

A Step-by-Step Framework to Mitigate Al Risk



By Rachel Sams

Lead Consultant and Editor

Resource Type: Articles, Risk eNews

Topic: Data Privacy, Tech Risk, Cybersecurity

I have some bad news.

You and your team members don't have all the skills you need to deal with Al risk.

Here's some slightly better news: No one else has all the skills they need to deal with AI, either! The technology's moving so fast even industry experts are hard pressed to keep up.

Now for the good news: You and your team members aren't experts on AI, but you are experts on your nonprofit's capabilities and challenges. You are experts on making it work with what you have. You probably do that every year with your budget, you do it every day when you provide services, and you can do it with artificial intelligence, too.

Al is a broad term for the use of computer science and data to enable problem solving. Machines or processes that use Al have been designed to respond to input similarly to the way humans would. Our accompanying article in this issue, "Three Very Human Qualities to Help You Mitigate Al Risk," offers a primer on qualities you and your team members should hone to help you navigate the world of artificial intelligence.

As you experiment with that, we offer a step-by-step framework to evaluate AI risk and make good decisions about when and how to use AI in your nonprofit—and when not to. The framework is designed to be flexible, not prescriptive. Every organization is different, and an AI use that might benefit one nonprofit's team, community, and clientele might feel very wrong to another nonprofit. Consider this a general guide for how to approach your organization's AI risk journey.

Because Al touches so many areas within nonprofits in an incredibly diverse sector, your progress through this Al risk framework may not be linear—you might sometimes work on two or three steps at once. Wherever your journey takes you, this framework can provide peace of mind in the fast-moving world of Al.

How to Begin

If your nonprofit is early in its AI journey, everything may feel scary. Some aspects of AI will likely continue to spook or worry your team over time, but as you build more experience with the technology, your confidence in assessing AI risk will grow. These steps will build a strong foundation for wherever your AI journey leads you.

Make sure your nonprofit has strong practices to center equity. This was a key piece of advice from

Sarah Di Troia, senior strategic advisor for project innovation at Project Evident, during a *Chronicle of Philanthropy* webinar about AI for nonprofits last year. Does your nonprofit have practices to center the voices of people who will be most affected by a change? Those practices are necessary to create AI experiments that avoid harm to communities and groups that have been marginalized, Di Troia said. Organizations that don't yet have that foundation aren't ready to experiment with AI and should focus first on establishing strong equity practices. If you don't know whether you have a strong equity foundation, some great questions to ask might include "How do voices get heard at this organization?" and "Whose input does this organization act on?" If you need help to assess that, Equity in the Center, the National Council of Nonprofits and the Building Movement Project, among others, offer resources for nonprofits working to center equity.

Empower humans in your transition to using AI. Allison Fine, Beth Kanter, and Philip Deng, writing in the *Stanford Social Innovation Review*, encourage nonprofits to take the time to explain to employees why your organization is considering AI use. Ask team members to share their hopes and fears about AI. What information do they need to understand its possibilities, limitations, and risks? Seek team members' input into how your organization should and shouldn't use AI. Provide access to webinars and training sessions to upskill your team. The *Chronicle of Philanthropy*, Microsoft and others hold frequent sessions on AI for nonprofit organizations, some of which are free.

Your team members may worry that using AI will lead to layoffs or job changes. They may or may not share those fears with you. If they do, acknowledge those concerns and help assess what employees need to feel more comfortable with the technology. Don't make promises you can't keep; you may not be able to guarantee that no one will ever lose their job because of AI at your nonprofit. And leaders should never promise that jobs within an organization won't change and evolve over time. Listen, emphasize your nonprofit's long-term goals for AI use, and ask employees how they'd like to participate in pursuing those goals.

Start small. Pilot an AI experiment. Ask yourself: What pain point could AI help solve in your organization? Choose one and monitor it closely before expanding your use of AI. Beth Kanter shared this strategy in a 2023 webinar on how nonprofits can tap into curiosity to overcome AI fears. Kanter urged nonprofits to consider: Where are the bottlenecks in your organization? How could AI address them? What safeguards will you put in place to minimize harm? Set clear guidelines about how you will evaluate your trial, what success would look like, and what would cause you to stop the experiment immediately.

Used well, AI can help your team members spend less time on repetitive processes and more time on bigpicture issues and building relationships—the things humans were born to do. But to use AI well, you must center people and your mission at every step: the community members you serve, the team members who serve them, and the change or impact your nonprofit intends to make in our world.

How to Lay the Groundwork for Next Steps

Once you've set your nonprofit's foundation for AI and completed a successful test, you'll be in a good position to start making some rules. You'll want to develop a code or policy to govern how you will make decisions about AI as your use expands.

Key things you'll want to consider including in your code:

The purpose of your AI policy. Why does your organization want to experiment with AI? What aspects of your nonprofit's work will this policy touch?

What kinds of AI use are encouraged, and within what parameters. What behaviors and uses fit with your nonprofit's mission and values?

What kinds of AI use are prohibited. What behaviors and uses will your nonprofit not allow under any circumstances?

How your organization will train, equip and educate team members to use AI. How will you work with your team to find out what skills they want to hone and help them do that within your budget?

How you will preserve the security and privacy of data in your AI use. What practices will you use to safeguard sensitive data? What security requirements will you have for AI services, vendors and products? What

will you do if data is breached? When and how will you use informed consent, and what opt-out options will you give constituents?

