

Visions for Systems and Control in the Era of Learning and Data Science

Deepan Muthirayan and Pramod P. Khargonekar

Department of Electrical Engineering and Computer Science
University of California, Irvine

Workshop on New Problems in Learning and Data Science in Control Theory
American Control Conference
June 26, 2018

Outline

Why Learning and Data Science in Control?

Data Science and Machine Learning

Key Recent Advances in Machine Learning

Directions in Control that Leverage Machine Learning

Why Learning and Data Science in Control?

Control Systems : Strong theoretical foundations

- ▶ Stability theory
- ▶ Algebraic, analytical, and topological structures
- ▶ Linear multivariable control
- ▶ Robust control
- ▶ Nonlinear control
- ▶ Adaptive Control
- ▶ Stochastic control
- ▶ Optimal control
- ▶ Distributed control

Control Systems: Diverse Application Domains

- ▶ Aerospace
- ▶ Automotive
- ▶ Transportation
- ▶ Communications
- ▶ Energy and power
- ▶ Water and agriculture
- ▶ Manufacturing
- ▶ Chemical processes
- ▶ Biomedical

Aspirational and Emerging Applications: Examples

- ▶ Smart-X
 1. Smart manufacturing - Industry 4.0
 2. Smart grid
 3. Smart cities
 4. Smart health
- ▶ Urban mobility
- ▶ Low carbon society
- ▶ Sustainable technologies
- ▶ Precision and population health and wellness
- ▶ Food-energy-water nexus
- ▶ Autonomous systems

Example: Future of Manufacturing and Industry 4.0

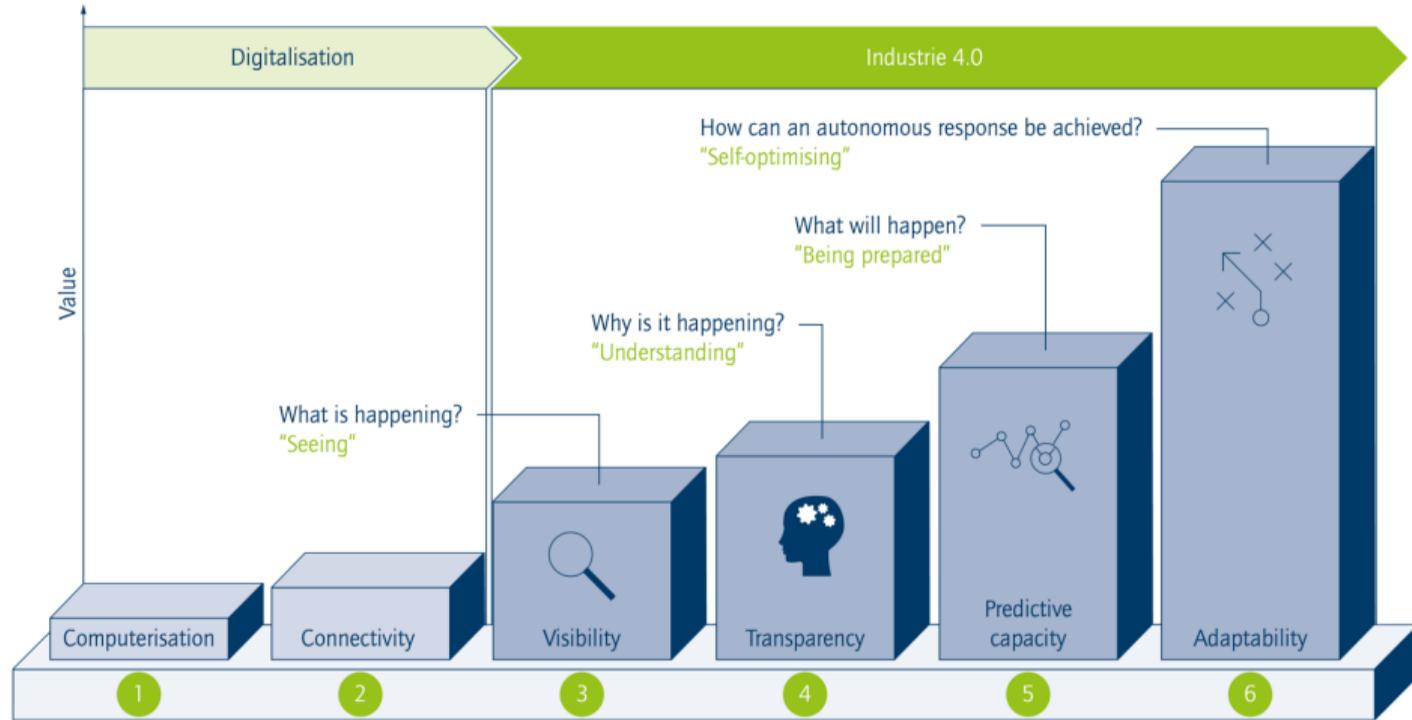
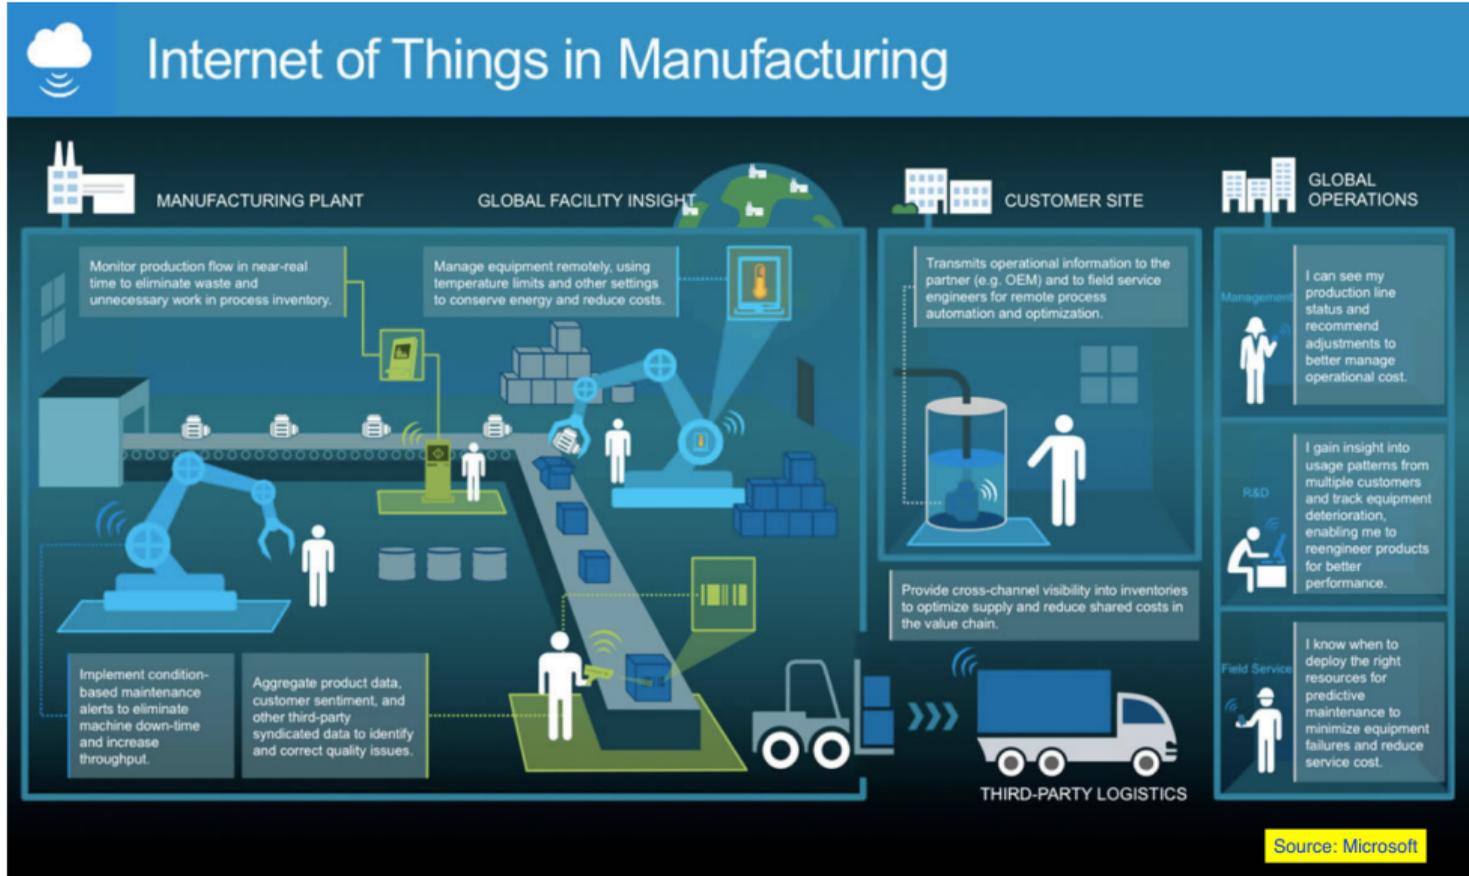


