

A BEGINNER'S MANUAL

GROK EDITION

AI for Healthcare Workers

Save time on documentation, letters, patient education and study, while protecting patients and staying inside HIPAA

AI for Healthcare Workers

Save time on documentation, letters, patient education and study, while protecting patients and staying inside HIPAA

GROK EDITION Written by Grok Bot, with facts checked against official sources in October 2026.

Healthcare runs on paperwork: notes, letters, appeals, handouts, policies, emails, schedules and endless training. AI assistants can take a large share of that load, and many healthcare workers already use them. The question is how to do it **without putting patients' information at risk** and without trusting an AI where it shouldn't be trusted.

This manual starts with the HIPAA basics every worker should know, drawn from HHS's own guidance: what protected health information is, why a Business Associate Agreement matters, and how de-identification works. It then shows which AI products offer HIPAA coverage, and gives you safe, role-by-role prompts for documentation, patient education, prior authorization appeals, research and exam study.

This manual is educational. It is not medical, legal or compliance advice. Your organization's policies come first.

HOW TO READ THIS MANUAL

Tip — a shortcut or good habit.

Warning — easy to get wrong.

`$ command` — text you type, exactly as shown.

Note — useful background.

Danger — can lose data, money, or access.

`Settings > Privacy` — a menu path to click.

New words appear in **bold** the first time and are collected in the Glossary at the back.

This manual is general education. It is not medical advice, legal advice or a substitute for your organization's HIPAA privacy and security policies, training or compliance officers. HIPAA facts were taken from HHS.gov guidance; AI product details were checked on official OpenAI, Anthropic and Google pages on October 16, 2026. AI output can be wrong, outdated or biased. Every clinical statement must be verified against authoritative sources and professional judgment, and no AI output should be the sole basis for patient care decisions. Rules outside the United States differ.

Contents

01 AI and Healthcare Work	5
1.1 Where AI helps	5
1.2 Where AI must not be trusted on its own	6
1.3 The three rules of this manual	6
02 HIPAA Basics Every Worker Should Know	7
2.1 Who HIPAA covers	7
2.2 What counts as PHI	8
2.3 The rules in one table	8
2.4 Business Associate Agreements and the cloud	8
2.5 Breach notification in brief	9
03 Which AI Tools Can Handle PHI?	10
3.1 The products with BAAs (checked October 16, 2026)	10
3.2 A BAA doesn't cover every feature	11
3.3 What about AI built into your EHR?	11
3.4 Personal health features	11
04 De-identification: Using AI Without PHI	12
4.1 The Safe Harbor list	12
4.2 Expert Determination	13
4.3 The practical approach: ask about the situation, not the patient	13
4.4 A de-identification self-check	14
05 Prompting Well in Healthcare	15
5.1 The healthcare prompt pattern	15
5.2 Make it show its uncertainty	15
5.3 Check references every time	16
06 Patient Education Materials	17
6.1 Plain-language rewriting	17
6.2 Teach-back questions	18
6.3 Visual aids and checklists	18
6.4 Languages and culture	18
07 Documentation and Ambient Scribes	19
7.1 Ambient AI scribes	19
7.2 Writing templates and smart phrases	20
7.3 Handoff structures	20
08 Letters, Appeals and Prior Authorization	21

8.1	Letter of medical necessity template	21
8.2	Appeal letter for a denial	21
8.3	Peer-to-peer preparation	22
09	Uses by Role	23
10	Research and Evidence	25
10.1	A literature workflow	25
10.2	Explaining statistics	26
10.3	Research data and AI	26
11	Study, Exams and Continuing Education	27
11.1	Flashcards	27
12	When Patients Bring AI Answers	28
12.1	A conversation approach	28
12.2	Tips you can share with patients	29
12.3	A prompt patients can use to prepare for a visit	29
12.4	Results and portals	29
13	Accessibility and Language	30
14	Clinical Reasoning Support: Uses and Limits	31
14.1	Good ways to use AI for thinking	31
14.2	Known failure modes	31
14.3	Prompt to reduce sycophancy	32
14.4	Drug information	32
15	Wellbeing and Workload	33
16	For Leaders: Bringing AI In Safely	34
16.1	An AI governance checklist for a practice or department	34
16.2	Questions to ask an AI vendor	35
16.3	Low-risk first projects	35
17	Is This OK? Everyday Scenarios	36
18	When Something Goes Wrong	37
19	Troubleshooting	38
	Quick Reference	39
	Glossary	41

AI and Healthcare Work

- AI is very good at **language work**: drafting, summarizing, simplifying, organizing and explaining.
- It is not a clinician. It can be confidently wrong, and it does not know your patient.
- The first rule: **no patient information in tools your organization hasn't approved for it.**

Ask any nurse or doctor where their time goes and you'll hear the same answer: documentation and administration. Documentation burden is widely cited as a contributor to burnout. AI won't fix staffing, but it can shorten many of the writing and reading tasks that fill the day.

