Technology Readiness Levels for Machine Learning Systems

Alexander Lavin[*]
Latent Sciences

Ciaràn M. Gilligan-Lee
Spotify

Alessya Visnjic
WhyLabs

Siddha Ganju
Nvidia

Dava Newman
MIT

Sujoy Ganguly
Unity AI

Danny Lange
Unity AI

Atılım Güneş Baydin
Oxford University

Amit Sharma
Microsoft Research

Adam Gibson
Konduit

Yarin Gal
Alan Turing Institute

Eric P. Xing
Petuum

Chris Mattmann
NASA Jet Propulsion Lab

James Parr
NASA Frontier Development Lab

Abstract

The development and deployment of machine learning (ML) systems can be executed easily with modern tools, but the process is typically rushed and means-to-an-end. The lack of diligence can lead to technical debt, scope creep and misaligned objectives, model misuse and failures, and expensive consequences. Engineering systems, on the other hand, follow well-defined processes and testing standards to streamline development for high-quality, reliable results. The extreme is spacecraft systems, where mission critical measures and robustness are ingrained in the development process. Drawing on experience in both spacecraft engineering and ML (from research through product across domain areas), we have developed a proven systems engineering approach for machine learning development and deployment. Our *Machine Learning Technology Readiness Levels (MLTRL)* framework defines a principled process to ensure robust, reliable, and responsible systems while being streamlined for ML workflows, including key distinctions from traditional software engineering. Even more, MLTRL defines a lingua franca for people across teams and organizations to work collaboratively on artificial intelligence and machine learning technologies. Here we describe the framework and elucidate it with several real world use-cases of developing ML methods from basic research through productization and deployment, in areas such as medical diagnostics, consumer computer vision, satellite imagery, and particle physics.

Keywords: Machine Learning; Systems Engineering; AI Ethics; Medical AI; Space Sciences

Introduction

The accelerating use of artificial intelligence (AI) and machine learning (ML) technologies in systems of software, hardware, data, and people introduces vulnerabilities and risks due to dynamic and unreliable behaviors; fundamentally, ML systems learn from data, introducing known and unknown challenges in how these systems behave and interact with 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 or broader products they'll be integrated within for deployment. A main concern is models are typically trained and tested on only a handful of curated datasets, without measures and safeguards for future scenarios, and oblivious of the downstream tasks and users. Even more, models and algorithms are often integrated into a software stack without regard for the inherent stochasticity[ii] and failure modes of the hidden ML components.

Other domains of engineering, such as civil and aerospace, follow well-defined processes and testing standards to streamline development for high-quality, reliable results. *Technology Readiness Level (TRL)* is a systems engineering protocol for deep tech [2] and scientific endeavors at scale, ideal for integrating many interdependent components *and* cross-functional teams of people. No surprise TRL is standard process and parlance in NASA [3] and DARPA [4].

---

[*]lavin@latentsci.com

[ii]Consider the massive effect random seeds have on deep reinforcement learning model performance [1].

For a spaceflight project there are several defined phases, from pre-concept to prototyping to deployed operations to end-of-life, each with a series of exacting development cycles and reviews. This is in stark contrast to common machine learning and software workflows, which promote quick iteration, rapid deployment, and simple linear progressions. Yet the NASA technology readiness process for spacecraft systems is overkill; we need robust ML technologies integrated with larger systems of software, hardware, data, and humans, but not necessarily for missions to Mars. We aim to bring systems engineering to AI and ML by defining and putting into action a lean *Machine Learning Technology Readiness Levels (MLTRL)* framework. We draw on decades of AI and ML development, from research through production, across domains and applications: for example, computer vision in medical diagnostics and consumer apps, automation in self-driving vehicles and factory robotics, tools for scientific discovery and causal inference, streaming time-series in predictive maintenance and finance.

In this paper we define our framework for developing and deploying robust, reliable, and responsible ML systems, with several real test cases of advancing models and algorithms from R&D through productization and deployment. Additionally, MLTRL prioritizes the role of AI ethics and fairness, and our systems AI approach can help curb the large societal issues that can result from poorly deployed and maintained AI and ML technologies, such as the automation of systemic human bias, denial of individual autonomy, and unjustifiable outcomes (see the *Alan Turing Institute Report on Ethical AI* [5]). The adoption and proliferation of MLTRL provides a common nomenclature and metric across teams and industries. The standardization of MLTRL across the AI industry should help teams and organizations develop principled, safe, and trusted technologies.

TRL 0 **First Principles** A stage for greenfield research.

TRL 1 **Goal-oriented Research** Moving from basic principles to practical use.

TRL 2 **Proof of Principle (PoP) Development** Active R&D is initiated.

TRL 3 **Systems Development** Sound software engineering.

TRL 4 **Proof of Concept (PoC) Development** Demonstration in a real scenario.

TRL 5 **Machine Learning "Capability"** The R&D to product transition.

TRL 6 **Application Development** Robustification of ML modules, specifically towards one or more use-cases

TRL 7 **Integrations** ML infrastructure, product platform, data pipes, security protocols

TRL 8 **Flight-ready** The end of system development.

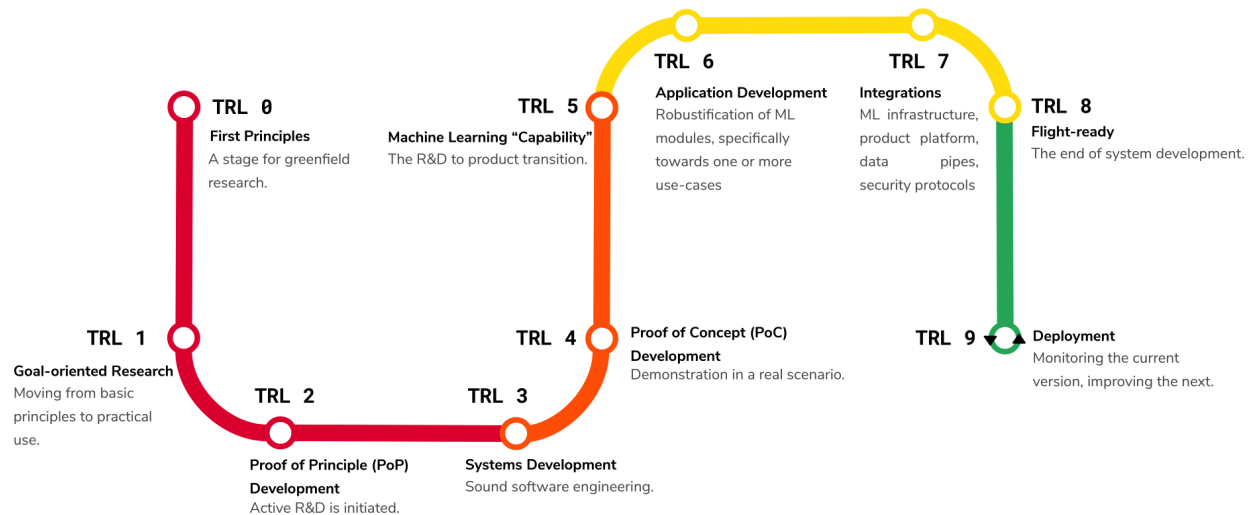**TRL 9** **Deployment** Monitoring the current version, improving the next.

Figure 1: MLTRL spans research through prototyping, productization, and deployment. Most ML workflows prescribe an isolated, linear process of data processing, training, testing, and serving a model [6]. Those workflows fail to define how ML development must iterate over that basic process to become more mature and robust, and how to integrate with a much larger system of software, hardware, data, and people. Not to mention MLTRL continues beyond deployment: monitoring and feedback cycles are important for continuous reliability and improvement over the product lifetime.

Results

MLTRL defines technology readiness levels (TRLs) to guide and communicate AI and ML development and deployment. A TRL represents the maturity of a model or algorithm[iii], data pipelines, software module, or composition thereof; a typical ML system consists of many interconnected subsystems and components, and the TRL of the system is the lowest level of its constituent parts. The anatomy of a level is marked by gated reviews, evolving working groups, requirements documentation with risk calculations, progressive code and testing standards, and deliverables such as TRL Cards (Figure 3) and ethics checklists. These components, as well as MLTRL metrics and methods, are discussed later in examples and in the Methods section.

[iii]Note we use "model" and "algorithm" somewhat interchangeably when referring to the technology under development. The same MLTRL process applies for e.g., a machine translation model or an algorithm for A/B testing.

MACHINE LEARNING TECHNOLOGY READINESS LEVELS

The levels are briefly defined as follows and in Figure 1, and elucidated with real-world examples later.

Level 0 - First Principles   This is a stage for greenfield AI research, initiated with a novel idea, guiding question, or poking at a problem from new angles. The work mainly consists of literature review, building mathematical foundations, white-boarding concepts and algorithms, and building an understanding of the data. For work in theoretical AI and ML there will not yet be data to work with (for example, a novel algorithm for Bayesian Optimization [7], which could eventually be used for many domains and datasets). The outcome is a set of concrete ideas with sound mathematical formulation, to pursue through low-level experimentation in the next stage. When relevant, we also expect conclusions about data readiness, including strategies for getting the data to be suitable for ML. To graduate, the basic principles, hypotheses, and research plans need to be stated, referencing relevant literature. The reviewer here is solely the lead of the research lab or team, for instance a PhD supervisor. With graduation, a *TRL Card* should be started to succinctly document the methods and insights thus far – this key MLTRL deliverable is detailed in the Methods section and Figure 3.

