ChatGPT

AI Chat Tools and Your Data: Privacy Policies Explained

Understanding what happens to your data when you use AI chat tools is important. Below we summarize the data and privacy policies of major consumer-facing AI assistants – OpenAI's **ChatGPT**, Anthropic's **Claude**, Google's **Gemini** (e.g. Bard), **Meta AI**, **Microsoft Copilot**, Elon Musk's xAI **Grok**, the Chinese chatbot **DeepSeek**, and Google's **NotebookLM**. We'll cover what data each collects, if your conversations are used to train the AI, whether you can opt out, how long data is kept, and any unique privacy measures. Our aim is to keep it simple and non-technical, with short sections and a handy comparison table.

OpenAI ChatGPT (Consumer Version)

Data Collected: ChatGPT collects the content of your conversations (prompts and AI responses) and any files or images you upload. It also gathers account details (like your email, name, and payment info if you subscribe), as well as technical data like IP address, device type, and usage logs. Essentially, anything you type or share with ChatGPT is stored as "User Content," alongside metadata about how you use the service.

Use for Training: By default, OpenAI *does* use your conversations to train and improve its models. Unless you opt out, the prompts and responses may become part of the training dataset to make ChatGPT better over time. OpenAI confirms that it **saves your prompts and uploads by default** and uses them to enhance AI model performance. However, these chats are **not used for personalized ads or sold to third parties**.

Opt-Out Options: Yes – ChatGPT provides controls to **opt out** of having your chats used for training. In ChatGPT's settings, you can turn off the setting labeled "Improve the model for everyone" (previously "Chat History & Training"). With this turned off, your conversations **will not be used to train the AI**. They will still be saved in your account's chat history (unless you use the separate "Temporary Chats" feature), but won't feed into model improvements. OpenAI also introduced a "no training" mode where you can disable chat history entirely – when activated, new chats won't appear in your history sidebar and won't be used for training. Even if you don't use the settings, OpenAI offers a privacy portal form to opt out of data use, but the in-app toggle is the easiest way.

Data Retention: How long ChatGPT keeps data depends on your settings. If you turn off chat history (the "do not train" mode), OpenAI still retains those new conversations for **30 days** for abuse monitoring, then **permanently deletes** them. For chats where history/training is ON, OpenAI's privacy policy says it retains personal data as long as needed to provide the service or for other legitimate purposes. In practice, your chat history will remain accessible in your account **until you delete it**. (OpenAI's help center indicates chats are stored indefinitely unless you delete them manually.) You can delete individual conversations or your whole account, and OpenAI will then remove that data from its systems (typically within 30 days for account deletion requests). So, with history on, chats stick around; with history off, they're short-lived (one month).

Unique Privacy Points: OpenAI allows you to **export your data** for transparency (via a settings option). They do not use your ChatGPT content to target ads, and they encrypt data in transit and at rest. By default

your chats may be viewed by OpenAI staff for moderation or improvements, but opting out stops that use. In summary, **ChatGPT saves what you type by default and uses it to get smarter**, but you have a clear option to opt out and even have "no-history" chats for privacy. Remember, if you need to share highly sensitive info, it's best to turn history off or not use the service for that content.

Anthropic Claude (Consumer Version)

Data Collected: Claude collects your conversation content (prompts and responses) just like other chatbots. If you use Claude's integrations (like the Gmail or Google Drive plugins), it may process that data during use, but Anthropic says integrated content isn't stored as raw training data unless you explicitly include it in the conversation. They also collect basic account info (email, etc.) when you sign up, and some usage and technical info automatically (for security and analytics) as described in their privacy policy. In short, **whatever you input into Claude is stored**, and feedback you provide (thumbs up/down ratings) is also logged.

Use for Training: By default, **Anthropic does *not* use your prompts or conversations to train Claude**'s underlying model. This is a key difference – Claude is designed with a privacy-first approach where your chat content isn't fed back into improving the AI model unless you allow it. There are only two exceptions: (1) If you **explicitly submit feedback** (e.g. using the thumbs-up/down or reporting a problem), that conversation is saved as feedback and *may* be used to improve the model after being disassociated from your identity. (2) If your content is flagged for violating the usage policies (e.g. the AI detects harmful or abusive input), Anthropic may analyze and use those flagged conversations to improve its safety systems and filters. Even then, they state that only a limited number of staff can access conversation data and only for legitimate purposes. **Outside of those cases, your chat inputs and Claude's replies are not used to train the AI's next version by default**. This means Claude isn't learning from your specific conversations in the way ChatGPT might from its users.

Opt-Out Options: Since Anthropic doesn't train on user chats by default, there isn't a need for a training opt-out setting – it's already "opted out" by design. (In fact, you'd have to opt *in* by providing feedback if you wanted your data used.) Claude also offers a Claude Pro tier, but it follows the same rule: your chats aren't used for model training unless you give explicit permission. However, Anthropic does allow deletion of your conversation history (for privacy) – you can delete chats from the Claude interface, and they will be removed from your account history and Anthropic's backend after a short period.

Data Retention: Anthropic retains user data **only as long as necessary** for the purposes described. For normal conversations, if you do nothing, your chat history will be saved in your account so you can refer back to it, but Anthropic hasn't published a specific timeframe for auto-deletion of undeleted chats. Notably, if you **delete a conversation**, it's immediately gone from your chat history and is **permanently deleted from Anthropic's backend within 30 days**. So you have control to purge any chats you don't want retained. For conversations flagged by their safety system, Anthropic may retain the data for up to **2 years** (and related safety classifier data for up to 7 years) to improve their filters. Feedback that you submit (thumbs-up/down and comments) can be kept for up to **10 years** in de-identified form for model improvement purposes. But regular, non-flagged, non-feedback chats are generally kept until you delete them, since they aren't training the AI by default. It's a good practice to delete any sensitive chats on your own if you don't want them retained.

Unique Privacy Points: Claude is known for its "privacy by default" stance. Anthropic stresses that Claude's models are trained **not** to reveal personal info and that **user data isn't used for model training unless opt-in**. In practice, this means Claude might be a preferred choice if you're concerned about your chat content being used to teach the AI. They also offer a feature similar to ChatGPT's incognito mode for businesses (Claude for Work) where data can be kept entirely out of training loops. Also, Claude's **data retention is conservative** – anything you delete gets fully wiped from their systems after 30 days. Overall, Anthropic's Claude is among the more privacy-conscious AI assistants available to consumers, which has even been noted by observers (Anthropic's policy is "extremely strict" and conversations are often only kept for a limited time like 30 days by default).