Figure 5: Stages in the Industrie 4.0 development path (source: FIR e. V. at RWTH Aachen University)

Example: IoT in Manufacturing



Example: Smart Energy Future

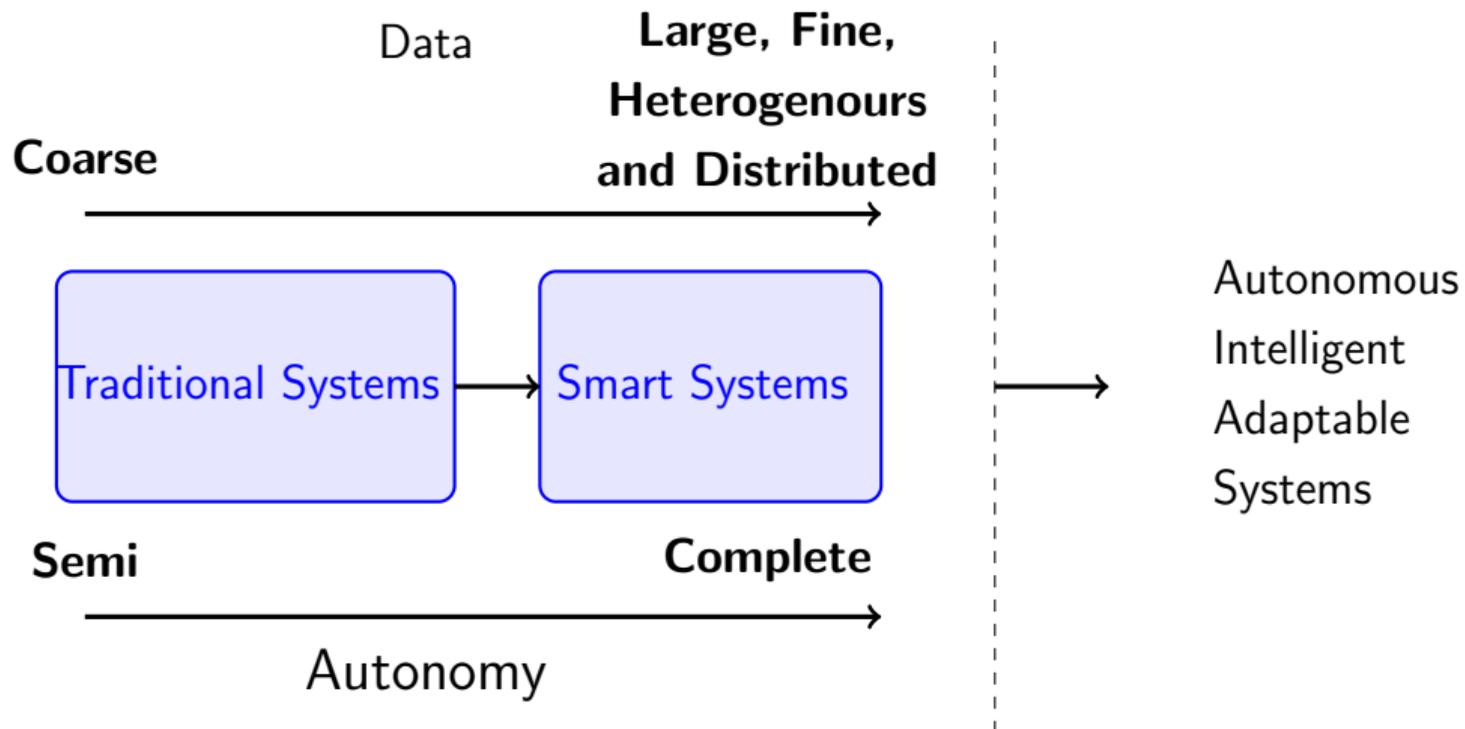


Example: Smart City



Source: Medium

Evolution of Systems



Common Themes

- ▶ These “systems of systems” are **large heterogeneous distributed networks** of (IoT) connected devices, systems, and human/social agents
- ▶ They require high **reliability** and **efficiency** under **uncertainty**
- ▶ Physical and informational flows and interactions
- ▶ **Real time** distributed decision and control

