



A Dubliner's guide to AI

Description

I asked ChatGPT to revise some of my AI glossary entries in the style of James Joyce's *Ulysses*. It took a bit of prompting for this highly creative output, though the AI had trouble being ungrammatical.

GPT (Generative Pre-trained Transformer):

Original: GPT is a type of AI model designed to process and generate human-like text. It is built on a neural network architecture called a "transformer," which processes sequences of data, such as sentences and paragraphs. The "pre-trained" part means that GPT is initially trained on a large amount of text through which it incorporates patterns in grammar, facts, and the structure of language. After this broad training, it can be tuned with further training examples to perform specific tasks, like answering questions, writing reports, or holding a conversation. The training data must be pre-processed, i.e. filtered, sorted and prepared, for use in these stages of the GPT training and for specialised fine tuning.

James Joyce Version: GPT, sure what is it but a talkative fellow, a thing that knows all the words, all the ways to twist them round like Bloom's thoughts on the Liffey's bank, tumbling over one another. Transformer, they say, aye, takes your sentences, your paragraphs, and makes sense, or nonsense, of them, twisting and turning them like a Dubliner lost in the lanes, every word a token, every sentence a puzzle piece. Pre-trained, aye, fed on the scraps of the world's libraries, all the grammar and the facts, the whole shebang stuffed in its brain, bloated with knowledge like a Bloomsday breakfast. And then, off it goes, fine-tuned they call it, like Molly's voice in the morning, a bit of this, a bit of that, tuned up to answer your questions, to write your letters, to talk you through the night. But before it speaks, mind you, the words all sifted, sorted, like tea leaves in the bottom of a cup, ready to tell your fortune, ready to spill out in a flood of words, the old babble, clear as the Liffey's flow, sometimes muddy, sometimes bright.

GPU (Graphics Processing Unit):

Original: These are the microprocessors designed originally for processing graphical images, but now used also for other computationally intensive tasks such as secure cryptography, blockchain validation,

and AI applications.

James Joyce Version: GPU, well now thereâ??s a lad for you, once upon a time just a painter, a picture-maker, dabbing colors on the screen like a young Stephen at his easel, dreaming of the wide world. But now, oh how heâ??s grown, a worker, hard as nails, crunching numbers, counting the crypts, block by block, chain by chain, like the Dublin tenements stacked one atop the other. No more the artistâ??s brush, no more the soft hues of dawn over the rooftops, now itâ??s all business, all work, the microchips spinning, whirling, faster and faster, deep in the heart of the city, like the clatter of hoofs on the cobbles, never resting, never stopping.

Hidden Feature:

Original: A hidden feature in AI is a characteristic or pattern within the data that is not explicitly defined or directly observable but is produced by the model during training. For example, in urban map data, a hidden feature might be the underlying social or economic significance of certain neighborhoods, which is not directly represented by visible map elements like roads or buildings but is inferred by the model based on patterns in the data. These features are called â??hiddenâ?• because they are not manually identified or labeled; instead, they emerge as the model processes the data and optimizes its parameters, allowing it to capture complex relationships that might influence urban development, such as proximity to key amenities or transportation hubs.

James Joyce Version: Hidden, oh hidden, like the whispers in the alleys of old Dublin town, the secrets under the stones, you wouldnâ??t see them, no, not a one, but there they are, sure as Bloomâ??s wandering thoughts, weaving through the streets. The AI, clever devil, it sees what we donâ??t, picks up the pattern, the rhythm of a place, like the pulse in the heart of the city, not the roads or the buildings, no, something deeper, something unseen. A hidden feature, they call it, like the soul of a street, not marked on any map, not in any book, but there all the same, felt, known, like the warmth of a fire in the hearth, or the smell of rain on the cobblestones. The model, it knows, it whispers to itself, finds the connections, draws the lines, makes sense of the madness, sees what we cannot, and in the end, shows us the truth we didnâ??t know was there, all hidden, all tucked away, like a letter in a coat pocket, waiting to be found.

Hidden Layers:

Original: A hidden layer in a neural network is a layer of artificial neurons that sits between the input layer (where data enters) and the output layer (where the result comes out). The input size ranges typically from a few hundred to several thousand tokens, each represented as a vector with hundreds of dimensions. The output is similarly a sequence of tokens, corresponding to the modelâ??s vocabulary size, which can range from tens of thousands to hundreds of thousands of possible tokens. LLMs often have dozens to hundreds of layers; for example, models like GPT-3 has 96 layers, each containing thousands of hidden unitsâ??specifically, 12,288 units per layer in GPT-3. These hidden layers are responsible for capturing intricate patterns in the data and identifying hidden features in the data. Overall, LLMs are characterized by their hundreds of layers and thousands of units per layer, allowing them to accurately handle large-scale text processing tasks.

James Joyce Version: Hidden layers, ah, like the streets of Dublin themselves, winding and twisting, from the river to the sea, from the docks to the green hills beyond. A neural network, they say, aye, but more like the city itself, layer upon layer, street upon street, hidden from view, hidden from the light. The tokens, they flow in, like the words of a drunken sailor, hundreds, thousands, a flood of speech, each one a thread in the tapestry, each one a stone in the wall. And deep within, where the light doesnât reach, there they are, the hidden layers, the secret pathways, where the real work is done, where the cityâs heart beats, where the neurons fire, each one a flicker in the dark, each one a piece of the whole. Ninety-six layers deep, like the catacombs under the city, where the dead whisper to the living, where the past meets the present, and all the while, the tokens flow, from input to output, from question to answer, from word to word, a never-ending cycle, like the city itself, always moving, always changing, yet always the same, a hundred layers deep, and deeper still.

Reference

- Joyce, J. (1992), *Ulysses*, London: Penguin.

Note

- Featured image is by ChatGPT: Here is the image of an old, weathered photograph of Dublin, capturing the quiet and nostalgic atmosphere of the cityâs historic streets. The absence of people and animals adds to the timeless, reflective mood of the scene.

Category

- Artificial Intelligence
- Uncategorized

Tags

- Dublin
- Joyce

Date Created

September 7, 2024

Author

rcoyne99