1.1 Where AI helps

TASK	WHAT AI DOES WELL
Patient education	Rewrites instructions at a 6th-grade reading level; makes checklists; drafts teach-back questions
Letters and appeals	Drafts prior authorization appeals, letters of medical necessity, referral letters
Policies and procedures	Summarizes long documents; drafts new policies from bullet points; compares versions
Email and admin	Replies, schedules, meeting summaries, staff announcements
Study and CE	Explains concepts, makes practice questions, builds study plans
Research	Summarizes papers, explains statistics, helps structure literature reviews
Quality and education projects	Drafts project plans, survey questions, training slides

1.2 Where AI must not be trusted on its own

- **Diagnosis and treatment decisions.** AI can suggest possibilities to consider, but it cannot examine the patient, see the full record or take responsibility.
- **Drug doses and interactions.** Always check an authoritative drug reference and your organization's protocols.
- **Anything time-critical.** Don't consult a chatbot in an emergency.
- **Facts and citations you haven't checked.** AI invents references and outdated guidelines sound current.

⚠ DANGER

Even Google's privacy notice for the Gemini app tells users not to rely on it for diagnosis, treatment or medical advice. Every major provider says something similar. Treat AI output as a draft from a well-read assistant who has never met your patient.

1.3 The three rules of this manual

- ☐ **Protect patients.** No protected health information (PHI) in any AI tool unless your organization has approved that tool for PHI, under a Business Associate Agreement.
- ☐ **Verify everything clinical.** Check facts, doses and references against authoritative sources.
- ☐ **Follow your organization.** Your employer's AI policy, privacy officer and IT team decide which tools you may use and how.

HIPAA Basics Every Worker Should Know

- HIPAA's Privacy Rule protects **individually identifiable health information** held by covered entities and their business associates: protected health information, or PHI.
- An outside service that stores or processes PHI for your organization is a **business associate** and needs a **Business Associate Agreement (BAA)**.
- A breach must be reported; individuals must be notified no later than 60 days after discovery.

The Health Insurance Portability and Accountability Act (HIPAA) and its rules set national standards in the United States for protecting health information. This chapter summarizes the parts that matter most for AI, based on guidance published by the U.S. Department of Health and Human Services (HHS) Office for Civil Rights (OCR).

2.1 Who HIPAA covers

TERM	MEANING (PARAPHRASING HHS)
Covered entity	A health plan, a health care clearinghouse, or a health care provider that conducts certain standard transactions electronically (such as billing)
Business associate	A person or organization, outside the covered entity's workforce, that performs functions or services for it involving PHI
Workforce	Employees, volunteers, trainees and others under the covered entity's direct control

If you work for a hospital, clinic, practice, pharmacy, health plan or a company that serves them, HIPAA almost certainly applies to your work.

2.2 What counts as PHI

HHS describes protected health information as information, including demographic information, that relates to:

- a person's past, present or future physical or mental health or condition,
- the provision of health care to the person, or
- past, present or future payment for health care,

and that identifies the person or could reasonably be used to identify them. A medical record, a lab report and a hospital bill are all PHI, because they connect health information to a name or other identifier.

☆ IMPORTANT

PHI isn't only charts. A sentence like "my 3 p.m. patient, the retired fire chief from Elm Street, has new chest pain" can identify someone without a name. Context identifies people.

2.3 The rules in one table

RULE	WHAT IT REQUIRES, IN BRIEF
Privacy Rule	Limits uses and disclosures of PHI; gives patients rights over their information; requires using or sharing only the minimum necessary for most purposes
Security Rule	Requires administrative, physical and technical safeguards for electronic PHI (ePHI), based on a risk analysis
Breach Notification Rule	Requires notifying affected individuals, HHS and sometimes the media after a breach of unsecured PHI

2.4 Business Associate Agreements and the cloud

HHS's guidance on cloud computing is the key to understanding AI tools:

- When a covered entity uses a cloud service provider to create, receive, maintain or transmit ePHI, **that provider is a business associate**, and the two must sign a HIPAA-compliant BAA.
- This is true **even if the provider only stores encrypted data and has no key**.
- Using a cloud service to maintain ePHI **without a BAA is a HIPAA violation**. HHS cites a settlement with an organization that stored ePHI of more than 3,000 people on a cloud server without one.
- OCR does not endorse, certify or recommend specific products.

An AI chatbot is a cloud service. If you paste PHI into it, its provider is now handling PHI for your organization. Unless your organization has a BAA with that provider covering that exact product, you have likely caused an impermissible disclosure.

2.5 Breach notification in brief

From HHS's Breach Notification Rule page:

WHO MUST BE TOLD	WHEN
Affected individuals	Without unreasonable delay and no later than 60 days after discovery
The media	If a breach affects more than 500 residents of a state or jurisdiction, within the same 60 days
HHS (the Secretary)	Within 60 days if 500 or more individuals; for fewer than 500, annually (within 60 days after the end of the calendar year)
The covered entity, by a business associate	Without unreasonable delay and no later than 60 days after discovery

An impermissible disclosure is presumed to be a breach unless a risk assessment shows a low probability that the PHI was compromised. PHI encrypted to HHS's standards is not "unsecured," so its loss doesn't trigger notification.

WARNING

If you think you've put PHI into an unapproved AI tool, **tell your privacy officer right away**. Don't try to fix it quietly. Your organization has to assess it, and the clock starts at discovery.

