

# THE AI LEARNING KILL-SWITCH

Turning Off AI Training in 6 LLMs



Targeting Users of Perplexity, ChatGPT, Grok, Gemini, Claude, Copilot

**Peter Serzo**

**Copyright © 2026 Peter Serzo**

All rights reserved.

No part of this book may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

This is a free giveaway edition.

# Table of Contents

- Table of Contents.....2**
- Introduction.....3**
- Terminology.....4**
  - DPA Availability by Provider.....5
- General Compliance Expectations.....6**
- Perplexity.....7**
  - Free/standard Pro: Turn it Off.....7
  - Enterprise Pro: Training is already off.....7
- GEMINI.....8**
  - Consumer/Personal accounts: Turn it off.....8
  - Business/Enterprise: Gemini for Google Workspace.....8
  - Important limitations and exceptions.....9
- Claude.....10**
  - Consumer: Free/Pro/Max.....10
  - Business/Enterprise.....10
  - Important caveats.....10
- Grok.....12**
  - Consumer (Web): Turn off Grok training on X.com.....12
  - Consumer (Mobile): Turn off Grok training in the X mobile app.....12
  - Business/Enterprise: Turn off Grok.....12
  - Important Caveats.....13
- Copilot.....14**
  - Free Plan: Turn off Copilot training.....14
  - Consumer (Windows): Turn off Copilot training.....14
  - Enterprise: Microsoft 365 Copilot in a firm.....14
- ChatGPT.....15**
  - Free/Consumer: ChatGPT turn off training.....15
  - Extra: Global opt-out via OpenAI privacy portal.....15
  - Business / enterprise.....15
- You are in Control.....16**
- ABOUT THE AUTHOR.....17**

# Introduction

After working with hundreds of organizations in my 30 plus career and the last three as the AI Enterprise Architect at a large bank I can safely say this: Your team isn't waiting for policy - they're already deep in the use of GenAI tooling. This is Shadow AI. Using unapproved AI tools.

From drafting client emails in ChatGPT to tweaking operating procedures with Copilot, GenAI crept into your firm like WiFi and smartphones did. These tools are fast, convenient, and quietly indispensable. You didn't approve it, but it's here, living in Excel, Outlook, Chrome, Google. Platforms like Quickbooks, Taxdome, and CRMs are baking it into their architectures as core functionality.

Employees use public models without realizing those tools might be learning from your firm's data (sensitive and/or internal). The same prompts that save an hour on a tax memo could help train the next version of a model with your company information inside it.

The Big 6 we are covering in this eBook: Perplexity, Grok, Gemini, ChatGPT, Claude, and Copilot will soon be essential business tools if they are not already. Unless you configure the LLMs **privacy and data-sharing settings** correctly, they will quietly absorb company IP, client information, compliance documents or whatever else it is fed..

The problem is you can't secure what you don't control. These tools live beyond your servers and update faster than your policies. Most firms don't have the luxury of a dedicated AI security team. But protecting your clients, your reputation, and your compliance obligations under GLBA, FTC, IRS, and CCPA demands action.

This short guide gives you the "what to click and why" for the key settings every organization should enable or disable to keep the Big 6 from training on your data.

No jargon, no coding, no panic. These settings work whether you are an enterprise or a "Free" or "Pro" user.

I want to stress: Configuration is only one small step in the journey to prioritize the security of your data. Sonareon ([Sonareon.com](https://sonareon.com)) can build what your company needs with our AI Framework. We cover everything from WISPs, to Audits, to governance, AI Policy and acceptable use. No matter how small or large, we collaborate with you on the tooling you decide to utilize.

## Terminology

To understand the settings of each GenAI tool we must first have a common language. All these vendors provide Terms of Service/Use that you agree to before using their tooling. While many of us don't read these (I have been guilty) it is important to understand common terms. The following is a table of some of them:

LLM	A type of artificial intelligence designed to understand and generate human language at scale.
Provider/GenAI	Generative AI: For the purpose of this document this is constrained to the big 6 LLMs mentioned in the Introduction.
DPA	Data Processing Addendum: the provider acts as a " <b>Processor</b> ," meaning they only handle data based on your organization's specific instructions. This is standard and covered next.
Ownership	You generally own the content you create.
Human in the loop	A human is actively involved in the review, decisioning and input (prompting) interactions with the LLM.
Prompt	This is the input you give to the LLM. It can be in the form of text, image, audio, or document.

