Non-Profits
Toolkit for AI Governance

INQ CONSULTING • (905)-215-7300 • INFO@INQ.CONSULTING

Table of Contents

Introduction

This toolkit is designed help non-profit organizations navigate the AI integration within their operations. It addresses critical considerations for the responsible use, governance, and deployment of AI technologies, ensuring that organizations can leverage these tools effectively while upholding ethical standards.

This toolkit should be tailored to fit your organization's size and needs. Some parts might seem unnecessary or out of reach, especially for smaller organizations. Many non-profits are just starting with AI and may use off-the-shelf solutions, making AI governance committees or structures less immediate or necessary to develop gradually.

What is AI?

An AI tool (or system) is an engineered tool or system that generates outputs such as content, forecasts, recommendations, or decisions for a given set of human-defined objectives.

This toolkit summarizes best practices for AI governance in the non-profit sector but recognizes that not all will be relevant for every organization. Smaller organizations should prioritize their key use cases and guide employees on responsible AI use through an AI Acceptable Use Policy. A comprehensive roadmap might not be needed, and instead of developing a full AI Ethics Committee, you can enhance existing vendor management and procurement processes to address AI-related issues.

Medium-sized organizations might establish a more robust vendor management system and a preliminary AI Ethics Committee, while more mature organizations can address all topics covered in this toolkit. Use this toolkit as an aspirational resource, not a mandatory one, and customize it to fit your organization's specific requirements.

This toolkit is structured around several key areas of AI governance, specifically:

1. **Roadmap for AI Governance Processes and Tools**: This section provides a structured approach to developing and implementing governance mechanisms that ensure AI technologies are used responsibly and in alignment with organizational goals.
2. **Guidance on Acceptable Use of AI Policy**: Outlines policies that define acceptable and unacceptable uses of AI, ensuring that AI applications are used ethically and effectively to further the nonprofit's mission.
3. **Guidance on Vendor Assessments for AI-Specific Risks**: Offers strategies for assessing potential AI vendors, focusing on the unique risks associated with AI systems to mitigate potential negative impacts on the organization.

4. **Guidance on the Terms of Reference (TOR) for an AI Ethics Committee**: Provides a framework and mandate for an AI Ethics Committee, detailing potential responsibilities and their role in overseeing AI initiatives within an organization.

5. **Guidance on AI Skills-building**: Encourages the development of AI skills programs to educate staff and stakeholders about AI, enhancing their understanding and ability to engage with AI tools effectively, for example, Microsoft's Digital Skills Center for Nonprofits.

6. **Guidance on Terminology**: Aids in understanding key AI-related terms and concepts, fostering clearer communication and deeper understanding among all organization members regarding AI technologies.

Developed in collaboration with INQ Consulting's experts in AI governance, this toolkit aims to empower non-profit organizations to adopt AI solutions that align with their values and enhance their impact, while managing risks and embracing opportunities presented by AI technology.

1. Roadmap for AI governance processes and tools

As non-profits increasingly adopt AI to enhance mission-driven activities, establishing a robust AI governance framework becomes an important part of the organization’s digital transformation process. This section of the toolkit provides a structured roadmap to guide non-profit organizations in developing effective governance processes and tools for AI. This roadmap is suitable for organizations of any size. However, smaller ones (e.g., under 50 employees) might find some parts overly complex or less beneficial for ensuring responsible AI use, while these sections are more valuable to medium to large organizations (e.g., over 200 employees). Use this resource as a helpful guide and reference, not as a strict set of requirements.

By considering this suggested roadmap, organizations should be able to assess whether their AI implementation are ethical, transparent, and aligned with their strategic values and the broader values of the communities they serve. The roadmap outlines critical steps from the initial assessment of current AI use to the continuous improvement of governance and oversight practices, ensuring that AI technologies are used responsibly and effectively to maximize impact while mitigating risks.

Step 1: Define Governance Objectives

Establish clear goals for AI governance, such as ensuring ethical usage, compliance with regulations, and alignment with the organization's mission and values. Developing explicit AI Values and principles can be helpful here.

Step 2: Assessment of Current AI Use

Conduct an inventory of current AI technologies and applications within the organization and evaluate existing processes and tools related to AI, noting areas lacking governance.

Step 3: Determine Future AI Use Cases

Start by assessing existing pain points and areas for improvement, considering existing AI capabilities and the expected developments in the future. Consider how the AI tool will be used and what value it will bring and consider whether you will be able to build the solution in-house or whether it would be more advantageous to procure the solution from a third party. Engage with the communities the AI tools are supposed to serve to understand a diverse set of perspectives. The resulting wish list should then be evaluated and prioritized based on:

- Alignment with the objectives of the non-profit organization
- Feasibility, resources, and ease of development vs. procurement
- Timeliness
- Risks
- Value and positive impact

The prioritization exercise likewise requires input from a broad spectrum of experts, including leadership, IT staff, data scientists, and external experts depending on the use case.