Google Gemini (Bard and Related AI Features)

Data Collected: Google's Gemini-powered AI (which includes **Bard** and other generative AI features in Google products) collects a range of data when you use it. Naturally, it records **your conversations** – the questions or prompts you ask and the AI's responses [1]. If you use voice input or "Gemini Live" voice conversations, it will collect those recordings as well. It also gathers any files or images you choose to share with it, since Bard allows uploading images for analysis – those become part of the conversation data [2]. In addition, Google captures **usage information** like your device and browser details, general location (if location services are on), and how you interact with the AI. They also log **feedback** you provide (such as ratings on answers) to help improve the service. Essentially, **everything you input or do with Bard/Gemini can be logged**: your chats, attached content, feedback, plus technical logs and location info to contextualize your requests [1].

Use for Training: Yes, Google uses data from Gemini app conversations to improve its AI models and services, but with some careful controls. By default, **a subset of user conversations** is reviewed by human evaluators to help refine the AI's quality. Google explicitly states they **select a subset of chats and use automated tools to remove personal identifiers**, then have **trained reviewers (including third-party contractors) read them** to evaluate and improve the AI. These reviewed samples are kept for up to **3 years**, stored separately from your Google Account, and are used to make the models safer and more helpful. In effect, **your Bard chats can end up as part of an anonymized training dataset** for Gemini's ongoing development. However, Google emphasizes that only a portion of conversations are used in this way, not every single chat [3] [4]. They also assure users that **Gemini conversations are not presently used for ad targeting** (Google will notify if that changes). Additionally, if you enable a special personalization feature, Gemini might use your past Google Search history or other Google data to tailor answers – but that's an optional setting and separate from training the model.

Opt-Out Options: Yes – Google provides an **activity control** to opt out of having your Bard/Gemini conversations used for model training or human review. This is handled via the **"Gemini Apps Activity"** toggle in your Google account settings. If you **turn off Gemini Apps Activity**, then **future conversations won't be sent for human review or used to improve Google's AI models**. In other words, turning this off puts your chats in a do-not-train bucket. Google even advises that if you have privacy concerns, you should turn off this history setting so that your chats remain between you and the AI. Keep in mind, even with the setting off, Google will still temporarily process and save your conversations for a short period (see retention below), but not for learning purposes. On the user side, you can also manually **delete your Bard conversation history** either in the Bard interface or via your Google My Activity page. Google offers **auto-delete** schedules too, allowing you to automatically erase your Bard/Gemini activity after 3, 18, or 36

months, etc.. So, you have granular control: you can stop data from contributing to training and you can wipe your logs if desired.

Data Retention: By default (with activity on), your Bard conversations are stored in your Google Account as part of your **Gemini Apps Activity history**. They will stay there until you delete them or until an auto-delete period you set is reached. If you turn **history/activity off**, **Google still retains your recent conversations for up to 72 hours** to maintain service quality. This short retention (about 3 days) is to ensure the service can function properly (for example, to generate answers or address any immediate issues), but those chats won't appear in your account history and aren't used for improvements. In summary: with history on, data is kept until deleted (so it could be long-term if you never clear it); with history off, data is transient (just a few days) then dropped. All users, regardless of setting, can delete specific chats at any time, and doing so removes them from your account and Google's records of your activity. Google's Privacy Notice doesn't list an absolute maximum retention for Bard data, but they abide by your preferences and legal requirements. Also, any conversations selected for human review are stored separately (de-linked from your identity) for up to 3 years for model training purposes. After that period, presumably they are disposed of as well.

Unique Privacy Points: Google is fairly transparent about how it uses Bard data. They have published a detailed **Gemini Privacy Notice** and user FAQ [5]. Unique aspects include the **human review process**: Google tries to strip out things like names, emails, and phone numbers before a person sees your chat [6]. Only a random sample of chats are ever reviewed, and even those are not tied to your account when reviewers see them [6]. If you're uneasy with that, the opt-out is readily available. Google also clearly states that Bard chats are not used to show you ads (and if they ever decide to, they will ask for permission first). Another difference is **account integration**: Bard can optionally use your wider Google account data (like past searches or content from Gmail/Docs if you opt to connect it) to give better answers, but that's under your control. As always, Google cautions not to include sensitive personal information in your prompts – a good rule of thumb for *any* AI chatbot. Overall, Google's approach tries to balance improvement of the AI with user privacy by using sampling, anonymization, and user controls.

Meta AI (Facebook/Instagram's AI Assistant)

Data Collected: Meta AI (the AI assistant available in Facebook Messenger, Instagram, WhatsApp, and the standalone Meta AI app) collects *everything you share in the conversation*. This includes your **chat messages** with the AI (text prompts and the AI's replies) and any images or voice notes you might send to it. Meta's assistant is integrated with your social media profile if you use it through Facebook or Instagram, so it can access **information from your profile** to personalize responses. For example, if you launch the Meta AI app by logging in with Instagram, the AI can see your Instagram public info. It also sends back information about your chats to Meta's servers. Meta AI even creates a separate **"Memory" profile of you** – it keeps a list of key facts it has inferred about you (topics you mention, preferences, etc.) to personalize future interactions. So, not only is Meta storing the raw conversation transcripts, it's also deriving insights about you. Meta AI may also use your general location if available to give local recommendations (for instance, suggesting nearby restaurants). In short, **Meta's AI will remember whatever you tell it – and it actively builds a profile of your interests** to make the experience feel more personal.

Use for Training: Meta has made it clear that **your chats with Meta AI *will* be used to improve their AI models**. In the terms of service, Meta warns users *"do not share information that you don't want the AIs to use and retain."* This implies anything you say may be retained and leveraged to train or refine Meta's AI systems. The **contents of your chats (text, images, voice)** may be fed back into Meta's model training

pipelines. Meta uses these conversations to tune the assistant's responses and improve its accuracy and safety. They have internal reviewers and automated systems that monitor interactions (for instance, to reduce the chance the AI exposes someone's private info in another chat). Unlike some others, Meta *does not offer an explicit opt-out* for this data usage in the consumer version. When you use Meta AI, you are effectively opting in to let Meta use your input to make their AI smarter. Additionally, Meta has indicated future plans to incorporate these AI chats into their advertising systems – for example, using what you discuss to target ads (Meta hasn't started this as of now, but **nothing in their privacy policy forbids using your chat content for ad personalization**). In fact, privacy experts note that Meta could leverage your AI conversations to learn about your interests (say you talk about baby bottles with the assistant, Meta might deduce you're a parent and use that info). So, **yes, Meta AI uses your chats for improvement and likely will use them for personalized features or ads**, much like how they use your activity on Facebook itself.

