nten

Generative AI and the Social Sector

POLICY NARRATIVE 2024

AN NTEN RESOURCE
by Nolwenn Godard &
Lindsey Washburn

What's Inside

Generative AI use in the social sector

In today's rapidly evolving digital landscape, generative artificial intelligence presents both transformative opportunities and significant challenges for all organizations, regardless of mission, size, budget, or other factors. As we embrace the potential of these powerful tools, it's crucial to establish a comprehensive policy that ensures their ethical, responsible, and effective use.

Social sector organizations are uniquely defined and driven by their mission and commitment to their stakeholders. GenAI has the potential to enhance operational efficiency and amplify impact, but it also poses risks that must be carefully managed. These include the potential misuse of personal identifiable information or reidentification of individuals, propagation of biases embedded within AI models, inadvertent harm that could arise from automated processes and decision-making, and increased security issues.

Whether the organization uses GenAI tools or not, it's highly probable that its employees do, sharing the organization's data or context on private GenAI tools. This can inadvertently introduce significant risks to the organization and its stakeholders. Organizations should develop and implement policies on GenAI, even if they do not directly utilize these tools. Employees should be educated on the responsible use of GenAI to uphold the organization's values and protect its intellectual property.

Organizational policies on GenAI establish standardized practices, mitigate risks, prevent unauthorized use, and guide incident response.

EXAMPLE OF UNAUTHORIZED DATA SHARING

An employee uses a third-party GenAI tool to draft a report and unintentionally uploads confidential client information. Without proper training and policies, this could lead to a data breach, damaging client trust and leading to potential legal repercussions.

While organizations may face different risks from GenAI, there are general principles and concerns that apply universally. An organization that provides services to 2SLGBTQIA+ youth must be vigilant about safeguarding sensitive information, while an animal rights group might focus on ensuring the integrity of data used in advocacy efforts. Despite these differences, the core issues of privacy, bias, transparency, and accountability remain consistent across all sectors.

If it's not carefully managed, GenAI could undermine the very mission of the organization. For example, a GenAI content creation tool might inadvertently generate biased or harmful narratives based on flawed training data, potentially damaging the reputation and credibility of the organization. Additionally, the mishandling of personally identifiable information could lead to breaches of trust, legal ramifications, and significant harm to the individuals we aim to support.

It's critical to make informed choices about the adoption and implementation of GenAI technologies. This involves understanding both the potential benefits and risks.

A GenAI Use Policy serves as a guiding framework for the responsible use of GenAI in the social sector. It emphasizes our collective responsibility to harness these technologies in ways that align with our ethical standards and missions.

We invite you to engage with this policy by tailoring it to your organization's mission, vision, and values. Here are some resources you can use to get started.

- **GenAI Use Policy template** for nonprofit organizations
- **Data Empowerment resources** for policies related to managing data collection and security

RISK SCENARIOS

- ✖ A nonprofit uses GenAI to analyze donor data and tailor fundraising strategies. While this can lead to more effective campaigns, it also necessitates stringent measures to protect donor privacy and ensure that AI-generated recommendations are free from bias.
- ✖ A foundation uses GenAI to expedite the review of grant evaluations. If the organization doesn't understand how the tool works or evaluate the answers produced, the foundation may inadvertently disadvantage certain applicants or perpetuate systemic inequalities.
- ✖ A youth action group uses GenAI flippantly because they lack policy guidance on appropriate use. The organization is unaware of the climate impact of GenAI (i.e. energy and water consumption) and wastes resources without proper guidance and evaluation of GenAI tools.

Join us in our commitment to ethical and responsible AI practices. We recognize that this technology changes quickly and we must also commit to continuously educating ourselves on the harms and benefits. This also means reviewing our policies and contracts with third-party vendors on a regular basis.

Remember: GenAI is a tool and it may not be the best tool for your objectives or organization more broadly. Encourage your organization to think critically about this technology and the results produced.

What makes GenAI different?