Don't forget to always check whether a non-AI solution may be equally as good or better to solve the identified issue or meet the needs. Misapplying AI can waste time, money, and diminish enthusiasm among employees, volunteers, and leaders for future AI initiatives.

Step 4: Accountability and Oversight Mechanisms

Consider setting up an AI Ethics cross-organization Committee (see Section IV) or board responsible for overseeing AI projects and initiatives. Implement review processes for new AI deployments, including ethical reviews and risk assessments.

Step 5: Risk and Impact Assessment

With the existing and near-term use cases in mind, obtain an overview of applicable laws and regulations and their compliance requirements. Identify foreseeable risks to users and to the non-profit that can arise from the capabilities or use of the AI system (see complementary presentation of AI risks). Engagement of affected stakeholders is key to surface possible long- and short-term risks.

It will be useful at this stage to be clear on whether the AI solution is supposed to be procured or developed in-house, as the risk factors will vary accordingly. In some instances, developing solutions in-house is not feasible or valuable. It will heavily depend on the size of the organization, capabilities, and what vendor solutions exist on the marketplace.

Current AI legislation (drafts) categorize AI systems by risk or impact. In the EU AI Act, for example, some high-risk use cases are prohibited while others are tied to strict compliance requirements. Familiarity with existing categories and categorization of the identified use cases can help future-proofing compliance efforts at this stage.

Step 6: Development of Ethical Principles, Guidelines, and Policies

Start by establishing a set of ethical principles for the organization, which can expand to guidelines and inform the creation of or refinement of policies on data privacy, algorithmic fairness, transparency, and accountability based on legal and technical requirements. Ensure alignment with existing enterprise norms and policies and, where possible, build on what is already in place rather than creating multiple, disparate documents. This integrated approach enhances coherence, streamlines governance practices, and leverages existing frameworks effectively.

Where applicable, establish guidelines for data handling, model training, deployment, monitoring, and continuous improvement that respect privacy, ensure security, and align with the established ethical principles.

Step 7: Develop Key Performance Indicators (KPIs)

Define specific, measurable KPIs to evaluate the effectiveness and impact of AI technologies and governance practices.

Ensure these indicators align with the organization's strategic goals and ethical standards, providing a clear benchmark for success and areas for improvement.

These KPIs might be basic at first, especially for organizations new to AI. However, it's crucial to set clear success criteria for any AI project to make sure it adds value and to apply lessons learned to future projects. For instance, an organization using embedded AI tools might want to measure productivity improvements before and after implementing the tool.

Step 8: Integration of AI Audit and Compliance Tools

Where possible, deploy tools that monitor AI systems for their performance and ethical adherence. Likewise, ensure automatic logging and sufficient timelines for record-keeping. Also, strive to audit AI systems to ensure they function as intended and comply with governance standards. Keep in mind, this may not always be possible or feasible. Undertaking audit and compliance efforts should be prioritized only in highly sensitive contexts and when staff have the appropriate knowledge and expertise to conduct these activities reliably.

Step 9: Skills and Capacity Building

Offer continuous training for staff on AI technologies and governance (see Section V). Emphasize building AI skills and ongoing learning to enhance productivity. Create resources to keep employees informed about the latest AI governance trends and technologies.

Step 10: Establish Feedback Loops

Create mechanisms for feedback from users and stakeholders on AI applications to continuously improve governance practices.

Step 11: Continuous Improvement and Adaptation

Regularly update the governance framework to adapt to new technologies, regulatory changes, evolving organizational needs, and stakeholder feedback. Stay informed about advancements in AI and governance standards to ensure ethical AI usage. Promote continuous learning and skill development within the organization.

Step 12: Reporting

Prepare reports on AI governance efforts for internal stakeholders and external regulators as needed.

Step 13: AI System Retirement and Decommissioning

Establish guidelines for the retirement and decommissioning of AI systems that are no longer efficient, ethical, or aligned with the organization's goals.

Develop processes for safely archiving data, transferring responsibilities, and ensuring that all dependencies are appropriately managed to minimize disruptions.

Evaluate the environmental and social impacts of decommissioning and ensure that the retirement of AI systems adheres to ethical disposal and data protection standards.

