Artificial Intelligence and Machine Learning

Vijay Gadepally
Jeremy Kepner, Lauren Milechin, Siddharth Samsi

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited. This material is based upon work supported by the Under Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Under Secretary of Defense for Research and Engineering.

Slide contributions from: Siddharth Samsi, Albert Reuther, Jeremy Kepner, David Martinez, Lauren Milechin

Outline

- **Artificial Intelligence Overview** ⇨
- **Machine Learning Deep Dives**
  - **Supervised Learning**
  - **Unsupervised Learning**
  - **Reinforcement Learning**
- **Conclusions/Summary**

What is Artificial Intelligence?

Narrow AI: The theory and development of computer systems that perform tasks that augment for human intelligence such as perceiving, classifying, learning, abstracting, reasoning, and/or acting

General AI: Full autonomy

AI and ML - 3
VNG 010720

Definition adapted from Oxford dictionary and inputs from Prof. Patrick Winston (MIT)

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

AI. Why Now?

Big Data

Source: DARPA/ Public domain

Compute Power

Source: DARPA/ Public domain

Machine Learning Algorithms

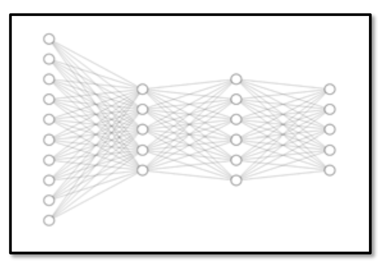

Convergence of High Performance Computing, Big Data and Algorithms that enable widespread AI development

AI and ML - 4
VNG 010720

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

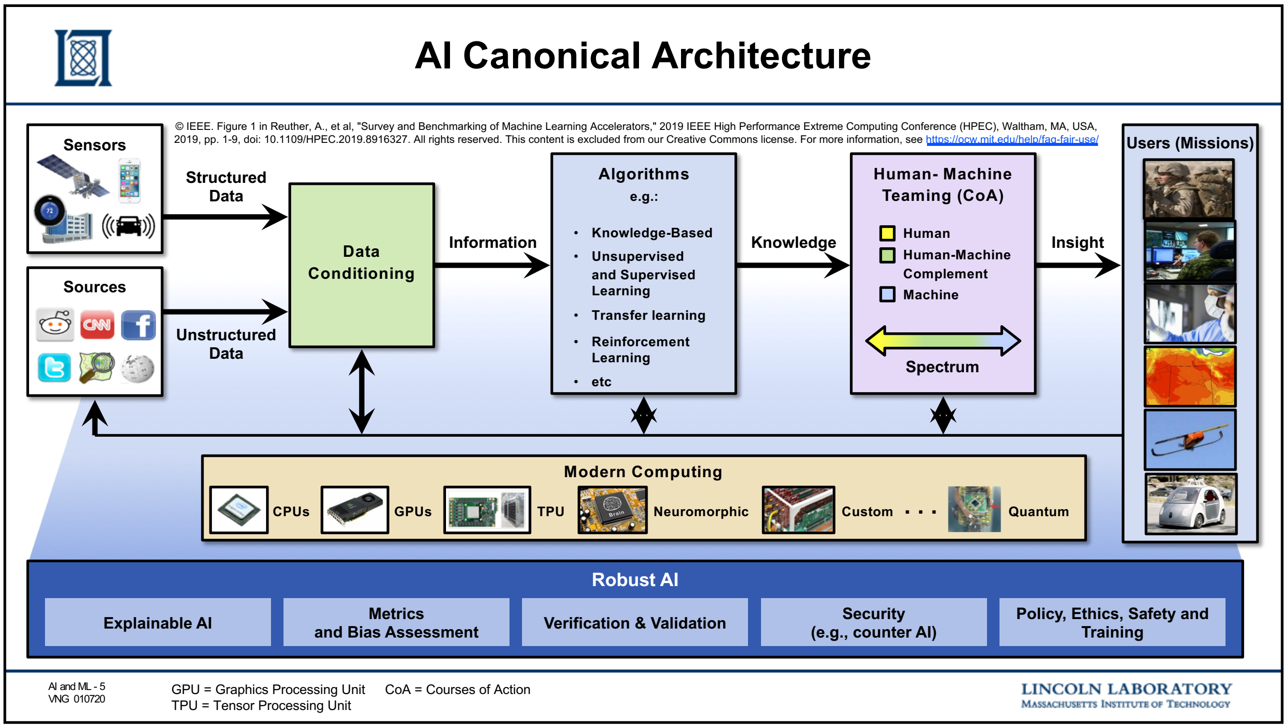
AI Canonical Architecture
© IEEE. Figure 1 in Reuther, A., et al, "Survey and Benchmarking of Machine Learning Accelerators," 2019 IEEE High Performance Extreme Computing Conference (HPEC), Waltham, MA, USA, 2019, pp. 1-9, doi: 10.1109/HPEC.2019.8916327. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/
Sensors
Sources
Structured Data
Unstructured Data
Data Conditioning
Information
Algorithms
e.g.:
• Knowledge-Based
• Unsupervised and Supervised Learning
• Transfer learning
• Reinforcement Learning
• etc
Knowledge
Human- Machine Teaming (CoA)
Human
Human-Machine Complement
Machine
Spectrum
Insight
Users (Missions)
Modern Computing
CPUs
GPUs
TPU
Neuromorphic
Custom
Quantum
Robust AI
Explainable AI
Metrics and Bias Assessment
Verification & Validation
Security (e.g., counter AI)
Policy, Ethics, Safety and Training
AI and ML - 5
VNG 010720
GPU = Graphics Processing Unit
TPU = Tensor Processing Unit
CoA = Courses of Action
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

