8 Myths and Truths about AI/ML and MOSA

Demystifying and Simplifying the AI/ML Domain for MOSA Implementations

Dr. Mark Vriesenga
Chief Engineer, Information Analytics
Global Engineering Fellow
BAE Systems, Mission Solutions

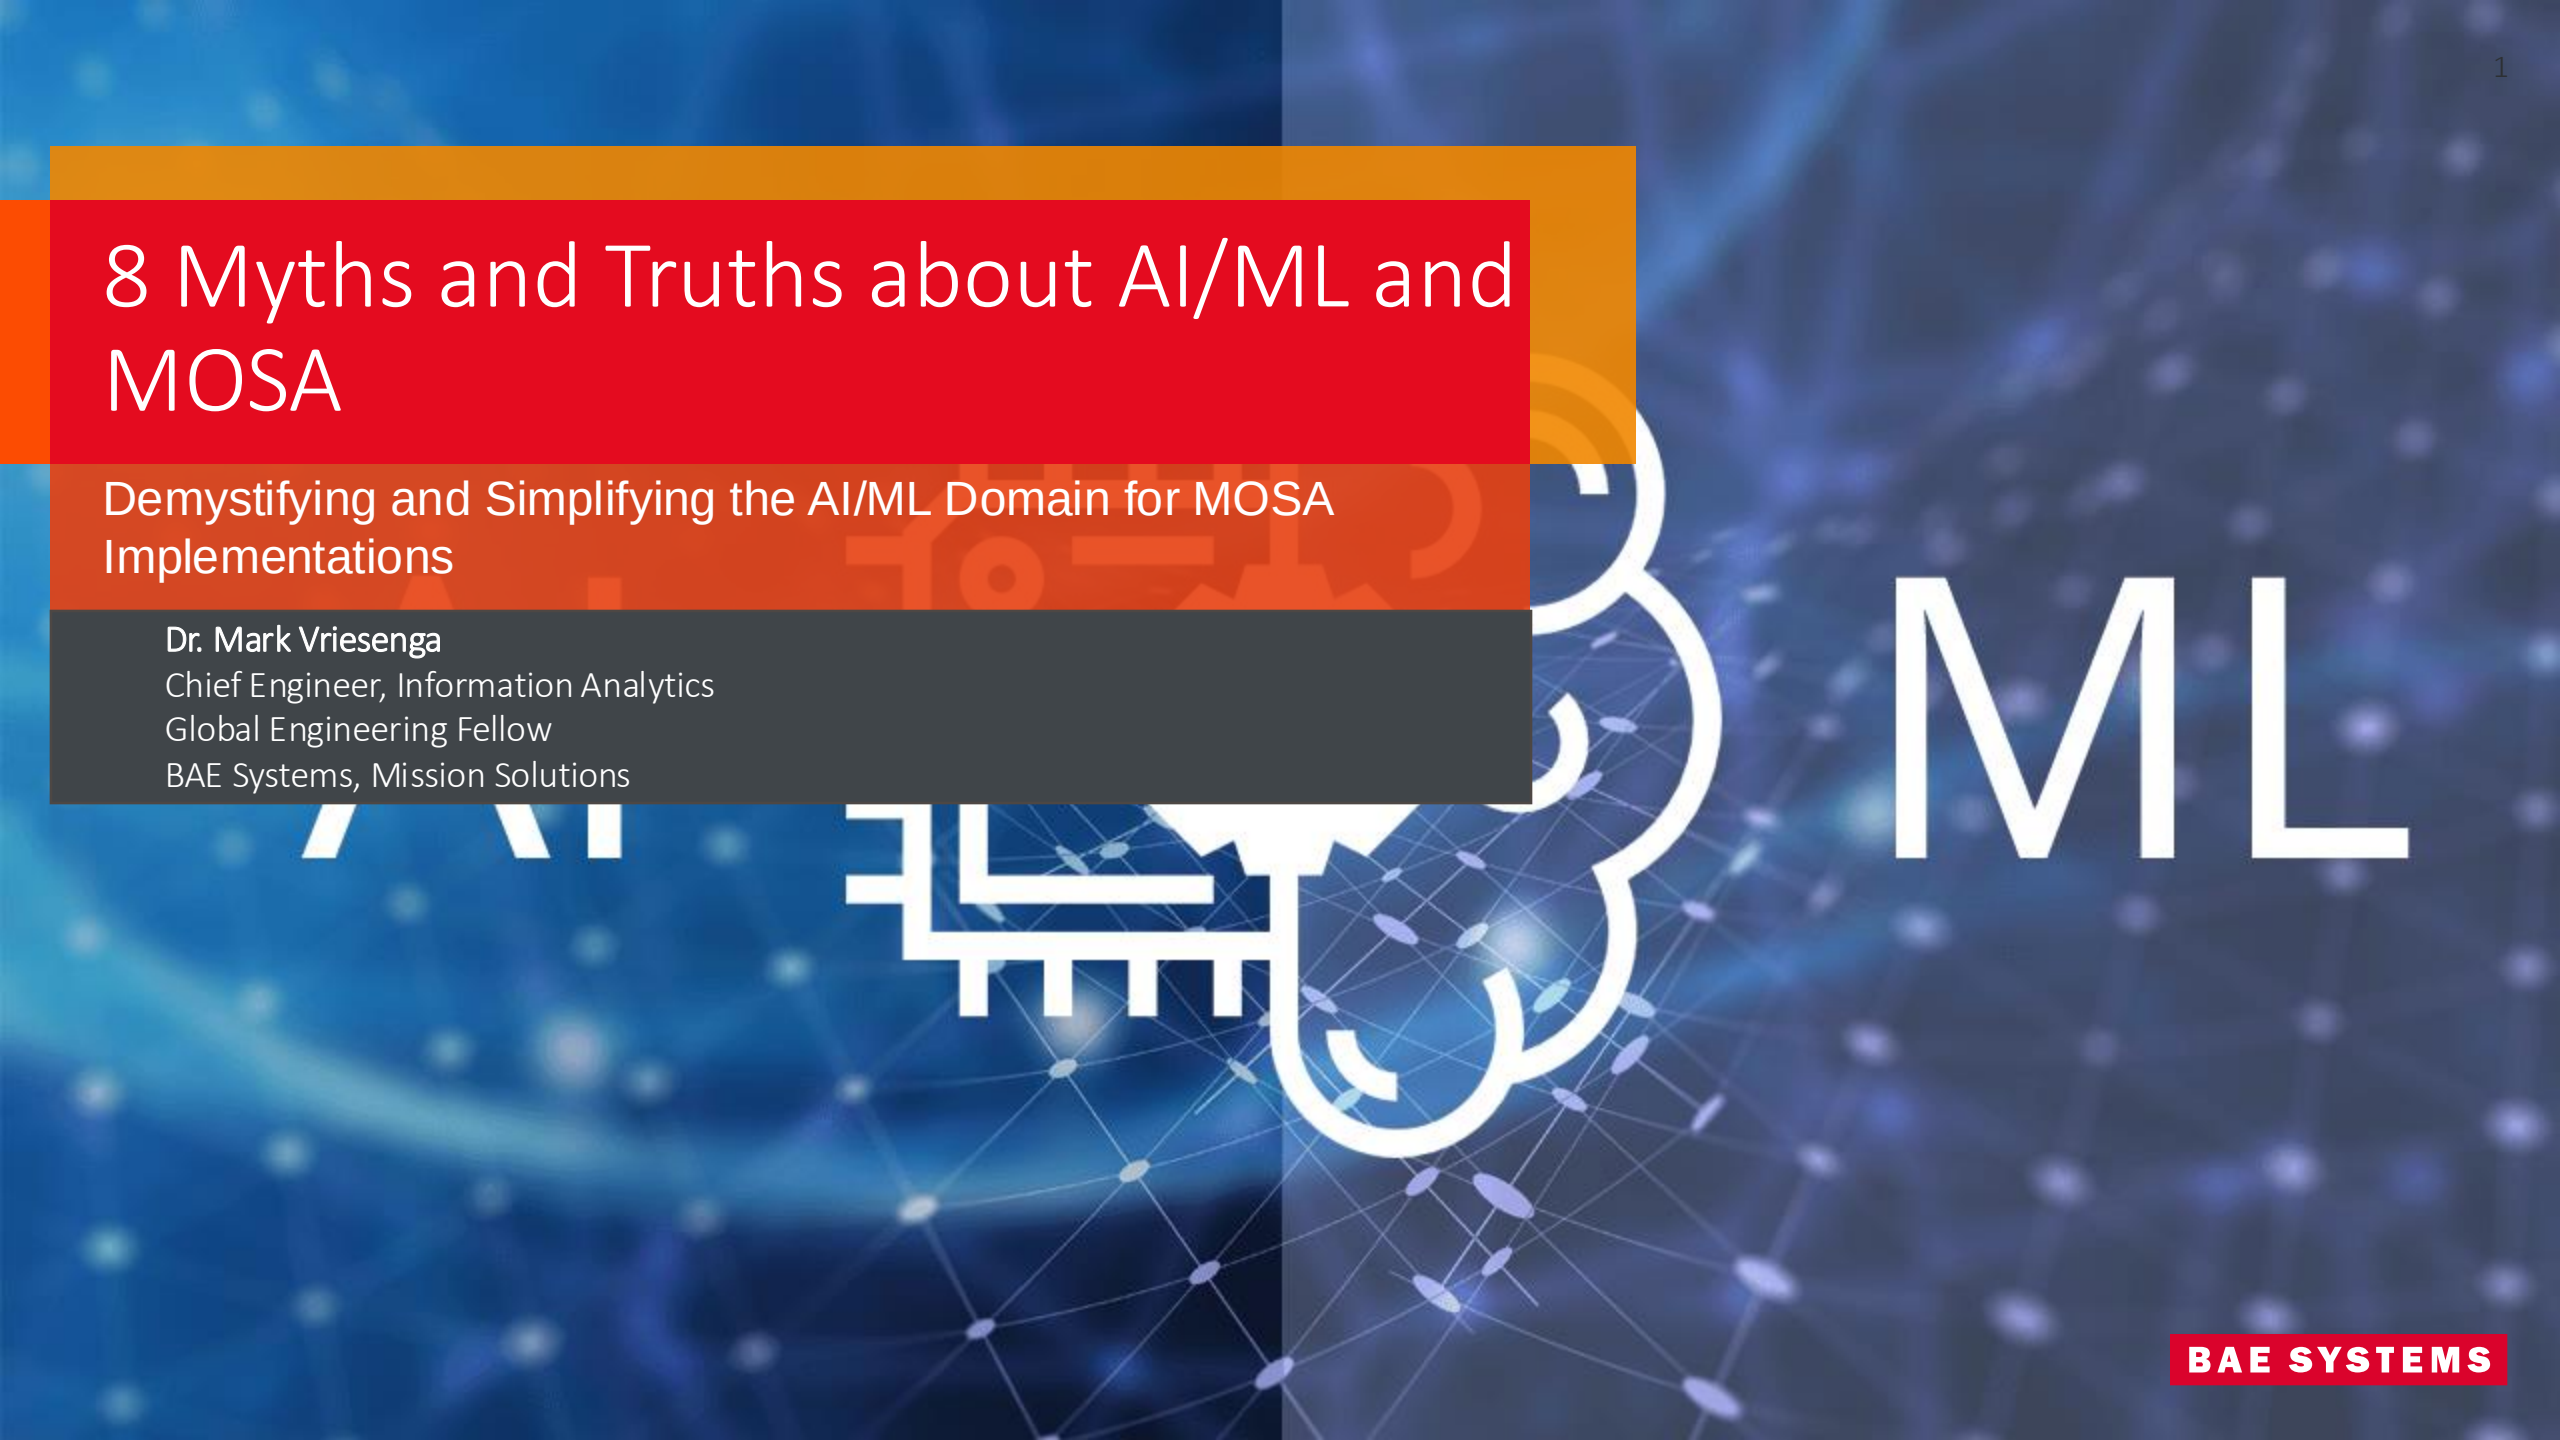

BAE SYSTEMS

Presenter Bio

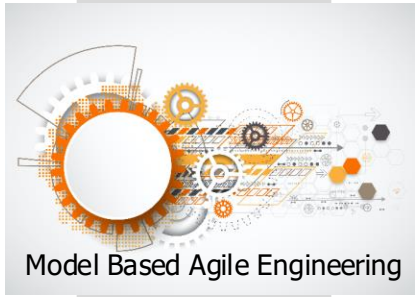

- **Dr. Mark Vriesenga**
  Chief Engineer, Information Analytics
  Global Engineering Fellow
  BAE Systems, Mission Solutions
- **Background & Experience**
  - 29 years at BAE Systems
    - Algorithm Developer (2 years)
    - Chief Engineer, Advanced Programs (5 years)
    - Business Development (3 years)
    - Strategic Program Capture (4 years)
    - I&S University Founder (6 years)
    - Cyber Resilience Capability Group (3 years)
    - Deputy Director, FAST LABS Cyber Technology (1 year)
    - Model-Based Agile Engineering Capability Group (MBAE CG)
  - Offensive & Defensive Cyber Security (15 years as a SME)

BAE SYSTEMS

Agenda

- Insight Principles
- 8 Myths and Truths
  - AI/ML is New Technology
  - AI and ML are the Same Thing
  - AI/ML applies to any Problem
  - AI/ML extrapolates to produce new Facts
  - AI/ML algorithm selection is Unimportant
  - Chat GPT is “Thinking”
  - AI/ML can be Objective and Unbiased
  - AI/ML can be applied to Mission Critical Functions
