Future of Cyber Security Enabled by AI

Dr. William Streilein

Cyber Analytics and Decision Systems

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Outline

- Background
  - Definition of Area
  - AI History Highlights
- Lay-of-the-Land
- AI for Cyber Security
- Summary

National Challenges and Role of AI

Intelligent Systems and Autonomy

Information Superiority

	National Challenges	*Role of AI in Augmenting Humans*
	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

<u>Narrow AI:</u>

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

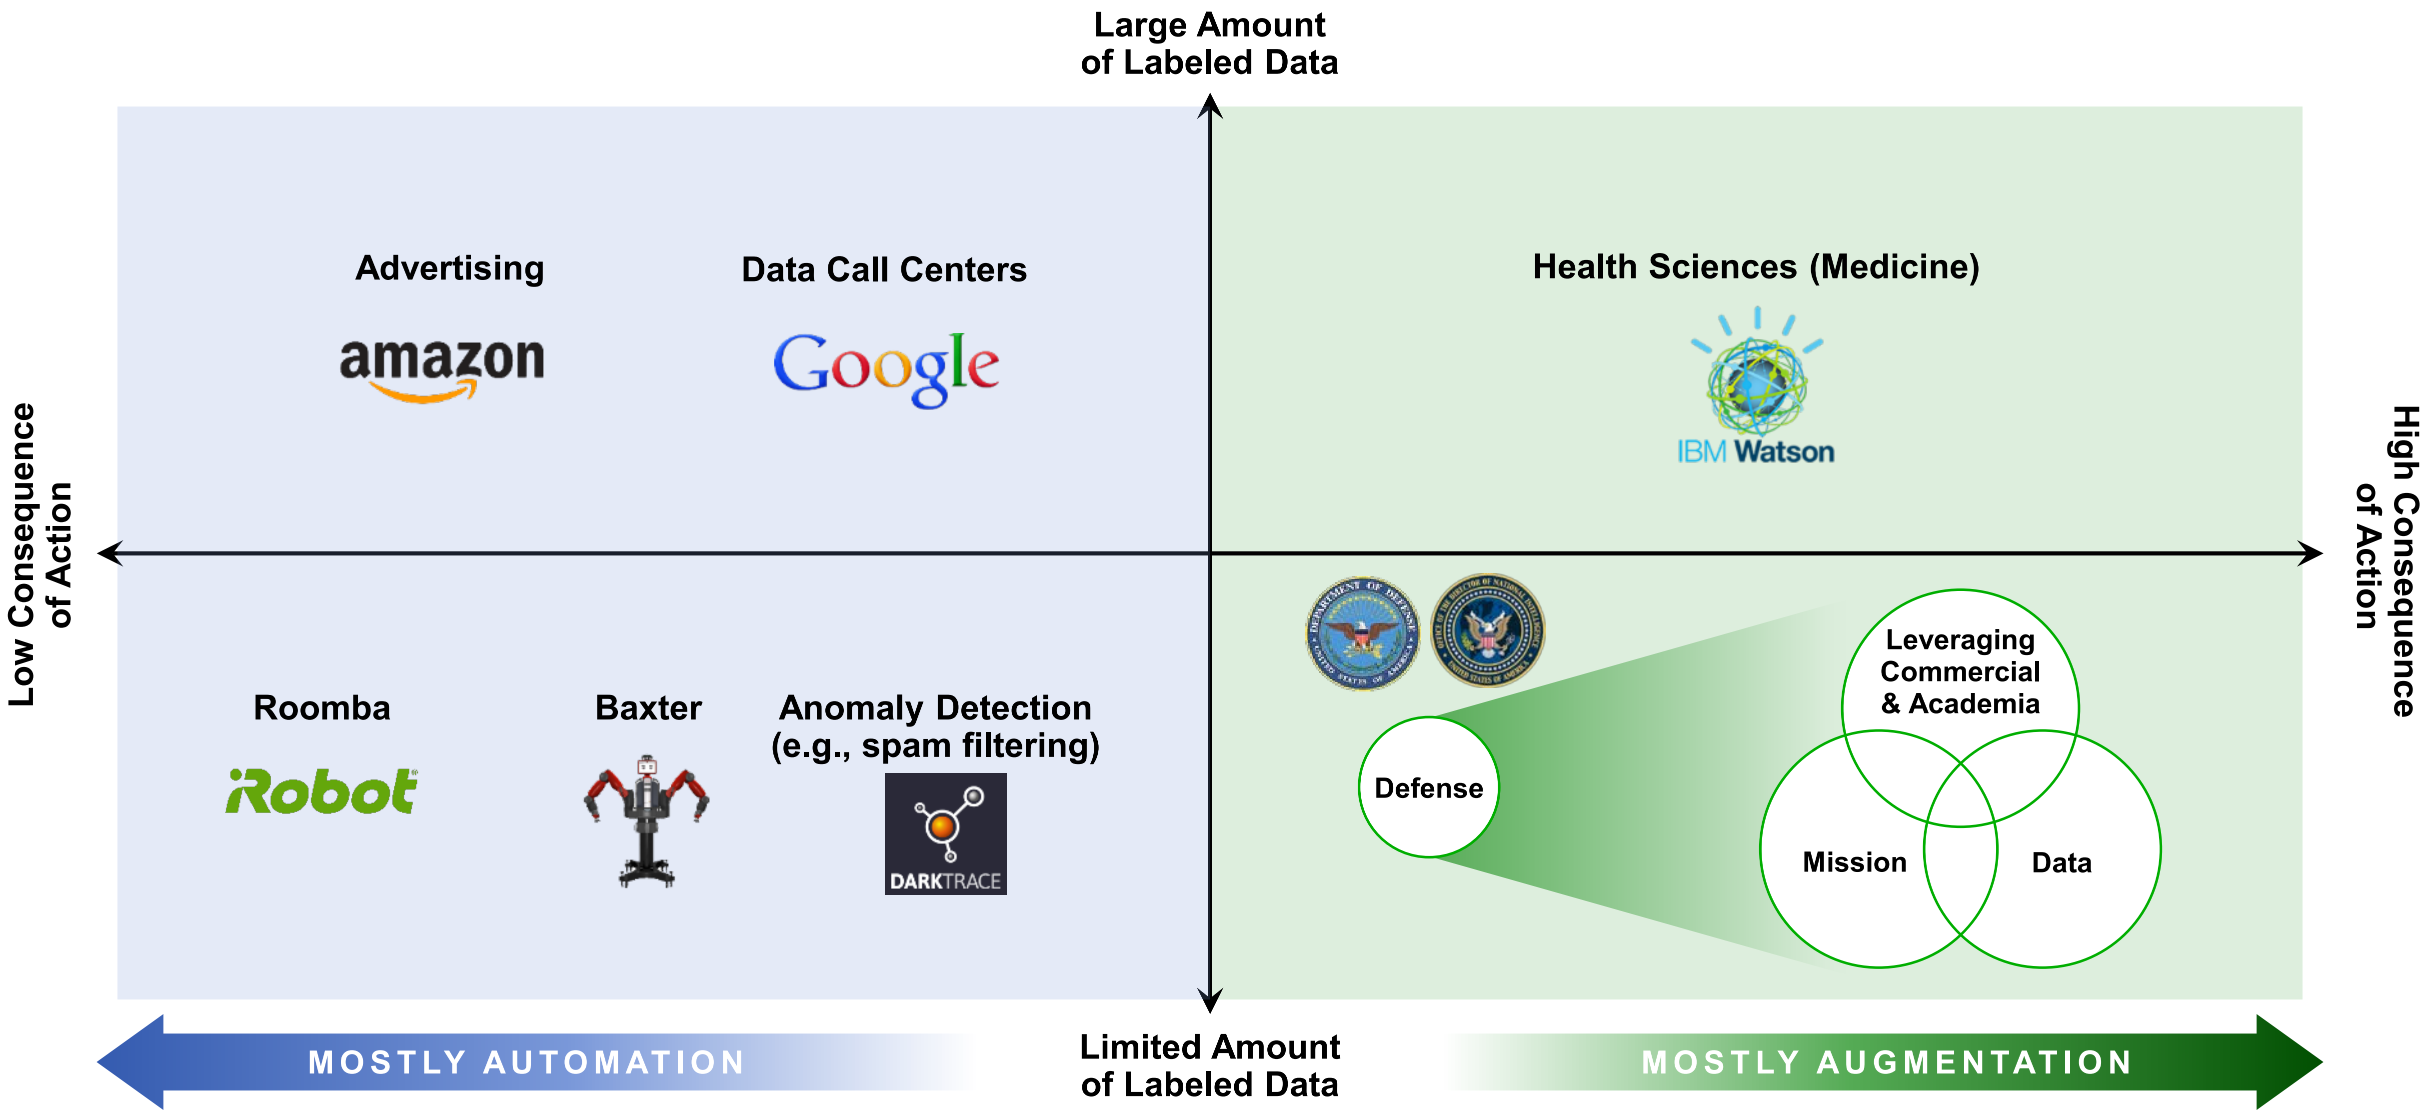

Select History of Artificial Intelligence

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

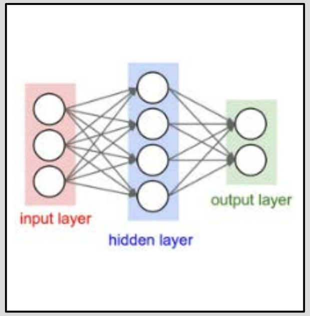

1982 - Expert Systems Pioneer DENDRAL project at Stanford

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

2005 - Google's Arabic and Chinese to English translation

2007 - DARPA Grand Challenge *("Urban Challenge")*

2011 - IBM Watson defeats former ***Jeopardy!*** champions *(Brad Rutter and Ken Jennings)*

2012 - Team from U. of Toronto (Geoff Hinton's lab) wins the ImageNet Large Scale Visual Recognition Challenge with deep-learning software

2014 - Google's GoogleNet Object classification at near human performance

2015 - DeepMind achieves human expert level of play on Atari games *(using only raw pixels and scores)*

2016 - DeepMind AlphaGo defeats top human Go player *(Lee Sedol)*

2016 - DARPA Cyber Grand Challenge

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

AI “Winters”

1974–1980	The first “AI winter”
1970s	Knowledge-based approaches
1980–88	Expert systems boom
1988–93	Expert systems bust; the second “AI winter”
1986	Neural networks return to popularity
1988	Pearl’s Probabilistic Reasoning in Intelligent Systems
1990	Backlash against symbolic systems; Brooks’ “nouvelle AI”
1995–present	Increasing specialization of the field
	Agent-based systems
	Machine learning everywhere
	Tackling general intelligence again?

Source: UNC Computer Science

The first AI winter 1974–1980

In the 70s, AI was subject to critiques and financial setbacks. AI researchers had failed to appreciate the difficulty of the problems they faced. Their tremendous optimism had raised expectations impossibly high, and when the promised results failed to materialize, funding for AI disappeared. At the same time, the field of connectionism (or neural nets) was shut down almost completely for 10 years by Marvin Minsky's devastating criticism of perceptron. Despite the difficulties with public perception of AI in the late 70s, new ideas were explored in logic programming, commonsense reasoning and many other areas.

Bust: the second AI winter 1987–1993

The business community's fascination with AI rose and fell in the 80s in the classic pattern of an economic bubble. The collapse was in the perception of AI by government agencies and investors – the field continued to make advances despite the criticism. Rodney Brooks and Hans Moravec, researchers from the related field of robotics, argued for an entirely new approach to artificial intelligence.

Source: Wikipedia, History of artificial intelligence

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

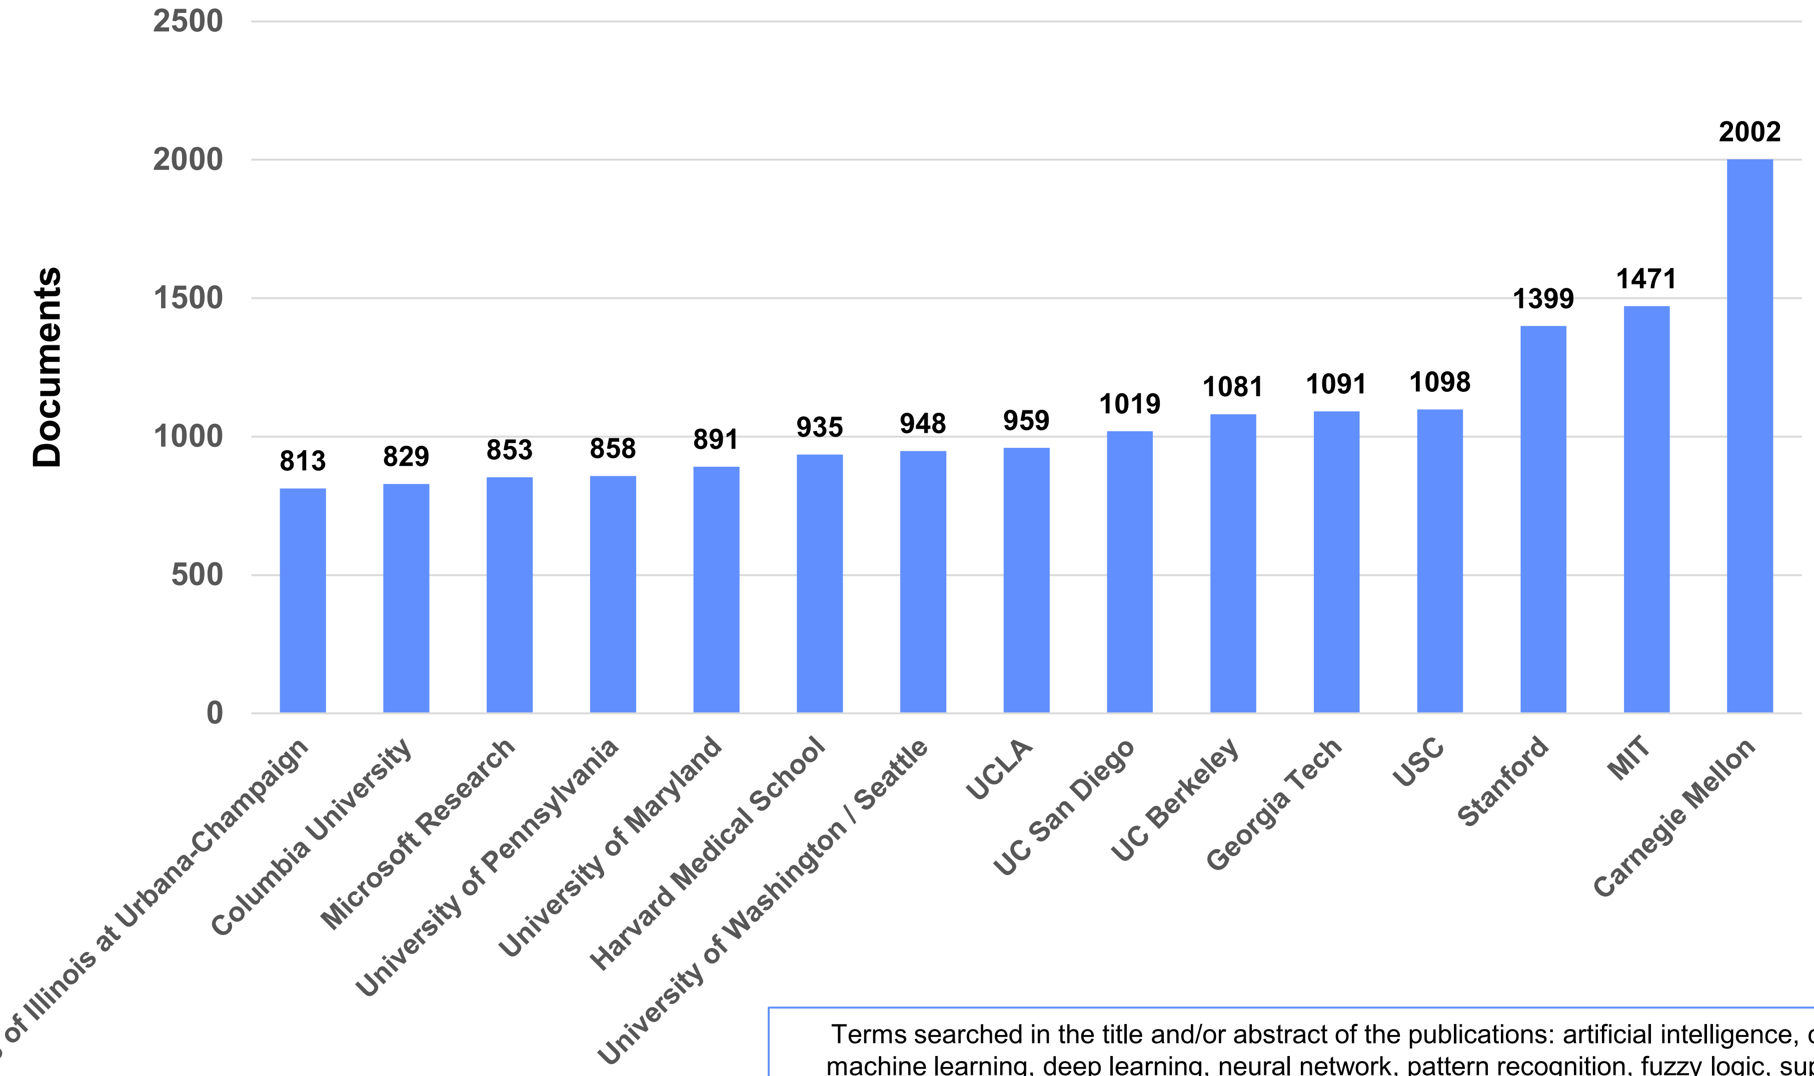

Terms searched in the title and/or abstract of the publications: artificial intelligence, cognitive computing, machine learning, deep learning, neural network, pattern recognition, fuzzy logic, support vector machine