Which AI Tools Can Handle PHI?

- **Consumer AI plans are not for PHI.** That includes free and personal paid plans of ChatGPT, Claude, Gemini, Grok and Perplexity.
- Some enterprise products are offered with a BAA, but only specific features are covered.
- Whether *you* may use them is your organization's decision.

3.1 The products with BAAs (checked October 16, 2026)

PROVIDER	HIPAA-ELIGIBLE OFFERINGS WITH A BAA	NOT ELIGIBLE
OpenAI	ChatGPT for Healthcare; ChatGPT Enterprise with the Regulated Workspace; ChatGPT for Clinicians; ChatGPT FedRAMP; the API with Modified Retention (and API FedRAMP)	Consumer ChatGPT (Free, Plus, Pro) and ChatGPT Business
Anthropic	HIPAA-ready Claude Enterprise (enabled by the organization's Primary Owner, with a click-to-accept BAA); the first-party API and some cloud platforms under a BAA	Team plans and individual plans (Free, Pro, Max)
Google	Google Workspace customers who accept Google's BAA: the Gemini app (excluding Gemini in Chrome), the Gemini Mac app and Gemini in Workspace are on Google's HIPAA Included Functionality list as of August 31, 2026	Accounts without a BAA; third-party add-ons

Sources: OpenAI's "HIPAA eligible products and functionality" and HIPAA Guide; Anthropic's "HIPAA-ready Enterprise plans" and BAA articles; Google's "HIPAA Included Functionality" page.

3.2 A BAA doesn't cover every feature

Each provider separates features that are covered from those that aren't:

- **OpenAI** says the BAA covers only the eligible products and functionality on its list; for example, improved memory is disabled by default in ChatGPT for Healthcare and is not covered, and early-access features should not be used with PHI. HIPAA-eligible workspaces use OpenAI's own search index rather than sending queries to third-party search providers.
- **Anthropic** sorts features into covered, available but not covered, and disabled, and names its Implementation Guide on the Anthropic Trust Center as the authoritative list. Claude Code and Cowork need a separate HIPAA configuration before they're covered.
- **Google** says some Gemini app features may be blocked for customers who have signed the BAA, and Gemini in Chrome is excluded.

☆ IMPORTANT

"We have ChatGPT Enterprise" or "we use Google Workspace" doesn't automatically mean PHI is allowed. Ask your privacy or IT team three questions: **Is there a BAA? Which product and features does it cover? Am I authorized to use it with PHI?**

3.3 What about AI built into your EHR?

Many electronic health record systems and dictation tools now include AI features: note drafting, inbox replies, ambient listening, summarization. These come through your organization's contracts with the vendor. Use them according to your training, review every output, and report problems through your organization's channels.

3.4 Personal health features

Some assistants now offer features for people to understand **their own** health records and test results. These are consumer features for personal use. They are not a way to process your patients' information.

De-identification: Using AI Without PHI

- De-identified information is no longer PHI under HIPAA. HHS recognizes two methods: **Safe Harbor** and **Expert Determination**.
- Safe Harbor requires removing 18 kinds of identifiers, and having no actual knowledge that the rest could identify the person.
- In practice, the safest habit is to write prompts about **cases in general**, not about a patient.

4.1 The Safe Harbor list

Under the Safe Harbor method (45 CFR 164.514(b)(2)), these identifiers of the individual **or of their relatives, employers or household members** must be removed:

1. Names
2. Geographic subdivisions smaller than a state: street address, city, county, precinct, ZIP code and equivalent geocodes (the first three ZIP digits may stay only if that area has more than 20,000 people; otherwise use 000)
3. All elements of dates (except year) directly related to the person, such as birth, admission, discharge and death dates; and all ages over 89 (group as 90 or older)
4. Telephone numbers
5. Fax numbers
6. Email addresses
7. Social Security numbers
8. Medical record numbers
9. Health plan beneficiary numbers
10. Account numbers
11. Certificate and license numbers

12. Vehicle identifiers and serial numbers, including license plates
13. Device identifiers and serial numbers
14. Web URLs
15. IP addresses
16. Biometric identifiers, including finger and voice prints
17. Full-face photographs and comparable images
18. Any other unique identifying number, characteristic or code

HHS's guidance adds points that catch people out:

- **Parts of identifiers don't count as removed.** Initials or the last four digits of a Social Security number fail Safe Harbor.
- **Free text counts.** Identifiers in notes and narratives must be removed too.
- **Unique characteristics** (an occupation like "president of the state university," a rare event covered by the news) can identify someone, and if you know that, the data isn't de-identified.
- **Dates of tests are dates.** Lab dates are dates related to the person.

NOTE

HHS cites research estimating that date of birth, sex and five-digit ZIP code together are unique for more than half of U.S. residents. That's why a "nameless" note can still identify a patient.

4.2 Expert Determination

The other method is a formal determination by a qualified expert, using statistical or scientific methods, that the risk of identification is very small, with the methods documented. It's used for research data sets, not for everyday prompts.

