Generative AI Policy Creation Guide

A thought-provoking, practical framework for building Generative AI policy that supports responsible adoption.

How to Use This Guide

An effective AI policy gives people the clarity, confidence, and guidance to use AI responsibly in their everyday work. Successful adoption depends on leadership, practical collaboration across the organization, and clear expectations that employees can apply. Complete this guide manually or use the LearnAIR AI Policy Builder digital assistant [click here] for a guided, AI-assisted experience.

Use this guide to:

- Build a new organizational AI policy
- Evaluate and strengthen an existing policy
- Identify decisions that must be made before policy language is drafted
- Align leadership to support responsible AI adoption

Disclaimer

This is a policy creation guide, not legal advice. Final policies should be reviewed by the organization's legal, privacy, cybersecurity, human resources, and compliance leaders.

Table of Contents

Phase 1: Prepare Before Writing 3
Review Existing Policies 3
Inventory Current AI Use 4

Phase 2: Make the Core Policy Decisions 5
Define Your AI Strategy 5
Determine Data Rules 5
Approve Tools & Accounts 6
Identify High-Risk Use Cases 6
Define Human Review Requirements 7

Phase 3: Build Your AI Policy 8
1. Purpose & Scope 8
2. Core Principles 9
3. Approved AI Tools 10
4. Data Classification & AI Input Rules 11
5. Acceptable & Encouraged AI Use 13
6. Prohibited Uses 14
7. Human Oversight & Output Review 16
8. Transparency & Disclosure Standards 17
9. Vendor and Third-Party AI 19
10. Incident Reporting & Response 20
11. Policy Review, Ownership & Governance 22

Prepare before writing

An AI policy should reflect how your organization actually works. Before drafting policy language, take time to understand your current AI landscape, identify the people who should help shape the policy, and review the governance already in place. This preparation creates a policy that is aligned with organizational goals and therefore more likely to be adopted.

Review Existing Policies

An AI policy should complement existing governance. Many of the controls needed for AI already exist in other organizational policies.

Review the following policies (if applicable) before drafting

Existing Policy	Relevant AI Considerations
Information Security	Data handling, approved software, authentication
Privacy Policy	Personal information, consent, retention
Acceptable Use Policy	Technology expectations
Data Classification	Public, confidential, restricted information
Records Management	Retention of prompts and AI outputs
Intellectual Property	Ownership of AI-generated work
Vendor Management	Third-party AI providers
Code of Conduct	Ethical AI use
Incident Response	Reporting AI-related issues

Questions to Consider

- Does another policy already address this issue?
- Will this AI policy conflict with existing guidance?
- Are existing approval processes sufficient for AI tools?

Inventory Current AI Use

Most organizations are already using AI in some capacity, whether formally approved or not. Understanding current usage provides the foundation for practical governance and helps identify opportunities, risks, and existing gaps.

Complete the following inventory

Question	Notes
Which AI tools are currently being used?	
Which departments use AI tools?	
What tasks are employees completing with AI?	
What information is being shared with AI tools?	
Which AI tools have been formally approved?	
Are any AI tools connected to internal systems?	
Which use cases involve customers or employees?	
Which use cases could impact legal, financial, or employment decisions?	

Watch For

- [ ] Personal AI accounts being used for company work (e.g., ChatGPT, Gemini)
- [ ] Browser extensions (e.g., Grammarly AI, Monica)
- [ ] AI meeting assistants (e.g., Otter.ai, Fireflies)
- [ ] AI coding assistants (e.g., GitHub Copilot, Cursor)
- [ ] Connected AI agents (e.g. ChatGPT Agents with API access)
- [ ] Automated workflows (e.g., Zapier, Make.com)
- [ ] Shadow AI (unapproved AI tools used without IT knowledge)

Make the Core Policy Decisions

Before drafting policy language, leadership should first agree on the organization's approach to AI. These decisions establish the direction, expectations, and governance model that the policy will ultimately document.

Below are practical questions to help think through these key considerations

Define Your AI Strategy

Every AI policy should support a larger organizational objective.

