

EXPLAINER

A Primer on A.I.

Terms and Concepts
Nonprofits Need to Know

WITH
SUPPORT
FROM



Microsoft

THE CHRONICLE
OF PHILANTHROPY.

Responsible AI for Accelerating Mission Outcomes

It's rare in human experience to know that you were present the moment that a new idea, a new technology, or a new way of doing things changed the world forever. This age of AI is one such moment, and the possibilities it is creating to transform how nonprofits do what they have always done—make the world a better place—are profound and exciting.

The Microsoft Tech for Social Impact team is focused on four areas where we see the greatest potential for AI to deliver concrete benefits for nonprofits of every size and degree of technical expertise:

- **Engaging with donors and supporters:** The ability of AI to analyze, assess, and predict donor behavior will help nonprofit employees be more effective by identifying donors who are most likely to contribute to campaigns and causes, at what level, and at what frequency.
- **Transforming program effectiveness:** AI-powered services including speech and language recognition will enable program participants to find what they need from nonprofits far more quickly and allow caseworkers and volunteers to identify and deliver services that lead to the best outcomes with far greater accuracy.
- **Improving mission outcomes:** From responding to natural disasters and humanitarian crises to addressing climate change, improving the health of communities, and expanding education opportunities for children and adults, AI will usher in a new era of effectiveness and success for nonprofits around the world.
- **Strengthening cybersecurity:** The worldwide rise in cyber threats poses new dangers for us all, but nonprofits can be particularly vulnerable because they often lack the resources to safeguard the data that makes it possible for them to operate. To truly tip the scale in favor of nonprofits, organizations need to be able to detect patterns and behaviors that are not obvious to the human eye and inform their response with the latest and most advanced security practices.

At Microsoft, we have an incredible responsibility to ensure that the AI capabilities we create and deliver are safe, secure, ethical, and serve society broadly and fairly. Our commitment to responsible AI dates back to 2017, when we brought together a team of researchers, engineers, and policy experts to develop a set of [AI principles](#). Since then, we created the Office of Responsible AI to coordinate responsible AI governance and launched our [Responsible AI Standard](#), which provides a clear framework for how we identify and measure the potential for harm and build controls in our systems to mitigate risk from the very outset.

This is just the beginning. We are committed to ensuring that nonprofits are at the forefront of the adoption of AI capabilities that can lead to greater success. To explore how your nonprofit can get started with this technology (and leverage our grants to do so), visit us at microsoft.com/nonprofits.

Sincerely,



Justin Spelhaug

Vice President & Global Head – Tech for Social Impact, Microsoft Philanthropies



Artificial intelligence has been getting lots of attention since the emergence of new, more advanced tools last year. When used well, A.I. has the potential to transform many aspects of nonprofits' operations — from using data for better decision making to streamlining rote tasks — so people can focus more on harder problems and relationship-building.

Yet few charities have embraced A.I., experts say. If your nonprofit is one of the holdouts, you may be struggling to understand the technology, how best to use it, and how to avoid important missteps.

To help you get up to speed, here's a glossary of key A.I. terms and concepts nonprofits should understand when exploring this technology, with smart tips and insights from a variety of experts. This is the first of several resources the *Chronicle* is developing to help you get a better handle on A.I. — and put it to work at your organization in a thoughtful way.

ARTIFICIAL INTELLIGENCE

“A machine that mimics human intelligence,” says Gayle Roberts, chief development officer at Larkin Street Youth Services, a nonprofit that fights youth homelessness in San Francisco. Roberts's team has been using A.I.-powered tools for a couple of years, including [Raise](#), a fundraising platform she says helped the organization net about \$2 million in major gifts last year.

A.I. has been around for decades in various forms, such as Netflix's personalized viewing recommendations, virtual assistants Siri and Alexa, and targeted advertising from retailers. But interest has exploded in recent months with the launch of new generative A.I. tools such as [ChatGPT](#).

“The commonality with whatever type of A.I. you're talking about is that you're always dealing with making predictions based on extremely large data sets and the weights of different variables and algorithms, and with statistics,” says Afua Bruce, founder of ANB Advisory Group, a consultancy that helps organizations use data and technology including A.I.

Many nonprofits have already been using A.I. but don't realize it because they didn't seek out those tools specifically or turn them on, says Amy Sample Ward, CEO of NTEN, an organization that trains nonprofits on the equitable use of technology. When email and messaging systems suggest quick replies or when donor databases recommend supporters to contact, they are using A.I., Sample Ward says.

It's best to assume A.I. was built into your technology tools, Sample Ward adds. “Look into your settings, figure out what you can and can't turn off, how you can turn it off, and how you want to use it,” because A.I. features are often enabled by default, and they “might not be what we actually want them to be.”

MACHINE LEARNING

A form of A.I. that can learn from data and perform specific tasks but doesn't simulate human intelligence. Sometimes the terms “machine learning” and “A.I.” are used interchangeably, but they aren't quite the same. It may be helpful to think of A.I. as the umbrella concept and machine learning as an approach or application under it.

Roberts points to Google's page-ranking system as an example of machine learning. The technology uses data to provide search results but doesn't mimic human intelligence to solve a more complex problem.

For most nonprofit workers, the details about how these terms differ don't really matter in practice, Bruce says. "The concerns are the same ethically about how bias is built into different data sets and things like that."

GENERATIVE A.I.

A new technology that analyzes large data sets to create original content — such as text, audio, images, or video — based on a human's prompt. In other words, generative A.I. algorithms predict what they think you're looking for and provide the best matches based on statistics, Bruce says. These tools may scan the internet for examples or use data fed to them.