Conclusion

By following the steps outline here, organizations can align their AI strategies with their mission goals and the broader values of the communities they serve, ultimately maximizing the positive impact of AI while mitigating associated risks. This proactive approach to AI governance will position non-profits to navigate the evolving landscape of technology with confidence and integrity.

2. Guidance on Acceptable Use of AI Policy

This section establishes guidance on how to define, implement, and enforce a policy that governs the internal usage of AI and, where applicable, informs the terms of use for AI-driven tools developed by your non-profit for social impact purposes. An Acceptable Use Policy (AUP) ensures that AI technologies are utilized ethically, transparently, and in alignment with your organization's mission, whether employed internally or made available to the public. By clearly defining acceptable use, this section helps prevent misuse, protects stakeholder interests, and supports the creation of trust in the technologies your organization deploys to further its social impact goals.

An AUP should contain the following elements:

1. Overview
2. Purpose
3. Scope
4. Definitions
5. Policy
    a. Acceptable Use of AI Systems
    b. Generative AI
    c. Unacceptable Use of AI Systems
    d. Unintentional Misuse
6. Enforcement
7. Review

Annex 1: Key Risks of AI Systems
Annex 2: Determining High-Impact AI Systems
Annex 3: Actions Required for High-Impact AI Systems

Where applicable, sections 5 and its subsections as well as section 6 could be subdivided into "internal use" and "use of AI systems made available to the public." Alternatively, terms of use for AI systems deployed by the non-profit could be developed separately and referred to in the AUP, insofar as these tools are used internally as well.

Acceptable and Unacceptable Use

Acceptable and unacceptable use of AI including generative AI will of course vary depending on many factors such as the purpose and capabilities of the AI system, the applicable laws and regulations, the data used to train the model, the size, composition, and vulnerability of the intended user base, the inherent risk, potential for misuse and its impact, mode of deployment and control retained over the system, presumed effectiveness of security measures, and ethical considerations.

Given these factors to consider, an important element of acceptable use is a prior assessment of the potential impact, whether that is in regard to a procured AI tool, or a system developed in-house (see the section on Annex III below).

A pre-determined list of acceptable and unacceptable AI applications will be informative. It will be helpful to make the list of unacceptable uses as exhaustive as possible to provide clarity to users. The list of acceptable uses need not be exhaustive, but clear guidelines should be included for approval processes for AI uses that are not included in the list of pre-approved systems, distinguishing between high- and low-impact AI systems (see section on Annex II below) and procurement and development.

When deciding on the acceptable and unacceptable use of AI in your organization, consider the security risks of using general-purpose AI tools available on the open internet. Evaluate how these tools can be used and implement necessary controls, such as prohibiting the input of personal or sensitive information. For internally deployed AI tools, ensure they are properly vetted before allowing the use of sensitive information.

Consider referring to other applicable internal policies and ensure that they are aligned with the AUP.

Unintentional Misuse

Unintentional use can be defined as inadvertently employing an AI System in a manner that results in or may result in adverse outcomes or harm, without any malicious intent. Consider requiring users to immediately report unintentional misuse and making available support and guidance regarding next steps rather than employing disciplinary measures.

Enforcement

Be transparent about the consequences following the breach of the AUP. Consider defining serious misconduct and indicating that dismissal, civil action, and criminal charges may follow it.

Annex 1: Key Risks of AI Systems

This Annex may set out, at a high-level, some risks associated with the use of AI Systems, including general-purpose AI Systems. You may draw from your risk assessment, highlighting risks that you have identified as relevant for your use case(s). Ensure making it clear that this Annex is for educational purposes only and not an exhaustive list of all possible risks.

Annex 2: Determining High-Impact AI Systems

If you distinguished between high- and low-impact AI systems in your approval process, this Annex should set out what high-impact AI systems are. You can rely on an impact assessment you conducted and/or draw inspiration from applicable (draft) legislation and widely recognized guidance materials.

Annex 3: Actions Required for High-Impact AI Systems

This Annex should set out the approval process, including the individuals responsible for requesting and granting approval, who needs to be consulted, and what documentation needs to be provided by the requestor.

It should also detail the considerations the person responsible for approval has to make and the form of documentation expected. You could consider drafting a template for the request and approval to facilitate the approval process.

Conclusion

By clearly defining what constitutes acceptable and unacceptable use, the AUP ensures that AI technologies are employed ethically and transparently, aligning with the organization's mission and the broader values it upholds. The inclusion of detailed annexes further aids organizations in understanding the risks associated with AI systems and the necessary steps for managing high-impact AI applications.