- Next Steps for AI/ML MOSA
- Questions?

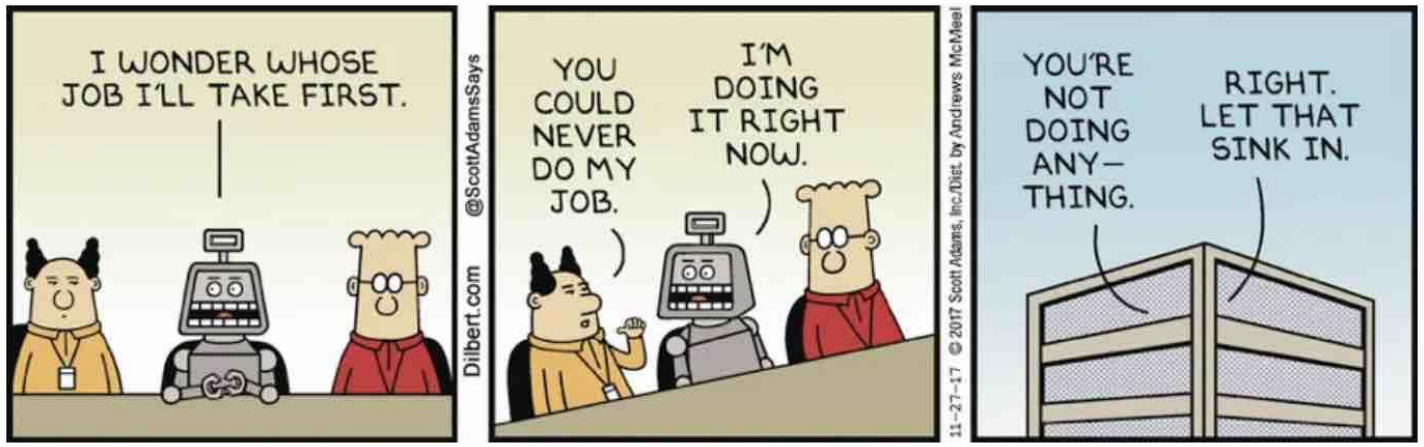

Image credit: Scott Adams

BAE SYSTEMS

Insight Principles

- Artificial Intelligence (AI) and Machine Learning (ML) are in vogue (again)
  - Customers are actively seeking answers on WHAT mission capabilities should be AI/ML enabled
- Advancements in AI/ML are driven by large-scale computing and abundant data
  - Cloud computing, GPU acceleration, and large data sets provide a novel compute foundation and enables historically compute-bound algorithms
- The true limits of what AI/ML CAN and CANNOT do is obscured by hype
  - Deep Fakes, Synthetic Art and Music, and Chat GPT interact at a human level and frequently “fool the observers”
- The merger of AI/ML and MOSA requires careful adjudication of stochastic and deterministic views, respectively

Insight Principles
Hype Cycle for AI/ML

- AI/ML is over 70 years old and remains an active area for research and development
- Deployed algorithms generally address well-bounded problems
  - Computer vision (food inspection, quality control)
  - Recommendation engines (online retailing)
  - Intelligence search (document classification, optimized search)
  - Data labeling (classification and labeling of data)
- "Big AI/ML" is likely 5-10 years out

