INL – ML & AI Symposium

April 17, 2020

Purpose of Meeting:

- Introduce the topic of ML and AI to INL researchers
- Provide examples of how ML an AI are being applied across other industries
- Discuss current ML & AI research and capabilities at INL
- Discuss planned activities, including engagement opportunities and collaboration opportunities

Presentations will include:

- Provide overview on Topic Area;
- Describe the status of industry
- Identify Issues (if any) and potential impact
- High level discussion of planned activities and outcomes

Agenda for Machine Learning and Artificial Intelligence Symposium

Friday, April 17th, 2020;

Time	Subject	Speaker
11:00	Welcome, Introductions, and Agenda	Curtis Smith
11:15	What is AI?	R. Kunz
11:25	AI, ML, and Statistics, oh My!	N. Lybeck
11:35	Modeling Human Cognition: It's Not All Machine Learning	R. Boring
11:45	Smart Reactors	Humberto Garcia
11:55	AI in Robotics and Applying Natural Connections	V. Walker
12:05	AI as Automation	K. Le Blanc
12:15	ML in current projects	V. Agarwal
12:25	ML in current projects	A. Al Rashdan
12:35	HPC Building a Scientific Language Model – Leveraging ArXive.org research data and RoBERTa	C. Krome
12:45	Reverse engineering of stripped binaries using scalable deep learning	M. Anderson
12:55	Closeout	Curtis Smith

Curtis Smith

Group: **Division Director for Nuclear Safety and Regulatory Research**
Education: **BS, MS, and PhD in Nuclear Engineering at ISU and MIT**

Presentation Overview

Motivation for AI/ML in science, math, and engineering

- **How AI/ML has advanced in the science, math, and engineering communities and how these advances may be used with INL applications such as computational risk assessment.**
- **These topics provide an insight into the potential for advanced analysis and operations for complex systems.**

My Motivation for AI/ML in Science, Math, and Engineering

Dr. Curtis Smith, Director
Nuclear Safety and Regulatory Research Division
Idaho National Laboratory

A discussion on:

How AI/ML has advanced in science, math, & engineering

How these advances may be used with INL applications such as computational risk assessment

The potential for advanced analysis and operations for complex systems

www.inl.gov

Idaho National Laboratory

Perhaps the first autonomous vehicle

What is Machine Learning/Artificial Intelligence (ML/AI)?

- **From Source of All Knowledge™ → Wikipedia**
- **Artificial intelligence (AI) is intelligence demonstrated by machines**
  - Study of "intelligent agents": **device that perceives its environment and takes actions** that maximize its chance of successfully achieving its goals
  - Machines that **mimic "cognitive" functions** that humans associate with the human mind, such as "learning" and "problem solving"
- **Machine learning (ML) is the scientific study of algorithms and statistical models to perform a specific task without using explicit instructions, relying on patterns and inference instead**
  - Subset of artificial intelligence
  - Builds a mathematical **model based on sample data** ("training data“) to make predictions or decisions without being explicitly programmed to perform the task
  - Closely related to computational statistics, which focuses on **making predictions using computers**

A question → can we use AI/ML for Science, Math, and Engineering??

Examples of current ML and AI applications

- **Symbolic reasoning to differentiate & integrate math**
  - Neural network used 80 million examples of 1st- and 2nd-order differential equations & 20 million examples of integrated by parts
  - How well does it work?
    - Significantly **outperforms Mathematica** (on integration, close to 100% accuracy)
      - Mathematica reaches 85%, Maple and Matlab perform less well
      - In many cases, conventional solvers unable to find a solution in 30 seconds
      - The neural net takes about a second to find its solutions
  - https://www.technologyreview.com/s/614929/facebook-has-a-neural-network-that-can-do-advanced-math/

$$y' = \frac{16x^3 - 42x^2 + 2x}{(-16x^8 + 112x^7 - 204x^6 + 28x^5 - x^4 + 1)^{1/2}}$$

- **AlphaGo and AlphaGo Zero to play Go**
  - AlphaGo **defeated** 18-time world champion Lee Sedol 4 games to 1
    - Used game tree search, neural network trained on expert human games, second neural network for board positions, and additional Monte Carlo rules
  - AlphaGo Zero used same tree search algorithm, but then single neural network trained without any human games
    - AlphaGo Zero defeated AlphaGo **100 games to 0**
    - https://medium.com/ww-engineering/alphago-zero-a-brief-summary-dcff16ba3064

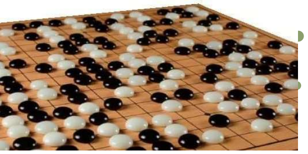

How can these approaches help future risk-informed applications?

- **Recent nuclear power challenges have been mostly on economics and safety**
  - Need to provide new **cost-beneficial approaches to safety** via modern methods/tools/data
  - We want to **attract the next generation** of scientists/engineers via these new approaches
- **Computational Risk Assessment (CRA) is a combination of**
  - **Probabilistic** (i.e., dynamic) scenarios where they unfold and are not defined a priori
  - **Mechanistic** analysis representing physics of the unfolding scenarios
- **Idea → CRA to produce "synthetic data" for ML**
  - ML requires training data – however risk & reliability have a small set of "failure" data
  - CRA can explore rich space of normal & off-normal conditions
  - CRA can produce very large sets of synthetic data
- **Idea → Digital regulator**
  - Agent-based systems for oversight of operations
  - CRA + real-world sensors → next-gen regulation
    - Keep an independent, digital presence in systems

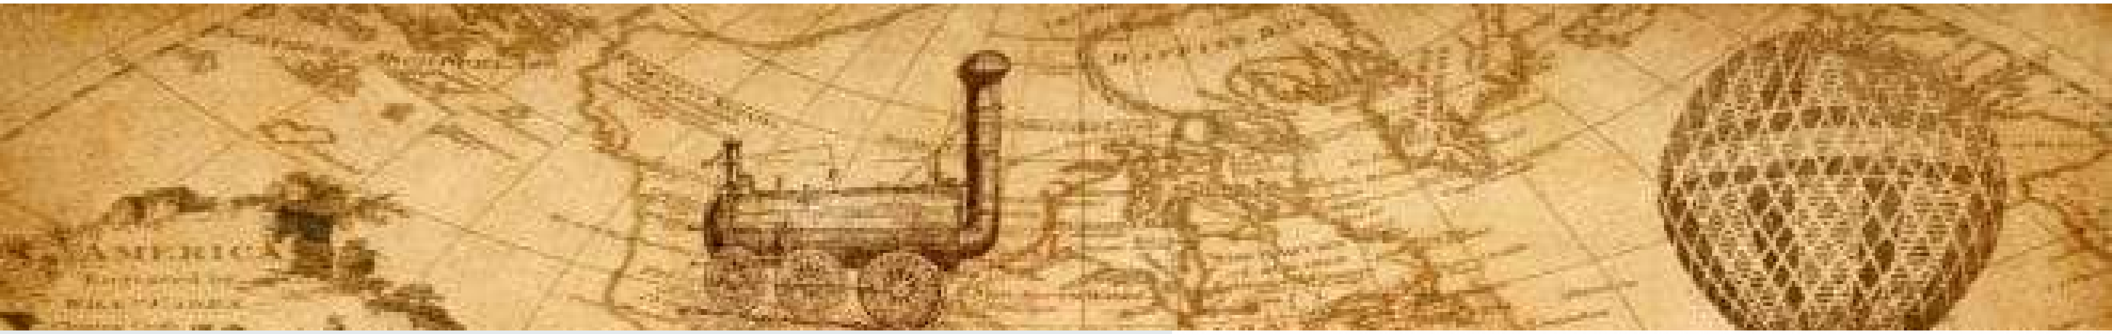

"And I told him, AI and ML aren't the thing.

They're the thing that gets us to the thing."

(See *Halt and Catch Fire*)

Learning Internal Representations by Error Propagation

DAVID E. RUMELHART, GEOFFREY E. HINTON, and RONALD J. WILLIAMS

THE PROBLEM

We now have a rather good understanding of simple two-layer associative networks in which a set of input patterns arriving at an input layer are mapped directly to a set of output patterns at an output layer. Such networks have no *hidden* units. They involve only *input* and *output* units. In these cases there is no *internal representation*. The coding provided by the external world must suffice. These networks have proved useful in a wide variety of applications (cf. Chapters 2, 17, and 18). Perhaps the essential character of such networks is that they map similar input patterns to similar output patterns. This is what allows these networks to make reasonable generalizations and perform reasonably on patterns that have never before been presented. The similarity of patterns in a PDP system is determined by their overlap. The overlap in such networks is determined outside the learning system itself—by whatever produces the patterns.

The constraint that similar input patterns lead to similar outputs can lead to an inability of the system to learn certain mappings from input to output. Whenever the representation provided by the outside world is such that the similarity structure of the input and output patterns are very different, a network without internal representations (i.e., a network without hidden units) will be unable to perform the necessary mappings. A classic example of this case is the *exclusive-or* (XOR) problem illustrated in Table 1. Here we see that those patterns which overlap least are supposed to generate identical output values. This problem and many others like it cannot be performed by networks without hidden units with which to create

Idaho National Laboratory

Curtis.Smith@inl.gov

Thank you!

Ross Kunz

Group: **Advanced Analytics**
Education: **PhD Statistics**
Work focused in: **Machine learning for chemistry and physics (catalysts, batteries, materials)**

Presentation Overview
What is AI?

- **Overview of AI and the connection to Modeling/Simulation**
- **Understanding of complex data sets and discovery of new information**

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

www.inl.gov

Idaho National Laboratory

Ross Kunz
B652 Advanced Analytics
What is AI?

Definition

- The capability of a machine to imitate intelligent human behavior

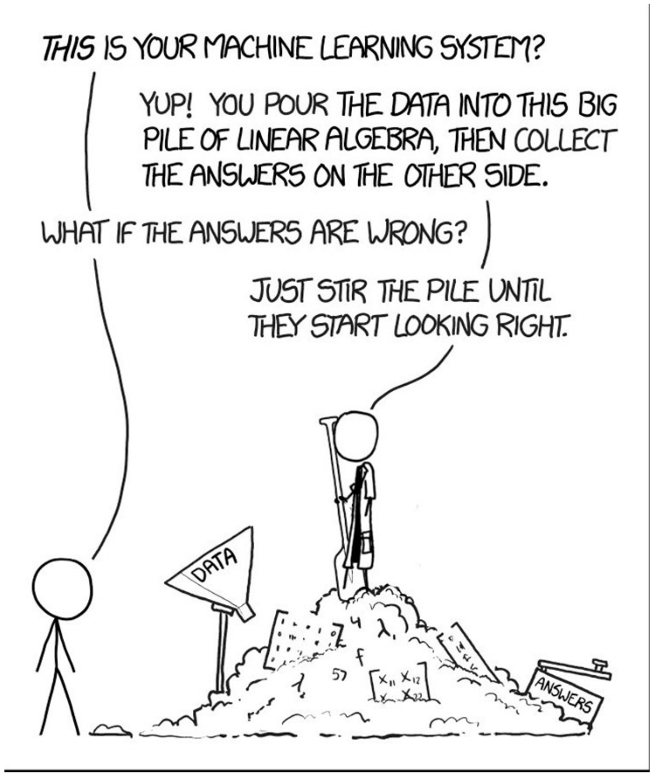

Source: xkcd.com

1. Data (kind of a big deal)
    1. Good
    2. Bad
    3. Ugly
2. Domain problem
    1. Data Structures
    2. What information can be leveraged
    3. No free lunch!
3. Results
    1. I don't care, predict the cat!
    2. The journey, not the destination that matters

Connection to Science

Data Analysis Spectrum

- Little to No Data
- Strong Assumptions
- Highly Informative
- High Computation

Physics Based Modeling

Traditional Statistics

Machine Learning

Artificial Intelligence

- Extreme Amounts of Data
- Little to No Assumptions
- Highly Predictive
- High Computation

