

Fighting Fire with AI

Why AI is Different...and How

BC AI Wildfire Symposium
October 12, 2018

Mark Crowley
Assistant Professor
Electrical and Computer
Engineering Department
University of Waterloo

Outline

- My Background
- Artificial Intelligence and Machine Learning, Big Data and all that...
- Concepts: Reinforcement Learning, Deep Learning, Deep RL, MCTS
 - Results: Learning Agent Based model of forest fire spread from satellite data
- Concepts: Surrogate Models and Parameter Optimization
 - Results: Fires Simulators and Weather Estimation
 - Results: Let Burn Analysis, Fire Treatment Planning
- Some other Relevant AI/ML Methods
- Future Challenges and Opportunities

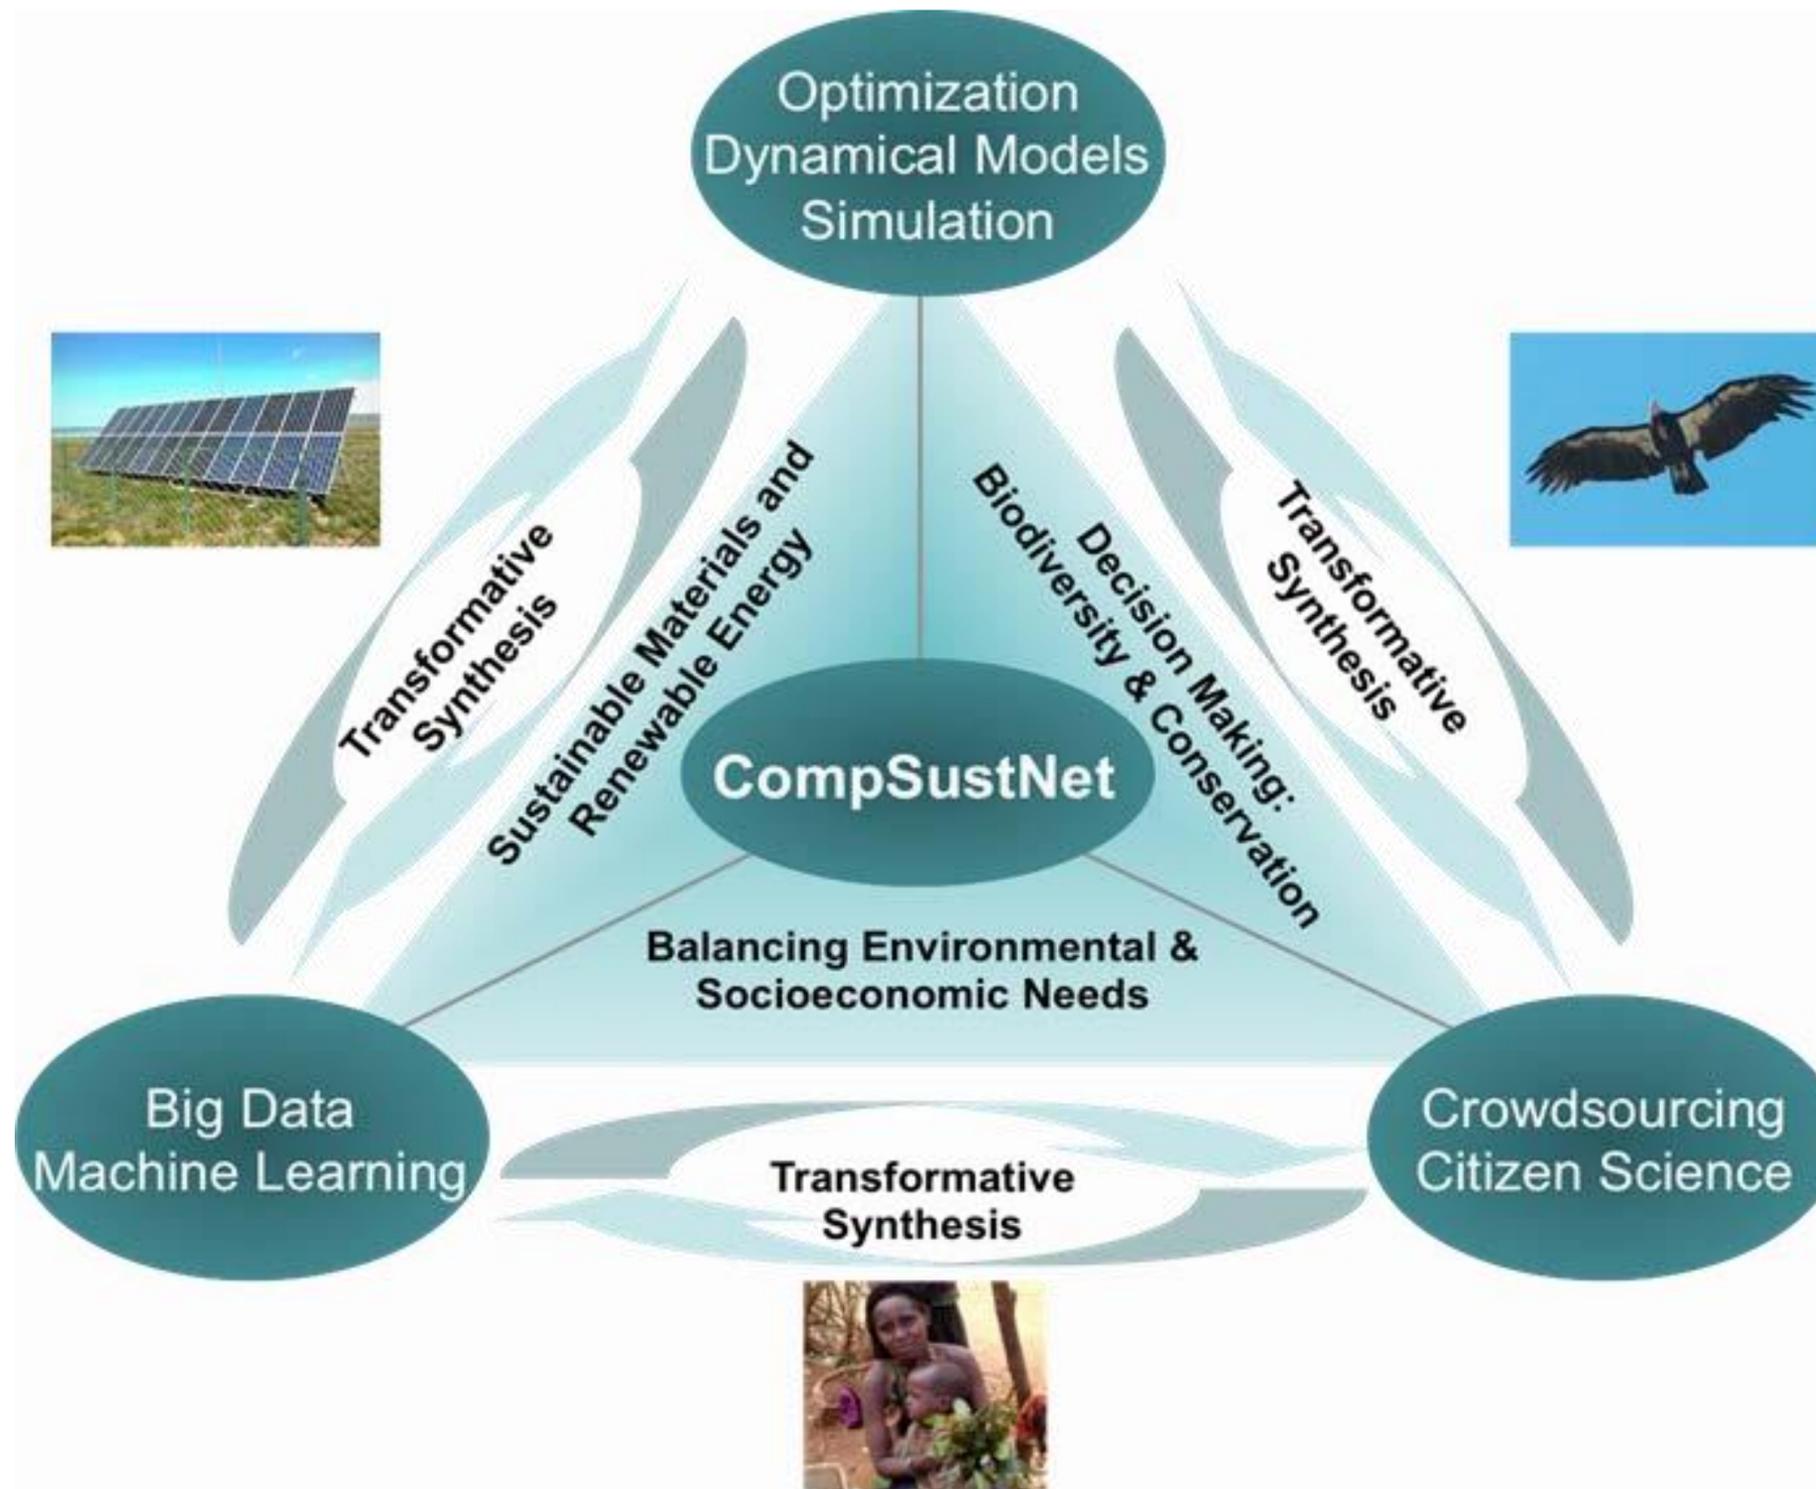
About Me

- PhD : Computer Science at the University of British Columbia
 - Inference in Bayesian networks and Cyclical Causal Models
 - Reinforcement Learning for Sustainable Forest Harvest Management
- Postdoc : Oregon State University
 - NSF Project on **Computational Sustainability** (<http://www.compsust.net>)
 - Forest Fire Planning – “Treat” vs “Let Burn” decision making
 - Developed PAC MDP Planning for Control of Invasive Species
- Assistant Professor : UWaterloo – ECE Department
 - My lab: UW ECE ML <https://uwaterloo.ca/scholar/mcrowley/lab>
 - **Waterloo.AI : Waterloo Artificial Intelligence Institute**
 - WICI: Waterloo Institute for Complexity and Innovation
 - Element AI: Faculty Research Fellow

Good Timing - Other Recent Related Initiatives

- United Nations Sustainable Development Goals (UN-SDGs)
 - 2nd AI For Good Summit 2018, Geneva – ITU/UN Initiative
 - Using AI to tackle Equality, Social, Economic, Infrastructure, Energy, Environmental issues
 - AI for Social Good Workshop, at the 2018 NIPS machine learning conference
 - Sustainable Development Network in Canada (hosted at UWaterloo)
- Other corporate and non-profit initiatives
 - AI4Good.org
 - Microsoft – AI for Earth (Azure cloud computing)
 - IBM Watson AI XPrize
- Computational Sustainability
 - Institute for Computational Sustainability (Cornel, OSU) NSF <\$10M, 2008>
 - Computational Sustainability Network NSF <2016>
 - Tracks at international AI conferences: AAI, IJCAI, ...
- Fire Workshops (and these are just the ones I've attended!)
 - Forest and Wildland Fire Management: a Risk Management Perspective. Banff, BIRS 2017
 - Wildland Fire Appropriate Response Workshop, Western U, Feb 2018.

Computational Sustainability



Find out more: <http://www.compsust.net/>

Research Vision

*To augment human **decision making** in complex domains and environments in a **dependable and transparent** way.*

Domains

(where does the complexity come from?)

- **Spatially Spreading processes in natural systems**
 - Fire, invasive species, floods
 - Medical images
- Classification and Anomaly Detection for Streaming Big Data
- Automotive
 - Driver Behaviour Learning
 - Autonomous Driving
 - Object classification and understanding

Methods

(how do we solve it?)

- Reinforcement Learning
- Deep Learning
- Ensemble Methods
- Feature Reduction/Extraction

Some of My Related Research

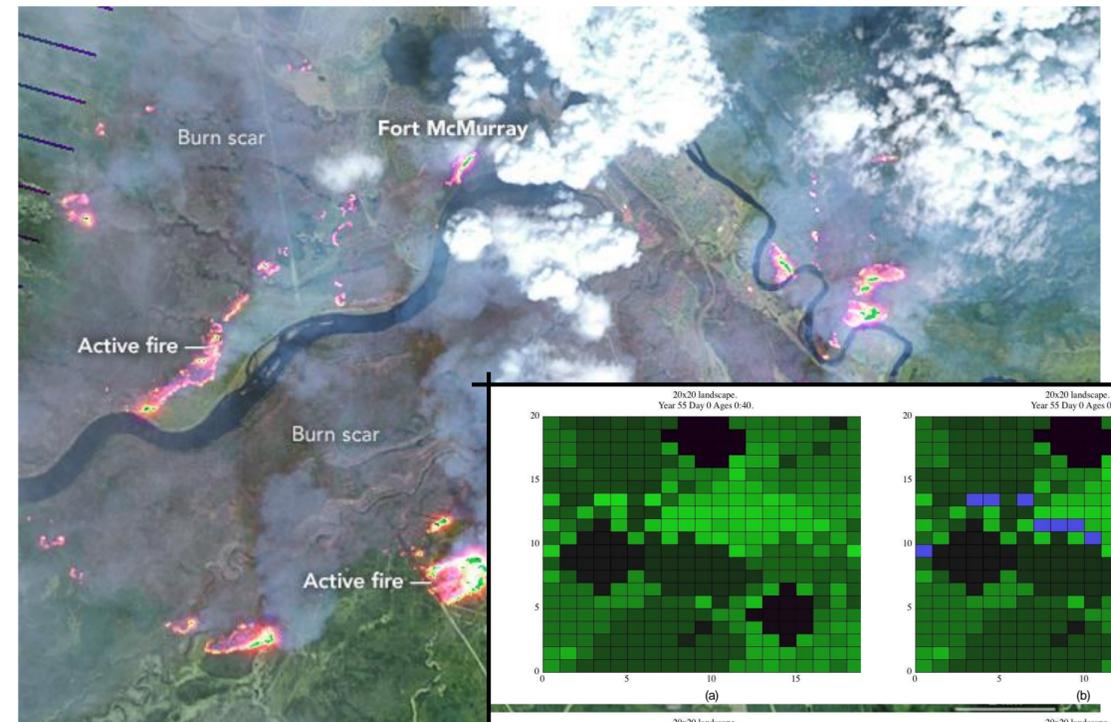
Invasive Species Control



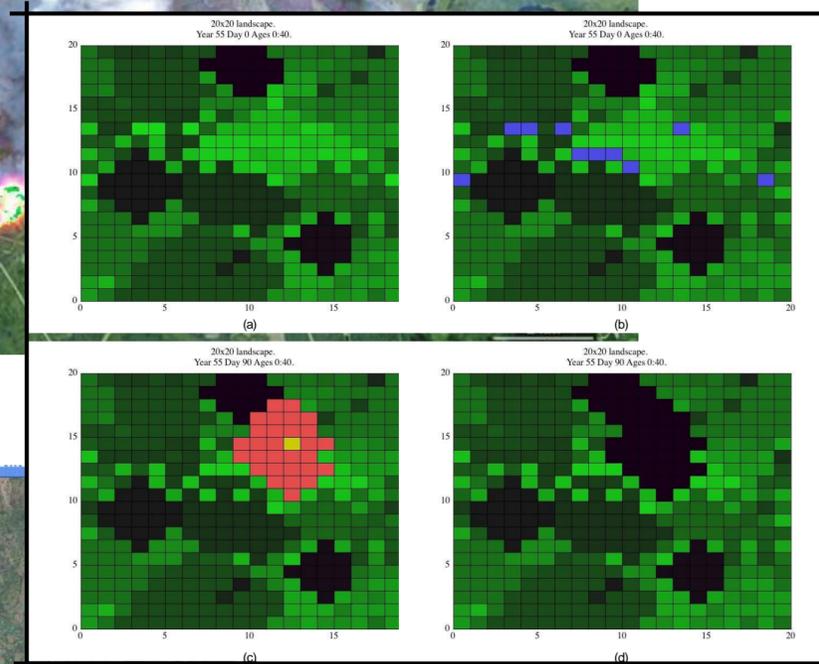
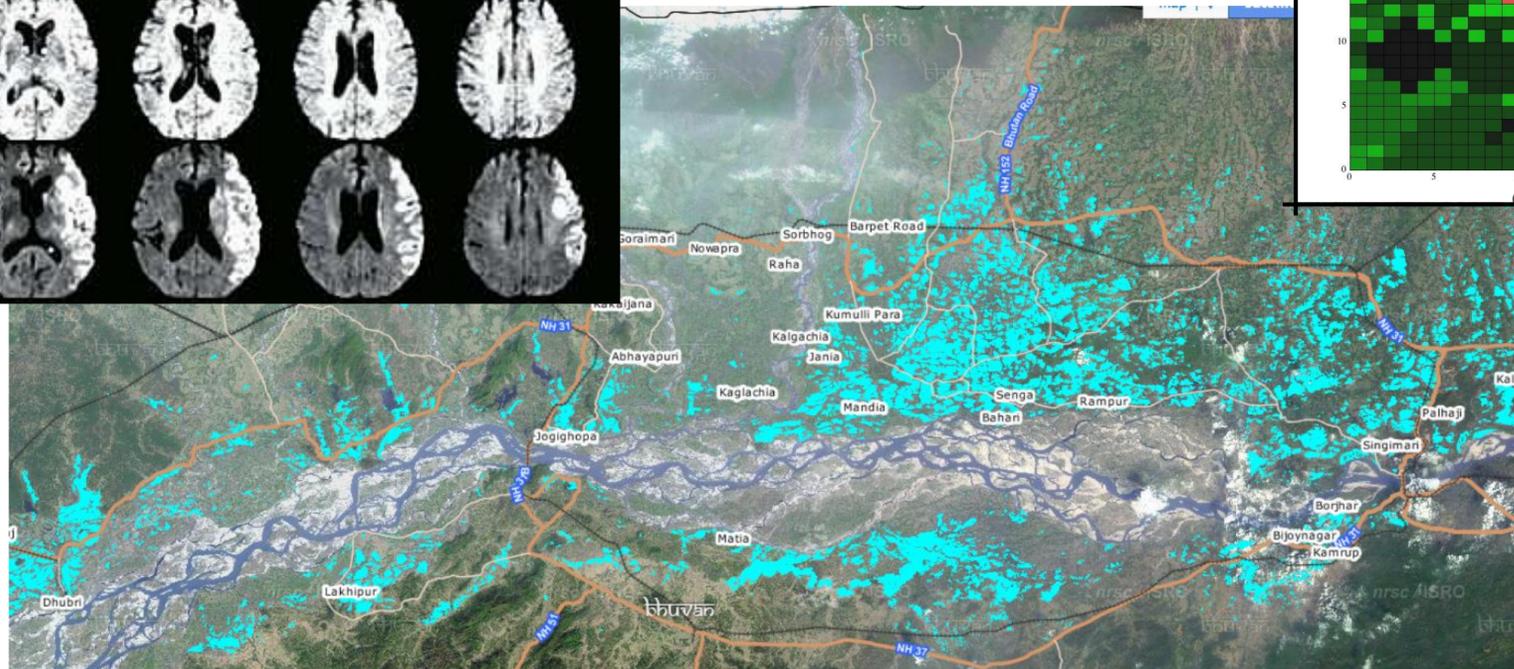
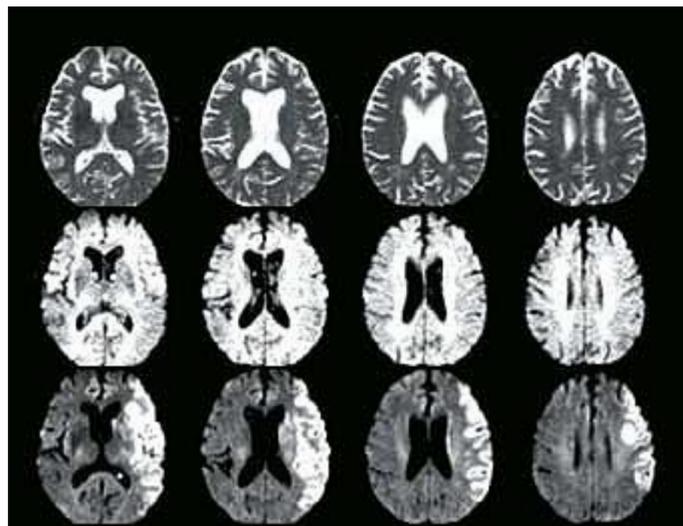
Sustainable Forest Management



Predicting and Preventing Forest Wildfires



Medical Imaging



Flood Prediction

Some AI Research on Forest Fire Domain

- **Fire Suppression:**

- Letting Fires Burn: given a fire that is occurring, decide to let it burn or not
- Trade off between suppression costs, cost of large fires, benefit of free fuel reduction on reducing future fires.

