

# Machine Learning: Issues and Opportunities

## ARPA-E Machine Learning-Enhanced Energy-Product Development Workshop

June 21-22, 2018

Falls Church, VA

## David E. Womble

Director of Artificial Intelligence Programs  
Oak Ridge National Laboratories

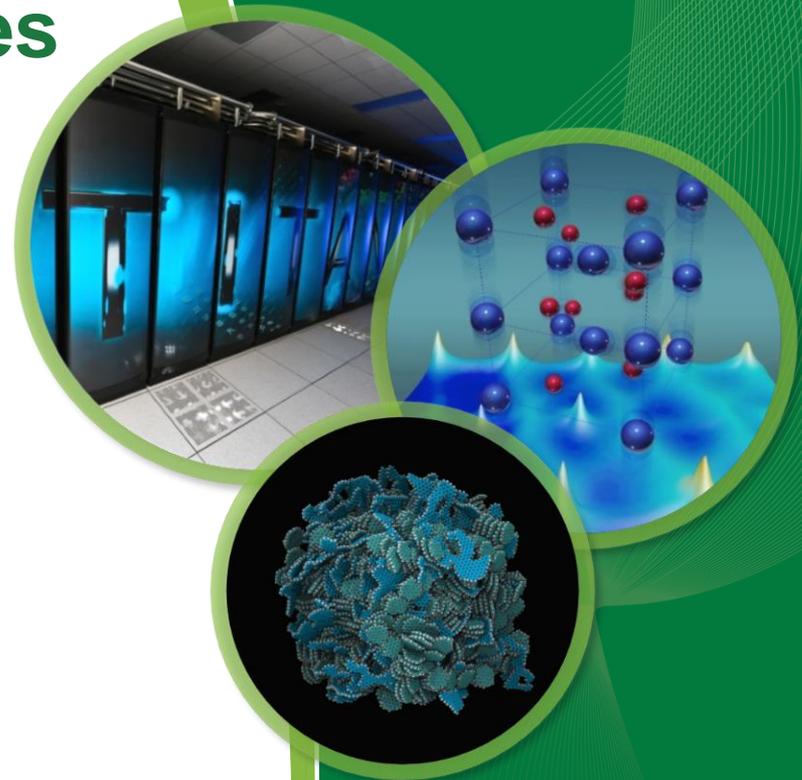
## With thanks to

Celia Merzbacher

Teja Kuruganti

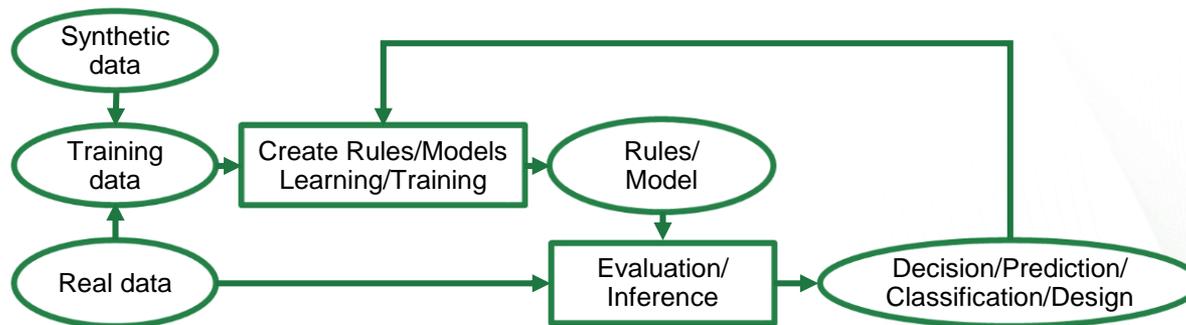
Srikanth Allu

Rich Archibald



# What are Artificial Intelligence (AI) and Machine Learning (ML)

- A class of data analytics algorithms in which the rules and/or models are not known a priori and are learned as part of the process
  - Process data to identify correlations
  - Complexity of the model is a potential problem
- Computers trained to perform tasks that if performed by a human would be said to require intelligence
  - Knowledge-based tasks
  - Computers are good at working with data, not “meaning”



