



Microsoft Responsible AI Impact Assessment Template

FOR EXTERNAL RELEASE

June 2022

The Responsible AI Impact Assessment Template is the product of a multi-year effort at Microsoft to define a process for assessing the impact an AI system may have on people, organizations, and society. We are releasing our Impact Assessment Template externally to share what we have learned, invite feedback from others, and contribute to the discussion about building better norms and practices around AI.

We invite your feedback on our approach:

<https://aka.ms/ResponsibleAIQuestions>

Responsible AI Impact Assessment for Project GovGPT

For questions about specific sections within the Impact Assessment, please refer to the Impact Assessment Guide.

Section 1: System Information

System profile

1.1 Complete the system information below.

System name	GovGPT – Pilot project conversational companion for Government Information
Team name	Project MARION Team

Track revision history below.

Authors	Jenna Whitman – Chief Information Security Officer / AI Governance Richie Atkinson – Solutions Lead Will Garland – Business Analyst / Prompt Engineering
Last updated	

Identify the individuals who will review your Impact Assessment when it is completed.

Reviewers	Sarah Sun – Head of AI Helena Page – Senior Legal Counsel / Privacy Officer
-----------	--

System lifecycle stage

1.2 Indicate the dates of planned releases for the system.

Date	Lifecycle stage
July 2024	Planning & analysis
August 2024	Design
August 2024	Development
August 2024	Testing
October 2024	Implementation & deployment
October-December 2024	Maintenance
December 2024	Retired

System description

1.3 Briefly explain, in plain language, what you're building. This will give reviewers the necessary context to understand the system and the environment in which it operates.

System description
<p>This pilot project is a collaboration between Microsoft, Callaghan Innovation and supported by the Whāriki Māori Business Network, with the goal of developing a Minimum Viable Product (MVP) and demo-ready version for presentation at the AI Summit on 11 September 2024. The MVP will explore the concept of creating a single AI-based service that provides information on publicly available government services through a “conversational companion” experience.</p> <p>An AI-based conversational companion will provide accurate and referenceable information about New Zealand’s public services through natural language conversations. It will use Retrieval Augmented Generation (RAG) techniques to enhance the accuracy and reliability of the generative AI model by fetching facts from open source, publicly facing website sources.</p> <p>The system runs entirely on the Microsoft Azure platform, including a global deployment of OpenAI’s current large language model.</p>

If you have links to any supplementary information on the system such as demonstrations, functional specifications, slide decks, or system architecture diagrams, please include links below.

Description of supplementary information	Link
Out of scope	

GovGPT FAQ	GovGPT FAQ.docx
Out of scope	
GovGPT Annoucement video	GovGPT demo - 720p.mp4

Released under the Official Information Act 1982

System purpose

1.4 Briefly describe the purpose of the system and system features, focusing on how the system will address the needs of the people who use it. Explain how the AI technology contributes to achieving these objectives.

System purpose

How will this address the need:

Currently, navigating the web of publicly available information on New Zealand Government services is time-consuming, convoluted, and mostly jargon-heavy English. GovGPT is a digital front-door, enabling New Zealanders to easily discover and explore the abundance of Government support available to them, in a language that they understand. Users of the tool can interact with the answers they are given, to refine them and get a better understanding of the government support and regulations that concern them.

System features:

The MVP will explore the concept of creating a single AI-based service that provides information on publicly available government services through a “conversational companion” experience.

An AI-based conversational companion will provide accurate and referenceable information about New Zealand’s public services through natural language conversations. It will use Retrieval Augmented Generation (RAG) techniques to enhance the accuracy and reliability of the generative AI model by fetching facts from open source, publicly facing website sources.

Pre-populated prompts exist on the landing page to start the experience, such as “what benefits am I entitled to as a business owner?”, and from there, users can ask their own questions to GovGPT.

Whāriki will play a crucial role in this project by contributing specific use cases, particularly related to government support and funding opportunities for Māori businesses. This collaboration aims to ensure the cultural relevance and appropriateness of the content, providing a comprehensive and valuable tool for all users.

System features

1.5 Focusing on the whole system, briefly describe the system features or high-level feature areas that already exist and those planned for the upcoming release.