3. Guidance on Vendor Assessments for AI-Specific Risks

Organizations should carefully assess potential AI vendors, focusing on the unique risks associated with AI systems. This section provides a strategic framework for conducting thorough vendor assessments, ensuring that partnerships align with ethical standards, compliance requirements, and the organization's mission. Ideally, the proposed framework should be integrated into existing vendor assessment processes.

Prioritize social impact outcomes by embedding AI risk and benefit analysis into AI procurement decisions

When procuring AI solutions, non-profit organizations focused on social impact will want to ensure that their procurement documentation emphasizes that public benefit is the central criterion for evaluating vendors. This entails defining clear requirements and success criteria related to the societal impact outcomes that prospective AI suppliers are required to demonstrate and meet.

An important aspect of successful procurement is risk awareness and mitigation. Upon evaluating potential AI-related risks internally, procurers should (a) demand that prospective AI suppliers develop risk management and remediation strategies to address the identified risks and (b) request the potential suppliers to articulate additional risks not previously considered. Procurers can use this to compare the organization's risk assessments with the risk management plans offered by AI suppliers, ensuring adequate remediation.

Be aware that those tasked with procuring AI solutions may lack the expertise to define requirements adequately, possibly overlooking critical product or service features. The aspects of AI solutions might be too complex to specify accurately, leading to cost overruns when organizations realize the need for additional features not included initially.

Embed relevant legislation and guidelines from professional bodies into the procurement process

Non-profits must be aware of and comply with existing legislation when procuring AI solutions. While it differs slightly from jurisdiction to jurisdiction, in general, applicable laws seek to govern the use of AI systems in a few key ways:

1. **Transparency:** Organizations using AI for automated decisions are typically required to explain how and why certain decisions are made.
2. **Notice and Disclosure:** Organizations should inform individuals and patients accordingly if they interact with an AI system, such as a medical chatbot or an AI Scribe.
3. **Right to Contest:** Where AI is used to make decisions about an individual's access to care or prioritization, they should be able to challenge decisions made.

4. **Data Protection:** Organizations must ensure that any data use, including AI, has appropriate protections depending on its sensitivity and de-identification requirements.

5. **High-impact AI Systems:** Although not yet passed into law, Canada's draft *Artificial Intelligence and Data Act* (AIDA) proposes that organizations deploying "high-impact" AI systems must establish AI governance mechanisms. These include conducting risk, impact, and bias assessments, testing mitigation measures, ensuring human oversight and maintaining written accountability frameworks detailing roles, responsibilities, and management of risks. Ongoing monitoring for bias, accuracy, and effectiveness are further requirements mandated by the act.

Table A below sets out in more detail the privacy and AI regulations applicable in Canada, or expected to be applicable soon, and their implications for procurers of AI systems. While the CPPA referenced in the table is not directly applicable to non-profits, alignment with this legislation will likely ensure responsible AI governance by non-profits as well.

Table A: Privacy and AI Regulations and Implications for AI Procurers

Note: Regulations are constantly evolving and may change. Readers should verify the latest AI regulations in their jurisdictions to ensure they are up to date.

Act and Provision	Implications for Procurers
Explainability	
The Digital Charter Implementation Act, 2022, Part I: Consumer Privacy Protection Act (CPPA) If the organization has used an automated decision system to make a prediction, recommendation, or decision about the individual that could have a significant impact on them, the organization must, on request by the individual, provide them with an explanation of the prediction, recommendation, or decision. (s. 63 (3)) The act goes on to outline that the explanation must indicate the type of personal information that was used to make the prediction, recommendation, or decision, the source of the information and the reasons or principal factors that led to the prediction, recommendation, or decision. (s. 63 (4)))	AI procurers should verify that their AI solutions meet the required explainability standards. This means ensuring the AI is sufficiently explainable to offer clear explanations for its predictions, recommendations, or decisions regarding individuals when needed. Suppliers must also disclose which personal information was used to make these decisions.
Quebec's Bill 3, *An Act respecting health and social services information and amending various legislative provisions* A body must inform the person concerned, at the latter's request (s. 65): - of the information used to render the decision; and - of the reasons and the principal factors and parameters that led to the decision.	