Level 1 - Goal-Oriented Research   To progress from basic principles to practical use, we design and run low-level experiments to analyze specific model/algorithm properties (rather than end-to-end runs for a performance benchmark score). The experiments, good results or not, and mathematical foundations need to pass a review process with fellow researchers before graduating to level 2; the panel for this gated review includes additional members of the research team. The application is still speculative, but through comparison studies and analyses we start to understand if/how/where the technology offers potential improvements and utility. Code is *research-caliber*: The aim here is to be quick and dirty, moving fast through iterations of experiments. Hacky code is okay, and full test coverage is actually discouraged, as long as the overall codebase is organized and maintainable. It is important to start semantic versioning practices early in the project lifecycle, which should cover code, models, *and* datasets. This is crucial for retrospectives and reproducibility, issues with which can be costly and severe at later stages. This versioning information and additional progress should be reported on the TRL Card (see for example Figure 3).

Level 2 - Proof of Principle (PoP) Development   Active R&D is initiated, mainly by developing and running in testbeds: simulated environments and/or surrogate data that closely matches the conditions and data of real scenarios – note these are driven by model-specific technical goals, not necessarily application or product goals (yet). An important deliverable at this stage is the formal research requirements document (with well-specified verification and validation (V&V) steps)[iv]. To graduate from the PoP stage, the technology needs to satisfy research claims made in previous stages, with the analyses well-documented and reproducible. Here is one of several *key decision points:* The R&D team considers several paths forward and sets the course: (A) prototype development towards Level 3, (B) continued R&D for longer-term research initiatives and/or publications, or some combination of A and B. We find the culmination of this stage is often a bifurcation: some work moves to applied AI, while some circles back for more research. This common MLTRL cycle is an instance of the non-monotonic *discovery switchback* mechanism (detailed in the Methods section).

Level 3 - System Development   Here we have checkpoints that push code development towards interoperability, reliability, maintainability, extensibility, and scalability. Code becomes *prototype-caliber*: A significant step up from research code in robustness and cleanliness. This needs to be well-designed, well-architected for dataflow and interfaces, generally covered by unit and integration tests, meet team style standards, and sufficiently-documented. Note the programmers' mentality remains that this code will someday be refactored/scrapped for productization; prototype code is relatively primitive with regard to efficiency and reliability of the eventual system. The Level 3 review includes teammates from applied AI and engineering. They'll focus on sound software practices, as well as version control for models and datasets. With the transition to Level 4 and proof-of-concept mode, the working group should evolve to include product engineering to help define service-level agreements and objectives (SLAs and SLOs) of the eventual production system.

Level 4 - Proof of Concept (PoC) Development   The aim is to demonstrate the technology in a real scenario. This stage is the seed of application-driven development; for many organizations this is the first touch-point with product managers and stakeholders beyond the R&D group. Thus TRL Cards and requirements documentation are instrumental

---

[iv]A *requirement* is a singular documented physical or functional need that a particular design, product, or process aims to satisfy. Requirements aim to specify all stakeholders' needs while not specifying a specific solution. Definitions are incomplete without corresponding measures for verification and validation (V&V). *Verification: Are we building the product right? Validation: Are we building the right product?*

in communicating the project status and onboarding new people. At this stage, quick proof-of-concept examples are developed to explore those application areas, and communicate the quantitative and qualitative results. Note the experiment metrics have most likely evolved from ML research to the applied setting: proof-of-concept evaluations should quantify model/algorithm performance (e.g., precision and recall) and computational costs (e.g., CPU vs GPU runtimes), and also metrics that are more relevant to the eventual end-user (e.g., number of false positives in the top-N predictions of a recommender system). We find this PoC exploration reveals specific differences between clean and controlled research data versus noisy and stochastic real-world data, issues that can then be targeted for further development. In review, we demonstrate the utility towards one or more practical applications (each with multiple datasets), taking care to communicate assumptions and limitations, and again reviewing data-readiness: evaluating the real-world data for quality, validity, and availability. The review also evaluates security and privacy considerations – defining these in the requirements document with risk quantification is a useful mechanism for mitigating potential issues (discussed further in the Methods section). AI ethics processes vary across organizations, but all should engage in ethics conversations at this stage, including ethics of data collection, and potential of any harm or discriminatory impacts due to the model (as the AI capabilities and datasets are known). MLTRL requires ethics considerations to be reported on TRL Cards at all stages, which generally link to an extended ethics checklist. The *key decision point* here is to push onward with application development or not. It is common to pause projects that pass Level 4 review, waiting for a better time to dedicate resources and/or pull the technology into a different project.

Level 5 - Machine Learning "Capability"   At this stage the technology is more than an isolated model or algorithm: it is a specific capability, such as producing depth images from stereo vision sensors on a mobile robot. In many organizations this represents a technology transition or handoff from R&D to productization. MLTRL makes this transition explicit, evolving the requisite work, guiding documentation, objectives and metrics, and team; indeed, without MLTRL it is common for this stage to be erroneously leaped completely, as shown in Figure 2. An interdisciplinary working group is defined, as we start developing the technology in the context of a larger real-world process – i.e., transitioning the model or algorithm from an isolated solution to a module of a larger application. By this stage the fidelity of technology has increased significantly, and it is straightforward for others in the organization to spin-up and run the ML capability; the knowledge and expertise cannot remain within the R&D team, let alone an individual ML developer. The verification and validation (V&V) measures and steps defined in earlier R&D stages (namely Level 2) must all be completed by now, and the product-driven requirements (and corresponding V&V) are drafted. Graduation from Level 5 should be difficult, as it signifies the dedication of resources to push this ML technology through productization. This transition is a common challenge in deep-tech, sometimes referred to as "the valley of death" because project managers and decision-makers struggle to allocate resources and align technology roadmaps to effectively move to Level 6, 7 and onward. MLTRL directly addresses this challenge by stepping through the technology transition or handoff explicitly.

Level 6 - Application Development   The main work here is significant software engineering to bring the code up to *product-caliber*: This code will be deployed to users and thus needs to follow precise specifications, have comprehensive test coverage, well-defined APIs, etc. The resulting ML modules should be robustified towards one or more target use-cases. If those target use-cases call for model explanations, the methods need to be built and validated alongside the ML model, and tested for their efficacy in faithfully interpreting the model's decisions – crucially, this needs to be in the context of downstream tasks and the end-users, as there is often a gap between ML explainability that serves ML engineers rather than external stakeholders [8]. The level review of this technology will mainly consider the code quality, the set of newly defined product requirements, system SLA and SLO requirements, data pipelines spec, and an AI ethics revisit now that we are closer to a real-world use-case. In particular, regulatory compliance is mandated for this gated review; the data privacy and security laws are changing rapidly, and missteps with compliance can make or break the project. Depending on the application, an ML model may not be deployable without restrictions; this typically means being embedded in a rules engine workflow where the ML model acts like an advisor that discovers edge cases in rules. The deployment setting(s) should be addressed thoroughly in the the product requirements document, as ML serving (or deploying) is an overloaded term that needs careful consideration. First, there are two main types: internal, as APIs for experiments and other usage mainly by data science and ML teams, and external, meaning an ML model that is embedded or consumed within a real application with real users. The serving constraints vary significantly when considering cloud deployment vs on-premise or hybrid, batch or streaming, open-source solution or containerized executable, etc. Even more, the data at deployment may be limited due to compliance, or we may only have access to encrypted data sources, some of which may only be accessible locally – these scenarios may call for advanced ML approaches such as federated learning [9] and other privacy-oriented ML [10]. These factors are hardly considered in model and algorithm development (although hardware choices typically are, such as GPU versus edge devices). It is crucial to make these systems decisions at Level 6–not too early that serving scenarios and requirements are uncertain, and not too late that corresponding changes to model or application development risk deployment delays or failures.

This marks a *key decision* for the project lifecycle, as this expensive ML deployment risk is common without MLTRL (see Figure 2).

Level 7 - Integrations   For integrating the technology into existing production systems, we recommend the working group has a balance of infrastructure engineers *and* applied AI engineers – this stage of development is vulnerable to latent model assumptions and failure modes, and as such cannot be developed solely by software engineers. The review should focus on the data pipelines and test suites; a scorecard like the ML Testing Rubric [11] is useful. We stress the need for tests that run use-case specific critical scenarios and data-slices – a proper risk-quantification table will highlight these. Additionally, a "golden dataset" should be defined to baseline the performance of each model *and succession of models*–see the computer vision app example in Figure 4–for use in the continuous integration and deployment (CI/CD) tests. These tests additionally help mitigate underspecification in ML pipelines, a key obstacle to reliably training models that behave as expected in deployment [12]. We additionally advocate for *metamorphic testing*: a software engineering methodology for testing a specific set of relations between the outputs of multiple inputs. When integrating ML modules into larger systems, a codified list of metamorphic relations [13] can provide valuable verification and validation measures and steps. It is important that quality assurance engineers (QA) play a key role here and through Level 9, and prioritize *data governance*: how data is obtained, managed, used, and secured by the organization. QA for AI would additionally need to cover auditing for downstream accountability of AI methods.

