

TOWARDS SYSTEMS AI & DECISIONS INTELLIGENCE

Alexander Lavin*

Latent Sciences & NASA Frontier Development Lab

Traditional software systems are deterministic by nature, performing as programmed. Artificial intelligence (AI) and machine learning (ML) technologies learn from data, shifting the paradigm to systems that are inherently dynamic, and introducing known and unknown challenges in how these systems respond to their environment. Currently the approach to building AI technologies is siloed: models and algorithms are developed in testbeds isolated from real-world environments, and without the context of larger systems they'll be integrated within for deployment. A systems engineering approach to AI is needed, where systems engineering describes the principled processes and organizational frameworks that enable the cohesion and synergy of complex, interdependent subsystems.

Systems AI is a relatively nascent discipline, marked by several case studies developed over years of AI/ML in practice [1, 2, 3], and a foundational systems engineering framework for AI: *Technology Readiness Levels for Machine Learning (TRL4ML)* [4, 5]. Derived from the robust processes and testing standards of spacecraft development, TRL4ML is an industry-proven systems engineering framework, designed for efficient and reliable AI/ML research, productization, and deployment. TRL4ML defines a scalable, universal Systems AI framework and lingua franca, which has led to its rapid adoption across the AI industry – for example, with Google Cloud and Brain teams for bringing deep learning research into production, NASA's AI accelerator lab² and JPL for developing robust ML models in the context of larger hardware and software systems, Unity AI towards consumer computer vision applications, and deep tech startups from medical research to causal inference in finance and healthcare [4].

Decision Intelligence (DI) is a relatively unknown discipline for understanding the use of data and ML at scale, and for analyzing cause and effect in decision making. That is, DI answers the question “If I take this action today, what will be the outcome tomorrow?” A DI model links actions to outcomes, and makes assumptions and predictions explicit so they can be modified with data, feedback, or domain expertise over time. A proper DI framework can also be used to coordinate domains that are often siloed in practice, such as simulation, game theory, complexity, predictive analytics, and so on [6].

For DI systems to incorporate AI and related technologies properly, they must utilize Systems AI, namely the principled processes of TRL4ML. For example,

- ML-specific test suites are prescribed in TRL4ML that consider various types of data distribution shifts, performance regressions (even with hidden feedback loops), model mis-calibration, uncertainty quantification, and more factors that are critical to decision models in practice.
- DI is a discipline for analyzing chains of cause and effect, and TRL4ML includes specific protocol for causal AI. For example, underspecification is a key failure mode for practical ML pipelines [7], and testing causal assumptions helps mitigate this risk.
- Ethics in the use of data and AI is of critical importance with decision making systems. Ethics reviews are included in several of the framework's gated reviews. Even more, TRL4ML “Cards”, a key deliverable in the process, explicitly communicate ethical considerations (amongst other information) to all stakeholders.

Coupled with these overarching disciplines is the field of *data and information fusion (IF)*: the process of combining data or information to develop improved estimates or predictions of component or system states. In the concrete case, this can be one or many sensor data sources, which provide discrete quantitative measurements. In the more abstract case, IF includes human information that may manifest as domain expertise, intuition, error, priors or bias. Any useful framework for Systems AI must consider IF challenges and should leverage proven methodology, for example the classic JDL data fusion process model [8]. The TRL4ML approach includes IF processes, as IF and AI share the challenges of working with data in complex systems (e.g. cascading dependencies: upstream processes that impact them, and downstream processes they affect). Interestingly, we can draw parallels between the components of IF and AI: low-level data ingestion and fusion can be equivocated with machine perception, comprehension in IF models

*lavin@latentsci.com

²The Frontier Development Lab: nasa.ai

is analogous to representation learning and predictive modeling, and both have components of human (or end-user) interaction.

Nonetheless there is much to be done to improve TRL4ML and DI methods. Our driving aim is to develop *AI for rational decision making under uncertainties*. TRL4ML prioritizes risk quantification and mitigation, but can do better to make uncertainty quantification a first-class citizen – same goes for DI models. More importantly, the community should continue the adoption of TRL4ML as a Systems AI framework: To drive the development of robust, reliable AI/ML systems, and also to gain feedback from a rich variety of domains and use-cases.

Keywords: Systems Engineering; Artificial Intelligence; Decisions Intelligence; Data & Information Fusion

References

- [1] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald C. Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. Software engineering for machine learning: A case study. *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 2019.
- [2] Eric Breck, Shanjing Cai, E. Nielsen, M. Salib, and D. Sculley. The ml test score: A rubric for ml production readiness and technical debt reduction. *2017 IEEE International Conference on Big Data (Big Data)*, pages 1123–1132, 2017.
- [3] Siddha Ganju, Anirudh Koul, Alexander Lavin, J. Veitch-Michaelis, Meher Kasam, and J. Parr. Learnings from frontier development lab and spaceml - ai accelerators for nasa and esa. *ArXiv*, abs/2011.04776, 2020.
- [4] Alexander Lavin, Scott Penberthy, Massimo Mascaro, Danny Lange, Sujoy Ganguly, Siddha Ganju, Ian Goodfellow, and James Parr. Technology readiness levels for machine learning. *Applied AI Letters (to appear)*.
- [5] Alexander Lavin and Gregory Renard. Technology readiness levels for machine learning systems. *ICML Workshop on Challenges Deploying ML Systems*, 2020.
- [6] Loric Pratt. *Link: How Decision Intelligence Connects Data, Actions, and Outcomes for a Better World*. Emerald Publishing Limited, 2019.
- [7] Alexander D’Amour, K. Heller, D. Moldovan, Ben Adlam, B. Alipanahi, Alex Beutel, C. Chen, Jonathan Deaton, Jacob Eisenstein, M. Hoffman, Farhad Hormozdiari, N. Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, M. Lucic, Y. Ma, Cory Y. McLean, Diana Mincu, Akinori Mitani, A. Montanari, Zachary Nado, V. Natarajan, C. Nielson, Thomas F. Osborne, R. Raman, K. Ramasamy, Rory Sayres, J. Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, V. Veitch, Max Vladymyrov, Xuezhi Wang, K. Webster, S. Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. Underspecification presents challenges for credibility in modern machine learning. *ArXiv*, abs/2011.03395, 2020.
- [8] L. Snidaro, J. G. Herrero, J. Llinas, and E. Blasch. Recent trends in context exploitation for information fusion and ai. *AI Mag.*, 40:14–27, 2019.