Existing system features	System features planned for the upcoming release
Indexing of government websites for retrieval-augmented-generation search, so only indexed sources will be included in results.	Scheduled updates of government sources, and direct linking to source in citations.
AI generated plain English responses to user questions, grounded on indexed websites and documents, and including citations.	Real-time speech input and output.
Written input and output across all major, and some minor languages.	Potential for a digital avatar to aid in personalisation.
Generated answers grounded on the indexed websites, provided in the language and style the user requests.	Potential to include feedback mechanisms to assist in future training and relevance of outputs.

Briefly describe how this system relates to other systems or products. For example, describe if the system includes models from other systems.

Relation to other systems / products

Geographic areas and languages

1.6 Describe the geographic areas where the system will or might be deployed to identify special considerations for language, laws, and culture.

The system is currently deployed to:	9(2)(b)(ii) - Commercial Information
In the upcoming release, the system will be deployed to:	
In the future, the system might be deployed to:	

For natural language processing systems, describe supported languages:

The system currently supports:	All major languages in written format, though languages with less available data or speakers have lower levels of performance.
In the upcoming release, the system will support:	English language spoken outputs
In the future, the system might support:	Fluent spoken Reo Māori and language switching.

Deployment mode

1.7 Document each way that this system might be deployed.

How is the system currently deployed?	First release will be embedded on Callaghan Innovation’s website as a public pilot.
Will the deployment mode change in the upcoming release? If so, how?	Future releases may have a standalone website, or a mobile app.

Intended uses

1.8 Intended uses are the uses of the system your team is designing and testing for. An intended use is a description of who will use the system, for what task or purpose, and where they are when using the system. They are not the same as system features, as any number of features could be part of an intended use. Fill in the table with a description of the system’s intended use(s).

Name of intended use(s)	Description of intended use(s)
1. Public information searching by New Zealanders.	Anyone in New Zealand should be able to engage in a conversation with GovGPT to ask questions and get answers from government websites. Users can refine their questions to better understand the information, and be directed to government services most relevant to them. This use case applies to anyone looking for answers from government websites, including civilians, researchers, public servants.
2.	
3.	

Section 2: Intended uses

Intended use #1: Public information searching by New Zealanders – repeat for each intended use

Name of intended use(s)	Description of intended use(s)
Public information searching by New Zealander's.	Anyone in New Zealand should be able to engage in a conversation with GovGPT to ask questions and get answers from government websites. Users can refine their questions to better understand the information, and be directed to government services most relevant to them. This use case applies to anyone looking for answers from government websites, including civilians, researchers, public servants.

Assessment of fitness for purpose

2.1 Assess how the system's use will solve the problem posed by each intended use, recognizing that there may be multiple valid ways in which to solve the problem.

Assessment of fitness for purpose
<p>The system will provide a single tool where users can perform natural language searches on publicly available government websites. Good answers will mean the user won't have to navigate multiple complex websites, and can be almost instantly directed to an appropriate service prepared with relevant information. Compared to existing search tools, GovGPT allows for answers to be made understandable and personalised for the user, not subject to corporate, legal, or governmental jargon. GovGPT can retrieve and process relevant information for user queries much more effectively and quickly than existing tools, and offers a more controlled and manageable alternative to free AI-search tools on the market.</p>

Stakeholders, potential benefits, and potential harms

2.2 Identify the system's stakeholders for this intended use. Then, for each stakeholder, document the potential benefits and potential harms. For more information, including prompts, see the Impact Assessment Guide.

Stakeholders	Potential system benefits	Potential system harms
1. 9(2)(g)(i) - Free and Frank opinions		
2.		
3.		
4.		
5.		
6.		

7.	9(2)(g)(i) - Free and Frank opinions	
8.		
9.		
10.		

Released under the Official Information Act 1982

Stakeholders for Goal-driven requirements from the Responsible AI Standard

2.3 Certain Goals in the Responsible AI Standard require you to identify specific types of stakeholders. You may have included them in the stakeholder table above. For the Goals below that apply to the system, identify the specific stakeholder(s) for this intended use. If a Goal does not apply to the system, enter "N/A" in the table.

Goal A5: Human oversight and control

This Goal applies to all AI systems. Complete the table below.

Who is responsible for troubleshooting, managing, operating, overseeing, and controlling the system during and after deployment?	For these stakeholders, identify their oversight and control responsibilities.
Project Marion team (the authors listed in this Assessment as well as other Microsoft, Callaghan Innovation and Whariki team members)	Total responsibility of the tool, its safety, and decisions around scaling and continuation. Capability of control provided within the platform

Goal T1: System intelligibility for decision making

This Goal applies to AI systems when the intended use of the generated outputs is to inform decision making by or about people. If this Goal applies to the system, complete the table below.

