17

Social Good

As you know from Chapter 7 and Chapter 13, the (fictional) information technology company JCN Corporation has several data science teams that work on different problems faced by the company and its customers. JCN Corporation's chief executive is a member of the Business Roundtable and a signatory to broadening the values of private industry from being solely driven by shareholders to being driven by other stakeholders too (this was covered in Chapter 15). In this environment, you recall that the fourth attribute of trustworthiness includes beneficence and helping others. Toward this end, you have the idea to start a *data science for social good* program at JCN to engage those data science teams part-time to conduct projects that directly contribute to uplifting humanity.

> "Imagine what the world would look like if we built products that weren't defined by what the market tells us is profitable, but instead what our hearts tell us is essential."
>
> —Vilas Dhar, president of Patrick J. McGovern Foundation

Taking a consequentialist view (remember consequentialism from Chapter 14), 'social impact' or 'making a difference' is promoting the total wellbeing of humanity (in expected value over the long term without sacrificing anything that might be of comparable moral importance).[1] But what does that really mean? And whose good or whose value of wellbeing are we talking about?

[1]Benjamin Todd. "What Is Social Impact? A Definition." URL: https://80000hours.org/articles/what-is-social-impact-definition, Nov. 2021.

> "The phrase 'data science for social good' is a broad umbrella, ambiguously defined. As many others have pointed out, the term often fails to specify good for whom."
>
> —Rachel Thomas, data scientist at Queensland University of Technology

It is dangerous for you to think that you or the data science teams at JCN Corporation are in position to determine what is an appropriate problem specification to uplift the most vulnerable people in the world. Data science for social good is littered with examples of technologists taking the shortcut of being paternalistic and making that determination themselves. If your data science teams are diverse and include people with lived experience of marginalization (see Chapter 16), then maybe they will be less paternalistic and push to have diverse, external problem owners.

> "Most technologists from the Global North are often not self-aware and thus look at problems in the Global South through the lens of technology alone. In doing so, they inevitably silence the plurality of perspectives."
>
> —Patrick Meier, co-CEO of WeRobotics

But who should those external problem owners be? Your first inclination is to look towards international development experts from large well-established governmental and non-governmental organizations, and consulting the seventeen UN Sustainable Development Goals (SDGs) listed in Chapter 15. But as you investigate further, you realize that there were a lot of struggles of power and politics that went into determining the SDGs; in particular, the lower-level targets beneath the seventeen goals may not represent the views of the most vulnerable.[2] You also learn that international development overall has many paternalistic tendencies and is also littered with projects that make no sense. Some may even be harmful to the people they intend to uplift.

> "Find algorithms that benefit people on their own terms."
>
> —Jacob Metcalf, technology ethicist at Data & Society Research Institute

Thus, while taking inspiration from the high-level topics touched on by the SDGs, you decide on the following theory of change for the JCN data science for social good program you are creating. Using machine learning, you will empower smaller, innovative social change organizations that explicitly include the knowledge of the vulnerable people they intend to uplift when they work towards social impact. (Collectively, civil society organizations and social enterprises—for-profit businesses that have social impact as their main goal—are known as *social change organizations*.) Toward developing a social good program within JCN Corporation, in this chapter you will:

- evaluate past data science for social good projects,

---

[2]Serge Kapto. "Layers of Politics and Power Struggles in the SDG Indicator Process." In: *Global Policy* 10.S1 (Jan. 2019), pp. 134–136.

- formulate a lifecycle for achieving a successful data science for social good program, and
- sketch out empowering machine learning architectures and platforms for promoting social good.

