

Understanding and Leveraging Today's Artificial Intelligence and Machine Learning




Michael Garris

Senior Scientist / National Institute of Standards and Technology


Co-Chair / National Science and Technology Council,
Subcommittee on Machine Learning and Artificial Intelligence



What's all the fuss?



Emerging Technology



Emerging Technology



Emerging Technology




Emerging Technology



Emerging Technology



Emerging Technology



Emerging Technology



What is AI?



What is AI?


- Think and act like humans
- Think and act rationally
- Sometimes mutually exclusive. 😊
- AI Problem Space Categories
 - Knowledge Representation
 - Perception
 - Logical Reasoning
 - Planning and Navigation
 - Prediction

What is AI?


- Field of Study
- Technology
 - System
 - Characteristic

AI Conceptualization (1 of 3)


Embodied



Embedded



AI Conceptualization (2 of 3)



* DNN = Deep Neural Network

AI Conceptualization (3 of 3)

Narrow AI	General AI
○ Application specific/ task limited	○ Perform general (human) intelligent action
○ Fixed domain models provided by programmers	○ Self-learns and reasons with its operating environment
○ Learns from thousands of labeled examples	○ Learns from few examples and/or from unstructured data
○ Reflexive tasks with no understanding	○ Full range of human cognitive abilities
○ Knowledge does not transfer to other domains or tasks	○ Leverages knowledge transfer to new domains and tasks
○ Today's AI	○ Future AI?

New Wave of AI

- Availability of Big Data
- Improved Machine Learning (ML) Algorithms
- More Powerful Computing
- Mobile Connectivity



What can ML/AI do?



ML/AI for ...

- Classification
- Function Approximation
- Prediction
- Control
- Simulation
- Anomaly Detection
- ...

ML/AI for Manufacturing (1 of 2)

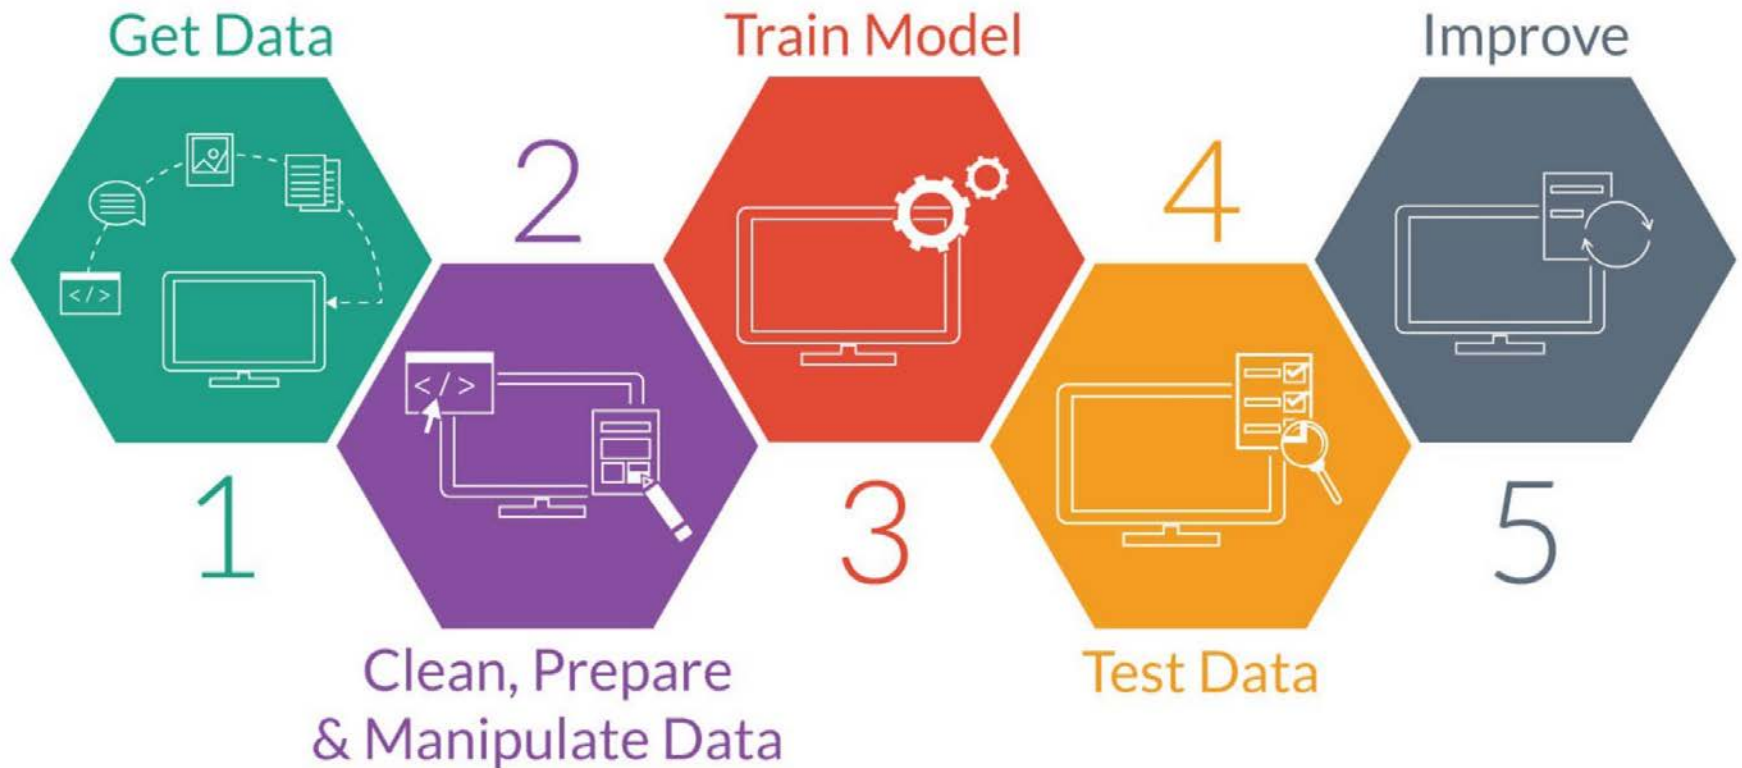
- Advanced Data Analytics*
 - Predictive Maintenance
 - Yield, Throughput, Resource Efficiency
 - End-to-End Systems Optimization
 - Including Integrated Supply Chain

* <http://www.mckinsey.com/business-functions/operations/our-insights/manufacturing-analytics-unleashes-productivity-and-profitability>

ML/AI for Manufacturing (2 of 2)


- Human Assistive Technologies
 - Automated Inspection and Quality Control
 - Collaborative Robots
 - Virtual and Augmented Reality
 - Training and enhanced context awareness
 - Workforce Knowledge Base
 - ...

Machine Learning Workflow



<https://machinelearning-blog.com/2017/11/19/fsgdhfju/>

1980's Neural Network



Today's (Deep) Neural Networks




Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

ML/AI Challenges

- Probabilistic
 - With inherent error rates and uncertainty
- Data Driven
 - With vulnerability to learning unwanted patterns of bias
- Extreme Dimensionality
 - Making decisions opaque (little explainability)
- Model Validation
 - Bound performance
 - Account for unforeseen outliers
 - How good is good enough?
 - Now introduce dynamic / continuous learning!?*

Summary: Today's ML/AI ...

- Powerful Tool
- Great Impact on Application Domains
- Challenges and Limitations
- Life Cycle Costs



Questions?

Michael Garris
(mgarris@nist.gov)

NIST