Designing a Partnership Framework in AI for Social Good

Caroline Trier, Lu Sevier
Booz Allen Hamilton
Trier_Caroline@bah.com, Sevier_Lu@bah.com

Abstract

While artificial intelligence (AI) has been heralded as a technology capable of solving unique problems, social good challenges are inherently structural and require the partnership of many stakeholders in order to apply AI for social good (AI4SG) in a sustainable and scaled manner. This paper explains current challenges in project implementation, surveys framework approaches, and contributes our differentiating lessons learned on scaling projects to problem domain-wide impact. The goal is to guide partnering organizations through challenges and identifying opportunities to accelerate the application of AI4SG.

1 Introduction

The rapid growth of AI in the last several years has revolutionized the way businesses and organizations function. Heralded as a technology capable of solving unique challenges in humanitarian crises, public health, and the environment, the value of this transformative technology has not yet been fully realized for mission-oriented organizations: non-profit, multi-lateral, and public organizations whose core function is social good. We define "Social good" projects as those aimed at increasing the efficacy and efficiency of social initiatives to improve access to opportunity, especially for communities of individuals for whom opportunities have historically been limited [Abebe & Goldner, 2018]. While many mission-oriented organizations recognize the opportunity, or hype, to help mission execution [Coulton *et al.*, 2015], they often lack enough familiarity with the technology to see the full potential within their own organization's data and if these benefits outweigh the risks and costs of AI adoption.

Social good challenges are inherently structural and AI solutions will only ever be a small part of addressing them [Moore, 2019]. Therefore, a fundamental need is the collaboration and coordination of government, non-profit, academia, and for-profit stakeholders in order to apply AI in a sustainable and scaled manner [Susha *et al.*, 2019]. In general, government and non-profit (mission-oriented) stakeholders specialize in domain expertise, while academia and for-profits provide the technical expertise. Both are needed to create effective and sustainable solutions.

There are various approaches that mission-oriented stakeholders and technical stakeholders have taken to form partnerships in AI4SG - from crowd sourcing through hackathons [DSB, 2020] to developing small volunteer-led initiatives [DataKind, 2020a], fellowships [DSSG, 2019], accelerators [DataKind, 2020b] or sponsoring grants and technology resources like Microsoft's AI for Good and Google's Impact Challenge [Google, 2018]. Unfortunately, many partnerships encounter barriers that prevent solution implementation, such as:

- Data access - conflicting or lacking legal provisions, regulatory compliance requirements, and privacy concerns
- Resource constraints – from difficult data collection and costly access to lacking data storage in a form useful for social good
- Partnership design - collaborators lacking experience in scoping projects that contain incentives for both parties while considering the long-term desired outcomes
- Project failure – due to the previous barriers and novelty of these partnerships, there are not many case studies to serve as inspiration [Susha *et al.*, 2019; Hager *et al.*, 2017].

This creates a need for a repeatable, sustainable approach to guide partnership development.

2 Related Work

When working on AI for Social Good (AI4SG) projects, it is essential to understand the challenges that prevent these solutions from being beneficial. They are often agnostic across data projects but are exacerbated by the lack of explanability and scale of AI [Tekisalp, 2020]. In order to mitigate these challenges, organizations experienced in AI4SG implementation have devised frameworks to guide partnership and project design.

2.1 Data Project Agnostic Challenges

Data access is one of the most listed challenges in implementing social good data projects [Susha *et al.*, 2019; Hager *et al.*, 2017; Niño *et al.*, 2017; Chui *et al.*, 2018; Tekisalp, 2020; DataKind, 2020]. Many governmental stakeholders exist in a field of varied regulatory compliance requirements on privacy and security. Non-profits may similarly be restricted or completely lack legal provisions [Susha *et al.*, 2019]. If the mission-oriented stakeholder needs a data partner, often they find themselves lacking reusable data infrastructure [Hager *et al.*, 2017] and face timeline delays negotiating conflicting legal provisions [Susha *et al*., 2019]. While many resources for good data sharing practices exist, it must be considered and built into any partnership approach.

Another challenge in implementing AI4SG projects is non-profit inexperience and limited capacity to evaluate, develop, procure and use AI-enabled solutions [Tekisalp, 2020]. This results in resource gaps in knowledge, data, and project costs. Use cases of AI4SG shared through conferences [ITU, 2019; Coulton *et al.,* 2015] have increased, but mission-oriented organizations still struggle in knowing how AI can be applied to their work [DataKind, 2020]. Even organizations with accessible data face challenges with data maturity, quality, and volume. Barriers also arise with the inherent costs of compute resources, cloud storage, and particularly AI expertise. Identifying funding systems, through either volunteerism, corporate social responsibility, or grants is essential [Susha *et al*., 2019]. For these projects to be sustainable, more resources need to be pooled to close the knowledge gap on AI.

