

Artificial Intelligence

Artificial Intelligence

The printable version is no longer supported and may have rendering errors. Please update your browser bookmarks and please use the default browser print function instead.

Lead Authors: *Barclay Brown, Tom McDermott, Rael Kopace, Ramakrishnan Raman, Ali Raz, and Kevin Robinson*

This article provides an overview and definitions of artificial intelligence and machine learning and their importance to and relationships with systems engineering.



Contents

Introduction

Machine Learning Foundations of AI

Systems Engineers and AI

Systems Engineering and AI: SE4AI and AI4SE

SE4AI: Addressing Practical Implementation
Challenges for AI

AI4SE: Leveraging AI to Advance State-of-the-Art in
Systems Engineering

Landscape of AI Applications in Industry and Systems

Acknowledgements

References

Works Cited

Primary References

Additional References

Introduction

Artificial Intelligence (AI) is perhaps best defined as the ability of a system to exhibit behavior, which if exhibited

by a human being, would be considered intelligent. It's the ability of a machine to process and learn from data, to recognize and understand patterns, to solve problems, and make decisions autonomously. The term "Artificial Intelligence" dates to the early 1950s where it was first introduced to mimic human-level intelligence capabilities in software and hardware systems—a goal which still is far from reachable in the near future despite the rapid technological advancement in the AI field . It is important for systems engineers to note that "AI" does not designate a specific technology—intelligent behavior can be implemented in a system in a number of ways, including conventional software or hardware logic. Modern systems that use AI fall into two major categories: systems where the decision-making logic is explicitly pre-programmed, and systems where the decision-making capability is learned from data. This latter category, known as Machine Learning (ML), is the predominant category for the modern paradigm shift in capabilities enabled by AI.

Machine Learning Foundations of AI

ML builds its foundation from statistics and computer science disciplines, accelerated by the advances in computational power and memory offered by modern hardware systems (i.e., CPUs and GPUs). The ever-increasing number of ML algorithms available today can be classified into three major categories of Supervised, Unsupervised, and Reinforcement learning.

- Supervised Learning is where labelled input-output datasets are available and ML algorithms learn the mapping between input-output pairs presented in the datasets (for example, learning from clearly labelled pictures of cars and tanks)
- Unsupervised Learning is where the labels in data sets are not available and the ML algorithms learn the grouping or clustering of input and outputs (for example, identifying common customer attributes, classifying spam email, or detecting fraudulent transactions)
- Reinforcement Learning is where an AI agent trains itself to make decisions to effect a positive change in the environment based on a defined reward function (e.g., training a robot to avoid obstacles by penalizing crashes)

The implementation approaches for all ML types require large datasets and can be based on statistical models or neural networks including Deep Neural Networks (DNNs). Irrespective of the implementation approaches and the type of ML, the workflow for deploying an ML-based AI solution remains largely common and is shown in Figure 1 (details on each of these steps can be found in the Digital Transformation book by Thomas Siebel).

Figure 1 A typical ML Workflow

Understanding the subtleties of statistical models and neural networks for ML implementation along with the related workflow will become imperative for systems engineers as they begin to incorporate ML solutions in systems and perform systems engineering activities for AI systems.

Systems Engineers and AI

Given the holistic nature of the systems engineering discipline, it is practical to think about AI and its role in the larger context of digital transformation, which can be defined as the confluence of big data, cloud computing, AI and Internet of Things (IoT). Digital transformation is enabling organizations to evolve their business and operational models. As such most organizations are collecting vast amounts of data pertaining to their business, operations, systems, customers, and personnel. Big data has created the conditions that necessitate the application of AI solutions supported by the technological advancements in computational power and memory to aid in the analysis of big data and the automation of labor intensive and time-consuming processes.

The ability to collaborate and communicate with data scientists, AI engineers, cloud computing specialists, application developers, and other experts is becoming a necessity for many systems engineers. Systems engineers must become sufficiently knowledgeable and familiar with as many elements of the digital transformation ecosystem as possible to best serve the needs of the organization and its mission in order to identify the areas of business or the systems that would most benefit from the application of AI. Understanding the benefits as well as the potential pitfalls of an AI development initiatives is also important. To that end, systems engineers must be able to identify the right requirements and architectures, expertise, partnering organizations, technology stack, development platforms, and other relatable elements. The roles and

responsibilities of systems engineers versus data scientists, data engineers, application developers, and AI engineers must be differentiated. The market is filled with an ever-growing number of AI platforms and it is virtually impossible for a non-data scientist to be able to master each of them. Moreover, the advances in the AI domain happen at such a rapid pace that nowadays many companies may choose to avoid investing in their own internal data science teams to develop AI solutions. Instead, they can purchase or license models that are ready for implementation from companies that have solved similar problems before. Additionally, there exist unified platforms that aim to standardize workflows and come equipped with the appropriate technology stack to collect, label, and feed data into supervised learning models or alternatively assist with the development of models.