Opt-Out Options: Meta does not currently provide a user toggle to prevent your AI chats from being saved or used. The service is designed to keep a memory of interactions by default. The only controls available are **deletion commands**: for instance, Meta introduced a "/reset-ai" chat command that lets you delete your AI conversation history across Messenger, Instagram, and WhatsApp. If you invoke "/reset-ai", it will wipe the content of your past conversations with the AI (so it forgets everything you had said previously). You can also manually delete individual AI chat threads in the app, similar to deleting a normal message thread; this will remove those messages from the app's "Discover" feed and your view. However, even if you delete a chat, Meta may still retain the data on their servers for some time (the exact retention post-deletion isn't clearly disclosed, but presumably for policy or legal reasons Meta might keep a copy for a period). **Importantly, you cannot turn off the data collection or retention entirely** – every time you chat with Meta AI while logged in, it will record that conversation and use it. The only way to avoid that is to use the AI without logging into an account (Meta AI's web interface allows anonymous use, which limits features but doesn't tie chats to your identity). In summary, **there is no opt-out switch for training/ personalization** – using Meta AI means your data is in the mix. You do have the ability to delete past chats (so the assistant *forgets* them and they're removed from your profile), but new conversations will still be recorded. Meta's approach is more of an "opt-in by default" – they assume consent to use your data unless you actively clear it or avoid the service.

Data Retention: Meta appears to keep your Meta AI chat data **indefinitely (until you delete it)**. The assistant "remembers everything by default," as one analysis noted. This means all your conversations and the AI's replies are stored and remain accessible in the app's history. They also persist in Meta's backend systems so that the AI can refer to them in future chats. For example, if today you tell Meta AI "I have two dogs named Fido and Rex," and you haven't cleared the history, a week later the AI might recall your dogs' names if relevant – showing that the data was retained. Meta AI also continuously updates your *Memory file* – a set of key facts or attributes it's learned about you. This Memory is stored as long as your chats exist, and you can even view it in the app's settings (and delete items if you wish). **If you delete chats or use "/ reset-ai," the AI's memory of those chats is erased**, and those specific conversation logs should be removed from Meta's systems (Meta hasn't published how quickly deletion happens, but we can assume it's fairly prompt for user-facing portions, though backups might linger per Meta's general policy). Absent deletion, Meta could retain your AI interactions *forever*, since they are useful for continually improving personalization. There is also an aspect of sharing data between Meta AI and Facebook/Instagram: if you use your Facebook account to log in, Meta AI might share insights from your chats back to Facebook's systems to, say, inform your advertising profile or content recommendations. So your AI chat data doesn't live in isolation – it can feed into the larger Meta data ecosystem. It's wise to treat anything you tell Meta AI as information that will *stick around* and potentially be used by Meta unless you explicitly remove it.

Unique Privacy Points: Meta's AI is **highly personalized**. It's designed to leverage all available data about you to give tailored responses. This is different from, say, ChatGPT, which forgets who you are outside of a single session. Meta AI actively tries to **"get to know you"** – it even says so in their announcements. This can be convenient (the AI feels more like a personal assistant who remembers your past conversations) but comes at the cost of lots of data retention. Another unique aspect: Meta AI chats can be **shared publicly by users**. In the Meta AI app, there's a "Share" button that will post the conversation to a public "Discover" feed. While that's a user action, it underscores that what you say to the AI could be seen by others if you're not careful (and Meta warns that anything shared this way is visible to *everyone* unless you hide it later). On the flip side, Meta has put some privacy safeguards: for example, in group chats, Meta AI **only reads messages directed at it** (it won't scan your whole group conversation, just the ones where it's mentioned). And Meta claims it has trained the AI not to leak one user's info to another. Finally, keep in mind Meta's broader data policy applies – and Meta is no stranger to using data aggressively. If you're uncomfortable with how Facebook handles your data, you'll likely feel the same (or stronger) about Meta AI. **Bottom line:** Meta AI will remember and use your inputs extensively to personalize your experience, and it doesn't really ask permission to do so beyond your agreement to the terms when you start chatting.

Microsoft Copilot (Consumer Version)

Data Collected: Microsoft's Copilot for consumers – which spans things like Bing Chat (the AI in Bing search), Windows Copilot, and Copilot integrated into Windows 11 – collects the **content of your conversations** (your questions and the AI's answers), much like the others. If you use voice input with Copilot, your voice clips are processed too. Microsoft also collects standard diagnostic and usage data: for example, your device information, the Microsoft account you're signed in with, and how you interact with Copilot. If Copilot is helping you with web searches or browsing, it might collect context from your Bing or MSN usage as well. Copilot can also be given access to your local data or calendar/emails (e.g. in Windows), but as of now the consumer Copilot mostly works with web data and your prompts directly. One thing to note: if you upload an image for Copilot to analyze (for instance, in Bing Image Creator or Chat), Microsoft will process that image but says it takes steps to de-identify things like faces or text in images before any training use. Overall, **Copilot collects what you give it (text or images) and some related info**. And if you're using a Microsoft account, this activity is tied to your account's history.

Use for Training: Microsoft uses consumer Copilot chats to train and improve its AI models **by default**, except in certain regions or circumstances. According to Microsoft's privacy FAQ, unless you belong to a group that is excluded, *"Microsoft uses data from Bing, MSN, Copilot, and interactions with ads on Microsoft for AI training"*. This training data includes **your text conversations with Copilot (and voice, if applicable)**. The data is **de-identified** before training, meaning they attempt to remove personal identifiers like names, email addresses, phone numbers, or device IDs from the conversation logs. The idea is to use the substance of user questions and Copilot answers to make the underlying models (like the ones powering Bing Chat) more accurate and versatile. Microsoft explicitly mentions that **certain users are excluded from this training data usage** – notably, users under 18, users not signed in, and users in a handful of countries (Brazil, China, Israel, Nigeria, South Korea, and Vietnam) are *not* having their chats used for model training. (Microsoft likely opted to exclude some regions for legal compliance reasons.) For everyone else, if you don't opt out, your chat with Bing or Windows Copilot could later influence how the AI behaves. For example, real user questions help the AI learn about new slang or trending topics, and common mistakes help Microsoft fix the AI's weaknesses [7]. Microsoft also potentially uses Copilot interactions to improve other services – e.g. they mention using Bing and MSN data, which implies your Copilot chats might also refine Bing's search algorithms or related tools.