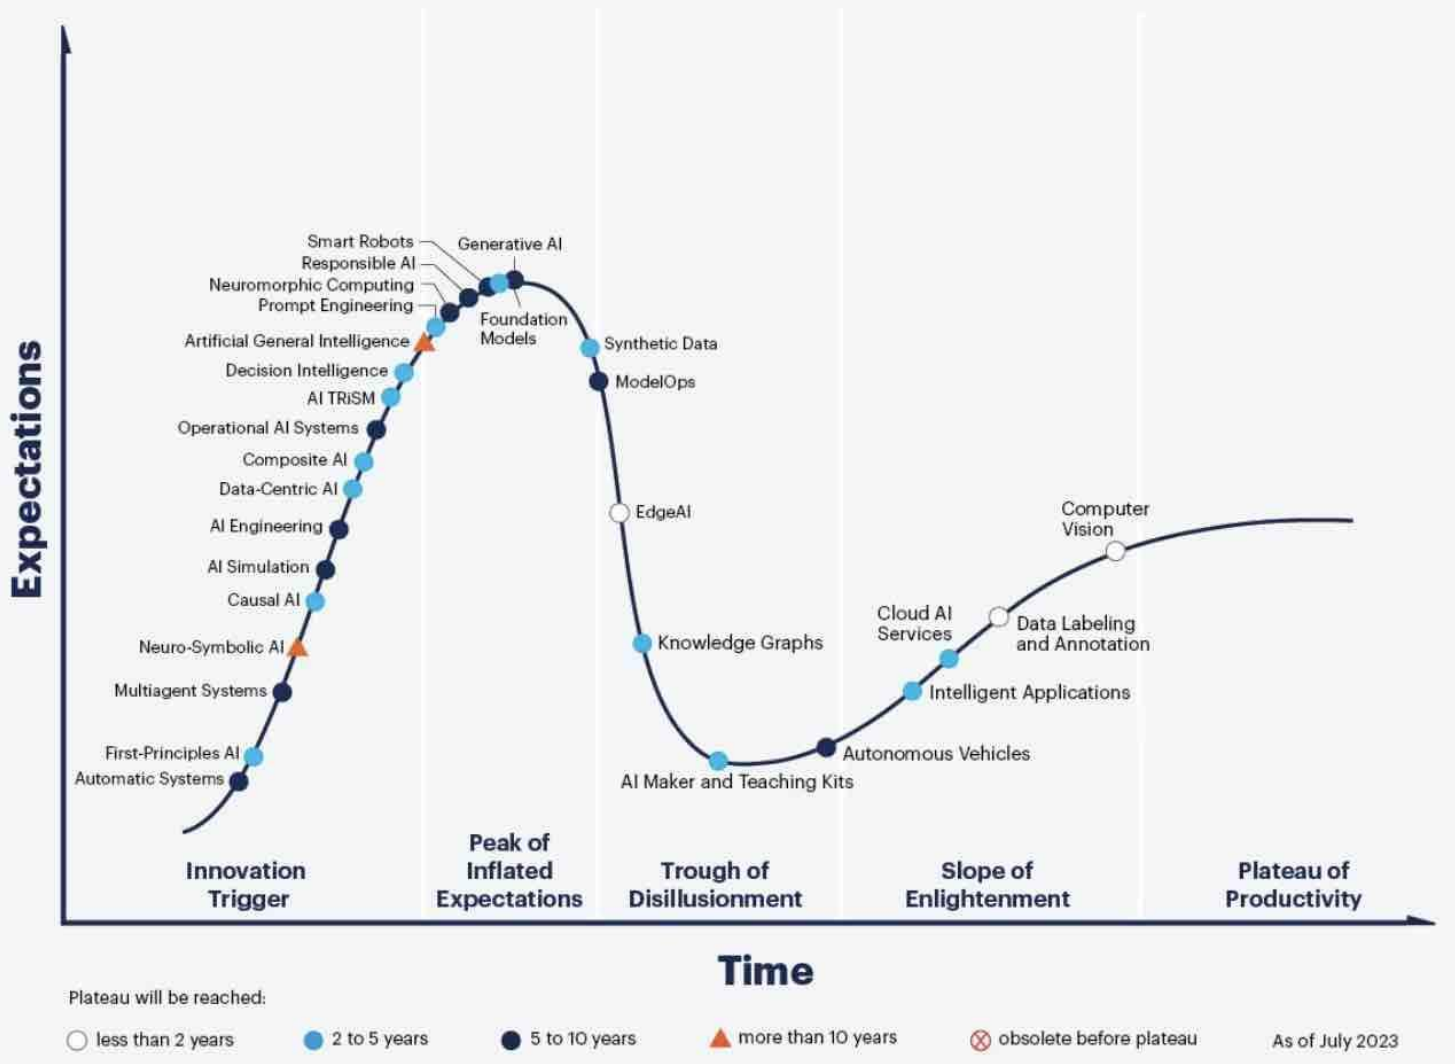

Insight Principles
AI/ML MOSA Challenges

- Historically, MOSA focuses on planned determinism to achieve modularity, openness, and scalability.
  - **Modularity**: This principle involves designing systems that are divided into interoperating modules. Each module performs a specific function and can be developed independently, allowing for easier upgrades, maintenance, and modifications.
  - **Openness**: At its core, this principle emphasizes using open standards and interfaces. This collaborative approach allows different modules to communicate and work together seamlessly, fostering a sense of unity and shared purpose in system design. It also enables components from other vendors to be integrated into the system, promoting diversity and innovation.
  - **Scalability**: This principle refers to the system's ability to handle a growing amount of work by adding resources. A scalable system can accommodate growth and change without a significant increase in complexity, providing a sense of adaptability and readiness for the future.
- MOSA establishes consistent engineering design principles across weapon systems to make components more easily removable, upgradeable, and interoperable.

Insight Principles
AI/ML MOSA Challenges

- The **stochastic nature of AI/ML** creates unique MOSA challenges.
  - **Integration Challenges** - Integrating AI/ML into MOSA-based systems is challenging due to complex algorithms and large datasets. Careful design is required to ensure seamless interaction with other system modules.
  - **Data Quality and Availability** - AI/ML models depend on high-quality training data, which can be challenging to ensure in a MOSA context due to data silos and incompatible formats. Establishing data governance and shared data repositories is crucial for successful AI/ML integration in MOSA.
  - **Scalability and Resource Constraints** - Scaling AI/ML within a MOSA framework can be resource-intensive, requiring significant computational power and storage. MOSA systems must handle these requirements for scalable AI/ML deployment, potentially using cloud-based solutions or optimized resource allocation.
  - **Cybersecurity Risks** - AI models can be vulnerable to data poisoning, model theft, and adversarial attacks. It is crucial to ensure the security and resilience of AI/ML components within the modular architecture through data encryption, access controls, and continuous monitoring.
  - **Coordination and Standardization** - Stakeholders must collaborate and set common standards and interfaces for AI/ML and MOSA integration to avoid fragmented efforts and ensure interoperability.
  - **Verification and Validation** - New testing and evaluation approaches are necessary to ensure trustworthiness.

8 Myths and Truths

- To address the AI/ML MOSA challenges, we begin by demystifying AI and ML technologies
- **Artificial Intelligence (AI)** – a family of algorithms deriving from the field of cognitive science and focused on emulating human thought
  - Computer vision and other sensing (audio, kinematic, …)
  - Sensory data understanding and decision-making
- **Machine Learning (ML)** – a family of algorithms deriving from the field of statistical data modeling and focused on
  - Data interpolation and generalization
  - Object classification
  - Relationship and trend detection