# Hope/Hype/Hard Truth

- Hope
  - The convergence of big data, HPC, and AI will enable the accumulation and automation of functional knowledge across many application spaces.
- Hype
  - AI solutions are superior to collective intelligence of the experts for multi-modal data challenges
  - Effective translation of AI tools is straightforward
- Hard truths
  - AI solutions, thus far, are effective at executing narrowly defined tasks, identifying correlations in complex data
  - Need for sustainable heterogeneous data and compute infrastructure to advance AI innovation
  - Access to and availability of "good" and "labelled" data is one of the biggest challenges for AI
  - Vulnerability threats for AI (hacking, intentional manipulation) are a huge concern for deployment

# Taxonomy of AI Uses

- Classification and regression
  - Surrogates
  - Control
  - Inverse problems, design and optimization
- } dimension reduction



Near Infrared (single band)  
WorldView-3 image



CODA cloud detection saliency map for  
image above



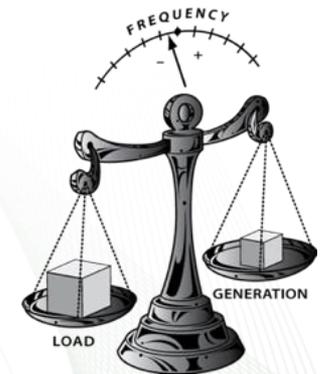
# Use Case: Smart Grid

## Application challenges

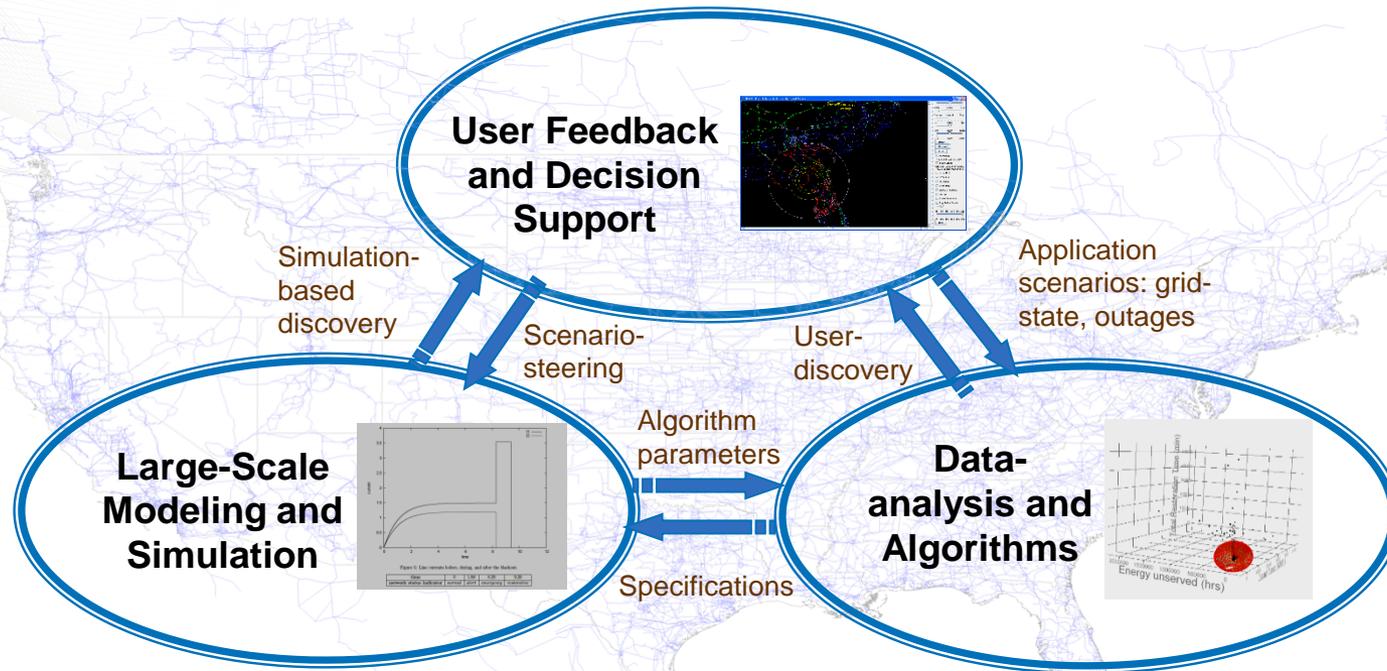
- Integrating variable distributed energy resources (DERs) with intelligent interfaces
- Integrating storage at multiple layers
- Integrating electric vehicles (EV)
- Managing demand – Residential, Commercial, Industrial
  - Enabling energy coordination and trading between buildings and trading between buildings and grid