Level 8 - Flight-ready   The technology is demonstrated to work in its final form and under expected conditions. There should be additional tests implemented at this stage covering deployment aspects, notably A/B tests, blue/green deployment tests, shadow testing, and canary testing, which enable proactive and gradual testing for changing ML methods and data. Ahead of deployment, the CI/CD system should be ready to regularly stress test the overall system and ML components. In practice, problems stemming from real-world data are impossible to anticipate and design for – an upstream data provider could change formats unexpectedly or a physical event could cause the customer behavior to change. Running models in shadow mode for a period of time would help stress test the infrastructure and evaluate how susceptible the ML model(s) will be to performance regressions caused by data. We observe that ML systems with *data-oriented architectures* are more readily tested in this manner, and better surface data quality issues, data drifts, and concept drifts – this is discussed later in the Beyond Software Engineering section. To close this stage, the review is a diligent walkthrough of every technical and product requirement, showing the corresponding validations, and the review panel is representative of the full slate of stakeholders. The *key decision* is go or no-go for deployment, and when.

Level 9 - Deployment   In deploying AI and ML technologies, there is significant need to monitor the current version, and explicit considerations towards improving the next version; for instance, performance degradation can be hidden and critical, and feature improvements often bring unintended consequences and constraints. Thus at this level, maintenance engineering (i.e. monitoring and update methods) takes over. Monitoring for data quality, concept drift, and data drift is crucial; no AI system without thorough tests for these can reliably be deployed. Similarly, if actuals [14] are available, continuous evaluation should be enabled. In many cases actuals come with a delay, so it is essential to record model outputs in a way that can enable evaluation down the road. Monitoring for data quality issues and data drifts is crucial to catch deviations in model behavior, particularly those that are non-obvious in end-performance. To enable monitoring, the ML pipeline should be instrumented to log system metadata, model metadata, and data itself. Data logging is unique in the context of ML systems – data logs should capture statistical properties of input features and model predictions. Logs must be monitored for anomalies, and identified anomalies should be accessible to all relevant stakeholders (i.e., applied and research engineers, product managers, and system operators). With monitoring data, concept, and model drifts, the logs are to be sent to the relevant applied *and* research engineers. The latter is often non-trivial, as the model server is not ideal for model "observability" because it does not necessarily have the right data points to link the complex layers needed to analyze and debug models. To this end, MLTRL requires the drift tests to be implemented at earlier stages well ahead of deployment. Again we advocate for data-first architectures rather than the software industry-standard design by services (discussed later), which aids in surfacing and logging the relevant data types and slices when monitoring AI systems. Regarding retraining and improving models, monitoring must be enabled to catch training–serving skew and let the team know when to retrain. Towards model improvements, adding or modifying features can often have unintended consequences, such as introducing latencies or even bias. To mitigate these risks, MLTRL has an *embedded switchback* here: any component or module changes to the deployed version must cycle back to Level 7 (integrations stage) or earlier. Additionally, for quality ML products, we stress a defined communication path for user feedback without roadblocks to R&D; we encourage real-world feedback all the way to research, providing valuable problem constraints and perspectives.

Notice MLTRL is defined as stages or levels, yet much of the value in practice is realized in the transitions: MLTRL enables teams to move from one level to the next reliably and efficiently, and provides a guide for how teams and objectives evolve with the progressing technology.
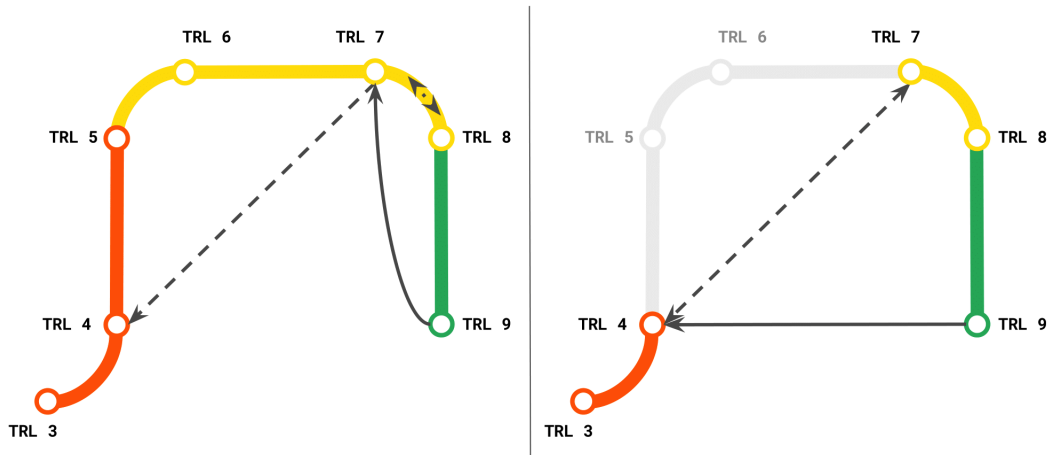
Figure 2: Most AI/ML projects live in these sections of MLTRL, not concerned with fundamental R&D – that is, completely using existing methods and implementations, and even pretrained models. In the left diagram, the arrows show a common development pattern with MLTRL in industry: projects go back to the ML toolbox to develop new features (dashed line), and frequent, incremental improvements are often a practice of jumping back a couple levels to Level 7 (which is the main systems integrations stage). At Levels 7 and 8 we stress the need for tests that run use-case specific critical scenarios and data-slices, which are highlighted by a proper risk-quantification matrix [15]. Cycling back to previous lower levels is not just a late-stage mechanism in MLTRL, but rather "switchbacks" occur throughout the process (as discussed in the Methods section and throughout the text). In the right diagram we show the more common approach in industry (*without* using our framework), which skips essential technology transition stages – ML Engineers push straight through to deployment, ignoring important productization and systems integration factors. This will be discussed in more detail in the Methods section.

Discussion

In practical AI and ML scenarios there are many challenges and risks that MLTRL aims to mitigate. Failure to address these can lead to unintended and often expensive consequences, such as medical misdiagnosis and biased decision-making. We now highlight some of the main challenges, discussed with MLTRL strategies in the context of real-world test-cases.

EXAMPLES

Human-in-the-loop neuropathology

Neuropathology, the microscopic examination of neurosurgical specimens for cancerous tissue, is a challenging, labor intensive task that can improve in accuracy and efficiency with recent computer vision methods. Naud & Lavin [16] present an interesting approach that leverages an unsupervised generative vision model to produce low-dimensional visualization of brain tissue types, which automatically identifies the specific tissue samples for further inspection by a medical expert. The project began at Level 0 with theoretical ML work on manifold geometries, and at Level 5 was directed towards development a medical decision-making tool.

Hidden feedback loops can be common and problematic in real-world systems influencing their own training data: over time the behavior of users may evolve to select data inputs they prefer for the specific AI system, representing some skew from the training data. In this case, selecting whole-slide images that are uniquely difficult for manual inspection, or even biased by that individual user. Similarly we see underlying healthcare processes can act as hidden confounders, resulting in unreliable decision support tools [18]. MLTRL helps by requiring ML-specific tests (in this case shadow testing and monitoring data invariants), and TRL Cards that make hidden feedback loops obvious to all (Figure 3).

Model availability can be limited in many deployment settings: for example, on-premises deployments (common in privacy preserving domains like healthcare and banking), or from the infrastructure's inability to scale to the volume of requests. This can severely limit the team's ability to monitor, debug, and improve deployed models. MLTRL mitigates this concern in several ways, mainly with the explicit transition from tech-driven requirements in R&D to product-driven requirements at Level 6; the research engineers coordinate with product engineers and managers, to
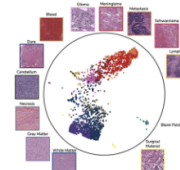
TECHNOLOGY NAME	Neuropathology Copilot v1.0		Model / alg details	The SP-VAE model runs unsupervised on neurological whole-slide images (WSI), producing a latent manifold that represents a hierarchical organization of tissue types. An medical expert identifies several data points to inspect. *Example visualization of the latent organization of brain tissue types.*
TRL	4 <link to previous cards>			
R&D OWNER / REVIEWER	A. Lavin / G. Renard			
PROD OWNER / REVIEWER	S. Wozniak / S. Jobs			
COMPONENT CODES	1.1, 4.2, 4.3			
TL;DR	Analyze WSI of brain tissue in 3 main steps: (1) unsupervised CV model produces Poincare manifold viz (Naud & Lavin '20), (2) domain expert selects data points, (3) U-Net classifier		Metrics, results	Classification accuracy >0.97 on the 5 main brain cancer types. Inference per WSI runs ~1.0s on 2-GPU. Full quantitative reports: < link to experiments wiki >
Data considerations	3 datasets have been used to train and validate the system: 1. Open dataset (Naud & Lavin '20) 2. Pilot dataset provided by BioLab, v1.0 3. Simulated datasets (w/ structured domain randomization), v2.3		Caveats, known edge cases, recommendations	Changing imaging sources will require retraining the full model (notably the SP-VAE annealing parameter). Whenever possible it is recommended that users provide feedback annotations. Non-tissue material is correctly flagged as anomalous.
Ethics	Note the demographics info on specific Dataset Cards. Datasets anonymized, pipeline runs w/o metadata. The Latent Sciences Ethics Checklist has been completed.		Key assumptions	The training and production images are equivalent, specifically from the exact same sensor(s).
			Intended use	The model must include human expert in the loop, and it has not yet been validated for other disease areas.