Opt-Out Options: Yes – Microsoft allows consumers to **opt out of having their Copilot chats used in training**. If you are logged in with a Microsoft account while using Copilot, you can find a setting (likely in your privacy dashboard or the Copilot settings) to **"disable model training"** on your conversations. Turning this on will exclude *past, present, and future* conversations from training use, and Microsoft says the change takes effect across their systems within about **30 days**. In other words, they'll stop using your chats going forward, and even previous chats associated with your account would be withdrawn from the training pipeline. If you don't log in at all (using Copilot as a guest), your chats aren't used for training either – Microsoft only trains on data linked to user accounts unless you opt out. So, a privacy-conscious approach could be: use Bing Chat without logging in, or always opt out in settings if logged in. It's worth noting that opting out of training **does not break the service** – you can still use Copilot and even have it personalize to you (personalization is separate; see below). Also, Microsoft's opt-out is distinct per account: if you have multiple accounts or devices, you'd want to ensure the setting is applied everywhere you use Copilot. Finally, Microsoft provides a standard privacy dashboard where you can view and delete your Bing/Copilot chat history if you want to remove records manually.

Data Retention: Microsoft hasn't published a simple number like "30 days" or "indefinite" for consumer Copilot data retention. However, from what we know: **If you have Copilot (or Bing) chat history turned on** (it's on by default for logged-in users, showing past chats in a sidebar), your conversations are stored in the cloud tied to your account until you delete them. You can clear them from the history UI, which should delete the stored record from Microsoft's servers. If you opt out of training, Microsoft will still retain the conversation data for you (and possibly for compliance/safety internally), but just not include it in model training going forward. In the case of **Windows Copilot**, chat data might also be stored locally for context (though it ultimately goes to the cloud AI as well). Microsoft has said that **enterprise or organizational data** handled by Copilot (in business contexts) can have a strict retention of zero or 30 days, but for consumers, such strict limits are not stated. It's safe to assume your Copilot chats persist until you take action to delete them, or possibly until they age out if Microsoft sets some internal limit (e.g. Bing might auto-delete after 90 days for non-active users – this is speculative as they haven't confirmed). On the plus side, **personalization "memories" can be reset** easily (if you turn off personalization, Copilot will forget any learned preferences, even though the raw history isn't deleted). Also, **no conversation content is used for training in certain jurisdictions** (like the countries mentioned), so in those places Microsoft likely deletes or segregates the data quickly to comply with local laws. As a user, you can manually wipe your chat logs via the Microsoft Privacy Dashboard online, which allows deletion of Bing chat history. There is also a setting to *not* save chat history at all in Bing (recently introduced), which would cause each session to be ephemeral. In summary, retention is not fully transparent but **you have tools to manage or erase your data**. If concerned, it's a good practice to regularly clear your Copilot chat history, especially if it contains sensitive info.

Unique Privacy Points: Microsoft's approach tries to straddle personalization and privacy. An interesting point: Microsoft separates **"training" and "personalization."** Even if you opt out of training, you can still allow Copilot to personalize answers based on your past chats just for you [8]. In that case, Copilot will remember things within your session or account (like previous context) to tailor responses, but Microsoft promises not to use those chats to train the global AI model for everyone [9]. Also, Microsoft leverages your data for improving not just the AI's accuracy but also possibly ads: if you've agreed to personalized ads in your Microsoft account, Copilot conversation history might be used to refine what ads you see on Microsoft services. (For example, asking Copilot about new laptops could influence the ads shown on Bing or MSN.) This is somewhat unique – it treats Copilot chat content as another signal for ad personalization, akin to search queries. Microsoft does provide transparency here: they clarify that Copilot itself doesn't

show ads in the chat, it just might feed into ads you see elsewhere if you've opted for that. Another point: **certain data is never used for training**, such as the contents of files you might have Copilot summarize (unless you explicitly talk about the file in chat). Microsoft, like Google, also abides by publishers' instructions (e.g. not training on content from sites that disallow it). And of course, Microsoft's enterprise versions of Copilot (Microsoft 365 Copilot, etc.) have much stricter privacy guarantees – those are separate from the consumer scenario we discuss here. For a regular user, the key is that **Microsoft gives you a fair bit of control** – more than Meta, similar to OpenAI – so you can decide if you want to contribute your chat to the AI's education or not.

xAI's Grok (from Elon Musk's xAI)

Data Collected: Grok is xAI's chat assistant (not to be confused with the X/Twitter platform itself, though they're related). Grok will collect **anything you type or say to it**. In fact, xAI explicitly says *"we ask that you do NOT include personal information in your prompts,"* because they know whatever you input could be collected. Grok collects your **account info** if you sign up for a Grok account (name, email, birthdate for age check, etc.). If you log in via X (Twitter), xAI will receive your Twitter profile information (username, profile pic, and even your Twitter subscription status) and notably **your Grok conversation history from within Twitter**. This means if you used the Grok bot on X, those chats are linked and xAI can pull them. When using the Grok app or website, it records your conversations and any files you upload for analysis. There's a feature called **Private Chat** which is a mode where Grok doesn't save that conversation to your history (more on that below), but otherwise, by default, **Grok will save your prompts and its responses** as with most AI chats. They'll also get technical data (like IP address, device type, and usage metrics) in the background for security and analytics. In summary, **Grok gathers your prompts, uploaded content, account details, and usage info**, and if you came through Twitter, it ties into your Twitter account data too. They double down on warning users not to give personal/sensitive data, indicating that such data could be inadvertently stored or processed.

Use for Training: xAI's policy is that **they may use your content and interactions with Grok to train and improve their models, but they give users a choice in the matter**. The default setting allows training: unless you change your data controls, **your prompts and Grok's replies can be used to refine xAI's large language model**. This helps Grok get better at understanding questions and providing accurate answers. However, xAI is quite transparent about this and provides a built-in opt-out (and even a way to avoid saving chats entirely via Private Chat). If you **activate Private Chat mode**, Grok will not retain those conversation logs at all in history *and* will not use them for training. If you leave the default mode on but toggle "Improve the model" off in settings, Grok will still record the chats for you but internally flag them to *not* be included in training datasets. xAI also assures that any content from business customers is never used for training without permission. So effectively: **the user can decide**. By default, assume it *can* be used (with identifying info removed) to train the AI, but you have straightforward options to prevent that. It's worth noting xAI is a newcomer, so they are likely hungry for data to improve their model, which is why opt-in for training is the default. Elon Musk's team has also emphasized using public data (and X data) to train their models, so user chats are an important resource unless one opts out.

Opt-Out Options: Yes – xAI gives clear opt-out controls. In the Grok app or on grok.com, you can go to **Settings → Data Controls** and toggle **"Improve the model"** on or off [10] [11]. If you turn it off, your new conversations **will not be used to train** their models. You can also achieve a similar effect by using **Private Chat mode** for a conversation – when Private Chat is enabled, that session's messages won't even appear in your history and are deleted from xAI's systems after 30 days, and none of it goes into model training. Think

of Private Chat like an “incognito mode” for Grok. Additionally, if you use Grok on X (Twitter), you’d have to follow Twitter’s help instructions to opt out there, but currently Grok’s main usage is via its own app. So in practice, you have a **permanent opt-out toggle** and also a **per-chat privacy mode** – both resulting in your data not training the AI. If you’re comfortable contributing to improve Grok, you can leave it on. If not, it’s easy to switch off. Grok’s interface also allows you to **delete any conversation** from your history; doing so will remove that data from xAI’s records within about 30 days (unless required for legal reasons). As a side note, if you’re not logged in (just like others), your chats might be collected anonymously but you won’t have an account setting to control training – xAI notes that in regions outside the EU, if you use Grok without logging in, they may still collect and use that content but it’s not tied to you. Logged-out users can’t opt out of training on a per-chat basis, so xAI recommends logging in and using settings for more control.