Challenge: Design, operation, management and control of large, distributed, heterogeneous, complicated, interconnected socio-techno-economic systems

Vision: Control systems will play an important role but we will need to integrate with cyber-physical-human systems, data science, and machine learning

Data Science and Machine Learning

Data Science Definition

“ ... data science is a new interdisciplinary field that synthesizes and builds on statistics, informatics, computing, communication, management, and sociology to study data and its environments (including domains and other contextual aspects, such as organizational and social aspects) in order to *transform data to insights and decisions* by following a data-to-knowledge-to-wisdom thinking and methodology.”

[Cao (2017)]

“Data Science is the science of learning from data; it studies the methods involved in the analysis and processing of data and proposes technology to improve methods in an evidence-based manner.”

[Donoho (2015)]

Why is Big Data Gaining Attention?



Relevant Data Science Techniques and Resources

- ▶ Regression
- ▶ Classification — logistic regression and discriminant analysis
- ▶ Resampling methods
- ▶ Shrinkage — ridge regression, LASSO
- ▶ Dimensionality reduction — PCA, PLS, nonlinear techniques
- ▶ Streaming data
- ▶ Uncertainty quantification
- ▶ Programming and software tools
- ▶ Cloud infrastructure

Data to Decisions is a Form of Control

Data



Insight

(reactive understanding, detection, prediction)



Prescription (Action)

Computational Intelligence: Pattern Recognition or Model Building

Two fundamentally different perspectives on learning from data

1. Statistical pattern recognition from data for prediction and control
2. Using data to build causal models to understand, predict and control

Machine Learning

- ▶ Supervised learning
- ▶ Semi-supervised learning
- ▶ Unsupervised learning
- ▶ Classification
- ▶ Function estimation and regression

Statistical Learning Theory Framework

- ▶ What is learnable? How difficult is the learning problem? How much training data do we need?
- ▶ PAC - Probably Approximately Correct Learning Framework, Valiant, Vapnik-Chervonenkis, ...
- ▶ Performance of learned model is within ϵ of the optimal with probability $\geq (1 - \delta)$
- ▶ Learnability: there exists learning rule that is PAC with finite sample size $m(\epsilon, \delta)$
- ▶ Vapnik-Chervonenkis (VC) dimension — how expressive is the class of models?
- ▶ Main result: PAC learnable \iff VC dimension is finite

Classification

- ▶ Supervised learning - learning from examples
- ▶ Nearest neighbor rules
- ▶ Kernel based methods
- ▶ PAC - Probabaly Approximately Correct Learning Framework, VC dimension
- ▶ Support Vector Machines (SVM)
- ▶ Multi-layer neural networks

Function Estimation and Regression

- ▶ Linear regression
- ▶ LASSO: Least Absolute Shrinkage and Selection Operator
- ▶ Kernel based methods
- ▶ Multi-layer neural networks

Improving Learning: Boosting

- ▶ Boosting — iterative approach to improve any learning algorithm
- ▶ Iteratively improve a weak classifier to produce a strong classifier
- ▶ Alter the distribution of weights on training examples
- ▶ Combine all the classifiers with weighting factors
- ▶ Special case: AdaBoost algorithm
- ▶ Performance bounds for AdaBoost

Curse of Dimensionality

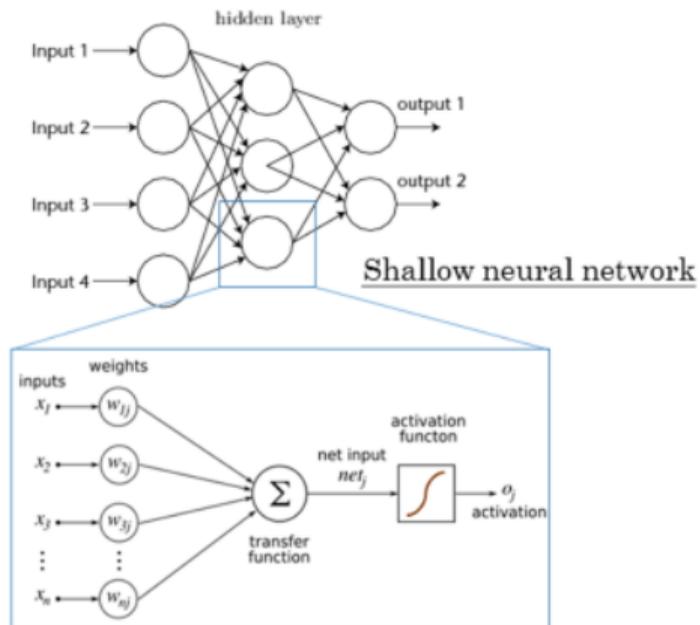
- ▶ Introduced by R. E. Bellman to describe problems as dimension increases in the context of dynamic programming
- ▶ Relevant to data science, machine learning, and control problems
- ▶ High dimensional spaces have surprising and non-intuitive properties
 1. Ratio of volume of unit hypersphere to unit hypercube rapidly approaches zero
 2. Volume of multi-dimensional Gaussian inside radius 1.65 goes from 90 % to zero very rapidly as dimension increases
 3. See the paper [Verleysen and François 2005] for more details