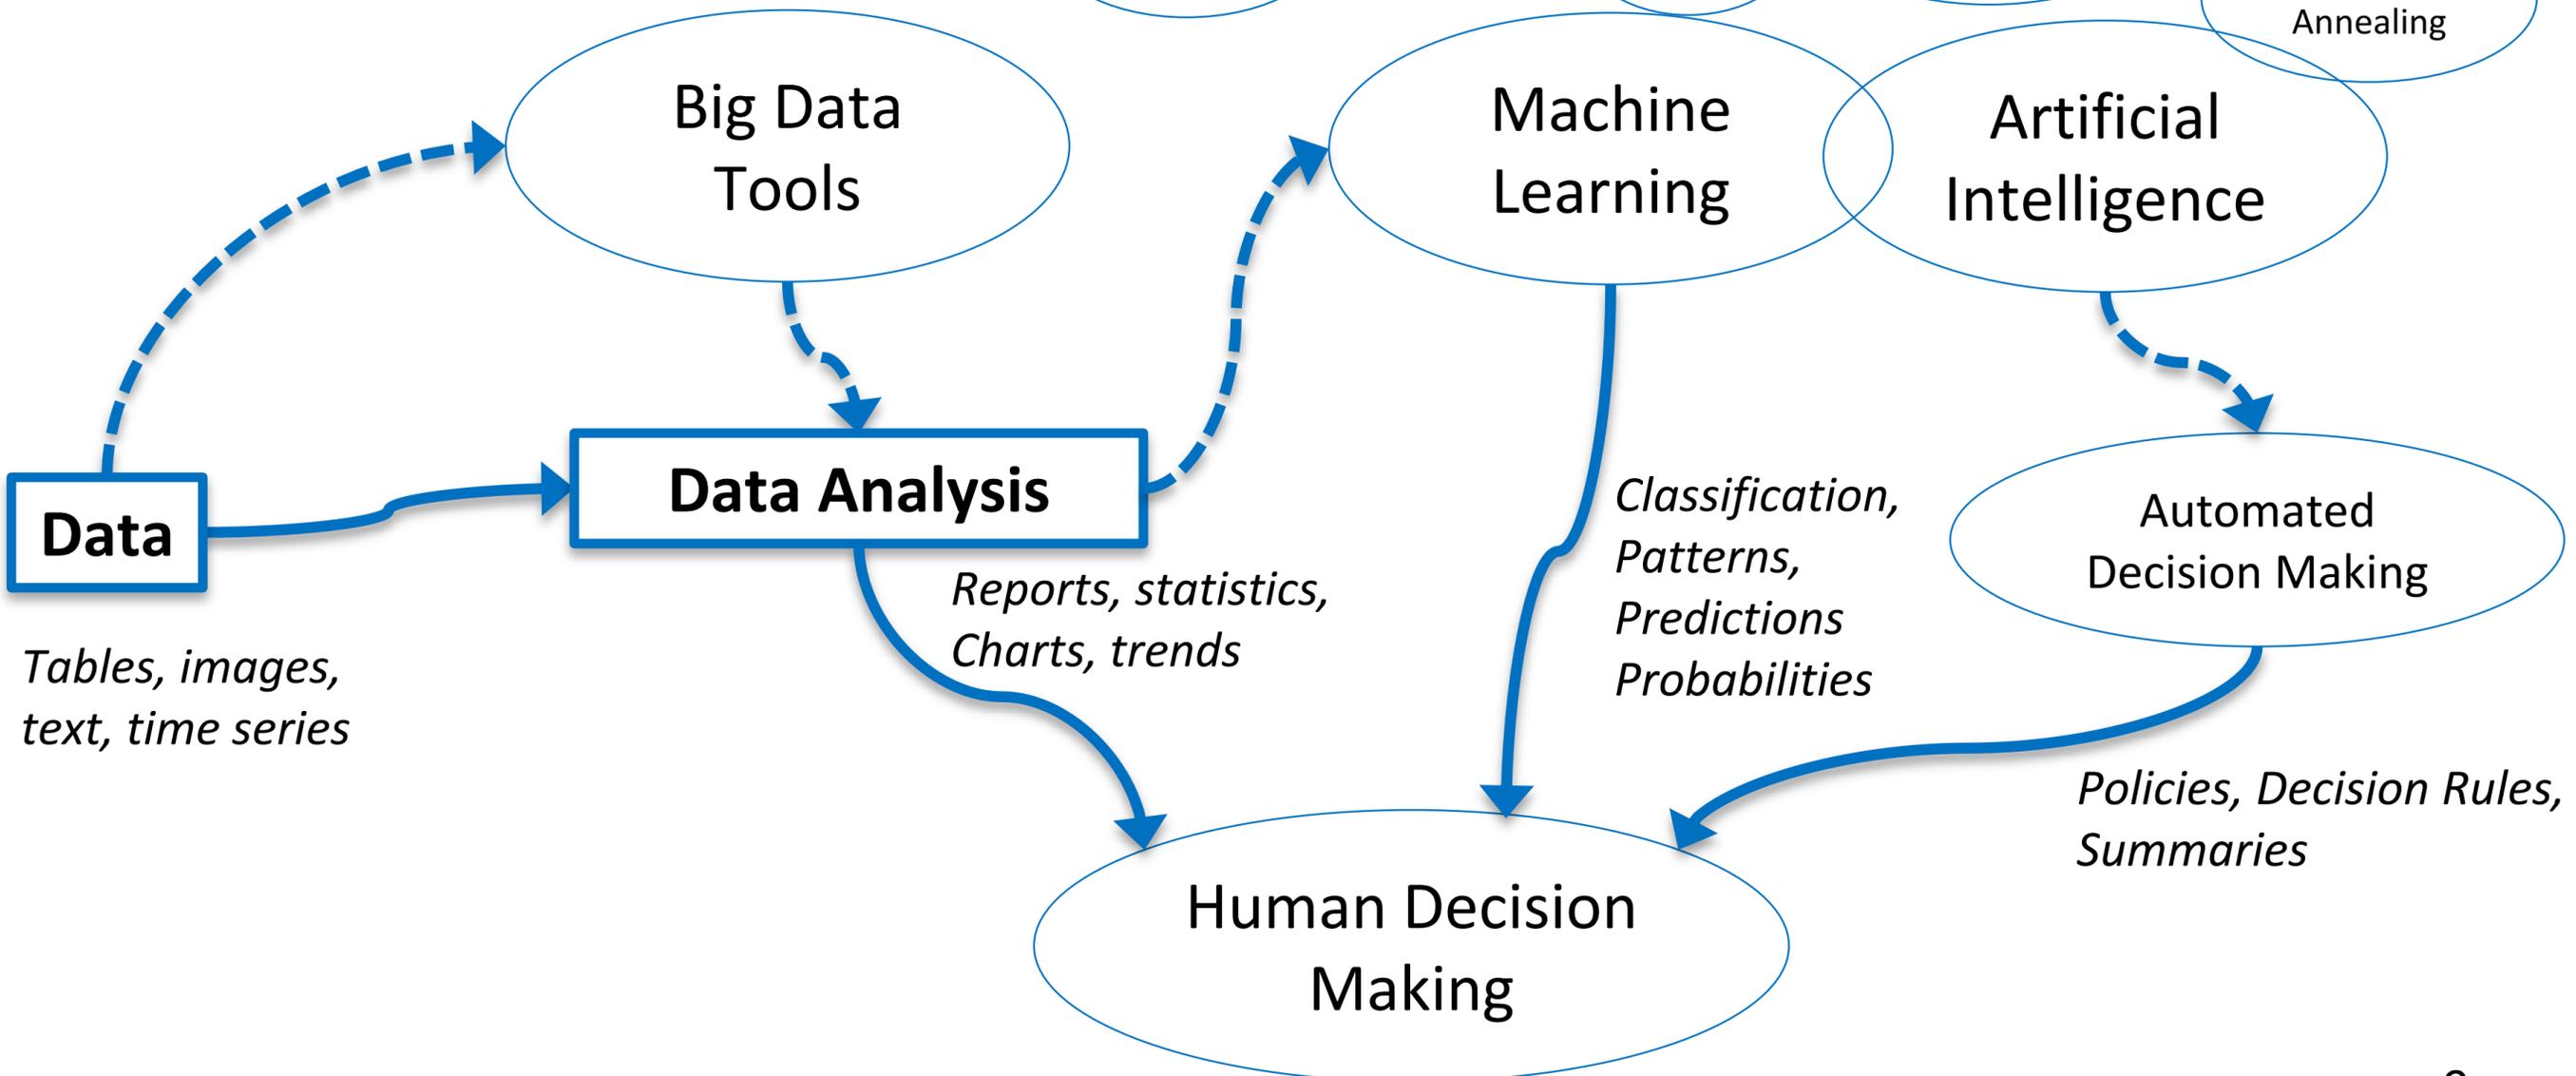
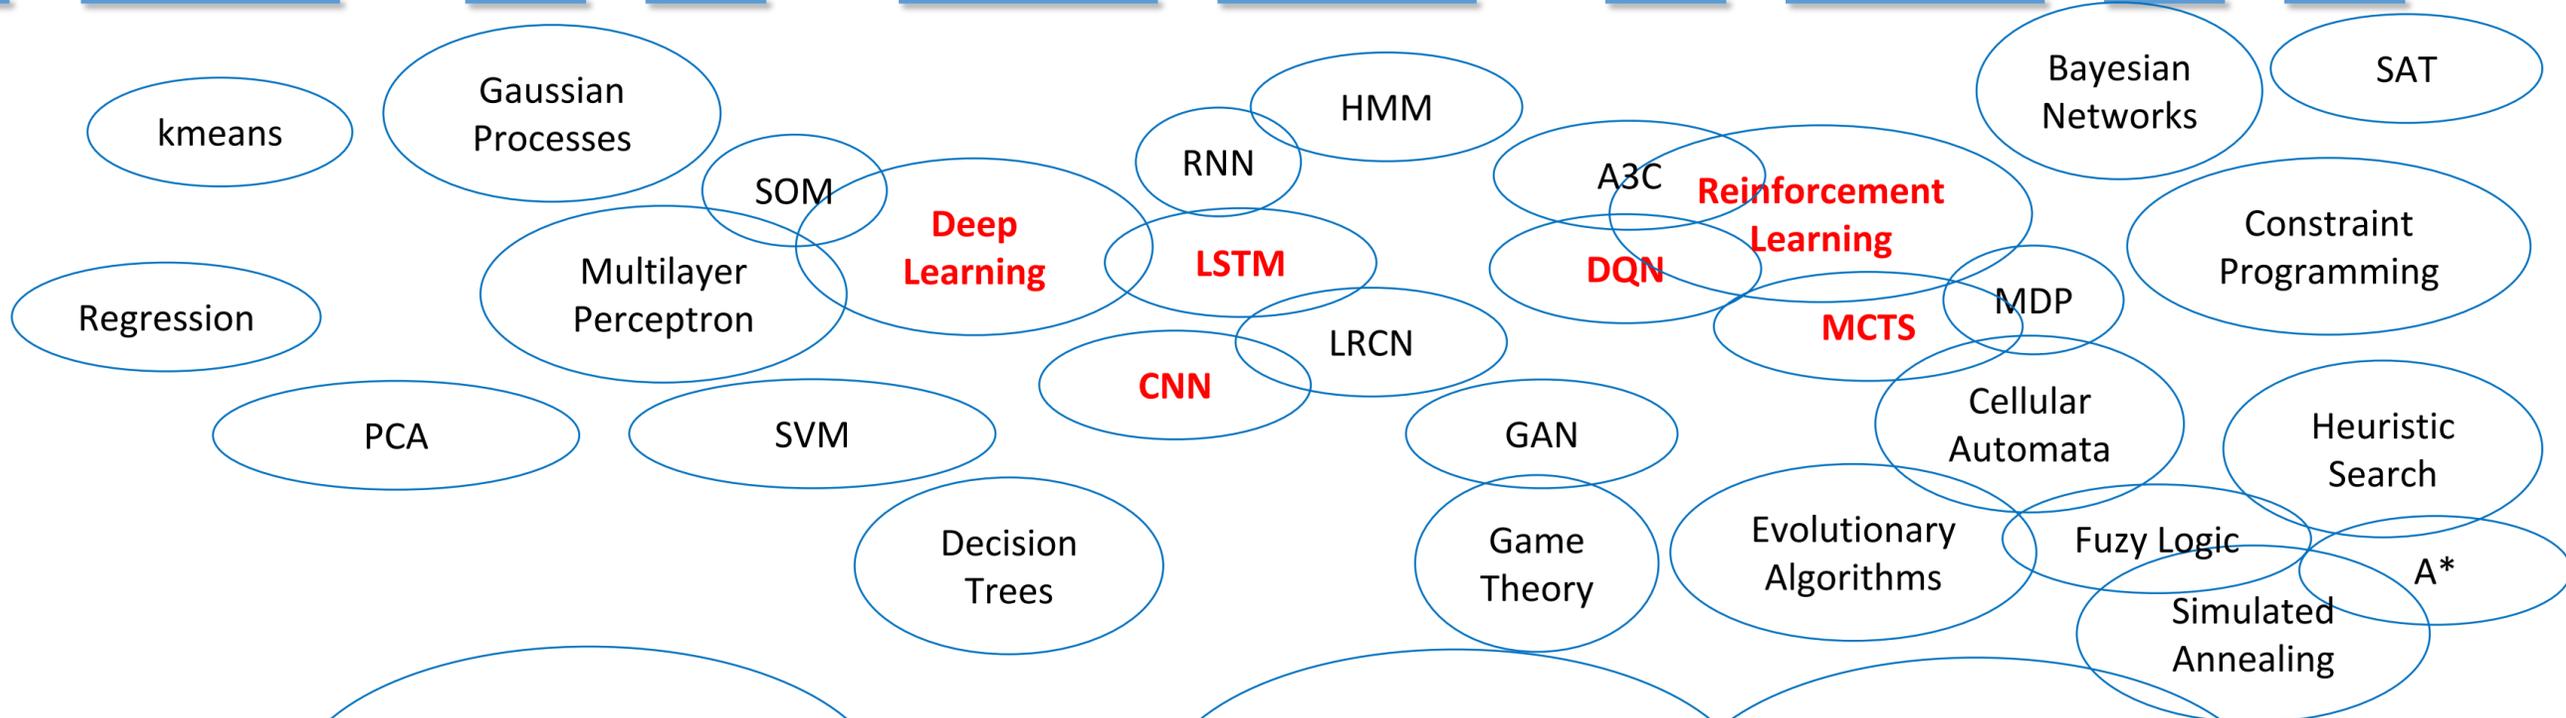
- **Fire Spread Dynamics:**

- Imagine fire is the agent on the landscape, each spread from cell to cell is an action.
- Learn a fire spread policy from data

- **Fuel Treatment:**

- find optimal treatment of fuels over time to reduce expected cost of catastrophic fires.