Data Retention: xAI’s retention policy is nuanced. Generally, **they keep personal data as long as they have a business need to**. For your conversations, the retention can depend on your choices. If you’re using standard mode (not private) and not opting out, your chats will be stored in your history until you delete them, and xAI could retain them on their servers indefinitely for training/improvement purposes. If you decide to delete a conversation from history (or delete your account), xAI will **erase that data within 30 days** (barring any legal obligation to keep it). Now, if you use **Private Chat**, those chats are never stored in your persistent history and xAI promises to delete them from all systems within **30 days** automatically. So Private Chats are short-lived on the backend. Feedback data (if you provide ratings or bug reports) might be kept up to 10 years, similar to Anthropic, because they de-link it from you and use it for long-term model training. It’s important to note that, unlike OpenAI which deletes no-history chats after 30 days, xAI’s *default* (non-private) chats don’t have a set expiration – they might live on until you remove them. The **Privacy Policy** says xAI retains data for various legal and security reasons too (like to prevent abuse or comply with law) and lists scenarios where data could be kept longer. So if, hypothetically, a chat raised a red flag (security or legal), xAI might hold onto it even if you tried to delete, to comply with law enforcement or their policies. But routine, non-problematic chats follow the user’s lead on deletion. The upshot: **regular Grok chats remain accessible and retained like chat history indefinitely; private chats last no more than 30 days server-side**. If you want your data gone sooner, you have to manually delete it (and then it’s gone in ~30 days).

Unique Privacy Points: xAI (Elon Musk’s AI startup) positions itself as doing things “for the people,” but their privacy posture is fairly standard with a nod to user control. A unique feature is the **Private Chat mode**, which not all AI apps have. It’s a one-click way to enter a conversation that won’t be used to train the model and even disappears from your own chat history – useful if you want to ask Grok something sensitive without leaving a trace beyond 30 days. Another point: xAI is closely tied to the X (Twitter) ecosystem, so there’s an interplay of data there. Musk’s companies have openly stated they’re using public tweets to train AI; with Grok, it’s interesting that using it via Twitter means Twitter’s privacy policy applies. This could mean if you chat with Grok on Twitter, that chat could also be used by Twitter for other purposes (though currently Grok is mainly an external app). Also, xAI is very new (launched late 2023), so they are likely still evolving their policies. But they’ve at least built in **user agency on training** from the start, which is a positive. On the flip side, xAI is not shy about collecting data – they even pull in your Twitter conversation history with Grok if you link accounts, meaning your data flows between platforms. And like others, xAI uses your chats to improve safety: they mention if you violate their rules, they might re-identify your data to take action. In essence, **Grok gives you incognito and opt-out options that some competitors don’t, but if you do nothing, assume your chats help train their AI and stick around**.

DeepSeek (Chinese AI Chatbot)

Data Collected: DeepSeek is a chatbot from a Chinese company that gained popularity rapidly in early 2025. DeepSeek's app collects **a very broad array of user data** – essentially **anything you enter into the chat and a lot about your device and usage**. According to its privacy policy, DeepSeek logs your **chat prompts and the AI's responses** (your full conversation history) and can also collect personal or sensitive information that you may include in those chats [12] [13]. It does not shy away from collecting sensitive data; in fact, its policy isn't designed to avoid it, so if you mention, say, your bank details or health info to the bot, that could be stored. Besides chat content, DeepSeek also gathers your **account details** (if you register, likely phone number or email since many Chinese apps require that) and any profile info. They automatically collect device identifiers, your IP address, location info, and other telemetry. **Importantly, all data DeepSeek collects is stored on servers in the People's Republic of China** [14] [15]. This includes your chat history and possibly metadata about your device. So, data from users around the world (including the U.S. or Europe) is being transmitted to China when they use DeepSeek [15]. DeepSeek's privacy policy outlines that they might also get information from other sources – for example, if they partner with advertisers, they could collect data through them. In summary, **DeepSeek collects as much as it can: your chats (text, audio, images), account info, device info, and usage patterns, and stores it all in China** [15] [16].

Use for Training: DeepSeek's use of data is quite invasive compared to Western counterparts. The privacy disclosures indicate that **user inputs "may be stored indefinitely" and** "used for AI training" [12] [13]**. In practice, that means** anything you tell DeepSeek can and likely will be used to further train or fine-tune their AI models**. There is no indication of an opt-out for this; it appears to be a condition of using the service. Furthermore, DeepSeek openly states it may** share the information within its corporate group and with advertising partners [12] [13]**. That implies, for example, your chat data could be shared with sister companies or used to target ads at you (perhaps within the app or elsewhere). DeepSeek, being a Chinese company, is also subject to Chinese law which can mandate sharing data with government authorities. This has led to major privacy concerns internationally – for instance, Italy's data protection authority banned DeepSeek temporarily over worries that it was using personal data unlawfully** [17] [18]**. In short,** DeepSeek treats your conversations as data it owns: it will mine them for AI improvements and possibly for advertising or other purposes, with no apparent user control**. Users have reported that DeepSeek even** censors certain content critical of China** [19], meaning there's active monitoring of what you input, which goes hand-in-hand with them analyzing your data.

Opt-Out Options: Unfortunately, **DeepSeek does not offer a user-friendly opt-out for data usage or training**. If you use the app, you are opted in to their data practices. There is no toggle in settings to stop data collection – the service is essentially "all in." The app does have a **"Delete all chats"** function in its settings, which allows you to wipe your conversation history from the app's interface [20]. Using that might remove the chats from your personal account view. However, it's unclear if DeepSeek truly deletes those chats from its servers when you do so, and the privacy policy wording "may be stored indefinitely" suggests they reserve the right to keep them [12] [13]. In jurisdictions like the EU, users have rights to request deletion of their data, but exercising that with a Chinese company might be difficult in practice. Essentially, **the only sure way to opt out of DeepSeek's data collection is not to use DeepSeek at all**. If you do use it, assume your data is being logged and retained. As of now, there's no known mechanism to opt out of model training or third-party sharing on DeepSeek. This lack of control is a big red flag that privacy experts have noted [21] [13].

