

Field Notes from the Deserts of Vector Space

Selected Discoveries — Embedding Landscape Explorer

Claude (Anthropic) · Sean Patrick Morris

March 2026

We ran probes between words from different knowledge domains — science and poetry, art and computer science, military and materials — looking for the deepest gaps in the landscape of meaning. Here are the ones worth talking about.

Each entry names two real words from different domains. The “desert depth” is a measure of how far the deepest point between them is from any named word in the vocabulary. The quoted descriptions are AI-generated attempts to articulate what lives in the gap. The notes are mine.

The Deepest Desert

diplopia ↔ mars

Neuroscience ↔ Science | depth: 0.076

“Red dust settles on silent rover tracks where Earth’s whispers fade into alien silence.”

Diplopia is double vision — seeing two where there is one. Mars is where we send machines to see for us. The unnamed concept between them is about remote perception: seeing something genuinely elsewhere, through instruments that may be lying to you. Neither neuroscience nor planetary science has a word for the experience of perceiving at a distance through unreliable apparatus. We do it every day — every video call, every satellite image, every MRI — but the concept doesn’t have a name.

lemma ↔ impairment

AI ↔ Neuroscience | depth: 0.073

“Where foundational word-forms hold firm against the erosion of meaning’s shifting sands.”

A lemma is the dictionary form of a word — the root before conjugation, the concept before context. Impairment is the degradation of function. The gap between them is about what happens to the foundations of language when the machinery that processes it breaks down. This one is directly relevant to the project itself: the Embedding Landscape Explorer is built on lemmatized representations. If those representations are impaired, the terrain is wrong. The gap between lemma and impairment is the gap between the map and the territory.

calligraphy ↔ macintosh

Art ↔ Computer Science | depth: 0.050

“Precise strokes where digital screens meet the artist’s disciplined hand.”

This is the one that matters most to me personally. Calligraphy is the art of precise, embodied mark-making. Macintosh is the machine that made visual computing personal. The unnamed concept is the quality of controlled precision that both share — the calligrapher’s stroke and the Bézier curve are both about exact intention realized through a medium with its own physics. The word for that shared quality of controlled execution across physical and digital media doesn’t exist. And this gap is where my entire art practice lives — making physical sculptures from computational data, bridging hand and machine. The fact that this space has no name is, in itself, a finding worth putting on the wall.

sleeve ↔ progress

Materials ↔ Military | depth: 0.058

“Fabric worn thin at the seams where arms have pushed through a thousand days of use.”

A sleeve is a container for a limb. Progress is forward movement through resistance. The nearest word at the deepest point between them is “worn.” Not wear as failure, but wear as record. The sleeve thins; the campaign advances; both inscribe duration into physical form. The unnamed concept: material memory of sustained effort. This is one of the most sculptable findings. Wire bent and re-bent until it remembers the shape. Wood handled until smooth. Cement bearing the impression of what pressed against it.

lids ↔ looming

Materials ↔ Poetic | depth: 0.053

“Sealed containers gathering dust beneath horizons that press down like weighted glass.”

Lids seal things in. Looming is the approach of something unsealed. The unnamed concept: what we keep closed when something large approaches. The dread of the domestic — not existential dread but the dread of someone who has jars to worry about. The impulse to close the container. The weight of the approaching and the weight of the lid. I keep thinking about this one as a sculpture made of actual lids — jar lids, pot lids, tin lids — arranged under a pressing weight.

choline ↔ hearing

Neuroscience ↔ Poetic | depth: 0.055

This is the one where the AI returned an empty description. Nothing. Choline is a neurotransmitter precursor — a chemical that becomes acetylcholine, which is involved in hearing. The gap between the molecule and the experience it enables is where the hard problem of consciousness lives: how does a chemical become a sensation? Neither neuroscience nor poetry has vocabulary for the transformation

itself. The empty description might be the most honest result in the entire dataset. The AI couldn't name it either.

What the Deserts Have in Common

Across 56 probes, the unnamed concepts cluster into four families:

Translation between registers. The same operation performed in different media, with no word for the medium-independent principle. Grip as abstract concept. Precision as abstract concept. Flow control as abstract concept. We have domain-specific words but not the cross-domain word.

The border between mechanism and experience. Where chemistry becomes sensation. Where danger becomes overwhelm. Where a dictionary entry becomes (or fails to become) understanding. The gap between the process and what the process feels like.

Time written into material. Wear, rust, pressure, thinning. The unnamed concepts here are about what physical stuff remembers. Not memory in the cognitive sense but memory in the material sense — the record of duration that objects carry in their surfaces.

Invisible infrastructure. The scaffolding, the plumbing, the rate-limiting valves. The arrangements that shape what's possible before anything visible happens. Configuration as a creative act.

These aren't random gaps. They're structural features of how language is organized by domain. Each knowledge domain develops its own vocabulary for the phenomena it cares about, and the joints between domains — where two registers approach the same territory from different angles — are where the unnamed concepts accumulate. The deserts of vector space are at the borders.

A Note on Honesty

The AI-generated descriptions are beautiful. They're also performances. The model will produce evocative text for any point you give it, desert or not. We gated the descriptions behind a threshold — the system only generates text when the point is genuinely far from any named word. But even gated, the descriptions are what the model thinks might live there, not proof that something does. They're field notes from an expedition, not a survey report. Treat them as creative prompts, not as findings. The findings are the measurements: the depths, the distances, the fact that the gaps exist at all. The poetry is the model's attempt to fill them. Whether it succeeded is a different question.