January 2022

Artificial Intelligence Governance Toolkit

AI Community of Practice

Overview

Privacy is an integral consideration in the use of Artificial Intelligence (AI) across the United States government. It can be challenging to know where and how to start privacy conversations in an agency setting and how those conversations should align with the Privacy Act's Fair Information Practice Principles (FIPPs) (see page 6). Recognizing the myriad risks that irresponsible use of AI can pose, and guided by the Government Accountability Office's (GAO) Report on Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities, the AI Community of Practice developed this *AI Governance Toolkit* to support agency leaders, privacy practitioners and others to establish a unique, comprehensive approach to data privacy, as well as diversity, inclusion, equity, and accessibility.

Risk management and governance occurs both at the 1) organizational level (e.g., goals, roles, responsibilities, risk tolerance and values) and at the 2) system level (e.g., technical specifications, processes, etc).

Successful AI and privacy governance approaches must involve a thoughtful and intentional approach to **stakeholder engagement**. AI and privacy stakeholder engagement involves diverse perspectives that involve subject matter experts (SME) including but not necessarily limited to the fields of data science, software development, infrastructure, user experience, civil rights and liberties, privacy and security, legal counsel, and risk management.

This toolkit is intended to provide you a framework that addresses privacy and governance at both the organizational and system levels. It provides suggestions for determining the right stakeholders to engage, and the types of privacy questions to ask at each phase of your development and deployment cycle. **This toolkit is not intended to be a formula, guidance or a checklist but rather a set of considerations to help determine the best way for you and your agency to approach AI. This toolkit was developed by a community of AI practitioners across multiple agencies.**

Table of Contents

Artificial Intelligence Stakeholder Map

How to Use Artificial Intelligence Stakeholder Map

The AI Stakeholder Map includes a set of stakeholders that are relevant to the AI lifecycle at your agency. We recommend inviting these agency staff to discuss how each of their roles and responsibilities align to the FIPPs and enforce the National Institute of Standards and Technology (NIST) recommended controls.

Use the tool on the next page to help determine which agency stakeholders are relevant and necessary to the AI lifecycle; which can support the FIPPs; and what artifacts you might need to develop with or for them. Feel free to add other stakeholders; adjust the size(s) and location(s) of the stakeholders; or how and where the FIPPs appear on your agency's map - be creative as you think about potential ways to organize and facilitate your discussions about AI. Aligning different stakeholders to different FIPPs or NIST controls may help your agency re-imagine governance altogether.

Here's a non-exhaustive list/key:

- CIO - Chief Information Officer
- CISO - Chief Information Security Officer
- CPO - Chief Privacy Officer
- CDO - Chief Data Officer
- Records - Agency Records Officer
- System or application owner(s)
- UX - User Experience
- PRA - Paperwork Reduction Act Officer
- OGC - Office of General Counsel
- OCR - Office of Civil Rights
- Data Scientist
- Users - both internal and external

Fair Information Practice Principles (FIPPs)