As a result, rather than mirroring the roles and responsibilities of data scientists and other AI domain professionals, it is beneficial for most systems engineering professionals to develop an understanding at various degrees of depth of the relevant technologies that support and are used to create AI applications including but not limited to the following:

- Data integration
- Relational and non-relational databases
- Cloud storage
- Enterprise architecture infrastructure and APIs
- ML learning frameworks
- Batch and online processing
- Executable environments
- Commonly used programming languages
- UI and data visualization tools

Systems Engineering and AI: SE4AI and AI4SE

At an early 2019 Future of Systems Engineering (FuSE) workshop hosted by the International Council on Systems Engineering (INCOSE), the terms AI for SE and SE for AI were first used to describe the dual transformation envisioned for both the SE and AI disciplines. The “AI4SE” and “SE4AI” labels have quickly become symbols for an upcoming rapid evolutionary phase in the SE community. SE’s Digital Engineering transformation will enable both transformation of SE

practices using AI for SE and drive the need for new systems engineering practices that support a new wave of automated, adaptive, and learning systems, termed SE for AI. AI4SE applies AI and ML techniques to improve human-driven engineering practices. This goal of “augmented intelligence” includes outcomes such as achieving scale in model construction and efficiency in designing space exploration systems. SE4AI applies SE methods to the design and operation of intelligent systems, with outcomes such as improved safety, security, ethics, etc.

Systems engineers have long sought to evolve engineered systems to include human-like intelligence, thinking, and autonomous decision-making. In many cases, the inclusion of intelligent capabilities means increasing non-determinism in the system’s behavior, exponentially increasing complexity in the design, verification, and validation of such systems. The initial approaches of inculcating intelligence predominantly involved human-in-the-loop for capturing knowledge from experts, codifying it in some form that is machine understandable (for instance, a “knowledge database”), and enabling the system to interpret the codified knowledge so as to make the required optimal decisions and exhibit the desired intelligent behavior. These initial approaches involved human-in-the loop for the entire gamut of activities from capturing the knowledge and codifying the knowledge to enabling (programming) the system to leverage the knowledge.

More recent contemporary approaches represent the next evolution of inculcating intelligence in systems, wherein the human gathers the required data and programs the system such that the system can learn/build the knowledge from the data provided by the human. Learning from data requires significant computing resources (processing power and memory) but recent advances have expanded the availability of large amounts of both source data and matching computing resources to enable intensive processing and learning from data. Recent ML algorithms are the result of this evolution. The system learns during the design and development phases and is expected to demonstrate limited forms of adaptive behavior.

In both approaches, a human is in the loop to validate the intelligence gained by the system, and to make a conscious decision to deploy the intelligence in the system. The next evolution is expected to be predominantly based on a system learning from the other systems it interacts with (machine-to-machine

learning) and building its intelligence from the interactions and from the ecosystem (cloud). Evolution in that direction is as seen by some of the recent progress made in reinforcement learning models and algorithms. The system learns by trial-and-error and once it meets performance expectations, can apply its behavior on a larger scale.

The augmented intelligence scenarios described above will need an underlying and synergistic foundation of both SE4AI and AI4SE which are briefly described in the following subsections.

SE4AI: Addressing Practical Implementation Challenges for AI

The increasing application of AI in systems presents challenges for both the systems engineering community and the AI community. Their common goal is to ensure that the future users of the system can be certain that the behavior and performance of that system is what is needed. That is, the AI system's behavior is verified through a set of activities that check its compliance against system requirements and validate it for fitness to meet user needs. The need for verification and validation presents a challenge for the engineering of AI systems as the failure modes observed in AI systems are different and may not be adequately addressed by traditional Systems Engineering life-cycle approaches. From a systems engineering lifecycle perspective, appropriate tailoring of conventional processes is needed to "engineer" a system that is intelligent—a system that learns during development and is designed for adaptive behavior. New approaches are required for developing requirements, evaluating when these intelligent systems are ready for operations, and for ensuring their adaptive behavior is safe and produces the desired outcomes. The data required for enabling a system to acquire the desired intelligence poses several challenges from being biased (inculcating a biased intelligence reflected in the system's behavior) to not being representative of the various use case scenarios envisaged for the system. The system's intelligence is only as good as the data that was used to train it. There are many considerations pertaining to the broader challenge of engineering safe and effective human interaction with intelligent systems. These include (a) leveraging cognitive strengths of humans and AI, (b) gaining trust, (c) explainable intelligence (d) identifying and understanding bias, and (e) dealing with uncertainty.