Leadership Discussion Questions

- Why are we adopting AI?
- Which business priorities should AI support?
- What level of experimentation are we comfortable encouraging?
- What responsibilities should always remain with people?

Determine Data Rules

Determine your data categories (e.g. internal, confidential, client information, financial information) & whether each can be used with AI.

Leadership Discussion Questions

- Which information should never be entered into AI?
- Should different roles have different permissions?
- Do customer contracts place additional restrictions on AI use?

Approve Tools & Accounts

Leadership should define which tools employees may use and how new tools are evaluated over time.

Tip: Since AI evolves quickly, maintain your approved tools as a living document rather than embedding detailed tool lists directly into the policy.

Identify High-Risk Use Cases

Not every use of AI carries the same level of risk. Organizations should identify activities that require additional oversight before the policy is written.

Consider Whether AI Will Be Used For

- ☐ Hiring and recruitment
- ☐ Employee performance reviews
- ☐ Legal or contractual decisions
- ☐ Financial approvals
- ☐ Customer eligibility decisions
- ☐ Healthcare or clinical support
- ☐ Automated customer communications
- ☐ AI agents that take actions on behalf of users
- ☐ Access to internal systems
- ☐ Other high-impact decisions

Leadership Discussion Questions

- Which AI uses require leadership approval?
- Where should AI assist rather than decide?
- Which decisions should always remain human?

Define Human Review Requirements

Human review requirements should reflect the level of access, autonomy, and potential impact involved.

Determine Where and to Whom Human Review Belongs

AI Use Case	Human Review Required?	Reviewer
Summarizing meetings and creating follow-up tasks		
Sending customer or partner communications		
Recommending decisions that affect employees or customers		
Updating CRM, HR, finance, or project records		
Initiating purchases, payments, refunds, or account changes		
Accessing confidential data through connectors or agents		
Publishing content or responding publicly on behalf of the organization		
Deploying software changes		
Escalating incidents or triggering security actions		

Leadership Discussion Questions

- Can the system act, or only recommend?
- Can the action be easily reversed?
- What data and systems can it access?
- Who remains accountable when an AI-supported workflow crosses departments?
- What should cause the system to pause and request human input?
- Does the system affect a person, payment, contract, record, or customer interaction?
- Does approval need to happen before the action, after the action, or only when risk thresholds are reached?

Build Your AI Policy

The sections below provide a recommended structure for your AI policy.

01 Purpose and Scope

Explain why the organization has an AI policy, what it hopes to accomplish, and who the policy applies to.

Keep this section focused on direction rather than rules. Employees should understand the organization's intent before reading the details of the policy.

Include

- Purpose of the policy
- Organizational AI vision
- Business objectives
- Who the policy applies to
- Defintions of the terminology employees need to understand the policy.
- Which technologies are covered
- Which business activities are included

02 Core Principles

Set the organization's shared expectations for responsible AI use. These principles should help employees make sound decisions when a tool, workflow, or risk is new.

Include

1. **Human Accountability:** People remain responsible for decisions, actions, and work produced with AI support.
2. **Privacy and Data Protection:** AI use must follow organizational requirements for confidential, personal, regulated, and client-owned information.
3. **Security:** Employees must use approved tools, accounts, integrations, and access controls.
4. **Accuracy and Reliability:** AI outputs should be reviewed and verified based on their intended use and level of risk.
5. **Fairness and Non-Discrimination:** AI should be evaluated for bias and should not create unlawful or unfair outcomes.
6. **Transparency:** AI use should be disclosed when it materially affects a person, decision, interaction, or public-facing output.
7. **Purposeful Use:** AI should be used for legitimate organizational purposes and within the boundaries of an employee's role and authority.
8. **Responsible Innovation:** Employees should be encouraged to explore valuable AI use cases while following established safeguards and approval processes.

03 Approved AI Tools

Identify the AI platforms, accounts, and connected services employees are authorized to use for business purposes.