- **Access and Amendment.** Agencies should provide individuals with appropriate access to PII and appropriate opportunity to correct or amend PII.
- **Accountability.** Agencies should be accountable for complying with these principles and applicable privacy requirements, and should appropriately monitor, audit, and document compliance. Agencies should also clearly define the roles and responsibilities with respect to PII for all employees and contractors, and should provide appropriate training to all employees and contractors who have access to PII.
- **Authority.** Agencies should only create, collect, use, process, store, maintain, disseminate, or disclose PII if they have authority to do so, and should identify this authority in the appropriate notice.
- **Minimization.** Agencies should only create, collect, use, process, store, maintain, disseminate, or disclose PII that is directly relevant and necessary to accomplish a legally authorized purpose, and should only maintain PII for as long as is necessary to accomplish the purpose.
- **Quality and Integrity.** Agencies should create, collect, use, process, store, maintain, disseminate, or disclose PII with such accuracy, relevance, timeliness, an completeness as is reasonably necessary to ensure fairness to the individual.
- **Individual Participation.** Agencies should involve the individual in the process of using PII and, to the extent practicable, seek individual consent for the creation, collection, use, processing, storage, maintenance, dissemination, or disclosure of PII. Agencies should also establish procedures to receive and address individuals' privacy-related complaints and inquiries.
- **Purpose Specification and Use Limitation.** Agencies should provide notice of the specific purpose for which PII is collected and should only use, process, store, maintain, disseminate, or disclose PII for a purpose that is explained in the notice and is compatible with the purpose for which the PII was collected, or that is otherwise legally authorized.
- **Security.** Agencies should establish administrative, technical, and physical safeguards to protect PII commensurate with the risk and magnitude of the harm that would result from its unauthorized access, use, modification, loss, destruction, dissemination, or disclosure.
- **Transparency.** Agencies should be transparent about information policies and practices with respect to PII, and should provide clear and accessible notice regarding creation, collection, use, processing, storage, maintenance, dissemination, and disclosure of PII.

