

Self-Assessment Tool with Guiding Questions for Responsible AI approach

The AI Policy Lab (AIPL) at Umeå University is committed to conducting fundamental, long-term, evidence-based research, ensuring that we provide transformative insights to diverse stakeholders nationally and globally. The Lab's goal is to foster an informed, responsible and sustainable integration of AI in society.

This tool is designed to support **internal dialogue** within organisations considering the adoption of AI. It focuses mainly not on compliance or technical detail, but on values, governance, ethics and broad organisational readiness (asking about why, who, what, how and where). The primary goal is to support decision-makers in answering **Question Zero**: "Under what conditions should an AI system be adopted, if at all?"

The tool is built for **group settings** (e.g., cross-department/unit meetings, workshops) that include technical, HR, legal and diverse user and stakeholder perspectives. It is not an audit, compliance checklist or a test. It is a **self-reflection instrument** that may be used iteratively to support organisational discussions.

Who Should Use This Tool?

- Primary audience: Cross-functional groups inside organisations (e.g., IT, legal/compliance, operations, ethics, HR, management), whenever relevant, including outside stakeholders.
- Ideally, the discussion is moderated by a neutral facilitator who guides the process, ensures that the conversation remains structured, and helps create space for diverse perspectives to be heard.
- Responses should be provided collectively by the group, reflecting the organisation's shared understanding and position rather than individual opinions.

We strongly encourage to include a diversity of perspectives, especially those who may be directly affected by the AI/technology in unintended ways (e.g., people with disabilities, minority groups, end-users, frontline workers). We also highly recommend that this assessment tool is used **before any procurement, development or deployment** decision is finalised. Doing so means avoiding additional costs to address challenges later down the line due to AI tools being adopted in suboptimal ways.

How to Use This Tool

1. Organise a meeting/workshop with a cross-unit group of people inside your organisation.
2. The sections (why, who, what, how, where) are meant to be explored as different aspects of responsible AI. However, they do not have to be addressed in order. Answers in one section often inform answers in another. You may find that exploring "What" raises new questions about "Why" or that decisions about "How" require you to revisit "Who". Feel free to move between sections as needed.
3. **Use this as a reflective, iterative process.** Initial responses can be revisited after pilot projects, external consultation or after hearing from stakeholders you had not initially considered.
4. Try to reach consensus, but record areas of disagreement as potential risks, blind spots or areas where additional input is required.

5. Document your responses in this form. It may be useful to return to the tool at later stages to adjust priorities or assess progress.
6. Consider integrating this tool with existing organisational processes (e.g., risk management, IT governance, privacy impact assessments, etc.).

This tool neither prescribes decisions, nor enforces outcomes. What you do with the answers and how responsibilities are assigned is part of your internal AI governance and accountability framework.

Time and Commitment

A full session using the tool may take between **1.5 to 3 hours**, depending on the number of participants and level of detail. Prioritised sections or rapid rounds can be used for shorter planning meetings.

Important disclaimer. If you plan to use the questions below (in original or modified form) within your organisation(s), we kindly ask to **cite** the original source below:

Titareva, T., Carli, R., Dahlgren Lindström, A., Dignum, V., Fabris, B., Ericson, P., & Tucker, J. (2026). <i>Self-Assessment Tool with Guiding Questions for Responsible AI Approach by AI Policy Lab (AIPL)</i> . Test version 3 from March 24, 2026. Umeå University.
--

If you have any questions or recommendations, please contact us at: contact@aipolicylab.se.

Section A: Why (Motivation)

This section aims to explore the reasons and motivations for the AI system's choice, as well as to understand the problem that the system is supposed to address.

A1: Why do you plan to adopt the AI system?

A2: What problem(s) is your organisation trying to solve with a new AI system (e.g., expected outcomes and functions)?

A3: What are the available alternatives to using AI system, including human, other technical non-AI solutions, etc.?

A4: Why do you think AI is necessary for your organisation's workflow? Please discuss your answer with each other.

If you have addressed this block of questions with your group, how confident are you that your unit/organisation has considered the main elements of responsible AI adoption **with respect to the motivation for an AI system's adoption** (1 - not confident at all, 5 - extremely confident)?

