Analytics Vidhya

Skills for Success

A Comprehensive and Industry-Aligned Curriculum

Table of Contents

Curriculum Overview

The GenAI Pinnacle program is meticulously crafted in response to the burgeoning demand in Generative AI technology, featuring elite instructors renowned for their profound expertise in the field.
Our comprehensive curriculum delves into Large Language Models, Stable Diffusion Models and the foundational elements of Data Science & Machine Learning, preparing you to excel in the rapidly evolving landscape of Generative AI.

- 200+ Hours of Learning
- 26+ Tools, Libraries & Frameworks
- 10+ Generative AI Projects
- 15+ Assignments to practice

Large Language Models

1: Getting Starting with LLMs

- **Evolution of NLP**
  - Evolution of NLP: Pre Deep Learning
  - Evolution of NLP: Deep Learning
  - Transformers

- **Introduction to LLMs**
  - What is a language model?
  - What is a large language model?
  - Understanding Foundational Model
  - Different types of LLMs: Continuing the text
  - Different types of LLMs: Instruction Tuned
  - Encoder LLMs, Decoder LLMs and Encoder Decoder LLMs

- **The Current State of the Art in LLMs**
  - The Current State of the Art in LLMs
  - What Next?

2: Building LLM applications using Prompt Engineering

- **How to build LLM Applications?**
  - Overview of different methods to build LLM applications

- **Introduction to Prompt Engineering**
  - Why do we need Prompt Engineering?
  - What is Prompt Engineering?

- **Prompt Engineering using ChatGPT API and Open Source LLM**
  - Introduction to ChatGPT API
  - Prompt Engineering with Open Source LLM

- **Understanding Different Prompt Engineering Techniques**
  - Few-shot Prompting
  - One-shot Prompting
  - Zero-shot Prompting
  - CoT Prompting (**Coming Soon**)
  - Self Consistency (**Coming Soon**)
  - ToT (**Coming Soon**)
  - Advanced Prompt Engineering Techniques (**Coming Soon**)

- **Building LLM App using Prompt Engineering**

3: Building Production Ready RAG systems using LlamaIndex

- **Introduction to RAG systems**
  - Why RAG systems?
  - What is a RAG system?
  - A brief overview of steps involved in a RAG Framework

- **Introduction to LlamaIndex**
  - What is a LlamaIndex?
  - Why LlamaIndex?

- **Components of LlamaIndex**
  - Overview of Different Components of LlamaIndex
  - Data Sources
  - Indexing
  - Embeddings
  - VectorDBs
  - LLMs
  - Retrieval
  - Decoupling chunks used for retrieval vs. chunks used for synthesis
  - Adding Metadata

- **Code: Build your first QA system on private data**

- **Evaluation of RAG systems**
  - Necessity of Evaluation
  - Retrieval Evaluation
  - Response Evaluation

- **Code: RAG + Evaluation**

- **Advanced approaches for powerful RAG system**
  - Router Query Engine, SubQuestionQuery Engine
  - Structured Retrieval for Larger Document Sets: Metadata filters + Auto Retrieval
  - Document Hierarchies + Recursive Retrieval
  - Finetuning Embeddings

- **Agents**

- **Building a powerful RAG system using LlamaIndex**

4: Finetuning LLMs

- **Introduction to Finetuning LLMs**
  - What is Finetuning LLMs?
  - Why Finetuning LLMs?
  - Research Papers and blogs for additional details

- **Understanding different finetuning techniques**
  - Feature Extraction
  - Retraining all parameters
  - The problem with traditional methods
  - Setting up the runpod instance
  - Code Part1: Finetuning NLP usecase using feature extraction and full finetuning
  - Code Part2: Finetuning NLP usecase using feature extraction and full finetuning
  - Research Papers and blogs for additional details

- **Introduction to Parameter Efficient Finetuning Techniques**
  - What is Parameter Efficient Finetuning Techniques?
  - How PEFT addresses the issues with Full Fine-tuning
  - Different types of PEFT methods

- **Prompt Tuning Techniques**
  - Prompt Tuning
  - Prefix Tuning
  - Code:Finetuning NLP usecase using Soft Prompting Techniques

- **Adaptor Modules**
  - (IA)3
  - LoRA
  - Quantization
  - QLoRA
  - Code:Finetuning NLP usecase using PEFT LoRA
  - Code: Finetuning instruction following LLM using PEFT

5: Training LLMs from Scratch

- **Training LLMs from Scratch**
  - Scaling laws
  - Parallel and Distributed computing strategies
  - Challenges involved in training LLMs
  - Best practices to train LLMs
  - How do you train LLMs from scratch?
  - Project: Continuing text
  - Project: Instruct optimized + RLHF

- **Evaluating LLMs**

Stable Diffusion Models

1: Getting Started with Stable Diffusion

- **Introduction to Stable Diffusion**
  - Introduction and fundamental concepts.
  - History and evolution of Stable Diffusion.
  - Detailed look at Stable Diffusion as a shining star in AI.
  - Insights into versions 2.0 and 2.1 of Stable Diffusion.
  - Paper review: Understanding CNNs through Fergus and Zeiler's work.

- **Overview of the Stable Diffusion Process**
  - Basic intuition behind Stable Diffusion.
  - Real-world examples showcasing Stable Diffusion.
  - Reverse diffusion explained.
  - The overarching purpose of Stable Diffusion in AI.