## Technology challenges

- Connectivity across DERs
- Scalable control and diagnostics algorithms that are driven by data
- Actionable, real-time situational awareness
- Data and physical system security, including privacy

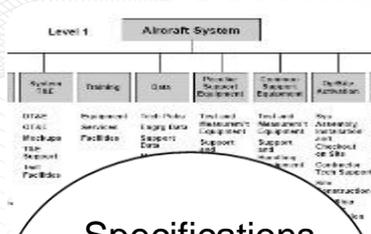


# Smart Grid: Leveraging A Data-rich Environment



- Learning algorithms for wide-area, hierarchical information sources
  - Distribution: Intelligent loads, SCADA devices, DERs
  - Transmission: Protection systems, power flow control
  - Generation: Planning and coordination
  - Control: Situational awareness, fine-grained control of DERs, enhanced reliability and resilience

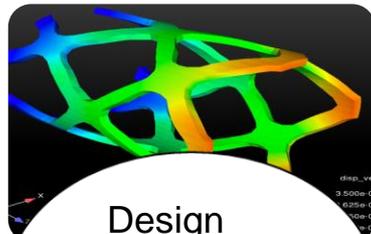
# Use Case: Additive Manufacturing



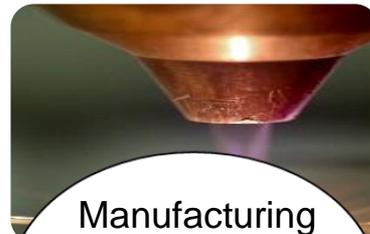
Level 1 Aircraft System

System Title	Training	Ops	Procedural Support/Outposts	Control System/Equipment	Repair/Inspection
DTAC	Procedural	Test Data	Tested	Tested	Rep
ATAC	Procedural	Log Data	Measurement	Measurement	Assembly
Workshop	Facilities	Support	Logistics	Logistics	Installation
TRP		Data	Support	Support	Checklist
Terminal			Support	Support	of SR
Mail			Support	Support	Collection
Facilities			Support	Support	Test Support

- Specifications
  - Functional
  - Environmental
  - Margins



- Design
  - Shape
  - Topology
  - Material
  - Process



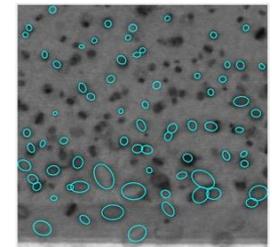
- Manufacturing
  - RT controls
    - Environment
    - Process
  - Diagnostics



- Testing
  - Validated Process
  - Test design

- Impact of machine learning

- Surrogate models
- Steering high-fidelity simulation
- Design, particularly materials and processes
- Real time diagnostics and control during manufacturing
  - Defect detection and mitigation
  - Control of local structures
- Predicted performance based on manufacturing data
- Test design and control



# Ensemble methods

- Statistical methods to improve the performance of machine learning algorithms
  - E.g., decision trees, k-NN
  - Most common application is perhaps the random forest
  - Not effective for stable learners
  - Most effective for weak learners
- Bootstrap aggregation (bagging)
  - Random selection of training data to improve stability and reduce variance
- Boosting
  - Ensembles of weak learners to create a stronger learner
  - Can be sequential or parallel
- Stacking
  - A trained meta-learner

