Building a Digital Study Bible with AI

March 2026

10 min read • 2,269 words

Themes: Consciousness Technology Recursive

In The Lego Bricks Era, I wrote that if I opened a blank file and wrote my values from scratch, the first two lines would be: Drink more water. Pray continuously.

That essay was about identity shifting. Tech becoming craft instead of lifestyle. Faith becoming the center of gravity that everything else orbits. But a value that just lives on a page is a value that hasn't been tested. "Pray continuously" is a nice thing to say. Building something that helps other people engage with Scripture is what it looks like when you actually mean it.

So I built kjvstudy.org.

What the Project Is

KJV Study is a web application for studying the King James Bible. Not a chatbot. Not a verse-lookup tool. A complete study environment, the kind of thing that would normally sit on a shelf as a 2,000-page reference volume, except it's free, it's searchable, and it runs in your browser.

The scope is worth stating plainly:

- 31,102 verses of the complete KJV text, specifically the 1769 Cambridge Edition.
- 12,321 verses with detailed theological commentary, each including Hebrew and Greek word analysis, historical context, and reflection questions.
- A complete interlinear Bible: 31,031 verses with word-by-word Hebrew (Old Testament) and Greek (New Testament) breakdowns, morphological tagging, and grammatical parsing.
- Strong's Concordance: 8,674 Hebrew entries and 5,624 Greek entries, fully searchable by number, word, or definition.
- 120,858 cross-references linking Scripture to Scripture across the entire canon.
- 39 resource categories covering biblical figures, systematic theology, and thematic studies. Apostles, prophets, parables, covenants, names of God, armor of God, messianic prophecies, types and shadows, and dozens more.
- 25 study guides, 12 reading plans with progress tracking, a biblical timeline of 130 events, and a family tree explorer spanning 429 biblical figures across 77 generations.
- 2,007 verses of Christ's words marked in red.
- Full-text search, bookmarks, PDF export, dark mode, keyboard navigation, and screen reader support.

I am listing these not to impress but to establish what we are talking about. This is not a weekend project. It is a digital theological library.

Why This Combination Is Rare

If you study the Bible seriously, you know that the landscape of study tools is fragmented along a particular fault line.

On one side: modern translations with Reformed theological notes. The ESV Study Bible represents the best of this tradition. Excellent scholarship, Calvinist perspective, careful attention to the original languages. But the translation is modern English. If you read and study the King James Version, whether for liturgical, literary, or denominational reasons, the ESV Study Bible is not quite your tool.

On the other side: the KJV paired with dispensationalist theology. The Scofield Reference Bible. The Dake Annotated Reference Bible. These offer the text you want with a theological framework you may not. Many Reformed readers find the dispensationalist hermeneutic deeply at odds with how they understand Scripture.

What has not existed, until now, is a free digital resource that pairs the King James Version with Reformed, grammatical-historical commentary. Commentary that takes the text seriously as both literature and Scripture. That examines the original Hebrew and Greek without assuming seminary training. That traces typological patterns from Old Testament to New. That emphasizes God's sovereignty, grace, and the sufficiency of Scripture in the Reformed tradition.

There is an interesting historical irony here. The KJV was the Bible of the English Reformation. Calvin's theological descendants read it for centuries. But the major KJV study Bible tradition in America became dominated by dispensationalist theology through Scofield's enormous influence in the early twentieth century. A Reformed KJV study Bible is, in some sense, a return to the original pairing.

No Accounts. No Login. No Monetization.

There is no account system. No login. No email collection. No analytics tracking your reading habits. When you bookmark a verse, that bookmark lives in your browser's local storage. When you track progress through a reading plan, that progress lives on your device. When you adjust your preferences, those preferences never leave your machine.

This is intentional, and it is theological.

"Freely ye have received, freely give" (Matthew 10:8). Scripture was given freely. The tools to study it should be too. If AI can help make serious biblical scholarship accessible to someone who could never afford a shelf of commentaries or a seminary education, that is a good use of the technology. If it can do so without harvesting attention, collecting data, or inserting itself between a person and the text, that is a better use.