Please circle: 1 2 3 4 5

Section B: Who (Stakeholder Mapping)

This section aims to map the possible (internal and external) stakeholders involved in the adoption process of the AI system. It should not consider only direct stakeholders (such as organisational leaders, employees and customers) but also acknowledge indirect stakeholders who may be affected (such as minorities, socially marginalised individuals and other vulnerable populations). It explores who is involved or impacted, either through direct or indirect interaction with the AI system or through social and ethical consequences of its use.

B1: Which stakeholders do you think would be impacted by your unit/organisation adopting the AI system?

B2: Which stakeholders have not been included in the previous answer (e.g., possibly impacted underrepresented communities, marginalised groups, less visible people, people with disabilities, younger/older people, others)?

B3: Does the AI system offer an opt out option for all impacted stakeholders?

- a) Yes, with alternative
- b) Yes, with no alternative
- c) No
- d) I don't know
- e) Other (please explain):

B4: Which stakeholders are asked for consent in case of the adoption of the AI system?

B5: Who currently does the work that is planned to be performed by the new AI system?

B6: Which stakeholders could benefit if the AI system is adopted and how?

B7: Which stakeholders could potentially experience any risks/harms after adopting the AI system and how (e.g., financial, labour, discrimination, etc.)?

If you have addressed this block of questions with your group, how confident are you that your unit/organisation has considered the main elements of responsible AI adoption **with respect to the stakeholder mapping phase** (1 - not confident at all, 5 - extremely confident)?

Please circle: 1 2 3 4 5

Section C: What (Type of an AI system)

This section aims to identify and define the type of AI system that is being considered (e.g., a purchased service, an internally developed solution, etc.). It also examines the potential risks and vulnerabilities associated with that choice (such as technical, ethical, privacy or data misuse). The questions explore what type of system is being considered and the possible consequences of that choice.

Technical/IT employees in your organisation are essential to answer the questions below.

C1: What kind of AI system/method/tool is your unit/organisation planning to adopt (e.g., system: a specific chatbot for client service, method: transformers, tool: specific large language model)?

C2: Are you planning to develop the AI system in-house, procure an existing system (e.g., existing large language model, predictive analytics tool, fraud detection tool, financial risk assessment tool, etc.) or procure a custom system (e.g., through outsourcing, consultants, etc.)?

- a) In-house
- b) Procure existing system
- c) Procure custom system
- d) Hybrid case (e.g., pre-trained model + custom interface)
- e) Other (please explain):

C3: Does this AI system, method, or tool fall into one of the following categories: generative AI (i.e., reproducing text, video, audio, etc.), predictive (i.e., making predictions about the future based on past patterns), categorising (e.g., image analysis, topic or sentiment analysis, etc.) or a mix of those?

- a) Generative
- b) Predictive
- c) Categorising
- d) I am not sure
- e) Other (please explain):

C4: Will the AI system be off the shelf or the one built for the specific problem?

- a) General purpose that we will adapt to our needs
- b) Specialised by someone else in a similar sector
- c) Specialised by us for our problem or built from scratch
- d) Other (please explain):

If you have addressed this block of questions with your group, how confident are you that your unit/organisation has considered the main elements of responsible AI adoption **with respect to the selection of the type of an AI system to possibly employ** (1 - not confident at all, 5 - extremely confident)?

Please circle: 1 2 3 4 5

Section D: How (The potential adoption process of an AI system)

This section aims to review the existing workflow processes that might be influenced by an AI system. Certain aspects need to always be considered: security, human oversight and transparency. Otherwise, if these elements cannot be ensured, the system should not be deployed. It looks into how the AI system could operate within the workflow in a responsible way.

Technical/IT employees in your organisation are essential to answer the questions below.

D1: How are you planning to deploy the AI system in your current context?

D2: How will the existing workflow change with the integration of the AI system (introducing new tasks, replacing technology(-ies), automating process(es), etc.)?

D3: How do you plan to monitor and analyse the new AI system's outputs and performance (e.g., using existing tools/systems, implementing new ones)?