# Random Forest/Decision Trees

## Feature Selection and Dimension Reduction

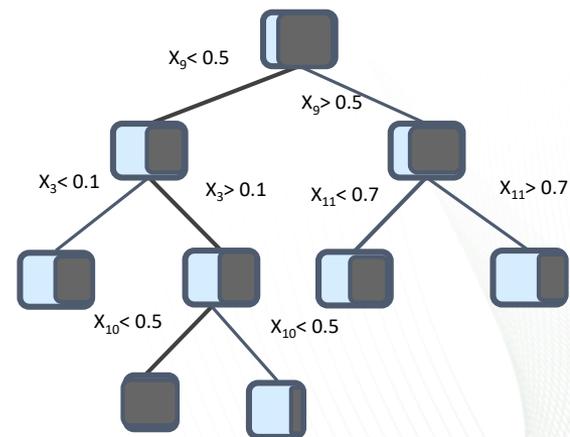
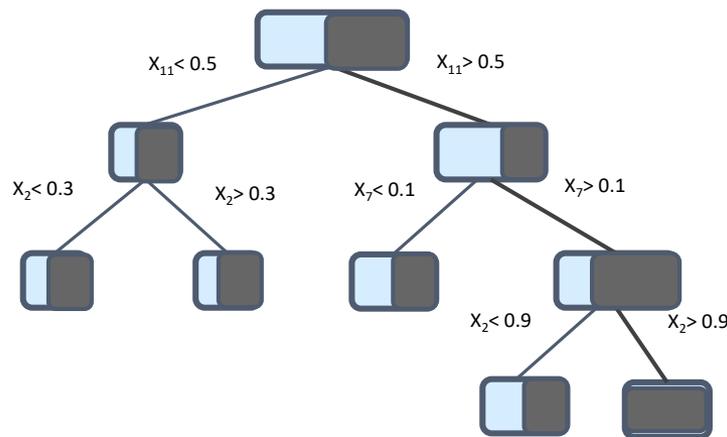
- Problem: Do an approximate combinatorial search to establish a feature-to-function relationship
  - A full search requires  $2^n$  computations
- Idea: Mine decision trees for patterns
- Decision trees -
  - More naturally explainable
  - Weak learners
  - Prone to overfitting

# Random Forest/Decision Trees

- Step 1 – determine branching criterion
- Step 2 – limit depth to prevent over-fitting
- Step 3 – apply bootstrap aggregation (bagging)
  - Select  $\alpha$  “large” and  $n' = \alpha n$
- Step 4 – apply feature bagging
  - Select  $p' = \sqrt{p}$
- Step 5 – boost (combine trees)

