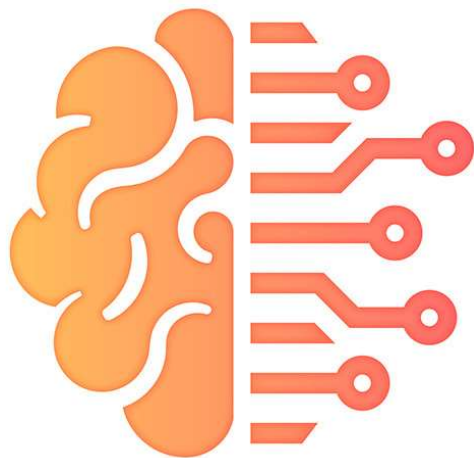


Nonprofits & NGOs

8 Steps Nonprofits Can Take to Adopt AI Responsibly

How organizations can engage with AI in smart and ethical ways right now.

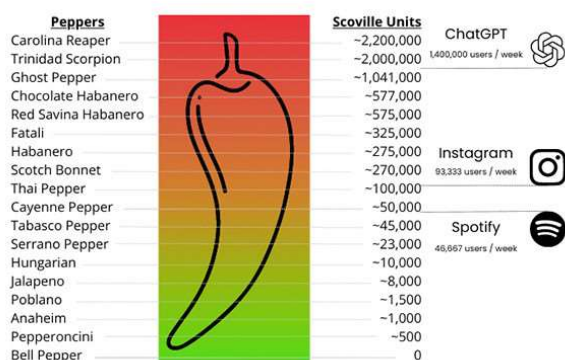
By Beth Kanter, Allison Fine & Philip Deng | Sep. 7, 2023



(Illustration by iStock/ilyast)

The breakneck rate at which people have adopted the generative AI writing tool ChatGPT has been astonishing. ChatGPT, which creates new text based on user prompts, surpassed 100 million monthly active users in just two months, according to [data compiled by UBS](#). Compare that to TikTok and Instagram, hit platforms that took nine and 30 months, respectively, to reach that same benchmark.

Using the [Scoville scale](#), a scientific measure of the spiciness of chili peppers, we can, somewhat unscientifically, analogize our way to understanding just how much of an outlier the ChatGPT phenomenon really is in the history of software applications.



This chart compares the relative

Generative AI is now in stage two of the [Gartner Hype Cycle](#), a graphic representation of innovation maturity used in the business sector. It has moved past stage one, “innovation trigger,” wherein the technology shows up everywhere all at once, and is now at its “peak of inflated expectations.” We have seen this cycle many times before with technologies like personal computers, mobile phones, and social media. However, there’s one difference

spiciness of chili peppers with the rate of user growth for some of the fastest-growing apps in history. (Image by Philip Deng)

between AI and previous ones: It's making decisions that only people could make until right now. This includes decisions about who to hire and what stories to tell, and it makes this moment both exciting and terrifying.

In the social sector, many nonprofits believe AI will solve all our work problems, instantly and painlessly. (It won't.) Others have outsized fears about it wreaking havoc and supplanting staff. Some organizations are jumping in early; others will wait. But one thing is clear: Nonprofit engagement with AI is necessary and requires thoughtful preparation.

Are you enjoying this article? Read more like this, plus *SSIR's* full archive of content, **when you subscribe.**

How Nonprofits Are Using AI Now

Fully integrating AI into the workplace will likely take years, as **a recent McKinsey study** indicates, but the process is already underway. Fundraisers are currently using it to help write donor thank-you notes, newsletters, grant proposals, and press releases. One development officer told us that using ChatGPT to help draft and edit these kinds of materials is saving her many hours of work a week. And a grant writer for a human services agency shared that she used ChatGPT to edit three different grant proposals for conciseness and all three were funded!

Everyday software products such as Google Workspace and Microsoft Office have also begun to incorporate generative AI functions to help users craft presentations, respond to emails, and schedule meetings. Nonprofits are just starting to use these tools, and efforts like Microsoft and Data.org's **Generative AI Skills Challenge**, a grant for generative AI training and upskilling to drive social impact, will no doubt accelerate adoption.

Other types of AI are taking hold in organizational life as well. Chatbots (conversational interfaces using **large learning language models**) like ChatGPT can answer questions from website users or provide internal information to nonprofit employees on demand. AI-powered screening tools are assessing potential employees and service recipients. And there are numerous **AI-driven fundraising tools** to help with donor prospecting and managing donor communication.

Handling AI With Care

Used carefully and skillfully, AI has the potential to save people thousands of hours of rote, time-consuming work a year, freeing them up to do the kinds of things only people can do: solve problems, deepen relationships, and build communities. And approaches like “co-boting,” which involves using AI tools to augment rather than replace human jobs, can bring out the best in both bots and people and yield the dividend of time.

