AI Risk & Governance Checklist

A practical review tool for responsible AI adoption, oversight, and implementation planning

This checklist helps organizations identify major AI risks and governance needs before piloting, implementing, or scaling AI-enabled workflows. It is designed for startups, research teams, commercialization groups, operations leaders, and organizations adopting AI in a changing regulatory, funding, and business environment.

The checklist is not a substitute for legal, privacy, cybersecurity, regulatory, or compliance advice. It is intended to help leadership, technical teams, and change management stakeholders ask the right questions, assign ownership, and decide what controls should be in place before broader AI use.

How to Use This Checklist

Review each section with the people responsible for leadership, users, data, technology, risk, compliance, and implementation. Mark each item Yes, Partial, or No. Items marked Partial or No should be carried into the risk register, governance role table, or decision summary before the AI use case moves forward.

Organization and AI Governance Context

Organization / Team	
Checklist Owner	
Date	
AI Tool / Use Case Reviewed	Example: chatbot, proposal assistant, operational model, customer support tool, analytics workflow, internal automation
Primary Users / Affected Teams	
Current Stage	Idea | Discovery | Pilot | Implementation | Scaling | Review

Best-Fit Uses for This Checklist

- Evaluating AI tools before pilot use or organization-wide rollout.
- Reviewing risks for AI-assisted proposal writing, research support, customer communications, operations, analytics, or internal automation.
- Preparing governance controls for AI adoption in a startup, university-linked venture, or scientific organization.
- Clarifying human oversight, data handling, approval requirements, and user training needs.
- Supporting change management planning when AI will alter workflows, roles, decisions, or documentation practices.

1. Strategic Fit and Use Boundaries

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
The AI use case is tied to a clear business, research, operational, or commercialization objective.	☐	☐	☐	
The organization has defined what the AI tool is intended to do and what it should not be used for.	☐	☐	☐	
The team understands whether the use case supports internal productivity, decision support, customer interaction, analysis, automation, or product development.	☐	☐	☐	
High-stakes decisions that require expert judgment have been identified and are not being fully delegated to AI.	☐	☐	☐	
Leadership has approved the intended scope, pilot size, and conditions for expansion.	☐	☐	☐	
The organization has identified whether the use case could affect employees, customers, funders, investors, patients, students, partners, or the public.	☐	☐	☐	

2. Data Privacy, Security, and IP

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
The organization has identified what data, documents, prompts, files, or system integrations the AI tool will use.	☐	☐	☐	
Sensitive, confidential, proprietary, customer, employee, health, grant, or regulated data has been identified before use.	☐	☐	☐	
The team understands whether data entered into the AI tool may be stored, trained on, shared, logged, or accessed by a vendor.	☐	☐	☐	
Intellectual property and confidentiality concerns have been reviewed for the intended use case.	☐	☐	☐	
Access controls are defined so only appropriate users can upload, view, export, or act on sensitive AI outputs.	☐	☐	☐	
Cybersecurity, vendor, and data retention concerns have been reviewed before implementation.	☐	☐	☐	

3. Accuracy, Validation, and Quality Control

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
The organization has defined how AI outputs will be reviewed, validated, or tested before use.	☐	☐	☐	
Users understand that AI outputs can be incomplete, inaccurate, outdated, biased, or misleading.	☐	☐	☐	
The team has identified which outputs require expert review before being shared or acted upon.	☐	☐	☐	
There is a process for documenting corrections, errors, rejected outputs, or quality concerns.	☐	☐	☐	
The organization has defined acceptable accuracy, quality, or performance thresholds for the use case.	☐	☐	☐	
AI outputs are not being treated as authoritative without appropriate human review and source verification.	☐	☐	☐	

4. Bias, Fairness, and Stakeholder Impact

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
The team has considered whether the AI use case could affect people differently across groups, roles, locations, or customer types.	☐	☐	☐	
The organization has reviewed whether biased data, assumptions, or outputs could create unfair or harmful results.	☐	☐	☐	
Users have a way to report concerns about biased, inappropriate, or harmful AI behavior.	☐	☐	☐	
The organization has considered whether AI use should be disclosed to users, customers, employees, or other stakeholders.	☐	☐	☐	
The team has considered whether the AI tool could unintentionally reduce transparency, accountability, or trust.	☐	☐	☐	
Stakeholder concerns about adoption, oversight, job impact, or decision quality have been discussed.	☐	☐	☐	