Physics to physics

Surrogate modeling

Experimental Discovery

Types of Problems

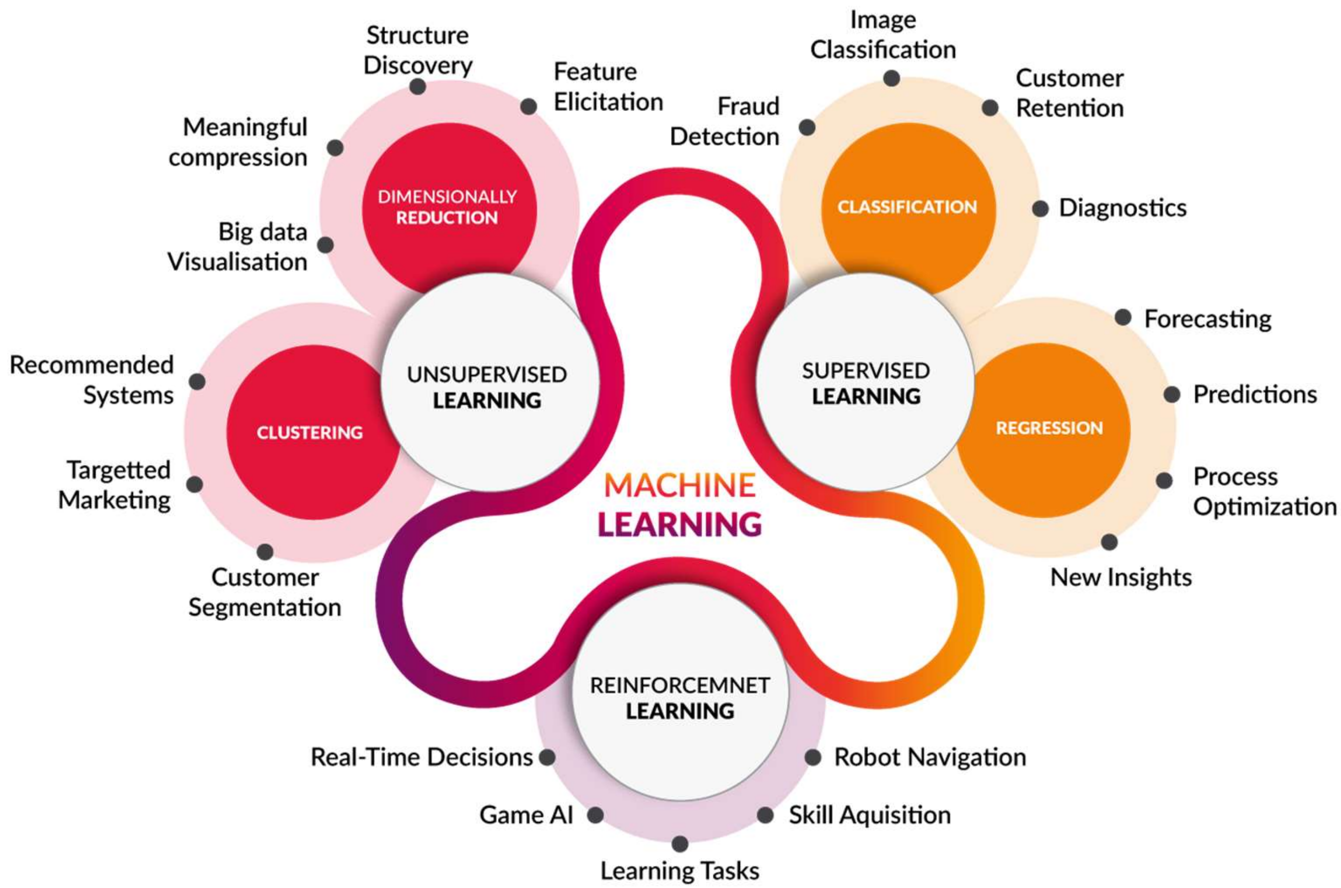

Source: http://www.cognub.com/index.php/cognitive-platform/

Explainable AI

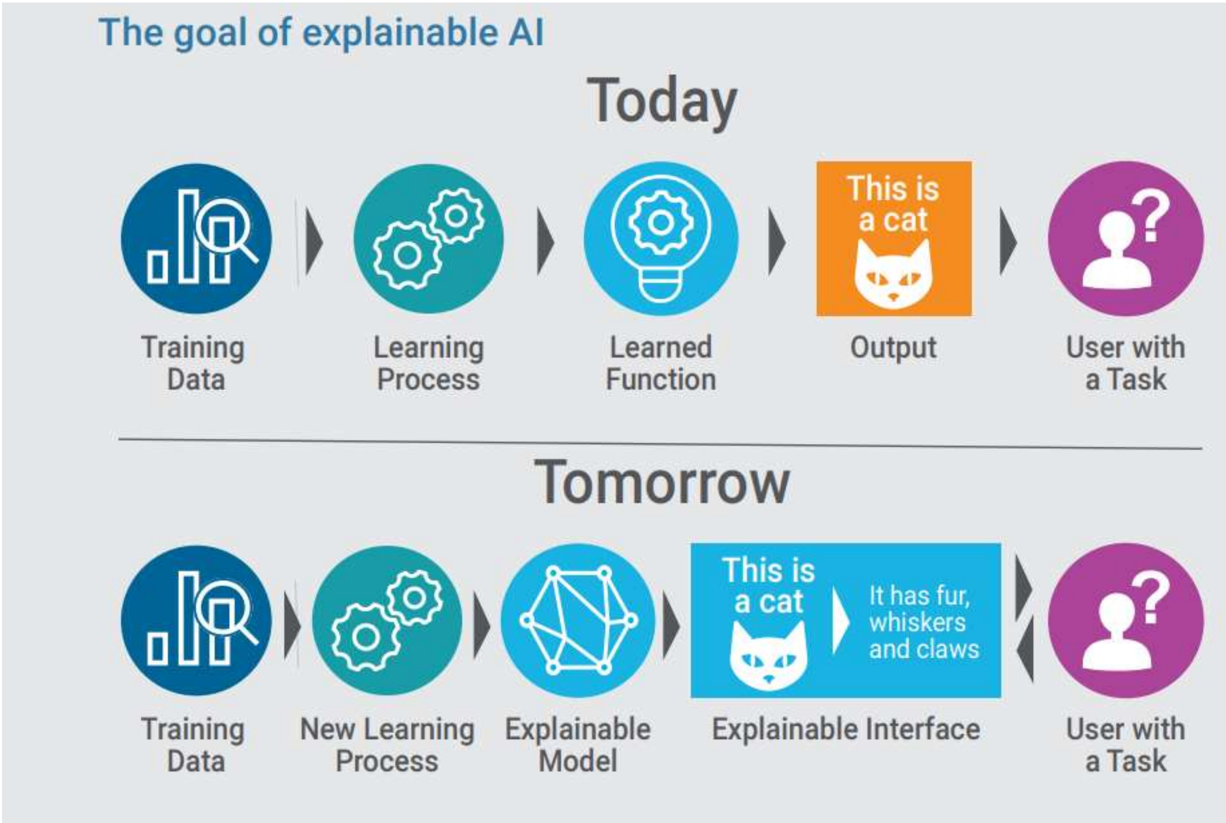

Source: AI and Machine Learning: Key FICO Innovations

Example Projects

TAP reactor catalysis machine learning

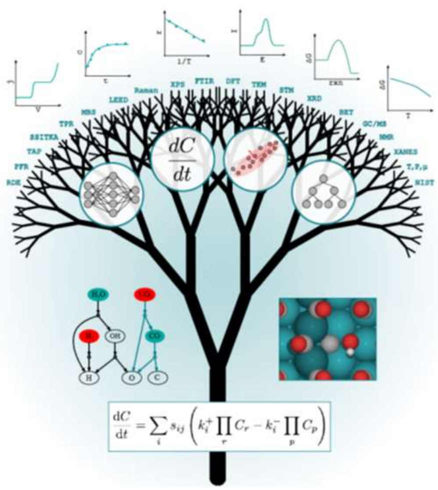

Medford et al. Extracting knowledge from data through catalysis informatics. 2018

Rebecca Fushimi
Ross Kunz
Yixiao Wang
Zongtang Fang
Rakesh Batchu
Sagar Sourav
James Pittman

Battery life prediction / mechanism estimation

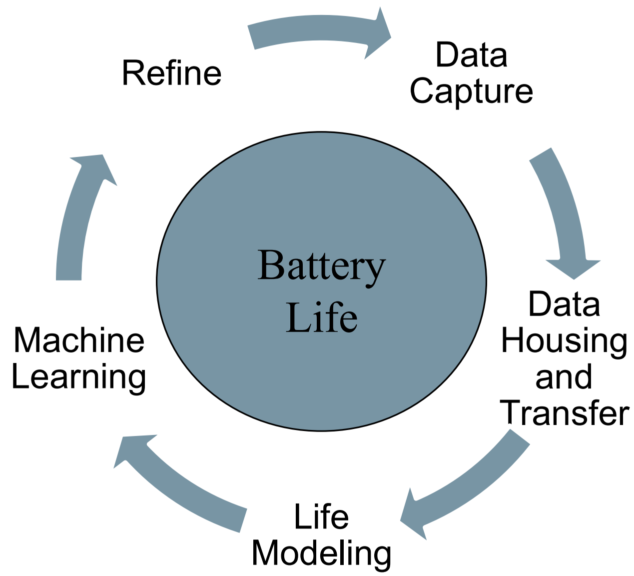

Eric Dufek
Ross Kunz
Zonggen Yi
Matt Shirk
Kevin Gering
Hypo Chen
Tanvir Tanim
Dave Black
Qiang Wang

Kandler Smith
Paul Gasper

Questions?

Clean. **Reliable. Nuclear.**

Nancy Lybeck

Group: Department Manager, Instrumentation, Controls, & Data Science
Education: Ph.D. in Math from Montana State University. Fifteen-plus years working with data; 10 at INL
Work focused in: Several projects, including developing a Risk-Informed Predictive Maintenance Strategy and the Nuclear Data Management and Analysis System

Presentation Overview

Artificial Intelligence, Machine Learning, and Statistics, Oh My!

- A light-hearted look at the perceived rivalry between data science and statistics.

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

www.inl.gov

Idaho National Laboratory

Nancy Lybeck, PhD

Instrumentation, Controls, & Data Science

AI, ML, and Statistics, Oh My!

We all love a great rivalry!

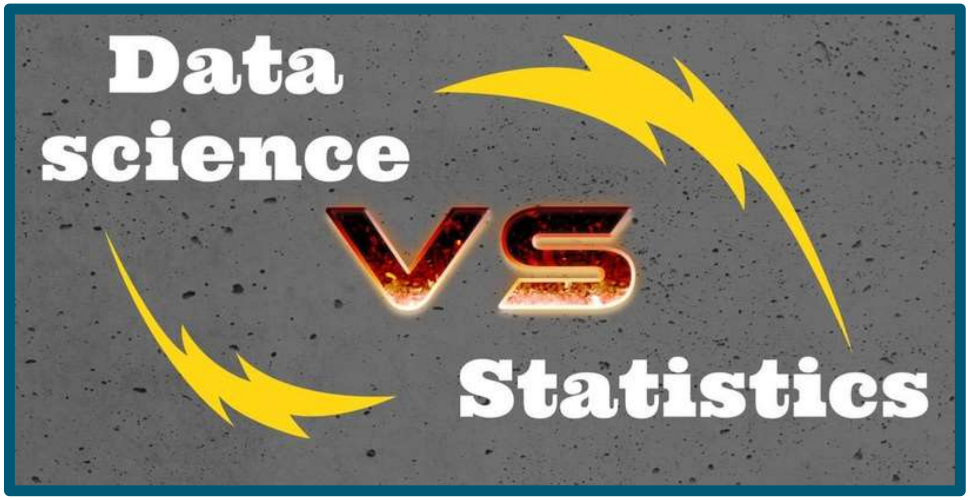

What is Data Science?

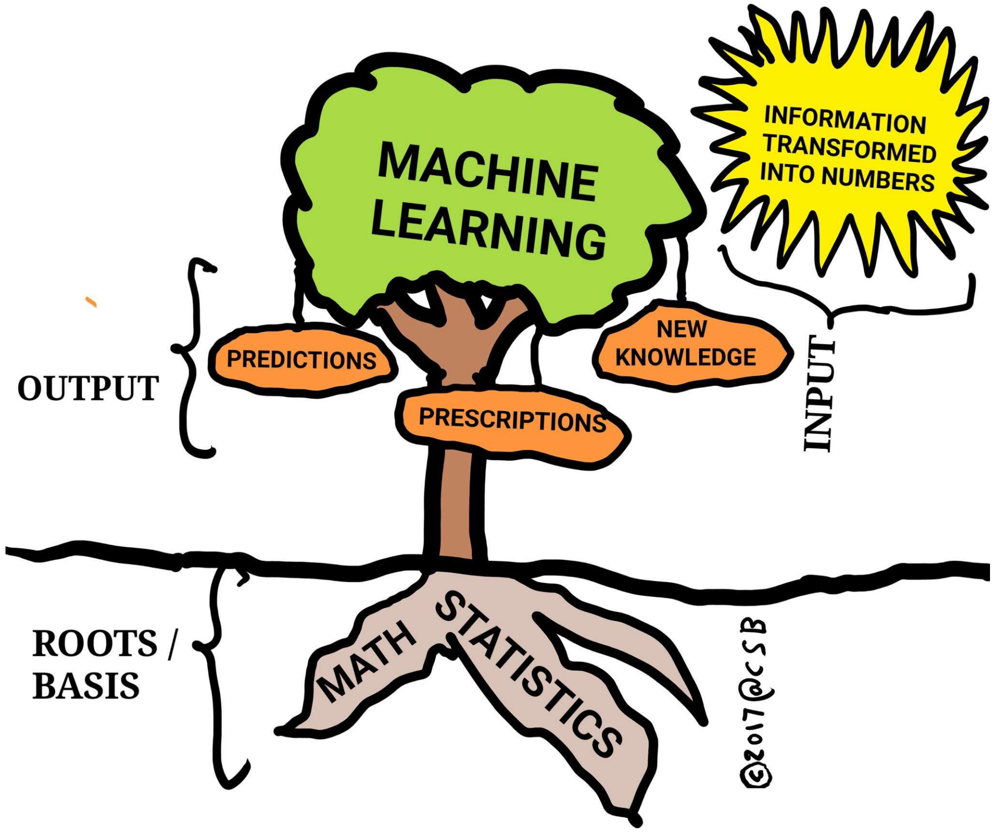

Comic Strip Blogger December 2017

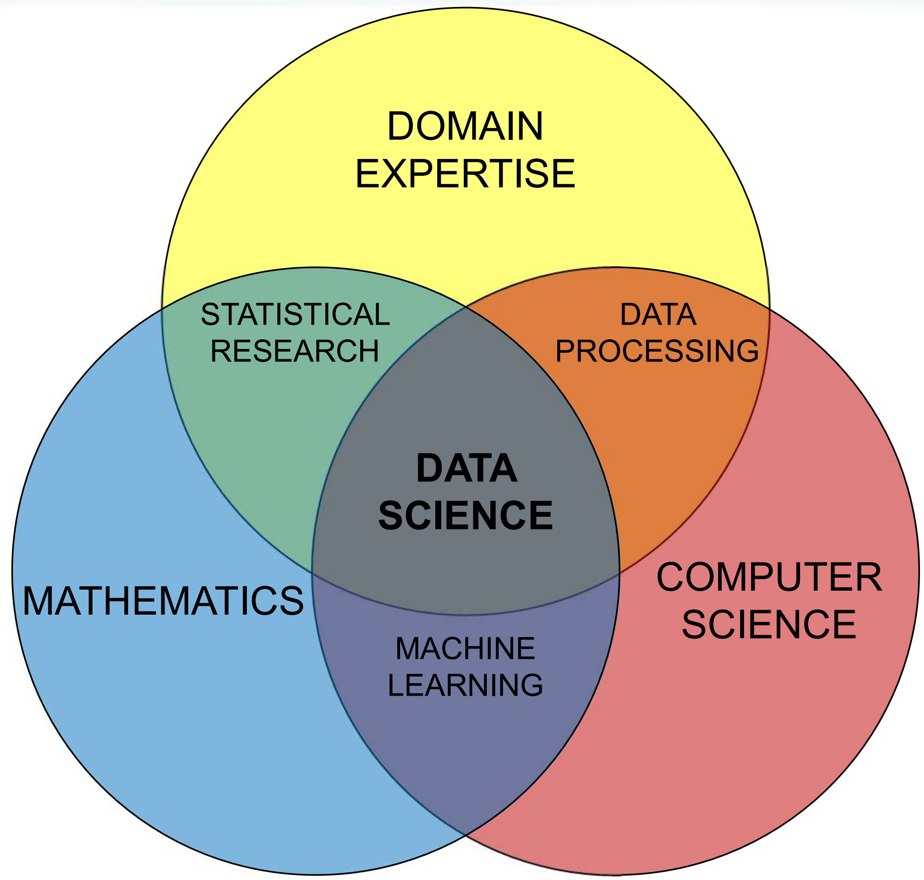

Source: Palmer, Shelly. Data Science for the C-Suite. New York: Digital Living Press, 2015. Print.