Figure 3: The maturity of each ML technology is tracked via *TRL Cards*, which we describe in the Methods section. Here is an example reflecting a neuropathology machine vision use-case [16], detailed in the Discussion Section. Note this is a subset of a full TRL Card, which in reality lives as a full document in an internal wiki. Notice the card clearly communicates the data sources, versions, and assumptions. This helps mitigate invalid assumptions about performance and generalizability when moving from R&D to production, and promotes the use of real-world data earlier in the project lifecycle. We recommend documenting datasets thoroughly with semantic versioning and tools such as *datasheets for datasets* [17].

work on model and deployment constraints. Typically this coordination does not happen–teams are often ignorant of the work and coordination done in Levels 4 through 7–leading to expensive consequences after deployment. This was not an issue with the neuropathology tool we developed.

Uncertainty quantification and propagation are quite valuable in many AI scenarios, yet not straightforward to implement in practice. In medical applications it is of critical importance, to provide measures of confidence and sensitivity, and for AI researchers through end-users. From the systems perspective, we suggest quantifying the uncertainties of components and propagating them through the system, which can improve safety. Probabilistic ML methods, rooted in Bayesian probability theory, provide a principled approach to representing and manipulating uncertainty about models and predictions [19]. For this reason we advocate strongly for probabilistic models and algorithms in AI systems. In this machine vision example, the MLTRL technical requirements specifically called for a probabilistic generative model to readily quantify various types of uncertainties and propagate them forward to the visualization component of the pipeline, and the product requirements called for the downstream confidence and sensitivity measures to be exposed to the end-user. Component uncertainties must be assembled in a principled way to yield a meaningful measure of overall system uncertainty, based on which safe decisions can be made [20]. See the Methods section for more on uncertainty in AI systems.

Costs of edge cases can be significant, sometimes risking expensive machine downtime, or medical failures. MLTRL helps identify edge cases and mitigate the risks in several ways: First with explicit risk metrics that are estimated for every technical and product requirement. The edge cases are communicated to all team members and stakeholders via TRL Cards (see Figure 3), and corresponding tests are prescribed (see Level 7). In some cases, the costs of edge cases are such that AI-based automation is not a safe or reliable option–the AI pipeline will incorporate human intervention or other decision logic to ensure these are handled reliably. Working with domain experts–a theme in MLTRL–encourages human-in-the-loop solutions when they are practical. Human-in-the-loop methods can vary across applications; automation scenarios are a spectrum. This "neuropathology copilot" product is designed for human-computer collaboration, not full automation. The post-deployment lifecycle aims to increase automation over time (with real-world validation and building trust), while always keeping medical experts in the loop. Notably this was a technical requirement defined early on, in Level 4.

End-user trust can be difficult to achieve, often preventing the adoption of ML applications, particularly in the healthcare domain. MLTRL has several mechanisms for improving the trust and confidence in AI and ML technologies, notably stakeholder involvement and communication. Gated reviews are useful for sharing progress with stakeholders that is most relevant to them (i.e. progress towards developing the goal) instead of showing only technical advances.
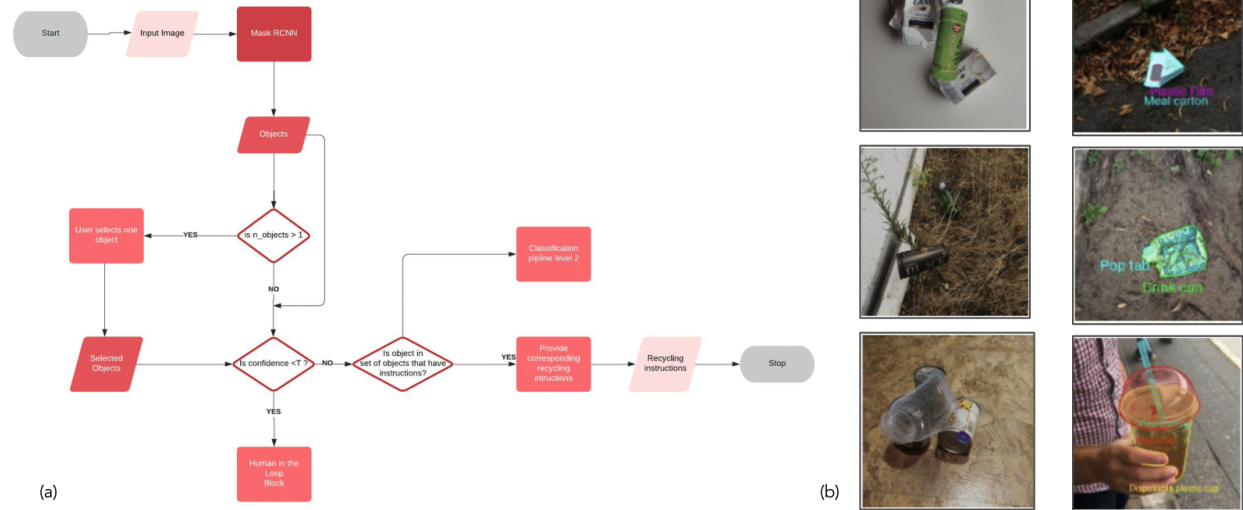
Figure 4: Computer vision pipeline for an automated recycling application (a), which contains multiple ML models, user input, and image data from various sources. Complicated logic such as this can mask ML model performance lags and failures, and also emphasized the need for R&D-to-product hand off described in MLTRL. Additional emphasis is placed on ML tests that consider the mix of real-world data with user annotations (b, right) and synthetic data generated by Unity AI's Perception tool and structured domain randomization (b, left).

Levels 3-5 are appropriate stages for engaging clinicians, medical researchers, and enterprise-level decision makers, which in this example payed dividends in building a usable application for neuropathologists. Unfortunately many healthcare AI projects delay this stakeholder involvement and incorrectly assume end-user trust is a given with good technical results.

Project roadmapping and budgeting is always challenging when ML technologies are integrated within product roadmaps and business objectives. The challenge is exacerbated in multidisciplinary projects such as this medical AI example, where not all stakeholders can assess arcane ML maturity definitions. Yet with MLTRL, all can understand that more time and budget is required to get from Level 4 to 9, and the process definitions make it clear how to progress. The dividends we have seen from integrating this into project management are rather significant.

Computer vision with real and synthetic data

Advancements in physics engines and graphics processing have advanced AI environment and data-generation capabilities, putting increased emphasis on transitioning models across the simulation-to-reality gap [21, 22, 23]. To develop a computer vision application for automated recycling, we leveraged Unity Perception[v], a toolkit for generating large-scale datasets for perception-based ML training and validation. Using Unity Perception we produced synthetic images to complement real-world data sources (Figure 4). This application exemplifies two challenges in ML product development that MLTRL helps overcome: handling mixed data sources and surfacing hidden performance degradation.

Multiple and disparate data sources are common in deployed ML pipelines, yet often ignored in R&D. In this example, we implemented pipelines and extended test suites to cover open-source benchmark data, real user data, and synthetic data. We had to assume bias for the real data sources due to class-imbalance and imperfect annotations, and implement additional CI/CD tests to monitor these in the deployed application – this was communicated to deployment engineers from TRL Cards at early R&D stages. Furthermore, given we were generating synthetic images (with structured domain randomization) in order to complement the real data, tests for different data distribution shifts were needed at multiple points in the classification pipeline, and thorough shadow tests were implemented ahead of deployment. Interestingly, we recognized the need for specific types of synthetic images early on because of the data-readiness evaluations prescribed by MLTRL in the earliest stages, ahead of actually developing any models. Typically data quality and availability issues present huge blockers because teams discover them too late in the game.

---

[v]github.com/Unity-Technologies/com.unity.perception

Hidden performance degradation can be challenging to detect and debug in real ML systems because gradual changes in performance may not be immediately visible. Common reasons for this challenge are that the ML component may be one step in a series. Additionally, local/isolated changes to an ML component's performance may not directly affect the observed downstream performance. We can see both issues in the illustrated logic diagram for our app (Figure 4). For example, a slight degradation in the initial CV model may not heavily influence the following user input. However, when an uncommon input image appears in the future the app fails. The monitoring tests prescribed at Levels 8 and 9 and the three main code quality checkpoints in the MLTRL process help surfacing the hidden performance degradation problem. ML pipelines often grow out of glue code, with poorly defined borders and abstractions between components. With code checkpoints we aim for well-architected software that minimizes these danger spots.

Model usage requirements can make or break an ML product. Memory usage, compute power requirements, hardware availability, network privacy and latency, etc. should be clearly communicated when working with data science teams who may only understand the statistics or ML theory behind a model, but not its system requirements and how it scales. MLTRL accomplishes this communication with requirements documentation and evolving working groups that mix expertise. Related to the deployment scenarios discussion earlier in Level 6, a higher bar should be employed for real-time systems or models that are hard to update, which was exactly the case in this consumer computer vision app.

