alteryx

17200 Laguna Canyon Road
Irvine, CA 92618
888.836.4274

March 13, 2024

The AI Verify Foundation
The InfoComm Media Development Authority

Re: Comments on the Proposed Model Governance Framework for Generative AI

Introduction

Alteryx welcomes the opportunity to provide input to the AI Verify Foundation and the InfoComm Media Development Authority (IMDA). Alteryx supports governments in leading collaboration with industry to establish guidance, standards, and policies to ensure the responsible development, deployment, and use of AI, and wishes to praise the AI Verify Foundation and IMDA for undertaking the development of this Model Governance Framework as an important effort in that arena.

Alteryx is a leading provider of data science and analytics automation software with a sizable base of customers and users in Singapore. The Alteryx AI Platform for Enterprise Analytics delivers easy end-to-end automation of data engineering, analytics, reporting, machine learning, and data science processes, enabling enterprises everywhere to democratize data analytics across their organizations for a broad range of use cases.

Alteryx develops and deploys sophisticated low-code/no-code analytics software, including both on-premise and cloud-based offerings, that make advanced analytics and artificial intelligence/machine learning (AI/ML) generated insights available to workers in diverse fields of business and government. Alteryx is also substantially invested in the responsible governance of AI, with all of our AI work governed by our Responsible AI Principles (https://www.alteryx.com/trust/ai-principles). Thus, as a developer, integrator, user, and deployer of responsible AI systems, Alteryx has a significant interest in the Model Governance Framework for Generative AI.

In general, Alteryx believes that the draft Framework provides a well-considered basis for approaching governance of generative AI systems. The Framework aligns quite closely with our Responsible AI Principles of Transparency and Explainability, Human Agency and Oversight, Trust and Accountability, Reliability and Safety, Fairness and Inclusion, and Empowering Social Good. We look forward to working with the AI Verify Foundation and IMDA to further strengthen the Framework and, to that end, we offer several specific comments below.

Shared Accountability

As the draft Framework rightly recognizes, “generative AI, like most software development, involves multiple layers in the tech stack.” It is best conceived of as a

17200 Laguna Canyon Road
Irvine, CA 92618
888.836.4274

value chain, with multiple organizations potentially contributing to the development and refinement of the AI system that ultimately interacts with an end user. Alteryx occupies a role in the center of this value chain, training, modifying, and integrating foundation models developed by other organizations into applications that may be used directly by end users, or that may be further refined by other enterprise customers.

Alteryx applauds the AI Verify Foundation and IMDA for recognizing the spectrum of activity along the AI value chain, and for avoiding a simplistic bifurcation of AI developers and deployers. We strongly agree that "Players along the AI development chain need to be responsible towards end-users, and the structural incentives should align with this need." We recommend that the Framework further elaborate on this concept of the AI development chain to consider a greater variety of roles. To ensure comprehensive responsibility, we suggest the Framework explicitly define a broader set of roles within the AI development chain (e.g., model developers, integrators, application developers, deployers) and delineate their specific responsibilities towards end-users. Greater specificity of roles and responsibilities will help in aligning structural incentives across the board.

The draft Framework notes that actors in the AI value chain "include model developers, application deployers, and cloud service providers (who often provide platforms on which AI applications are hosted)." In a footnote, it recognizes that "the generative AI development chain is complex, and that application developers and application deployers can sometimes be two different parties," but then erases that distinction by deciding to use "the term 'application deployers' to refer to both application developers and deployers."

Within the AI value chain, developers of foundation models clearly take actions that fundamentally shape any AI system based on such a model, and shape the capabilities and risks associated with the model. Often, however, AI application developers or actors, who otherwise integrate foundation models into broader systems, take measures – including training the model with new data sets, adjusting model weights, prompt engineering, incorporating privacy and security controls, and creating user interfaces – that also impact the capabilities and risks associated with the model. Before an AI system is deployed to interact with an end user, there may be two, three, or more organizations that modify the original model. The Governance Framework should explicitly recognize this complexity and establish best practices that take it into account.

At the forefront of responsible AI practices lies the principle of transparency, and the provision of accurate technical information at each step along the value chain. AI integrators and application developers depend upon such technical documentation to make the best decisions to mitigate risk as they further develop an AI system. Alteryx applauds the AI Verify Foundation and IMDA for urging the standardization of transparency around AI safety measures, but the Framework should move beyond its focus on AI model developers to address the importance of transparency throughout the value chain. If organizations are held accountable for providing transparent, accurate technical documentation along each step of the value chain, the result will be a far

17200 Laguna Canyon Road
Irvine, CA 92618
888.836.4274

healthier AI ecosystem in which risks are comprehensively identified and may be effectively mitigated.

Accessing Quality Data

Alteryx strongly agrees with the Draft Framework's acknowledgement that "Data is a core element of model development. It significantly impacts the quality of the model output. Hence, what is fed to the model is important and there is a need to ensure data quality, such as through the use of trusted data sources."

We support the Framework's recommendations on addressing concerns about copyrighted data and personal or sensitive information, as well as on the importance of data governance. Data quality must extend beyond trusting the source of data, however, and include the integrity of the data itself. Alteryx often talks about the "4 Cs" of data quality: the consistency, conformity, completeness and currency of the data. Consistency means data is statistically valid and internally coherent, and considers whether there are extreme values, outliers or anomalies. Conformity refers to acceptable standards and patterns to which data must adhere. Completeness indicates that all necessary data has been included and there are no missing values. And currency requires validating that the data is up-to-date and has been refreshed regularly. These principles do not apply to every scenario, but represent a good starting point for interrogating data quality.