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

Top 14 Patent Holders in AI Per Country (2011–2016)

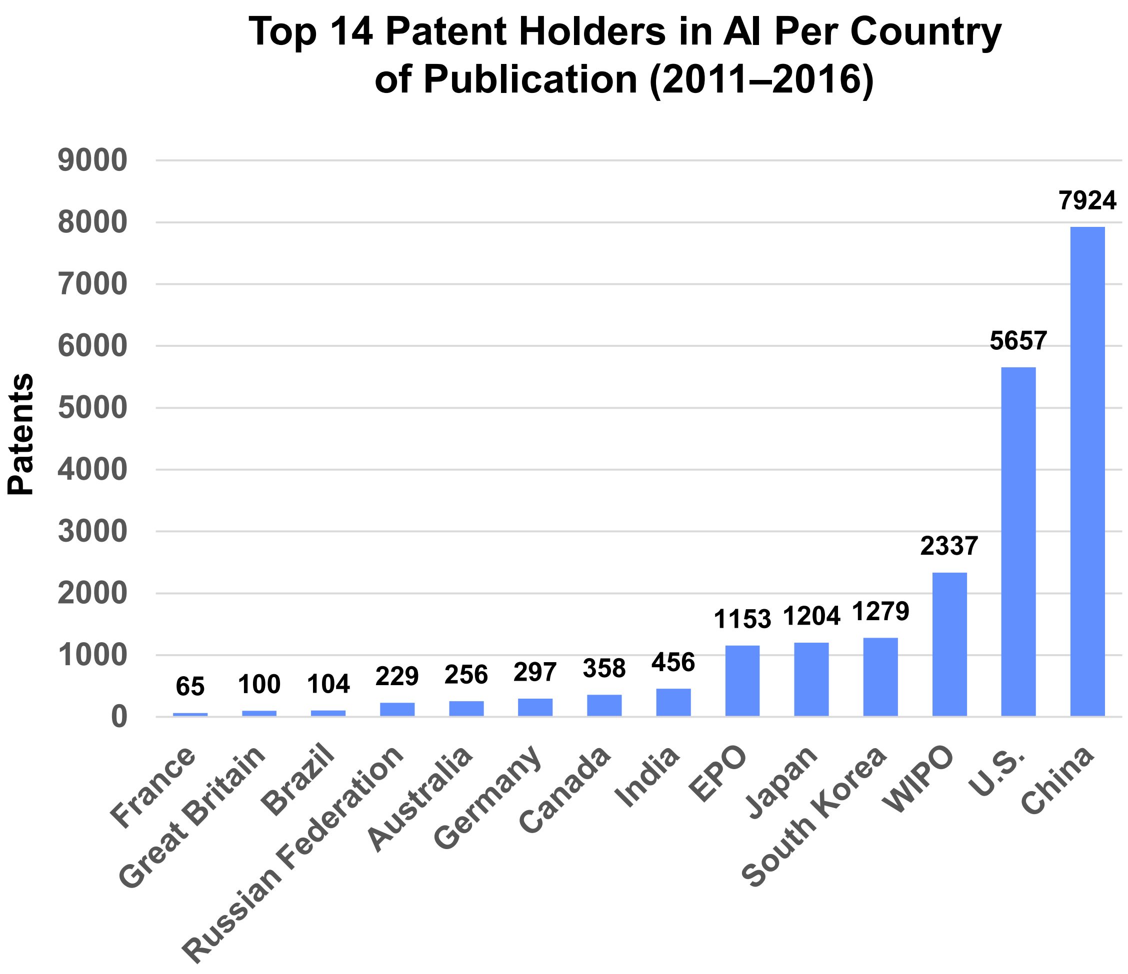

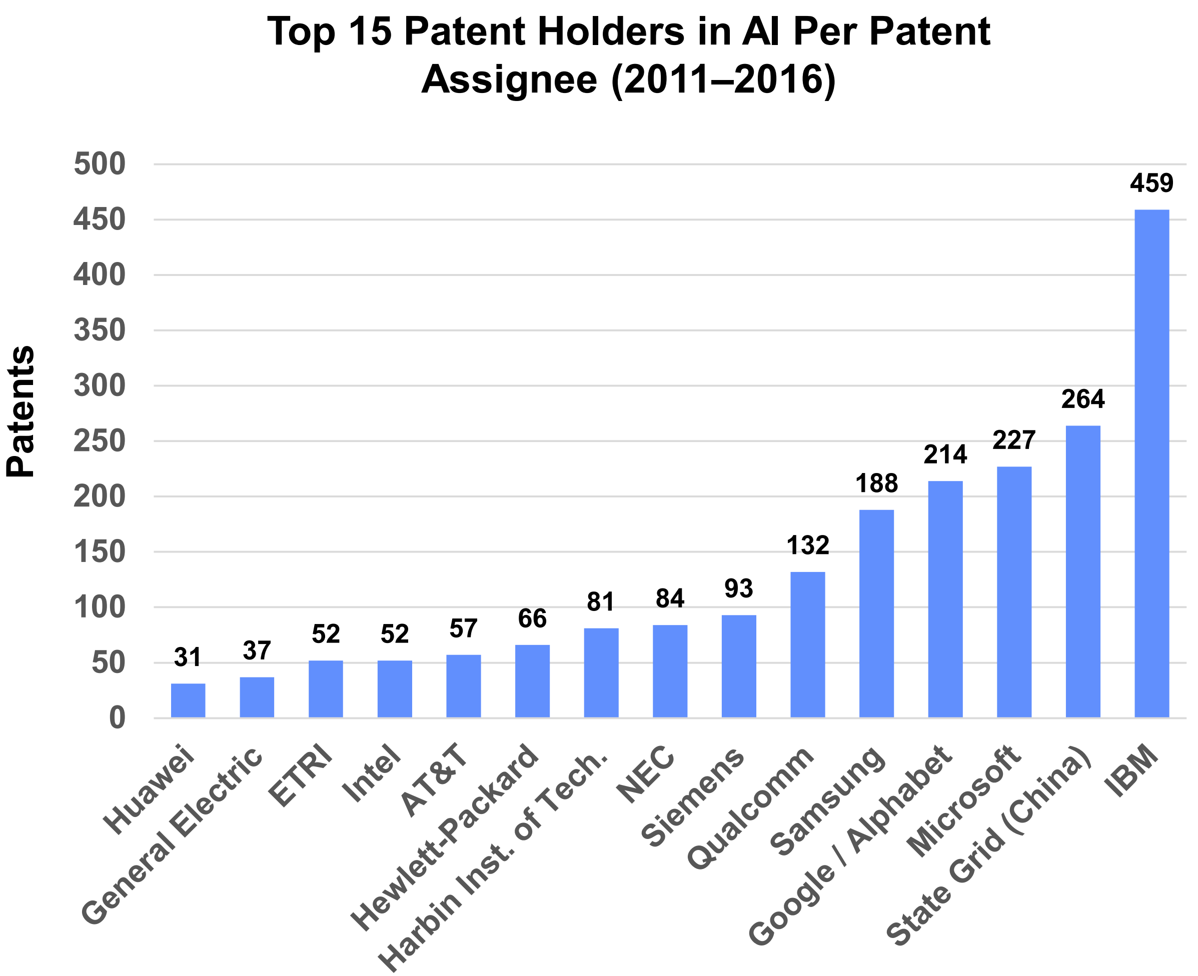

Terms searched in the title and/or abstract of the patent record: artificial intelligence, cognitive computing, machine learning, deep learning, neural network, pattern recognition, fuzzy logic, support vector machine

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

EPO = European Patent Organization
ETRI = South Korea Elec. Tech. Research Institute

China is Putting a Major Investment into AI

Chinese Government has indicated plans to invest $150B over 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

Outline

- **Background**
- **Lay-of-the-Land**
  - **AI Canonical Architecture**
  - **Summary of Study Outreach**

