

2026

The Risks and Rewards of Artificial Intelligence: A Guide for Nonprofits

Melaina Lenzi

University of Washington Tacoma

Abstract

The following research paper aims to provide a preliminary guide for nonprofits considering AI implementation in their service delivery. The paper evaluates emerging lab and field research studies to assess the potential positive and negative impacts of AI use on various business segments. This research studies investigated AI use on labor productivity, accounting automation, donor segmentation, and data privacy. Emerging results indicate nonprofits may experience greater labor productivity gains from AI implementation, as well as improved accounting accuracy and donor management. However, AI use comes with the risk of AI “hallucinations,” data privacy violations, amplifying biases, and large-scale issues, such as negative environmental impacts and asymmetrical global development. This paper also presents emerging frameworks for AI governance (as of March 2026), including the EU AI Act, the California Consumer Privacy Act, and the National Institute of Standards and Technology’s AI Risk Management. Given the research and the high-stakes nature of nonprofit work, this paper recommends the following: 1) human oversight is crucial for mitigating AI risks; 2) protection for sensitive and private data should be a priority at every level and should be anonymized or redacted if using open-sourced AI models; in fact, 3) nonprofits should to opt for enterprise AI models that offer greater data protection to protect their donors and program participants; 4) any creative content generated by AI is at risk of copyright infringement; therefore, nonprofits may want to carefully consider how and when they are using AI-generated content for front-facing materials until greater defined laws are established; and 5) AI risk management includes incident reporting and regular audits to monitor for AI errors.

Keywords: AI, Nonprofits, AI Benefits, AI Risks, Ethics, AI Implementation, AI Governance, Donor Segmentation

Introduction

AI models offer nonprofits the opportunity to increase productivity and improve their resource management and service delivery. However, nonprofits face risks when using AI, including amplifying discrimination and biases, compromising data privacy and security, compliance complexities and liability concerns, and inaccurate information. This paper investigates research on AI use in the workplace, evaluates the findings for a nonprofit setting, presents methods for AI risk mitigation, and suggests policies and procedures for nonprofit AI adoption. This paper aims to act as a decision guide for nonprofits in the first wave of AI implementation. While conducting the literature review, it became apparent that research on AI in the workplace (and specifically for nonprofits) was an emerging topic for study at the time of composition. Therefore, the findings below extrapolate on available research and apply it to a nonprofit context. The research here shows AI use may offer nonprofits increased productivity, that is, in terms of time saved accomplishing routine tasks, automated accounting, and improved donor management, potentially resulting in better resource allocation. However, this paper warns that AI comes with potential risks that may negatively impact nonprofit service delivery, as well as broader societal implications. These risks may influence some nonprofits' decision of adopting the technology early. This paper presents the research so nonprofit leadership can make the best decision for their operational needs, recommends supervised AI systems for nonprofits, and highlights areas where nonprofits may need to exercise caution when using AI.

Terminology

The field of AI falls under the computer science discipline focused on creating systems that can perform human-like intelligence, most commonly learning, reasoning, and problem-solving. Toloka, a Dutch multinational data service company specializing in generative AI

development and large language models (LLMs), explains the difference between AI, machine learning, LLMs, and generative AI. Machine learning (ML) is a subset of AI focused on algorithms that can learn the patterns of training data and make inferences about new data. As such, these systems can use pattern recognition to perform tasks without being explicitly programmed for each task (Toloka Team, 2026). Machine learning enables AI tools, such as, but not limited to, AI chatbots, task automations, or AI agents, to improve over time. LLMs are a type of deep learning model trained on massive amounts of text data to process and generate human language. Common examples are Open AI's ChatGPT, Microsoft's Copilot, and Google's Gemini. Lastly, generative AI (genAI) are AI systems capable of creating new content such as text, images, code, audio, or video (Toloka Team, 2026). LLMs are a type of genAI. For nonprofits, ML systems could streamline processes, manage repetitive tasks, and increase operational impact.

1. AI Benefits for Nonprofits: Promise & Potential

Nonprofits often employ small teams and face challenges to meet service needs. Recent years have presented increased uncertainty into the nonprofit sector. In Sage's 2024 Nonprofit Technology Trends Report, 56% of surveyed nonprofits reported revenue remained the same or lower from the previous year (Miller, 2024). Continued economic uncertainty and higher competition for funding has caused many nonprofits to face staffing shortages and donor management challenges. Emerging AI technology offers nonprofits the opportunity to improve resource allocation, decision-making, and efficiency. AI machine learning, writing tools, and automation can streamline workflow, such as decreasing time spent writing reports, press releases, or emails, to create more time for strategic and creative thinking.

1.1 Increased Productivity

AI consultants for nonprofits like Sigma Forces or Daxko Operations claim generative AI and AI writing tools have large, untapped potential for nonprofits and can increase productivity, often measured as time saved completing tasks. Daxko claims using LLMs and AI writing tools to support community outreach and grant writing will drastically cut down correspondence time (Arora, 2024), while Sigma Forces claimed AI automation saves 15-20 hours on administrative tasks per week (Valenti, 2025). While some of these claims have been supported by industry reports, academic research shows varied results from generative AI use on time saved. To evaluate the potential of AI for nonprofits on increased productivity, this paper evaluates three different research studies.

In 2023, MIT researchers conducted a randomized control trial in a pre-registered online experiment, where 453 college-educated professionals were assigned occupation-specific, incentivized writing tasks to test the productivity effects of generative artificial intelligence technology, specifically ChatGPT (GPT-3.5 model). After being randomly divided in half, where the treatment group used ChatGPT and the control group did not, participants were asked to write press releases, short reports, analysis plans, and delicate emails (Noy & Zhang, 2023). Productivity was measured by the time taken to complete a task and the quality of the work, evaluated by blinded experienced professionals working in the same occupation as participants (Noy & Zhang, 2023). Evaluators were incentivized to grade outputs carefully, and each piece of output was seen by three evaluators (Noy & Zhang, 2023). Researchers found, “in the treatment group, time taken on the posttreatment task dropped by 11 min relative to the control group, who took an average of 27 min ($P < 0.001$). Average evaluator grades increased by 0.45 SDs ($P < 0.001$)” (Noy & Zhang, 2023, “Productivity”). Their results show that ChatGPT raised

productivity, where the average time taken to complete tasks increased by 40% and output quality rose by 18% (Noy & Zhang, 2023).¹

In a similar study found in a working paper, researchers at the Federal Reserve Bank of St. Louis tested AI's impact on work productivity to evaluate how much work is assisted by generative AI.² Researchers ran a survey-based study, asking respondents how much they used AI technology in the previous week or previous month to measure the intensity of usage: days per week and time per day when used (Bick et al., 2025). Survey responses were gathered in August 2024 and again in November 2024 to compare results. Bick et al. (2025) asked workers from the November survey who had used genAI in the previous week to estimate how much time genAI saved them that week (p. 21). Researchers found genAI saved 34.2% of workers an hour or less, 25.4% reported two hours, 19.9% reported 3 hours, and 20.4% reported four hours or more. To measure productivity, researchers computed the share of working hours saved as the ratio of time saved last week to hours worked last week for each AI user. They found an average time savings of 5.2% of work hours, or about 2.1 hours for a 40-weekly-hour worker. When computed for all workers, including non-users, genAI, on average, saved workers 1.4% of total hours or about half an hour. These findings indicate there may be time-saving benefits to generative AI use in the workplace. However, Bick et al. (2025) also found that the time saving effect from genAI use varies across occupations. Nonprofit organizations (NPOs) were not

¹ This research was found in a working paper. This is likely due to the emerging nature of AI research in the workplace at the time of composition. Nonetheless, this author wants to ensure transparency on the research for all reading this paper and those using it as a guide on AI adoption. This author encourages further research into AI and its effects on the workplace.