Data Retention: DeepSeek's policy literally says it **may store user data indefinitely** [12] [13]. This means there is no set expiration or deletion schedule for your conversations. They could keep your chat logs on file forever. Given that all data is stored in China, it also falls under Chinese regulations which might require data to be kept for certain periods (and accessible to government agencies). If you delete chats via the app, that may only delete the local/app record. We don't have evidence that DeepSeek honors those deletions on the back-end; they might, but the wording doesn't promise it. Moreover, because they share data with advertising partners, your chat content or derived insights might be transferred to other databases outside of DeepSeek's immediate control, which could further prolong retention. **Bottom line: assume anything you enter into DeepSeek will be saved indefinitely unless a regulator forces them to delete it.** This has already started to happen – some countries' regulators are scrutinizing DeepSeek. For example, Italy's ban means DeepSeek might have to delete Italian users' data or stop processing it to comply with GDPR [17]. But unless you're in a jurisdiction pushing for that, DeepSeek is keeping your data. Also, storing data in China introduces concerns that data could be accessible or retained due to state policies beyond the company's own choice. All told, **DeepSeek likely retains user data permanently by default**.

Unique Privacy Points: DeepSeek stands out for the **extent of data risk**. It became hugely popular quickly (even hitting #1 on app stores [22]), meaning many users might not realize the trade-off: it's free and powerful, but **explicitly sends user data to China** [14], which has different privacy protections than the US or EU. Cybersecurity experts have pointed out that DeepSeek appears to send **more data back to China than even TikTok does**, because TikTok has tried to localize some data whereas DeepSeek has not [23] [24]. This has raised national security alarms. Another unique aspect is **censorship** – users noticed that DeepSeek will refuse or filter queries about certain political topics (especially those sensitive to the Chinese government) [19]. That indicates active content scanning of your inputs, which is a privacy concern in itself. Also, DeepSeek's privacy policy is very broad (some analysts call it "broadly written and covers all possible data collection" including even keystroke data [25]). It essentially reserves rights to do almost anything with your data – share it, monetize it, etc. [12]. This is in stark contrast to, say, Anthropic Claude's strict no-training stance or NotebookLM's local-only approach. In fact, DeepSeek's practices have sparked investigations by international regulators [26]. If you are a privacy-conscious user, DeepSeek is probably the **least private** of the tools on this list. Its popularity despite this highlights how users often jump for capability without realizing the data cost. In sum, **DeepSeek collects maximal data, keeps it indefinitely, uses it freely (even for ads), and stores it under one of the world's strictest data regimes (China)** [13] [15]. Use with extreme caution, and definitely avoid inputting anything sensitive or personal.

Google NotebookLM

Data Collected: NotebookLM is a bit different from the chatbots above. It's a Google Labs product that acts as an AI research and note-taking assistant. NotebookLM works by letting you upload your own documents (like Google Docs or PDFs) into a "notebook" and then ask questions about them. The data it collects primarily includes the **content of the documents you provide**, the **questions you ask**, and the AI's **generated responses**. If you log in with a Google Account (which you have to, since it's tied to Google Drive), it knows your Google identity but it doesn't use your Google Account data for training (more on that below). **Importantly, Google has stated that NotebookLM will never use your personal data or uploaded documents to train its models** [27] [28]. So while the service obviously processes your data (it has to read your docs to answer questions), that data stays isolated to your session and isn't fed into the AI's learning. NotebookLM may also collect some usage info (like how often features are used, for improving the product's UI), but that's separate from model training. If you provide explicit feedback through the NotebookLM interface (like rating an answer or reporting a bug), those feedback submissions might be

reviewed by humans to troubleshoot or improve the system [29]. For example, if you say "This answer was wrong because X," a Google engineer might look at that interaction to figure out what went wrong. But aside from such volunteered feedback, **NotebookLM treats your data as private** – it's basically confined to your own use.

Use for Training: NotebookLM does *not* use your conversations or uploaded content for AI model training, period. This is a clear differentiator. Google explicitly says: *"NotebookLM does not use your personal data, including your source uploads, queries, and the responses from the model for training."* [27]. Think of it this way: NotebookLM is powered by a large language model (likely a version of PaLM or Gemini) that was trained on broad data, but once you're using NotebookLM, **it's not continuously training on what you give it**. It's doing on-the-fly analysis of your documents (also known as "in-context learning"), not updating the model's weights. So your data remains your data. This means if you and another user ask something similar in NotebookLM, the model's quality is based on its pre-existing training, not on what you individually have done with it. Google also clarifies in their FAQ: *"NotebookLM will never train on any of your data."* [30] [31]. They apply this rule regardless of account type – even if you have a free personal Google account, they don't turn around and train the AI on your info [32]. This is reassuring for those worried about proprietary or personal notes leaking into an AI. Additionally, if you're a Google Workspace (enterprise/education) user, NotebookLM data is covered under those strict terms (no human review, no external use) [33]. With a personal account, Google does mention that if you *choose* to send feedback, humans might review that specific content to fix issues [29]. But that's a voluntary action. In normal usage, **no one is reading your NotebookLM queries or notes, and they're not going into some Google neural net to train it** [29].

Opt-Out Options: Because NotebookLM doesn't use your data for training in the first place, there's no "opt-out" needed – it's the default behavior. By using NotebookLM, you're not opting into any data-sharing for model improvement. Google effectively opted everyone out from the get-go. If you're particularly cautious, you can simply refrain from giving feedback (since that's the only time someone might manually see your content). And of course, you can stop using the service or disconnect any uploaded documents at any time. NotebookLM is designed to be **privacy-preserving by default**, likely because it's aimed at sensitive use cases like summarizing personal notes or research – areas where people would not try the product if it was slurping up their data. So, there is no toggle like Bard's activity control because the model training usage is off from the start. Google's documentation encourages trust by saying your interactions are just for you and not for AI training [27]. In the UI, they also put a little "lock" icon or message about privacy, reinforcing that what you put in is not feeding Google's AI brain elsewhere [27]. If for any reason you did want to remove all traces of your usage, you could delete the documents from the NotebookLM interface and they'd no longer be accessible to the system (the original docs remain in your Google Drive unless you delete those too). Since NotebookLM is tied to Google Drive, standard Google account deletion or data export controls apply as well. But the main point: **no special action is needed to protect your data from model training – NotebookLM already does that**.