Transparency and Notice of Interaction	
The Digital Charter Implementation Act, 2022, Part I: Consumer Privacy Protection Act (CPPA) An organization must make readily available, in plain language, information that explains the organizations' policies and practices put in place to fulfill its obligations under this Act. (s. 62 (1)). Section 62(2) provides detail on how to satisfy the transparency requirement under s. 62(1).	AI procurers should clearly inform individuals in plain language when interacting with an AI system. Additionally, there should be a straightforward process for end-users to update or correct their personal information. Non-profits must collaborate with AI suppliers to guarantee the solutions can fulfill these obligations.
Quebec's Law 25, *An Act to modernize legislative provisions as regards the protection of personal information* Section 12.1 requires organizations to inform individuals when their personal information is used to render a decision based exclusively on automated processing of such information, no later than at the time the individual is informed of the decision itself. Section 12.1 also requires that upon request, organizations must inform individuals about whom such a decision has been made of the personal information used to render the decision; of the reasons and the principal factors and parameters that led to the decision; and the right to have the personal information used to render the decision corrected.	
Quebec's Bill 3, *An Act respecting health and social services information and amending various legislative provisions* A body that uses information it holds to render a decision based exclusively on automated processing of the information must inform the person concerned accordingly, not later than at the time it informs the person of the decision. (s. 65)	
De-identification of Data	
The Digital Charter Implementation Act, 2022, Part I: Consumer Privacy Protection Act (CPPA) An organization that de-identifies personal information must ensure that any technical and administrative measures applied to the information are proportionate to the purpose for which information is de-identified and the sensitivity of the personal information (s. 74) With very few exceptions enumerated in the CPPA, an organization must not use information that has been de-identified alone or in combination with other information, to identify an individual.	AI suppliers should be required to implement measures that comply with data security, privacy, and retention laws. This includes protecting information through de-identification and anonymization, when necessary, in line with regulatory standards. Suppliers must adopt technical and administrative safeguards appropriate to the sensitivity of the data powering their AI solutions. Furthermore, data no longer needed for the AI system must be disposed of per data retention guidelines.
Quebec's Law 25, *An Act to modernize legislative provisions as regards the protection of personal information* The only mention of de-identification related to AI concerns the ability to use data without consent. When data is for purposes consistent with those for which it was collected (s. 12 para. 2 (1)) or that its use is necessary for study or research purposes or for the production of statistics and if the information is de-identified (s. 12 para. 3) then said data can be used without consent. See section 4.2 for details on this exception and the definitions of "consistent purposes" and "de-identified" information.	

Quebec's Bill 3, *An Act respecting health and social services information and amending various legislative provisions* A body must not keep the information it holds beyond the time required to achieve the purposes for which it was collected or used. (s. 16) At the end of the preservation period for the data, a body holding information must destroy or anonymize it. (s. 111) Information is anonymized if it is, at all times, reasonably foreseeable in the circumstances that it irreversible no longer allows the person concerned to be identified, even indirectly. (s. 111) Such anonymization must be carried out according to generally accepted best practices and according to the criteria and terms determined by a regulation made under section 73 of the Act respecting Access to documents held by public bodies and the Protection of personal information.	
Contestability and Redress	
The Digital Charter Implementation Act, 2022, Part I: Consumer Privacy Protection Act (CPPA) Nothing specific to AI applications; however, an individual may make a complaint or request information from an organization with respect to its compliance with the CPPA. The organization must respond to any complaint or request it receives. (s. 73 (1))	Depending on the jurisdiction, AI suppliers must bake in the capability for end-users to correct information about them promptly and simply.
Quebec's Law 25, *An Act to modernize legislative provisions as regards the protection of personal information* Law 25 appears to be silent with respect to requiring consent for AI decision-making; but it does however allow the individual with an individual right to "to submit observations to a member of the personnel of the organization who is in a position to review the decision made by automated means".	
Quebec's Bill 3, *An Act respecting health and social services information and amending various legislative provisions* A body must inform the person concerned, at their request of the right of the person concerned to have the information used to render the decision rectified. The person concerned must be given an opportunity to submit observations to a member of the body's personnel or a professional practicing his or her profession within the body who is in a position to review the decision. (s. 65)	

Requirements for High-impact AI Systems	
The Digital Charter Implementation Act, 2022, Part III: Artificial Intelligence and Data Act ("AIDA") Section 11 outlines requirements for those managing the operations of a high-impact AI system, which includes AI systems deployed in a healthcare context. A person who manages the operations of a high-impact system must, a. ensure that the requirements set out in paragraphs 10(1)(a) to (h) are met and keep the records referred to in paragraphs 10(2)(a) to (c), if there are reasonable grounds to believe that any of the acts described in paragraphs 10(1)(a) to (h) have not been accomplished in respect of the system; b. establish measures to identify, assess and mitigate the risks of harm or biased output that could result from the use of the system, in accordance with the regulations; c. carry out tests of the effectiveness of the mitigation measures that are established under paragraph (b), in accordance with the regulations; d. ensure that humans are, in accordance with the regulations, overseeing the system's operations; e. establish measures allowing users to provide feedback on the system's performance, in accordance with the regulations; f. in the time and manner prescribed by regulation, publish on a publicly available website a plain-language description of the system that includes the following information: o how the system is being used; o the types of output that it generates; o the mitigation measures established under paragraph (b); any other information that may be prescribed by regulation.	AI procurers must collaborate with suppliers to ensure compliance with the operational requirements of high-impact AI systems. These requirements encompass record-keeping, development of risk mitigation strategies, testing and validating the effectiveness of these strategies, implementing human oversight (human-in-the-loop), setting up mechanisms for user feedback on system performance, and disseminating information through user-friendly platforms like websites. This information should detail the AI system's usage, the types of outputs it produces, and the risk mitigation measures in place.