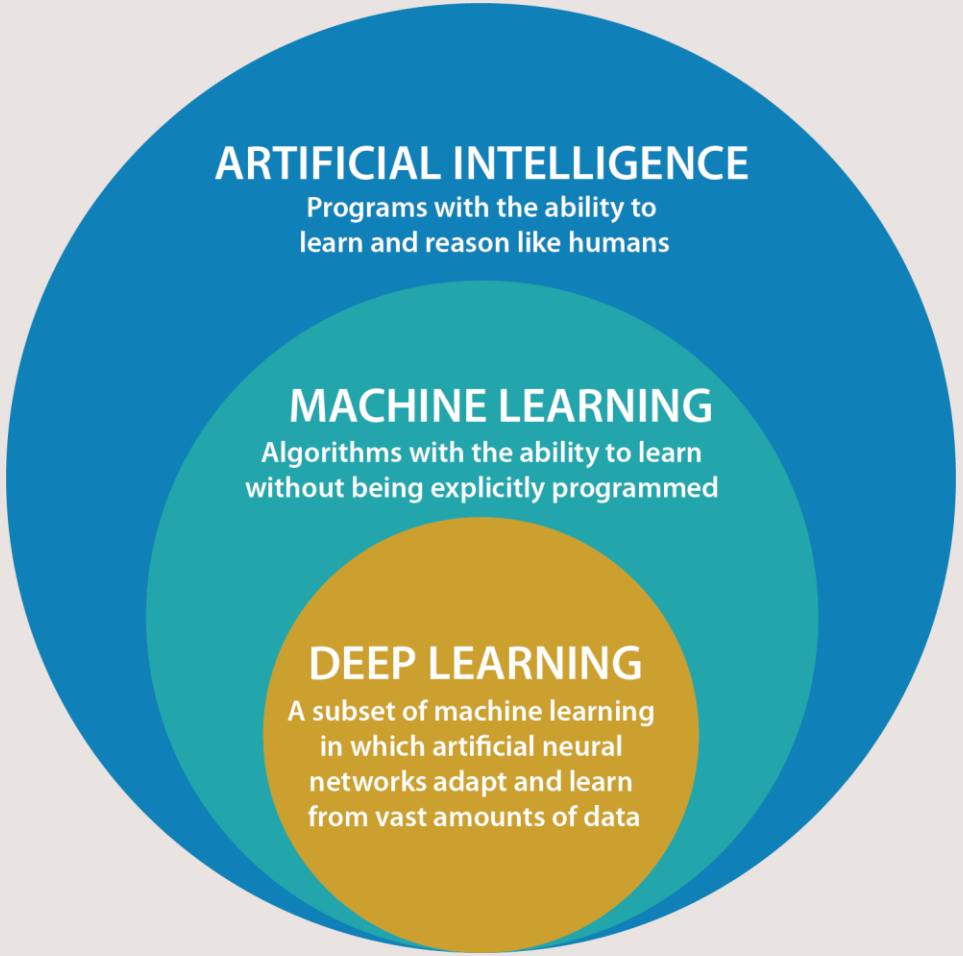

Myth #1: AI/ML is New Technology

- AI/ML began in the 1950's and has **progressed in waves**.
- The current wave of AI/ML is fueled by the **availability of computing and data**.
  - Most of the math and algorithms are a decade old!
- **System-of-systems AI/ML** is an emerging new discipline focusing on smartly combining algorithms based on their actual abilities.

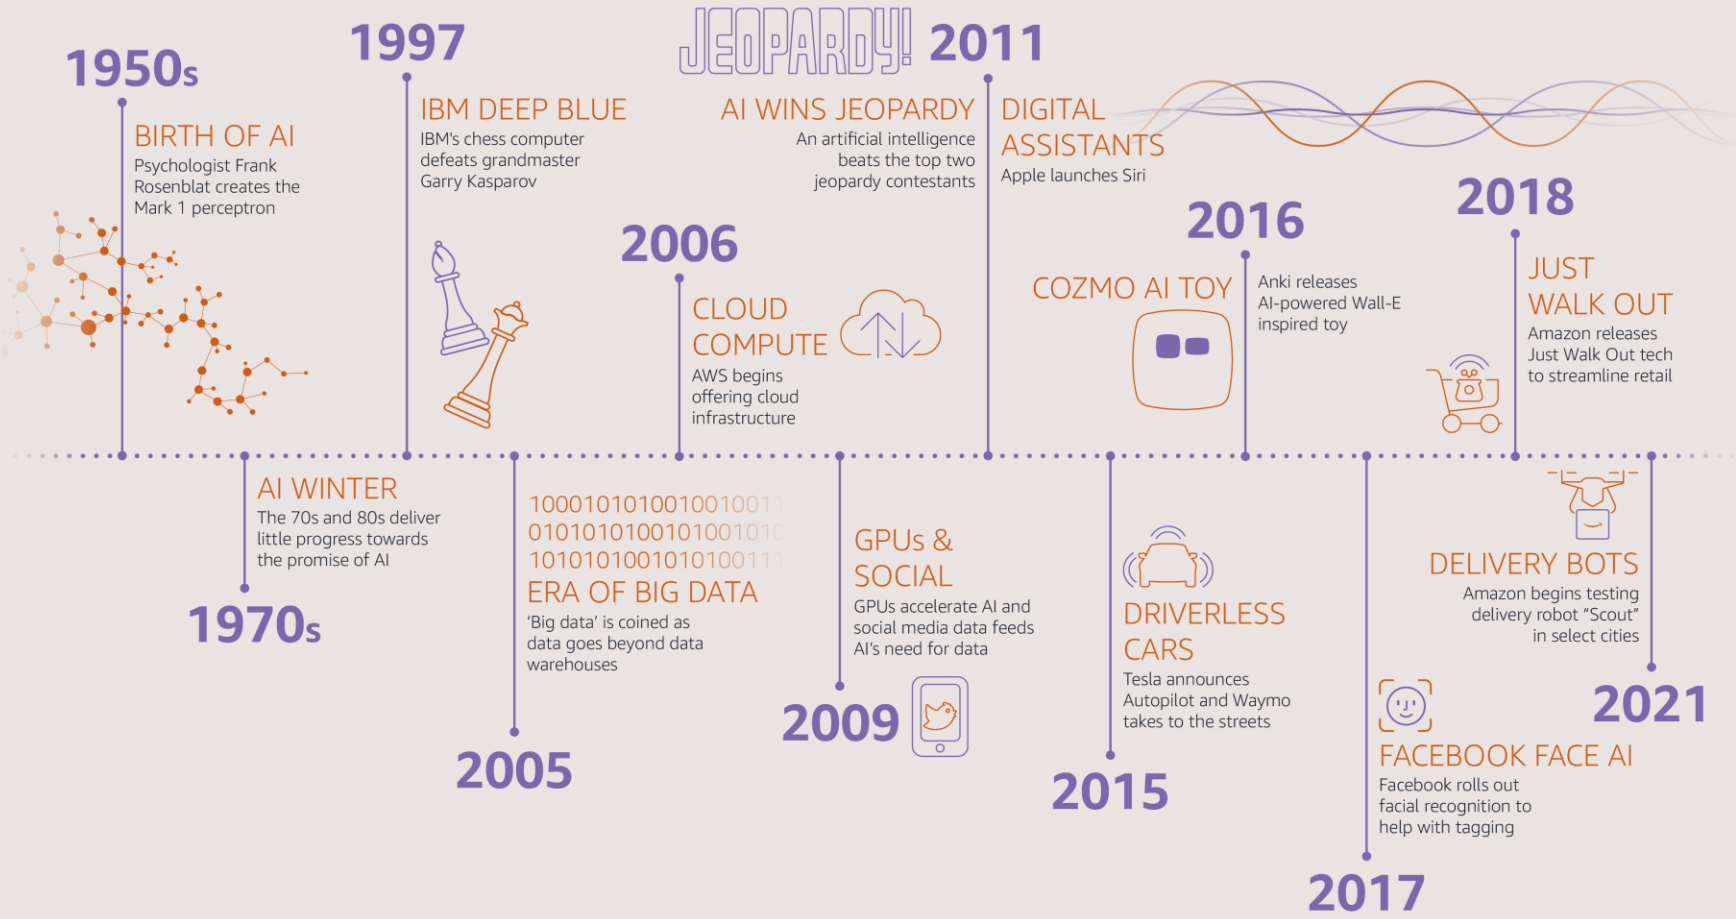

Myth #2: AI and ML are the Same Thing

- Artificial Intelligence and Machine Learning come from very different research domains.
- AI derives from cognitive science and focuses on algorithms representing human vision, reasoning, and cognition.
  - E.g., Expert systems, decision trees,
- ML derives from statistical data modeling and focuses on fitting mathematical models to provided data sets.
  - E.g., Data clustering, data regression, neural networks

	AI	ML
1	AI allows a machine to simulate human intelligence to solve problems	ML allows a machine to learn autonomously from previous data
2	The goal is to develop an intelligent system that can perform complex tasks	The goal is to build machines that can learn from data to increase the accuracy of the output
3	AI uses technologies in a system so that it mimics human decision-making	ML uses self-learning algorithms to produce predictive models
4	AI works with all types of data: structured, semi-structured, and unstructured	ML can only use structured and semi-structured data
5	AI systems use logic and decision trees to learn, reason, and self-correct	ML systems rely on statistical models to learn and can self-correct when provided with new data

BAE SYSTEMS

Myth #3: AI/ML Applies to any Problem

- AI/ML algorithms consist of two parts.
  - A mathematical model used to represent the underlying data
  - An optimization algorithm used to build, tune (train), and execute the model
- The AI/ML models' ability to 'reason/create/innovate' depends on how well the model represents the underlying data.
- Generally, AI/ML algorithms can be trained to 90-95% accuracy, after which they begin to memorize the data and lose their generalization ability.

