AI Privacy & Cost Cheat Sheet for Screenwriters

The question that matters on an **unpublished script**: *"if I paste my pages in, will they be used to train the model?"* Most **consumer** plans train by default — switch it off; **Business / Team / Enterprise** never do. Read each row: **your plan → is it private? → what to do → what it costs.** *(Opt-out steps repeated on page 2.)*

PROVIDER	PLANS — IN ORDER	TRAINS ON YOUR WORK?	WHAT TO DO TO KEEP YOUR SCRIPT PRIVATE	COST / MONTH — SAME ORDER
ChatGPT OPENAI · US	**Free · Go · Plus · Pro**	✗ Trains by default	Settings → Data Controls → turn off **"Improve the model for everyone"** — or use a **Temporary Chat**.	$0 · $8 · $20 · $100/$200 Pro has two usage tiers, same models
	Business / Enterprise	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts.	~$20–25 / user
Claude ANTHROPIC · US	**Free · Pro · Max**	✗ Trains by default	Settings → Privacy → turn off the **model-training toggle**. (Off = 30-day retention; on = 5 years.) You are auto opted-*in* on all three — you must actively decline.	$0 · $20 · $100/$200
	Team / Enterprise / API	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts. Safety-flagged chats can still be used — see p.2.	$25 / $125 per seat
Gemini GOOGLE · US	**Free · AI Plus · AI Pro · AI Ultra**	✗ Trains by default	**myactivity.google.com/product/gemini** → turn off **"Keep Activity"** (was 'Gemini Apps Activity') — or use a **Temporary Chat**. Keep Activity is ON on every consumer tier; humans may review a sample.	$0 · $4.99 · $19.99 · $100/$200 Ultra restructured May 2026
	Workspace (business)	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts.	from ~$14 / user
Copilot MICROSOFT · US	**Free · Microsoft 365 Premium**	✗ Trains by default	Account → Privacy → turn off **"Training on conversation activity."**	$0 · $19.99 Premium replaced Copilot Pro, which stopped renewing 1 Aug 2026. EU: €9.99 / €12.99 Family
	Microsoft 365 Copilot (work)	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts. Enterprise Data Protection. In the EU, also disable **flex routing** — see p.2.	$30 / user
Le Chat MISTRAL · EU	**Free · Pro**	✗ Trains by default	Privacy → turn off **"Allow your interactions to be used to train our models"** — or, on Pro, switch on **No Telemetry Mode**. EU-based: **easiest EU/GDPR route**, see p.2. Mistral's consumer terms are ambiguous — set the toggle regardless.	$0 · €14.99 €5.99 students
	Team / Enterprise / API	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts. No-training default; EU-first (some US subprocessors).	business pricing
Perplexity US	**Free · Pro · Max**	✗ Trains by default	Settings → Preferences → turn off **"AI data retention."**	$0 · $20 · $200
	Enterprise Pro · Enterprise Max	✓ Private — never trains	Nothing — private by contract. Safe to paste scripts. Only Enterprise **Max** adds configurable retention and audit logs.	~$40 · ~$325 / user
Grok XAI / X · US	**All consumer plans** Lite · SuperGrok · Heavy	✗✗ High risk	Settings → Data Sharing → off — but guarantees are weak. **Avoid for unpublished IP.** Trains on your chats *and* your X posts.	$10 · $30 · $300
	Business / Enterprise	✓ Private — never trains	Nothing — private by contract. No-training guarantee — new, and the only safe way to use Grok.	from ~$30 / seat
Meta AI LLAMA · US	**All consumer plans**	✗✗ High risk	**Don't paste scripts.** No usable opt-out. (Llama is open — self-host it, see p.2.) EU training legally contested; complaints open in 11 member states.	Free
DeepSeek CHINA	**All consumer plans**	✗✗ High risk	**Don't paste scripts.** (Open weights — self-host for full privacy.) Data stored in China; blocked in Italy, ruled unlawful in Berlin.	Free

✗ **Trains by default** = on by default; you can opt out (see 'What to do') ✓ **Private** = never trains, safe to paste scripts ✗✗ **High risk** = avoid for unpublished work Prices monthly, **August 2026**, as published by each provider.

How to lock it down — and which to choose

✓ Turn off training (step by step)

If you only do one thing: open data/privacy settings and switch the "improve / train the model" toggle to **OFF**. It's buried but takes seconds.

ChatGPT: Settings → Data Controls → *Improve the model for everyone* = Off; or click **Temporary Chat** per session.