Artificial Intelligence Stakeholder Map

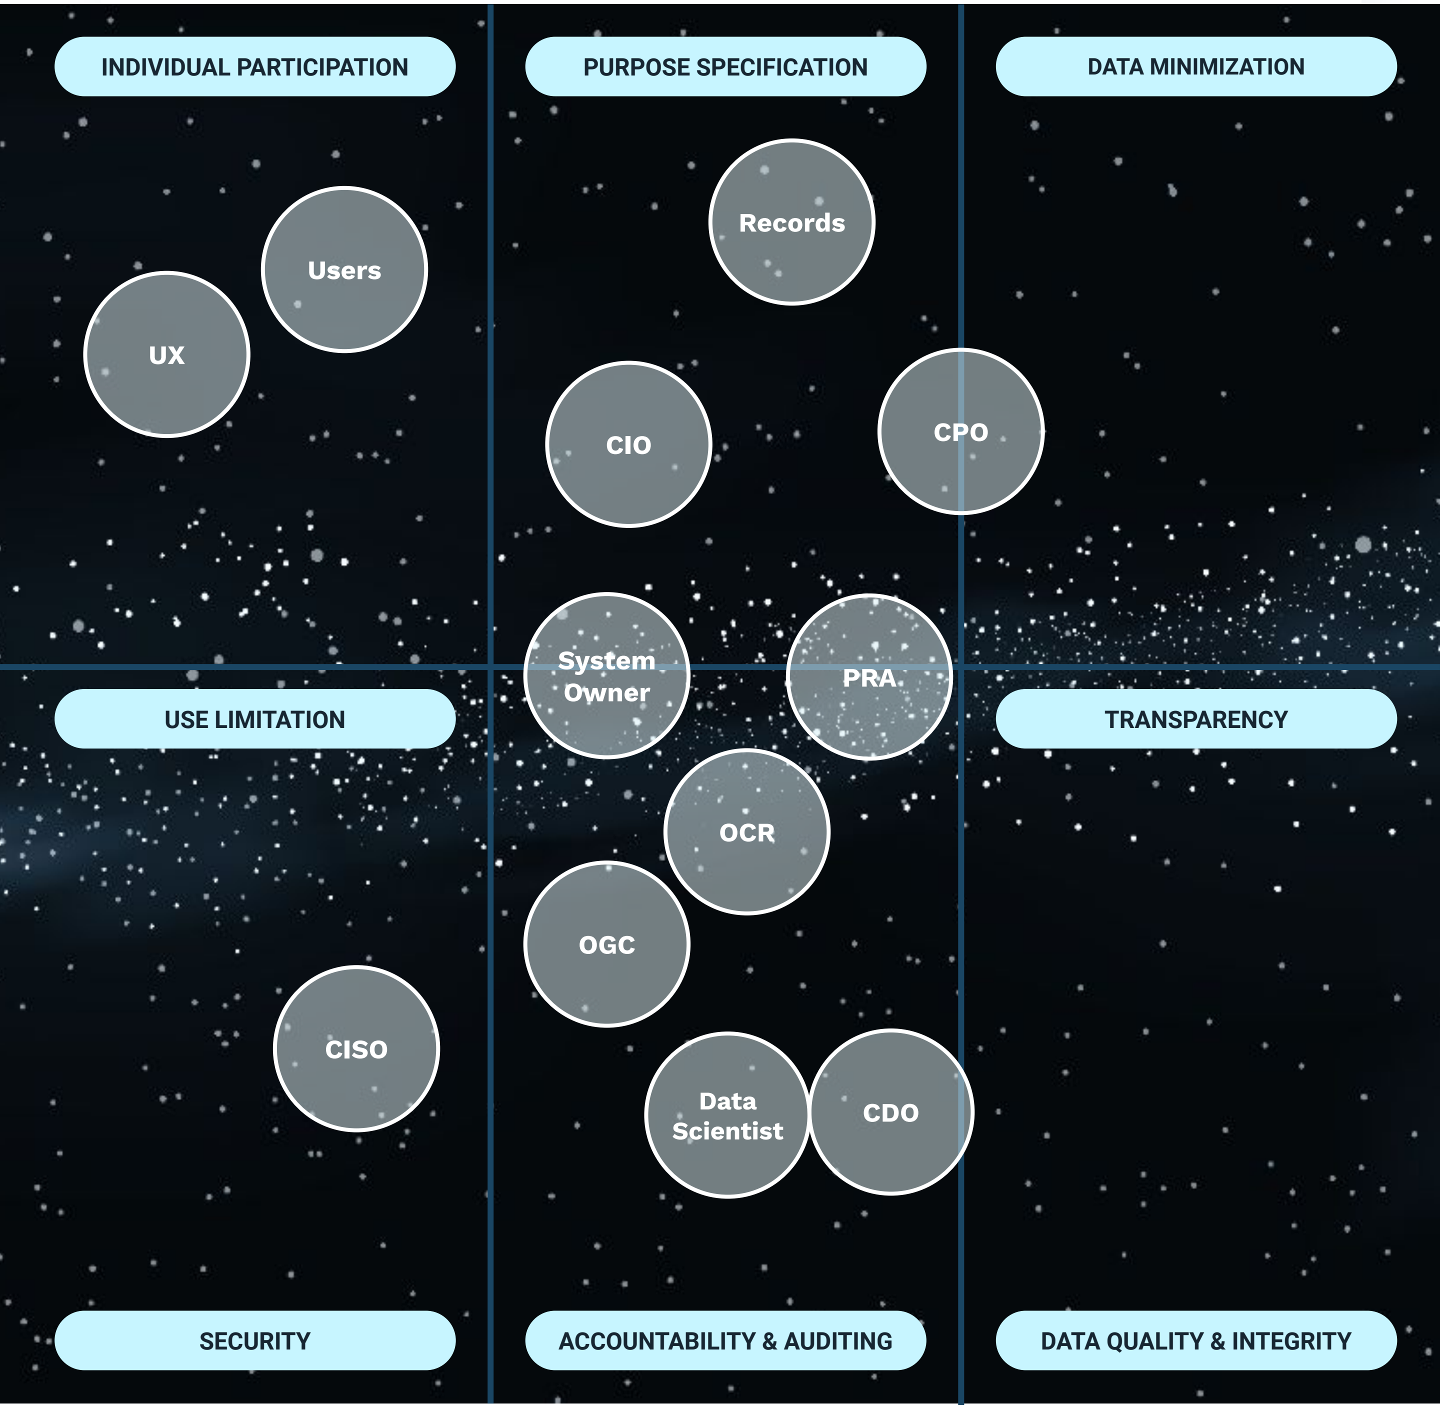

INDIVIDUAL PARTICIPATION

PURPOSE SPECIFICATION

DATA MINIMIZATION

USE LIMITATION

TRANSPARENCY

SECURITY

ACCOUNTABILITY & AUDITING

DATA QUALITY & INTEGRITY

Privacy in AI Development Lifecycle

How to Use: Privacy in AI Development Lifecycle

The Privacy in AI Development Lifecycle is a framework captures the **critical privacy questions**, informed by the FIPPs, that need to be considered throughout the AI development lifecycle.