- **AI for Cyber Security**

- **Summary**

AI Canonical Architecture

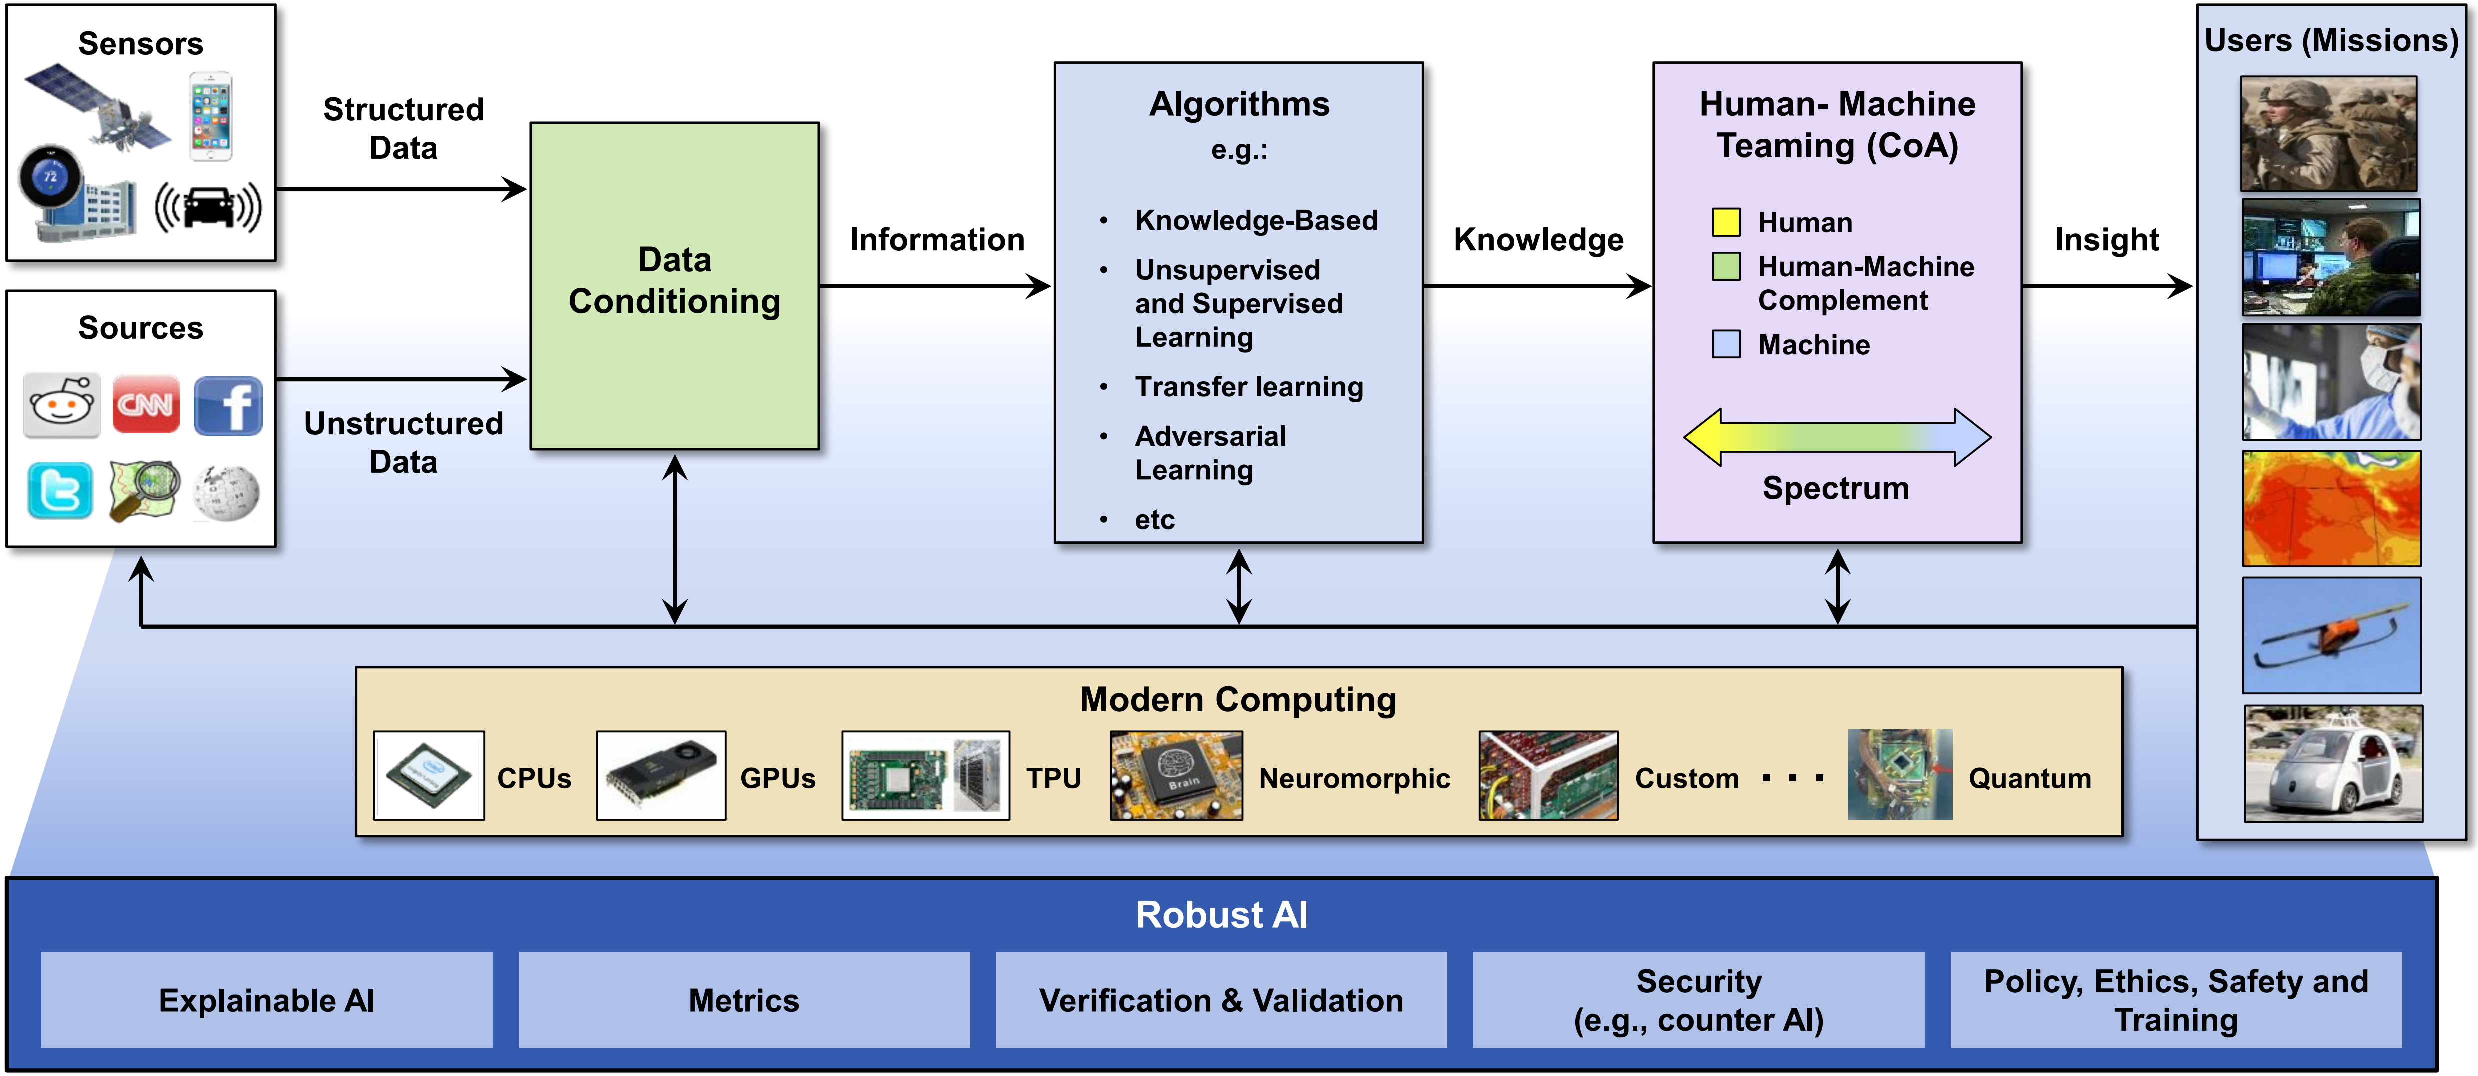

GPU = Graph Processing Unit
TPU = Tensor Processing Unit
CoA = Courses of Action

Four Components of Machine Learning Solutions

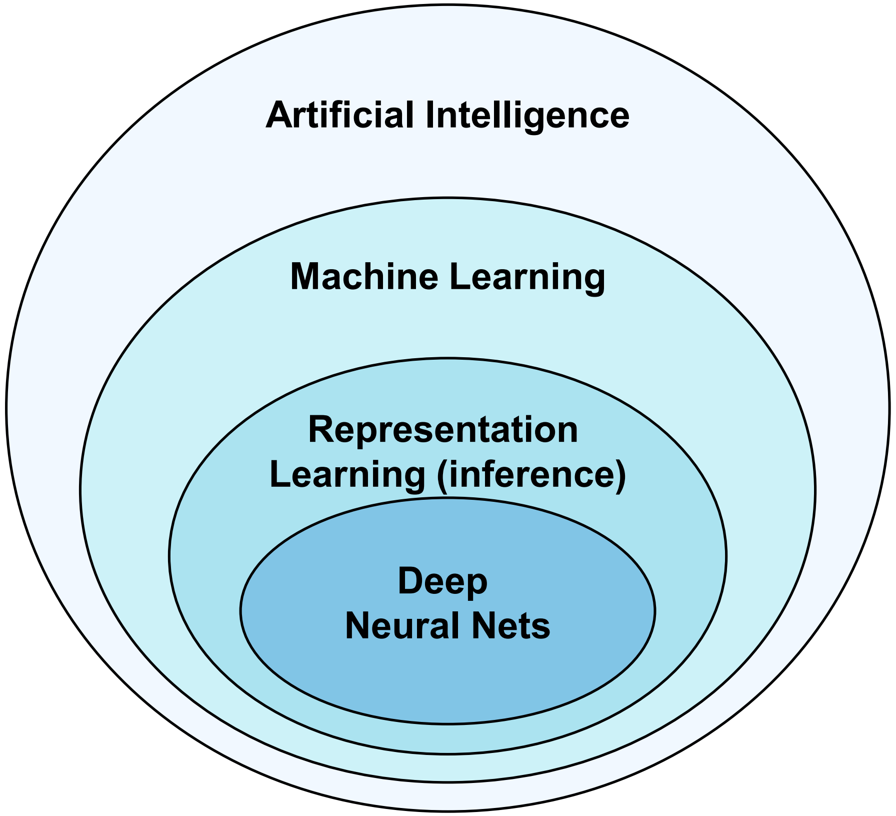

Image Adapted From: “Deep Learning”
I. Goodfellow, et.al., 2016 MIT Press

1. Define a Problem

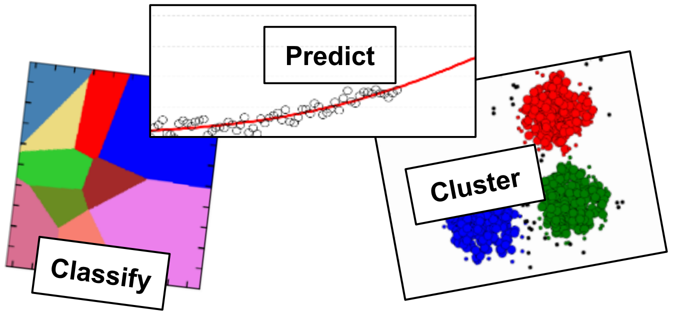

2. Gather Data

3. Create Full Train/Test Solution Pipelines

4. Provide Computation that Makes the Solution Feasible

National Security Specific Concerns Compared To Commercial Sector

Commercial Sector	National Security
High dimensionality	High dimensionality
Large volume	Large volume
Known truth / continual development	Unknown truth / not tolerant to errors
Mild consequences of decisions	Large consequences of decisions
Past is representative of future	Past does not always represent future
Competitive environment	Adversarial environment
Mostly consumer users	Today requires sophisticated users
Explainability is not the largest issue	Trust / explainability is core
Quantifiable success ($$)	Harder to measure success

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

Spectrum of Commercial Organizations in the Machine Intelligence Field

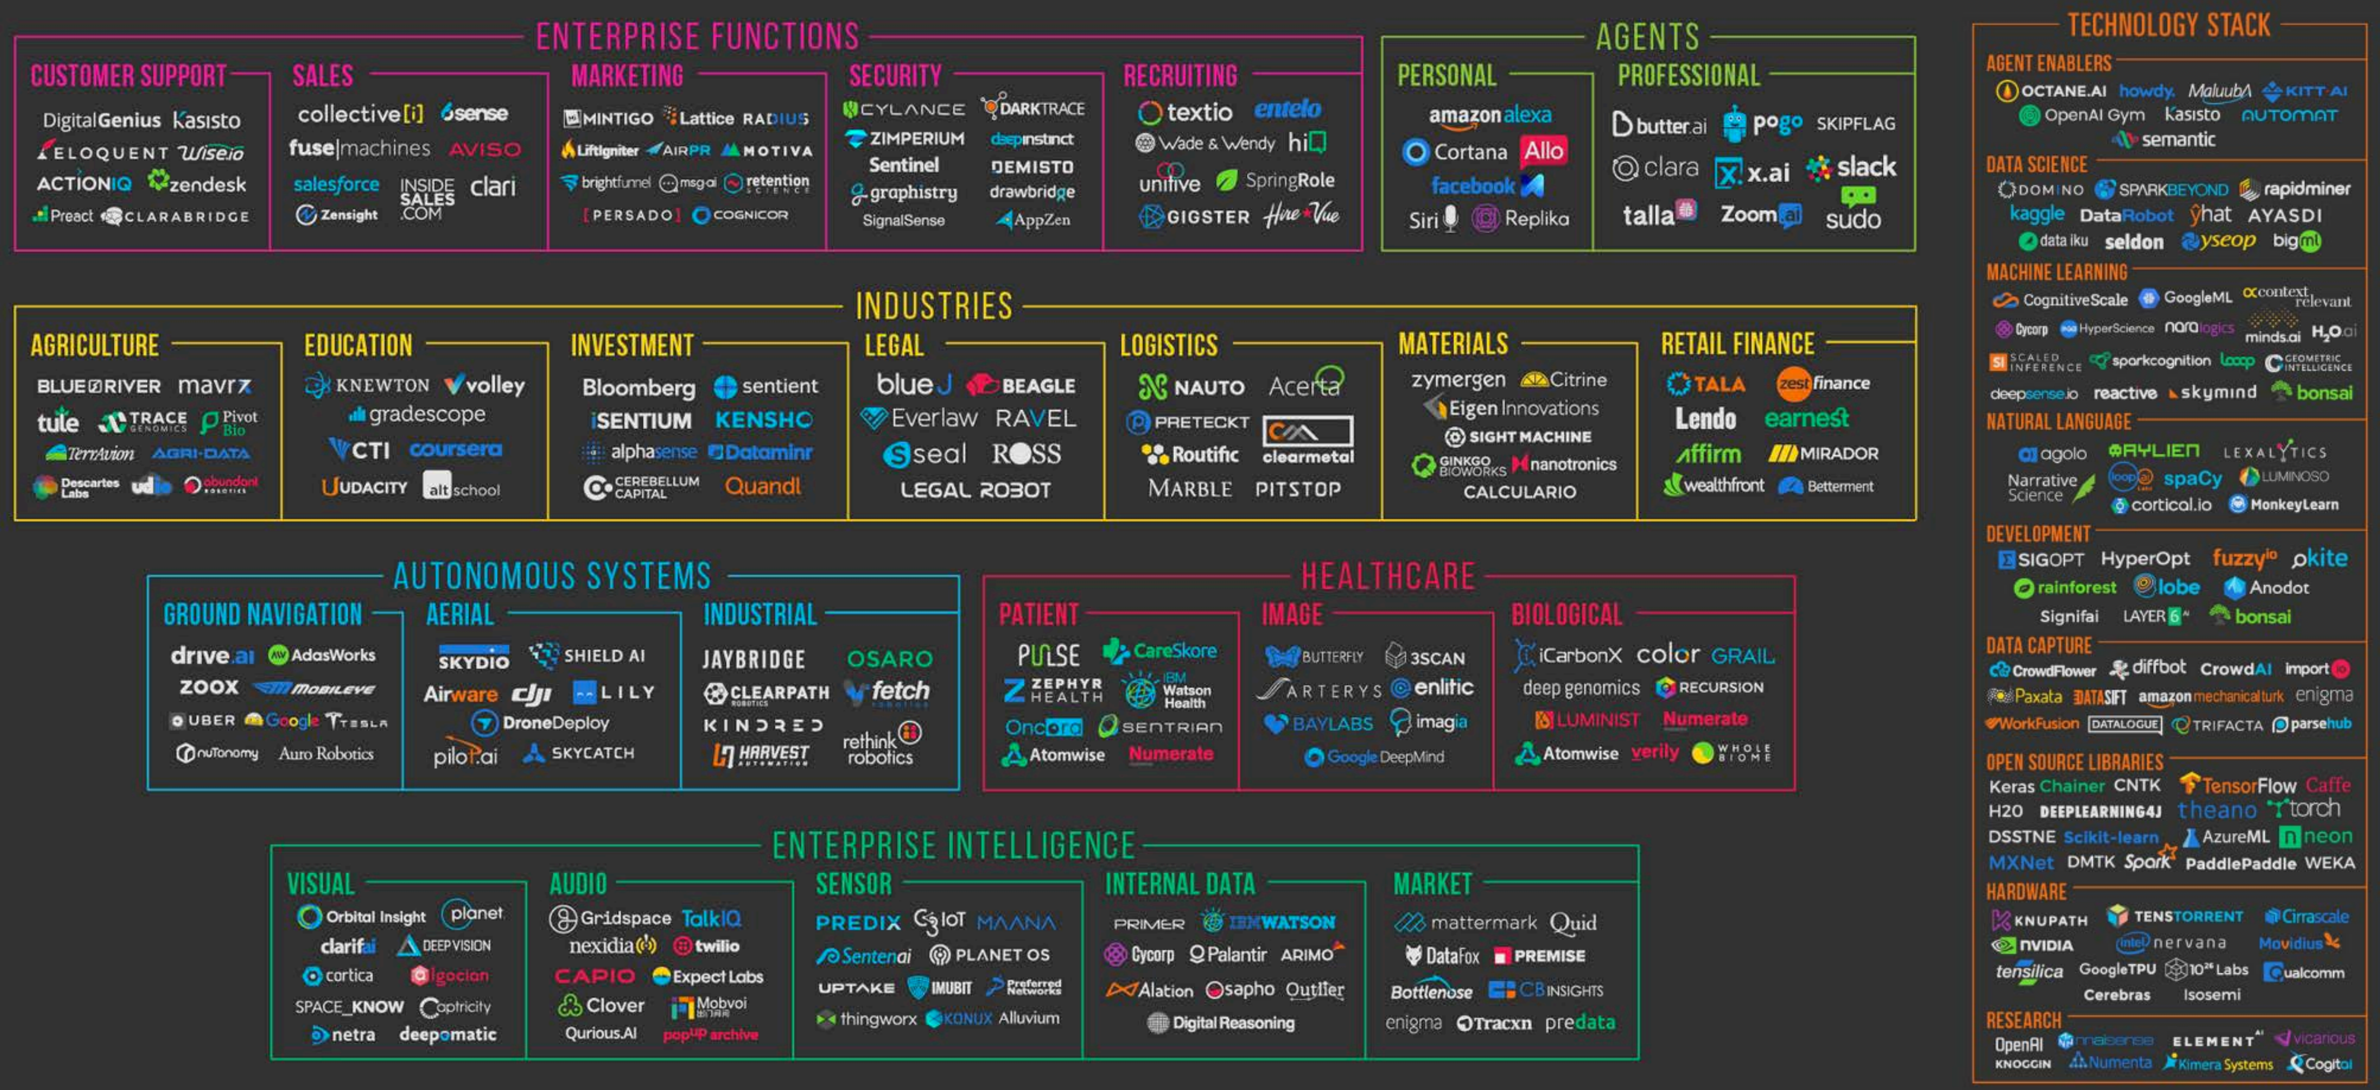

Source: Shivon Zilis, 2016, https://www.oreilly.com/ideas/the-current-state-of-machine-intelligence-3-0

Outline

- **Background**
- **Lay-of-the-Land**
  - **AI Canonical Architecture**
  - **Summary of Study Outreach**
- **AI for Cyber Security**
- **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

Qntfy

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

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

Outline

- Background
- Lay-of-the-Land
- AI for Cyber Security
  - Background
  - Findings and Recommendations
- Summary

Cyber Security: Critical Threat Surfaces

Global Trends

- Vulnerable users
- Insider threats (both malicious and unwitting insiders)
- Proliferation of devices, e.g. IoT
- Mission success in the presence of cyber attacks
- Attacker use of automation / AI