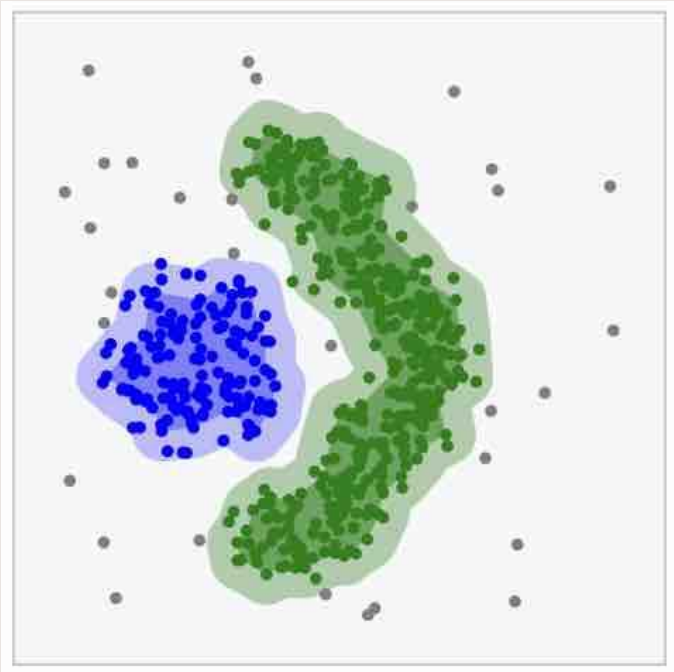

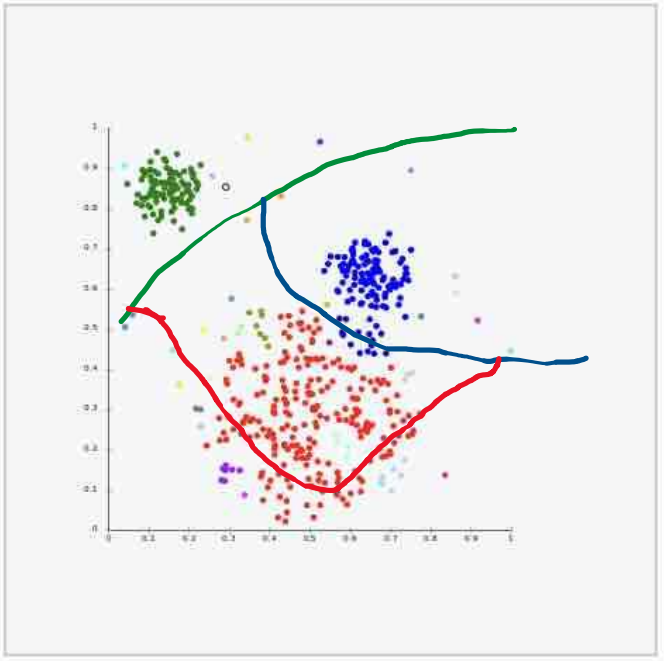

Myth #4: AI/ML Extrapolates to Produce New Facts

- As a general rule, "AI/ML algorithms are most reliable when operating inside the convex hull of their training data."
- AI/ML algorithms may have degraded performance on data they have not "seen" before.
  - During the model fitting process, data is used to tune and optimize the math model.
  - In cases where no data is available, there is nothing to fit to, so the model values are unconstrained and may be unreliable.
- If you train an AI/ML on puppies, it will learn puppies; if you show it a lion it may call it a "really big puppy."

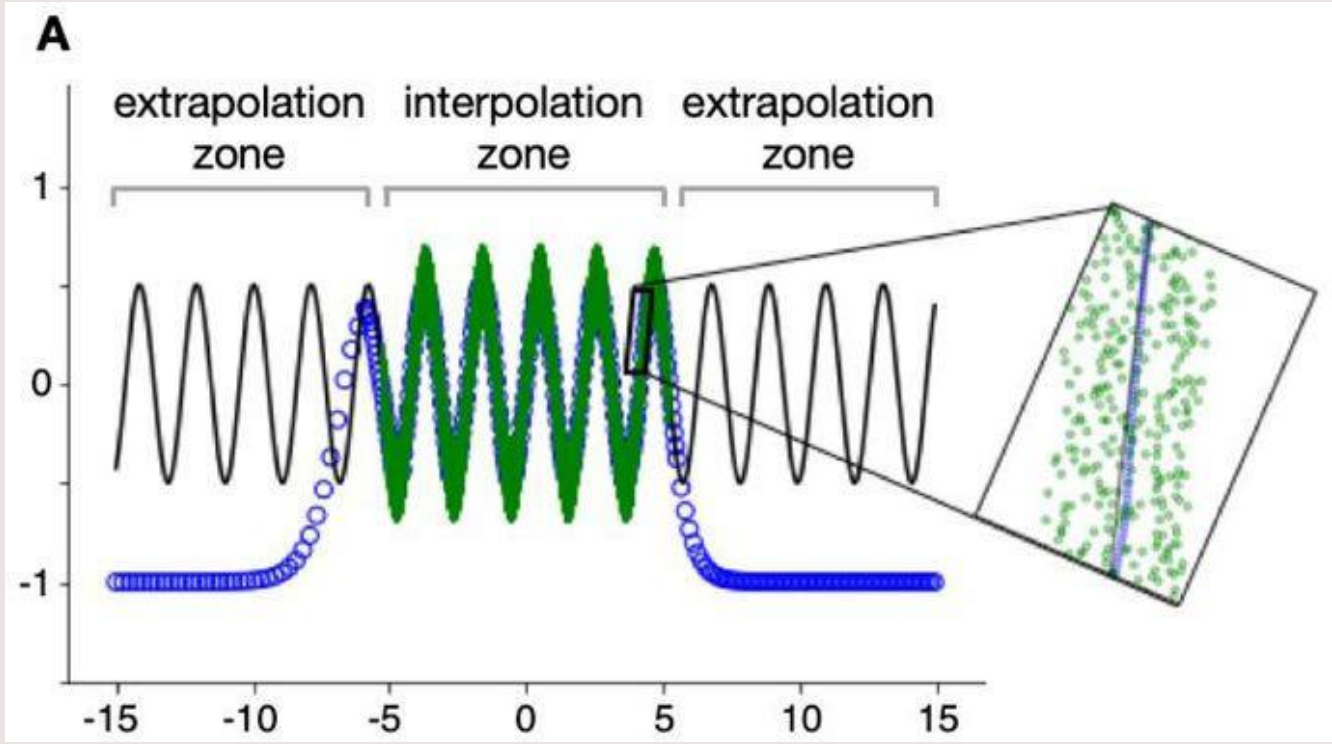

Myth #5: AI/ML Algorithm Selection is Unimportant

- Today, **hundreds of AI/ML algorithms are available** each having specific strengths and weaknesses.
- Choosing the wrong algorithm for the job significantly jeopardizes the success of AI/ML efforts.
- **Selecting the right algorithm is critical.**
  - AI or ML algorithm
  - Specific algorithm class (e.g., rule-based, Bayesian, clustering, piecewise linear, artificial neural network (ANN), GAN, …)
  - Algorithm topology (e.g., number of layers and nodes on an ANN)
  - Degrees of freedom versus available training data (e.g., 10x-100x data per node in an ANN)