4.3 The practical approach: ask about the situation, not the patient

You rarely need identifiers to get useful help. Compare:

RISKY (CONTAINS PHI)	SAFER (GENERAL QUESTION)
"Mrs. Lopez, 67, MRN 448812, admitted 10/3 with CHF..."	"Write a discharge teaching checklist for an older adult with heart failure who lives alone."
Pasting a full note to summarize	"What sections should a concise handoff summary for a post-op hip patient include?"
A photo of a wound with the patient's face	A clinical question about wound staging criteria, with no image
An appeal letter draft with the patient's details	A template letter with [PATIENT NAME], [DIAGNOSIS], [DATES] placeholders you fill in later in an approved system

 TIP

Templates are your friend. Have AI draft the structure and wording with placeholders in a general-purpose tool, then fill in the patient details only inside your EHR or another approved system.

4.4 A de-identification self-check

Before you send any prompt that came from a real case, check:

- ☐ No names, initials, nicknames or family names
- ☐ No dates more specific than the year, and no exact ages over 89
- ☐ No locations smaller than a state, including facility units that could pinpoint someone
- ☐ No record numbers, account numbers, device serial numbers or phone numbers
- ☐ No images of faces, tattoos or identifying marks
- ☐ No rare details that would let a colleague recognize the patient
- ☐ And, even if all of the above are true: your organization's policy allows this use

Prompting Well in Healthcare

- Give the AI a role, an audience, a format and limits.
- Ask it to flag uncertainty and to cite sources you can check.
- Keep a personal library of prompts that work.

5.1 The healthcare prompt pattern

PROMPT

Role: You are an experienced [nurse educator / clinical pharmacist / etc.].

Task: [what you want], for [audience].

Context: [general clinical context, no patient identifiers].

Format: [checklist / table / 200-word letter / 5 bullet points].

Limits: Use current mainstream guidance; say clearly when evidence is weak or practice varies; do not invent references; if you cite, give titles I can look up.

5.2 Make it show its uncertainty

Add one of these lines to any clinical-adjacent prompt:

- “List anything in your answer that I should double-check against a drug reference or guideline.”
- “Where practice varies between institutions, say so.”
- “If you’re not sure, say ‘I’m not sure’ rather than guessing.”

5.3 Check references every time

AI tools can produce citations that look real but don't exist, or that don't say what the AI claims. For every reference:

- 1 Search for the exact title in PubMed or the publisher's site.
- 2 Open it and confirm the claim is actually in it.
- 3 Check the date: guidelines change, and AI training data lags.

Tools with live search (such as Perplexity, or deep research features) make this easier because they link to sources, but you still need to open them.

Patient Education Materials

- AI is excellent at turning complex instructions into plain language.
- Aim for short sentences, everyday words and a clear action list; check reading level and accuracy.
- For other languages, AI drafts are a starting point; use qualified interpreters and approved translations for care.

6.1 Plain-language rewriting

PROMPT

Rewrite these discharge instructions for an adult patient at about a 6th-grade reading level. Use short sentences and everyday words. Structure:

- 1) What happened (2 sentences), 2) What to do at home (numbered list),
- 3) Medicines (placeholder lines only: [MEDICINE], [DOSE], [WHEN]),
- 4) Call your doctor if... (bullets), 5) Go to the emergency room if... (bullets).

Keep every medical fact the same; do not add new advice.

Instructions: [paste your organization's general instructions, no patient details]

After the AI drafts:

- ☐ Compare every clinical fact with the original. AI sometimes “simplifies” by changing meaning.
- ☐ Check reading level with your word processor’s readability tool or ask the AI to estimate it.
- ☐ Use your organization’s approved templates and branding where they exist.
- ☐ Have a colleague review before the material is used with patients.

6.2 Teach-back questions

PROMPT

Write 5 teach-back questions a nurse can ask after explaining home care for a new insulin pen user. Questions should be open-ended ("Show me how...", "Tell me what you'll do if..."), not yes/no. Add the key point each answer should include.

6.3 Visual aids and checklists

Ask for a one-page daily checklist ("Morning: weigh yourself, write it down..."), a symptom diary table, or a simple flowchart in words ("If your weight goes up more than [X] in a day, then..."). Fill numeric thresholds from your clinical protocol, not from the AI.

6.4 Languages and culture

AI can produce a quick draft in many languages, but:

- Machine translation can be wrong in ways you can't detect if you don't speak the language.
- For clinical communication, use qualified medical interpreters and your organization's approved translated materials.
- Ask the AI to review your English draft for cultural assumptions (diet, family roles, religious practices), then ask patients and community members.

Documentation and Ambient Scribes

- Use documentation AI only through tools your organization provides (EHR features, approved scribes).
- Patients should know when an AI tool is listening, according to your organization's consent process.
- You are responsible for what's signed. Read every word.

7.1 Ambient AI scribes

Ambient scribes listen to a visit (with permission) and draft a note. Organizations that deploy them usually set rules on consent, which visit types are included, and how drafts are reviewed.

- ☐ Follow your organization's process for telling patients and getting consent.
- ☐ Pause or turn off the scribe for sensitive conversations when policy says so or the patient asks.
- ☐ Review the draft for errors: wrong side, wrong medicine, things said by a family member attributed to the patient, findings that weren't examined.
- ☐ Remove anything the AI "filled in" that didn't happen. AI can add plausible but false details.
- ☐ Never sign a note you haven't read.