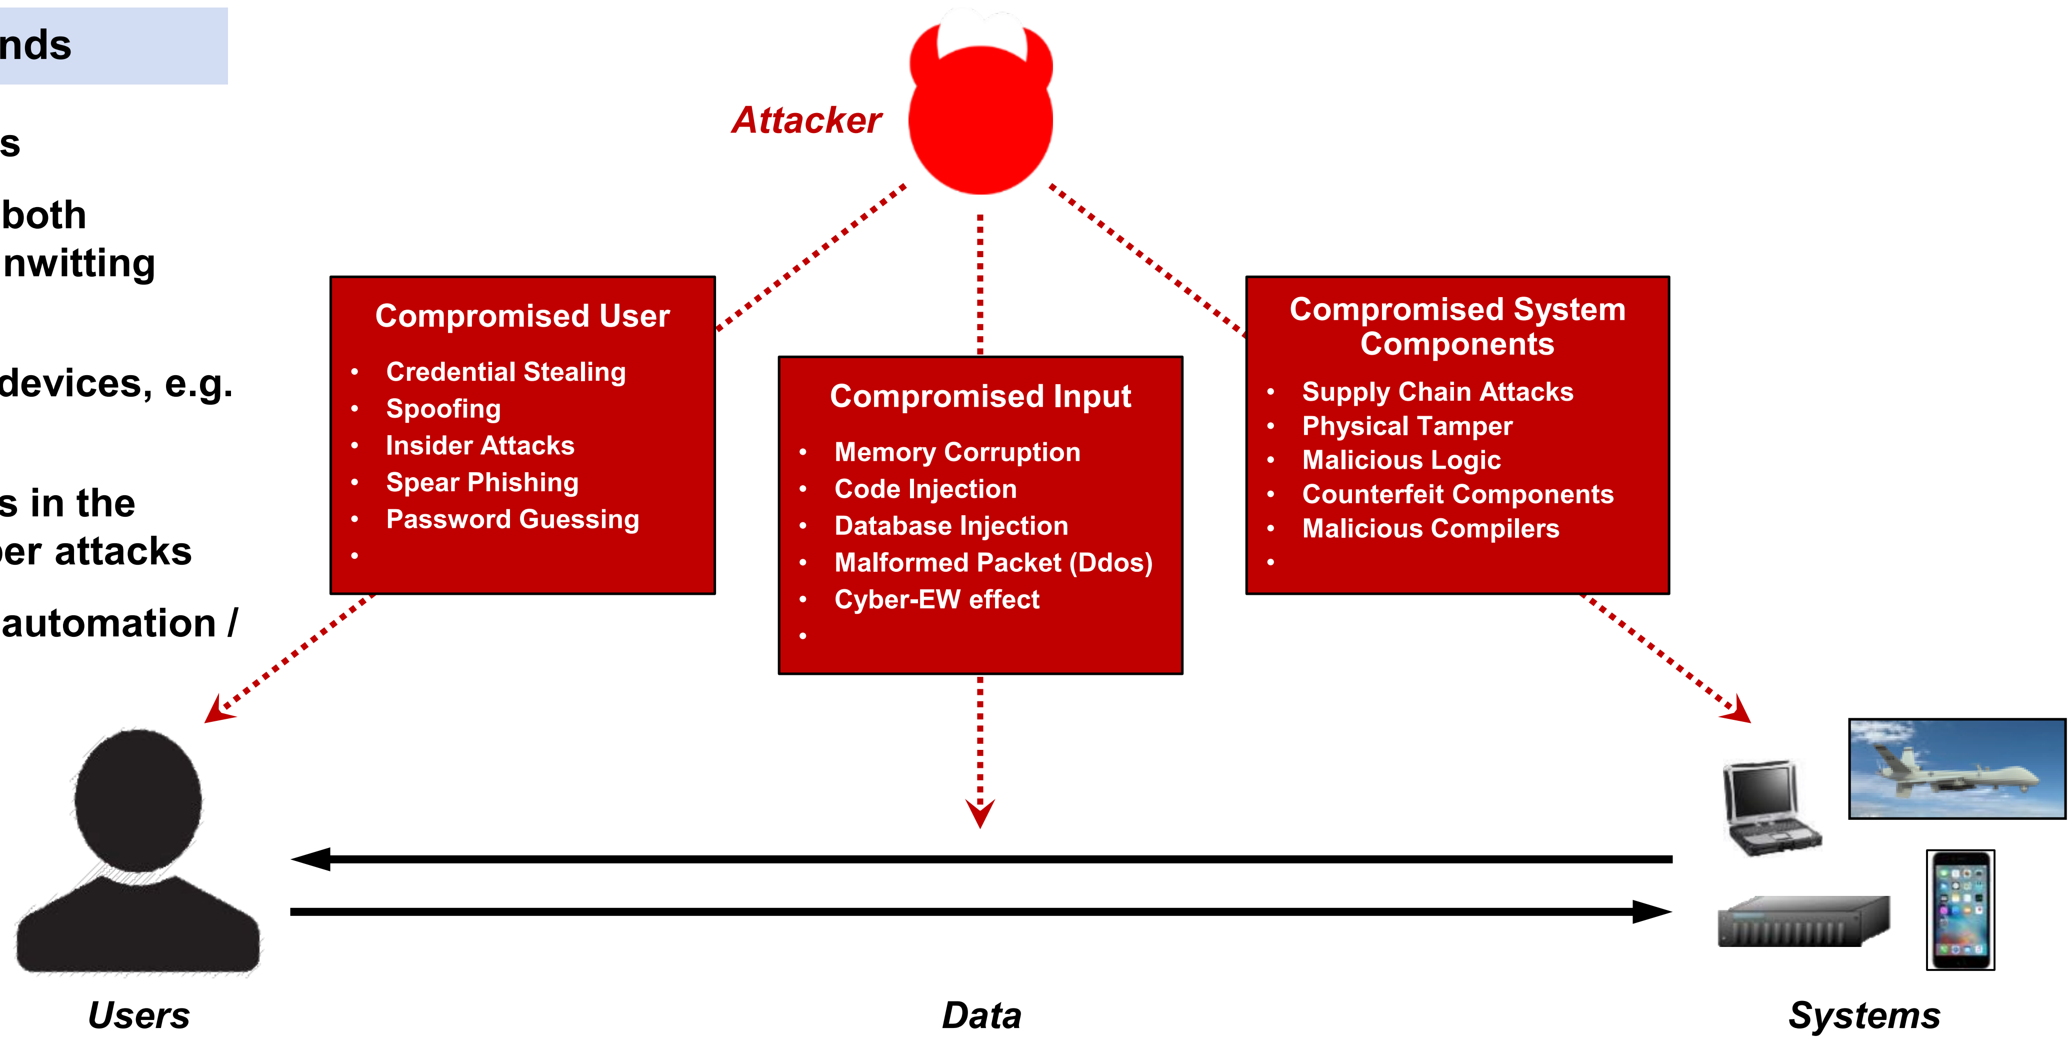

Global Trend: Sophisticated Attacks More Easily Accomplished with Automation

Growth of Threat

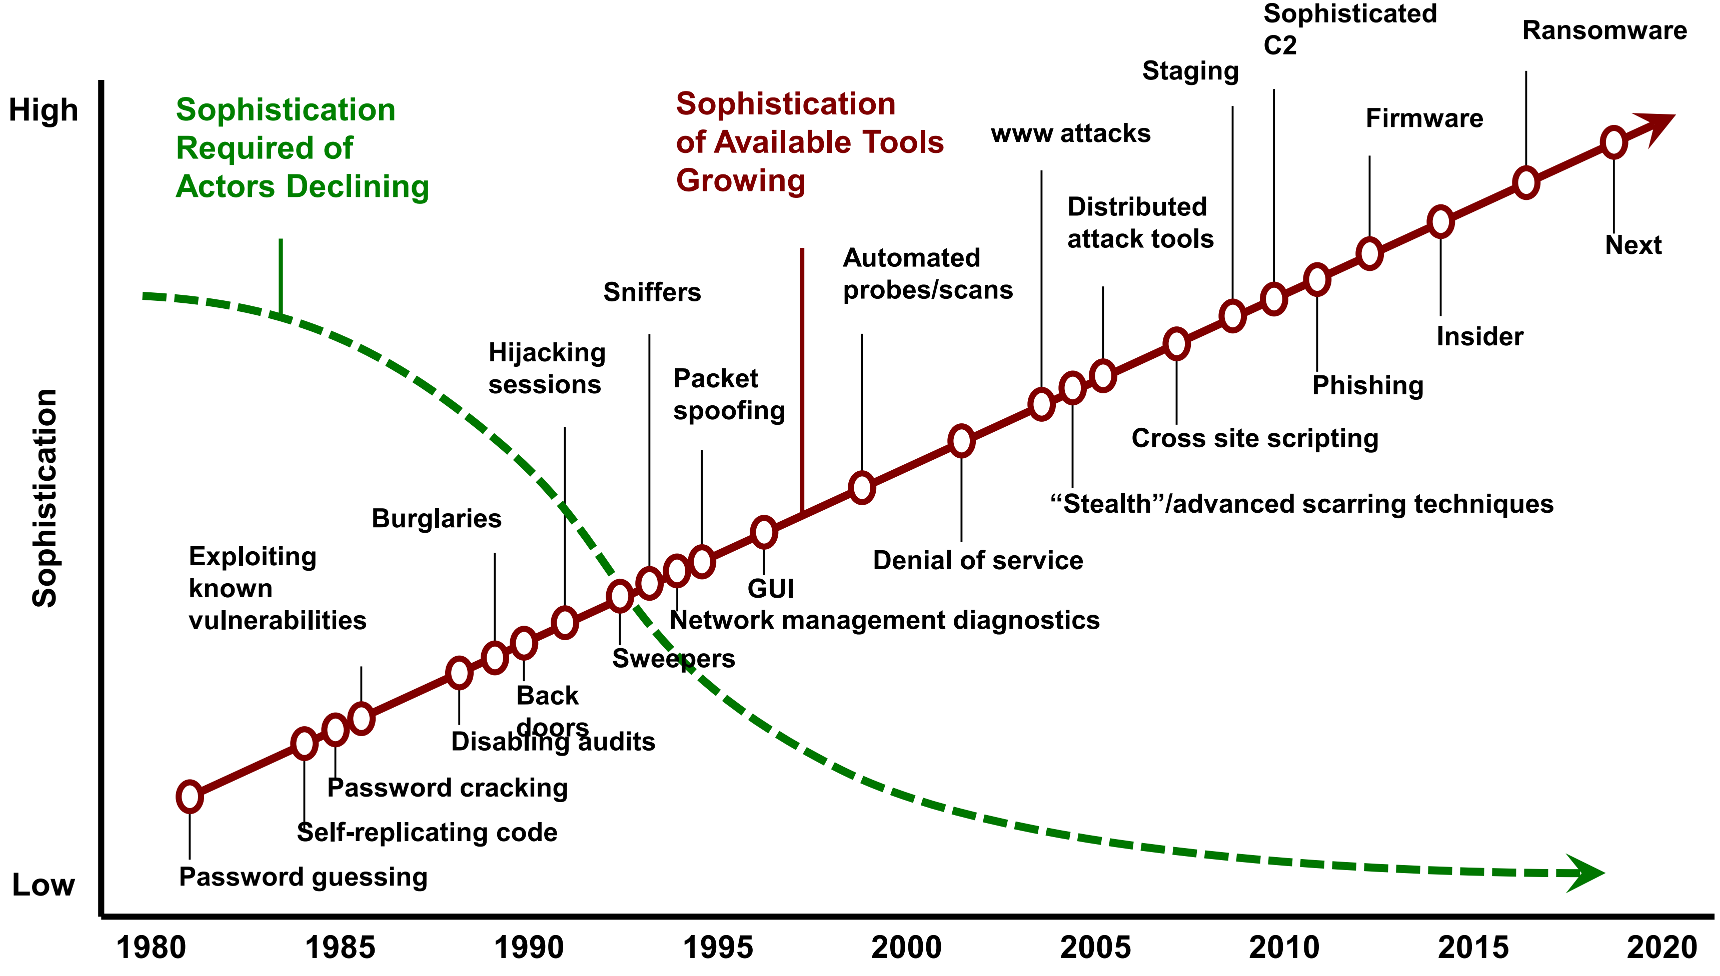

NOTEWORTHY FACTS

- 250K new malware programs are registered each day
- There were 357M new email malware variants in 2016 - 36% more new variants than in 2014.
- There were 463M new variants of ransomware in 2016 - 36% more new variants than in 2015.
- 99 days to detect compromise - adversary gains access in 3
- Internet of Things and Cloud are hot targets (e.g. Mirai botnet) – 2 min to compromise
- Projected cyber attack costs in 2019: $2.1T

Sources: https://www.symantec.com/security-center/threat-report
http://expandedramblings.com/index.php/cybersecurity-statistics/
http://www.nato.int/docu/Review/2016/Also-in-2016/cyber-defense-nato-security-
https://www.fireeye.com/blog/threat-research/2017/03/m-trends-2017.htm
https://www.ag-test.org/en/statistics/malware

The Cyber Battleground

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

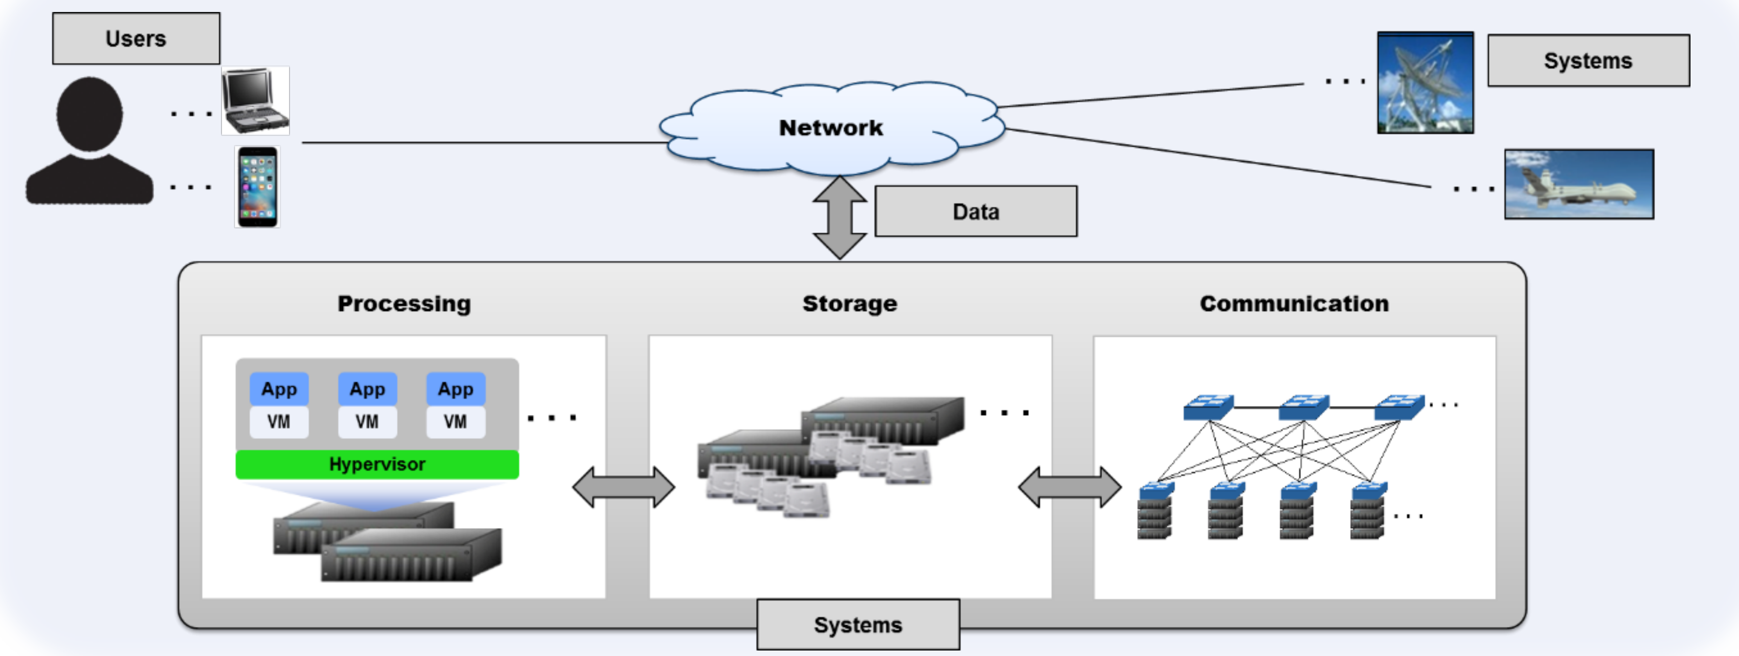

Defense Stages	Identify (Recon)	Protect	Detect	Respond	Recover
Impact	Focused defense	Deflect attacks	ID new attacks	Stop attacks	"Mission" fight through