Partnerships also struggle with communication hurdles and coordination shortfalls. Mission-oriented stakeholders and technical stakeholders often have different domain languages, making communication around problem definition and AI solution proposals difficult. This manifests as stakeholder demand for project management training from experienced organizations at the outsets of partnerships [DSSG, 2018a] and complaints of unusable solutions when projects end without a sustainment plan [Chui *et al.*, 2018]. AI exacerbates this through hype surrounding the term [Moore, 2019]. Finally, lack of coordination amongst stakeholders leads to one-off projects and redundancy of work.

2.2 AI Project Specific

There are many challenges unique to AI. Explainability, fragility or lack of generalizability, and potential to reinforce bias are a few that have hindered its implementation [Mannarswamy & Roy, 2019]. Researchers are working to design frameworks that address these risks [Floridi *et al.*, 2020]. AI practitioners often struggle with buy-in and resources to properly address these challenges [Cavello, 2020]. Application of AI to social good only increases the imperative, due to the heightened consequences on marginalized populations the solutions seek to help [Hager *et al.*, 2017]. Ultimately, projects fail without reaching implementation into the day to day use of mission-oriented organizations.

2.3 Frameworks for Partnering

We draw on the work of surveys that seek to find themes in one-off use case approaches and from AI4SG technical practitioners, such as Data Science for Social Good (DSSG) and DataKind, who draw on experience with hundreds of data projects. Our findings indicate consistent phases to AI4SG framework design [DSSG, 2018a; DataKind, 2020; Niño *et al.*, 2017; Susha *et al.*, 2019].

Partnership Discovery. Discovery starts with looking at a partner organization's mission, and what is preventing them from successfully fulfilling it. The organization can utilize data to improve execution but must consider the amount and type of data readily available. Moreover, the organization must be invested and see data as a critical mission resource worth maintaining. At this point, AI4SG practitioners attempt to mitigate data access challenges around regulations and legal agreements. Some offer example agreements [DSSG, 2018b] and security assurances [DSSG, 2018c]. Without access, potential for partnership is severely lowered but no resources have been lost in investing.

Problem Scoping. Stakeholders then scope a potential project through conversational or workshop methods [DataKind, 2020]. A major consideration is discussing data maturity, including quality, granularity, volume, gaps, maintenance, storage, and security [Haynes, 2016]. The stakeholders also discuss business cases that AI can be applied to by turning the mission into measurable goals and actions which inform model metrics optimization. Using AI to improve quality and efficiency of a specific critical or oft-occurring action or service is a good initial path [DataKind, 2020]. Common types of analysis include description, detecting events or anomalies, prediction, optimization, and behavioral change. At this point there is some consideration given for anticipated challenges with the AI: ethical considerations such as the privacy, transparency, discrimination/equity, and accountability issues around the project [DSSG, 2018a].

Project Implementation. Project implementation can take several paths including crowd sourcing events for ideation, prototype model development with feedback and refinement, handoff and maintenance agreements [DataKind, 2020]. What remains constant is the need for project champions in each stakeholder organization to lead momentum, clear communication, and manage expectations. The mission-oriented organization champions represent the model's end users, ensuring the model is understandable, trustworthy, and usable. The AI4SG practitioners are technical and project management experts in AI. They make AI concepts accessible while developing solutions [Niño *et al.*, 2017]. These experts usually come from academia or private industry, but for sustainment purposes, should knowledge transfer to technical teams within the mission organization. Motivation of stakeholders shape cost sharing and steps taken. Costs are often offset by grants and volunteer labor sourced by the AI4SG practitioner organization. As a result, implementation steps must include periodic milestones to account for volunteer turn and shifts in momentum [Niño et al., 2017]. Recently, these organization have noted the need for better inter-project coordination of resources and outcome scaling [Tekisalp, 2020]. For example, DataKind launched an accelerator with Microsoft in order to encourage sharing between projects in a similar problem domain [DataKind, 2020].