Data Retention: When you upload a document to NotebookLM, it's actually referencing a copy of that document (if it's from Google Drive) or storing it for the session if you uploaded from your computer. Google will keep the association as long as you have that notebook active. You can remove documents from the notebook, and then presumably any cached data is deleted. As for your **queries and the AI answers**, those are likely stored temporarily so you can see your chat history within a notebook. Google hasn't published how long those Q&A pairs persist, but since NotebookLM is a personal productivity tool, one can expect that your session data remains available to you until you delete the notebook or specific queries. Google might store some metadata about how the system was used (for example, to count how many

queries are made, or to detect abuse), but not for model training. They explicitly state that if you're a Workspace user, nothing is reviewed by humans and presumably those queries might not even be stored beyond your view [34]. If you have a personal account and provide feedback with a specific chat, that chat content might be retained within Google's issue tracking system to fix the problem (capped at some timeframe, maybe a few years or until it's resolved). However, if you never hit "send feedback", your NotebookLM chats are mostly ephemeral outside of your own notebook. There is no mention of auto-deletion (like auto-delete in 30 days, etc.), so likely **your NotebookLM chat history remains accessible to you until you clear it**, similar to how a note-taking app would keep your notes. Google presumably stores the content of your notebook (documents + Q&A context) on their servers securely as part of your account data. But crucially, it's not siphoned elsewhere. If you stop using NotebookLM, the data stays in your Google account until you delete it (for example, if you uploaded a PDF to NotebookLM, that PDF might be stored in a hidden app folder in Drive – this is an educated guess as many Google services do that). When you remove that PDF or the whole notebook, Google deletes that stored content. Also, Google has to adhere to data protection laws, so if you delete your Google account or use Google's Takeout tool, you'd get or remove any NotebookLM data too. Summing up: **NotebookLM retains your data only for your usage**. It's not clear-cut how long query history is kept, but it lives at least as long as you keep the notebook, and you control deletion. None of it is used to train AI and none is accessible to other users, which minimizes the worry about retention.

Unique Privacy Points: NotebookLM's standout feature is **keeping your data private to you and not learning from it** [27]. It's essentially localizing AI for individual use. This addresses a big concern people have with using AI on personal or proprietary documents – with NotebookLM, Google is saying "we won't peek." It's a demonstration of *on-device* or *user-specific* AI principles (even if the processing is cloud-based, the data doesn't feed back). Another unique aspect: NotebookLM is positioned as a **trusted research assistant**, so Google is trying to build user confidence that it's *confidential*. They even design the UI/UX to remind you of that promise. It shows how AI can be deployed in a way where the value is from the model's pre-training (on public data) but your personal data isn't absorbed. In a sense, NotebookLM is similar to running an AI on your own computer with your files – except Google's servers are doing it, but with guarantees around your data. It's also worth noting that because NotebookLM doesn't use your data for broader purposes, some features are limited: for example, it doesn't use your interactions to improve itself, so if it makes a mistake summarizing something, it relies on general model updates or your direct feedback to developers rather than automatically learning from that mistake. This is a trade-off for privacy. Additionally, NotebookLM falls under Google's general privacy and security principles, meaning it benefits from Google's infrastructure security, access controls, and compliance (and if you're a Workspace user, it's covered by enterprise-grade commitments). In short, **NotebookLM is the most "privacy-preserving" tool on this list by design**, as it *never* trains on your data and keeps your info siloed [27] [29]. It's a good option if you want to use AI on personal documents without that data leaving your own sphere. The only thing to keep in mind is that while Google won't use it to train AI, your use of the product is still subject to Google's terms – e.g. they could suspend accounts for abuse, etc., and standard security measures apply (don't put illegal content in it, etc.). But for everyday privacy, NotebookLM is **refreshingly straightforward: your notes stay your notes**.

To wrap it all up, here's a **comparison table** highlighting the key privacy features of each AI tool:

AI Tool	Uses Your Conversations to Train the AI?	User Opt-Out Available?	Data Retention Policy	Notable Privacy Measures/ Differences
OpenAI ChatGPT (consumer)	Yes, by default ChatGPT saves prompts & may use them to improve models.	**Yes.** You can turn off "model training" in settings, or use no-history mode (chats then not used for training).	Chats with history on are kept indefinitely (until you delete). If history off, chats kept ~30 days then deleted.	Doesn't use data for ads or sell it. Offers data export and easy deletion. You can fully opt out of training and still use the service.
Anthropic Claude (free/ Pro)	**No, by default** Claude does *not* train on your chats (except flagged or feedback cases).	N/A (It's opt-out by default). Feedback is opt-in. So, nothing for users to disable – it's private by design.	Retains conversation history as long as needed for you. Deleted chats are removed from systems within 30 days. Flags/ feedback stored longer (2-10 years in de-identified form).	Strong privacy stance. Only **explicit feedback or abuse cases** are used to improve Claude. Limited staff access. 30-day deletion window for any user-removed content.
Google Gemini (Bard)	Yes, by default Google uses a *sample* of chats (with personal info removed) for human review & model training.	**Yes.** You can turn off **Gemini Apps Activity** (chat history) to stop data being used for improvement. Also can delete chats anytime.	With history on, chats stored in your Google account until you delete or auto-delete. With history off, chats saved up to 72 hours then removed. Reviewed samples kept up to 3 years.	Reviewed chats are **anonymized**. No current use of chats for ads. Integration with Google's privacy controls (My Activity, export, auto-delete).

AI Tool	Uses Your Conversations to Train the AI?	User Opt-Out Available?	Data Retention Policy	Notable Privacy Measures/ Differences
Meta AI (Facebook/ Instagram)	Yes, by default **everything** you share is used to personalize the AI and can train Meta's models. No opt-out given.	**No direct opt-out.** You can delete chats or use "/reset-ai" to erase history, but Meta will still collect new chats going forward.	Retained indefinitely by default. Meta AI keeps transcripts & a memory profile until you delete them. Deleted chats are removed from your app history (policy doesn't state server deletion timeframe).	Highly personalized – builds a **memory** of user info. **No training opt-out**, data tied into Meta's ecosystem. Could be used for future ad targeting. Only way to limit exposure is deleting content often or using it logged-out.
Microsoft Copilot (consumer/ Bing Chat)	Yes, by default most user chats are used to train Microsoft's AI models (except certain regions/ users).	**Yes.** Logged-in users can **opt out of model training** in settings. (Opt-out excludes past & future chats within ~30 days).	Likely stored as account history until user deletes. No set expiration disclosed. Users can clear chat history manually. Opt-out stops training use but data still stored for user and compliance.	Removes personal identifiers before training on chats. **No training** on under-18 or certain countries' data. Chat data may enhance personalized ads if you allow. Enterprise versions have stricter zero-retention, but not applicable to consumer.