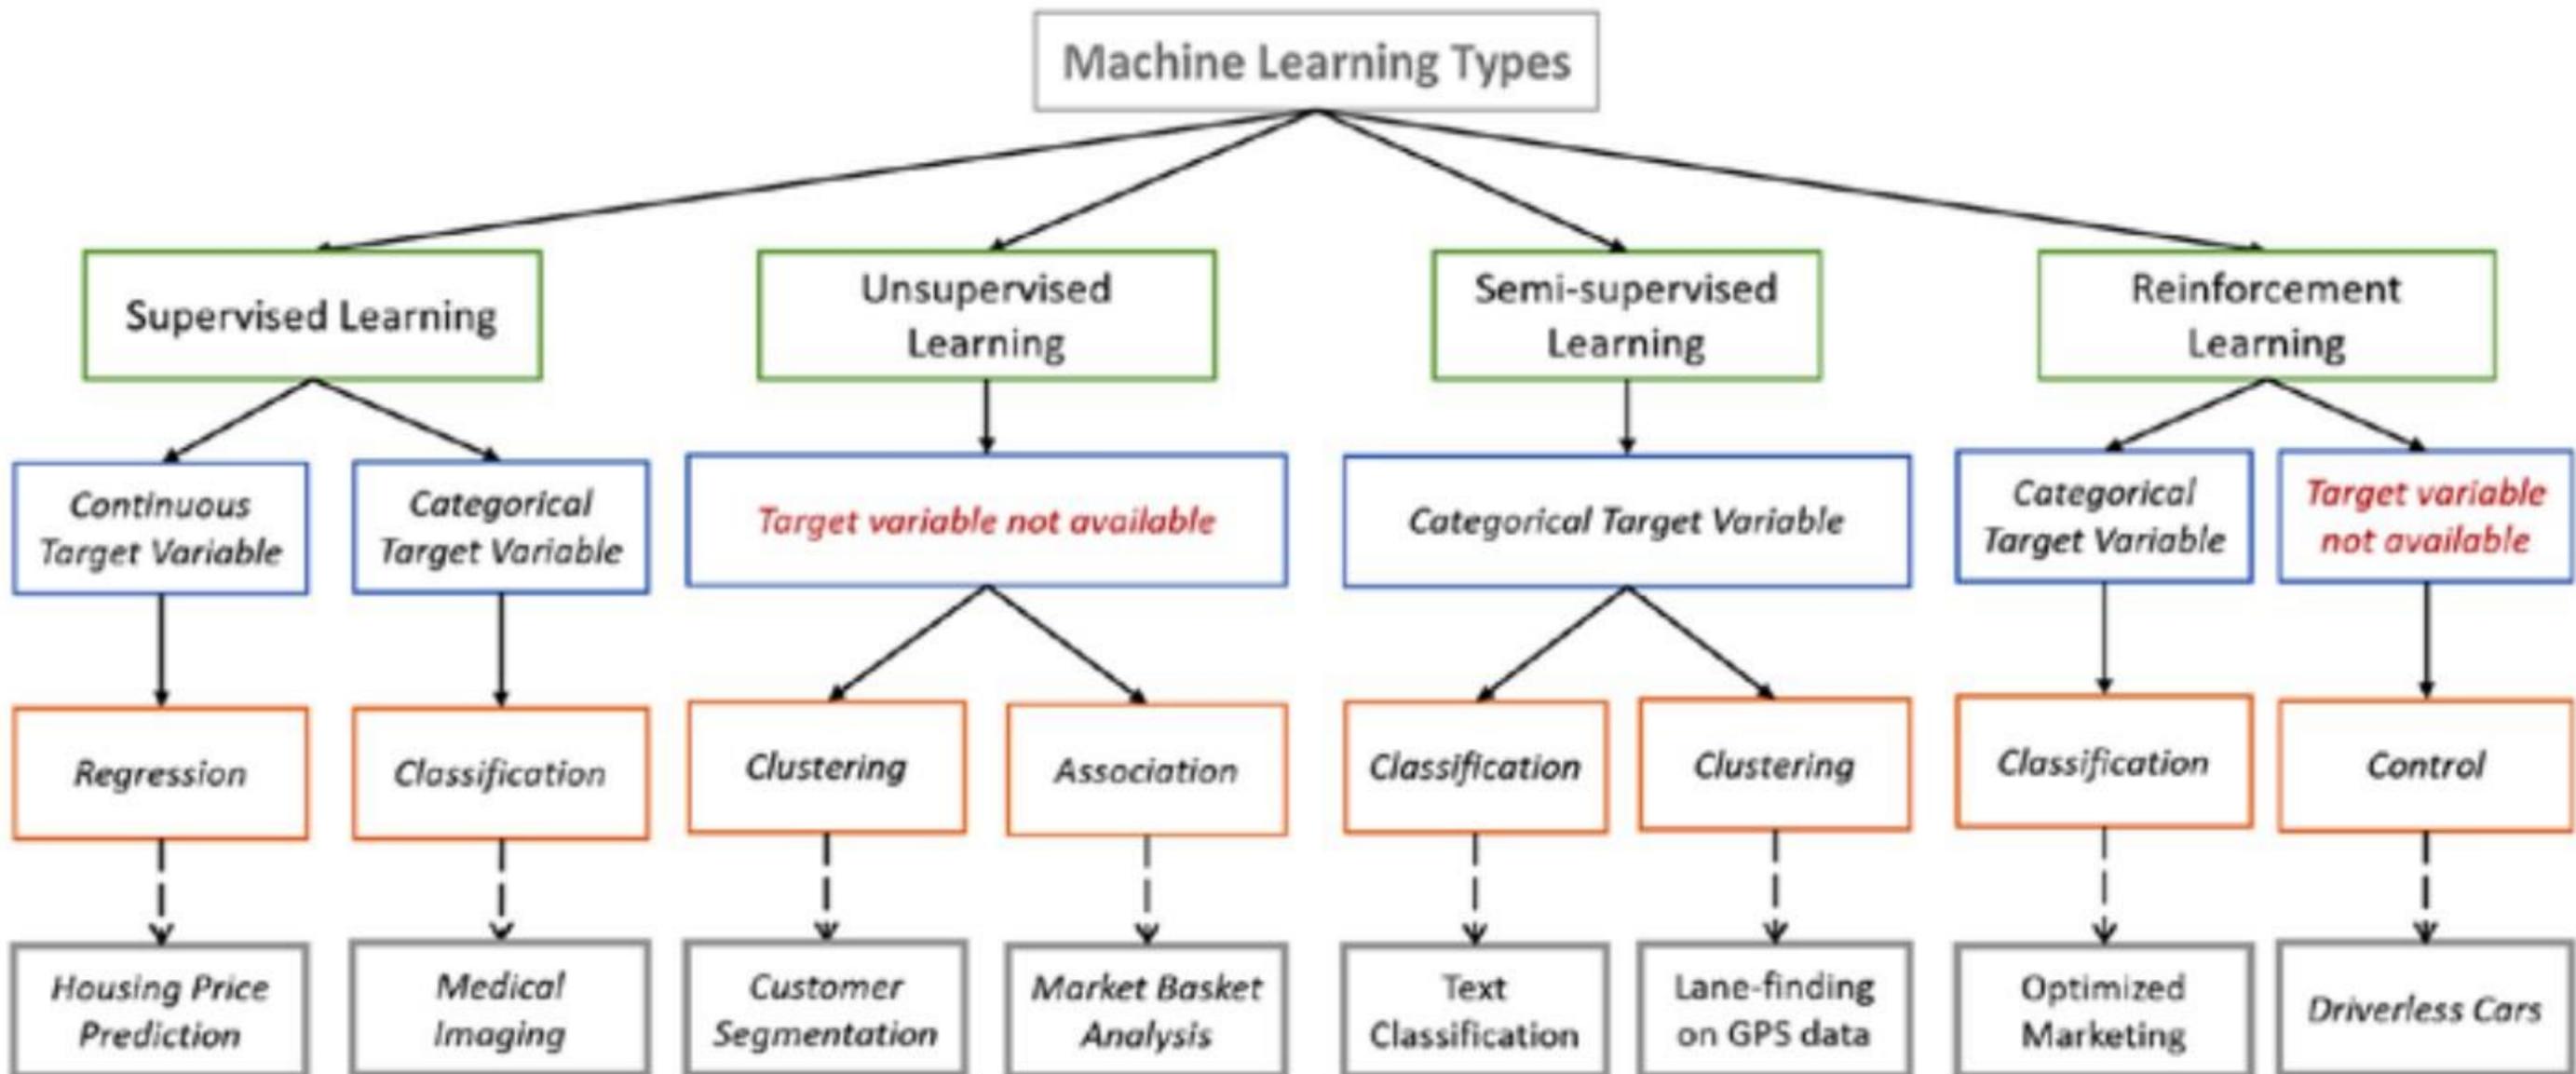
Data, Big Data, Machine Learning, AI, etc, etc,...



Major Types of Machine Learning

"Detect patterns in data, use the uncovered patterns to predict future data or other outcomes of interest"

– Kevin Murphy, "Machine Learning: A Probabilistic Perspective", 2012



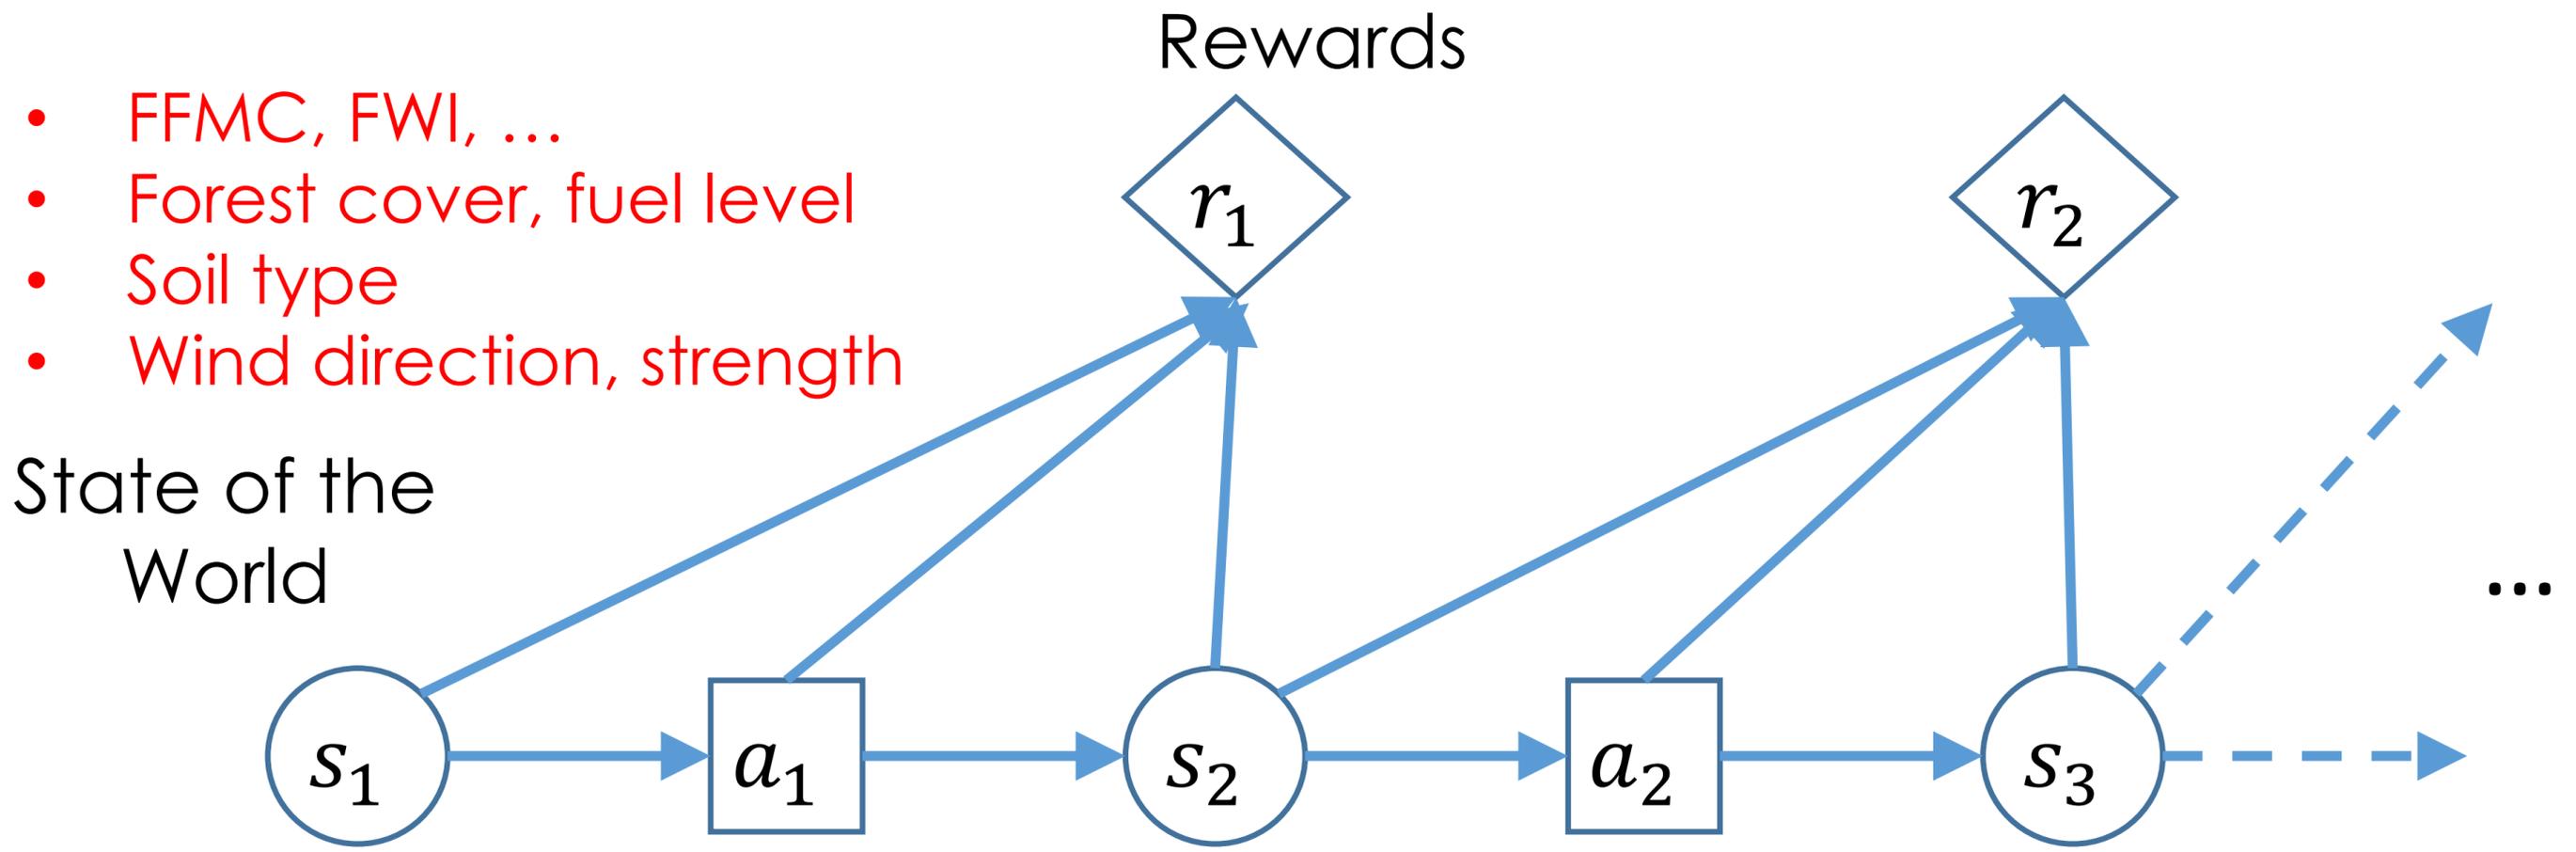
Defining Your Objectives

- ▶ In Regression/Prediction/Classification
 - ▶ Objective function is often *low error, high accuracy, high recall*
 - ▶ Hard if you don't know the true answer
- ▶ In Decision Making we need something else:
 - ▶ **Value function:** Expression of what situations or properties are preferred or required
 - ▶ Can be formulated as **Costs/Rewards** of taking particular actions
- ▶ Then we try to search for, or directly compute, a **policy which maximizes those values**
- ▶ What does that look like?

Markov Decision Process (MDP)

- FFMC, FWI, ...
- Forest cover, fuel level
- Soil type
- Wind direction, strength

- suppression cost
- damage, area burned



Actions

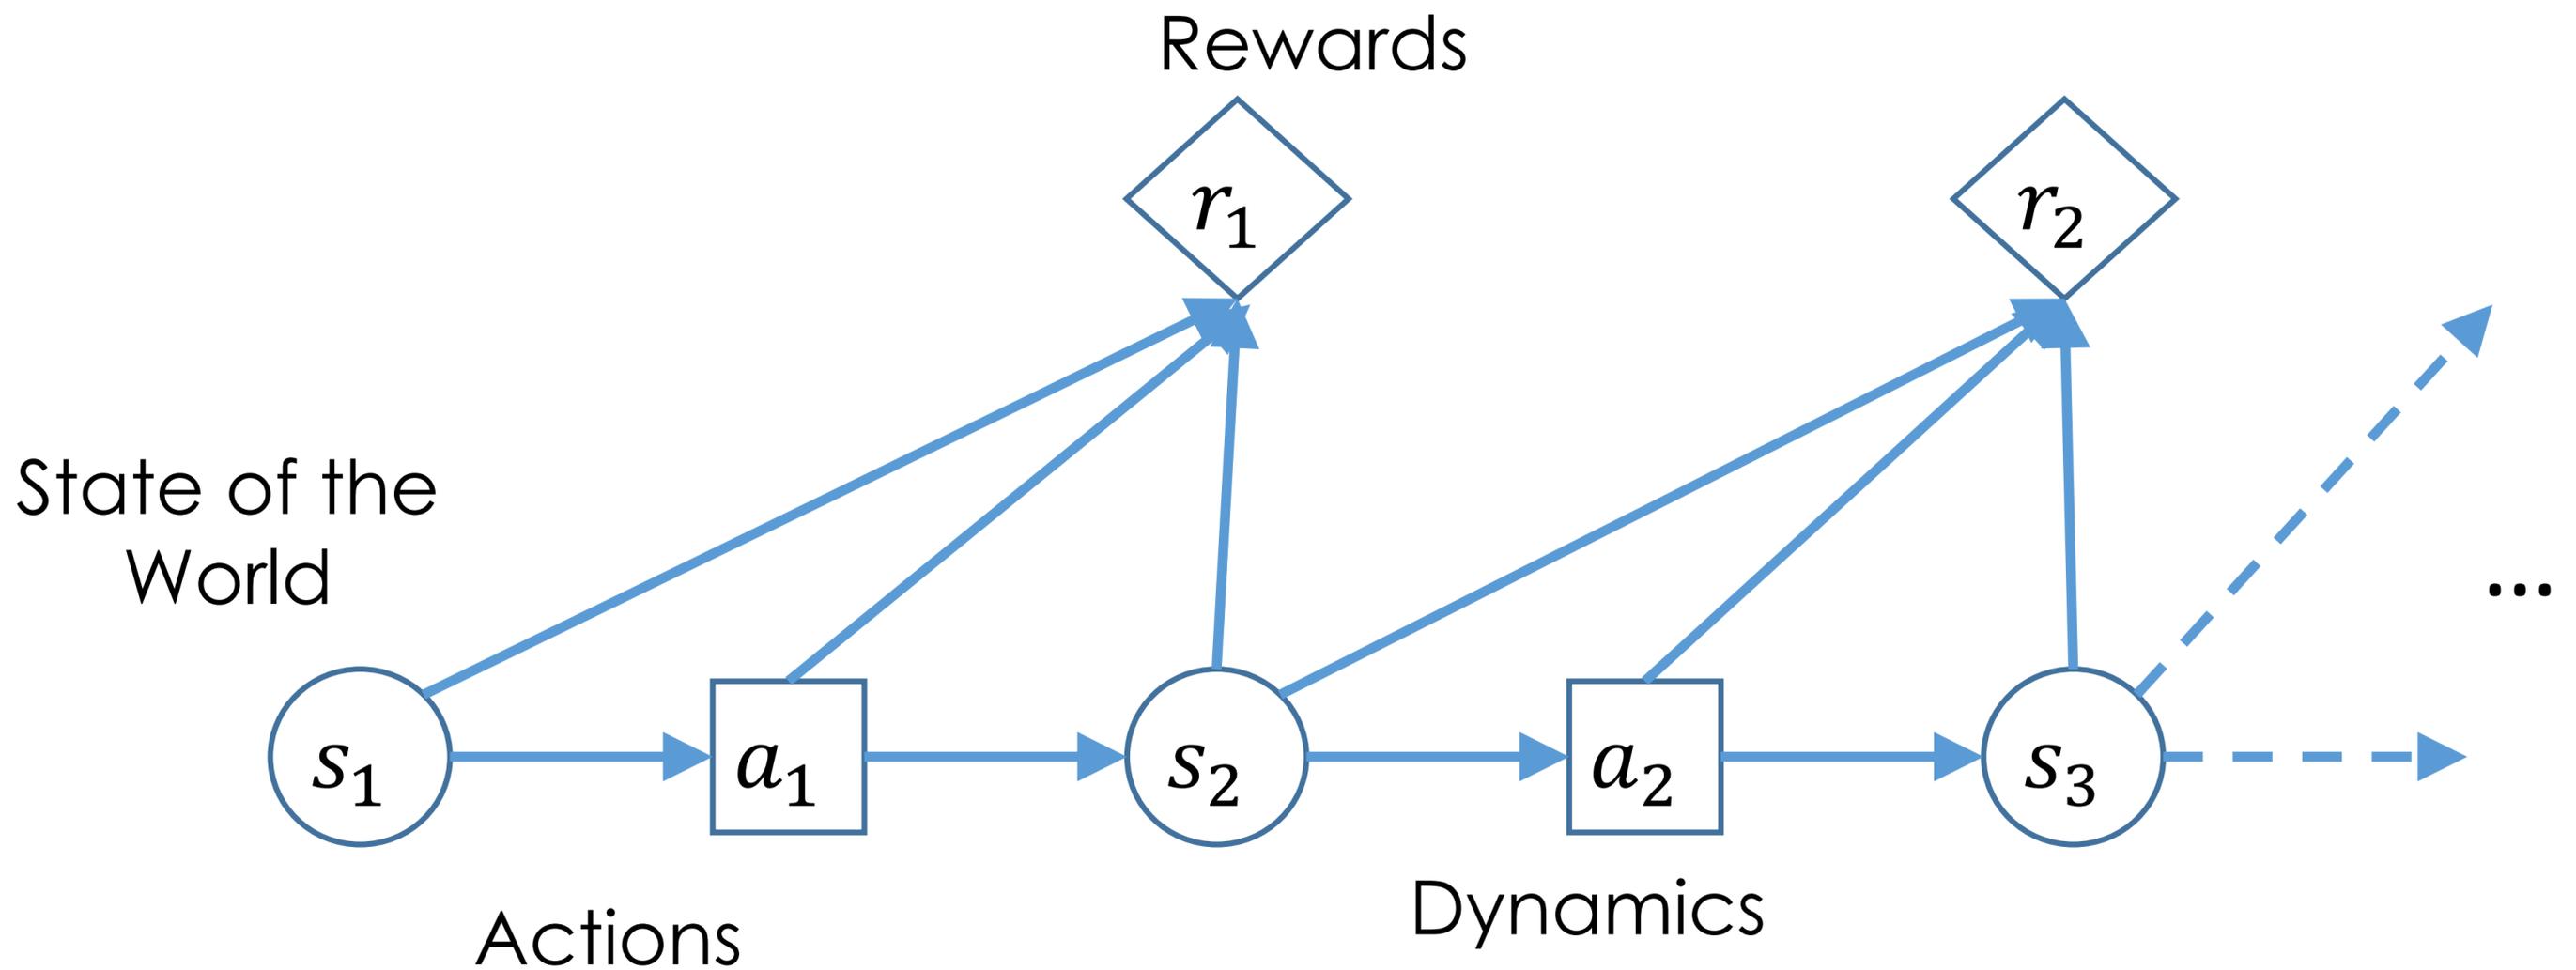
- Let Burn/Suppress
- Fireline (x,y,shape,...)