- **Understanding Core Components of Stable Diffusion**
  - Exploring UNet architecture.
  - Understanding the role of latents in Stable Diffusion.
  - Two-part series on CLIP: its role and application in Stable Diffusion.

- **Summarizing Stable Diffusion**
  - How different components integrate within Stable Diffusion.
  - Future directions and what lies ahead in Stable Diffusion.

- **Understanding Variational Autoencoders**
  - Introduction and working principles of VAEs.
  - Applications and significance of VAEs in the context of Stable Diffusion.

- **In-depth Understanding of Stable Diffusion - Part I**
  - Understanding the forward process in Stable Diffusion.
  - Exploring the noise schedule and its role.
  - Coding the forward process.
  - Reverse process demystified.

- **In-depth Understanding of Stable Diffusion - Part II**
  - Positional encoding in detail.
  - Loss functions specific to Stable Diffusion.
  - Components of the diffusion model explained.

- **Implementing DDPM from Scratch**
  - Understanding the problem statement.
  - Step-by-step guide to implementing DDPM from scratch across multiple videos.

2: Mastering Methods and Tools of Stable Diffusion

- **Understanding DALL·E 2**
  - The Future Roadmap of DALL·E 2.
  - Understanding Text Conditioning in DALL·E 2.
  - Delving into OpenAI's CLIP and its integration with DALL·E 2.
  - Overview of the Training Process for DALL·E 2.
  - Introduction to Glide, a component of DALL·E 2.
  - Various Training Pathways for DALL·E 2.

- **Steps Involved in Training Stable Diffusion**
  - Training from Scratch vs. Finetuning.
  - Strategies for Training Stable Diffusion at Scale.
  - Discussing the Ethical Implications in AI training.
  - A series of deep dive hands-on videos covering the setup, auto encoders, schedulers, text and positional embedding, experiments with embeddings, and textual inversion in Stable Diffusion.

- **Mastering stability.ai and its tools**
  - Introduction to Stability.AI.
  - Utilizing Stability.AI tools effectively.
  - Reviewing Deployment Options for SD WebUI.
  - Using Automatic1111's WebUI on RunPod's GPU environment.

- Hands-on sessions focusing on the installation, setup, generation, and settings within the SD WebUI.

- **Prompt Engineering Concepts for Stable Diffusion**
  - Analyzing Prompts and Understanding the Prompts Matrix.
  - Delving into Prompt Strengths and Weights.
  - Techniques for Prompt Editing and Blending.
  - Utilizing XYZ Plots and exploring Dream Studio for enhanced prompt engineering.

3: Advanced Stable Diffusion Techniques

- **Exploring Img2Img Techniques**
  - Overview of the Img2Img course.
  - Understanding Img2Img through categorical intuition.
  - In-depth exploration of image in-painting and its various forms and factors.
  - Image editing techniques within Img2Img.
  - Hands-on tutorials covering vanilla prompt-based transformation, image in-painting, and sketch and in-paint sketch in Img2Img.

- **CIVITAI Introduction and Applications**
  - Introducing CIVITAI and its relevance in Stable Diffusion.
  - Exploring the extensions of Stable Diffusion and their applications in CIVITAI.

- **Delving into MidJourney**
  - The storm and story behind MidJourney.
  - Drawing inspiration from MidJourney and its unique aspects.
  - Practical guide on how to use MidJourney effectively.
  - Exploring alternatives to MidJourney and their comparative benefits.

- **InstructPix2Pix and ControlNet**
  - Reviewing the InstructPix2Pix paper, providing an overview and setup guidance.
  - Introducing ControlNet, its setup, and its significance in Stable Diffusion.
  - Exploring ControlNets OpenPose and its applications.
  - Techniques for upscaling images using Stable Diffusion and ControlNets.

Basics of Data Science & Machine Learning

- **Python for Data Science**
  - Basics of Python
  - Explore Python for Data Science
  - Important libraries and functions in Python
  - Reading files and manipulating data in Python
  - Working with data frames, lists, and dictionary

- **Data Exploration and Statistical Inference**
  - Use Matplotlib and Seaborn for data visualization
  - Creating charts to visualize data and generate Insights
  - Univariate and Bivariate analysis using python
  - Perform Statistical Analysis on real-world datasets
  - Build and Validate Hypothesis using statistical tests
  - Generating useful insights from the data

- **Basics of Machine Learning**
  - Learn Important Machine Learning concepts
  - Perform data cleaning and Preprocessing
  - In-depth understanding of the Basics of prediction models
  - k-NN, Linear Models, Decision Tree
  - Math Behind Each Machine Learning Algorithm

- Advanced algorithms like Bagging and Boosting, SVM, Naive Bayes
- Hyperparameter Tuning to improve model
- Solving real-world business problems using Machine Learning

- **Fundamentals of Deep Learning**
  - Essential Concepts of Deep Learning
  - Working on Neural Network from Scratch
  - Activation Functions and Optimizers for Deep Learning
  - Understand Deep Learning architectures (MLP, CNN, RNN, LSTM, GRUs)
  - Explore Deep Learning Framework, PyTorch
  - Learn to tune the hyperparameters of Neural Networks
  - Build Deep Learning models to tackle real-life problems
  - Use Transfer Learning for training models

Shaping Tomorrow with Generative AI

Enroll now

Contact us:


08068342847 | 08046107668


training_support@analyticsvidhya.com


analyticsvidhya.com/genaipinnacle