GenAI is distinct from previous forms of AI and requires a different approach for social sector organizations in implementing and evaluating this technology. GenAI builds on advances in conventional AI and uses large quantities of data to output unique written, audio, and/or visual content in response to freeform text requests from its users and programmers.

GenAI tools have the capacity to produce entirely new content instead of simply regurgitating inputted data. Unlike conventional AI systems designed for specific tasks, GenAI models are designed to be flexible and multifunctional.

GenAI technology operates using foundation models, which are large-scale machine learning models with versatile capabilities. These models are trained on vast datasets that span the entire internet, serving as the basis for applications that address specific business, policy, or social needs. Developing foundation models requires significantly more computing power and human resources than conventional AI development.

GenAI models utilize human-generated content as part of their training data, allowing them to respond to free-text queries with human-like output. However, despite their ability to produce coherent and natural-sounding responses, there is no guarantee that the output is accurate. Many widely available GenAI models were designed to showcase their potential rather than solve specific problems, leading to possible inaccuracies, fabrications, inappropriate content, and biases.

These models demonstrate the unprecedented power of GenAI. Enterprise models continue to improve, closely mimicking human writing, drawing, and speaking. However, the rapid development and widespread availability of GenAI have also accelerated policy, business, and social risks, making them more urgent than those posed by previous AI technologies.

GenAI products are already available as standalone applications and are being integrated into many other consumer-facing technology products, such as chatbots on websites. These tools offer unprecedented capabilities to create engaging content, automate routine tasks, and streamline communications.

Conventional AI models, on the other hand, are usually designed for specific tasks and are often limited by the scope of the inputted training data as well as the

technical expertise of the programmer. Model training is the process by which AI models ingest input datasets to learn the underlying patterns within the data and produce predictions for the context that the model was trained on.

Conventional AI is already widely used in products across various sectors. Some examples of conventional AI include:

- **Robotic process automation:** Automating routine administrative tasks, such as data entry and processing, to improve efficiency.
- **Fraud detection tools:** Identifying and preventing fraudulent activities in financial transactions and other operations.
- **Image classification systems:** Categorizing and tagging images for better organization and accessibility.
- **Recommendation engines:** Personalizing content and service recommendations to enhance user experience.
- **Interactive voice assistants:** Providing voice-activated assistance and information retrieval for users.

While conventional AI has significantly contributed to operational efficiencies and improved service delivery, GenAI offers a new level of adaptability and creativity. For the social sector, this means the ability to generate tailored content, develop innovative fundraising campaigns, and engage more effectively with diverse audiences. However, it also necessitates a thoughtful approach to ensure these tools are used ethically and responsibly, safeguarding against biases and unintended consequences.

Risks with GenAI are also different than with previous versions of AI:

- **Lower barriers:** GenAI is easier and cheaper to use.
- **Faster impact:** GenAI tools work quickly and on a larger scale.
- **Wider reach:** More systems and processes are affected by GenAI.
- **Increased exposure:** Larger, diverse training datasets make GenAI more vulnerable to bad actors.
- **Higher complexity:** GenAI's complex technology involves many producers and consumers.
- **New risks:** These are novel risks arising from GenAI's unique ability to generate high-quality outputs in various forms, such as text, images, audio, and video – also referred to as "deepfakes".

GENAI CAN SUMMARIZE VAST AMOUNTS OF DATA QUICKLY, IDENTIFYING KEY TRENDS AND INSIGHTS THAT MIGHT BE MISSED THROUGH MANUAL ANALYSIS.

By understanding the distinct capabilities and risks of GenAI, social sector leaders can make informed decisions about using GenAI, better harness its potential to advance their missions and amplify their impact while maintaining their commitment to ethical practices and social responsibility.

New and amplified risks of GenAI tools for the social sector

As generative artificial intelligence tools become increasingly integrated into the operations of nonprofits and foundations, it is essential to understand the new and amplified risks these technologies present.

