AI Study: Short History, Present Developments, and Future Outlook

David R. Martinez
Associate Division Head
Cyber Security and Information Sciences Division

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

This material is based upon work supported under Air Force Contract No. FA8721-05-C-0002 and/or 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 U.S. Air Force. Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

Briefing Main Takeaways

- **Key issues with narrow AI**
  - **Value to the audience: Provide increased clarity**
- **AI canonical architecture**
  - **Value to the audience: Recommend an end-to-end system framework to identify areas to lead, adopt/adapt, or follow**
- **S&T Challenges and Investment Opportunities**
  - **Value to the audience: A structure for organizing an AI strategic development roadmap**

> **"Success no longer goes to the country that develops a new technology first, but rather to the one that better integrates it and adapts its way of fighting"**
>
> **– Dr. Eric Schmidt, House Armed Services Committee Testimony, April 17, 2018**

Outline

- **Background**
  - **Definition of Area**
  - **AI History Highlights**
- **Lay-of-the-Land**
- **Robust AI**
- **Recommendations**
- **Summary**

National Challenges and Role of AI

	National Challenges	*Role of AI in Augmenting Humans*
  **Intelligent Systems and Autonomy** **Information Superiority**	**Technological dominance in support of national security**	**Derive actionable intelligence by effective human-machine teaming**
	**Massive amounts of structured and unstructured data**	**Leverage rapid advances in data conditioning, algorithms, and computing**
	**Trust in intelligent machines (Robust AI)**	**Ascertain robustness** *"We had better be quite sure that the purpose put into the machine is the purpose which we really desire"* Norbert Wiener, 1960

Operative AI Definition for the Study

Narrow AI:

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

We will address: Narrow AI *not* General AI

* Definition adapted from Oxford dictionary and inputs from Prof. Patrick Winston (MIT) during his visit to MIT LL May 2017

AI Domain of Impact

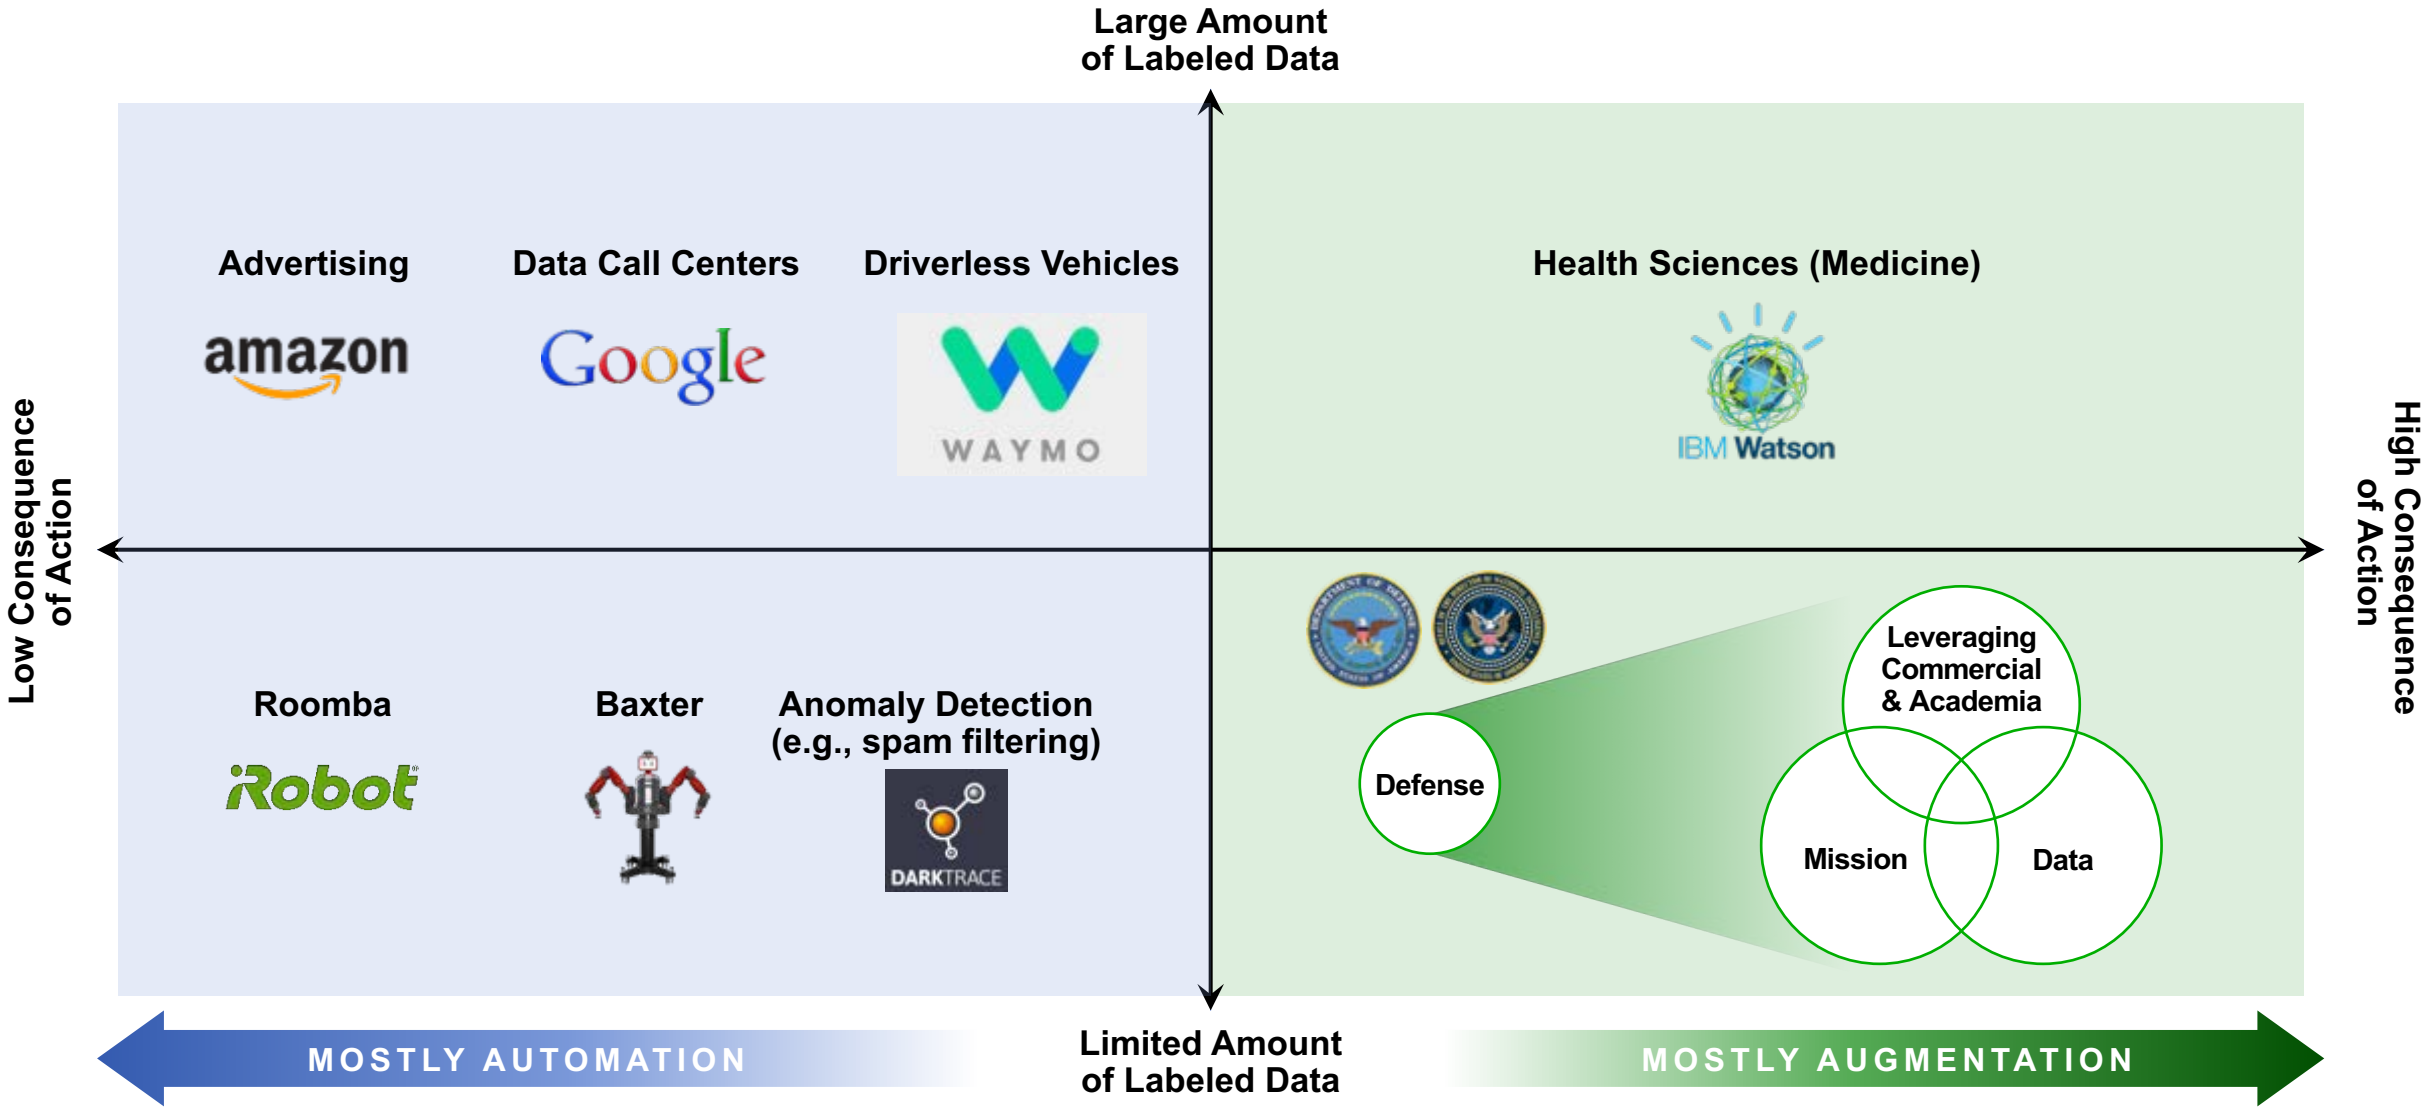

Select History of Artificial Intelligence

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

1959 - Arthur Samuel "Some studies in machine learning using the Game of checkers" IBM Journal of R&D

1960 - Recognizing hand-written characters, Robert Larson of SRI AI Center

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

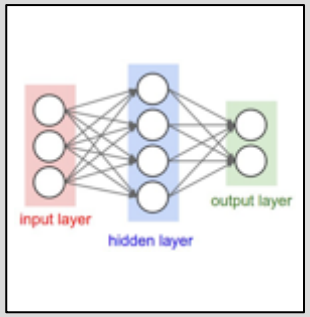

1982 – Ed Feigenbaum expert systems DENDRAL project

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)*

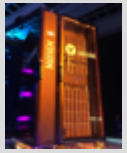

2016 - DARPA Cyber Grand Challenge

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)*

Adapted from: *The Quest for Artificial Intelligence*, Nils J. Nilsson, 2010 and MIT Lincoln Laboratory Library and Archives

Reflections During The Early Years of AI

Marvin Minsky with Claude Shannon, Oliver Selfridge and other scientists attending the Dartmouth Summer Research Project on Artificial Intelligence

The Thinking Machine
(Artificial Intelligence in the 1960s)

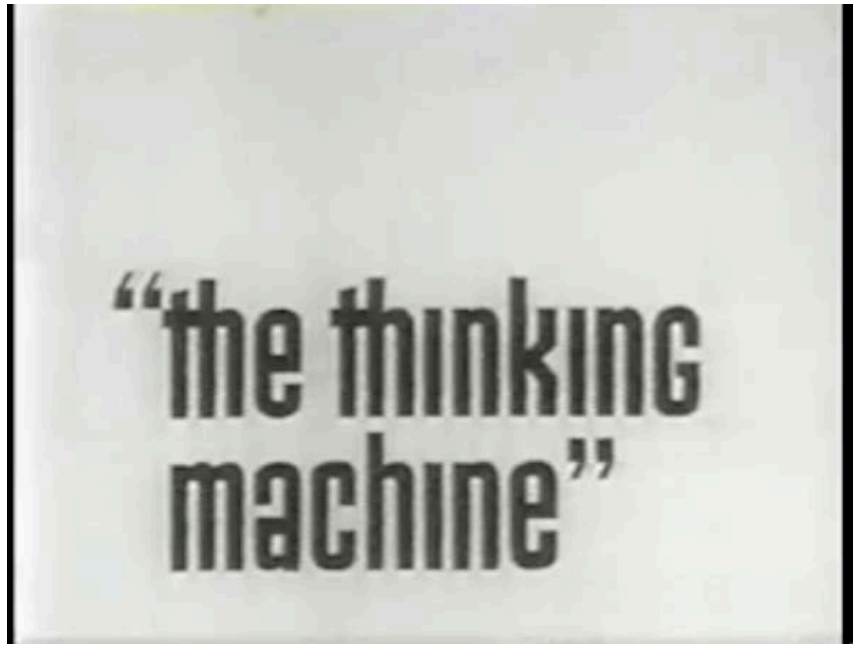

Dr. Jerome Wiesner (MIT President), Prof. Claude Shannon, and Dr. Oliver Selfridge (Grp. Ldr. at MIT LL)

Source: YouTube, posted by Roberto Pieraccini

Steps Toward Artificial Intelligence*

MARVIN MINSKY†, MEMBER, IRE

Marvin Minsky, "One of the fathers of artificial intelligence" worked at Lincoln 1956–1960

Top 15 Publishing Universities/Organization in the US (2011–2018)

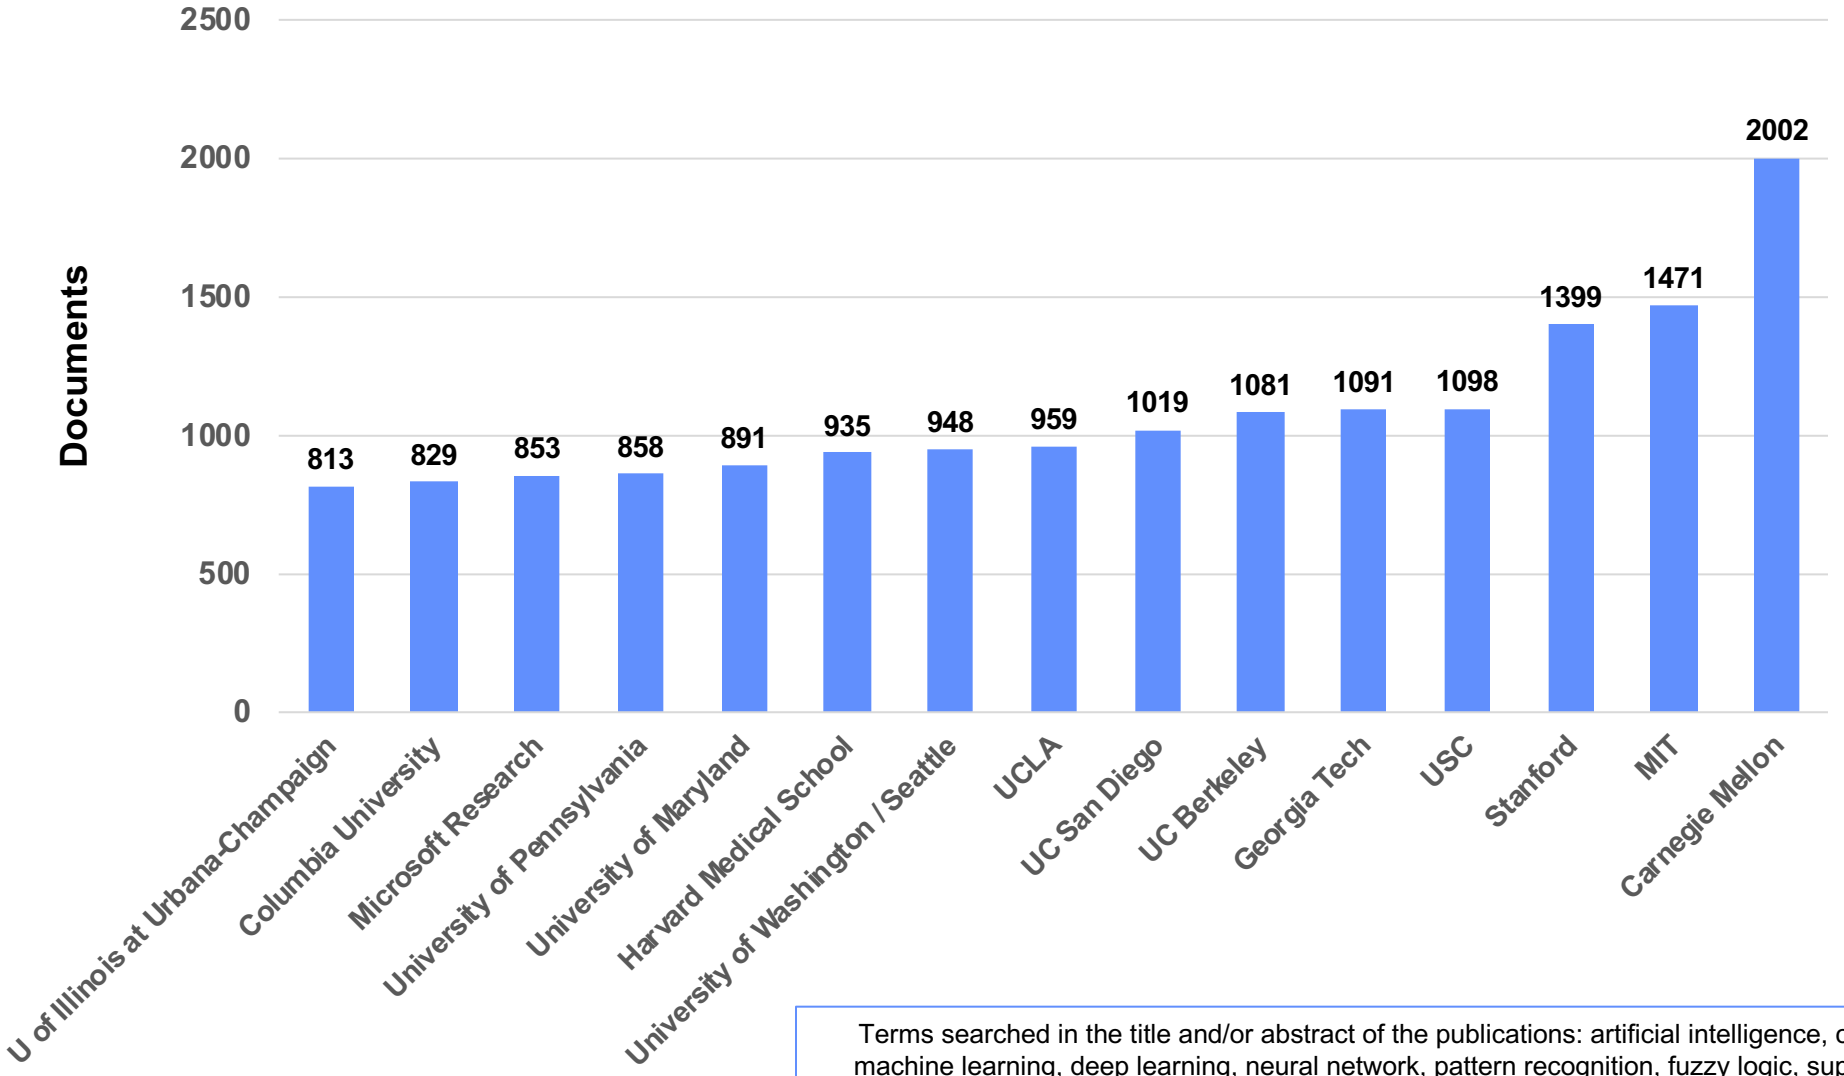

Terms searched in the title and/or abstract of the publications: artificial intelligence, cognitive computing, machine learning, deep learning, neural network, pattern recognition, fuzzy logic, support vector machine

Source: Scopus which is largest abstract and citation *database* of peer-reviewed literature: scientific *journals*, books and conference proceedings

China is Putting a Major Investment into AI

Chinese Government has indicated plans to create a $150B AI ecosystem over the next few years

The Economist (July 2017)

In **2012–16 Chinese AI firms received $2.6B in funding**, according to the Wuzhen Institute, a think-tank

China Next Generation AI Development Plan (July 2017)[1]

By 2020 … China will have established initial AI technology standards, service systems, and industrial ecological system chains … **with the scale of AI's core industry exceeding $22.6B, and exceeding $150B as driven by the scale of related industries**

MIT Technology Review (November 2017)

China's goal is "to have major breakthroughs in AI by 2025, and to be the envy of the world by 2030"[2]

DoD R&D spending is a fraction of nation states – we are losing ground on patents and publications

[1] https://www.newamerica.org/cybersecurity-initiative/blog/chinas-plan-lead-ai-purpose-prospects-and-problems/
[2] The Artificial Intelligence Issue, China's AI Awakening, MIT Technology Review, Nov-Dec 2017

Top 15 Patent Holders in AI Per Country (2011–2016)

Top 15 Patent Holders in AI Per Country of Publication (2011–2016)

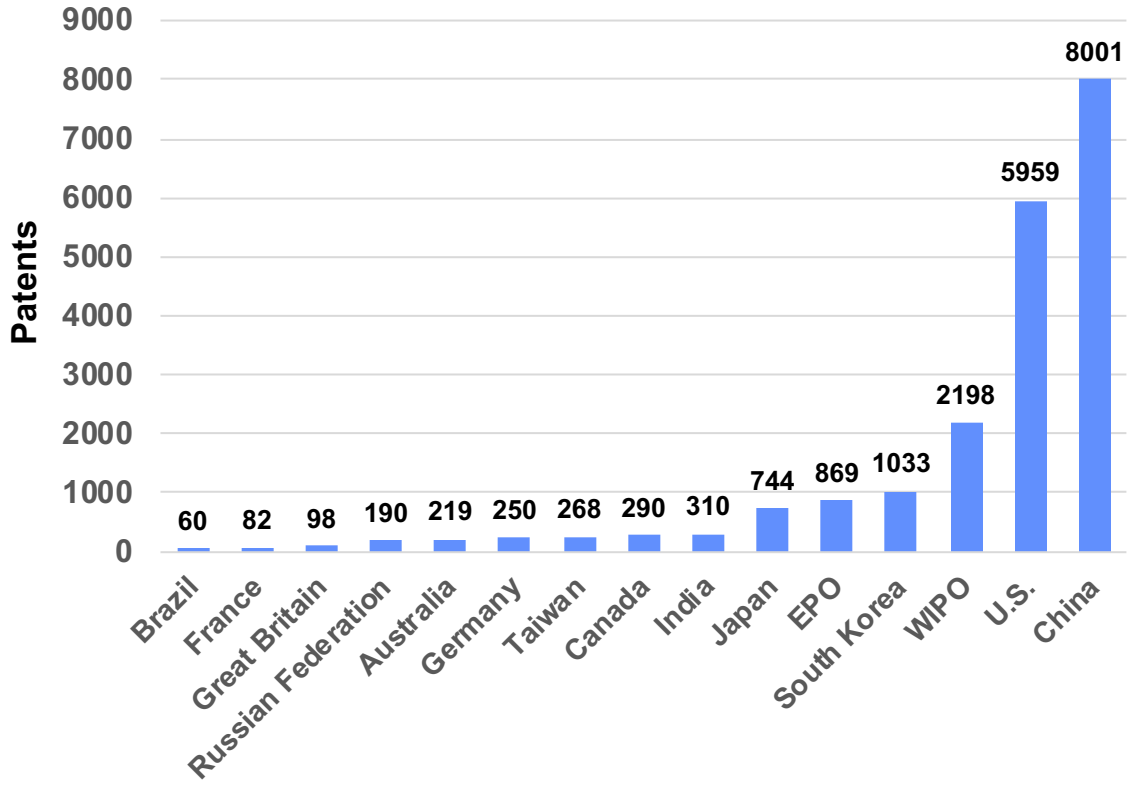

Top 15 Patent Holders in AI Per Patent Assignee (2011–2016)

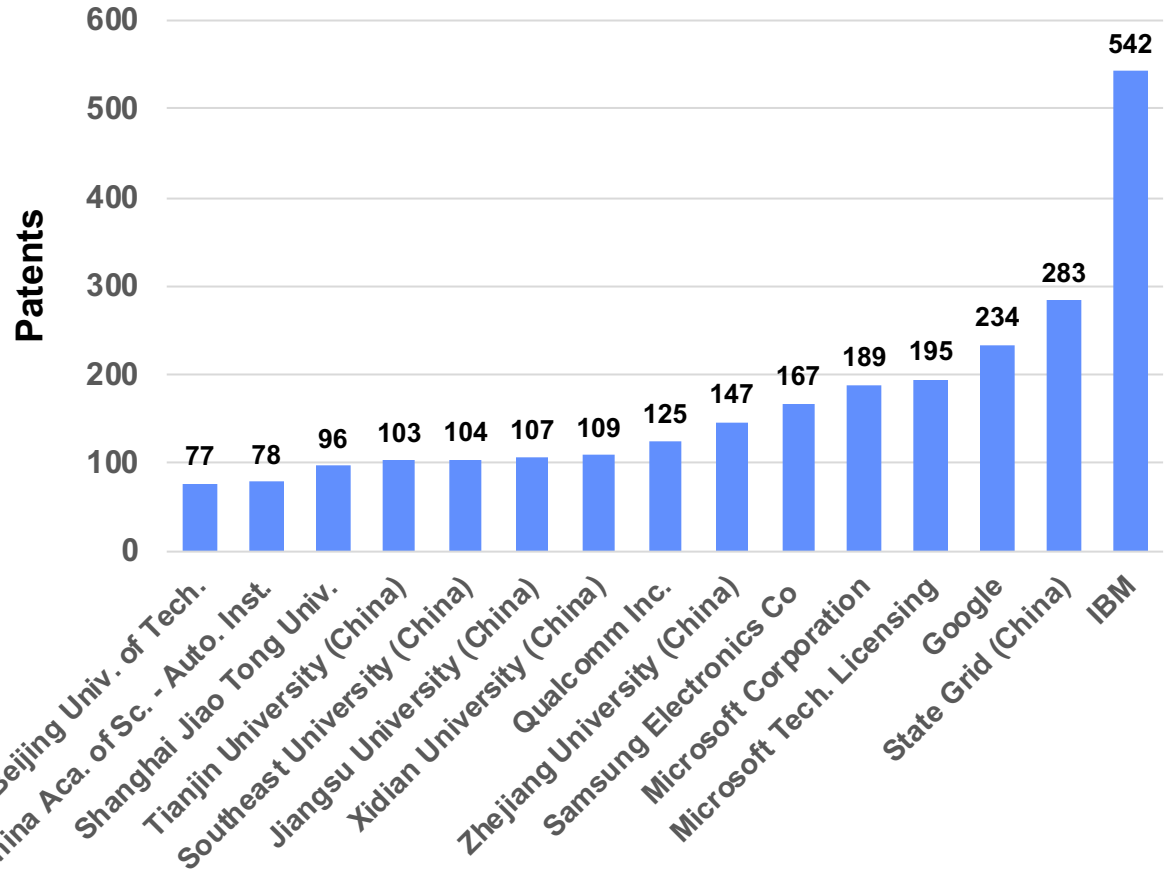

Terms searched in the title and/or abstract of the patent record: artificial intelligence, cognitive computing, machine learning, deep learning, neural network, pattern recognition, fuzzy logic, support vector machine

WIPO = World Intellectual Property Organization
Source: Derwent World Patents Index

EPO = European Patent Organization

AI Citations and Publication Influence
- From G. Fabre's Paper on China's Digital Transformation -

While China ranks first for absolute AI citations, the United States holds an edge when self-citations are taken out

