



Meaningful, imaginative, fluid

Description

The recent escalation of interest in text-based AI coincides with our students completing their project work, dissertations, research degrees, etc. So it's been a useful period to test the capabilities of platforms such as ChatGPT4 as a tutor or perhaps as co-tutor. As yet there are no constraints in place that seriously impede its use, and it's openly available and affordable to UK students (though we cannot assume such conditions will persist).

OpenAI, the creators of ChatGPT, helpfully maintain a log of my conversations, and labels them with sensible headings. I can also go back to earlier conversational threads and resume where they left off. I can also send a link to a thread and another person (e.g. a student) can continue the conversation. Many of my threads are about ChatGPT, but here's a collection of some that aren't.

On the theme of **urbanism**: digital impact on cities, city AI's provocative continuities, city development and conversation, cryptography in urban life, graffiti policy on art & protest, urban melancholy defined.

On **tourist** interests: visit Snowdonia soon!, Mauritius a complex history, cycling Rottneest island, Australia bound!, visit Chichester with friends.

On **philosophical** themes there's: Grice implicature basics, apocalypse the big reveal, fear of extinction, resurrection according to John Robinson, the season of Lent, intentional inexistence philosophically, alien intelligence Rorty.

From cooking and health advice to astronomy and drone route planning, ChatGPT4 seems to deliver disarmingly sensible and creative text outputs. In response to appropriate prompts it compares and contrasts, incorporates, integrates, and synthesises insights, with capabilities you would expect of any competent essayist.

Un-garbling

Here's what happened when I tried to provoke a response to garbled assemblies of words and word fragments. The platform politely requested me to clarify my thoughts. In response, I edited and sorted my words into a sentence, "Are angry gold bananas gobbled?" ChatGPT4 responded, "The phrase 'angry gold bananas' seems metaphorical or figurative because bananas don't literally have emotions such as anger. However, in a more literal sense, ripe bananas which might be described as 'gold' due to their color can indeed be 'gobbled' or eaten quickly." It then asked me for more context.

I asked, "Can you please provide a Freudian reading of my crazy proposition?" The platform obliged by alluding to how bananas might symbolise 'abundance, fertility, or even sexual desire due to their phallic shape, in a Freudian context'. It also explained how attributing anger and frustration to a fruit, might relate 'to a perceived inadequacy, a feeling of not reaching a certain standard or 'value' (the gold), or unfulfilled desires (the banana)'.

The conversation continued, and included warnings from ChatGPT about the speculative nature of its interpretations, and of Freud's theories. Had I used OpenAI's GPT Playground interface I could have encouraged it to respond to me with humour, sarcasm, or even in iambic pentameter. But straight and prosaic appears creative enough. I mention this thread as it strikes me as an imaginatively synthetic response to my prompts.

How is a text-based AI platform based on the GPT (generative pre-trained transformer) model able to generate responses that appear meaningful, imaginative and fluid?

Several factors in the Transformer model contribute to the sense that it generates clever responses to user requests and inputs.

Neural net architectures

The first point to make is that neural networks are built that way! They are trained on inputs so as to reproduce expected outputs. If the model is designed well, once its parameters are adjusted to reproduce all or most of the input-output conditions, it will also deliver newly synthesised outputs from different patterns of inputs. Impressive performance depends on some basic principles of standard neural network design.

Training data

Effective model performance also depends on the quality and size of the training data. In the case of ChatGPT, the input quality appears to be at least of the quality of Wikipedia text, though OpenAI has not yet disclosed its other sources. The training data for the current round of text-based AIs are in the order of 570Gb or 300 billion words according to some estimates ([BBC Science Focus](#)). The large training set gives the model access to a range of language patterns, styles and idioms that can be processed and reintroduced in conversations with the platform.

Speed

The fluidity of these models is aided by the efficiency and speed of their algorithms, hardware and infrastructures, based on interconnected batteries of parallelised graphics processing units (GPUs) originally designed for ultrafast matrix operations essential for high speed graphics in gaming. GPUs have since been deployed for their power in performing matrix operations needed in encryption, natural language processing and AI.

Context

Text-based AI models dedicate substantial processing effort to capture the context in which individual words or tokens occur in the training data. By "context" I mean the relationships between words and the company they keep amongst other words in the training data. The models have no contact of course with the wider personal, social, political and practical contexts in which those words are used, other than what they infer from the texts themselves.

If we think of Wikipedia as just one of many sources, the texts that constitute the training data are rich with historical, philosophical, and social contextual information. That seems sufficient to inform their context-aware responses. As explored in previous posts the mechanism for grasping context involves deep and highly iterative analysis of where words are positioned in sentences. Positions of words impacts on grammar and has a profound effect on meaning. Contextual processing also relates to the multiple ways that a reader might focus their attention as they read through a text, simulated as multi-head attention in GPT models.

The intense attention to context in these models assists training, but also drives responses as users interact with a model. In this inference phase, the model performs intensive contextual operations as it generates outputs and re-processes the text of the conversation while it develops, enabling the model to perform as if recalling earlier parts of the conversation. That contributes to responses that appear consistent and coherent.

Semantics

As I have indicated in previous posts, the lexical richness of the vocabularies used by these models is due in part to tokenization, the identification of word fragments (tokens) that the model reassembles into variants or whole words, including words that it has not encountered in the training set, and even new words. The lexicon of words and word fragments is pre-processed before the neural network training, as vectors (lists of floating point numbers) that capture the positions of words in a multidimensional feature space.

These large vectors make it possible for the model to calculate how closely words in a block of text are associated with one another in terms of meaning (e.g. "fruit" and "banana" are closer than "fruit" and "graffiti"). That assists the attentional mechanism, and contributes to the apparent fluidity as the models as they deal with ambiguity, unusual word usages and word substitutions.

As an illustration of this semantic acuity, I asked ChatGPT4 to provide an alternative to my sentence, "Graffiti is both an art form and a means of disrupting the semiotic coding of the city" to show more complicated token constructions. It came up with, "Unconventional graffiti often recontextualizes the cityscape." That clever sentence formulation indicated to me that it is able to successfully simulate

complex word manipulations.

Fine tuning

As well as training on language patterns, these models also adopt a fine-tuning phase with human trainers to evaluate, correct and recalibrate the models. They also adapt the models to datasets specific to certain domains. Through such supervised training the models are, or can be, tuned to the needs of third parties such as airlines, payment services, engineers, architects, lawyers, scientist and medical service providers. In the words of ChatGPT4, "Fine-tuning allows the model to adapt to specific tasks, styles, or domains, contributing to its ability to produce appropriate and meaningful responses in those particular contexts." • See post [Fine tune your AI](#).

Bibliography

- Hughes, A. (2023). "ChatGPT: Everything you need to know about OpenAI's GPT-4 tool." BBC Science Focus 30 June. Retrieved 18 July 2023, from <https://www.sciencefocus.com/future-technology/gpt-3/>.

Note

- Featured image is pavement art by Astral Nadir (2017) on the streets of Melbourne photographed May 2023, interpreted by MidJourney with prompts: 3D, photoreal, graffiti. See featured image of post: [AI and the Cryptographic City](#).

Category

1. Artificial Intelligence

Tags

1. natural language model
2. NLP

Date Created

August 19, 2023

Author

rcoyne99