Basic AI Usage Guidelines

This is a set of basic AI usage guidelines provided by the AI Advisory Board Summit Group. The guidelines address the most pressing questions that we received from the survey on AI usage across the TM organization earlier this year. The purpose of these guidelines is to provide technical guidance and information that can be used to inform more specific AI usage policies across the TM organization.

We anticipate that the guidelines will need to be updated regularly as the technology is constantly evolving.
Please reach out with feedback or suggestions at ai@tm.org.

Executive Summary

- **Confidential information** should not be entered into personal AI tools (ChatGPT, Gemini, Claude, Grok, Perplexity, etc.).
- **Enterprise accounts** (e.g., Gemini in Google Workspace) and **Siri AI** are the recommended solutions for confidential use cases. But when any AI tool searches the web, confidential information can leak through the search queries.
- **Self-hosted open-source models** (e.g., via Ollama on a Laptop) are the safest choice when working with highly confidential information.
- **Privacy preserving solutions are available** for common tasks such as Audio Transcription and Translation.
- **AI agents** should not send, publish or delete anything without human review, and should only access confidential data through a business / enterprise account.
- **AI-generated text** has recognizable patterns, which can influence how it is received and can be reduced when desired.

---

1. Leading AI Providers and Models

Below is an overview of major providers of general-purpose AI models.

- **OpenAI: ChatGPT** — http://chatgpt.com
- **Google: Gemini** — https://gemini.google.com/
- **Anthropic: Claude** — http://claude.ai
- **xAI: Grok** — https://grok.com/
- **Apple Siri AI –** https://www.apple.com/apple-intelligence/

A. Access With Personal Account (Free or Paid)

The standard way to interact with these AI models is via a personal account with options for a Free or Paid Tier (cost often between 20 and 200$/month).

Quality

- **Free Tier:** High-quality general-purpose models
- **Paid Tier:** Very high quality; access to the best models
- Free tiers may include usage limits

Privacy and Security

Low (Free and Paid Tiers)

- Privacy is *not* preserved
- Confidential data should not be shared
- Usage logs are stored and accessible to provider employees
- Data may be used for model training unless opted out.

Note on training: Allowing non-confidential content to be used for model training can help models accurately represent our knowledge, but **there is no guarantee** that the model will be actually trained on our data, even when we allow it.

Recommended Examples

- Drafting emails or public-facing content
- Creating marketing materials
- Brainstorming non-confidential ideas

Not Recommended Examples

- Summarizing confidential meeting notes
- Translating confidential course materials

B. Access With Business / Enterprise Account

Providers

- **Google Workspace AI** — https://workspace.google.com/solutions/ai/

    - Non-profit discounts: https://workspace.google.com/learning/google-workspace-for-nonprofits-with-gemini
- **Claude for Business – Team, Enterprise (20 seats min)**
    - Non-profit details: https://claude.com/solutions/nonprofits
- **OpenAI for Business** — https://openai.com/business/
    - Non-profit details: https://help.openai.com/en/articles/9359041-openai-for-nonprofits

Quality

- **Very High**
  Enterprise accounts generally include the provider's strongest models.

Privacy and Security

Medium to High (varies by provider)

- Designed for confidential data
- Data ownership remains with the customer
- Data encrypted in transit and at rest
- Logs are automatically deleted (e.g., after 30 days)
- Data is not used for model training by default
- Rare exceptional human review may occur (e.g., legal requirements)

Note:

- **Google Workspace Gemini** currently provides the strongest enterprise privacy guarantees.
- OpenAI, Claude and Grok have weaker guarantees and may review content suspected of violating terms of service.

Recommended Examples

- Summaries of confidential internal meetings
- Translating internal but not extremely sensitive materials

Not Recommended Examples

- Extremely confidential content (e.g., TM and Advanced Technique instructions)

C. Apple Siri AI

Apple's assistant, built into recent Apple devices. Requests are handled on the device or on Apple's Private Cloud Compute servers.

Availability

- Beta in macOS and iOS 27 since 14 September 2026, in English, with more languages coming October.
- Must be switched on under Siri settings with "Try Siri AI (Beta)"
- Requires iPhone 15 Pro, iPhone 16 or later, or an iPad or Mac with M1 or later
- In the EU, available on Mac but not on iPhone or iPad

Quality

High

Privacy and Security

High

- Data sent to Private Cloud Compute is not stored or accessible to Apple.
- Data is not used for model training unless you opt in. To check, go to Settings (System Settings on Mac) > Privacy & Security > Analytics & Improvements and make sure "Improve Siri & Apple Intelligence" is off
- **Privacy is lower when answering from web search**, as anonymized search queries leave Private Cloud Compute (same as Section 3.A)

Recommended Examples

- Everyday drafting, summarizing and questions on a personal device
- Confidential material, except when Siri answers from web search

Not Recommended Examples

- Extremely confidential content (e.g., TM and Advanced Technique instructions)

---

2. Open Source AI Models for Self-Hosting

Open-source models can run locally on personal devices (e.g., MacBooks with Apple Silicon) or private servers.

Leading Self-Hosting Platforms

- **Ollama** — http://ollama.com
- **Jan.ai** — https://www.jan.ai/

Quality

- **Medium**
  Not as strong as top proprietary models, but highly capable for many tasks.

