Reducing the cognitive load of AI changes

Andrew · 1 October 2026 · 2 min read

When reviewing large amounts of AI-generated code, I often find that the LLM chooses terms for abstractions that do not always map to my own choices. For instance, what it may call a `MutationIntent` might personally be more natural to me as an `EditRequest`. Because the LLM's choice of words is not my ideal choice, I have to do a mental lookup of what it means every time I see it, which adds cognitive load.

This may seem like a small friction, but the cognitive load accumulates when considering how dozens of new terms interact in unfamiliar code. I can only hold a finite number of these semantic lookups in my head before I start misinterpreting how things work.

To minimize this, before review, I post-process AI changes with this prompt:

```
Please review the changes and extract any unconventional or bespoke terms used
for abstract objects, processes and concepts. Create a temporary markdown file
with each term, its meaning and why it was chosen, and some proposed alternative
terms for it. I will then use this markdown file to confirm the term choice or to
provide my own custom term. You will then incorporate any changes. The goal here
is to map your language choices to my own language choices so that I can
understand the concepts more easily.
```

I then go through and confirm term choices. The AI does find-and-replace everywhere, including documentation. The resulting code is much easier to review because now it's written more like it came from my brain.