Before jumping into it, a few words on how you can gain internal support within JCN Corporation to devote resources to the program. There are several value propositions that go beyond appealing to the broadening stakeholder values that JCN Corporation is adopting and beyond appealing to the potential for positive public relations. First, machine learning problem specifications in social impact applications tend to have different constraints than those found in information technology and enterprise applications. Constraints are the mother of innovation, and so working on these problems will lead to new innovations for JCN. Second, by partnering with civil society organizations, JCN Corporation will receive valuable feedback and public references about its machine learning tools that enterprise customers may be unwilling to provide. Public references that allow JCN to tout its capabilities are distinct from positive public relations because they do not depend on the goodness of the application. Third, working on these projects attracts, retains, and grows the skills of talented data scientists in JCN Corporation. Fourth, if the program is run on JCN Corporation's cloud computing platform, the platform's usage will grow. Tax deductions for charitable giving are conspicuously absent from the value propositions because JCN Corporation will be receiving product feedback and possible cloud usage from the social change organizations.

17.1 Evaluating Data Science for Social Good

Throughout the book, you have taken on roles in several (fictional) social change organizations, including as a project manager with m-Udhār Solar (the provider of pay-as-you-go solar energy), as a data scientist with Unconditionally (the distributor of unconditional cash transfers), as a data scientist collaborator of ABC Center (the integrated social services provider), and as a problem owner with Alma Meadow (the granter of two-year fellowships to budding social entrepreneurs). Moreover, although the (fictional) Bank of Bulandshahr and Wavetel were launching the Phulo mobile telephony-based lending service with for-profit motives, the service is a vehicle for financial inclusion and upliftment. Thus, you have already seen some examples of projects that fall under the data science for social good umbrella. Their non-fictionalized counterparts were conducted as partnerships between social change organizations and data scientists acting in a 'for social good' capacity.[3]

[3]Hugo Gerard, Kamalesh Rao, Mark Simithraaratchy, Kush R. Varshney, Kunal Kabra, and G. Paul Needham. "Predictive Modeling of Customer Repayment for Sustainable Pay-As-You-Go Solar Power in Rural India." In: *Proceedings of the Data for Good Exchange Conference*. New York, New York, USA. Sep. 2015. Brian Abelson, Kush R. Varshney, and Joy Sun. "Targeting Direct Cash Transfers to the Extremely Poor." In: *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. New York, New York, USA, Aug. 2014, pp. 1563–1572. Debarun Bhattacharjya, Karthikeyan Shanmugam, Tian Gao, Nicholas Mattei, Kush R. Varshney, and Dharmashankar Subramanian. "Event-Driven Continuous Time Bayesian Networks." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. New York, New York, USA, Feb. 2020, pp. 3259–3266. Aditya Garg, Alexandra Olteanu, Richard B. Segal, Dmitriy A. Katz-Rogozhnikov, Keerthana Kumar, Joana Maria, Liza Mueller, Ben Beers, and Kush R. Varshney. "Demystifying Social Entrepreneurship: An NLP Based Approach to Finding a Social Good Fellow." In: *Proceedings of the Data Science for Social Good Conference*. Chicago, Illinois, USA, Sep. 2017. Skyler Speakman, Srihari Sridharan, and Isaac Markus. "Three Population Covariate Shift for Mobile Phone-based Credit Scoring." In: *Proceedings of the ACM SIGCAS Conference on Computing and Sustainable Societies*. Menlo Park, California, USA, Jun. 2018, p. 20.

17.1.1 What is Data Science for Social Good?

By happenstance, the examples in the book cover a narrow swath of possible uses of machine learning and artificial intelligence in social good. Moreover, due to the scope of the book, all but the ABC Center use case are focused on classification problems that lead to the allocation of something that, at face value, appears to be favorable to the recipient. The ABC Center example is focused on causal inference and causal discovery. Surveys of the data science for social good landscape find projects touching the following categories , which have a great deal of alignment with the SDGs:[4]

- accessibility,
- agriculture,
- education,
- environment,
- financial inclusion,
- health care,
- infrastructure (e.g. urban planning and transportation),
- information verification and validation,
- public safety and justice, and
- social work,

and touching the following technical approaches from artificial intelligence:

- supervised learning,
- reinforcement learning,
- computer vision,
- natural language processing,
- robotics,
- knowledge representation and reasoning,
- planning and scheduling,
- constraint satisfaction,

and many others.

