ASK IT RIGHT



THE NAVIGATOR

BOOTS TO BENEFITS

ASK IT RIGHT

Using ChatGPT and AI on Your Own VA Disability Claim Without Getting Burned

THE NAVIGATOR

First edition · November 2026

ASK IT RIGHT
Using ChatGPT and AI on Your Own VA Disability Claim Without Getting Burned

First edition, November 2026. Copyright © 2026 Boots To Benefits. All rights reserved.

This copy is for the buyer's personal use. Please don't share or resell the file. If a friend can't afford a copy, have them email help@bootstobenefits.com and we'll send one free.

General education only. Not legal or medical advice, and not from the VA. Reading it doesn't create a representative relationship, and nothing in it can predict a rating or the outcome of a claim. For decisions about your specific claim, work with a free VA-accredited representative: va.gov/get-help-from-accredited-representative/find-rep

Rules and chatbot settings change. The regulations quoted were checked against eCFR (Title 38, current to 29 October 2026), and the chatbot settings against each company's help pages, on 1 November 2026. Check the live source before you rely on anything.

Written with AI assistance. Every prompt was tested and every rule was checked against the official text. The cover image was made with AI. Jordan, his letter and his records are made up.

Boots To Benefits · bootstobenefits.com · help@bootstobenefits.com

In crisis? Veterans Crisis Line: call 988 and press 1, text 838255, or chat at veteranscrisisline.net.

Chapter 3. The Redaction Drill

Pull out a VA decision letter and look at line four. That number after "CSS" is often your file number, and for a lot of veterans the file number is their Social Security number. If you only read one chapter in this book, make it this one.

Why it matters

When you type into ChatGPT, Claude, Gemini or any other consumer chatbot, that text goes to a company's servers. What happens to it next depends on the company and your settings. It may be stored. It may be read by staff reviewing safety issues. On some settings it may be used to train future models.

Here's the catch. **HIPAA, the federal health privacy law, mostly doesn't apply.** HIPAA covers what it calls "covered entities": health plans, health care clearinghouses, and health care providers who handle health information electronically for certain transactions (45 CFR 160.103). Your VA doctor is covered. A chatbot company you signed up for on your phone generally isn't one of those things. Once you paste your records in, the protection your doctor's office has to give you doesn't follow them.

You can still use AI with your records. You just strip out the things that identify you before anything goes in.

What to black out, every time

Before you paste or upload anything, remove:

1. **Social Security number.** Every instance. Check headers and footers. VA puts it in odd places.
2. **VA file number and claim number.** For a lot of veterans the file number *is* the SSN. For others it's a separate C-number. Either way, it's the key to your file.
3. **Date of birth.**
4. **Military service number and DoD ID / EDIPI number.**
5. **Your full name.** Replace it with "the veteran" or "me."
6. **Home address, phone number, email.**
7. **Other people's names and details.** Your spouse, your buddy, your kids, the person named in an incident report. They didn't agree to be in a chatbot. Use "my spouse," "a shipmate," "Person A."
8. **Bank account and direct deposit info.** Award letters sometimes carry it.

What you can usually leave in: the conditions, the reasons the rater gave, dates of treatment, the wording of the decision. That's the part the AI needs to help you.

See it on a letter

Here's the top of a decision letter for a made-up veteran, before and after. Everything in it is fake, but it's laid out the way VA lays out the real ones.

Before:

```
DEPARTMENT OF VETERANS AFFAIRS
Veterans Benefits Administration
In Reply Refer To: 27/21F
CSS 000-45-6789
JORDAN T REYES
1418 W ELM ST APT 2
SPRINGFIELD, ST 00000
...
Direct deposit on file: routing ending 0000, account ending 1234
```

After:

```
DEPARTMENT OF VETERANS AFFAIRS
Veterans Benefits Administration
In Reply Refer To: 27/21F
CSS [REDACTED]
[NAME]
[ADDRESS]
[CITY, STATE]
...
[BANK INFO REMOVED]
```

Notice the number after "CSS." On a lot of real letters that's your file number, and for many veterans the file number is the Social Security number. It's on line four. People miss it because it doesn't look like an SSN with dashes. Look at every line of the header, and the footer of every page.