Major Challenges to Cyber Security

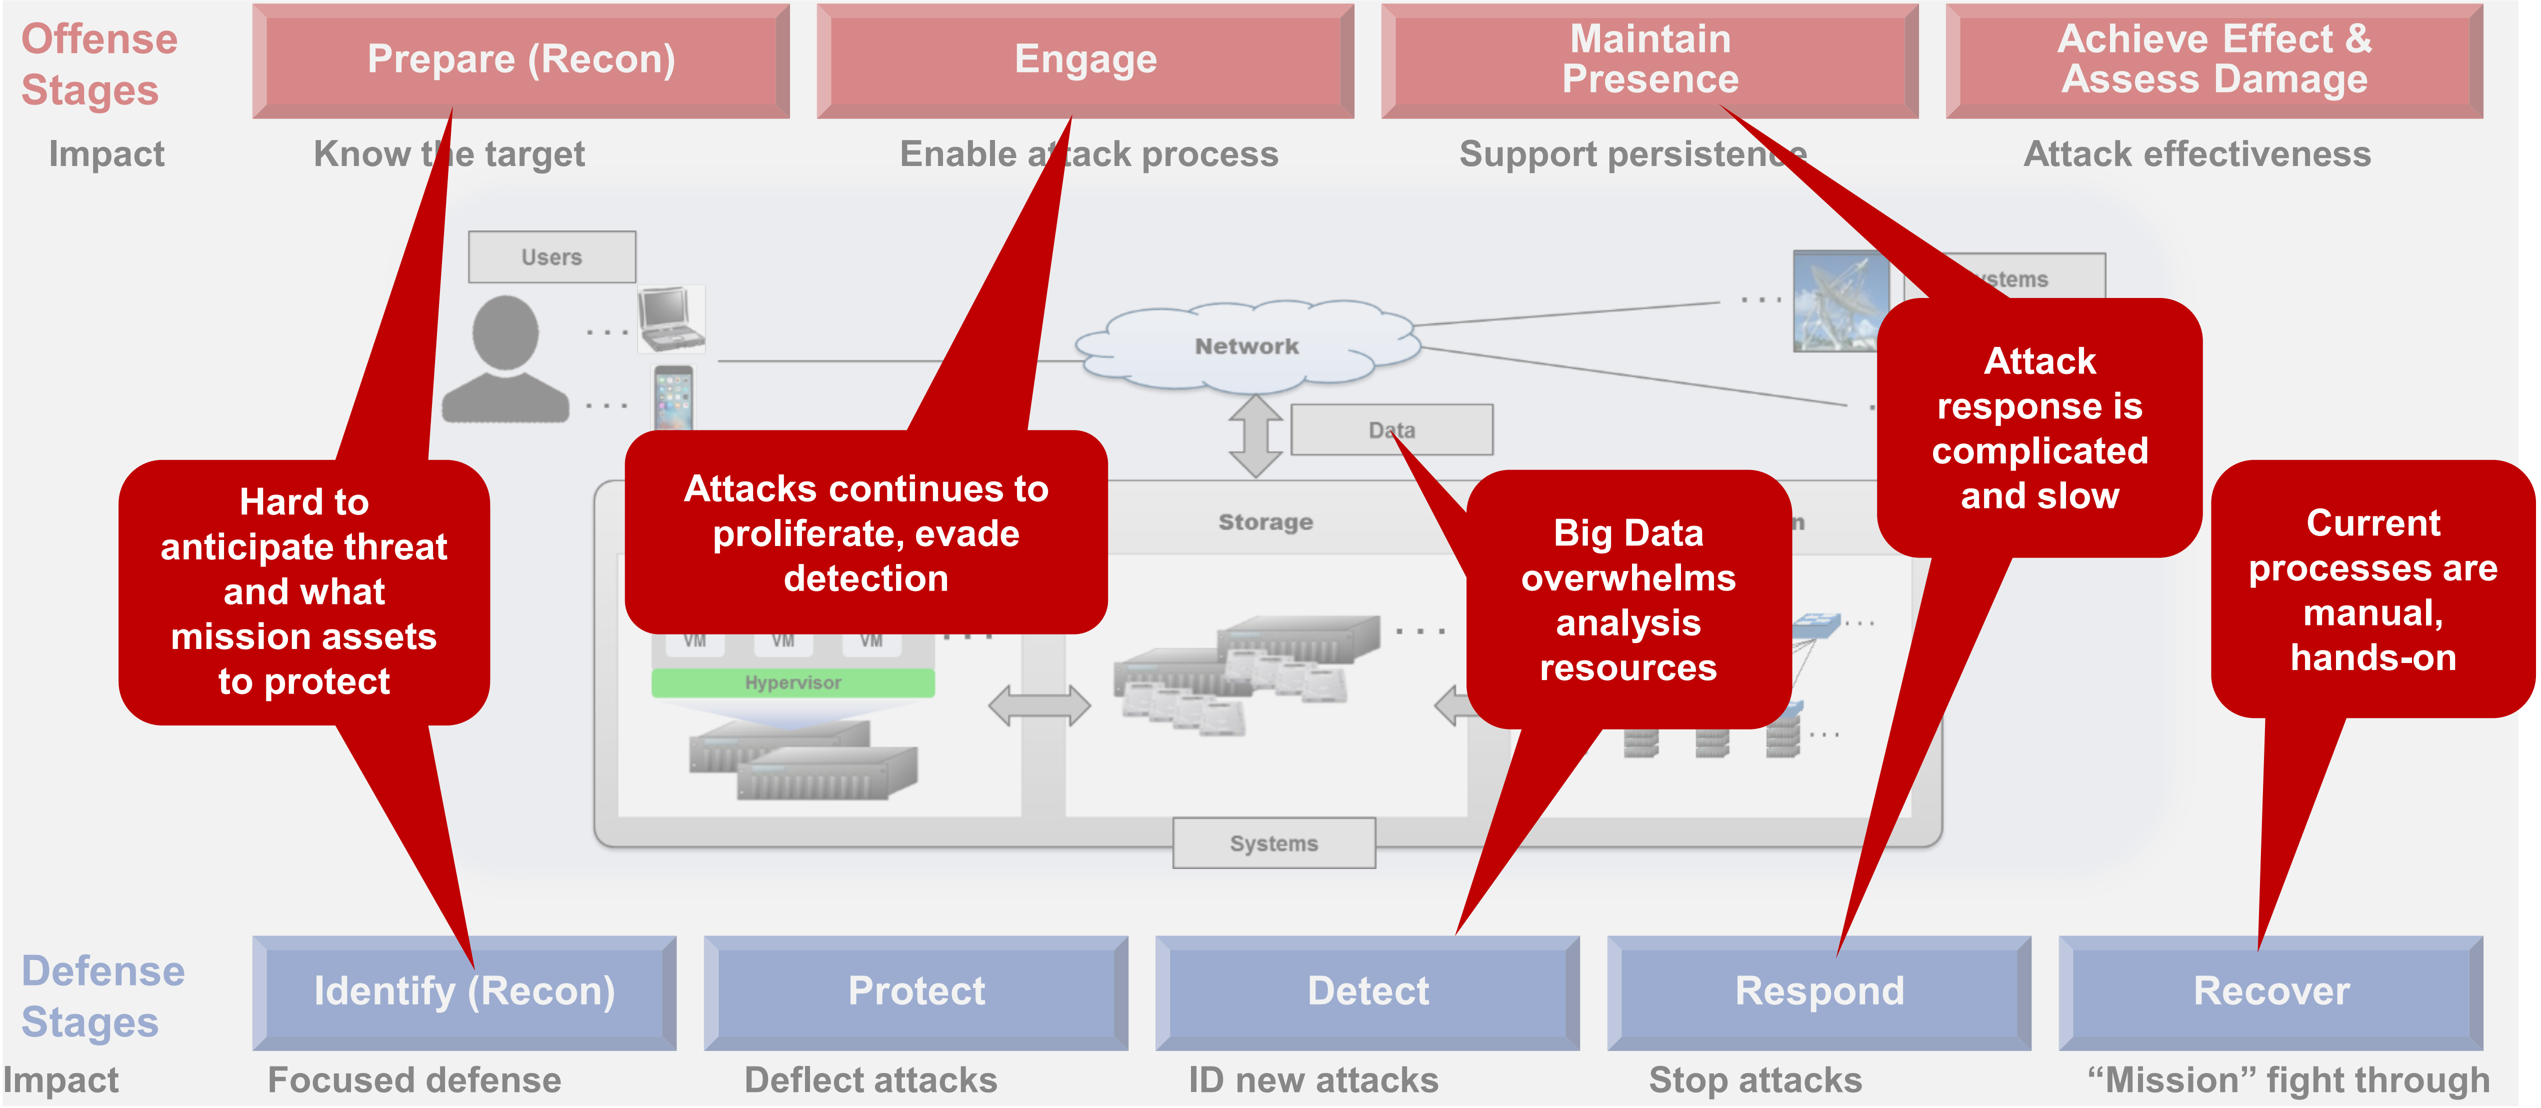

COA – Course of Action

AI Winter for Cyber

- **"Outside the Closed World: On Using Machine learning for Network Intrusion Detection", (Sommer, Paxson, 2010)**
- **Described significant challenges with applying machine learning (e.g. AI) to cyber attack intrusion detection**
- **Caused 'AI for Cyber Winter' that forced people to abandon AI-based approaches in cyber security altogether**
- **Paxson later acknowledged impact was not as intended (AICS16, 2016)**
  - **Only applied to intrusion detection**
  - **Other cyber security aspects are amenable to AI**

Outside the Closed World:
On Using Machine Learning For Network Intrusion Detection

Robin Sommer
International Computer Science Institute, and
Lawrence Berkeley National Laboratory

Vern Paxson
International Computer Science Institute, and
University of California, Berkeley

***Abstract*—In network intrusion detection research, one popular strategy for finding attacks is monitoring a network's activity for *anomalies*: deviations from profiles of normality previously learned from benign traffic, typically identified using tools borrowed from the machine learning community. However, despite extensive academic research one finds a striking gap in terms of actual deployments of such systems: compared with other intrusion detection approaches, machine learning is rarely employed in operational "real world" settings. We examine the differences between the network intrusion detection problem and other areas where machine learning regularly finds much more success. Our main claim is that the task of finding attacks is fundamentally different from these other applications, making it significantly harder for the intrusion detection community to employ machine learning effectively. We support this claim by identifying challenges particular to network intrusion detection, and provide a set of guidelines meant to strengthen future research on anomaly detection.**

***Keywords*-anomaly detection; machine learning; intrusion detection; network security.**

I. Introduction

Traditionally, network intrusion detection systems (NIDS) are broadly classified based on the style of detection they are using: systems relying on *misuse-detection* monitor activity with precise descriptions of known malicious behavior, while *anomaly-detection* systems have a notion of normal activity and flag deviations from that profile.[1] Both approaches have been extensively studied by the research community for many years. However, in terms of actual deployments, we observe a striking imbalance: in operational settings, of these two main classes we find almost exclusively only misuse detectors in use—most commonly in the form of signature systems that scan network traffic for characteristic byte sequences.

This situation is somewhat striking when considering the success that machine-learning—which frequently forms the basis for anomaly-detection—sees in many other areas of computer science, where it often results in large-scale deployments in the commercial world. Examples from other domains include product recommendations systems such as used by Amazon [3] and Netflix [4]; optical character recognition systems (e.g., [5], [6]); natural language translation [7]; and also spam detection, as an example closer to home [8].

In this paper we set out to examine the differences between the intrusion detection domain and other areas where machine learning is used with more success. Our main claim is that the task of finding attacks is fundamentally different from other applications, making it significantly harder for the intrusion detection community to employ machine learning effectively. We believe that a significant part of the problem already originates in the premise, found in virtually any relevant textbook, that anomaly detection is suitable for finding *novel* attacks; we argue that this premise does not hold with the generality commonly implied. Rather, the strength of machine-learning tools is finding activity that is *similar to something previously seen*, without the need however to precisely describe that activity up front (as misuse detection must).

In addition, we identify further characteristics that our domain exhibits that are not well aligned with the requirements of machine-learning. These include: *(i)* a very high cost of errors; *(ii)* lack of training data; *(iii)* a semantic gap between results and their operational interpretation; *(iv)* enormous variability in input data; and *(v)* fundamental difficulties for conducting sound evaluation. While these challenges may not be surprising for those who have been working in the domain for some time, they can be easily lost on newcomers. To address them, we deem it crucial for any effective deployment to acquire deep, semantic *insight* into a system's capabilities and limitations, rather than treating the system as a black box as unfortunately often seen.

We stress that we do *not* consider machine-learning an inappropriate tool for intrusion detection. Its use requires care, however: the more crisply one can define the context in which it operates, the better promise the results may hold. Likewise, the better we understand the semantics of the detection process, the more operationally relevant the system will be. Consequently, we also present a set of guidelines meant to strengthen future intrusion detection research.

[1] Other styles include *specification-based* [1] and *behavioral* detection [2]. These approaches focus respectively on defining allowed types of activity in order to flag any other activity as forbidden, and analyzing patterns of activity and surrounding context to find secondary evidence of attacks.

Sommer, Paxson paper shifted cyber security research focus from AI to secure methods

Cyber AI Start-up Landscape

However, the market is now flush with companies leveraging AI for cyber

Outline

- **Background**
- **Lay-of-the-Land**
- **AI for Cyber Security**
  - **Background**
  - **Findings and Recommendations**
- **Summary**

AI Canonical Architecture

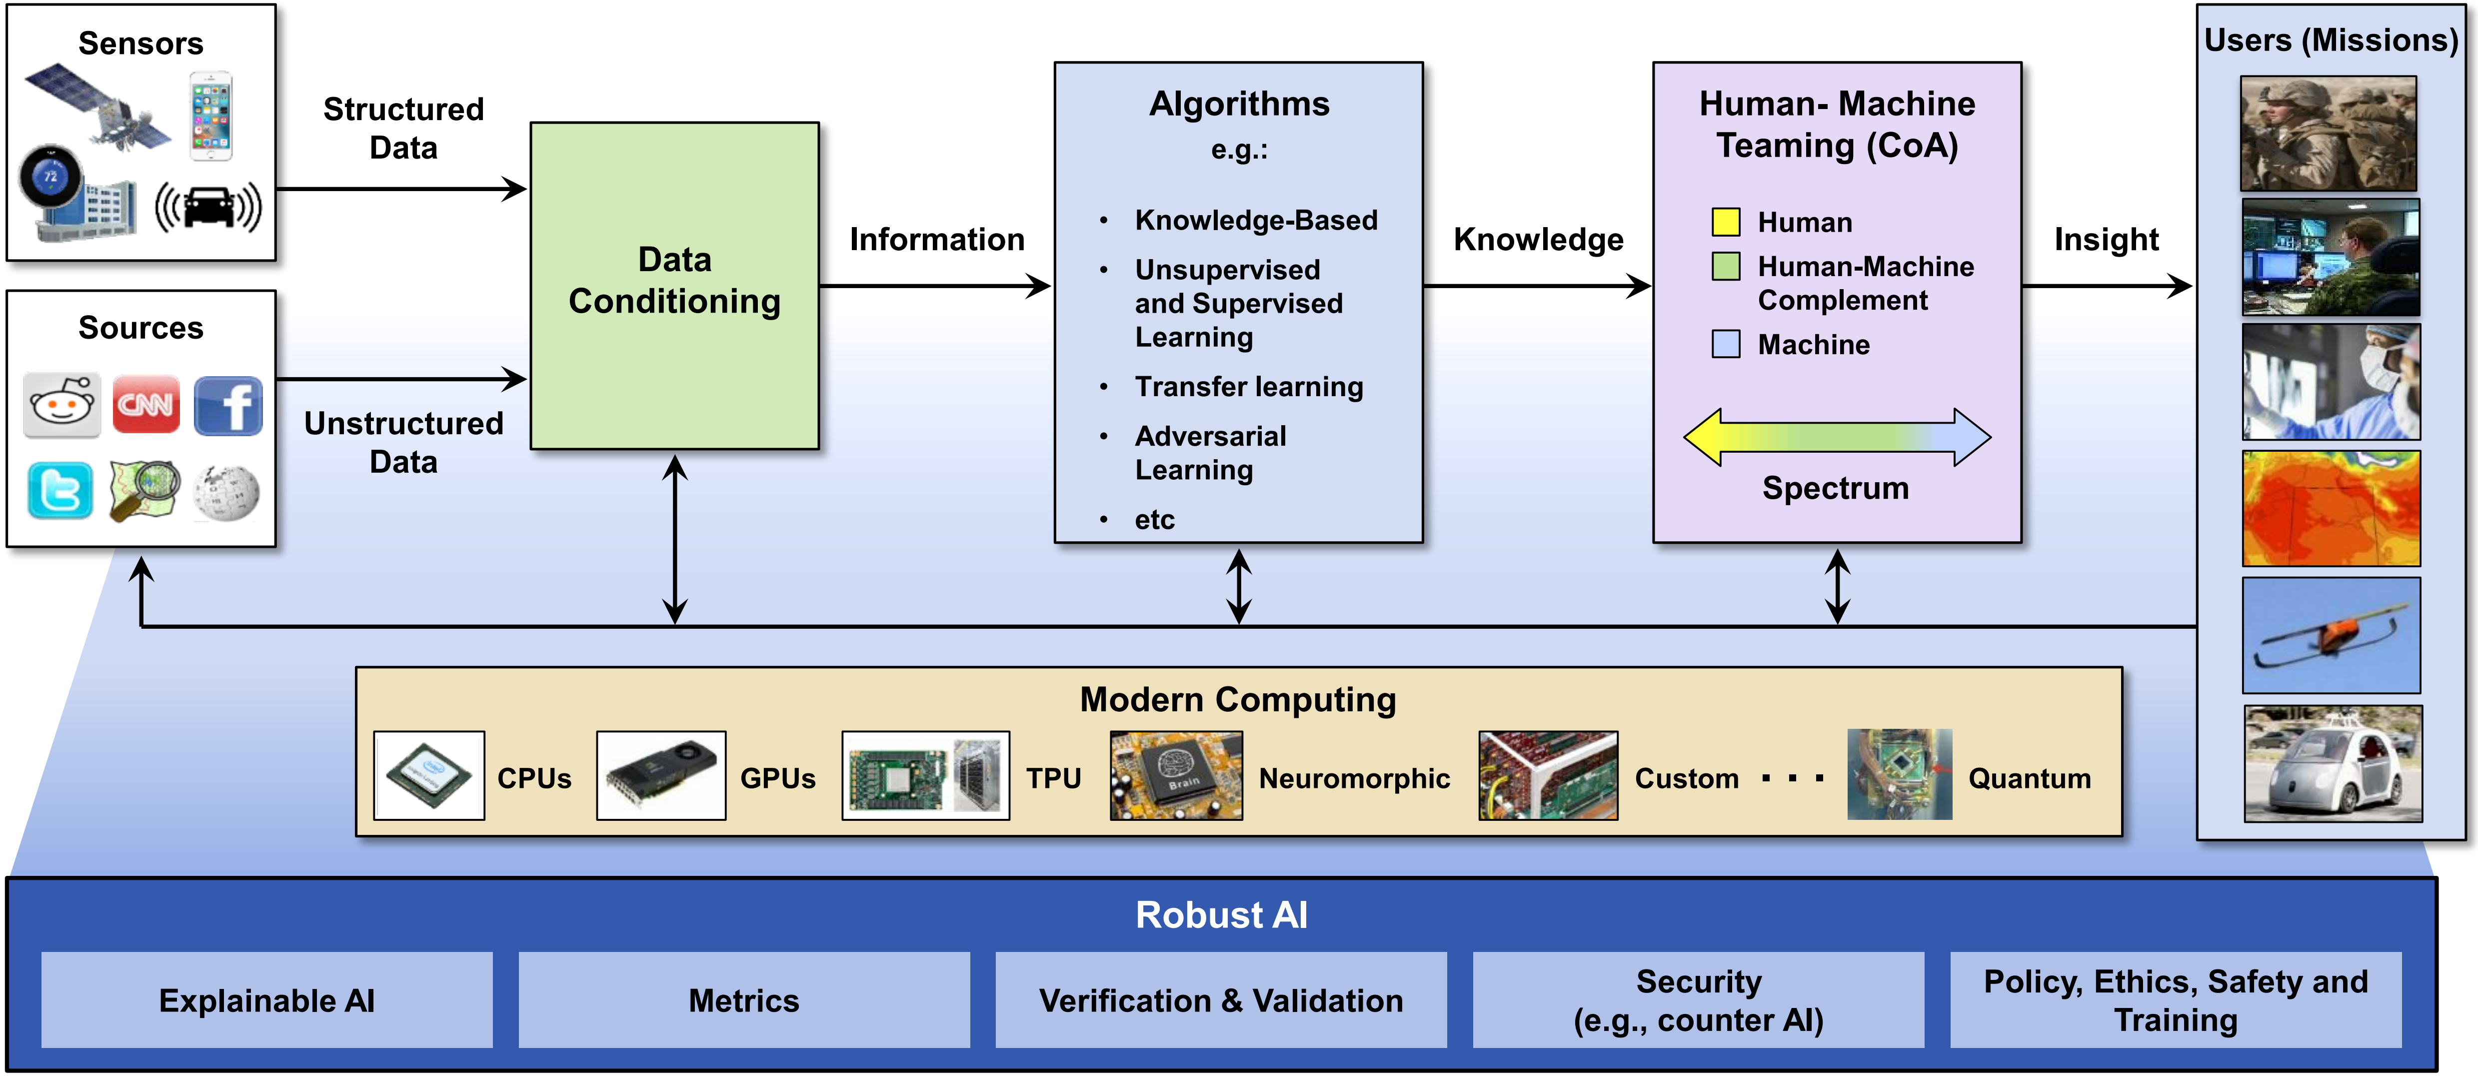

GPU = Graph Processing Unit
TPU = Tensor Processing Unit
CoA = Courses of Action

Unstructured and Structured Data

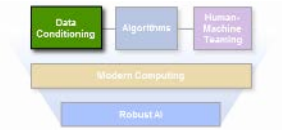

Structured Data Types

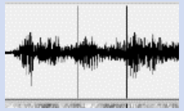
Speech


Sensors


Network Logs


Metadata

Unstructured Data Types


Social Media


Human Behavior

Reports

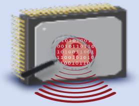
Side Channel

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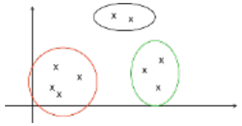 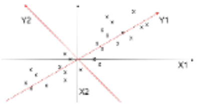	• Limited machine learning • Dimensionality Reduction • Clustering/Pattern Recognition • Outlier Detection
Data Labeling 	• Initial data exploration • Highlight missing or incomplete data • Reorient sensors/recapture data • Look for errors/biases in collection