Include

Consider addressing the following:

- Approved AI platforms and services
- Organizational account and licensing requirements
- Personal vs. company-managed accounts
- Approved browser extensions, integrations, and connectors
- AI agents and autonomous workflows
- Data access permissions for each tool
- Tool ownership and administration
- Process for requesting new AI tools
- Approval criteria and decision-making authority
- Review cadence for approved technologies

Example Language

"Approved AI platforms include Microsoft 365 Copilot and Claude Enterprise for authorized business use. Employees may only use organization-approved accounts for business activities. Requests for new AI tools, integrations, connectors, or autonomous agents must be submitted to the AI Governance Team. Evaluations will consider security, privacy, data handling, business value, and system compatibility. Approved technologies are maintained in the Approved AI Tools Register and reviewed regularly."

04 Data Classification & AI Input Rules

Define what types of information employees may share with AI systems and under what conditions.

Include

Consider addressing the following:

- Organizational data classifications (e.g., Public, Internal, Confidential, Restricted)
- Client-owned and customer data
- Personally Identifiable Information (PII)
- Protected Health Information (PHI), where applicable
- Financial, legal, and regulated information
- Intellectual property and trade secrets
- Authentication credentials, passwords, and API keys
- Data minimization and de-identification practices
- AI input restrictions by data classification
- Guidance for situations where employees are unsure

04 Data Classification & AI Input Rules

Example Language

"Information may only be shared with AI systems based on its data classification and approved organizational controls.

- Public data: Approved AI tools.
- Internal data: Organization-approved AI platforms using company-managed accounts.
- Confidential data: Enterprise AI platforms approved for confidential data.
- Restricted: data No external AI systems unless explicitly authorized.

When in doubt, classify information at the more restrictive level."

05 Acceptable & Encouraged AI Use

Suggest where AI can support employees, improve work, and create value across the organization. This section should give people clear examples of appropriate AI use cases so they can adopt AI with confidence.

Ideas to Include

- Exploring ideas, testing assumptions, and comparing possible approaches
- Turning unstructured notes, meetings, or source materials into organized plans, reports, and next steps
- Breaking complex work into smaller tasks and identifying opportunities for AI support
- Researching, synthesizing, and organizing information
- Drafting, refining, and adapting content for different audiences and formats
- Preparing for meetings by generating agendas, summaries, follow-up actions, and discussion points
- Analyzing data, identifying trends, and recommending next steps
- Building reusable workflows, role-based AI assistants, and approved low-risk automations
- Retrieving and synthesizing knowledge from approved organizational sources
- Supporting project planning, task management, prioritization, and recurring operational work
- Supporting learning, onboarding, and knowledge transfer

06 Prohibited AI Use

Define the AI activities, behaviors, and practices that are prohibited because they create unacceptable legal, security, privacy, ethical, or operational risk.

Include

- Entering restricted, regulated, or unauthorized data into AI systems
- Using unapproved AI tools, accounts, integrations, or agents for business purposes
- Allowing AI to make consequential decisions without appropriate human oversight
- Circumventing organizational security, privacy, legal, or compliance controls
- Sharing passwords, authentication credentials, API keys, or other sensitive access information with AI
- Creating deceptive, fraudulent, discriminatory, or misleading content
- Deploying AI agents or automations beyond their approved permissions or authority
- Using AI in violation of applicable laws, regulations, customer contracts, or organizational policies

06 Prohibited AI Use

Example Language

"AI may not be used to: (a) make hiring, promotion, termination, compensation, or other consequential employment decisions without appropriate human review and approval; (b) access, process, or share restricted, regulated, or unauthorized information in violation of organizational data handling requirements; (c) impersonate employees, customers, or other individuals, or generate deceptive or misleading content; (d) bypass established security, privacy, legal, or compliance controls; (e) deploy AI agents, automations, or connected workflows that can take consequential actions without the required governance, approvals, or human oversight; or (f) operate in violation of applicable laws, regulations, customer agreements, or organizational policies."

07 Human Oversight & Output Review

Define how employees should review, validate, and remain accountable for AI-assisted work. This section establishes expectations for verifying AI outputs, approving AI-supported actions, and ensuring appropriate human judgment throughout the work process.