Broadly there are three key challenges for the systems engineering of intelligent systems:

1. **New failure modes not previously experienced in the engineering of systems.** The AI community recognizes that there are five main failure modes that cause the AI systems to not behave safely and as expected. These new failure modes include negative side effects, reward hacking, scalable oversight, unsafe exploration, and distributional shift .
2. **The unpredictability of performance due to non-deterministic and evolving behavior.** ML systems initially learn from predetermined data and through the activity of validation, system engineers check the compliance of the system performance against its intended purpose, captured in a Systems Specification. The challenge with AI, and specifically ML, is predicting the performance and behaviour of the AI algorithm on unseen data in future operations. ML systems exhibit non-deterministic performance, with the performance of some systems evolving as the system learns (changes performance) during operations. This presents challenges in validating system compliance before the system enters operations.
3. **Lack of trust and robustness in future systems performance.** System validation is based on a basic four step approach: obtaining results from a validation test, comparing the measured results to expect results, deducing the degree of compliance, and deciding on the acceptability of this compliance. A key aspect in deciding the acceptability of compliance is expert judgement. Expert judgement requires an understanding of the result as compared to the relevance of the context of use, and therefore the results need to be explainable. Explainable behaviour of AI Systems is problematic, and therefore determining a level of trust and robustness in future systems performance is challenging.

AI4SE: Leveraging AI to Advance State-of-the-Art in Systems Engineering

The next evolution of SE practices is expected to be predominantly driven by AI technologies assisting systems engineers in the engineering development

lifecycle activities. AI4SE addresses how AI can enhance the systems engineering lifecycle for engineered systems, across the various lifecycle phases including concept development, requirements, architecture design, implementation, integration, verification, validation, and deployment. Enhancing and assisting systems engineering processes, methods, and tools, with tangible impacts on the quality of the engineered system as well as on the cycle time for the various life cycle activities, would be some of the primary focus areas of AI4SE. For instance, AI technologies can be leveraged to advise system architects on the various architecture and design decisions options based on intelligence built from collective prior experience of decisions made in earlier systems. A second example is leveraging AI technologies to assist in arriving at various corner test cases during verification. Historical systems engineering life cycle data would be the predominant drivers for the use of AI for such applications. A third example would pertain to the simulation aspects, where digital and synthetic environments (e.g., digital twins) are leveraged to understand various lifecycle operation scenarios and providing better insights to systems engineers on understanding the implications of architecture design decisions on the engineered systems. However, some forms of distinction are being envisaged in the community between AI4SE as against automation, digitization and digital linking of various SE life cycle artifacts and work products. While automation and digitization activities may be enhanced by AI, these activities may still be dominated by conventional software, with outcomes focused on significant reduction in life cycle time and efficiency on managing scale.

Landscape of AI Applications in Industry and Systems

Over the last few decades, ML-based solutions have become ubiquitous in technical and social systems and have enabled new business opportunities and even new business models. The rapid commercialization of AI is bringing profound changes to all markets. Capabilities are advancing and it is becoming easier to develop and implement AI solutions given the growing number of AI tools and platforms. Organizations in every industry segment are increasingly realizing that AI is key to market leadership and that they all have processes that are suited for AI. Early adopters of AI have noticed significant tangible benefits and given the lower barriers of entry that exist nowadays, others are making

significant investments to accelerate their AI adoption.

AI is being integrated into the fabric of business and systems and AI solutions are being deployed to solve a wide spectrum of use cases at every layer of the enterprise. The following table provides examples of AI applications across the different fields and sectors.

