SparkAI Consortium@SDSC

Summer 2024 Project Brief: "GovernAI: An Assessment of Evolving Industry Policies and Practices"

INTRODUCTION

SPARKAI is the new, industry focused AI research consortium at the San Diego Supercomputer Center, UC San Diego. This brief presents SPARK's summer 2024 project, "GovernAI", and concludes with an invitation to those interested to join and contribute to the project.

CONTEXT

Future progress and commercialization of AI promises to be rapid. The opportunities are huge, as are the uncertainties and business risks, posing critical governance challenges. A review of over 80 reports, policy statements, and corporate disclosures on AI governance point to several common themes. First, the overwhelming majority list typical challenges that any new, disruptive technology poses. Not unique to AI, but perhaps AI will compound them. Second, the majority describe principles, policies and procedures inherited from previous governance frameworks – in the words of one expert, will AI just require "an amped up current governance process?" Third, "explainable," or "responsible," or "trusted" AI will require a considerably larger group of stakeholders. How large, who they are, and how they will work together is mentioned but not detailed. Fourth, where will the benefits of "best in class" governance be seen and measured in the numbers? Risk reduction, financial performance, avoidance of litigation, reputational benefits? Finally, AI governance will not be free. The required investments in people, time, money and organizational processes are necessary, but what will they cost? To the extent the above challenges are context specific, then greater attention to defining, explaining and measuring that context is warranted.

GOALS

The project goal is to develop a systematic methodology to address the above knowledge gaps, and to discern what is unique about AI in comparison with other disruptive technologies.[1] An overarching goal is to deconstruct AI governance into discernible segments addressing first, the technical inputs, possibilities and constraints of AI (a moving target); second, to address the role of key stakeholders (for example, the corporate roles involved, including new ones); and third, to investigate use case examples of early, middle and late adopters, and the reasons involved. Context-specific characteristics of adoption strategy, by industry, by AI application, will be incorporated.

METHODOLOGY

Initial data collection will be interview based, along with secondary (text) analysis of firm and government disclosures (mandated and voluntary). A typical company or government interview will be supplemented by an analysis of disclosures. For example, an executive interview in company X, will be supplemented by SEC filings, company disclosures, and analyst comments. All interviews will be confidential and anonymized. Any quotes or reference to an interview or company will require written permission. Briefings will be circulated to all sponsors for review for an agreed-upon period of time, typically 3-6 months. Following approval, a summary document will be produced and distributed. The project will proceed in stages that mirror a stylized management decision. For example, a decision scenario involving the potential benefits and risks of long-term storage of AI data and analytics is shown in Figure 1.

TIMELINE AND OUTPUTS

The overall timeline for this study will be broken into quarterly segments, totaling 9-12 months, contingent on funding. Each quarterly segment will produce a full report-in on preliminary insights from interview and secondary analysis. Necessarily limited in the first quarter (July-September), but expanding and with greater detail as the project unfolds. Monthly project updates (zoom calls) will be scheduled. Milestone dates for project report-ins and presentations will align with the SPARK 2024 calendar. Looking ahead, we will plan a detailed briefing for all SPARK consortium members at our Nov 13 Board Meeting (in person in San Diego). All project participants will be invited to zoom project briefings starting in August, and the summary report in the Fall.

---

[1] OMB's final Cloud Smart guidelines (2019) were based on three strategy pillars, including modernizing security practices and automation, improvements in procurement, and modernizing workforce training and recruitment for the new skills required for cloud adoption. ed in a previous, data technology disruption, cloud services adoption and "Cloud Smart" transformation, new procedures, management roles, data governance practices, measurements and review, including audit functions, were needed.

FOR THOSE INTERESTED IN KNOWING MORE

We are grateful to Solix Technologies for providing seed funding to enable the startup of this summer project. For those individuals and/or companies interested in knowing more, and potentially contributing to the project in some capacity (interviews, use cases, underwriting a research workshop or event), please contact your local Solix representative or the principal investigator.

All individuals, companies, and public organizations contributing to the project will be invited to join interim briefings, project reports, and receive the final project report.

QUESTIONS, COMMENTS AND SUGGESTIONS

Please direct to James (Jim) Short, Principal Investigator, at jshort@ucsd.edu

FIGURE 1: A stylized management decision scenario involving long-term storage of AI artifacts*

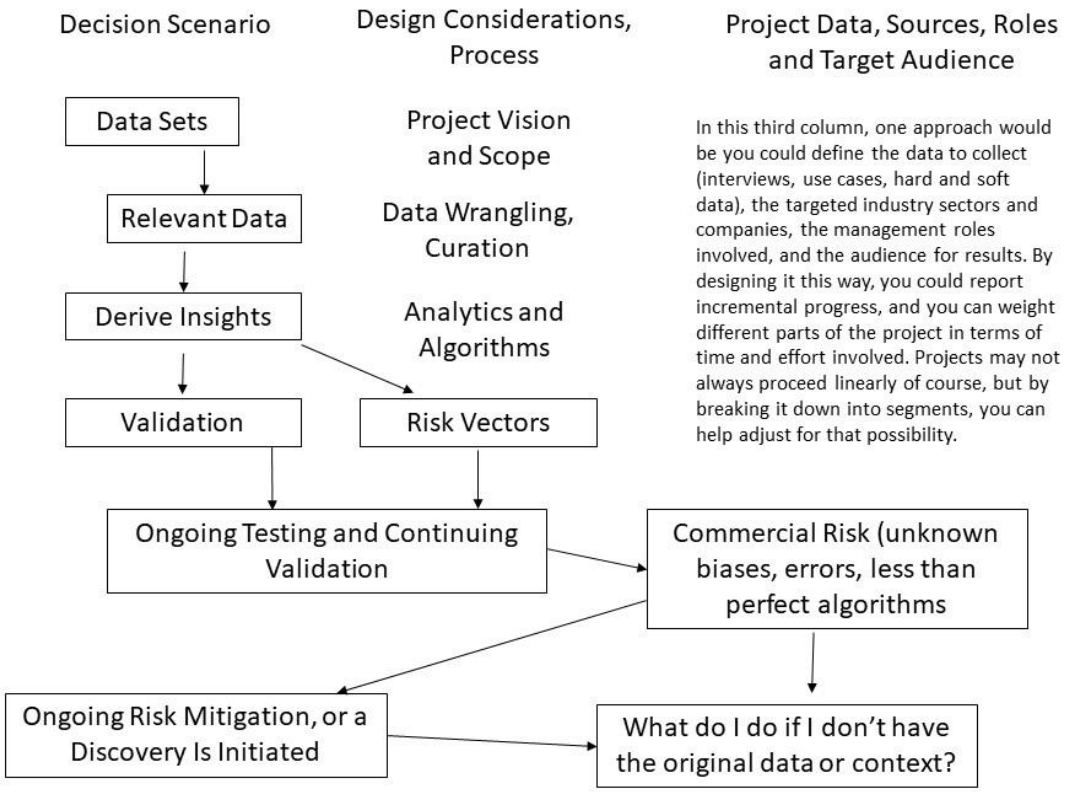

*with thanks to Barry Rudolph, SPARKAI Board Member, for helping to co-develop this decision scenario