Accelerating scientific discovery with massive particle physics simulators

Computational models and simulation are key to scientific advances at all scales, from particle physics, to material design and drug discovery, to weather and climate science, and to cosmology [24]. Many simulators model the forward evolution of a system (coinciding with the arrow of time), such as the interaction of elementary particles, diffusion of gasses, folding of proteins, or evolution of the universe in the largest scale. The task of inference refers to finding initial conditions or global parameters of such systems that can lead to some observed data representing the final outcome of a simulation. In probabilistic programming [25], this inference task is performed by defining prior distributions over any latent quantities of interest, and obtaining posterior distributions over these latent quantities conditioned on observed outcomes (for example, experimental data) using Bayes rule. This process, in effect, corresponds to inverting the simulator such that we go from the outcomes towards the inputs that caused the outcomes. In the "Etalumis" project [26] ("simulate" spelled backwards), we are using probabilistic programming methods to invert existing, large-scale simulators via Bayesian inference. The project is as an interdisciplinary collaboration of specialists in probabilistic machine learning, particle physics, and high-performance computing, all essential elements to achieve the project outcomes. Even more, it is a multi-year project spanning multiple countries, companies, university labs, and government research organizations, bringing significant challenges in project management, technology coordination and validation.

Integrating with legacy systems is common in scientific and industrial use-cases, where ML methods are applied with existing sensor networks, infrastructure, and codebases. In this case, particle physics domain experts at CERN are using the SHERPA simulator [27], a 1 million line codebase developed over the last two decades. Rewriting the simulator for Bayesian inference is infeasible, and the ML experts would need significant onboarding to gain working knowledge of the codebase. Defining simulator-integration constraints upfront as requirements (which is often overlooked in early R&D levels), and working hand-in-hand with CERN scientists, led to an elegant solution: a probabilistic programming execution protocol was developed to reroute random number draws in the stochastic simulator codebase (or any codebase in any language for that matter) to the probabilistic programming system, thus keeping the legacy codebase intact. This mechanism allows the probabilistic programming system to intelligently direct the stochastic choices in the simulator to achieve a desired outcome, and in the process inferring the conditions that lead to any given observed outcome of the system. A more invasive method that modifies SHERPA would not have been acceptable.

Coupling hardware and software architectures is non-trivial when deploying ML at scale, as performance constraints are often considered in deployment tests well after model and algorithm development, not to mention the expertise is often split across disjoint teams. The Etalumis test case takes this to the extreme, where high-performance computing infrastructure was needed throughout the project lifecycle, specifically working with NERSC engineers and their Cori supercomputer. It was instrumental to define computing requirements and corresponding V&V tests at multiple levels – early experimentation (before Level 4), proof-of-concept prototyping, and system deployment – to progressively scale to the goal supercomputing deployment scenario. Typical ML workflows that follow simple linear progressions [6, 28] would not enable ramping up in this fashion. Additionally, the natural "embedded switchback" from Level 4 to 2 (see the Methods section) provided an efficient path toward developing an improved, amortized inference method–i.e., using a computationally expensive deep learning based inference algorithm to train only once, in order to then do fast, repeated inference in the SHERPA model. Leveraging cyclic R&D methods, the Etalumis project could iteratively improve inference methods without stalling the broader system development, ultimately producing the largest scale posterior inference in a Turing-complete probabilistic programming system.

Interpretability is often a desired feature yet difficult to deliver and validate in practice. Particularly in scientific ML applications such as this, mechanisms and tooling for domain experts to interpret predictions and models are key for usability. Understanding this end-user need early in the project is another reason why the model-based approach afforded by probabilistic programming was chosen. That is, when we work with a simulator codebase by interfacing with it at the level of probabilistic sampling, the simulator is not used as a black box. This is in contrast to alternative approaches, such as sampling labeled synthetic data from a simulator, discarding the simulator, and training deep learning models with this synthetic data (such as the main methods used recently in computer vision applications). In the probabilistic generative model setting we've defined, Bayesian inference gives results that are interpretable because they include exact locations and processes in the model that are associated with each prediction. Starting in R&D stages with ML methods that are inherently interpretable, we are well-positioned to deliver interpretable interfaces for the end-users later in the project lifecycle.

Causal inference, from laboratory domain to deployment

Understanding cause and effect relationships is crucial for accurate and actionable decision-making in many settings, from healthcare and epidemiology, to economics and government policy development. Unfortunately, standard machine learning algorithms can only find patterns and correlation in data, and as correlation is not causation, their predictions cannot be confidently used for understanding cause and effect. Indeed, relying on correlations extracted from observational data to guide decision-making can lead to embarrassing, costly, and even dangerous mistakes, such as concluding that asthma reduces pneumonia mortality risk [29]. Fortunately, there has been much recent development in a field known as causal inference that can quantitatively make sense of cause and effect from purely observational data [30]. The ability of causal inference algorithms to quantify causal impact rests on a number of important assumptions that must be carefully deliberated over during their development and training. The MLTRL framework makes transparent the need to carefully document and defend these assumptions, thus ensuring the safe and robust creation, deployment, and maintenance of causal models. We elucidate this with recent work by Richens et al. [31], developing a causal approach to computer-assisted diagnosis which outperforms previous purely machine learning based methods.

Grounded assumptions and proofs are explicitly required in the earliest stages of MLTRL. Unlike purely predictive models, causal inference requires an extensive modeling and identification exercise before any models can be fit, wherein one specifies the cause-and-effect relationships between relevant variables and decides whether it is theoretically possible to estimate the target effect from data. A very important assumption underlying the performance of the diagnostic model was the causal relationships between the diverse set risk factors, diseases, and symptoms included in the model. To learn these relations, doctors and healthcare professionals were consulted to employ their expansive medical domain knowledge which was robustly evaluated by additional independent groups of healthcare professionals. The MLTRL framework ensured this issue is dealt with and documented correctly, as such knowledge is required to progress from Level 1; failure to do this has plagued similar healthcare AI projects [32]. Another vital component of building causal models is whether the causal question of interest is *identifiable* from the causal structure specified for the model. In this medical diagnosis example, identification was crucial to establish, as the causal question of interest, "would the observed symptoms not be present had a specific disease been cured?", was highly non-trivial. Again, MLTRL ensures this vital aspect of model building is carefully considered, as a mathematical proof of identifiability would be required to graduate from Level 0.

Adjusting for and monitoring confounding bias is an important aspect of causal model performance which is not present in standard machine learning algorithms, and often overlooked in the context of broader AI systems. The standard approach is to employ propensity score-based methods to remove such bias when the target decision is binary, and use multi-stage ML models adhering to the assumed causal structure [33] for continuous target decisions (and high-dimensional data in general). However, the quality of bias adjustment achieved in any specific instance with propensity-based matching methods needs to be checked and documented, with alternate bias adjusting procedure required if appropriate levels of bias adjustment are not achieved [34]. As above, MLTRL ensures transparency and adherence to this important aspect of causal model development, as without it a project cannot graduate from Level 2. Even more, MLTRL ensures tests for confounding bias are developed early-on and maintained throughout later stages to deployment. Still, in many cases, it is not possible to completely remove confounding in the observed data. TRL Cards offer a transparent way to declare specific limitations of a causal ML method.

Sensitivity analysis is often limited to R&D experiments or a post-hoc feature of ML products. MLTRL requires this throughout the lifecycle as components of ML test suites and gated reviews. In the case of causal ML, sensitivity analysis is the last important check a causal model must go through to be fit for deployment, to ensure that its predictions are robust to potential violations of the form of the causal structure such as the presence of missing or unobserved confounders. Best practice is to employ sensitivity analysis for this robustness check [35]. MLTRL ensures this check is highlighted and adhered to, and no model will end up graduating Level 2–let alone being deployed–unless it is passed.

PoC-level tests for causal models must go beyond that of typical ML models. *Consistency*, for example, is an important property of causal methods that informs us whether the method provably converges to the true causal graph in the limit of infinite sample size. Quantifying consistency in the test suite is critical when datasets change from controlled laboratory settings to open-world, and when the application scales. And PoC validation steps are more efficient with MLTRL because the process facilitates early specification of the evaluation metric for a causal model in Level 2. Causal models cannot be validated by standard held-out tests, but rather require randomization or special data collection strategies to evaluate their predictions [36, 37]. Any difficulty in evaluating the model's predictions will be caught early and remedied.

AI for planetary defense

The CAMS (Cameras for Allsky Meteor Surveillance) project [38], established in 2010 by NASA, uses hundreds of off-the-shelf CCTV cameras to capture the meteor activity in the night sky. Initially, resident scientists would each night perform the triangulation of tracks or streaks of light in the night sky, captured by two or more cameras, and compute a meteor's trajectory, orbit, and lightcurve. Each solution was manually classified as a meteor or not (i.e., planes, birds, clouds, etc). In 2017, a project run by the Frontier Development Lab, the AI accelerator for NASA and ESA, aimed to automate the data processing pipeline and filtration of meteors in the CAMS project [39, 40]. The automation led to orders of magnitude improvements in operational efficiency of the system, and coordination with open-source contributors led to four-fold global expansion of the data capture network.