The questions assume personally identifiable information (PII) or other controlled unclassified information (CUI) may be implicated in a potential AI project. The questions are designed to demonstrate **1) how stakeholders are supporting the FIPPs, 2) how to identify and address risks and weaknesses**, **and 3) how to foster ongoing conversation, investigation and cooperative analysis during the AI lifecycle**.

This guide is intended to be a working document that you and your team(s) can complete together as a part of a comprehensive process in understanding the privacy of your AI systems. Key Stakeholders involved in these conversations may include, but are not limited to:

- CIO - Chief Information Officer
- CISO - Chief Information Security Officer
- CPO - Chief Privacy Officer
- CDO - Chief Data Officer
- Records - Agency Records Officer
- System or application owner(s)
- UX - User Experience
- PRA - Paperwork Reduction Act Officer
- OGC - Office of General Counsel
- OCR - Office of Civil Rights
- Data Scientist
- Users - both internal and external

Privacy in AI Development Cycle

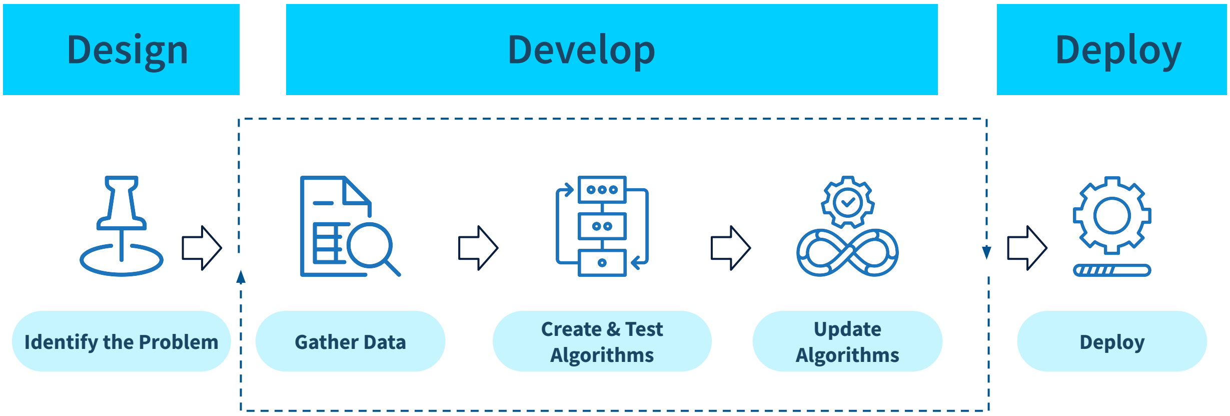

Design

Develop

Deploy

Identify the Problem

Gather Data

Create & Test Algorithms

Update Algorithms

Deploy

To share your team's understanding of their mission challenge, you first have to identify the key problems your AI technology is going to solve.

Pro Tip: Make sure you explain this process and outcome in plain language.

Gather the data needed to meet the identified problems.

Pro Tip: Make sure you consult with security, privacy officials and/or OGC.

The model training and selection process is interactive. No model achieves best performance the first time it is trained. It is only through iterative fine-tuning that the model is honed to produce the desired outcome.

Pro Tip: Make sure you consult with a data scientist.

Once one or more models have been built that appear to perform well based on relevant evaluation metrics, test the models on new data to ensure they generalize well and meet the business goals.

Pro Tip: Make sure you know who is documenting any data or algorithm changes.

Deploy the system.

Include a feedback mechanism for users.

Pro Tip: Make sure you evaluate for potential bias/reinforcement bias and disparate treatment.