However, we’re concerned that organizations eager to collect on that dividend of time are diving into the use of incredibly powerful technology without enough consideration of the human or ethical ramifications. In 2022, the National Eating Disorders Association (NEDA), the largest nonprofit organization supporting individuals and families affected by eating disorders, **became the poster child** for rash and reckless use of AI. The COVID-19 pandemic created a surge in calls to its hotline, and the employees and 200 volunteers who monitored it could barely keep up with the demand. In response, the organization implemented Tessa, an AI-driven chatbot that responded to queries from people suffering from eating disorders, *and* promptly handed out pink slips to hotline staff. Worse still, NEDA failed to adequately supervise or control the chatbot, which began dispensing harmful advice to callers.

Generative AI raises a host of ethical questions and complexities. Large learning language models like ChatGPT are trained on content from the Internet, including sites like Wikipedia, Twitter, and Reddit. Because this data includes overrepresented, white supremacist, misogynistic, and ageist views, the tools trained on it can further **amplify biases and harms**. ChatGPT can also provide false information, or what researchers call “hallucinations,” responding to user prompts or questions with very persuasive but inaccurate text. Another danger is that fundraising staff using any of the public, free models might accidentally enter private donor information into the model, effectively giving other users access to it. There are increasing concerns and **lawsuits about intellectual property** because many generative AI tools have incorporated copyrighted text and artwork from the Internet into their databases without permission. And there are risks related to ethical storytelling. The Toronto-based nonprofit Furniture Bank, for example, switched to AI-generated but realistic-looking images meant to evoke sympathy in its 2022 holiday campaign, **raising many ethical questions** alongside donations.

Adopting ChatGPT also brings reputational risks, especially if people don’t carefully review the text it generates before sharing it; results can be insensitive and offensive. A KFC restaurant in Germany, for example, recently sent out an AI-generated tweet suggesting that people commemorate **Kristallnacht** with a KFC meal. The text translated to: “It’s a memorial day for the Night of Broken Glass! Treat yourself with more tender cheese on your crispy chicken. Now at KFCheese!” Vanderbilt University’s

Peabody College in Tennessee meanwhile deferred to ChatGPT to **write a “heartfelt” letter** to students about coming together as a community to cope with a mass shooting at Michigan State University.

Using AI ethically is not a technical challenge but a leadership imperative, and its adoption must be deeply human-centered. Leaders need to develop guidelines and offer training to ensure that everyone understands what it does, and is careful and thoughtful about its use. This includes not only the people in charge of the technology but all staff. Imagining that staff members aren’t using generative AI because your nonprofit either hasn’t developed guidelines or has forbidden its use is foolish. Recent **research from Fishbowl** shows that 43 percent of working professionals in many fields are using generative AI and that 70 percent of them are doing it without their boss’s knowledge. (Maybe the boss is using it, too!)

8 Steps to Adopting AI Responsibly

We urge organizations to take a variety of steps toward the smart, ethical adoption of AI right now. Rather than a one-and-done checklist, the considerations below provide a framework for ongoing exploration, experimentation, and growth.

1. Be knowledgeable and informed. Nonprofit leaders must understand how the AI tools staff are using or plan to integrate into their work can affect operations and outcomes, and rather than rushing to experiment and failing fast, think through potential risks and rewards before diving in. Developers are pitching AI as a kind of magical fairy dust, but it’s important to be clear on what each tool does and doesn’t do. For example, using a chatbot on the frontline to dispense mental health information or therapy is potentially dangerous, because clinicians are still trying to figure out how to **effectively deploy** them, and **under what conditions**.

2. Address anxiety and fears. Perhaps the greatest fear about AI is that it will take away jobs. In a recent PwC **survey** of 52,000 people across 44 countries and territories, almost a third of respondents were concerned that AI or other technologies would make them redundant within the next 3 years. However, the likely reality, according to the World Economic Forum job report, is that jobs **will change**, not disappear, and that it will happen over time, not immediately. Leaders can mitigate the anxiety staff have about changes to their jobs and work by having open and honest conversations about the use of the technology and the steps the organization is taking to ensure that its use is values-aligned and human-centered. They can also offer guarantees that no employees will be laid off as a result of the organization adopting AI and that AI will augment and support staff work.

3. Stay human-centered. Before adopting AI, nonprofits should create a written pledge explaining that AI will be used only in human-centered ways. It should state that people will always oversee the technology and make final decisions on its use, in ways that don't create or exacerbate biases. As Amy Sample Ward, CEO of the **Nonprofit Technology Enterprise Network** (NTEN), told us, "I think one of the simplest and most important guidelines is that tools should not make decisions. That's been a core part of NTEN's internal approach." Organizations should pay special attention to instances where AI intermediates human interaction, such as collecting client information to **assess mental health conditions or providing counseling services**, where risks to human well-being are greatest. This includes thinking through workflows for co-boting: What will people always oversee, and where and how can bots augment that work?