## DPA Availability by Provider

Provider	DPA Status	Where to Find It / How to Sign
Gemini	Standard	Included in the <b>Google Cloud Data Processing Addendum (CDPA)</b> for Workspace and Google Cloud users.
ChatGPT	Standard	Available for <b>ChatGPT Enterprise, Business, and API</b> users. It is often a click-through agreement in the admin console.
Claude	Standard	Automatically incorporated into Anthropic's <b>Commercial Terms of Service</b> (API and Claude for Work).
Copilot	Standard	Covered under the massive <b>Microsoft Products and Services Data Protection Addendum (DPA)</b> for all M365 Business/Enterprise tiers.
Perplexity	Available	Perplexity has a formal DPA for <b>Enterprise Pro</b> users. They explicitly state that data under this DPA is not used for model training.
Grok (xAI)	Available	xAI provides a DPA for <b>Enterprise/API</b> customers. It governs personal data processed on behalf of the customer.

## General Compliance Expectations

- Do not put PII data in prompts
- The big 6 should be in your written information security program (WISP) limiting what the organization can utilize
- Your AI Usage Policy should spell out the guidelines, permissions, customer facing impact, user responsibilities and roles at a minimum

# Perplexity

Perplexity's switch is called AI data retention (for consumer tiers) and there's a separate hard guarantee for Enterprise Pro where training is off by design.

## Free/standard Pro: Turn it Off

On web or mobile:

1. Log into Perplexity.
2. Click your profile icon → Settings.
3. Go to the Privacy or Data & privacy section.
4. Find AI data retention (or similar wording about using your data to improve AI).
5. Turn AI data retention off to stop your prompts being retained and used for training/fine-tuning.

What this does:

- Prevents your future prompts/threads from being used to train or fine-tune Perplexity's models (beyond minimal, short-term operational logging).
- May also shorten how long your data is stored. Older threads can be deleted after a defined window.

If you want to go further, manually delete past threads that contain sensitive content so they aren't available for any future use.

## Enterprise Pro: Training is already off

For Enterprise Pro:

- Perplexity explicitly states that enterprise data is never used to train or fine-tune their AI models.
- Their contracts with third-party model providers (OpenAI, Anthropic, etc.) also forbid those providers from training on Perplexity enterprise data.
- Admins can configure retention (for example, 30 days) after which threads are purged.

So on Enterprise Pro, you do not need to flip a training toggle; the "no training on your data" stance is built into the plan.

# GEMINI

Think of Gemini like three different garages with the same logo: consumer Gemini, Gemini Advanced, and Gemini for Workspace/Vertex. Same brand, very different data and PII handling expectations.

Gemini's "training switch" lives in one place: Gemini Apps Activity. Turn that off, and you stop your chats being used to improve Google's models going forward.

## Consumer/Personal accounts: Turn it off

On the web:

1. Go to the Gemini app ([gemini.google.com](https://gemini.google.com)) and make sure you're signed into the right Google account.
2. Click your profile picture.
3. Click on the Activity icon (clock with an arrow) or the link to Gemini Apps Activity. This takes you to your Google Account activity page for Gemini.
4. At the top, click Turn off for Gemini Apps Activity (sometimes labeled "Keep Activity" in newer UI).
5. Confirm when prompted. You can optionally choose to delete existing Gemini Apps Activity so past chats/uploads aren't kept for training.

On Android / iOS Gemini app:

1. Open the Gemini app.
2. Tap your profile picture in the top-right.
3. Tap Gemini Apps Activity.
4. Tap Turn off, then confirm; you can also select "Turn off and delete activity" if you want to wipe past history as well as stop future training.

What this does:

- Stops Google from using your future Gemini chats and uploads to improve machine-learning technologies and products.
- Stops saving new conversation history in Gemini. You lose cross-session personalization and won't see your past chats list going forward.
- You can still use Gemini normally; it just won't log and reuse your conversations for product improvement.

## Business/Enterprise: Gemini for Google Workspace

When you use Gemini in Google Workspace (Docs, Sheets, Gmail, Chat, Gemini app under your org domain), data is treated under the Workspace DPA and enterprise privacy commitments.

Google states that your Workspace Gemini content is not used to train generative models outside your domain, and is not used for other customers.

Interactions stay within your organization, with strict data access controls that prevent cross-user/session leakage and respect existing Drive/Gmail permissions.

