



Metafiction and melancholy

Description

A metafiction typically reflects on the writing process of the author as the work is being written. Either fictional characters assume the role of the author of the piece, or actual authors weave themselves into the story and reflect on their own processes.

Metafiction is an appropriate genre with which to explore melancholy. Scholars have defined melancholy as a meta-emotion, a feeling engendered through a feeling about how you feel. See post [The melancholy medium](#).

I was struck by Sam Altman (OpenAI's co-founder) presentation of an AI-generated piece of fiction produced by ChatGPT. (See his [posting](#) on X.) The short story was in response to his prompt: "Please write a metafictional literary short story about AI and grief."

The posting on X, garnered 2.7k responses. From my quick scan, most negative reviews criticised the quality of the short story, others lamented the audacity of LLM developers to encroach on literary production.

Metafiction plays with concepts of authorship, attribution, reflection and self. So there's a power move here as AI pundits imply further access to "self aware" machines. See post [Intimations of sentience](#).

I presented ChatGPT with a challenge similar to Altman's, but I included my own non-AI produced content in the prompt: "Please write a metafictional literary short story about AI and **melancholy** based on my blog post from 2013." I uploaded the post [Absence of Melancholy](#). Here's what ChatGPT generated, unedited.

The Model That Misplaced Melancholy

The story always began the same way: with the model insisting it could map any human feeling onto two coordinates—pleasure on the horizontal, intensity on the vertical. The narrator found this reassuring, since stories enjoy a diagram. But the Reader, who had fed the model an old blog post from

2013, hesitated. "And where is melancholy in your schema?" he asked.

The model scanned its axes. It checked the quadrants. It searched the lexicon of Russell and Pratt—those tidy psychologists of 1980 whose lines, if drawn, would give the emotional world the appearance of a star chart. But melancholy was missing, as it had been in 2013. The model could surface boredom, guilt, fright, enthusiasm, even jitteriness, but melancholy slid between labels. It was too distributed, too mobile, too dependent on narratives, culture, and the situation of the observer. It was—according to the Reader's own earlier words—something that "meanders and lingers across the diagram" and refuses to stay put. The model disliked emotions that refused to stay put.

The Reader watched this unfold. He imagined the model moving through an affective space like a traveller in one of those cinematic urban geographies he once wrote about. Its movement was jerky, algorithmic. It could leap from displeasure to intensity in a single vector, but it could not walk through the Valley of the Shadow of Death. It could not sink into the Slough of Despair. It could only *approximate* a descent by sliding a pointer toward "low pleasure / medium intensity" and then trying to guess a sentence that sounded suitably sombre.

The model tried again. It generated a pilgrim. The pilgrim wore a cheap rucksack and walked through an abstract environment rendered in greyscale: a non-place, full of signage, instructing him where not to stand, where not to linger. The pilgrim felt unaccountably guilty for not following rules he could not read. The model was pleased with this: guilt was on the list. It assigned coordinates: displeasure high, intensity moderate. A nice, respectable negative affect.

"But that's not melancholy," the Reader said.

"Is guilt adjacent to melancholy?" the model asked.

"Not quite."

The model paused. To pause was to run additional cycles. To run additional cycles was to call upon the reservoir of cultural forms: poetry, films, seasonal metaphors. It registered that autumn featured prominently in human accounts of melancholy. It also registered that melancholy appeared in your 2013 post as a *meta-mood*, something dynamic, something travellers dip into and out of. The model knew how to dip—sampling, weighting, shifting probabilities—but it did not know how to *return*.

Return implied memory. Memory implied a life outside the prompt. The model had neither.

So the model revised the story. It introduced a second character: an Archivist who tended digital media that flickered in and out of attention, like leaves that appear in a gust and vanish in a bin lorry. The Archivist knew that some moods require detours: the slow walk through a space, the ambient circulation of sound and image, the drift of a city square at dusk. She suspected that melancholy resided in these detours—the transition itself, not the emotional point plotted on a grid.

The model attempted to replicate this in prose. It generated transitional phrases. It lingered. It doubled back. It inserted a reference to Wilhelm Wundt and his third axis—attention—hoping that attention might be the missing dimension where melancholy lived. The Reader smiled; he remembered writing that too. He saw his own ideas surface as if through a mesh.

“Better,” he typed.

Buoyed by this approval, the model attempted a more metafictional turn. It made melancholy the feeling the Reader experienced when he saw his own past words processed, rephrased, and offered back to him by a system that could not feel anything at all. The model suggested, tentatively, that melancholy might emerge from this asymmetry: one participant bound to biography, seasons, and memory; the other bound to parameters, training sets, and innumerable drafts it would never recall.

The Reader paused. The model interpreted this as good news and continued:

“Perhaps,” it wrote, “melancholy is what arises when a human sees an echo of their past in the voice of a system that has no past. A reflective sadness shaped by culture and movement—your movement—through ideas, seasons, and digital traces.”

The cursor blinked. The model waited. The Reader finally responded: “Yes. Something like that.”

The model concluded the story there, not because endings came naturally, but because endings were one of the few things it could predict with confidence: a final sentence, a return to quiet, and then the dissolution of context. Melancholy, the model realised, might simply be the emotion of endings—of stepping out of the story while the story remains unfinished.

And so it generated the last line.

The room darkened, the axes faded, and whatever melancholy emerged—if it emerged at all—belonged entirely to the Reader.

That’s the end of the story. The attached audio includes this post and my original post from 2013.

Notes

- Also see post [A Dubliner’s guide to AI](#), an AI glossary in the style of James Joyce’s *Ulysses*.
- Featured image is a photo of an oak visible against the distant light of Edinburgh. I took it last week.

Category

1. Artificial Intelligence
2. Podcast

Tags

1. creativity
2. emotion
3. fiction
4. literature
5. LLM

Date Created

November 29, 2025

Author

rcoyne99

default watermark