Include

- Human accountability for AI-assisted work and decisions
- Output verification before information is relied upon or shared
- Risk-based review requirements
- Approval requirements for AI-supported actions
- Escalation procedures when AI outputs are uncertain or potentially harmful
- Expectations for monitoring AI agents and automated workflows
- Documentation requirements for high-impact AI-assisted decisions

Don't Forget The Human Part™

Human oversight is more than reviewing AI output. It includes verifying important information, approving consequential actions, monitoring autonomous workflows, and remaining accountable for the final outcome regardless of how the work was completed.

08 Transparency & Disclosure Standards

Define when AI use should be communicated to employees, customers, partners, regulators, or the public. Clear disclosure standards build trust while helping the organization meet legal, contractual, and ethical obligations.

Include

- AI-generated or AI-assisted public content
- Customer-facing AI interactions
- AI-generated recommendations used in consequential decisions
- AI-generated images, audio, or video
- AI meeting recording and transcription
- Internal documentation of AI-assisted work, where appropriate
- Situations requiring consent before AI is used
- Regulatory or contractual disclosure requirements

08 Transparency & Disclosure Standards

Example Language

"The organization will disclose the use of AI whenever required by applicable law, contractual obligations, or internal policy. Employees must not represent AI-generated or AI-assisted work as solely human-created when disclosure is required.
Public Communications: AI-generated or AI-assisted content will be disclosed when required by law, contract, or organizational policy.
Meeting Recording & AI Notetakers: Participants will be notified before AI systems record, transcribe, or summarize meetings when required by law, customer agreement, or organizational policy.
Customer Interactions: Customers will be informed when interacting primarily with an AI system or when AI materially supports customer-facing communications, where required.
AI-Generated Media: AI-generated or AI-edited images, audio, and video used externally must be clearly identified when required by law or when omission could reasonably mislead.
High-Impact Decisions: The use of AI in employment, financial, legal, healthcare, safety, or other high-impact decisions must be documented and retained in accordance with applicable governance and recordkeeping requirements."

09 Vendor & Third-Party AI

Define the requirements for evaluating, approving, and managing third-party AI systems, vendors, and service providers. This section helps ensure external AI technologies meet the organization's security, privacy, legal, and governance standards before they are introduced into the business.

Include

- Vendor evaluation and approval process
- Security, privacy, and compliance requirements
- Data processing and data usage agreements
- AI model transparency and vendor documentation
- Contractual requirements and service-level expectations
- Data residency and retention requirements
- Ongoing monitoring and periodic vendor reviews
- Governance requirements for embedded AI features and third-party agents

10 Incident Reporting & Response

Define how AI-related incidents are identified, reported, investigated, contained, documented, and resolved. The objective is to minimize business, legal, privacy, security, and operational impact while improving future AI governance.

Include

- What constitutes an AI incident
- Incident severity levels
- Reporting timelines and channels
- Investigation responsibilities
- Containment and corrective actions
- Response requirements
- Documentation and record retention
- Incident severity and response table (see example below)

10 Incident Reporting & Response

Example Language

"Employees must promptly report suspected or confirmed AI-related incidents through the organization's designated reporting channels. AI incidents include unauthorized data disclosure, security breaches, regulatory violations, deployment of unapproved AI systems, harmful AI outputs, or other material business, legal, privacy, or operational risks.
Critical Incidents: Report immediately and escalate to Security, Legal, and Executive Leadership.
High-Severity Incidents: Investigate within 24 hours, contain the impact, and notify appropriate stakeholders.
Medium-Severity Incidents: Document, review, and remediate operational issues through the responsible department.
Low-Severity Incidents: Document minor errors, AI hallucinations, and policy questions, and address them through coaching or additional training.
All AI incidents must be documented, investigated, and retained in accordance with organizational governance and record retention requirements."

11 Policy Review, Ownership & Governance

Define who owns the AI policy, how governance decisions are made, when the policy is reviewed, and how changes are approved and communicated. This section ensures the policy evolves alongside changes in technology, regulations, business operations, and organizational risk.