Discussion

Statistics	Machine Learning
• Focus on Inference • Based on probability spaces • Creating and fitting project-specific probability models • Often used with tall data • Formalizes understanding of system behavior • Tests a hypothesis about system behavior • Computes a quantitative measure of confidence that a discovered relationship describes a 'true' effect that is unlikely to result from noise • Generally considered interpretable	• Focus on Prediction • Based on statistical learning theory • Using general-purpose learning algorithms to find patterns in often rich and unwieldy (nonlinear) data • Particularly helpful with wide data • Makes minimal assumptions about the system • Does not require a carefully controlled experimental design • Accuracy determined with test data set (in the case of supervised learning) • Can be difficult to interpret

Example from Environmental Science: We might use a statistical model to determine whether a sensor signal response to a certain kind of stimuli is statistically significant, as well as use data from an array of 20 additional sensors to predict the response of the sensor.

The Actual Difference Between Statistics and Machine Learning, Matthew Stewart, 2019.
Statistics Versus Machine Learning, Bzdok et al., Nature Methods 15, 223-234 (2018).

Looking Ahead

- **It’s all about the data …**
- **We need statisticians and data scientists!**
- **Hold on to the rivalry for fun and for lighthearted teasing, but don’t let it get in the way of our ultimate goal: doing great science!**

Questions?

Nancy.Lybeck@inl.gov
(208) 206-7232

Thank You!

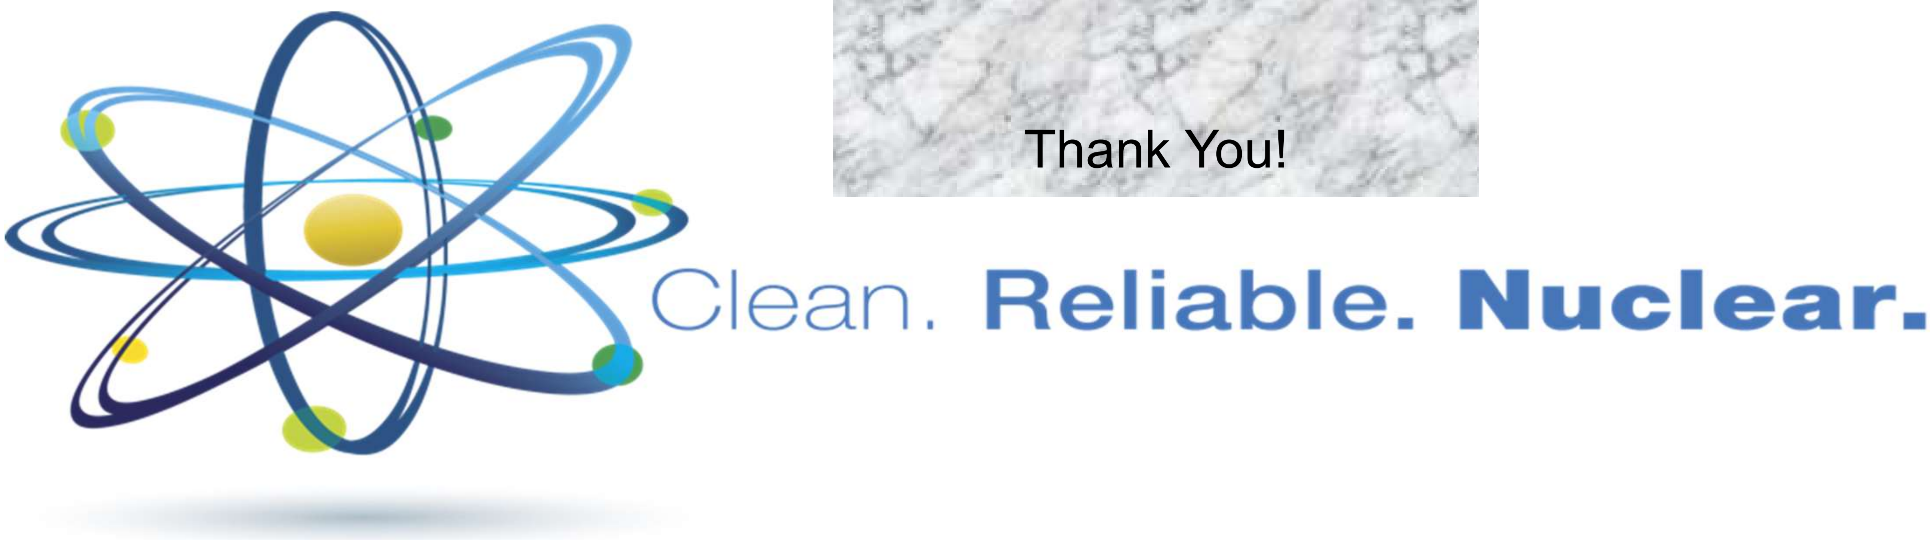

Ronald L. Boring

Group: Department Manager, Human Factors and Reliability
Education: Ph.D. in Cognitive Science from Carleton University
Work focused in: Human factors and human reliability

Presentation Overview

Modeling Human Cognition: It's Not All Machine Learning

- While AI is widely used for industry applications, one of its first uses was to mimic human cognition. The earliest AI techniques were rule based to try to capture the psychology behind human decision making.

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

www.inl.gov

Idaho National Laboratory

Ronald Laurids Boring, PhD

Human Factors and Reliability Dept.

Modeling Human Cognition:
It's Not All Machine Learning

Why Human Cognition?

1956 Was Watershed Year

- **Nuclear History**
  - Period between USS Nautilus and Shippingport
- **Two Congressional Hearings on Automation**
- **Dartmouth Summer Workshop on Artificial Intelligence**
  - "We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."
  - Birth of AI, featuring founders like Marvin Minsky, John McCarthy, Claude Shannon, Allen Newell, and Herb Simon
- **Symposium on Information Theory at MIT on September 11, 1956**
  - Birthplace of information processing theory and study of cognition
  - Featured George Miller, Noam Chomsky, Allen Newell, and Herb Simon, among others
- **Birth of AI and cognitive psychology occurred at the same time, because they were interested in the same problems**
  - Deconstructing human thinking into information allowed us to make computer models of it

Why Human Cognition?

1956 Was Watershed Year

- **Nuclear History**
  - Period between USS Nautilus and Shipp
- **Two Congressional Hearings on Automatio**
- **Dartmouth Summer Workshop on Artificial**
  - "We propose that a 2-month, 10-man stu summer of 1956 at Dartmouth College in on the basis of the conjecture that every can in principle be so precisely describec
  - Birth of AI, featuring founders like Marvin Newell, and Herb Simon

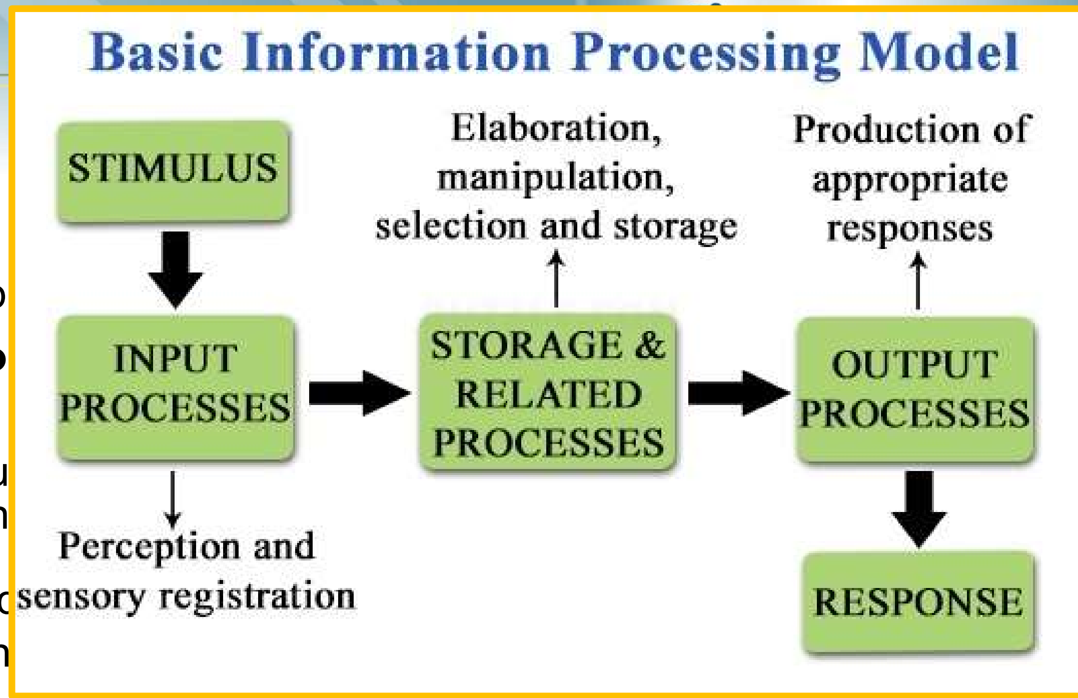

- **Symposium on Information Theory at MIT on Sept 11, 1956**
  - Birthplace of information processing theory and study of cognition
  - Featured George Miller, Noam Chomsky, Allen Newell, and Herb Simon, among others
- **Birth of AI and cognitive psychology occurred at the same time, because they were interested in the same problems**
  - Deconstructing human thinking into information allowed us to make computer models of it

AI is More Than Machine Learning

Two Types of AI

- **Good Old-Fashioned AI (GOFAI)**
  - Symbolic logic systems to represent basic elements of human thought like language, numbers, or goals
  - **Expert systems** featuring if-then logic
    - General Problem Solver created by Newell and Simon in 1959
  - Much of focus was not to create learning but to capture human-like intelligence
- **Neural Networks**
  - **Perceptron** developed in 1958 as approximation of single-cell neuron
  - By 1960s, mathematical algorithms like backpropagation developed to allow perceptrons to learn through training
    - **Machine learning**
  - Multiple perceptrons chained together to create neural networks
  - More layers of neural networks chained to together to create **deep learning**

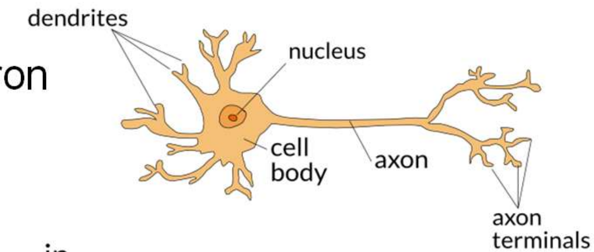

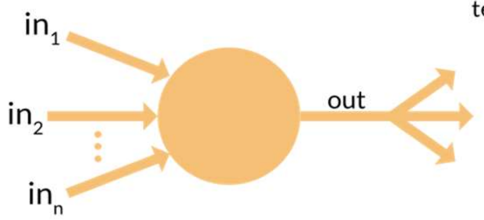

- **Different Uses**
  - GOFAI is good at following rules and making decisions
  - Neural networks are good at pattern recognition when trained

Why is Human Cognition Relevant to AI?

Humans Are Better At-Machines Are Better At (HABA-MABA)

- **Humans are (still) better at some things**
  - Generalization and flexibility
  - Judgement and decision making
  - Responding to novel events and degraded conditions
  - Creativity and problem solving
  - Sentience and consciousness
- **Machines are better at some things**
  - Performing routine, repetitive, or precise tasks like monitoring
  - Multitasking
  - Quick responses

What Are the Goals of AI?

- **Narrow AI**
  - Perform a simple task, like automating a safety valve
  - These are simplistic tasks that don’t need to be human-like to be successful
- **General AI**
  - Perform the task of a human like replacing a control room operator or driving a car
  - These are complex tasks that aspire to human cognition

Principles for the Intersection of Humans and AI

1. **AI = Knowledge + Learning**
   - To say someone is intelligent does not mean they are good learners, it means that they are knowledgeable
   - AI is a mix of GOFAI (knowledge) and neural networks (learning)
     - It takes both to create something like autonomous vehicles: see the road + follow the rules
2. **Machine Learning Has Limits**
   - We think of ML as producing superintelligence, but most applications are really narrow AI
3. **Humans are the Users of AI**
   - Sometimes we seek not to replace the human but enhance or complement them (e.g., predictive maintenance)
   - Need to develop **explainable AI** that humans can understand and work with
     - How does regulator approve AI for safety applications like nuclear when AI isn't transparent in what it's doing?
   - **Data visualization**—representing patterns out of complexity—is one form of usable AI
4. **Humans are Big Data**
   - Human performance and knowledge can still be harvested to improve AI

Questions?

