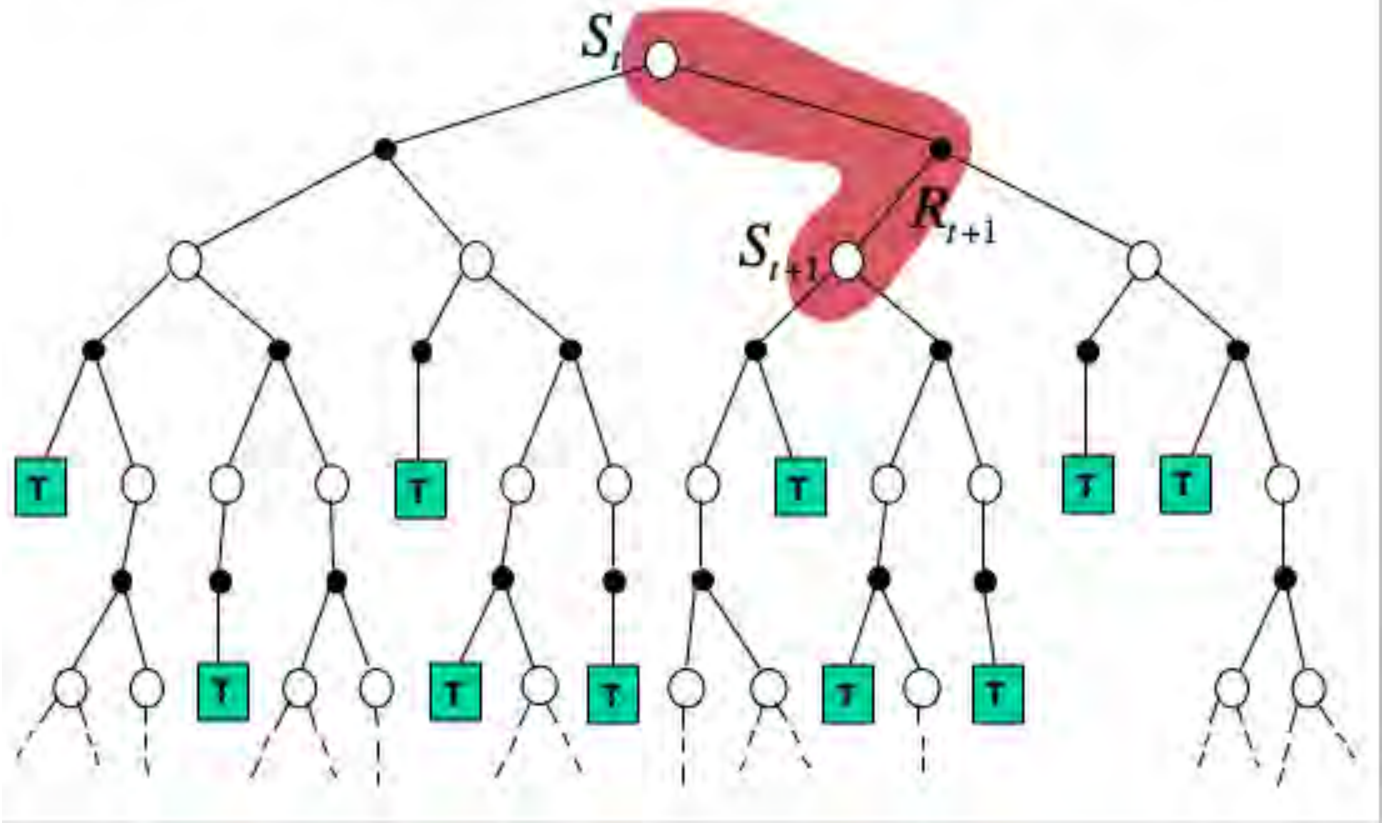

Application: Route Planning

- Planning a route for a trip on map
  - Distance, traffic
  - Road network known
  - Shortest travel time, avoid congestion

- Planning a route for robot navigation
  - With or without map
  - Perception as input

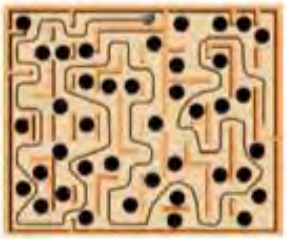

McGill

Application: Traffic signal control

- **Background**
  - Traffic lights control traffic flow at intersections.
  - Affects throughput, delay, waiting time, etc
- **Traditional methods**
  - Fixed-time intervals for red-yellow-green
  - Traffic model-based methods
- **Road network**
  - Multiple intersections: control at one intersection has impact on neighboring intersections.