5. Human Oversight and Decision-Making

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
A human owner is responsible for the AI use case, not just the technology vendor or software tool.	☐	☐	☐	
The organization has defined who reviews AI outputs before they are used in decisions, communications, proposals, reports, or operations.	☐	☐	☐	
Escalation steps are defined when an AI output appears incorrect, risky, biased, incomplete, or outside approved scope.	☐	☐	☐	
Users know when they are allowed to rely on AI and when they must seek expert review or leadership approval.	☐	☐	☐	
Decision rights are clear for approving pilot launch, workflow changes, user access, and broader deployment.	☐	☐	☐	
The organization has defined who can pause, revise, or discontinue AI use if risks emerge.	☐	☐	☐	

6. Adoption, Training, and Change Management

Governance / Risk Item	Yes	Partial	No	Notes / Owner / Next Step
Users have been trained on the intended use, limitations, risks, and acceptable practices for the AI tool.	☐	☐	☐	
The AI-enabled workflow has been explained clearly to people affected by the change.	☐	☐	☐	
The organization has created or planned documentation, SOPs, job aids, or usage guidance.	☐	☐	☐	
User feedback will be collected during pilot or early implementation.	☐	☐	☐	
The adoption plan considers trust, usability, role changes, communication needs, and support requirements.	☐	☐	☐	
Success metrics include adoption, workflow impact, quality, risk reduction, and user confidence, not only technical performance.	☐	☐	☐	

AI Risk Register

Use this table to record major risks identified during the checklist review. Rate probability and impact as Low, Medium, or High, then define the mitigation step and owner.

Risk	Category	Probability	Impact	Mitigation / Control	Owner	Status
Sensitive data entered into non-approved tool	Privacy / Security					
AI output used without review in high-stakes decision	Accuracy / Oversight					
Users do not trust or adopt the new workflow	Adoption					

AI Governance Roles and Decision Rights

Clarify who owns major decisions before deployment or wider use. For smaller organizations, one person may hold several roles, but the responsibilities should still be explicit.

Governance Role	Primary Responsibility	Assigned Person / Team	Decision Authority / Notes
Executive Sponsor	Sets priority, approves resources, resolves major conflicts.		
AI Use Case Owner	Owns the business need, workflow fit, and adoption plan.		
Technical Lead	Reviews model/tool fit, integrations, performance, and technical feasibility.		
Data / Security Reviewer	Reviews data access, confidentiality, privacy, security, and record handling.		
Human Oversight Lead	Defines review steps, escalation points, and acceptable use boundaries.		
Training / Change Lead	Coordinates user training, communications, documentation, and feedback.		

AI Governance Policy Starter Checklist

Use this section to identify what written guidance or SOPs may be needed before broader AI use. This is not a legal review; it is a practical planning prompt for governance, risk reduction, and responsible adoption.

Policy / SOP Area	Needed?	Current Status	Notes / Next Step
Acceptable AI Use Policy	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Data Privacy and Confidentiality Guidance	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Human Review and Approval Procedure	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Prompting and Output Documentation Guidance	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Prohibited Uses and High-Risk Use Restrictions	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Vendor / Tool Review Procedure	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Training and User Acknowledgment Process	☐ Yes ☐ No	Not started | Draft | Active | Needs update	
Incident, Error, or Concern Reporting Procedure	☐ Yes ☐ No	Not started | Draft | Active | Needs update	

Governance Decision Summary

Complete this summary after reviewing the checklist. The goal is to determine whether the AI use case is ready to proceed, needs additional controls, or should be paused until major issues are addressed.

Overall Governance Status	☐ Ready to proceed ☐ Proceed with controls ☐ Needs revision ☐ Pause / do not proceed
Top 3 Risks	
Required Controls Before Launch	
Training / Communication Needed	
Approvals Required	
Next Review Date	

ScienceDocs AI and Change Management Support

ScienceDocs helps organizations plan responsible AI adoption, evaluate readiness, align stakeholders, define governance controls, and connect AI implementation with commercialization, grant strategy, optimization engineering, workflow redesign, and organizational change. Ready to get started? Contact ScienceDocs to discuss AI risk, governance, and implementation planning for your organization.