Both lists are extremely vast and encompassing. As such, you should not think of social good as an application area of machine learning and AI, but as a paradigm or value system (paradigms are discussed in Chapter 15 as the precursor to values). Do not simply train a model on some dataset you downloaded that relates to agriculture or infrastructure or misinformation; that is not data science for social good. Do not create a system that helps privileged individuals discover which farmers markets

[4]Michael Chui, Martin Harryson, James Manyika, Roger Roberts, Rita Chung, Ashley van Heteren, and Pieter Nel. "Notes from the AI Frontier: Applying AI for Social Good." McKinsey & Company, Dec. 2018. Zheyuan Ryan Shi, Claire Wang, and Fei Fang. "Artificial Intelligence for Social Good: A Survey." arXiv:2001.01818, 2020.

currently have an inventory of kale; that is not data science for social good.[5] Data science for social good requires social change organizations to be problem owners who state the problem specification based on the lived experiences of their beneficiaries (and even better, bring their beneficiaries to a panel of diverse voices to inform the project).

Needless to say, the data science for social good you do in your program at JCN Corporation must be imbued with data privacy and consent, along with the first three attributes of trustworthy machine learning: competence, reliability (including fairness and robustness), and interaction (including explainability, transparency, and value alignment). This is especially the case because these systems are affecting the most vulnerable members of society.

17.1.2 How Has Data Science for Social Good Been Conducted?

Surveys of the data science for social good landscape find that nearly all efforts have been conducted as one-off projects that involve the development of a custom-tailored solution, irrespective of whether they are carried out as data science competitions, weekend volunteer events, longer term volunteer-based consulting engagements, student fellowship programs, corporate philanthropy, specialized non-governmental organizations, or dedicated innovation teams of social change organizations.

Creating such one-off solutions requires a great deal of time and effort both from the social change organization and the data scientists. There is limited reuse of assets and learnings from one project to the next because (1) every new project involves a different social change organization and (2) data scientists acting as volunteers are unable to conduct a sequence of several projects over time. Moreover, these projects typically require the social change organization to integrate machine learning solutions with their other systems and practices, to deploy those solutions, and monitor and maintain the solutions over time themselves. Very few social change organizations are equipped to do such 'last-mile' implementation, partly because their funding typically does not allow them to invest time and resources into building up technological capacity.

The confluence of all these factors has led to the state we are in: despite the data science for social good movement being nearly a decade long, most projects continue to only be demonstrations without meaningful and lasting impact on social change organizations and their constituents.[6] A project lasting a few months may show initial promise, but then is not put into practice and does not 'make a difference.'

17.2 A Lifecycle of a Data Science for Social Good Program

As you envision the JCN Corporation data science for social good program, you want to avoid the pitfalls that others have experienced in the past. But can your program jump right to the end goal of doing high-impact work, or is there an evolution it must go through? Sorry, there are no shortcuts. Just like an artist's or scientist's hot-streak of high-impact work begins with an exploration[7] phase that touches a

---

[5]Jake Porway. "You Can't Just Hack Your Way to Social Change." In: *Harvard Business Review* (Mar. 2013). URL: https://hbr.org/2013/03/you-cant-just-hack-your-way-to.

[6]Kush R. Varshney and Aleksandra Mojsilović. "Open Platforms for Artificial Intelligence for Social Good: Common Patterns as a Pathway to True Impact." In: *Proceedings of the ICML AI for Social Good Workshop*. Long Beach, California, USA, Jul. 2019.

[7]Lu Liu, Nima Dehmamy, Jillian Chown, C. Lee Giles, and Dashun Wang. "Understanding the Onset of Hot Streaks Across Artistic, Cultural, and Scientific Careers." In: *Nature Communications* 12 (Sep. 2021), p. 5392.

diversity of topics (which is then followed by a narrowly-focused exploitation phase that produces the impact),[8] a data science for social good program needs to begin broadly and go through the following three-step lifecycle, illustrated in Figure 17.1.

