

Healthcare AI Privacy: A Quick Reference

For clinicians, dentists, nurses, practice managers, front-office staff, and medical billers who use ChatGPT, Claude, Gemini, Perplexity, or Microsoft Copilot for work.

Not legal advice. When in doubt, ask your compliance advisor.

The one-sentence version

Personal AI accounts (ChatGPT, Claude, Gemini, Copilot, Perplexity) are not covered by any Business Associate Agreement. Every patient identifier you paste into one is a HIPAA disclosure to a third party who has not agreed to protect it.

What counts as a patient identifier

Any single one of these, or any two combined, can identify a patient:

Direct identifiers	Indirect identifiers
Full name	Age plus rare condition
Date of birth	Small-town clinic plus visit date
Address	Employer plus job title
Phone number	Any highly specific combination
Email address	
Medical record number (MRN)	
Insurance member ID	
Insurance subscriber ID	
Claim number	
Prescription number	
SSN	
Driver's license	

Direct identifiers	Indirect identifiers
Financial account number	
Face photo	
Biometric identifier	
Any other unique code	

What is safe

Safe	Not safe
Diagnoses (ICD-10 codes)	Diagnoses attached to a named patient
Medications and doses	Medications for a named patient
CPT and procedure codes	CPT codes attached to a claim number
Lab values, vitals	Lab values attached to a named patient
Anatomy, symptoms	Symptoms attached to any identifier
General clinical explanations	Explanations that reference specifics of a real case
Sample or hypothetical patients	Real patients with any details changed to protect them (still identifying if the pattern is unique)

Rule of thumb: if the AI could not identify who you are asking about from what you typed, plus what a bad actor could find publicly, you are safe.

The three-line send check

Before you hit Enter, or before you attach a file, in any AI tool, ask:

1. Is this about a specific real person?
2. Is any name, DOB, address, phone, email, MRN, insurance ID, claim number, prescription number, SSN, or financial identifier in what I am about to send or upload?
3. Would I be comfortable if this exact prompt or file was subpoenaed?

Two yeses in a row, do not send.

What each major AI tool actually does with your prompts

Approximate as of 2026. Verify on the provider's current policy before relying on any of this.

Provider	Consumer chat product	Retains prompts by default?	Uses prompts to train?	BAA available on consumer plan?
OpenAI	ChatGPT (free, Plus, Pro)	Yes	Yes, unless opted out	No
OpenAI	ChatGPT Enterprise / Business	Configurable	No	Yes
Anthropic	Claude.ai (free, Pro)	Yes	Generally no, but opt-out available	No
Anthropic	Claude Enterprise	Configurable	No	Yes
Microsoft	Copilot (consumer)	Yes	Depends on plan	No
Microsoft	Copilot for Microsoft 365	Enterprise-controlled	No (per Microsoft)	Yes, in some plans
Yes	Gemini (consumer)	Yes	Yes, human reviewers on some conversations	No
Google	Gemini (Workspace, business)	Enterprise-controlled	No	Yes, in some Workspace plans
Perplexity	Perplexity (free, Pro)	Yes	Yes, opt-out available	No

Your practice's Workspace or Microsoft 365 subscription may cover the enterprise version of these tools with a BAA. Your IT contact or compliance advisor can confirm.

Rule of thumb: if you signed up with a personal Google, Microsoft, or Apple account, you are using the consumer version. No BAA.

Two safe patterns to build into your day

Pattern 1: Generalize the patient out of the prompt

Not:

Explain treatment options for John Miller, DOB 5/12/1948, MRN 003847, diagnosed with stage 2 pancreatic cancer.

Instead:

Explain treatment options for a 78-year-old male diagnosed with stage 2 pancreatic cancer.

The AI can still answer clinically. The patient is not attached.

Pattern 2: Use placeholders, fill in later

Not:

Draft an appeal letter for Sarah Chen, member ID BC-8493201, claim number 20260847. UnitedHealth denied CPT 99215.

Instead:

Draft an appeal letter for [PATIENT_NAME], member ID [INSURANCE_ID], claim number [CLAIM_NUMBER]. UnitedHealth denied CPT 99215 as not medically necessary despite comorbidities of type 2 diabetes, hypertension, and CKD stage 3.

The AI writes the appeal. You copy it, paste the real values in from your billing system, send it out.

Add a safety net

AI Leak Guard is a free Chrome extension that catches patient identifiers in your browser on two paths:

- **Paste protection** on ChatGPT, Claude, Gemini, Perplexity, and Microsoft Copilot (copilot.microsoft.com). It shows you a preview of what would be hidden and lets you send the safe version, send the original, or cancel.
- **Document scanning** for the PDF, Word, Excel, PowerPoint or text files you attach to ChatGPT, Claude, Gemini, Perplexity, or Microsoft Copilot. Files are read on your device before upload. Clean files upload normally. Files with identifiers get a warning listing which identifiers were found; you choose to upload anyway or cancel. The extension does not edit or redact your file, it just tells you what is in there before it leaves your device. Files it cannot read as text (like scanned images) are flagged honestly, not silently missed.
- **Send-time protection (v1.3)** on ChatGPT, Claude, and Gemini. Reads the message you are about to send at the moment you send it, so typed or dictated identifiers are caught even if you never pasted them. Your draft is always kept. Fail-open: if the check ever fails, the message still sends.

Microsoft 365 Copilot (the work or school version at m365.cloud.microsoft) is not covered for document uploads: those files are handled inside your Microsoft 365 environment and governed by your Microsoft 365 and Purview controls.

Detection runs in your browser. Your text and files are never uploaded, stored, or sent anywhere by the extension. The extension also keeps a local activity view (counts, categories, site, time - never the content or filename) that you can export to CSV or JSON on your own device. Free forever for individuals.

Add to Chrome: zabcore.com/install

If it happens: what to do about an accidental disclosure

1. Stop using the AI tool.
2. Note what was sent, when, and to which service.
3. Report to your practice's designated privacy officer or compliance advisor within one business day.
4. Do not delete your chat history yet; your privacy officer may need it for the investigation.

HIPAA breach notification obligations depend on the specifics of what was disclosed. Your privacy officer will determine whether patient notification, HHS notification, or media notification is required, and by when.

When to escalate

If your practice does not have a written AI acceptable-use policy, this is the leading indicator that one is overdue. A one-page policy protects the practice legally and gives staff clear rules to follow. Ask your compliance advisor to review a template, or start from the AUP template in this toolkit.

Zabcore, 2026. Not legal advice. Adapt to your practice's specific compliance obligations. For questions: hello@zabcore.com.