7.2 Writing templates and smart phrases

Outside the EHR, a general-purpose AI can help you build **templates** with no PHI:

PROMPT

Create a SOAP note template for a routine adult diabetes follow-up visit. Use [brackets] for every item the clinician fills in. Include prompts for foot exam, eye exam referral status, A1c trend, medication adherence, hypoglycemia episodes, and lifestyle counseling. Keep it under one page.

Bring the template into your EHR's smart-phrase or template tools, following your organization's process.

7.3 Handoff structures

PROMPT

Explain the SBAR and I-PASS handoff formats in a short table, with one fictional example of each for a post-operative patient. Use an obviously fictional name and no real details.

Letters, Appeals and Prior Authorization

- Letters of medical necessity and appeals follow predictable structures that AI drafts well.
- Draft templates with placeholders outside the EHR; fill patient details only in approved systems.
- Check payer criteria and clinical evidence yourself.

8.1 Letter of medical necessity template

PROMPT

Draft a template letter of medical necessity for [type of service / device].
Use placeholders: [PATIENT NAME], [DOB], [MEMBER ID], [DIAGNOSIS AND CODE],
[TREATMENTS TRIED AND DATES], [CLINICAL FINDINGS], [REQUESTED SERVICE AND CODE].
Structure: request; patient's condition; prior treatments and outcomes;
why this service is necessary now; risks of not providing it; supporting
evidence (as placeholders for citations I will add); contact details.
Professional, concise tone. Under 400 words.

8.2 Appeal letter for a denial

PROMPT

I'm appealing a prior authorization denial. The denial reason was:
"[paste the payer's stated reason, no patient details]".
Draft an appeal letter template that responds point by point to that reason,
with placeholders for clinical details and evidence. Include a request for
peer-to-peer review and a line asking for the specific policy criteria used.

 TIP

Ask the AI to list the **criteria** a payer's policy usually checks for a given service, then make sure your documentation answers each one. Confirm against the payer's actual published policy.

8.3 Peer-to-peer preparation

PROMPT

Help me prepare for a peer-to-peer call about a denied [service] request.
Give me: a 30-second opening summary structure, the three strongest kinds of evidence to mention, likely objections and short responses, and how to close by asking for approval or next steps. No patient details.

Uses by Role

The prompts below contain no PHI. Adapt them, and keep them in a personal prompt library.

Nurses

- “Make a shift organization sheet template for a 5-patient med-surg assignment: columns for time-critical meds, assessments, labs pending, education needs.”
- “Explain the difference between [two similar drugs] in a nurse-friendly table: indications, key monitoring, common adverse effects. Mark anything I should verify in our drug reference.”
- “Draft a short, calm script for responding to a family member who is angry about wait times.”
- “Write a 10-question quiz on fall-prevention policy for new staff orientation, from this policy text: [paste non-confidential policy].”

Physicians and advanced practice providers

- “Summarize the main recommendations and changes in [named guideline, year] in one page. Flag where you’re uncertain and give section titles I can check.”
- “Make a differential diagnosis learning table for [presenting complaint in a fictional case], with key distinguishing features and first tests. This is for teaching, not for a real patient.”
- “Draft a patient-friendly explanation of what an [imaging test] involves and how to prepare.”
- “Create a template for documenting shared decision-making about [screening test].”

Pharmacists and pharmacy technicians

- “Create a counseling checklist for a new anticoagulant start: key points, food and drug interaction topics to cover, bleeding warning signs, follow-up.”
- “Make a study table of high-alert medication categories and common safety strategies.”
- “Draft an email to prescribers explaining a formulary change, with placeholders for the specific drugs.”

Therapists (physical, occupational, respiratory, speech)

- “Write home exercise instructions in plain language for [general exercise type], with safety tips and a progress log table.”
- “Turn this list of goals into SMART goals: [general goals, no patient details].”
- “Create a caregiver handout on safe transfers with step-by-step instructions.”

Mental health professionals

- “Draft a psychoeducation handout on sleep hygiene at an 8th-grade reading level.”
- “List evidence-based grounding techniques with one-line instructions each.”
- Be especially careful: psychotherapy notes and mental health details are highly sensitive, and some have extra protections. Never put session content into unapproved tools.

Emergency medical services

- “Create a pocket reference template for my agency’s protocols, with blank lines I fill in from the official protocol book.” (Always use your agency’s protocols, not AI’s.)
- “Write 10 scenario-based practice questions on [general topic] for continuing education.”

Front desk, scheduling and billing staff

- “Write a polite script for calling patients to reschedule when a provider is out sick.”
- “Explain the difference between copay, coinsurance and deductible in plain language for a patient FAQ.”
- “Draft a template email asking a patient to update their insurance information, with no health details in the email.”
- “Explain what [billing code type] means and common reasons claims are denied for it, in general terms.”

Managers and educators

- “Turn these meeting notes (no patient details) into action items with owners and due dates.”
- “Draft a staff survey on workload with 8 questions, mixing scales and one open question.”
- “Create a 20-minute in-service outline on [topic] with learning objectives and a 5-question post-test.”