Claude: Settings → Privacy → *model-training toggle* = Off (keeps retention to 30 days).

Gemini: myactivity.google.com/product/gemini → *Keep Activity* = Off (was 'Gemini Apps Activity'); or use a **Temporary Chat**.

Copilot: Account → Privacy → *Training on conversation activity* = Off.

Le Chat: Privacy → *Allow your interactions to be used to train our models* = Off (needed on Free **and** Pro); Pro and above can also switch on **No Telemetry Mode**.

★ What to actually use

Best for protecting an unpublished script (ranked):

1. A **local open model** (Ollama / LM Studio) — 100% private, free.
2. A **Business / Team / Enterprise / API** seat on ChatGPT, Claude, Gemini-for-Workspace or M365 Copilot — contractually never trained, with a data-processing contract.
3. **Mistral** on Team/ Enterprise — EU data residency, no-training by default; the easiest *EU-native* paperwork. (Pro needs the opt-out above.)
4. A personal account with **training OFF + temporary chat** — OK for everyday drafting, *not* for confidential pages.

For a production company, school or writers' room: a shared **Team plan** (~$25–30/seat) gives every writer a no-train account + central control + a DPA — far better than policing personal toggles.

GDPR for EU writers — it's the *tier*, not the brand

1. Free / personal accounts are never GDPR-clean. Every provider (ChatGPT, Gemini, Claude, Mistral, Copilot, Perplexity) trains on your input by default *and* gives you no data-processing contract — don't put scripts or real people's data into one.

2. Mistral is the easiest EU option, not the only one. ChatGPT, Gemini, Claude and Microsoft Copilot are equally GDPR-usable on their **Business / Enterprise / API** tiers — each offers a **DPA** + the **EU-US Data Privacy Framework + Standard Contractual Clauses**. **Lean on the SCCs:** the Framework survived its first court challenge but is under appeal at the EU's top court (decision expected 2027), and a June 2026 US Supreme Court ruling on FTC independence has EU regulators formally asking Brussels to review it. Google, Microsoft and OpenAI also offer **EU data-residency regions**; Claude's EU residency is only via AWS or Google, not first-party. GDPR compliance is a legal setup, not a postcode.

3. Safest of all: run a model locally (Mistral / Llama via Ollama) — no transfer, no training, no GDPR exposure. Ideal for unpublished work.

4. Two layers of risk: GDPR bites when your text names *real people*; confidentiality / IP risk applies to *every* unpublished script — so use a no-train plan and switch the opt-out on regardless.

5. For any company or school: sign the **DPA (AVV)** + run a **DPIA (DSFA)** before rollout; on Microsoft, disable **"flex routing"** (on by default since 17 Apr 2026 for EU/ EFTA tenants: Copilot → Settings → View all → Flex routing) to keep inference in the EU Data Boundary.

Hard no: DeepSeek — blocked in Italy and declared unlawful by Berlin's regulator (though Apple and Google refused to delist it, so it's still downloadable in much of the EU) — and **Grok** on consumer plans (live EU inquiries). **Meta AI** is legally *contested*: complaints are open in 11 member states. Avoid all three for confidential work.

⚠ Other red flags

'Safety review' carve-outs: Anthropic's policy (effective July 2026) lets *flagged* chats train even after opt-out (inputs and outputs kept up to 2 years, safety scores up to 7, and it may trigger ID/age verification) — opting out is not absolute. **Deleted ≠ gone:** ChatGPT Temporary Chats aren't trained on but are kept ~30 days, and a provider under a legal hold can be ordered to preserve logs you 'deleted' — OpenAI's New York Times order now covers only a defined Apr–Sep 2025 set and *excludes* EEA / Swiss / UK conversations, but the precedent stands. **Watch the tier wording:** 'Pro' rarely means private (it's still a consumer plan) — only Business / Team / Enterprise / API carry the no-training contract.

● The gold standard: run it on your own machine (free, fully private)

Open-weight models you can download and run offline — nothing ever leaves your laptop, so there is no training and no retention to worry about:

Tools: Ollama or LM Studio (free, point-and-click). **Good models to run:** Llama (Meta), Mistral, Qwen, DeepSeek — the *downloaded* versions are private even though the hosted apps are not. **Trade-off:** needs a reasonably powerful computer (ideally 16 GB+ RAM / Apple Silicon or a decent GPU) and is a little less capable than the top cloud models — but for sensitive IP it's unbeatable on privacy.