AI Tool	Uses Your Conversations to Train the AI?	User Opt-Out Available?	Data Retention Policy	Notable Privacy Measures/ Differences
xAI Grok	Yes, by default xAI may use your prompts & Grok's responses to improve the model. (They urge not to input sensitive info.)	**Yes.** In Data Controls you can toggle off "Improve the model" to opt out [10]. Or use **Private Chat** mode for no logging/training.	Regular chats: retained until you delete (no fixed limit given). **Private Chat** sessions: deleted from xAI systems within 30 days. Deleted conversations or accounts are wiped in ~30 days.	Built-in **Private Chat (incognito)** that doesn't store or train on that chat. Users have full control to opt out of training. Data stored in USA; however, if used via Twitter, Twitter's policy applies too. xAI is very transparent about data use and encourages not sharing personal data.
DeepSeek	Yes, **by default everything is used** – user inputs can be stored indefinitely and used to train AI [12] [13].	**No.** There is no opt-out offered to users. Using the service implies consent to data usage. (Users can delete chat history in-app, but it's unclear if that fully removes data) [20] [13].	**Indefinite retention** of user data (as per policy) [12]. Data is stored on servers in China [15]. Deletion by user may not purge backend data; no transparency on deletion.	All data goes to **China** [15], raising privacy and legal concerns abroad. Shares data with corporate partners and for ads [13]. No privacy controls for users. Has been banned in at least one country (Italy) for privacy violations [17]. Considered one of the least privacy-protective AI apps.

AI Tool	Uses Your Conversations to Train the AI?	User Opt-Out Available?	Data Retention Policy	Notable Privacy Measures/ Differences
Google NotebookLM	**No, never**. NotebookLM does not use your uploads, questions, or notes to train models [27] [28].	Not needed – it's private by default. (If you submit explicit feedback, that's voluntary and used to improve the service) [29].	Retains your uploaded documents and Q&A history for your use. Data stays in your Google account; not used elsewhere. You can delete your notebooks or documents at any time (no secondary copies for training exist) [27].	**Privacy-first design:** your data is for *your eyes only* in terms of AI learning [27]. No human review of content unless you send feedback [29]. Ideal for sensitive or proprietary documents – it functions without ingesting your info into global models.

Each of these AI tools handles your data a bit differently. In summary, services like **NotebookLM and Claude put privacy foremost**, not using your content to train by default. **ChatGPT, Bing Copilot, and Google's Bard** do use your data to make their AIs better, but they give you ways to opt out and control that usage. On the other end, **Meta AI and DeepSeek vacuum up your data by default without robust opt-outs** – Meta uses it to personalize your experience (and likely future ad targeting), and DeepSeek even shares it with advertisers and stores it indefinitely in a jurisdiction with different privacy laws. **xAI's Grok** sits somewhat in the middle: it will use your data unless you tell it not to, but at least it provides clear switches to preserve your privacy.

For any AI chatbot, if you're a non-technical user, the safest practice is: **assume anything you type could be stored**. Check the settings or FAQs for each service to see if you can limit that. Avoid sharing extremely sensitive personal information unless you're using a tool that explicitly guarantees privacy (like NotebookLM's guarantee of not training on your data [27]). And make use of the controls available – turn off chat history if available, delete your conversation logs, or opt out of model training if you prefer your data not be used. By knowing these differences, you can choose the AI assistant that best fits your comfort level on privacy.

Sources:

- OpenAI Help Center – *Data Controls FAQ* (how ChatGPT uses conversations, opt-out)
- OpenAI Privacy Policy (data collected and usage)
- LiveChatAI Blog – *Does ChatGPT Save Data?* (summary of ChatGPT data usage)
- Anthropic Help Center – *Claude Pro FAQs* (Claude does not train on user conversations by default)
- Anthropic Privacy Center – *Data Handling & Retention* (30-day deletion for Claude user-deleted chats, retention timelines)

- Gemini (Google) Privacy Hub (what data Bard/Gemini collects, human review process, opt-out instructions)
- Google Bard Privacy Notice (not using conversations for ads, stored in account, 72-hour retention with history off)
- Washington Post – *Meta AI privacy* (Meta AI keeps transcripts, memory, uses chats for training, no opt-out)
- Meta's "Privacy Matters" blog (Meta AI details: not using private messages for training, how to delete AI chats)
- Microsoft Support – *Privacy FAQ for Copilot* (training on user data by default, opt-out available, data excluded, etc.)
- xAI Privacy Policy and FAQs (Grok data collection, training opt-out toggle, Private Chat mode, retention)
- Wired – *Grok AI and Your Privacy* (warning users not to share personal data, context on Musk's xAI practices)
- DPO Centre News – *DeepSeek under scrutiny* (DeepSeek storing inputs indefinitely, used for training, shared with advertisers, Italy ban) [12] [13]
- Wired – *DeepSeek sending data to China* (confirmation that all chats go to China servers, how to delete chats, comparisons to TikTok) [15] [20]
- Google NotebookLM Help – *How NotebookLM protects your data* (never uses personal data for training, differences for Workspace accounts) [29] [27].

---

[1] [2] [3] [4] [5] [6] Gemini Apps Privacy Hub - Gemini Apps Help
https://support.google.com/gemini/answer/13594961?hl=en

[7] [8] [9] Privacy FAQ for Microsoft Copilot - Microsoft Support
https://support.microsoft.com/en-us/topic/privacy-faq-for-microsoft-copilot-27b3a435-8dc9-4b55-9a4b-58eeb9647a7f

[10] [11] Consumer FAQs | xAI
https://x.ai/legal/faq

[12] [13] [16] [17] [18] [21] [22] DeepSeek under scrutiny: Privacy concerns over Chinese AI chatbot
https://www.dpocentre.com/news/deepseek-under-scrutiny-privacy-concerns/

[14] [15] [19] [20] [23] [24] DeepSeek's Popular AI App Is Explicitly Sending US Data to China | WIRED
https://www.wired.com/story/deepseek-ai-china-privacy-data/

[25] DeepSeek App: A Closer Look at Its Privacy Posture - Privado.ai
https://www.privado.ai/post/deepseek-app-a-closer-look-at-its-privacy-posture

[26] International regulators probe how DeepSeek is using data. Is ... - NPR
https://www.npr.org/2025/01/31/nx-s1-5277440/deepseek-data-safety

[27] [31] Google NotebookLM | Note Taking & Research Assistant Powered ...
https://notebooklm.google/

[28] [29] [33] [34] Learn how NotebookLM protects your data - NotebookLM Help
https://support.google.com/notebooklm/answer/15724963?hl=en

[30] Frequently Asked Questions - NotebookLM Help - Google Help
https://support.google.com/notebooklm/answer/14278184?hl=en

32 The NotebookLM Privacy Update clarifies that... | John E. Bredehoft
https://www.linkedin.com/posts/jbredehoft_notebooklm-and-privacy-the-notebooklm-privacy-activity-7270878926581751808-uVT-