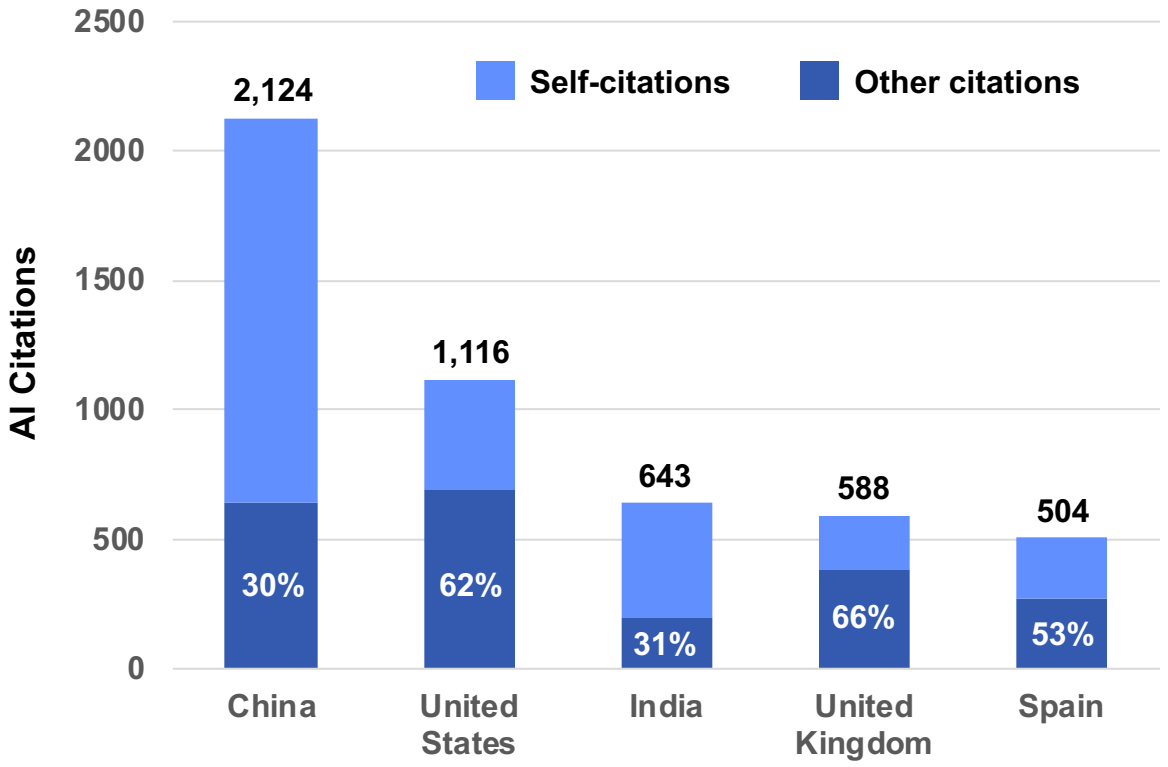

Publication influence (H-index[2])

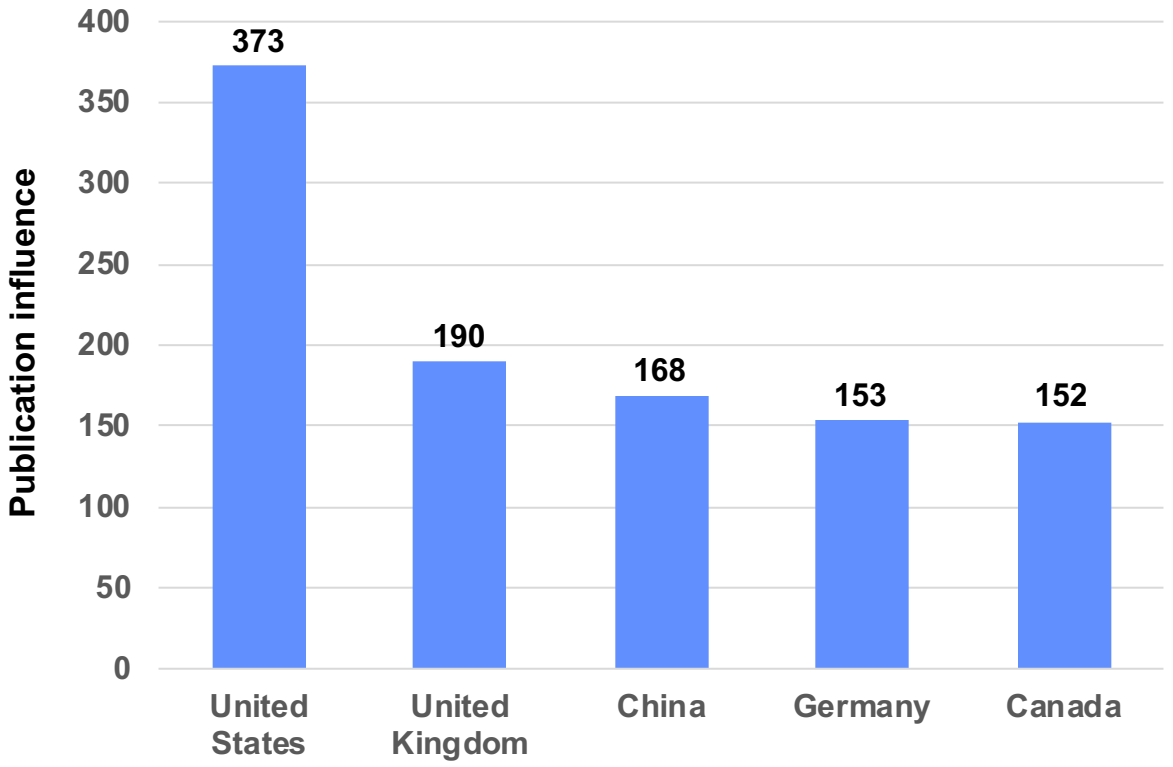

1 Self-citation occurs when a journal cites another article published in the same journal.
2 The H-index ranks both the productivity of scholars and the citation impact of their publications. A higher H-index number indicates more publications that are widely cited.

Source: SCImago Journal Rank 2015; McKinsey Global Institute analysis
Publication: G. Fabre, "China Digital Transformation, Why is AI a Priority for Chinese R&D?", HAL archives-ouvertes, June 2018

Association for the Advancement of AI (AAAI 2019)
- Statistics -

Papers by Country

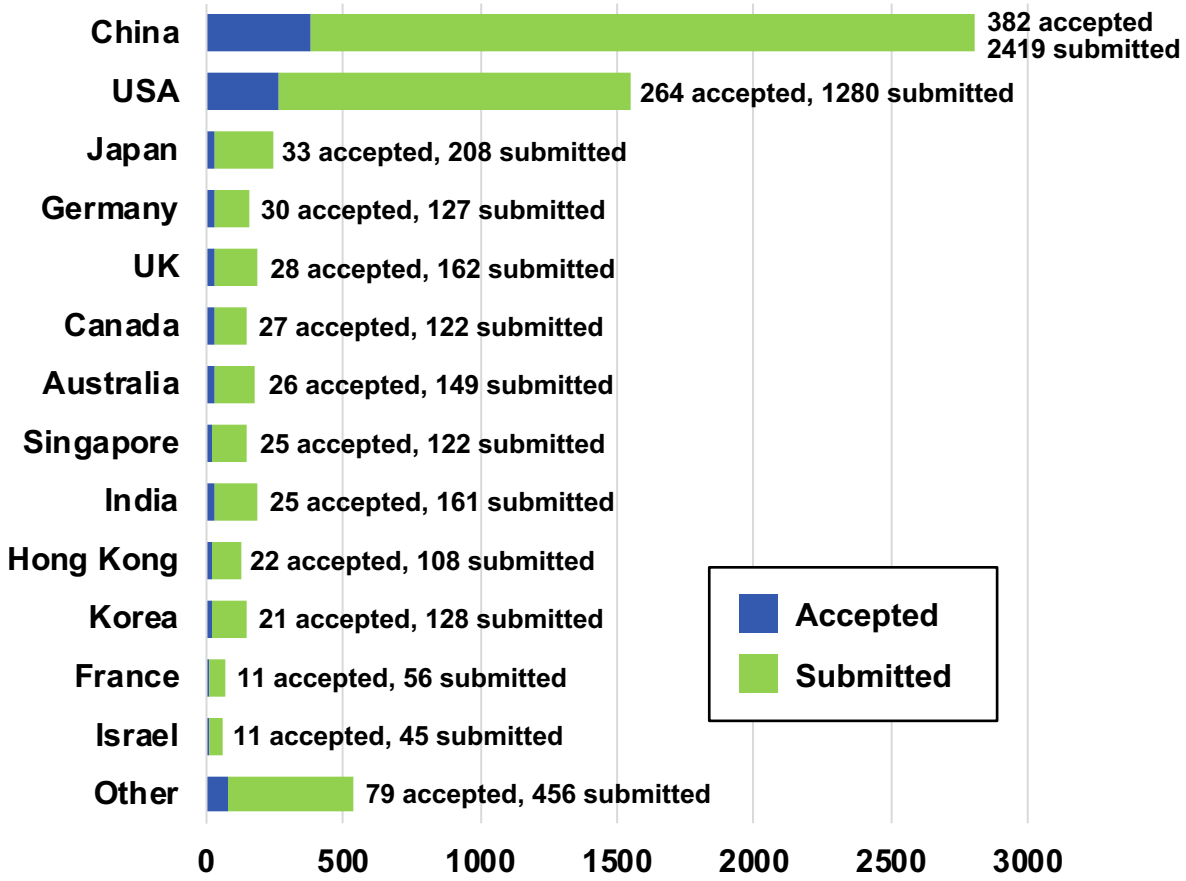

Acceptance Rate by Country

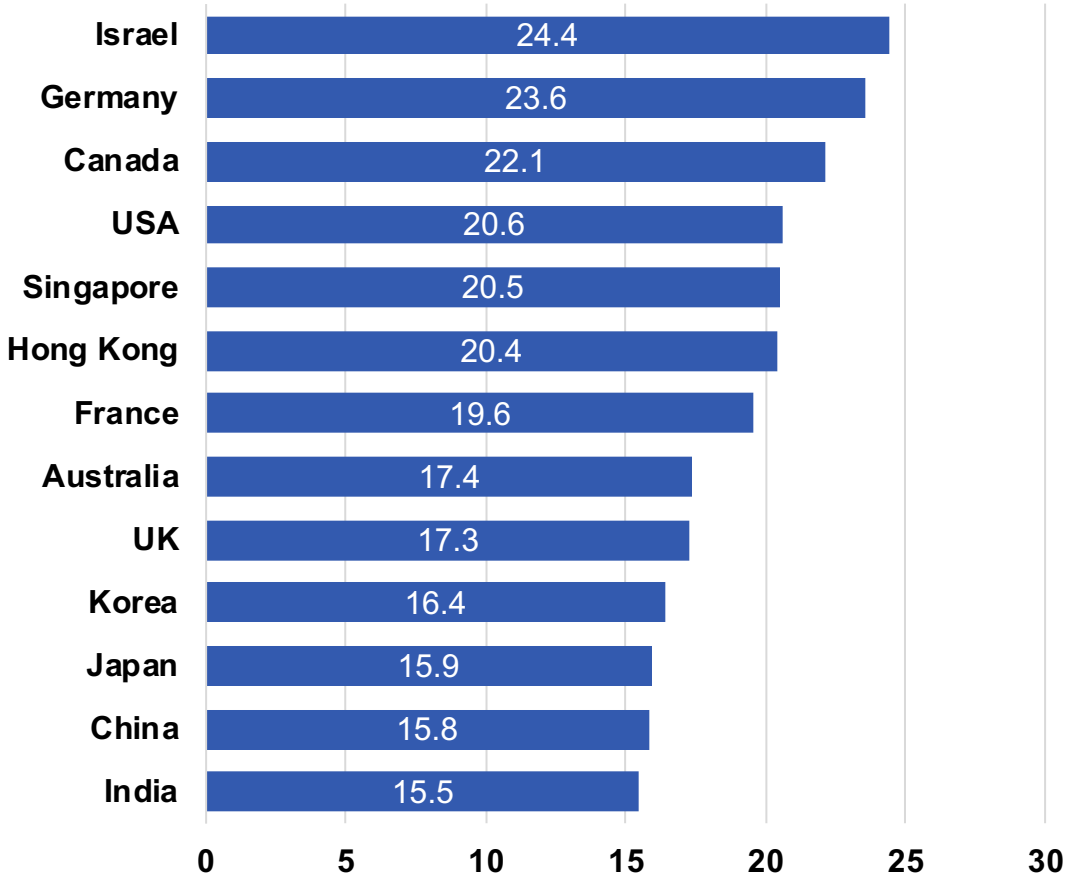

Outline

- Background
- Lay-of-the-Land
  - AI Canonical Architecture
  - Summary of Study Outreach and Highlights
- Robust AI
- Recommendations
- Summary

AI Canonical Architecture

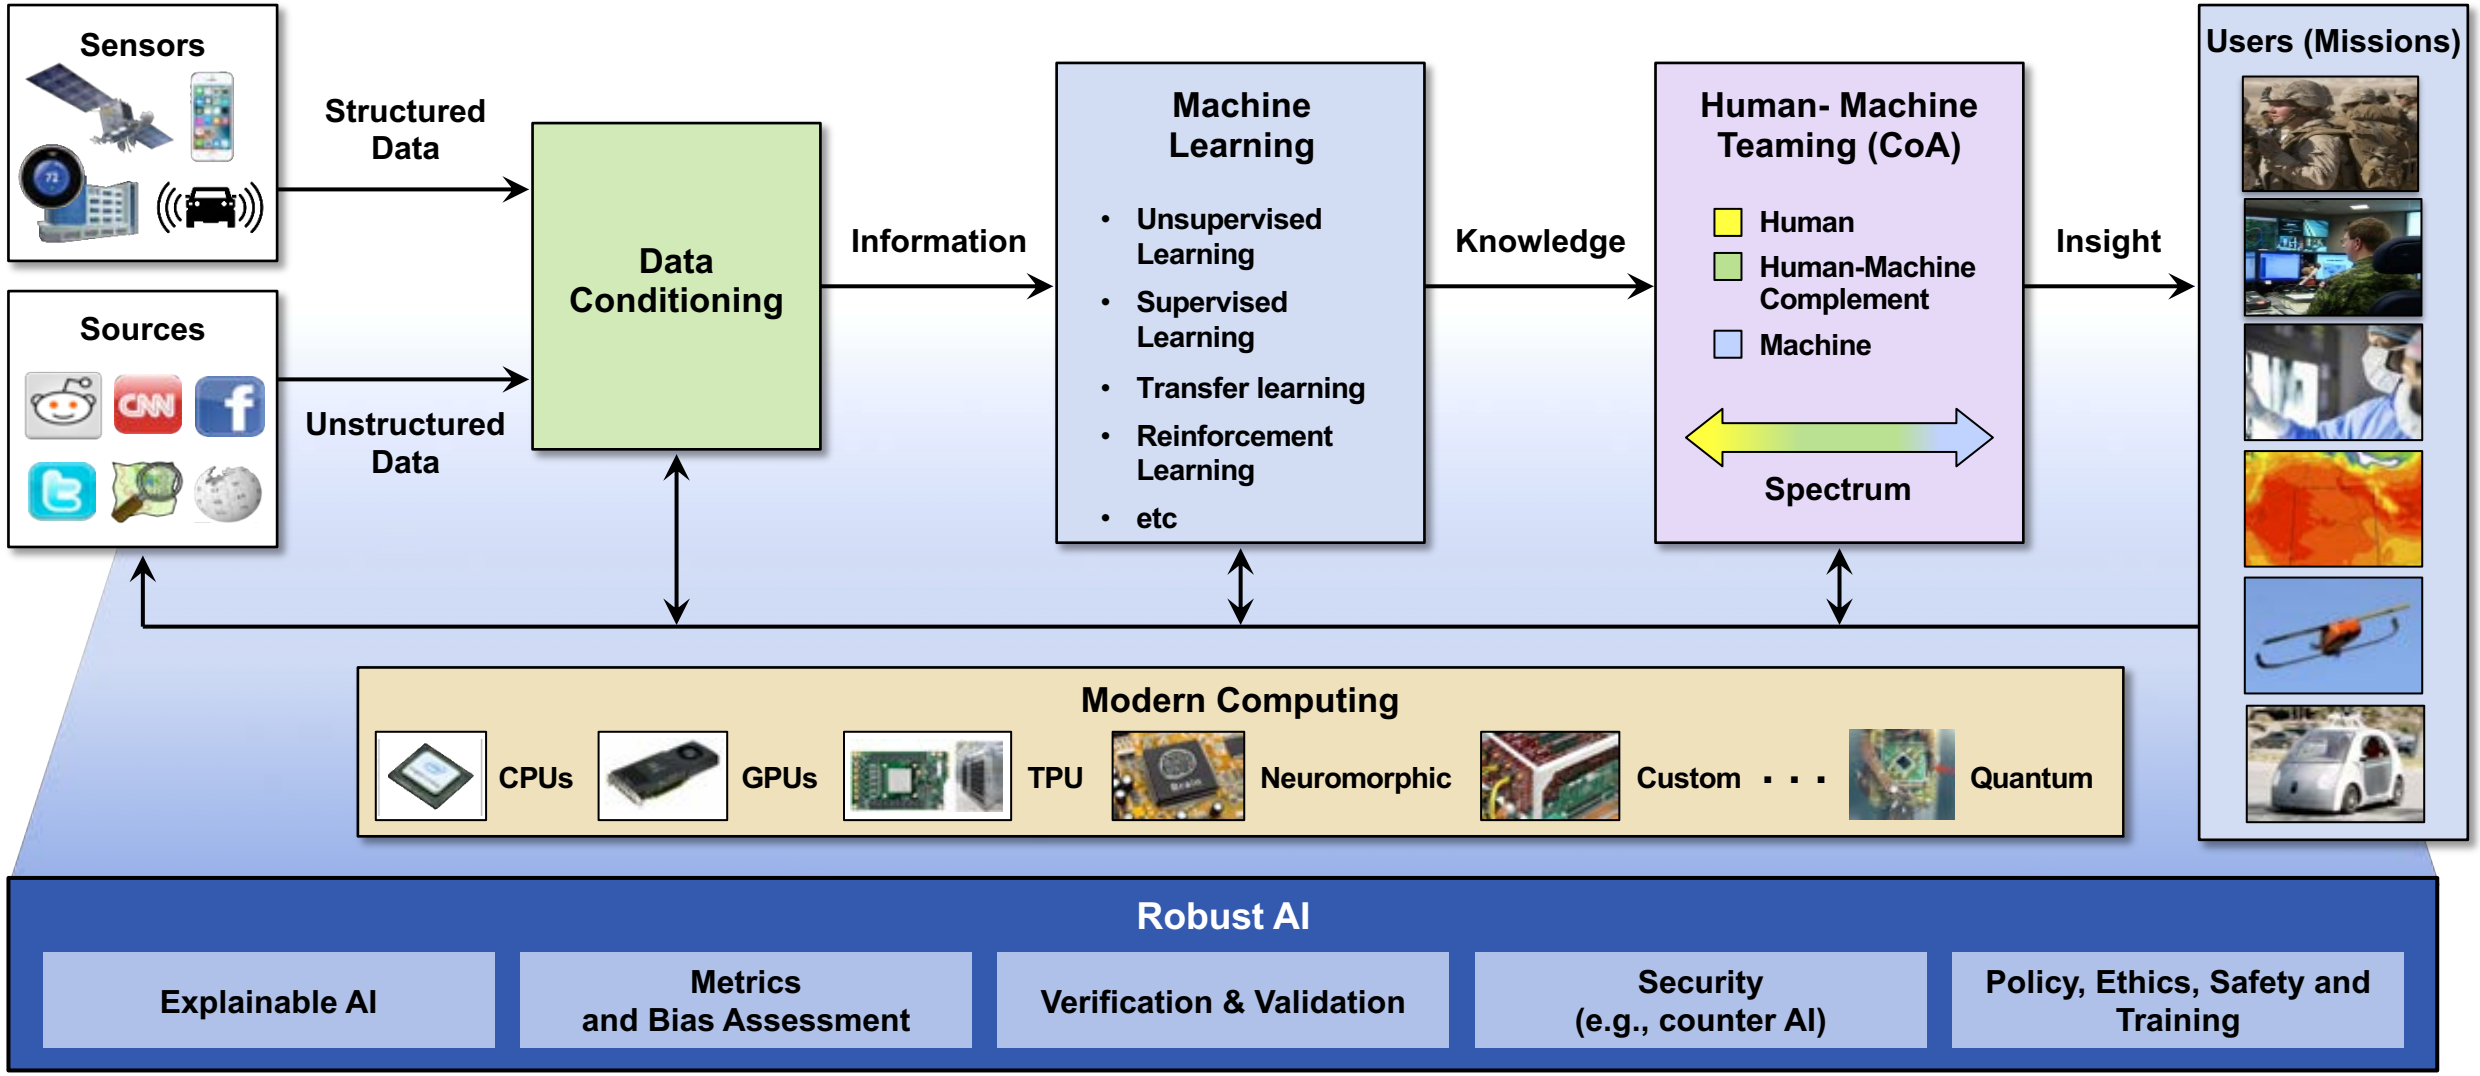

GPU = Graph Processing Unit
TPU = Tensor Processing Unit
CoA = Courses of Action

AI Stack from Dean Andrew Moore
- Former School of Computer Science at CMU -

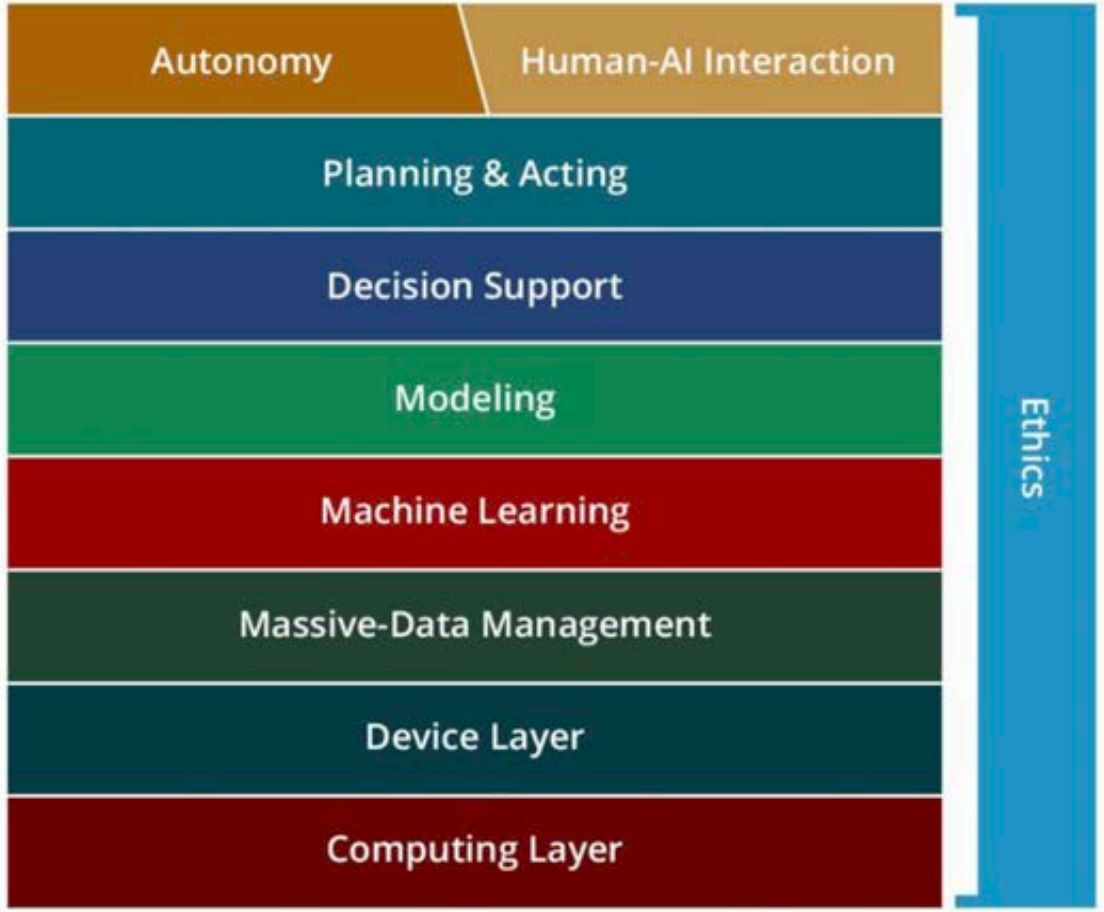

Figure 1. The AI Stack as envisioned and defined by Andrew Moore, Dean of the School of Computer Science, Carnegie Mellon University.

National Security Specific Concerns Compared To Commercial Sector

Commercial Sector	National Security
High dimensionality	High dimensionality
Large volume	Large volume
Mostly using enterprise cloud computing	Need for both cloud and tactical computing
Human-machine teaming is tolerant to errors	Human-machine teaming must be robust
Abundant amounts of labeled data	Limited amounts of labeled data
High capacity datalinks	Low capacity/intermittent datalinks
Competitive environment	Adversarial environment / countermeasures
Mostly consumer users	Today requires sophisticated users
Explainability is not the largest issue	Trust / explainability is core

AI will be a technological enabler (i.e., data and algorithm warfare) against: radical extremists, terrorists, and peer nations to defend our homeland and abroad

Four Components of Machine Learning Solutions

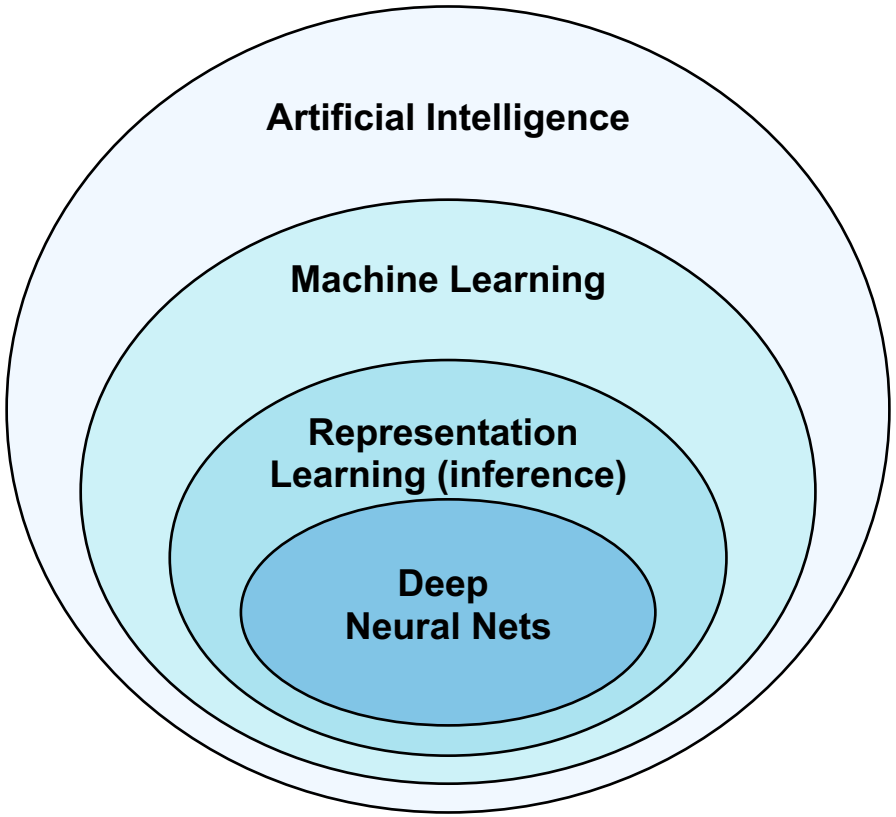

Image Adapted From: "Deep Learning"
I. Goodfellow, et.al., 2016 MIT Press

1. Define a Problem

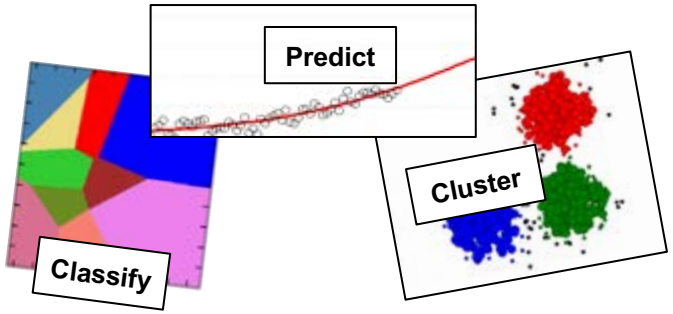

2. Gather Data

3. Create Full Train/Test Solution Pipelines

4. Provide Computation that Makes the Solution Feasible

Eric Schmidt, Former Executive Chairman of Alphabet Talk at MIT, May 2017

Where is Research Needed?

