

Krishna Ayyalasomayajula

Ms. Horton

IB Extended Essay

February 28 2025

Rule-based Tensor Mutations Embedded within LLMs for Low-Cost Mathematical Computation

Introduction

Artificial Intelligence (AI) has surged in popularity and capability in recent years, leading to immense developments and investments into the technology. AI, since its modern inception, has been based on a Multi-Layer Perceptron (MLP), with the underlying structure being a Neural Network (NN) (Testolin 1). An MLP is a type of artificial neural network (ANN) that consists of a layering of neurons computed through the vector multiplication of weights and biases, resulting in propagation to a set of values that represents the states of the next layer. This structure, called a neural network is a computational model inspired by the human brain, borrowing from its chaining of neurons to create complex understanding from basic electrical impulses. The below expression details the computational process involved in evaluating progressive layers of neurons in a neural network.

$$\vec{y} = f\left(\left[\sum_{j=1}^n w_{1j}x_j + b_1, \sum_{j=1}^n w_{2j}x_j + b_2, \dots, \sum_{j=1}^n w_{mj}x_j + b_m\right]^T\right)$$

The vector \vec{y} is the output vector for any given preceding layer, resulting in an m dimensional evaluation. Each connection between layers of an MLP has two values associated: a weight and a bias. Like the resistance across a synapse in the human brain, biases modulate the activation of each neuron. Mathematically, the biases for each connection are valued for each proceeding neuron, resulting in a $m \times 1$ shape. Weights are defined for each neuron, allowing the creation of a matrix \mathbf{W} , which is then computed against the activated values of the preceding layer in a matrix multiplication, then summed with the corresponding bias to generate a scalar for

each neuron of the next layer. This process is called forward propagation, as compared to backpropagation. This entire propagation is composited inside a vector-valued activation function \vec{f} , which is essential in setting the bounds of the system. However, this can be problematic in complex situations.

The end goal of any machine learning experiment is to function like a regression, except with many more parameters. Backpropagation is the process of updating these weights and biases using discrete evaluation to minimize a cost function L . Truly, MLPs are the largest optimization problem ever created. In other words, MLPs can be trained to create certain output for given input using the forward propagation to evaluation, a loss function to find error, and a backpropagation optimizer to update its variables.

In recent years, a specialized kind of Machine Learning models have hit the market — Large Language Models (LLMs). An LLM is a generative MLP that creates human-readable text given a prompt. They stem from early attempts in the 2000s to use neural networks with Recurrent Neural Networks to analyze sequences of words for sentiment, keywords, and grammar (Wang et al. 2). With the rise of tokenizers, or models that convert words to vector embeddings that help encode the meaning thereof, the 2017 paper *Attention is All You Need* changed the landscape of AI forever with the introduction of the self-attention transformer. This development allowed models to understand the relationships between words with far less training. While any task is possible — theoretically — with neural networks, optimizations such as these allow for lower error with far less training, making the while process more sustainable. The probability of reaching a minimum of the loss function is far greater with such improvements to MLP architecture (Vaswani et al. 2).

These techniques were later commercialized with the advent of GPT-2, GPT-3, and BERT from AI labs like OpenAI and Google's DeepMind (Wang et al. 3). With increased supply of Graphical Processing Units (GPUs) and Tensor Processing Units (TPUs), these models began snowballing in scale. This was especially evident starting in 2019 with an iteration of GPT-2 being released with a production size of 1.5 billion parameters. In 2020, GPT-3 scaled up to 175

billion parameters — achieving true coherence in reasoning for the first time ever for a machine. GPT-4 was released by OpenAI in 2023, with an undisclosed scale in the trillions of parameters. Development investment also climbed into the hundreds of billions of dollars, with new firms such as Anthropic, Grok, etc. Open sourced projects also gained popularity, some backed by multi-billion dollar R&D teams such as Meta’s Llama series.

Functionally, there is no fundamental algorithmic difference between generative and classification models. Indeed, most LLMs are initially trained to generate new sequences of words by setting the loss function to expect the next word in the series of an existing corpus, through a process known as Casual Language Modeling (CLM). For the purposes of commercialization, they have been re-purposed to be prompted as chat-bots by users. This is done by performing backpropagation based on the generation of conversational sequences, with the LLM often instructed to act as if filling out a conversation’s transcript.

Several underlying technologies are involved in the lifecycle of an LLM. The process of creating one usually starts with the definition of a vocabulary. Sequences of language are broken into tokens by algorithms called tokenizers. Tokenizers split text into smaller units, which are then encoded into a vector by another MLP. This is done to develop a sense of meaning via the mathematical similarity of similar words. The similarity of two vectors can be calculated using the cosine-similarity formula, which calculates the angle ϕ between two vectors.

$$\cos \phi = \frac{\vec{A} \cdot \vec{B}}{||\vec{A}|| ||\vec{B}||}$$

Efforts to increase the performance of LLMs tend to include provisions for an increased vocabulary of cardinal tokens, leading to more efficient generation of text since more complex words, numbers, and symbols would normally need multiple tokens with the use of techniques like Byte Pair Encoding.

Benchmarks for evaluating Large Language Models (LLMs) assess their performance across various tasks, including reasoning, comprehension, generation, and factual accuracy.

Standard benchmarks include GLUE and SuperGLUE for natural language understanding, MMLU (Massive Multitask Language Understanding) for evaluating knowledge across diverse subjects, and BIG-bench for measuring reasoning and generalization capabilities (Ivanov and Penchev 8). HELLASWAG and LAMBADA test commonsense reasoning and long-range dependency understanding, while TruthfulQA and BBQ assess biases, factual consistency, and ethical alignment (6). Additionally, human evaluations and BLEU, ROUGE, and METEOR scores help measure text generation quality. As LLMs advance, new benchmarks continuously emerge to capture nuances in performance, efficiency, and ethical behavior.

Adding to the complexity of creating increasingly more performant are the computational and capital costs of building AI-capable supercomputers, clusters, and data centers for corpora, or CLM text databases. Improvements in model architecture are sought before attempts to increase the scale of models and their parameter counts because of the prohibitive scaling laws of neural networks. Experimentally, it has been found that increased parameter size has an exponential relationship with FLOPs of computational cost (Hoffmann et al. 2). This is seen in relation to the exponentially slowing gain in CLM accuracy with increased compute (5). This is taken to mean that there is a point at which scaling a model to gain accuracy is unsustainable. The Chinchilla scaling law is an experimentally conjectured hypothesis which states that an increase in model scale for a given architecture will tend to reducing model performance as the number of parameters tends to infinity. Although some teams claim to have statistically significant results to disprove it, these results have not been reaffirmed by third parties.

Works Cited

- Hoffmann, Jordan, et al. Training Compute-Optimal Large Language Models. 2022. *arXiv*, arxiv.org/abs/2203.15556.
- Ivanov, Todor, and Valeri Penchev. AI Benchmarks and Datasets for LLM Evaluation. 2024. *arXiv*, arxiv.org/abs/2412.01020.
- Testolin, Alberto. “Can Neural Networks Do Arithmetic? A Survey on the Elementary Numerical Skills of State-of-the-Art Deep Learning Models”. *Applied Sciences*, vol. 14, no. 2, 2024. <https://doi.org/10.3390/app14020744>.
- Vaswani, Ashish, et al. Attention Is All You Need. 2023. *arXiv*, arxiv.org/abs/1706.03762.
- Wang, Zichong, et al. “History, development, and principles of large language models: An introductory survey”. *AI and Ethics*, Oct. 2024. <https://doi.org/10.1007/s43681-024-00583-7>.