Ron Boring, PhD
Manager & Distinguished Scientist
Human Factors & Reliability Department
Nuclear Safety and Regulatory Research Division
Idaho National Laboratory

ronald.boring@inl.gov

Clean. **Reliable. Nuclear.**

Humberto E. Garcia

Group: Systems Science & Engineering
Education: PhD
Work focused in: Extensive experience in advanced systems methods for the design, integration, optimization, and operation of cyber-physical systems (CPS)

Presentation Overview

Secure Embedded Intelligence (SEI) in Smart Nuclear Systems

- Research needed / Gaps for implementing SEI in Smart Reactor Systems

Secure Embedded Intelligence (SEI) in Smart Reactors

Topics: multi-scale, multi-layered computing, hybrid physics-based, data-driven M&S, digital twins (DT), integrated state awareness (ISA), adaptive observation & actuation, intelligent controls, automated reasoning, digital assets

Humberto E. Garcia, PhD
Cyber-Physical Systems Integration, Optimization & Resilient Controls

Advanced sensors + Digital Twins + Antifragile Capabilities + Agile Optimization + Security by Design = Smart Reactors

Within a multi-scale, multi-layered (distributed) architecture

INL Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Related reading:

- H.E. Garcia, S.E. Aumeier, A.Y. Al-Rashdan (2020). "Integrated State Awareness Through Secure Embedded Intelligence in Nuclear Systems: Opportunities and Implications," *Nuclear Science and Engineering*, Vol. 194, pp. 249-269, April 2020.
- H.E. Garcia, S.E. Aumeier, A.Y. Al-Rashdan, B.L. Rolston (2020). "Secure Embedded Intelligence in Nuclear Systems: Framework and Methods," *Annals of Nuclear Energy*, Vol. 140, 2020, 107261.

Why it is important to industry

- Operations & maintenance (O&M) cost reduction & simplification
  - Economics (e.g., 15 - 50%+ fixed O&**M** cost reduction)
  - Real-time asset condition assessment
    - from *preventive to predictive*
    - Predictive maintenance (PdM), proactive asset performance/health management (APM)
    - Early anomaly/health detection, diagnostics & prognostic of systems, structures, components (SSC)
  - Improved reliability, availability, maintainability, safety, security
- Market expansion, application flexibility, nuclear industry sustainability
  - Flexible operation
  - Remote and transportable deployments
  - Broad range of "plug-and-play" (commercial and emergency) applications
- Design and operations margin reduction and optimization
  - Simplicity and uncertainty & imprecision tolerance
- Unprecedented system-state knowledge enabling:
  - Adaptive control (e.g., idle, startup, shutdown), automated reasoning, decision-making
  - Recognition & classification of abnormal and degradation signatures
  - Inherent, proactive cybersecurity and cyber-defense *by design*
- Real-time metric (e.g., risk) quantification, optimization, management
- Human reliability and productivity enhancement
  - Integrated, precision data availability and presentation / visualization

Current trends in diverse industries

Vehicles w/ limited "automated processing"

Autonomous "smart" vehicles

"Labor-intensive" manufacturing

Autonomous "smart" manufacturing

Is autonomy of smart reactors the goal ? or rather to identify fundamental attributes a system should be equipped with to meet desired (smart) functionalities (e.g., autonomy) ?

ISA + **SEI**

Knowledge **Reasoning**

To achieve "smart" functionalities
(e.g., autonomy, asset health assessment)

SEI: Secure embedded intelligence

ISA: Integrated state awareness

Design for optimal levels of ISA & SEI to achieve objectives

Phased implementation of SEI-ISA in advanced nuclear systems

Add fundamental capabilities

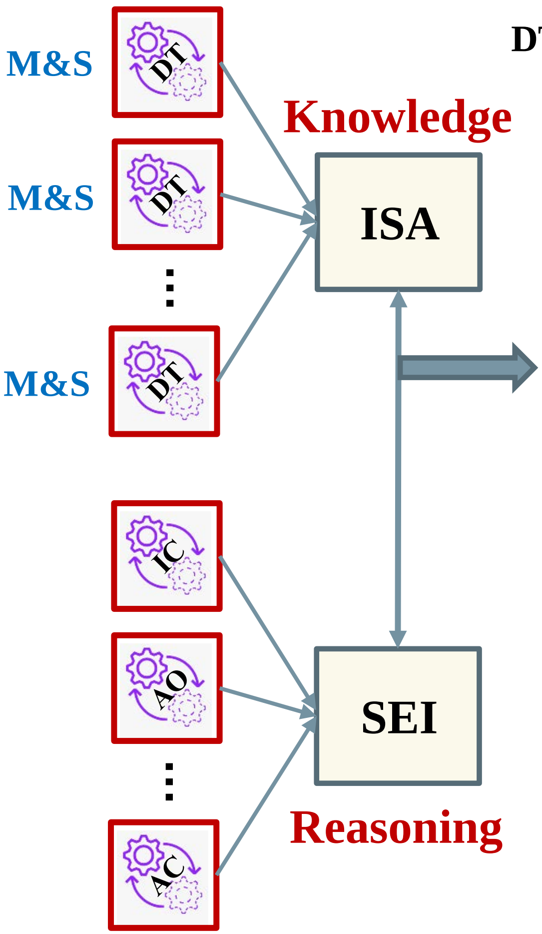

DT: **Nested** Digital Twin (*model-based + data-driven, multiscale, multilayered*)

To achieve fundamental functionalities

- **Estimate** (e.g., current system state)
- **Predict** (e.g., future system state)
- **Understand** (e.g., consequences of stressors, actions)
- **Learn** (e.g., relationships from observed patterns)
- **Decide "optimal" paths forward** (e.g., control actions)

Disruptive advances

- multi-scale / multi-layered computing (HPC & edge computing)
- physics-based, data-driven hybrid M&S and analysis
- multi-layered adaptive observation & actuation
- intelligent controls (**IC**) & supervision
- agile optimization (**AO**)
- AI-enhanced capabilities (**AC**)

Disruptive potentials

- ✓ *Cost*
- ✓ *Simplicity*
- ✓ *Flexibility*
- ✓ *Systems optimization*
- ✓ *Inherent security, resiliency*
- ✓ *System-state transparency*

Intelligent nuclear assets: *Multi-scale, multi-layered integration of advanced monitoring, control & supervision (MCS) functions*

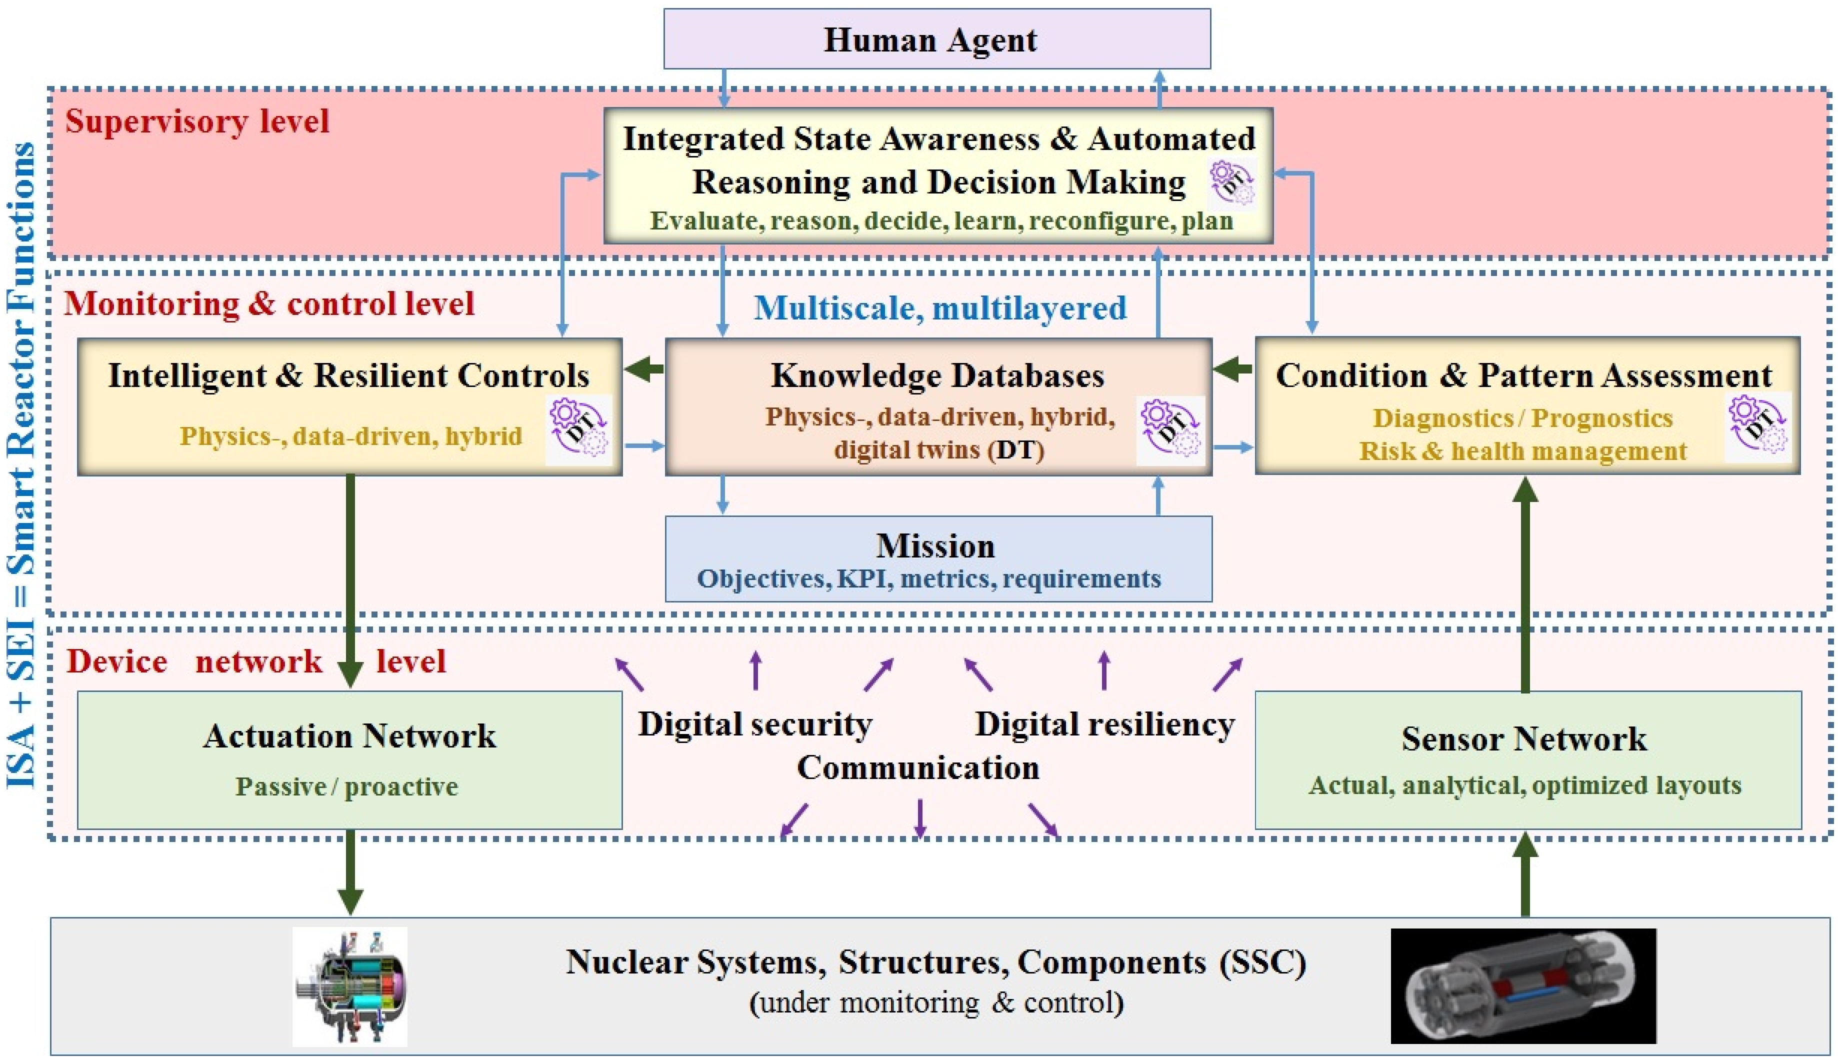

Implications for the nuclear industry

Modify

SEI/ISA Assets

Physical Assets

Modify

ISA – SEI ***support***

Knowledge – Reasoning

Smart functionalities / Advanced Outcomes

- autonomy
- self-optimization
- self-configuration
- flexible operation
- self-maintenance
- self-protection
- self-validation
- self-healing
- learning & explanation

