

Managing AI Personalization Responsibly

Why instructions, memory, chat history and connected services should be reviewed regularly.

Personalization at a glance

Personalization becomes safely useful only when instructions, memory, history, connected services, and privacy controls are reviewed separately.




Figure 1: Personalization consists of several separate layers.

Important note: The illustrations are simplified explanatory graphics, not direct screenshots from a personal account. Features, labels and switches may vary by provider, plan, country, age, device, account type and staged rollout. This guide complements the product-specific how-tos and does not replace reviewing the current settings in your own account.

As of 6 August 2026 | English | Audience: personal users

Contents

- 1. What personalization actually includes
- 2. Why control matters
- 3. Benefits and limitations
- 4. Risks and common misunderstandings
- 5. What data should be stored long term
- 6. Comparison of ChatGPT, Copilot, Claude and Gemini
- 7. Initial setup and regular review
- 8. Decision guide for common situations
- 9. Checklist and sources

Five core rules

- **Personalization is context, not truth.** A relevant answer can still be wrong, incomplete or one-sided.
- **Less is often better.** Store only information that is stable, useful and low sensitivity.
- **Memory, history and training are separate.** Turning off one control does not automatically delete data elsewhere.
- **Keep project information inside the project.** Do not overload global memory with short-lived tasks.
- **Control is recurring.** Review settings after major updates and at least quarterly.

Goal of this guide: you should always be able to explain what information an AI system uses about you, why it uses it, and how to change or remove it.

1. What personalization actually includes

Personalization is not one switch. Modern AI assistants combine several layers: response style, explicit instructions, saved memory, past chats, connected apps, and separate privacy and training controls. Providers document these as distinct features that can be managed separately [1-9].

Six separate control layers

One control does not replace another. For example, disabling memory does not automatically erase chat history or disconnect data sources.

Response style	Personality, warmth, length, formatting	changes presentation
Direct instructions	Language, role, preferred workflow	explicit and global
Saved memory	Long-term facts and preferences	must be reviewed and corrected
Past chats	Context from earlier conversations	may reuse outdated assumptions
Connected services	Email, files, calendar, browser or app context	expands the data scope
Privacy & training	Storage, deletion, activity and model improvement	configure separately

Figure 2: The main personalization layers and what they affect.

Especially important: “memory off” does not automatically mean “chat history deleted”. Likewise, “do not use for model training” is not the same as “no personalization”. Review each area separately.

2. Why control matters

When configured well, personalization saves time and makes responses more consistent. When left unchecked, it can carry outdated assumptions, incorrect memories or unnecessary personal details into new conversations. Convenience should therefore be balanced with transparency and data minimization.




Figure 3: Benefits and risks need to be balanced deliberately.

What proper control improves

- Transparency: you know why a response feels personal.
- Correctness: outdated or incorrect memories are discovered earlier.
- Privacy: fewer unnecessary details are stored long term.
- Consistency: global and project-specific instructions conflict less often.
- Independence: you can recognize when a recommendation may be shaped by earlier preferences.

3. Risks and common misunderstandings

- **Incorrect memories:** The system may infer inaccurately what is important long term. Review stored information actively.
- **Outdated context:** Location, role, devices, goals or preferences may change while the memory remains.
- **Confirmation bias:** A heavily personalized assistant may continue existing views too comfortably. Ask explicitly for counterarguments and alternatives.
- **Global overreach:** Very detailed global instructions can unnecessarily constrain new tasks or conflict with project instructions.
- **Connected data sources:** Email, files, calendars, browser context or other apps expand available context and therefore the need for review.
- **Deletion confusion:** Deleting a chat, memory or activity may have different effects. Providers document these areas separately [1-9].
- **Shared devices:** Without separate accounts, preferences and private information from different people can become mixed.

Check question: Would I be comfortable if this information unexpectedly appeared in another answer six months from now? If not, it should not be stored as persistent personalization.

4. What data should be stored long term

The following traffic-light model is not a legal classification; it is a practical decision aid. Be especially cautious with health, financial, family and employment information. Personal data about other people should generally not be stored as your own memory.

What information belongs in personalization?

A practical classification for personal accounts. When uncertain, store less and provide the information only in the current chat.




Figure 4: Practical data traffic light for persistent personalization.

Temporary or incognito chats

ChatGPT, Claude and Gemini provide modes that avoid persistent memory; Microsoft separates personalization, history and training through different controls depending on the Copilot surface. These modes reduce later reuse but do not automatically mean that no technical processing or short-term retention occurs [3, 4, 6, 9].

