IBM **Engineering**

KUGLER MAAG CIE
besser mit uns

Challenges in Systems Engineering of Intelligent and Autonomous Systems

By
Sky Matthews, IBM
Dr. Michael Fäustle, Kugler Maag Cie

IBM

Challenges in Systems Engineering of Intelligent and Autonomous Systems

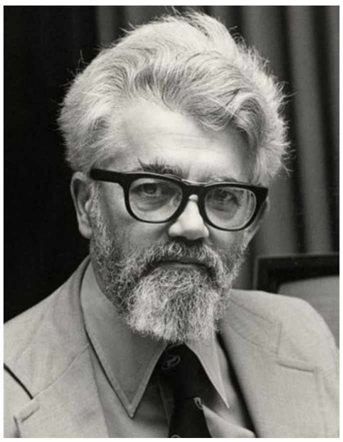

John McCarthy
September 4, 1927 – October 24, 2011

Abstract

John McCarthy, one of the "founding fathers" of artificial intelligence, coined the term Artificial Intelligence (AI) in 1955[1], and organized the famous Dartmouth conference in Summer 1956 that started AI as a field. It took over 60 years but AI has now reached the mainstream, starting a revolution that will change the way we design and engineer products using AI. Incorporation of AI into products brings several major challenges that will require us to adapt existing engineering processes. These challenges are driven by the complexity in understanding and predicting behavior of intelligent systems in a wide variety of circumstances, and the need to adapt to the new technologies and approaches driving intelligent and autonomous systems. Current processes tend to treat artificial intelligence models as just another piece of software, but these are fundamentally different from traditional software and existing methods are insufficient for dealing with these differences. This paper provides an introduction to the challenge ahead and establishes a framework for evolving our existing systems engineering processes to adapt to intelligent and autonomous systems. Subsequent papers will detail specific process areas and how AI will impact these along with suggestions for ways to adapt current methods.

Contents

What characterizes an "intelligent system"?

AI is notoriously difficult to define precisely. For our purposes there are two strongly defining characteristics that we think make AI-based algorithms fundamentally different than traditionally programmed algorithms:

1. Learning – the system learns to recognize patterns from data and is capable of continuously refining itself by learning from additional data.
2. Self-direction – the system makes choices that direct its learning without specific programming by a human. This means that in addition to learning, the choices made during the creation of the algorithm were of the system's own making, potentially without any human supervision.

In other words, the system makes its own inferences or decisions on the criteria and features that define the outputs of its AI-based algorithms. In autonomous systems, the decisions the system makes are determined by the outputs of those algorithms. And the system tunes its AI-based algorithms over time as it learns from "experience" (additional data inputs).

Impact of AI in automotive design and engineering

We are entering an era where technological change enabled by AI will force us to re-examine our traditional methods for designing and engineering automotive systems. This will be further forced by likely emergence of new regulatory frameworks that will govern AI-based systems, such as the European Commission guidelines on ethics in autonomous vehicles[2].

Challenges introduced by AI

1. **Functional behavior includes learning components**
   Most of our engineering approaches assume functionality that is controlled by predictable mathematical models (physics of engine or suspension components), or by human coded and explainable software algorithms. Engineers made explicit, traceable, reviewable and explainable design decisions. Learning components build up models from large amounts of training data to classify inputs and make decisions. The exact behavior of these models and how they come up with their outputs is largely opaque and cannot be easily explained, reviewed, inspected or debugged.
2. **Behavior continuously adapts and evolves over time**
   As more data is captured, the AI models are generally re-trained with this additional data, and the newer models may be deployed out in to existing products. Today, this training is centralized. In future, some products may undergo further continuous learning to better to adapt to their local environments and users. The continuous change of these models will create further challenges for testing and certification. Additional techniques such as evolutionary AI (which borrows concepts and techniques from genetics) are likely to make behavioral adaptation even more mainstream in future.
3. **Lack of human review and accountability**
   The most advanced and most common form of AI model in use in vehicle object recognition systems and other AI-based components today are deep learning models in which there are multiple "hidden" layers of neural networks, wired together[3]. It is extremely difficult to see exactly what part of the model contributes to the outcome. There is work underway to make these systems more explainable (see https://www.ibm.com/watson/explainable-ai), however, it is currently extremely difficult to enable humans to review AI models in the same way other engineering artifacts are reviewed today and may remain so. They can only really be reviewed through testing of outputs against test data, and inspection of training data and processes, which places even greater importance on testing processes.
4. **AI Training and Bias**
   The quality and accuracy of an AI model will depend entirely on the quality and extent of training data available. Engineering processes must address the need for massive amounts of training data, and in many cases may require extensive review to ensure accuracy and completeness, and additional work may be needed to manually label the training data[4]. Even with large volumes of high quality data, training AI-based systems will still face complexities such as bias due to inherent issues in the systems from which the data is extracted. The notion that AI systems can be biased has been well covered in the media[5] in recent months. Bias is generally caused by imbalances in training data, and because the systems we are employing AI in are systems and processes that are deeply connected in our society, and that therefore already reflect the natural biases at work in our society. The AI learns the biases inherent in the systems it is being embedded in to. There are emerging tools, such as IBM's Watson Openscale[6] that can help identify and mitigate the effect of bias in AI models. We must take conscious steps to mitigate this natural effect in the engineering of AI-based systems.
5. **Limited "awareness" (context for decision-making)**
   Human decisions are both blessed and cursed with a large amount of sensory input along with a lifetime's worth of "common sense" knowledge and other rules of the world. This gives us a critical ability to reason in the face of new situations or inadequate data. As drivers, we are continuously evaluating these learned rules as we drive (often subconsciously) – for example, hearing a sudden surge in volume of the engine of a vehicle beside us is an input we may subconsciously register as a cue that the vehicle is about to accelerate. Additional awareness in the form of current news, weather, local environmental knowledge, etc. gives us further context that informs our decision making. Today, most AI-based systems are fundamentally simple pattern sense-and-respond machines and are working with far more limited inputs and baseline knowledge than humans. When we replace human decision-making with AI in a system, we must acknowledge what