² This study lists the most reported generative AI technology used by survey respondents were ChatGPT (32% of all respondents), Gemini (19%), and Microsoft Copilot (15%) (Bick et al., 2025).

explicitly represented in the study, but the findings may, nonetheless, be useful for nonprofit consideration.³

Figure 1. The Impact of Generative AI on Work Productivity by Career Sector




Figure 1 shows the share of workhours spent using AI to share of the time saved of hours worked. *Source: The Federal Reserve Bank of St. Louis. (2025) The Impact of Generative AI on Work Productivity.*

Above, the graph plots different occupations' shares of previous week's hours spent using AI against the time saving shares of hours worked in previous weeks to represent which occupations

³ Figure 1 is included in this paper because nonprofits often have departments or roles within the organization that are represented in the study and seen in Figure 1. For example, NPOs oftentimes have a business and finance department that handles accounting, budgeting, and payroll. Nonprofits also have administration or management positions where AI could be helpful in decreasing the time it takes to complete routine tasks, as discussed in the study. Additionally, some nonprofits may be specifically aimed at the other sectors listed, such as arts, entertainment, and sports; health practitioners; or education. For these reasons, this paper included this study for nonprofit consideration.

get the most time saved from AI use. The graph indicates that generative AI may be more useful and increase productivity for nonprofits in their management or finance departments.

In another recent working paper, Brynjolfsson et al. (2023) used Difference in Differences (DiD) regression testing, a statistical technique used to estimate the causal effect of a treatment by comparing the changes in outcomes over time between a treatment group and control group, to test the effect of generative AI on productivity across different experience levels (p.13). Researchers divided customer service agents into two groups: agents who were never given access to the AI tool during the sample period (control group), and agents who were eventually given access and observed before and after intervention (treatment group). Group two was divided into pre-AI observations, that is, before agents were given access to the treatment, and post-AI observations, after treatment, to test whether AI access increased productivity, measured as customer service issues resolved per hour (Brynjolfsson et al., 2023). The study found access to AI tools increases productivity for customer support agents by 14%, including a 34% improvement for novice and low-skilled workers, but with minimal impact on experienced and highly skilled workers (Brynjolfsson et al., 2023, p. 2). These research findings may not generalize across all firms and production processes; however, nonprofits could implement these findings into their operations by integrating ML systems to teach lower skilled, part-time, or volunteer participants tacit donor and program knowledge from skilled, long-term employees. For example, long-term employees and nonprofit administrators could input their successful donor, fundraising, or program support interactions into a ML system as training data. Then, new employees with less internal knowledge of the organization could be given access to that ML system to provide more consistent service delivery.

Overall, AI adoption and increased productivity vary across occupation and employee experience. While quantitative research is emerging, as seen with the working paper above, it appears that evidence supporting nonprofit AI adoption or increased productivity claims are from a combination of survey-based research, industry reports, or AI consultant firms. Industry reports and consultant firms may have higher internal biases for supporting AI adoption. Survey-based research has its role, but findings can vary as the responses are subjective. Therefore, it appears nonprofits may see positive results from partnering with AI to save time on routine tasks, but to a lesser degree than AI consultants may claim.

1.2 Accounting Automation

AI automation may offer more support for nonprofits, specifically in their financial operations. Using AI to streamline accounting operations could improve accounting efficiency, reduce errors, and optimize resource allocation. According to Mohammad et al. (2020), AI in accounting may increase speed, maintain accuracy, and reduce material misstatements in accounting, indicating AI may be more useful and reliable in this field. For NPOs facing limited resources, AI automation in accounting may offer overall lower operating costs (Amneh et al., 2020, p.484). Similarly, Azman et al. (2021) tested AI automation and ML systems' ability to identify and correctly classify invoices from scans of electronic, hand-written, and machine-printed receipts. They found AI automation out-performed human accountants with upwards of 90% accuracy (Azman et al., 2021, p. 228). In fact, AI automation and ML systems in bookkeeping management made the system more reliable, improved bookkeeping systems for small and medium enterprises, and simplified the initial processes of in-house financing in every business sector (Azman et al., 2021, p. 229). This indicates nonprofits can optimize AI automation and ML systems in their bookkeeping process to streamline routine tasks like data

entry, automate invoice processing, and reduce accounting errors for overall increased operation and lower costs. The potential higher productivity and lower cost gains from AI use in accounting may allow nonprofits to redirect resources towards other areas to increase efficiency or service delivery.

1.3 Donor Segmentation

Fundraising and donor management are two areas most responsible for successful nonprofit service delivery; however, applying effective fundraising techniques and understanding donor behavior are perhaps the hardest tasks for nonprofit leaders. According to recent sector report, 2026 is the second consecutive year of donor decline at a rate of -0.8%, and converting new donors to repeat givers is a continued pain point (Association of Fundraising Professionals, 2026). Optimizing fundraising data analytics may therefore offer greater returns. However, the literature reveals various data analytics go unused due to lack of expertise, low financial budgets, and insufficient awareness of data analytics capabilities in the nonprofit sector (Alsolbi, 2023, p. 14). AI machine learning in turn may offer nonprofits additional support in these areas. ML systems can analyze fundraising and donor data, helping nonprofits increase donor retention and increase net donation revenue.

A recent paper, published in *Journal of Nonprofit Innovation*, explored various ML techniques—supervised learning, unsupervised learning, semi-supervised, and reinforcement—to observe ML's ability to enhance nonprofit operations (Holzer, 2024). Supervised learning is a type of machine learning technique that uses labeled data sets to train AI models to identify patterns. Unsupervised learning uses machine learning algorithms to categorize unlabeled data sets. Semi-supervised combines the previous two methods, while reinforcement learning has an AI agent learn through trial and error by receiving feedback on its actions. Holzer (2024)

simulated donor data using Python⁴; this simulated data was then used to train a supervised learning algorithm, called an artificial neural network (p. 48). An artificial neural network is a mathematical model designed after biological neurons, where stacked units learn pattern-recognizing weights and biases to map informational inputs to determined outputs (Lee, 2026). Eighty percent of the simulated data was input into the model with the known target or prediction results; the remaining 20% was reserved to compare known results to the model's predictions to evaluate the model's accuracy (Holzer, 2024, p. 49). Results demonstrated that the supervised learning algorithm, the neural network model, was able to predict repeat donors with 86% accuracy. For nonprofits looking to increase the number of repeat donors, machine learning systems may provide insights into donor behavior, enabling nonprofits to identify and prioritize fundraising efforts.