Dynamics

$$T(s_1, a_1) \rightarrow \text{distribution}(s_2)$$

- Deterministic or stochastic
- Fire spread
- Upcoming Weather
- Fuel/soil cycles
- Human behaviour

Markov Decision Process (MDP)



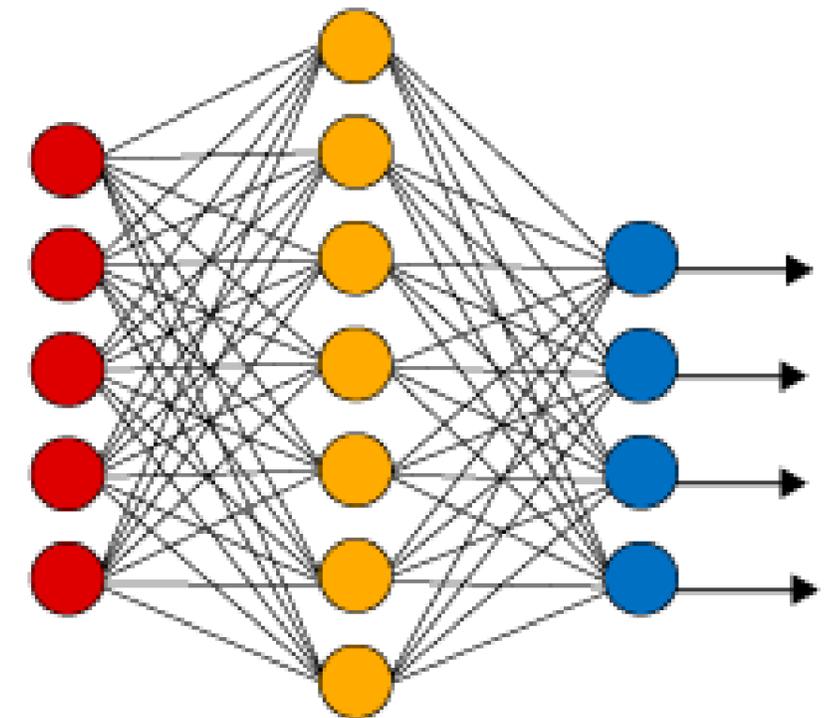
Which part of this picture do you know?
Which can you estimate?
Which do you *need* to know?

Reinforcement Learning as an MDP

- Reinforcement Learning is learning the policy $\pi(a|s)$ for an MDP when you **do not have access** to the full definition of:
 - the rewards $r(s^t, a^t, s^{t+1})$
 - AND/OR the dynamics $T(s^t, a^t, s^{t+1})$
- Training must be carried out interactively :
 1. Commit to action using latest (or some) policy
 2. Find out the next state and reward from the world/simulator/environment
 3. Improve your policy
 4. Repeat
- A.K.A :
 - Reinforcement Learning [Sutton and Barto 1998]
 - In some fields **Approximate Dynamic Programming** (ADP) [Powell 2007, 2009] is discussed which is essentially the same as RL

Basic Three Layer Neural Network

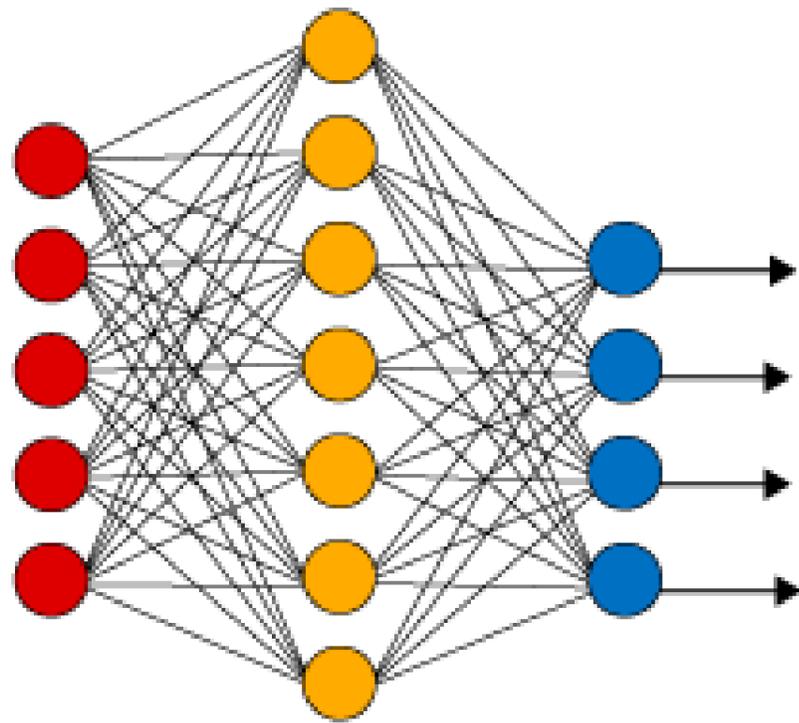
- Input Layer
 - vector data, each input collects **one feature**/dimension of the data and passes it on to the (first) hidden layer.
 - **Measurements, aggregate stats, computed indices, pixels**
- Hidden Layer
 - Each hidden unit computes a weighted sum of all the units from the input layer
- Output Layer
 - Each output unit computes a weighted sum of all the hidden units and passes it through a **threshold function**.
 - **Target variables, predictions, class labels, images**



input hidden output

So What is Deep Learning?

Simple Neural Network

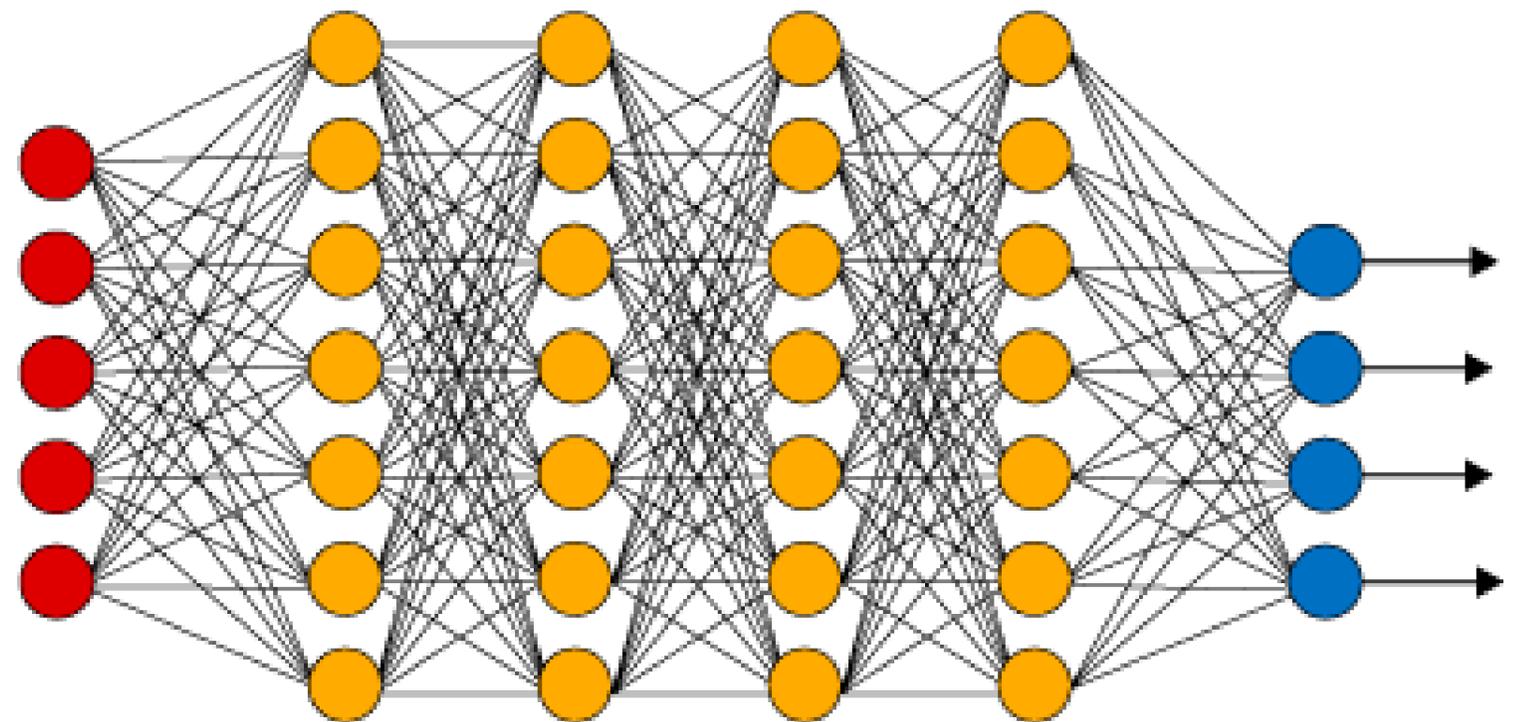


● Input Layer

● Hidden Layer

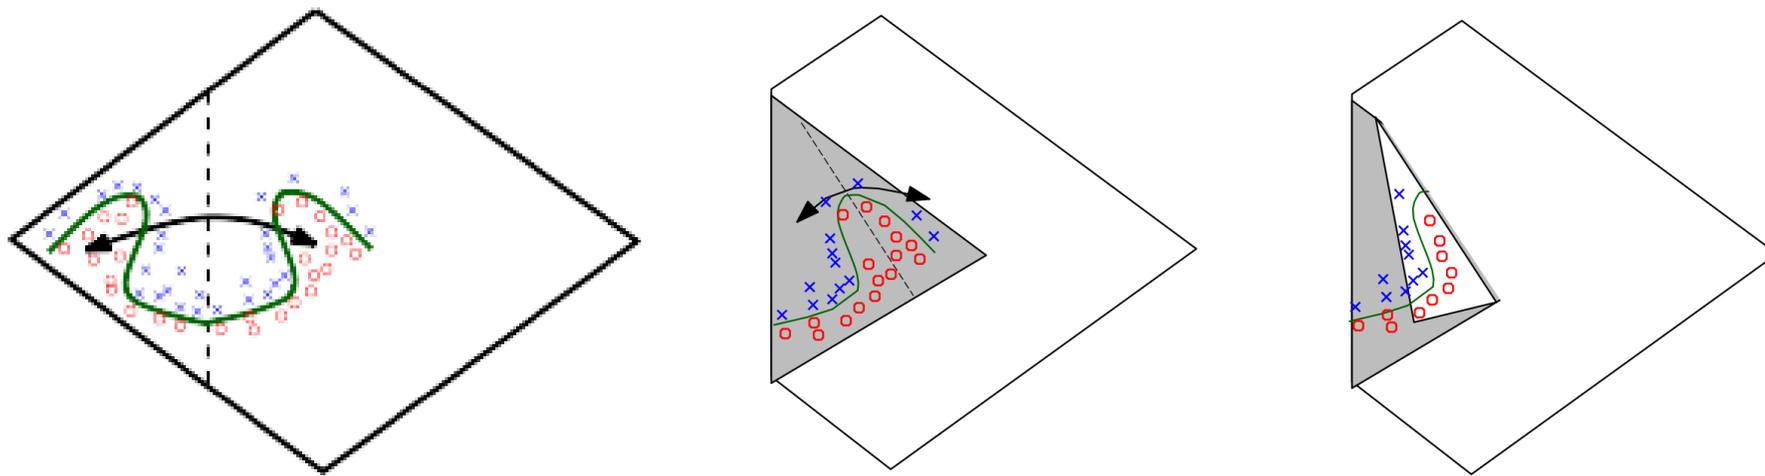
● Output Layer

Deep Learning Neural Network



- Hackernoon

Why Go Deep?



(Goodfellow 2016)

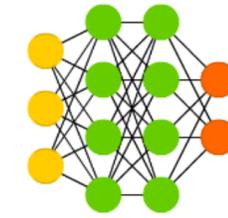
Using Hidden units, each hidden layer increases power, exponential advantage of additional layers.

Neural Networks

©2016 Fjodor van Veen - asimovinstitute.org

-  Backfed Input Cell
-  Input Cell
-  Noisy Input Cell
-  Hidden Cell
-  Probabilistic Hidden Cell
-  Spiking Hidden Cell
-  Output Cell
-  Match Input Output Cell
-  Recurrent Cell
-  Memory Cell
-  Different Memory Cell
-  Kernel
-  Convolution or Pool

Deep Feed Forward (DFF)



Perceptron (P)



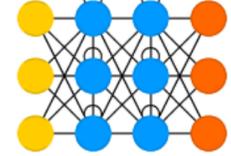
Feed Forward (FF)



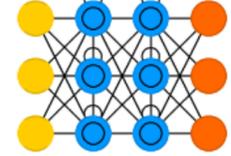
Radial Basis Network (RBF)



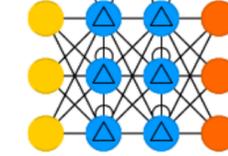
Recurrent Neural Network (RNN)



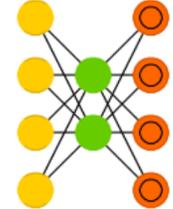
Long / Short Term Memory (LSTM)



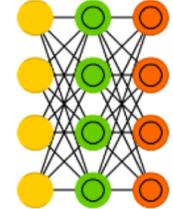
Gated Recurrent Unit (GRU)



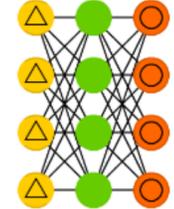
Auto Encoder (AE)



Variational AE (VAE)



Denosing AE (DAE)



Sparse AE (SAE)



Markov Chain (MC)



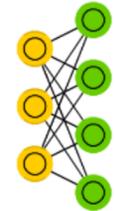
Hopfield Network (HN)



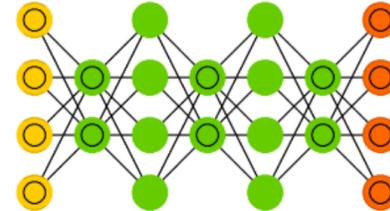
Boltzmann Machine (BM)



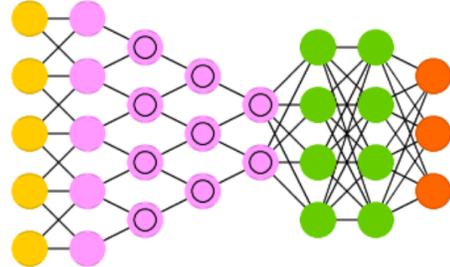
Restricted BM (RBM)



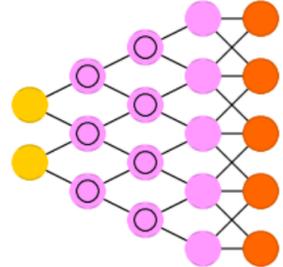
Deep Belief Network (DBN)



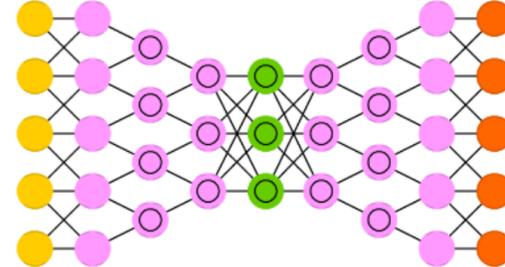
Deep Convolutional Network (DCN)



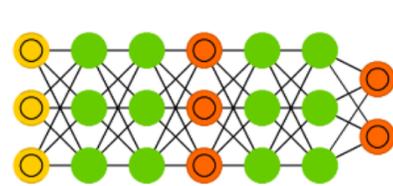
Deconvolutional Network (DN)



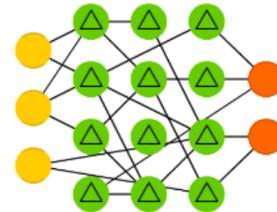
Deep Convolutional Inverse Graphics Network (DCIGN)



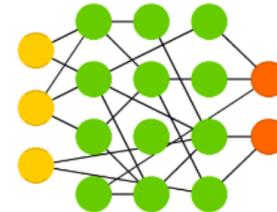
Generative Adversarial Network (GAN)



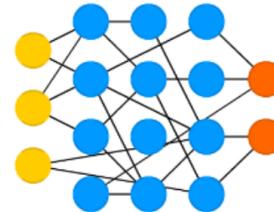
Liquid State Machine (LSM)



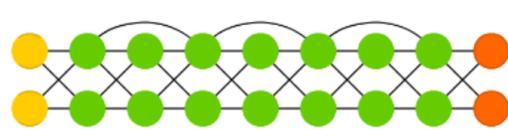
Extreme Learning Machine (ELM)



Echo State Network (ESN)



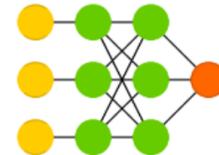
Deep Residual Network (DRN)



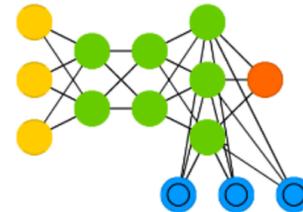
Kohonen Network (KN)



Support Vector Machine (SVM)



Neural Turing Machine (NTM)