level of contextual awareness is used by the human decision maker and how to compensate for this. Autonomous systems are another factor that adds to the challenge of proving safe and secure development processes to auditors and assessors. OEMs and suppliers are well advised to select and set up their processes and tools to cope with this challenge already today [footnote ref to Daimler white paper]

Additional challenges for AI in automotive systems:

- Specificity of AI model to vehicle – modern automobile architectures have variability allowing for the manufacture of thousands of different options and configurations. Each of these may subtly change the vehicle dynamics and behavior, making it all the more difficult to prove that a particular AI model will work adequately in each specific vehicle configuration.
- Testing and validation – testing is an already complex process for automotive systems, proceeding through many different phases of development, from completely virtual testing based on models, to testing of prototype components and vehicles, to testing of final production vehicles. The immense complexity and variability of situations for autonomous vehicles adds exponential complexity to testing. Furthermore, it is nearly impossible to say with certainty when "adequate" coverage has been achieved. Finally, there are many proven techniques for identifying and testing boundary conditions for traditional systems (e.g. Boundary Value Analysis[7]). Such techniques are generally built on the assumption that it is possible to identify well-defined ranges of valid inputs or create mathematical models that represent good approximations of the behavior of the final systems. It is difficult to maintain confidence that such techniques have the same validity for AI-based systems, particularly in autonomous vehicles where the range of possible inputs is practically infinite. New and emerging standards for functional safety in autonomous systems such as SOTIF (ISO PAS 21448[8]) will place more requirements on manufacturers to address this validation complexity.
- Compliance and certification – government regulations often require that a manufacturer is able to demonstrate compliance to standards (such as ASPICE or ISO 26262) prior to delivery of a new vehicle model, and compliance is often validated through a certification process. There are very few regulations today to deal specifically with autonomous systems, but that will change as those systems become more prevalent, and as product liability potentially puts even greater onus on manufacturers and suppliers to ensure adequate governance over the engineering process for these intelligent systems.

All of these challenges will require adaptation of our current Systems Engineering (SE) methods and best practices. Our existing SE methodologies, for example as found in the INCOSE Systems Engineering Handbook[9] or the Harmony Agile MBSE Deskbook[10], have little or no accommodation for AI-based systems and are largely silent on the challenges described above. While current methods are still appropriate for identifying and selecting architectural alternatives based on stakeholder needs, and for capturing, modeling and mitigating critical risks, there are significant enhancements required to ensure that these methods adequately capture and communicate the special challenges, along with the increased opportunities, posed by AI technologies.