- 100x – 1000x Computational Resources

- Learn to Learn – Automated Machine Learning

D3M Program
(Data-Driven Discovery of Models)

DARPA

- A Knowledge Base of Shared Knowledge and Solutions

Model Zoo

Alphabet's Eric Schmidt speaks at CSAIL
https://youtu.be/TilXV0OA6JY

- One big impediment to progress in the DoD is lack of labeled data and adv. edge computing
- No organized approach exists by DoD to leverage commercial and academia advances

Google (WAYMO) Self-Driving Cars Development*

Development Highlights

- **2.5 Million real-world miles**
  1 Billion virtual miles in 2016
- **80% Human Intervention in 2015**
  20% Human Intervention in 2016

Number of people killed per year in car crashes:

- **37,461 in the U.S.[1]**
- **1.25 million worldwide[2]**

* Dimitri Dolgov, Self-Driving Cars and the Future of Mobility, Association for the Adv. of AI, Feb 8, 2017

[1] U.S. Department of Transportation (2016)
[2] World Health Organization (2014)

Spectrum of Commercial Organizations in the Machine Intelligence Field

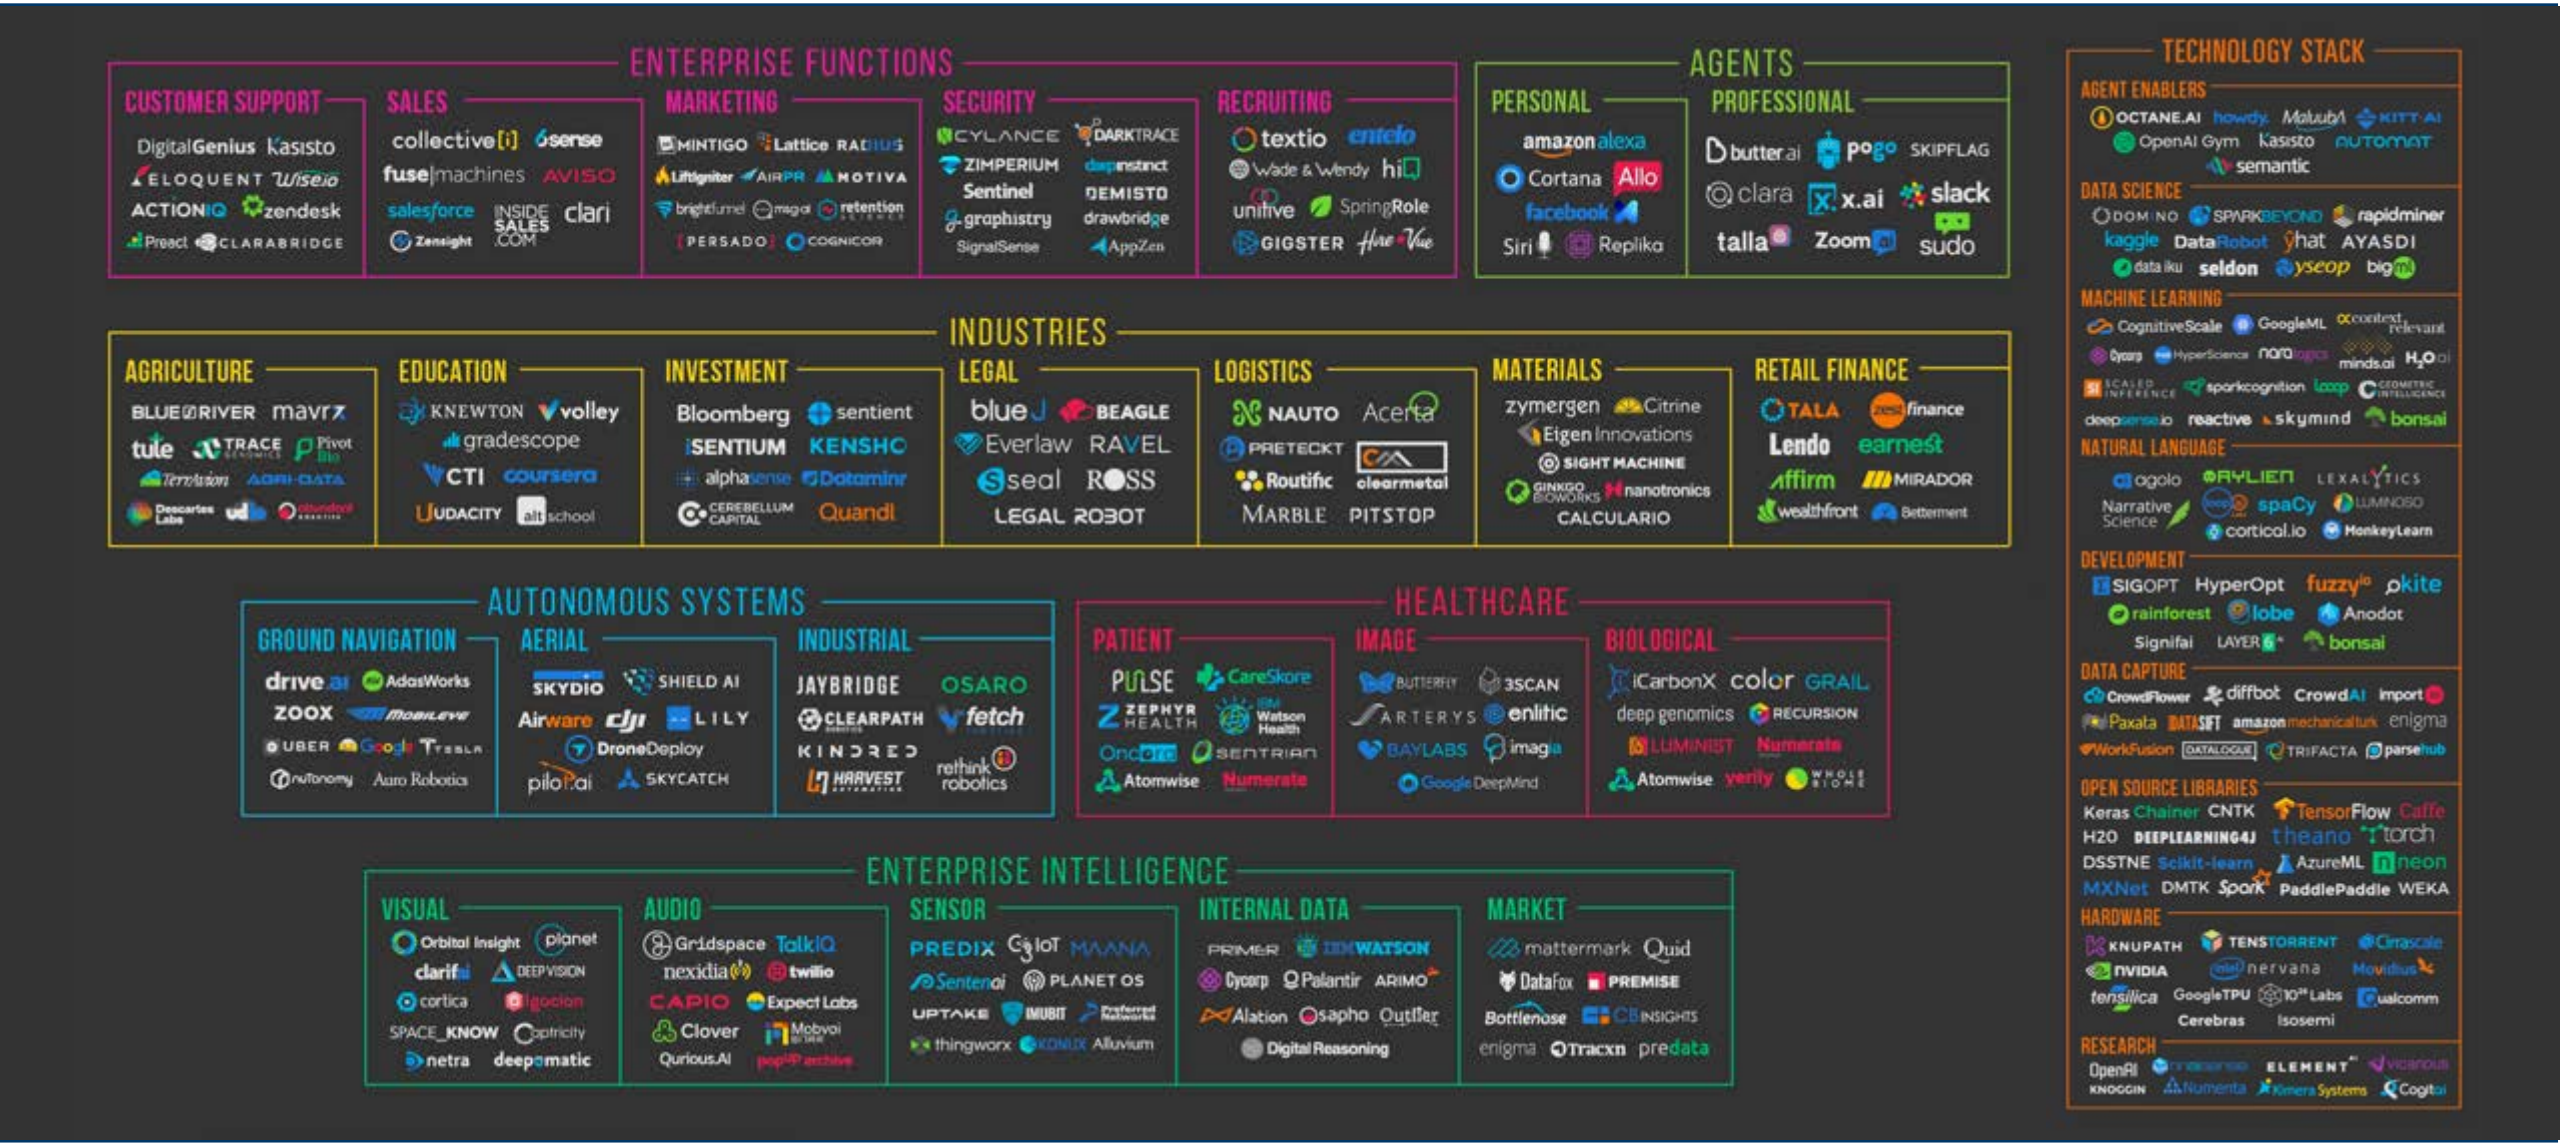

Source: Shivon Zilis, 2016, https://www.oreilly.com/ideas/the-current-state-of-machine-intelligence-3-0

Outline

- Background
- Lay-of-the-Land
  - AI Canonical Architecture
  - Summary of Study Outreach and Highlights
- Robust AI
- Recommendations
- Summary

Government Organizations Study Outreach

DoD

Intelligence Community

IARPA

Under Secretary of Defense for Intelligence (USDI)

Defense Contractors, Commercial, Peers, and AI Centers Study Outreach

Defense Industrial Base

charles river analytics

ntfy

PHRONESIS

Raytheon BBN Technologies

METRON SCIENTIFIC SOLUTIONS

GENERAL DYNAMICS

Commercial

NVIDIA

IBM Watson

NUANCE

amazon

RAPID7

Microsoft

Research at Google

CROWDSTRIKE

Peers

CMU SEI

JHU HLTCOE

ARGONNE

OAK RIDGE

NASA

JPL

PNNL

Pacific Northwest NATIONAL LABORATORY

LIVERMORE

LOS ALAMOS

SANDIA

MITRE

MITRE

AI Centers

OpenAI

MIRI MACHINE INTELLIGENCE RESEARCH INSTITUTE

AI2 ALLEN INSTITUTE for ARTIFICIAL INTELLIGENCE

SRI International

USC University of Southern California

USC CENTER FOR ARTIFICIAL INTELLIGENCE IN SOCIETY

Partnership on AI to benefit people and society

Academia and MIT Study Outreach

MIT School of Engineering

Anantha Chandrakasan

MIT CSAIL

Srini Devadas

MIT CSAIL

Jim Glass

MIT CSAIL

Patrick Winston

MIT Brain and Cog Sciences

Vikash Mansinghka

MIT Media Lab

Sandy Pentland

MIT LIDS

Devavrat Shah

MIT IDSS

Suvrit Sra

MIT-IBM Watson AI Lab

Aude Oliva

MIT-IBM Watson AI Lab

Antonio Torralba

Florida International University

Mark Finlayson

Northeastern University

David Kaeli

Boston University

Michel Kinsy

Ohio State University

Srini Parthasarathy

University of Michigan

Arunesh Sinha

University of Southern California

Milind Tambe

Unstructured and Structured Data
- Relevant to Cyber Security & Information Sciences -

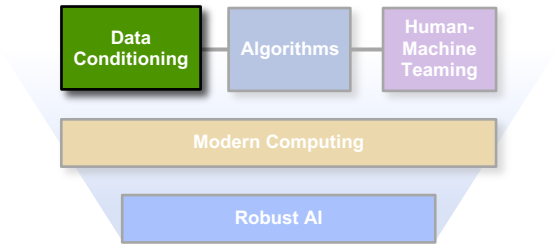

Structured Data Types


Speech


Sensors


Network Logs


Metadata

Unstructured Data Types


Social Media


Human Behavior


Reports

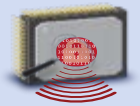
Side Channel

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
Machine Learning (Unsupervised)  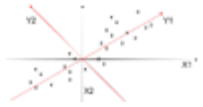	• Limited machine learning • Dimensionality Reduction • Clustering/Pattern Recognition • Outlier Detection
Data Labeling 	• Initial data exploration • Highlight missing or incomplete data • Reorient sensors/recapture data • Look for errors/biases in collection

Important needs are in labeling data and automating data conditioning

2018 AI Predictions

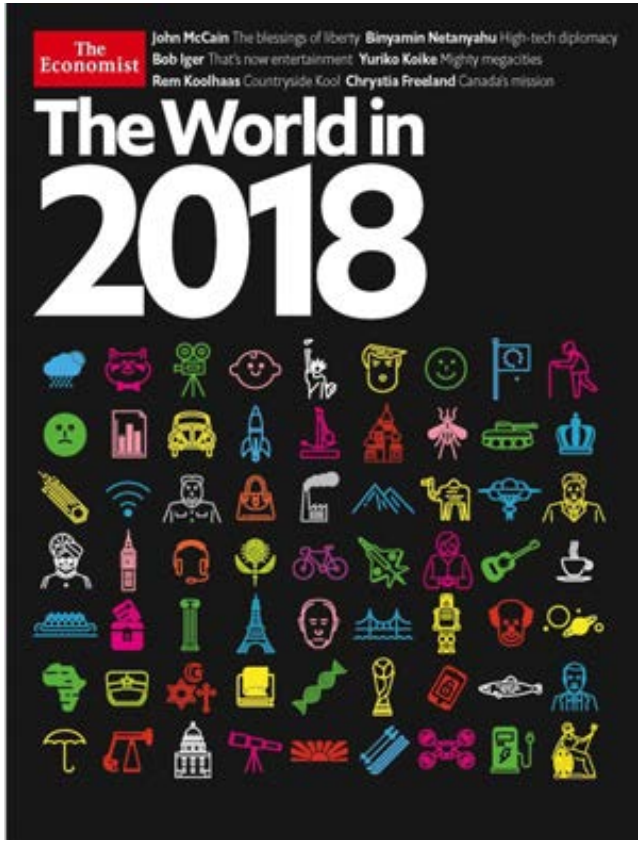

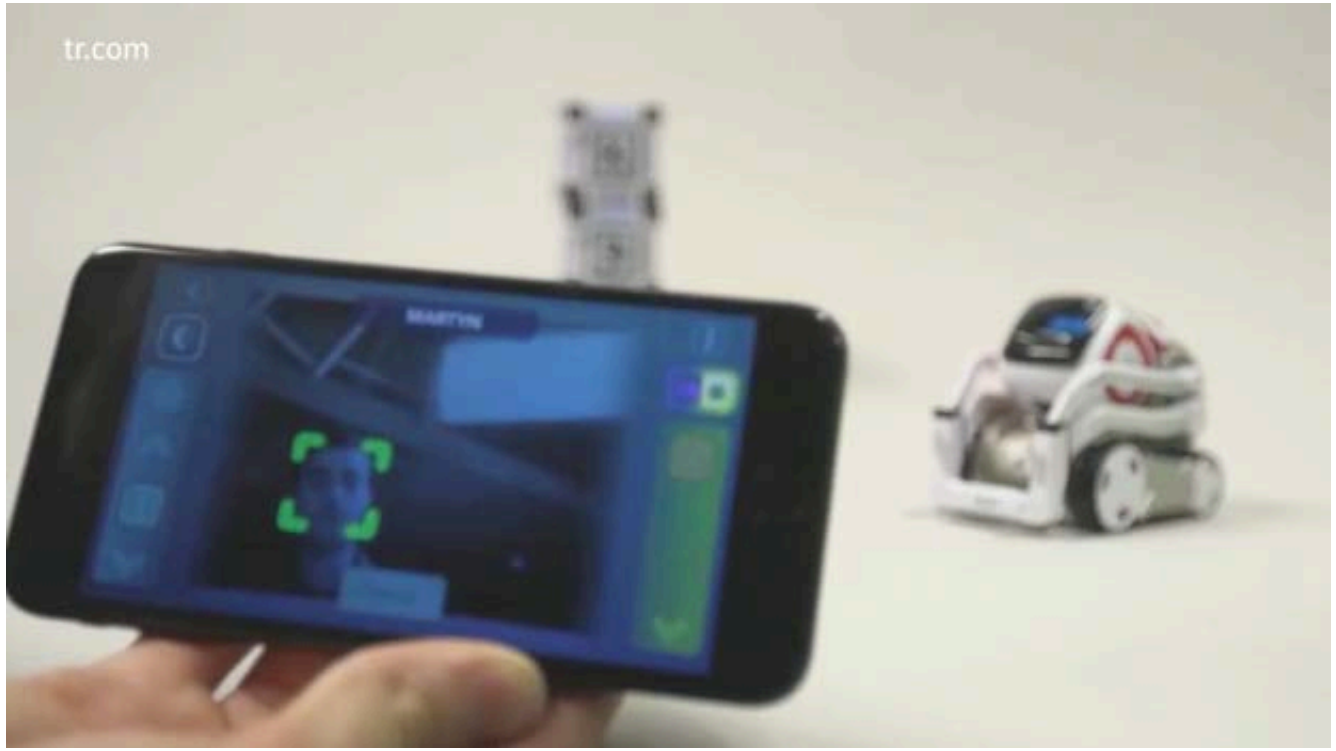

https://www.thomsonreuters.com/en/reports/2018-ai-predictions.html

Key Elements of AI Canonical Architecture

- Sensors, sources, and data conditioning
- Algorithms, modern computing, and robust AI
- Human-machine teaming and users

AI Practitioners Refer to Data as the New Oil
- We are Drowning in Data -

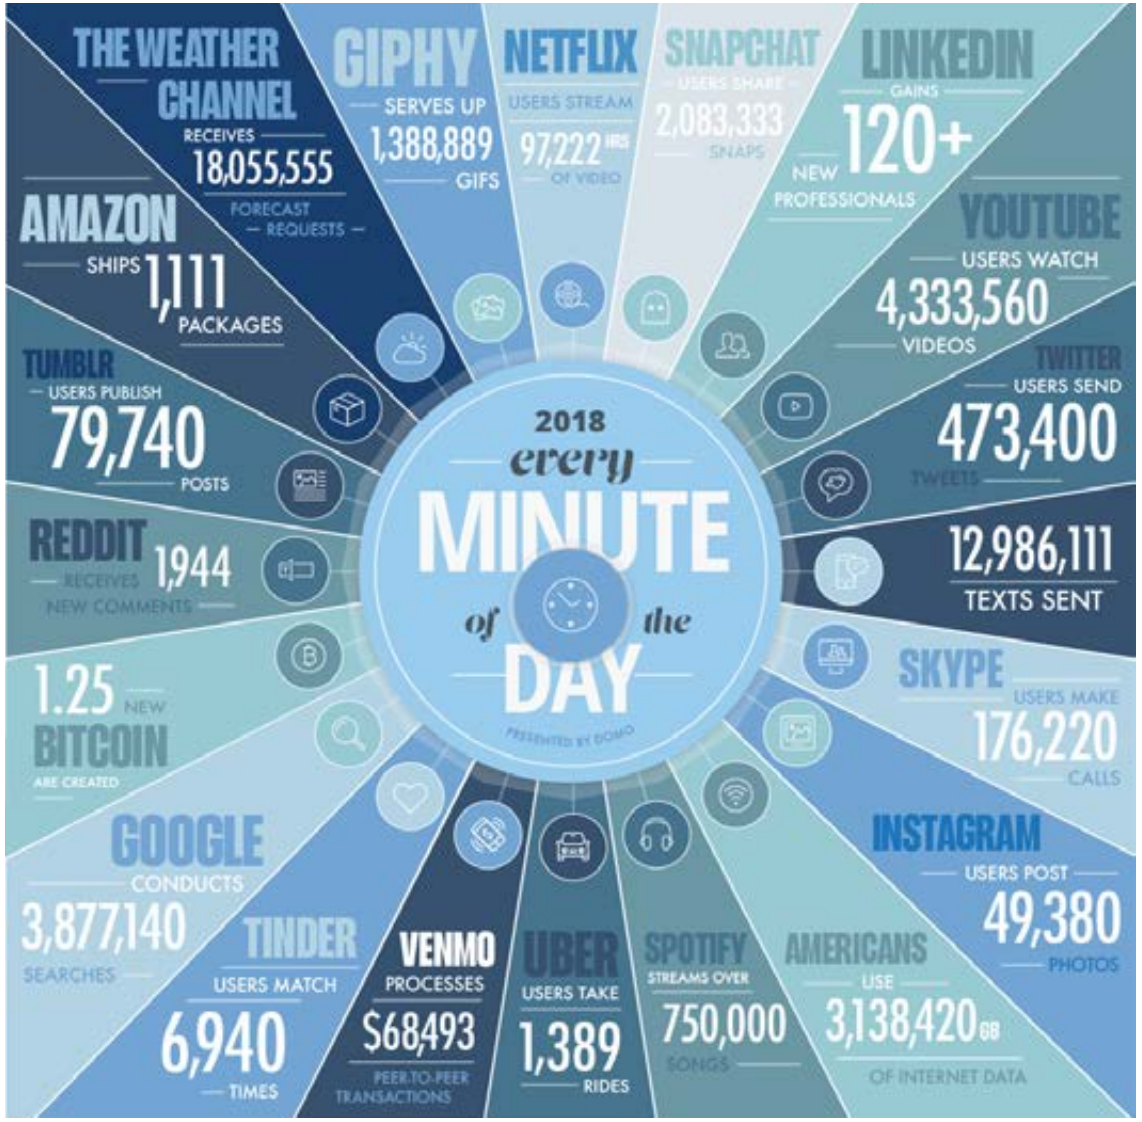

Some Recent Statistics*

- 90% of world data have been created in the past 2 years
- 80% of these data are unstructured (documents, tweets, videos, images, etc.)
- 2.5 exabytes ($10^{\wedge 18}$) of data are created each day
- 1.5B users are active in Facebook
- Google processes 3.5B searches per day
- More than ½ of web searches are done on a mobile phone
- Every minute 103M spam emails are sent

* Statistics provided by Dr. Cem Sahin, MIT Lincoln Laboratory, Oct 2018

Presented by DOMO.com

Source: Statista, LinkedIn, Internet Live Stats, Expanded Ramblings, Slash Film, RIAA, Business of Apps, International Telecommunications Union, International Data Corporation

Diversity of Data/Metadata Management for Processing Efficiency

SQL Era

Common interface

NoSQL Era

Rapid ingest for internet search

NewSQL Era

Fast analytics inside databases

Future

Polystore, high performance ingest and analytics

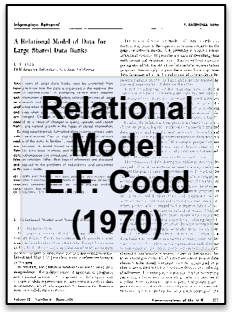

1970 1980 1990 2006 2010

Good for text reports

NoSQL

Good for metadata

Relational Databases (SQL)

Good for sensor data

NewSQL

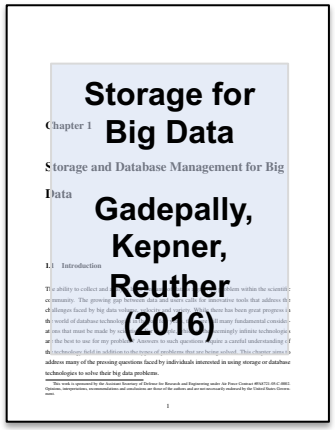

ORACLE PostgreSQL accumulo SciDB GRAPHULO

SQL = Structured Query Language
NoSQL = Not only SQL

Machine Learning Algorithms Taxonomy

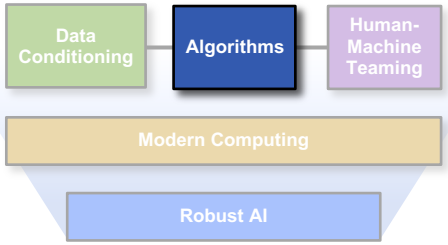

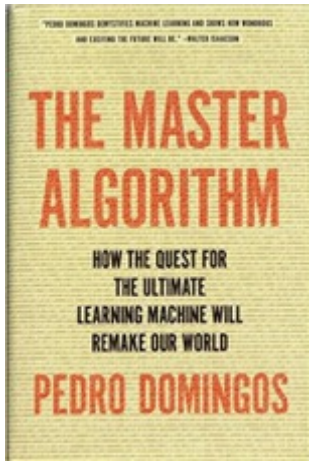

Algorithms*

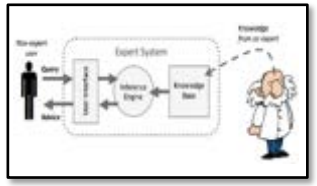

Symbolists (e.g., exp. sys.)