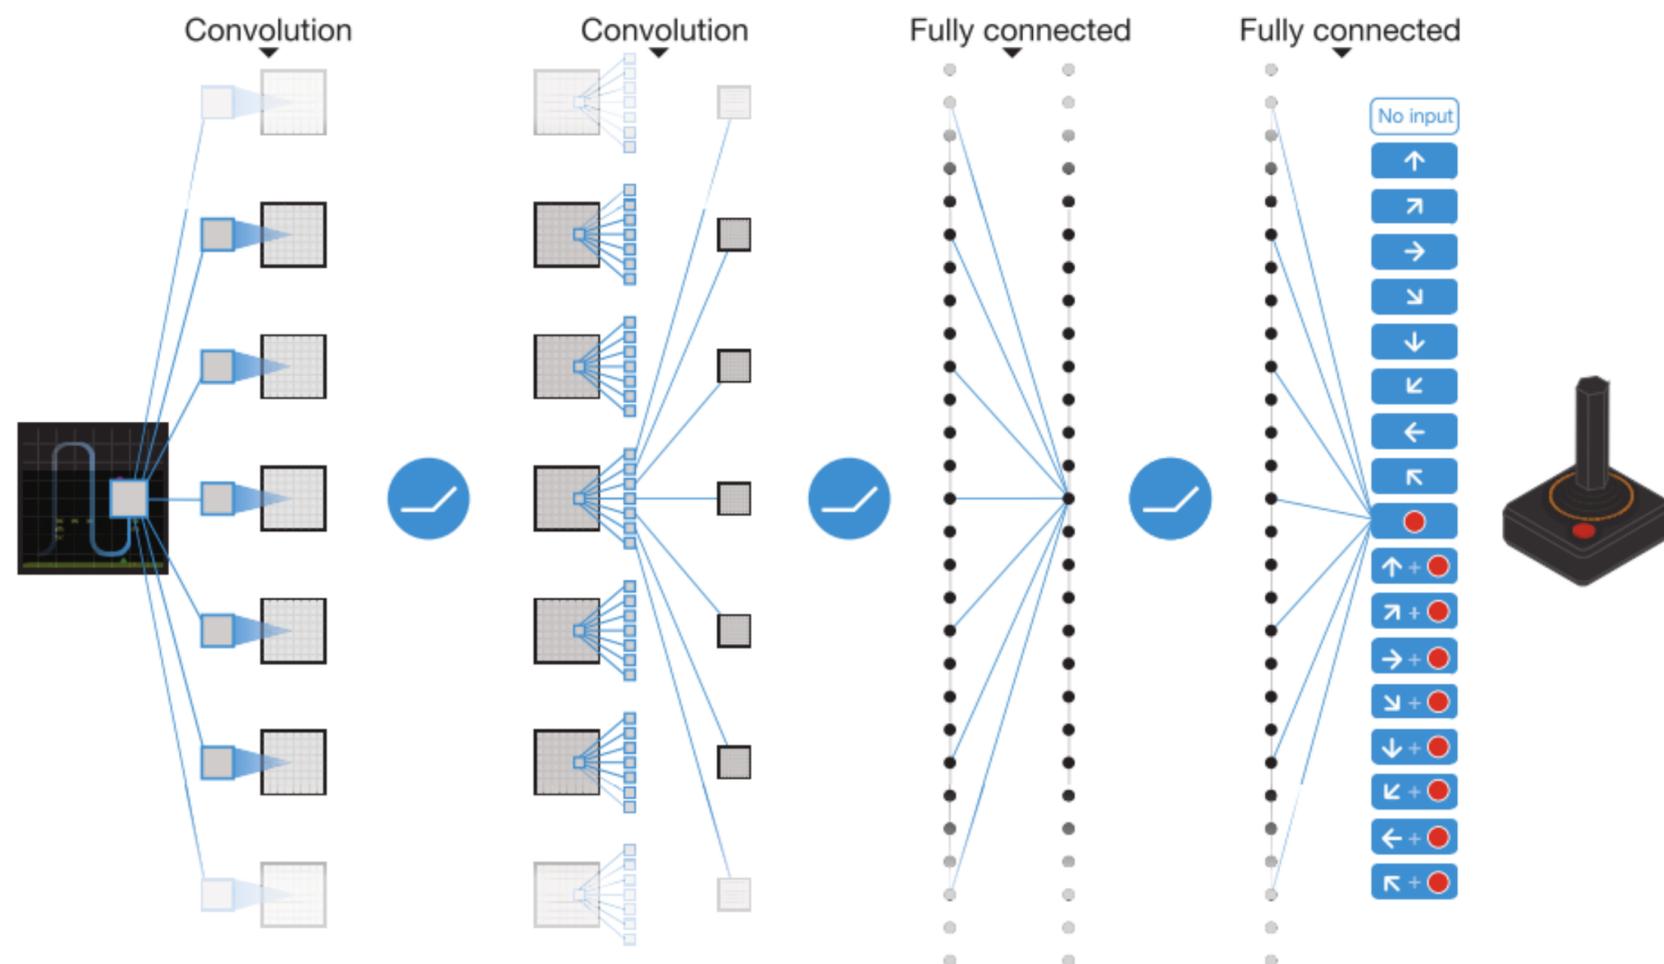


Some Types of Deep Neural Networks

- RBM: Restricted Boltzman Machines (RBM) - bidirectional deep models. Older method but still useful for building consistent belief models.
- DeepRL: Deep Reinforcement Learning
- RNN: Recurrent Neural Networks (RNN) - allow links from outputs back to inputs, over time, good for time series learning
- LSTM: Long-Term Short-Term networks – a particular way of learning from time series data by controlling forgetting
- LRCN: Long-term Recurrent Convolutional Networks
- CNN: Convolutional Neural Networks – More efficient structure when input data are images
- GAN: General Adversarial Networks - train two networks at once in competition to improve robustness

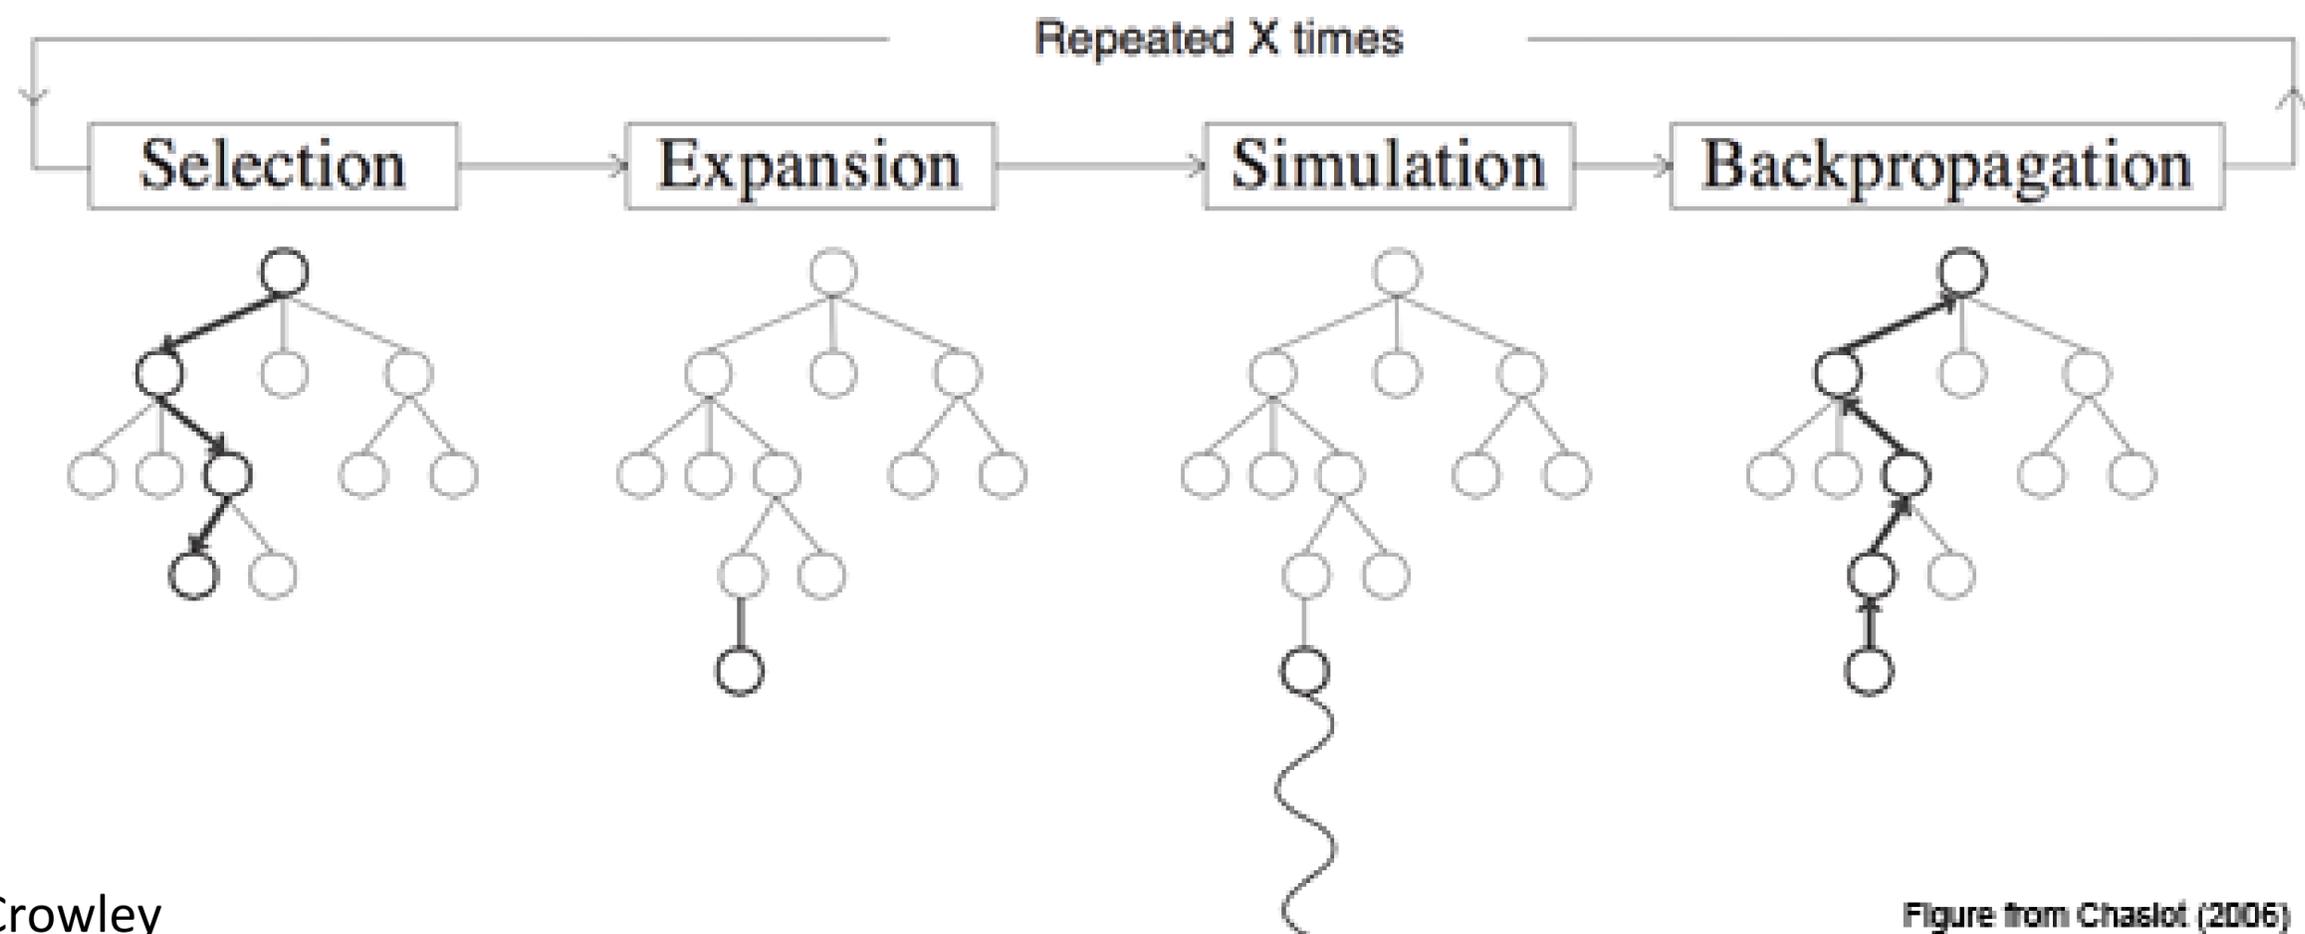
Deep Learning, Deep RL, DQN, A3C, etc...

- Can also use **CNNs + Fully Connected Deep Network** for learning a representation of a policy
- Flurry of advances since 2014 by Google DeepMind and others applying Deep Learning to RL algorithms
- Deep Q-Learning – **DQN**
- Asynchronous Advantage Actor-Critic - **A3C**
- Many algorithms since then trying to provide a better way to learn the value function with DNNs
 - Alpha Go – RL + human training
 - Alpha Zero – RL + MCTS search + playing itself (Go, Chess)



Monte-Carlo Tree Search

- Optimizing a policy using tree search using simulator to try new sets of actions and resulting states
- Available when you don't know:
 - the transition dynamics or instantaneous rewards/values
- But when you do know:
 - The available actions at any moment and
 - The final “winning/losing” conditions



Prediction of Spatially Spreading Processes

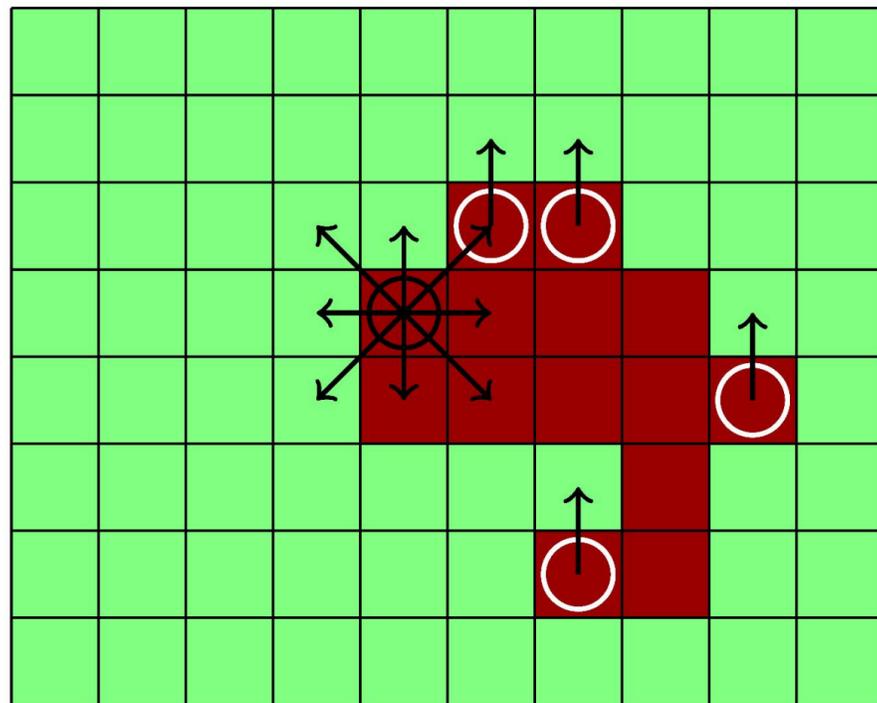
Using Forest Wildfires as a demonstration domain:

- **RLDM 2017** - The idea of using RL for learning dynamics from image data, comparison of classical RL algorithms with DQN [Subramanian and Crowley (2017)]
- **Frontiers in ICT Journal 2018** - Compared to broader range of RL methods Also looked at Gaussian Processes for a fully supervised comparison Noticed a tradeoff between MCTS and A3C [Ganapathi Subramanian and Crowley (2018)]
- **CAI Paper 2018** - proposed a new algorithm MCTS-A3C to take advantage of strengths of both.

Goal: Demonstrate that learning an agent-based style model be done from raw data via RL, and can provide comparable results to other methods.

Problem setup

- Using satellite images from two large forest wildfires in Northern Alberta : Richardson 2011, For McMurray 2016.
- Used just publically accessible low rez images from USGS, (Landsat)



(a) Schematic of the state and actions



(b) Raw Color Image

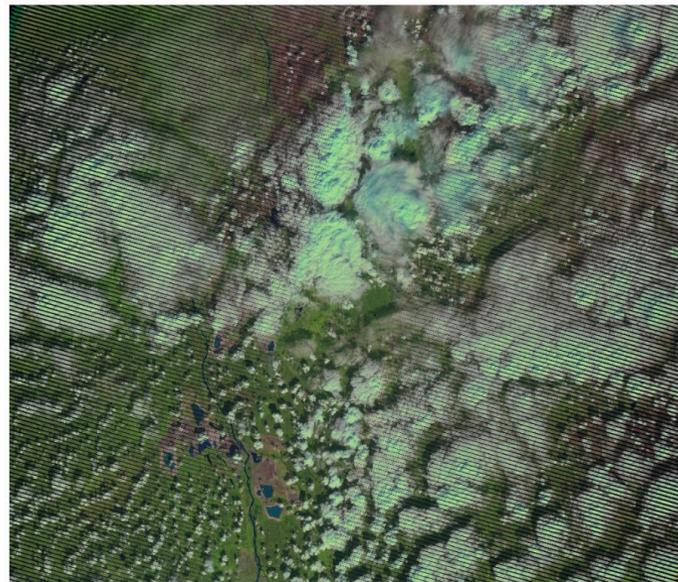


(c) Thermal Image

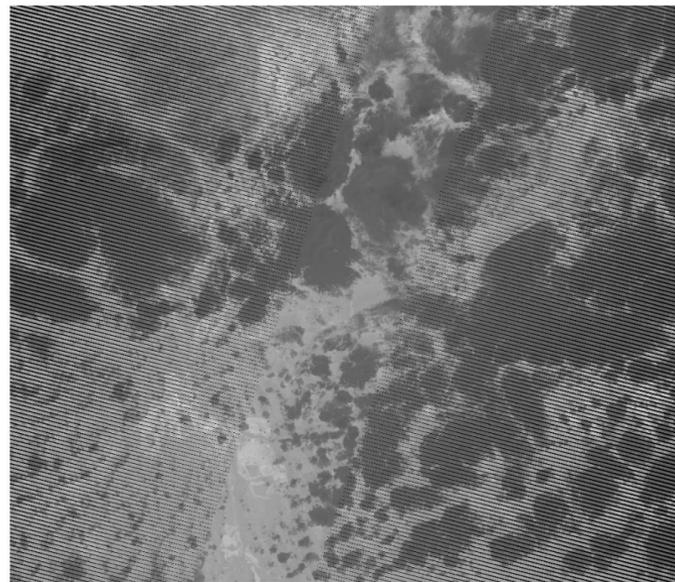
RL algorithms considered

- **DQN**: Neural Networks are used to represent the Q values in this algorithm as opposed to simpler tabular representation of the Q values [Mnih et al. (2013)].
- **DQN with PER**: Modifies DQN by targeting maximum learnable experiences [Schaul et al. (2016)].
- **A3C**: A global network interacts with group of worker agents who have their own environment and network parameters. Both a value estimate and a policy is used. Value estimate is used to update the policy [Mnih et al. (2016)].
- **MCTS**: Monte Carlo Tree Search. Hypothetical simulation rollouts from any state to compute reward estimate. Involves Selection, Expansion, Simulation and Back Propagation [Browne et al. (2012)].