Similarly, in a 2021 study, researchers combined a natural field experiment and a causal machine learning technique with open data to increase a charity's fundraising effectiveness (Cagala et al., 2021). Data in this study was gathered using two separate, comprehensive data sets: the administrative records provided by the Catholic Church and Google Maps. Google Maps API provided geospatial information on economic and cultural facilities (such as restaurants, supermarkets, medical facilities, and churches) near each individual's residence. Once combined, researchers used a causal machine learning system to identify heterogeneity in donors preferences to identify donors who would most likely become repeat givers. The fundraising instrument was asking for a small unconditional gift with an accompanying solicitation letter. Donor information was then split into a warm list, donors who had previously donated once or more before the experiment, and a cold list, donors who had never donated. To

⁴ An overview of python code used in this study to simulate donor behavior created by Holzer (2024) can be found at https://github.com/JustinHolzer/ML_for_Nonprofits

maximize fundraising potential, researchers used machine learning to prioritize solicitations that would likely cover the cost of the fundraising instrument (this process is referred to as “optimal targeting”). Researchers found that the optimal targeting rule in the warm list allowed charities to substantially increase donations. In the first year after the experiment, the average net donation, less costs, under the optimal targeting rule was 13.8% higher than if everyone on the warm list had received an ask. The machine learning optimal targeting technique was able to identify individuals who would likely become a repeat donor and whose expected donation would exceed the solicitation costs, thus resulting in higher net fundraising revenue. The study shows machine learning may allow nonprofits to increase both fundraising revenue and donor retention.

However, this study also demonstrated another tool for nonprofits to optimize machine learning: the use of publicly available information. This study emphasized geospatial characteristics due to their availability to charities and ability to act as proxies for income and other socioeconomic characteristics (Cagala et al., 2021). These findings indicate nonprofits may use publicly available data to teach machine learning systems to target donors who are likely to give again, thus lowering fundraising costs and increasing net fundraising returns.

1.4 Addressing Access Barriers to Machine Learning: Using Open Data

AI adoption and machine learning may increase spending for nonprofits already facing limited resources. One way for nonprofits to bridge the gap between costs and optimizing AI is using open data. Open data refers to “data that is freely available for use, reuse, and distribution by anyone, without restriction” (UNESCO, n.d.).

In another study, open data was used to evaluate machine learning ability to predict donor behavior. Schetgen et al. (2021) used several supervised learning techniques on Facebook data to evaluate its value in predicting first-time donation behavior. Researchers used data simplification

techniques to cope with the great number of Facebook pages and groups. They input the data into seven different machine learning techniques (logistic regression, k-nearest neighbors, bagged trees, random forest, adaboost, extreme gradient boosting, and artificial neural networks) (Schetgen et al., 2021). These are types of supervised learning algorithms. This process helped identify which combination of data reduction and AI modeling techniques were best at accurately predicting donation behavior when using social media data.

Schetgen et al. (2021) found that Facebook pages and categories of Facebook pages were the most important data types. The most important predictors included age, education, residence, materialism, consumption, and interest in nonprofits. To assess predictive performance, researchers applied two distinct measures widely used to evaluate the performance of classification algorithms and AI models: the receiver-operating characteristic (ROC) curve and the area under the curve (AUC), along with the top decile lift (TDL). The ROC curve graphically represents the relation between the true positive rate, that is, how often the algorithm or AI model correctly predicts the positive observation, and the false positive rate, or how often the model incorrectly predicts a negative observation as positive. In this experiment, a positive observation is when the model correctly predicts donation behavior using Facebook data across all observations. The AUC can be interpreted as the probability that a random positive observation gets a higher score than a random negative observation; a score of 0.5 indicates that the predictive model does not perform any better than a random classifier, whereas a value equal to 1 means that the predictive model is perfect (Schetgen et al., 2021, pp. 6-7). The AUC measures the performance over the entire range of predictions, whereas the TDL focuses on the predictions with the highest likelihood to become a donor. TDL measures how much better a predictive model is at identifying positives, in this case correctly predicting those with the

highest likelihood to become a donor, in comparison to random selection (Schetgen et al., 2021, p. 7).

The study found their top performing combination was singular value decomposition (data simplification method) and logistic regression (supervised learning technique). This combination resulted in an AUC of 0.72 and a TDL of 3.33 (Schetgen et al., 2021, p. 7). This means their top performing algorithm model is moderately able to distinguish between donors and non-donors, while a TDL of 3.33 suggests that, within the top 10% of potential donors, their model identifies 3.33 times more donors likely to donate than a random model (Schetgen et al., 2021).

These findings indicate Facebook data may be used as a valuable, free data source that can be leveraged by machine learning. Nonprofits can use publicly available data, such as Facebook data or geospatial data, for machine learning algorithms to identify, acquire, and predict who is likely to become a donor, and which ones are more likely to donate again.

1.5 Summary: Nonprofits & AI Benefits

Overall, artificial intelligence may offer nonprofits better resource allocation resulting in lower operating costs and higher material gains. The time-saving ability of AI may increase productivity in nonprofit tech, administration, or financial departments. AI automation may particularly offer financial departments lower overall costs through AI bookkeeping to reduce errors and streamlined invoice processing. Machine learning predictive models can organize, analyze, and predict donor behavior to identify major donors and those likely to give again. For nonprofits, this processing power could improve donor management through improved donor targeting. However, AI use is accompanied by potential risks, specifically for nonprofits where

information accuracy, compliance, data security, and ethics are high concerns. The following sections will discuss the potential concerns nonprofits face with AI use.

2. AI Risks for Nonprofits: Concerns & Considerations

AI use comes with concerns around information reliability, data security, and liability issues. Nonprofits deciding to implement AI systems or those already using AI without organization-specific policies may want to consider these risks. The literature reveals AI-related risks largely come from generative AI use. The following sections will explain the concerns of generative AI and how it may impact nonprofit organizations.