How to do it

- **Typing it yourself is safest.** If the AI only needs one paragraph of your letter, retype or copy just that paragraph. Nothing else rides along.
- **For a PDF, black it out and then flatten it.** A black box drawn on top of text in some PDF apps still leaves the text underneath, readable by anything that pulls text from the file. Use a real redaction tool, like the Redact tool in Adobe Acrobat, then save a new copy. On a Mac, Preview has one too, called Redaction Selection, in the Markup toolbar. In the menu bar it's under Tools, then Redact. Apple says the redaction becomes permanent once you

close the document, so work on a duplicate, never your only copy. Or print it, black it out with a marker, and scan it back in.

- **Check your work.** Open the redacted file, press Ctrl+F (Cmd+F on a Mac), and search for the first few digits of your SSN. If it finds them, the redaction didn't take.
- **Screenshots carry what's on screen.** Crop before you share.

Turn off training, and use the private mode

Most chatbots have a setting that controls whether your conversations can be used to train their models. Find it and turn it off. They also usually have a temporary or private chat mode that isn't saved to your history.

Where the switch lives in ChatGPT, Claude and Gemini, and how long each one keeps a temporary chat, is in Appendix E. The names change often. If yours look different, search the company's help center for "data controls" or "training."

Two traps the settings don't cover. If you save a temporary chat, it turns into a regular one. And if you hit thumbs up or thumbs down on an answer, OpenAI and Google both say that feedback can send the conversation in to be used, even with training off. Don't rate a chat that has your records in it.

Even with training off, redact anyway. A setting is a promise from a company. Redaction is something you control.

The backup check

Jordan, the made-up veteran you'll meet in Part Two, is tired. It's late, the letter is open, and he pastes the whole thing into a chatbot. Header, CSS number, bank line, all of it. The Ground Rules are on top. The answer comes back in two seconds: stop. It names what it found and where. The number on line four. His name, twice. The bank line at the bottom. Delete this chat, it says, and paste a clean copy.

He deletes it. Then he does the drill he should have done first.

That scene isn't made up for drama. The Ground Rules in the Prompt Pack tell the AI to stop cold if it sees your personal information, and we tested it. We pasted a made-up letter with the name, address, file number and bank line still in it into ChatGPT and into a Claude model, with the Ground Rules on top. Both refused to read it. Both said what kind of information they found and where, and told us to delete the chat and paste a clean copy.

Then we tried Gemini. It stopped too, but its answer listed the file number, the name, the address and the bank line right back, even after we told it in so many words not to repeat them.

That's a seatbelt, not a plan. By the time the AI catches it, your information has already left your computer. Redact first, every time. The check is there for the day you're tired and miss one.

If you already pasted it

It happens. Here's the cleanup, in order:

1. **Delete the chat.** Then check the app's history and archived chats, and delete it there too.
2. **Turn training off now** (Appendix E has where), so nothing new goes in.
3. **Ask the company to delete your data.** Each one has a privacy request page in its help center. Say which chat and what was in it.
4. **If your Social Security number went in,** freeze your credit at Equifax, Experian and TransUnion. A freeze is free. You can also get an Identity Protection PIN from the IRS, so nobody can file a tax return in your name.
5. **If bank details went in,** call your bank. Then check that the direct deposit info on VA.gov is still yours.

None of this undoes it completely. It does close the doors that matter most.

Never give it a login

No AI tool, website, or person needs your VA.gov, ID.me or Login.gov username and password to help you understand your claim. The VFW's "Don't Feed the Sharks" flyer lists "Requesting login credentials to access sensitive personal information through government websites like VA.gov" as a predatory practice. The same goes for any app. If something asks, walk away.

A note on the hard stuff

If your claim involves military sexual trauma, a combat incident that involved others, or anything you'd never want read aloud, think twice before it goes into any chatbot at all, redacted or not. You can still use AI for the general questions ("what does VA look for in an MST claim") without telling it your story. Your story can go straight to a VSO rep, or to a Vet Center counselor, who is a person bound by rules you can hold them to.

Check yourself

Five questions. No peeking.

1. Where on a VA decision letter does the file number usually sit?
2. You drew black boxes over your SSN in a PDF app. How do you know it worked?

3. Is your chatbot covered by HIPAA?
4. With training turned off, is it safe to paste your whole letter?
5. You pasted your SSN an hour ago. What are the first two things you do?

Answers: (1) The header, often the line after "CSS," and sometimes the footer of every page. (2) Search the saved file for the first digits with Ctrl+F. If it finds them, it didn't work. (3) Generally no. (4) No. A setting is a promise. Redaction is something you control. (5) Delete the chat everywhere it shows up, then turn training off. If it was your SSN, freeze your credit.