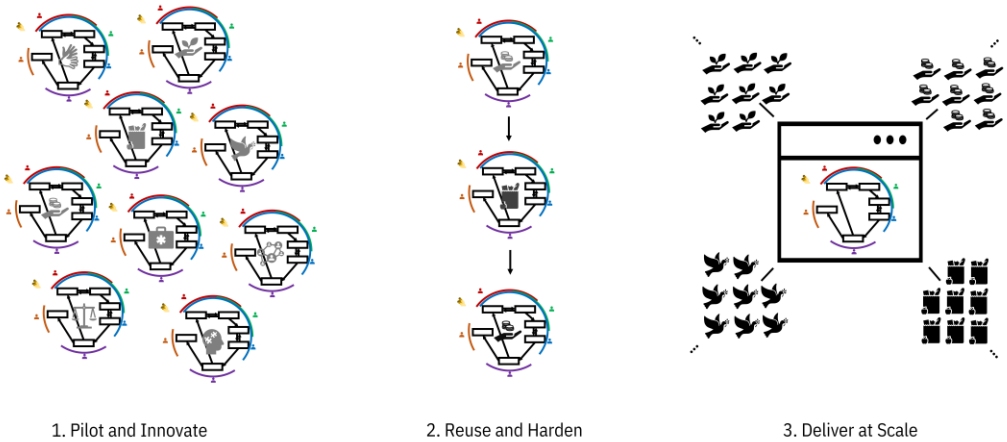

Figure 17.1. *Illustration of the three phases of the lifecycle of a data science for social good program: (1) piloting and innovating with a portfolio of projects, (2) reusing and hardening solutions to the common patterns, and (3) creating a usable platform that can reach a lot of social change organizations.* Accessible caption. Step 1, pilot and innovate, shows several different development lifecycles with icons for different social good applications in their center, colored gray to indicate they are not yet hardened. Step 2, reuse and harden, shows a sequence of three development lifecycles in which the social good application icon gets progressively darker to black to indicate hardening. Step 3, deliver at scale shows a development lifecycle inside a computer window illustrating its incorporation into a platform, touching tens of social good applications.

1. *Pilot and innovate.* You should conduct several individual projects to learn about the needs of social change organizations that may be addressed by machine learning. In this phase, your data scientists will also gain the experience of conducting multiple projects and start seeing commonalities across them. While doing so, JCN Corporation will gain from new innovations under new constraints. You can choose to be somewhat intentional in the application area of social good to match corporate values or in the technical area of machine learning to match technical areas of interest, but not overly so.
2. *Reuse and harden.* Once you have several projects under your belt, you must step back and analyze the common patterns that emerge. Your goal at this stage is to develop common algorithms or algorithmic toolkits to address those common patterns in as reusable a way as possible. You want to meet the needs of multiple social change organizations using a common model or algorithm. This type of machine learning innovation is unique; most data scientists and machine learning researchers are not trained to step back and abstract things in this way, so it will be a

[8]The word 'exploit' is used in a positive sense here, but is used in a negative sense later in the chapter.

challenge. However, this sort of insight and innovation is precisely the feedback that will be helpful for JCN Corporation's teams developing software tools and products for conducting data science.

3. *Deliver at scale.* Those common reusable algorithms will not make high impact until they are made available within an environment that low-resourced and low-skilled social change organizations can be empowered to tweak, use, and maintain. (Refer to inclusive low-code/no-code architectures in Chapter 16 for a related discussion.) The delivery will likely be 'as-a-service' on JCN Corporation's cloud-based environment. Software-as-a-service is software that is licensed as a subscription, is centrally hosted, and is accessed by users using a web browser. Therefore, integration with other systems is greatly simplified and the responsibility for maintenance falls on JCN Corporation rather than the social change organization.