2.1 Unreliable or Inaccurate Information

AI gets information wrong. Generative AI systems often generate what experts call “hallucinations” or confidently present false, nonsensical, or fabricated information, such as made-up facts, sources, or images that aren’t grounded in reality or its training data (IBM, 2023). According to Brynjolfsson et al. (2023), “generative AI tools face significant challenges in real-world applications. At a technical level, popular LLM-based tools, such as ChatGPT, have been shown to produce false or misleading information in unpredictable ways, generating concern about their ability to be reliable in high-stakes situations” (p. 8). Nonprofits rely on accurate information and public trust. Inaccurate information from generative AI risks backtracking, rework, and damaging public and donor relationships through quality failure. For example, Deloitte recently faced criticism for producing a report with AI-generated errors, including a fabricated quote from a federal court judge and references to nonexistent academic research papers (Dhanji, 2025; Karp, 2025; McGuirk, 2025; Shibu, 2026).

Furthermore, Brynjolfsson et al. (2023) raised concerns on AI accuracy outside of the lab: “while LLM models often perform well on specific tasks in the lab, the types of problems

that workers encounter in real-world settings are likely to be broader and less predictable” (p. 8). This raises concerns for both AI’s ability to provide accurate assistance in every circumstance and whether workers will be able to distinguish cases where AI tools are likely to be less accurate. This concern is compounded when considering AI use may be used to meet labor skill gaps, indicating thorough training on proper AI use and oversight may be vital for high-stakes work. Brynjolfsson et al. (2023) found the efficacy of AI technologies is likely to depend on how they interact with existing workplace structure and may have more limited effects in practice due to the need for complementary organizational investments, skill development, or business process redesign (p. 8). For nonprofits facing limited resources and labor shortages, investing in proper AI training may outweigh the reward. Because generative AI technologies are only beginning to be used in the workplace, little is currently known about their impacts, thus results may vary across individual nonprofits, and careful consideration may be required.

With that in mind, AI may not be saving workers as much time as they think. Humlum and Vestergaard (2025) used Difference in Differences (DiD) to estimate null effects on earning and recorded hours, examining the early labor market impacts of AI chatbots by linking large-scale representative adoption surveys to administrative labor records in Denmark (p. 1). With their confidence intervals ruling out changes larger than 3%, the authors found AI users do not experience better outcomes in the labor market (Humlum & Vestergaard, 2025, pp. 14-15). For users who report using AI tools daily and report substantial benefits, researchers were unable to find any net changes in earnings or hours (Humlum & Vestergaard, 2025, p. 16). These findings indicate AI may not be saving users as much time as they estimate.

2.2 Privacy and Security Risks

Nonprofits are often entrusted with sensitive data, referred to as personally identifiable information (PII), such as names, addresses, social security numbers, and financial details of both donors and beneficiaries. Data leaks can occur when using open-source AI tools, such as OpenAI's ChatGPT or Google's Gemini. Open-source models are available for public use and modifications. Using open-sourced AI models present data leak risks if sensitive information or PII is not anonymized or redacted before being inputted. The major concern here is once the AI system uses the information for training, it becomes nearly impossible to retrieve or delete the information from the AI model (Sadoian, 2025) and can appear in future outputs. This raises serious data privacy concerns. Specifically for nonprofits, data leaks from open-sourced AI could damage relationships with program participants and donors. Enterprise AI, subscription-based models, offer nonprofits greater protection from data leaks. Enterprise AI subscriptions allow data encryption, regulation compliance, and more control over data usage. Nonprofits seeking AI adoption may want to opt for enterprise AI models to mitigate risks.

2.3 Liability

Liability and legal regulations regarding AI use remain largely unsettled. That said, there have been recent liability concerns around AI generated content that nonprofits may want to consider. LLMs rely on vast amounts of data to generate content, including millions of copyrighted works from images, text, music, and other intellectual property. Generative AI does not make truly original work but instead rearranges its learned data in new or different ways. Therefore, using AI-generated content for marketing, website copy, visual materials, or published writing comes at risk of copyright infringement (Shields, 2025). Those who use copyrighted material, whether aware of the infringement or not, are liable (Shields, 2025). Small

nonprofits are particularly vulnerable here as the cost of infringement lawsuits are costly and may compromise small nonprofit longevity.

Additionally, according to Kelley Kronenberg, attorneys at law, AI generated content also causes ownership dilemmas. Under U.S. copyright law, protection extends only to works created by human authors. Thus, if the majority of the content was AI-generated with minimal human input, that material may not qualify for copyright protection at all. Moreover, the scope of protection for AI-assisted content remains unclear. AI-generated drafts or outlines that are significantly edited, revised, or added to with human creativity may or may not qualify for copyright protection. That is, nonprofits could be on either end of a copyright lawsuit with AI-generated content.

2.4 Summary: Nonprofits & AI Risks

Overall, AI use in nonprofits presents legitimate legal concerns. Currently, AI legislation is still developing and compliance remains unclear. Nonprofits could face data leaks and data privacy violations from improper use of open-sourced AI models resulting in lawsuits, reputational damage, or potential harm to partners. Copyright laws around AI-generated content are unclear about what is and is not protected. The law has yet to catch up with new technology. In the meantime, nonprofits may face copyright violations or infringement lawsuits using AI-generated content. The issue of liability presents a related concern: the lack of AI regulation.

3. Legal and Existing Frameworks

Currently, at the time of writing this paper, there is a general lack of regulation and governance frameworks available for AI management. AI regulation is in the beginning stages of development, and the status of nonprofit regulation is still uncertain. As of March 2026, there are no U.S. Federal regulations on AI and limited state laws (Balidemaj, 2024). As AI continues to

develop, more frameworks for AI governance and development will continue to emerge. This paper highlights the EU AI Act as the world-leading AI development legislation, the state of California's Consumer Privacy Act, and the AI framework developed by National Institute of Standards and Technology (NIST) for its method of AI governance across sectors, AI models, and AI systems. This section stands as a starting point for nonprofits. Additional framework resources for nonprofits appear in Appendix B.

3.1 EU AI Act

The AI Act in the European Union is the first regulation on artificial intelligence that aims to promote the development of trustworthy AI while protecting fundamental rights and public safety. It established four risk levels of AI systems. The highest level is unacceptable risk; this level entirely prohibits any AI systems that manipulate users, exploit vulnerability, or enable mass biometric surveillance (EU Artificial Intelligence Act, 2024). Nonprofits are unlikely violators of this level, especially smaller ones. That said, this ban on AI technology is currently only active in the EU. Therefore, nonprofits should be aware of the potential risks of AI systems and be prepared to exercise judgement as this technology develops and potentially impacts their communities.

The next level is high-risk AI systems, which are AI systems that profile individuals (i.e., automated processing of personal data to assess various aspects of a person's life, such as work performance, location, health, etc.). The EU AI Act requires high-risk AI systems to meet standard risk, privacy, and security measures, such as establishing a risk management system, maintaining data governance, and deploying human oversight. Nonprofits would likely fall under this category since they often record and maintain personal identifiable information of their donors and program participants. Previous sections highlighting potential privacy and security

risks regarding AI emphasizes the importance of taking early intervention steps. Nonprofits may want to consider proactive measures to protect nonprofits community members.

