Semantic Meaning Without Syntactic Representation

A Critical Framework for the Embedding Landscape Explorer

Claude (Anthropic) · Sean Patrick Morris

March 2026

The Short Version

There are ideas that exist but don't have words. Not in the "I can't quite express how I feel" sense, but in a measurable, computational, point-to-it-on-a-map sense. This project finds them.

What We Did

Language models — the things that power ChatGPT and Claude and everything else you've been arguing about at dinner — work by turning words into numbers. Specifically, they turn each word into a long list of numbers (a "vector") that captures what the word means based on how it's used across billions of sentences. Words that are used in similar contexts end up with similar numbers. "King" and "queen" are close together. "King" and "luminescence" are far apart.

The important thing is that this space of numbers is continuous. There's a vector for "king" and a vector for "luminescence," and there are vectors for every point in between. Those in-between points don't correspond to any word. But the model treats them as meaningful — they have neighbours, they have properties, they sit in a region of the space that encodes something. There's a concept there. It just doesn't have a name.

So we built a tool that turns this mathematical space into a landscape — actual 3D terrain you can walk through with keyboard controls, like a video game set in the geography of meaning. High ground is where lots of words cluster together. Valleys are where meaning thins out. And deserts are the places where the space encodes something but no word lives there.

Then we went looking for what lives in the deserts.

What We Found

We loaded vocabulary from nine different knowledge domains — science, poetry, materials, AI, neuroscience, computer science, art, military, and a big bucket of shared words that appear across multiple domains. 8,735 words total. We built the terrain. We mapped the deserts. Then we ran probes between words from different domains and looked for the deepest gaps.

The first thing we found is that the deserts are shallow. With nearly nine thousand words in the vocabulary, most of the space is spoken for. The deepest desert we found is only 7.6% of the maximum possible

distance from a named word. Language has colonized this territory pretty thoroughly. The deserts aren't Saharan expanses — they're hairline cracks between tiles. Narrow fractures where two bodies of meaning almost touch but don't quite meet.

The second thing we found is that the most interesting cracks are between domains. Not between "maybe" and "perhaps" — those are basically the same word and there's nothing between them. But between "calligraphy" (art) and "macintosh" (computer science), there's a gap. Between "sleeve" (materials) and "progress" (military), there's a gap. Between "lemma" (the dictionary form of a word, from AI) and "impairment" (neuroscience), there's a gap.

These gaps aren't random. They have structure. When we asked an AI to describe what might live in each gap, it produced things like:

> *"Fabric worn thin at the seams where arms have pushed through a thousand days of use."*

That's the gap between "sleeve" and "progress." The unnamed concept: duration made visible through material wear. Not wear as failure, but wear as record. The sleeve thins; it doesn't break. Progress accumulates; it doesn't arrive. There's a concept there about the slow inscription of effort into physical form, and English doesn't have a word for it.

> *"Precise strokes where digital screens meet the artist's disciplined hand."*

That's the gap between "calligraphy" and "macintosh." The unnamed concept: the quality of controlled precision that a calligrapher's brush and a vector path share. Not "digital art" — that phrase exists. The specific quality of disciplined execution across physical and digital media. No word for it.

So What?

Three things.

For language: Linguists have argued for a century about whether the limits of our vocabulary are the limits of our conceptual world. This project doesn't settle that argument, but it provides a new kind of evidence. The mathematical structure of meaning — as learned from how people actually use words — extends beyond the vocabulary available to name it. The concepts exist in the data. The words don't.

For AI: The unnamed regions of embedding space are where language models generate novel combinations, hallucinate, and produce outputs that surprise their operators. Understanding the shape of what's named and what isn't is relevant to interpretability — figuring out what these models "know" versus what they're making up. The deserts are where the guardrails aren't.

For art: I make physical sculptures from computational data. The sculpture is the data, not an illustration of it. When a cross-domain probe identifies a desert between two concepts, the resulting sculpture materializes that specific absence. The physical object fills the lexical desert. It becomes the syntactic

representation that was missing. That's a stronger claim than "data visualization" or "AI art." The AI is the sensor, not the artist. The work is the act of naming — through physical form — what the embedding space reveals and language has not yet said.

The Thesis

Stated formally: semantic meaning exists in embedding space that lacks syntactic representation in natural language.

Stated informally: there are real ideas hiding in the cracks between different ways of talking about the world, and if you build the right tools, you can find them, measure them, and make things out of them.

Stated as a painting title: are there deserts in vector space?

Yes. Shallow ones. Full of things that don't have names yet.

Caveats

This project was built in about a week with Claude Code, a head cold, and a vocabulary pulled from an old shoebox. The taxonomy is rough. The "desert depth" measurements are relative to a specific vocabulary at a specific moment — add more words and the deserts shift. The AI-generated descriptions of unnamed concepts are creative prompts, not ground truth. They're what the model thinks might live there, which is interesting but not authoritative.

The embedding space itself (GloVe 300d) is frozen in time — trained on a corpus that stopped growing years ago. A different model would produce a different terrain. The deserts are properties of the relationship between the vocabulary and the model, not properties of language itself. Whether the unnamed concepts are genuinely "real" or just artefacts of the geometry is a philosophical question this project raises but does not settle.

But the cracks are there. You can measure them. And some of what lives in them is worth picking up and looking at.