Outlook for future Systems Engineering methods and processes

There are several steps that Systems Engineers can take to start embracing the incorporation of AI in to product design:

1. Recognize that AI has a role not just in the system being designed, but in the design process itself. AI will augment the skills of engineering teams and assist the decision-making process. The first applications of AI in engineering processes are already visible, for example, IBM's Requirements Quality Assistant helps practitioners write better requirements using AI-trained natural language processing to assess the quality of written requirements[11].
2. Distinguish "learning" components from regular software components in the System architecture. This is the first step to ensuring that design reviews and quality verification steps appropriate for an AI-based component are used, and that all stakeholders are clear on where and how AI is being applied.
3. As noted above, AI depends heavily on data inputs so the System Engineering artifacts should include a greater emphasis on data as a first class element of the System architecture, the data flow that contributes to both the training of an AI model, and the inputs that the model uses in operation.
4. Ensure that the System Engineering process and system requirements explicitly capture any necessary ethics policies and procedures, such as assessment for bias, data privacy, and functional safety.
5. Account for the needed skills and expertise during the assessment and review of all aspects of AI in the system. Many organizations do not have a large enough experience base with AI yet to contain this in-house. There are many ways to supplement in-house skills with outside AI expertise as well as to build out training programs to provide in-house employees with valuable AI skills.
6. Finally, ensure you have a clear approach to closed-loop monitoring of AI-based systems in operation to gain immediate insight in to their real-world performance, user feedback, training deficiencies, and other unexpected issues so that these can be addressed rapidly.

Summary

There is little doubt that we are in the midst of a rapid transformation of many of our complex systems through the adoption of AI. We believe that the methods and processed used in engineering these systems need to be enhanced to address the significant challenges that AI technology poses. By adopting some of the recommendations listed above in our existing Systems Engineering methods, engineering organizations can ensure that appropriate discipline is applied, and issues and risks of AI are being addressed and mitigated. In addition, engineering organizations should begin looking forward by considering what tools and methods will be needed to support these challenges of engineering for intelligent and autonomous systems.

We will have more to say on specific issues and best practices outlined here in future white papers, with several planned topics including:

– Managing the lifecycle of AI models
– Planning for future applications of AI in the Engineering lifecycle
– Enhancing model-based engineering methods for intelligent and autonomous systems
– Testing and verification of intelligent and autonomous systems
– Managing compliance practices for intelligent and autonomous systems

About IBM Engineering
As products become more sophisticated, **with increasing integration of software and ubiquitous IoT connectivity**, the effort to engineer them increases exponentially, with a greater number of requirements, more dependency on modeling, greater testing, and increasing collaboration between teams. IBM Engineering Lifecycle Management (ELM) is an end-to-end systems development solution that empowers engineers and their teams to more easily manage requirements, workflow, test management and systems modeling. **ELM is the foundation for enabling engineering teams to embrace and leverage AI, analytics, and digital twin/digital thread in their development processes**. As the **leading solution** for engineering teams — particularly in highly regulated industries — ELM capabilities enable these teams to scale across multi-variant designs, systems-of-systems, geo-dispersed teams and supplier networks, while managing compliance and regulatory requirements...all to ultimately deliver products at greater speed, lower cost, and higher quality.

About KUGLER MAAG CIE
In everyday business, it all boils down to success. By ensuring the corporate strategy is implemented methodically and professionally, we help foster front-line innovation. With our consultation, your R&D successfully puts your ideas and innovations on the road.

As the leading consulting firm in automotive electronics development, Kugler Maag Cie provides both management consulting and process excellence. We design your process-oriented R&D organization adaptively to market-needs in a dynamic business environment.

Our experts pioneered the deployment of Agile development methods in Automotive R&D. Since 2010 we accompany the agile transformation of your corporation on all organizational levels and promote the genesis of an agile mind-set. Of course, your projects and processes will be consistently with such industry standards as Automotive SPICE and Functional Safety (ISO 26262). Our Automotive Security experts foster awareness within your organisation of comprehensive, end-to-end safeguards.

Looking for consultancy on Automotive electronics R&D? With us, you'll find it a lot easier.

Footnotes

1. https://en.wikipedia.org/wiki/John_McCarthy_(computer_scientist)
2. http://europa.eu/rapid/press-release_IP-19-1893_en.htm
3. https://medium.com/cracking-the-data-science-interview/the-10-deep-learning-methods-ai-practitioners-need-to-apply-885259f402c1
4. https://towardsdatascience.com/artificial-intelligence-get-your-users-to-label-your-data-b5fa7c0c9e00
5. https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/
6. https://www.ibm.com/cloud/watson-openscale
7. https://en.wikipedia.org/wiki/Boundary-value_analysis
8. https://www.iso.org/standard/70939.html
9. https://www.incose.org/products-and-publications/se-handbook
10. http://merlinscave.info/Merlins_Cave/Tutorials/Entries/2017/9/13_Harmony_aMBSE_Deskbook.html
11. https://www.ibm.com/us-en/marketplace/requirements-quality-assistant

IBM Corporation
New Orchard Road
Armonk, NY 10504
Produced in the United States of America, 2019