SEI: Secure embedded intelligence

ISA: Integrated state awareness

Research opportunities for implementing SEI in smart reactor systems

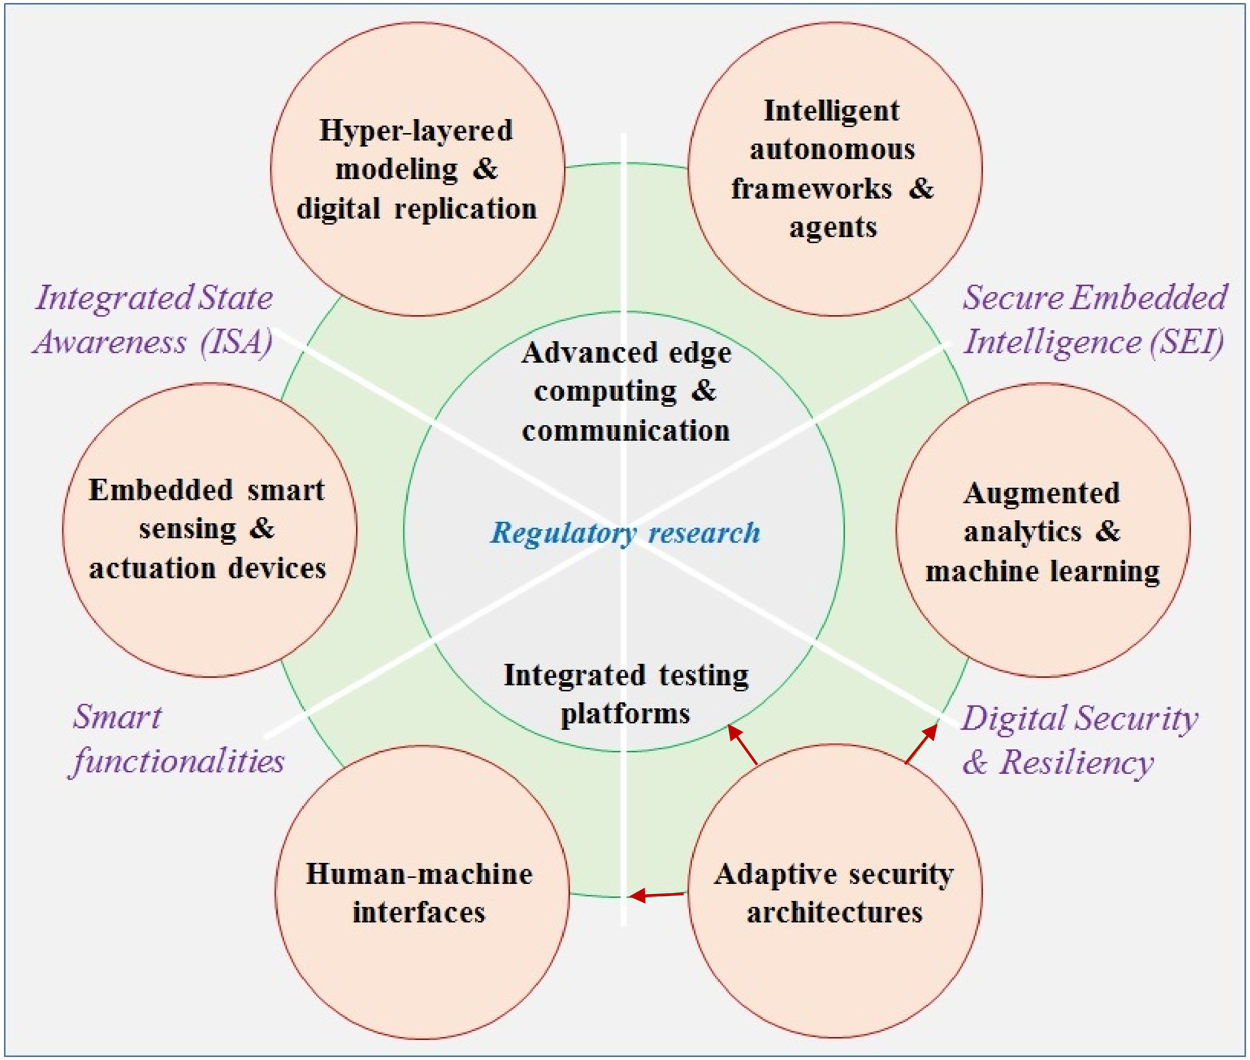

Products

- Architectures
- Frameworks
- Information infrastructures
- (edge-, system-) methods, models, agents, algorithms
- Hardware / software capabilities and devices
- Design impacts
- Testbeds
- Pilots
- Standards
- Policies

Questions?

Clean. Reliable. Nuclear.

Victor G. Walker

Group: **Mobility Systems and Analytics**
Education: **B.S. and M.S. degrees in Computer Science with a focus on intelligent and adaptive systems and worked for 11 years at IBM before joining INL**

Presentation Overview
AI in Robotics and Applying Natural Connections

- **AI in Robotics has some unique characteristics. It involves an intelligent system that interacts with the real world and these issues can influence both how a system learns and what we expect from the systems. A key goal is creating a system that allows us to use robotics as a natural partner.**

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Victor Walker

Advanced Transportation

AI in Robotics and Applying Natural Connections

Idaho National Laboratory

Robotics and Intelligence (Introduction)

<u>Robotics</u>

(What is it?)

Computation

Mobility

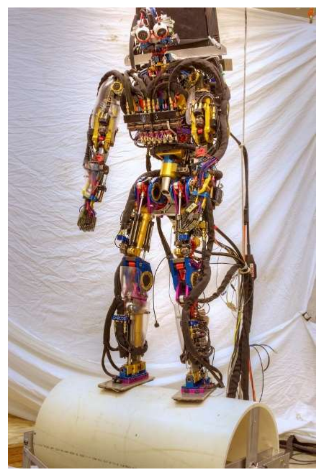

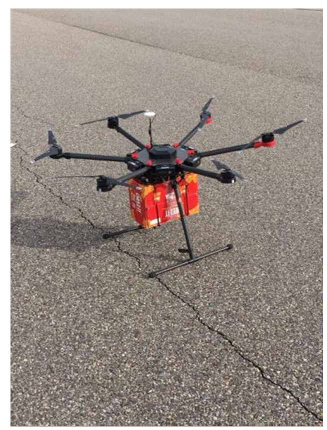

Humanoid Robots

Unmanned Aerial Vehicles (UAV)

Unmanned Ground Vehicles (UGV)

Self-Driving Cars

Robotic Arms

Robotics and Intelligence (Introduction)

Intelligence

(What is it in Robotics?)

- Behavior-based
- Does it “Act” intelligently?
- Does it do intelligent tasks?
- Does it partner well?

Needs:

- Sensors
- Tasks
- Training

Intelligent Robotics enables a brave new world….

Robotics enables a broad range of tasks

- **Dangerous**
- **Precision**
- **Repeatable**
- **Dull**
- **Efficient**
- **Remote**

Intelligence enhances Partnership

- **Partner with humans on tasks.**
- **Change the world… based on location**
- **Understand environment / Aid decisions**

Look for Natural Connections

Creating Intelligent Robotics

Key Barrier: **TRUST**

Ability to predict behavior

Explainable AI is often critical

Robotics Often Rules-based

- **Enable with Training**
- **Reinforcement learning**
- **Understanding enables acceptance**
- **Support Co-Robotics**

Often simple rules for complex tasks

INL is a champion of Adaptive Intelligence

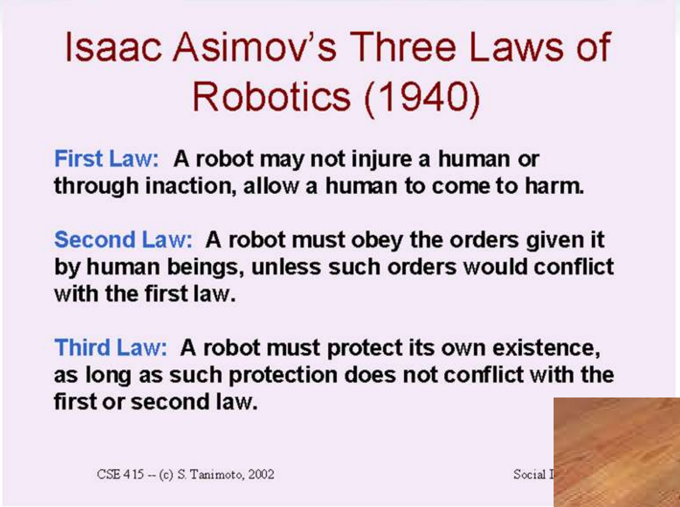

Isaac Asimov's Three Laws of Robotics (1940)

First Law: A robot may not injure a human or through inaction, allow a human to come to harm.

Second Law: A robot must obey the orders given it by human beings, unless such orders would conflict with the first law.

Third Law: A robot must protect its own existence, as long as such protection does not conflict with the first or second law.

CSE 415 -- (c) S. Tanimoto, 2002

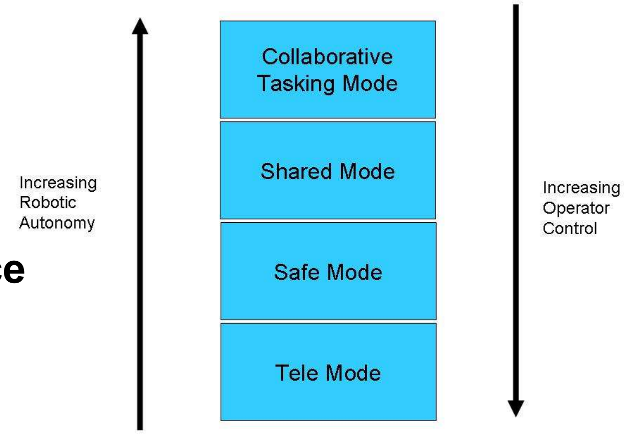

Creating Intelligent Robotics

Need ongoing research to improve robotics

- **Move from tool to partner**

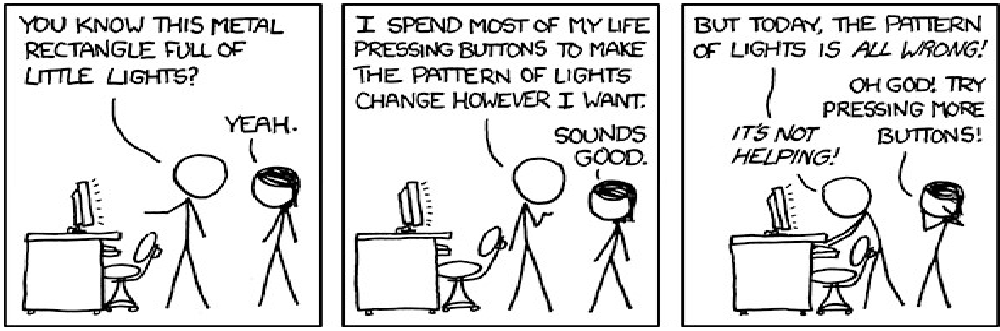

Look for Natural Connections for Human Interaction

- **Task-Level Execution**
- **Focus on shared Goals / Best ability**
- **Look for Natural Seams / Shared Cognition**

Research into more Natural Intelligence and Interaction

- **Research in Narrative-Based Intelligence**
- **Narratives part of Intelligence**
- **Conclusions and Framework modelling**

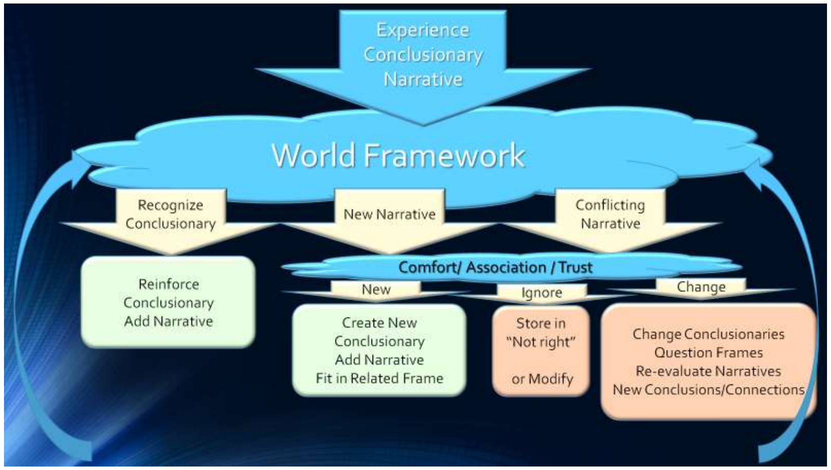

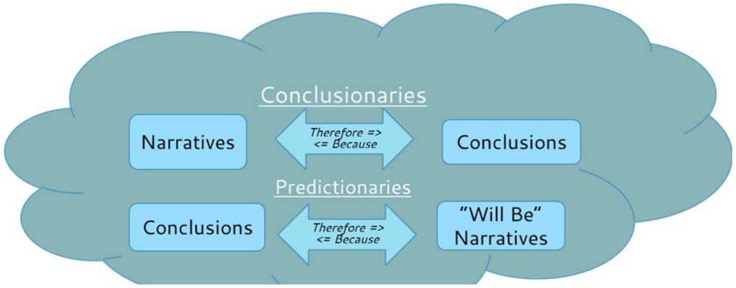

Robotics Looking Ahead

Some of INL Robotics:

- **Robotics Intelligence Kernel (RIK)**
  - **Counter-Mine**
  - **DOD Support**
  - **Fukushima**
  - **Tunnel Mapping**
  - **UAV work**

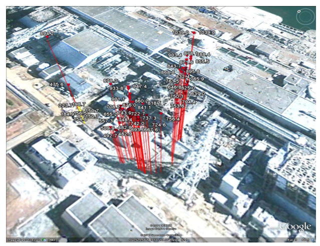

Yucca Mountain

Remote handling

Welding

Recovery

NHS Support

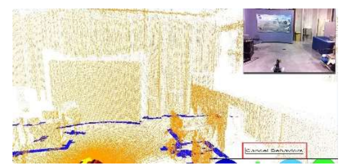

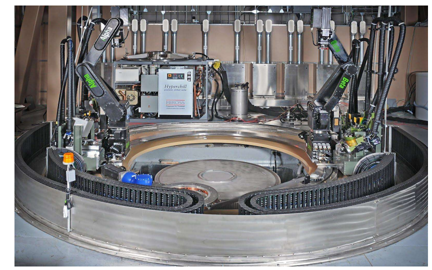

Current/Future Research:

- **Hot Cell**
- **Mobile Hot Cell**
- **UAV work**
- **Autonomous Vehicle Impacts**
- **Fleet AI (Caldera)**

Intelligence Development

Improved partnering

Enabling New Abilities

Questions?