Regular gated reviews are essential for making efficient progress while ensuring robustness and functionality that meets stakeholder needs. In this example, reviews stressed the significance of numerical improvements and comparison to existing baselines, and helped identify and overcome issues with data imbalance–night sky activity is governed by very few meteors relative to non-meteors. The evolving panel of reviewers at different stages of the project was essential for covering a variety of verification and validation measures – from open-source code quality to A/B testing a novel web tool[vi]. The Frontier Development Lab, in developing AI solutions for NASA, ESA, and many sponsor organizations, churns out batches of successful AI projects largely because of its weekly, well-defined and scored reviews with expert peers and stakeholders [41].

Quantifiable progress was critical for verifying the tech-readiness of AI models and supporting tools in the context of the larger CAMS system, which had started development years earlier with other teams in partner organizations. At the component level, quantifying robustness with ML testing suites (using scoring measures like that of the ML Testing Rubric [11], and devising a checklist based on metamorphic testing [13]) was crucial for integrating open-source contributions into the main AI subsystem, and checking off required verification measures. For integrating the AI subsystem into the larger CAMS system, the use of technology readiness levels provided a clear metric for the maturity of the AI technologies, making for clear communication and efficient project integration. Without MLTRL it is difficult to have a conversation, let alone make progress, towards integrating AI and data subsystems and components.

Flight-proven at TRL-9, the AI components in CAMS require continual monitoring for model and data drifts, such as weather changes, smoke, and cloud patterns that affect the view of the night sky. The AI pipeline is largely automated with CI/CD, runs regular regression tests, and production of benchmarks. Manual intervention can be triggered when needed, such as sending low confidence meteors for verification to scientists in the CAMS project. The team also regularly releases the code, models, and web tools on the open-source space sciences and exploration AI toolbox, SpaceML[vii]. Releasing a robust, flight-proven AI pipeline to open-source accelerated adoption by the AI community and citizen scientists to further test and improve the CAMS tooling.

BEYOND SOFTWARE ENGINEERING

There are several key areas where machine learning (ML) development is unique from software engineering (SWE). For instance, the behavior of ML systems is learned from data, not specified directly in code. The data requirements around ML (i.e., data discovery, management, and monitoring) adds significant complexity not seen in other types of SWE. Not to mention an array of ML-specific failure modes; for example, models that become mis-calibrated due to subtle data distributional shifts in the deployment setting, resulting in models that are more confident in predictions than they should be. Plus, ML opens up new threat vectors across the whole deployment workflow: for example, a poisoning attack to contaminate the training phase of ML systems, or membership inference to see if a given data record was part of the model's training. These factors and others suggest common software architectural patterns and management practices do not lend themselves to AI systems. There are many benefits to using a *data-oriented architecture (DOA)* [32] with the data-first workflows and management practices prescribed in MLTRL. DOA aims

---

[vi]cams.seti.org/FDL
[vii]spaceml.org

to make the data flowing between elements of business logic more explicit and accessible with a streaming-based architecture rather than the micro-service architectures that are standard in software systems. One specific benefit of DOA is making data available and traceable by design, which helps significantly in the ML logging challenges and data governance needs we discussed in Levels 7-9.

More generally, ML codebases have all the problems for regular code, plus ML-specific issue at the system level, mainly as a consequence of added complexity and dynamism. The resulting entanglement, for instance, implies that the SWE practice of making isolated changes is often not feasible – Scully et al. [42] refer to this as the "changing anything changes everything" principle. Another result is that ML systems almost necessarily increase the technical debt; package-level refactoring is generally sufficient for removing technical debt in software systems, but this is not the case in ML systems.

RELATED WORKS

A recent case study from Microsoft Research [28] similarly identifies a few themes describing how ML is not equal to software engineering, and recommends a linear ML workflow with steps for data preparation through modeling and deploying. They define an effective workflow for isolated development of an ML model, but this approach does not ensure the technology is actually improving in quality and robustness. Their process should be repeated at progressive stages of development in the broader technology lifecycle. If applied in the MLTRL framework, the specific steps of the ML model workflow–that is, people, software, tests, objectives, etc.–evolve over time and subsequent stages. Gelman et al. [43] go in depth to describe a more iterative process for Bayesian ML. These recommended workflows do not consider an ML model in the context of larger systems; in many projects, the typical model development workflows they describe occur within each stage of MLTRL. Also related to our work, Google teams have proposed ML testing recommendations [11] and validating the data fed into ML systems [44]. For NLP applications, typical ML testing practices struggle to translate to real-world settings, often overestimating performance capabilities. An effective way to address this is devising a checklist of linguistic capabilities and test types, as in Ribeiro et al. [45]–interestingly their test suite was inspired by metamorphic testing, which we suggested earlier in Level 7 for testing systems AI integrations. A survey by Paleyes et al. [32] go over numerous case studies to discuss challenges in ML deployment. They similarly pay special attention to the need for ethical considerations, end-user trust, and extra security in ML deployments. On the latter point, Kumar et al. [46] provide a table thoroughly breaking down new threat vectors across the whole ML deployment workflow (some of which we mentioned above). These works, notably the ML security measures and the quantification of an ML test suite in a principled way – i.e., that does not use misguided heuristics such as code coverage – are valuable to include in any ML workflow including MLTRL, and are synergistic with the framework we've described in this paper. These analyses provide useful insights, but they do not provide a holistic, regimented process for the full ML lifecycle from R&D through deployment. An end-to-end approach is suggested by Raji et al. [47], but only for the specific task of auditing algorithms; components of AI auditing are mentioned in Level 7, and covered throughout in the review processes.

Sculley et al. [42] go into more ML debt topics such as undeclared consumers and data dependencies, and go on to recommend an ML Testing Rubric as a production checklist [11]. For example, testing models by a canary process before serving them into production. This, along with similar shadow testing we mentioned earlier, are common in autonomous ML systems, notably robotics and autonomous vehicles. They explicitly call out tests in four main areas (ML infrastructure, model development, features and data, and monitoring of running ML systems), some of which we discussed earlier. For example, tests that the training and serving features compute the same values; a model may train on logged processes or user input, but is then served on a live feed with different inputs. In addition to the Google ML Testing Rubric, we advocate *metamorphic testing*: a SWE methodology for testing a specific set of relations between the outputs of multiple inputs. True to the checklists in the Google ML Testing Rubric and in MLTRL, metamorphic testing for ML can have a codified list of metamorphic relations [13].

In domains such as healthcare there have been the introduction of similar checklists for data readiness – for example, to ensure regulatory-grade real-world-evidence (RWE) data quality [48] – yet these are nascent and not yet widely accepted. Applying AI in healthcare has led to developing guidance for regulatory protocol, which is still a work in progress. Larson et al. [49] provide a comprehensive analysis for medical imaging and AI, arriving at several regulatory framework recommendations that mirror what we outline as important measures in MLTRL: e.g., detailed task elements such as pitfalls and limitations (surfaced on TRL Cards), clear definition of an algorithm relative to the downstream task, defining the algorithm "capability" (Level 5), real-world monitoring, and more.

D'amour et al. [12] dive into the problem we noted earlier about model mis-calibration. They point to the trend in machine learning to develop models relatively isolated from the downstream use and larger system, resulting in underspecification that handicaps practical ML pipelines. This is largely problematic in deep learning pipelines, but we've also noted this risk in the case of causal inference applications. Suggested remedies include *stress tests*–empirical

evaluations that probe the model's inductive biases on practically relevant dimensions–and in general the methods we define in Level 7.

CONCLUSION

We've described *Machine Learning Technology Readiness Levels (MLTRL)*, an industry-hardened systems engineering framework for robust, reliable, and responsible machine learning. MLTRL is derived from the processes and testing standards of spacecraft development, yet lean and efficient for ML and software workflows. Examples from several organizations across industries demonstrate the efficacy of MLTRL for AI and ML technologies, from research and development through productization and deployment, in important domains such as healthcare and physics. Even more, MLTRL establishes a much-needed lingua franca for the AI ecosystem. As a next step we aim to provide open-source tools (namely templates for TRL Cards and ethics checklists) for more teams and organizations to adopt MLTRL. Our hope is that our systems framework is adopted broadly in AI and ML organizations, and that "technology readiness levels" becomes common nomenclature across AI stakeholders – from researchers and engineers to sales-people and executive decision-makers.

Methods

Gated reviews

At the end of each stage is a dedicated review period: (1) Present the technical developments along with the requirements and their corresponding verification measures and validation steps, (2) make key decisions on path(s) forward (or backward) and timing, and (3) debrief the process[viii]. As in the gated reviews defined by TRL used by NASA, DARPA, et al., MLTRL stipulates specific criteria for review at each level, as well as calling out specific *key decision points* (noted in the level descriptions above). The designated reviewers will "graduate" the technology to the next level, or provide a list of specific tasks that are still needed (ideally with quantitative remarks). After graduation at each level, the working group does a brief post-mortem; we find that a quick day or two pays dividends in cutting away technical debt and improving team processes. Regular gated reviews are essential for making efficient progress while ensuring robustness and functionality that meets stakeholder needs. There are several important mechanisms in MLTRL reviews that are specifically useful with AI and ML technologies: First, the review panels evolve over a project lifecycle, as noted below. Second, MLTRL prescribes that each review runs through an AI ethics checklist defined by the organization; it is important to repeat this at each review, as the review panel and stakeholders evolve considerably over a project lifecycle. As previously described in the levels definitions, including ethics reviews as an integral part of early system development is essential for informing model specifications and avoiding unintended biases or harm [50] after deployment.