Important needs are in labeling data and automating data conditioning

Data Conditioning: The Open-Source Intelligence Opportunity Big Data Boom

- Open source data is growing exponentially
  - 2.8B Internet users
  - 2B smartphone users
  - Commercial satellites and imagery coming online
- Data are rich with information about systems, users, organizations, relationships, events
- Data can be used to enrich information from classified sources

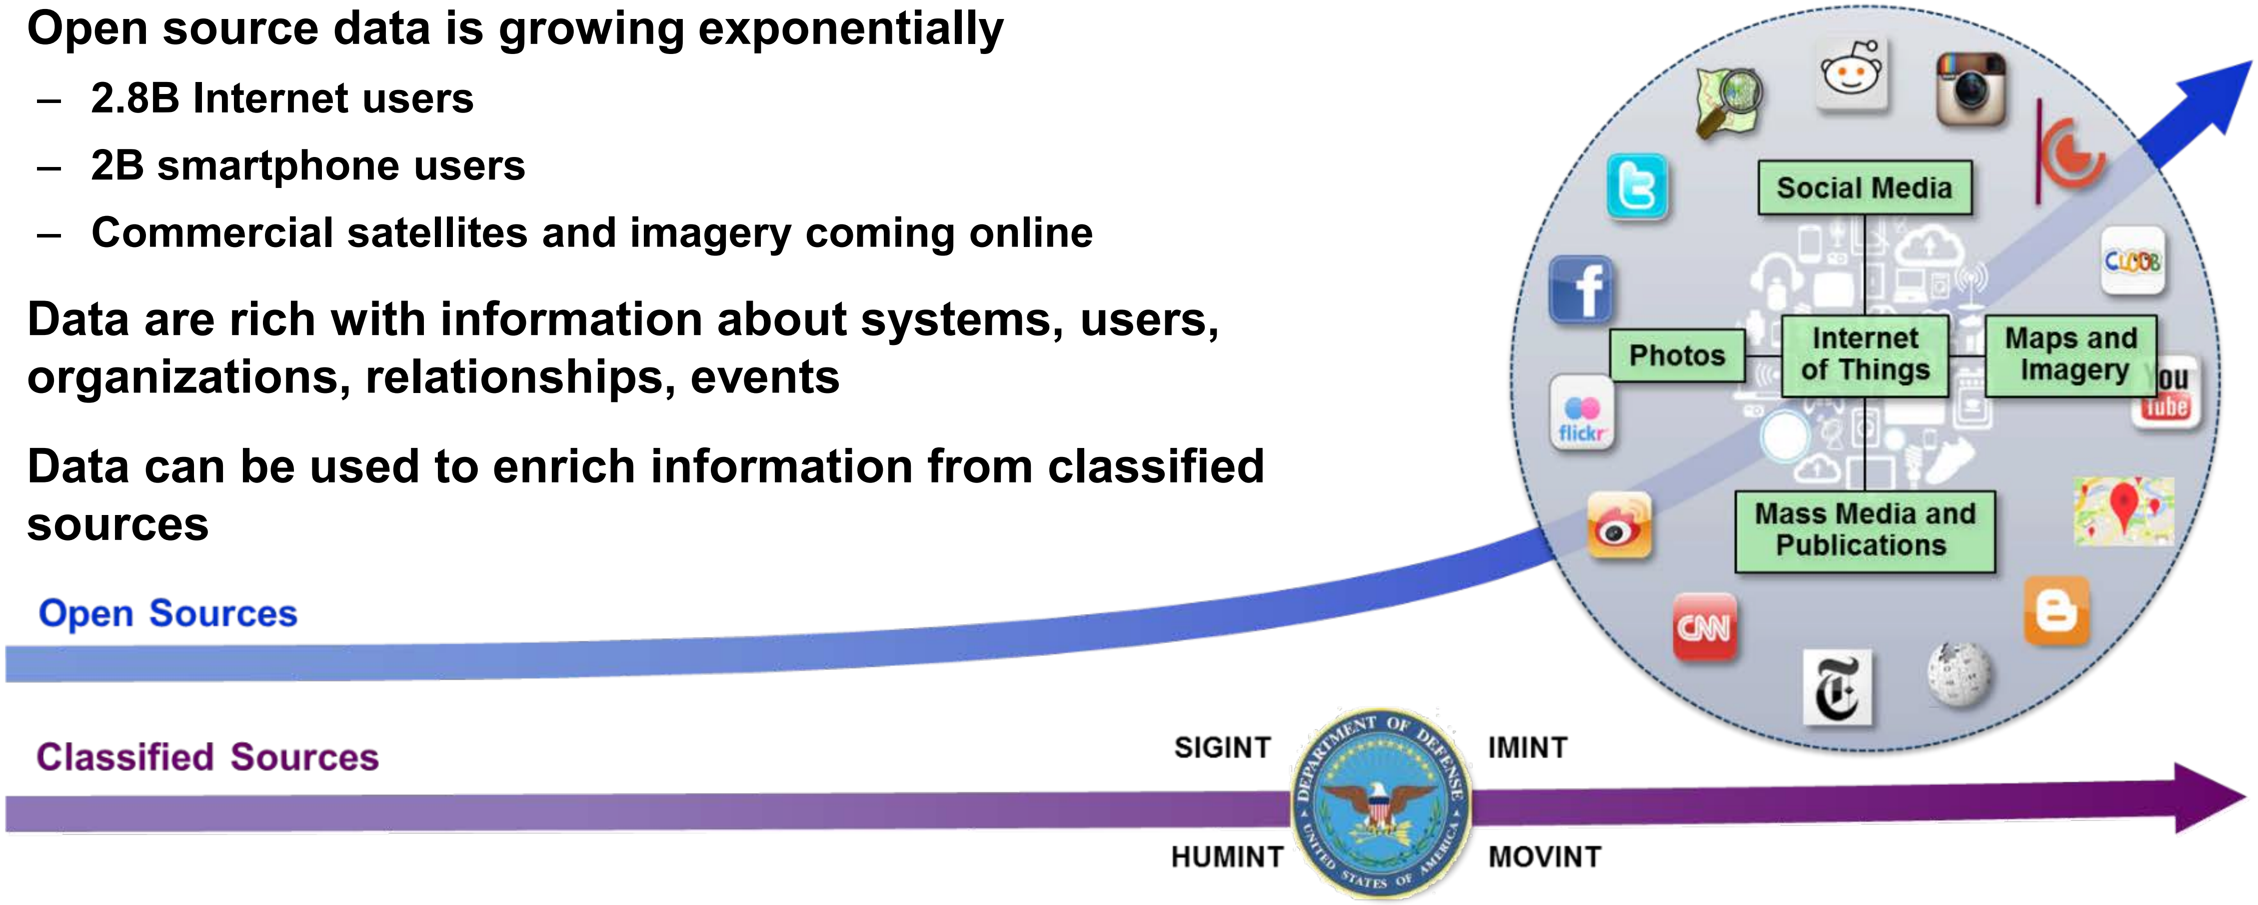

Finding #1: Cyber data is voluminous and is both structured and unstructured

Source: Internet Trends Report 2015 KPCB

Large Truth-marked Cyber Datasets Are Hard to Find

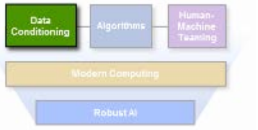

- **Continued need for commercial and government Enterprises to share data from incidents**
- **Some databases exist but are not easy to use or widely accessible**
- **Very little cyber data is truth-marked**
- **Much academic research still leverages antiquated datasets**

Cybersecurity Information Sharing Act of 2015

May 2016 Volume 11, Issue 5

From the Desk of Thomas F. Duffy, Chair

We've all heard talk of the Cybersecurity Information Sharing Act, but what does it really mean? We hope that this newsletter is a quick cheat sheet that highlights the key takeaways, as well as provide resources for additional information if you'd like to conduct a deeper dive into the topic.

Table 1 Summary of popular datasets in the intrusion detection domain [30]

Data source	Dataset name	Abbreviation
Network Traffic	DARPA 1998 TCPDump Files	DARPA98
	DARPA 1999 TCPDump Files	DARPA99
	KDD99 Dataset	KDD99
	10% KDD99 Dataset	KDD99-10
	Internet Exploration Shootout Dataset	IES
User behavior	Unix User Dataset	UNIXDS
System call sequences	DARPA 1998 BSM Files	BSM 98
	DARPA 1998 BSM Files	BSM 99
	University of New Mexico Dataset	UNM

Finding #2: Lack of ground truth for cyber inhibits algorithm application to DoD problems

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

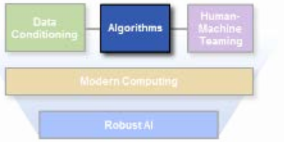

Machine Learning Algorithms Taxonomy

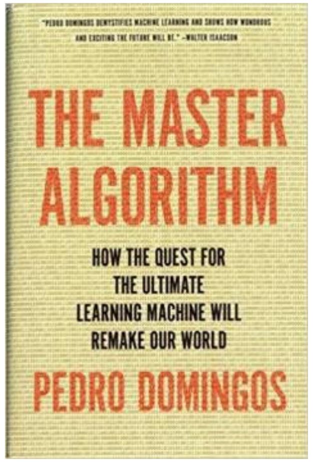

Algorithms*

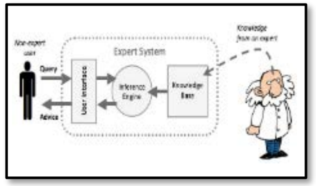

Symbolists (e.g., exp. sys.)

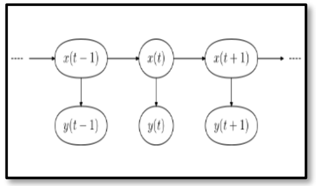

Bayesians (e.g., naive Bayes)

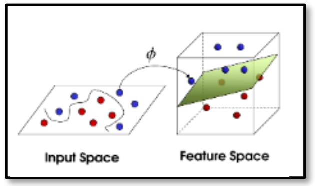

Analogizers (e.g., SVM)

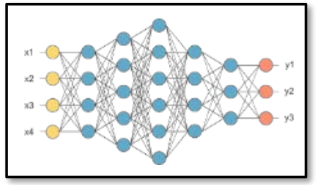

Connectionists (e.g., DNN)

Evolutionaries (e.g., genetic programming)

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

Machine Learning Applied to Classifiers

Training Data → Model

Iterate

Cross-Validation Data → Model Assessment

Test Data → Classifier → Output

→ Vision Deep CNN

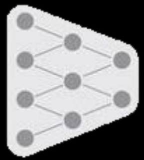

→ Language generating RNN

→ A group of people shopping at an outdoor market.

There are many vegetables at the fruit stand.

CNN = convolution neural network
RNN = recurrent neural network

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

AI Algorithms for Cyber

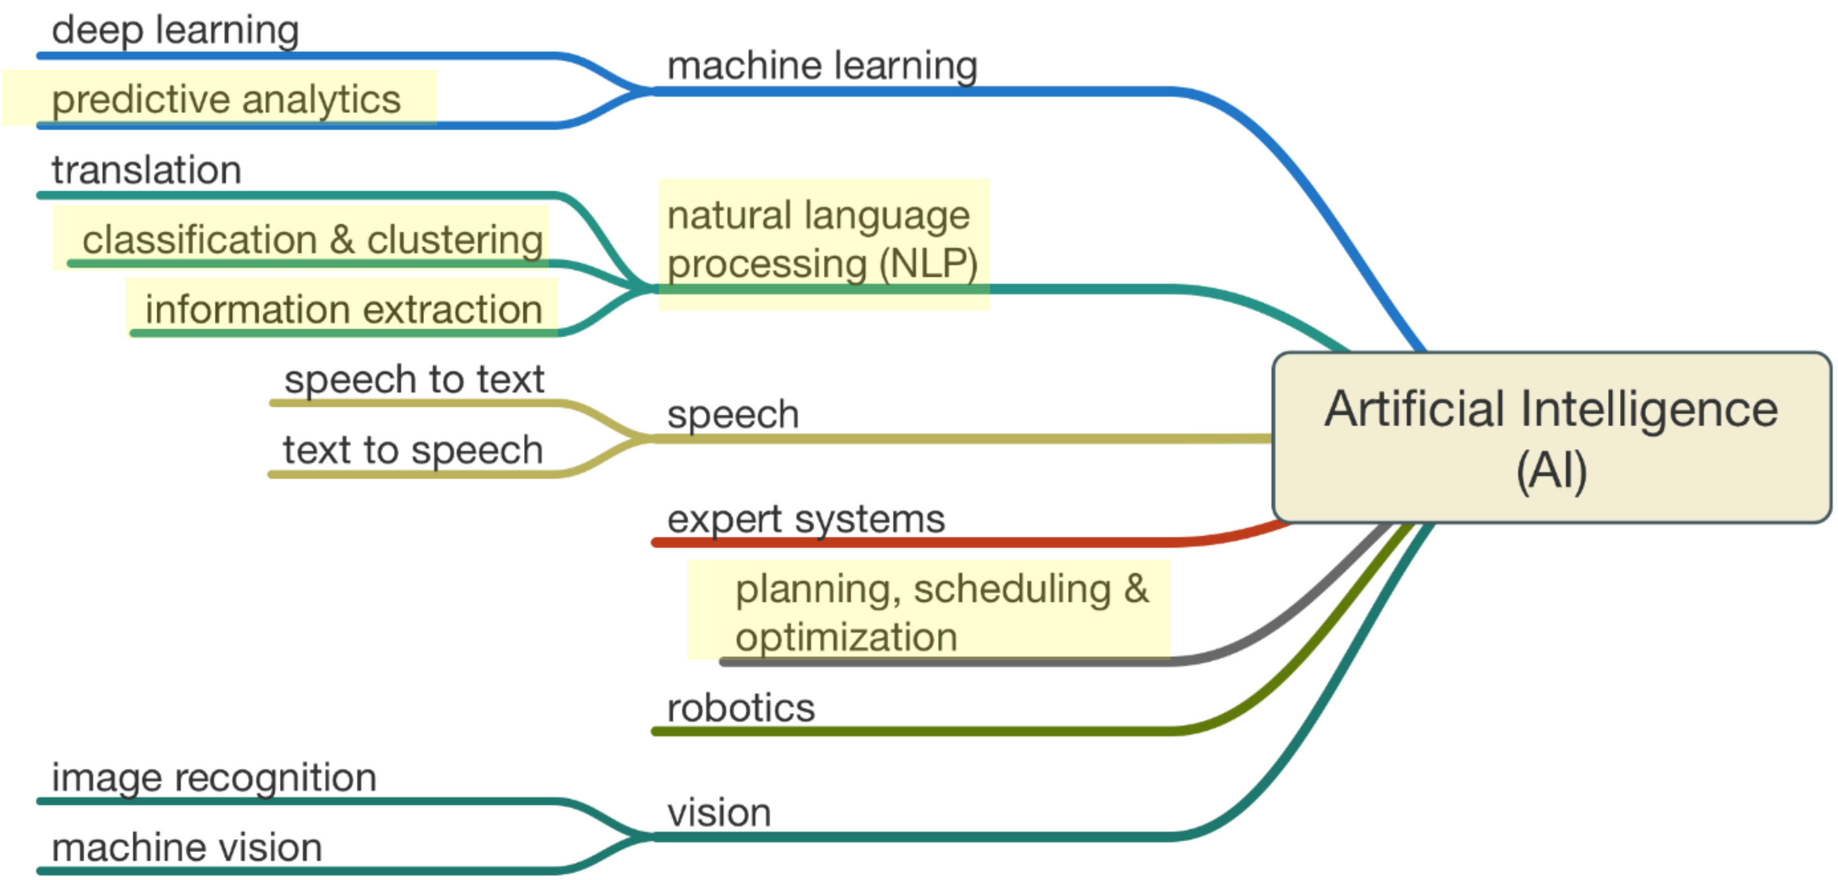

Finding #3: **Many algorithms exist which can be applied to cyber security**

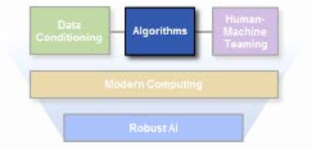

Artificial Intelligence Can Help

Offense Stages: Prepare (Recon) | Engage | Maintain Presence | Achieve Effect & Assess Damage