Alteryx believes that investment in ensuring data quality is core to AI governance. Automating data preparation tasks can not only make these investments far more efficient, but can help identify data quality problems early in the process.

In addition, the Framework notes the need for "concerted upskilling of the workforce," and Alteryx would argue this upskilling is particularly important in relation to effectively leveraging data for generative AI. We see a critical need to ensure public and private sector workforces have sufficient capabilities to work with data in the ways necessary to support a healthy AI ecosystem, including ensuring well-prepared, quality data is available for model development and training, establishing strong data governance mechanisms, and identifying data issues, such as copyright or privacy challenges. A safe, secure, and innovative AI ecosystem will depend on a data literate workforce.

Crucially, efforts to prepare a diverse workforce for responsible adoption of AI technologies must approach upskilling broadly, rather than simply seeking to recruit more highly skilled data scientists. Data scientists and other specialists can help solve the most challenging technical problems, but basic data and AI literacy will be needed across the workforce to harness AI's potential. Organizations seeking to develop and adopt generative AI must invest in building a workforce capable of working responsibly with data and AI as a central element of their approach to generative AI governance. In view of our commitment to responsible AI, we recognize our own responsibility to train and maintain a data literacy workforce, and to build products on the basis of quality data, given that there will be a range of data literacy levels among our users.

17200 Laguna Canyon Road
Irvine, CA 92618
888.836.4274

Alteryx is keen to be part of the solution to these workforce challenges around the world and, in fact, we are already working in Singapore to do so. Our "SparkED" Analytics Education Program[1] provides foundational training in the essentials of working with data to both university students and individual mid-career learners at no cost. Independent learners of all backgrounds have used the training gained through SparkED to build new careers and take new steps in their existing ones; students and educators are engaging through the SparkED program at over 800 universities and other academic institutions in more than 50 countries. We are already partnering with institutions of higher education such as Temasek Polytechnic, Ngee Ann Polytechnic, Singapore Institute of Technology, and Nanyang Technological University, and we are eager to partner with the Government of Singapore to maximize access to upskilling opportunities like Sparked across Singapore's current and future workforce.

Security and Incident Reporting

The draft Framework rightly prioritizes trust and security in generative AI systems, as well as the need for incident reporting. Strong privacy and security protections are essential to building the trust of users, and the public in general, for AI systems, and should be built into the foundations of the AI development process. Likewise, incident reporting can help ensure that common challenges are identified and addressed collectively.

The draft draws parallels between established practices for cybersecurity and cyber incident reporting and nascent thinking about security and incident reporting for AI. Particularly with regard to security, it takes a cautious stance, suggesting the need for additional insight around what may be needed to adapt current approaches. We appreciate that caution.

Existing governance approaches to cyber security and security incident reporting provide a strong foundation for building trust in AI systems, and we believe governments and industry should work together to apply those existing standards, laws, and best practices to AI systems. We should avoid the temptation to build new frameworks for AI simply because the technology is new.

For security incident reporting, more thinking is needed to illuminate whether there are classes of reportable incidents in AI systems that would not be covered under existing cyber security incident reporting guidelines. For example, Singapore's *Cyber Security Act* defines a "cybersecurity incident" as "an act or activity carried out without lawful authority on or through a computer or computer system that jeopardizes or adversely affects its cybersecurity or the cybersecurity of another computer or computer system." The U.S. National Institute of Standards and Technology has an even broader definition: "An occurrence that actually or potentially jeopardizes the confidentiality, integrity, or

[1] More information about the SparkED program is available at: https://www.alteryx.com/sparked.

17200 Laguna Canyon Road
Irvine, CA 92618
888.836.4274

availability of an information system or the information the system processes, stores, or transmits or that constitutes a violation or imminent threat of violation of security policies, security procedures, or acceptable use policies."[2]

Similarly, existing frameworks for cybersecurity governance may not contemplate every specific risk associated with AI but may nonetheless prove adequate for addressing those risks. Software development best practices, such as threat modeling and attack surface mapping, will help developers of AI systems identify potential risks, including AI-unique attack vectors, and mitigate them during development.

As the draft Framework recognizes, some "refinements may be needed" to existing cybersecurity practices "given the unique characteristics of generative AI," but we should prioritize application of existing tools and practices in the near-term, while working to identify any potential gaps in existing approaches. In the immediate future, building on the existing corpus of policies, regulations, standards, and best practices addressing privacy and cybersecurity, including incident reporting, will generate earlier, and better, returns rather than seeking to build wholly new approaches.

Conclusion

AI technologies offer tremendous potential benefits for our economy, security, government, and society. At Alteryx, we deliver innovations that enable organizations to combine sophisticated AI systems with powerful and trusted data connection, data cleansing, and data analytics products that improve user efficiency and productivity, improve analytic quality, and drive innovation. Yet, we are also cognizant of the risks of AI developed or deployed irresponsibly. For that reason, we look forward to working with the government in support of a practical, risk-based, and agile governance framework that drives accountability and safety across the AI ecosystem. We look forward to continuing this important dialogue.

Sincerely,

Tommy Ross
Head of Global Public Policy | ALTERYX

[2] National Institute of Standards and Technology Computer Security Resource Center Glossary, https://csrc.nist.gov/glossary/term/incident#:~:text=An%20occurrence%20that%20actually%20or,security%20procedures%2C%20or%20acceptable%20use.