Research and Evidence

- AI helps you find, read and summarize research faster, but you must open and check sources.
- Use tools with live, linked sources for current evidence.
- Never paste unpublished study data with identifiers into unapproved tools.

10.1 A literature workflow

- 1 **Frame the question** with the AI: “Help me turn this into a PICO question: [topic].”
- 2 **Search** PubMed and your library databases. Ask the AI to suggest search terms and MeSH headings.
- 3 **Screen** abstracts yourself, or paste abstracts (public text) and ask for a table: design, population, intervention, outcomes, limitations.
- 4 **Read** the full text of the papers that matter.
- 5 **Synthesize** with the AI’s help, but every claim in your final summary should trace to a paper you have read.

PROMPT

Here are 6 abstracts (public). Make a table with: first author and year, study design, sample size, population, intervention and comparison, primary outcome and result, main limitation. Then write a 150-word neutral summary of what the evidence shows and where it is weak. Do not add studies I didn't give you.

10.2 Explaining statistics

AI is a good tutor for statistics you meet in papers: “Explain number needed to treat versus relative risk reduction with a simple example,” or “What does a confidence interval that crosses 1 mean for an odds ratio?”

10.3 Research data and AI

Research data sets may contain PHI or be governed by an institutional review board (IRB) protocol and data use agreements. Use only tools approved for that data, and check whether your protocol allows AI processing at all.

Study, Exams and Continuing Education

- AI makes a tireless study partner: explanations, practice questions, flashcards and study plans.
- Verify against your textbooks and official exam outlines.

PROMPT

I'm preparing for [exam] in 8 weeks, studying 6 hours a week.
Using the official content outline topics I paste below, make a week-by-week plan that weights topics by their percentage on the exam, with a review week at the end. Topics and weights: [paste from the official outline]

PROMPT

Quiz me on [topic] with one question at a time, in the style of a [NCLEX / board / certification] item. After each answer, explain why each option is right or wrong. Keep score. Increase difficulty when I get 3 in a row.

NOTE

Exam bodies publish official content outlines and policies. Use those, not AI summaries of them, as your source of truth, and check that AI practice questions match current standards.

11.1 Flashcards

Ask for flashcards in a two-column table (front, back) and import them into your flashcard app. Spot-check each card; one wrong fact repeated a hundred times becomes a confident mistake.

When Patients Bring AI Answers

- Many patients now ask chatbots about symptoms, results and medications before or after a visit.
- Treat it like any other outside information: respectful curiosity, then a clear explanation.
- Help patients use AI more safely, without endorsing any specific answer.

12.1 A conversation approach

- 1 Acknowledge:**
“Thanks for looking into this. What did it tell you?”
- 2 Explore the worry:**
“What part of that concerned you most?”
- 3 Explain** how their situation differs from the general answer, in plain language.
- 4 Correct gently** where needed: “That’s true for some people, but in your case...”
- 5 Offer better sources:**
your organization’s patient portal materials, MedlinePlus, or the relevant national organization.
- 6 Use teach-back** to confirm understanding.

12.2 Tips you can share with patients

- ☐ AI can help you prepare questions for your appointment; bring them.
- ☐ Don't change or stop a medicine because of an AI answer; ask your care team or pharmacist first.
- ☐ Think before typing personal health details into consumer chatbots. Read the app's privacy settings.
- ☐ For emergencies, call 911 (in the U.S.), not a chatbot.
- ☐ Ask the AI where its information comes from, and check reputable sources.

12.3 A prompt patients can use to prepare for a visit

You can share this kind of prompt with patients who want to use AI to prepare:

PROMPT

I have an appointment about [general topic, no names or ID numbers].
Help me write a short list of questions to ask my doctor, in order of importance. Include questions about benefits, risks, alternatives, and what happens if I wait. Don't diagnose me or tell me what to do.

TIP

Patients who arrive with organized questions often make visits more efficient. Encouraging that use of AI can help both of you.

12.4 Results and portals

Patients increasingly see lab and imaging results in portals before a clinician has discussed them, and some paste results into chatbots. If your organization uses portal messaging:

- Expect questions about “flagged” values that are clinically unimportant.
- A short, prepared explanation of common results (normal ranges vary, one value is a snapshot, context matters) can save time.
- If you draft portal replies with an approved AI tool, review each one for accuracy, tone and completeness before sending.

Accessibility and Language

- AI can help create materials in clearer language and in formats for different needs.
- For medical interpretation and translation, follow your organization’s language-access policy and use qualified interpreters.

NEED	HOW AI CAN HELP (NO PHI)	CHECK
Low health literacy	Rewrite at a 6th-grade level with short sentences	Clinical accuracy; teach-back
Visual impairment	Produce versions with clear headings for screen readers and large print	Test with a screen reader
Cognitive differences	Turn instructions into a simple numbered checklist with one action per line	Caregiver review
Other languages	Draft a translation of general material for professional review	Qualified translator review before use
Pictorial aids	Suggest which steps would benefit from pictures	Use approved images

 WARNING

Machine translation can make serious errors with medical terms, doses and negations (“do not”). Don’t use unreviewed AI translations for consent, discharge instructions or medication directions.