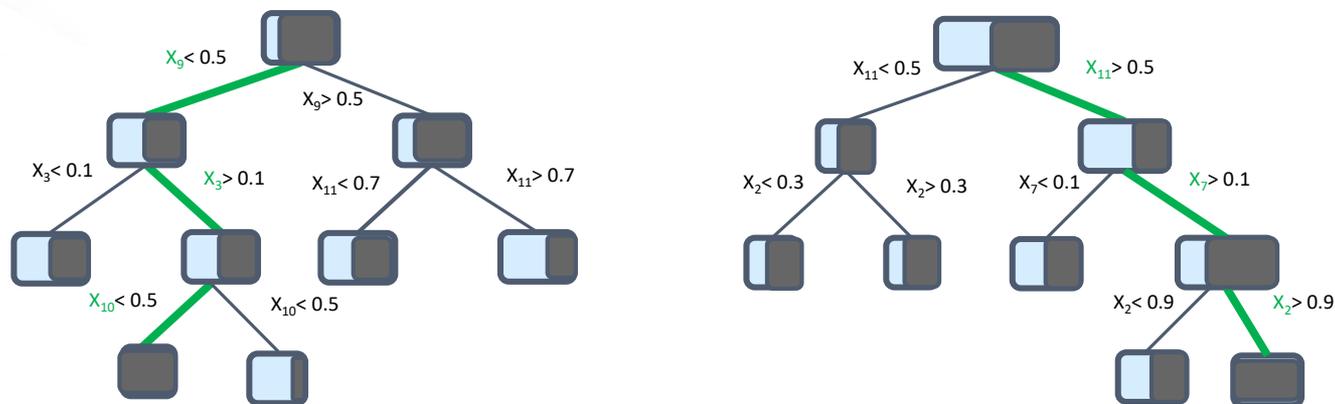
Decision tree

Random Forest

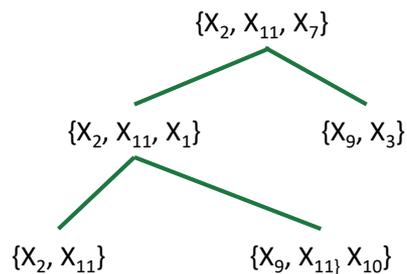


# Random Forest/Decision Trees

- Step 6 – Identify branching patterns and select feature sets

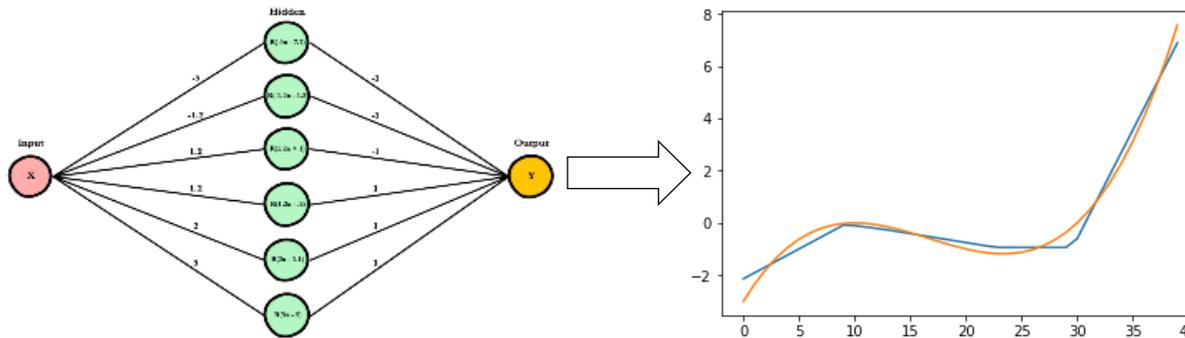


- Step 7 – create new RF branching on selected sets



# Neural Networks

- A NN is simply a function approximation, and a NN with a single hidden layer can approximate any function

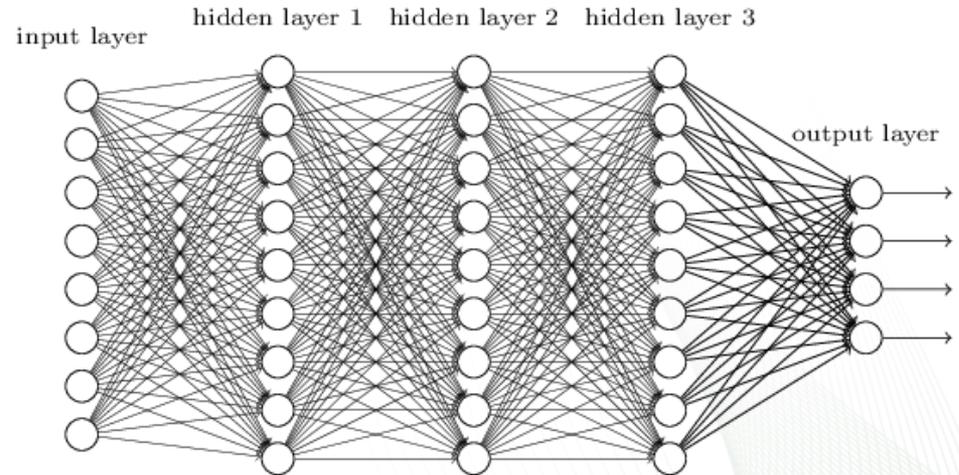
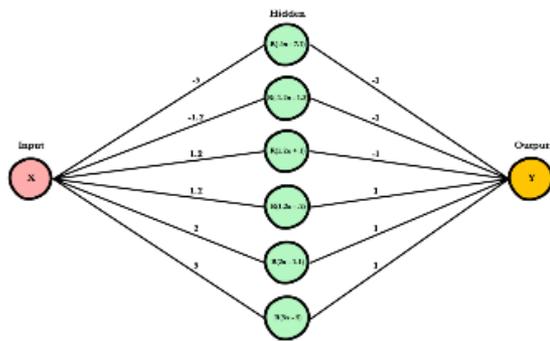