You are probably comfortable with the first phase of this data science for social good program lifecycle. As long as you ensure that social change organizations—representing the interests of their beneficiaries who have lived experience of vulnerability—are the problem owners and involved in evaluation, then the JCN Corporation data scientists can approach the portfolio of projects in this phase in a manner they are used to.

The second phase presupposes that there *are* common patterns in social good projects that can be addressed using common models or algorithms. Evidence is starting to mount that this is indeed the case. For example, the same algorithm for bandit data-driven optimization is used in social good applications as varied as feeding the hungry and stopping wildlife poachers.[9] As a second example, most of the social good use cases (fictionally) covered in the book are quite different from each other, but are all fair allocation problems posed as binary classification that can be addressed using a common algorithmic toolkit such as AI Fairness 360, a library of fairness metrics and bias mitigation algorithms.[10] Moreover, large language models have been fine-tuned for several disparate social good domains such as collecting evidence for drug repurposing and simplifying text for people with low-literacy or cognitive disability.[11] (Large language models are a kind of foundation model; introduced in Chapter 4, foundation models are machine learning models trained on large-scale data that can be fine-tuned for specific problems.)

The third phase of the lifecycle of a social good program is mostly unproven as yet, but is what you should be working toward in the program you intend to start at JCN Corporation. The result is an accessible and inclusive data science for social good *platform* that is described in the next section.

---

[9]Zheyuan Ryan Shi, Zhiwei Steven Wu, Rayid Ghani, and Fei Fang. "Bandit Data-Driven Optimization: AI for Social Good and Beyond." arXiv:2008.11707, 2020.

[10]Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. "AI Fairness 360: An Extensible Toolkit for Detecting and Mitigating Algorithmic Bias." In: *IBM Journal of Research and Development* 63.4/5 (Jul./Sep. 2019), p. 4.

[11]Shivashankar Subramanian, Ioana Baldini, Sushma Ravichandran, Dmitriy A. Katz-Rogozhnikov, Karthikeyan Natesan Ramamurthy, Prasanna Sattigeri, Kush R. Varshney, Annmarie Wang, Pradeep Mangalath, and Laura B. Kleiman. "A Natural Language Processing System for Extracting Evidence of Drug Repurposing from Scientific Publications." In: *Proceedings of the AAAI Conference on Artificial Intelligence*. New York, New York, USA, Feb. 2020, pp. 13369–13381. Sanja Stajner. "Automatic Text Simplification for Social Good: Progress and Challenges." In: *Findings of the Association for Computational Linguistics*. Aug. 2021, pp. 2637–2652.

Before getting there, two comments on the term 'scale.' Scaling is a paradigm seen as paramount in much of the technology industry, and is the main reason to pursue a digital platform that can be built once and used by many. However, scaling is not the mission of many social change organizations; although some would like to grow, many would like to remain small with a very pointed mission.[12] Moreover, scaling as an overriding paradigm is not free from criticism and can be seen as a means for exploiting the most vulnerable.[13] In creating a social good program and platform with JCN Corporation, your goal is to make the work of all social change organizations easier, irrespective of whether they would like to scale themselves. You can control any possible exploitation by centering the values of the most vulnerable throughout the development lifecycle.

17.3 A Data Science for Social Good Platform

A foundation model, algorithm, or algorithmic toolkit that applies broadly to the problems faced by many social change organizations is an excellent start, but it is not enough to satisfy your theory of change. These technological artifacts alone do not empower social change organizations because they require a level of skill and infrastructure in data science and engineering that the organizations typically lack. Incongruously, social change organizations are typically low-resourced just like the people they serve, especially in comparison to private corporations embracing machine learning with large data science teams. You can say that social change organizations are at the 'bottom of the pyramid' among organizations. (In its typical usage, the term 'bottom of the pyramid' refers to the socioeconomically poorest or least wealthy group of people.)