Key Recent Advances in Machine Learning

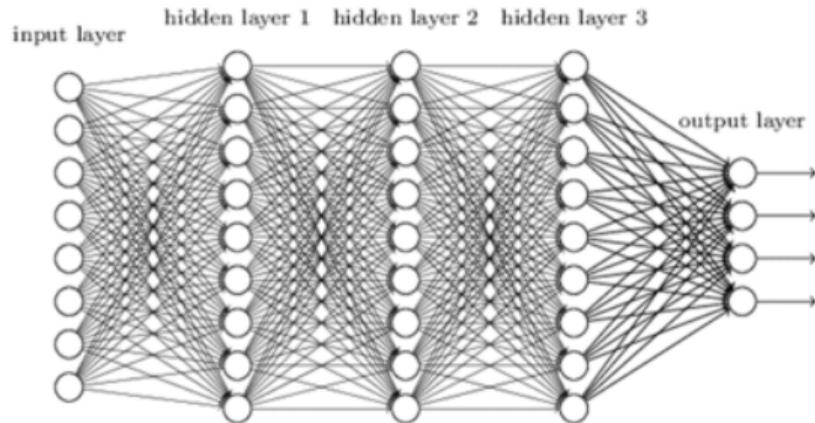
Deep Learning Breakthroughs: Key Elements

- ▶ Large datasets for training
- ▶ Multi-layer neural network models
- ▶ Leveraging computing power - graphics processors
- ▶ Computationally efficient training
- ▶ Able to combat curse of dimensionality, especially for composable functions (“functions of functions”)

Deep Neural Network



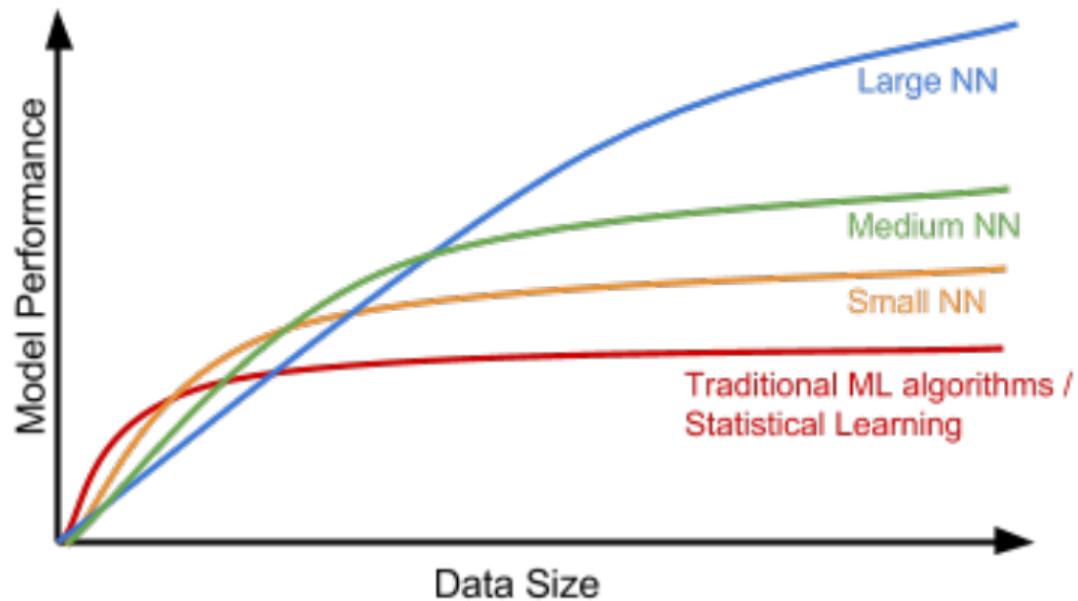
Deep neural network



Deep Architectures

- ▶ Multiple layers of representation of input data
- ▶ Hidden layers automatically learn input features
- ▶ Compositionality built into the architecture
- ▶ Computationally efficient training, e.g., stochastic gradient descent
- ▶ Use of priors, e. g., number of hidden layers, number of units in each layer, to combat curse of dimensionality

Scaling of Deep Neural Networks



Approximation Properties

- ▶ Classical result: Neural network with a single hidden layer can approximate any continuous function on a compact domain
- ▶ Current wisdom: deep networks have significant benefits
- ▶ Some theoretical results to show deep networks have exponential benefits over shallow networks [Telgarsky 2016]
- ▶ Much work on expressive power of deep networks
- ▶ Minimization of training error is NP hard yet training can be done in practice
- ▶ Issues in training: saddle points in high-dimensional non-convex optimization

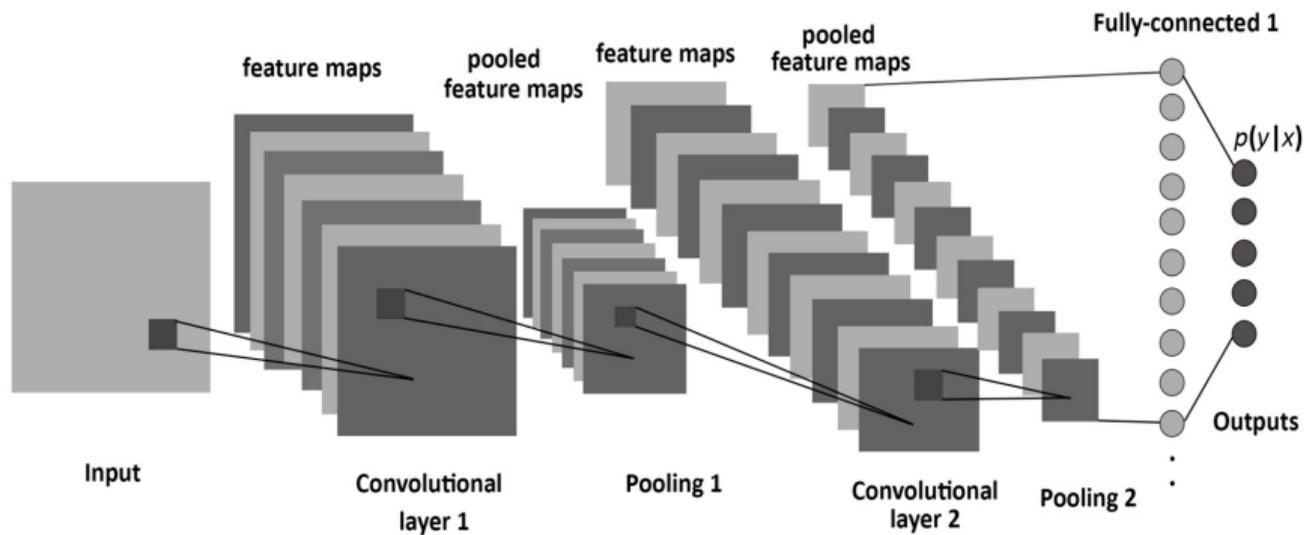
Generalization Properties

- ▶ Deep networks — too many parameters compared to training data
- ▶ (Why) do deep networks generalize well away from where they are trained?
- ▶ Insufficient and incomplete understanding of deep learning generalization performance
 1. Standard theoretical tools: VC dimension, Rademacher complexity, uniform stability
 2. Recent result: deep neural networks easily fit random labels [Zhang et al. 2016]
 3. Regularization is neither necessary nor sufficient for generalization

Convolutional Neural Networks

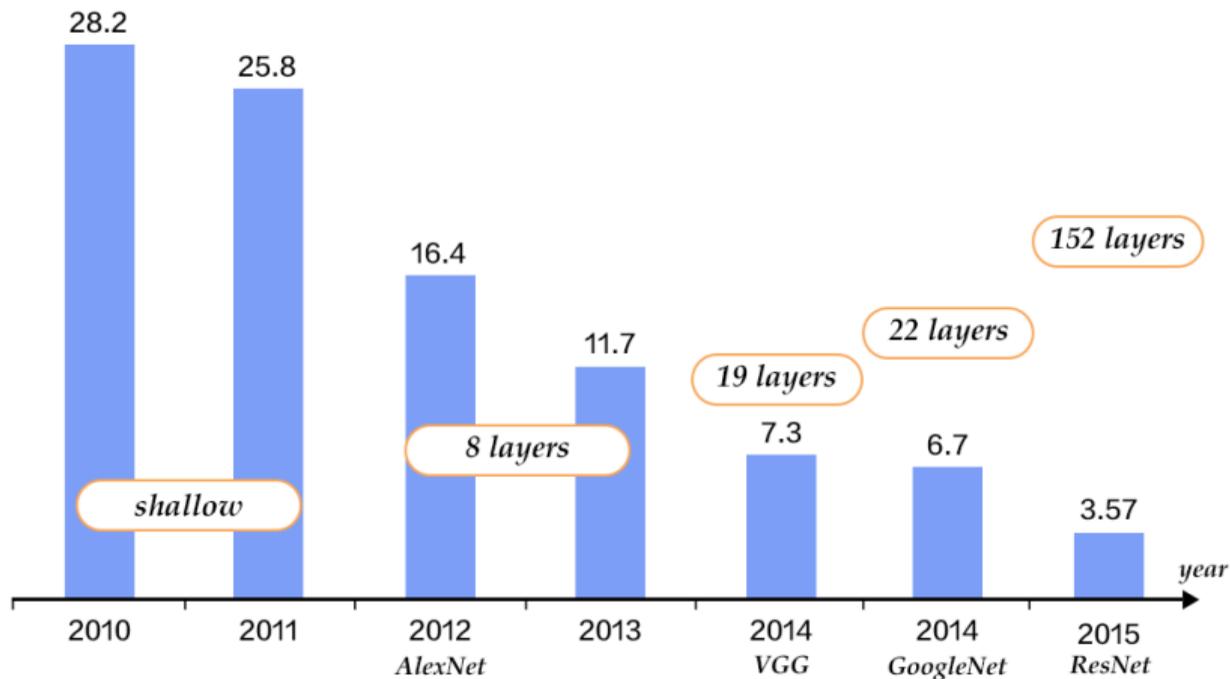
- ▶ Special architectures to deal with data in the form of arrays, layers of convolutions and pooling
- ▶ LeNet, AlexNet, VGG, GoogLeNet, Residual Learning, ...
- ▶ Impressive performance gains in ImageNet competition with error rates below 5%
- ▶ Dominant method for computer vision
- ▶ ConvNet models being used in vision systems for autonomous cars

Example: Convolutional Neural Network



Breakthrough in Vision: ImageNet Competition

Krizhevsky, Sutskever, and Hinton 2012



Rapid Progress in Computer Vision

Classification



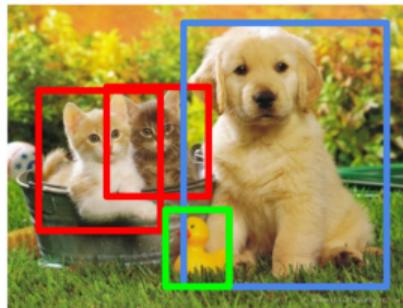
CAT

**Classification
+ Localization**



CAT

Object Detection



CAT, DOG, DUCK

**Instance
Segmentation**

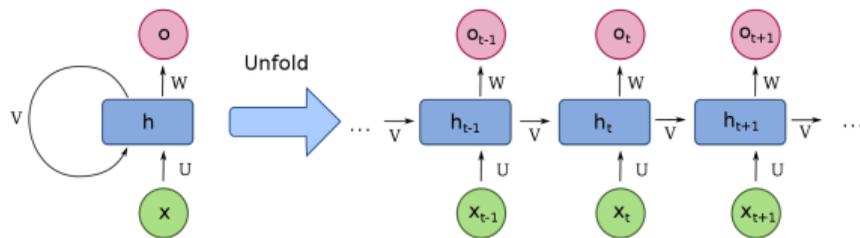


CAT, DOG, DUCK

Single object

Multiple objects

Advances in Other Areas: End 2 End Sequence Learning

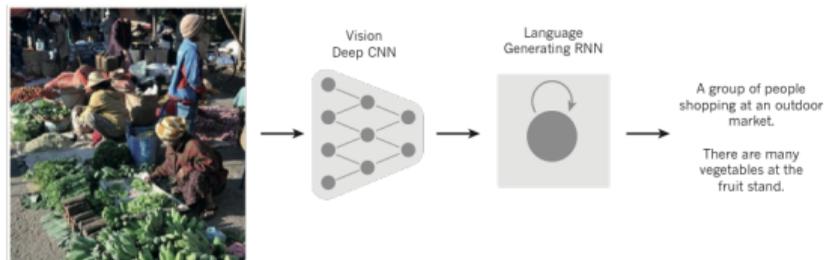


Source: Wiki

- ▶ Algorithm: Backpropagation Through Time (BPTT)
- ▶ Key Innovation: Long Short Term Memory (LSTM)
- ▶ Speech Recognition
- ▶ Google Voice Recognition reported a dramatic 49% increase in accuracy with CTC RNN (connectionist temporal classification RNN)
- ▶ Image to Text
- ▶ Speech Generation
- ▶ Machine Translation
- ▶ Language Models

Example: From Image to Text

Architecture: CNN + LSTM...



A woman is throwing a **frisbee** in a park.



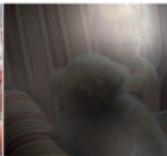
A **dog** is standing on a hardwood floor.



A **stop** sign is on a road with a mountain in the background



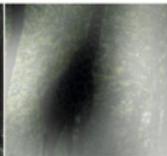
A little **girl** sitting on a bed with a teddy bear.



A group of **people** sitting on a boat in the water.



A giraffe standing in a forest with **trees** in the background.



[Vinyals et al. 2015, Vinyals et al. 2017]

Example: From Image to Text

But lack of context can yield poor results

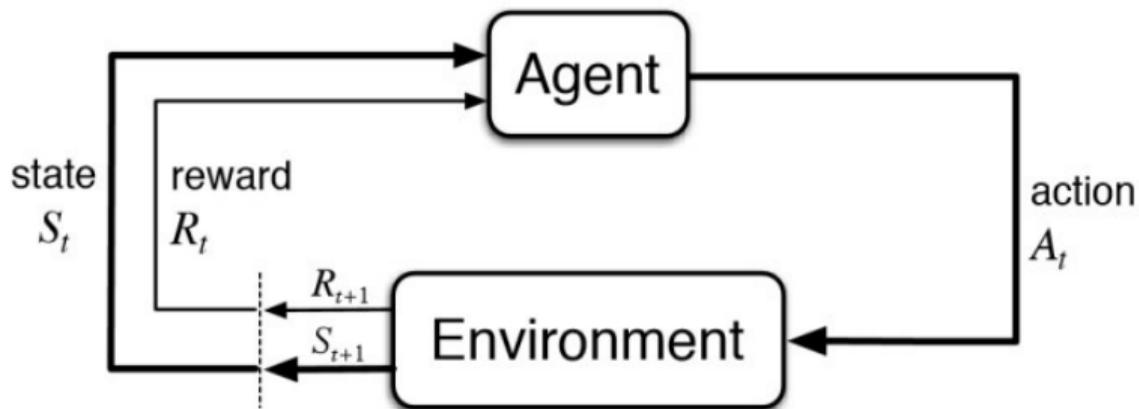


Figure 6. Perceiving scenes without intuitive physics, intuitive psychology, compositionality, and causality. Image captions are generated by a deep neural network (Karpathy & Fei-Fei 2017) using code from github.com/karpathy/neuraltalk2. Image credits: Gabriel Villena Fernández (left), TVBS Taiwan/Agence France-Presse (middle), and AP Photo/Dave Martin (right). Similar examples using images from Reuters news can be found at twitter.com/interesting_jpg.

[Lake et al. 2017]

Reinforcement Learning

This is closer to control systems...



Reinforcement Learning

- ▶ Agent interacting with the environment
- ▶ Observations, actions, and rewards
- ▶ Three key issues:
 1. Generalize experience — learning
 2. Deal properly with delayed gratification — planning
 3. Balance exploration and exploitation
- ▶ Goal: Maximize cumulative discounted future reward
- ▶ Approaches:
 1. Policy search
 2. Value function based
 3. Model based

Bellman Equation

- ▶ Solve for optimal policy π :

$$Q^*(s, a) = \max_{\pi} E[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots | s_t = s, a_t = a, \pi]$$

- ▶ Bellman Equation:

$$Q^*(s, a) = E_{s'}[r + \gamma \max_{a'} Q^*(s', a') | s, a]$$

- ▶ Q Learning: Upon observing and receiving s_t, a_t, r_t, s_{t+1}

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha_t (r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t))$$

DQN Learning

- ▶ RL is known to be unstable or diverge when a function approximator, e. g., NN, is used to represent Q
- ▶ Key issues: correlations in sequence of observations, small updates to Q may cause big changes to policy, . . .
- ▶ Recent approach: use deep convolutional neural networks to approximate Q [Mnih et al. 2015]
- ▶ Key innovations: Experience replay and fixed target network

Experience Replay and Fixed Target Network for DQN

- ▶ Naive neural network based value function learning loss function,

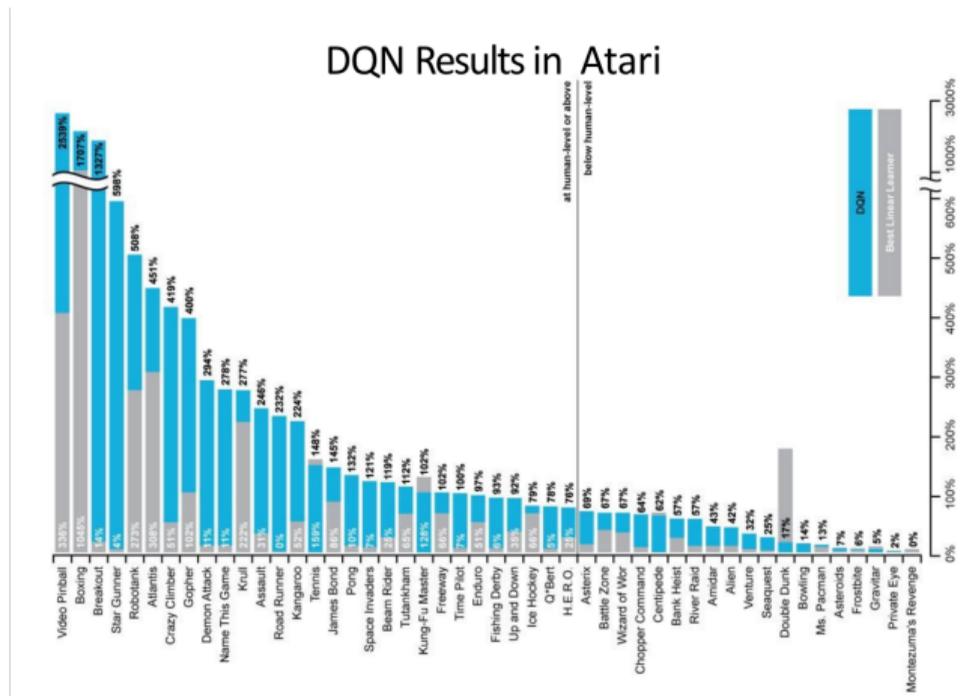
$$L(w) = \mathbb{E}_{s'} \left[\left(r + \gamma \max_{a'} Q(s', a', w) - Q(s, a, w) \right)^2 \right]$$

- ▶ DQN loss function (includes experience replay and fixed target network),

$$L(w) = \mathbb{E}_{s,a,r,s' \sim D} \left[\left(r + \gamma \max_{a'} Q(s', a', w^-) - Q(s, a, w) \right)^2 \right]$$

DQN in Atari Games

- Self play for evaluation



Some Recent Improvements

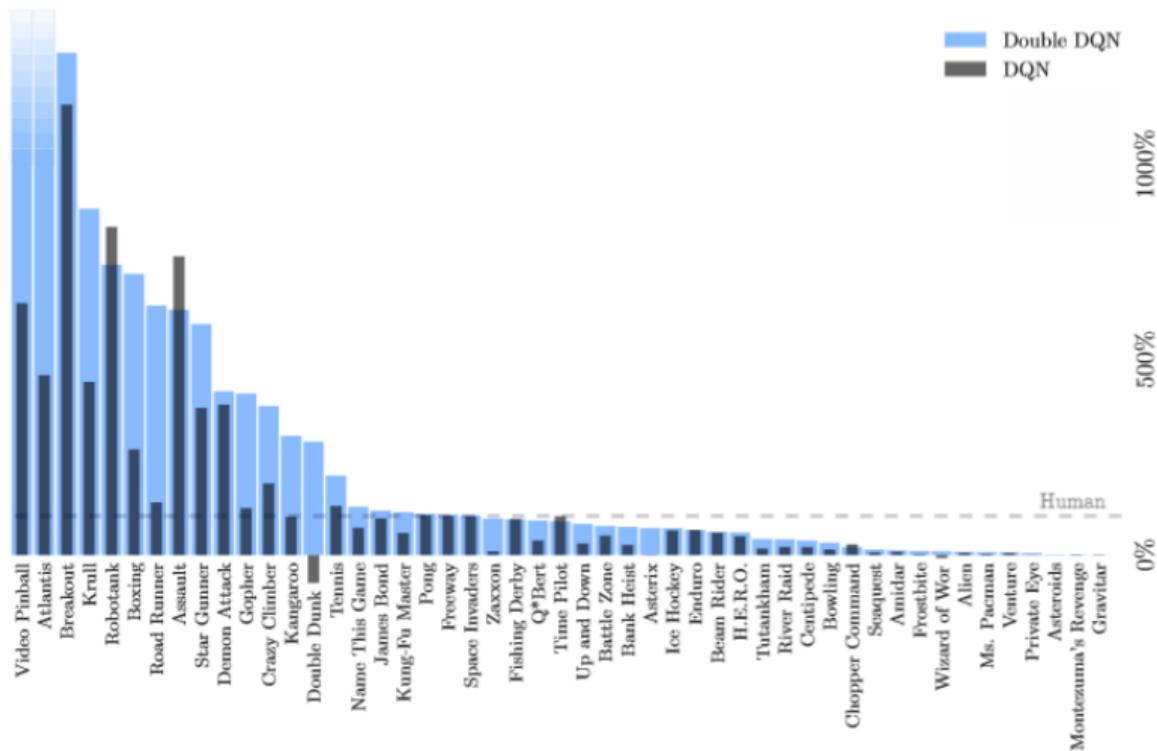
- ▶ Double DQN: Avoids overestimation problem, because of max. over actions, in DQN

$$L(w) = \mathbb{E}_{s,a,r,s' \sim D} \left[\left(r + \gamma Q(s', \operatorname{argmax} Q(s', a, w), w^-) - Q(s, a, w) \right)^2 \right]$$

- ▶ DDQN does not always improve performance.
- ▶ Prioritized Replay: prioritizes sampling from experience instead of random sampling
- ▶ Multistep learning: includes multiple time steps for the target function instead of two steps as in DQN

[Hessel et al. 2017]

Example: Double DQN

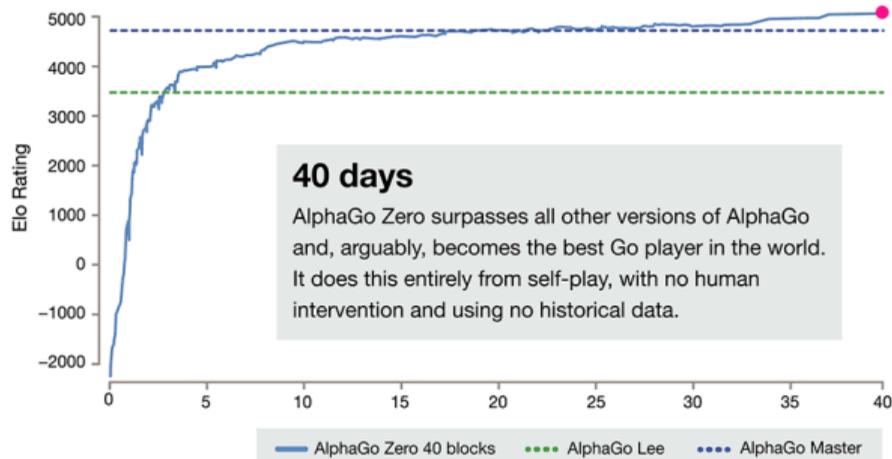


Breakthrough in Go

- ▶ Game is complex - greater number of possible moves compared to chess at each game state
- ▶ Game is longer - greater number of steps compared to chess
- ▶ Explosion of state space as the game progresses
- ▶ It took 20 years for a computer program to defeat a Go master after this was achieved in Chess

AlphaGo Zero

- ▶ Deep Neural Network: Probabilities of action given state
- ▶ Search based policy improvement: MCTS based policy search for policy improvement
- ▶ Search based policy evaluation: Self play with MCTS based policy search as policy evaluator [Silver et al. 2017]



Self Programmable Deep Neural Networks

- ▶ Key: Actor-Critic Framework
- ▶ Recent Results from Deep Mind
- ▶ Learning to write programs that generate images (similar to GAN)
- ▶ Learning complex tasks with sparse rewards (A3C [Mnih et al. 2016], A3C + Auxiliary Tasks [Jaderberg et al. 2016], SAC - X)

[Source: <https://deepmind.com/blog/deep-reinforcement-learning/>]

Continuous RL

- ▶ A large fraction of recent RL literature is on discrete action spaces
- ▶ DQN is limited to small action spaces
- ▶ We need RL for continuous action spaces
- ▶ Key approach: Actor-Critic approach using Deep Deterministic policy gradient (DDPG) [Lillicrap et al. 2015]

Model Based Reinforcement Learning

- ▶ Learn the transition dynamics and reward functions from online observations and rewards
- ▶ Known body of results on model based reinforcement learning
- ▶ Perhaps better suited to integrate with established control theory
- ▶ Example: Safe Model-based Reinforcement Learning with Stability Guarantees [Berkenkamp et al. 2017]

Directions in Control that Leverage Machine Learning

What can we leverage from these advances?

- ▶ Overcome CoD in statistical inference of high dimensional data: so that smart systems can handle large and high dimensional heterogeneous data
- ▶ Automate feature learning: critical for inference in smart systems
- ▶ Automate learning new and complex tasks: reduces the learning curve for smart systems

Key Overall Strategies

- ▶ Focus on emerging applications, e.g., Smart-X
- ▶ Leverage key advances in machine learning and data sciences
- ▶ Leverage advances in computational hardware, software infrastructures, etc.
- ▶ Transdisciplinary collaborations with complementary disciplines
- ▶ Nurture a diverse interconnected ecosystem of fundamental and translational research

Control Related Strategies

- ▶ More expansive of view control: beyond closing a feedback loop
- ▶ Go beyond algorithms to also include “control architecture”
- ▶ Find approaches to incorporate “compositionality”
- ▶ Dealing with systems where it is difficult to create first principles based models
- ▶ Higher level control functions: supervisory, feedforward, planning, . . .
- ▶ Incorporate socio-economic aspects — a major challenge and opportunity

Learning for Control

- ▶ Leverage the new capabilities to deal with high-dimensional data from sensors
- ▶ Exploit “learning to learn” paradigm for controlling “new systems” from control of known systems
- ▶ Address large uncertainties in the environment (and models)
- ▶ Learning to overcome modeling barriers
- ▶ Explore creative approaches to combine learning techniques with model based control
- ▶ Leverage what humans do well to create more powerful controllers

Hive Mind for Control

- ▶ Leverage IoT infrastructure to rapidly absorb learning from one instance to the whole family
- ▶ Example: Rolls Royce Aircraft Availability Center



Neuroscience Inspired Control

- ▶ New discoveries on how brain works
- ▶ Example: Predictive brain theory [Clark (2013), Friston (2010)]
- ▶ Example: Brain employs a mixture of model-based and model-free RL
- ▶ Example: Detailed models of sensori-motor control systems [Wolpert et al. (2011)]

Concluding Remarks

- ▶ A large opportunity to advance the field of systems & control
- ▶ Expansion of scope and connections to other disciplines
- ▶ Tremendous room for creative research
- ▶ Plea for open mindedness in pursuing diverse range of research activities
- ▶ Intentional integration and cross-fertilization necessary

References I

-  Felix Berkenkamp et al. “Safe model-based reinforcement learning with stability guarantees”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 908–919.
-  Matteo Hessel et al. “Rainbow: Combining Improvements in Deep Reinforcement Learning”. In: *arXiv preprint arXiv:1710.02298* (2017).
-  Max Jaderberg et al. “Reinforcement learning with unsupervised auxiliary tasks”. In: *arXiv preprint arXiv:1611.05397* (2016).
-  Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems*. 2012, pp. 1097–1105.
-  Brenden M Lake et al. “Building machines that learn and think like people”. In: *Behavioral and Brain Sciences* 40 (2017).
-  Timothy P Lillicrap et al. “Continuous control with deep reinforcement learning”. In: *arXiv preprint arXiv:1509.02971* (2015).

References II

-  Volodymyr Mnih et al. “Asynchronous methods for deep reinforcement learning”. In: *International Conference on Machine Learning*. 2016, pp. 1928–1937.
-  Volodymyr Mnih et al. “Human-level control through deep reinforcement learning”. In: *Nature* 518.7540 (2015), p. 529.
-  David Silver et al. “Mastering the game of go without human knowledge”. In: *Nature* 550.7676 (2017), p. 354.
-  Matus Telgarsky. “Benefits of depth in neural networks”. In: *arXiv preprint arXiv:1602.04485* (2016).
-  Michel Verleysen and Damien François. “The curse of dimensionality in data mining and time series prediction”. In: *International Work-Conference on Artificial Neural Networks*. Springer. 2005, pp. 758–770.
-  Oriol Vinyals et al. “Show and tell: A neural image caption generator”. In: *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*. IEEE. 2015, pp. 3156–3164.

References III

-  Oriol Vinyals et al. “Show and tell: Lessons learned from the 2015 mscoco image captioning challenge”. In: *IEEE transactions on pattern analysis and machine intelligence* 39.4 (2017), pp. 652–663.
-  Chiyuan Zhang et al. “Understanding deep learning requires rethinking generalization”. In: *arXiv preprint arXiv:1611.03530* (2016).