Examples of generative A.I. include OpenAI's [ChatGPT](#) and Microsoft's [Bing Chat](#), text-to-text applications that can write text and answer questions; [DALL-E 2](#), which creates images and art; and Meta's [Make-a-Video](#), which produces videos from text or static images.

CHATBOT

An A.I. tool that uses data to respond to a human's prompt through a chat. Not all chatbots, also known simply as bots, are generative A.I. Most people have been using bots powered by more traditional A.I. for years, experts say.

For example, if you asked questions on a cable company's website a few years ago, a bot probably responded, Bruce says. Unlike generative A.I., those kinds of less advanced bots work by "listening for key words and giving key words back," she says, rather than analyzing information to make predictions and create something that didn't exist before.

People often think of chatbots as being conversational, but you could

also use them simply to offer a different way of navigating your nonprofit's website, Sample Ward says. "It doesn't need to be something that's making recommendations or telling people what to do or making decisions." Instead, a bot could just serve as a tool that people ask how to find something on your site; it would answer with the URL to that page.

CHATGPT

A bot that analyzes data scanned from the internet in 2021 to write original responses to prompts. The A.I. company OpenAI launched this tool in late November 2022. The newest and most powerful iteration, [GPT-4](#), is available through a paid subscription, while the older version is free.

"Anybody doing content development or editing should be using [GPT-4] yesterday," Roberts says, though it's important to create thoughtful policies for how employees can use this technology. There are nearly unlimited possibilities, she says, such as helping draft or edit grant proposals.

BING CHAT

A free bot developed by Microsoft that scrapes data from the internet to respond to prompts. Because GPT-4's knowledge is limited to information that was on the internet in 2021, Bing is more useful for helping with online research, she says, while GPT-4 is better as an assistant for creating and editing content.

HALLUCINATIONS

False information or lies that generative A.I. fabricates. "These systems are very eager to give you an answer," Roberts says. "They make stuff up a lot."

To avoid this problem, use ChatGPT as an editorial assistant, rather than a writer, she suggests. “Start with original content and use it to generate revisions or research off of that,” she adds. “That starts you from a place of actual fact so it doesn’t result in hallucinations.”

Keep in mind that these tools may also make up citations they reference, Sample Ward says, so be sure to check for accuracy when using them for research.

LARGE LANGUAGE MODEL

“The technical term for a lot of text data,” Bruce says, which any generative A.I. tool needs to operate. “ChatGPT’s scan of the internet was building a large language model,” she says.

A.I. ALGORITHM

“The engine of A.I. is an algorithm,” says Sarah Di Troia, senior strategic adviser of product innovation at Project Evident, a consultancy that helps organizations use data to measure and improve their impact. “That’s what’s doing the work. And the data is the fuel — it makes it go.”

You could also think of an algorithm as the rules an A.I. tool follows, as Sample Ward puts it. But you don’t need to have a perfect understanding of this or any other A.I. terms, Sample Ward adds. It’s more important to understand what they mean for your organization specifically. For example, donor databases often have A.I. algorithms turned on by default, such as making outreach recommendations based on predictions of supporters’ behavior. “Unless you’ve told it how to predict those donors, it’s choosing how to [do so],” Sample Ward says. “And that’s likely built on bias in whatever data sets it was using.”

Even if you don’t have a strong grasp of the technical language, you should feel free to ask vendors about the algorithms their products use so you can make sure they align with your nonprofit’s values — and turn off any that don’t. “You don’t need to know literally the code of the algorithm to say, ‘Hey, this isn’t how we want to be having our data used’ or whatever else,” Sample Ward says.

- **Black Box:** If you can’t test an algorithm or review the data it uses, that algorithm is called “black box,” Di Troia says. Most commercial algorithms fall into this category. So, ask questions about the vendor’s values and its algorithm, including how it was built, what data it was trained on, how the algorithm and data were tested for bias, and whether you can do testing.

- **Open Source:** An algorithm that you can test and train on your own data, Di Troia says. These tools are typically free. [GitHub](#) is an example of a platform where you can find open-source algorithms.

PREDICTIVE ANALYTICS

A system that uses A.I. to make predictions based on data, such as analyzing previous giving behavior to forecast donations. Predictive analytics has “some of the greatest potential” for nonprofits, Di Troia says, because it can help them segment people they interact with — such as program participants or donors — more easily and effectively. That enables nonprofits to offer a level of customization that retailers have been using with consumers for years, she says.

For example, Netflix has more than 2,000 taste profiles for users, Di Troia says, which the company uses to provide tailored home screens and

movie recommendations. Predictive analytics makes that possible, she says, by identifying segments that are “in some ways, more complex than you can do with the naked eye.”

BIAS IN A.I.

Bias can show up in the coding used to program an A.I. algorithm or in the data that fuels it, Di Troia says. And sometimes bias is built into an algorithm indirectly. For example, Di Troia knows of a nonprofit that used an algorithm to evaluate microloan applications, and the tool was weighting nonretail experience as more valuable than retail experience. “It didn’t say, ‘Don’t give microloans to women,’ but women are more likely to have retail work experience than other types of experience that were more

highly favored in this algorithmic model for lending,” she says.

Di Troia also offers an example of how data itself can be biased. If you feed an algorithm data about a population and that information is underrepresentative of certain groups, she says, the responses the tool provides will reflect the majority.

You should assume any data set is biased, says Sample Ward, who adds: “It’s not about finding this clean, pure, beautiful, perfect data set out in the world that doesn’t have bias. It’s being able to understand which biases are in there so that you can intentionally address them.”

For example, if you know a data set is limited by time or demographics, you could look for more information from other sources to fill in gaps.