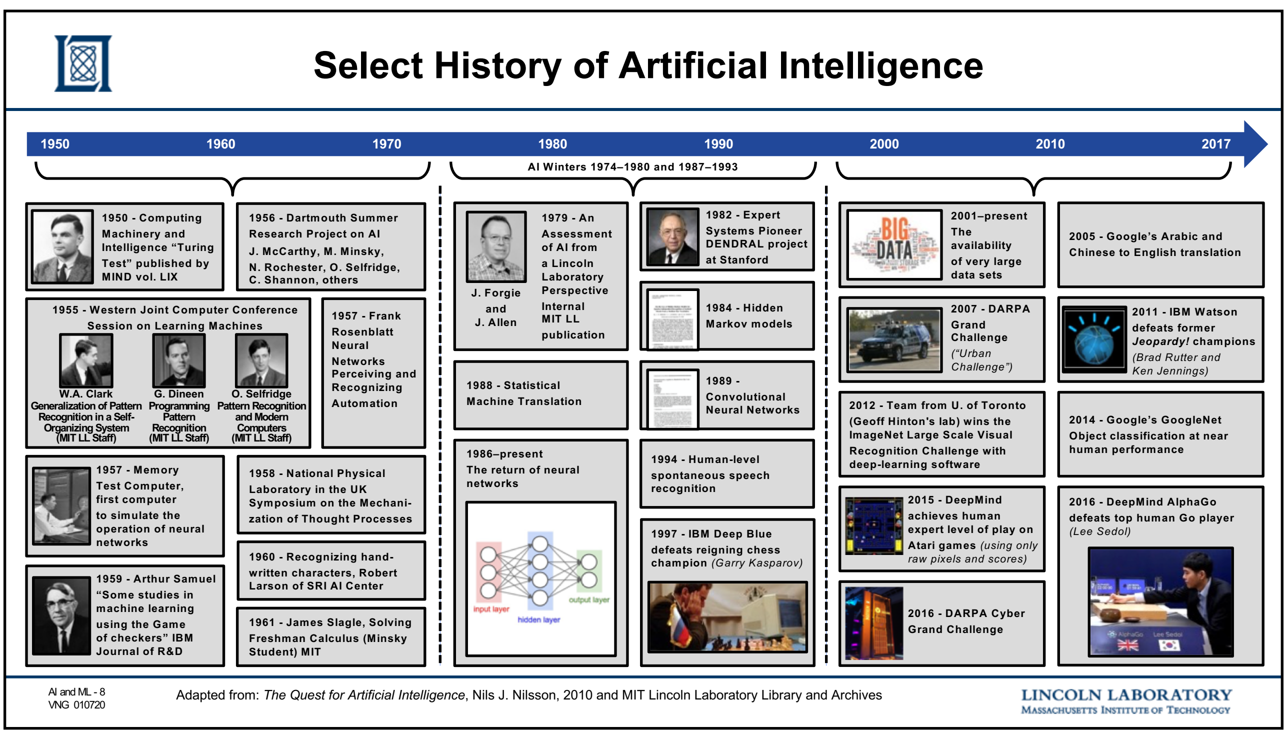
Select History of Artificial Intelligence
1950
1960
1970
1980
1990
2000
2010
2017
AI Winters 1974–1980 and 1987–1993
1950 - Computing Machinery and Intelligence "Turing Test" published by MIND vol. LIX
1956 - Dartmouth Summer Research Project on AI
J. McCarthy, M. Minsky, N. Rochester, O. Selfridge, C. Shannon, others
1955 - Western Joint Computer Conference Session on Learning Machines
W.A. Clark Generalization of Pattern Recognition in a Self-Organizing System (MIT LL Staff)
G. Dineen Programming Pattern Recognition (MIT LL Staff)
O. Selfridge Pattern Recognition and Modern Computers (MIT LL Staff)
1957 - Frank Rosenblatt Neural Networks Perceiving and Recognizing Automation
1957 - Memory Test Computer, first computer to simulate the operation of neural networks
1958 - National Physical Laboratory in the UK Symposium on the Mechanization of Thought Processes
1960 - Recognizing hand-written characters, Robert Larson of SRI AI Center
1959 - Arthur Samuel "Some studies in machine learning using the Game of checkers" IBM Journal of R&D
1961 - James Slagle, Solving Freshman Calculus (Minsky Student) MIT
J. Forgie and J. Allen
1979 - An Assessment of AI from a Lincoln Laboratory Perspective
Internal MIT LL publication
1988 - Statistical Machine Translation
1986–present The return of neural networks
input layer
hidden layer
output layer
1982 - Expert Systems Pioneer DENDRAL project at Stanford
1984 - Hidden Markov models
1989 - Convolutional Neural Networks
1994 - Human-level spontaneous speech recognition
1997 - IBM Deep Blue defeats reigning chess champion (Garry Kasparov)
2001–present The availability of very large data sets
2007 - DARPA Grand Challenge ("Urban Challenge")
2012 - Team from U. of Toronto (Geoff Hinton's lab) wins the ImageNet Large Scale Visual Recognition Challenge with deep-learning software
2015 - DeepMind achieves human expert level of play on Atari games (using only raw pixels and scores)
2016 - DARPA Cyber Grand Challenge
2005 - Google's Arabic and Chinese to English translation
2011 - IBM Watson defeats former Jeopardy! champions (Brad Rutter and Ken Jennings)
2014 - Google's GoogleNet Object classification at near human performance
2016 - DeepMind AlphaGo defeats top human Go player (Lee Sedol)
AI and ML - 8
VNG 010720
Adapted from: The Quest for Artificial Intelligence, Nils J. Nilsson, 2010 and MIT Lincoln Laboratory Library and Archives
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Artificial Intelligence Evolution

* Waves adapted from John Launchbury, Director I2O, DARPA

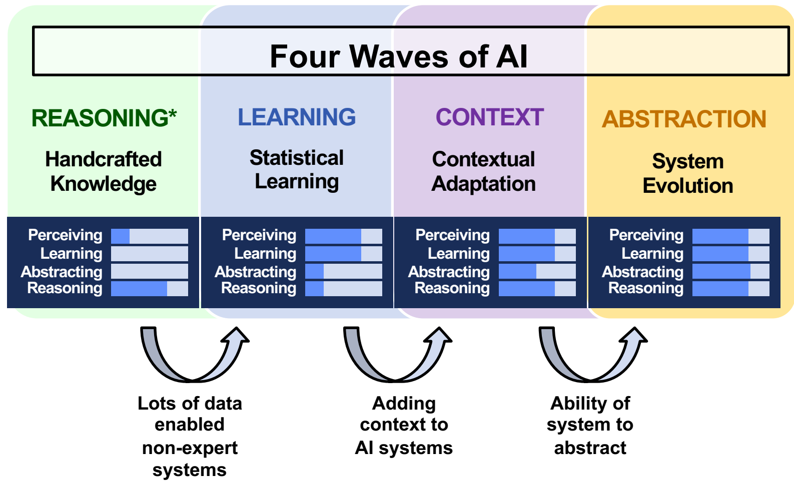

Spectrum of Commercial Organizations in the Machine Intelligence Field

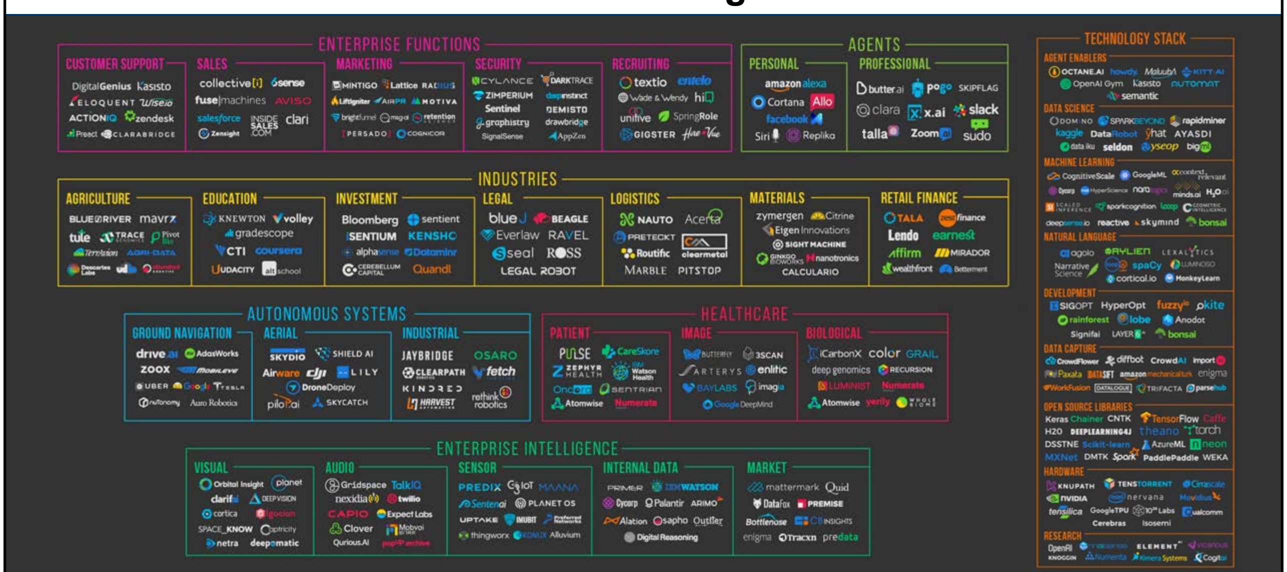

Source: Shivon Zilis, 2016, https://www.oreilly.com/ideas/the-current-state-of-machine-intelligence-3-0