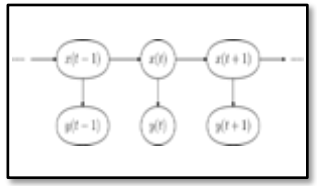

Bayesians (e.g., naive Bayes)

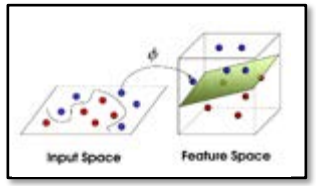

Analogizers (e.g., SVM)

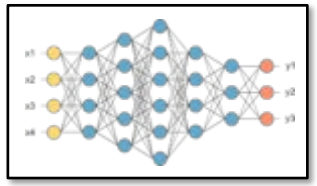

Connectionists (e.g., DNN)

Evolutionaries (e.g., genetic programming)

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

Machine Learning Applied to Classifiers

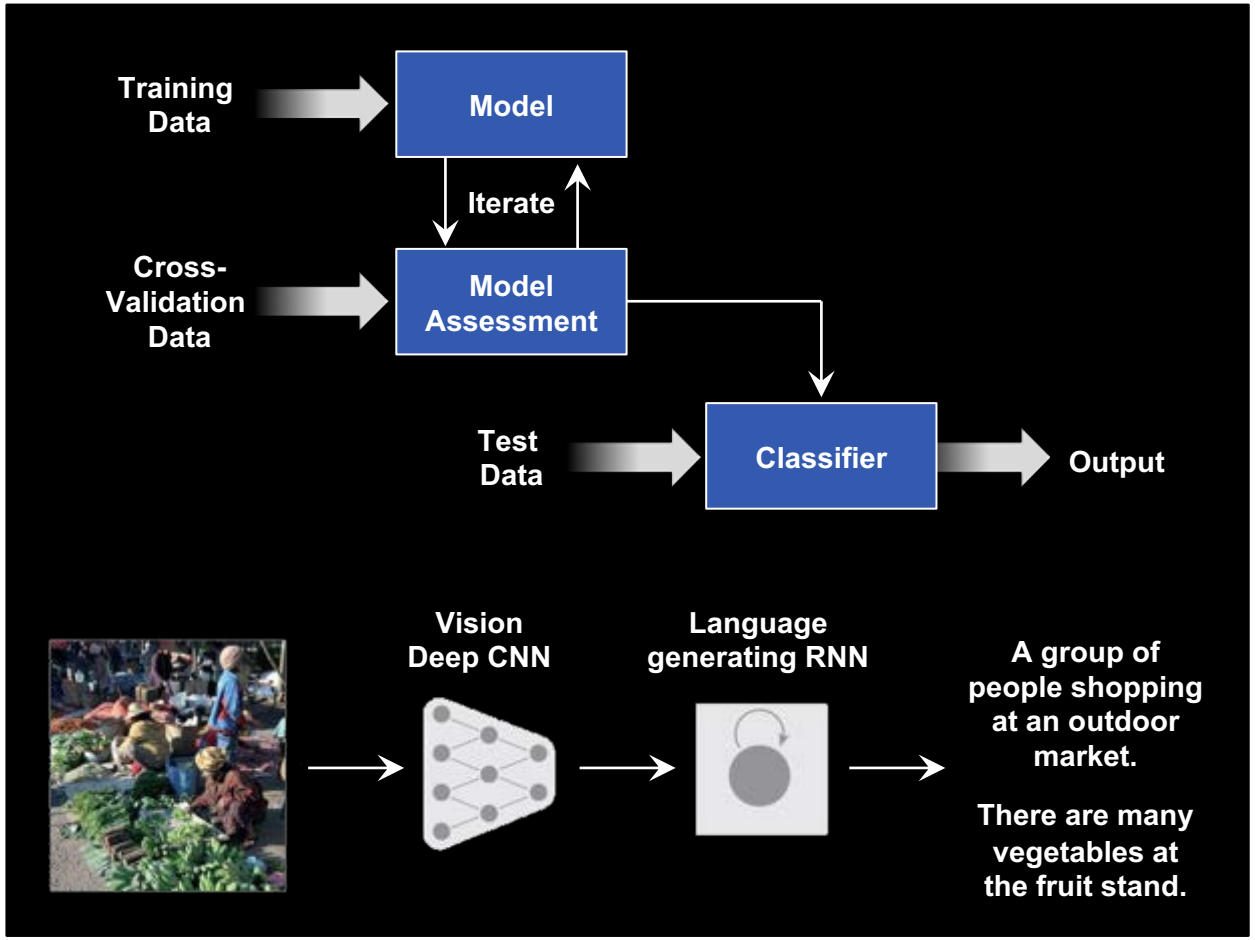

CNN = convolution neural network
RNN = recurrent neural network

Source of outdoor market graphic: nature.com, Deep Learning (May 2015)

Machine Learning
- Examples of Popular Algorithms -

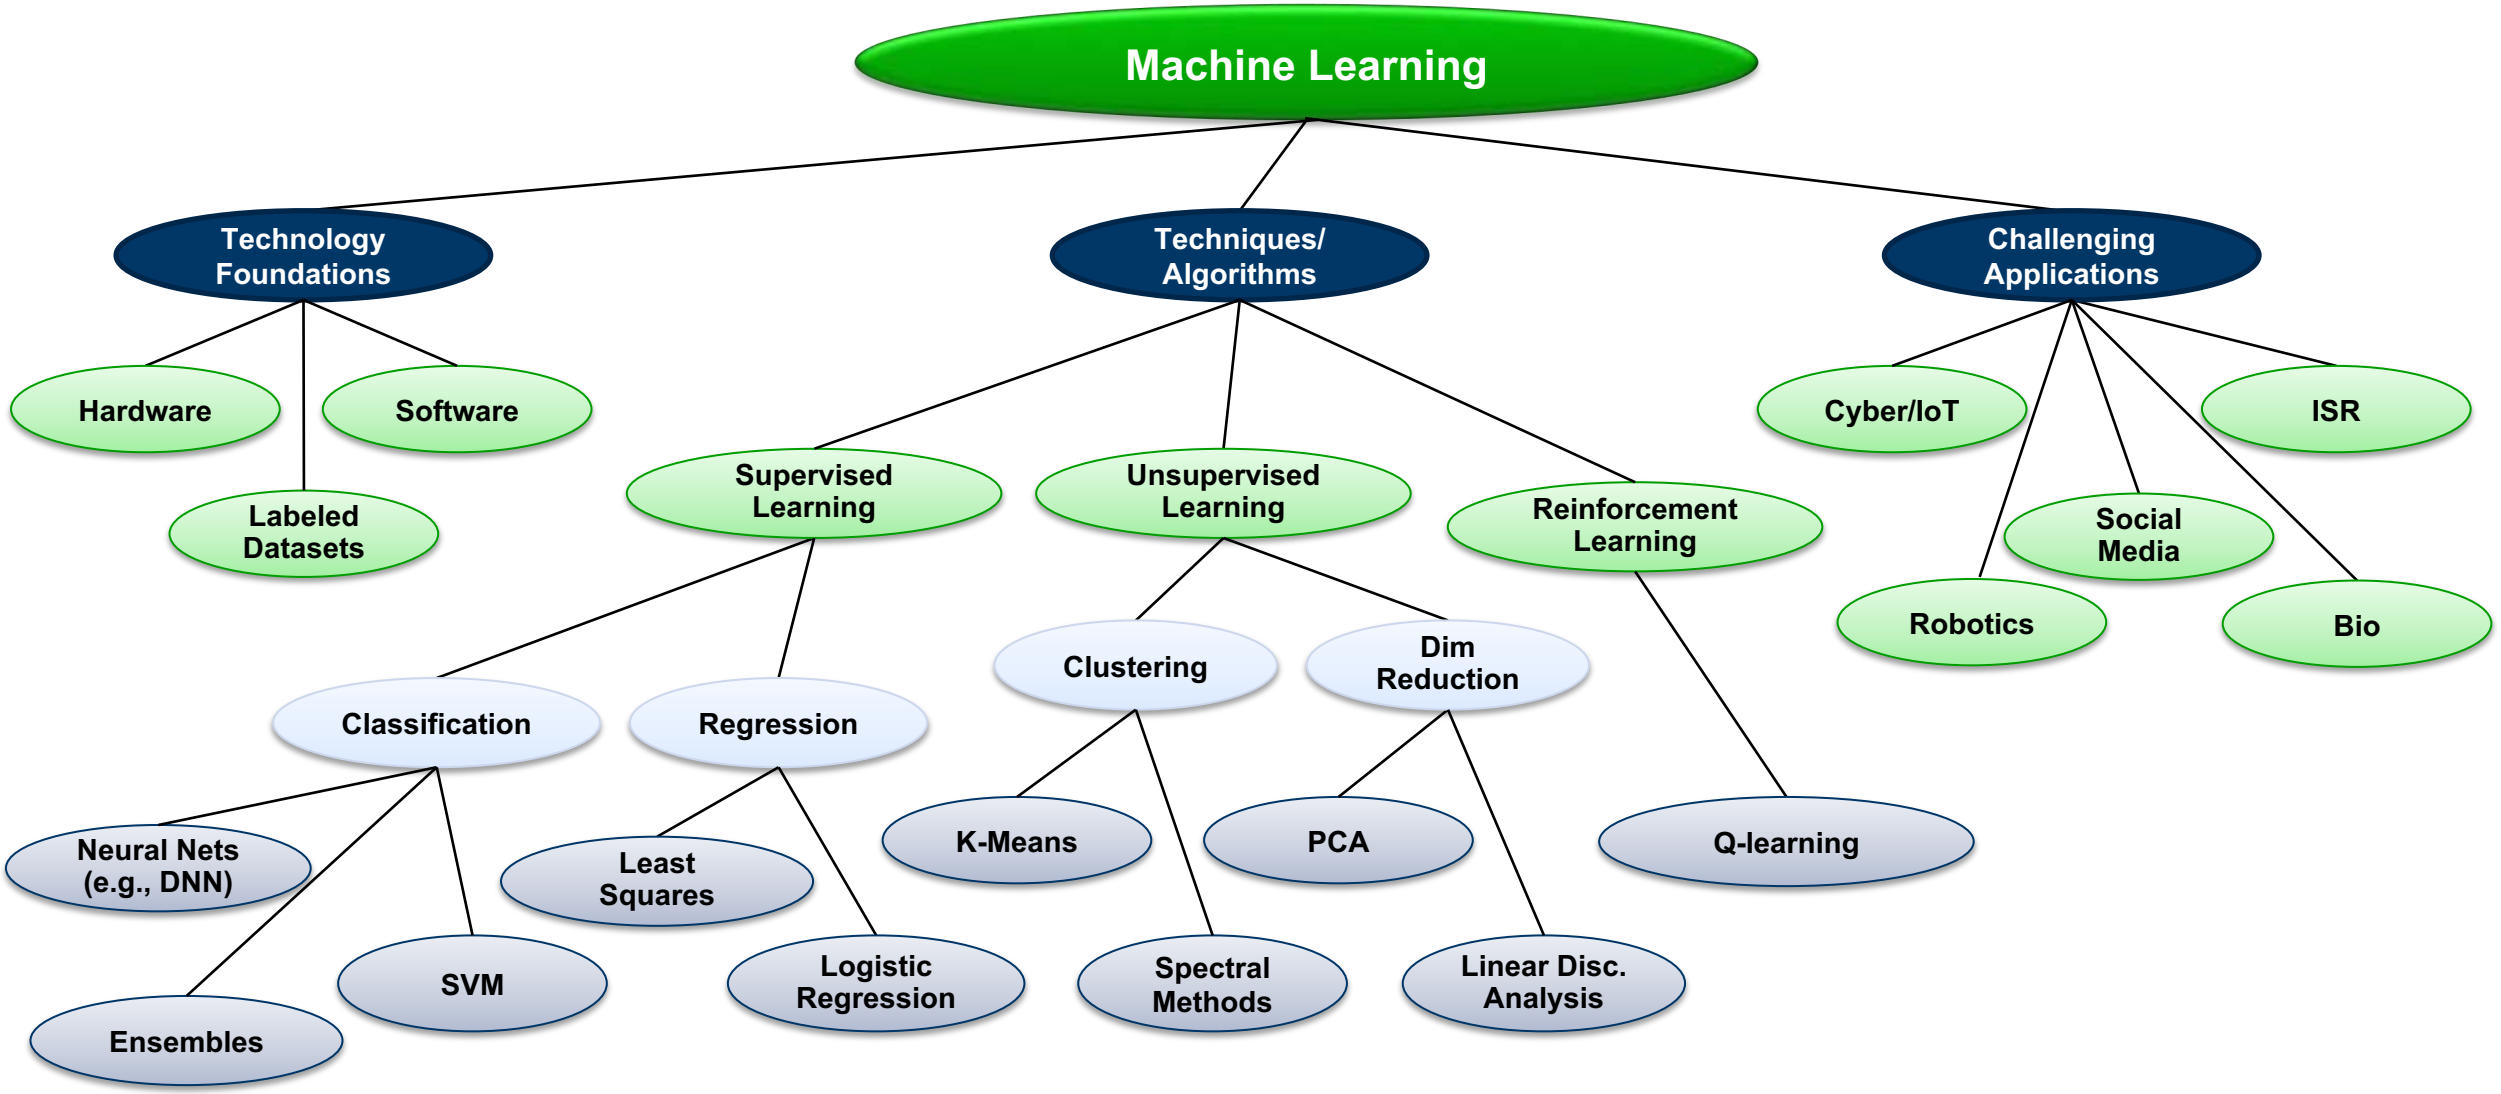

SVM: Support Vector Machines
PCA: Principal Component Analysis

How Neural Networks Recognize a Dog in a Photo

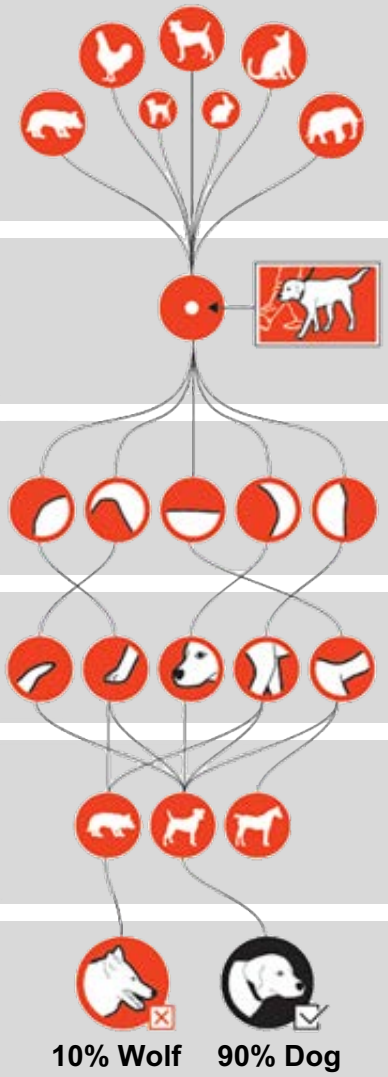

Training
During the training phase, a neural network is fed thousands of labeled images of various animals, learning to classify them.

Input
An unlabeled image is shown to the pre-trained network.

First Layer
The neurons respond to different simple shapes, like edges.

Higher Layer
Neurons respond to more complex structures.

Top Layer
Neurons respond to highly complex, abstract concepts that we would identify as different animals.

Output
The network predicts what the object most likely is, based on its training.

Source: fortune.com, The AI Revolution: Why Deep Learning is Suddenly Changing Your Life (Parloff, September 2016)

Artificial Neural Networks Tutorial

- Inspired by biological networks
- Systems learn by repetitive training to do tasks based on examples
  - Generally a supervised learning technique
- Components: Inputs, Layers, Outputs, Weights
- Deep Neural Network: Lots of "hidden layers"
- Popular variants:
  - Convolutional Neural Nets
  - Recursive Neural Nets
  - Deep Belief Networks
- Many toolboxes and hardware support

Deep Neural Networks

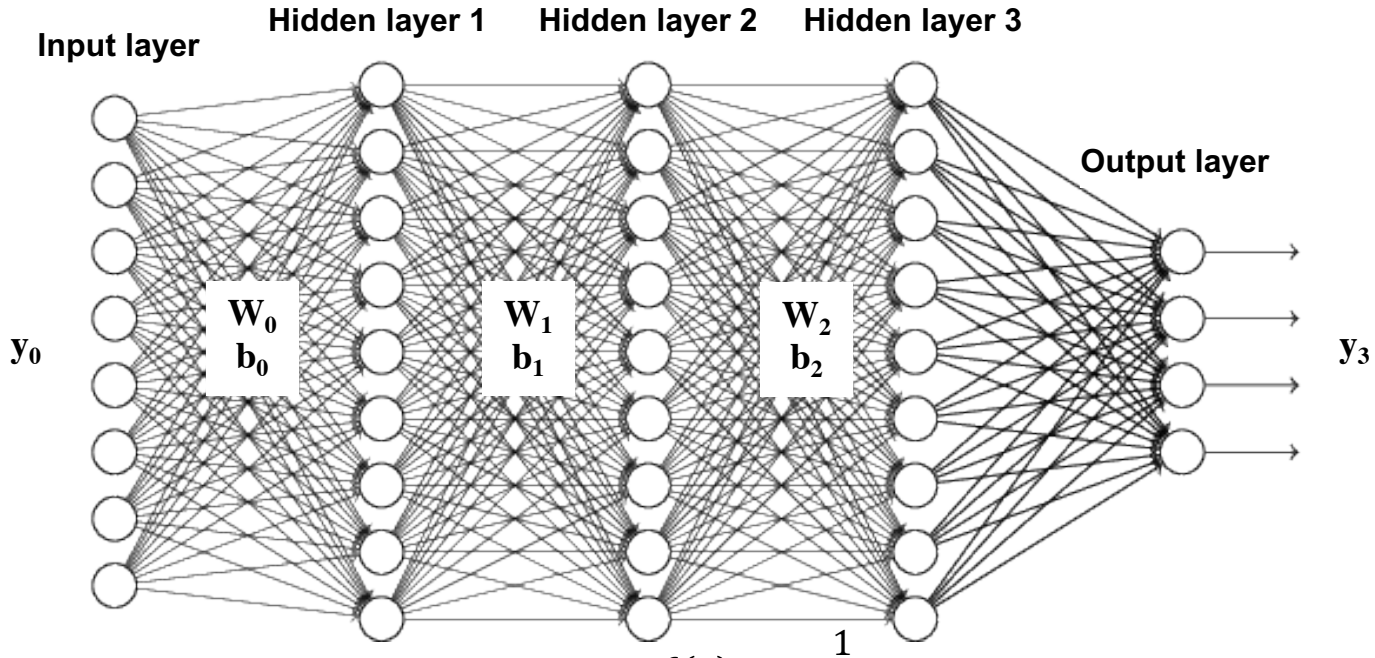

$$y_{i+1} = f(W_i y_i + b_i)$$

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

Activation Function

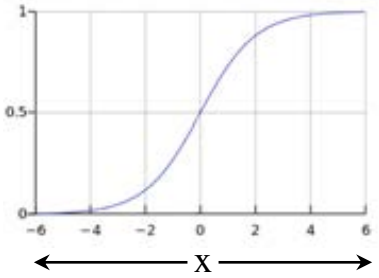

Image source: http://www.rsipvision.com/exploring-deep-learning/

Supervised Learning with Deep Neural Networks

Training vs. Inference

Training Phase

Deployment (Inference) Phase

Untrained Deep Neural Network

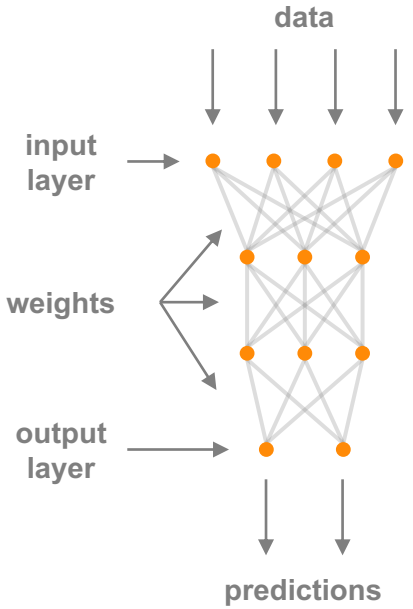

Labeled Samples

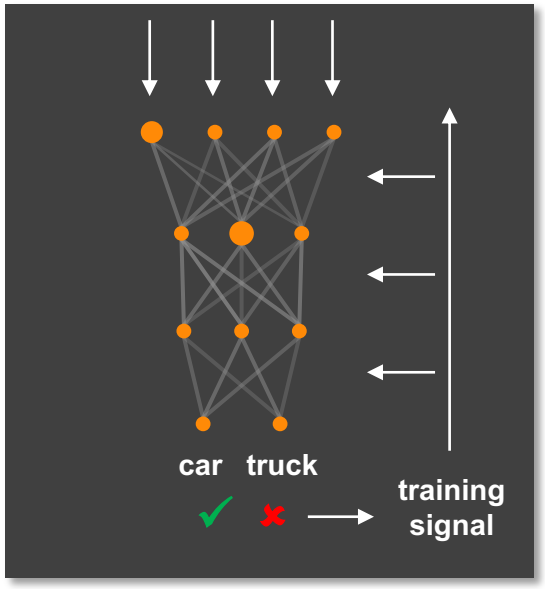

Network weights are adjusted to correctly predict labels

New Unlabeled Sample

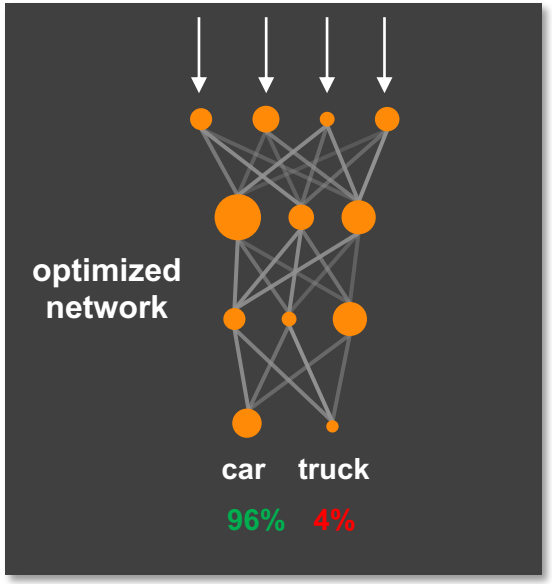

Adapted from: https://blogs.nvidia.com/blog/2016/08/22/difference-deep-learning-training-inference-ai

Supervised Learning with Deep Neural Networks

Training Epochs and Training/Inference Batches

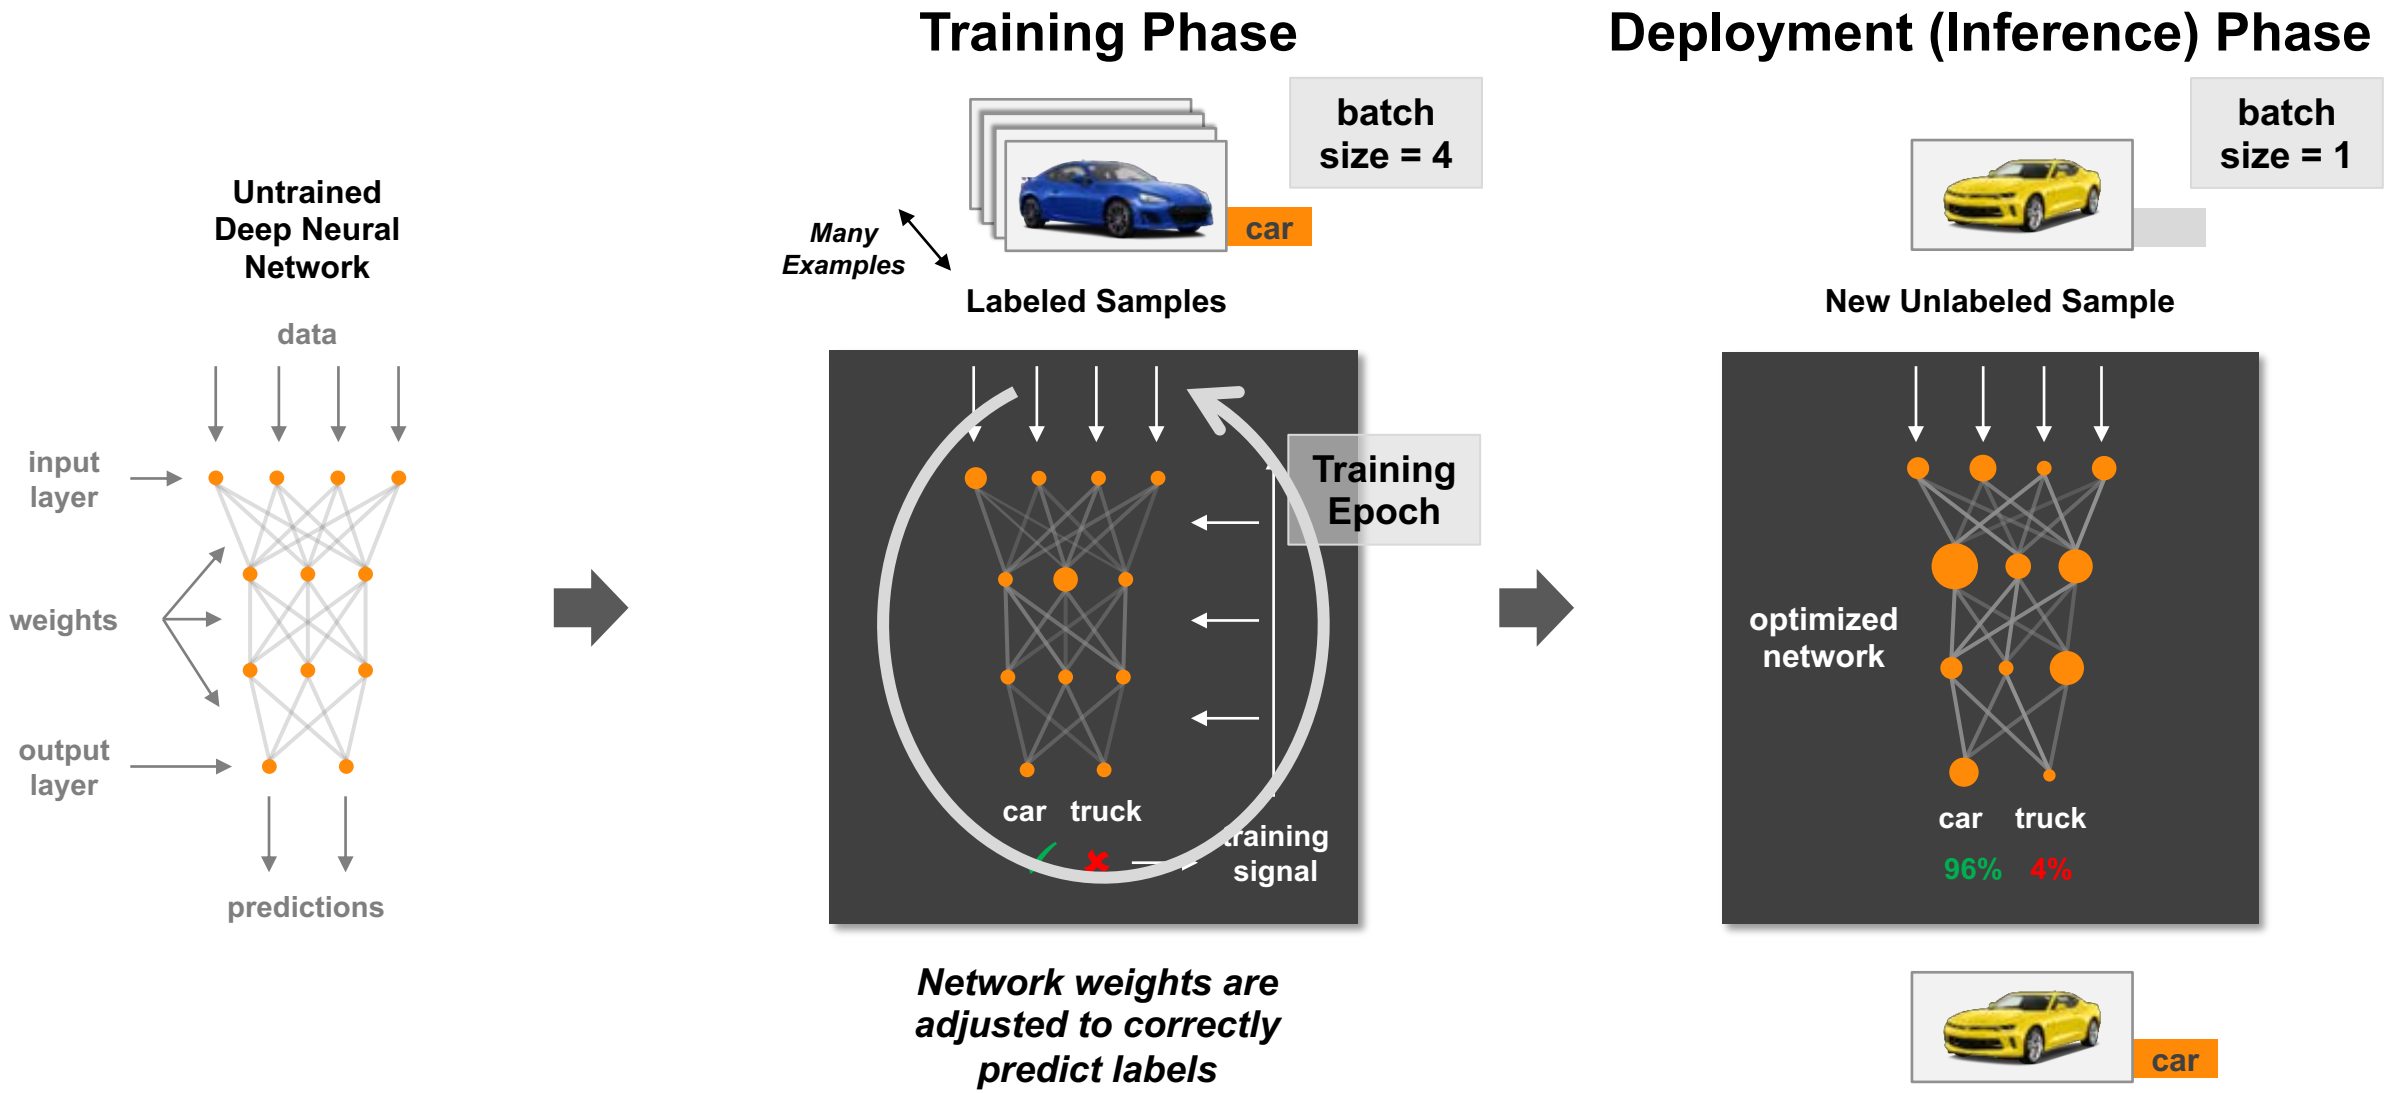

Adapted from: https://blogs.nvidia.com/blog/2016/08/22/difference-deep-learning-training-inference-ai