Domain	General examples of AI usage
Sales	Price optimization, forecasting, performance management, dynamic recommendations (think Amazon, Netflix making product and movie recommendations) [Ref: https://hbr.org/2018/07/how-ai-is-changing-sales]
Security	Breach risk prediction, incident response, early identification and classification of cyber threats [Ref: https://www.ibm.com/case-studies/cargills-bank-td]
Anti-fraud	Identification of fraudulent behavior and transactions [Ref: https://www.fico.com/blogs/5-keys-using-ai-and-machine-learning-fraud-detection]
HR	Candidate assessment, screening time reduction, skills to jobs alignment [Ref: https://www.forbes.com/sites/jeannemeister/2019/01/08/ten-hr-trends-in-the-age-of-artificial-intelligence/?sh=426c0c7d3219]
Marketing	Programmatic advertising for target audiences, behavior analysis, interactive marketing through chatbots [Ref: https://www.forbes.com/sites/forbesagencycouncil/2019/08/21/how-artificial-intelligence-is-transforming-digital-marketing/?sh=46ae38e021e1]
Personal Assistant	Control smart home objects, interact with the web and provide answers to questions, manage calendar
Smart tools	Smart thermostats, doorbells, home security system, baby monitor, real time language translation [Ref: https://www.pcmag.com/news/the-best-smart-home-devices-for-2020]
Finance	Automated financial trading and trade validation to protect from obvious errors and irrational price swings, identification and management of risk based on user history, classification of loan and credit applications [Ref: https://builtin.com/artificial-intelligence/ai-finance-banking-applications-companies]
Healthcare	Assisted diagnostic medical imaging, management and accuracy check of electronic medical records, disease risk prevention, personalized health management [Ref: https://www2.deloitte.com/cn/en/pages/technology-media-and-telecommunications/articles/global-ai-development-white-paper.html]
Education	Virtual adaptive teaching and learning, curriculum personalization, virtual tutor [Ref: https://medium.com/towards-artificial-intelligence/artificial-intelligence-in-education-benefits-challenges-and-use-cases-db52d8921f7a]
Autonomous Driving	Real time sensor data processing, dynamic path planning, auto health monitoring, route optimization, real time traffic monitoring [Ref: https://www.embedded.com/the-role-of-artificial-intelligence-in-autonomous-vehicles/]
Retail	Smart inventory, demand forecasting, smart logistics, unmanned store, automated customer service inquiries through chatbots [Ref:]
Manufacturing	Quality and safety checks, improving yield and performance, failure mode prediction, predictive maintenance, adaptive design, energy saving, production forecasting [Ref: https://www2.deloitte.com/cn/en/pages/technology-media-and-telecommunications/articles/global-ai-development-white-paper.html]
Media	Personalized media content, content discovery, real time fact checking, identifying false content [Ref: https://neoteric.eu/blog/10-use-cases-of-ai-in-manufacturing/]
Government and Legal	Non-compliant behavior and tax evasion, policy checks, citizen assistance chatbots, emergency and disaster resource identification and monitoring, automated due diligence on background information, legal analytics, content generation [Ref: https://emerj.com/ai-sector-overviews/ai-in-law-legal-practice-current-applications/]
Agriculture	Pest monitoring, crop return prediction, optimizing yield, monitoring and managing closed environment, weather forecasting [Ref: https://www.intel.com/content/www/us/en/big-data/article/agriculture-harvests-big-data.html]
Logistics	Predictive supply chain network management, demand and capacity planning, route optimization [Ref: dhl.com/content/dam/dhl/global/core/documents/pdf/glo-core-trend-report-artificial-intelligence.pdf]
Oil and Gas	Oil seep detection with robots, precision drilling, temperature and pressure monitoring, predictive maintenance [Ref: https://www.sparkcognition.com/top-4-ai-applications-oil-gas-industry/]

Acknowledgements

The article is provided by INCOSE's AI working group and includes summarized sections from the AI Primer for Systems Engineers which is currently under development. The AI Primer for Systems Engineers will provide a broader view of AI discipline and its implication for Systems Engineering. The AI working group would like to acknowledge (in alphabetical order) Barclay Brown, Tom McDermott, Rael Kopace, Ramakrishnan Raman, Ali Raz, and Kevin Robinson for their contributions to this article.

References

Works Cited

Ammanath, B., Jarvis, D. and Hupfer, S., 2021. Thriving in the era of pervasive AI. [online] Deloitte Insights. Available at: <<https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>> [Accessed 23 April 2021].

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.

Jordan, Michael I. "Artificial intelligence—the revolution hasn't happened yet." Harvard Data Science Review 1.1 (2019).

Sharma, M., 2021. Navigating the New Landscape of AI Platforms. [online] Harvard Business Review. Available at: <<https://hbr.org/2020/03/navigating-the-new-landscape-of-ai-platforms>> [Accessed 23 April 2021].

Siebel, Thomas M. Digital transformation: survive and thrive in an era of mass extinction. Rosetta Books, 2019.

TechRepublic. 2021. 10 ways data and analytics will impact businesses. [online] Available at: <<https://www.techrepublic.com/article/10-ways-data-and-analytics-will-impact-businesses/>> [Accessed 23 April 2021].

Primary References

Siebel, Thomas M. Digital transformation: survive and thrive in an era of mass extinction. Rosetta Books, 2019.

Additional References

None.

< Previous Article | Parent Article | Next Article >

SEBoK v. 2.7, released 31 October 2022

Retrieved from
"https://sandbox.sebokwiki.org/index.php?title=Artificial_Intelligence&oldid=65958"

This page was last edited on 10 October 2022, at 08:30.