Procurers must ensure that AI systems comply with these principles and adhere to procurement guidelines that vary from sector to sector. Designing a compliance checklist like the one provided in Table B will streamline the process and ensure compliance.

Table B: Checklist for AI Procurers

Dimension	Questions for AI procurers to consider
1. Privacy	1. Will your organization provide any data that could contain personal or sensitive information to the AI supplier? • If so, a detailed review of the supplier's privacy program and data governance policies should be conducted. 2. Has the supplier performed an AI liability assessment? 3. Has the AI system been trained on personal information?

	4. Does the supplier make any representations about the use of personal information collected during the use of the AI system by the non-profit? 5. How will the personal information collected or used throughout the lifecycle of the AI system be safeguarded? 6. Will the solution leverage privacy enhancing technologies to safeguard user personal information?
2. Bias and Non-discrimination	7. Has the supplier described possible sources of bias in the data used to train and develop the AI system(s)? If so, have they clearly articulated the steps taken to remediate identified bias? 8. Has the supplier described mechanisms to prevent biased or incomplete data from being used as an input to the AI system(s)? 9. Has the supplier adequately described bias checking procedures, including requirements for periodic bias reviews of the AI system(s) to identify and remediate potential abnormalities? 10. Has the supplier conducted a third-party audit of their AI system(s) to ensure it is free from bias?
3. Explainability	11. Are the explanations provided by the AI system consistent with our organization's expectations? 12. Has the supplier provided guidance and explanations on how the results from the AI system(s) should be interpreted? 13. Has the supplier highlighted the key inputs to their AI system(s) and how they affect the outputs? 14. Is the supplier willing to use additional technology solutions (sometimes referred to as XAI "explainable AI tools"), to increase the explainability of the AI system(s)?
4. Transparency and Knowledge Transfer	15. Is the supplier willing to provide documentation on any of the following aspects? • How to correctly implement and use the AI system(s) • Intended, unintended, or restricted uses for the AI system(s) • Relevant performance criteria, definitions, and metrics • Specifications for (in)appropriate inputs and data for use by the AI system(s) • Installation and use instructions • Known risks and limitations of using the system • Malicious or inappropriate uses that may emerge under conditions of foreseeable misuse • Risk mitigation procedures, tools, or practices that can be made available to our organization
5. Safety and Security	16. Has the supplier included human oversight and intervention mechanisms? If not, is our organization equipped to perform these duties? 17. Has the supplier obtained any cybersecurity certification (e.g., ISO 27001, SOC 2)? 18. Has the supplier assessed the AI system(s) to ensure they are resilient against malicious attempts by unauthorized users to exploit system vulnerabilities? 19. Has the supplier described a disaster recovery and continuity plan should the AI system(s) be rendered unavailable?

6. Accuracy and Effectiveness	20. Has the supplier conducted a third-party audit of the AI system(s) to validate their accuracy and effectiveness? If so, obtain and analyze a copy of the findings report. 21. Has the supplier adequately articulated steps taken to ensure the AI system(s) are resilient against errors, faults, or inconsistencies that may occur within the system or its operating environment, especially if the AI system(s) interact with natural persons or other systems? 22. Has the supplier described the key performance indicators and minimum performance metrics that the AI system(s) must meet? 23. Has the supplier described an approach to actively monitor and track performance of the AI system(s) to identify degraded performance, accuracy, or effectiveness? 24. Has the supplier provided evidence that relevant experts were involved in the development and testing of the AI system(s)? 25. Has the supplier identified potential data gaps or shortcomings and described how they were remediated? 26. Has the supplier described a re-training plan for the AI system(s)?
7. User Experience	27. Has the supplier clearly articulated how the proposed AI system(s) will lead to benefits for users? 28. Has the supplier acknowledged and understood your organizational values and principles that are guiding the development and use of the proposed AI system(s)? 29. Does the supplier allow for redress in instances where individuals may be negatively affected? 30. How are updates to the algorithms and data handling practices conducted and communicated to users? 31. Will there be an audit trail that records modifications to the system’s operation and performance?
8. Accountability	32. Has the supplier provided guidelines for end-users on how to effectively integrate the AI solution into practice? 33. Has the supplier implemented safeguards or mechanisms to prevent stakeholders from becoming overly reliant on the tool, which could lead to complacency? 34. Does the supplier have a track record of providing reputable services in the non-profit sector? 35. Can the supplier demonstrate that the AI system(s) were developed in a diverse, multidisciplinary team? 36. Can the supplier demonstrate that the AI system(s) being procured have been designed in alignment with responsible AI practices and standards (e.g., adhering to ISO/IEC 42001)? 37. Has the supplier implemented any of the following? • AI ethics policy or principles • AI risk management policy or procedure • Ethical AI development procedures • AI impact/risk assessment procedures • AI acceptable use policy • AI incident response or management procedure