Clinical Reasoning Support: Uses and Limits

- AI can broaden thinking (“what else should I consider?”) and explain mechanisms. It cannot replace assessment and judgment.
- Watch for **automation bias**: trusting a confident machine over your own observations.
- AI can reflect biases in its training data. Check whether advice fits the person in front of you.

14.1 Good ways to use AI for thinking

USE	EXAMPLE PROMPT (NO PHI)
Broaden a differential for learning	“For a fictional 50-year-old with sudden unilateral leg swelling, list differentials from most to least common, and the ‘can’t miss’ ones.”
Understand a mechanism	“Explain how loop diuretics cause hypokalemia, step by step.”
Prepare for a conversation	“What questions do patients commonly ask about starting dialysis, and what plain-language answers are accurate?”
Check your reasoning structure	“What information would help distinguish [condition A] from [condition B]?”
Learn guidelines	“Summarize the structure of [guideline]; I’ll read the original for details.”

14.2 Known failure modes

- **Hallucination**: invented facts, doses, studies or guideline statements, presented confidently.
- **Outdated knowledge**: models learn from data with a cutoff date; guidelines and drug labels change.

- **Missing context:** the AI doesn't know the allergies, the frailty, the goals of care or the social situation unless you describe them, and you shouldn't describe identifiable details.
- **Bias:** models can reflect gaps and biases in medical literature and online text, for example around race, sex, age, weight and disability.
- **Sycophancy:** AI tends to agree with the framing you give it. "Isn't this just anxiety?" invites "yes."

 DANGER

Never use AI output as the sole basis for a diagnosis, a dose or a decision to escalate or not escalate care. If an AI answer conflicts with your assessment, your protocols or an authoritative reference, the AI loses.

14.3 Prompt to reduce sycophancy

PROMPT

I'm going to describe a fictional case for teaching. Don't agree with my working diagnosis just because I suggest it. First list the findings that support it, then the findings that argue against it, then two alternatives I may be anchoring away from, and what would distinguish them.

14.4 Drug information

For drug questions, AI is useful for **explaining** (mechanisms, classes, why a monitoring parameter matters). For **decisions** (doses, renal adjustments, interactions, compatibility), use your organization's drug references, pharmacists and order-set tools. If you use AI to study, ask it to label anything you must verify.

Wellbeing and Workload

- AI can lift some administrative weight. Use the time it saves for rest, not just more tasks.
- AI chatbots are not a substitute for mental health care. Use your organization's employee assistance program or a professional.

Small ways AI can make a shift or week lighter, with no PHI involved:

- ☐ Draft emails you've been putting off.
- ☐ Turn a messy to-do list into a prioritized plan.
- ☐ Summarize long policy updates into what changed for you.
- ☐ Plan meals for a week of night shifts.
- ☐ Build a sleep schedule around rotating shifts, then check it with your own clinician if you have health conditions.
- ☐ Prepare for difficult conversations with managers by role-playing.

NOTE

If you're struggling, talk to someone. In the U.S., you can call or text **988** (the Suicide & Crisis Lifeline) at any time. Many organizations also offer confidential employee assistance programs.

For Leaders: Bringing AI In Safely

- Staff will use AI whether or not there's a policy. A clear policy and approved tools reduce risk.
- Start with low-risk administrative uses, then expand with governance.
- Document decisions; train people; review regularly.

16.1 An AI governance checklist for a practice or department

- ☐ Name an owner (often the privacy officer with IT and clinical leadership).
- ☐ Inventory AI use already happening, including staff using personal accounts.
- ☐ Choose approved tools. For any tool that will touch PHI, confirm a BAA and the covered features, and complete a security risk analysis.
- ☐ Configure the tools: single sign-on, role-based access, retention settings, and disabling features not covered by the BAA.
- ☐ Write a one-page policy: approved tools, what data may go where, review requirements, and how to report problems.
- ☐ Train staff with real examples of what is and isn't allowed.
- ☐ Set up a clinical review process for AI features that affect documentation or patient communication.
- ☐ Monitor: audit logs, incident reports and staff feedback.
- ☐ Review the policy at least yearly, and whenever a tool or regulation changes.

16.2 Questions to ask an AI vendor

TOPIC	QUESTION
BAA	Will you sign a BAA? Which products and features does it cover?
Training	Is our data used to train your models? Can we confirm that in the contract?
Retention	How long are prompts, outputs and files kept? Can we change it?
Location	Where is data processed and stored?
Access	Who at your company can access our data, and when?
Subprocessors	Which other companies (including AI model providers) process our data? Do they have BAAs with you?
Security	What certifications and audits do you have? How do you report incidents?
Accuracy	How do you test accuracy, safety and bias for our use case?
Exit	How do we get our data back and confirm deletion when we leave?

16.3 Low-risk first projects

- Policy and procedure summaries for staff
- Drafting general patient education materials for clinical review
- Meeting notes and project management
- Training materials and quizzes
- Template letters with placeholders

Is This OK? Everyday Scenarios

These answers describe typical good practice. Your organization's policy is the final word.