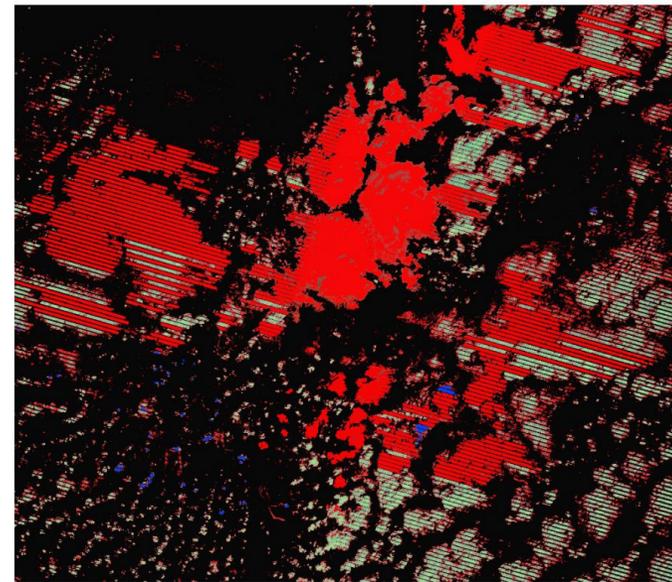
Results



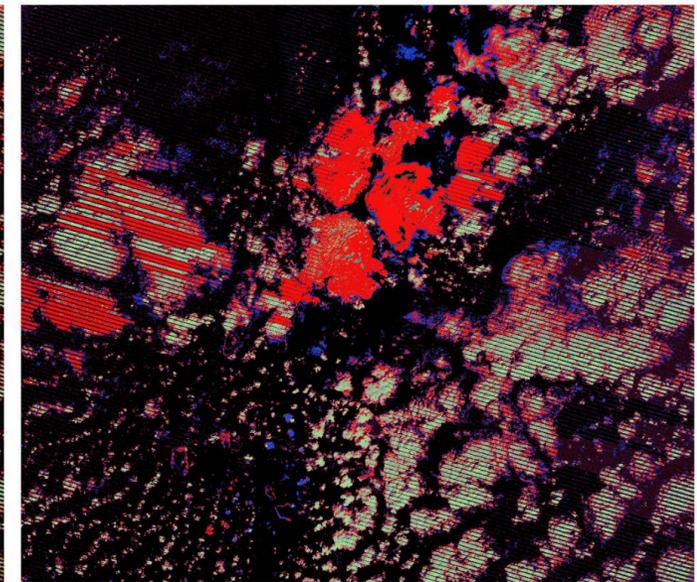
(a) Satellite Image of August 11



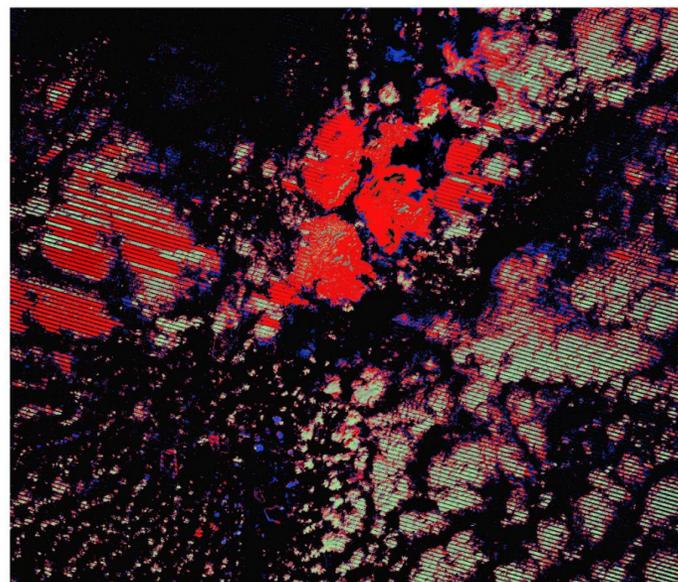
(b) Thermal Image of August 11



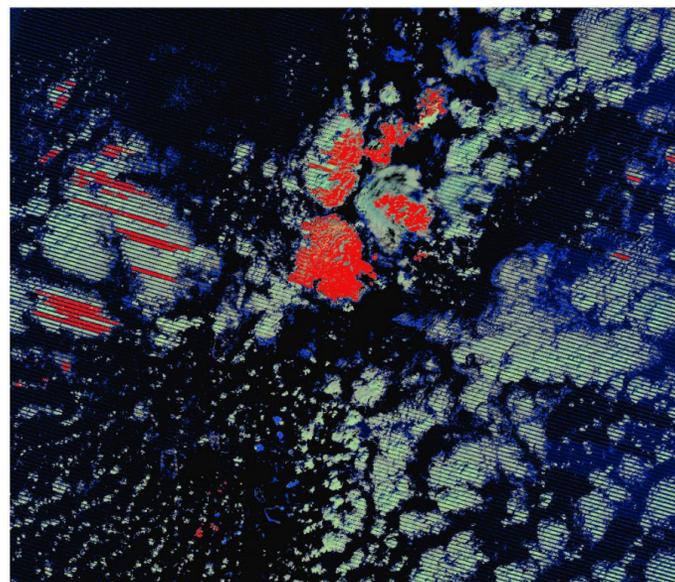
(c) Gaussian Process



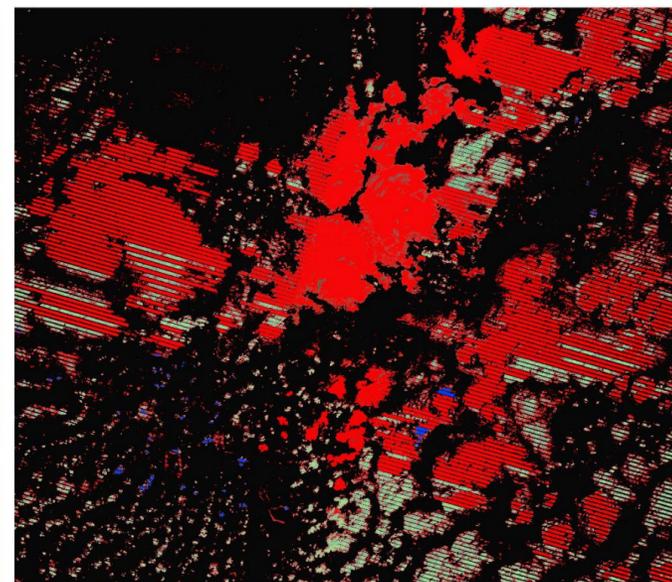
(d) Value Iteration



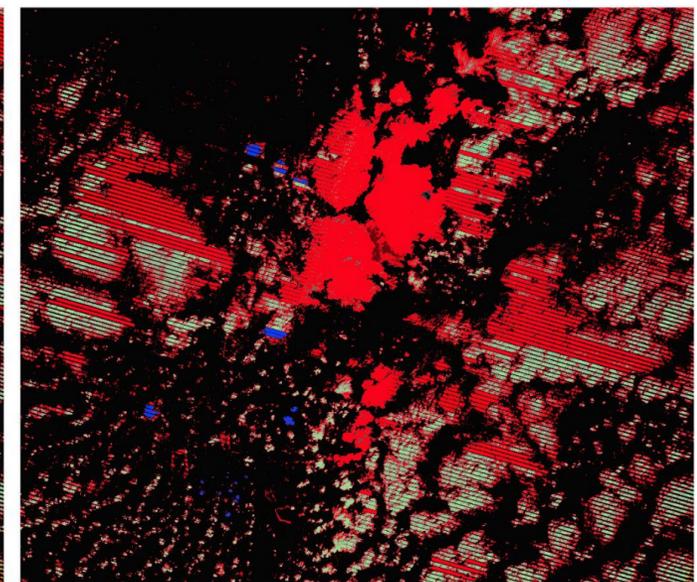
(e) Policy Iteration



(f) Q Learning



(g) MCTS



(h) A3C

Results

Method	(C)	(D)	(E)	(F)
GP	60.5%	47.9%	45.3%	20.5 %
V.I	88.5%	68.4%	30.1%	6.4%
P.I	89.3%	67.8%	35.8%	8.9%
Q.L	84.2%	61.4%	26.4%	5.3%
MCTS	65.3%	55.7%	49.7%	5.8%
A3C	90.1%	81.8%	50.8%	13.4%

Table 3: Average Accuracy of each algorithm trained on the Richardson Fire but applied on the Fort McMurray fire for different time durations.

Farsite Forest Fire Simulator Platform

While I was at OSU we created a forest wildfire simulation framework:

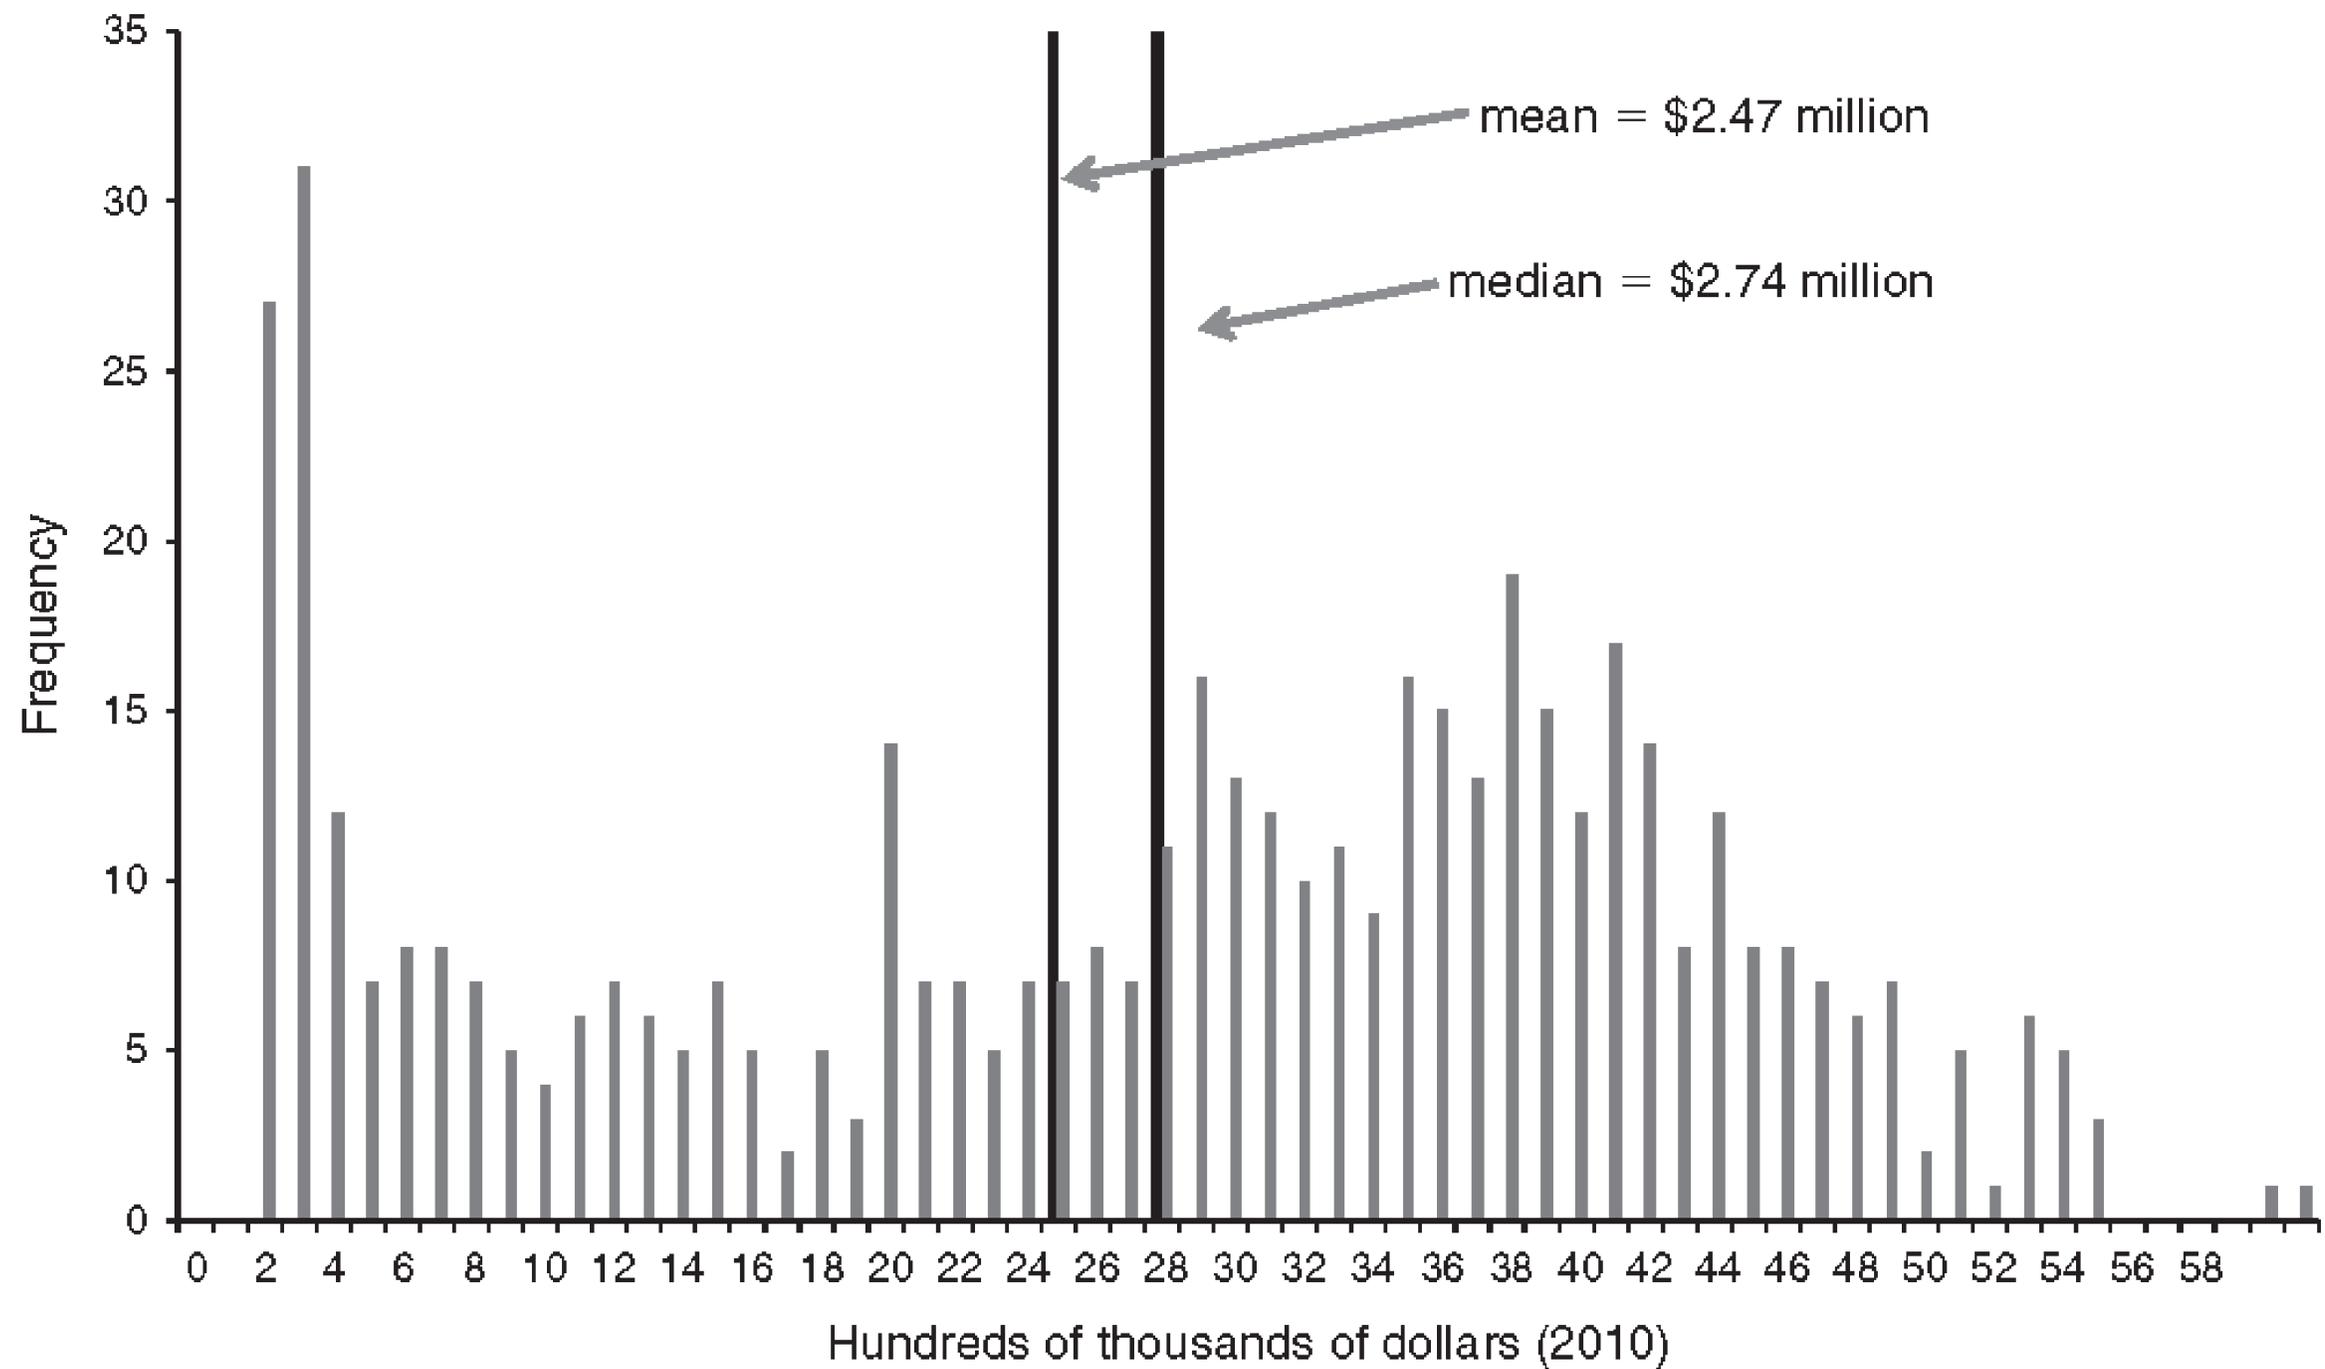
- Utilized the **Farsite** fire spread simulator (used by US Forest Service, similar to Prometheus)[Finney, 1998]
- Used a simple model of the spatial distribution of lightning strikes (based on historical data)
- Added a fire duration model [Finney et al.,2009]
- Combined with a high-resolution FVS forest growth simulator [Dixon, 2002].
- Weather simulated by resampling from the historical weather time series observed at a nearby weather station.