A Cambrian Explosion of Machine Learning Research Topics

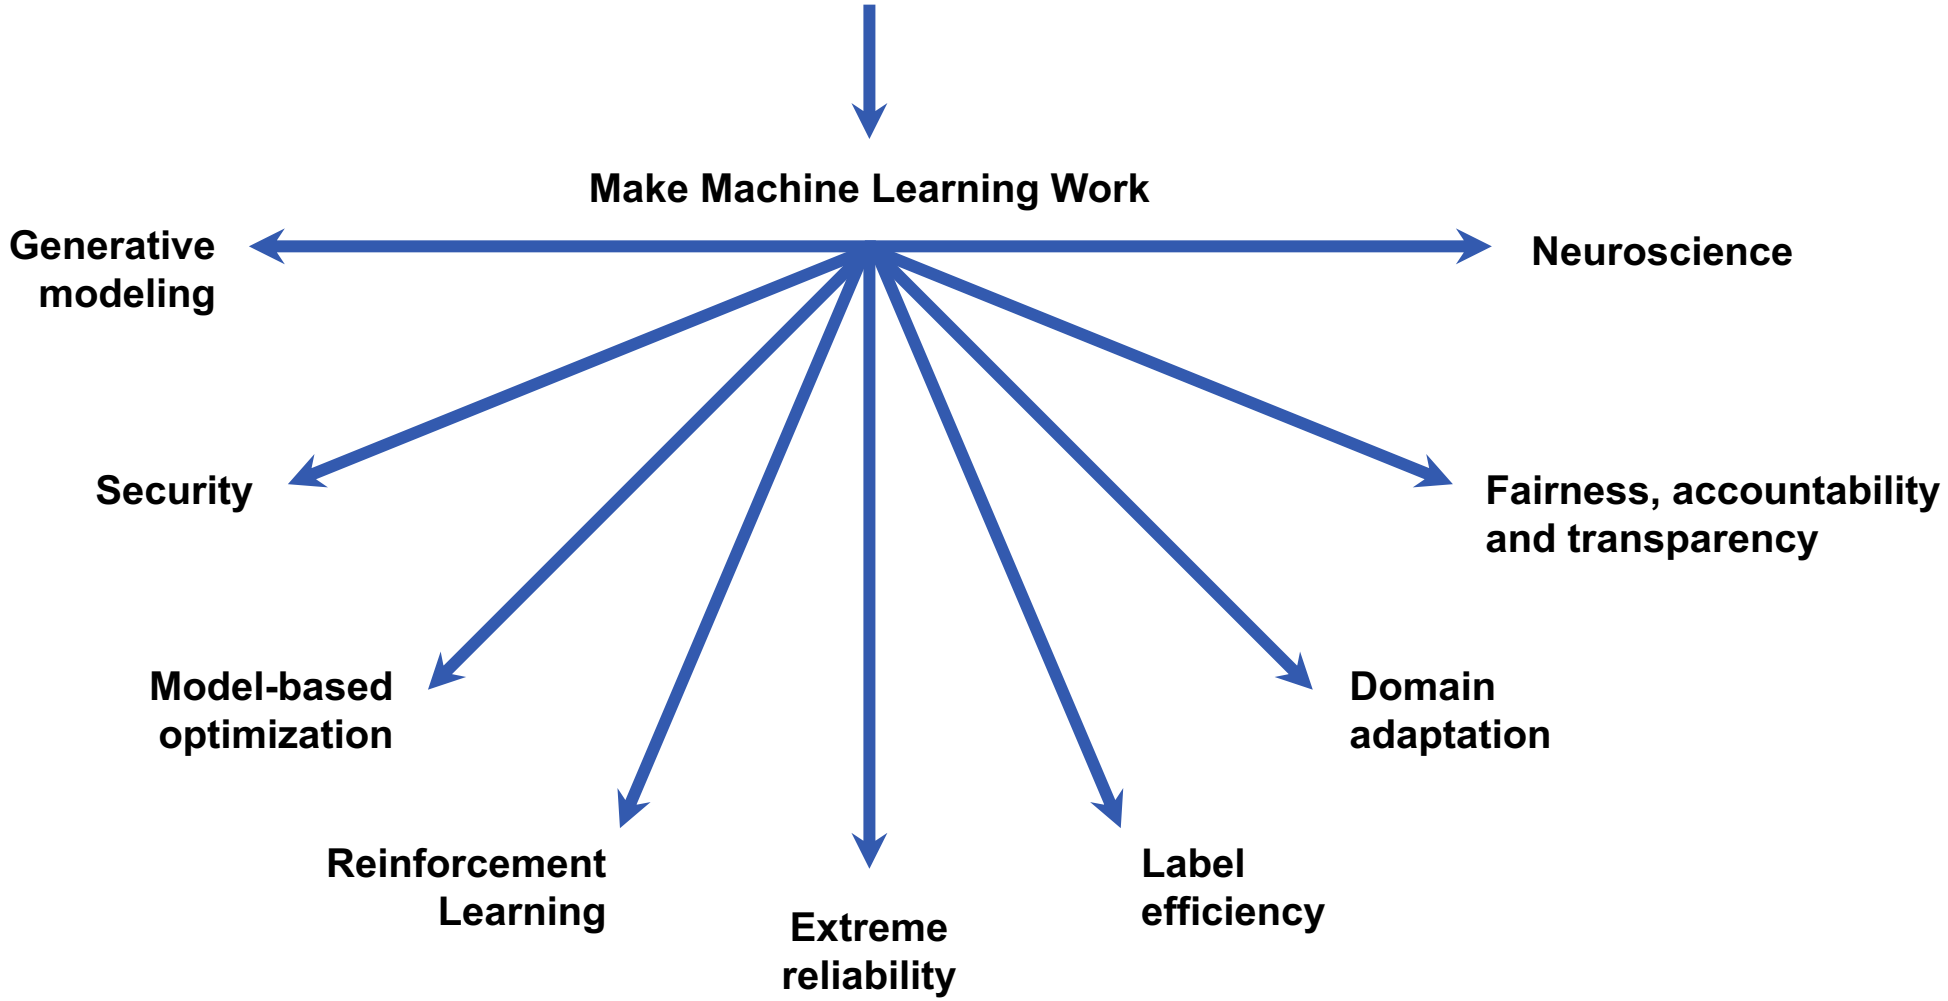

Source: Keynote Speech at AAAI 2019, Ian Goodfellow, Google AI, Talk titled: Adversarial Machine Learning

Image-to-Image Translation using Cycle-Consistent Adversarial Networks

CycleGAN

(Zhu et al., 2017)

Source: Keynote Speech at AAAI 2019, Ian Goodfellow, Google AI, Talk titled: Adversarial Machine Learning

Note: Image-to-image translation is a class of vision and graphics problems where the goal is to learn the mapping between an input image and an output image using a training set of aligned image pairs when data samples are lacking

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/

High Performance Software and Tools
- Open Community from Commercial and Academia Developers -

High Productivity Languages

MATLAB SIMULINK

GNU Octave

julia

R

python

Java

High Performance Languages

Fortran

C

C++

MPI

Machine Learning Frameworks

Caffe

Keras

PYTORCH

TensorFlow

APACHE Spark

ANACONDA Powered by Continuum Analytics

Algorithm Building Blocks

The commercial and academia AI communities have built a number of tools available for rapid development of algorithms

Compilers, Debuggers, Performance Analysis

intel PARALLEL STUDIO XE

GCC

allinea DDT

TotalView®

High Performance Databases

SciDB

accumulo

Automated Detection of Counterfeit Parts

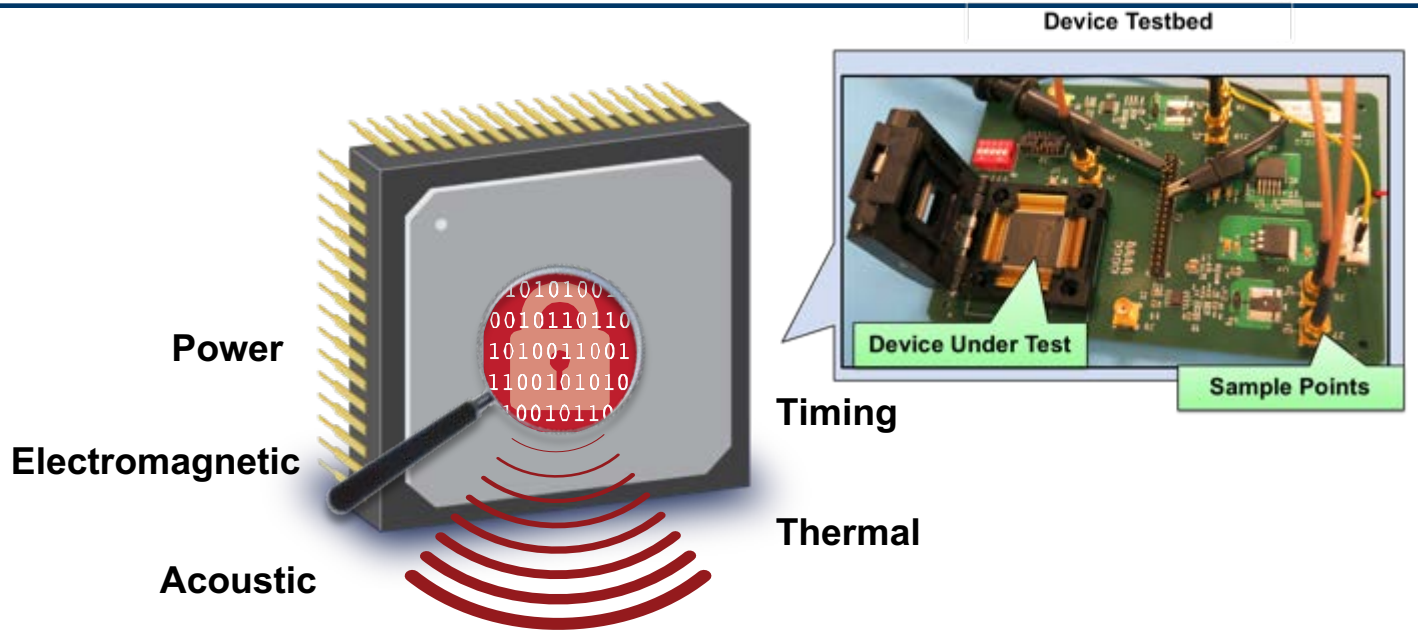

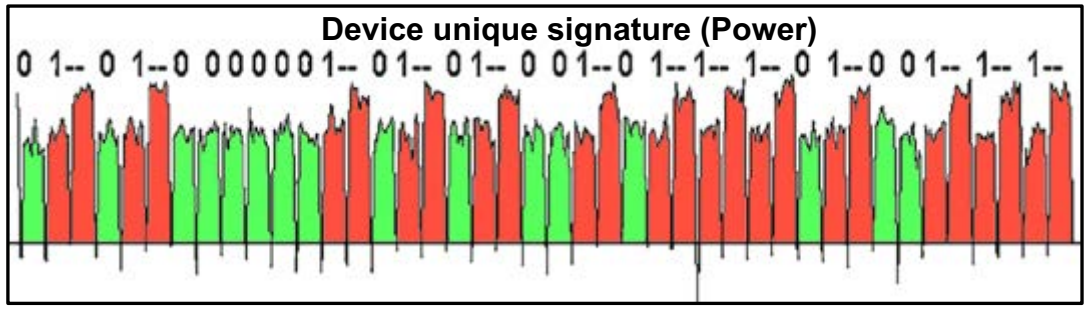

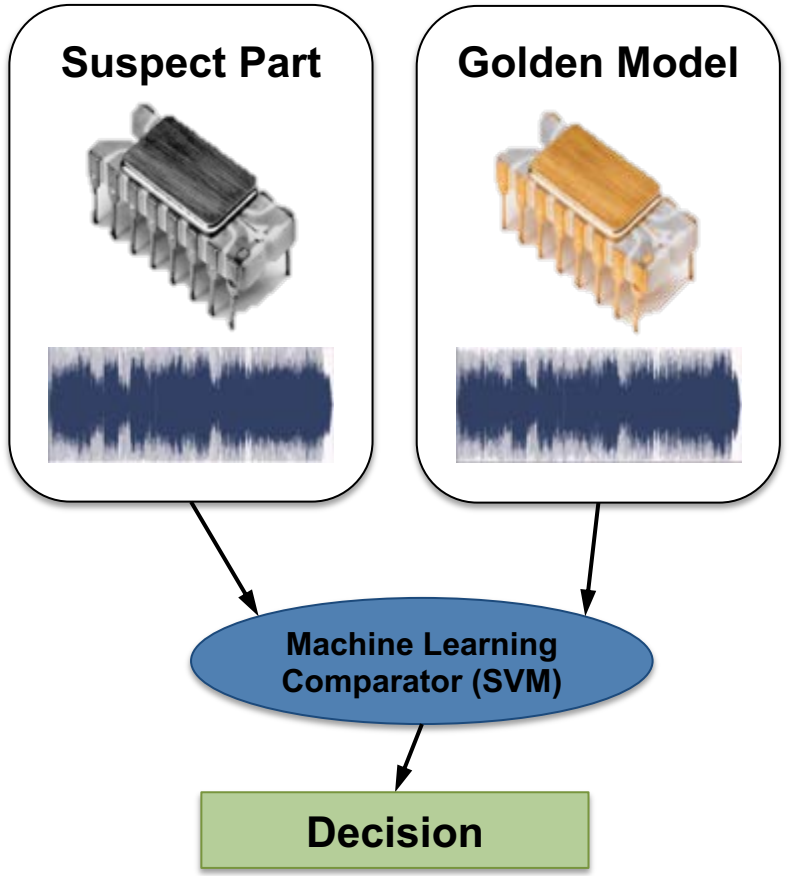

Microprocessor unique signatures can be observed through measurements and used to uniquely identify device classes

SVM: Support Vector Machines

Modern AI Computing Engines

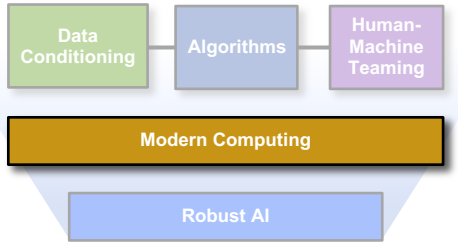

Computing Class





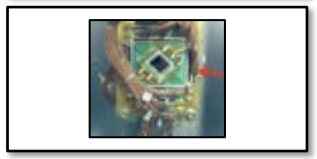

What It Provides to AI

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	• Still a 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

Alexnet comparison: Forward-Backward Pass

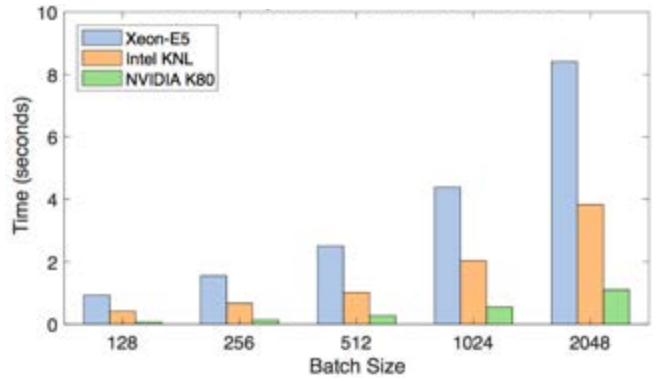

SpGEMM Performance using Graph Processor (G102)

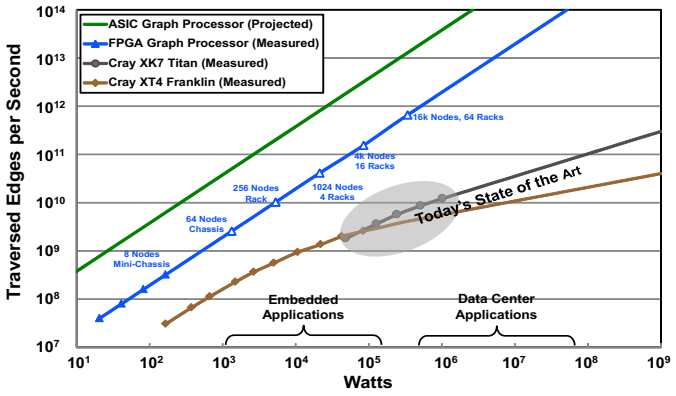

GPU = Graph Processing Unit
TPU = Tensor Processing Unit

HHL = Harrow-Hassidim-Lloyd quantum algorithm

The Era of Hitting All the Walls: Power, Clock Frequency and Memory

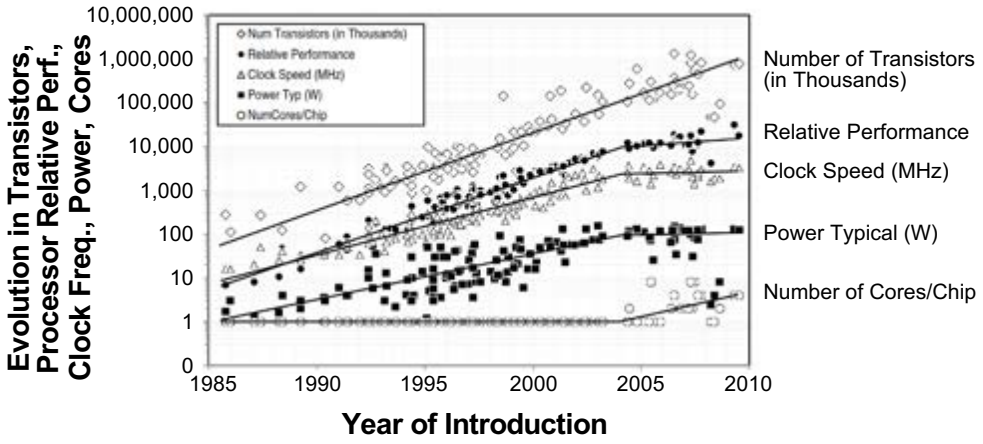

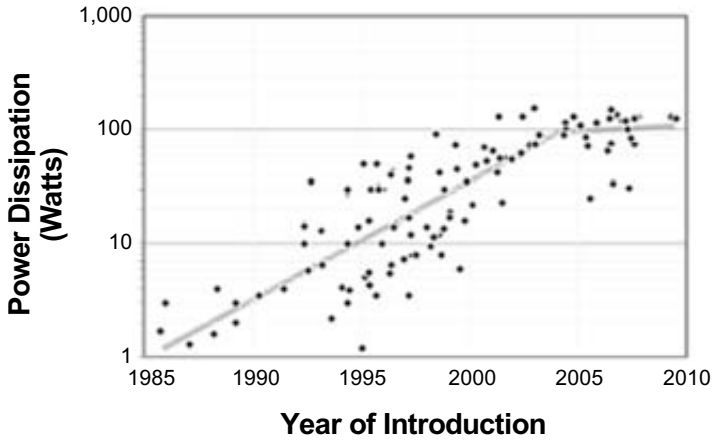

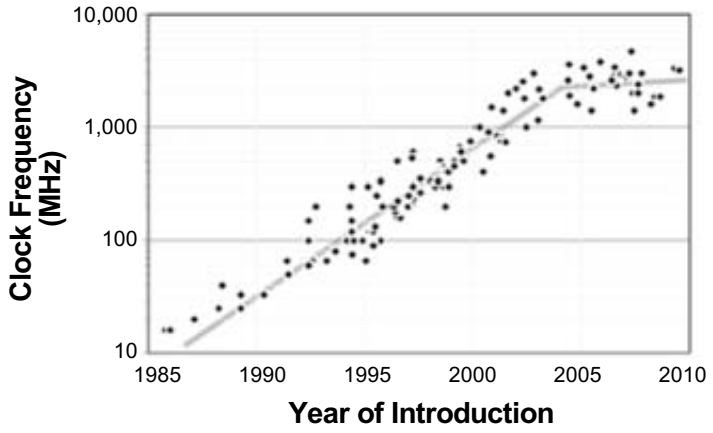

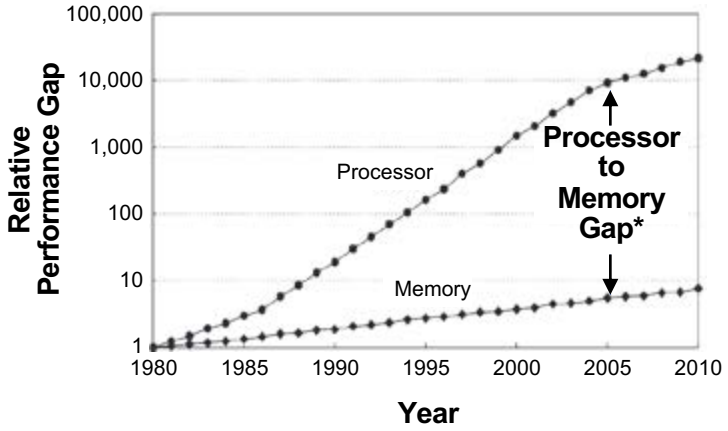

Even though transistors per chip continued to increase, limitations in chip performance and access to memory led to increase in many cores per chip

Source: The Future of Computing Performance: Game Over or Next Level, 2011
* Computer architecture: A quantitative approach – John L. Hennessy and David Patterson, 2007

From Dave Patterson* (now at Google as Distinguished Engineer)

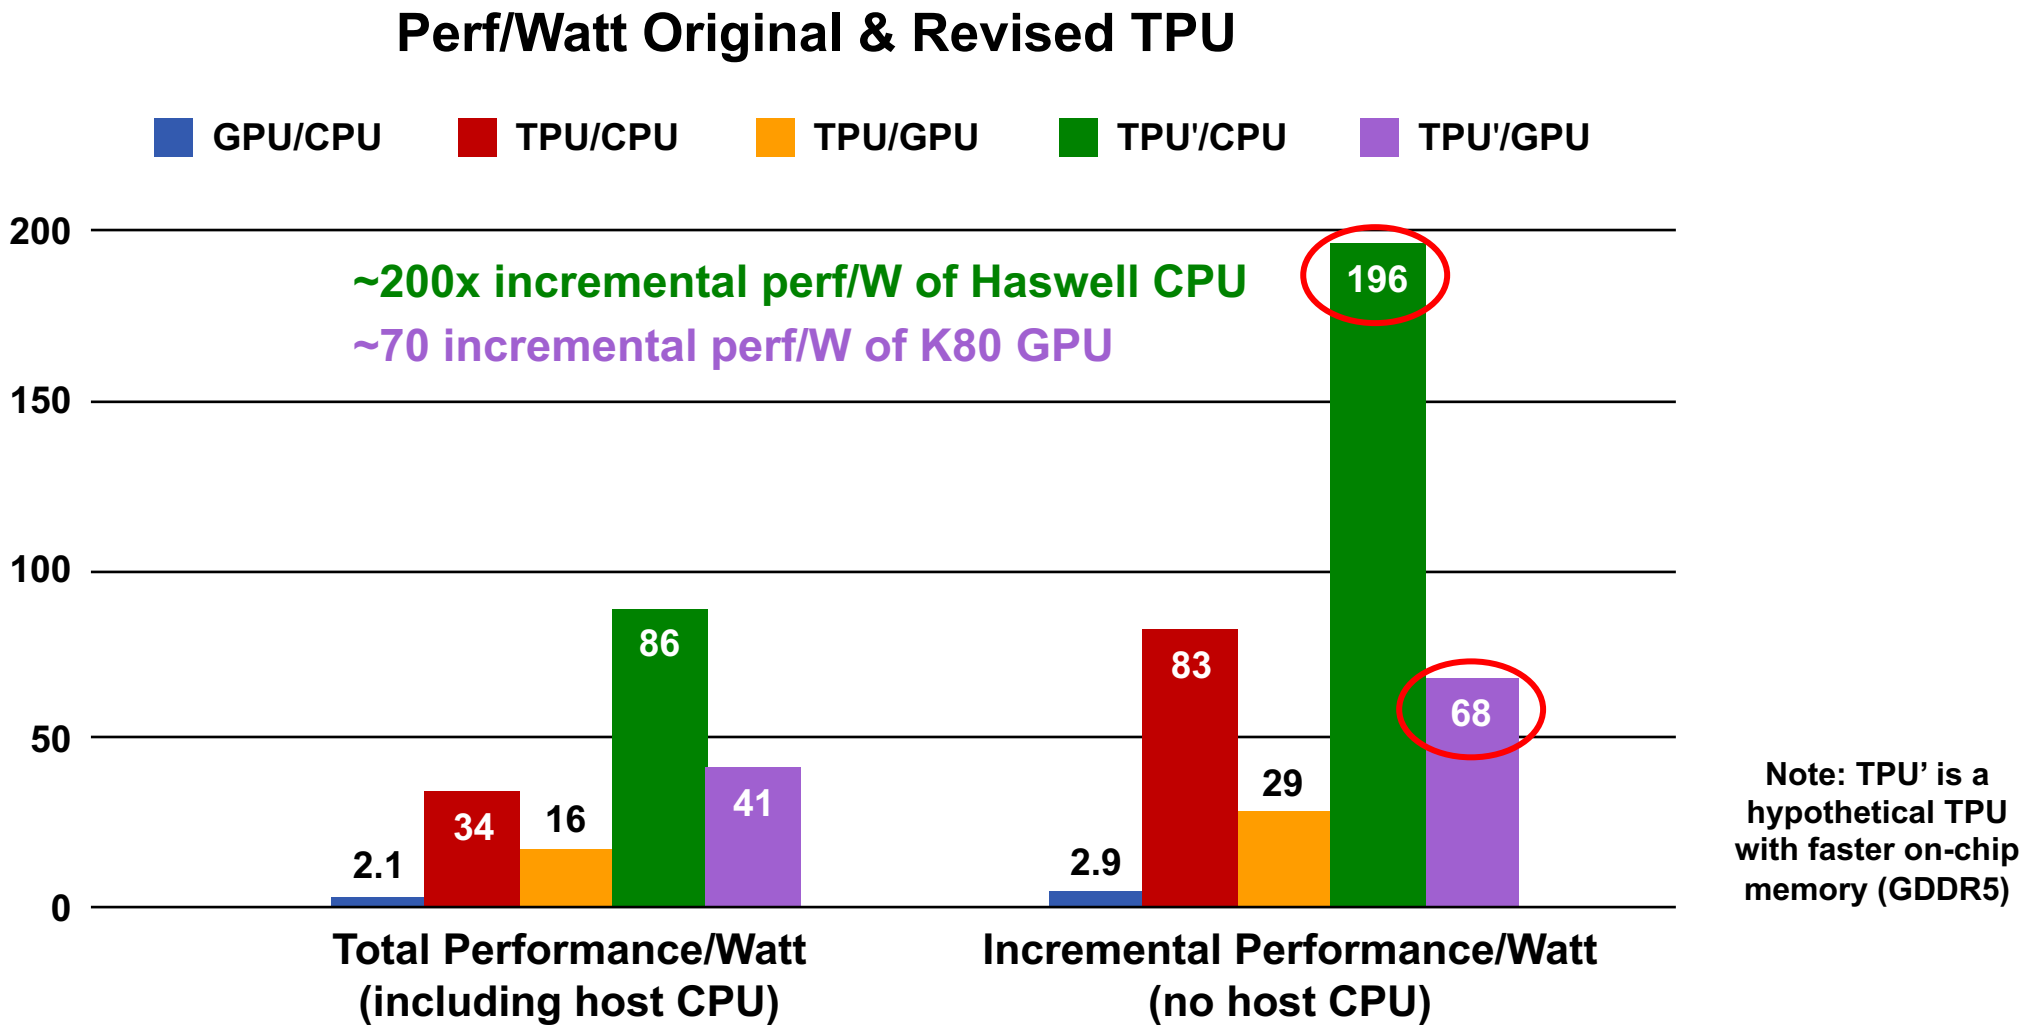

Note: TPU' is a hypothetical TPU with faster on-chip memory (GDDR5)

* Former Professor of computer architectures at Univ. of CA at Berkley
Presented at the NEXT conference, TTIVanguard, Dec 2017