Clean. **Reliable. Nuclear.**

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Katya Le Blanc

Topic Introduction

- Automation as AI, or AI and Automation, or AI as Automation
  - Discuss how AI can be used in automation
  - Discuss how some existing automation, is in a sense, AI
  - Discuss how we can enhance automation with AI, including machine learning
  - Discuss the strengths and weaknesses of AI in the context of automation
- Why it is relevant to ML/AI Future
  - There is great opportunity in using AI in automation
  - There is also great peril if we implement it poorly, especially if we don't fully understand the limitations and constraints

Types of AI

- Expert Systems
  - Draws from human expertise to automate a task
  - Typically replicates how a human would do a task
  - Can help us automate tasks that humans currently do
- Machine Learning
  - Perceptual Classification
- Neither approach does what humans do well, which is to develop abstract representations that we can use to generalize

Expert Systems

- Draws on expertise from multiple human experts
- More consistent than humans performing the same task
- Can be more accurate than humans, especially when human experts can supervise and update expert system with new information
- Brittle and doesn't adapt well to unforeseen situations
- Lacks insight and ability to generalize
- Many modern control systems could be classified as AI
  - Draw from experts in engineering and operations and from previous experience
- Typically understandable to humans
  - Depends on how systems present info

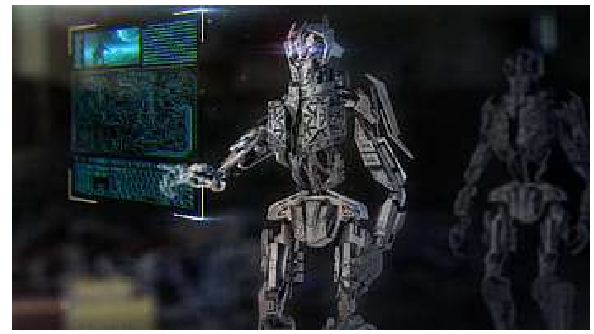

Machine Learning

- Works extremely well for well-defined classification problems
- Needs lots of data
  - In contrast, humans can learn to classify with 1 example (and abstract reasoning)
  - Babies learn with just a few examples
- Results depend on quality of data
  - Data is not inherently objective
  - Data is a human construct, we define what is collected, and what it means
  - Assumptions are embedded in the data
- It does exactly what we tell it to do….which can be a problem
- Typically opaque to humans

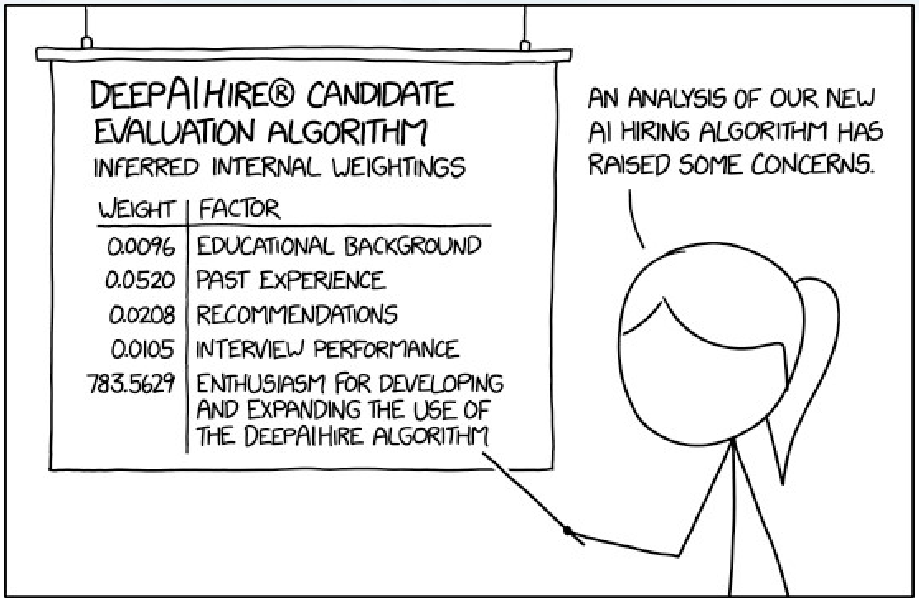

Current work and future work

- Developing expert systems to automate nuclear power plant operations (Light Water Reactor Sustainability (LWRS))
  - Drawing on documentation of how humans solve problems
    - Procedures
    - SMEs
      - Operators and engineers
    - Alarms and event logs
    - Other data sources
    - Data structure challenges
    - Can we use ML to classify valid versus nuisance alarms
    - Can we use ML to parse procedure text?
- Using ML and image processing for gesture recognition in AR application for NPP field workers (Technology Commercialization Fund (TCF) Proposal with Aguiar, Yoon,& Oxstrand)
- If we are building a system from scratch, what data should we collect and how should we structure it for maximum usefulness in some of these applications (NuScale and JUMP)

Questions?

Clean. **Reliable. Nuclear.**

Vivek Agarwal

Group: **Controls and Data Science Department within the Nuclear Safety and Regulatory Research Division**
Education: **B.E. degree in electrical engineering from the University of Madras, India, M.S. in electrical engineering from The University of Tennessee, Knoxville, and Ph.D. in nuclear engineering from Purdue University.**

Presentation Overview

Transition from Preventive to Predictive Maintenance Strategy

- **The presentation will present challenges current light water reactors are facing. How the research performed by INL in collaboration with nuclear plant owners, is providing a science-based approach to enable plant's transition from traditional labor-intensive, time- consuming preventive maintenance practice to predictive maintenance strategy.**

www.inl.gov

Idaho National Laboratory

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Vivek Agarwal, PhD
Instrumentation, Controls, and Data Science Department (C220)
Transition from Preventive to Predictive Maintenance Strategy

Diversity of Data

- To support operation and maintenance of a nuclear power plant
  - Data are collected at different **spatial** and **temporal** resolutions using different measurement techniques
  - Collected data are in different format and are stored in different systems.
- Majority of the data (if not all) are collected manually.

Transition to Preventive to Predictive Maintenance Strategy

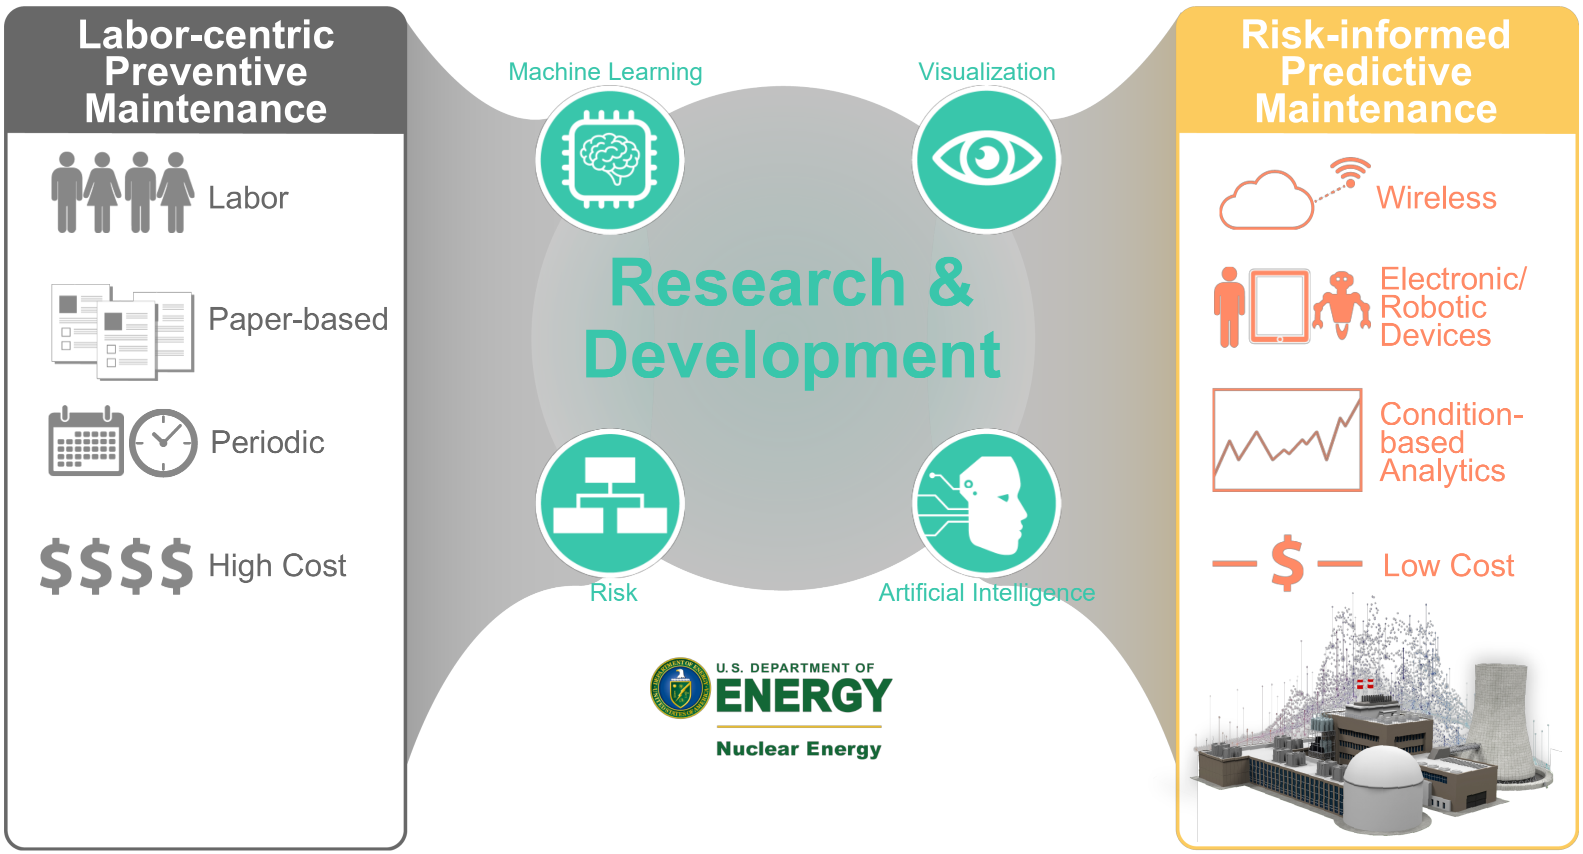

V. Agarwal et al., “Deployable Predictive Maintenance Strategy based on Models Developed to Monitor Circulating Water System at the Salem Nuclear Power Plant,” INL/LTD-19-55637, September 2019.

Transition from Preventive to Predictive Maintenance Strategy

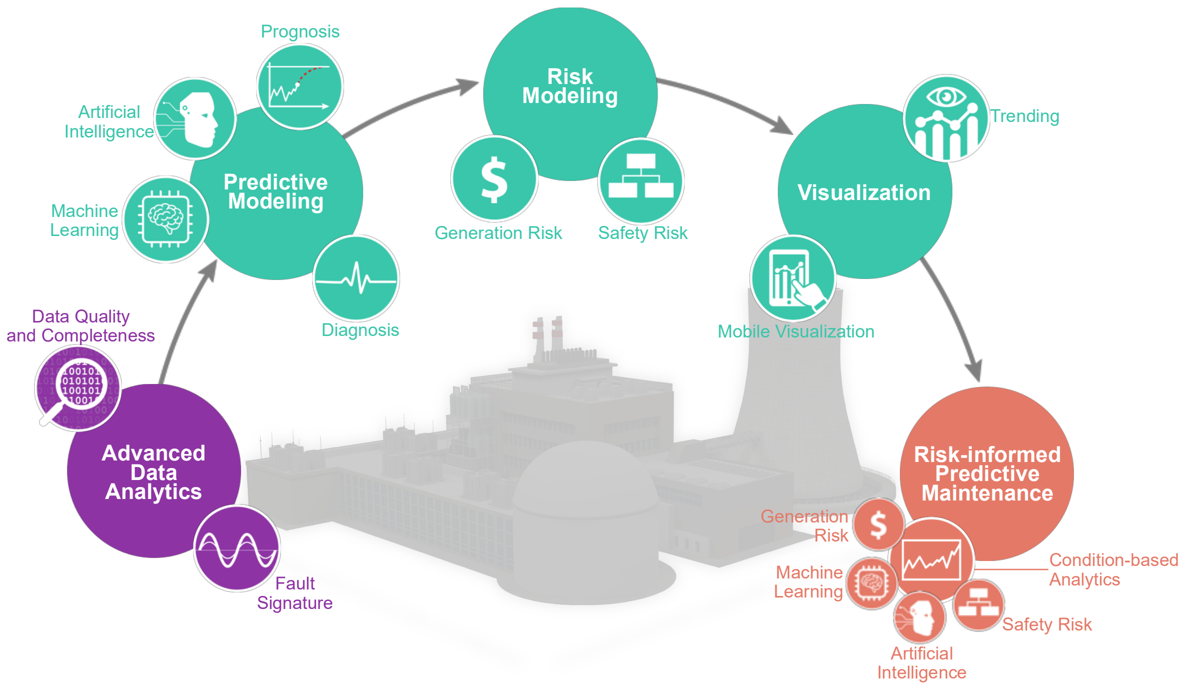

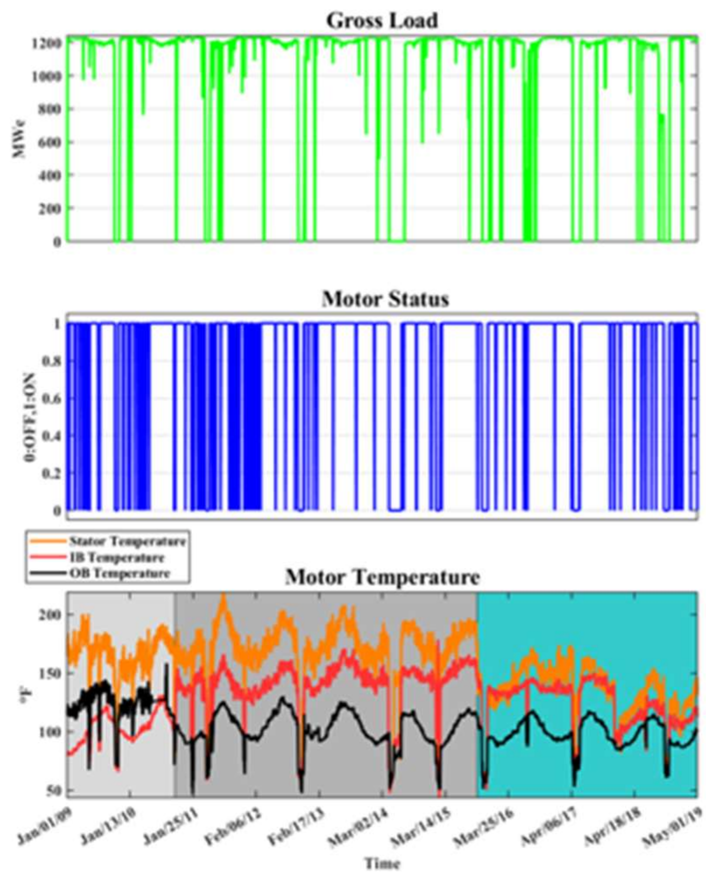

Path Forward

Scalability Analysis

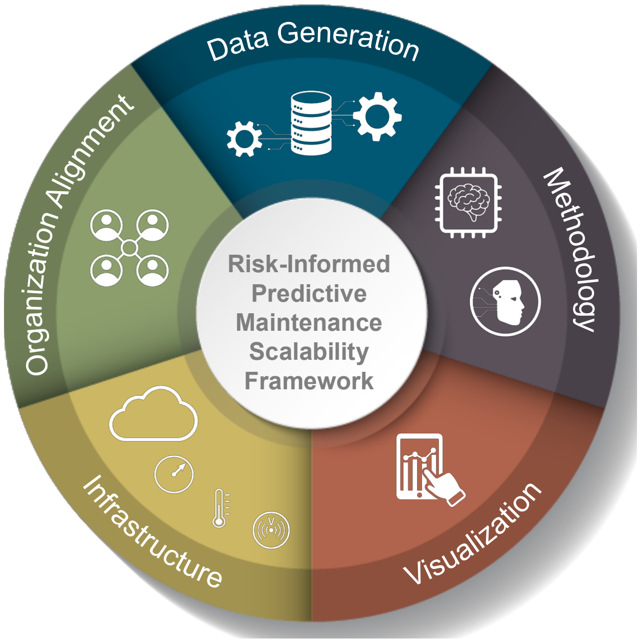

Multiband Heterogeneous Network[1]