Using the Forest Fire Simulator

The group at OSU has used this system for various results including descriptive, predictive and prescriptive analysis:

- Let-Burn analysis for different economic scenarios [Houtman et al. (2013)]
- Economic Analysis of Fire suppression amongst multiple managers [Lauer, 2017]
- Policy Optimization for Fire Suppression

Future Cost Savings for Let-Burn Policy



Fire Treatment Research

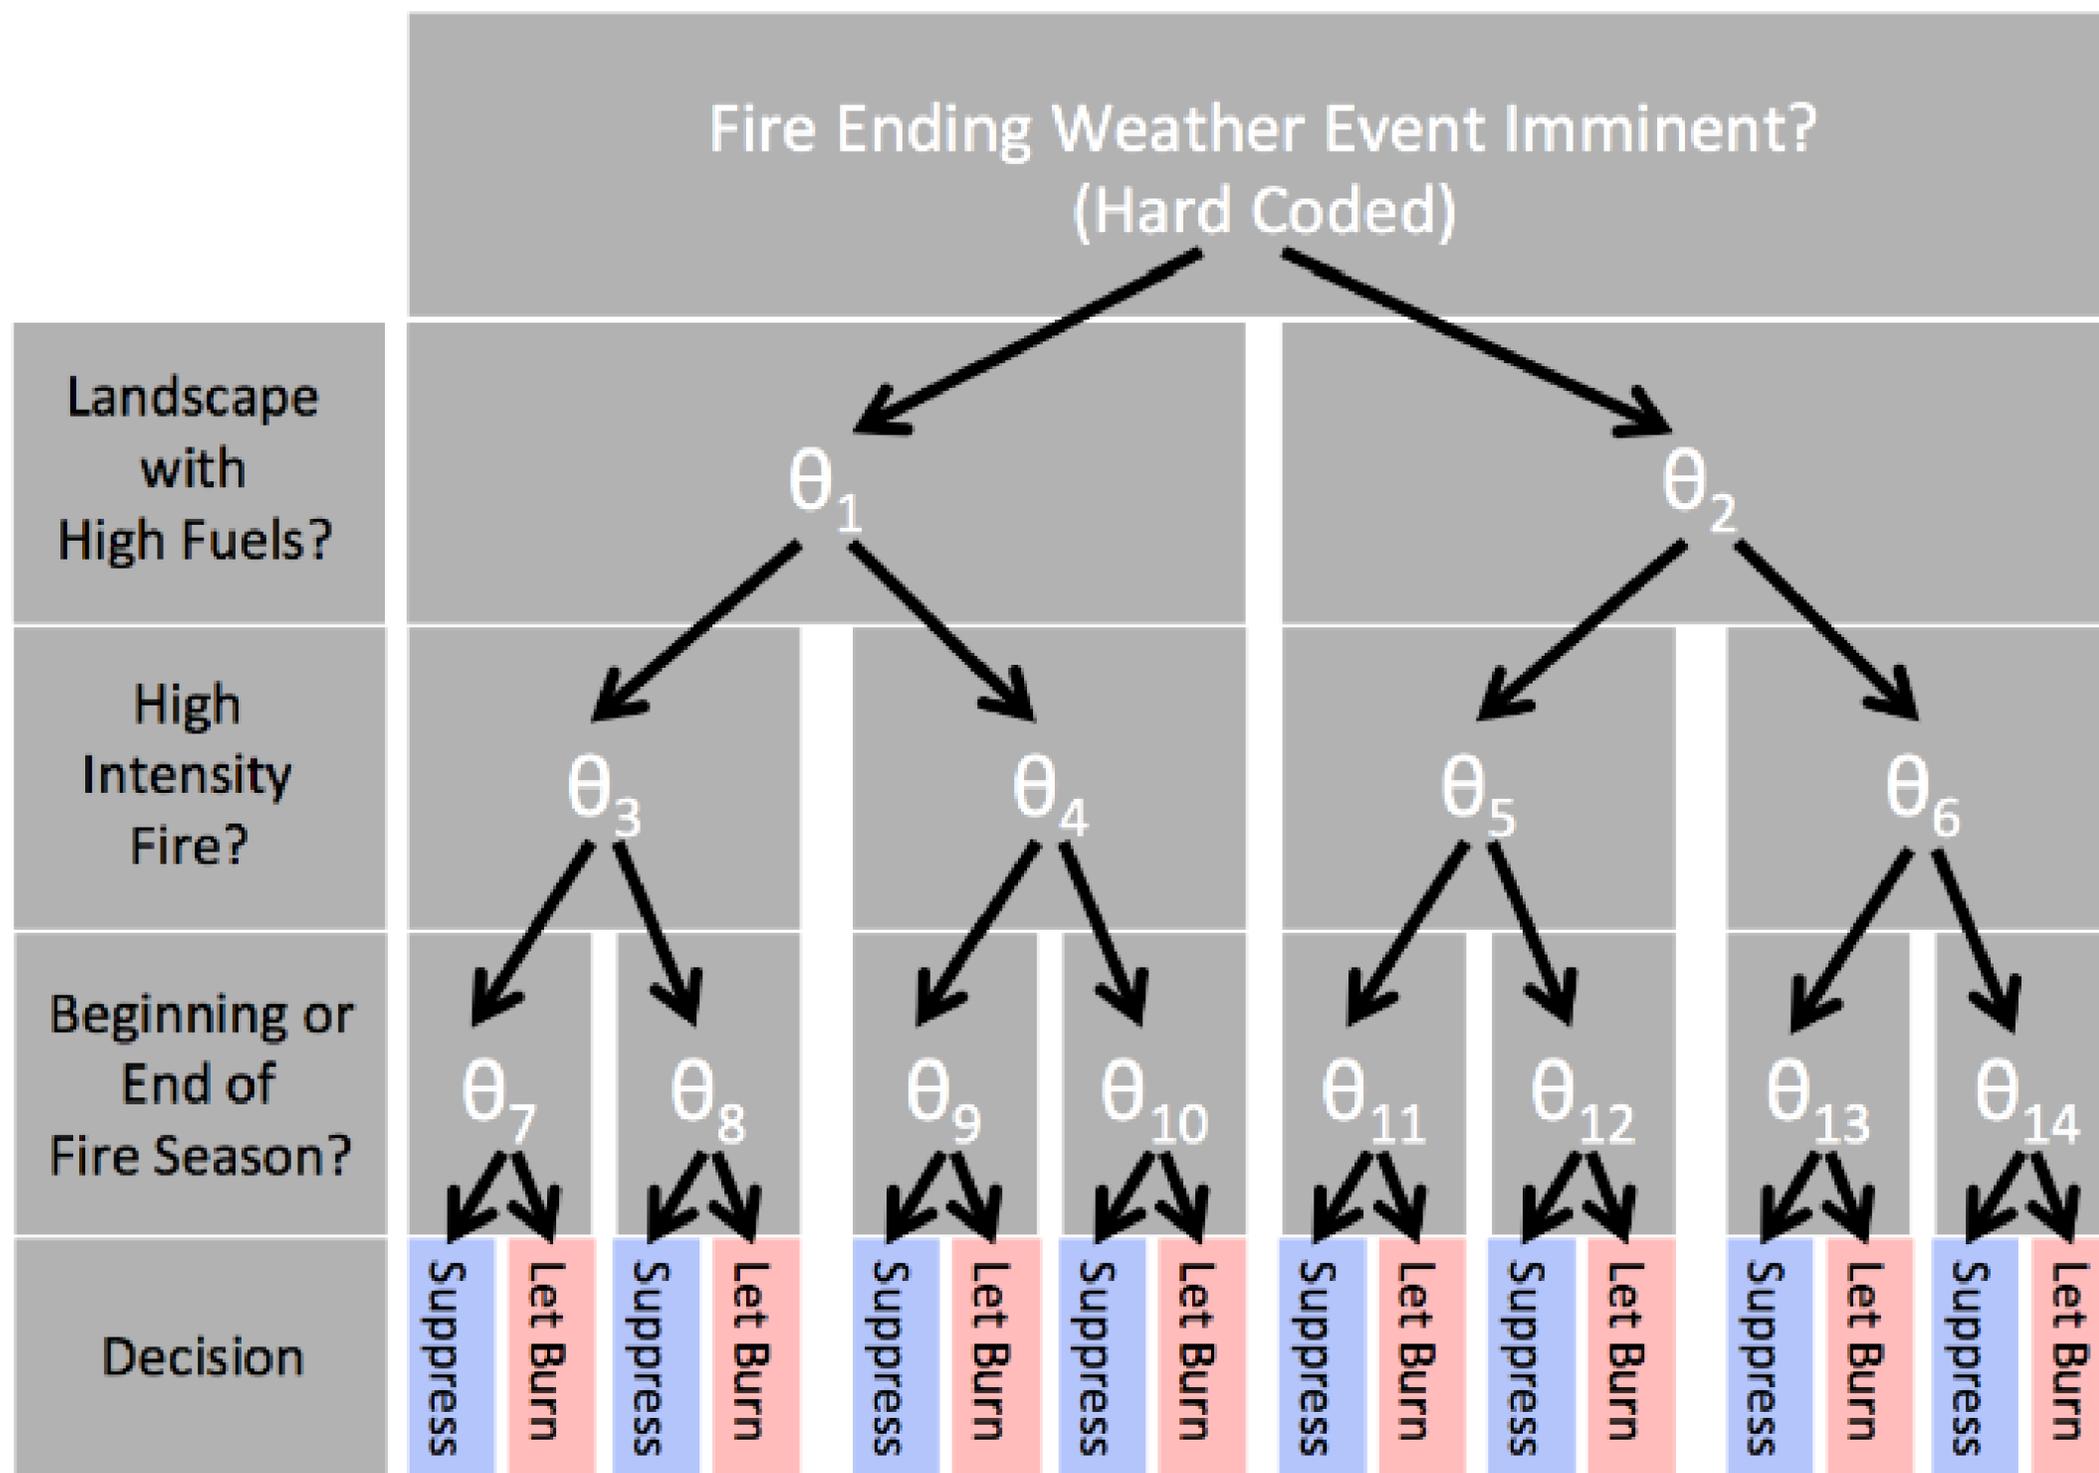
- Using RL (or ADP in their work) [Lauer, 2017] use a stand-based value function $Q_j(s, a)$
- The goal, for each stand j :
 - choose the best action a for that stand based on experience.
- In new work they also formulate a **multi-agent (game theory)** version of this where each agent is a different land manager and they operate on their own cells.
 - The $Q_j^c(s, a)$ is now learned for each agent c and their actions are optimized as a Nash equilibrium.

[C. J. Lauer, C. A. Montgomery, and T. G. Dietterich, “Spatial interactions and optimal forest management on a fire-threatened landscape,” *For. Policy Econ.*, vol. 83, pp. 107–120, Oct. 2017.]

Policy Optimization for Fire suppression

- Problems with the simulator:
 - Too slow for interaction → learn *surrogate model*
 - Surrogate model learning – constructs new trajectories by pasting together parts of old ones
- [McGregor, 2017] Define fire suppression problem as an MDP, solve via direct policy search RL
 - Actions: suppress or let burn for given wildfire
 - Policy defined as function with parameters π_{θ}
- Optimization of policy parameters
 - Sequential Model-based Optimization for general Algorithm Configuration (**SMAC**) [Hutter, 2010]
 - Iterative search for new θ values, model as random forest
- Multiple reward setups considered
 - Components – suppression cost, timber values, deviation from restoration target, air quality (burn days), recreation target

Parameterized Decision Tree Policy

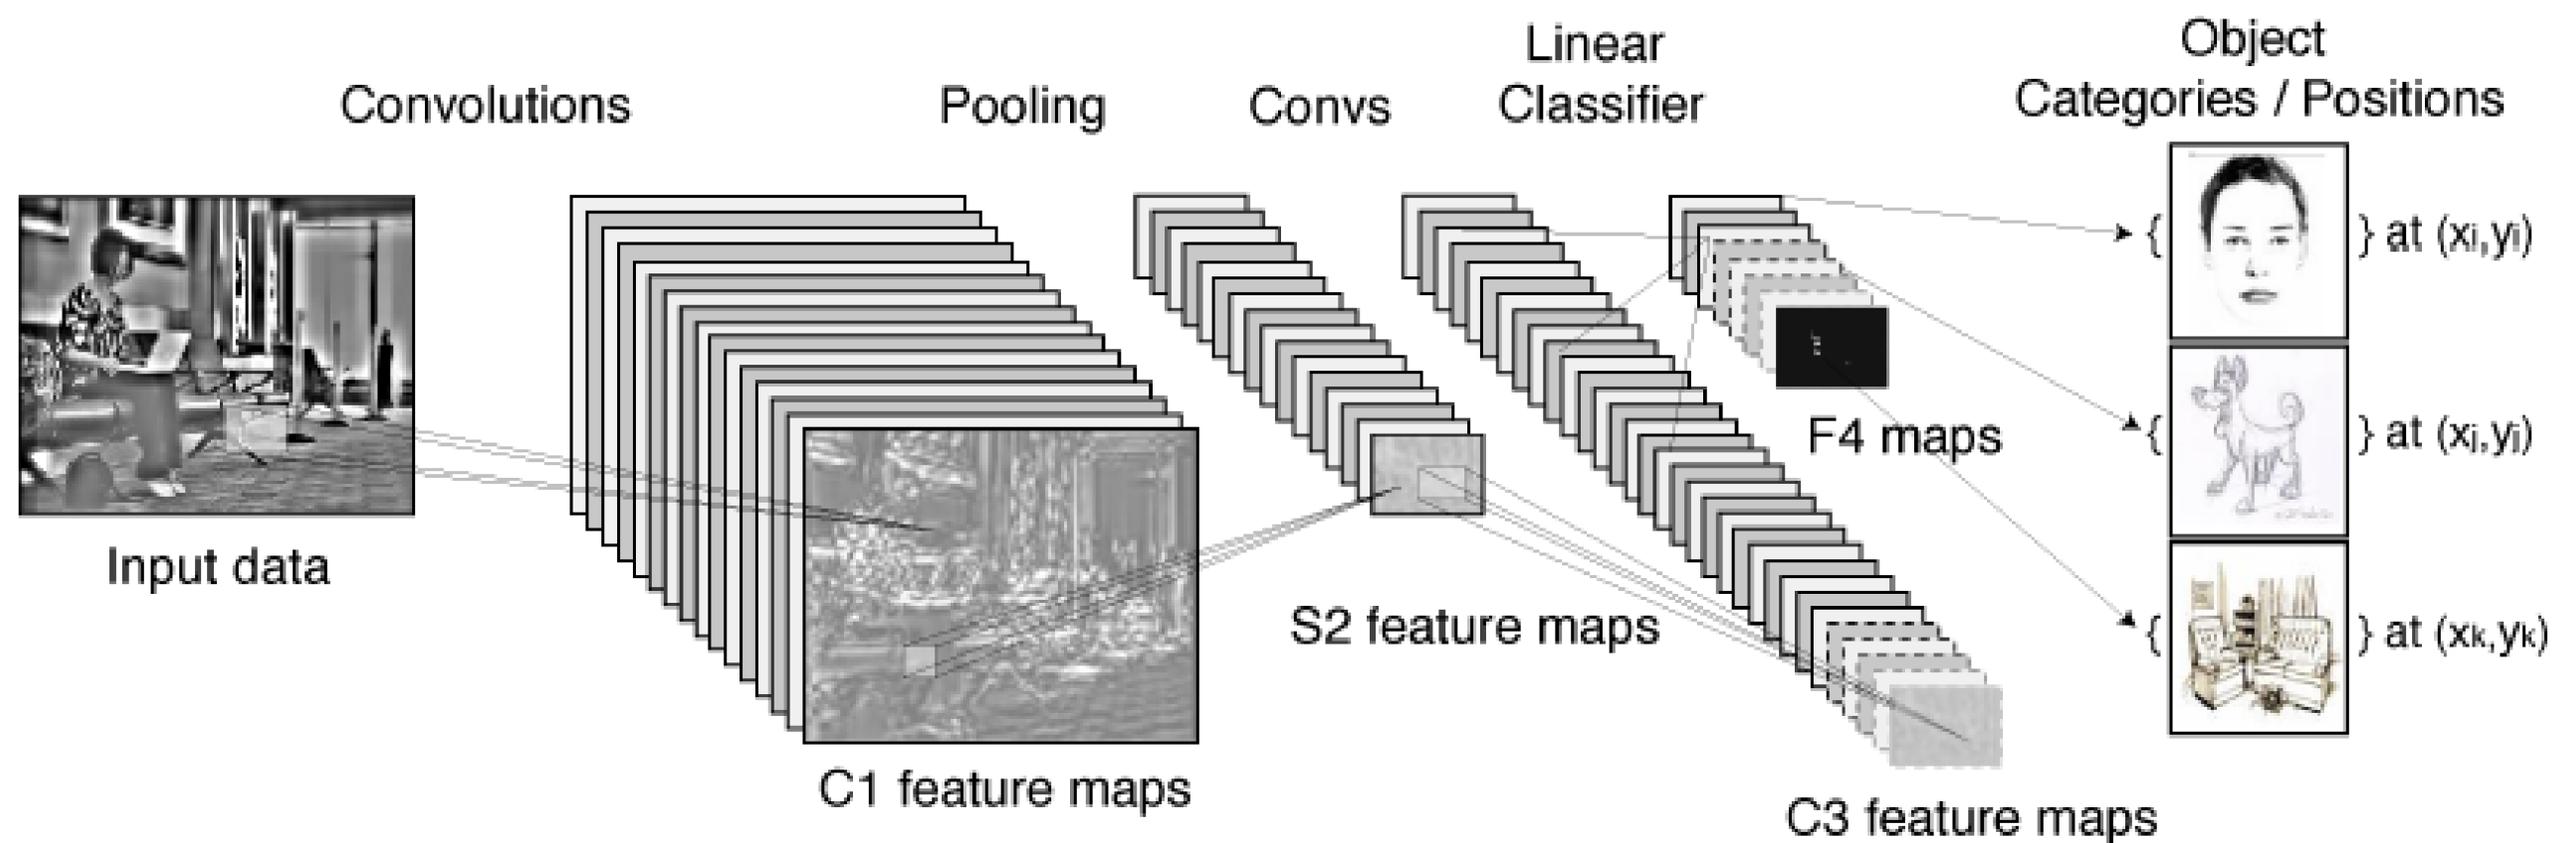


[McGregor, 2017]

Other Relevant AI/ML Methods

- CNNs – For classification of image or spatial data
- LSTM – neural networks with recurrent links across time
- LRCN – learning encoder/decoder to predict entire future frames of a ‘movie’
- GANs – Generative models for learning transferable patterns

Convolutional Neural Networks (CNNs)