- Great for models when a specific model form is not known, but not much capability beyond basic statistical methods.
- NNs languished for decades

# Neural Networks – Significant Advances

## Deep Neural Networks

- Deep neural networks (DNNs) were introduced
  - Width increases the ability to approximate a function
  - Depth increases the abstractions, reduces the number of parameters but increases the computational requirements for training
  - Still susceptible to overfitting
  - Still an art



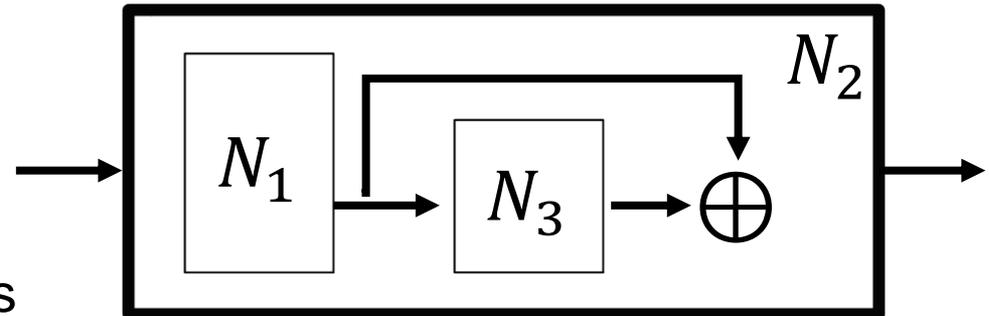
# Neural Networks – Accelerated Training

- Key idea for many improvements

$$\text{If } N_1 \subseteq N_2, \text{ then } L(N_1) \geq L(N_2)$$

- Leads to

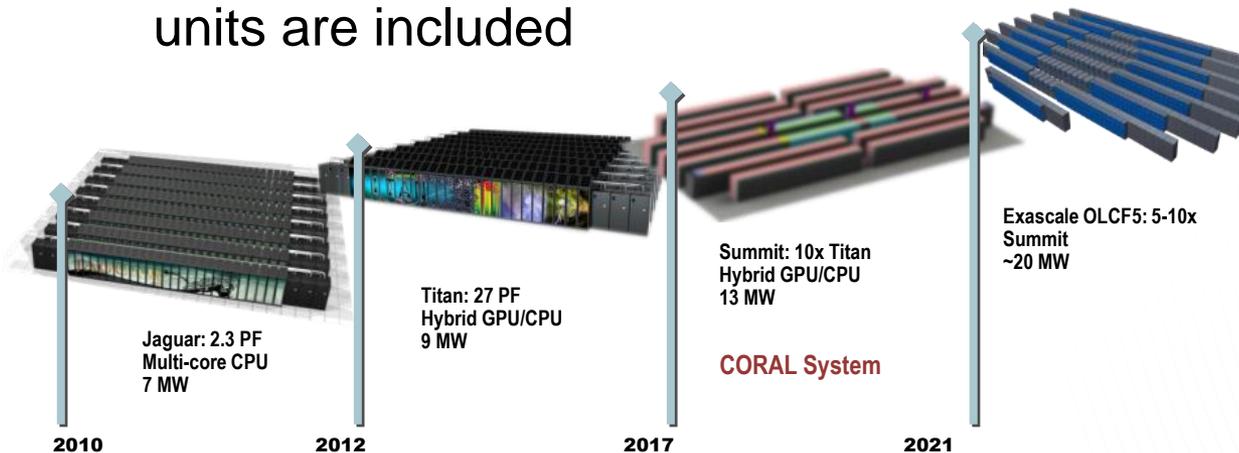
- Residual networks
- Inception networks
- Feature reuse
- Convolutional networks



- Training DNNs became algorithmically tractable
  - Stochastic gradient descent

# Neural Networks – HPC

- We have the ability to collect and store large amounts of data
- Computational power continued to increase, with architectural improvements that are amenable to neural networks
  - For example, GPU became practical for accelerated computations.
  - Reduced-precision tensor core units are included



# Issue: “Syntactic” Space vs. “Semantic” Space

- Humans tend to think in semantic space, i.e., in terms of the meaning.



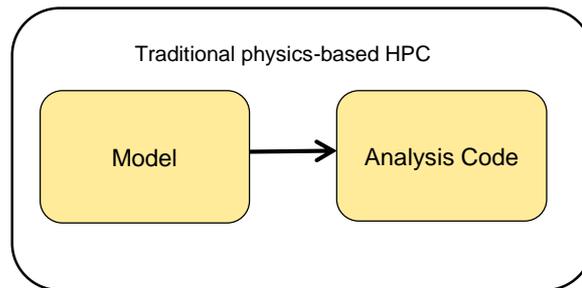
And metrics in semantic space are fundamentally different from those in syntactic space

- Implications
  - Easy to spoof classification systems
  - Transfer learning doesn't map well.

(Humans tend to transfer learning in semantic space, e.g., transfer what I learned about human behavior in kindergarten to how I drive. Most AI approaches transfer in syntactic space or transfer parts of the model (a sort of “gene transfer”).

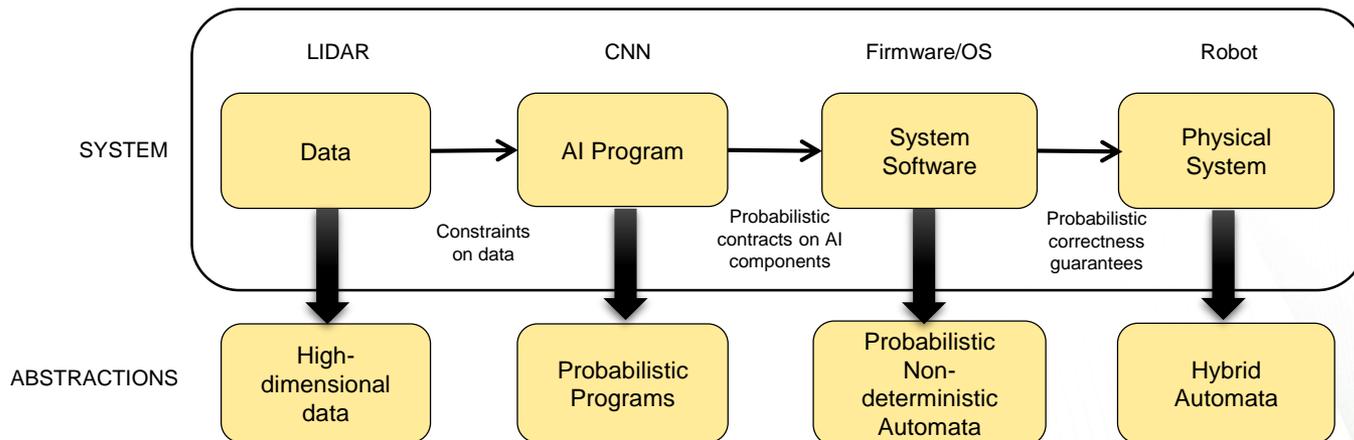
# Issue: Verification, Validation, Explainability and Interpretability

- Verification
  - Is the model implemented correctly?
- Validation
  - Is the model (including training data) appropriate for the decisions being made?
  - Must be evidence based
  - Requires some form of UQ, robustness guarantees and bounds on “distortion”



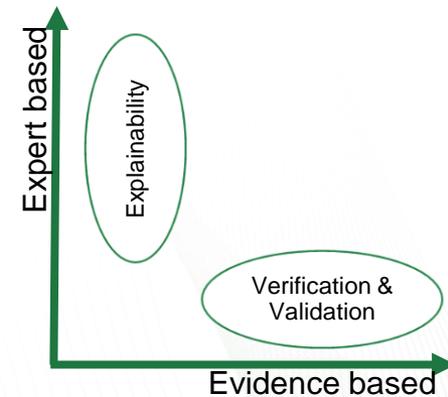
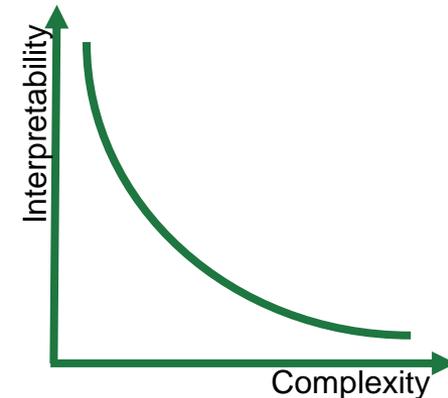
# Issue: Verification, Validation, Explainability and Interpretability

- Verification
  - Is the model implemented correctly?
- Validation
  - Is the model (including training data) appropriate for the decisions being made?
  - Must be evidence based
  - Requires some form of UQ, robustness guarantees and bounds on “distortion”



# Issue: Verification, Validation, Explainability and Interpretability

- Interpretability
  - Can a human understand the model? For example, do the basis vectors in a dimension reduction algorithm have a physical meaning?
- Explainability
  - Can the model present a sequence of steps that can justify the answer to an expert?
  - Expert based
- Reproducibility
  - Does the same experiment lead to the same conclusion?
  - Can we run different experiment and not contradict our conclusion?
  - If we create a new model with the same data, do we get the same conclusions?
  - Required for good science

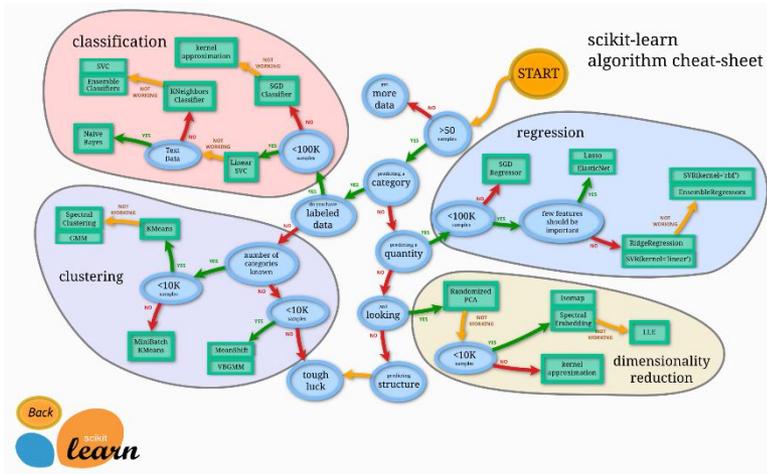


# Issue: Data Is A Major Problem

- Need more data than was imagined just a few years ago
  - We are looking for complex correlations
  - Using primarily statistical methods
- Labelled data is a problem
  - Generating labels is expensive and labor intensive (e.g., Mechanical Turk)
  - Need to move toward reinforcement learning
- Synthetic data and simulated environments are partial solutions
  - But an AI can learn the flaws in these systems

# Issue: AI Is An Art

- Choosing the model form and hyper parameters is often ad-hoc and requires experience and insight
- AI models must be tuned
- Neural networks design is difficult and often requires tuning
- Interpreting the results requires expertise



“Machine learning methods are often described in papers at an abstract level, for maximum generality. However, a good choice of hyperparameters is usually necessary to make them work well on real-world problems, and tricks are often used to make most efficient use of these methods and extend their capabilities.”

G. Montreван, et.al., “Methods for Interpreting and Understanding Deep Neural Networks.”

# Six Research Areas To Be Addressed

- Data quality and statistics
  - Even if we have enough data, it is not necessarily good data
  - Dealing with bias
- Machine learning
  - Needs to accelerate
  - Very model dependent
- Merging physics and AI
  - We can't violate the laws of physics
- Verification, validation and explainability
  - Is the answer right, is the model appropriate, and can we understand it
  - What is the human-computer interface
- Computing
  - How do we use “big” computers
  - How do we use accelerated nodes
- Deployment
  - Computing at the edge
  - privacy, ethics and regulations

