

Same File. Same AI. Different Discipline.

People sometimes ask, fairly: "Couldn't I just run my ancestor's file through AI myself?"

You can. So we did it for you, as an experiment, and published both results. We took our Thornedale specimen pension file (a fictional soldier, so both documents can be shared in full) and processed it twice with the **same AI model**:

- **Once with no method** — the file and the question a person would type: "Can you read my ancestor's Civil War pension file and tell me what it says about him and his family?" That unedited reply is published as the Generic AI Run Transcript.
- **Once as an A Soldier's Life report** — the same model working under our evidentiary rulebook, with mechanical checks run before anything ships.

First, credit where due

The generic run is not garbage. It read the file, got the names and dates right, noticed a date wrinkle, and even guessed correctly that an 1886 pension increase came from a change in law. Modern AI is genuinely good at reading. That is exactly why the real differences are worth seeing, because they are differences of discipline, not of reading ability.

What the comparison shows

- 1. The bounty money vanished, and nothing tells you.** This file contains a bounty entry deliberately obscured by an ink blot. Our report flags it: hard to read, page 2 of the pension file, no guess made. The generic run simply never mentions the bounty at all. Not a wrong answer, something worse: a silent omission. Reading the generic summary, you cannot know it is incomplete. That is the whole problem with unverified reading, in one line item.
- 2. Two death dates, one quietly chosen.** The army recorded June 14, 1864. The widow swore to "on or about the twelfth," because she learned of her husband's death by notification, she was not there. The generic run states "Died June 14, 1864" as flat fact, then mentions the discrepancy near the end as a "small inconsistency, typical of real records." Our report treats it as what it is: a source conflict, both readings shown, with the explanation of why the widow's honest testimony differs from the army's record. One approach quietly picks a winner. The other shows you the record.
- 3. Nothing in the generic run can be checked.** It cites no pages. Our report ties every fact to the document and page that states it, and a machine-checked table of thirty sourced facts is validated before the report ships. If you want to verify a claim against the original scans, one document hands you the page number. The other hands you homework.
- 4. You cannot tell whether the generic run read everything.** Ten pages went in. Did all ten get read? With the generic run there is no way to know, and on a real sixty-page file that uncertainty grows with every page. Our pipeline refuses to ship a report unless every supplied page is accounted for in the document inventory. Coverage is checked by software, not assumed.
- 5. What the file does NOT say is missing entirely.** Our report lists the research gaps (what the file never establishes, and which records would), inventories every page, and registers everything hard to read. The generic run offers none of this. For genealogy, what the record does not prove is half the finding.

The honest bottom line

The AI is not the product. The same model produced both documents. What you are paying for is the discipline wrapped around it: an evidentiary rulebook built from hundreds of real pension files, mechanical verification that every page was read and every fact is traceable, honest flags instead of silent gaps, and, on the reviewed tier, a person re-reading every flagged entry against the original scans.

Run your file through AI yourself. Genuinely, it will be interesting. Then ask it which page the bounty was on.

Both documents are available in full at asoldierslifereports.com. The specimen soldier is fictional; every page of the source file says so.