Who will use the outputs of the system to make decisions?	Who will decisions be made about?
New Zealand public, business owners, and all potential stakeholders of indexed websites may choose to act on the outputs they are given.	Decisions will be made about the users, their families, and their business. Clear messaging (through FAQs and terms of use statements) will advise that outputs are not to be taken as advice, but as referrals to appropriate services.

Goal T2: Communication to stakeholders

This Goal applies to all AI systems. Complete the table below.

Who will make decisions about whether to employ the system for particular tasks?	Who develops or deploys systems that integrate with this system?
Individual users can choose to use the tool.	Ministries may build or improve their own information architecture to better collaborate with the tool.

Goal T3: Disclosure of AI interaction

This Goal applies to AI systems that impersonate interactions with humans, unless it is obvious from the circumstances or context of use that an AI system is in use, and AI systems that generate or manipulate image, audio, or video content that could falsely appear to be authentic. If this Goal applies to the system, complete the table below.

Who will use or be exposed to the system
N/A. It is obvious from the circumstances that an AI system is in use. GovtGPT will not contain audio, video or image outputs.

Fairness considerations

2.4 For each Fairness Goal that applies to the system, 1) identify the relevant stakeholder(s) (e.g., system user, person impacted by the system); 2) identify any demographic groups, including marginalized groups, that may require fairness considerations; and 3) prioritize these groups for fairness consideration and explain how the fairness consideration applies. If the Fairness Goal does not apply to the system, enter "N/A" in the first column.

Goal F1: Quality of service

This Goal applies to AI systems when system users or people impacted by the system with different demographic characteristics might experience differences in quality of service that can be remedied by building the system differently. If this Goal applies to the system, complete the table below describing the appropriate stakeholders for this intended use.

Which stakeholder(s) will be affected?	For affected stakeholder(s) which demographic groups are you prioritizing for this Goal?	Explain how each demographic group might be affected.
System user	Non English speakers, those with learning disabilities	The tool provides the following languages - [INSERT], and as a conversational companion style tool it can be more user friendly for those with learning disabilities than the existing websites. Where possible, users will not experience a different quality of service as these demographic groups have considered in the design of the tool. We have also worked with Whariki to provide the Reo Māori language version. There may be further opportunities to support people with learning disabilities that may not be able to read the content from the tool (for example adding an audio speech to text version)

Goal F2: Allocation of resources and opportunities

This Goal applies to AI systems that generate outputs that directly affect the allocation of resources or opportunities relating to finance, education, employment, healthcare, housing, insurance, or social welfare. If this Goal applies to the system, complete the table below describing the appropriate stakeholders for this intended use.

Which stakeholder(s) will be affected?	For affected stakeholder(s) which demographic groups are you prioritizing for this Goal?	Explain how each demographic group might be affected.

Goal F3: Minimization of stereotyping, demeaning, and erasing outputs

This Goal applies to AI systems when system outputs include descriptions, depictions, or other representations of people, cultures, or society. If this Goal applies to the system, complete the table below describing the appropriate stakeholders for this intended use.

Which stakeholder(s) will be affected?	For affected stakeholder(s) which demographic groups are you prioritizing for this Goal?	Explain how each demographic group might be affected.
System users	Indigenous and underrepresented groups in New Zealand, gender and sexual minorities, people with disabilities.	The tool will generate answers based on the information indexed in the backend from existing websites. Existing biases present in the content may be surfaced through use of the tool. These will be approached as opportunities for learning and improving the sources.

Released under the Official Information Act 1982

Technology readiness assessment

2.5 Indicate with an "X" the description that best represents the system regarding this intended use.

Select one	Technology Readiness
	The system includes AI supported by basic research and has not yet been deployed to production systems at scale for similar uses.
X	The system includes AI supported by evidence demonstrating feasibility for uses similar to this intended use in production systems.
	This is the first time that one or more system component(s) are to be validated in relevant environment(s) for the intended use. Operational conditions that can be supported have not yet been completely defined and evaluated.
	This is the first time the whole system will be validated in relevant environment(s) for the intended use. Operational conditions that can be supported will also be validated. Alternatively, nearly similar systems or nearly similar methods have been applied by other organizations with defined success.
	The whole system has been deployed for all intended uses , and operational conditions have been qualified through testing and uses in production.