4. Use data safely. The nonprofit sector can and should raise the bar on the safe and ethical use of data. Leaders must answer yes to questions like: Is our AI tool using written materials by authors who have consented to be included in the data set? And are we giving our subscribers, donors, and other supporters opportunities to be "**forgotten**" as the European Union recommends? *The Smart Nonprofit*, which two of us coauthored, devotes an entire chapter to how leaders can start to answer these questions and understand the privacy issues related to different tools they are thinking of adopting.

5. Mitigate risks and biases. The cybersecurity field uses a process called "threat modeling" to envision what could go wrong when moving an AI-use-case concept from prototype to pilot to full implementation. Nonprofits should deploy a similar risk-based planning approach to ensure that AI use is safe and unbiased, including discussing and identifying worst-case scenarios. The nonprofit **Best Friends Animal Society**, for example, wanted to use a chatbot to help potential cat adopters select cats during its Black Cat Adoption Week campaign. However, during the testing stage, it discovered that it would be easy for the chatbot to repeat racial slurs or inappropriate sexual innuendos. Even with a controlled pilot, it was almost impossible to train the bot on what not to say. In addition, the programmers had favorite cats to recommend, creating bias, so the team ultimately decided to pursue other strategies.

6. Identify the right use cases. Nonprofits should begin by using AI to solve exquisite pain points, bottlenecks, and problems that are stopping other things from happening. Recognizing that it was spending enormous amounts of time doing prospect research, for example, the **National Aquarium** in Baltimore development staff turned to an AI-driven **donor database that fully integrated** wealth screening data, in part by assisting the process of copy-pasting data and research from different sources. Other pain points might include frustration and wasted time searching for documents or fielding the

same questions from constituents. Tasks that are both time-consuming and extremely repetitive are often good candidates for AI automation or augmentation.

7. Piloting the use of AI. Organizations should never use AI at scale without careful consideration and controlled experiments. Leaders should start with a prototype test that's small, time-limited, and assessed by staff, including how the use of AI affected their jobs, their relationship with the organization, and their time. In addition to measuring the impact on staff and external stakeholders, organizations need to carefully and thoroughly check the accuracy of results and determine any bias before rollout. For example, the nonprofit **Talking Points** developed an app that uses AI for translating conversations between teachers and non-English speaking parents but began using it only after extensive user testing to ensure that the translation into different languages was accurate.

8. Job redesign and workplace learning. As AI frees up staff time and allows them to take on other job functions, job descriptions will need updating and some staff will need upskilling so that they can effectively oversee and reap the productivity benefits of tools such as ChatGPT. Organizations should create a shared playbook with best practices that includes guidance on things like how to craft prompts. It should also establish ground rules, such as staff can use AI output only as a first draft, must manually fact-check responses, and not share confidential information with public AI models. At NTEN, staff can use any analysis, machine learning, or other tech systems to gather or provide some perspective on data or information, but they must serve only as one point of analysis. NTEN staff makes all the decisions and takes all the actions. Finally, a healthy workplace requires human connection, so professional skills development must include exercising emotional intelligence, empathy, problem formation, and creativity.

AI, generative and otherwise, will likely find its way into every aspect of organizational life over the next few years. It will fundamentally change what people do in organizations and what technology does. To prepare, senior leaders of nonprofits should begin to create robust ethical and responsible use policies right now to ensure that their organizations stay human-centered and in charge of the technology. Responsible, thoughtful engagement will not only reap productivity benefits and improve the work experience for nonprofit staff, but also help nonprofits serve clients better and attract more donors.

Support *SSIR*'s coverage of cross-sector solutions to global challenges.

Help us further the reach of innovative ideas. **Donate today.**

*Read more stories by **Beth Kanter**, **Allison Fine** & **Philip Deng**.*



Beth Kanter is an internationally recognized thought leader and trainer in digital transformation and well-being in the nonprofit workplace. She is the coauthor of the award-winning *Happy Healthy Nonprofit: Impact without Burnout* and coauthor with Allison Fine of *The Smart Nonprofit*.



Allison Fine is a trailblazing force in the realm of technology for social good. Fine currently serves as the president of Every.org, an organization that enables every nonprofit to accept any kind of donation payment method online without charging platform or direct transaction fees.



Philip Deng is the cofounder and CEO of **Grantable**, a world-class grant-writing assistant for any organization, powered by AI. He devised the prototype from 15 years of experience working with grant-funded organizations in the nonprofit sector on three continents.

DOI: 10.48558/vtxc-gr37

If you like this article enough to print it, be sure to subscribe to SSIR!

Copyright © 2023 Stanford University.

Designed by Arsenal, developed by Hop Studios