McGill

Application: Vehicle repositioning

- **Setup**
  - Grid world, system-wide repositioning
  - Independent driver policy – less coordination
- **Agent**
  - CNN-based, input: vehicles status, projected supply distribution, future demand
  - Action: neighboring grid
  - Reward: weighted sum of fulfillment – pick-up distance
  - Drivers execute actions in order.
- **Training**
  - Simulation based on NYC taxi data
  - Double DQN
- **Results**
  - Reduced unfulfilled requests by 20% compared to RHC.

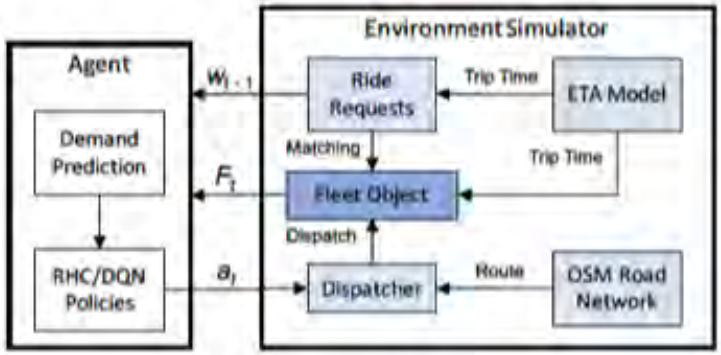

Cf Oda et al, 2018

Application: StreetLearn

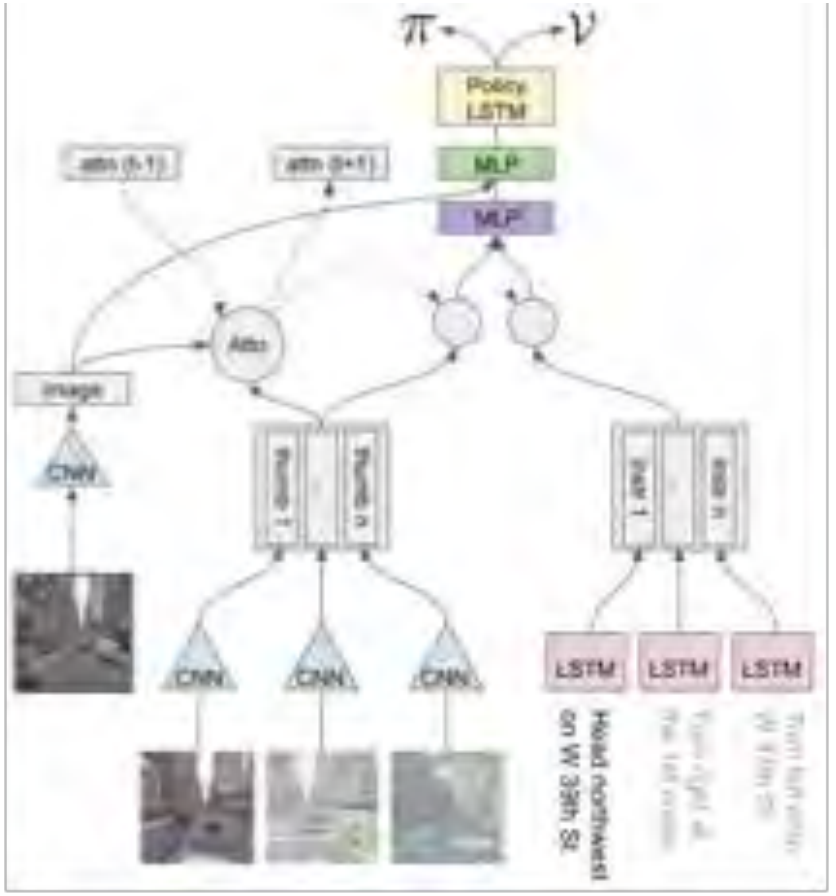

Cf. Hadsell et al, 2018, 2019

McGill

Opportunities and Challenges

- AI methodology is becoming very mature
- But prediction vs causal mechanism is still a open problem
- Training in simulation vs deployment
- Safety / risk management need to be incorporated
- Ethical considerations need to be incorporated