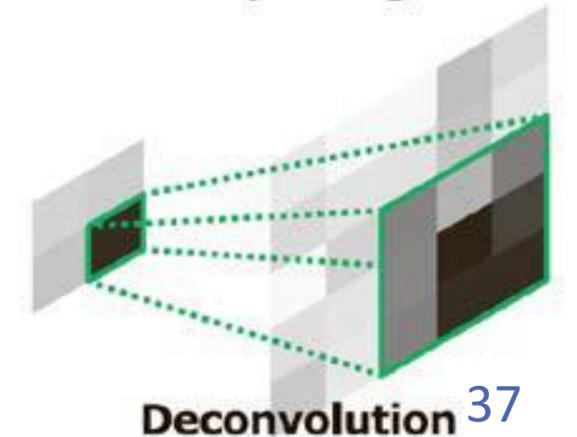
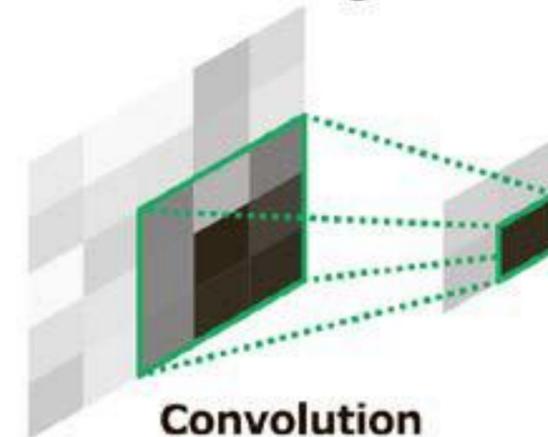
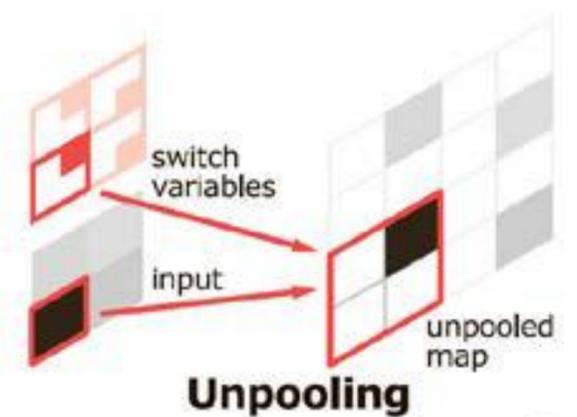
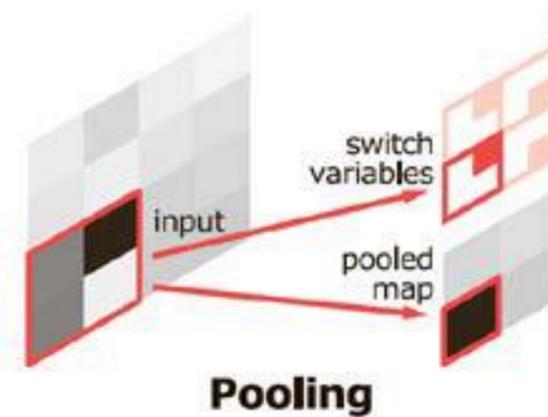
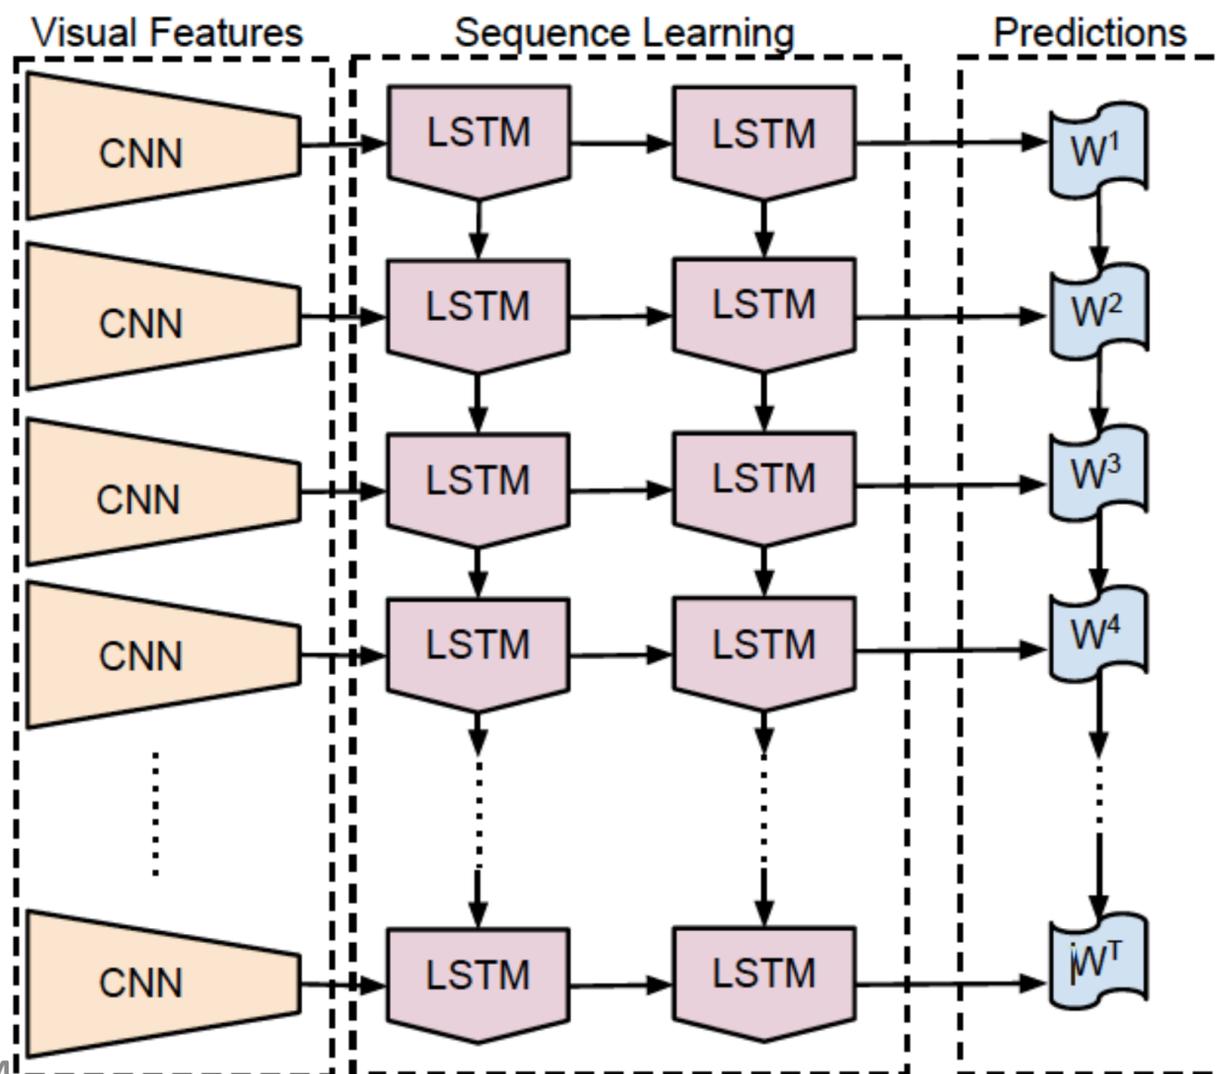


LSTMs

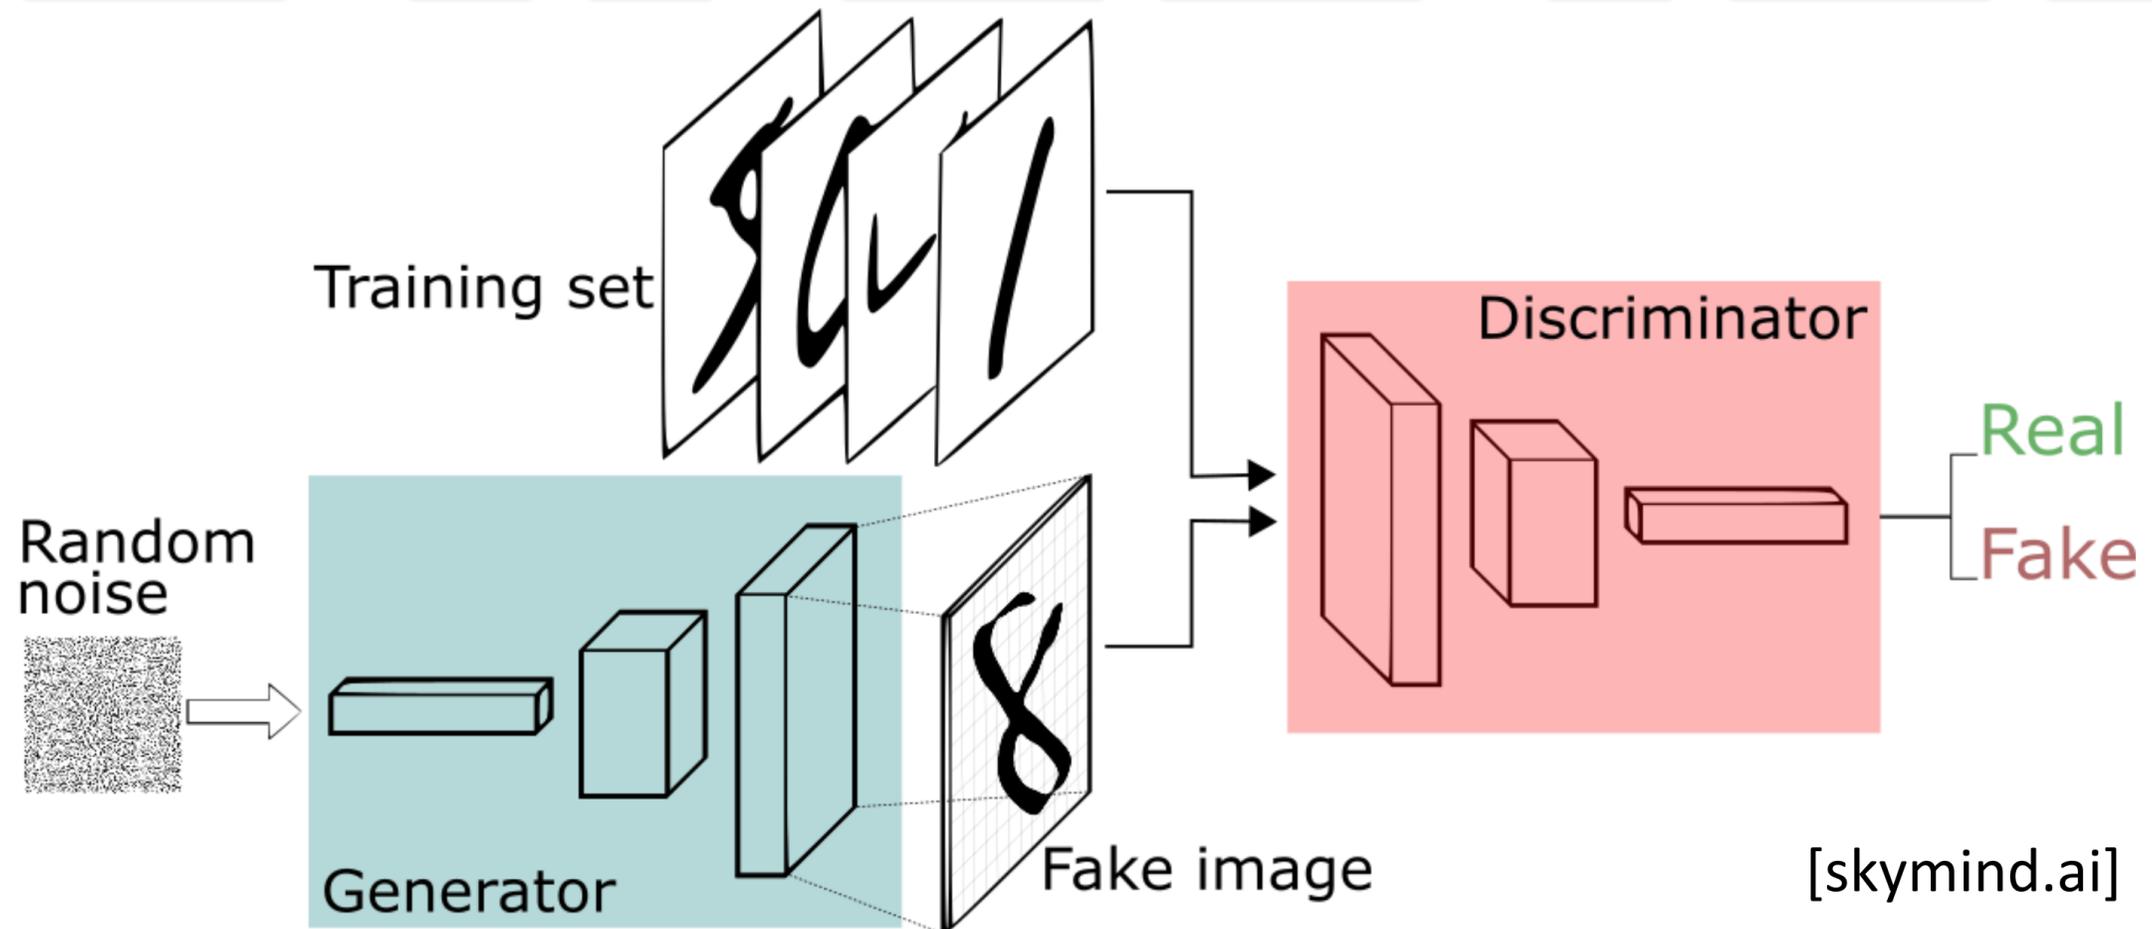
- Long-Term Short-Term networks
- Particular implementation of Recurrent Neural Networks
- integrate strategically remembered particular information from the past
- formalizes a **process for forgetting information** over time.
- useful if you need to learn patterns over time

Long-term Recurrent Convolutional Networks (LRCN)

- ▶ Combine the abilities of CNNs and LSTMs into an end-to-end architecture
- ▶ Treat satellite images as videos, use LRCN to learn generative model for prediction.
- ▶ Ideally we only need to get the spreading process right, not the entire image. Could also do this for series of treatment images if available.

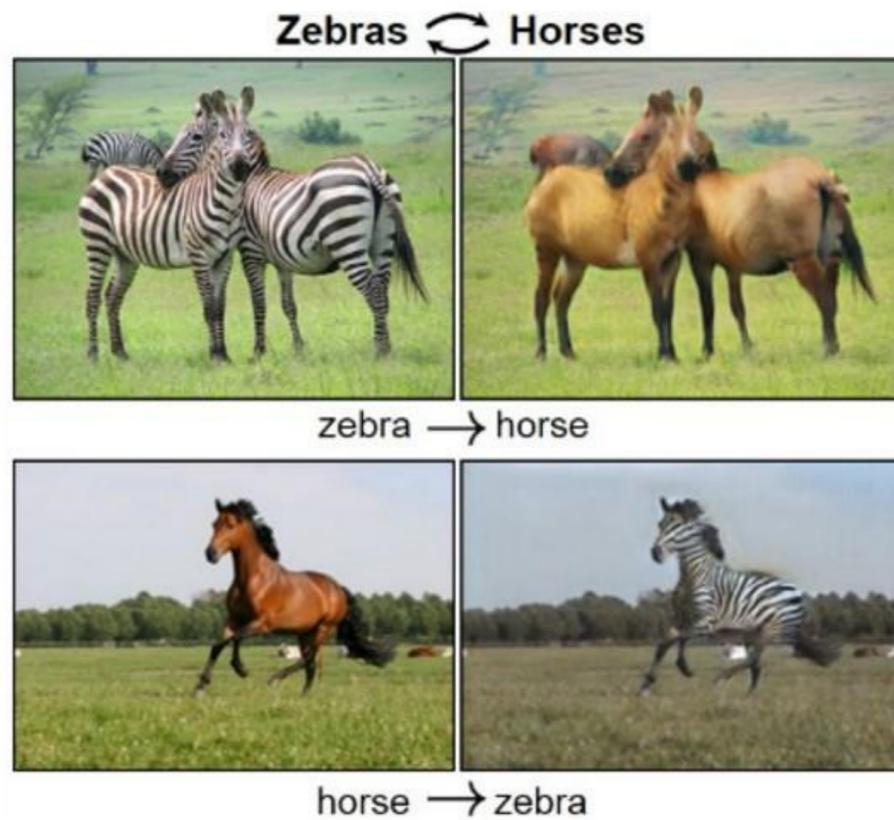


Generative Adversarial Network (GAN)



- One network produces/hallucinates new answers (generative)
- Second network distinguishes between the real and the generated answers (adversary/critic)

Generative Adversarial Network (GAN)



- How can this approach to learn and understand learned models of
 - Forest fire spread?
 - Changing fire risk conditions?
 - Previous firebreak/treatment strategies

Future Work, Interesting Questions

- Forest fire prediction requires additional information consisting of fire-fighting intervention but how to obtain and integrate this data into learning?
- Can the Deep RL policy representations be tailored more closely to this kind of spatially spreading problem?
- How can we fuse high resolution satellite image data with non-image data at different resolutions?
- Agent Based Models can be learned, how can we combine existing “top-down” knowledge about processes with “bottom-up” machine learning coming from all the data?

Conclusions

- **AI/ML are powerful tools for:**
 - Learning patterns from data
 - Supporting human decision making
 - Representing Values
 - Processing Massive Datasets
- **Forest Wildfires are a huge challenge**
 - Large datasets, complex dynamics
 - Challenging decision and value tradeoffs
- **AI/ML methods are underused in the field**
 - The field is changing quickly and tasks which used to be “impossible” are now routine, it’s time to re-evaluate what is possible
 - We’ve only scratched the surface
- **What I think is needed:**
 - High quality data with labels
 - Well defined values/preferences/rewards
 - Concrete targets
 - More ability to work with practitioners to answers the questions they actually need answered

Thank You

UNIVERSITY OF
WATERLOO



WATERLOO.AI
WATERLOO ARTIFICIAL INTELLIGENCE INSTITUTE



Mitacs
Globalink



References

- McGregor, R Houtman, C Montgomery, R Metoyer, TG Dietterich. “*Fast Optimization of Wildfire Suppression Policies with SMACS*”. Arxiv 1703.09391. 2017.
- Christopher J. Lauer, Claire A. Montgomery, Thomas G. Dietterich. “*Spatial interactions and optimal forest management on a fire-threatened landscape*” In *Forest Policy and Economics*, Volume 83, Pages 107-120. 2017.
- Crowley. “Using Equilibrium Policy Gradients for Spatiotemporal Planning in Forest Ecosystem Management” *IEEE Transactions on Computers: Special Edition on Computational Sustainability*, 2013.
- Houtman, Montgomery, Gagnon, Calkin, Dietterich, McGregor and Crowley. “Allowing a wildfire to burn: Estimating the effect on future fire suppression costs”. *International Journal of Wildland Fire*, 2013.
- K. Hall, M. Alkaee Taleghan, H. J. Albers, T. Dietterich, and M. Crowley, “Managing Invasive Species in a River Network,” in *Third International Conference on Computational Sustainability*, 2012.

References

- Crowley and Poole. Policy Gradient Planning for Environmental Decision Making with Existing Simulators. AAAI 2011.
- V. Mnih, A. Puigdomènech Badia, M. Mirza, A. Graves, T. Harley, T. P. Lillicrap, D. Silver, K. Kavukcuoglu, K. Com, and G. Deepmind, “Asynchronous Methods for Deep Reinforcement Learning.” 2016.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. a Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.

Waterloo.AI — The Waterloo AI Institute

- Joint initiative of Faculty of Engineering and Faculty of Mathematics



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING



UNIVERSITY OF WATERLOO
FACULTY OF MATHEMATICS

- Includes over 100 researchers in AI:
 - management sciences, electrical and computer, mechanical and mechatronics, systems design, civil and environmental, and chemical engineering
 - computer science, statistics and actuarial sciences, combinatorics and optimization
 - public health, health systems, biology, chemistry, earth and environmental sciences, and physics and astronomy, economics, accounting and finance

Waterloo.AI — Our Approach

OPERATIONAL ARTIFICIAL INTELLIGENCE

FOR ANY BUSINESS, ANYWHERE, ANYTIME



- We take a complementary, problem-driven approach
 - based on foundational AI that addresses operational problems and enables new products and services
- Our goal
 - bring operational AI to the manufacturing, service, consumer, transportation, finance, agriculture, healthcare, and public sectors

Waterloo.AI – Our Advantage

OPERATIONAL ARTIFICIAL INTELLIGENCE

FOR ANY BUSINESS, ANYWHERE, ANYTIME



- Multidisciplinary research teams including mathematicians, computer scientists, and engineers
- Established expertise in collaborating with industry and developing real-world solutions to commercial challenges

Waterloo.AI — Planned Activities

- Organization of workshops and seminars in various areas of AI involving industrial partners, government and research labs at UW
- Short courses throughout the year with tutorials in pertinent fields of machine learning and AI (reverse co-op)
- Short term and long term research projects through funding sponsorship of partners
- Competitions and datafests involving teams composed of faculty members and students

Waterloo.AI — Foundational Research Areas

- **Machine learning, statistical learning**
- Natural language processing
- **Computer vision**
- **Probabilistic models, knowledge discovery, and knowledge representation**
- **Multi-agent systems** and game theory
- Health Informatics
- **Optimization and decision making**
- Trust modeling
- Affective computing and sentiment analysis
- Human-computer interaction
- Neuroscience



A Short History

- ▶ 40's Early work in NN goes back to the 40s with a simple model of the neuron by McCulloch and Pitt as a summing and thresholding devices.
- ▶ 1958 Rosenblatt in 1958 introduced the Perceptron, a two layer network (one input layer and one output node with a bias in addition to the input features).
- ▶ 1969 Marvin Minsky: 1969. Perceptrons are 'just' linear, AI goes logical, beginning of "AI Winter"
- ▶ 1980s Neural Network resurgence: Backpropagation (updating weights by gradient descent)
- ▶ 1990s SVMs! Kernels can do anything! (no, they can't)