Impact: Know the target | Enable attack process | Support persistence | Attack effectiveness

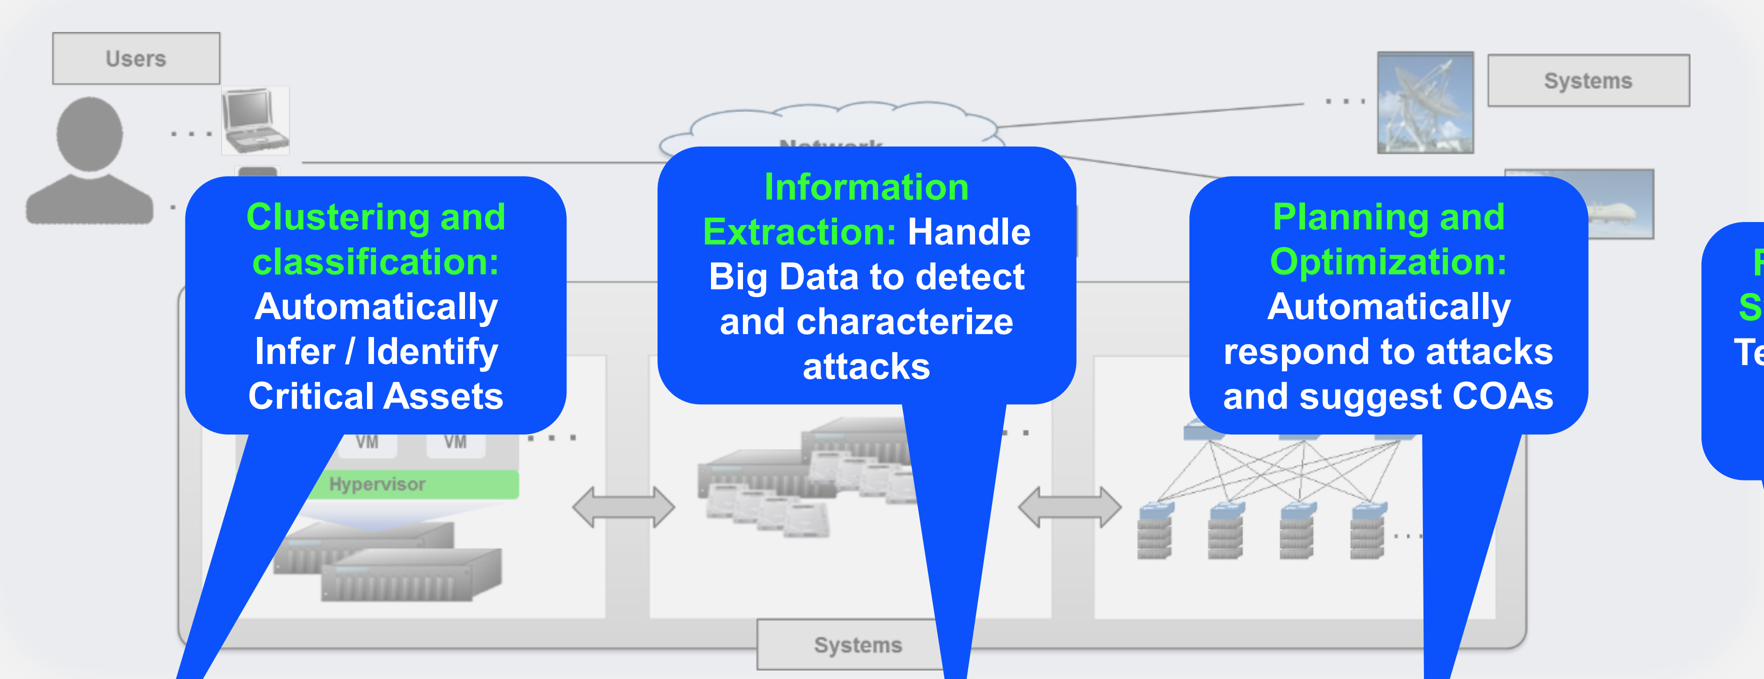

Predictive Analytics: Monitor and anticipate the threat

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

AI Resources Are Available on Line

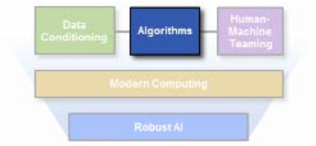

- **Open source toolkits allow users to leverage machine learning easily**
- **Commercial companies build business on AI tool kits that can be applied easily**
- **A Knowledge Base of Shared Knowledge and Solutions**

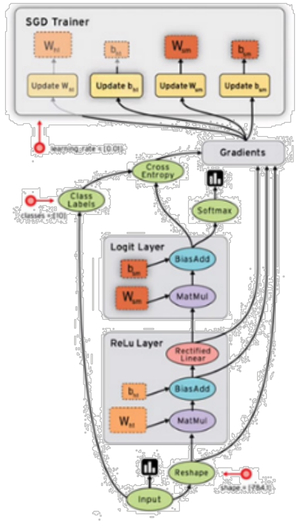

ReLuPlex

TensorFlow

Microsoft Azure

Apache Singa

Finding #4: Academia, commercial sectors are advancing algorithms and AI capabilities

Finding #5: Peer organizations are benefiting from open source communities

Human Machine Teaming

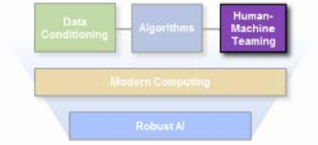

Study Finds Cyberthreat Data Overwhelming to Security Workers

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

Finding #6: Cyber security data overwhelms overworked analysts

Finding #7: The US / DoD faces serious workforce shortages in cyber security expertise

Human Machine Teaming: Automated Cyber Decision Making

- **CASCADE - Cyber Adversarial SCenario modeling and Automated Decision Engine**
  - Dynamically quantifies risk in the face of an adaptive adversary
  - Considers mission context to selection optimal course of action (COA)
  - Prototype applied to configuration of network segmentation defense

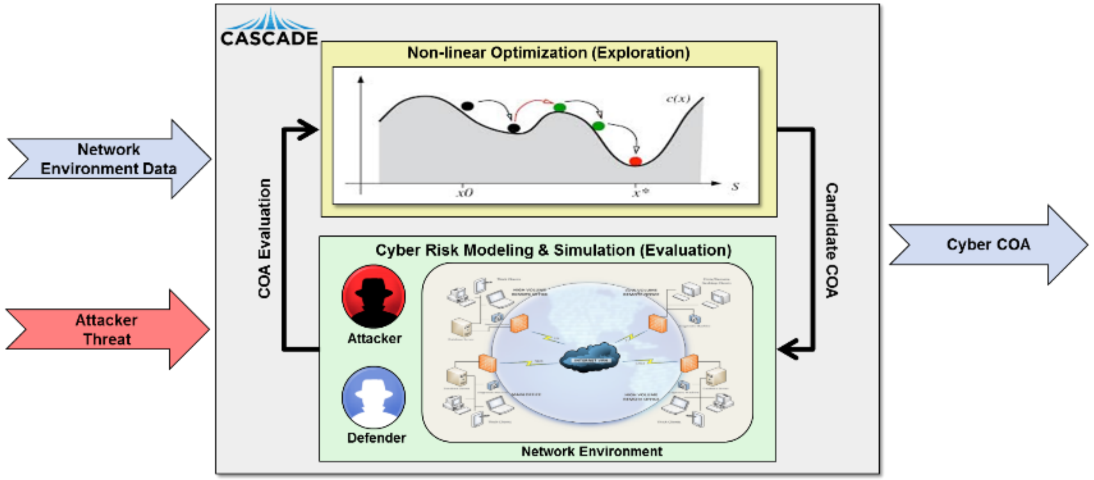

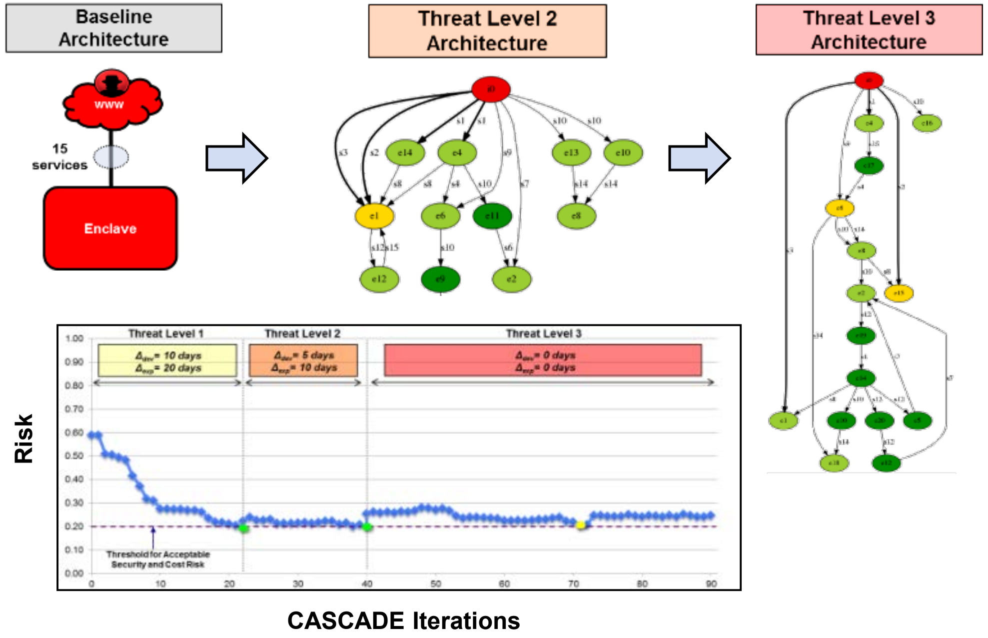

Robust AI: Engendering Trust

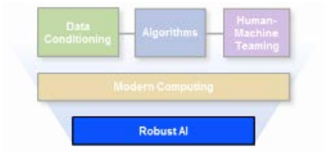

Confidence Level vs. Consequence of Actions

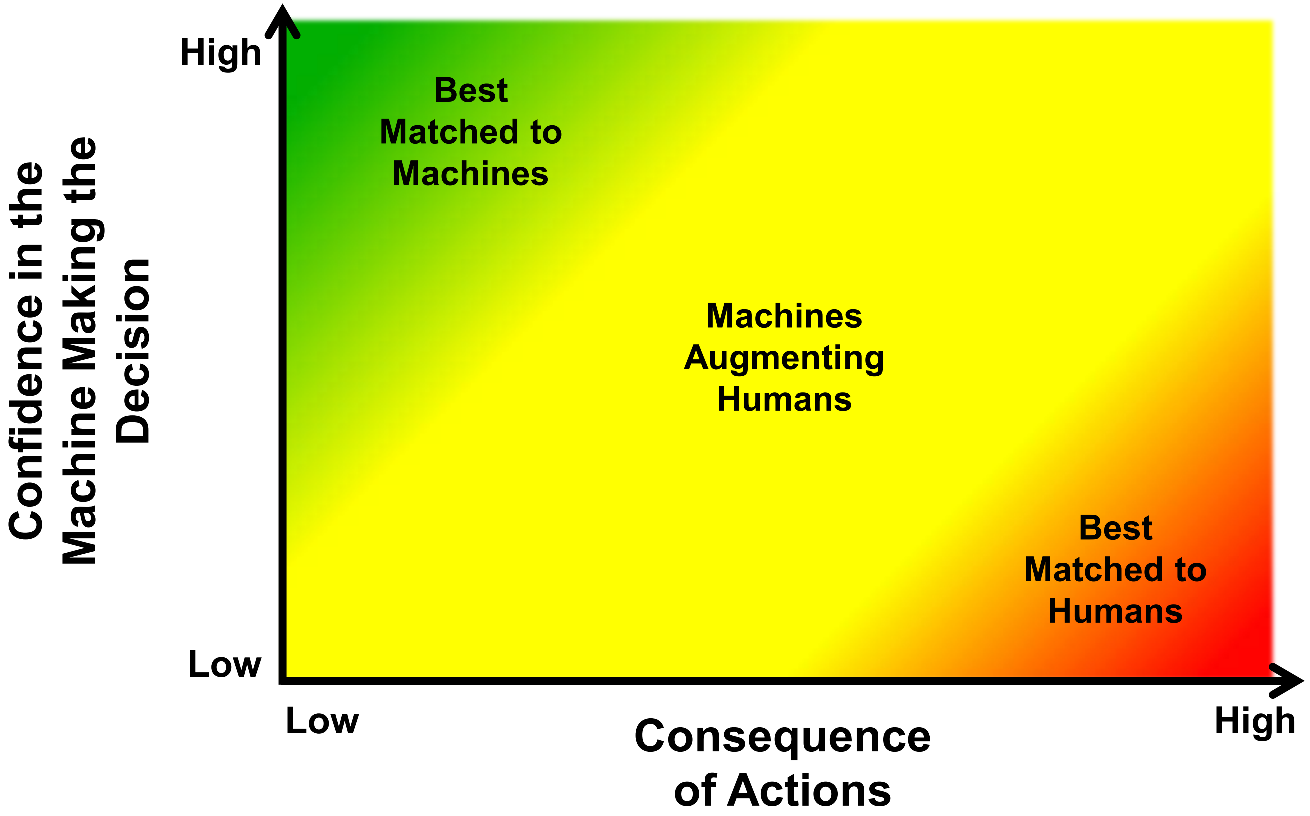

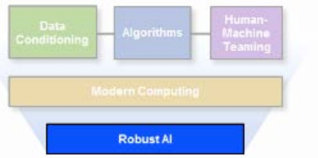

AI for Cyber must be Robust as Well

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

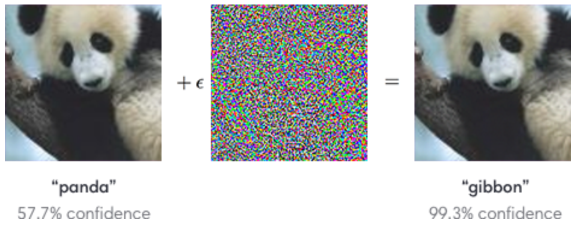

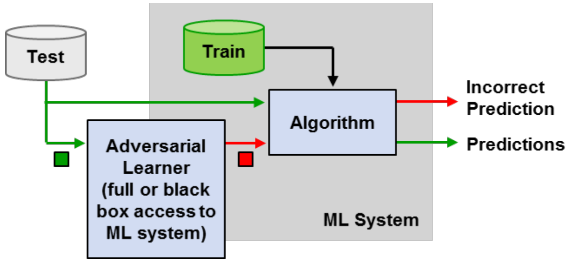

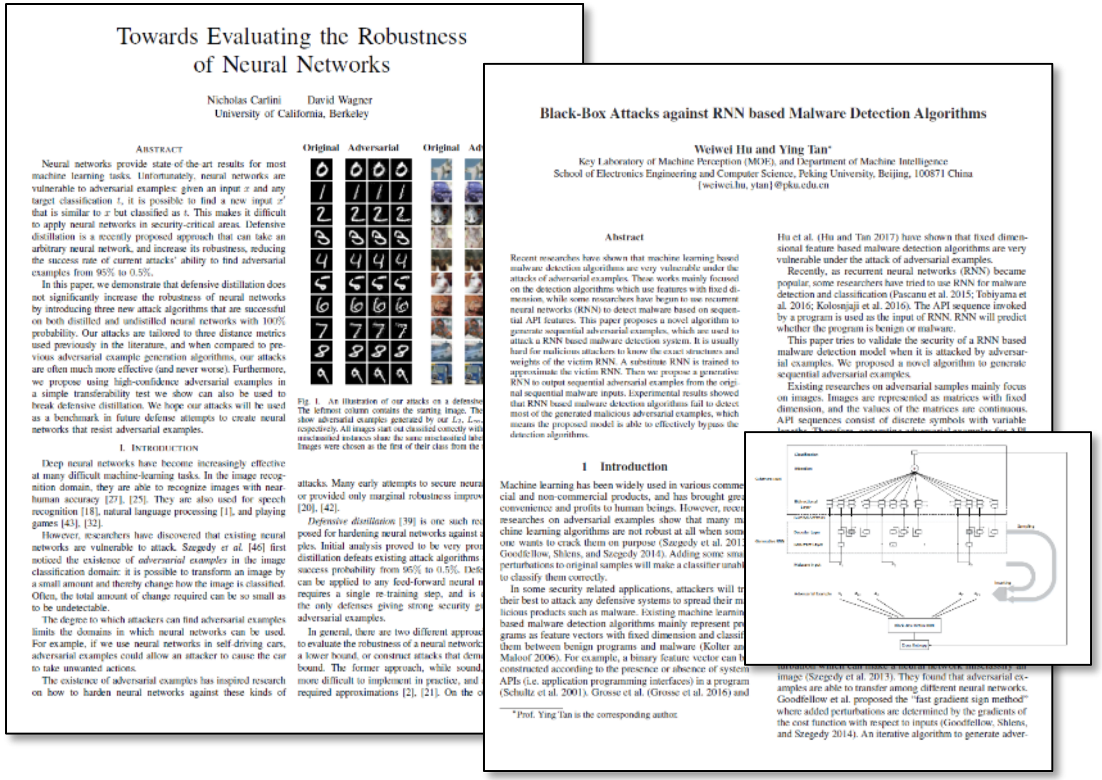

- **Cyber examples are appearing in literature demonstrating capabilities**
  - **Malware evades detection**
  - **Nefarious connections hidden by noise**
  - **Etc ..**