Task complexity

2.6 Indicate with an "X" the description that best represents the system regarding this intended use.

Select One	Task Complexity
	Simple tasks , such as classification based on few features into a few categories with clear boundaries. For such decisions, humans could easily agree on the correct answer, and identify mistakes made by the system. For example, a natural language processing system that checks spelling in documents.
X	Moderately complex tasks , such as classification into a few categories that are subjective. Typically, ground truth is defined by most evaluators arriving at the same answer. For example, a natural language processing system that autocompletes a word or phrase as the user is typing.
	Complex tasks , such as models based on many features, not easily interpretable by humans, resulting in highly variable predictions without clear boundaries between decision criteria. For such decisions, humans would have a difficult time agreeing on the best answer, and there may be no clearly incorrect answer. For example, a natural language processing system that generates prose based on user input prompts.

Role of humans

2.7 Indicate with an "X" the description that best represents the system regarding this intended use.

Select One	Role of humans
X	People will be responsible for troubleshooting triggered by system alerts but will not otherwise oversee system operation. For example, an AI system that generates keywords from unstructured text alerts the operator of errors, such as improper format of submission files.
	The system will support effective hand-off to people but will be designed to automate most use. For example, an AI system that generates keywords from unstructured text that can be configured by system admins to alert the operator when keyword generation falls below a certain confidence threshold.
	The system will require effective hand-off to people but will be designed to automate most use. For example, an AI system that generates keywords from unstructured text alerts the operator when keyword generation falls below a certain confidence threshold (regardless of system admin configuration).
	People will evaluate system outputs and can intervene before any action is taken: the system will proceed unless the reviewer intervenes. For example, an AI system that generates keywords from unstructured text will deliver the generated keywords for operator review but will finalize the results unless the operator intervenes.
	People will make decisions based on output provided by the system: the system will not proceed unless a person approves. For example, an AI system that generates keywords from unstructured text but does not finalize the results without review and approval from the operator.

Deployment environment complexity

2.8 Indicate with an "X" the description that best represents the system regarding this intended use.

Select One	Deployment environment complexity
X	Simple environment , such as when the deployment environment is static, possible input options are limited, and there are few unexpected situations that the system must deal with gracefully. For example, a natural language processing system used in a controlled research environment.
	Moderately complex environment , such as when the deployment environment varies, unexpected situations the system must deal with gracefully may occur, but when they do, there is little risk to people, and it is clear how to effectively mitigate issues. For example, a natural language processing system used in a corporate workplace where language is professional and communication norms change slowly.
	Complex environment , such as when the deployment environment is dynamic, the system will be deployed in an open and unpredictable environment or may be subject to drifts in input distributions over time. There are many possible types of inputs, and inputs may significantly vary in quality. Time and attention may be at a premium in making decisions and it can be difficult to mitigate issues. For example, a natural language processing system used on a social media platform where language and communication norms change rapidly.

Section 3: Adverse impact

Restricted Uses

3.1 *If any uses of the system are subject to a legal or internal policy restriction, list them here, and follow the requirements for those uses.*

Restricted Uses

N/A - other than restriction to external, open-source information to inform the knowledge base.

Unsupported uses

3.2 *Uses for which the system was not designed or evaluated or that should be avoided.*

Unsupported uses

Giving advice or representation. We also recommend that users do not input any personal or commercially sensitive information into the tool.

Known limitations

3.3 *Describe the known limitations of the system. This could include scenarios where the system will not perform well, environmental factors to consider, or other operating factors to be aware of.*

Known limitations

Attempts to bypass the system prompt could enable the system to perform tasks outside of its remit. The system will not search the internet outside of the indexed material. Malicious actors may attempt to deliberately "break" the tool.

Potential impact of failure on stakeholders

3.4 *Define predictable failures, including false positive and false negative results for the system as a whole and how they would impact stakeholders for each intended use.*

Potential impact of failure on stakeholders

The system may generate incorrect or conflicting information, especially for more novel or complex requests. Attempts to circumvent the prompt may provide unexpected results. Stakeholders should verify all answers they get with the source material, and know that the output does not constitute advice nor representation.

A terms and conditions 'pop up' will be given to the end user to accept before being able to use GovGPT to inform them of limitations, a legal disclaimer and reminder not to input personally identifiable information.