NEW RISKS

1. **Misinformation and disinformation:** GenAI can generate content that appears credible but is factually incorrect or misleading. This can lead to the spread of misinformation or disinformation, which can harm public perception and trust in social sector organizations.
2. **Privacy violations:** The use of GenAI tools in processing and analyzing personal data can increase the risk of privacy breaches. Sensitive information about beneficiaries, donors, or staff may be inadvertently exposed or misused.
3. **Unintended consequences in automated decision-making:** AI-driven decision-making processes can lead to unintended and potentially harmful consequences. For example, an automated system might deny services or resources to individuals based on flawed criteria or biases embedded in the algorithm.
4. **Intellectual property issues:** GenAI can create new content, such as text, images, or music, which may inadvertently infringe on existing intellectual property rights. Social sector organizations must navigate the complexities of copyright and ownership in AI-generated outputs.
5. **Inadvertent breach of contracts:** If third-party information (from customers, vendors, or business partners) is submitted to a GenAI tool as input, it is likely unauthorized by the information owner and could constitute a breach of contract. Most agreements restrict the use of such information to the provision of services or contract performance. Using it as input for a GenAI tool often exceeds these limits. Additionally, input data might be used to train the AI models, posing further risks. Some GenAI tools allow opting out of using input data for training, but not all. Even with enterprise GenAI tools, careful consideration and due diligence are necessary when sharing information across departments or agencies.

AMPLIFIED RISKS

1. **Bias and discrimination:** As humans have explicit and implicit biases built into our society, GenAI has the capacity to perpetuate and even amplify existing biases present in training data. This can result in discriminatory outcomes, such as biased hiring practices, inequitable service delivery, or skewed research findings that fail to represent diverse populations accurately.

DEVELOP AND ADHERE TO ETHICAL GUIDELINES FOR AI USE THAT ALIGN WITH THE ORGANIZATION'S MISSION AND VALUES, INCLUDING PROCUREMENT, USAGE, AND REQUIRED TRAINING FOR STAFF.

2. **Reputational damage:** The misuse or failure of GenAI tools can lead to significant reputational damage. For instance, a chatbot providing insensitive responses to marginalized individuals or an AI-generated report containing errors can erode trust and credibility.

3. **Security vulnerabilities:** The complexity of GenAI systems can introduce new security vulnerabilities. These tools may be susceptible to hacking, data manipulation, or other cyber threats, posing risks to the integrity and confidentiality of organizational data.

4. **Dependency and loss of human expertise:** Overreliance on GenAI tools can lead to a diminished emphasis on human expertise and judgment. This dependency may weaken the capacity for critical thinking and reduce the overall resilience of the organization in addressing complex issues.

MITIGATING THE RISKS

To mitigate these risks, social sector organizations should adopt a proactive and comprehensive approach, including:

1. **Robust data governance:** Implement strict data governance policies to protect sensitive information and ensure compliance with privacy regulations. Regularly audit data practices and AI systems to identify and address potential vulnerabilities.

2. **Bias mitigation strategies**

    a. Use diverse and representative datasets to train AI models. Employ bias detection and correction techniques (such as generating new, unbiased training data) and involve diverse stakeholders in the development and evaluation of AI systems.

    b. Analyze impact on vulnerable communities:

        i. Develop guidelines to analyze the impact that adopting GenAI tools may have on vulnerable communities, ensuring that these tools do not exacerbate existing inequalities.

ii. Test and audit products and GenAI outputs specifically for impact on individuals of different regions, income levels, races, ethnicities, genders, ages, religions, abilities, and sexual orientations.

3. **Transparency and accountability:** Ensure transparency in AI decision-making processes by documenting how AI tools are used and making this information accessible to stakeholders. Establish clear accountability mechanisms to address any adverse outcomes promptly.

4. **Ethical AI frameworks:** Develop and adhere to ethical guidelines for AI use that align with the organization's mission and values, including procurement, usage, and required training for staff. Regularly review and update these guidelines to reflect evolving best practices and societal expectations.