## Important limitations and exceptions

- Even with Gemini Apps Activity off, Google may retain some data briefly (for around 72 hours) for security, abuse detection, and reliability.
- Gemini Live audio/video has its own checkbox like “Improve Google services with your audio and Gemini Live recordings”; uncheck that too if you use Live.
- Gemini API / Vertex AI: public docs and community threads indicate there is currently no general UI toggle to opt out of product-improvement use for some Gemini API scenarios; you rely instead on Google Cloud’s enterprise data-use terms and should avoid sending unnecessary PII in those contexts
- For regulated environments, Google pushes you toward Workspace Gemini and Vertex AI, which inherit enterprise security, compliance certifications, and admin controls
- Gemini Advanced (paid consumer) - Gemini Advanced is still an individual consumer service; it adds capabilities and storage, not a fundamentally different data-controller relationship

## Claude

Claude recently pulled a “terms of service plot twist”: consumer chats now feed training by default unless you flip a switch. Here’s how to flip it.

### Consumer: Free/Pro/Max

The control is the “Help improve Claude” toggle.

On web:

1. Sign in at [claude.ai](https://claude.ai).
2. Click your user icon (bottom-left).
3. Click Settings.
4. Go to Privacy → Privacy settings.
5. Find “Help improve Claude” (or “You can help improve Claude”).
6. Turn this switch off to opt out of using your chats and code for training.

On mobile app:

1. Open Claude and log in.
2. Tap the menu (stacked lines).
3. Tap the Settings gear.
4. Tap Privacy.
5. Toggle “Help improve Claude” to off.

You can also opt out when the “Updates to Consumer Terms and Policies” popup shows: uncheck/turn off the “You can help improve Claude” option before hitting Accept.

## Business/Enterprise

Anthropic has a strict policy of **not training** its foundational AI models on data from its business and enterprise customers. Under the **Claude for Work**, **Claude Enterprise**, and **API** (including Amazon Bedrock and Google Cloud Vertex AI) terms, your prompts and outputs are excluded from training pipelines by default.

This creates a significant “privacy wall” compared to consumer accounts (Free, Pro, and Max), which are subject to a policy where data may be used for training unless the user explicitly opts out. For business users, this protection is legally reinforced through a **Data Processing Addendum (DPA)** and the **Commercial Terms of Service**, ensuring that proprietary company data remains private and sandboxed.

## Important caveats

- Opt-out is not retroactive: data already used for training cannot be removed from models.

- Opt-out affects future chats (and old chats only if you reopen them after opting in). Deleting a conversation keeps it out of training.
- If you do nothing when prompted, your chats and code are treated as training data and may be stored for up to five years. Opting out keeps you on the shorter retention regime.
- Business / enterprise offerings (Claude for Work, Gov, Education, API via Bedrock or Vertex) are not subject to this consumer training change. Those chats are excluded from model training under their own terms.

# Grok

Grok is basically “on” for training by default, and the kill switch lives in your X privacy settings, not in the Grok UI itself.

## Consumer (Web): Turn off Grok training on X.com

1. Open X (x.com) in a browser and log into your account.
2. Click More in the left sidebar (three-dot “More” menu).
3. Click Settings and privacy.
4. Go to Privacy and safety.
5. Find Data sharing and personalization, then click Grok (or “Grok & third-party collaborators”).
6. In Data sharing, uncheck or toggle off:
  - “Allow your posts as well as your interactions, inputs, and results with Grok (and X AI) to be used for training and fine-tuning.”
7. (Recommended) Click Delete conversation history to remove existing Grok chats from their systems (subject to up-to-30-day retention for security/legal).

## Consumer (Mobile): Turn off Grok training in the X mobile app

1. Open the X app on your phone.
2. Tap your profile icon to open the side menu.
3. Tap Settings & Support, then Settings and privacy.
4. Tap Privacy and safety.
5. Tap Grok (or “Grok & third-party collaborators”).
6. Turn off the toggle for “Allow your posts as well as your interactions, inputs, and results with Grok to be used for training and fine-tuning.”
7. Optionally tap Delete conversation history to wipe existing Grok chats.

## Business/Enterprise: Turn off Grok

Grok Business and Grok Enterprise the policy is a “zero data training” commitment, meaning your business data, prompts, and outputs are never used to train their foundational models. This is a standard feature across both paid tiers, ensuring that proprietary information remains isolated from the public training corpus used for consumer versions.

For companies with extreme security needs, xAI also offers an **Enterprise Vault**, which provides an isolated data plane and customer-managed encryption keys to further sandbox sensitive inputs.