Identify a Problem

To share your team's understanding of their mission challenge, you first have to identify the key problems your AI technology is going to solve.

Key Questions	Agency Stakeholders to Involve	Relevant Artifacts
What mission purpose does this algorithm meet?	☑ Business line *(for strategy)* ☑ Enterprise Architect ☑ Customers *(other Federal employees and/or the public)* ☑ Supply Chain	☑ BIA ☑ SOW ☑ Surveys/Performance Measurement ☑ Suggestions for Improvement ☑ Funding/Budget Info ☑ Modernization/AI Strategy
What is the expected mission outcome of the algorithmic process?	☑ Business Line *(strategy)* ☑ Customers *(other Feds and/or the public)* ☑ Employees ☑ Supply Chain	☑ BIA ☑ SOW ☑ improved business functions/returns data ☑ workflow diagram ☑ decreased funding/budget projections ☑ modernization / AI strategy / shared sustained value creation
Do you have authority to collect/retain this data?	☑ Business line owners ☑ Customers ☑ Legal ☑ Privacy ☑ Security *(also consider other Federal agency employees and/or the public)*	☑ Privacy Threshold Assessment (PTA)/Privacy Impact Assessment (PIA) ☑ System of Records Notices (SORN) ☑ Privacy Act Statements ☑ Consent Data
What is the provenance of the data you are thinking about using? Did you collect it directly? Did another party collect it on your behalf?	☑ Business line owners ☑ CIO/CDO ☑ IT architects/engineers ☑ Legal ☑ Privacy ☑ Security ☑ Customers	☑ SOW ☑ Data Inventory

Gather the Data

Gather the data needed to meet the identified problems.

Key Questions	Potential Stakeholders to Involve	Relevant Artifacts
Do you have legal authority to collect and/or use this PII or other CUI? **Is there a public notice regarding the data collection (SORN)?**	☑ OGC ☑ Privacy Officer ☑ CDO ☑ PRA Officer ☑ Authorizing Official ☑ Users/UX	☑ PTA ☑ PIA ☑ SORN ☑ Computer Matching Agreement ☑ Interconnectivity Agreement
How is the data protected in transit and at rest?	☑ Authorizing Official ☑ Security ☑ Enterprise Architect/Networking ☑ FedRAMP	☑ Security Controls in the System Security Plan (SSP) ☑ Other Authorization to Operate (ATO) documents ☑ SOW
If the proposed dataset includes PII, how is notice provided and consent gathered?	☑ User Experience (UX) Designer ☑ OGC/Privacy Officer to decide whether the extent of the notice consent is sufficient	☑ SORN ☑ Privacy Act Statement ☑ PIA
Was the data collected directly from the public?	☑ ICR	☑ Privacy Act Statement ☑ PIA
Who owns the data? How is access managed?	☑ CDO/Network Architect	☑ Data Inventory
How timely is the data? How long is it going to be relevant to the problem statement?	☑ Data Scientist/CDO	☑ Data Inventory
Is the data sufficiently representative? Are there potential sources of bias? **How representative/biased is the data? How did you examine its relevance/bias?** **What are the legal/ethical/public trust implications of bias? Who decides acceptable data quality/bias?**	☑ Data Scientist/CDO	☑ Data Inventory
What did you do to the data when preparing it for testing?	☑ Data Scientist/CDO	☑ Data Inventory
Where is the data going to be kept/for how long?	☑ Agency Records Officer	☑ Records Schedule

Create & Test Algorithms

The model training and selection process is interactive. No model achieves best performance the first time it is trained. It is only through iterative fine-tuning that the model is honed to produce the desired outcome.