Neural Network Processing Performance

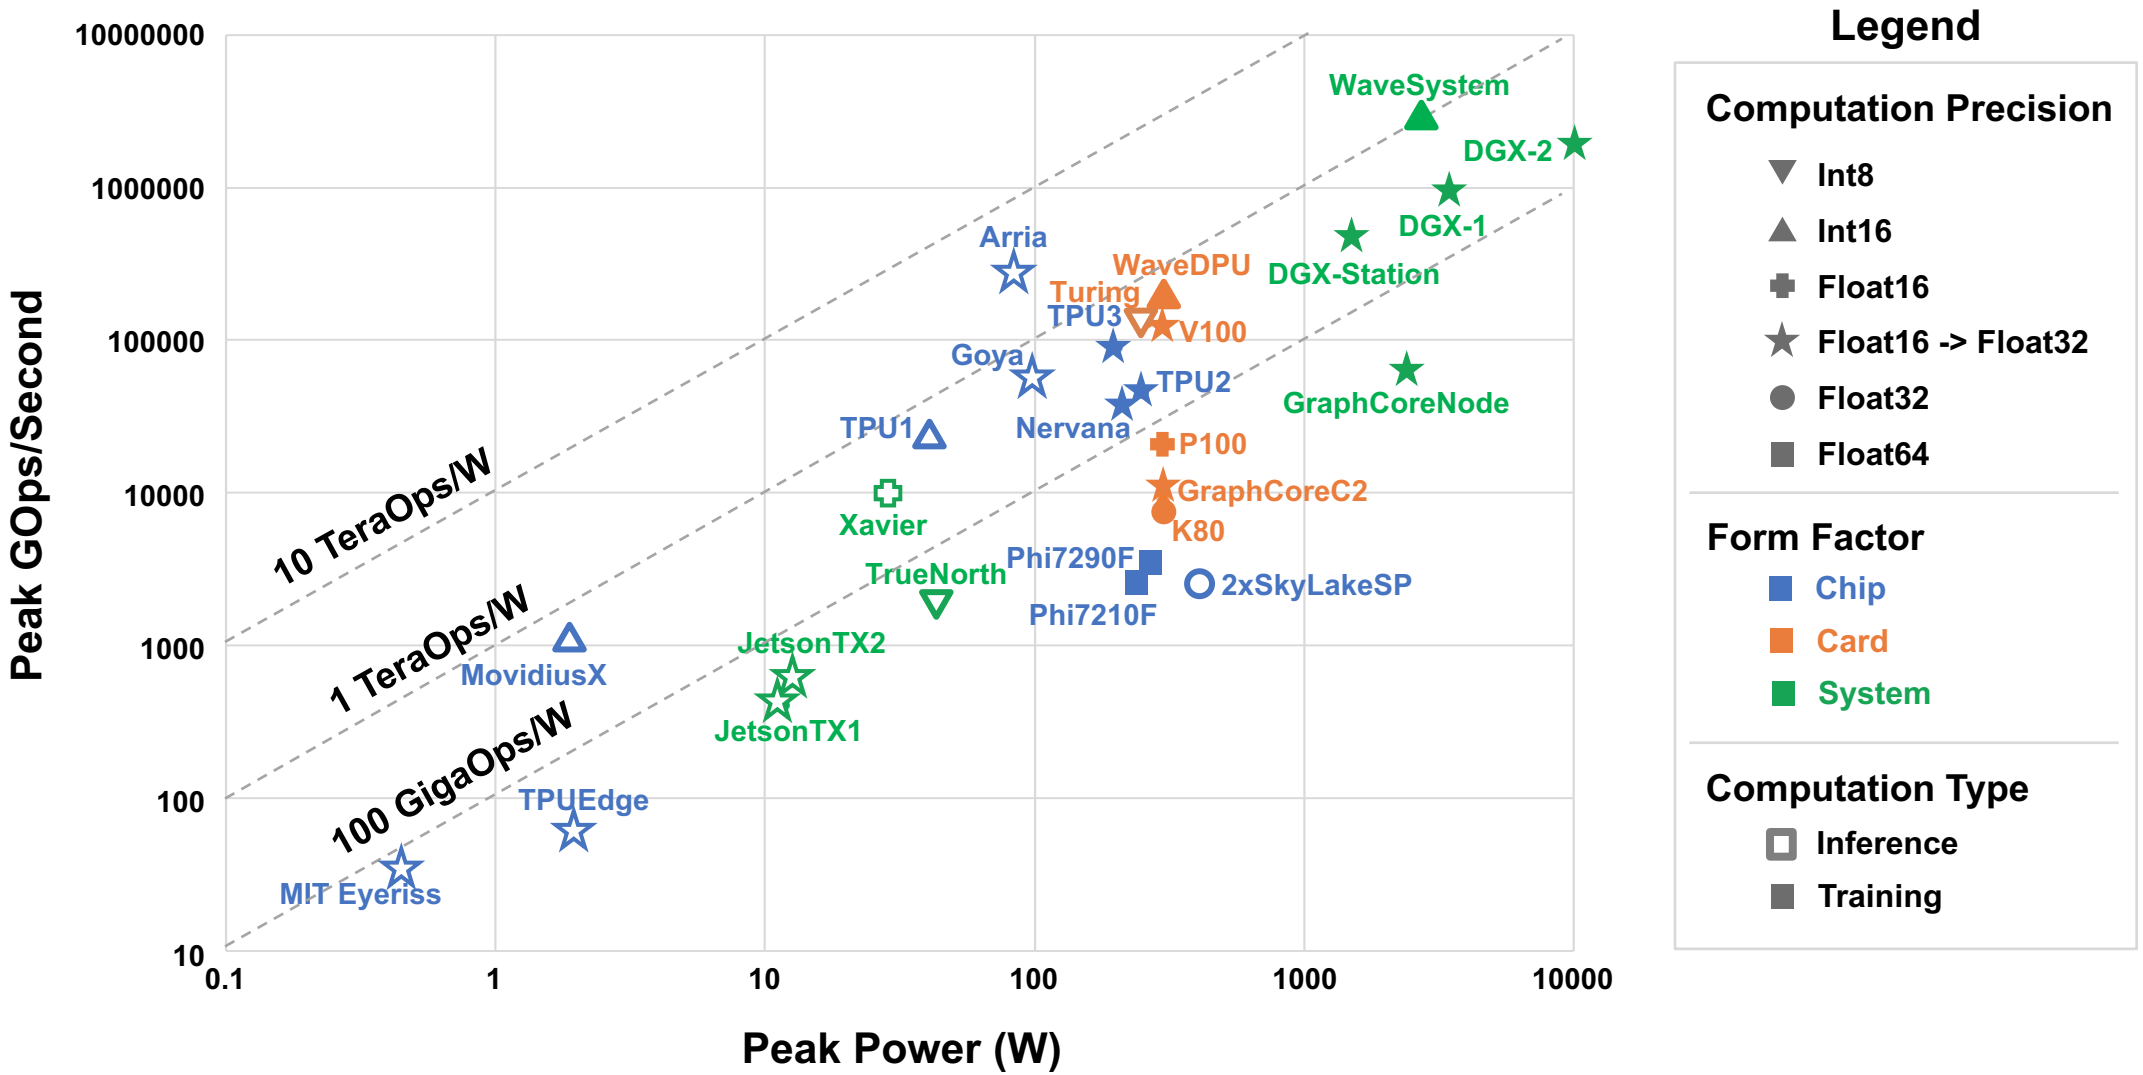

Knowledge-to-Insight: Enable Real-Time (or Near Real-Time) Actions

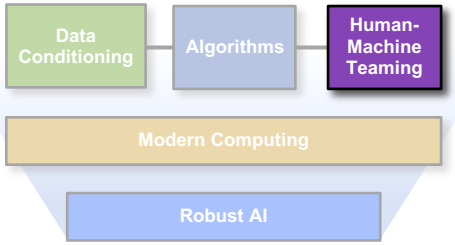

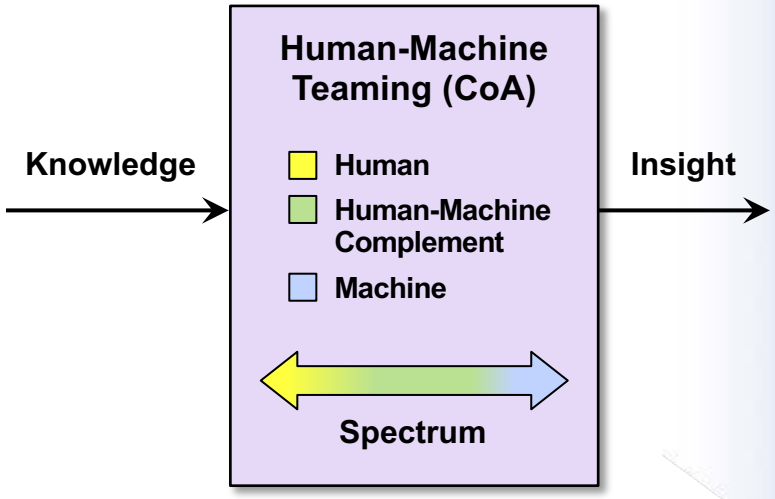

Human-Machine teaming will consist of intelligent assistants enabled by artificial intelligence

http://www.milkeninstitute.org/events/conferences/global-conference/2016/panel-detail/6182

Moderator: Alexandra Suich, U.S. Technology Editor, the Economist
Speakers: Guruduth Banavar, Vice President and Chief Science Officer, Cognitive Computing, IBM
Stuart Russell, Professor, Electrical Engineering and Computer Sciences, University of California, Berkeley; Vice Chair, World Economic Forum Council on AI and Robotics
David M. Siegel, Co-Chairman, Two Sigma

Human-Machine Teaming

Study Finds Cyberthreat Data Overwhelming to Security Workers

A recent Ponemon report shows that organizations neglect to share essential cyberthreat data with board members and C-level executives.

Challenge: Cybersecurity and Big Data

"By 2018 the United States alone faces a shortage of 140,000 to 190,000 people with analytical expertise and 1.5 million managers and analysts with the skills to understand and make decisions based on the analysis of big data."[1]

[1] McKinsey&Company (May 2011). "Big data: The next frontier for innovation, competition, and productivity." Available at: http://www.mckinsey.com/insights/MGI/Research/Technology_and_Innovation/Big_data_The_next_frontier_for_Innovation

DIRECTORATE FOR EDUCATION AND HUMAN RESOURCES 4

- **Cyber security data overwhelms overworked analysts**
- **The US / DoD faces serious workforce shortages in cyber security expertise**

Artificial Intelligence Application to Cyber Security

Offense Stages	Prepare (Recon)	Engage	Maintain Presence	Achieve Effect & Assess Damage
Impact	Know the target	Enable attack process	Support persistence	Attack effectiveness

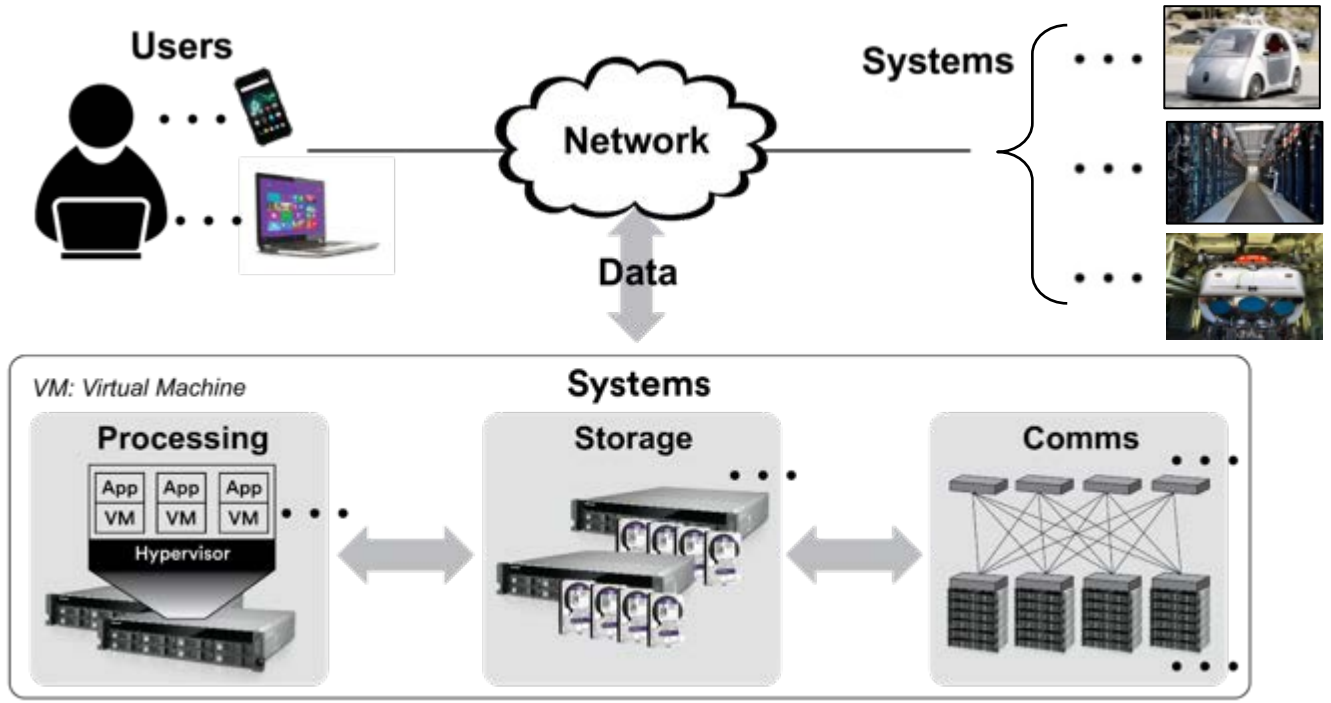

Defense Stages	Identify (Recon)	Protect	Detect	Respond	Recover
Impact	Focused defense	Deflect attacks	ID new attacks	Stop attacks	"Mission" fight through

COA – Course of Action
NLP – Natural Language Processing

Artificial Intelligence Application to Cyber Security

Offense Stages	Prepare (Recon)	Engage	Maintain Presence	Achieve Effect & Assess Damage
Impact	Know the target	Enable attack process	Support persistence	Attack effectiveness

Predictive Analytics: Monitor and anticipate the threat

Clustering and classification: Automatically Infer / Identify Critical Assets

Information Extraction: Handle Big Data to detect and characterize attacks

Planning and Optimization: Automatically respond to attacks and suggest COAs

Recommender Systems & NLP: Team w/ humans to enact solutions

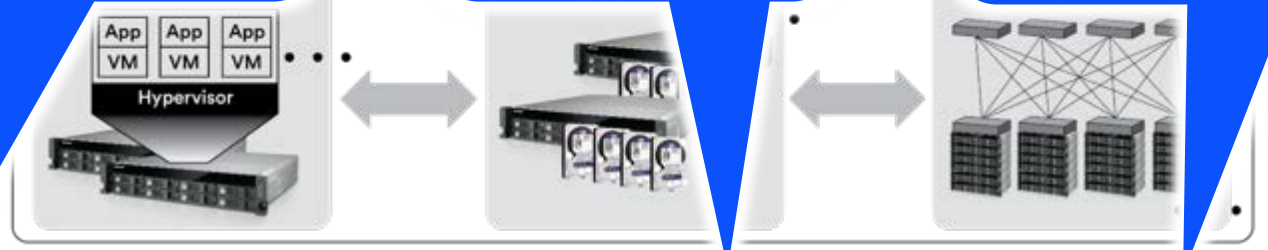

Defense Stages	Identify (Recon)	Protect	Detect	Respond	Recover
Impact	Focused defense	Deflect attacks	ID new attacks	Stop attacks	"Mission" fight through

COA – Course of Action
NLP – Natural Language Processing

Machine Learning to Reducing Cyber Analysts Workload

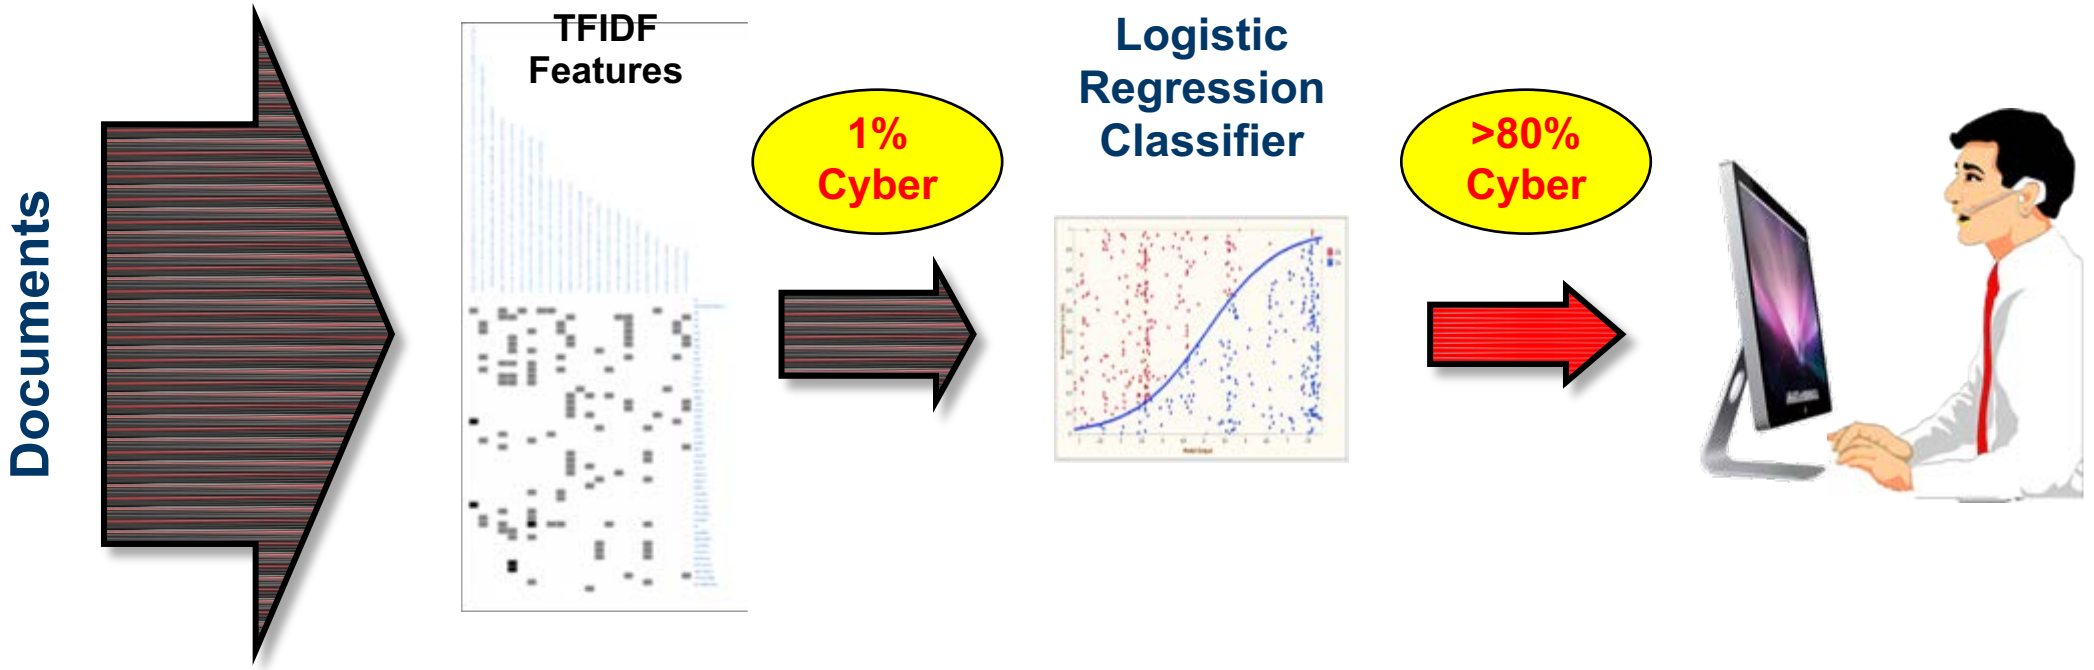

- **Feature Generation**
  - Source-dependent extraction/processing
  - Stemming (hack, hacker, hacks, hacking)
  - Term Frequency Inverse Document Frequency (TFIDF)
- **Logistic Regression Classifier**
  - Rapid classification and training
  - Robust in face of mislabeling
  - Good initial performance

Machine Learning to Prioritize Vulnerabilities for Protection

- Current practice classifies the severity of vulnerabilities by a fixed formula from simplest to worse ("CVSS score")
- Natural Language Processing (NLP) associates free-form descriptions of vulnerabilities with critical exploits observed in practice

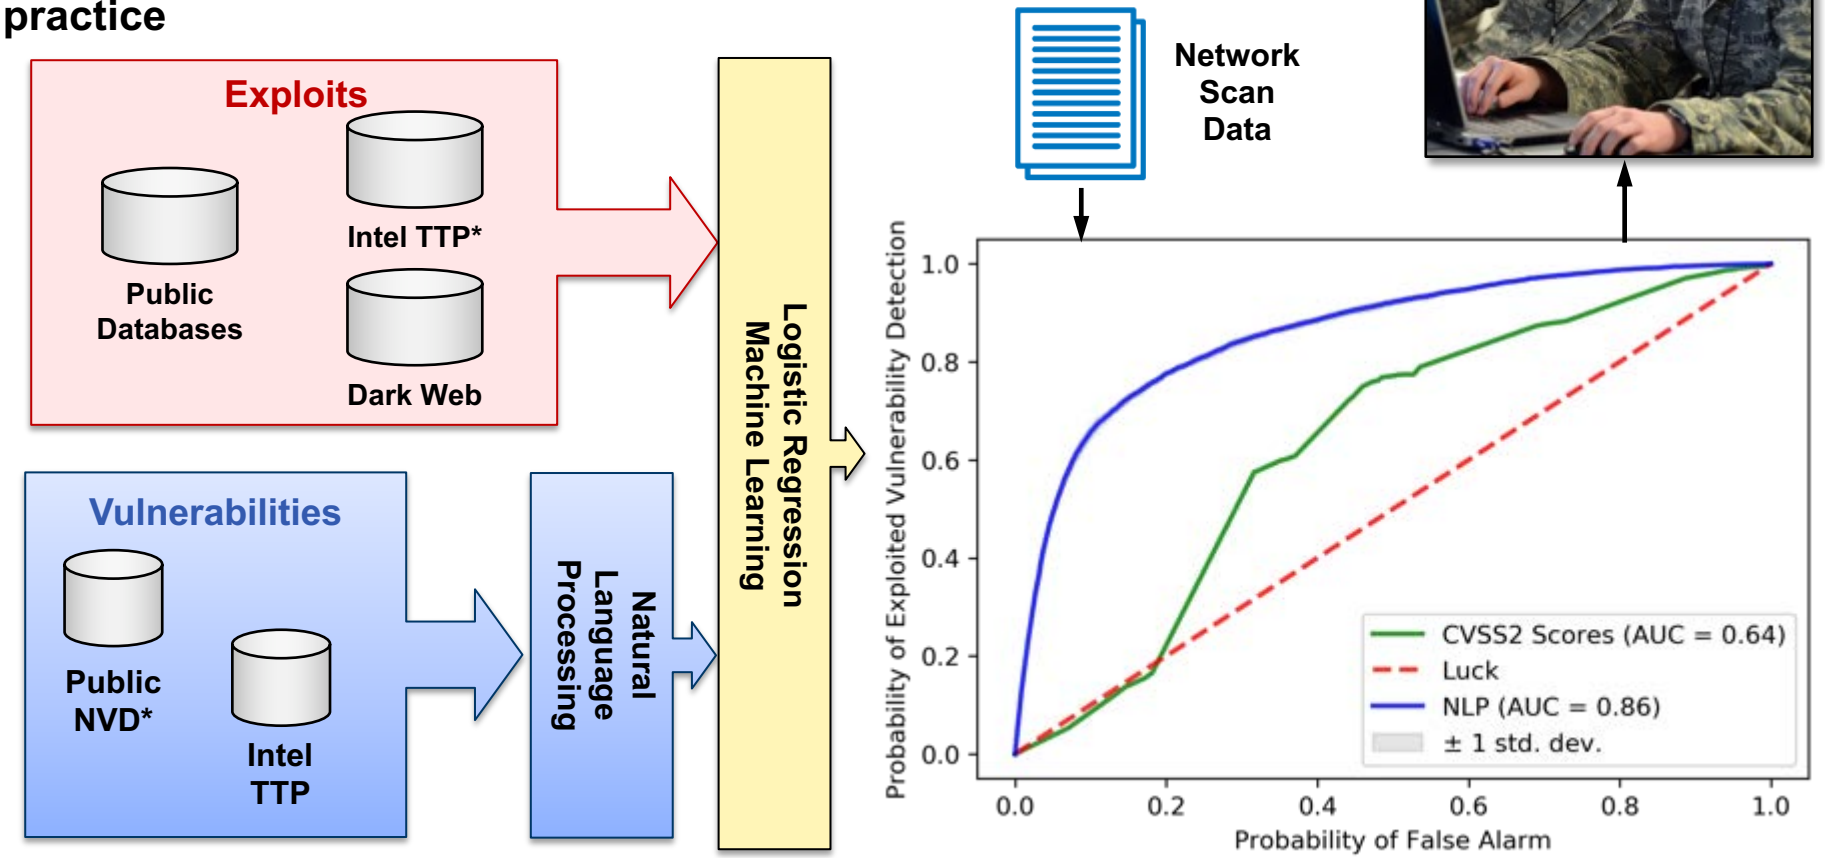

- NLP processing gives better results than CVSS scores in predicting vulnerabilities with exploits
  - Most vulnerabilities are never exploited
- Allows for easy customization according to particular threats and TTPs
- Approach potentially can lead to cyber analysts higher effectiveness

CVSS: Common Vulnerability Scoring System (a rule based approach)
NVD: National Vulnerability Database
TTP: Tactics, Techniques, and Procedures

Outline

- Background
- Lay-of-the-Land
  - AI Canonical Architecture
  - Summary of Study Outreach and Highlights
- Robust AI
- Recommendations
- Summary

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	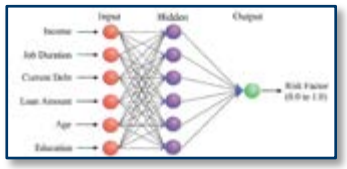	Explainability, dimensionality reduction, feature importance inference
Validation & Verification	Algorithms need to meet mission specifications		Robust training, "portfolio" methods, regularization
Security	System vulnerable to adversarial action (both cyber and physical)		Model failure detection, red teaming
Policy, Ethics, Safety, and Training	Unwanted actions when controlling heavy or dangerous machinery		Risk sensitivity, robust inference, high decision thresholds