Conclusion

This section equips non-profits with the necessary tools to effectively assess AI vendors, ensuring that chosen partners align with ethical standards, compliance requirements, and the organization’s mission. By systematically evaluating AI-specific risks, non-profits can make informed decisions that safeguard their interests and enhance their ability to achieve impactful outcomes.

4. Guidance on the TORs for an AI Ethics Committee (Mandate of the Committee)

The establishment of an AI Ethics Committee is crucial for ensuring that non-profit organizations uphold ethical standards in the development and application of AI. This section provides guidance on creating the Terms of Reference (TORs) for such a committee, outlining its mandate, responsibilities, and operational framework.

Purpose of the AI Ethics Committee

The primary purpose of the AI Ethics Committee is to oversee and guide the ethical deployment of AI technologies within the organization. This includes reviewing AI projects for ethical concerns, advising on AI policy, and ensuring compliance with both internal standards and external regulations.

Key Responsibilities

1. **Ethical Review and Approval**: The committee is responsible for reviewing all AI initiatives to ensure they align with ethical guidelines and organizational values. This involves assessing potential risks and benefits, considering the impact on various stakeholders, and providing approval or recommendations for modifications.

2. **Policy Development**: The committee should play a central role in developing and updating policies related to AI use. This includes creating guidelines that address data privacy, security, fairness, and transparency.

3. **Training and Awareness**: It is imperative that the committee oversees the development of training programs for employees about AI ethics. This helps ensure that all team members are aware of ethical considerations and can identify potential issues in their work.

4. **Monitoring and Compliance**: The committee must monitor ongoing AI projects to ensure they continue to comply with ethical standards post-implementation. This includes regular audits and adapting policies as technologies or circumstances evolve.

5. **Stakeholder Engagement**: Engaging with diverse stakeholders, including beneficiaries, donors, and technical experts, is essential to gather a wide range of perspectives and insights on the ethical use of AI.

Composition of the Committee

- **Diverse Expertise**: The committee should include members with diverse backgrounds and expertise, including ethics, law, technology, and the specific sectors relevant to the organization’s mission.

- **Independent Members**: Including external experts or academics can provide an impartial perspective and enhance the credibility of the committee.

Operational Framework

- **Regular Meetings**: The committee should meet regularly, with the frequency of meetings adjusted based on the volume and complexity of AI projects under review.
- **Transparent Procedures**: All procedures and decisions should be documented transparently to build trust and accountability within and outside the organization.
- **Dynamic Adaptation**: The TORs should allow for flexibility to adapt to new ethical challenges and technological developments as AI evolves.

Conclusion

The TORs for an AI Ethics Committee should provide a clear, structured framework for overseeing the ethical use of AI within a non-profit organization. By defining the purpose, responsibilities, and operational procedures of the committee, non-profits can ensure that their AI initiatives are both innovative and aligned with ethical standards. This proactive approach is vital for maintaining trust and integrity in the rapidly evolving field of artificial intelligence.

5. Guidance on AI Literacy and Skills Training

Digital and AI literacy pose significant challenges for organizations aiming to integrate AI technologies and tools into their operations. However, effective skills training is crucial for the successful implementation of AI. Transitioning to an AI-driven organization requires time and often involves overcoming the challenges posed by existing cultural norms and practices.

This section provides actionable guidance on how non-profits can develop a robust training and education curriculum that fosters responsible AI adoption. By building a foundation of trust through education, non-profits can navigate the complexities of AI integration more smoothly, ensuring that all team members are equipped and confident to leverage AI technologies effectively.

Additionally, employing change management techniques is essential to facilitate this transition. Change management involves strategies such as clear communication of the benefits and impacts of AI, engaging stakeholders throughout the process, providing continuous support and training, and addressing resistance to change. These techniques help ensure that the shift to AI is well-received, and that staff are prepared to embrace new technologies and processes.