3 Our Approach

Women in Data Science (WiDS) is a gender-inclusive Booz Allen Hamilton program that seeks to hire, retain, and support diverse tech talent. Built on three pillars – Education, Sponsorship, and Social Good – the program has made executing AI4SG partnerships a core offering. From July 2018 through July 2020, WiDS partnered with nine mission-oriented organizations in domains such as health, environmental conservation, housing, education, and gender equality. Through our experiences, we developed and refined a framework, the Social Good Pipeline (SGP), that evolved the same *Partnership Discovery* and *Problem Scoping* methods of Section 2.3.

Similar to phases described in *Project Implementation*, we use a combination of open crowd sourcing events and closed working groups. The outcome of *Problem Scoping* is a concise measurable goal with a possible methodology approach, considered a strategy document. We then enter *Concept Development* by hosting a crowd sourcing hackathon. This approach allows us to quickly ideate methodology, find which shows the most potential, builds community momentum and provides the mission-oriented organization a quick set of demonstratable models. These rough prototypes offer the first round of feedback on understandability and model trustworthiness for the mission-oriented organization. *Solution Design* follows. In this phase we continue the refine, test, feedback cycle. Depending on the availability of funds, momentum, and ethical risk, consideration for mitigating potential bias, testing for fragility and usability is given till we reach a minimally viable product. Finally, *Solution Application* releases the model for usage, and we share the story.

3.1 Differentiators

Our iterations led to unique differentiators in risk tolerance, culture, and sustainable project scaling.

Risk Tolerance. The multi-phase format provides flexibility for various pathways to impact. Phases may be substituted, lengthened, paused or accelerated based on desired robustness, momentum and resource availability. Red flags, such as unexpected cost surges, revoked data access, lack of investment, or discussions that reveal a better methodology than AI, require off-ramps. We design phases to each conclude with a progressive outcome – literature research; AI strategy; prototype; documentation, pipelines and a hardened model – alternatively, recommendations on data collection or digital tools, or a technology roadmap to continue prototype development. When projects off-ramp, the mission-oriented organization still gains the experience of observing the stages, routine hurdles, and typical workstreams of an AI implementation process, increasing their AI literacy and building internal capacity. We save resources and attention for additional partnerships and iterate on the framework.

Culture. WiDS's risk tolerance and flexibility comes from a culture not only motivated by AI4SG but our mission to advance diversity in data science. Our partnership champions are not external volunteers, as criticized in Niño [2017]. Since our SGP leaders are responsible for setting the goals of the partnership, serving as the main contacts for the partner organization and driving the solutioning approach, we recognize their expertise and experience by fully funding their efforts. Through their participation in the SGP pipeline, leaders gain valuable technical and project management skills, and enable our firm to recruit and retain talent. Their work is in alignment with the firm's business goals and core values and is elevated by leadership. The result is motivated and dedicated expertise. This type of culture is not possible in a grants-based scenario that primarily rewards completion through publication.

Sustainable Project Scaling. There has been increasing demand for a holistic approach that brings together key organizations around common issue-areas [DataKind, 2020]. We start with a scale partner in mind during *Partnership Discovery* who will benefit from the AI development along with the primary mission-oriented organization. They participate in our open *Problem Scoping* phase and provide feedback on scope and applicability to their own use cases after prototyping phases. We have designed the partnership framework to fit within our public sector consulting business model to create a sustainable funding mechanism and accelerate the AI solution's application across the problem domain. Our scale partners are public sector organizations who increasingly request "show, don't tell" contract proposals. This has increased demand for rapid prototyping. The SGP is designed to have a low entry barrier, meaning the first couple phases provide many mission-oriented organizations with their first experience in AI while incurring little overhead. This stands in contrast to the grant program design by other companies and results in the pipeline's wide opening. As promising projects progress, we gain resource support from internal prototyping teams. These stories are then included in proposals to support our business model. This approach could be similarly used by for-profits in the product market and social entrepreneurs. We encourage academics to partner and access this alternative funding stream as well.

4 Conclusion

This paper analyzes partnership frameworks for applying artificial intelligence to solve social good challenges. It explains current challenges in partnership project implementation, and surveys key aspects of framework structures. We offered three contributions to improve AI4SG implementation: risk tolerance through flexibility, a motivated culture, and sustainable project scaling.

References

[Abebe & Goldner, 2018] Rediet Abebe and Kira Goldner. Mechanism Design for Social Good. *AI Matters*. 4(3):27-34, October 2018.

