



Decision Guide for Considering AI Use

Is AI right for your organization?

Non-profit organizations and government agencies face ongoing pressure from boards, elected officials, vendors, funders, and others to incorporate artificial intelligence (AI) into their operations. But, the technology often isn't the right fit, and adopting it can threaten your resources and reputation. This guide will help you think through AI before committing your organization to it.



What will you use AI to do?

Companies selling AI make lots of promises about what the technology can do. You should have in mind a specific goal for how it will improve your services and the lives of the people you serve or are accountable to. Current evidence establishes that AI often does not increase workplace efficiencies and could result in large-scale disaster, particularly for vulnerable people.



Is AI a legitimate way to do that?

Using AI to cut public benefits, accuse low-income people of fraud, decide eligibility for services, or deliver sensitive medical or legal information has failed repeatedly with devastating consequences to real people. These uses don't effectively serve the public and are not legitimate ways to make decisions or administer programs. Consider whether AI for your specific use reflects the public good and your values.



Are you prepared to vet AI?

Companies may sell you an off-the-shelf AI system, customize an off-the-shelf system for your use, or build you something entirely from scratch. In any case, you must be prepared to consider the full range of issues with what is offered.

- **Your specific gains.** What specific gains or efficiencies do you hope to get from the AI system?
- **Non-AI alternatives.** How is the AI better than human-based practices you currently use or could implement?
- **The seller's track record.** What's the AI seller's track record on projects like this? Does the technology do what the seller claims? Can you independently verify that?
- **Your technical expertise.** Do you have technical expertise on staff to understand and interrogate what the AI seller promises throughout the development process? If not, can you hire the needed expertise?
- **Your capacity.** Developing an AI project will require extensive staff time and resources to manage properly. Do you have the capacity needed?
- **The data the AI is based on.** Will you provide the data the seller uses to train the AI system? If not, where does the data come from? Is the sharing and use of such data permitted? Do you have full access to it? How will you ensure that the AI system doesn't use bad, biased, or otherwise harmful data?
- **The criteria the AI uses.** Will you decide what criteria the AI system will use to make decisions or produce other outputs? If not, what criteria will be used and why are those criteria chosen instead of others? Do you have full access to the criteria? Has that criteria been meaningfully tested?
- **The potential harms to real people.** How will you project the impact the AI system has on your operations and the people whom you serve? How will it work for people with language or disability access needs? How will you ensure that the AI system doesn't produce harmful or inaccurate results? How will you monitor for such harms?
- **Cyber and data security.** Does the AI create new vulnerabilities that hackers can exploit? How will you ensure sensitive information is not exposed?



Are you prepared to vet AI? (Continued)

- **The contract language.** How is success or failure defined? Does the contract include relevant targets or goals for the desired gains or efficiencies? Does the contract allow you to ask the seller to make needed changes? Does the contract allow you to hold the seller accountable with lawsuits, public disclosure of the data and criteria, evaluations, etc.?
- **Integrating and implementing.** How will the AI be integrated into your existing technology systems and workflows? How will staff be trained?
- **Legal requirements.** Are you aware of and complying with various laws and ethical requirements that might apply to your use of AI?



Will AI degrade the necessary human touch of your services?

Non-profit and government agencies often interact with people in crisis where empathy and connection are vital to building trust, getting a full understanding of the situation, and providing a service effectively. AI usage could undermine your ability to do this. AI could also change how you measure success by inappropriately emphasizing strict number counts over quality.



How will you involve the communities you serve in the process of developing and using the AI?

Members of the public, particularly low-income or otherwise marginalized people, should not be experimented on or be forced to submit to AI. If the AI will affect them directly, they should have a formal consultation role—not just a performative show—with the ability to offer input, ask questions, get answers, and decide whether to support the AI use or not. To make participation meaningful, consider scheduling, childcare, transportation, opportunity costs, language and disability access, comfort level using technology, and plain-language explanations.



What specific accountability mechanisms will make this project go right for the communities you serve?

Political, market, and legal accountability mechanisms are often broken or inapplicable with respect to government or non-profit services, particularly when low income people are affected. What will ensure this goes right?



How will you integrate your staff into the process of developing and using the AI?

AI use will affect how staff do their jobs, their workloads, the satisfaction they feel, their skillset, and their interactions with the people you serve. Staff should have a formal consultation role with the ability to ask questions, get answers, and decide whether to support the AI use or not.



What risk are you taking on by adopting AI systems?

When you use AI, consider the specific risks to resources, reputation, staff morale, and the good will of the communities you serve.




Once AI is being used, do you have the expertise and staff capacity to provide ongoing oversight?

You will need to monitor AI use to determine if it is meeting your goals, deviating from projections, or causing unanticipated harms. Make sure you have the expertise and staff capacity needed to do so effectively.




Once AI is being used, can you still perform the same function in a non-AI way?

The AI system may not work appropriately or as planned and may need to be changed or dismantled. Or, people subject to the AI may need or prefer a non-AI interaction. You must have viable ways both of overriding the AI and performing the same function without it.



Once AI is being used, are you acting transparently and accountably to the people you serve and the broader public?

You will need to share the results of the AI with respect to the goals you set out, identify shortcomings and unanticipated consequences, solicit ongoing feedback, and consider changing it.



Given all the complexities involved, is AI still likely to lead to a net gain in the efficiency or quality of your services?

The efficiency gains you hope to achieve are likely to be reduced by the resources you use to properly develop, implement, oversee, and address any fallout from using AI (for example, complaints, questions, or appeals). Consider whether you'll actually come out ahead.



Are you willing to accept the hidden harms of using AI?

Workers in other countries must toil in awful conditions to provide the data for companies building AI systems. Data centers that store and process information are correlated with spikes in energy prices, often drain water from places with little to spare, and otherwise harm the environment. Using AI perpetuates the endless gathering of data about us, which further entrenches corporate power and can be used by governments for authoritarian ends. Consider these harms, too.