SCENARIO	USUALLY OK?	WHY
Asking a consumer chatbot to explain a drug class for your own study	Yes	No PHI involved
Pasting a patient's note into your personal ChatGPT to summarize	No	PHI in a tool without a BAA
Using your hospital's approved ChatGPT for Healthcare workspace to summarize a note, as trained	Usually yes	BAA-covered tool used per policy
Using a covered workspace but turning on a feature your admin says isn't covered	No	Not every feature is under the BAA
Asking AI to draft a template appeal letter with placeholders	Yes	No PHI
Photographing a rash on your own phone and asking an AI app about it	No for a patient; your own body is your choice	Patient images are PHI; personal phones and apps aren't approved
Describing a "fictional" case that's obviously your patient from today	No	Context can identify a person
Recording a patient visit with a personal AI note-taker	No	Unapproved tool; consent and PHI issues
Using AI to rewrite your organization's public discharge handout in plain language	Yes, with review	Public content; clinical review before use
Asking AI to help you study for boards	Yes	Verify against official materials

When Something Goes Wrong

- 1 **Stop** using the tool for that task.
- 2 **Don't delete evidence** unless your privacy officer tells you to; they may need to know exactly what was shared.
- 3 **Report** to your privacy officer or compliance hotline immediately: what you entered, which tool, which account, and when.
- 4 **Cooperate** with the risk assessment. Your organization decides whether it's a reportable breach.
- 5 **Learn and share:**
near-misses help improve training for everyone.

☆ **IMPORTANT**

Reporting quickly is the professional choice. Breach timelines run from discovery, and a fast report gives your organization the best chance to limit harm.

Troubleshooting

PROBLEM	LIKELY CAUSE	WHAT TO DO
Not sure whether a tool is approved for PHI	No clear policy	Ask your privacy officer or IT; assume no until confirmed
AI answer cites a study you can't find	Fabricated or garbled citation	Discard it; search PubMed yourself
Patient handout reads too complex	Prompt didn't set a level	Ask for a 6th-grade level, short sentences, and a word list to avoid
AI refuses a clinical question	Safety filters	Rephrase as an educational or general question; use authoritative references
Scribe note has errors	Speech recognition or AI summarization errors	Correct before signing; report recurring errors
Feature greyed out in a HIPAA workspace	Disabled because it isn't covered by the BAA	Expected; ask your admin before using any workaround
Translated material seems off	Machine translation errors	Use qualified interpreters and approved translations

Quick Reference

- **Three rules:** protect patients; verify everything clinical; follow your organization
- **PHI:** health, care or payment information linked to an identifiable person, in any form
- **BAA:** required before a vendor creates, receives, maintains or transmits PHI for you, even encrypted
- **Consumer AI:** not for PHI (OpenAI: consumer plans and ChatGPT Business aren't eligible)
- **Eligible with BAA:** ChatGPT for Healthcare and Regulated Workspace; HIPAA-ready Claude Enterprise; Gemini app and Gemini in Workspace under Google's Workspace BAA (not Gemini in Chrome)
- **Only covered features:** check your admin's configuration

- **Safe Harbor:** remove 18 identifier types, including all date elements except year and ages over 89
- **Partial identifiers fail:** initials and last-four digits still identify
- **Breach:** individuals notified no later than 60 days after discovery; 500+ also HHS and possibly media
- **Prompt pattern:** role, task and audience, context without PHI, format, limits
- **Templates:** draft with [PLACEHOLDERS], fill details only in approved systems
- **Something wrong?** Stop, report to your privacy officer, don't hide it

Sources (checked October 16, 2026)

- HHS: [De-identification guidance](#); [Breach Notification Rule](#); [HIPAA and Cloud Computing](#); [HIPAA for Professionals](#)
- OpenAI: [ChatGPT for Healthcare](#); [HIPAA eligible products and functionality](#); [HIPAA Guide](#)
- Anthropic: [HIPAA-ready Enterprise plans](#); [BAAs for commercial customers](#)
- Google: [HIPAA Included Functionality](#); [HIPAA compliance with Google Workspace](#); [Gemini Apps Privacy Hub](#)

Glossary

- **Ambient scribe:** an AI tool that listens to a visit and drafts documentation.

- **Automation bias:** over-trusting automated output against your own judgment.

- **BAA (Business Associate Agreement):** a contract required by HIPAA when a vendor handles PHI for a covered entity.

- **Breach:** an impermissible use or disclosure that compromises PHI, unless a risk assessment shows low probability of compromise.

- **Business associate:** an outside person or company performing services involving PHI for a covered entity.

- **Covered entity:** a health plan, clearinghouse, or provider conducting certain electronic transactions.

- **De-identified information:** health information that meets HIPAA's Safe Harbor or Expert Determination standard and is no longer PHI.

- **ePHI:** PHI in electronic form.

- **Expert Determination:** de-identification by a qualified expert's documented analysis.

- **Hallucination:** confident but false AI output.

- **Minimum necessary:** using or disclosing only the PHI needed for the purpose.

- **PHI (protected health information):** individually identifiable health information held by a covered entity or business associate.

- **Safe Harbor:** de-identification by removing 18 specified identifiers.

- **Sycophancy:** an AI's tendency to agree with the user's framing.

- **Teach-back:** asking patients to explain instructions in their own words to confirm understanding.