TRL "Cards"

In Figure 3 we succinctly showcase a key deliverable: *TRL Cards*. The model cards proposed by Google [51] are a useful development for external user-readiness with ML. On the other hand, our TRL Cards are more information-dense, like datasheets for medical devices and engineering tools. These serve as "report cards" that grow and improve upon graduating levels, and provide a means of inter-team and cross-functional communication. The content of a TRL Card is roughly in two categories: project info, and implicit knowledge. The former clearly states info such as project owners and reviewers, development status, and semantic versioning–not just for code, also for models and data. In the latter category are specific insights that are typically siloed in the ML development team but should be communicated to other stakeholders: modeling assumptions, dataset biases, corner cases, etc. With the spread of AI and ML in critical application areas, we are seeing domain expert consortiums defining AI reporting guidelines – e.g., Rivera et al. [52] calling for clinical trials reports for interventions involving AI – which will greatly benefit from the use of our TRL reporting cards. We stress that these TRL Cards are key for the progression of projects, rather than documentation afterthoughts. The TRL Cards thus promote transparency and trust, within teams and across organizations.

Risk mitigation

Identifying and addressing risks in a software project is not a new practice. However, akin to the MLTRL roots in spacecraft engineering, risk is a "first-class citizen" here. In the definition of technical and product requirements, each entry has a calculation of the form $risk = p(failure) \times value$, where the value of a component is an integer $1 - 10$.

[viii]MLTRL should include regular debriefs and meta-evaluations such that process improvements can be made in a data-driven, efficient way (rather than an annual meta-review). MLTRL is a high-level framework that each organization should operationalize in a way that suits their specific capabilities and resources.

Being diligent about quantifying risks across the technical requirements is a useful mechanism for flagging ML-related vulnerabilities that can sometimes be hidden by layers of other software. MLTRL also specifies that risk quantification and testing strategies are required for sim-to-real development. That is, there is nearly always a non-trivial gap in transferring a model or algorithm from a simulation testbed to the real world. Requiring explicit sim-to-real testing steps in the workflow helps mitigate unforeseen (and often hazardous) failures. Additionally, comprehensive ML test coverage that we mention throughout this paper is a critical strategy for mitigating risks *and* uncertainties: ML-based system behavior is not easily specified in advance, but rather depends on dynamic qualities of the data and on various model configuration choices [11].

Non-monotonic, non-linear paths

We observe many projects benefit from cyclic paths, dialing components of a technology back to a lower level. Our framework not only encourages cycles, we make them explicit with "switchback mechanisms" to regress the maturity of specific components in an AI system:

1. *Discovery switchbacks* occur as a natural mechanism – new technical gaps are discovered through systems integration, sparking later rounds of component development [53]. These are most common in the R&D levels, for example moving a component of a proof-of-concept technology (at Level 4) back to proof-of-principle development (Level 2).
2. *Review switchbacks* result from gated reviews, where specific components or larger subsystems may be dialed back to earlier levels. This switchback is one of the "key decision points" in the MLTRL project lifecycle (as noted in the Levels definitions), and is often a decision driven by business-needs and timing rather than technical concerns (for instance when mission priorities and funds shift). This mechanism is common from Level 6/7 to 4, which stresses the importance of this R&D to product transition phase (see Figure 2 (left)).
3. *Embedded switchbacks* are predefined in the MLTRL process. For example, a predefined path from 4 to 2, and from 9 to 4. In complex systems, particularly with AI technologies, these built-in loops help mitigate technical debt and overcome other inefficiencies such as noncomprehensive V&V steps.

Without these built-in mechanisms for cyclic development paths, it can be difficult and inefficient to build systems of modules and components at varying degrees of maturity. Contrary to traditional thought that switchback events should be suppressed and minimized, in fact they represent a natural and necessary part of the complex technology development process – efforts to eliminate them may stifle important innovations without necessarily improving efficiency. This is a fault of the standard monotonic approaches in AI/ML projects, stage-gate processes, and even the traditional TRL framework.

It is also important to note that most projects do not start at Level 0; very few ML companies engage in this low-level theoretical research. For example, a team looking to use an off-the-shelf object recognition model could start that technology at Level 3, and proceed with thorough V&V for their specific datasets and use-cases. However, no technology can skip levels after the MLTRL process has been initiated. The industry default (that is, without implementing MLTRL) is to ignorantly take pretrained models, run fine tuning on their specific data, and jump to deployment, effectively skipping Levels 5 to 7. Additionally, we find it is advantageous to incorporate components from other high-TRL ranking projects while starting new projects; MLTRL makes the verification and validation (V&V) steps straightforward for integrating previously developed ML components.

Evolving people, objectives, and measures

As suggested earlier, much of the practical value of MLTRL comes at the transition between levels. More precisely, MLTRL manages these oft neglected transitions explicitly as evolving teams, objectives, and deliverables. For instance, the team (or working group) at Level 3 is mostly AI Research Engineers, but at Level 6 is mixed Applied AI/SW Engineers mixed with product managers and designers. Similarly, the review panels evolve from level to level, to match the changing technology development objectives. What the reviewers reference similarly evolves: notice in the level definitions that technical requirements and V&V guide early stages, but at and after Level 6 the product requirements and V&V takeover – naturally, the risk quantification and mitigation strategies evolve in parallel. Regarding the deliverables, notably TRL Cards and risk matrices [15] (to rank and prioritize various science, technical, and project risks), the information develops and evolves over time as the technology matures.

Quantifiable progress

By defining technology maturity in a quantitative way, MLTRL enables teams to accurately and consistently define their ML progress metrics. Notably industry-standard OKRs and KPIs [54] can be defined as achieving certain readiness

levels in a given period of time; this is a preferable metric in essentially all ML systems which consist of much more than a single performance score to measure progress. Even more, meta-review of MLTRL progress over multiple projects can provide useful insights at the organization level. For example, analysis of the time-per-level and the most frequent development paths/cycles can bring to light operational bottlenecks. Compared to conventional software engineering metrics based on sprint stories and tickets, or time-tracking tools, MLTRL provides a more accurate analysis of ML workflows.

Communication and explanation

A distinct advantage of MLTRL in practice is the nomenclature: an agreed upon grading scheme for the maturity of an AI technology, and a framework for how/when that technology fits within a product or system, enables everyone to communicate effectively and transparently. MLTRL also acts as a gate for interpretability and explainability–at the granularity of individual models and algorithms, and more crucially from a holistic, systems standpoint. Notably the DARPA XAI[ix] program advocates for this advance in developing AI technologies; they suggest interpretability and explainability are necessary at various locations in an AI system to be sufficient for deployment as an AI product, otherwise leading to issues with ethics and bias.

Robustness via uncertainty-aware ML

How to design a reliable system from unreliable components has been a guiding question in the fields of computing and intelligence [55]. In the case of AI/ML systems, we aim to build reliable systems with myriad unreliable components: noisy and faulty sensors, human and AI error, and so on. There is thus significant value to quantifying the myriad uncertainties, propagating them throughout a system, and arriving at a notion or measure of reliability. For this reason, although MLTRL applies generally to AI/ML methods and systems, we advocate for methods in the class of probabilistic ML, which naturally represent and manipulate uncertainty about models and predictions [19]. These are Bayesian methods that use probabilities to represent *aleatoric uncertainty*, measuring the noise inherent in the observations, and *epistemic uncertainty*, accounting for uncertainty in the model itself (i.e., capturing our ignorance about which model generated the data). In the simplest case, an uncertainty aware ML pipeline should quantify uncertainty at the points of sensor inputs or perception, prediction or model output, and decision or end-user action – McAllister et al. [20] suggest this with Bayesian deep learning models for safer autonomous vehicle pipelines. We can achieve this sufficiently well in practice for simple systems. However, we do not yet have a principled, theoretically grounded, and generalizable way of propagating errors and uncertainties downstream and throughout more complex AI systems – i.e., how to integrate different software, hardware, data, and human components while considering how errors and uncertainties propagate through the system. This is an important direction of our future work.

References

[1] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *AAAI*, 2018.

[2] Arnaud de la Tour, Massimo Portincaso, Kyle Blank, and Nicolas Goeldel. The dawn of the deep tech ecosystem. Technical report, The Boston Consulting Group, 2019.

[3] NASA. Nasa systems engineering handbook. 2003.

[4] United States Department of Defense. Defense acquisition guidebook. Technical report, U.S. Dept. of Defense, 2004.

[5] D. Leslie. Understanding artificial intelligence ethics and safety. *ArXiv*, abs/1906.05684, 2019.

[6] Google. Machine learning workflow. https://cloud.google.com/mlengine/docs/tensorflow/ml-solutions-overview. Accessed: 2020-12-13.

[7] B. Shahriari, Kevin Swersky, Ziyu Wang, R. Adams, and N. D. Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104:148–175, 2016.

[8] Umang Bhatt, Alice Xiang, S. Sharma, Adrian Weller, Ankur Taly, Yunhan Jia, Joydeep Ghosh, Ruchir Puri, José M. F. Moura, and P. Eckersley. Explainable machine learning in deployment. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020.

[9] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37:50–60, 2020.