5. **Conducting risk analysis:** Analyze potential threats and vulnerabilities arising with the use of GenAI, especially focusing on sensitive areas like data privacy and community trust.

6. **Evaluating regulatory impacts:** Continuously evaluate the potential impact of GenAI on regulatory issues, ensuring compliance with existing laws and advocating for policies that protect the interests of the communities served.

7. **Continuous education and training of the workforce:** Invest in ongoing education and training for staff and volunteers to build AI literacy and ensure they are equipped to use AI tools responsibly and effectively, ensuring they understand both the capabilities and limitations of these technologies.

GENAI OFFERS A NEW LEVEL OF ADAPTABILITY AND CREATIVITY. FOR THE SOCIAL SECTOR, THIS MEANS THE ABILITY TO GENERATE TAILORED CONTENT, DEVELOP INNOVATIVE FUNDRAISING CAMPAIGNS, AND ENGAGE MORE EFFECTIVELY WITH DIVERSE AUDIENCES.

By understanding and addressing these new and amplified risks, social sector organizations can harness the potential of GenAI while upholding their values, mission and safeguarding the communities they serve. This balanced approach will enable the social sector to leverage AI's capabilities responsibly, driving innovation and impact while maintaining trust and integrity.

Potential benefits of GenAI usage

When used ethically and transparently, GenAI has the potential to improve access to and utilization of social sector programs, as well as their services and their operations (such as personalized service delivery, enhanced data analysis, and efficient content generation). It can be used to empower the communities it serves.

By understanding these themes, social sector leaders can better navigate the complexities of GenAI, leveraging its advantages while mitigating its risks to advance their missions and enhance their impact.

1. **Summarization and classification:** GenAI can analyze hundreds of millions of data points simultaneously, creating comprehensive summaries of any collection of artifacts, whether they are in text, audio, or video format. As GenAI learns, it can also categorize and classify information by topic, format, tone, or theme. It can thus improve the performance, capacity, and efficiency of ongoing work, research, and analysis. Example use cases include:

    a. **Sentiment analysis:** GenAI can analyze text data from surveys, social media, and other sources to gauge public sentiment. This helps nonprofits understand how their initiatives are perceived and identify areas for improvement. Use GenAI to recommend improvements for processes and service delivery. This can help nonprofits understand public experience and improve their programs and communication to better serve their communities.

    b. **Data integration and summarization:** GenAI can integrate data from various sources, providing a comprehensive view of the information available. It can summarize vast amounts of data quickly, identifying key trends and insights that might be missed through manual analysis. For example, it can summarize meetings, research, and public outreach documentation, leveraging GenAI to find key topics, conclusions, action items, and insights without needing to read everything word for word. This saves time and ensures critical information is highlighted.

    c. **Pattern recognition:** By analyzing data, GenAI can identify patterns and correlations that are not immediately apparent. This can help nonprofits understand complex relationships between different variables and factors affecting their operations and impact. For example, GenAI identifies patterns in donor behavior, such as the types of campaigns

that attract the most donations, the demographic characteristics of repeat donors, and the optimal times for sending fundraising emails.

d. **Predictive analytics:** GenAI can use historical data to predict future trends and outcomes. This can assist nonprofits in forecasting demand for services, identifying potential challenges, and planning strategically for the future. For example, using historical donation data, GenAI predicts future donation trends and identifies potential high-value donors. It also forecasts the potential impact of upcoming campaigns based on past performance.

SOME GENAI TOOLS ALLOW OPTING OUT OF USING INPUT DATA FOR TRAINING, BUT NOT ALL.

e. **Resource optimization:** By analyzing data on resource usage, GenAI can recommend ways to optimize the allocation of funds, staff, and other resources, ensuring that the nonprofit operates as efficiently as possible. For example, based on the analysis, GenAI recommends strategies for optimizing campaign efforts, such as focusing on specific donor segments, adjusting campaign timing, and enhancing engagement through personalized communication.

f. **Custom queries:** GenAI can answer specific business questions by processing relevant data and generating detailed, context-specific responses. Nonprofit staff can ask GenAI questions in plain language and receive comprehensive answers supported by data analysis. For example, the nonprofit staff can ask specific questions, such as "Which demographic group is most likely to respond to our email campaigns?" or "What factors contribute to higher donation amounts?" GenAI can provide detailed answers backed by data.