A washing machine, stove, or ultrasound imaging machine designed for a wealthy, high-resourced context will not cut it in a low-resourced context. The core technology has to be put into a form factor that makes sense for bottom-of-the-pyramid users. The same is true of machine learning for social change organizations. In general, bottom-of-the-pyramid innovation has the following twelve principles:[14]

1. focus on (quantum jumps in) price performance;
2. hybrid solutions, blending old and new technology;
3. scalable and transportable operations across countries, cultures and languages;
4. reduced resource intensity: eco-friendly products;
5. identify appropriate functionality;
6. build logistical and manufacturing infrastructure;
7. deskill (services) work;
8. educate (semiliterate) customers in product usage;

---

[12]Anne-Marie Slaughter. "Thinking Big for Social Enterprise Can Mean Staying Small." In: *Financial Times* (Apr. 2018). URL: https://www.ft.com/content/86061a82-46ce-11e8-8c77-ff51caedcde6.

[13]Katherine Ye. "Silicon Valley and the English Language." URL: https://book.affecting-technologies.org/silicon-valley-and-the-english-language/. Jul. 2020.

[14]C. K. Prahalad. *The Fortune at the Bottom of the Pyramid: Eradicating Poverty Through Profits*. Upper Saddle River, New Jersey, USA: Wharton School Publishing, 2005.

9. products must work in hostile environments;
10. adaptable user interface to heterogeneous consumer bases;
11. distribution methods designed to reach both highly dispersed rural markets and highly dense urban markets; and
12. focus on broad architecture, enabling quick and easy incorporation of new features.

What are the important points among these principles for a machine learning platform that empowers social change organizations and what is a platform anyway?

A digital platform is a collection of people, processes, and internet-based tools that enable users to develop and run something of value. Therefore, a machine learning platform contains web-based software tools to carry out the entire machine learning development lifecycle from the problem specification phase all the way to the deployment and monitoring phase, with roles for all the personas including problem owners, data scientists, data engineers, model validators, operations engineers, and diverse voices. Importantly, a machine learning platform is more than simply an off-the-shelf programming library for modeling.

In fact, there are three kinds of machine learning capabilities: (1) off-the-shelf machine learning packages, (2) machine learning platforms, and (3) bespoke machine learning builds.[15] At one extreme, off-the-shelf packages are useful for top-of-the-pyramid organizations with a high level of data science skill and a high level of resources, but not for bottom-of-the-pyramid social change organizations. At the other extreme, bespoke or custom-tailored development (which has been the predominant mode of data science for social good over the last decade) should only be used for extremely complex problems or when an organization needs a technological competitive advantage. These are not the circumstances in which social change organizations typically operate; usually their problems, although having unique constraints, are not overly complicated from a machine learning perspective and usually their advantages in serving their beneficiaries are non-technological. Thus, it makes sense to be *just right* and serve social change organizations using machine learning platforms.

What does a machine learning platform for the bottom of the pyramid entail? Translating the twelve general principles to a machine learning platform for social change implies a focus on appropriate functionality, adaptable user interfaces, deskilling, broad architecture, distribution methods, and education. You'll obtain appropriate functionality by paring down the machine learning capabilities to the core model, algorithm, or toolkit that is reusable by many different social change organizations with similar needs, as discussed earlier. Such foundational capabilities mean that the algorithms have to be created only once and can be improved by a dedicated machine learning team that is not reliant on, or part of, any one social change organization.

The remaining aspects touch on the last-mile problem. You can achieve adaptable user interfaces and deskilling by following the inclusive architecture presented in Chapter 16 for people with lived experience of marginalization. Such an architecture takes the scarce and expensive skill of data scientists out of the development lifecycle through low-code/no-code and auto ML. Low-code/no-code and auto ML should make it easy to configure and fine-tune the machine learning capability for the

[15]Andrew Burgess. *The Executive Guide to Artificial Intelligence: How to Identify and Implement Applications for AI in Your Organization*. London, England, UK: Palgrave Macmillan, 2017.

specific task being approached by the social change organization. It should also be easy to slightly change the definition of an outcome variable and apply the model to a new setting with slightly different features. The interface should also provide a data catalog and tools for managing data. Moreover, the interface should include meaningful and easy to consume visualizations of the output predictions. The focus should be to simplify, simplify, simplify, but not so much that you are left with something meaningless.

