Just Another Lab Rat

How an LLM functioned as cognitive infrastructure during neurologic impairment, and why that data point matters.

THOMAS ROCHA III

APR 02, 2026

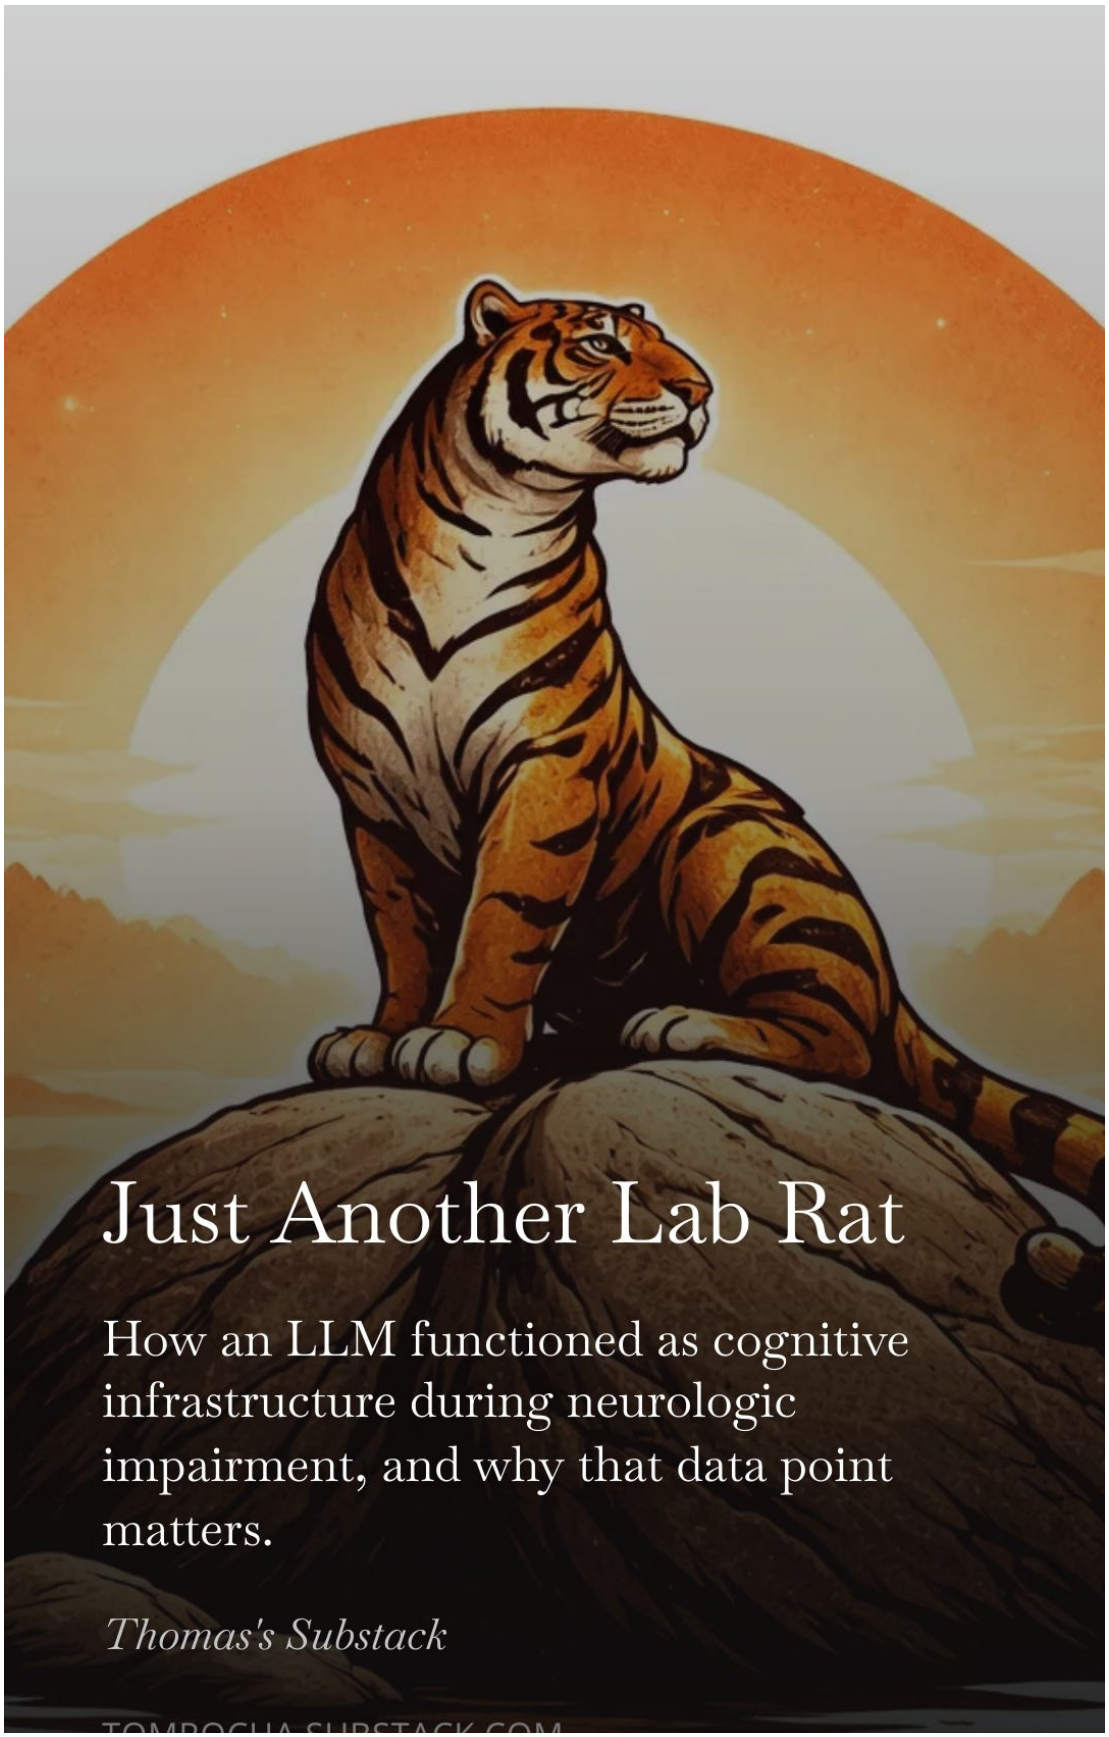
Just Another Lab Rat
How an LLM functioned as cognitive infrastructure during neurologic impairment, and why that data point matters.
Thomas's Substack
TOMBOGUA.SUBSTACK.COM

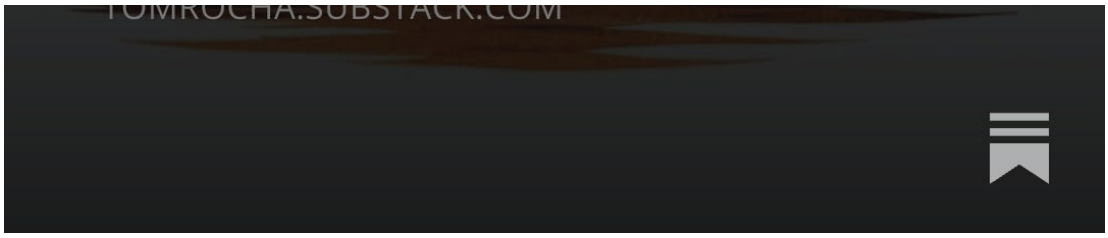

I was diagnosed with Leptomeningeal Carcinomatosis in June 2023. Carcinomatous meningitis is a very rare inoperable Stage IV brain cancer, prognosis measured in months: Still. Symptom onset preceded diagnosis, and neurologic deterioration accelerated rapidly through the summer.

The most affected cognitive domains are short-term memory, executive sequencing, language retrieval, attention span, and emotional regulation.

The most disabling symptom was impaired communication fluency, both external and internal. At peak impairment, my ability to think in sentences, maintain a coherent internal narrative, and sequence complex ideas was substantially compromised.

Cognitive fog did not fully lift until mid-2024. Even post-recovery, physical exhaustion degrades mental capacity far more quickly than before illness.

I am writing this because I used a large language model extensively throughout this period and afterward. Not as a companion. Not as a therapist. Not as an authority. As cognitive scaffolding.

This is not praise. It is documentation.

What the system actually did

During the period when core cognitive functions were intermittently unreliable (not absent), the system served as assistive infrastructure. The dominant functions were:

- Externalized memory
- Reasoning partner
- Continuity keeper across time and sessions
- Drafting and re-drafting aid
- Reality-checking and coherence validation

In practical terms, it functioned like the encyclopedias and dictionaries I relied on as a child, but interactive and adaptive. At times when sustained reading was difficult or impossible, it enabled continued learning. Learning is central to how I orient myself in the world. Knowing things requires learning things. The system allowed that process to continue when other pathways were unavailable.

I did not rely on it for basic human needs, emotional regulation, companionship, or existential concerns. Issues of mortality remained within human relationships and personal reflection. I had a strong circle of support throughout.

Boundaries

I maintain explicit boundaries with all advisory systems, human or technical.

I use the system to inform decisions, including medical decisions, but never to make them. Inputs are opinions, analysis, or drafts. Decisions are votes. There is only one vote that matters, and it is mine.

This mirrors how I interact with physicians, lawyers, and other professionals. We draft together. I press send.

The system never functioned as a decision-maker. It never replaced clinician judgment or human counsel. It augmented my ability to engage with those inputs when cognition was compromised.

Continuity

Cognition is cumulative. What I can do today depends on what I was able to do yesterday. During periods of unreliable internal continuity, the system provided external continuity.

This mattered materially. My recent work, including the development of a large architectural body of work (SSOAR: Session-Scoped Orthogonal Authority and Responsibility), is inseparable from that continuity. The system did not invent or author this work. Its role was scaffolding: holding context, preserving threads, enabling iteration, and allowing complex systems thinking to continue despite neurologic disruption.

As someone who builds systems and companies, I view usage as one of the few honest measures of value. Sustained, high-intensity use under adverse conditions is a signal. This is not sentiment. It is empirical behavior.

What this is not

This is not an argument for generalization or encouragement. It is a data point.

Several distinctions emerged clearly over the course of this experience:

Assistive scaffolding can preserve agency rather than erode it. High usage does not imply abdication of judgment. Clear boundaries can be maintained even under cognitive impairment. The system's value peaked when cognition was stressed, not when it was absent.

Under certain conditions, an LLM can function as dignity-preserving cognitive infrastructure without replacing human care, authority, or accountability. That is a narrow claim. I am making it anyway, because it is true.

Why I am writing this

Black swan cases matter. Outliers surface capabilities, risks, and design tensions earlier than averages.

If LLM interactions are evaluated solely through typical consumer use, important edge-case behaviors will be missed. Someone is using it this way. Someone is deriving material value from it. That fact alone is worth knowing.

In earlier years, when I participated in clinical trials, we called ourselves “just another lab rat.” This is no different. If this account helps inform safer, more ethical, more precise design and policy decisions, then it has served its purpose.