When and how you will disclose your nonprofit's use of AI to internal and external audiences. How will you document and communicate your use of AI?

What measures you will take to ensure accuracy and mitigate or avoid bias in your use of Al. How will you create safeguards to reduce the risk of plagiarizing from published material? Some options to consider:

- fact-checking any work created with generative Al against official sources, such as government data or your own;
- running a Google search on 250 words of any Al-created or assisted work to help gauge whether it's been previously published;
- using generative AI to iterate and improve your team's own work, rather than to create new work from scratch.

Also consider: How will you educate your team about potential Al biases and check for them? How will you ensure that a diverse group of constituents reviews all Al-generated work? What will you do if bias is found in Al-generated work, before or after the fact?

What consequences will result from intentional or unintentional violations of the policy. What responsibilities do team members have to report suspected violations? How should they report them? What will happen if someone made an honest mistake in their use of AI? What if someone intentionally misled others about their use of AI, or used it in a malicious way?

NRMC has created a draft Artificial Intelligence Policy, available to subscribers of our My Risk Management Policies product, which is available at www.myriskmanagementpolicies.org and costs \$179, or \$29 for Affiliate Members. Other organizations, including nonprofit consultants Roundtable Technology and Tech Soup, offer draft Al policy templates or guides for developing them, which can provide additional insight on what to consider and include in your organization's policy.

Take the time to evaluate any template or prototype Al policy against your mission and your nonprofit's unique opportunities, challenges, and constituencies. What parts of the policy might work well for you? What parts might not be a fit or need reworking?

Once you establish your policy, revisit it regularly. You may need to make changes and additions as you uncover new challenges and benefits of Al use. That's okay. A simple, flexible structure will provide a consistent backbone for your Al use as technology rapidly evolves and allow you to make changes along the way.

How to Advance

If you have your equity approach in place, you've sought feedback from your team on AI and started helping them upskill, and you have a pilot project in mind, here are some great next steps on your journey.

Keep humans in charge. Design any AI experiments to ensure team members regularly check the results of any AI processes you use. Review any content created by AI before distributing it. Create a checklist of what team members should review AI processes and content for. Spell out what you do and don't want AI-created work to contain.

Don't adopt AI tools you don't understand. Ask your vendors questions. Then ask questions about their answers. Keep asking "And then what?" or "What could happen next?" to uncover the unknowns in any AI tool you're considering. Ask your team: What's the best that could happen if we used this? What's the worst? What are some things that might or might not happen? Until you can explain the technology you want to use to someone who has no baseline knowledge of it, you're not ready for an AI experiment.

Here are some ethical questions to help you evaluate AI technology options, adapted from Automotive World.

- How do the systems you're considering collect data?
- How diverse and how credible are their data sources? How relevant are they to the context in which you want to use them? How do systems use the data? Where do they store it? Who can see it? How does this benefit employees, people receiving services, and the organization? Who could use this data for harm

and how?

- To avoid harm: Can you limit the data your nonprofit collects? Can you collect only statistics, without gathering any personal information or identifiable data? Can you avoid sending data to the cloud?
- What safeguards do vendors have in place to guard against cyberbreaches of their technology? How do those reflect the sensitivity level of the data they're storing?

How to Iterate

Once you've set up your Al approach and crafted your policy, it's important to revisit it regularly. The technology and the ethical frameworks around generative Al are evolving fast. Here are some ways to ensure your team continues to improve at managing Al risk.

Take time to evaluate. Set aside time quarterly with key team leaders to review these questions:

- How well did your Al use help you maximize your ability to deliver on your mission?
- What challenges arose in your use of AI this quarter? How did you meet them?
- How do you want to change your use of Al in the next quarter—scale back, scale up, or maintain? What
 new safeguards do you need to put in place to ensure the integrity of your Al work for staff and
 constituents?

Keep learning. Block time on your calendar to learn about AI. Attend webinars by experts, read what practitioners are writing, or join discussion groups about AI in nonprofit associations you belong to. Ask your team members each quarter what additional AI training they need to do their jobs well. Work with vendors or partners to see if you can obtain training pro bono or at a nonprofit discount.

Reminders for the Journey

The journey to understand AI risk will likely continue throughout our lifetimes and beyond. Remember that it's a marathon, not a sprint, and pace yourself.

No matter which step of this framework best reflects where your organization is on its AI risk journey, a few key principles can help you continue to learn and grow.

Don't wait. You and your team members may be frightened or intimidated by what you're hearing about Al. Al does have plenty of concerning capabilities. But the more you educate yourselves, the more confidence you'll gain in your ability to make good decisions about Al as individuals and as a team.

Don't rush. On the flip side, don't scramble to unveil a splashy Al initiative or replace a lot of your nonprofit's processes with Al at once. Build your foundation of equity, center humans, and start small.

Have fun. We're at a groundbreaking moment in the evolution of artificial intelligence. Al will continue to evolve, but this moment will never come again. Laugh with your team members at Al's quirks and hallucinations and give yourself the time and space to think about what you're learning. It's amazing what humans can do with a little time and space.

Rachel Sams is Lead Consultant and Editor at the Nonprofit Risk Management Center. As she researched AI risk for nonprofits over the past year, she experienced every stage of this framework multiple times. Reach her with thoughts and questions about how to manage artificial intelligence risk at rachel@nonprofitrisk.org or (505) 456-4045.