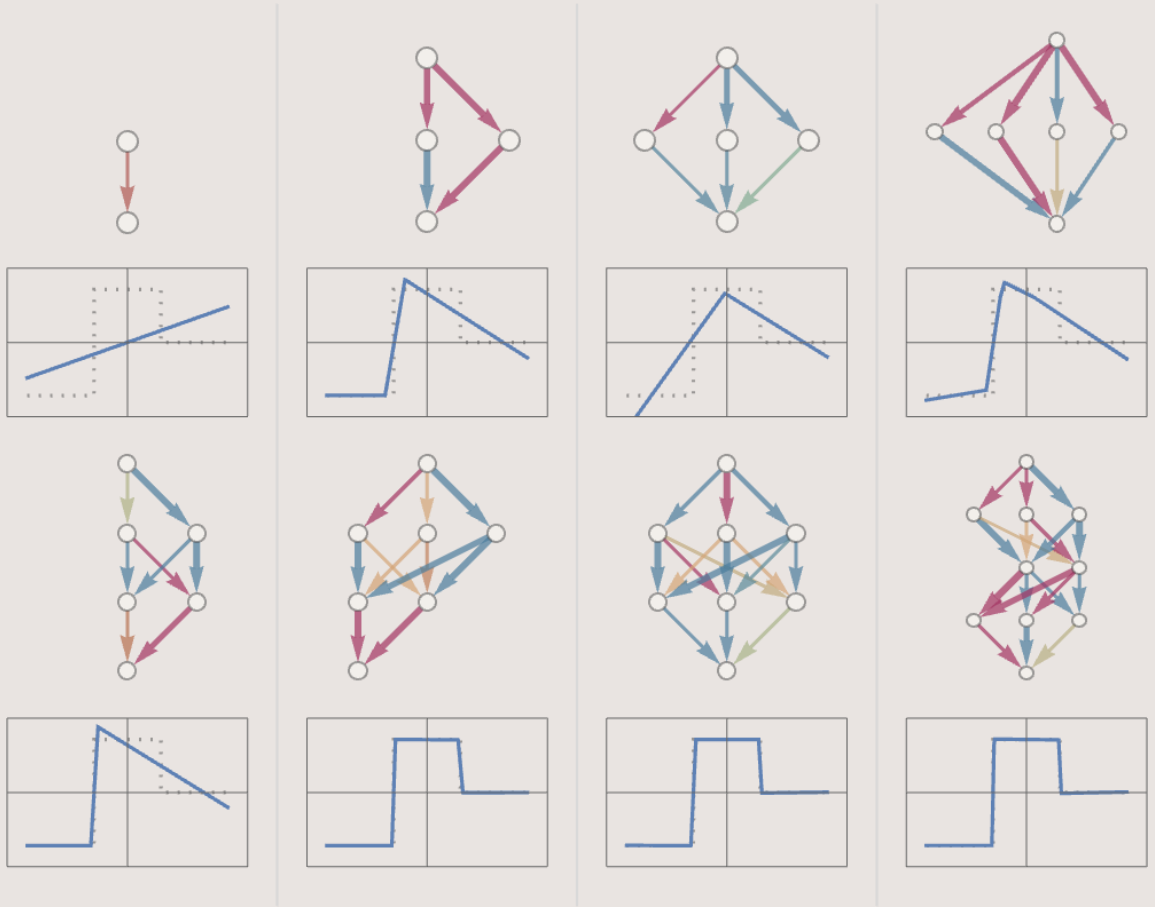

BAE SYSTEMS

Myth #6: Chat GPT is “Thinking”

- Chat GPT is a very large statistical model of English language documents (natural language processing).
- Based on how the algorithm works, the same set of seed words will always generate the same output (i.e., it is deterministic).
- The concept of temperature gives the appearance of ‘human creativity’ (0.0 = purely factual, 1.0 = highly generative).
- In the generative modes, GPT fibs leading to factually incorrect outputs and confusion between concepts (e.g., crocodile and alligator).

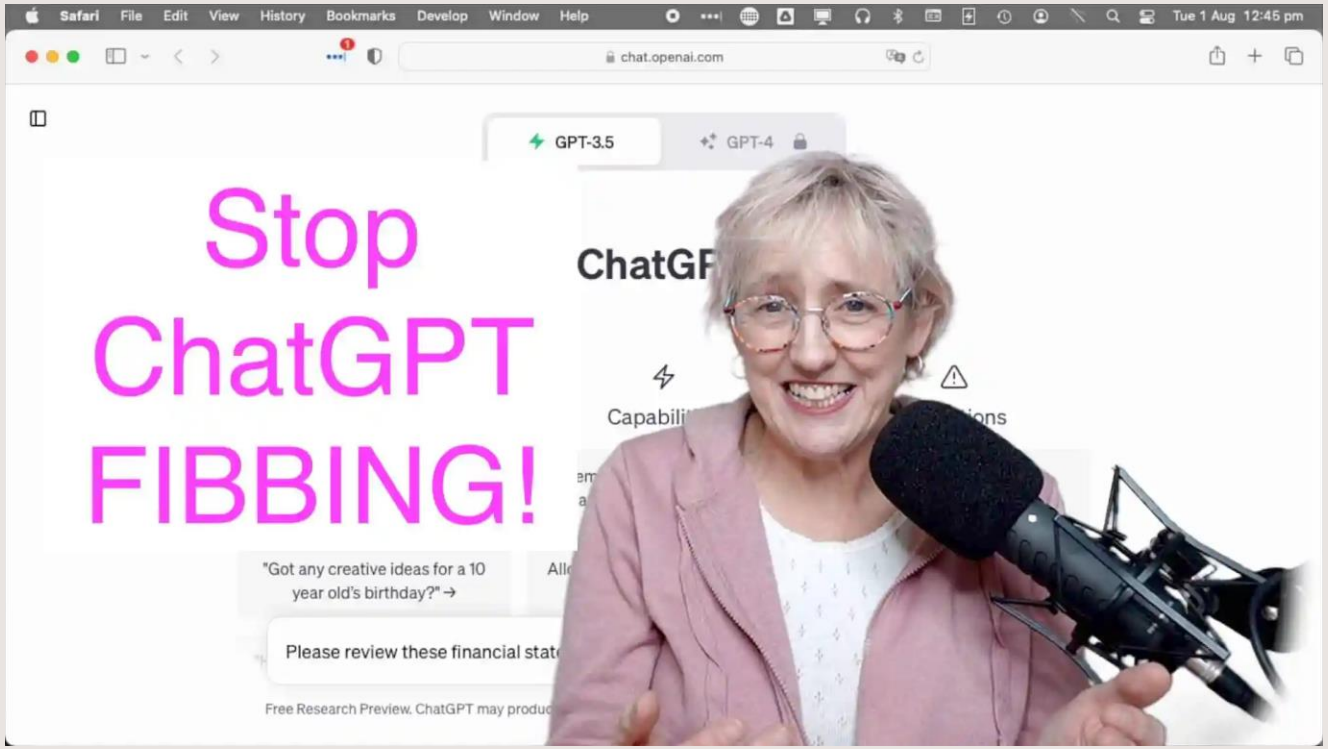

https://laurelpapworth.com/how-to-stop-chatgpt-from-lying-tutorial-on-temperature-parameter/

BAE SYSTEMS

Myth #7: AI/ML can be Objective and Unbiased

- AI/ML algorithms are based on statistical mathematics and expert rule modeling which inherently models biases in data.
  - In truth, all human cognition is based on bias; this is how the brain works.
- Biases in the math model are deduced from observations in the dataset during the training phase of development.
  - If the data set is factually correct, the resulting biases will be factually correct.
  - If the dataset is poisoned, the resulting biases may be factually incorrect.
- NOTE: When the poisoning is small relative to the training data set size, the AI/ML may generalize around poisoned data because it is statistically insignificant.

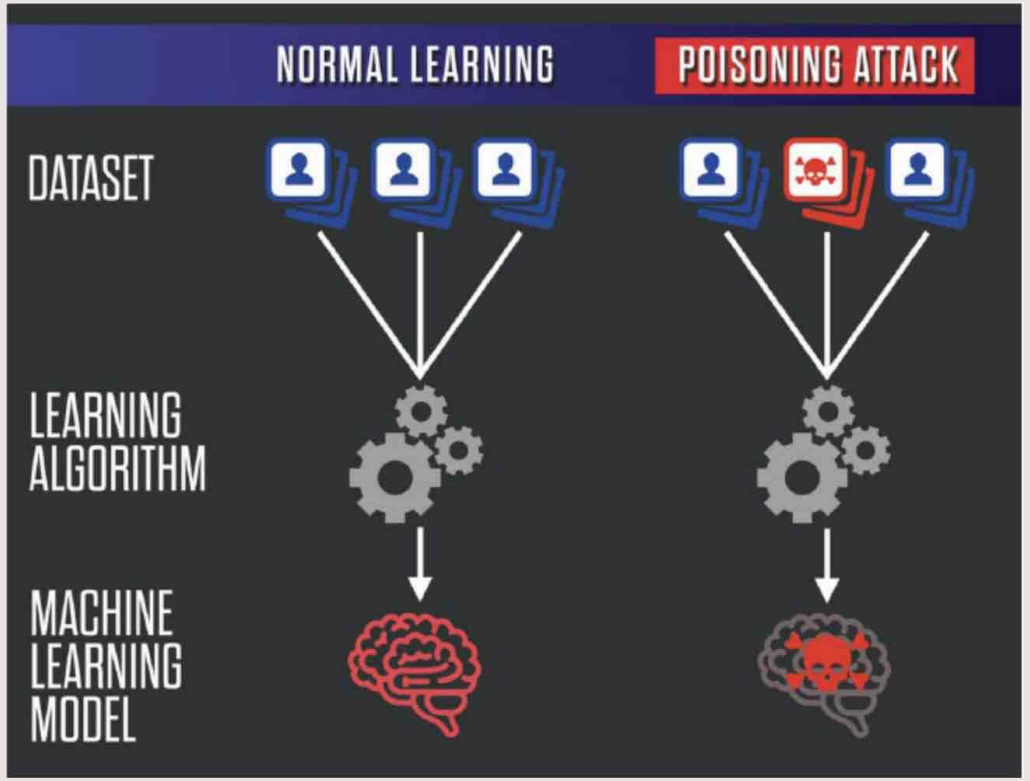

BAE SYSTEMS

Myth #8: AI/ML is Safe for Any Application

- Many AI/ML algorithm have **low explainability** and we cannot directly understand how decisions are made
  - Careful consideration and design is necessary to ensure accuracy, safety, and security in mission critical applications
- SoS AI/ML engineers must **carefully weigh the consequences of making wrong decisions**
  - Low-consequence applications: object classifiers, anomaly detectors, feedback control generators, etc.
  - Higher-consequence applications: missile seekers, autonomous weapon systems, commercial autodrive systems