5. Comparison of the four AI systems

The underlying logic is similar, but labels and switches differ. The key is not to memorize a specific menu term, but to ask the same six control questions in every system.

System	Direct instructions	Memory / past chats	Temporary / private	Additional context	Key checks
ChatGPT	Personality, characteristics and custom instructions	Saved memories and reference to past chats	Temporary Chat	Projects, files and connected services	Review memory, custom instructions and model training separately
Microsoft Copilot	Personalization and, depending on version, custom instructions	Memory and conversation history	availability varies by surface	Microsoft services, Edge context, files and apps	Personalization, training, history and Context Clues
Claude	Profile and project instructions plus styles	Chat search and memory	Incognito Chat	Projects, skills, files and connectors	Global versus project-specific instructions
Gemini	Saved information / instructions	Memory from past chats	Temporary Chat	Google apps, activity and connected services	Gemini Apps Activity, memory and connected apps

The product-specific guides show the current menu paths. This foundation guide explains why each area should be reviewed separately.

Six control questions for every system

- What global instructions are active?
- What has been stored about me long term?
- May the assistant use past chats?
- Which apps, files or account data are connected?
- Can my content be used to improve models?
- How do I delete or restrict each of these layers?

6. Initial setup and regular review

The initial setup usually takes ten to fifteen minutes. Start deliberately with a small number of stable instructions. Add memory only when its practical benefit is clear.

The personalization control cycle

Personalization is not a one-time setup. Review it after major updates and at least quarterly.



Figure 5: Personalization as a recurring control cycle.

- Set language, form of address, desired length and formatting.
- Open saved memories and review every item.
- Enable use of past chats only when continuity is wanted.
- Restrict connected apps and permissions to what is necessary.
- Review privacy and model-training controls separately.
- Test with three tasks: a factual question, a how-to and a controversial comparison.
- Record the date and key switches or take a screenshot for your own reference.
- Review after major product updates and at least every three months.

7. Decision guide for common situations

Which mode fits the task?




Figure 6: Choosing between global, project-specific and temporary use.

- **Everyday assistant:** Configure language, tone and format globally; keep memory small and review it regularly.
- **Ongoing project:** Use project instructions, project files and project-specific chats; avoid storing details globally without need.
- **Health, finance or legal topics:** Provide only necessary information, prefer temporary modes, verify independently and do not replace professional advice.
- **Shared device:** Use separate accounts, sign out and review saved memories.
- **Creative or public content:** Persistent style preferences can be useful; private biography and sensitive details are better left out.

8. Personal checklist

✓ Check

- ☐ I know which global instructions are active.
- ☐ I have read every saved memory and corrected inaccurate information.
- ☐ I know whether past chats may be used for new answers.
- ☐ I know which apps, files and account data are connected.
- ☐ I reviewed model training separately from personalization.
- ☐ I use temporary or project-specific chats when global storage is unnecessary.
- ☐ I do not store passwords, access credentials or confidential third-party data.
- ☐ I verify controversial or important answers beyond my personal preferences.
- ☐ I repeat the review after major updates and at least quarterly.

9. Official sources

This guide is based on the following official help pages. Feature names and availability may change, so consult the current provider help when configuring a specific account.

- [1] OpenAI - Memory FAQ: <https://help.openai.com/en/articles/8590148-memory-faq>
- [2] OpenAI - Data Controls FAQ: <https://help.openai.com/en/articles/7730893-data-controls-faq>
- [3] OpenAI - Temporary Chat FAQ: <https://help.openai.com/en/articles/8914046-temporary-chat-faq>
- [4] Microsoft - Copilot privacy controls: <https://support.microsoft.com/en-us/microsoft-copilot/microsoft-copilot-privacy-controls>
- [5] Microsoft - Conversation history in Microsoft Copilot: <https://support.microsoft.com/en-us/microsoft-copilot/conversation-history-in-microsoft-copilot>
- [6] Anthropic - Understanding Claude's personalization features: <https://support.anthropic.com/en/articles/10185728-understanding-claude-s-personalization-features>
- [7] Anthropic - Use Claude's chat search and memory: <https://support.anthropic.com/en/articles/11817273-using-claude-s-chat-search-and-memory-to-build-on-previous-context>
- [8] Google - Get personalization in Gemini Apps: <https://support.google.com/gemini/answer/16598623?hl=en>
- [9] Google - Gemini Apps Privacy Hub: <https://support.google.com/gemini/answer/13594961?hl=en-GB>

Core message: a well-personalized AI is not the one that knows as much as possible about you, but the one whose knowledge you understand, limit and review regularly.