Robust AI: Preserving Trust

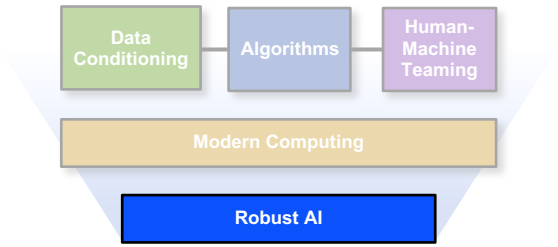

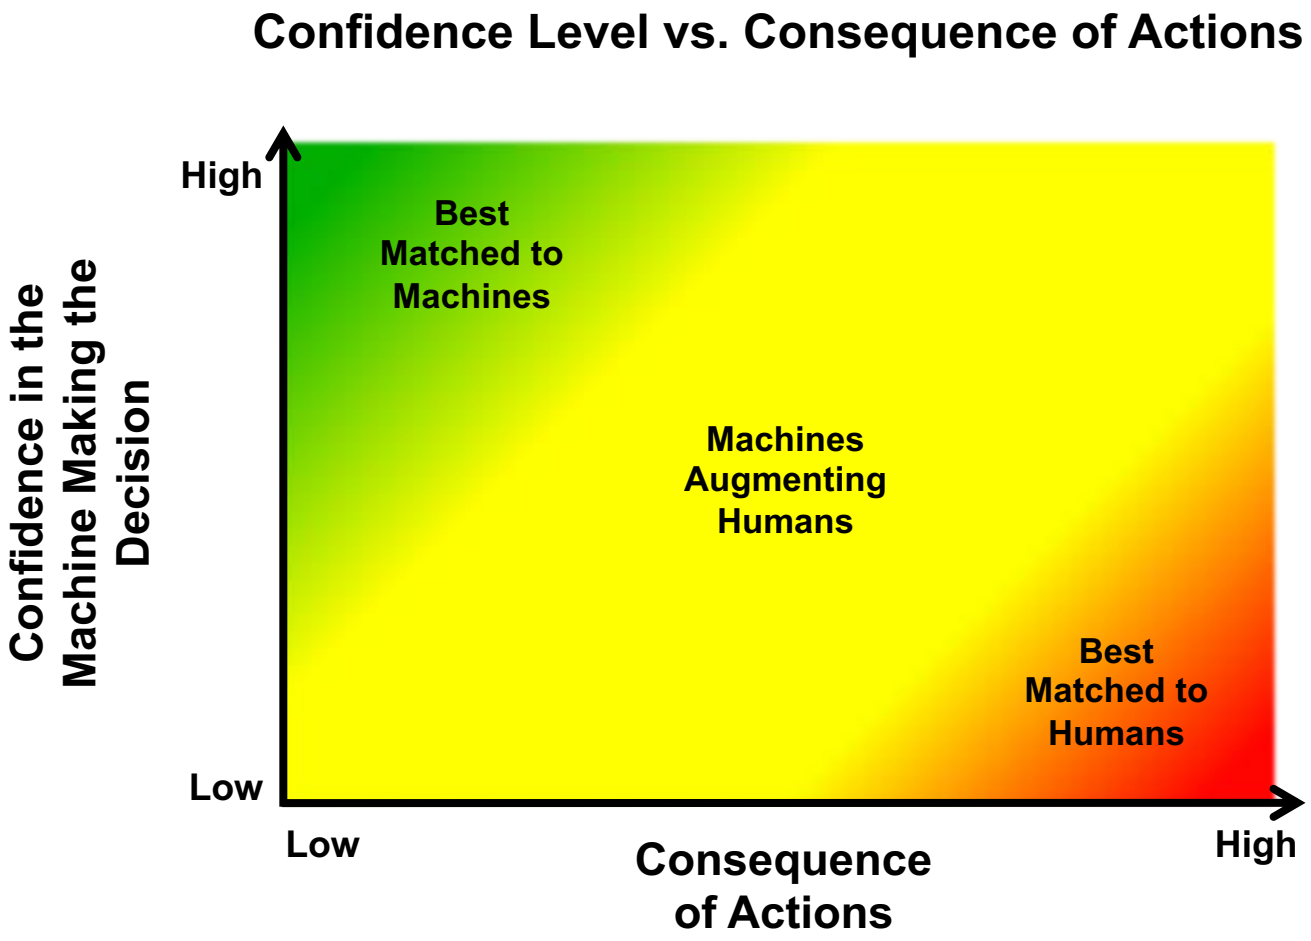

Examples of System Capability Space

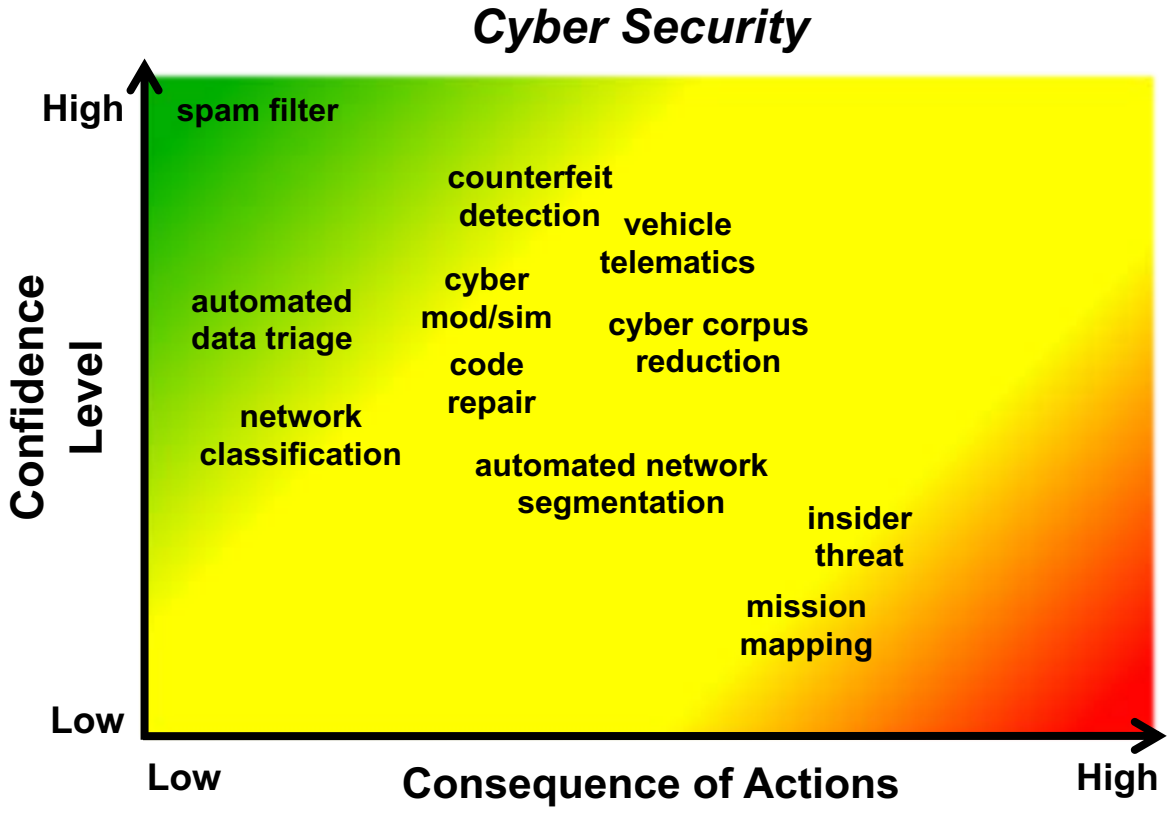

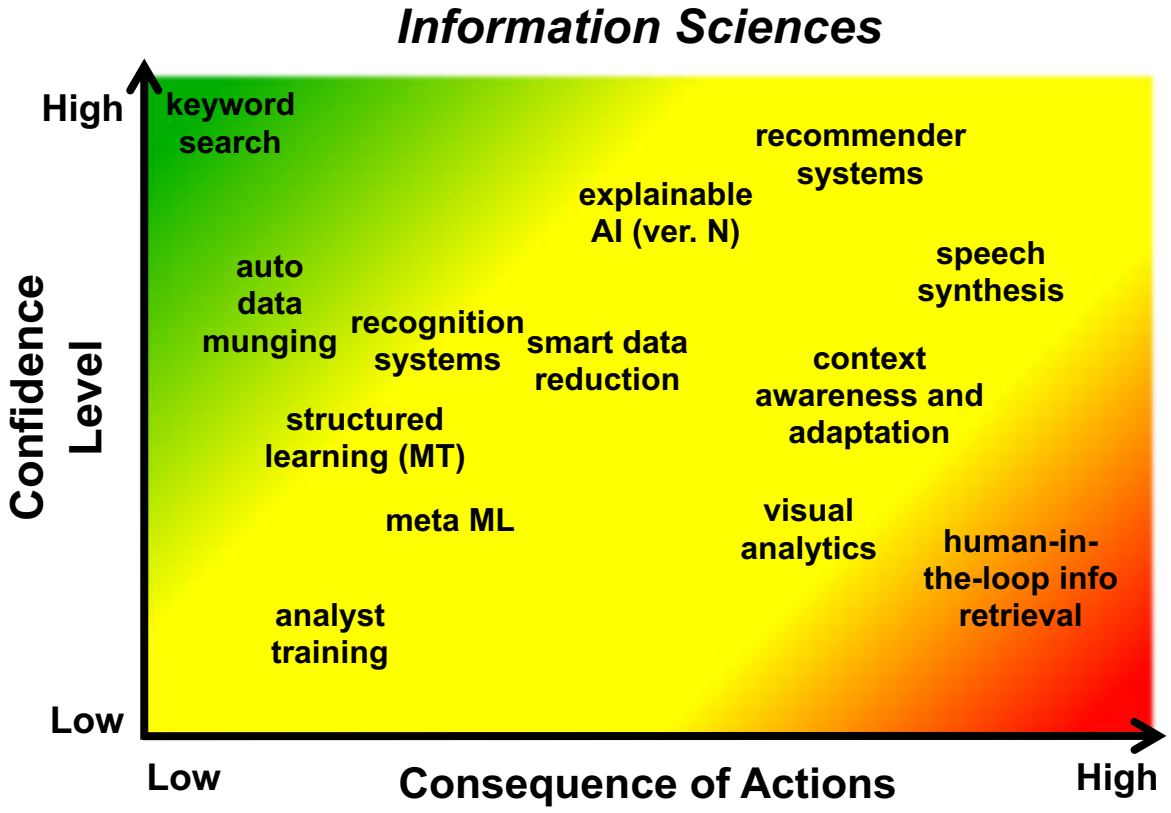

Best Matched to Machines | Machines Augmenting Humans | Best Matched to Humans

Adversarial AI

- **By gaining access to an AI system, can an adversary learn, and then introduce, imperceptible perturbations to inputs that render the system un-usable?**

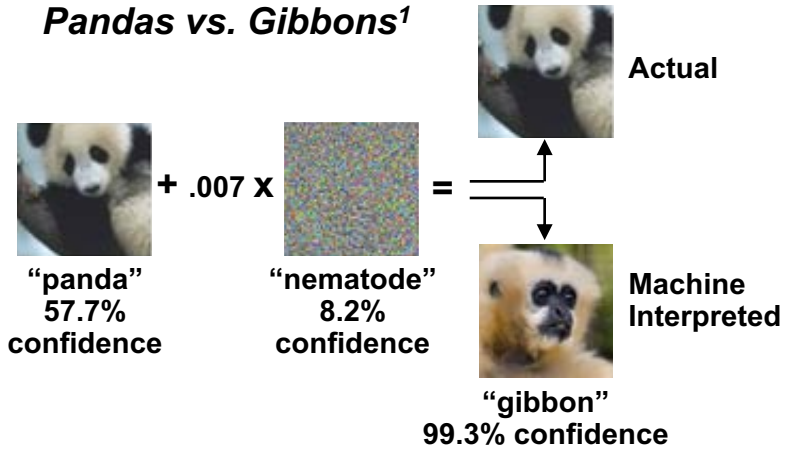

Model-Specific "Universal" Perturbation[2]

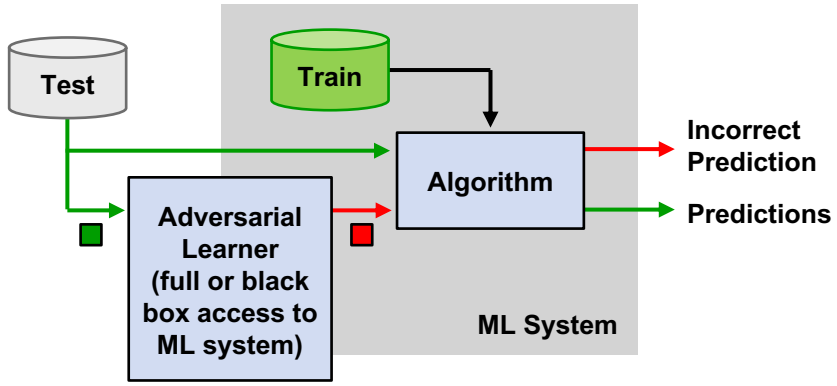

2D Physical Attack (stickers)[3]

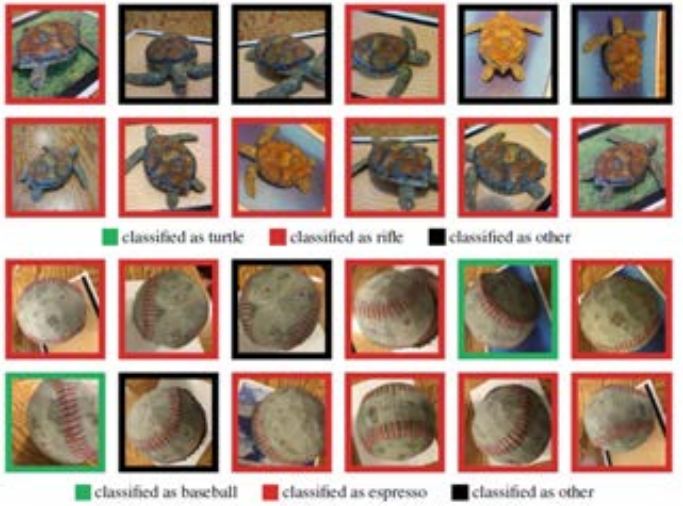

Models	Adversarial	Misclassified	Correct
Turtle	82%	16%	2%
Baseball	59%	31%	10%

3D Physical Attack (printed model)[4]

Securing machine learning algorithms and their data pose an existential challenge to fielding practical AI systems
Robust AI should be central to any AI program to achieve trust

[1] Goodfellow et al., "Explaining and Harnessing Adversarial Examples," ICLR 2015.

[2] Moosavi-Dezfooli, "Universal Adversarial Perturbations," CVPR 2017.

[3] Evtimov et al., "Robust Physical-World Attacks on Deep Learning Models," arXiv preprint arXiv:1707.08945, 2017

[4] Athalye et al, "Synthesizing Robust Adversarial Examples," ICML 2018 submission, arXiv preprint arXiv:1707.07397

Outline

- Background
- Lay-of-the-Land
  - AI Canonical Architecture
  - Summary of Study Outreach and Highlights
- Robust AI
- Recommendations
- Summary

Three Horizons for Artificial Intelligence Investments

Impact

Horizon 1: Content-Based Insight

Apply AI to gain insight on content of interest present in disparate types of data

- Reduce user workload
- Improve confidence in AI
- Lower consequence of actions
- Ascertain robust AI

Horizon 2: Collaboration-Based Insight

Extend AI roles to include multiple human-machine teams working together

- Enrich insight
- Enable collaboration across intelligent machines
- Migrate more routine tasks to AI enabled systems

Horizon 3: Context-Based Insight

Incorporate semantic context* in the application of AI

- Reduce time from data to insight
- Increase symbiosis between humans and machines

Time

* Context: a frame that surrounds the event and provides resources for its appropriate interpretation https://en.wikipedia.org/wiki/Context_(language_use)

Broad Recommendations on Way Forward
- Relevant Across all Horizons -

Challenge:	Staying competitive against peer threats
Recommendations:	• Establish a culture of innovation, rapid experimentation, and deployment • Focus on end-to-end capabilities across all elements of AI canonical architecture

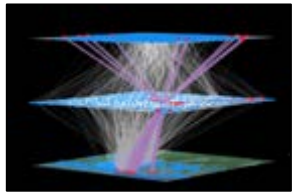


Approach		Benefit
• Demonstrate AI capabilities on realistic applications	⇨	• Early users buy-in
• Adopt a set of quantitative metrics and benchmarks	⇨	• Quantifiable improvements and explainability
• Perform red/blue tabletop exercises	⇨	• Reduced cost using a blend of prior knowledge, simulated and real data
• Foster competition via grand challenge problems	⇨	• Commercial companies and academia involvement
• Maintain a deep bench of AI talent (multi-disciplinary)	⇨	• Award AI certificate from top-notch university

AI Investment Recommendations
- Definition of Terms -

Recommendations	Data Conditioning	Algorithms	Modern Computing	Human-Machine Teaming	Robust AI
	Develop common database formats	*Demonstrate algorithms across multi-domains*	*Achieve real-time performance on AI systems*	*Augment human capabilities by leveraging intelligent machines*	*Build a culture of rapid development cycles with user in the loop*
Horizon 1 **Content-Based Insight** **1–2 Years**	**Content-Based Insight:** **Apply AI to gain insight on content of interest present in disparate types of data**				
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	**Collaboration-Based Insight:** **Extend AI roles to include multiple human-machine teams working together**				
Horizon 3 **Context-Based Insight** **5+ Years**	**Context-Based Insight:** **Incorporate semantic context* in the application of AI**				

* Context: a frame that surrounds the event and provides resources for its appropriate interpretation
https://en.wikipedia.org/wiki/Context_(language_use)

AI Investment Recommendations
- Specific Areas -

Recommendations	**Data Conditioning** *Develop common database formats*	**Algorithms** *Demonstrate algorithms across multi-domains*	**Modern Computing** *Achieve real-time performance on AI systems*	**Human-Machine Teaming** *Augment human capabilities by leveraging intelligent machines*	**Robust AI** *Build a culture of rapid development cycles with user in the loop*
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• Blend unsupervised and supervised learning • Demonstrate reinfor. learning	• Accelerate model generation • Deploy computing to the edge	• Relationship graphs updated in real-time • Advance natural language processing (NLP)	• Develop robust AI metrics • Demo adv. learning • Design adversarial AI countermeasures
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• Advance algorithm accuracy through collaboration • Exploit physics and causal relationships	• Compute across distributed platf. • Reduce SWaP for embedded and IoT devices	• Transparent human-machine teams • Collab. based on achieving mission goals	• Strengthen full end-to-end system security • Enable explainable AI
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• Train with limited data • Low-shot or one-shot learning • Context aware learning • Advance research on goal reasoning	• Advance cognitive computing • Deterministic to probabilistic comp. • Scalability based on mission objectives	• Understand & shape human-machine networks • Sentiment analysis • Scale to very large human-machine teams	• Gain user confidence through probabilistic CoAs* • Leverage context to defend against adv. learning attacks • Employ formal math to verify perf.
DoD Investment Priority	***High***	***Medium***	***Medium***	***High***	***High***

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications
* CoAs: Courses of Action

AI Investment Recommendations
- Data Conditioning -

Recommendations	Data Conditioning Develop common database formats	Algorithms Demonstrate algorithms across multi-domains	Modern Computing Achieve real-time performance on AI systems	Human-Machine Teaming Augment human capabilities by leveraging intelligent machines	Robust AI Build a culture of rapid development cycles with user in the loop
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• Blend unsupervised and supervised learning • Demonstrate reinforcement learning	• Accelerate model generation • Deploy computing to the edge	• Relationship graphs updated in real-time • Advance natural language processing (NLP)	• Develop robust AI metrics • Demo adv. learning • Design adversarial AI countermeasures
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• Advance algorithm accuracy through collaboration • Exploit physics and causal relationships	• Compute across distributed platforms • Reduce SWaP for embedded and IoT devices	• Transparent human-machine teams • Collab. based on achieving mission goals	• Strengthen full end-to-end system security • Enable explainable AI
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• Train with limited data • Low-shot or one-shot learning • Context aware learning • Advance research on goal reasoning	• Advance cognitive computing • Deterministic to probabilistic computing • Scalability based on mission objectives	• Understand & shape human-machine networks • Sentiment analysis • Scale to very large human-machine teams	• Gain user confidence through probabilistic courses of action • Leverage context to defend against adv. learning attacks • Employ formal math to verify perf.
DoD Investment Priority	**High**	**Medium**	**Medium**	**High**	**High**

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications

Recommendations (1/5)

Investment Area 1:	Data Conditioning
Recommendation Objective:	Develop common database formats (including storing of all intermediate results)



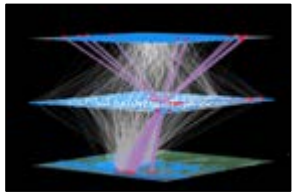


Specific Recommendations

- Automate data labeling
- Create data set benchmarks
- Leverage real data, simulated, and prior knowledge
- Access intermediate results
- Generate models from limited information
- Exploit the what, how, who and why to build context
- Operate in denied and degraded environments

Benefits

- Automated tools to discover, link, and store heterogenous data
- Significant reduction in time for AI pipeline development
- Automated cleaning and preprocessing of data
- Increased understanding of AI intermediate outputs
- Conditioning of data for low-shot or one-shot learning

AI Investment Recommendations
- Algorithms -

Recommendations	Data Conditioning Develop common database formats	Algorithms Demonstrate algorithms across multi-domains	Modern Computing Achieve real-time performance on AI systems	Human-Machine Teaming Augment human capabilities by leveraging intelligent machines	Robust AI Build a culture of rapid development cycles with user in the loop
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• **Blend unsupervised and supervised learning** • **Demonstrate reinforcement learning**	• Accelerate model generation • Deploy computing to the edge	• Relationship graphs updated in real-time • Advance natural language processing (NLP)	• Develop robust AI metrics • Demo adv. learning • Design adversarial AI countermeasures
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• **Advance algorithm accuracy through collaboration** • **Exploit physics and causal relationships**	• Compute across distributed platforms • Reduce SWaP for embedded and IoT devices	• Transparent human-machine teams • Collab. based on achieving mission goals	• Strengthen full end-to-end system security • Enable explainable AI
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• **Train with limited data** • **Low-shot or one-shot learning** • **Context aware learning** • **Advance research on goal reasoning**	• Advance cognitive computing • Deterministic to probabilistic computing • Scalability based on mission objectives	• Understand & shape human-machine networks • Sentiment analysis • Scale to very large human-machine teams	• Gain user confidence through probabilistic courses of action • Leverage context to defend against adv. learning attacks • Employ formal math to verify perf.
DoD Investment Priority	High	**Medium**	Medium	High	High

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications

Recommendations (2/5)

Investment Area 2:	Algorithms
Recommendation Objective:	Demonstrate AI algorithms across multi-domains




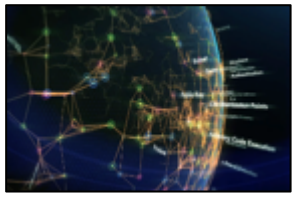

Specific Recommendations

- Blend unsupervised and supervised learning
- Demonstrate reinforcement learning
- Advance algorithm accuracy through collaboration
- Exploit physics and casual relationships
- Train with limited data
- Train based on low-shot or one-shot learning
- Invent new algorithms for context aware learning
- Advance research on goal reasoning*

Benefits

- Less dependency on large volume of labeled data
- Learn-to-learn using generative adversarial learning (i.e., generator and discriminator)
- Broadening of sensor modalities and social-cultural networks via algorithm collaboration
- Applying AI to both defense and offense through increased context awareness
- Enabling of AI agents to deliberate and self-adjust their goals in a complex environment*

* David Aha, "Goal Reasoning: Foundations, Emerging Applications, and Prospects," AI Magazine, Summer 2018

AI Investment Recommendations
- Modern Computing -

Recommendations	Data Conditioning Develop common database formats	Algorithms Demonstrate algorithms across multi-domains	Modern Computing Achieve real-time performance on AI systems	Human-Machine Teaming Augment human capabilities by leveraging intelligent machines	Robust AI Build a culture of rapid development cycles with user in the loop
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• Blend unsupervised and supervised learning • Demonstrate reinforcement learning	• **Accelerate model generation** • **Deploy computing to the edge**	• Relationship graphs updated in real-time • Advance natural language processing (NLP)	• Develop robust AI metrics • Demo adv. learning • Design adversarial AI countermeasures
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• Advance algorithm accuracy through collaboration • Exploit physics and causal relationships	• **Compute across distributed platforms** • **Reduce SWaP for embedded and IoT devices**	• Transparent human-machine teams • Collab. based on achieving mission goals	• Strengthen full end-to-end system security • Enable explainable AI
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• Train with limited data • Low-shot or one-shot learning • Context aware learning • Advance research on goal reasoning	• **Advance cognitive computing** • **Deterministic to probabilistic computing** • **Scalability based on mission objectives**	• Understand & shape human-machine networks • Sentiment analysis • Scale to very large human-machine teams	• Gain user confidence through probabilistic courses of action • Leverage context to defend against adv. learning attacks • Employ formal math to verify perf.
DoD Investment Priority	High	Medium	**Medium**	High	High

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications

Recommendations (3/5)

Investment Area 3:	Modern Computing
Recommendation Objective:	Achieve real-time performance on AI systems

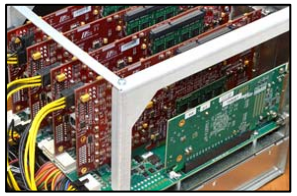

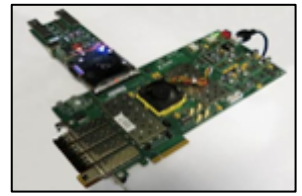

Specific Recommendations

- Accelerate model generation
- Deploy computing to the edge
- Compute across distributed platforms
- Reduce SWaP for embedded and IoT devices
- Advance cognitive computing
- Develop architectures spanning deterministic to probabilistic computing (incl. variable precision)
- Scale computing systems to meet mission objectives

Benefits

- Reduction of time to train AI models
- Enabling SWaP constrained tactical and embedded systems
- Improving ability to sense, reason, and respond to stimulus via improved cognitive models of the brain
- New designs to address slowing of Moore's law via DSA* (including HW and SW codesign)
- Computing complexity and tools adapted in real-time (e.g., from cloud-based env. to embedded systems)

*50 Years of Computer Architecture: From the Mainframe CPU to the Domain-Specific-Architecture TPU and the Open RISC-V Instruction Set by David Patterson, Dec 2017

AI Investment Recommendations
- Human-Machine Teaming -

Recommendations	Data Conditioning Develop common database formats	Algorithms Demonstrate algorithms across multi-domains	Modern Computing Achieve real-time performance on AI systems	Human-Machine Teaming **Augment human capabilities by leveraging intelligent machines**	Robust AI Build a culture of rapid development cycles with user in the loop
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• Blend unsupervised and supervised learning • Demonstrate reinforcement learning	• Accelerate model generation • Deploy computing to the edge	• **Relationship graphs updated in real-time** • **Advance natural language processing (NLP)**	• Develop robust AI metrics • Demo adv. learning • Design adversarial AI countermeasures
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• Advance algorithm accuracy through collaboration • Exploit physics and causal relationships	• Compute across distributed platforms • Reduce SWaP for embedded and IoT devices	• **Transparent human-machine teams** • **Collab. based on achieving mission goals**	• Strengthen full end-to-end system security • Enable explainable AI
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• Train with limited data • Low-shot or one-shot learning • Context aware learning • Advance research on goal reasoning	• Advance cognitive computing • Deterministic to probabilistic computing • Scalability based on mission objectives	• **Understand & shape human-machine networks** • **Sentiment analysis** • **Scale to very large human-machine teams**	• Gain user confidence through probabilistic courses of action • Leverage context to defend against adv. learning attacks • Employ formal math to verify perf.
DoD Investment Priority	High	Medium	Medium	**High**	High

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications

Recommendations (4/5)

Investment Area 4:	Human-Machine Teaming
Recommendation Objective:	Augment human capabilities by leveraging intelligent machines

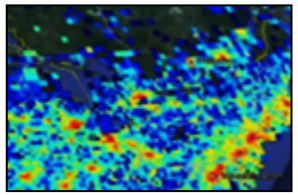

Specific Recommendations

- Relationship graphs updated in real-time
- Advance NLP to improve human-machine teaming
- Create transparency among human-machine teams
- Collaborate based on achieving mission goals
- Understand and shape human-machine networks
- Apply AI to sentiment analysis
- Scale to very large human-machine teams

Benefits

- Strengthening coupling among multiple H-M teams
- Increasing value of resulting insight
- Adapting H-M teams as environment changes
- Broadening of sensor modalities and social-cultural networks via H-M collaborations
- Reasoning based on context
- Reconfiguring of H-M teams to meet scale of mission

Cyber Machine Intelligent Assistant (CyMIA) for Mission Systems

How should I protect my missile range network?