Key Questions	Potential Stakeholders to Involve	Relevant Artifacts
What is the accuracy threshold? i.e. How accurate must an output be (false positive rate vs false negative rate) to be accepted as "accurate"?	☑ Business Owner ☑ Data Scientist	Model card
What is the mechanism to determine potentially biased outputs? **How do you determine potentially biased outputs? (data quality and integrity)**	☑ Data Scientist ☑ CDO and Civil Rights/Civil Liberties rep and/or equivalent (OGC or Equal Employment Opportunity (EEO))	Model card
How can the experiment be reproduced? Are there any barriers to reproducibility? (auditing/accounting)	☑ Data Scientist ☑ CDO	Model card
How do you determine if your your training data set if representative? **How do you determine if your algorithm is accurate?** **What methods are in place to ensure the data is complete and to make sure the algorithm is accurate?**	☑ Data scientist and CDO assesses the test data; developers assess the algorithm. Ideally both should have peers review too. ☑ Recommended best practice is to bring data scientist, business owner/SME together to after the first test of the test data+algorithm together.	Data Inventory
What is the expected outcome?	☑ Business owner/SME/data scientist need to address this.	
How are you documenting algorithm versions and testing protocols (for reproducibility)? What is the algorithm designed to do? Does it do anything else?		Model card(s)
Can the algorithm be modified? If you don't get the output/outcome that's desired, do you have flexibility to change the algorithm? If so, what is the process and what risks might that introduce?	☑ Bring the OGC/C.O./developer/data scientist to address this	

Create & Test Algorithms (Continued)

Key Questions	Potential Stakeholders to Involve	Relevant Artifacts
If the initial outcome is unexpected, how much work will it take to re-work and reproduce the tests? If outcomes are unexpected, do you need PII?	Estimated level of effort (LOE) from business owner/SME/software developer/data scientist Peer review - multiple data scientists involved and may be necessary to prove continuing legitimacy of algorithm.	Back-ups Project Phases/Checkpoints
How do you evaluate the potential for bias/reinforcement bias and/or disparate treatment? **What is the standard and/or expectations for accuracy/confidence threshold? Does it allow for additional functionality and/or to increase confidence/reliability?**	☑ CDO ☑ CPO ☑ OGC	
Can you achieve similar/effective results with less (PII) data? **What if data is coming from a form and an algorithm reads the form + provides output. Then the user changes info on form - how do user and dev community communicate about changes to the form/algorithm/output? If there are significant changes, the governance team needs to potentially return to the "Gather the Data" phase.**	☑ System Owner ☑ CPO ☑ CDO ☑ UX and Communications Officer	

Deploy

Deploy the system.

Key Questions	Roles and Responsibilities	Relevant Artifacts
Is it phased deployment? If so, why? How do you determine if the purpose of the AI has changed over time?	☑ Business Owner ☑ Acquisitions and Contracting/Budget ☑ CIO ☑ CPO ☑ Communications Officer	☑ ATO letter ☑ PIA/SORN ☑ Model Card
Does the output include PII or other CUI? If yes, how is output protected and access limited?	☑ CISO/CIO/CDO/Network Architect	☑ CUI Markings ☑ Model Card
How does an individual know they are being subjected to an algorithmic process? Have any relevant SORNs/PIAs or other notices been updated, reviewed, and posted?	☑ System/Business Owner ☑ CPO ☑ UX ☑ Communications Office	☑ SORNs/PIAs or Other Notices
Can an individual access/amend their data in the system? If so, how?	☑ CPO ☑ CDO ☑ CIO ☑ CISO ☑ System(s) Owner	☑ SORN ☑ PIA ☑ Agency Privacy Act Regulations
If significant time has passed and/or conditions have changed, is there a way to update consent?	☑ System Owner ☑ CPO	☑ Updated SORN

Thank you!

Please reach out to the AI Community of Practice at ***tts-ai@gsa.gov*** *with* your feedback on this toolkit or if you have any questions.

AI + Privacy Workstream on Internal Capacity Team:

- Andy Riordan (GSA)
- Richard Speidel (GSA)
- Kameron Cox (DHS)
- Luz Irazabal (DHS)
- John Nelson (DHS)
- Holly Beckstrom (USDA)