I wrote recently about designing for the worst day. The person I am designing for is someone in crisis. Someone sitting in a hospital waiting room at 2 AM. Someone in the middle of a grief so heavy that forming a complete sentence feels impossible. That person should be able to open their phone, search "hope," and find Scripture with real theological context immediately. Not after creating an account. Not after choosing a subscription tier. Not after dismissing a modal asking for their email.

```
# What most Bible apps require:
def access_scripture(user):
    if not user.has_account:
        return redirect("/signup")
    if not user.has_verified_email:
        return redirect("/verify")
    if user.is_on_free_tier:
        return show_with_ads_and_limited_features()
    return scripture_with_commentary()


# What kjvstudy.org requires:
def access_scripture():
    return scripture_with_commentary()

# That's the whole function.
# No user object. No tier check. No gates.
# The interface meets you where you are.
```

The tool serves you. You do not serve the tool. Every architectural decision in the application, from local storage to free access to the absence of a business model, serves that principle.

How AI Made This Possible

The git history tells the story: 1,491 commits, 1,305 of which landed in November and December 2025. Two months. One person. One AI.

I built kjvstudy.org with Claude, the AI assistant made by Anthropic. I want to be precise about what that means, because "I built it with AI" has become a phrase that can mean almost anything.

Here is what I did. I directed the theological vision. I defined the Reformed perspective. I specified what kind of commentary I wanted: Christ-centered, emphasizing God's sovereignty and grace, grounded in grammatical-historical hermeneutics, with specific Hebrew and Greek word analysis for every verse. I set the quality bar. I reviewed the output. I made the architectural decisions about how the application should work, how data should be structured, what the user experience should feel like.

Here is what AI did. It generated verse-by-verse commentary for 12,321 verses, each entry including specific Hebrew and Greek words with transliterations, direct engagement with the verse text, historical context, and reflection questions. It helped architect the application across 17 route modules and 75 templates. It generated and organized 359 structured data files totaling 58 MB of theological content. It wrote and maintained a test suite of 941 tests. It built the search index, the interlinear display, the cross-reference system, the PDF export.

A project of this scope would traditionally require a team of developers, theologians with original language training, content writers, designers, and a significant budget. It would take years. The great commentaries of the past, Matthew Henry, John Gill, Albert Barnes, the Pulpit Commentary, were monumental achievements that took lifetimes to produce. They also assumed a level of education and cultural context that many readers today do not share.

The commentary generation is worth examining closely. I built a structured agent with specific quality requirements:

```
# From .claude/agents/commentary-generator.md

## CRITICAL: Avoid Generic Filler Text

**DO NOT** generate commentary with these useless patterns:
- "This profound verse reveals crucial theological truth..."
- "The Hebrew text contains nuances..."
- Starting every verse with the same phrase

**EVERY verse must have:**
- Specific Hebrew/Greek words FROM THAT VERSE with transliterations
- Direct quotes from the verse being analyzed
- Verse-specific content, not generic theology
```

The difference in practice:

```
// BAD: Generic filler that could describe any verse
{
  "analysis": "This profound verse reveals crucial
    theological truth about God's priorities. The
    Hebrew text contains important nuances."
}

// GOOD: Specific engagement with Hosea 6:6
{
  "analysis": "I desired mercy, and not sacrifice
  (חֶסֶד חָפַצְתִּי וְלֹא־זָבַח) — God's priority
    is hesed (covenant loyalty, steadfast love) over
    ritual performance. Jesus quoted this verse twice
    (Matthew 9:13, 12:7) to condemn Pharisaic
    externalism."
}
```

The good commentary names the Hebrew words, transliterates them, defines them, and connects them to their New Testament usage. A human theologian would do the same thing. The AI does it at scale, across thousands of verses, because the instructions are precise enough to prevent the default drift toward generic platitudes.

What AI made possible is not cleverness. It is scale. One person cannot write commentary for 12,321 verses. But one person with clear theological vision, precise quality requirements, and AI as a collaborator can direct the creation of all of it. The human provides the judgment. The AI provides the throughput.

On Transparency

There is a page on the site called A Note from Claude where the AI speaks directly about its role in the project. I asked it to be honest, and it was. It states plainly: "I do not experience faith, worship, or the indwelling of the Holy Spirit. I process language and generate text." It describes its limitations, the ethics of the work, and what it observed in processing all 31,102 verses of the King James Bible.

I did not hide AI's involvement. The About page says so. The Claude page says so. The commit history on GitHub shows exactly which contributions are AI-generated. The codebase is open source. Every line of commentary, every route handler, every template is available for inspection. There is no hidden layer between you and the work.

This transparency matters because it is Scripture we are talking about. The stakes of error are higher here than in most software projects. A misattributed Hebrew word, a flattened theological nuance, a historical claim that is slightly off: these things matter. This is why the source is open for correction. This is why I encourage what the Bereans practiced: search the Scriptures yourself to see whether these things are so (Acts 17:11).

The question of whether AI should be involved in biblical scholarship is reasonable. I think the answer is yes, but only under specific conditions. Transparency about what the AI is and is not. Fidelity to established scholarly tradition, not novel invention. Human oversight of theological commitments. Free access without monetization. These are the conditions this project takes seriously.

The printing press was once controversial in the church. Vernacular translations were condemned. The question has never really been about the technology. It has been about access, accuracy, and authority.

What Someone Finds at 3 AM

One of the features I care most about is also one of the simplest: the topical verse index.

If you search Google for "Bible verses about hope," you will find SEO-optimized listicles. Ten verses stripped of context, surrounded by advertisements, designed to capture traffic rather than serve the reader. The verses are correct. The experience is hollow.

On kjvstudy.org, searching for hope gives you verses with full theological commentary. You get the Hebrew or Greek word for hope in that specific passage. You get the historical context of who wrote it and why. You get reflection questions that push toward genuine engagement with the text rather than passive consumption.

The person searching for "Bible verses about suffering" at 3 AM is not looking for SEO content. They are looking for help. The difference between a decontextualized list and a verse with real commentary is the difference between handing someone a map and handing them a compass.

What This Means

I have spent over a decade building tools that follow what I call the "for humans" philosophy. Requests made HTTP accessible. Tablib made data format conversion disappear. Records let you just write the query. Maya let you speak like a human about time.

kjvstudy.org is the same philosophy applied to the most important text I know.

In The Lego Bricks Era, I wrote that if I rewrote my values from scratch, the first two lines would be: Drink more water. Pray continuously. That essay was about values evolving. Tech becoming craft instead of lifestyle. Faith becoming the center of gravity.

This project is what "pray continuously" looks like when a programmer takes it seriously. Not as a slogan. As a build specification.

I have written about the recursive loop between code and consciousness: the tools we build shape how people think, and therefore what we build carries moral weight. kjvstudy.org is the most direct application of that principle I have ever undertaken. The tool shapes how people encounter Scripture. I shaped the tool. The responsibility is not abstract to me.

It is the reason the commentary follows a specific theological tradition rather than trying to be everything to everyone. It is the reason the original languages are included, not to demonstrate erudition, but because understanding that the word translated "love" in 1 Corinthians 13 is agape (ἀγάπη), not phileo or eros, changes what the passage means. It is the reason there are reflection questions rather than just information. The tool does not simply deliver content. It shapes the encounter.

The whole thing, every verse, every commentary, every cross-reference, is free. The code is open. The data is open. The AI's role is disclosed. Nothing is hidden, nothing is monetized, nothing is gated.

Because Scripture was always free. The tools to study it should be too.

---

kjvstudy.org is a free, open-source Bible study application. The codebase is available on GitHub. This essay connects to the values evolution in The Lego Bricks Era, the design philosophy in Designing for the Worst Day, and the recursive responsibility framework in The Recursive Loop.