The next two levels are limited-risk and minimal or no-risk AI systems. Limited-risk AI systems are systems that pose little to no threat to people's safety, privacy, or fundamental rights, such as chatbots or generated content. The act requires these systems to meet standard risk, privacy, and security measures and provide transparency. That is, users must be informed when interacting with AI systems or AI-content. Lastly, minimal or no-risk AI, like spam filters or video games, are largely unregulated under this act, but are subject to general safety, privacy, and transparency rules.

3.2 California State Laws on Data Privacy

While there are multiple US states that have implemented data privacy laws, field experts, policymakers, and AI consultants most frequently reference the California Consumer Privacy Act (CCPA). According to California's Attorney General, the CCPA establishes data processing obligations and provides consumer rights. Consumer rights under the CCPA include the right to know about the personal information a business collects about them and how it is used and shared, the right to delete their personal data, the right to opt-out of the sale or sharing of their personal information, the right to correct inaccurate personal information that a business has about a consumer, the right to limit the use and disclosure of sensitive personal information collected about a consumer, and lastly, the right to non-discrimination for exercising their CCPA rights (California Department of Justice, 2024).

3.3 NIST AI RMF

The National Institute of Standards and Technology (NIST, 2023) provided a robust approach in their AI Risk Management Framework (AI RMF). The report emphasizes the risks

and potential harmful impacts of AI use, including data security violations, inherited bias, and using inaccurate information. To mitigate these concerns, the AI RMF establishes four core functions: Govern, Map, Measure, and Manage (NIST, 2023, p. 20). According to the NIST (2023), the Govern function includes fostering a culture of risk management within all organization departments; it involves outlining processes that anticipate, identify, manage, and assess the risks an AI system can pose (NIST, 2023, p. 23). The Map function “establishes the context to frame risks related to an AI system” (NIST, 2023, p. 24) and is considered the basis for the Measure and Manage functions. The Measure function uses a variety of techniques and methodologies (quantitative, qualitative, mixed methods, etc.) to analyze and assess AI risk (NIST, 2023, p. 28). The final function, Manage, includes allocating risk resources for each function (NIST, 2023, p. 31).

Figure 2: NIST AI RMF Core



Figure 2 shows the NIST AI RMF Core functions: Govern, Map, Measure, and Manage. Governance is designed to inform and be integrated throughout the other three functions.

Source: The National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*.

Ultimately, the NIST AI RMF advocates prioritizing a culture of risk management with AI. While much of the specific actions in the NIST AI RMF are directed at AI developers, the AI RMF Core is a helpful guide offering a clear structure to AI management. For nonprofits, AI management could require AI training before using AI with internal data, identifying when AI may get information wrong or be ineffective for certain tasks, and developing a “Measure” function, like an audit process to better manage AI-assisted workflow. The report emphasizes the importance of clear regulatory requirements for employees and suggested aligning regulations based on business objectives (p. 24). For nonprofits pursuing integrating AI, this report emphasizes the importance of employee support and AI training. Therefore, nonprofit leadership may need to develop AI training before fully integrating AI systems.

3.4 Summary: Nonprofits & Legal Frameworks

Nonprofits face uncertainty around compliance as legislation tries to catch up with the new technology. Nonprofits seeking AI integration may want to proactively adhere to recommended regulation and AI governance models as state-, federal-, or country-specific legislation develops. Proactive adherence is an act of good faith and builds trust with nonprofit communities. It supports data privacy and security for their members and sets an example for broader business cultural norms.

4. Ethical Considerations & Frameworks

Ethics and ethical AI are a major concern among researchers. There is a wealth of literature on the topic, ranging from AI’s impact on the labor market and social outcomes to larger concerns like AI’s potential effect on legal and governmental frameworks. However, it appears researchers’ concerns and predictions stem from a consensus within the AI field: there are no ethical standards for AI development. Therefore, building ethical AI systems becomes

increasingly difficult. However, ethics and ethical responsibility is a large part of nonprofit operation. Ethical conduct remains a foundational value of nonprofit work and strengthens trust between nonprofits and their community members. Ethical violations hurt not only a nonprofit's reputation and community relationships but may harm the communities nonprofit aim to support. Damaged trust may lead to decreased financial support and thereby limit a nonprofit's ability to operate.

4.1 Amplify Discrimination and Bias

One major ethical concern regarding AI systems is integrated bias amplifying discrimination. Siau and Wang (2020) explain, “human bias, such as gender bias and race bias, may be inherited by AI” (p. 81). AI agents are trained by humans using human datasets, therefore existing biases may be learned by AI systems. Once AI models learn biased data, the bias will become an ongoing problem (Siau & Wang, 2020, p. 81). For example, Obermeyer et al. (2019) ran a study examining racial bias in healthcare algorithms by using data from 2013 to 2015 from a live, scaled healthcare algorithm deployed nationwide to compare the algorithm's decisions for Black versus white patients (“Data and analytic strategy”). Researchers found evidence of substantial disparities in program screening. For instance, at the same level of algorithm-predicted risk, Black people had significantly more illness burden than their white counterparts yet received a lower risk score. These individuals were put on par with healthier white people in terms of costs and were less likely to qualify for extra care as much as white patients with the same needs. This example indicates that AI algorithms can inherit and amplify human bias and negatively impact marginalized communities.

Nonprofits often support marginalized communities; therefore, if nonprofits use AI algorithm systems (that could be unknowingly discriminatory) in their decision-making process

or to scan program participants, some applicants could be scored unfairly and not receive the support they need. Ultimately, this could impact how nonprofit resources are allocated and to whom, indicating a need for human supervision over AI system. For this reason, nonprofits may want to reconsider using AI algorithms in decision-making or applicant screening, unless specifically developed for their organization or mission.

4.2 Large-scale Ethical Concerns and Implications

Nonprofits are ethical actors within the broader context of society. Therefore, when considering adopting AI, some nonprofits may have an interest in developing a greater awareness of the larger or more existential concerns of AI use. AI implications resulting in ethical concerns include environmental impacts and asymmetrical global development. While these concerns may not be of particular interest to all nonprofits, they highlight that AI use has far-reaching implications. As ethics and AI are a broad literary topic, this section illuminates selected concerns around AI use but should not be considered as a fully realized analysis of these concerns. Rather, this section provides a short list of arguments being discussed by experts, to be used by nonprofits to gain awareness and begin further investigation if interested in a particular issue. This section will cover researchers' concerns about AI impacts on democracy, the environment, and international development.

4.2.1 AI and Democracy

There are many sources that discuss potential AI impacts on democracy and civil rights. For example, Siau and Wang (2020) state unethical AI outputs fragment the truth and result in a loss of trust and societal support, which can dent the strengths of democracies (p. 82). Similarly, Summerfield et al. (2024) evaluated how democracies may be weakened by AI systems, and found AI systems, like LLMs, are able to write persuasive messages on political issues, may

exacerbate polarization by generating replies that flatter the user's perceptions and isolate them from ideologically opposing views, can hallucinate or cite poorly sourced material, and can be used to mislead voters (Summerfield et al., 2024). In this study, researchers tested the accuracy of leading proprietary and open-source models on practical queries about electoral participation in the US, finding that over half of replies were inaccurate (Summerfield et al., 2024, p. 9). These risks to voter persuasion, wider polarization, misinformation, and misuse could erode democracy and the integrity of free elections. That said, Summerfield et al. (2024) recognize where AI systems could enhance or reimagine how democracies might work better through greater information access, strengthening public discourse, and helping people find common ground (p. 11). If optimized ethically, AI could support greater political collaboration. Nonprofits who target civil rights, equitable social outcomes, or legal services may be particularly interested in these potential risks and rewards of AI use. The impact of AI on democracy and civil rights is still largely developing. For interested parties, further investigation into researchers' concerns, predictions, and potentials is recommended.

4.2.2 AI and the Environment

Researchers also raise concerns about AI's impact on the environment. According to Vries-Gao (2026), "the carbon footprint of AI systems alone could be between 32.6 and 79.7 million tons of CO₂ emissions in 2025, while the water footprint could reach 312.5–764.6 billion L" (p. 1). For reference, this means AI systems have approximately the same carbon footprint range as New York City (52.2 million tons of CO₂ emissions in 2023), and the water footprint may be in the same range as the entire global annual consumption of bottled water (446 billion L). That said, researchers evaluating the environmental impact of Information and Communication Technology (ICT) data centers and AI struggle to source data due to a lack of

reporting transparency in the tech sector. Therefore, the accuracy of these estimates varies; the actual footprint of AI systems could very well be larger or smaller. For environmentally minded nonprofits, using AI poses a conflict between values and potentially increased service delivery. Nonprofits facing this decision may want to evaluate this decision with board members, employees, and their community to decide what plan of action is right for them.

4.2.3 Asymmetrical Global Development

The last concern is AI's potential impact for developing economies, which may be of greater interest to international nonprofits. According to Maghsoudi et al. (2025), AI adoption and societal implications vary significantly worldwide, shaped by differing priorities, socio-economic realities, and governance capacities ("Abstract"). Researchers analyzed global public discourse on AI and its alignment with the United Nations' Sustainable Development Goals (SDGs) using natural language processing and a graph-based, multi-label classification framework. Their results revealed "fundamental disparities in how artificial intelligence is conceptualized and operationalized within sustainable development frameworks across different economic contexts" (Maghsoudi et al., 2025, p. 17). Their research found AI's transformative potential is universally acknowledged across nations, but AI deployment and impact vary drastically depending on a nation's economic position in the global landscape. The financial challenge was cited as the most significant barrier to equitable AI deployment. This indicates that while there are many SDGs that AI is well-suited to support, the lack of adequate funding for developing economies to implement AI solutions could widen the gap between national economies, raising serious global equity concerns. More research exists on this topic, often calling for greater research into how AI systems may amplify social inequality to mitigate potential harm (Hagerty & Rubinov, 2019). International nonprofits interested in AI adoption

may want to explore ways to bridge this gap and how to implement AI in ways that mitigates harm to vulnerable populations.

Overall, research shows the two sides of AI: one of great potential to meet world problems, and another of serious ethical concerns. As AI continues to develop, researchers are calling for intentional and transparent AI development and carefully considered implementation to prevent widening the global development gap, ethical violations, and degradation of democracy or the environment. Nonprofits can play an influential role in this call to action by integrating researchers' concerns into their own AI use, advocating for ethical AI, and preparing to support their communities as AI implementation continues.

5. Discussion

Nonprofit work can be considered high stakes because they often work with marginalized communities in sectors that directly impact people's lives, including education, healthcare, legal services, and more. Nonprofit decisions on AI adoption will clearly depend on each nonprofit as each has their own mission and values. That said, the following provides considerations for nonprofits pursuing AI adoption.

5.1 Human Oversight

Human oversight mitigates risks related to AI use, maintains safety and transparency, and prevents errors. The research found AI is subject to providing false, misleading, or at times, biased information. Inaccurate information puts nonprofits at risk of lawsuits, reputational damage, and misallocating resources. From a legal position, nonprofits could be held responsible for using AI-generated content and face copyright lawsuits from visual design artists, marketing companies, or other nonprofits who originally created collateral service materials that was then used by generative AI. The risk of copyright infringement demonstrates nonprofits' need to

reconsider how they should use genAI when creating front-facing content. Copyright laws only cover human-created works, leaving regulation on AI-assisted work unclear. Until greater legal regulations are developed, it may be advisable for nonprofits to limit their use of generative AI in creative or communication work, such as promotional or program material, blog posts, or grant writing. Use of inaccurate information could result in both ethical and reputational concerns for nonprofits. In a recent Deloitte financial report, an AI's hallucination caused a costly lawsuit and a misallocation of resources (Karp, 2025; Shibu, 2026). For nonprofits, a similar mistake could result in reputational damage with donors, board members, and their program partners.

Therefore, human oversight over AI should be mandatory within nonprofit operations. Proper employee training is also advised before integrating AI systems to mitigate and spread supervision over operation. Each employee could then provide oversight and critical judgement of AI outputs at each step of collaboration to ensure adherence to privacy and security requirements, legal compliance, and the organizations' mission and values. Lastly, employees, at all levels, should be aware of when AI tools may be inaccurate, biased, or unequipped to tackle certain workplace tasks.

5.2 Dealing with Data: Preparation & Training

The research finds improper use of AI could lead to data privacy leaks. This is a serious legal and ethical concern for any firm. The high-stakes nature of nonprofit work may require greater attention to proper data preparation when using AI systems. The research recommends restricting open-sourced AI models and instead opting for enterprise AI models that offer data encryption, regulation compliance, and more control over data usage. To further support data privacy, it is recommended to anonymize or redact sensitive, confidential, or personal identifiable information before inputting information into AI systems.

5.3 Transparency

Transparency around AI is where nonprofits can lead the way through proactive compliance and disclosure. In terms of ethics, researchers find a lack of ethical standards in AI training and are concerned about AI inheriting gender or racial biases, thereby amplifying systemic discrimination. This is where nonprofits may play a unique role in AI adoption. The current lack of AI legislation, transparency in the tech industry, and development of ethical AI puts nonprofits in the position to demand ethical standards in AI development as AI consumers and set the example for AI transparency as a firm. Emerging AI regulation, such as the EU AI Act or the California Consumer Privacy Act, provides nonprofits the opportunity to consider proactive compliance. This approach can build trust with communities and influence AI workplace standards in other industries. Clear disclosure of when and how AI is used with front-facing content builds organizational trust. This act of transparency exemplifies integrity and excellence. Doing so may also proactively protect nonprofit community members.⁵

5.4 Audits and AI training

For nonprofits, AI incident reporting and regular audits may be a necessary part of AI adoption. The research has discussed the potential risks and concerns of AI systems, including amplifying discrimination and biases, data privacy and security risks, compliance complexities, liability concerns, and inaccurate information. Audits and incident reporting can work together to mitigate these risks. Audits provide a systematic process for nonprofits to ensure AI systems are meeting their organization's standards. This is one method of AI risk management. Incident

⁵ For nonprofits looking for examples of AI disclosure, this author found WIRED's AI policy to be a helpful guide: "[How WIRED Will Use Generative AI Tools](#)" (WIRED, 2023).

reporting is another. Through logging and reporting when AI outputs are false, biased, or hallucinated, nonprofit can better determine how and when they want to continue using AI in ongoing operation. This can inform their training of proper AI use. According to the NIST AI RMF, risk management is a crucial pillar in AI governance. Quarterly user-input audits and evaluations based on company capacity could offer another solution to AI-related risks. Another step nonprofits can take in AI risk management is opting out of AI model training, turning off memory, or routinely deleting AI memory.

5.5 Limitations

This paper encountered several challenges. The first and foremost was the limited amount of quantitative research studies on AI use, specifically for the nonprofit sector. For this reason, much of the studies selected for consideration and their findings were used to make inferences on their applications for nonprofits. However, it must be recognized these inferences may not generalize to nonprofit operations. Further quantitative research on AI in the workplace and specific uses for nonprofits is recommended.

6. Summary

Ultimately, this paper recommends the following for nonprofits due to the high-stake nature of their work.

1. Human oversight is crucial for mitigating AI risks.
2. Data protection of sensitive and private data should be anonymized or redacted when using open-sourced AI models. Nonprofits likely should opt for enterprise AI models that offer greater data protection, due to the high-stakes nature of their work.

3. Any creative content generated by AI is at risk of copyright infringement. Nonprofits may need to carefully consider how and when they are using generative AI in front-facing content.
4. AI risk management includes incident reporting and regular audits to monitor for AI errors.
5. NPOs may want to consider the ethical implications of AI before implementing it into their service delivery.

Conclusion

As AI adoption continues to increase across sectors, nonprofits are faced with a choice. AI systems may support nonprofit operations by reducing the time it takes to complete administrative, writing, or operational tasks, thus increasing their productivity and allowing more time dedicated towards strategic or creative thinking. The evidence supports that AI can reduce errors and improve efficiency in accounting operations, as well as improve fundraising results through predicting donor behavior. Improvements in these areas may allow nonprofits to better allocate their resources and therefore better support their communities and meet donor needs.

However, AI use in nonprofit work is not without risks and concerns. The research covers a broad range of potential AI risks to address and illuminate the varying concerns of nonprofit decision makers. AI risks range from information inaccuracies to larger societal concerns like environmental damage or asymmetrical global development. This research highlights these concerns in hopes to provide evidence-based analysis on AI use in the workplace for nonprofit leadership to evaluate AI adoption with clear eyes. Overall, this research highlights the importance of human oversight over AI systems in the high-stakes work of nonprofit service

delivery. If nonprofits have or are in the process of adopting AI, this paper highly recommends paying close attention to proper employee training and proactive AI risk management.

References

- Acemoglu, D., & Autor, D. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*. 128 (6). 2188-2244. doi.org/10.1086/705716
- Alsolbi, I. N. (2023). Leveraging potentials of big data for better decision-making and value creation in nonprofit organizations (Order No. 30830589). [Doctoral dissertation, University of Technology Sydney]. <https://opus.lib.uts.edu.au/handle/10453/172435>
- Arora, R.N., (2024, July 1). *The beginner's guide to AI for nonprofits*. Daxko. daxko.com/insights/the-beginners-guide-to-ai-for-nonprofits
- Association of Fundraising Professionals. (2026). "Quarterly fundraising report". *Fundraising Effectiveness Project*. <https://publications.fepreports.org/>
- Azman, N.A., Mohamed, A., Jamil, A.M. (2021). Artificial intelligence in AI automated Bookkeeping: A value-added function for small and medium enterprises. *International Journal on Informatics Visualization*. 5(3). 224-230. joiv.org/index.php/joiv/article/view/669/353
- Babina, T., Fedyk, A., Xi He, A., Hodson, J. (2022). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*. 151. doi.org/10.1016/j.jfineco.2023.103745
- Balidemaj, V. (2024). Artificial intelligence integration, concerns, benefits, and the need for ethical policies for community foundations and nonprofit organizations. *Clemson University. All Dissertations*. 3697. www.proquest.com/docview/3122677290?%20Theses&fromopenview=true&pq-origsite=gscholar&sourcetype=Dissertations%2
- Bick, A., Blandin A., Deming, D. (2025, February 27). The rapid adoption of generative AI. *Federal Reserve Bank of St. Louis*. Working Paper 2024-027. <https://doi.org/10.20955/wp.2024.027>
- Brynjolfsson, E., Li, D., Raymond, L. (2023). Generative AI at work. *National Bureau of Economic Research*. Working Paper No. w31161. DOI 10.3386/w31161
- Cagala, T., Glogowsky, U., Rincke, J., Strittmatter, A. (2021). Optimal targeting in fundraising: A casual machine-learning approach. *Cornell University*. arxiv.org/abs/2103.10251
- California Department of Justice. (2024, March 13). *California Consumer Privacy Act (CCPA)*. Office of the Attorney General, California Department of Justice. www.oag.ca.gov/privacy/ccpa#heading5d

- Dhanji, K. (2025, October 6). Deloitte to pay money back to Albanese government after using AI in \$440,000 report. *The Guardian*. <https://www.theguardian.com/australia-news/2025/oct/06/deloitte-to-pay-money-back-to-albanese-government-after-using-ai-in-440000-report>
- EU Artificial Intelligence Act. (2024, February 27). *High-level summary of the AI Act*. EU Artificial Intelligence Act. <https://artificialintelligenceact.eu/high-level-summary/>
- Hagerty, A. & Rubinov, I., (2019). Global AI ethics: A review of the social impacts and ethical implications of artificial intelligence. *Cornell University*. arxiv.org/abs/1907.07892
- Holzer, Justin (2024). Machine learning for nonprofit organizations. *Journal of Nonprofit Innovation*. Vol.4(2). scholarsarchive.byu.edu/joni/vol4/iss2/6
- Humlum, A., & Vestergaard, E., (2025). Large language models, small labor market effects. *Becker Friedman Institute*. Working Paper No. 2025-56. papers.ssrn.com/sol3/papers.cfm?abstract_id=5219933
- IBM. (2023, September 1). *What are AI hallucinations?* IBM Think. ibm.com/think/topics/ai-hallucinations
- Karp, P. (2025, August 22). Academics raise alarm over suspected AI use in Deloitte report. *Financial Review*. <https://www.afr.com/companies/professional-services/academics-raise-alarm-over-suspected-ai-use-in-deloitte-report-20250822-p5mp0f>
- Lee, F. (2026). *What is a neural network?* IBM Think. <https://www.ibm.com/think/topics/neural-networks>
- Maghsoudi, M., Mohammadi, N., Bakhtiari, M., (2025, December). Artificial intelligence and sustainable development: Public concerns and governance in developed and developing nations. *Cleaner Environmental Systems*. *Science Direct*. Vol 19. doi.org/10.1016/j.cesys.2025.100340
- McGuirk, R. (2025, October 7). *Deloitte to partially refund Australian government for report with apparent AI-generated errors*. The Seattle Times. www.seattletimes.com/business/deloitte-to-partially-refund-australian-government-for-report-with-apparent-ai-generated-errors/
- Miller, D (2024, November 5). *The non-profit's AI journey - from tactical improvements to strategic advantage*. Diginomica. <https://diginomica.com/non-profit-ai-journey-tactical-improvement-strategic-advantage>
- Mohammad, S. J., Hamad, A. K., Borgi, H., Thu, P. A., Sial, M. S., Alhadidi, A.A. (2020). How artificial intelligence changes the future of accounting industry. *International Journal of Economics and Business Administration*. 8(3), 478-488. DOI: 10.35808/ijeba/538

- National Institute of Standards and Technology. (2023). *Artificial intelligence Risk Management Framework. (AI RMF 1.0)* (Report No. NIST AI 100-1). U.S. Department of Commerce. [//doi.org/10.6028/NIST.AI.100-1](https://doi.org/10.6028/NIST.AI.100-1)
- Noy, S. & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*. 381(6654). 187-192. DOI: 10.1126/science.adh2586
- Obermeyer, Z., Powers, B., Vogeli, C., Mullainathan, S., (2019). Dissecting racial bias in an algorithm used to manage health of populations. *Science*, 366(6464), 447-453. <https://www-jstor-org.offcampus.lib.washington.edu/stable/26843725>
- Sadoain, L. (2025, December 2025). *The shadow AI data leak problem no one's talking about*. UpGuard. upguard.com/blog/shadow-ai-data-leak
- Schetgen, L., Bogaert, M., & Van den Poel, D. (2021). Predicting donation behavior: Acquisition modeling in the nonprofit sector using Facebook data. *Science Direct*. 141(113446). doi.org/10.1016/j.dss.2020.113446
- Shibu, S. (2026). This government paid Deloitte \$1 million for report that included fake AI citations. *MSN*. <https://www.msn.com/en-us/money/companies/this-government-paid-deloitte-1-million-for-a-report-that-included-fake-ai-citations/ar-AA1R9x6F>
- Shields, T. (2025, June 26). *When AI content creation becomes a legal nightmare: The hidden risks every business owner must know*. Kelly Kronenberg. kelleykronenberg.com/blog/when-ai-content-creation-becomes-a-legal-nightmare-the-hidden-risks-every-business-owner-must-know/
- Siau, K., & Wang, W. (2020). Artificial intelligence (AI) ethics: Ethics of AI and ethical AI. *Journal of Database Management*. 31(2). 74-87. DOI:10.4018/JDM.2020040105
- Summerfield, C., Argyle, L., Bakker, M., Collins, T., Durmus, E., Eloundou, T., Gabriel, I., Ganguli, D., Hackenburg, K., Hadfield, G., Hewitt, L., Huang, S., Landemore, H., Marchal, N., Ovadya, A., Procaccia, A., Risse, M., Schneier, B., Seger, E., ... Botvinick, M., (2024). How will advanced AI systems impact democracy? *Cornell University*. 1-25. arxiv.org/abs/2409.06729
- Toloka Team. (2026, March 3). *Difference between AI, ML, LLM, and generative AI*. Toloka. toloka.ai/blog/difference-between-ai-ml-llm-and-generative-ai/
- UNESCO. (n.d.) Open data. *UNESCO*. <https://www.unesco.org/en/query-list/o/open-data>
- Valenti, K. (2025, September 29). *How nonprofits are using AI for greater social impact in 2025*. Sigma Forces. www.sigmaforces.com/post/nonprofits-ai-social-impact-2025

Vries-Gao, A. (2026, January 9). The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Cell Press. ScienceDirect*. 7(1).
doi.org/10.1016/j.patter.2025.101430

Appendix A

AI was not used to write, edit, or revise any of the material in this paper. AI did assist in finding 3 sources, as seen in the link below. The link below is the AI chat used to locate sources about AI use and nonprofit fundraising. Restrictions were provided to ensure AI did not write any of the research or analysis above. The sources provided in the chat were reviewed by this author for credibility and relevance to the research. Ultimately, three sources provided by AI were used in this paper:

- 1) “Optimal targeting in fundraising: A causal machine learning approach” by T. Cagala, U. Glogowsky, J. Rincke, and A. Strittmatter;
- 2) “Predicting donation behavior: Acquisition modeling in the nonprofit sector using Facebook data” by L. Schetgen, M. Bogaert, and D. Van den Poel;
- 3) “Machine learning for nonprofit organizations” by J. Holzer.
<https://chatgpt.com/share/69b6523a-4448-800b-88a6-3f13977c4b58>

Appendix B

List of AI Frameworks for Interested Nonprofits

This list includes peer-reviewed resources about AI integration and AI governance as it was found throughout the research process. This list is by no means exhaustive. Also note, this list does not contain industry or consultant resources on AI governance. It was the consideration of this author that nonprofits could easily find and access such information, whereas the research below may be harder to access or locate. The author hopes this list can offer nonprofits a credible and reliable starting point for AI governance and adoption.

Brendel, A. B., Mirbabaie M., Lembcke, T., Hofeditz, L., (2021). Ethical management of artificial intelligence. *Sustainability*, 13(4), 1-18. DOI:10.3390/su13041974

Corrêa, N.K., Galvão, J., Santos, J.W., Pino, C.D., Pinto, E.P., Barbosa, C., Massman, D., Mambrini, R., Galvão, L., Terem, E., Oliveira, N., (2023). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Cornell University*. 1(10). arxiv.org/abs/2206.11922?utm

Hernández, E.G. (2024). Towards an ethical and inclusive implementation of artificial intelligence in organizations: A multidimensional framework. *Cornell University*. arxiv.org/abs/2405.01697?utm

Perron, B.E., Goldkind, L., Qi, Z., Victor, B.G., (2025). Human services organizations and the responsible integration of AI: Considering ethics and contextualizing risk(s). *Computers and Society*, *Cornell University*. arxiv.org/abs/2501.11705?utm