Data is Critical To Breakthroughs in AI

Year	Breakthroughs in AI	Datasets (First Available)	Algorithms (First Proposed)
1994	Human-level read-speech recognition	Spoken Wall Street Journal articles and other texts (1991)	Hidden Markov Model (1984)
1997	IBM Deep Blue defeated Garry Kasparov	700,000 Grandmaster chess games, aka "The Extended Book" (1991)	Negascout planning algorithm (1983)
2005	Google's Arabic- and Chinese-to-English translation	1.8 trillion tokens from Google Web and News pages (collected in 2005)	Statistical machine translation algorithm (1988)
2011	IBM Watson became the world Jeopardy! champion	8.6 million documents from Wikipedia, Wiktionary, Wikiquote, and Project Gutenberg (updated in 2010)	Mixture-of-Experts algorithm (1991)
2014	Google's GoogleNet object classification at near-human performance	ImageNet corpus of 1.5 million labeled images and 1,000 object categories (2010)	Convolutional neural network algorithm (1989)
2015	Google's Deepmind achieved human parity in playing 29 Atari games by learning general control from video	Arcade Learning Environment dataset of over 50 Atari games (2013)	Q-learning algorithm (1992)
Average No. of Years to Breakthrough:		**3 years**	**18 years**

Source: Train AI 2017, https://www.crowdflower.com/train-ai/

AI Canonical Architecture

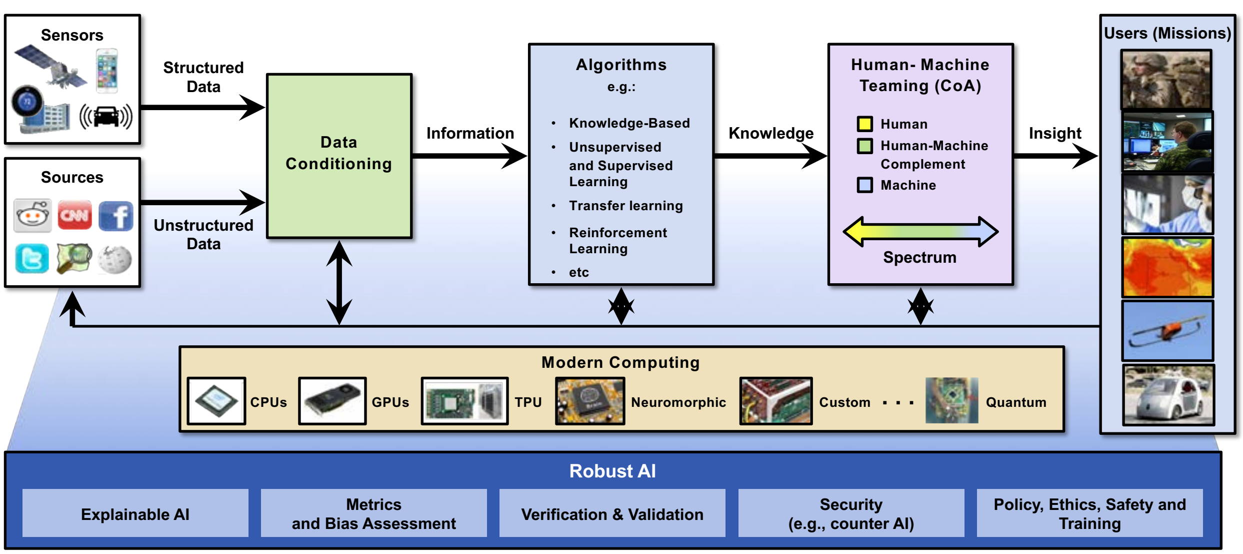

GPU = Graphics Processing Unit CoA = Courses of Action
TPU = Tensor Processing Unit

Unstructured and Structured Data

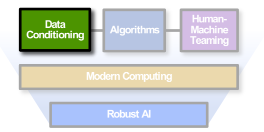

Structured Data Types

Speech | Sensors | Network Logs | Metadata

Unstructured Data Types

Social Media | Human Behavior | Reports | Side Channel

Data Conditioning/Storage Technologies

- Data to Information -

Technologies	Capabilities Provided
Infrastructure/Databases (ORACLE, accumulo)	• Indexing/Organization/Structure • Domain Specific Languages • High Performance Data Access • Declarative Interfaces
Data Curation	• Unsupervised machine learning • Dimensionality Reduction • Clustering/Pattern Recognition • Outlier Detection
Data Labeling	• Initial data exploration • Highlight missing or incomplete data • Reorient sensors/recapture data • Look for errors/biases in collection

Often takes up 80+% of overall AI/ML development work

Machine Learning Algorithms Taxonomy

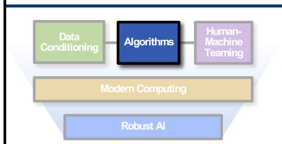

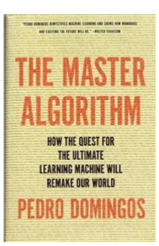

Algorithms*

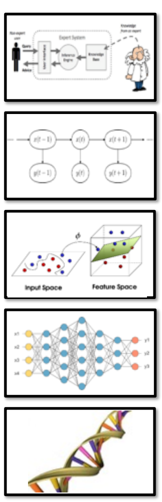

- Symbolists (e.g., exp. sys.)
- Bayesians (e.g., naive Bayes)
- Analogizers (e.g., SVM)
- Connectionists (e.g., DNN)
- Evolutionaries (e.g., genetic programming)

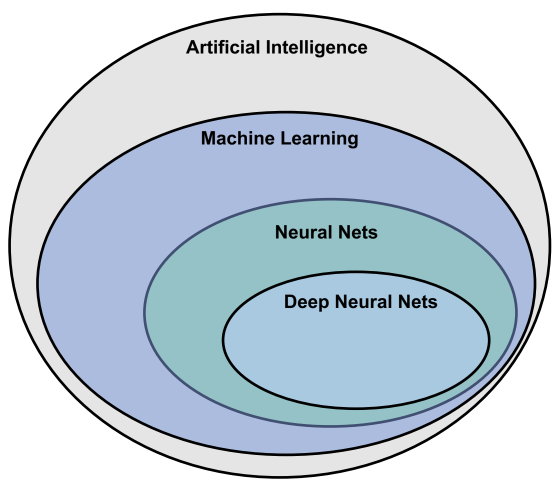

* "The Five Tribes of Machine Learning", Pedro Domingos

DNN = Deep Neural Networks
SVM = Support Vector Machines
Exp. Sys. = Expert Systems

Image Adapted From "Deep Learning" by Ian Goodfellow, Yoshua Bengio and Aaron Courville

Modern AI Computing Engines

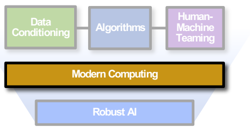

Computing Class

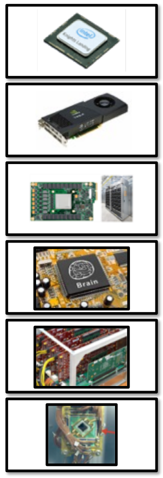

Computing Class	What It Provides to AI
CPU	• Most popular computing platform • General purpose compute
GPU	• Used by most for training algorithms (good for NN backpropagation)
TPU	• Speeds up inference time (domain specific architecture)
Neuromorphic	• Active research area
Custom	• Ability to speed up specific computations of interest (e.g. graphs)
⋮	
Quantum	• Benefits unproven until now • Recent results on HHL (linear system of equations)

