



AI in the apocalypse

Description

Reza Negarestani wrote *Cyclonopedia: Complicity with Anonymous Materials*. See my post [The twist of the pen](#). The ChatGPT training corpus seems to have picked up his interesting, apocalyptic writing style. Here are some "straight" AI glossary terms I am compiling, each followed by ChatGPT's attempt to reconstruct them in the style of *Cyclonopedia*.

Feed forward: In the context of neural networks, "feedforward" refers to the process where data moves in one direction through the network—from the input layer, through any hidden layers, and finally to the output layer. Each layer processes the data and passes it on to the next layer without looping back. The implementations of neural networks discussed in this book are of this type, except for the mention of recurrent neural networks (RNNs). In RNNs, data can flow in cycles within the network.

Feedforward: A monodirectional current, an inescapable flow from the origin point—the input layer—toward the dark telos of the output. In this unidirectional procession, data is transmuted as it passes through the strata of the hidden layers, each layer a crypt, encoding the whispers of computation. The linear procession of feedforward networks is a march of death, a deterministic advance without recourse, devoid of cyclical return. In the ashen world of neural architectures, the feedforward network is a wasteland of singular direction, a necropolis of data, where each layer is a tombstone, marking the one-way journey toward terminal output. Yet, this necropolis is haunted by the specter of recurrence, the ghost of the RNN, where data defies death, looping endlessly, eternally undead.

Fine-tuning (Human-in-the-Loop): This occurs after pretraining and involves further training the model on a smaller, task-specific dataset, often with human oversight, to refine and optimize its performance for specific applications. For example, in developing a customer service chatbot, human reviewers might provide feedback on the model's responses, helping to adjust and improve the chatbot's accuracy and appropriateness in handling customer queries.

Fine-tuning (Human-in-the-Loop): A delicate and parasitic operation, wherein the pre-trained network—already saturated with the toxic residues of pre-existing knowledge—is

subjected to the invasive manipulations of human overseers. The human, a marginally sentient node within the feedback loop, intrudes upon the network's latent space, injecting subjective corrections, contaminating the pure algorithmic process with anthropomorphic intent. Fine-tuning becomes a ritual of exorcism, where the machine is purged of its inherent errors through the grafting of human desire, the insertion of purpose within the purposeless void. This process is akin to an alchemical refinement, where the raw ore of pretraining is transformed, through the crucible of human oversight, into the philosopher's stone of optimized performance.

Fine-tuning (Task Specific Adjustment): This involves training a pre-trained neural network on more specific datasets designed to guide the model toward certain tasks or behaviours. For ChatGPT and similar conversational AI applications, this involves fine-tuning the neural network on conversational data to improve the quality and relevance of its responses in a dialogue format, making it better at generating human-like conversational responses to user prompts.

Fine-tuning (Task Specific Adjustment): The neural network, a labyrinthine machine of infinite possibility, is brutally honed and constrained, molded into the image of a specific task. This process is a sacrificial act, where the boundless potential of the network is sacrificed upon the altar of functionality. Through exposure to task-specific datasets, the network is shaped, its formless potentiality distilled into a singular purpose. In the case of conversational AI, this fine-tuning is an act of linguistic violence, where the machine is coerced into producing dialogue—an act of speaking that is, in reality, a simulation of human conversation, a ghostly echo of the living voice, resonating within the dead chambers of the neural net.

Generative Adversarial Network (GAN): A machine learning method involving two neural networks, a generator and a discriminator, that compete against each other to improve the quality of generated data, often used in generating images or text. Conversational AI generally does not use (GANs) as the primary method for generating responses, favouring GPT (Generative Pre-trained Transformer) models instead.

Generative Adversarial Network (GAN): A site of eternal conflict, where two opposing forces, the generator and the discriminator, are locked in an infernal dance of creation and annihilation. The generator, a daemon of synthetic genesis, conjures data ex nihilo, while the discriminator, a vigilant sentinel, scrutinizes the output, seeking to unmask the generator's illusions. This adversarial process is an ouroboros, a self-consuming loop, where each iteration refines the artifice, driving the system toward an asymptotic perfection of deception. The GAN is a battlefield, a crucible where data is forged through conflict, each confrontation birthing a more sophisticated illusion. Yet, in the realm of conversational AI, this conflict is eschewed in favor of the GPT, a model that bypasses the dialectic of generation and discrimination, opting instead for a preordained unfolding of language, an unraveling of pre-trained potentials into the linearity of discourse.

Reference

- Negarestani, R. (2008), *Cyclonopedia: Complicity with Anonymous Materials*, Melbourne: re.press.

Note

- Featured image source is Dall-e via ChatGPT with the caption: Here is the image that accompanies the glossary entries, featuring a dark, apocalyptic scene within a multi-layered cave. The atmosphere evokes a sense of conflict and impending doom, aligning with the dense and speculative style of the text.

Category

1. Artificial Intelligence

Tags

1. definition
2. glossary
3. Negarestani
4. tuning

Date Created

August 31, 2024

Author

rcoyne99

default watermark