[Cavello, 2020] B. Cavello. PAI Launches Interactive Project to Put Ethical AI Principles into Practice. News, 2020. https://www.partnershiponai.org/pai-launches-interactive-project-to-put-ethical-ai-principles-into-practice/

[Chui *et al.*, 2018] Michael Chui, Martin Harrysson, James Manyika, Roger Roberts, Rita Chung, Pieter Nel, and Ashley van Hetern. Notes from the AI Frontier: Applying Artificial Intelligence for Social Good. Discussion Paper, 2018. https://www.mckinsey.com/featured-insights/artificial-intelligence/applying-artificial-intelligence-for-social-good

[Coulton *et al.*, 2015] Claudia J. Coulton, Robert George, Emily Putnam-Hornstein, and Benjamin de Haan. Harnessing Big Data for social good: A grand challenge for social work. In *Grand Challenges for Social Work Initiative Working Paper No. 11,* pages 2-21, Cleveland, OH, July 2015. American Academy of Social Work and Social Welfare.

[DataKind, 2020a] DataKind. DataCorps. https://www.datakind.org/datacorps

[DataKind, 2020b] DataKind. Applying AI to Societal Challenges in US Cities: DataKind & MicrosoftAI Accelerator Final Report. 2020. https://tinyurl.com/y767fvwu

[DSB, 2020] Data Science Bowl. Accessed May 2020: https://datasciencebowl.com/

[DSSG, 2018a] Data Science for Social Good. Data Science Project Scoping Guide. 2018. http://www.datasciencepublicpolicy.org/home/resources/data-science-project-scoping-guide/

[DSSG, 2018b] Data Science for Social Good. Legal Agreements. 2018. http://www.datasciencepublicpolicy.org/home/resources/legal-agreements/

[DSSG, 2018c] Data Science for Social Good. Data Security. 2018. http://www.datasciencepublicpolicy.org/home/resources/data-security/

[DSSG, 2019] Data Science for Social Good. Accessed May 2020: http://www.dssgfellowship.org/

[Floridi *et al.*, 2020] Luciano Floridi, Josh Cowls, Thomas C. King, and Mariaosaria Taddeo. How to Design AI for Social Good: Seven Essential Factors. *Sci Eng Ethics*. April 3, 2020. doi:10.1007/s11948-020-00213-5

[Google, 2019] Google. Accelerating Social Good with Artificial Intelligence: Insights from the Google AI Impact Challenge. September 2019. Accessed May 2020: http://services.google.com/fh/files/misc/accelerating_social_good_with_artificial_intelligence_google_ai_impact_challenge.pdf

[Hager *et al.*, 2017] Gregory Hager, Ann Drobnis, Fei Fang, Rayid Ghani, Amy Greenwald, Terah Lyons, David C. Parkes, Jason Schultz, Suchi Saria, Stephen F. Smith, and Milind Tambe. Artificial Intelligence for Social Good. Workshop Report. Washington, DC, June 2016. Computing Community Consortium.

[Haynes, 2016] Lauren Haynes. Introducing the Data Maturity Framework. Blog, 2016. http://www.dssgfellowship.org//2016/04/28/introducing-the-data-maturity-framework/

[ITU, 2019] ITU. AI for Social Good Summit. Accessed May 2020: https://itu.foleon.com/itu/aiforgood2019/home/

[Mannarswamy & Roy, 2019] Sandya Mannarswamy and Shourya Roy. Evolving AI from Research to Real Life - Some Challenges and Suggestions. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18),* pages 5172-5179, Stockholm, Sweden, July 2018.

[Moore, 2019] Jared Moore. AI for Not Bad. In Proceedings of the 13th Ingernational AAAI Conference on Web and Social Media, Munich, Germany, June 2019. Frontiers in Big Data.

[Niño *et al.*, 2017] Mikel Niño, Roberto V. Zicari, Todor Ivanov, Kim Hee, Naveed Mushtaq, Marten Rosselli, Concha Sánchez-Ocaña, Karsten Tolle, José Miguel Blanco, Arantza Illarramendi, Jörg Besier, and Harry Underwood. Data Projects for "Social Good": Challenges and Opportunities. World Academy of Science, Engineering and Technology International Journal of Humanities and Social Sciences. 11(5):1094-1104, 2017.

[Susha *et al*., 2019] Iryna Susha, Åke Grönlund, and Rob Van Tulder. Data driven social partnerships: Exploring an emergent trend in search of research challenges and questions. *Government Quarterly.* 36(1):112-128, January 2019.

[Tekisalp, 2020] Lale Tekisalp. Beyond Hype and Innovation: AI for Social Good. Event Report, 2020. https://www.partnershiponai.org/beyond-hype-and-innovation-ai-for-social-good/