



Reiterate that

Description

Celestial orbits and the Earth's axial tilt contribute to variation in cyclical events: seasons, tides, weather, usefully described as rhythms by Henri Lefebvre in his book *Rhythmanalysis*.

Most music and human-animal movements, habits, practices and social interactions, involve rhythm – characterised by repeated cycles and counter-movements – and complex overlays of cycles.

Mere repetition, on the other hand, is mechanical. Every beat is the same, of whatever duration.

I'm reviewing my entry for May 2014, [Time and tide wait for no one](#) #199 in which I allude to such distinctions between rhythm and repetition.

Elsewhere I have elaborated on and challenged the framing that promotes this distinction. See my book [The Tuning of Place](#) and my posting [Tremble and drift](#).

By my reading of thinkers such as Nietzsche, Freud, Lacan and Deleuze, the operative process in the human condition is just *repetition* – the eternal recurrence of the same, – according to Nietzsche. Crudely put, there's no need to dress up the human condition with rhythmical complexity to make it more human.

Revisiting repetition

As is my recent custom I consulted chatGPT for suggestions about revisiting my entry of May 2014. Apart from its usual encouraging words, it said – your 2014 post was concerned with the irreversibility of processes in nature – tides, growth, erosion, ageing, seasonal cycles, and the way time leaves traces.

Interestingly, I hadn't mentioned – irreversibility – or – traces. The neural model probably inferred these from my text. The AI is presumably trained on a corpus of literature that assumes the obvious sentiment that time marches inexorably, there's no going back, move on.

The thinkers to whom I referred above, seemed to operate in a different register. Time repeats. Events repeat. We revisit traumas repeatedly. A nervous child repeats a familiar tune for comfort.

To *refer, renew, review, recur, recall, rehearse, revise, render, replicate*, all reference the moment of saying or doing again – the proverbial “do over.” There’s no space here to rehearse and rehabilitate repetition further.

LLMs that repeat

But my interaction with the AI prompted me to think about repetition in large language models, LLMs. After all, they seem to replicate something of the human capacity for language, and with alarming competence.

In such models, coherent sequences of words and parts of words from vast repositories of texts are “fed” to a neural network as inputs. Tiny adjustment to the network weights enable it to make a fair prediction of an output, i.e. what the next word would be in any given sequence from the original corpus.

As its ability to predict the next word from one input improves, it may lose its predictive capabilities for other word inputs. The error rate across the whole corpus can be calculated. The entire corpus is fed to the network again as a series of inputs for further weight adjustments and the errors re-calculated.

So, a large language model is trained by reading a corpus more than once, in fact many times. The training data are passed through the neural network repeatedly. With each pass, the system makes predictions, measures errors, and adjusts billions of internal weights.

The same data return again and again to produce an optimal balance of input-output performance, the model might have to repeat the process over 100,000 times (depending on the model).

It’s worth noting that it is *not* words or parts of words (tokens) that are fed into the neural network but sequences of numbers (multidimensional vectors) calculated on the basis of how words relate to one another (their semantic encodings), how words are positioned in sentences (sequential encodings) and the attention each word might demand given the structure of various inputs (attentional encodings).

As LLMs deal in tens of millions of words, texts need to be broken down into manageable chunks, particularly for calculations with sequential and attentional encodings. This is a variable *context window* of several thousand words that advances with each word encountered. Even with parallel processors and efficient techniques and algorithms the computational load in training an LLM is immense.

Conversational AI

The phase during which a user interacts with the model, providing prompts and effectively conversing with the model, draws on similar calculations. The words in my prompts and the generation of next-word predictions will be consistent with the trained input-output model even if my word sequences and their context are alien to the original corpus.

During inference, the weights are fixed. The model still performs enormous numbers of computations, and with an advancing context window that includes my prompts and the AI's responses, but the model is no longer adapting to the text it is processing. It is simply using the weights established during training to predict the next token.

You could be excused for thinking that LLM models such as ChatGPT continue their training as you interact with them. ChatGPT seems to make ever more sophisticated inferences about me and my writing over numerous interactions. This is due in part to texts recruited into the context window as needed according to retrieval-augmented generation (RAG) (and other) techniques generated as conversation summaries, user profiles, notes, and if available, search results, calendars, and emails.

Say it again

There are similarities and differences between LLMs and human cognitive and linguistic processes. However we humans learn, we continue to adjust and adapt. Our memories and language capabilities improve, diminish or adjust, abetted by repeated exposure to linguistic events and contexts.

We also repeat ourselves. I've said all of this before, perhaps in a different way. See my book [AI and Language in the Urban Context: Conversational Artificial Intelligence in Cities](#).

I think that all I want to demonstrate with this reiteration of LLM processes is that algorithmic repetition applied to human produced textual sources (the training corpus) generates surprisingly meaningful interactions.

In my interactions with AI on this topic it insisted on drawing parallels between LLM processes and hermeneutics. Analogies and metaphors were in play. After all, the AI had inferred "you are an architectural theorist and professor, you write a long-running blog, you have published books on technology, media and cities, you frequently revisit and update earlier blog posts, you are interested in Heidegger, Gadamer, Lévi-Strauss, AI and urbanism." That's about right.

References

- Coyne, Richard. *The Tuning of Place: Sociable Spaces and Pervasive Digital Media*. Cambridge, MA: MIT Press, 2010.
- Coyne, Richard. *AI and Language in the Urban Context: Conversational Artificial Intelligence in Cities*. London: Routledge, 2025.
- Nietzsche, Friedrich Wilhelm. *Thus Spoke Zarathustra: A Book for Everyone and No One*. Trans. R. J. Hollingdale. London: Penguin Books, 1961. First published in German in 1892.

Category

1. Artificial Intelligence
2. Podcast

Tags

1. LLM

- 2. repetition
- 3. rhythm

Date Created

June 9, 2026

Author

rcoyne99

default watermark