Privacy and Security

Very High

- All data stays on your device or server
- Ideal for highly confidential content
- No external logs, no third-party access

Recommended Examples

- Summaries of confidential meeting notes
- Translating confidential information
- Drafting internal sensitive material

Not Recommended Examples

- Public-facing content where maximum output quality is required and confidentiality is not an issue

---

3. AI Tools for Specific Applications

A. AI-Powered Web Search

These services combine AI reasoning with real-time web search:

- ChatGPT (with browsing)
- Gemini (with search)
- Claude (with search)
- Grok (web-integrated)
- Perplexity — http://perplexity.ai
- You.com — https://you.com/search

Privacy and Security

Low

- Any confidential data may be leaked through search queries
- Treat as equivalent to typing sensitive content into Google Search
- **Entering confidential information should be avoided**

Recommended Examples

- General research
- Public information lookup
- Fact checking and recency checks

B. Audio Transcription

MacWhisper

- MacOS app: https://goodsnooze.gumroad.com/l/macwhisper
- Can run fully on-device
- **Privacy:** Very High (on-device), High (cloud mode)

Turboscribe

- Website: https://turboscribe.ai/
- **Privacy:** High (data encrypted, processed on server)

C. Translation

- ChatGPT, Gemini, Claude, Grok — decent quality but **low privacy with personal accounts**
- **DeepL** — https://www.deepl.com/
    - Free Tier: Low privacy
    - Paid Tier: High privacy (suitable for confidential data)

---

4. AI Agents and Connectors

Some AI tools do not only generate text, but access your data or act on your behalf — reading files, browsing websites, sending emails, or running commands on your computer.

A. Connectors for Reading and Summarizing

Linking ChatGPT, Gemini or Claude to Gmail, Drive, Slack or Microsoft 365 so it can read and summarize your own material.

Privacy and Security

Follows the account it runs under (see Section 1)

- Low with a personal account — confidential data can be exposed without anyone entering it into a prompt
- Medium to High with a business / enterprise account, with the same limits as Section 1.B
- Low if the connector agent can also do web search (same as Section 3A), independent of personal/business account
- The agent sees everything the connected account can see — connect single folders or mailboxes, not everything
- Hidden instructions in a document, email or webpage can redirect it (prompt injection)

B. Agents That Take Actions

Browser agents that click and fill in forms, and desktop agents such as OpenClaw (https://openclaw.ai/) that run on your own machine.

Privacy and Security

Low

- Actions can be irreversible — emails sent, files deleted, content published
- Prompt injection becomes far more serious, as hidden instructions can trigger actions rather than only influence text
- Desktop agents run with system-level access to files, terminal and logged-in accounts
- Business or enterprise terms do not reduce these risks

Note on supervision: Agents should not act on their own. A person should review and approve anything that is sent, published, purchased or deleted.

Recommended Examples

- Non-confidential routine tasks that you watch while they run
- Drafts and proposed actions that you review and carry out yourself

Not Recommended Examples

- Letting an agent send, publish or delete anything without review
- Running a desktop agent on a device with access to confidential material

---

5. AI Writing Patterns

AI-generated text has recognizable patterns, and many readers now notice them. The same text can be received differently once it is assumed to have been written by AI. How much this matters depends on the material and the audience. The guidance below is not intended to hide the use of AI, but to help the text read well for the people it is written for.

Common Patterns

- Contrast pairs: "This is not a delay. This is a decision."
- Two short sentences placed side by side: "Clear. Simple."
- Image pairs that give no actual instruction: "less a hammer, more a scalpel"
- Openings that clear the throat: "Here's the thing." "Let me be clear."
- Sentences that comment on their own importance: "And that matters."
- Groups of three: "faster, simpler, clearer"
- A range where a number would do: "5 to 10 minutes" instead of "7 minutes"
- Endings that repeat what was just said: "In short…" "At the end of the day…"
- Recurring vocabulary: delve, leverage, robust, seamless, foster, landscape, underscore
- Headings, bullet points and bold phrases in short texts that would read better as prose
- Even paragraph lengths and an unvarying rhythm
- Frequent em dashes

Practical Steps for Pattern Reduction

- Start from your own notes and use AI to edit, rather than generating from scratch
- Provide a sample of your own writing as a reference for tone
- Ask for plain language. Referring to ASD-STE100, the aviation manual writing standard, works as a shorthand, though maybe not for personal or creative texts
- Add Zinsser's principles of simplicity, brevity, clarity and humanity, so the result does not become mechanical
- Keep a list of the patterns above and have the AI check each draft against it
- Double check the first and last paragraphs, which are usually the most generic
- Replace general examples with specific ones only you would know
- Vary sentence length, and read the text aloud before sending
- Start a new chat when one gets long, as the writing drifts back to these patterns

Further Reading

- Ruben Hassid, patterns that give away AI writing — https://www.linkedin.com/posts/ruben-hassid_i-absolutely-hate-ai-writing-and-my-entire-share-7493618680224788481-S8nJ/
- Charlie Guo, The Field Guide to AI Slop — https://www.ignorance.ai/p/the-field-guide-to-ai-slop
- The AI penalty and disclosure paradox, research on how AI-assisted communication is perceived — https://www.sciencedirect.com/science/article/pii/S2949882126000551