2. **Personalization and diversification of outputs:** GenAI's capacity to learn makes it easier for nonprofits to design services and products that are responsive to the diverse needs of their communities, across different geographies and demographics. GenAI solutions can recommend ways to display complex information in a way that resonates with various audiences or highlight relevant information from multiple sources tailored to individual needs. Example use cases include:

   a. **Targeted outreach:** Apply GenAI to service data to identify specific groups or subsets of participants who may benefit from additional outreach, support services, and resources based on their circumstances and needs.

3. **Communication, translation and accessibility:** GenAI can create unique content in a variety of formats. Based on a single prompt, a GenAI solution can easily construct a video or image that a user can refine.

These products can be in multiple languages, allowing nonprofits to make their videos, recordings, and other documents more accessible and inclusive. These translated outputs can be refined through a quality control process to ensure accuracy and inclusivity before reaching the intended audience.

Accessible communications are a critical part of ensuring that nonprofit services can meet people where they are. The ability to meet the varying communication needs of persons with disabilities and reach individuals in their primary languages is a priority for improving service delivery. Example use cases include:

a. **Accessible educational materials:** Use GenAI to help experts convert educational materials into formats like audiobooks, large print text, or braille documents. GenAI can also generate captions for video materials, making information more accessible for those with visual, hearing, or learning disabilities.

b. **Multilingual translations:** Leveraging GenAI to help experts translate websites, public documents, policies, forms, and other materials into various languages spoken within the community. This expands access to important information and services for non-native English speakers.

4. **Software coding:** GenAI's summarization, classification, and translation features make it a powerful tool for coders and the developer community at large. GenAI can generate code in multiple computing languages and translate code from one language to another. This can improve operations if an organization's system is using code that is written in an obsolete language. Moreover, GenAI has the potential to explain and categorize unfamiliar or uncertain code so that the organization can better understand the exact technical architecture of their applications.

THE MISUSE OR FAILURE OF GENAI TOOLS CAN LEAD TO SIGNIFICANT REPUTATIONAL DAMAGE.

5. **Data analysis, insights and forecasting:** Without specific training or pre-set rules, GenAI models can analyze multiple datasets to find meaningful insights for users. The conversational aspects of GenAI solutions can empower workers with a range of technical expertise to ask questions in plain language to uncover findings relevant to their work. Significantly, organizations can also use GenAI solutions to ask data-driven questions that are important to their missions. Example use cases include:

a. **Cyber protection systems:** GenAI-powered systems can rapidly analyze network activity logs, identify anomalies and threats, generate explanations of attacks, and propose remediation actions. This enables security teams to detect and respond to sophisticated cyberattacks in real-time before major damage occurs.

6. **Environmental sustainability:** Incorporating GenAI into operations can enhance organizational environmental sustainability measures. Use GenAI to analyze and optimize the allocation of resources, such as office supplies, transportation, and energy use, ensuring that operations are conducted as sustainably as possible.

   However, we must note that AI technology systems require a large amount of resources including electricity and water. As users and as a sector we can advocate that AI technology providers invest in mitigations and efficiencies to address environmental impacts of the servers and infrastructure needed to run their products.

By leveraging these advanced capabilities, social sector organizations can enhance their operational efficiency and effectiveness, ensure sustainable practices, and make data-driven decisions that align with their mission, values, and the diverse needs of the communities they serve. This balanced approach enables social sector organizations to maximize their impact while maintaining a commitment to ethical and environmentally responsible practices, promoting equity, inclusion, and accessibility.