D4: How will you ensure the security of your organisation's and your clients' data (compliance with privacy and security requirements)?

D5: How will impacted users (both internal and external) be able to see and understand the steps by which the AI system arrived at the conclusions?

D6: What measures or strategies are implemented to identify and mitigate potential negative consequences of the AI system's deployment?

D7: Who holds responsibility and authority for overseeing these risk mitigation measures and responding if issues arise?

D8: What systems and structures do you have in place for addressing complaints or mistakes related to the AI system's outputs, including the protection of "whistleblowers"?

If you have addressed this block of questions with your group, how confident are you that your unit/organisation has considered the main elements of responsible AI adoption **with respect to the potential adoption process** (1 - not confident at all, 5 - extremely confident)?

Please circle: 1 2 3 4 5

Section E: Where (Location and control of an AI system)

This section explores where the AI system is hosted, managed and controlled.

Technical/IT employees in your organisation are essential to answer the questions below.

E1: Where will the AI system run and data be stored (e.g., locally on your server or cloud, someone's else server/cloud, inside the EU or outside)?

E2: Where is the AI system's provider based?

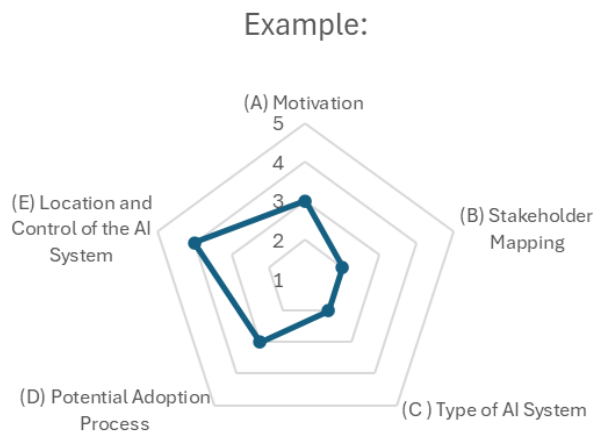
E3: Where does the training data originate from?

If you have addressed this block of questions with your group, how confident are you that your unit/organisation has considered the main elements of responsible AI adoption **with respect to the location and control for an AI system potentially adopted** (1 - not confident at all, 5 - extremely confident)?

Please circle: 1 2 3 4 5

Please draw lines in the diagram below corresponding to the numbers you selected in the last question of each section, following the example provided below. The resulting visual profile helps you quickly identify strengths and areas that may need to be strengthened or reconsidered in order to support responsible AI adoption.


This diagram can also be used to compare results across different contexts. For example, you may compare outcomes between different teams or stakeholder groups, across different AI applications, or across different moments in time (e.g., before and after implementing improvements). These comparisons can help track progress, identify gaps in perceptions or practices and guide targeted actions for improvement.



This example reflects the following scores in the previous questions:

(A) Motivation (Why?)	3
(B) Stakeholder Mapping (Who?)	2
(C) Type of AI System (What?)	2
(D) Potential Adoption Process (How?)	3
(E) Location and Control of AI System (Where?)	4

Draw lines in your organisations's diagram:



Section F: Concluding questions

This section aims to provide reflective questions before you take the final decision whether to adopt or not an AI system in your unit/organisation and context, as well as to consider whether your current processes are set for responsible AI technology adoption.

F1: Have you tested the AI system in a pilot? What did you learn?

F2: If yes, what would you adjust before the full roll-out?

F3: What do you think is the balance between the harms and benefits of the potential adoption of the AI system?

F4: What could be your next steps based on the responses you provided above?

Important disclaimer. If you plan to use the questions from this tool within your organisation(s) (in original or modified form), we kindly ask to **cite** the original source below:

Titareva, T., Carli, R., Dahlgren Lindström, A., Dignum, V., Fabris, B., Ericson, P., & Tucker, J. (2026). *Self-Assessment Tool with Guiding Questions for Responsible AI Approach by AI Policy Lab (AIPL)*. Test version 3 from March 24, 2026. Umeå University.