

August 13, 2026

Write for people

Something I've noticed recently is that I am no longer able to read and parse PR titles and descriptions across the internet. I don't think I stopped being able to read or understand code, but now that we are automatically generating more and more explanatory artifacts instead of writing them, navigating codebases as a human is becoming an exercise in futility.

I'm not the only one:



Brian Scanlan
@brian_scanlan

X.com

Counts of "load-bearing" references in pull requests in the intercom rails monolith:

2011–2023: 0

2024: 1

2025: 1

2026:

Jan: 0

Feb: 0

March: 2

April: 36

May: 63

June: 122

July: 188

3:00 PM · 8/12/26 · **26K** Views



14



33



517



26



We are already constantly drowning in tech jargon, much of which comes to us presented through breathless news headlines that make us feel bad for not knowing what the term was five seconds ago.

Tokens are just groups of characters that make up words and sentences. A fully agentic workflow is just [a model working in a loop with external tool calls. A tool is a program the model can call, like](#)

[bash](#). Sandboxes are (in theory, not so much in practice lately) just compute environments that provide OS-level isolation and security. Post-training is just the process of taking a raw model trained to predict the next tokens (group of characters) and continuing to training it towards a specific objective. An objective is just a human goal like summarization, or solving math problems. Reinforcement learning from human feedback is just post-training a model by having humans rank the model's answers and continuing to train the model to produce the best answers.

Some jargon is a good, necessary shortcut between people in the same field at the same level of understanding.

But where it starts to get dicey is when you have both jargon and long text explanations being generated by machines, for people at different levels of understanding. "Bumped dependencies" is good. Everyone knows we are updating packages. "Performed a scheduled dependency refresh as part of ongoing maintenance practices. Minor and patch-level version bumps were applied across the dependency graph, including transitive dependencies where applicable" is bad.

You didn't read the thing when you generated it, I won't read it when I'm reading it. We are generating adjectives and adverbs and throwing them at each other. The output is the result of a model trained on the entire corpus of technical internet documentation with the training objective being completing a sentence but not with the objective of human understanding.

We should strive for simplicity. If we can't, it could be because we don't understand the problem ourselves yet enough to compress it. That's cool. Understanding things is hard. [Naming things is hard, for the same reason](#). But it should be a little harder. Like 10% harder.

We should be [more tired than the model](#), or at least a little less verbose.

Don't miss what's next. Subscribe to Normcore Tech:

Email *

you@example.com

SUBSCRIBE

Running local models is good now [Older →](#)



Powered by [Buttondown](#), the easiest way to start and grow your newsletter.