## Important Caveats

- Turning training off stops X/xAI from using your posts + Grok interactions for model training and fine-tuning going forward.
- Making your X account private also keeps future tweets out of Grok's training crawl, but the explicit Grok toggle is the more direct control.
- xAI still keeps some data for security, abuse detection, and legal reasons, and anything trained before you opted out is not pulled back out of existing models.

# Copilot

Copilot has its own model-training toggle.

## Free Plan: Turn off Copilot training

The free Copilot plan doesn't include any controls to turn off training or adjust how your data is used. Consumer-tier services generally bundle data-use policies into the product with no switches to change them, and those kinds of privacy guarantees only appear in business or enterprise offerings. For the free plan, everything is governed by Microsoft's published privacy statement, which is the definitive source for how data is handled.

[Microsoft Privacy Statement – Microsoft privacy](#)

## Consumer (Windows): Turn off Copilot training

1. **Open Copilot from Windows:** Click your Profile Icon at the bottom left.
2. **Access Privacy Settings:** Click on your **Account Name/Email**, then select **Settings, Privacy** from the dropdown menu.
3. **Toggle Off Training:** Look for the section labeled **Model training on text** and **Model training on voice**. Switch both of these toggles to **Off**.

Optional cleanup:

- In the same Privacy area, choose Export or delete history and delete Copilot app / M365 / Windows activity history so past chats aren't kept around.
- Turn off Personalization and memory and click Delete memory so Copilot forgets what it has stored about you.

## Enterprise: Microsoft 365 Copilot in a firm

- Microsoft 365 Copilot with Commercial Data Protection is already designed so your tenant's data is not used to train Microsoft's foundation models for other customers; it stays within your tenant boundary.
- As an admin, you control whether Copilot is available at all via the M365 admin center or Group Policy:
  - Admin portal → Org settings → Microsoft 365 Copilot → turn off per app or per user.

# ChatGPT

For ChatGPT, the “training switch” lives under Data controls. Turning it off stops future chats being used to train models.

## Free/Consumer: ChatGPT turn off training

On desktop or mobile app:

1. Open ChatGPT and log in.
2. Click your profile (bottom-left on desktop, top-right or menu on mobile).
3. Click Settings (or “Settings & Data”).
4. Go to Data controls.
5. You’ll see one of these options, depending on your UI:
  - Chat history & training → toggle Off, which both:
    - stops saving new chats in history, and
    - stops using them for model training.
  - or Improve the model for everyone → toggle Off to stop your conversations being used to train models while keeping history if available

Important behavior:

- When Chat history & training is off, new chats are not used for training and also won’t appear in your history (it’s session-only).
- This only affects future conversations; it doesn’t untrain anything already learned from past chats.

## Extra: Global opt-out via OpenAI privacy portal

OpenAI also lets you make a broader request:

1. Go to [privacy.openai.com](https://privacy.openai.com) and sign in.
2. Submit a request under “Messages used to train our models”.
3. Specify that you want to opt out of having your messages used for model training.

## Business / enterprise

- ChatGPT Teams / Enterprise / Business / API: OpenAI states that business data and API data are not used for training by default under those plans, so there’s no per-user toggle needed for that purpose.
- You still manage retention and access (SSO, role-based access, org policies), but training on your org’s data is already off by contract

## You are in Control

AI is powerful.

AI is transformative.

AI is everywhere.

But you decide what it learns from your organization and you.

This book gives you the switches, the levers, and the clarity to stay in control.

Your data.

Your rules.

Your intelligence – not the machine's.

I want to stress this is one small component in your organization's journey to protecting your most valuable asset: Customer Data.

There are many other factors like AI Use Policy, training, roles, AI governance and more.

Do not hesitate to reach out to me if you would like to discuss implementing the

Sonareon AI Framework or having us create a WISP. Contact me here:

[pserzo@sonareon.com](mailto:pserzo@sonareon.com)

Thank you

Peter Serzo

## ABOUT THE AUTHOR



Oldsmobile Customer Assistant. Software Developer to Enterprise Architect. Peter (with a great team) put Comerica Bank in the cloud with SharePoint. Onboarded the company onto AI and developed the AI assessment. Peter is also an author and speaker. However, he is most proud of building partnerships, listening and understanding business pain reducing the friction of processes and technology. You can find out more about him at [Sonareon.com](https://sonareon.com) or view his digital resume at [molescout.com](https://molescout.com).