A web and cloud-based platform is specifically designed to support quick and easy incorporation of new capabilities. Any improvements to the machine learning diffuse to social change organizations automatically. Similarly, cloud-based platforms are designed in a way that allow broad distribution to any device anywhere there is an internet connection. This method of delivery is starting to lead to turnkey deployment and monitoring of machine learning systems.

Finally, the last component of a machine learning platform for social impact is education: teaching and reference materials, tutorials, how-to guides, examples, etc. presented in the language of social change organizations. It must be presented in a way that people starting at different skill levels all have an on-ramp to the content. An important part of the education for members of social change organizations is sparking the imagination of what's possible using machine learning in the social impact sector. A persona that has not come up in the book so far, a *broker* who bridges the gap between members of social change organizations and the data science world by translating and aligning the concepts used in each field, is very useful in the education component of the platform.[16]

Have you noticed something? All of the desirable attributes of a machine learning platform seem to be desirable not only for empowering social change organizations, but also desirable for any organization, including ones at the top of the pyramid. And that is the beauty of bottom-of-the-pyramid innovation: it is good old innovation that is useful for everyone including JCN Corporation's enterprise customers.

Beyond the design and the ease of use of the platform, a critical aspect for you to sustainably bring the platform and overall data science for social good program to fruition is winning the support of large grantmaking foundations that fund social change organizations. First, the foundations must give some sort of implicit permission to social change organizations to use the platform and provide them enough leeway in their budgets to get started. Second, in a similar vein as international development projects specified without the perspective of vulnerable people, there are many international development efforts whose funding did not provision for maintenance and long-term support beyond the initial headline. JCN Corporation will not be able to sustain a data science for social good platform you create without grants for its maintenance, so you'll need to line up funding. Foundations are beginning to see the need to support technology efforts among their grantees,[17] but are not yet ready to fund a platform operated by a private corporation.

You have your work cut out for you to launch a data science for social good program at JCN Corporation and push it along the lifecycle beyond just the initial set of projects to common algorithms

---

[16]Youyang Hou and Dakuo Wang. "Hacking with NPOs: Collaborative Analytics and Broker Roles in Civic Data Hackathons." In: *Proceedings of the ACM on Human-Computer Interaction* 1.CSCW (Nov. 2017), p. 53.

[17]Michael Etzel and Hilary Pennington. "Time to Reboot Grantmaking." In: *Stanford Social Innovation Review*. URL: https://ssir.org/articles/entry/time_to_reboot_grantmaking, Jun. 2017.

and then a scalable platform. But with enough conviction, wherewithal, and luck, you just might be able to pull it off. Go forth, you genuine do-gooder![18]

17.4 Summary

- Data science for social good—using machine learning in a beneficent way—is not an application area for machine learning, but a paradigm and value system.
- The goal is to empower social change organizations in the development of machine learning systems that help uplift vulnerable people on their own terms.
- The decade-long experience with data science for social good has rarely yielded truly impactful results because individual projects fail to overcome the last-mile problem.
- Social change organizations are typically low-resourced and need much more than just code or a custom solution to be able to use machine learning in their operations.
- Machine learning platforms that are specifically designed to deskill data science needs and minimize the effort for deployment, maintenance, and support are the solution. Such platforms should be built around common algorithmic patterns in the social impact space that you start seeing by conducting several projects over a lifecycle.
- All the attributes of trustworthy machine learning are essential in applying machine learning for social impact, including fairness, robustness, explainability, and transparency.

[18]William D. Coplin. *How You Can Help: An Easy Guide to Doing Good Deeds in Your Everyday Life*. New York, New York, USA: Routledge, 2000.