Finding #8: Adversarial attacks can limit effectiveness of cyber AI solutions

Finding #9: Vulnerable Cyber AI detection, classification algorithms can lead to incorrect behavior

Robust AI: Evaluating Classifier Performance using Human Expert Judgment

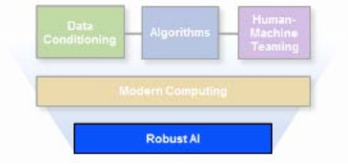

- **AI and ML currently applied mostly to simple, low consequence problems**
  - **We want to transition use to hard, high consequence problems**
- **Methods are needed to evaluate AI classifiers that leverage expert judgement**
  - **Must be robust to inter-rater dependence and variability**

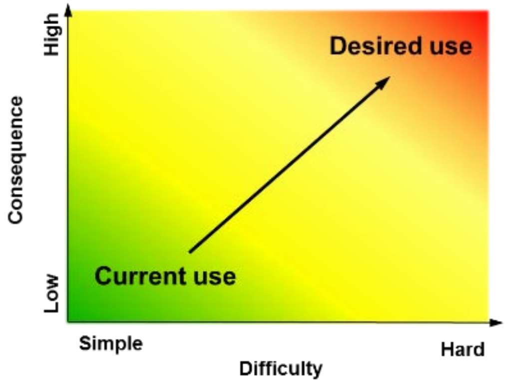

Scenario: Fake News Detection

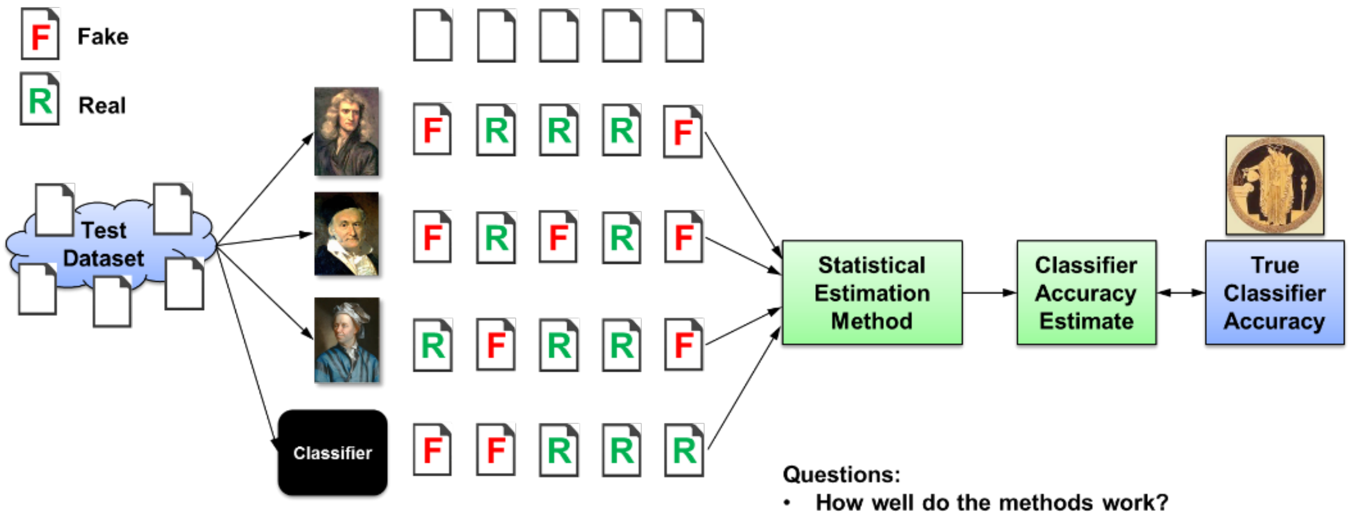

Questions:
- How well do the methods work?
- What factors impact the quality of the estimate?

Expert Dependence

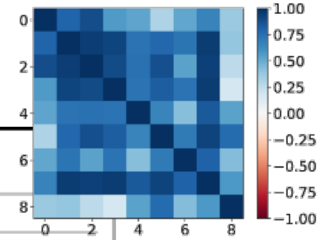

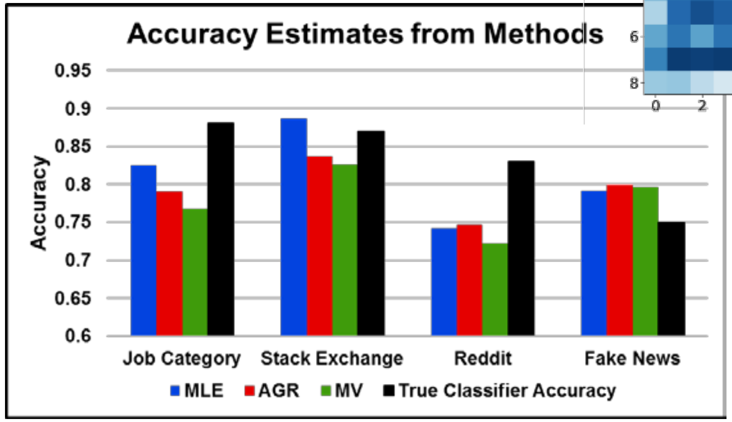

MLE: Maximum Likelihood Estimation
AGR: Agreement
MV: Majority Vote

Modern AI Computing Engines

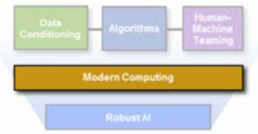

Computing Class

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)
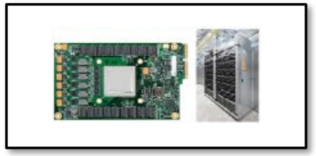	TPU	• Speeds up inference time (domain specific architecture)
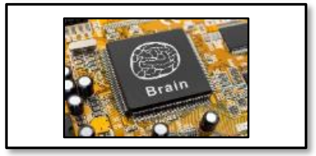	Neuromorphic	• Still a research area
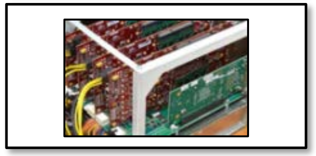	Custom	• Ability to speed up specific computations of interest (e.g. graphs)
	⋮	
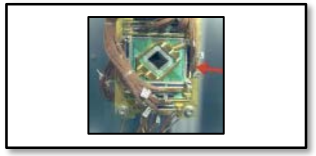	Quantum	• Benefits unproven until now • Recent results on HHL (linear system of equations)

Selected Results

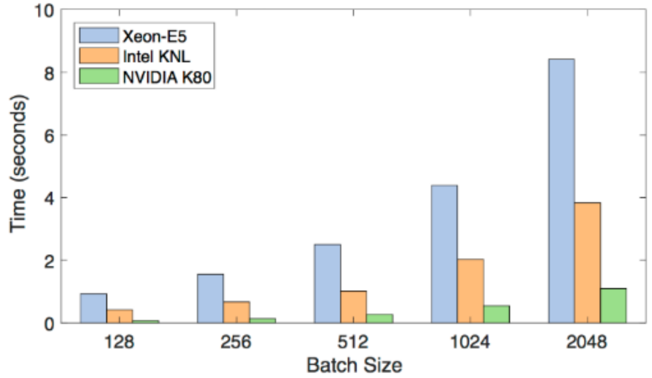

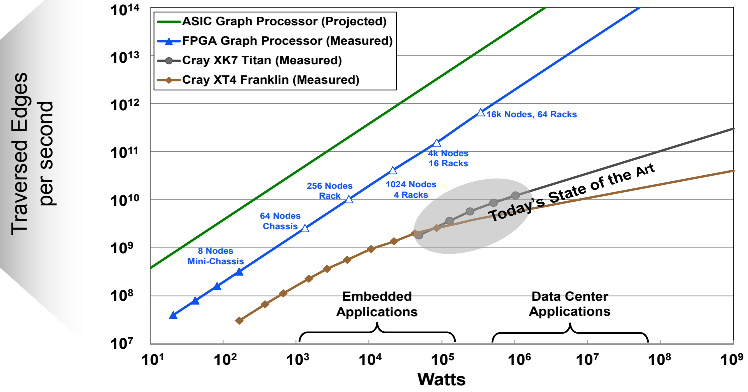

GPU = Graph Processing Unit
TPU = Tensor Processing Unit

HHL = Harrow-Hassidim-Lloyd quantum algorithm

Outline

- Background
- Lay-of-the-Land
- AI for Cyber Security
  - Background
  - Highlights, Findings and Recommendations
- Summary

Summary of Study Findings & Recommendations

	Findings	Recommendations
Data	1. Cyber data is voluminous and is multi-domain, structured and unstructured 2. Lack of ground truth for cyber inhibits algorithm application to DoD problems	✓ Lead the way in Cyber Big Data conditioning by leveraging expertise in Big Data collection, creation and curation to support AI for Cyber
Algorithms	3. Many algorithms exist which can be applied to cyber 4. Academia, commercial sectors are advancing algorithms and AI capabilities 5. Peer organizations are benefiting from open source communities	✓ Engage with academic community to maintain awareness of and influence where possible, leverage open-source toolkits and libraries to jumpstart DoD mission capabilities
Human Machine Teaming	6. Declining human resource environment creates opportunity to help 7. Recommender systems at core of much commercial AI success	✓ Automate and augment cyber tasks of data triage, correlation, leveraging active learning to improve AI solutions, capitalize on analyst cyber expertise
Robust AI	8. Adversarial attacks can limit effectiveness of cyber AI solutions, leading to incorrect behavior 9. Promising work in 'proving' AI behavior is appearing in academia	✓ Lead the way in robust AI for cyber for DoD applications, leveraging and applying recent work in academia

LL Cyber AI-related Workshops and Symposiums

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

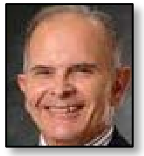

Dave Martinez

Neal Wagner

New Orleans, Louisiana • February 2, 2018

Theme: Applications of AI to Internet of Things

Keynotes

Sal Stolfo
Professor of Computer Science
Dept. of Computer Science,
Columbia University

Trung Tran
Laboratory of Physical Sciences,
University of Maryland,
Baltimore County

Graph Exploitation Symposium

- Brings together leading experts from universities, industry, and government to explore the state of the art and define a future roadmap in network science

Dedham, Massachusetts • April 23-25, 2018

Technical Co-Chairs

Ben Miller

Rajmonda Caceres

Chairs

Bill Streilein

Sanjeev Mohindra

POC: Ben Miller, bamiller@ll.mit.edu

Summary

- U.S. needs to regain AI leadership by strategically partnering with small and large commercial companies plus academia
- Potential for major impact remains for DoD applications
  - Although there is a lot of activity in community, only pockets of cyber success exist
- Transfer of algorithms to DoD mission is challenging

- Demonstrated achievements in applying AI to cyber
  - Fluent in Big Data architectures and databases
  - Cyber discussion detection, traffic characterization, counterfeit detection
- Focus should be on unique areas of expertise, connection to mission
  - Mission process and data requirements
  - Adapting latest algorithms to mission needs
  - Developing robust AI solutions

AI Bibliography List (few selected set)

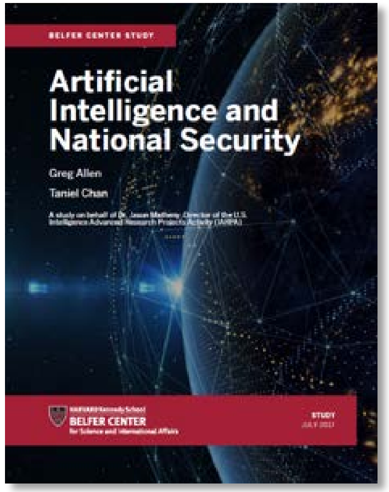

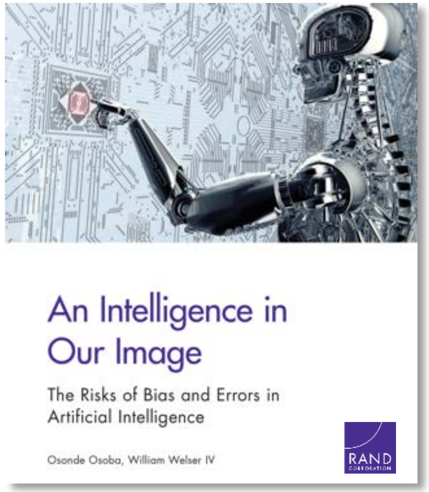

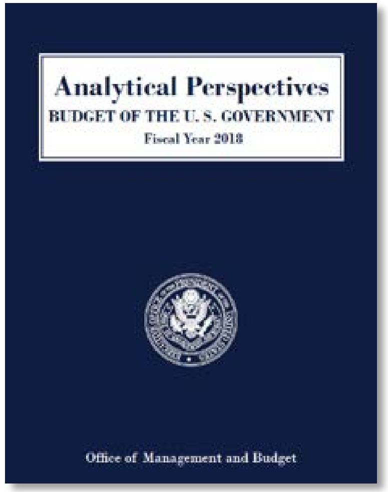

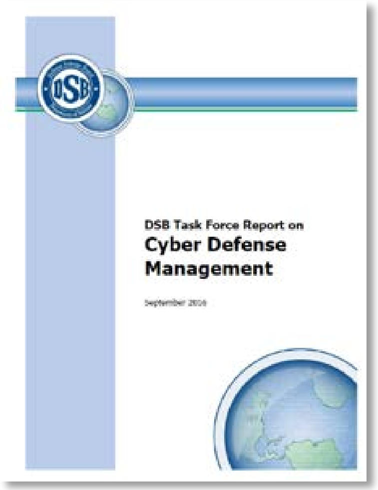

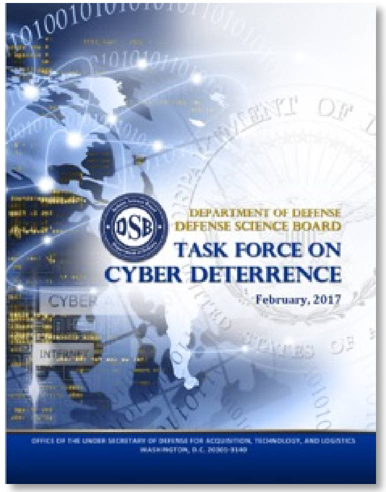

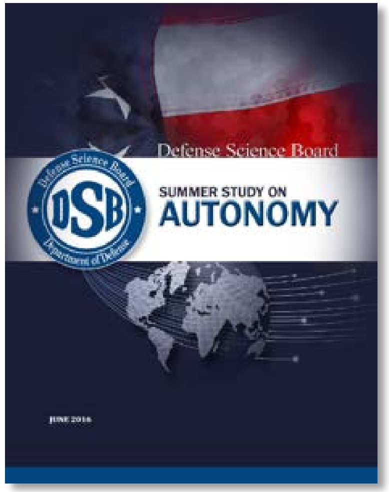

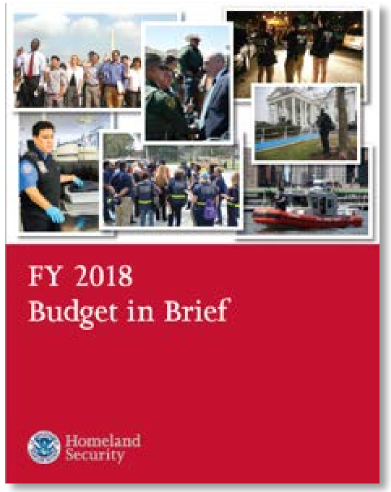

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY