Beyond Industrial AI:
The Path to Actionable Intelligence

Dr. Michael Sharp

michael.sharp@nist.gov

Disclaimer:

Commercial equipment and materials might be identified to adequately specify certain procedures. In no case does such identification imply recommendation or endorsement by the U.S. National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

Outline

- What is IAI?
  - How does information become intelligence?
  - Picking the right tool for the job
- State of Development
  - Where has IAI excelled?
  - Where has it fallen short?
- Lessons Learned
  - Understanding and managing IAI
  - Common mistakes
- The Path to Action
  - Where is autonomy appropriate?
  - Understanding risks
- What Comes Next?
  - Designing with humans in mind!
  - Community driven results!

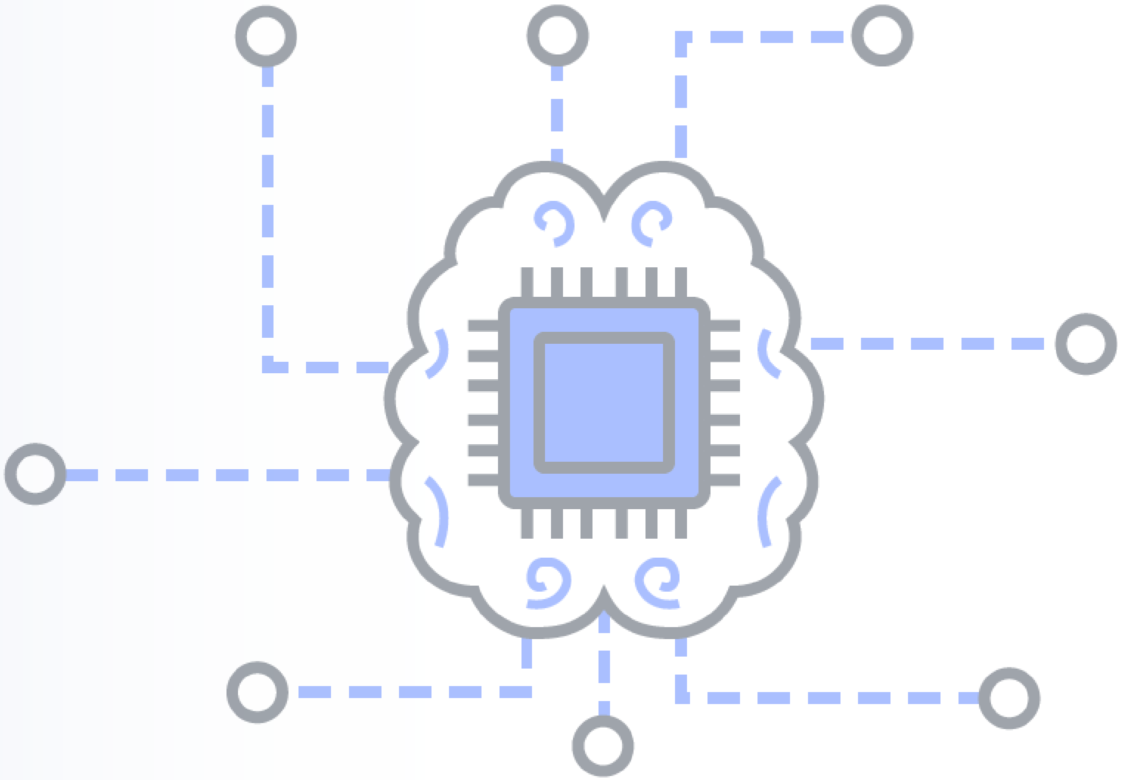

How Does Information Become Intelligence?

- Contextualizing Knowledge and Experience into Informed Actions
  - Record Data
  - Observe Trends / Behaviors
  - Create Rules - Explicit or Implicit
  - Enact Decisions Leading to Actions
- Intelligence is the ability to enact context driven decisions

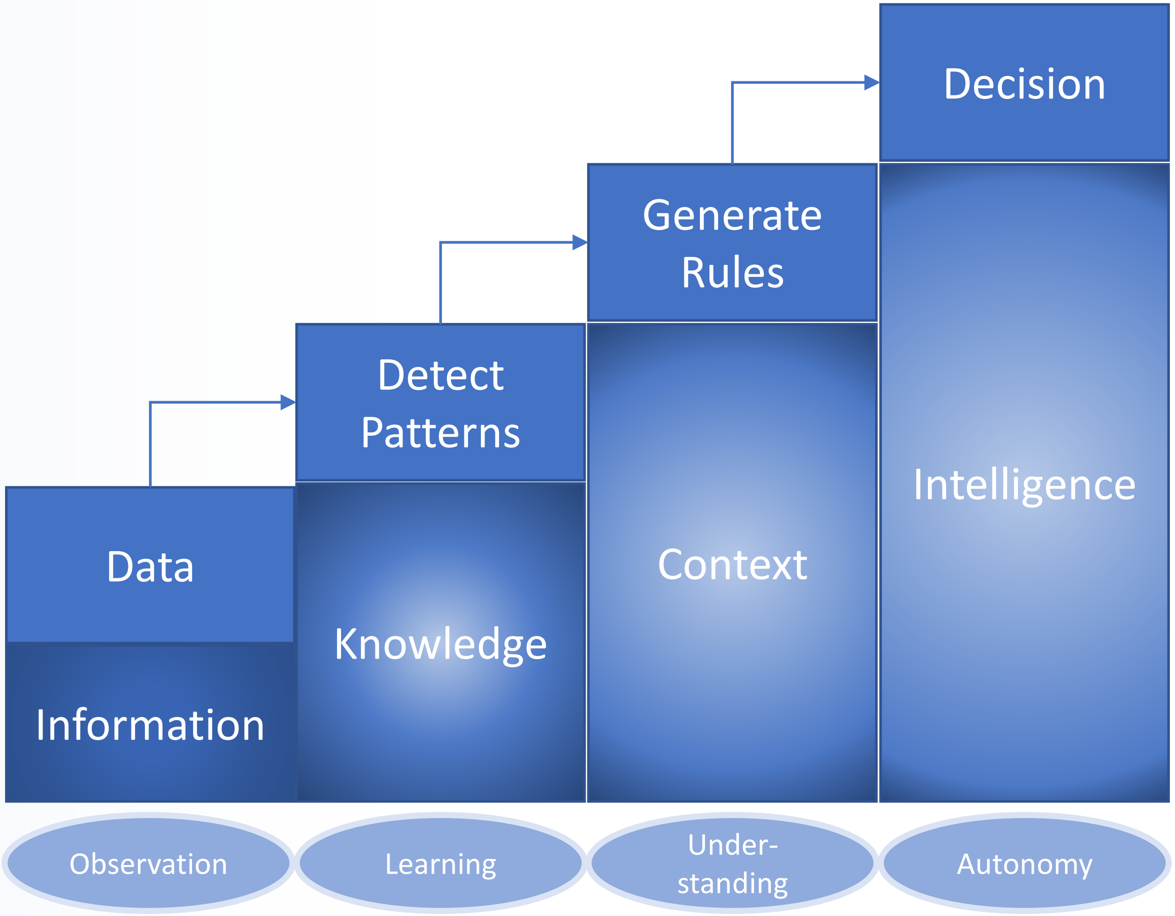

National Institute of Standards and Technology
U.S. Department of Commerce

Where Does Industrial AI Fit in to This?

- **IAI is a Subset of Artificial Intelligence Applied in Industrial Setting**
  - It is many of the same Tools and Algorithms used by non-physical industries and consumer products
    - IAI differentiates itself by having unique Expectations, Limitations, and Capabilities derived from utilizing domain specific assets
- **Is it Machine Learning?**
  - Sometimes, but not always
    - Many algorithms perform decision making or evaluation tasks using rigid rules supplied by human agents without updating or evolving
- **Will it Replace Human Operators?**
  - In almost every case, the answer is ***no***
    - IAI may alter the requirements or lessen the burden on a human agent, but it can not replace the final information synthesis and inspired action initiative that is available to human intelligence
      – i.e. *the ability to understand and imagine creative responses to any situation*

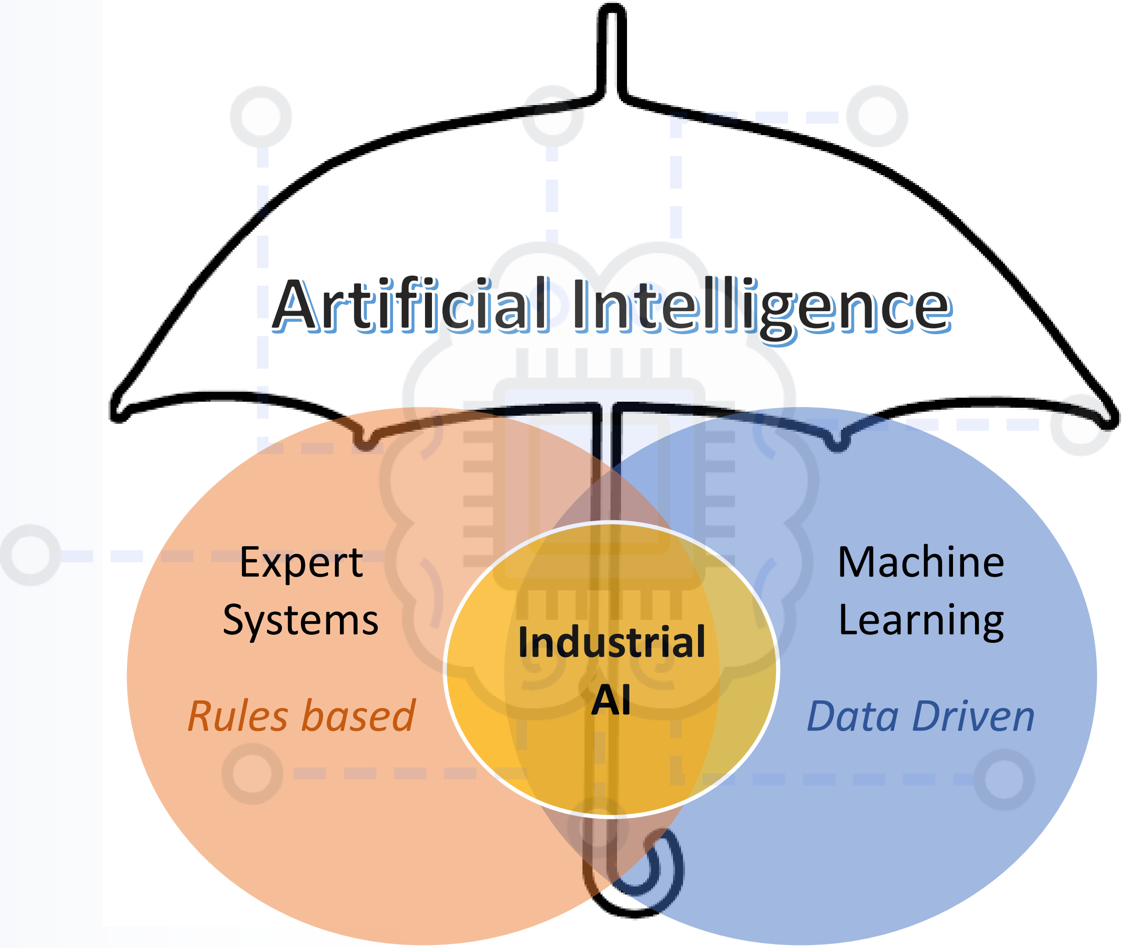

National Institute of Standards and Technology
U.S. Department of Commerce

What Distinguishes Industrial AI?

- IAI must function in **Real World** *Engineering* and *Industrial* applications

1. IAI Systems and Models are made to solve a ***known problem*** or provide some ***explicit benefit***
   - Because of this **Evaluation** is a necessary and continual process. A solution that worked last year may not meet the needs of today.
2. Practical use for the end user is a higher priority than philosophical elegance
   - Solutions that fulfill requirements and are ***easier to understand, verify, and maintain are preferred*** unless there is known reason not to
3. Justifications for modeling choices must come from the greater ***context of the application***.
   - Solutions are not applied in a void and thus must not be developed or evaluated agnostic to the domain and specific application.

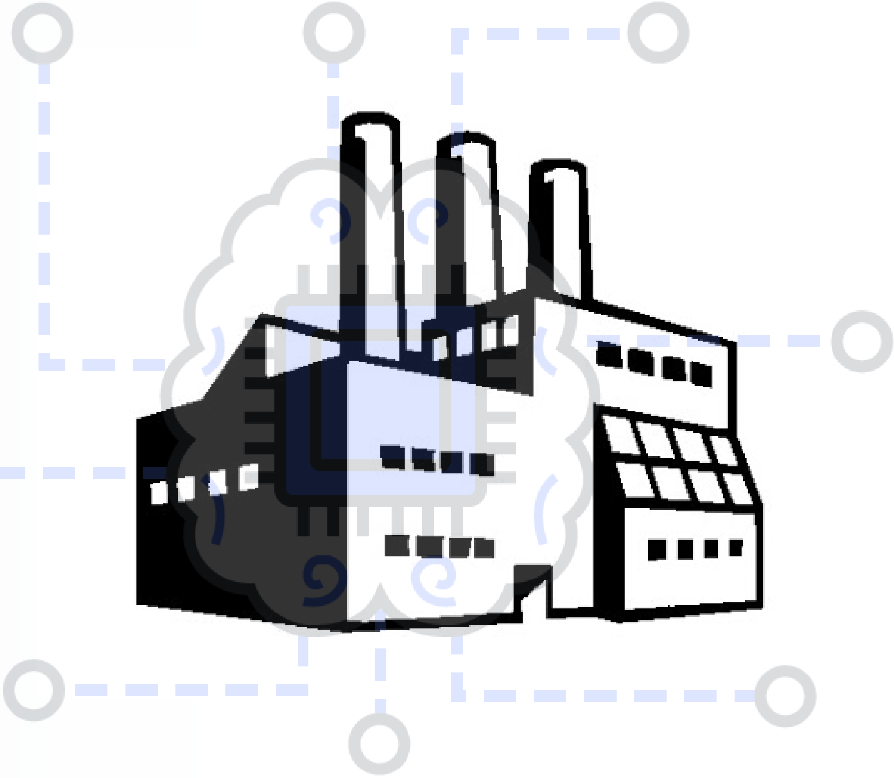

What are unique aspects of Industrial AI?

- Industrial AI Faces a Unique Set of Advantages and Disadvantages
  - **Data Availability**
    - + **Wealth of Data Sources** with unique / eccentric attributes
    - - **Restricted Usability** due to comparatively few labeled exemplars, asynchronicity, and lack of compatibility or accessibility
  - **Domain Application Assets**
    - + **Leverageable Rules and Domain Expertise** from years of experience to help create and test IAI
    - - **Complex Systems and Interactions** that with more importance of edge cases than found in many digital industries
  - **IAI Tool Use**
    - + **Wide Array of Available Tools** from open-source developers and 3rd party solutions
    - - **Expertise Gap** for developing, implementing, and evaluating IAI Tools – Especially for SMEs

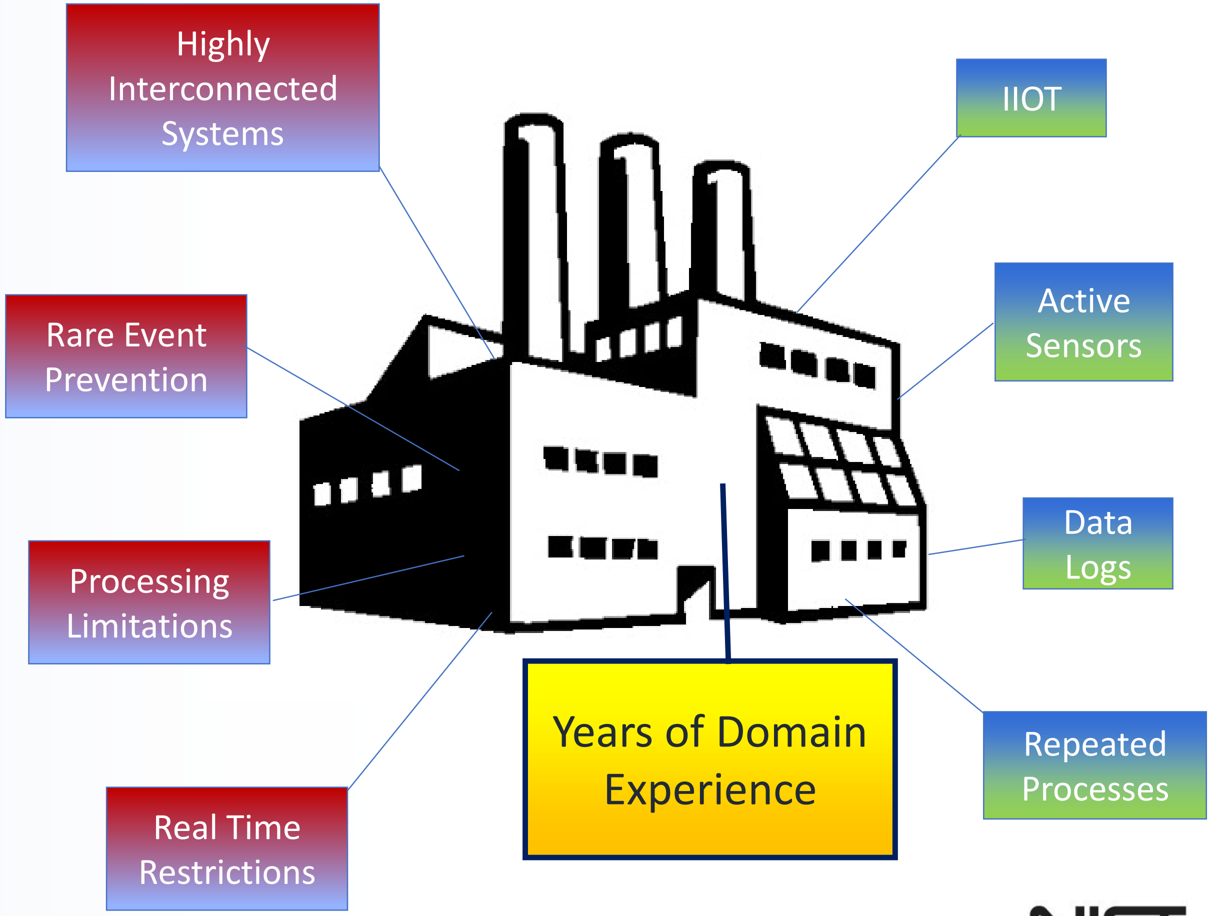

NIST National Institute of Standards and Technology U.S. Department of Commerce

Picking the Right Tool for the Job!

- **No single tool can manage every problem**
- Things to consider:
  - **Application and Domain Assets and Assumptions**
    - Will this tool work in my environment?
    - Ex. Time Series Analysis vs Point Estimators
  - **Form of Available Input**
    - Does the tool work better on / need certain types of inputs?
    - Ex. Convolutional Networks with Images
  - **Volume of Relevant Data**
    - Can the tool process in-situ at speeds I need?
    - Ex. Simulation-based tools may be slow
  - **Decision Support vs Decision Making** (i.e. Level of Autonomy)
    - How involved is human decision making?
    - Ex. Bayes Nets offer stepwise explainablity in decisions
  - **Required Fidelity**
    - Is the tool appropriately complex for the task?
    - Ex. "Go" vs "No Go" applications may not need high precision models
  - **Available Methods for Testing, Evaluation, Verification, and Validation**
    - If you can not test it ***in a way you trust***, *you cannot rely on it*

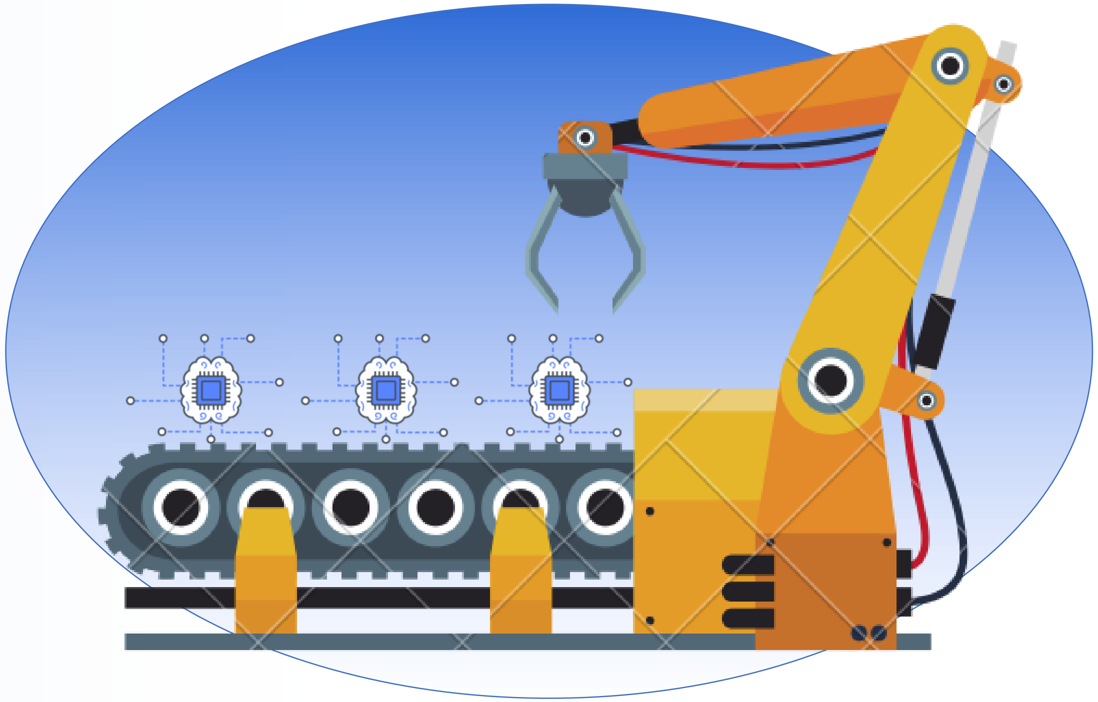

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Picking the Right Tool

- Information Discovery
  - Problem Area Identification
  - Document Relevance and Retrieval
- System Evaluation
  - Condition Monitoring
  - Optimized Process Control
- Product and Process Design
  - Automated Error Change Requests
  - Production Cost Estimation
- Training and Knowledge Transfer
  - Real-time Manuals and Instruction
  - Policy Capture and Evaluation

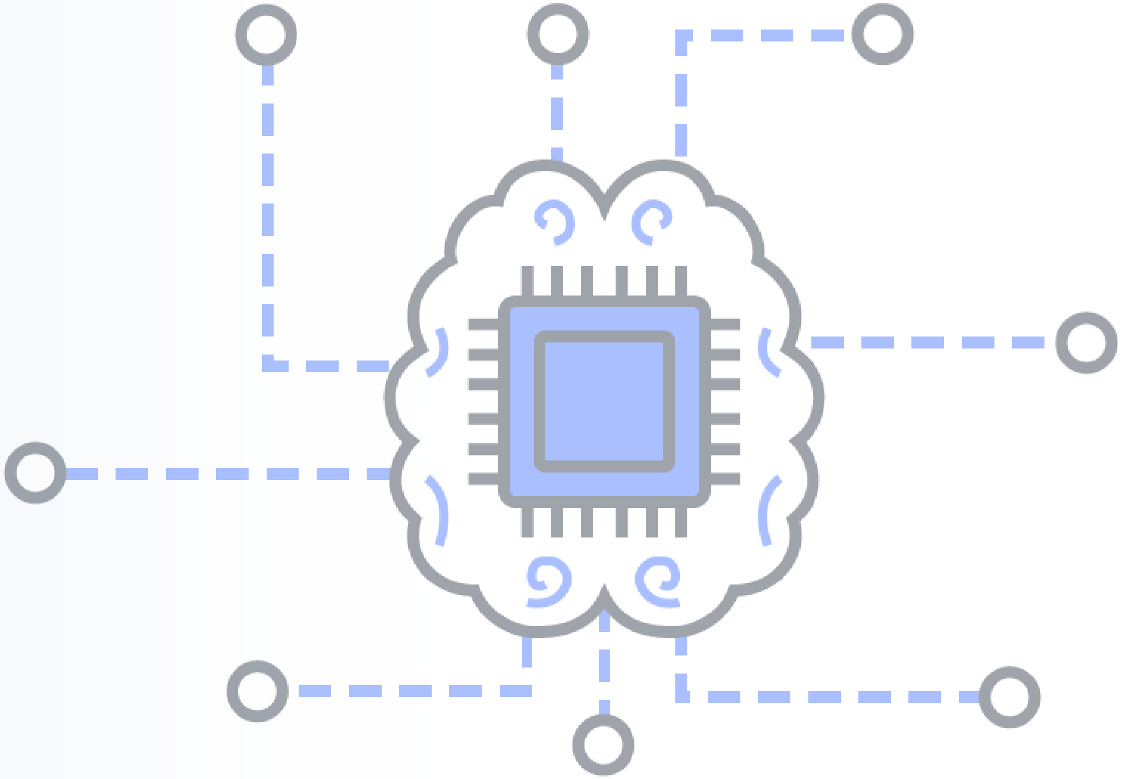

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Picking the Right Tool

- Information Discovery
    - Problem Area Identification
    - Document Relevance and Retrieval
- System Evaluation
    - Condition Monitoring
    - Optimized Process Control
- Product and Process Design
    - Automated Error Change Requests
    - Production Cost Estimation
- Training and Knowledge Transfer
    - Real-time Manuals and Instruction
    - Policy Capture and Evaluation

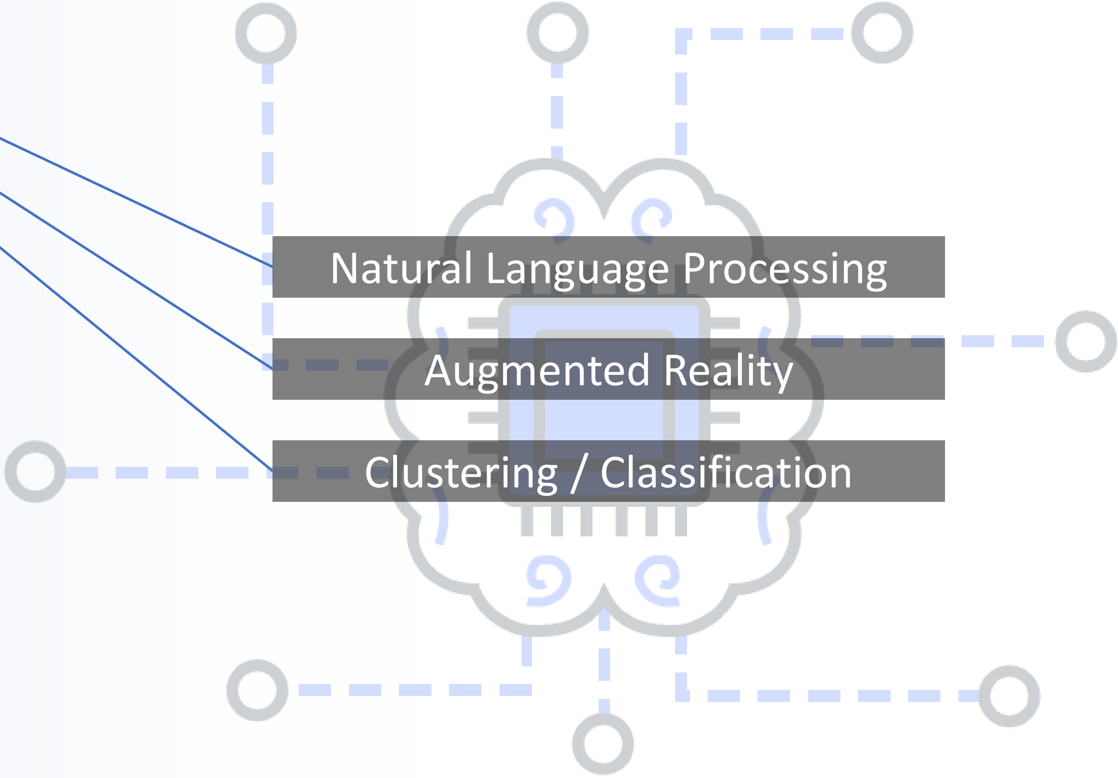

Picking the Right Tool

- Information Discovery
  - Problem Area Identification
  - Document Relevance and Retrieval
- System Evaluation
  - Condition Monitoring
  - Optimized Process Control
- Product and Process Design
  - Automated Error Change Requests
  - Production Cost Estimation
- Training and Knowledge Transfer
  - Real-time Manuals and Instruction
  - Policy Capture and Evaluation

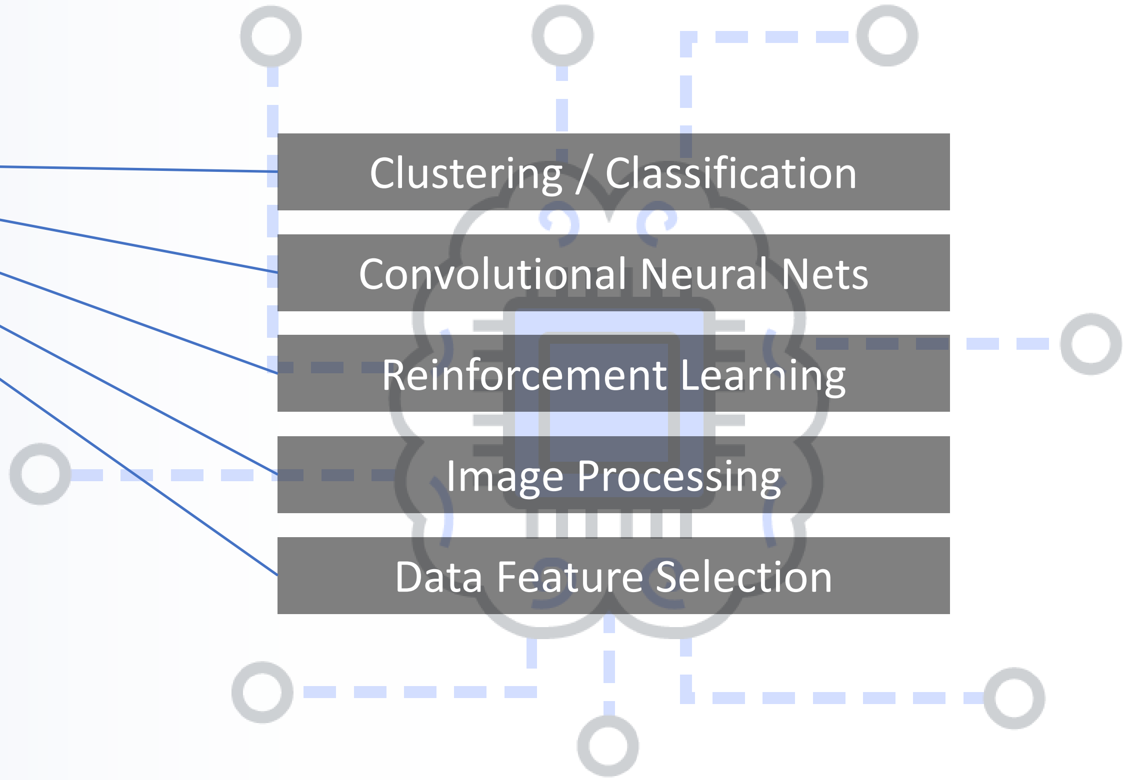

Picking the Right Tool

- Information Discovery
  - Problem Area Identification
  - Document Relevance and Retrieval
- System Evaluation
  - Condition Monitoring
  - Optimized Process Control
- Product and Process Design
  - Automated Error Change Requests
  - Production Cost Estimation
- Training and Knowledge Transfer
  - Real-time Manuals and Instruction
  - Policy Capture and Evaluation

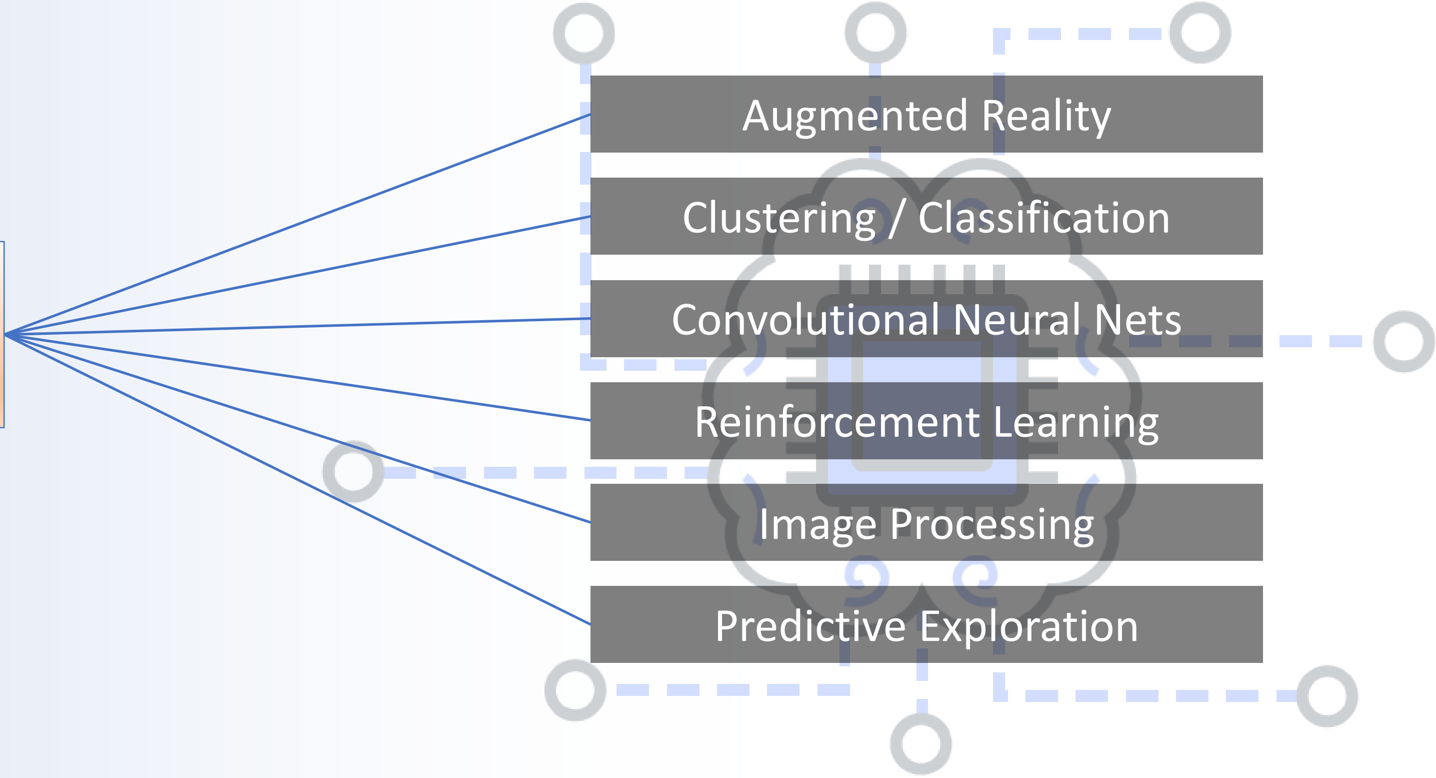

Picking the Right Tool

- Information Discovery
  - Problem Area Identification
  - Document Relevance and Retrieval
- System Evaluation
  - Condition Monitoring
  - Optimized Process Control
- Product and Process Design
  - Automated Error Change Requests
  - Production Cost Estimation
- Training and Knowledge Transfer
  - Real-time Manuals and Instruction
  - Policy Capture and Evaluation

Augmented Reality

Clustering / Classification

Natural Language Processing

Reinforcement Learning

Image Processing

What are some active developments of IAI?

- Success Stories
  - ***Information Discovery***
  - ***System Evaluation***
  - ***Product and Process Design***
  - ***Training and Knowledge Transfer***

Where has IAI already been successful?

- Success Story – Information Discovery : Problem Area Identification
  - **Goal –** Use maintenance work orders to uncover problem areas in a facility (originally automotive – now expanded)
  - **Challenge** – High volume of documents filled with fragmented language, misspellings, site specific jargon that neither computer or human can easily sort through
  - **Method** – Use NLP to assist humans in annotating words and phrases, then analyze relationships of recurring concepts

Raw Text from Maintenance Documents

"Hyd leak at saw atachment"	"Replaced seal in saw attachment but still leaking – Reapirs pending with ML"
"HP coolant pressure at 75 psi"	"Bad Gauge / Low pressure lines cleaned ou"
"Major hydraulic leak at Sp#6 horseshoe"	"Repaired horseshoe seals"

Technical manufacturing documents contain misspellings; jargon; abbreviations; short sentences making out of the box NLP solutions difficult

Nestor: NLP toolkit for manufacturing

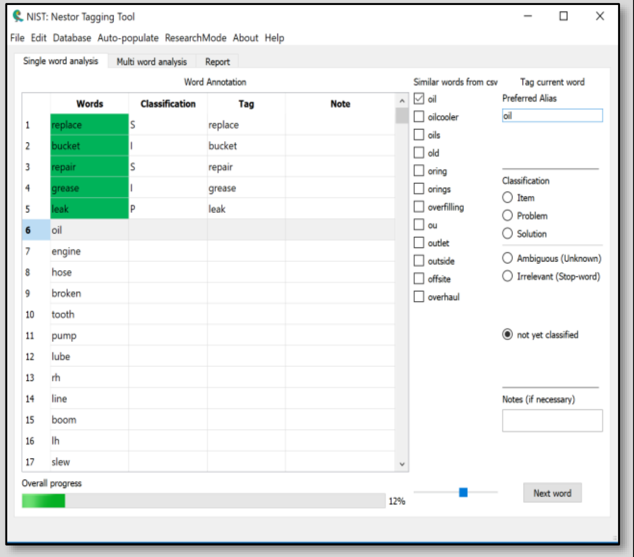

Nestor uses NLP to extract terms in order of importance for user classification

Cleaned and Annotated Text for Analysis

Item(s)	Problem(s)	Solution(s)
Hydraulic; Saw attachment; Seal	Leak	
High Pressure Coolant; Gauge; Low Pressure Line	Broken; Low Pressure	Cleaned
Hydraulic; SP#6 Horseshoe; Seal	Leak	

The output of Nestor automatically contextualizes manufacturing documents for analysis to improve operational decision making

https://www.nist.gov/services-resources/software/nestor

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Where has IAI already been successful?

- Success Story – System Evaluation: Robotic System Condition Monitoring System (CMS)
  - **Goal –** Identify degradation in a linked, coordinated system (Robotic Arm)
  - **Challenge** – Coordinated movements mask single element problems, non-repetitive movements and interactions with unknown external factors complicate transient analysis
  - **Method** – Utilize single and group-based models to alert users familiar with daily operations about trending or outlier anomalies. Give users about to "reset" what expected normal is as operations change.

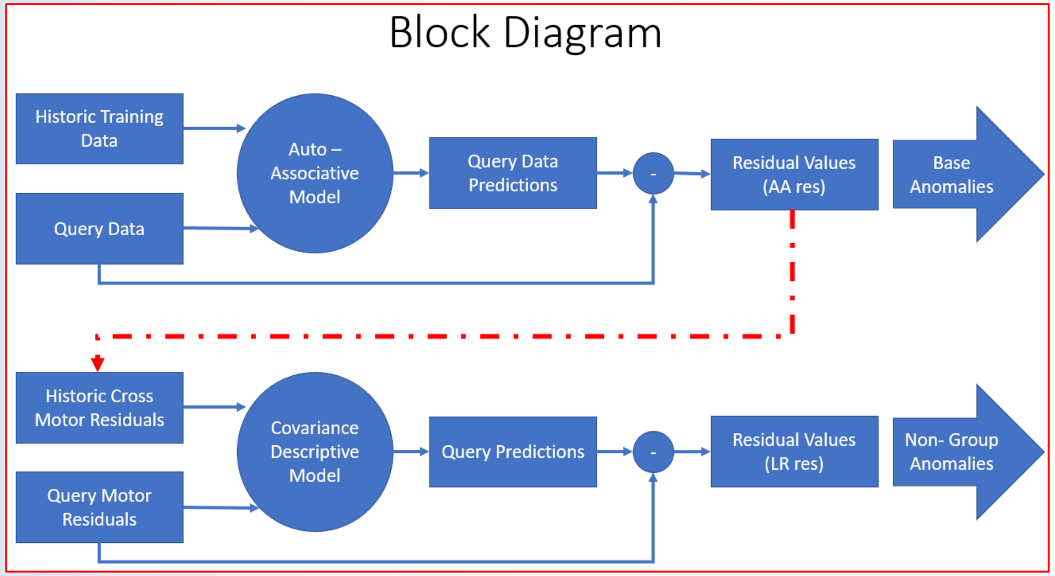

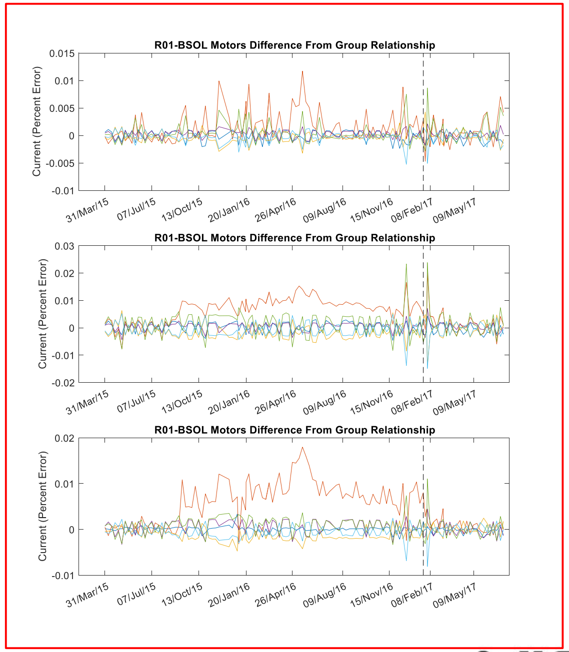

National Institute of Standards and Technology
U.S. Department of Commerce

Where has IAI already been successful?

- Success Story – Product and Process Design : Autonomous Error Change Requests (ECR)
  - **Goal –** Rapidly achieve solutions to product design conflicts (Ex. Inaccessible hole placement)
  - **Challenge** – High volume of minor (easy for human) changes burden human designers. Yet the high complexity of requirement interactions makes guaranteeing optimal solutions computationally prohibitive.
  - **Method** – Use rapid space exploration algorithms (ex Genetic Algorithms) to evaluate suitability and propose possible solutions to human designers to confirm or iterate

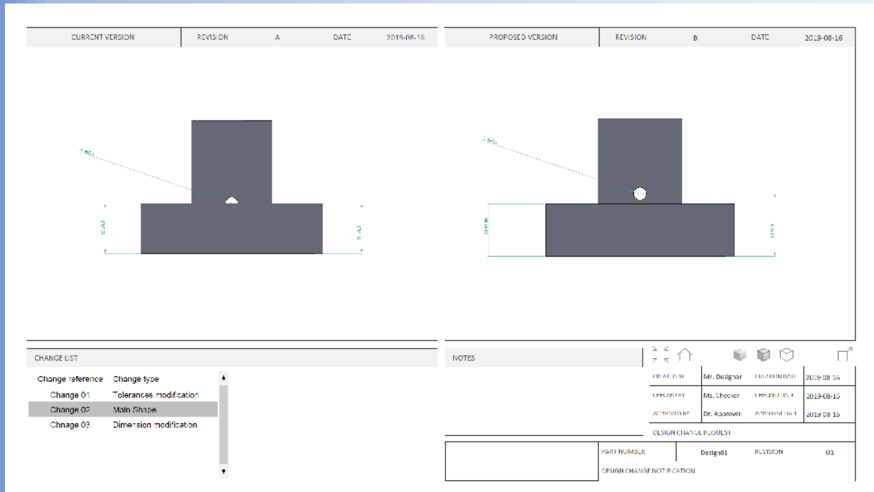

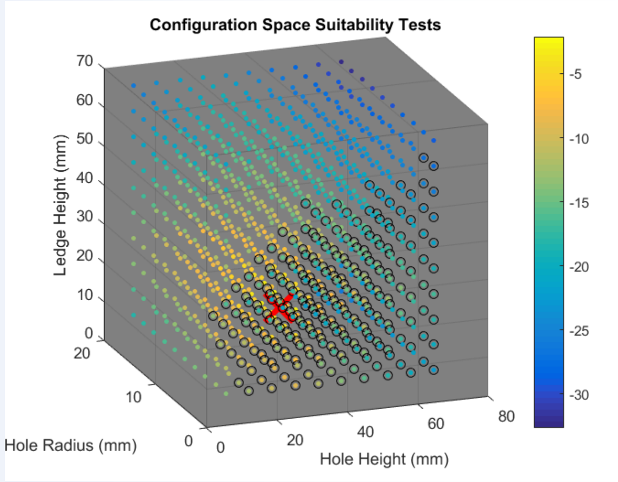

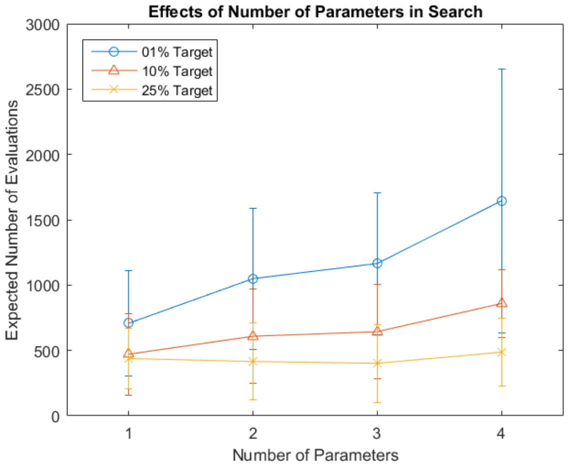

National Institute of Standards and Technology
U.S. Department of Commerce

Where has IAI already been successful?

- Success Story – Augmented Intelligence – Industrial Augmented Reality (AR)
  - **Goal –** Use computer vision to inform a worker in real time of critical information on products or assets
  - **Challenge** – To visualize product or asset data in real-time mapping on complex shapes and difficult materials (Ex. reflective Aluminum)
  - **Method** – Use existing standards (STEP AP242) and existing AR software to create full annotation overlay images with product information

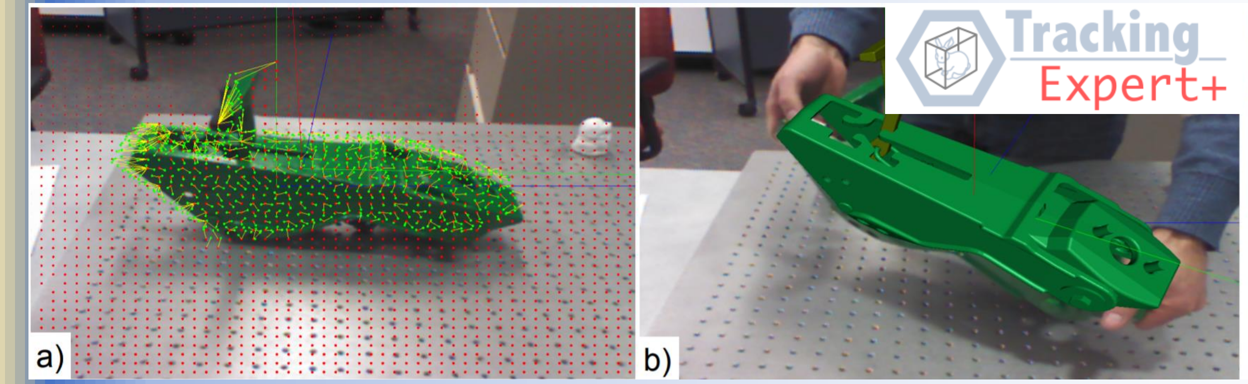

An Open-Source Solution for Augmented Reality Object Detection and Pose Tracking

https://github.com/usnistgov/TrackingExpertPlus

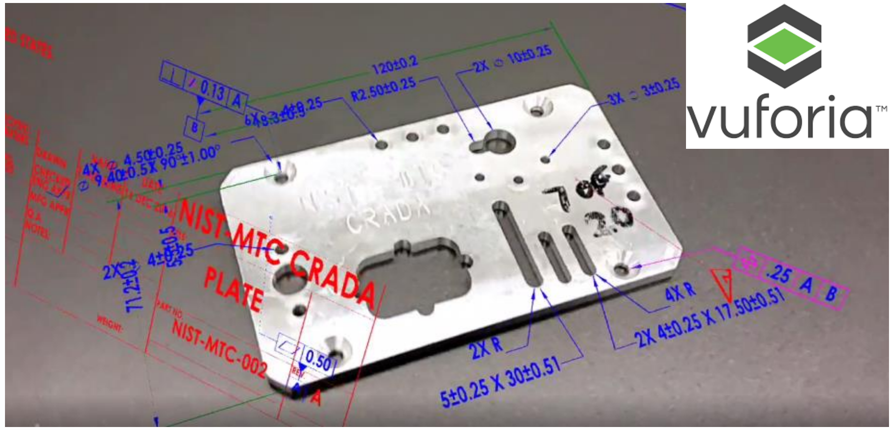

Spatially orientating product manufacturing Information from native standard models in AR

For more info contact:
william.bernstein@nist.gov

Misapplications of IAI :

When IAI looks good at first, but doesn't hold up to closer inspection

- It is easy to be mislead by early results
  - Why did it look good?
  - What went wrong?
  - How could it have been better?

National Institute of Standards and Technology
U.S. Department of Commerce

Misapplications of IAI : Using NLP to Estimate Repair Time

- Application
  - Predicting Time to Repair Based on Language in Maintenance Request
- Positive Indicator
  - NLP identification of Words with strong correlation to repair time (R > .7)
- Negative Results
  - LASSO driven linear model to predict repair time
    - Mean error ~21 hours (excluding outliers)

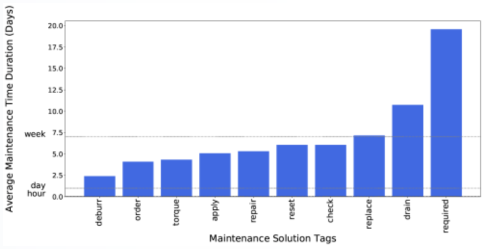

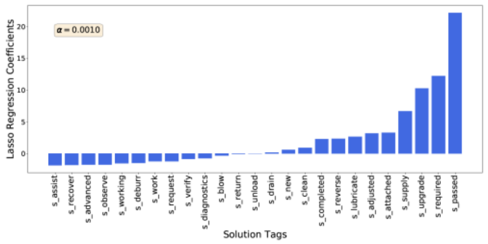

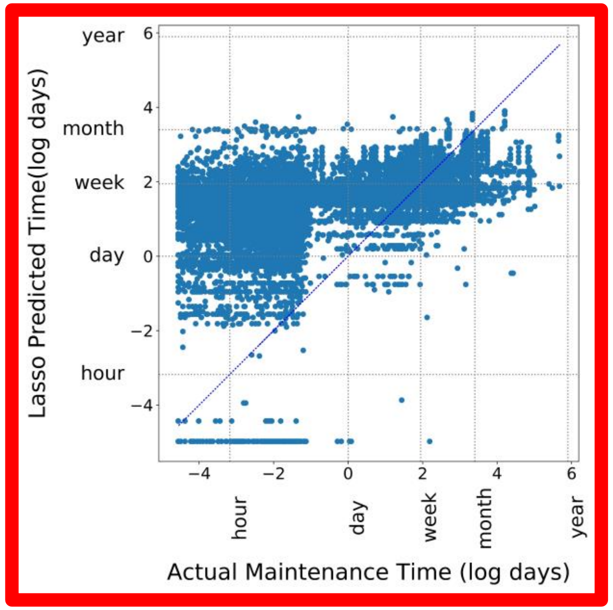

National Institute of Standards and Technology
U.S. Department of Commerce

Misapplications of IAI : Using NLP to Estimate Repair Time

- What went wrong?
  - Failed to account for uncertainty of data
    - Failed to account for complexity and influence of information not available in text input
  - Faulty Assumption
    - Linear addition of words
      - “Fix Left Headlight” = 5 min
      - “Fix Right Headlight” = 45 min
- What was happening?
  - Identification of words that loosely indicate ‘Long job’ vs ‘Short job’
  - Inadvertent Classifier: ~65% Accuracy
- What should have been done?
  - Identification of important phrases and concepts
  - Identify only broad categorical or conditional trends

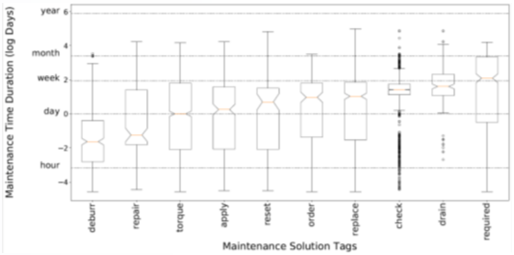

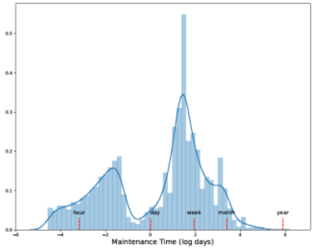

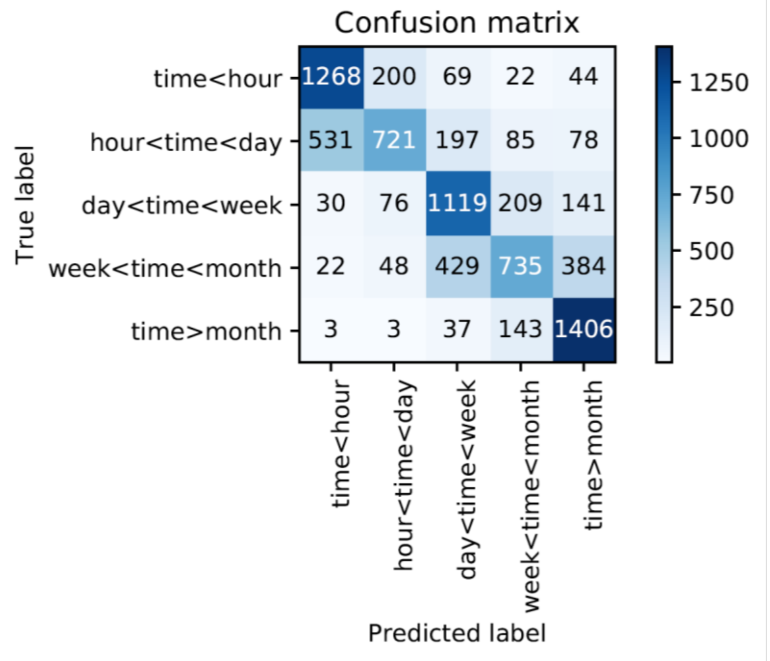

National Institute of Standards and Technology
U.S. Department of Commerce

Misapplications of IAI :

Using Deep NNs to Identify Faulty Products

- Application
  - Let big data techniques determine salient features for anomaly identification during battery production
- Positive Indicator
  - Quality prediction model developed
  - Fantastic Testing Metrics!
    - False Positive indications 0% total production
    - False Negative indications <20% total production
    - True Positive = 100%!!!
- Negative Results
  - Failed to maintain consentient results with extended data sets

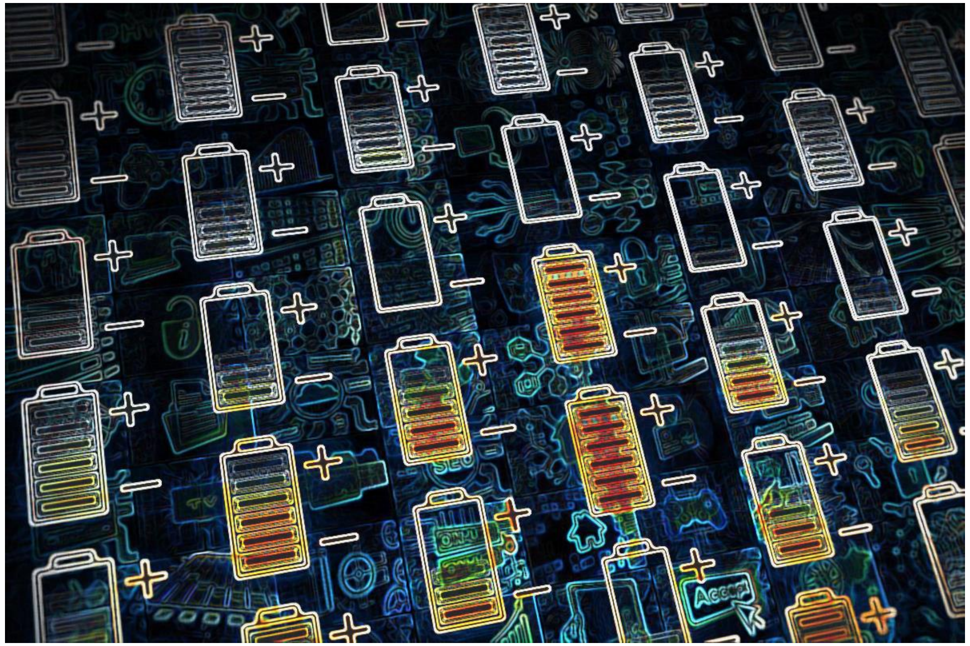

National Institute of Standards and Technology
U.S. Department of Commerce

Misapplications of IAI :

Using Deep NNs to Identify Faulty Products

- What went wrong?
  - Data Imbalance
    - Less than 0.001% of the development data was labeled “faulted”
  - Did not check P significance of the trial
- What was happening?
  - Randomly selecting ~20% (the false positive rate) of the data has a greater than 1 in 5 chance containing all faulted products
- What should have been done?
  - Rebalance development data to ensure more reliance in learning “bad” examples
  - Compare performance to simpler models

Misapplications of IAI :

Using Industrial AI to Emulate Unknown Physics

- Application
  - Using deep encoders to implicitly reproduce a transfer function to monitor physical phenomena
- Positive Indicators
  - Ability to recreate "similarly shaped" frequency spectrum to those captured in training
  - 100% classification of operating condition!
- Negative Results
  - Inability to generalize across operating conditions
  - Can not account for "quick" events in the data

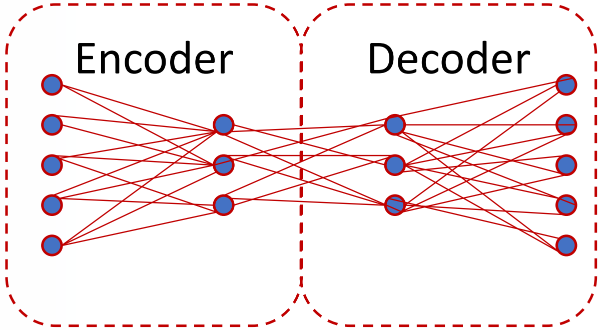

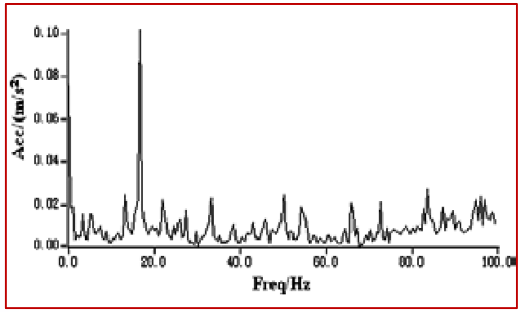

Misapplications of IAI :

Using Deep Encoders
to Emulate Unknown Physics

- What went wrong?
  - Trained on too few operating conditions
  - Had no physically backed “ground truth” target
  - No quantified measures for success
- What was happening?
  - Model was outputting an “average” rather than active spectrum
  - Broad operating conditions were easily distinguished
  - Assuming the model would learn what was “important”, without means to test that
- What should have been done?
  - Better “ground truth” goal posts
  - Broader testing on “unseen” operating conditions

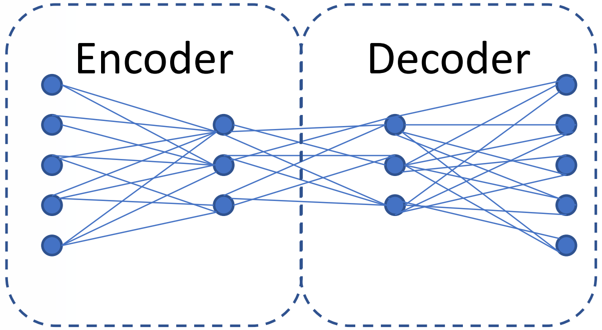

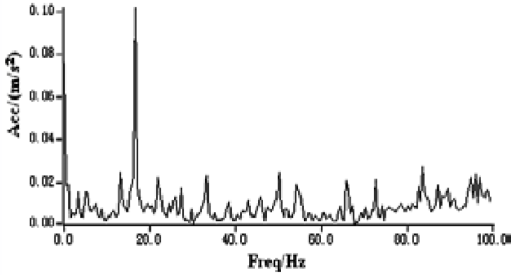

Lessons Learned: Common Themes

- Faulty Assumptions
  - Data
  - Model
  - Physical Setup
- Lack of Interrogation
  - Assuming output is correct or "good"
  - No context for performance metrics
- Loosing Sight of the End Goal
  - Not asking how the tool will be used
  - Neglecting to verify usefulness of tool
  - Ignoring real world limitations / requirements
- **Trying to remove Human knowledge from the model**
  - Developing purely data driven models for the sole reason of "elegance" or "simplicity"
  - Ignoring domain knowledge or well-established heuristics during development and testing

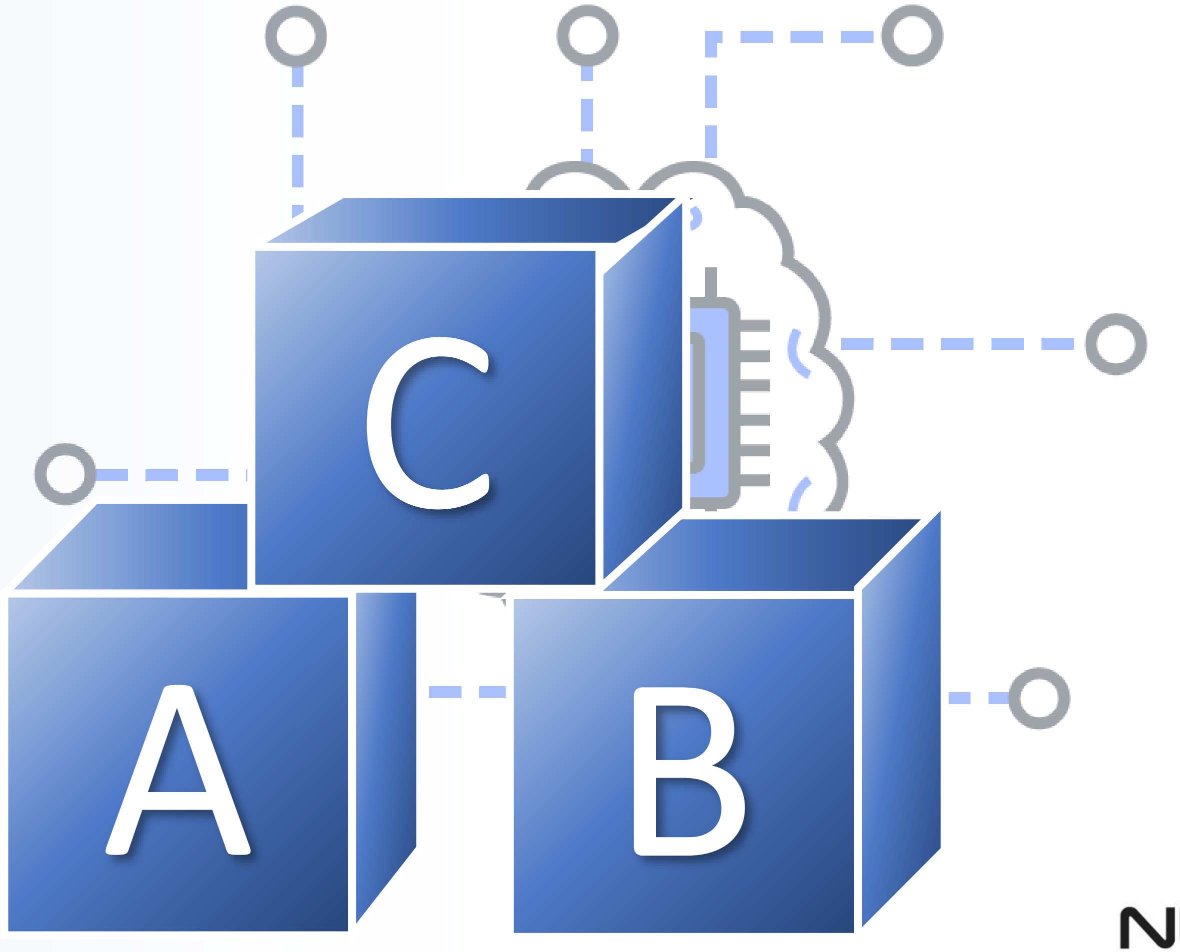

National Institute of Standards and Technology
U.S. Department of Commerce

Lessons Learned: Understanding the basics of your IAI

- Many misapplications of IAI could be avoided by remembering the four fundamental aspects of an IAI
  - The Need
    - Why is an IAI tool needed?
  - The Setup
    - What are my real-world limitations / assumptions?
  - The Model
    - What are my digital limitations / assumptions?
  - The User
    - How will this tool be used?

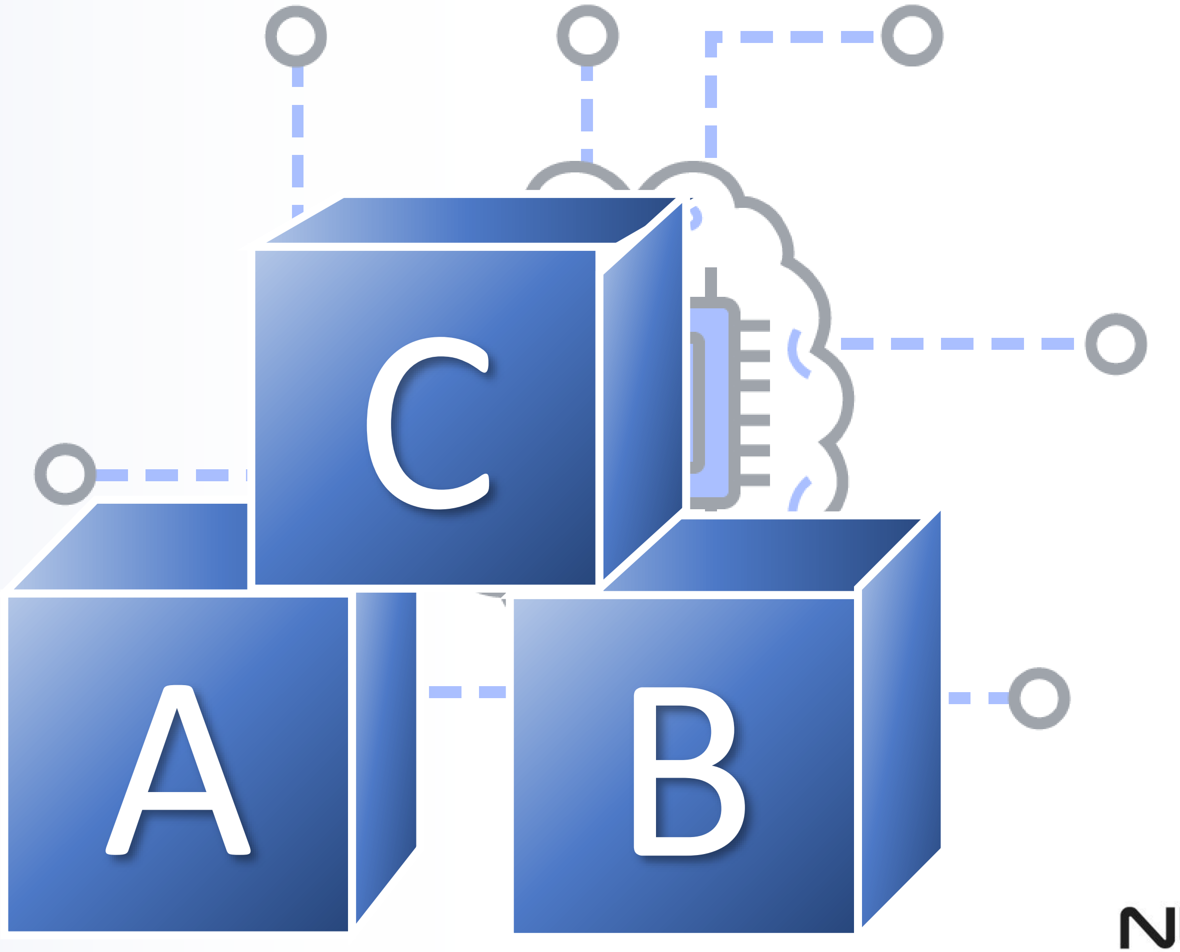

IAI Acumen : Ask about the basics

	ASK:	To Avoid:
Identify the Need	Why is AI needed for this problem?	• ***Needless Application of AI to a previously well understood problem*** • Producing output that provides no valuable contribution to real world users
Inspect the Setup	Has this technology been proven in my domain?	• Failure due to Complications of Scale: Timing, Volume, etc. • "False Equivalency" expectations of performance
	Is the training data relevant and appropriate to the real-world task?	• Trivial Output from Unrealistic Data • Unintended Bias from Imbalanced Data • Excessive Training Requirements from Superfluous Data
	Does the training data provide sufficient coverage / characterization?	• Increased Uncertainty from Sparse Data • Inability to Generalize from Lack of Variation • Narrow Windows of Reliable Output from Incomplete Data

IAI Acumen : Ask about the basics

	ASK:	To Avoid:
Inspect the Model	Is the model optimizing/training for the things I think it is?	• Under/Overcomplexity of Setup • Trivial Results • Misleading Results from Solving the "Wrong Problem"
	Has the model had sufficient training / why did it stop training?	• Bad Generalization from Training 'Stagnation' • Early Stoppages from 'Timeouts' • Hypersensitivity to Input and Inconsistent Output
	Are any model / domain assumptions being violated?	• Physically Meaningless Results • Wrong Model / Hyperparameters Selection • Trivialization or Overcomplication by Ignoring Application Domain

IAI Acumen : Ask about the basics

Inspect the Output

ASK:	To Avoid:
Are the relationships being captured sensible?	• Random Occurrence Modeling • Influence From Strong Outliers • Correlation vs Reverse Causation

Is this model over fitting data?	• Inability to Generalize due to
	  • Excessive Complexity
	  • ***Testing Data Too Similar to Training***
	  • Learning Superfluous Patterns

Am I looking at ***all*** the appropriate performance metrics?	• Unrealistic Expectations Based on Misleading Metrics or Incomplete Metrics • Using System Outside Qualified Bounds • Misinterpretation of Output

National Institute of Standards and Technology
U.S. Department of Commerce

10 Common Pitfalls of IAI Modeling

1. **The 'No value added' problem**
   - ***Needless Application of AI to a previously well understood problem***
   - Producing output that provides no valuable contribution to real world users
2. **Proving / Development on Systems that are Not (Sufficiently) Comparable**
   - Doesn't work at your application scale - Timing, Volume, etc.
   - Relying on application simplifications during development not present in your real-world use case
3. **Inappropriate Training Sets**
   - Data unrealistic / unavailable to real-world application
   - Biased / Imbalanced Data
4. **Inadequate Training Sets**
   - Incomplete / Too Sparse
   - Too Few Exemplars – NOTE: Lots of Data ≠ Lots of Examples
5. **Wrong "Question" Setup**
   - Ignorance of Domain – Overcomplication / Misspecification of Setup
   - Assumption / Simplification that Trivializes Results
   - Optimizing the wrong or an incomplete set of goals (Bad Cost function)
6. **Premature Training Termination**
   - Not Enough Cycles
   - Bad Scaling
7. **Violation of Model Assumptions**
   - Wrong Model /Bad Hyperparameters
   - Physically Impossible Solutions
   - Outside of Training Bounds
   - Memory vs Memoryless Conflicts
8. **Identification of Spurious Correlations**
   - Random Occurrence
   - Strong Outliers
   - Correlation vs Reverse Causation
9. **Over Fitting Data**
   - Overly Complex Model
   - ***Testing Data Too Similar to Training***
   - No Relevant Patterns Captured – TEST: Is it better than random guessing?
10. **Misinterpreting Performance Metrics**
    - Over Emphasizing Incomplete or Misleading Metrics
    - Omitting Key Metrics
    - Not Providing Bounds for Qualified Use

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

The Path to Action

- Now that we know how to check and trust IAI, How do we know where to deploy it?

- Not everything needs the same levels of autonomy
  - Zero Automation
    - Highly singular events – "One-off tasks"
    - Low volume / infrequent events
  - Human in the Loop
    - High Importance tasks
    - Events with lots of exceptions or outliers – "Edge-cases matter"
  - Full Autonomy
    - Highly repetitive or computer coordinated operations
    - Trivial events with low consequences for mishandling

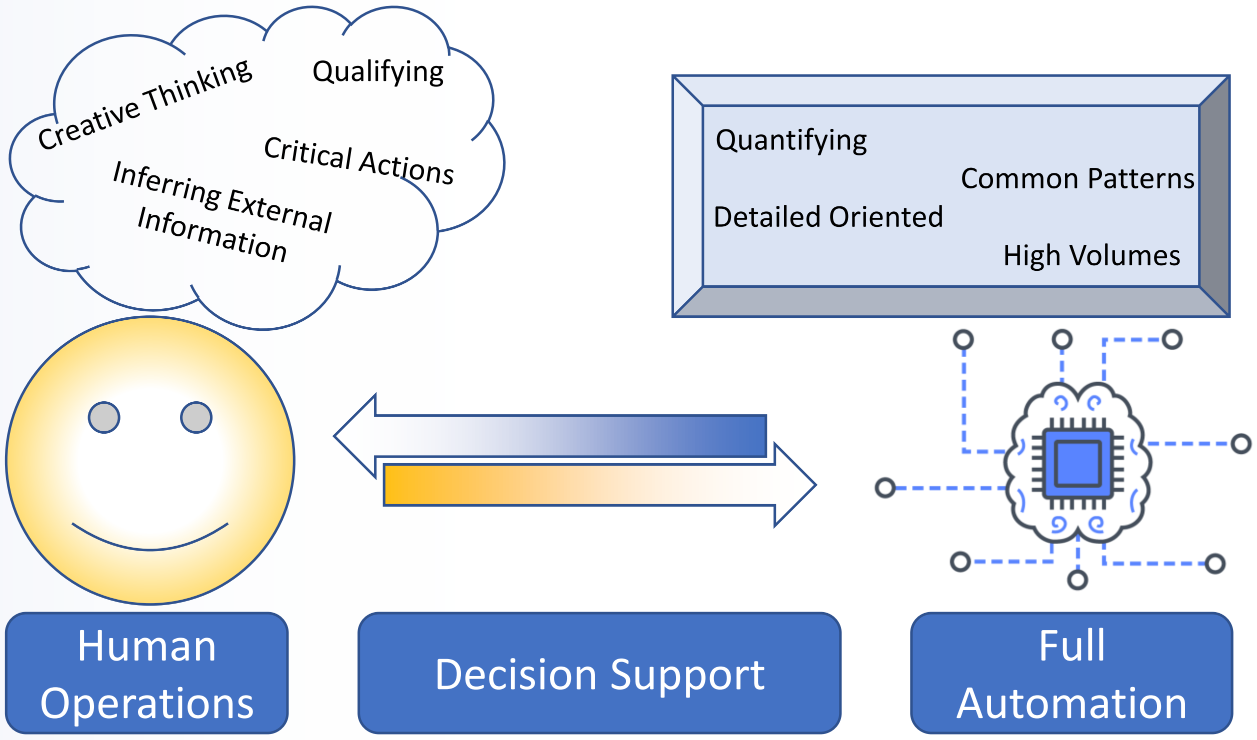

The Value of IAI – Can it lower risks?

- The most ubiquitous factor in decision making for industry is mitigating risk
  - Risk = Outcome * Frequency
  - Not all risks are bad
- To understand the value of an IAI you must understand the risks
  - Risks from the problem
  - Risks from the IAI
  - Risks from competing solutions
- Risk can be translated to many important metrics
  - Human / Environmental Safety
  - Production Output
  - Money Lost / Saved / Earned

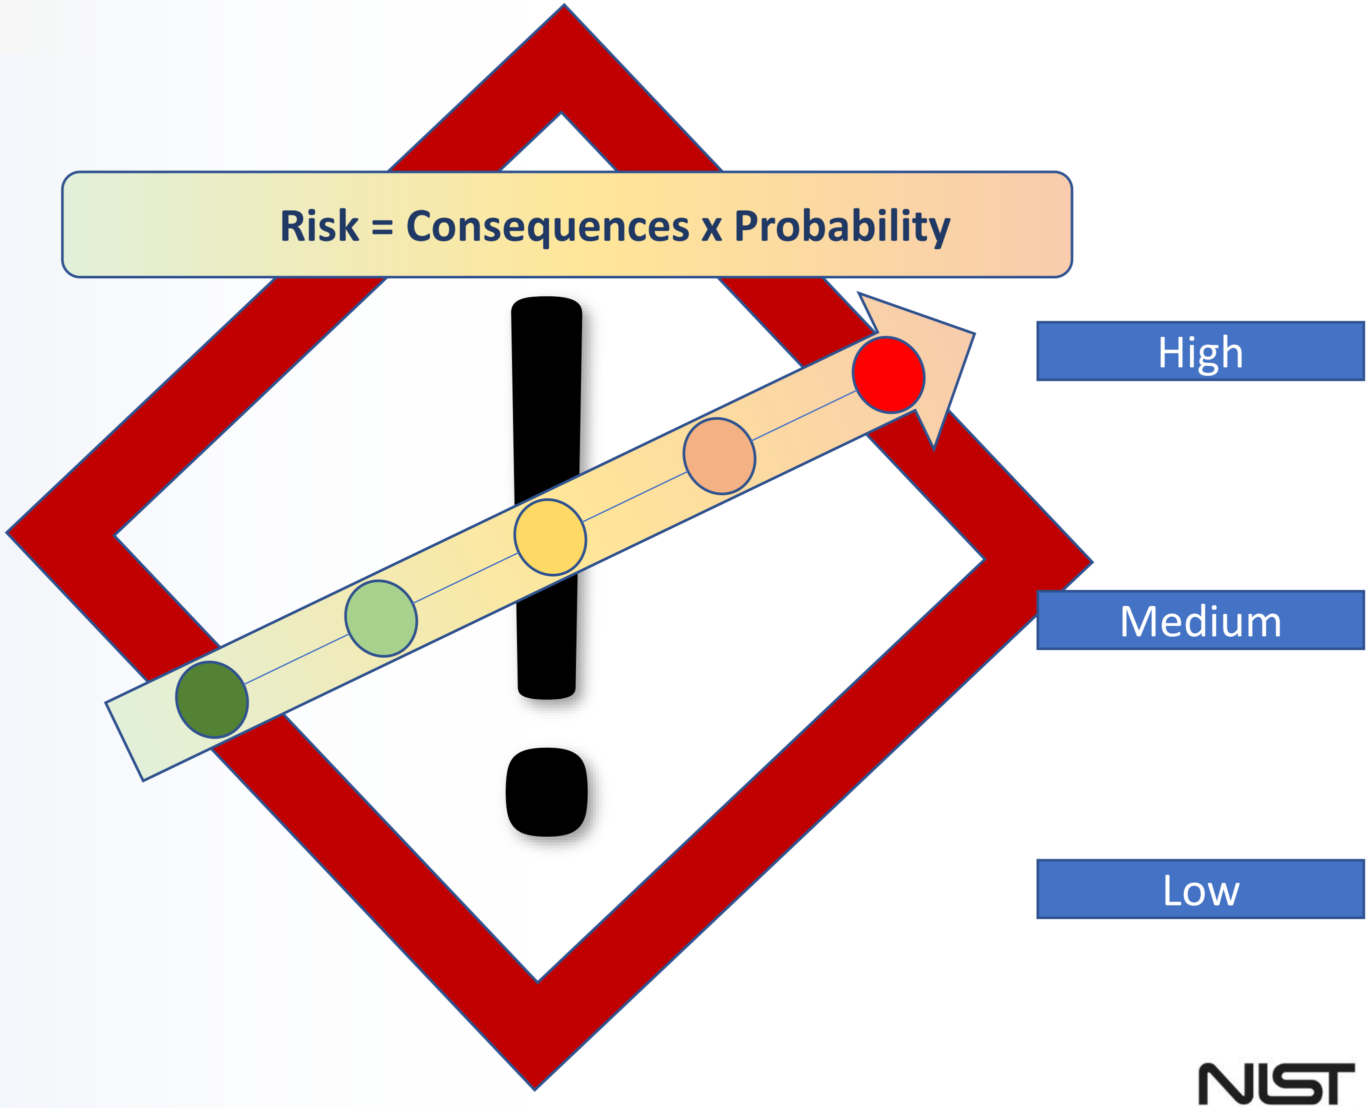

Is IAI Right for the Job? – Look Ahead

- When Considering IAI, Evaluate the Potential Risks
  - What are potential costs ?
  - What are the possible positive and negative outcomes?
  - What is the expected frequency or probability of each scenario?
- Translate these scenarios into meaningful metrics to help make informed decisions
  - Expected Production Quality
  - ROI
  - Facility Efficiency
  - etc

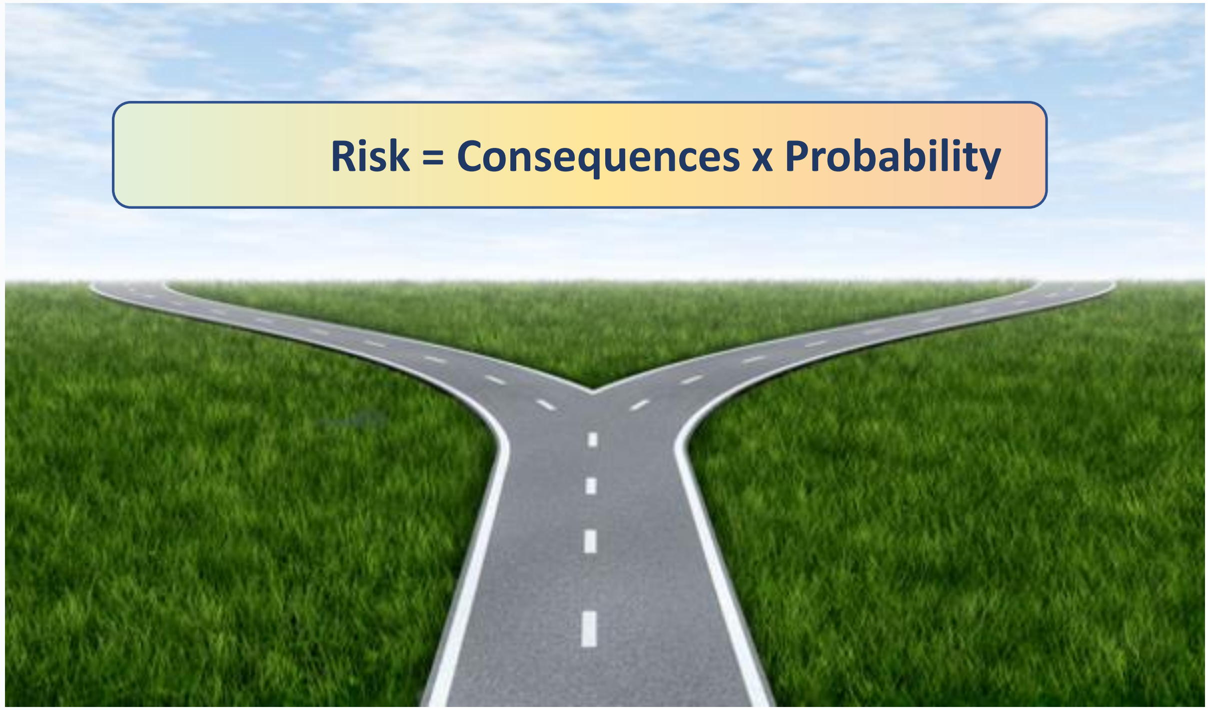

Barriers to IAI Use

- Needs:
  - Better State of the Practice Survey
    - Define broad classes of domain challenges
      - Coalesce common cross domain themes to coordinate research efforts
    - Define testing & evaluation best practices
      - Intuitive and Meaningful Metrics
  - Develop community driven domain-centric focus groups
    - More Domain Experts Comfortable with AI
    - Creation of Domain Specialized Plans and Tools
    - Ex. TLP - COI
  - Data Handling and Communication Standards
    - Data Provenance
    - Formatting
    - Meta Data and Storage

What Comes Next?

- Community driven results!
  - More Domain Experts Comfortable with AI
    - Stakeholder Education
    - Foster Trust with Measures of Effective Use
  - Do Not wait for new AI tools
    - Find the tools that already exist and can work for your needs
  - Lots of work!
    - Open-Source Tools
    - Public / Benchmark Data Sets
    - Industry Leaders Understanding the Capabilities and Limitations of IAI Tools

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Designing with Humans in Mind

- **IAI is *a Tool* that works *for* You!**
  - Like any tool, IAI works best when used by people who understand the Problem – Not the person who makes the tools
  - No one tool is right for every job – evaluating the efficacy of an IAI tool requires someone with intimate knowledge of both the problem and the application

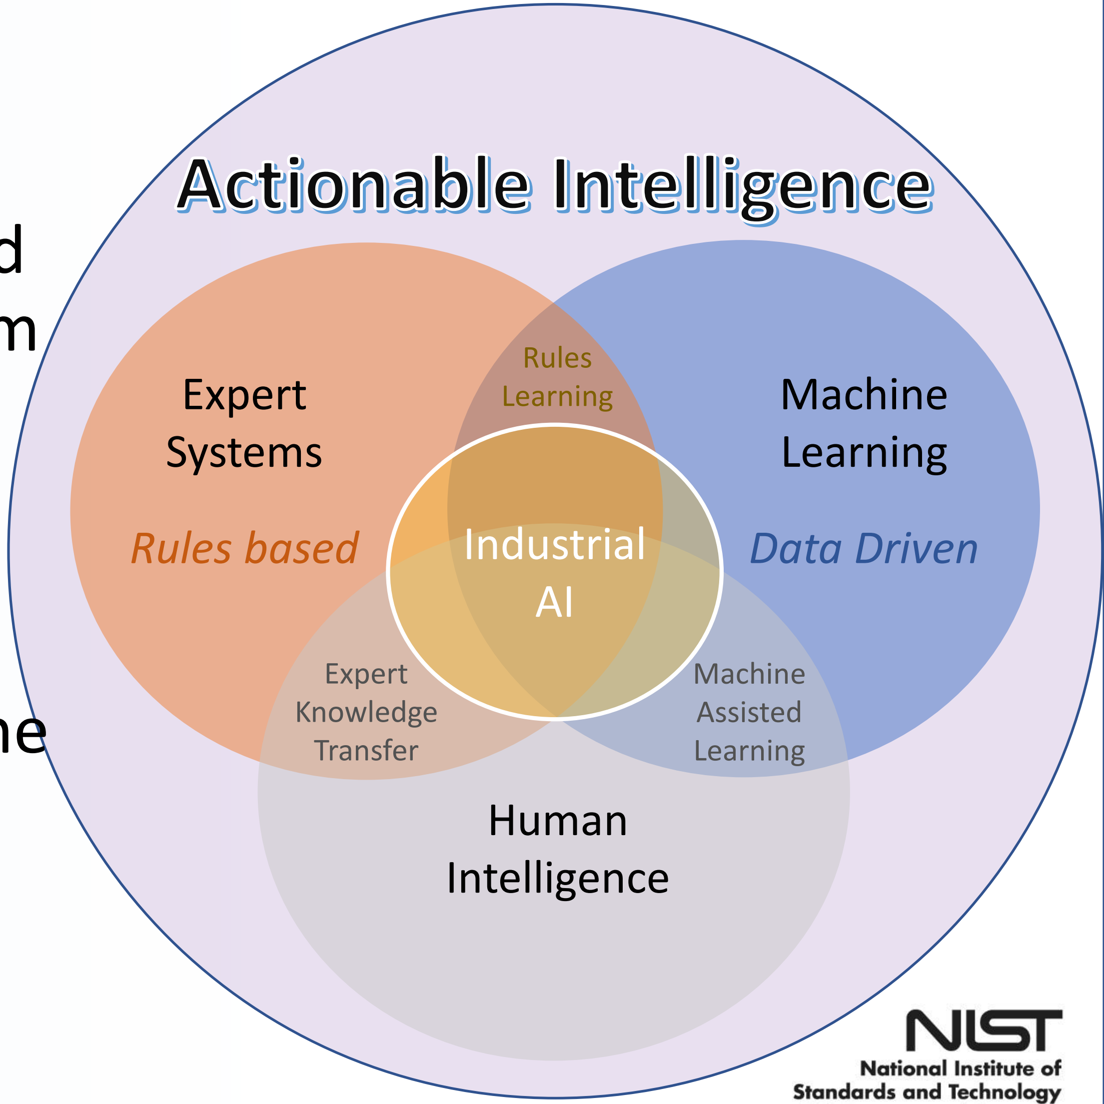

National Institute of Standards and Technology
U.S. Department of Commerce

Designing with Humans in Mind

- **IAI *is a Tool* that works *with* You!**
  - The barriers to learning and using IAI tools are lowering every day
  - Expert AI researchers and developers are creating new and better tools every day promising to solve challenges in any or every domain

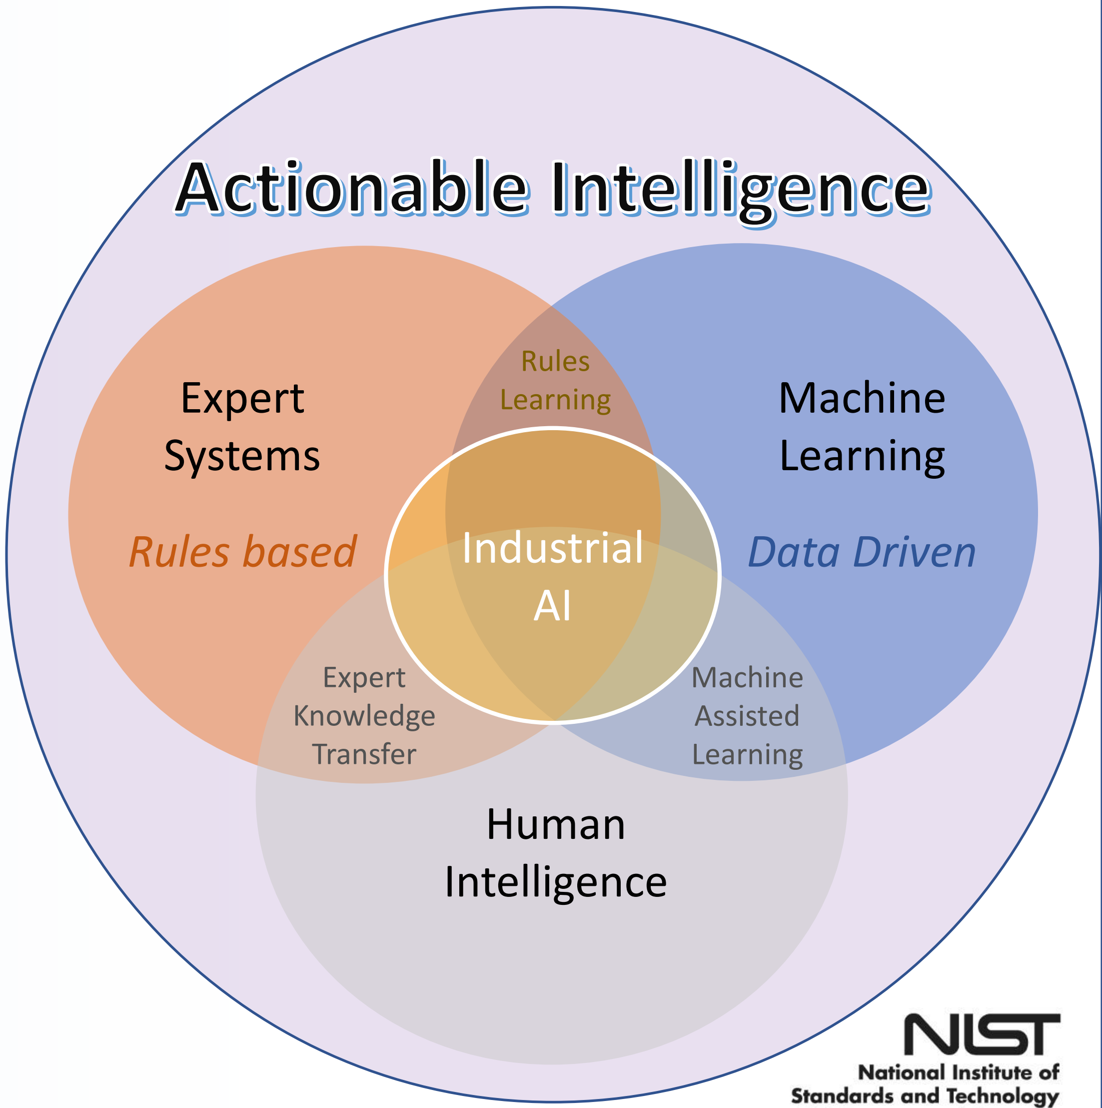

National Institute of Standards and Technology
U.S. Department of Commerce

Designing with Humans in Mind

- **IAI is *a Tool* that *needs* You!**
  - Industry needs more domain experts and industry practitioners to take the lead in using these tools in meaningful and productive ways
  - It is easier to train a domain expert to use / evaluate an IAI tool than to train an AI expert on the intricacies of a domain challenge

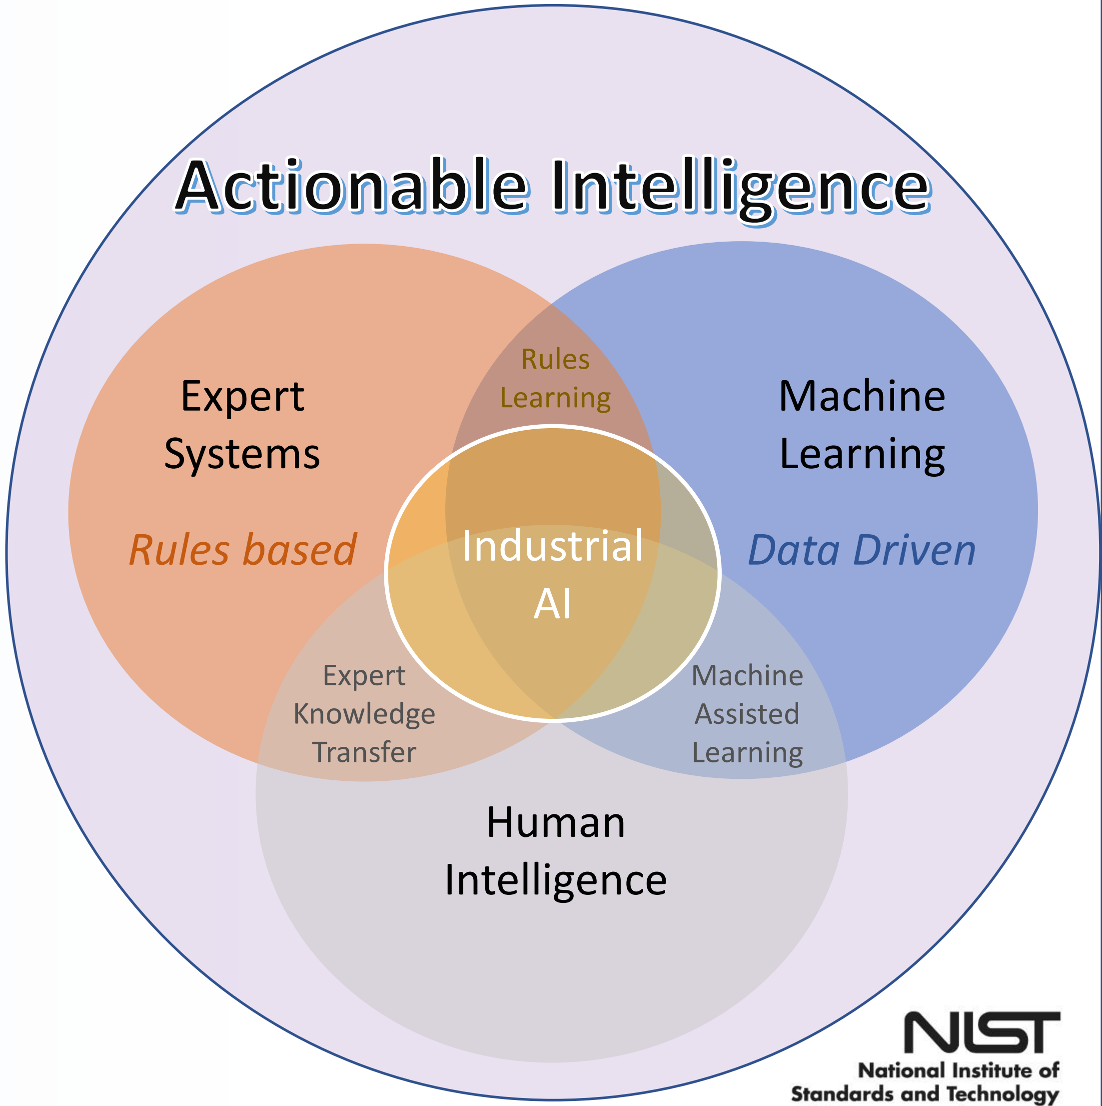

National Institute of Standards and Technology
U.S. Department of Commerce

The Road Ahead

- Active Research Areas at NIST
  - Technical Language Processing
    - Develop best practices guidelines on how to tailor tools for engineering text-based data
      - Capitalize on unique needs and restrictions of technical language
    - https://www.nist.gov/el/technical-language-processing-community-interest
  - Industrial AI Performance Management
    - Qualify and evaluate AI/ML tools and methods for industry
      - Domain aware testing
      - Intuitively informative metric generation
    - Identify, examine, and alleviate barriers to adoption
      - Trust building
      - Identification of open-source tools
      - Increasing usability
  - Open-Source Toolkits
    - TrackingExpert+
      - Apply AI/ML for object detection and Augmented Reality pose tracking in industrial scenes
      - https://github.com/usnistgov/TrackingExpertPlus
    - NESTOR
      - Applies AI/ML to annotate and tag maintenance work order data for structured data extraction
      - https://www.nist.gov/services-resources/software/nestor

National Institute of Standards and Technology
U.S. Department of Commerce

Thank you for attending!

Contact: michael.sharp@nist.gov

Dr. Michael E. Sharp is a Reliability Engineer at the National Institute of Standards and Technology (NIST) located in Gaithersburg, MD. He received a B.S (2007), M.S. (2009), and Ph.D. (2012) in Nuclear Engineering from the University of Tennessee, Knoxville, TN, USA. His research interests include signal analytics, machine learning, artificial intelligence, optimization, and natural language processing. Michael has worked on a wide array of projects including image processing for elemental material recognition, navel reliability monitoring, and manufacturing robotics diagnostic monitoring. He currently works with the NIST Systems Integration Division for Smart manufacturing.

National Institute of Standards and Technology
U.S. Department of Commerce