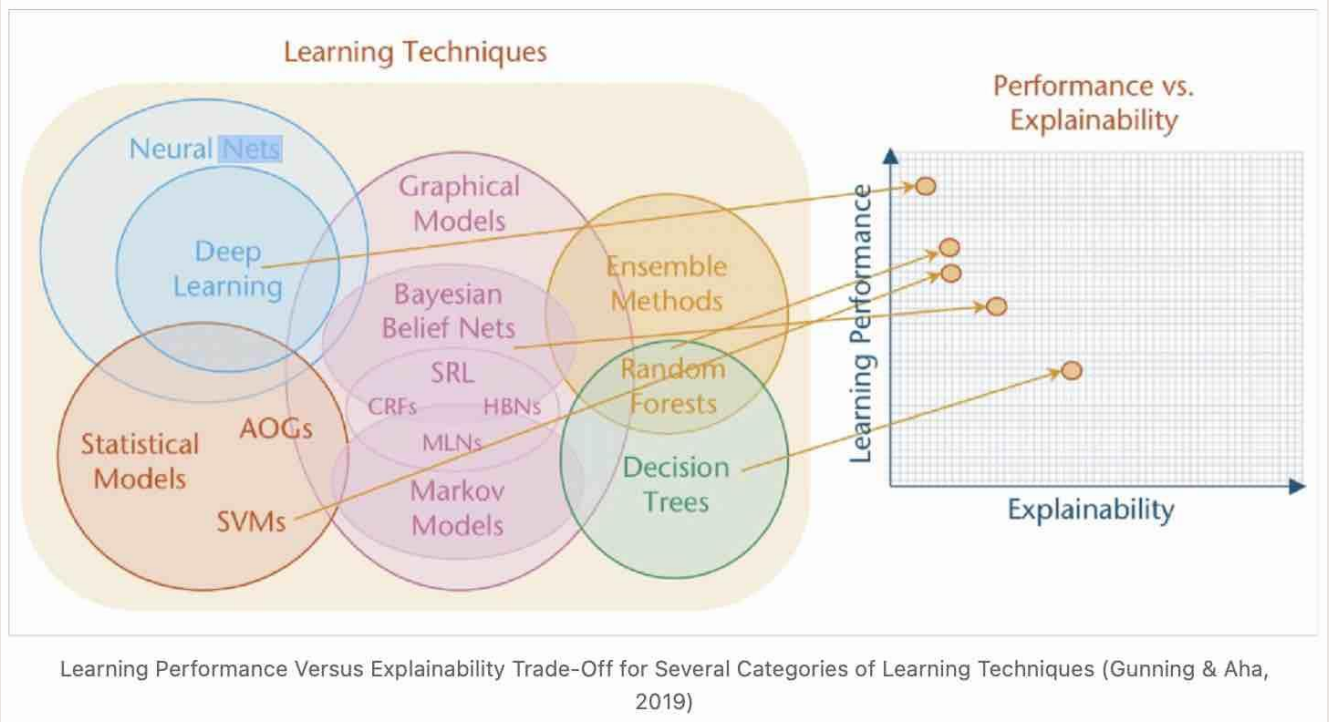

Learning Performance Versus Explainability Trade-Off for Several Categories of Learning Techniques (Gunning & Aha, 2019)

BAE SYSTEMS

Myth #9: AI/ML will Eventually Learn to Function like the Human Brain

- AI/ML today **barely mimics the functions of an insect brain**; there are many brain functions that today's AI/ML algorithms cannot replicate.
- It is more accurate to state "AI/ML is a tool for optimizing some types of control systems and machine functions".
  - When a closed-form algorithm exists to solve a system problem, AI/ML is not the best choice.
  - When a solution is not apparent and we have sufficient observations (data), an AI/ML algorithm may find an acceptable (and sometimes unexplainable) solution.

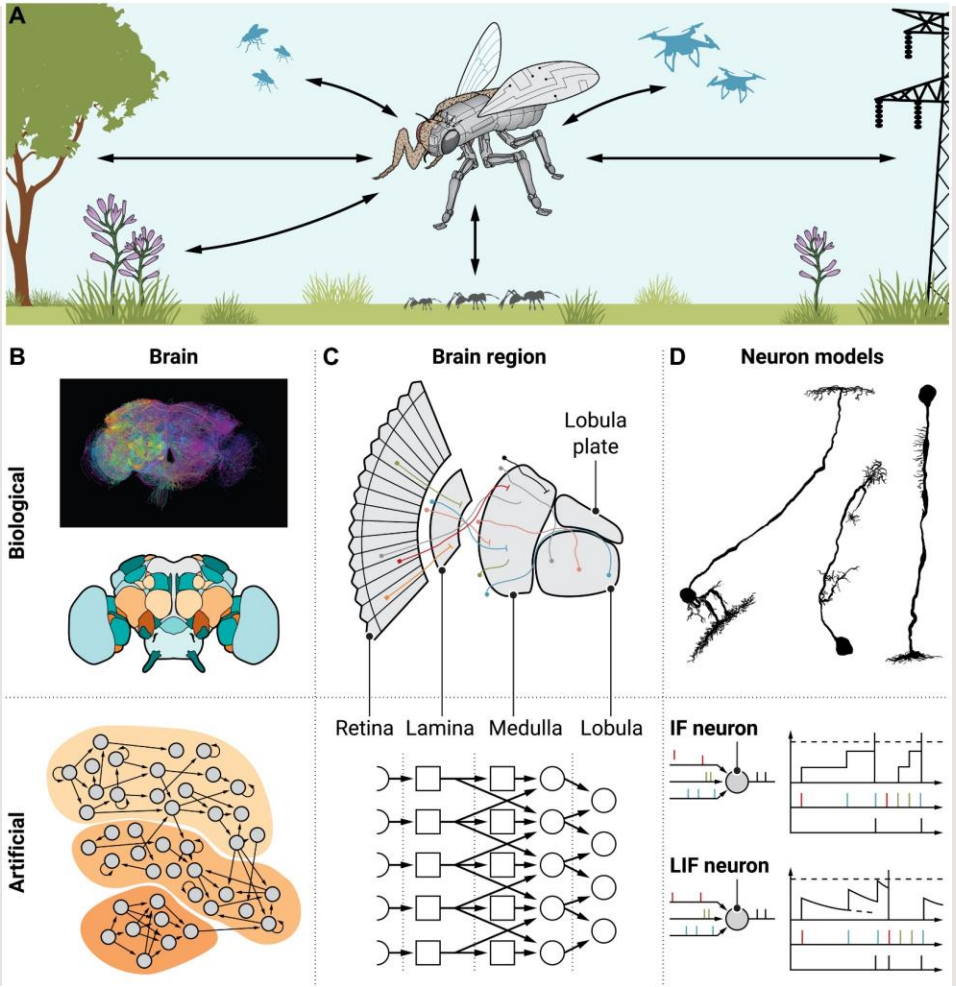

Myth #10: AI/ML is Best Implemented on Digital Computers

- There is **no existence proof that intelligence can be implemented on digital computers** leading many researchers to conclude that digital computing is the wrong machine.
- The **human brain is not a digital computer** and yet.
  - Has more intelligence than digital AI/ML systems
  - It uses a fraction of the power
  - Uses a fraction of the space
- **Neuromorphic computing** is exploring alternative computing platforms that are better suited to developing intelligent systems.