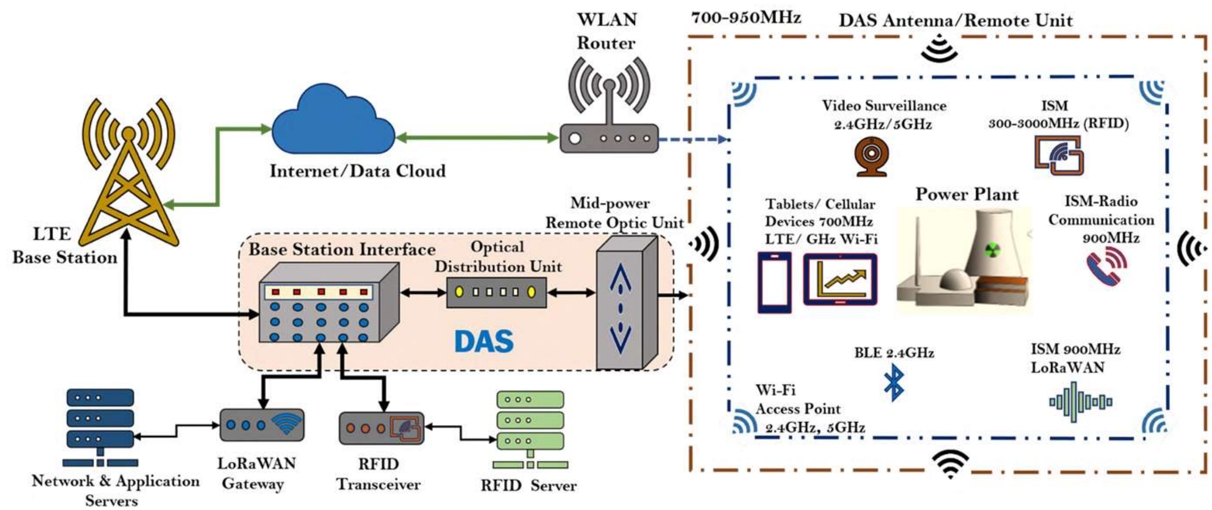

Scalability of developed approach across

- Same plant asset across the fleet and
- Different plant assets at the same plant site

Multiband Heterogeneous Network

- low power to high power, low-frequency to high-frequency, and short-range to long-range communication regimes

[1]Koushik, M., and V. Agarwal, "A Multi-Band Heterogeneous Wireless Network Architecture for Industrial Automation: A Techno-Economic Analysis," INL/EXT-19-55830, September 2019.

End Vision

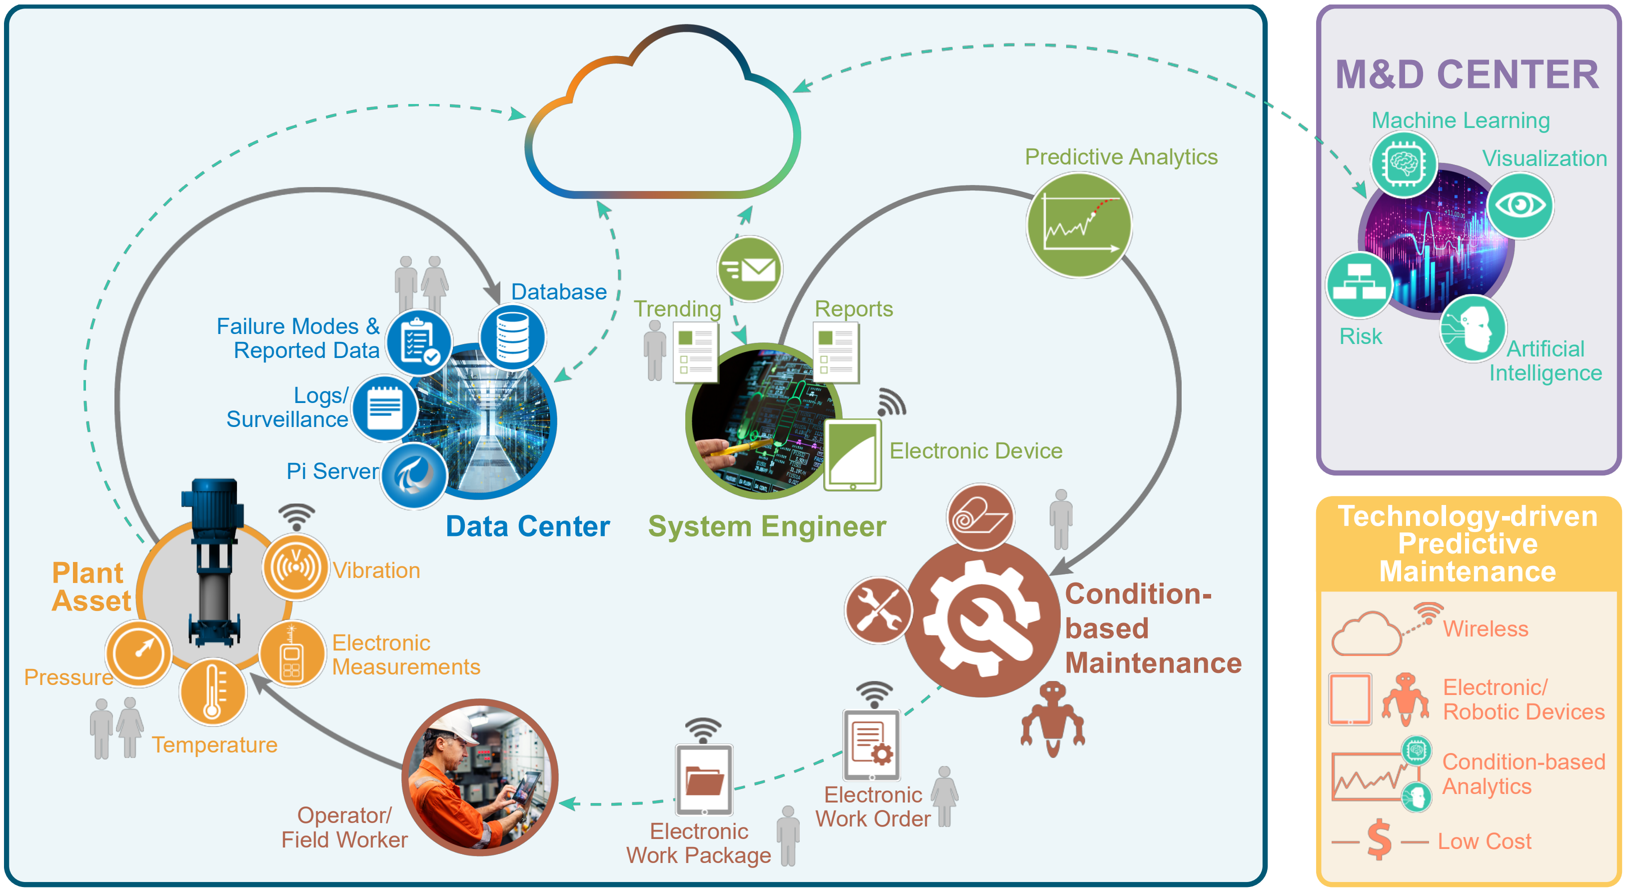

Database
Failure Modes &
Reported Data
Logs/
Surveillance
Pi Server
Data Center
Trending
Reports
Electronic Device
System Engineer
Predictive Analytics
Plant
Asset
Vibration
Electronic
Measurements
Pressure
Temperature
Operator/
Field Worker
Electronic
Work Package
Electronic
Work Order
Condition-
based
Maintenance
M&D CENTER
Machine Learning
Visualization
Risk
Artificial
Intelligence
Technology-driven
Predictive
Maintenance
Wireless
Electronic/
Robotic Devices
Condition-based
Analytics
Low Cost

Acknowledgments

Idaho National Laboratory

- James A. Smith
- Koushik A. Manjunatha
- Vaibhav Yadav

PKMJ Technical Services

- Mathew Mackay
- Francis Lukaczyk
- Michael Archer
- Nicholas Goss

Public Service Enterprise Group, Nuclear LLC

- Palas Harry

Questions?

Clean. **Reliable. Nuclear.**

Ahmad Al Rashdan

Group: **Instrumentation. Controls and Data Science Factors and Reliability**
Education: **Ph.D. in nuclear engineering from Texas A&M University, a M.Sc. in information technology and automation systems from Esslingen University of Applied Science in Germany, and a B.Sc. in mechanical engineering from Jordan University of Science and Technology.**

Presentation Overview

Machine Learning & Artificial Intelligence Symposium

- **Applications of Machine Learning in Automating Current Nuclear Operations and Work Processes**

www.inl.gov

Idaho National Laboratory

Applications of Machine Learning in Automating Current Nuclear Operations and Work Processes

April 17, 2020

Ahmad Al Rashdan, Ph.D.
Instrumentation, Controls, and Data Science

Machine Learning & Artificial Intelligence Symposium

Idaho National Laboratory

Motivation

LWRS
LIGHT WATER REACTOR SUSTAINABILITY

Machine Learning in a Nuclear Power Plant

Automate human activities (of visual, physical, analytical nature):

– Visual Physical Analytical

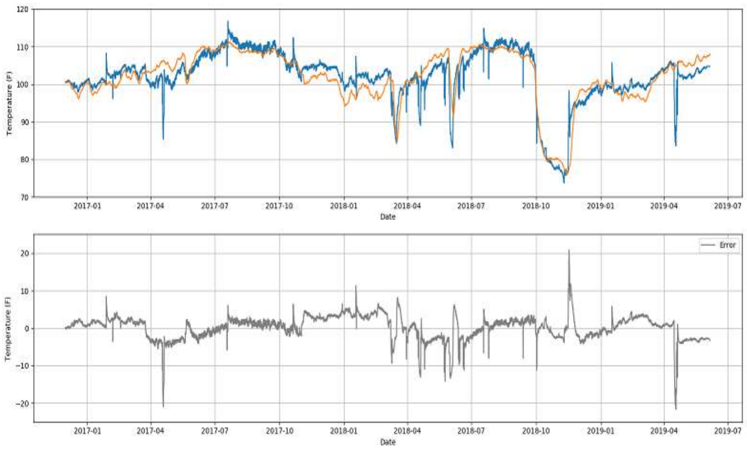

How? perform work autonomously, faster, more frequently, more accurately, or perform tasks that a human can't perform.

Why? Cost savings while sustaining safe and secure operations

Types of Applications

Collection

Analysis

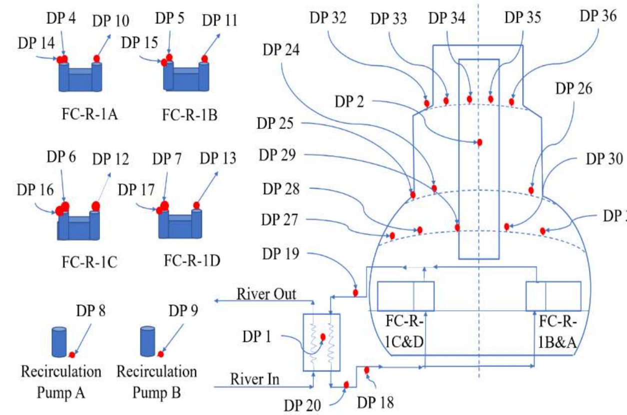

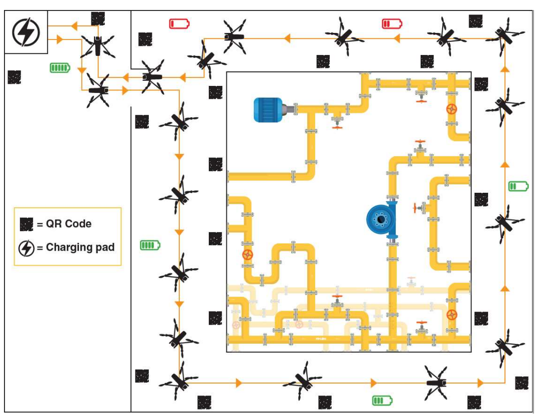

How does this advance ML/AI as a science?