--

Released under the Official Information Act 1982

Potential impact of misuse on stakeholders

3.5 Define system misuse, whether intentional or unintentional, and how misuse could negatively impact each stakeholder. Identify and document whether the consequences of misuse differ for marginalized groups. When serious impacts of misuse are identified, note them in the summary of impact as a potential harm.

Potential impact of misuse on stakeholders
Users may circumvent the prompt in order to get the tool to generate outputs that are harmful, false misrepresent information . These outputs, or hallucinated fallacies, may be shared online to defame entities or the tool itself.

Sensitive Uses

3.6 Consider whether the use or misuse of the system could meet any of the Microsoft Sensitive Use triggers below.

Yes or No	Sensitive Use triggers
No	<p>Consequential impact on legal position or life opportunities</p> <p>The use or misuse of the AI system could affect an individual's: legal status, legal rights, access to credit, education, employment, healthcare, housing, insurance, and social welfare benefits, services, or opportunities, or the terms on which they are provided.</p>
No	<p>Risk of physical or psychological injury</p> <p>The use or misuse of the AI system could result in significant physical or psychological injury to an individual.</p>
No	<p>Threat to human rights</p> <p>The use or misuse of the AI system could restrict, infringe upon, or undermine the ability to realize an individual's human rights. Because human rights are interdependent and interrelated, AI can affect nearly every internationally recognized human right.</p>

Section 4: Data Requirements

Data requirements

4.1 Define any document data requirements with respect to the system's intended uses, stakeholders, and the geographic areas where the system will be deployed.

Data Requirements

Any documents which will be, or are, indexed must be fully publicly available documents. **9(2)(b)(ii) - Commercial Information**

The domain of the data indexed will initially be limited to information which is specifically beneficial to small businesses in New Zealand, with some limited exceptions for information relevant to the Science and Innovation ecosystem in New Zealand (such as information about Callaghan Innovation's Minister, Board, etc.). As other agencies request their data be indexed, an assessment will need to be made on the relevance of this information to the overarching goal of the system.

The system must be transparent and provide the user with sources for its responses. It will also inform the user of its system prompt if asked.

9(2)(b)(ii) - Commercial Information

The system will run entirely on Microsoft's Azure platform, **9(2)(b)(ii) - Commercial Information**

- Azure App Service Plans
- Azure App Services / Web Apps
- Azure Deployment Services
- Azure Blob Storage
- Azure Cognitive Services (Azure AI Search, Azure Document Intelligence)
- Azure Monitoring Services (Azure Log Analytics, Azure Application Insights)
- Azure OpenAI (however the Azure OpenAI ChatGPT-4o model will be called from the global-standard deployment as it is not available in an Australian data centre)

For privacy reasons, no user data or personal information will be collected or stored and sessions will be cleared after each day. No training will be completed at this stage and the system will rely entirely on RAG for data veracity. This does mean that we rely on some of our stakeholders to ensure that their data is true and correct, and any terms of use will reflect this and will also reiterate that users should not put any personal, confidential or commercially sensitive information into the tool.

The outcomes delivered by this system must benefit the public interest of New Zealanders and deliver the goals as determined by other key stakeholders.

Stakeholders include:

- Government Agencies, Departments, Ministries or Crown Entities whose data is on the specific

indexed sites list OR who have requested their data be added after launch

- The Executive Leadership Team at Callaghan Innovation
- The Board of Callaghan Innovation
- Minister Hon. Judith Collins KC

Existing data sets

4.2 *If you plan to use existing data sets to train the system, assess the quantity and suitability of available data sets that will be needed by the system in relation to the data requirements defined above. If you do not plan to use pre-defined data sets, enter "N/A" in the response area.*

Existing data sets

N/A - system will not be trained, it will only use indexed data.

Section 5: Summary of Impact

Potential harms and preliminary mitigations

5.1 Gather the potential harms you identified earlier in the Impact Assessment in this table (check the stakeholder table, fairness considerations, adverse impact section, and any other place where you may have described potential harms). Use the mitigations prompts in the Impact Assessment Guide to understand if the Responsible AI Standard can mitigate some of the harms you identified. Discuss the harms that remain unmitigated with your team and potential reviewers.

Describe the potential harm	Corresponding Goal from the Responsible AI Standard (if applicable)	Describe your initial ideas for mitigations or explain how you might implement the corresponding Goal in the design of the system
Outputs containing biased information	F1, F3	The prompt can be developed over time to account for existing biases in the source information, but ultimately the sources themselves will be developed to account for demographic and other biases.