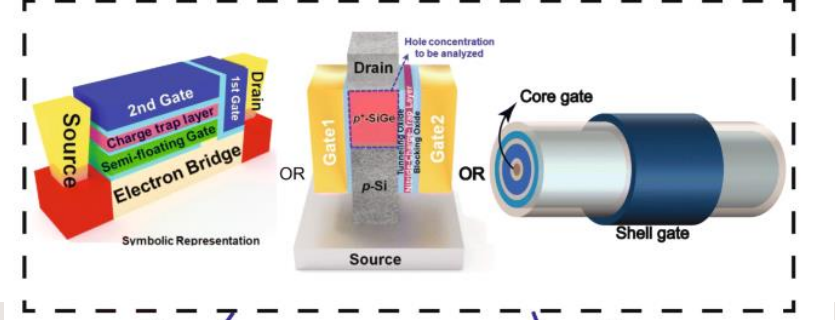

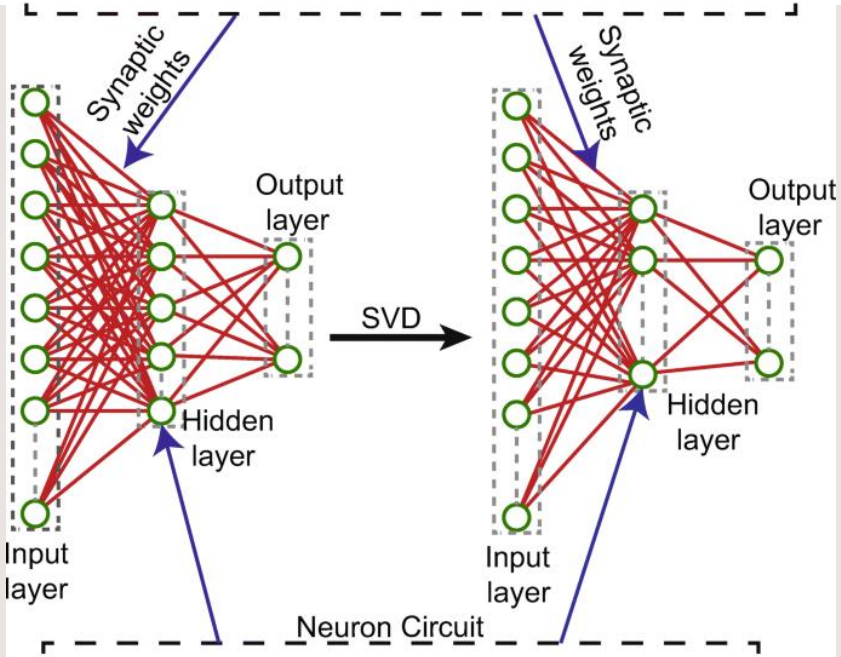

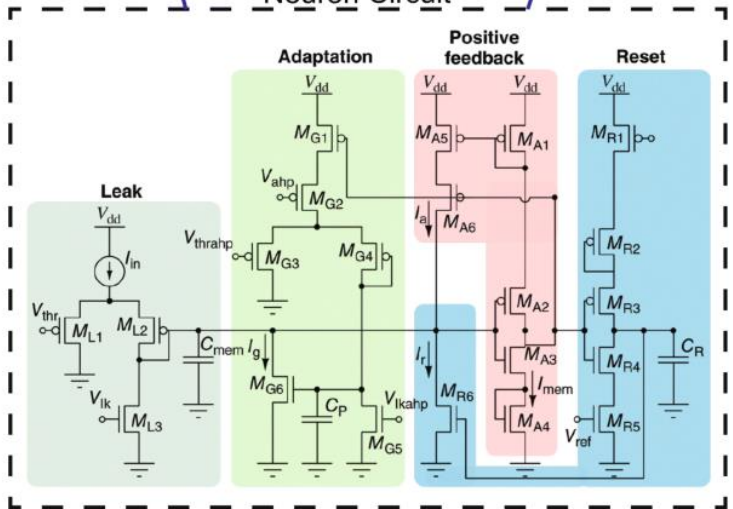

Next Steps for AI/ML MOSA

- Integrating AI/ML-enabled components into future weapon systems is a certainty.
- At BAE Systems, we continue to explore the relationship between non-deterministic AI/ML technologies and deterministic MOSA design principles.
- Architecture and design strategies are needed for implementing AI/ML MOSA across product scales.
  - Cloud IT → Weapons Platforms → Embedded Hardware Devices
- More work is required on this topic!

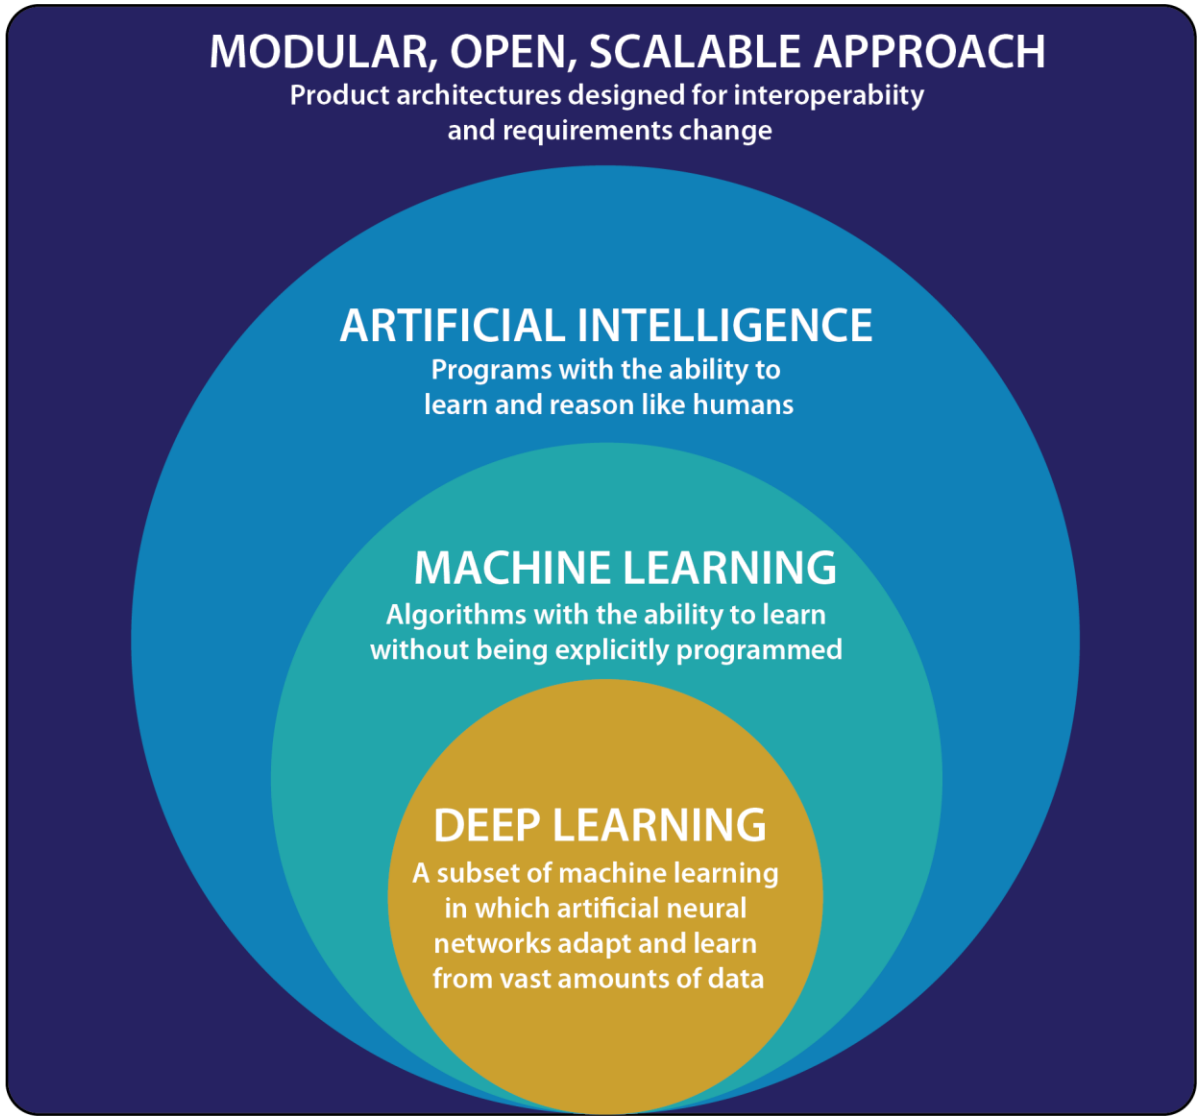

BAE SYSTEMS

Questions?