[ix]DARPA Explainable Artificial Intelligence (XAI)

[10] T. Ryffel, Andrew Trask, M. Dahl, Bobby Wagner, J. Mancuso, D. Rueckert, and J. Passerat-Palmbach. A generic framework for privacy preserving deep learning. *ArXiv*, abs/1811.04017, 2018.

[11] Eric Breck, Shanqing 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.

[12] 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.

[13] Xiaoyuan Xie, Joshua W. K. Ho, C. Murphy, G. Kaiser, B. Xu, and T. Chen. Testing and validating machine learning classifiers by metamorphic testing. *The Journal of systems and software*, 84 4:544–558, 2011.

[14] A. Botchkarev. A new typology design of performance metrics to measure errors in machine learning regression algorithms. *Interdisciplinary Journal of Information, Knowledge, and Management*, 14:045–076, 2019.

[15] N. Duijm. Recommendations on the use and design of risk matrices. *Safety Science*, 76:21–31, 2015.

[16] Louise Naud and Alexander Lavin. Manifolds for unsupervised visual anomaly detection. *ArXiv*, abs/2006.11364, 2020.

[17] Timnit Gebru, J. Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, H. Wallach, Hal Daumé, and K. Crawford. Datasheets for datasets. *ArXiv*, abs/1803.09010, 2018.

[18] P. Schulam and S. Saria. Reliable decision support using counterfactual models. In *NIPS 2017*, 2017.

[19] Zoubin Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521:452–459, 2015.

[20] Rowan McAllister, Yarin Gal, Alex Kendall, Mark van der Wilk, A. Shah, R. Cipolla, and Adrian Weller. Concrete problems for autonomous vehicle safety: Advantages of bayesian deep learning. In *IJCAI*, 2017.

[21] J. Tobin, Rachel H Fong, Alex Ray, J. Schneider, W. Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017.

[22] Arthur Juliani, Vincent-Pierre Berges, Esh Vckay, Yuan Gao, Hunter Henry, M. Mattar, and D. Lange. Unity: A general platform for intelligent agents. *ArXiv*, abs/1809.02627, 2018.

[23] Stefan Hinterstoißer, Olivier Pauly, Tim Hauke Heibel, Martina Marek, and Martin Bokeloh. An annotation saved is an annotation earned: Using fully synthetic training for object instance detection. *ArXiv*, abs/1902.09967, 2019.

[24] K. Cranmer, J. Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117:30055 – 30062, 2020.

[25] Jan-Willem van de Meent, Brooks Paige, H. Yang, and Frank Wood. An introduction to probabilistic programming. *ArXiv*, abs/1809.10756, 2018.

[26] Atilim Günes Baydin, Lei Shao, W. Bhimji, L. Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, Mingfei Ma, X. Zhao, P. Torr, V. Lee, K. Cranmer, Prabhat, and F. Wood. Etalumis: bringing probabilistic programming to scientific simulators at scale. *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019.

[27] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter. Event generation with sherpa 1.1. *Journal of High Energy Physics*, 2009:007–007, 2009.

[28] 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.

[29] R. Ambrosino, B. Buchanan, G. Cooper, and Marvin J. Fine. The use of misclassification costs to learn rule-based decision support models for cost-effective hospital admission strategies. *Proceedings. Symposium on Computer Applications in Medical Care*, pages 304–8, 1995.

[30] J. Pearl. Theoretical impediments to machine learning with seven sparks from the causal revolution. *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 2018.

[31] Jonathan G Richens, C. M. Lee, and Saurabh Johri. Improving the accuracy of medical diagnosis with causal machine learning. *Nature Communications*, 11, 2020.

[32] Andrei Paleyes, Raoul-Gabriel Urma, and N. Lawrence. Challenges in deploying machine learning: a survey of case studies. *ArXiv*, abs/2011.09926, 2020.

[33] V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, Christian L. Hansen, Whitney K. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *Econometrics: Econometric & Statistical Methods - Special Topics eJournal*, 2018.

[34] T. Nguyen, G. Collins, J. Spence, J. Daurès, P. Devereaux, P. Landais, and Y. Le Manach. Double-adjustment in propensity score matching analysis: choosing a threshold for considering residual imbalance. *BMC Medical Research Methodology*, 17, 2017.

[35] Victor Veitch and Anisha Zaveri. Sense and sensitivity analysis: Simple post-hoc analysis of bias due to unobserved confounding. *NeurIPS 2020, arXiv preprint arXiv:2003.01747*, 2020.

[36] D. Eckles and E. Bakshy. Bias and high-dimensional adjustment in observational studies of peer effects. *ArXiv*, abs/1706.04692, 2017.

[37] Yanbo Xu, Divyat Mahajan, Liz Manrao, A. Sharma, and E. Kiciman. Split-treatment analysis to rank heterogeneous causal effects for prospective interventions. *ArXiv*, abs/2011.05877, 2020.

[38] P. Jenniskens, P.S. Gural, L. Dynneson, B.J. Grigsby, K.E. Newman, M. Borden, M. Koop, and D. Holman. Cams: Cameras for allsky meteor surveillance to establish minor meteor showers. *Icarus*, 216(1):40 – 61, 2011.

[39] S. Zoghbi, M. Cicco, A. P. Stapper, A. J. Ordonez, J. Collison, P. S. Gural, S. Ganju, J.-L. Galache, and P. Jenniskens. Searching for long-period comets with deep learning tools. In *Deep Learning for Physical Science Workshop, NeurIPS*, 2017.

[40] Peter Jenniskens, Jack Baggaley, Ian Crumpton, Peter Aldous, Petr Pokorny, Diego Janches, Peter S. Gural, Dave Samuels, Jim Albers, Andreas Howell, Carl Johannink, Martin Breukers, Mohammad Odeh, Nicholas Moskovitz, Jack Collison, and Siddha Ganju. A survey of southern hemisphere meteor showers. *Planetary and Space Science*, 154:21 – 29, 2018.

[41] 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.

[42] D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, and Dan Dennison. Hidden technical debt in machine learning systems. In *NIPS*, 2015.

[43] Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Burkner, and Martin Modrak. Bayesian workflow. *ArXiv*, abs/2011.01808, 2020.

[44] Eric Breck, Marty Zinkevich, Neoklis Polyzotis, Steven Euijong Whang, and Sudip Roy. Data validation for machine learning. 2019.

[45] Marco Túlio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. Beyond accuracy: Behavioral testing of nlp models with checklist. In *ACL*, 2020.

[46] R. Kumar, David R. O'Brien, Kendra Albert, Salomé Viljöen, and Jeffrey Snover. Failure modes in machine learning systems. *ArXiv*, abs/1911.11034, 2019.

[47] Inioluwa Deborah Raji, Andrew Smart, Rebecca White, M. Mitchell, Timnit Gebru, B. Hutchinson, Jamila Smith-Loud, Daniel Theron, and P. Barnes. Closing the ai accountability gap: defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020.

[48] R. Miksad and A. Abernethy. Harnessing the power of real-world evidence (rwe): A checklist to ensure regulatory-grade data quality. *Clinical Pharmacology and Therapeutics*, 103:202 – 205, 2018.

[49] D. B. Larson, Hugh Harvey, D. Rubin, Neville Irani, J. R. Tse, and C. Langlotz. Regulatory frameworks for development and evaluation of artificial intelligence–based diagnostic imaging algorithms: Summary and recommendations. *Journal of the American College of Radiology*, 2020.

[50] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366:447 – 453, 2019.

[51] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019.

[52] Samantha Cruz Rivera, Xiaoxuan Liu, A. Chan, A. K. Denniston, and M. Calvert. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the spirit-ai extension. *Nature Medicine*, 26:1351 – 1363, 2020.

[53] Z. Szajnfarber. Managing innovation in architecturally hierarchical systems: Three switchback mechanisms that impact practice. *IEEE Transactions on Engineering Management*, 61:633–645, 2014.

[54] H. Zhou and Y. He. Comparative study of okr and kpi. *DEStech Transactions on Economics, Business and Management*, 2018.

[55] J. Neumann. Probabilistic logic and the synthesis of reliable organisms from unreliable components. 1956.

[56] Alexander Lavin and Gregory Renard. Technology readiness levels for ai & ml. *ICML Workshop on Challenges Deploying ML Systems*, 2020.

Acknowledgements

Author contributions statement

A.L. conceived of the original ideas and framework, with significant contributions towards improving the framework from all co-authors. A.L. initiated the use of MLTRL in practice, including the neuropathology test case discussed here. C.G-L. contributed insight regarding causal AI, including the section on counterfactual diagnosis. C.G-L. also made significant contributions broadly in the paper, notably in the Methods descriptions. Si.G. contributed the spacecraft test case, along with early insights in the framework definitions. A.V. contributed to the definition of later stages involving deployment, as did A.G. Both E.X. and Y.G. provided insights regarding AI in academia, and Y.G. additionally contributed to the uncertainty quantification methods. Su.G. and D.L. contributed the computer vision test case. A.G.B. contributed the particle physics test case, and significant reviews of the writeup. A.S. contributed insights related to causal ML and AI ethics. D.N. provided valuable feedback on the overall framework, and contributed significantly with the details on "switchback mechanisms". J.P. contributed to multiple paper revisions, and to deploying the systems ML methods broadly in practice for Earth and space sciences. C.M. additionally provided feedback on the overall methods. All co-authors discussed the content and contributed to editing the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to A.L.