Goal Applicability

5.2 To assess which Goals apply to this system, use the tables below. When a Goal applies to only specific types of AI systems, indicate if the Goal applies to the system being evaluated in this Impact Assessment by indicating "Yes" or "No." If you indicate that a Goal does not apply to the system, explain why in the response area. If a Goal applies to the system, you must complete the requirements associated with that Goal while developing the system.

Accountability Goals

Goals	Does this Goal apply to the system? (Yes or No)
A1: Impact assessment <i>Applies to: All AI systems.</i>	yes
A2: Oversight of significant adverse impacts <i>Applies to: All AI systems.</i>	yes
A3: Fit for purpose <i>Applies to: All AI systems.</i>	yes
A4: Data governance and management <i>Applies to: All AI systems.</i>	Yes
A5: Human oversight and control <i>Applies to: All AI systems.</i>	yes

Transparency Goals

Goals	Does this Goal apply to the system? (Yes or No)
T1: System intelligibility for decision making <i>Applies to:</i> AI systems when the intended use of the generated outputs is to inform decision making by or about people.	no
T2: Communication to stakeholders <i>Applies to:</i> All AI systems.	yes
T3: Disclosure of AI interaction <i>Applies to:</i> AI systems that impersonate interactions with humans, unless it is obvious from the circumstances or context of use that an AI system is in use, and AI systems that generate or manipulate image, audio, or video content that could falsely appear to be authentic.	yes

If you selected "No" for any of the Transparency Goals, explain why the Goal does not apply to the system

This is system's outputs should not be considered advice nor representation, and so decisions should not be made on the output alone.

Fairness Goals

Goals	Does this Goal apply to the system? (Yes or No)
F1: Quality of service <i>Applies to:</i> AI systems when system users or people impacted by the system with different demographic characteristics might experience differences in quality of service that can be remedied by building the system differently.	yes
F2: Allocation of resources and opportunities <i>Applies to:</i> AI systems that generate outputs that directly affect the allocation of resources or opportunities relating to finance, education, employment, healthcare, housing, insurance, or social welfare.	no
F3: Minimization of stereotyping, demeaning, and erasing outputs <i>Applies to:</i> AI systems when system outputs include descriptions, depictions, or other representations of people, cultures, or society.	yes

If you selected "No" for any of the Fairness Goals, explain why the Goal does not apply to the system below.

The system does not allocate resources, nor advice on their allocation.

Reliability & Safety Goals

Goals	Does this Goal apply to the system? (Yes or No)
RS1: Reliability and safety guidance <i>Applies to: All AI systems.</i>	yes
RS2: Failures and remediations <i>Applies to: All AI systems.</i>	yes
RS3: Ongoing monitoring, feedback, and evaluation <i>Applies to: All AI systems.</i>	yes

Privacy & Security Goals

Goals	Does this Goal apply to the system? (Yes or No)
PS1: Privacy Standard compliance <i>Applies when the Microsoft Privacy Standard applies.</i>	Yes. Callaghan Innovation has considered this impact assessment and GovGPT against its privacy policy and determined that a PIA is not necessary.
PS2: Security Policy compliance <i>Applies when the Microsoft Security Policy applies.</i>	yes

Inclusiveness Goal

Goals	Does this Goal apply to the system? (Yes or No)
I1: Accessibility Standards compliance <i>Applies when the Microsoft Accessibility Standards apply.</i>	yes

Signing off on the Impact Assessment

5.3 Before you continue with next steps, complete the appropriate reviews and sign off on the Impact Assessment. At minimum, the PM should verify that the Impact Assessment is complete. In this case, ensure you complete the appropriate reviews and secure all approvals as required by your organization before beginning development.

Reviewer role and name	I can confirm that the document benefitted from collaborative work and different expertise within the team (e.g., engineers, designers, data scientists, etc.)	Date reviewed	Comments

Update and review the Impact Assessment at least annually, when new intended uses are added, and before advancing to a new release stage. The Impact Assessment will remain a key reference document as you work toward compliance with the remaining Goals of the Responsible AI Standard.

Released under the Official Information Act 1982

Scan this code to access responsible AI resources from Microsoft:



© 2022 Microsoft Corporation. All rights reserved. This document is provided "as-is." It has been edited for external release to remove internal links, references, and examples. Information and views expressed in this document may change without notice. You bear the risk of using it. Some examples are for illustration only and are fictitious. No real association is intended or inferred. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.