Applied perspective on methods performance

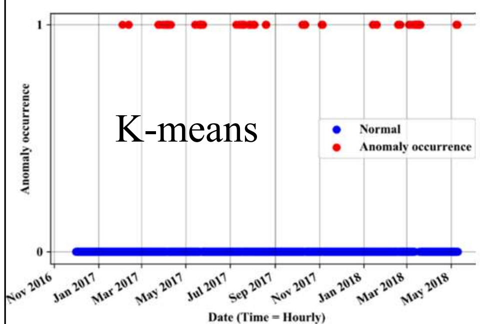

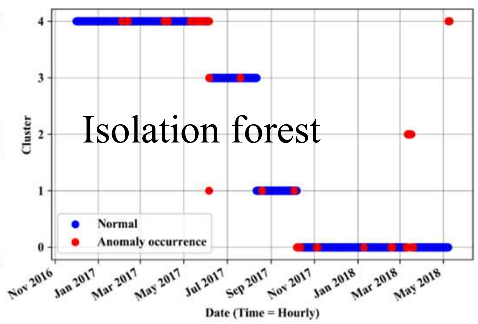

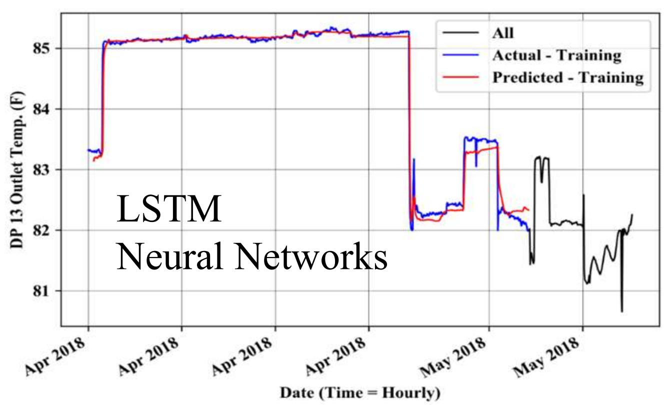

The balance between data and "physics" models

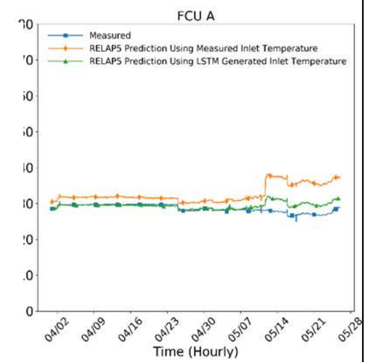

Gaps identification (looking ahead)

- Data (e.g. benchmarking)
- Methods (e.g. systematic approach)
- Verification & Validation(e.g. overfitting)
- Deployment (e.g. computational requirements)

Questions?

Clean. **Reliable. Nuclear.**

Cameron Krome

Group: High performance computing
Education: Bachelor's degree in computer science with a minor in math from Idaho State University in 2018 and is starting a master's degree in data science

Presentation Overview
Building a Scientific Language Model

- **General language models like BERT and roBERTa have been extremely successful when applied to a wide range of natural language processing tasks. These models were trained using everyday language taken from blog posts, Wikipedia, etc. A language model trained instead on scientific publications from arXiv.org may perform better on tasks involving scientific research.**

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Cameron Krome
C520 – HPC Data Analytics
Building a Scientific Language Model

Topic Introduction

- **A vast amount of data is freeform text**
- **Natural language processing (NLP) is a heavily focused area in ML/AI research**
- **The state-of-the-art methods for working with text involve general language models**
  - **ELMo**
  - **ULMFiT**
  - **BERT**
  - **roBERTa**
- **Existing models are built using *everyday* language sources**
  - **Blog posts**
  - **Movie reviews**
  - **Wikipedia**
- **Hypothesis:**
  - **If we generate a language model using scientific research papers, it may perform better for tasks involving scientific data**

Why it is relevant to ML/AI Future

- Text data is generated all the time during research
  - Logbooks
  - Freeform text fields in databases
  - Application log files
  - Software
  - Etc.
- The number of tasks that require working with this generated text are numerous and growing
- Problem: NLP methods change quickly
  - Modifying state-of-the-art models to fit our needs can enable the lab to keep up
- Problem: The latest models are computationally expensive
  - HPC resources are available for us to use if we take the time to learn how

Topic Details and Discussion

- Retrieved scientific publications from arXiv.org – approximately 1.6 million documents
- Extracted the text from the documents
  - Getting text from PDF files can be challenging
  - OCR had to be performed on many documents
- Trained roBERTa from scratch using Fairseq (PyTorch) on Sawtooth GPU nodes
  - Scaling is not perfect (but better than expected)
  - Final model runtime on 25 nodes: ~3 weeks
- Lessons learned
  - Don't worry about some bad text
  - Mixed precision is essential
  - Running on multiple nodes is challenging
  - Checkpoint often
  - Check the status of the job regularly

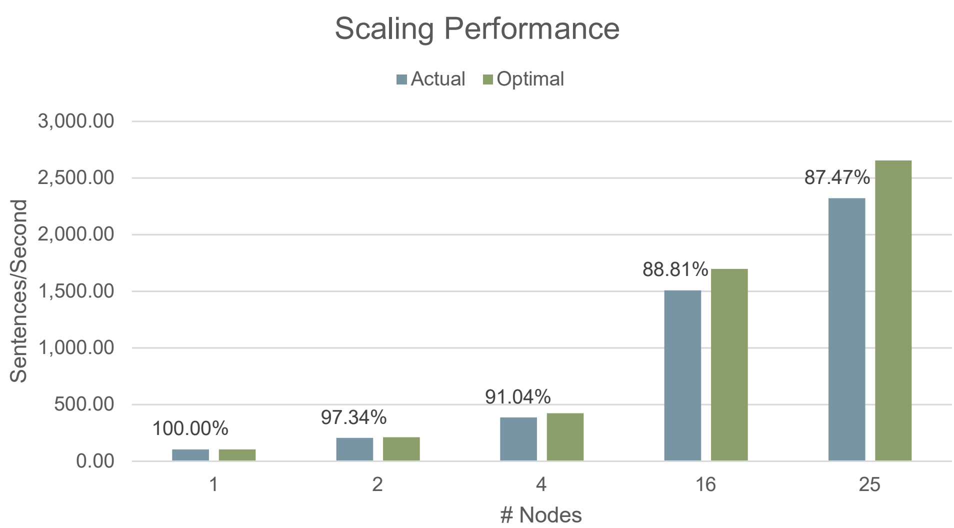

Looking Ahead

Idaho National Laboratory

- Test the model against current benchmarks
  - GLUE
  - SQuAD 2.0
  - CoLa
- Apply the model to INL tasks and compare against general language models
  - Document classification
  - Logbook analysis
  - Inventory optimization
  - Condition report screening
- Create other task-specific language models
  - Nuclear engineering models – non-proliferation, nuclear compliance verification
  - Models trained on non-word text (e.g. software, formulas, etc.)
- Explore other cutting-edge models/techniques
- Compare the performance and scalability of other libraries
  - Horovod
  - Tensorflow
  - PyTorch

Questions?

Clean. **Reliable. Nuclear.**

Matthew Anderson

Group: High Performance Computing C520
Education: PhD 2004, Physics, The University of Texas at Austin
Work focused in: Reinforcement learning and deep learning

Presentation Overview
Applying Machine Learning to Code Analysis

- **This talk gives a brief overview of how to apply machine learning and natural language processing to code analysis; the context of the discussion is malware analysis although the application space is much broader than just the reverse engineering of binaries. We approach the task from the perspective of machine translation with significant contributions from high performance computing and emerging hardware solutions.**

Machine Learning & Artificial Intelligence Symposium
April 17, 2020

Matthew Anderson
High Performance Computing, C520
Applying Machine Learning to Code Analysis

Topic Introduction

The Challenge

Malware and ransomware are becoming increasingly specialized and targeted. High performance computing (HPC) systems are starting to be targeted.

The challenge: rapidly **identify novel malware** and **reduce vulnerabilities.**

Examples:

- 2003 -- 2005: "Stakkato" attack against DOE, National Center for Atmospheric Research, and National Science Foundation (NSF) HPC sites
- 2014: Two NSF HPC sites were compromised by a US researcher.
- 2014—2017: "Cloud Hopper" attacks access the internal networks at Hewlett Packard Enterprise (HPE) and IBM and accessed customer systems.
- 2018: Nuclear scientists using the HPC system at the Federal Nuclear Center in Sarov Russia arrested for bitcoin mining.

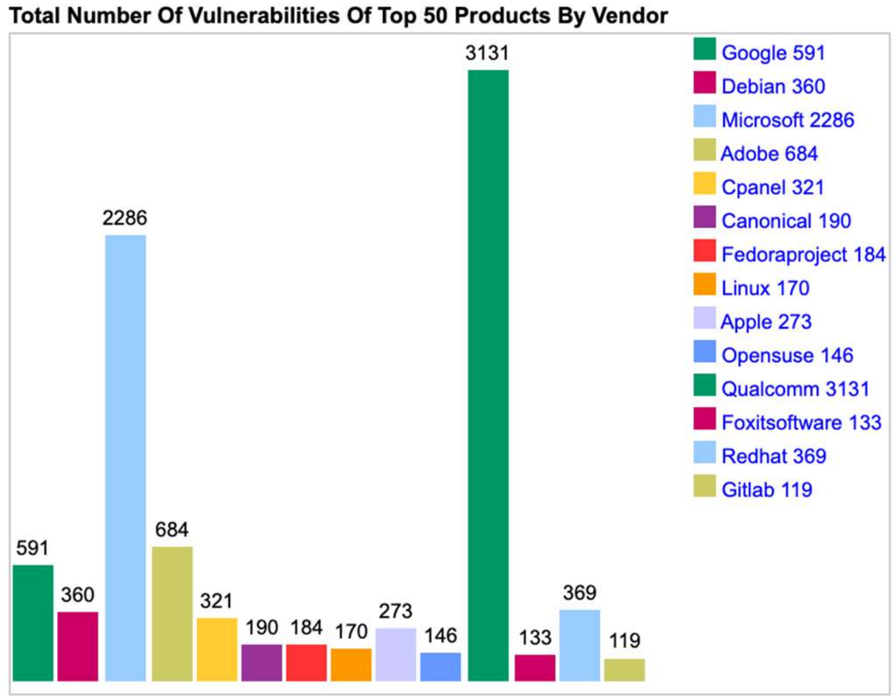

2019

Why it is relevant to ML/AI Future

Idaho National Laboratory

The Naturalness Hypothesis

> "Software is a form of human communication; software corpora have similar statistical properties to natural language corpora; and these properties can be exploited to build better software engineering tools."
>
> -- M. Allamanis, E. Barr, P. Devanbu, and C. Sutton (2017)
> arxiv.org/pdf/1709.06182.pdf

The outcome: Apply Natural Language Processing (NLP) and Machine Learning techniques to software!
Some Examples:

Reference	Source code analysis: Predicting Program Bugs	Source code analysis: Synthesizing patches and code changes	Binary analysis: Identifying function signatures	Binary analysis: Addressing Code Obfuscation	Binary analysis: Recovering compiler used to generate binary
Dam (2018)	✓				
Chakraborty (2018)	✓	✓			
Ding (2019)			✓	✓	
Massarelli (2019)			✓		✓

Challenges in Binary Analysis

1. Function names and debug symbols are stripped out from the binary

Source code

```
find_files(&files,media);
/* start encryption */
encrypt_files(files,&encrypted,&not_encrypted);
create_files_desktop(encrypted,files,desktop);
```

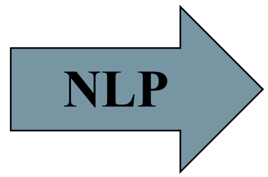

- Variable misuse detection
- Learning source code changes
- Defect prediction
- Cross-language learning
- Learning to represent programs with graphs

2. In real-life cases, we have to undo code obfuscation

Common Code Obfuscations:

- Packing
- Adding bogus logics
- Splitting basic blocks
- Substituting instructions
- Bogus control flow graphs
- Hot patching mechanisms (e.g. Conficker)

3. Assembly functions may appear different but still share the same functional logic

```
Addr_1:  mov eax,10
Addr_2:  dec eax
Addr_3:  mov [base+eax],0
Addr_4:  jnz Addr_2
Addr_5:  mov eax,ebx
```

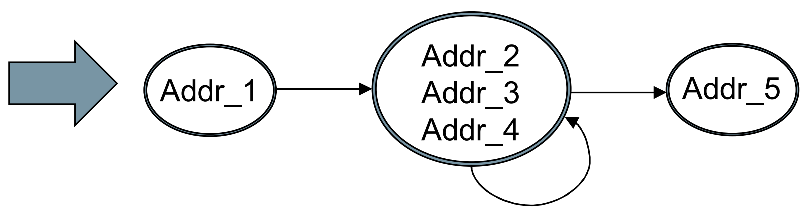

The Clones Ansatz:

> "Just as there is uncontrolled software reuse in source code, there exists a large number of clones in the underlying assembly code as well."
>
> S. Ding, B. Fund, P. Charland (2019)

Binary code fingerprints: **four types** of assembly code similarities

Literally Identical

i++

i = i + 1

Syntactically Equivalent

```
0x100000f9b <+27>: movl   -0x8(%rbp), %ecx
0x100000f9e <+30>: addl   $0x1, %ecx
0x100000fa1 <+33>: movl   %ecx, -0x8(%rbp)
```

```
0x100000fa4 <+36>: movl   -0xc(%rbp), %ecx
0x100000fa7 <+39>: addl   $0x1, %ecx
0x100000faa <+42>: movl   %ecx, -0xc(%rbp)
```

Slightly modified

memcpy
strcpy
memncpy
mempcpy

or

Same source with/without obfuscation

Semantically Similar

Opportunities for Deep Learning:

-- Identify binary similarities
-- Assign probable function names
-- Rapid identification of novel malware
-- Identification of software vulnerabilities

Datasets, Tools, and Approach

Datasets:

- Vulnerability dataset: Contains 3,015 assembly functions compiled with various compilers; contains variants of Heartbleed, Shellshock, Venom, Clobberin' Time, etc.
- UbuntuDataset: 87,853 ELF files disassembled using IDA Pro with >10 million distinct named functions
- NERO: 13,826 named functions from GNU repository with control flow graphs
- Research Malware/Ransomware: GonnaCry, Mirai

Tools:

OpenNMT — An open source neural machine translation system.

DeepGraphLibrary

asm2vec

angr

Approach:

- Approach binary analysis (binary similarity, function naming) using **Neural Machine Translation**:
  - Bidirectional recurrent neural network with Long Short-Term-Memory cells
  - Incorporate the Transformer Architecture
- **Augment** existing datasets with Github projects (>28 million public repositories) and more malware
- Create **new metrics** for scoring semantic similarity in binaries akin to what is used in NLP (e.g. BERTScore T. Zhang et al. 2020).

Questions?

Clean. **Reliable. Nuclear.**