Selected Results

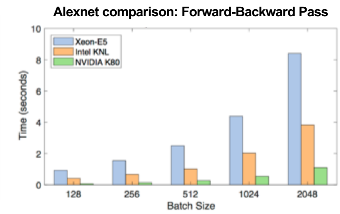

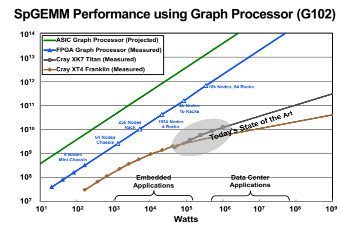

GPU = Graphics Processing Unit
TPU = Tensor Processing Unit
HHL = Harrow-Hassidim-Lloyd quantum algorithm

Neural Network Processing Performance

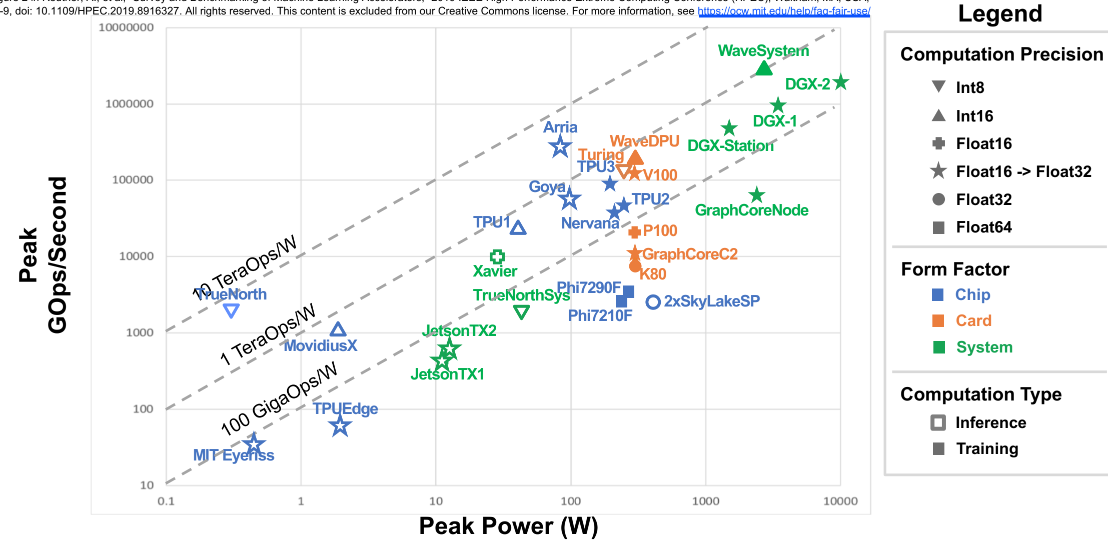

Reuther, Albert, et al. "Survey and Benchmarking of Machine Learning Accelerators." *arXiv preprint arXiv:1908.11348* (2019).

Robust AI: Preserving Trust

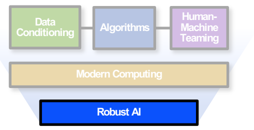

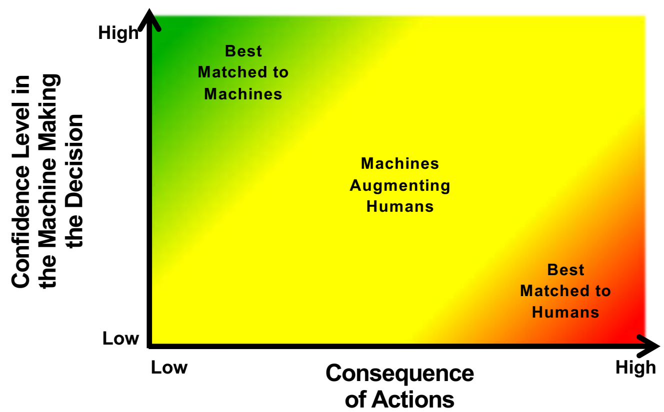

Importance of Robust AI

Robust AI Feature	Issue	Example	Solutions
Explainable AI	User unfamiliarity or mistrust leads to lack of adoption		Seamless integration, model expansion, transparent uncertainty
Metrics	Unknown relationship between arbitrary input and machine output	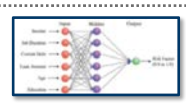	Explainability, dimensionality reduction, feature importance inference
Validation & Verification	Algorithms need to meet mission specifications		Robust training, "portfolio" methods, regularization
Security	System vulnerable to adversarial action (both cyber and physical)	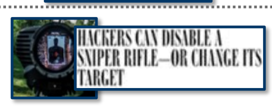	Model failure detection, red teaming
Policy, Ethics, Safety, and Training	Unwanted actions when controlling heavy or dangerous machinery		Risk sensitivity, robust inference, high decision thresholds

Human-Machine Teaming

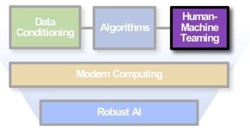

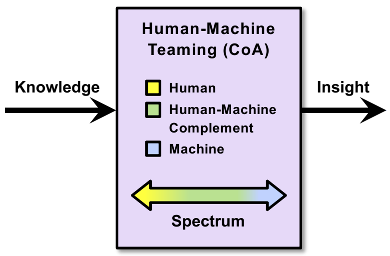

Human-Machine teaming will consist of intelligent assistants enabled by artificial intelligence

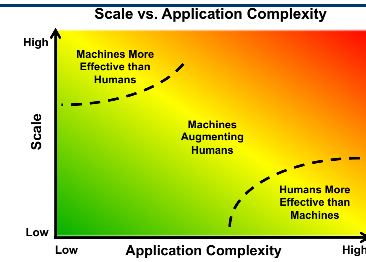

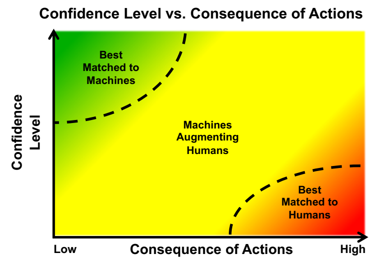

Critical Element of AI: Understanding how humans and machines can work together for applications

Outline

- Artificial Intelligence Overview
- Machine Learning Deep Dives
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
- Conclusions/Summary

What is Machine Learning?

- Machine Learning
  – Study of algorithms that improve their performance at some task with experience (data)
  – Optimize based on performance criterion using example data or past experience
- Combination of techniques from statistics, computer science communities
- Getting computers to program themselves
- Common tasks:
  – Classification
  – Regression
  – Prediction
  – Clustering
  – …

Traditional Programming vs. Machine Learning

Traditional Programming

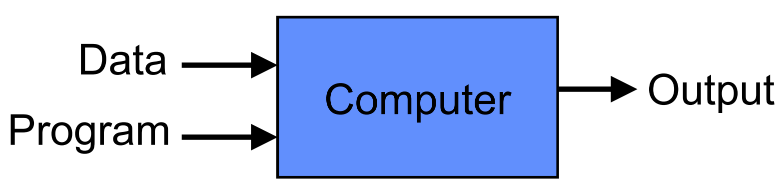

Machine Learning

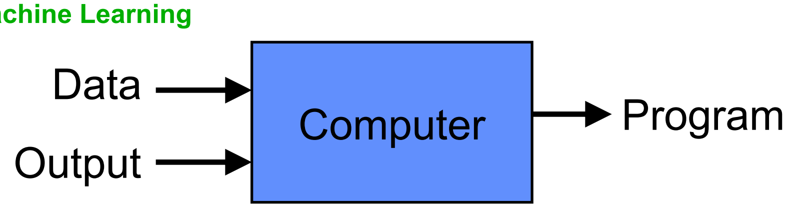

Source: Pedro Domingos, University of Washington

Machine Learning Techniques

Supervised (Labels)	Unsupervised (No Labels)	Reinforcement (Reward Information)

AI and ML - 23
VNG 010720

Machine Learning Techniques

Supervised (Labels)	Unsupervised (No Labels)	Reinforcement (Reward Information)

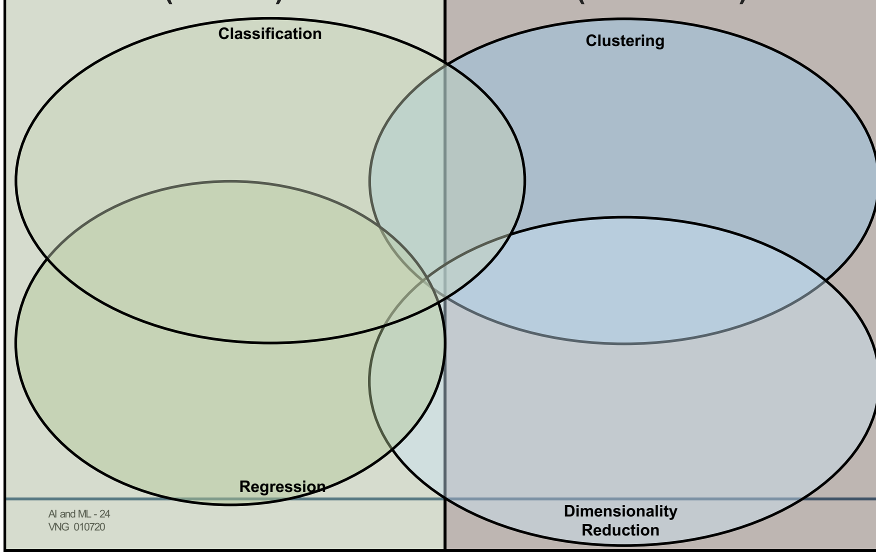

AI and ML - 24
VNG 010720

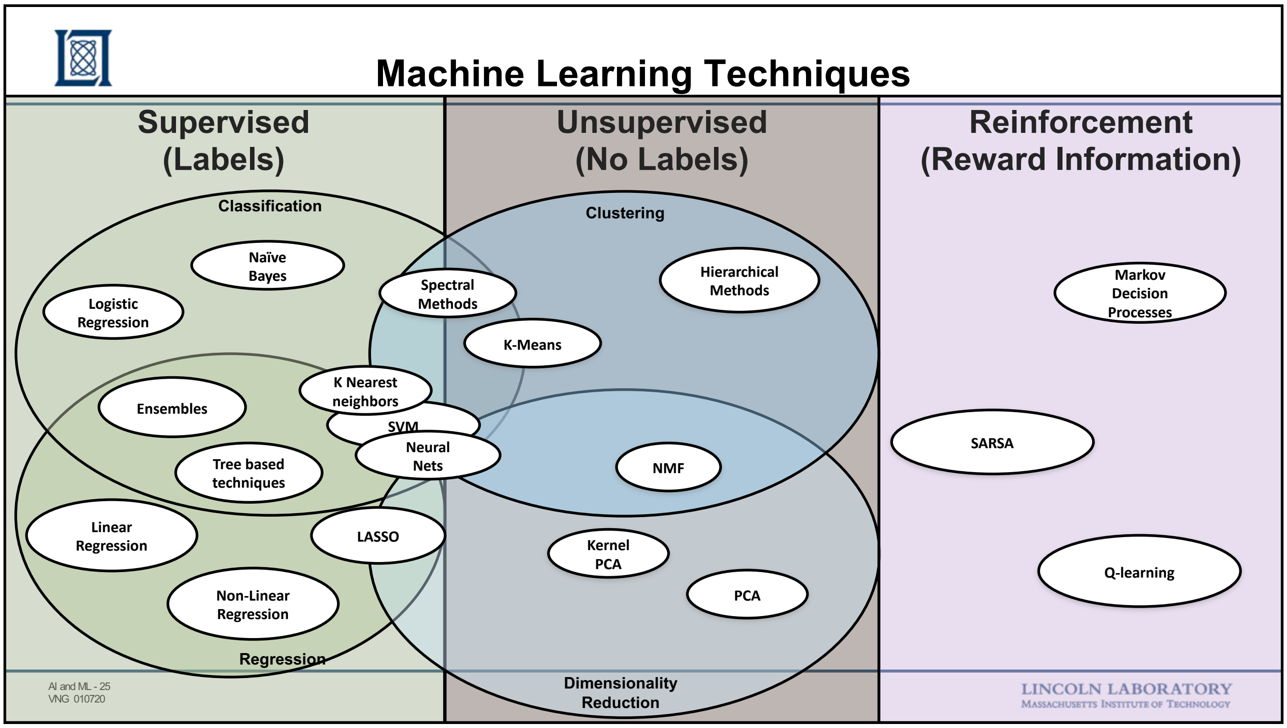
Machine Learning Techniques
Supervised (Labels)
Unsupervised (No Labels)
Reinforcement (Reward Information)
Classification
Naïve Bayes
Logistic Regression
Ensembles
K Nearest neighbors
SVM
Neural Nets
Tree based techniques
Linear Regression
LASSO
Non-Linear Regression
Regression
Clustering
Spectral Methods
Hierarchical Methods
K-Means
NMF
Kernel PCA
PCA
Dimensionality Reduction
Markov Decision Processes
SARSA
Q-learning
AI and ML - 25
VNG 010720
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

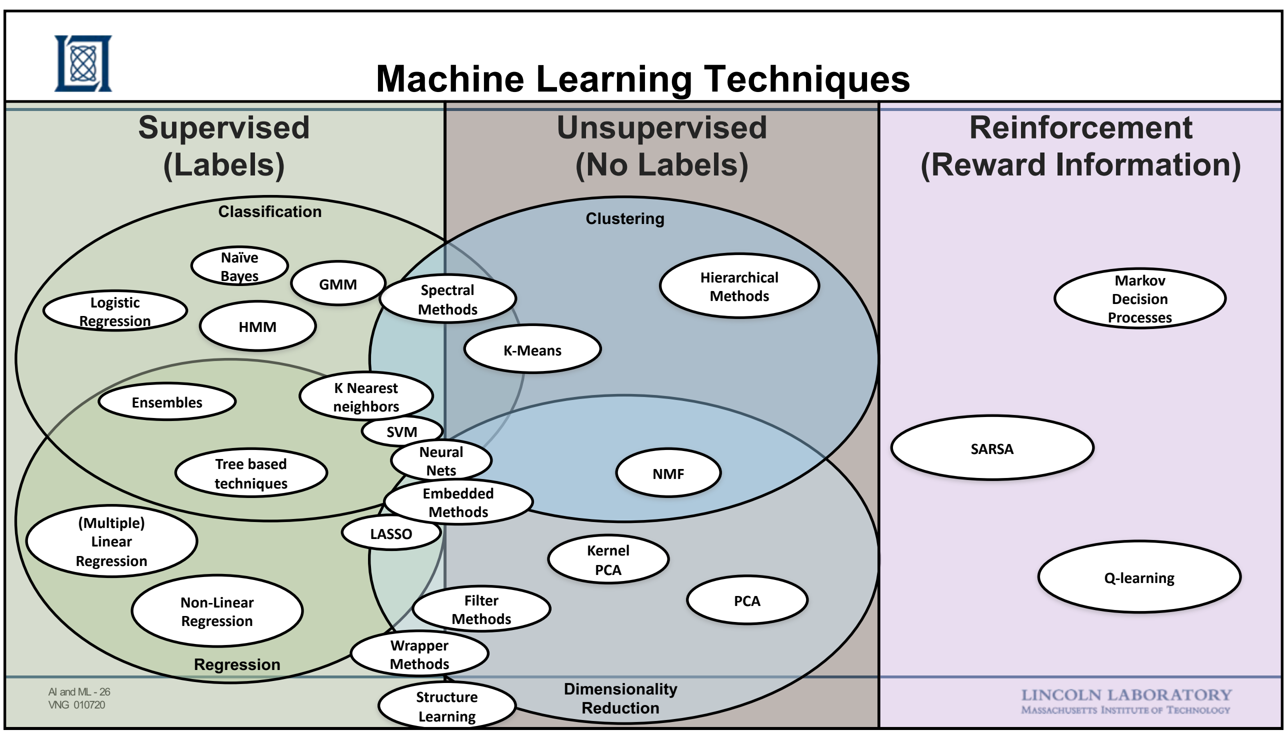
Machine Learning Techniques
Supervised (Labels)
Unsupervised (No Labels)
Reinforcement (Reward Information)
Classification
Naïve Bayes
GMM
Logistic Regression
HMM
Ensembles
K Nearest neighbors
SVM
Neural Nets
Tree based techniques
Embedded Methods
(Multiple) Linear Regression
LASSO
Non-Linear Regression
Filter Methods
Wrapper Methods
Regression
Structure Learning
Clustering
Spectral Methods
Hierarchical Methods
K-Means
NMF
Kernel PCA
PCA
Dimensionality Reduction
Markov Decision Processes
SARSA
Q-learning
AI and ML - 26
VNG 010720
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Common ML Pitfalls

- Over-fitting vs. Under-fitting
- Bad/noisy/missing data
- Model selection
- Lack of success metrics
- Linear vs. Non-linear models
- Ignoring outliers
- Training vs. testing data
- Computational complexity, curse of dimensionality
- Correlation vs. Causation

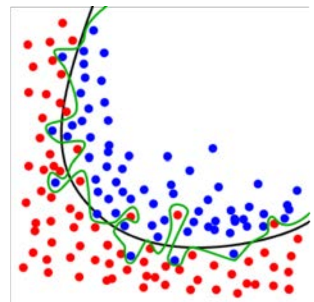

Outline

- Artificial Intelligence Overview
- Machine Learning Deep Dives
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
- Conclusions/Summary

Supervised Learning

- **Starting with labeled data (ground truth)**
- **Build a model that predicts labels**
- **Two general goals:**
  - **Regression: predict continuous variable**
  - **Classification: predict a class or label**
- **Generally has a training step that forms the model**

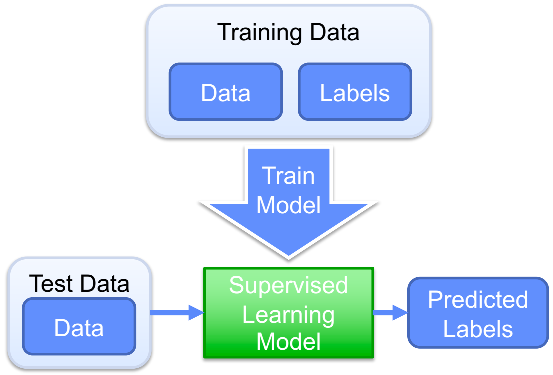

Artificial Neural Networks

- **Computing systems inspired by biological networks**
- **Systems learn by repetitive training to do tasks based on examples**
  - **Generally a supervised learning technique (though unsupervised examples exist)**
- **Components: Inputs, Layers, Outputs, Weights**
- **Deep Neural Network: Lots of "hidden layers"**
- **Popular variants:**
  - **Convolutional Neural Nets**
  - **Recursive Neural Nets**
  - **Deep Belief Networks**
- **Very popular these days with many toolboxes and hardware support**

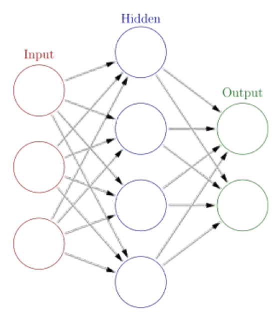

Deep Neural Networks

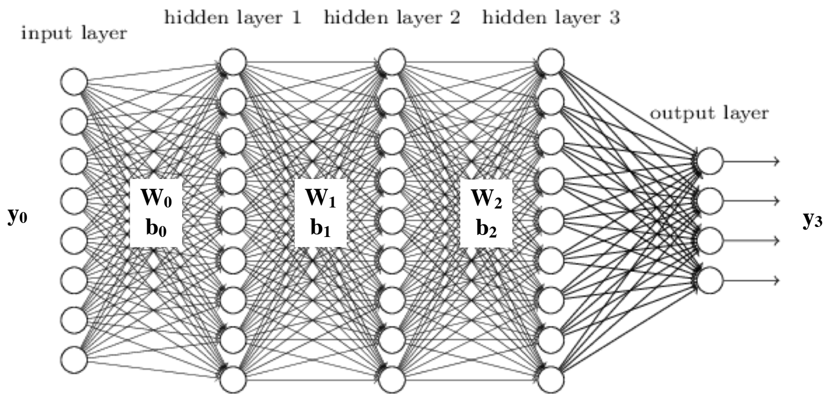

$$y_{i+1} = f(W_i\, y_i + b_i)$$

Image source: http://www.rsipvision.com/exploring-deep-learning/

Components of an Artificial Neural Network

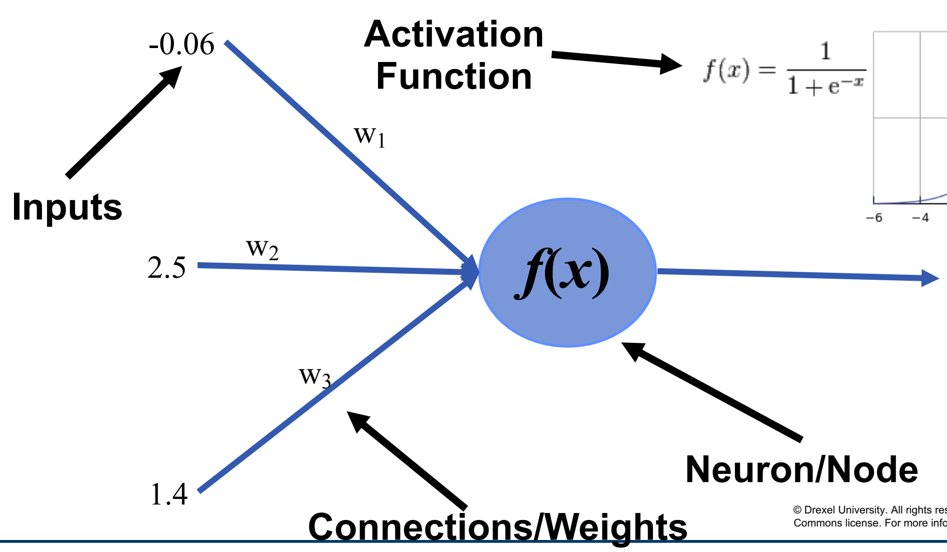

$$f(x) = \frac{1}{1 + e^{-x}}$$

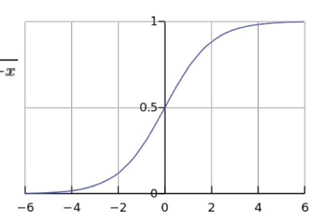

Image source: https://www.cs.drexel.edu/~greenie/cs510/CS510-17-08.pdf

Components of an Artificial Neural Network

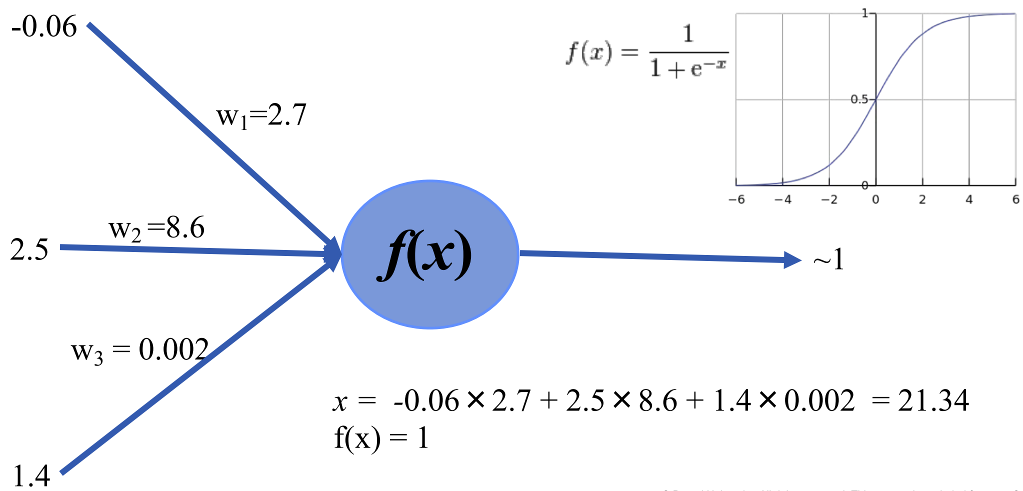

Image source: https://www.cs.drexel.edu/~greenie/cs510/CS510-17-08.pdf

Common Activation Functions

- **Step Function:** $f(x) = \begin{cases} 0, x < 0 \\ 1, x \geq 0 \end{cases}$

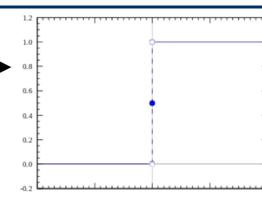

- **Sigmoid Function:** $f(x) = \frac{1}{1+e^{-x}}$

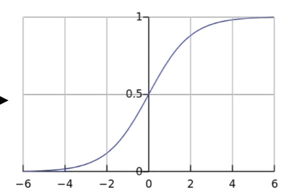

- **Tanh Function:** $f(x) = tanh(x)$

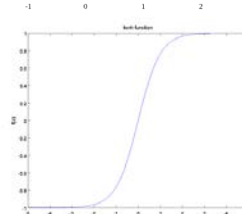

- **Rectified Linear Unit (ReLU):** $f(x) = max(0, x)$

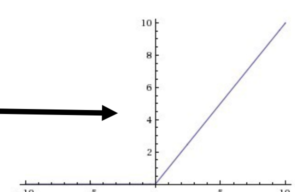

Neural Network Landscape

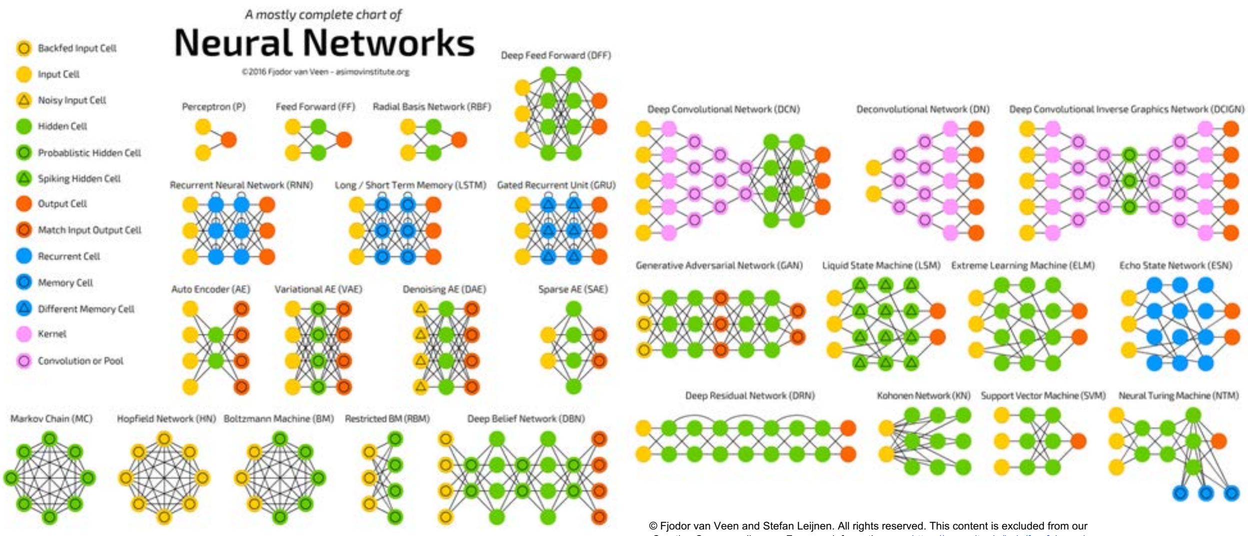

Image Source: http://www.asimovinstitute.org/neural-network-zoo/

Neural Network Training

***Key Idea*: Adjusting the weights changes the function represented by the neural network (*learning = optimization in weight space*).**

Iteratively *adjust weights* to reduce *error* (difference between network output and target output)

Weight Update

- ***perceptron training rule***
- ***linear programming***
- ***delta rule***
- ***Backpropagation***

Real neural network architectures can have 1000s of input data points, hundreds of layers and millions of weight changes per iteration

Neural Network Inference

- **Using the trained model on previously "unlabeled" data**

Neural Network Learning: Decision Boundary

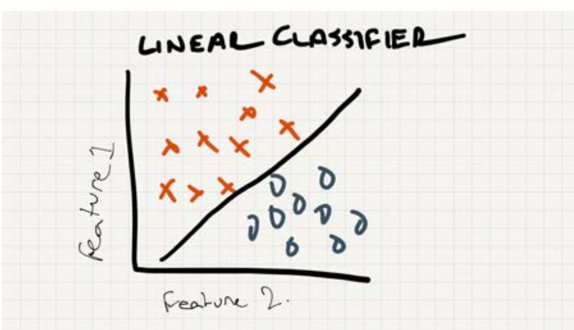

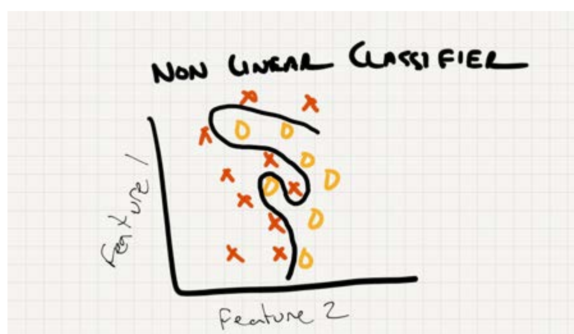

Designing a Neural Network

- Designing a neural network can be a complicated task.
- Many choices:

- Depth (number of layers)
- Inputs (number of inputs)
- Type of Network:
  - Convolutional Neural Network
  - Deep Feedforward Neural Network
  - Deep Belief Network
  - Long/Short Term Memory
  - …

- Types of layers:
  - MaxPool
  - Dropout
  - Convolutional
  - Deconvolutional
  - Softmax
  - Fully Connected
  - Skip Layer
  - …

- Training Algorithm
  - Performance vs. Quality
  - Stopping criteria
  - Performance function
- Metrics:
  - False positive
  - ROC curve
  - …

Outline

- Artificial Intelligence Overview
- Machine Learning Deep Dives
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
- Conclusions/Summary

Unsupervised Learning

- Task of describing hidden structure from unlabeled data
- More formally, we observe features $X_1, X_2, \ldots, X_n$ and would like to observe patterns among these features.
  - We are not interested in predcition because we don't know what an output *Y* would look like.
- Typical tasks and associated algorithms:
  - Clustering
  - Data projection/Preprocessing
- Goal is to discover interesting things about the dataset: subgroups, patterns, clusters?

More on Unsupervised Learning

- There is no good simple goal (such as maximizing certain probability) for the algorithm
- Very popular because techniques work on unlabeled data
  - Labeled data can be difficult and expensive
- Common techniques:
  - Clustering
    - K-Means
    - Nearest neighbor search
    - Spectral clustering
  - Data projection/preprocessing
    - Principal component analysis
    - Dimensionality Reduction
    - Scaling

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Clustering

- **Group objects or sets of features such that objects in the same cluster are more similar than those of another cluster**
- **Optimal clusters should**
  - **Minimize intra-cluster distance**
  - **Maximize inter-cluster distance**
- **Example of intra-cluster measure**
  - **Squared error se**

$$se = \sum_{i=1}^{k} \sum_{p \in c_i} \| p - m_i \|^2$$

where m_i is the mean of all features in cluster c_i

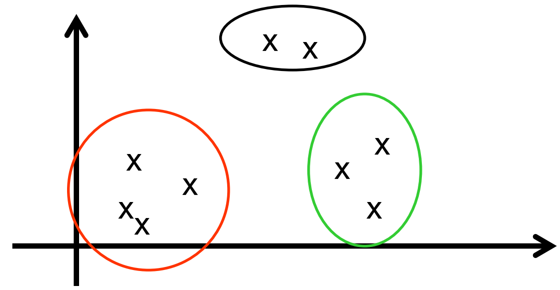

AI and ML - 43
VNG 010720

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Dimensionality Reduction

- **Process of reducing number of random variables under consideration**
  - **Key idea: Reduce large dataset to much smaller dataset using only high variance dimensions**
- **Often used to simplify computation or representation of a dataset**
- **Typical tasks:**
  - **Feature Selection: try to find a subset of original variables**
  - **Feature Extraction: try to represent data in lower dimensions**
- **Often key to good performance for other machine learning techniques such as regression, classification, etc.**
- **Other uses:**
  - **Compression: reduce dataset to smaller representation**
  - **Visualization: easy to see low dimensional data**

AI and ML - 44
VNG 010720

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Neural Networks and Unsupervised Learning

- Traditional applications of neural networks such as Image classification fall into the realm of supervised learning:
  - Given example inputs x and target output y, learn the mapping between them.
  - A trained network is supposed to give the correct target output for any input stimulus
  - Training is learning the weights
- Largely used to find better representations for data: clustering and dimensionality reduction
- Non linear capabilities

Example: Autoencoders

- Neural network architecture designed to find a compressed representation for data
- Feedforward, multi layer perceptron.
- Input layer number of features = output layer number of features
- Similar to dimensionality reduction but allows for much more complex representations

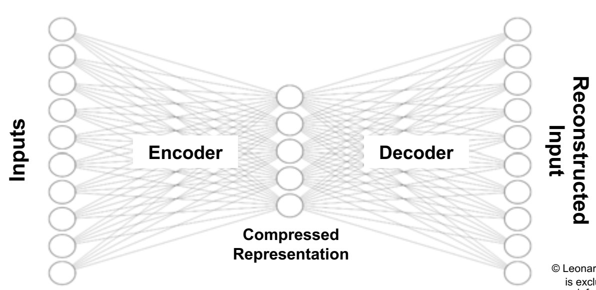

Example: Replicator Neural Network

- Conceptually, very similar to autoencoders
- Used extensively for anomaly detection (looking for outliers)
- Example architecture

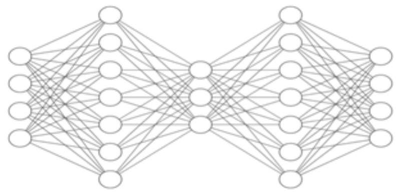

- Salient differences from an autoencoder: Step Activation Function, Inclusion of dropout layers
  - Step activation squeezes the middle layer outputs into a number of clusters
  - Dropout layers help with overfitting

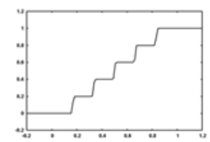

Step Activation Function

Hawkins, Simon, et al. "Outlier detection using replicator neural networks." *International Conference on Data Warehousing and Knowledge Discovery*. Springer, Berlin, Heidelberg, 2002.

Outline

- Artificial Intelligence Overview
- Machine Learning Deep Dives
  - Supervised Learning
  - Unsupervised Learning
  - ➡ Reinforcement Learning
- Conclusions/Summary

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Reinforcement Learning

- What makes reinforcement learning different from other machine learning paradigms?
  - There is no supervisor, only a reward signal
  - Feedback is delayed, not instantaneous
  - Time really matters (sequential, often inter-dependent data)
  - Agent's actions affect the subsequent data it receives
- Example: Playing Atari game
  - Rules of the game are unknown
  - Learn directly from interactive game-play
  - Pick actions on joystick, see pixels and scores

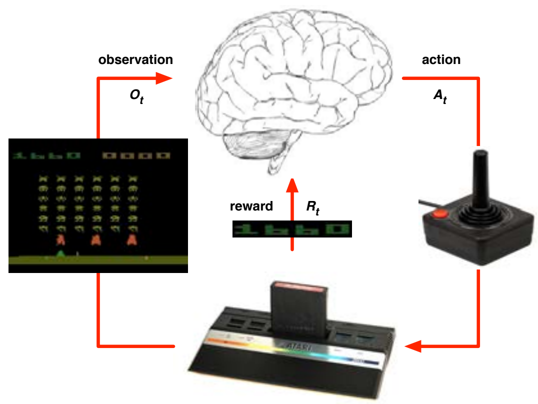

Image Source: "Introduction to Reinforcement Learning" by David Silver (Google DeepMind)

Other Reinforcement Learning Examples

- Fly stunt maneuvers in a helicopter
  - \+ reward for following desired trajectory
  - − reward for crashing
- Defeat the world champion at Backgammon
  - +/− reward for winning/losing a game
- Manage an investment portfolio
  - \+ reward for each $ in bank
- Control a power station
  - \+ reward for producing power
  - − reward for exceeding safety thresholds
- Make a humanoid robot walk
  - \+ reward for forward motion
  - − reward for falling over

Outline

- Artificial Intelligence Overview
- Machine Learning Deep Dives
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
- Conclusions/Summary

Summary

- Lots of exciting research into AI/ML techniques
  - This course looks at a number of relatively easy strategies to mitigate these challenges
- Key ingredients for AI success:
  - Data Availablity
  - Computing Infrastructure
  - Domain Expertise/Algorithms
- Large challenges in data availability and readiness for AI
- MIT SuperCloud platform (next presentation) can be used to perform the heavy computation needed
- Further reading:
  - "AI Enabling Technologies: A Survey." https://arxiv.org/abs/1905.03592

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY