

TRANSFORMING THE HEALTHCARE INDUSTRY WITH LARGE LANGUAGE MODELS

The healthcare industry is a complex and data-rich field that requires constant innovation to improve patient care, enhance diagnostic accuracy, and streamline administrative processes.

In recent years, large language models (LLMs) powered by artificial intelligence (AI) have emerged as a groundbreaking technology with the potential to revolutionize the health care landscape. These sophisticated models, such as OpenAI's GPT-3, have the ability to process and generate human-like text, making them invaluable tools for transforming various aspects of the healthcare industry.



However, off the shelf language models are built for multiple purposes. It is true that these LLMs have gained significant attention and utility across various fields, including natural language processing (NLP), conversational AI, and content generation.

However, when it comes to leveraging LLMs, there are two primary approaches:

1. Building your own model.
2. Utilizing existing models developed by organizations like OpenAI.

What if you could build an LLM model for your own organizational targets and needs?

Here are the key differences between using off the shelf models vs building your own LLM models.

Expertise, Knowledge and Resources: Creating a large language model from scratch requires extensive resources, expertise of AI developers and data scientists, MLOps team, and computational power.

It involves training the model on massive datasets, fine-tuning it through multiple iterations, and optimizing its performance. This process demands substantial time, expertise and computational resources, including high-performance hardware and storage systems.

However, custom build models are substantially more powerful for your use cases than using off the shelf models. The good news is that the BigRio team has extensive knowledge about Reinforcement Learning (RL) and Natural language processing (NLP) which are the basis of LLMs.

Time to Deployment: Developing a large language model from scratch is a time-consuming process. It involves several stages including data collection, preprocessing, training, and optimization. The time required for this process can vary significantly depending on the scale of the model, available resources, and expertise.

The BigRio team has extensive experience developing such language models. Therefore, we could substantially cut the time to development of building such models.

Continuous Updates and Improvements: Organizations that develop their own LLMs have the advantage of full control over the model's development, updates, and fine-tuning. They can adapt the model to their specific needs, incorporate domain-specific knowledge, and continuously improve the model over time.

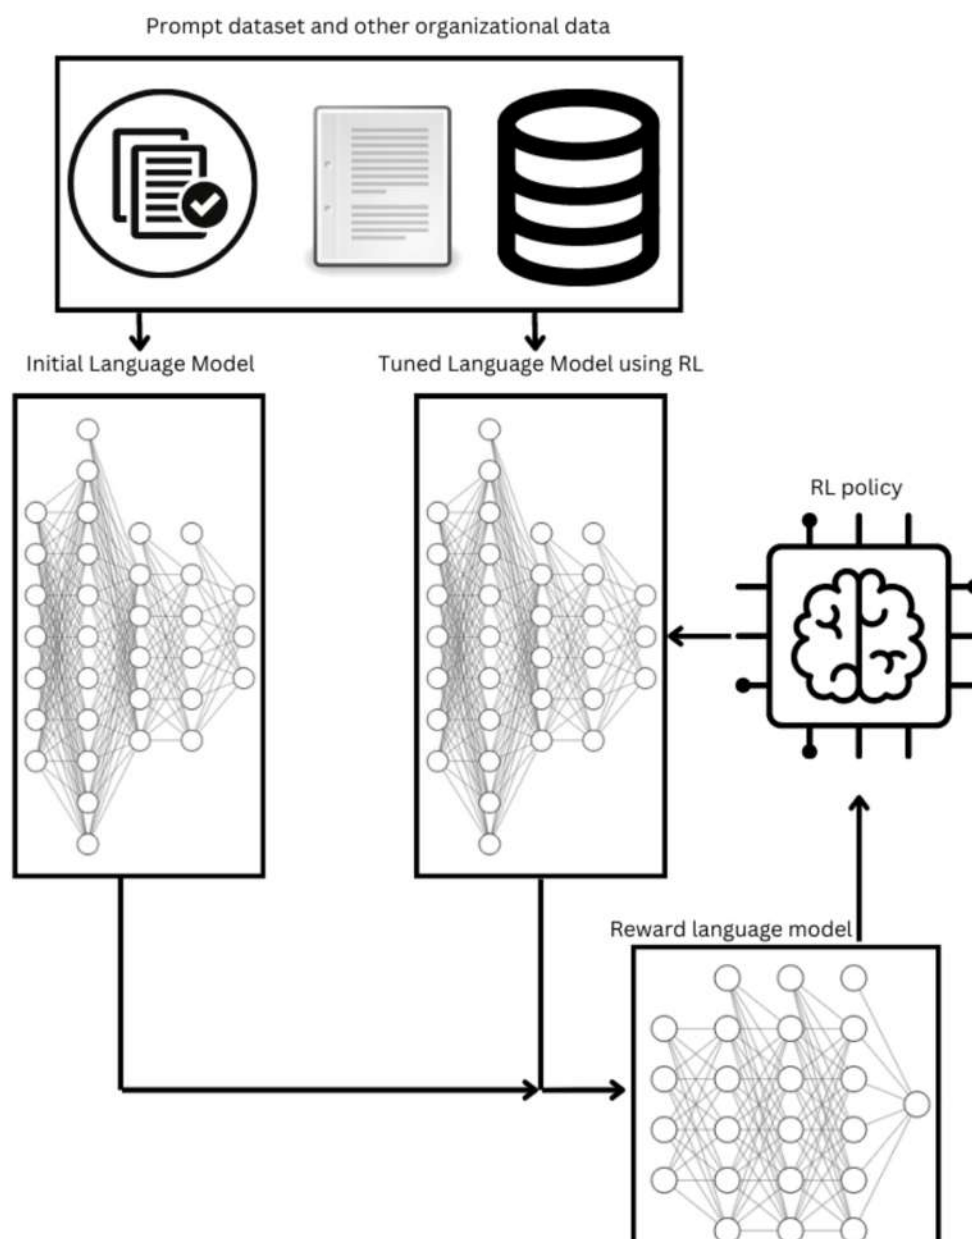
BigRio team has extensive experience with successfully implementing best practices of MLOps. Our experience can help organizations to perform continuous updates and improvements.

Why is now a good time to build your own LLM?

For a considerable time, training a language model using reinforcement learning seemed impossible due to engineering and algorithmic challenges. However, several organizations have managed to achieve success by employing a policy-gradient RL algorithm called Proximal Policy Optimization (PPO) to fine-tune some or all of the parameters of an initial language model.

Since fine-tuning an entire model with billions or even trillions of parameters is prohibitively expensive, the parameters of the language model are frozen. PPO has been a well-established algorithm with extensive documentation, making it a suitable choice for scaling up to distributed training in this context. Notably, the advancements in RL have facilitated the update process for such large-scale models using familiar algorithms.

To comprehend the fine-tuning task as an RL problem, we can define the policy as a language model that takes a prompt as an input and generates a sequence of text or probability distributions over text. The action space of this policy encompasses all tokens corresponding to the language model's vocabulary, while the observation space comprises the distribution of potential input token sequences, which is also considerably large based on prior RL applications. The reward function for this RL problem combines the preference model and a constraint on policy shift. The graph here shows the overall complexity of such LLMs.



Having the ability to build such LLMs tailored to your organizational needs could open the door for many opportunities. Let's explore some of the use cases in more detail.

Enhancing Clinical Decision-Making: One of the most promising applications of LLMs in health care is supporting clinical decision-making. These models can analyze vast amounts of medical literature, research papers, and patient data to provide evidence-based recommendations for diagnosis, treatment plans, and personalized medicine. By leveraging the collective knowledge of medical experts, LLMs can assist physicians in making more informed decisions, improving patient outcomes, and reducing diagnostic errors.

Natural Language Processing in Electronic Health Records (EHRs): Electronic Health Records (EHRs) have greatly improved the accessibility and organization of patient information. However, the abundance of unstructured data within EHRs poses challenges in extracting meaningful insights. LLMs equipped with advanced natural language processing (NLP) capabilities can parse through these records, extracting relevant information, and converting unstructured data into structured data. This transformation allows for better data analysis, risk prediction, and early detection of patterns that may aid in disease prevention and population health management.

Empowering Patient Engagement and Education: Health care providers often struggle to communicate complex medical information to patients in a clear and understandable manner. LLMs can bridge this gap by generating patient-friendly explanations of medical conditions, treatment options, and medication instructions. By simplifying medical jargon, LLMs can empower patients to actively participate in their own care, leading to improved adherence to treatment plans and better overall health outcomes.

Efficient Administrative Processes: The healthcare industry involves numerous administrative tasks such as claims processing, appointment scheduling, and insurance documentation. LLMs can automate and streamline these processes, reducing the administrative burden on health care providers and improving overall efficiency. For example, chatbots powered by LLMs can handle basic inquiries, schedule appointments, and provide information about insurance coverage, freeing up staff to focus on more complex tasks.

How could BigRio help more? Building your own LLM model is a powerful tool. So, there will be many ethical considerations and challenges. While the potential benefits of LLMs in health care are substantial, ethical considerations and challenges must be addressed. Patient privacy, data security, and bias in AI algorithms are critical concerns that need to be carefully managed. Striking the right balance between automation and human intervention is also important to ensure that LLMs are used as



tools to augment and support health care professionals rather than replacing them. The BigRio team has an outstanding track record of addressing these challenges and concerns.

As the saying goes, "With great power comes great responsibility" and BigRio proudly takes this on for our clients. Please email us at info@bigr.io to learn more about how we can collaborate and continue to move the needle forward using AI in healthcare!