AI governance is a continuous process. Assign clear ownership, define review triggers beyond a calendar date, and ensure every significant change is documented, communicated, and supported through employee education.

Example Language

"The AI Governance Committee is responsible for maintaining this policy and ensuring it remains aligned with the organization's strategic objectives, regulatory obligations, and risk tolerance. The policy will undergo a formal review at least annually and may be reviewed sooner following significant AI incidents, changes in applicable laws or regulations, the adoption of new AI technologies, material business or organizational changes, audit findings, or emerging security or compliance risks. All revisions must be reviewed and approved through the organization's established governance process. Policy versions will be documented through version control with a published change history. Employees will be notified of material updates and, where appropriate, required to acknowledge revised policy requirements."

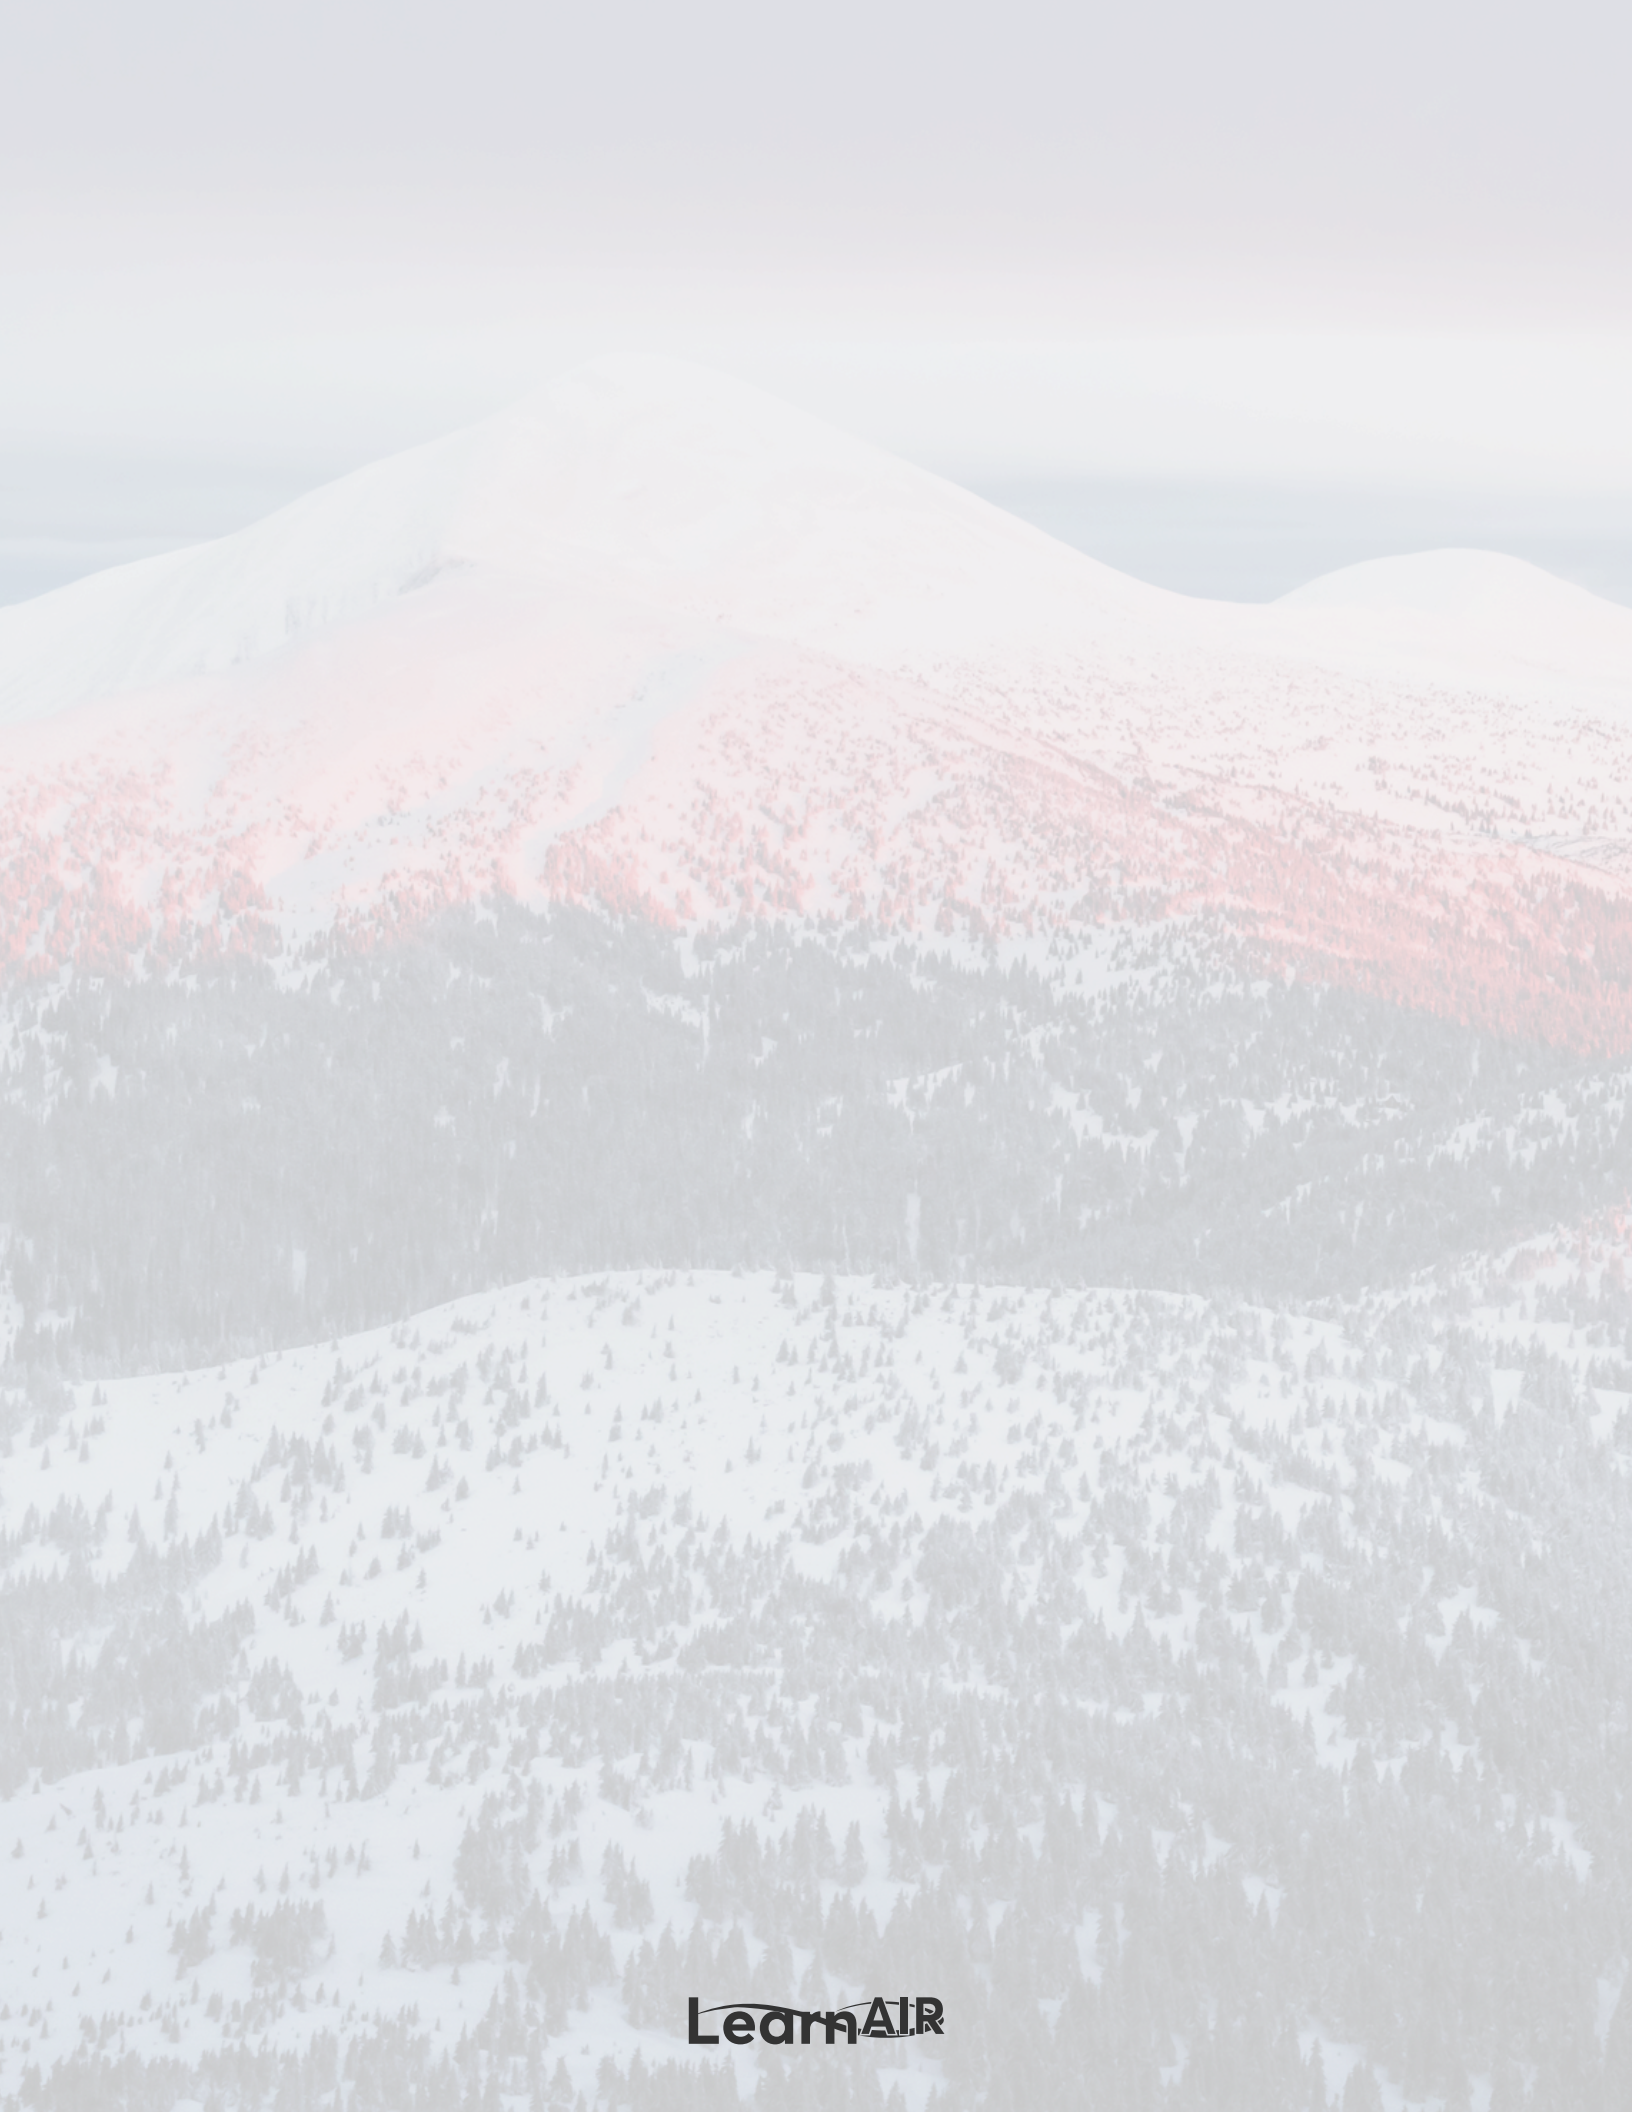
LearnAIR