Understanding AI Skills Needs

Begin by assessing the existing level of AI understanding among your organization's members. Surveys, interviews, and focus groups can help identify knowledge gaps and areas where further education is needed.

Based on the assessment, define clear learning objectives and outcomes for different roles within the organization, ensuring that everyone has the necessary knowledge to interact with AI tools relevant to their work.

Proposed Modules: Non-profits should prioritize multi-disciplinary spaces for colleagues from different departments to come together to learn about and discuss AI (reflecting diverse perspective and experiences). Organizations just getting started may want to start with public resources, such as Microsoft's Digital Skills Center for Nonprofits. Additional module ideas include:

- An Introduction to Artificial Intelligence: Provide an overview of key AI concepts and the differences between AI, machine learning, and deep learning. Seek to understand existing and future potential applications of AI and explore emerging market trends and advancements.

- Ethical Considerations for AI: Explore the ethical implications of AI on society and communities that your non-profit serves. Understand the principles of responsible AI and how to address biases and discrimination in AI algorithms.

- Developing an AI Strategy: Provide practical guidance on developing an AI strategy aligned with organizational goals and objectives and assessing organizational readiness for AI adoption.
- Effective AI Governance and Risk Management: Explore how to establish required governance frameworks for AI projects and how to ensure compliance with relevant regulations and legal considerations. Review the potential risks of AI and how to identify and mitigate them effectively.
- AI Project Management and Implementation: Review best practices for how to responsibly design, develop, deploy, and manage AI-powered projects. Discuss the stages of the AI life cycle and the best practices surrounding AI system performance and ongoing monitoring.
- AI and Organizational Transformation: Understand the organizational impact of AI adoption and identify workforce implications and skills required for AI integration. Explore strategies for upskilling and reskilling employees and change management strategies.

There is, of course, a space for department specific training. For example, there are different training considerations for the IT team than the legal team. But the more cross-functional collaboration on AI, the more exposure each department has to the other and the easier it is to develop a common understanding of terms, concepts, and approaches to responsible AI. This leads to faster innovation and adoption.

Sources for curriculum development include:

- Microsoft Philanthropies nonprofit skilling
- Nonprofit training, courses & resources | Microsoft for Nonprofits
- NIST Trustworthy & Responsible Artificial Intelligence Resource Center. https://airc.nist.gov/Home
- OECD Resources on AI. https://oecd.ai/en/resources
- Partnership on AI. https://partnershiponai.org/resources/
- AI Now Institute. https://ainowinstitute.org/

6. Guidance on Terminology

Effective communication within an organization about AI relies on a shared understanding of key terms and concepts. Establishing a common vocabulary, particularly one that aligns with international standards like ISO/IEC 22989:2022(E), ensures that discussions about AI across different departments are clear and productive. This section highlights the importance of such standardized terminology and proposes definitions for salient AI concepts, drawing from the international standards.

Importance of Standardized Terminology

Uniform terminology helps in mitigating misunderstandings and clarifies the scope and implications of AI technologies within organizational strategies. It aids in:

- Ensuring clarity in policy documentation and compliance requirements.
- Facilitating effective cross-departmental communication and collaboration on AI projects.
- Aligning with global standards to enhance credibility with international partners and stakeholders.

Key AI Concepts and Definitions:[1]

1. **Artificial Intelligence (AI)**: A discipline engaged in research and development of mechanisms and applications of AI systems (Definition 3.1.3)

2. **AI Tool (or AI System)**: An engineered tool or system that generates outputs such as content, forecasts, recommendations or decisions for a given set of human-defined objectives.

3. **Machine Learning**: A process of optimizing model parameters through computational techniques, such that the model's behaviour reflects the data or experience.

4. **Model Parameter**: An internal variable of a model that affects how it computes its outputs.

5. **Neural Network**: A network of one or more layers of neurons connected by weighted links with adjustable weights, which takes input data and produces an output.

[1] Sources: ISO/IEC 22989:2022(E) (https://lnkd.in/eEcmKivP); iapp's Key Terms for AI Governance (https://iapp.org/resources/article/key-terms-for-ai-governance/); NIST Glossary (https://csrc.nist.gov/glossary)

6. **Algorithmic or Computational Bias**: A systematic error or deviation from the true value of a prediction that originates from a model's assumptions or the input data itself.
7. **Data Governance**: A set of processes that ensures that data assets are formally managed throughout the enterprise. A data governance model establishes authority and management and decision-making parameters related to the data produced or managed by the enterprise.