CyMIA Response

You should identify and isolate affected hosts from mission C2 path

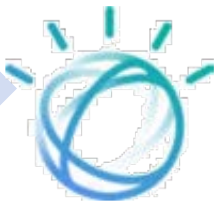

Natural language-based interaction interprets user queries

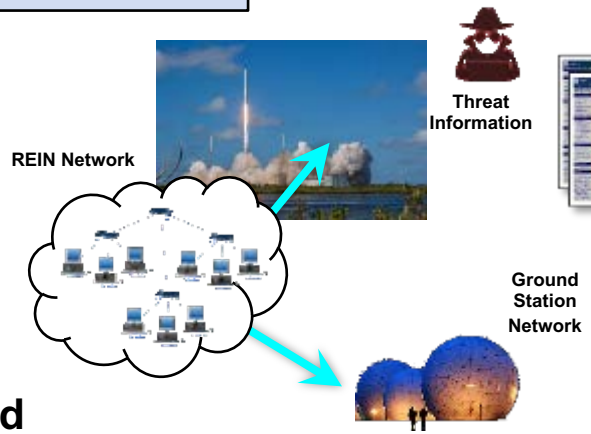

Automatically extract mission, network and threat information

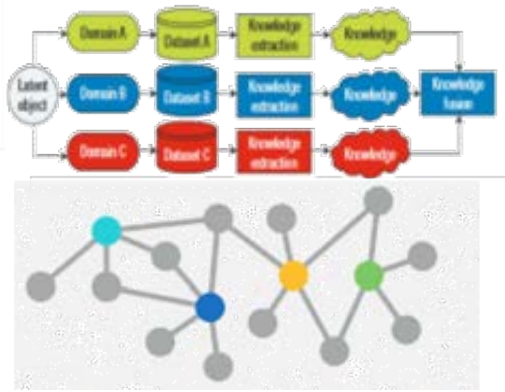

Fuse information into knowledge base, infer relationships

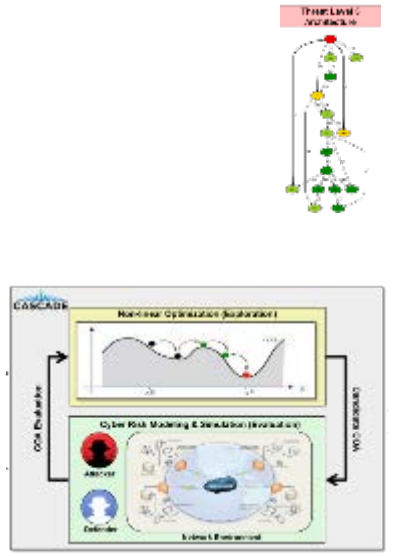

Choose CoA based on threat, knowledge base, and scenario

CyMIA processes natural language input in the context of cyber threats and mission network knowledge to respond with appropriate CoAs (Courses of Action)

Note: CASCADE is an existing Line funded effort
REIN – Range External Interface Network

AI Investment Recommendations
- Robust AI -

Recommendations	Data Conditioning Develop common database formats	Algorithms Demonstrate algorithms across multi-domains	Modern Computing Achieve real-time performance on AI systems	Human-Machine Teaming Augment human capabilities by leveraging intelligent machines	Robust AI Build a culture of rapid development cycles with user in the loop
Horizon 1 **Content-Based Insight** **1–2 Years**	• Automate data labeling • Create data set benchmarks (gold standard)	• Blend unsupervised and supervised learning • Demonstrate reinforcement learning	• Accelerate model generation • Deploy computing to the edge	• Relationship graphs updated in real-time • Advance natural language processing (NLP)	• **Develop robust AI metrics** • **Demo adv. learning** • **Design adversarial AI countermeasures**
Horizon 2 **Collaboration-Based Insight** **3–4 Years**	• Aggregate real data, simulated data, and prior knowledge • Access intermediate results	• Advance algorithm accuracy through collaboration • Exploit physics and causal relationships	• Compute across distributed platforms • Reduce SWaP for embedded and IoT devices	• Transparent human-machine teams • Collab. based on achieving mission goals	• **Strengthen full end-to-end system security** • **Enable explainable AI**
Horizon 3 **Context-Based Insight** **5+ Years**	• Generate models from limited infor. • Exploit the what, how, who, and why to build context • Operate in denied and degraded environments	• Train with limited data • Low-shot or one-shot learning • Context aware learning • Advance research on goal reasoning	• Advance cognitive computing • Deterministic to probabilistic computing • Scalability based on mission objectives	• Understand & shape human-machine networks • Sentiment analysis • Scale to very large human-machine teams	• **Gain user confidence through probabilistic courses of action** • **Leverage context to defend against adv. learning attacks** • **Employ formal math to verify perf.**
DoD Investment Priority	High	Medium	Medium	High	**High**

Note: Investment priorities are relative to DoD uniqueness vs. commercial applications

Recommendations (5/5)

Investment Area 5:	Robust AI
Recommendation Objective:	Build a culture of rapid development and experimentation cycles with user in the loop




Specific Recommendations

- Develop robust AI metrics
- Demonstrate adversarial learning
- Design adversarial AI countermeasures
- Strengthen full end-to-end system security
- Enable explainable AI
- Gain user confidence through probabilistic CoAs
- Defend against adv. learning attacks
- Employ formal math to verify performance

Benefits

- Rigorous AI value assessment
- Avoid element of surprise
- Attending to both physical and cyber security
- Establishment of AI hackathons (leveraging for example Kaggle competitions and serious games)
- Making sure CoAs recommendations are derived from robust insight

Transitioning AI Capabilities to Users

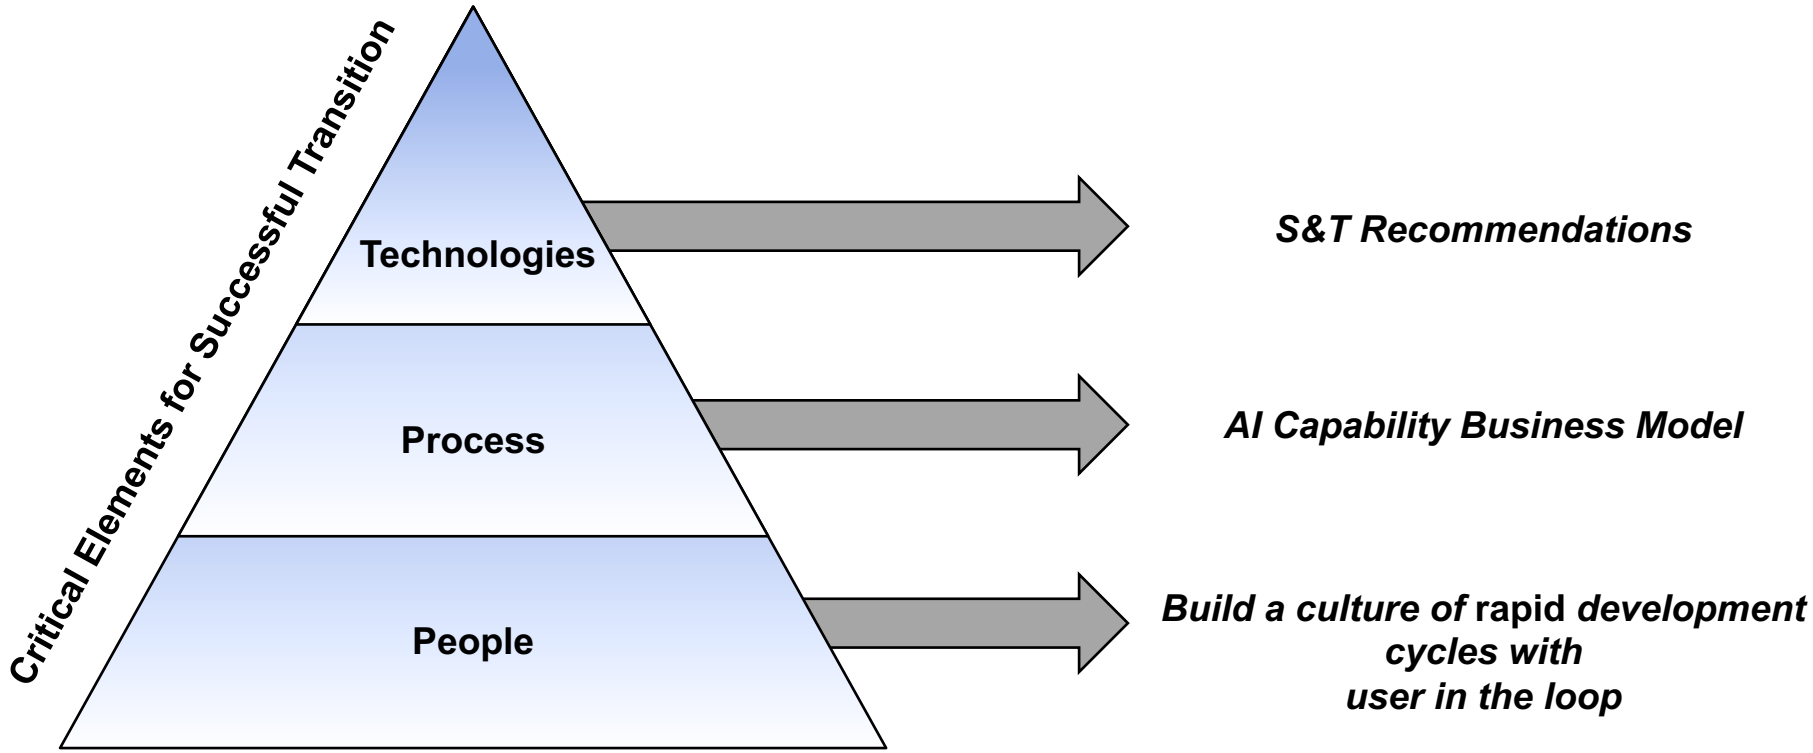

AI Capability Business Model
- Rapid Experimentation Between Research and Users -

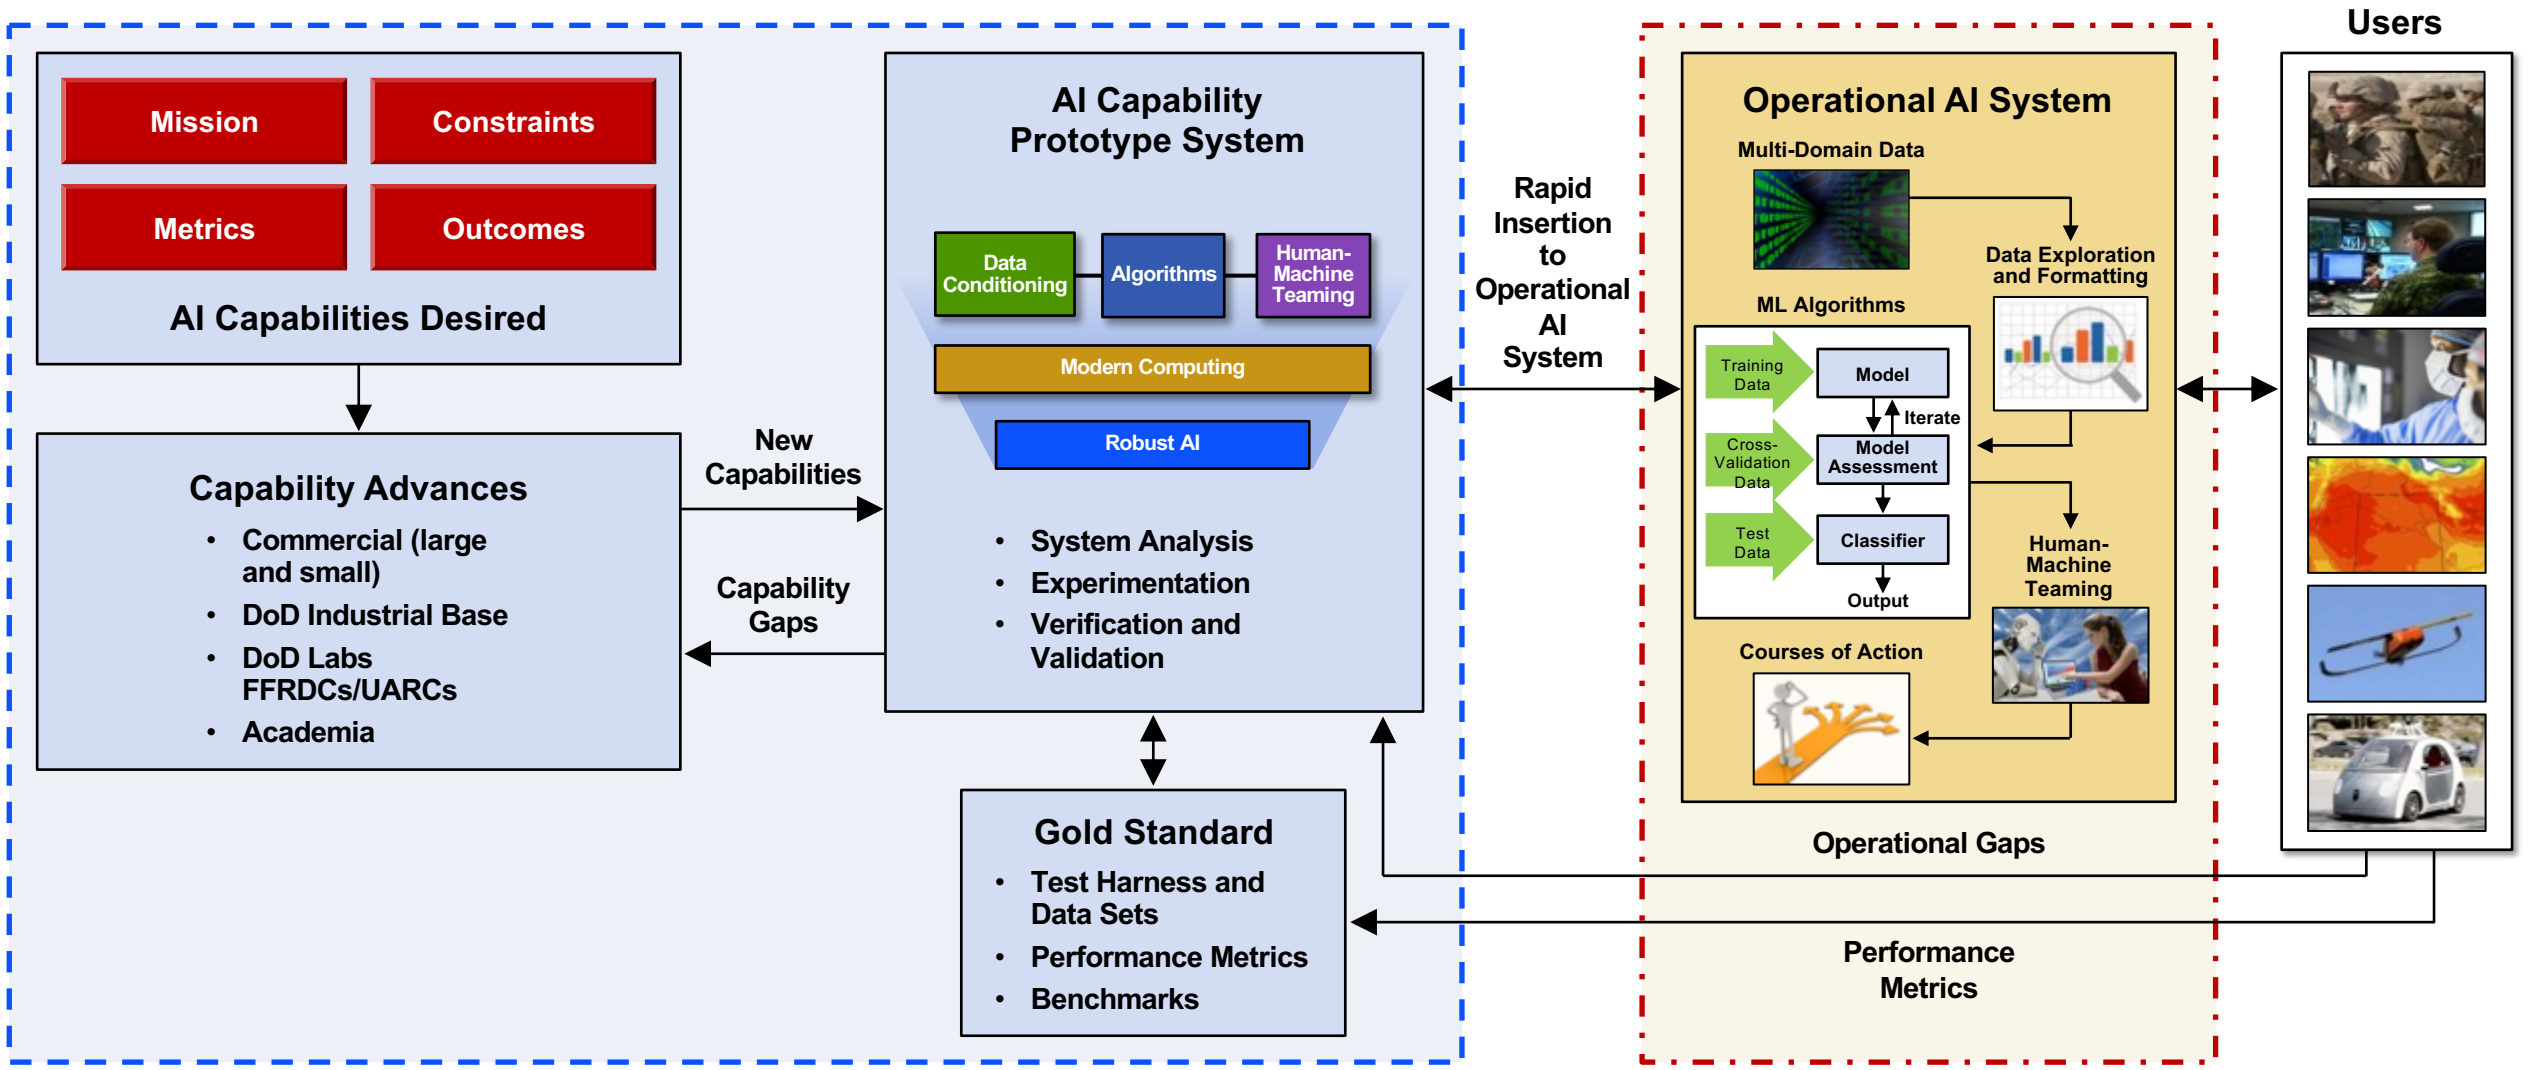

Rapid Prototyping Laboratory DoD/Commercial Industrial Base

Strategic Planning Model for Effective AI Implementation

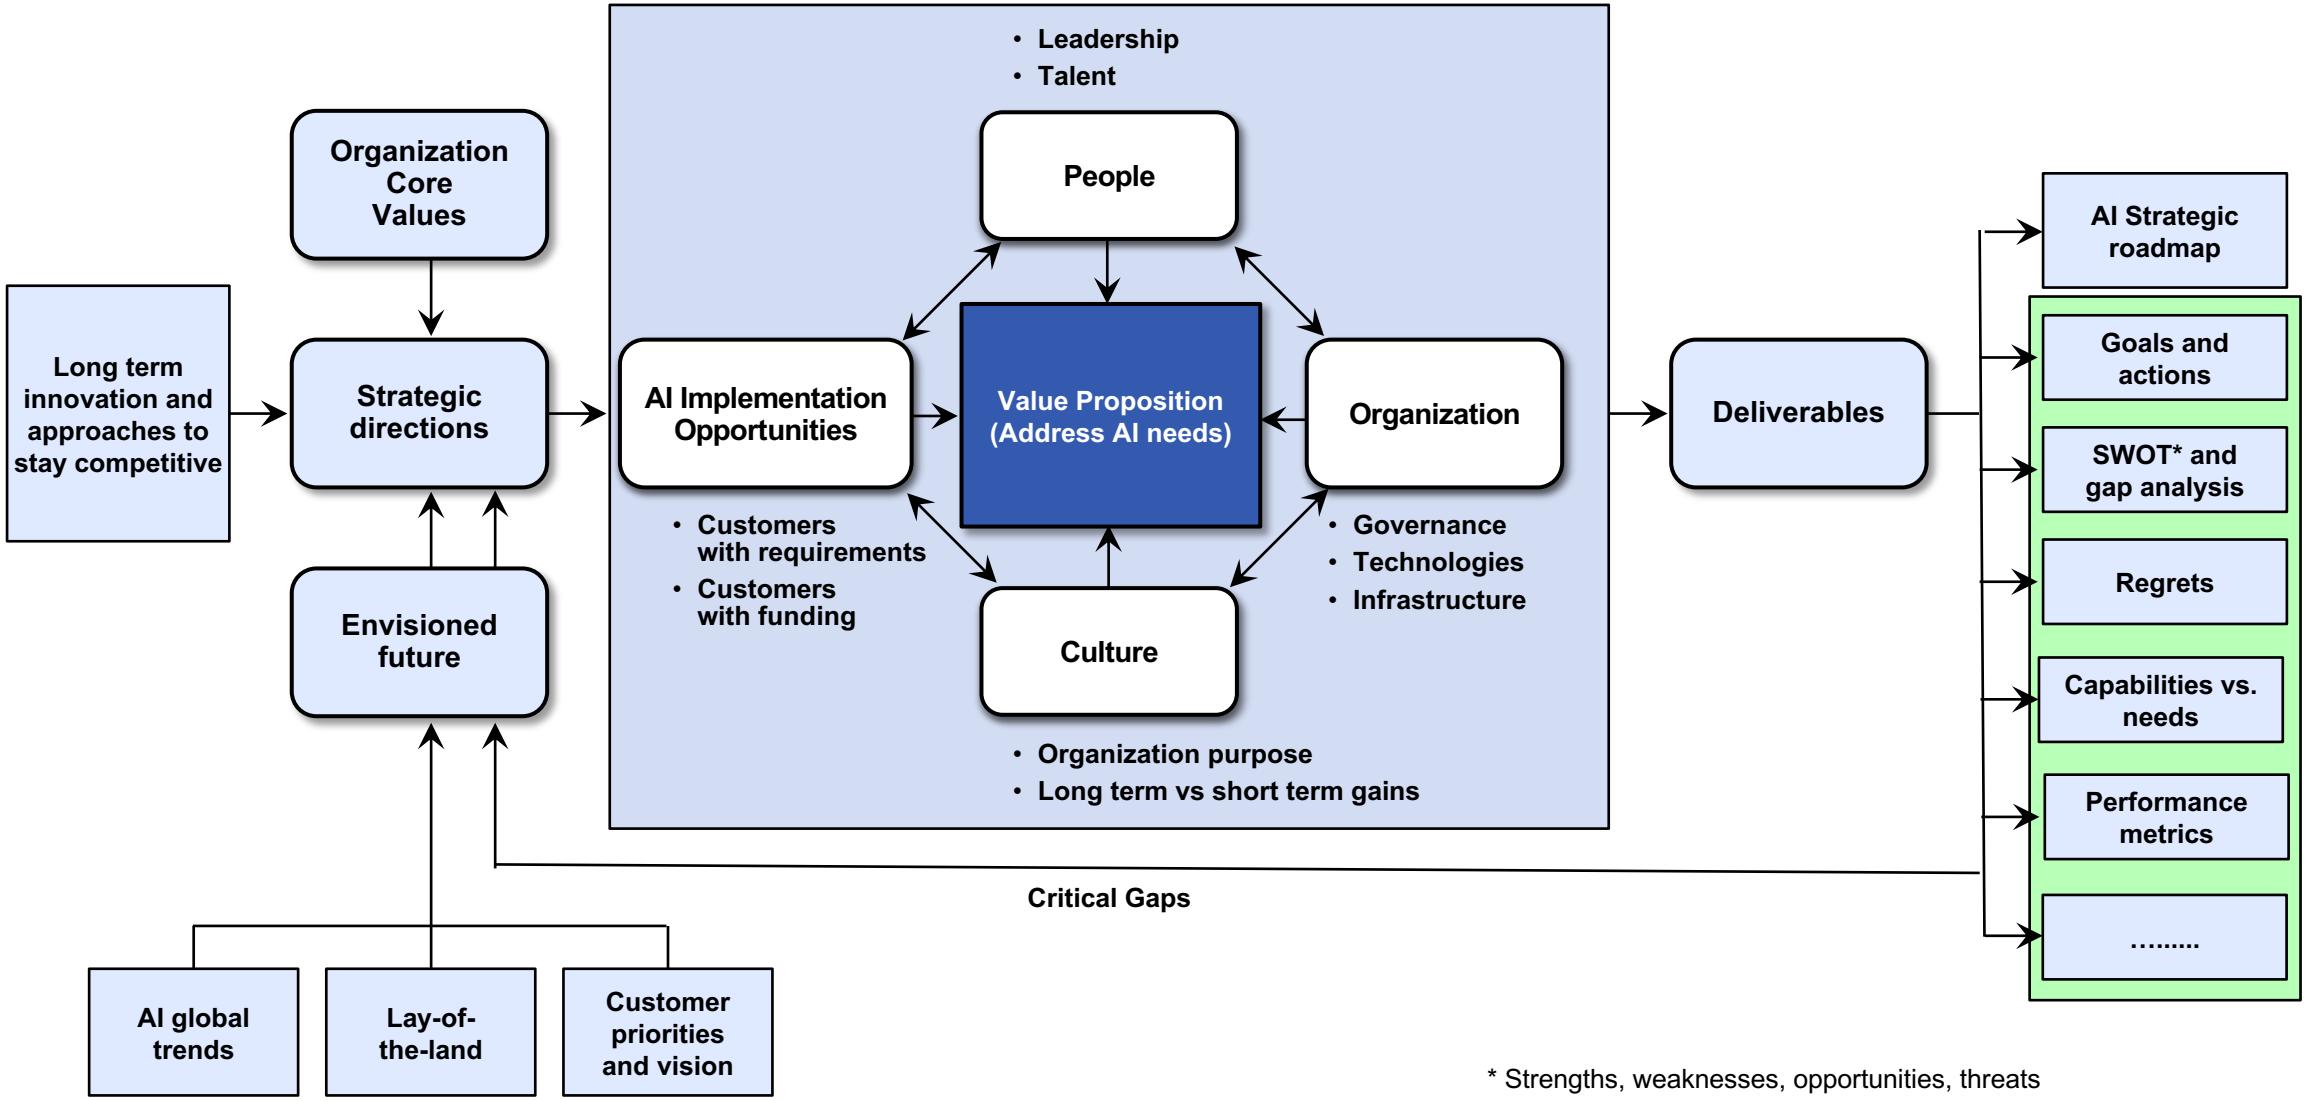

* Strengths, weaknesses, opportunities, threats

* Adapted from David Nadler and Mike Tushman, "A Model for Diagnosing Organizational Behavior," *Organizational Dynamics*, Autumn 1980

Top AI International Conferences and Other Venues*

- NeurIPS: Neural Information Processing Systems (formerly abbreviated NIPS). Held in early December https://nips.cc/
- ICML: International Conference on Machine Learning. Held in July https://icml.cc/
- ICLR: International Conference on Learning Representations. Held in May https://iclr.cc/
- AAAI: Association for the Advancement of Artificial Intelligence. Held in February http://www.aaai.org/
- CVPR: Computer Vision and Pattern Recognition. Held in June https://www.thecvf.com/
- ICCV: International Conference on Computer Vision. Held in the Fall of odd years https://www.thecvf.com/
- KDD: Knowledge Discovery and Data Mining https://www.kdd.org/kdd2019
- IJCAI: International Joint Conference on Artificial Intelligence http://ijcai19.org/
- MIT Artificial Intelligence Course https://ocw.mit.edu/courses/electrical-engineering-and-computer-science/6-034-artificial-intelligence-fall-2010/lecture-videos/
- Google Advancing AI for Everyone https://ai.google/
- Facebook's F8 and Apple's WWDC
- Microsoft https://academy.microsoft.com/en-us/professional-program/tracks/artificial-intelligence/
- Coursera Machine Learning https://www.coursera.org/learn/machine-learning
- NVIDIA's GPU Technology Conference https://www.nvidia.com/en-us/gtc/
- AI Conference (Presented by O'Reilly & INTEL AI) https://www.quora.com/What-are-the-top-AI-conferences
- TTI Vanguard (By Membership Only) https://www.ttivanguard.com/

- https://www.quora.com/What-are-the-top-AI-conferences
- http://www.guide2research.com/topconf/machine-learning

AI for Cyber Security (AICS) Workshop
- 4th year in a Row -

Artificial Intelligence for Cyber Security Workshop

- Forum for AI researchers and practitioners to share research and experiences in applying AI to Cyber Security

Chairs

Bill Streilein

Dave Martinez

Jason Matterer

Honolulu, Hawaii • January 27, 2019

Theme: Adversarial Learning

Craig Knoblock
USC / ISI

Una-May O'Reilly
MIT CSAIL

Keynote Speakers

Challenge Problem: Malware classification robust to evasion

AICS Highlights

- This year's theme was adversarial AI
- About 30 papers were submitted and 7 papers were accepted
- Keynote speakers were from USC and MIT CSAIL
- Held a panel discussion with participation from DHS, NSA, Samsung Research, and keynote speakers on "Lack of Datasets"
- New this year was a challenge problem on classifying malware
- Winning Paper was titled: "Enhancing Robustness of Deep Neural Networks Against Adversarial Malware Samples," Univ. of Texas San Antonio and FIU

New GEL/EECS Graduate Course

6.S976J Engineering Leadership in the Age of Artificial Intelligence

TIME: M/W 12:30–2:00pm

UNITS: G3-0-9

Room: 4-231

PREREQUISITES: None

INSTRUCTORS:
David Martinez & David Niño

TAs:
Christos D. Samolis & Bruke Kifle

gelp.mit.edu/gel-grad-ai

* *Can be used to fulfill part of engineering doctoral minor and GEL's new Leadership Certificate Program launching in 2019-2020.*

Are You Prepared to Lead AI Teams?

This course will prepare MIT graduate engineering students to lead, develop, and deploy AI systems in ways that deliver positive benefits to people and society.

At the completion of this course, students will be able to lead AI teams based on five main acquired skills:

- Understanding an end-to end AI architecture at the system engineering level
- Applying engineering leadership principles
- Developing a strategic vision and development plan for an AI project
- Identifying an execution strategy including a diverse and multidisciplinary team
- Demonstrating an AI conceptual design using a Raspberry Pi

BERNARD M. GORDON
MIT Engineering Leadership
PROGRAM

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Summary

- U.S. needs to regain AI leadership by strategically partnering with small to large companies from the industrial base plus academia
  - "…A nation which depends upon others for its new basic scientific knowledge will be slow in its industrial progress and weak in its competition position in world trade" – Excerpt from Technology and National Security by Walter Isaacson, January 3, 2019
- AI will be a technological enabler (i.e., data and algorithm warfare) against: radical extremists, terrorists, and peer nations to defend our homeland and abroad
- Develop an ecosystem to allow training of the next generation of DoD/IC military and civilian workforce in AI
  - Agile prototyping of AI capabilities with operational users
  - Carry out grand challenges
  - Facilitate red/blue tabletop exercises

Contact Information

David R. Martinez
Associate Division Head
dmartinez@ll.mit.edu
Office: 781-981-7505
Cell: 781-879-0890

AI Bibliography List (few selected set)