



GENERATIVE AI TRAINING



OVERVIEW:

Welcome to the Generative AI course at Dicecamp, where you will delve into the cutting-edge world of artificial intelligence that creates new content, designs, and solutions. This comprehensive course is designed to equip you with the knowledge and skills necessary to harness the power of Generative AI in various applications, from art and music to advanced problem-solving in technology and business.

- **Introduction to Generative AI:** Understand the fundamentals of Generative AI, including its history, evolution, and significance in the modern AI landscape.
- **Key Concepts and Techniques:** Learn about key techniques such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and other state-of-the-art models that drive generative processes.
- **Practical Applications:** Explore real-world applications of Generative AI in industries like entertainment, healthcare, finance, and more. Discover how AI can generate art, music, text, and even design new products.
- **Hands-on Projects:** Engage in hands-on projects that allow you to apply generative models to create unique outputs. Gain practical experience in training and fine-tuning generative models.
- **Tools and Frameworks:** Get acquainted with popular tools and frameworks used in Generative AI, such as TensorFlow, PyTorch, and other AI development environments.
- **Ethical Considerations:** Discuss the ethical implications of Generative AI, including issues related to originality, copyright, and the potential for misuse.
- **Future Trends:** Stay ahead of the curve by exploring future trends and advancements in Generative AI. Understand the direction in which this exciting field is heading.

WHAT ARE THE TAKEAWAYS?

- Foundational Understanding of Generative AI
- Understand the differences between generative and discriminative models.
- Understand the theoretical underpinnings of neural networks and deep learning as they apply to generative models.
- Learn how to collect, clean, and preprocess data for training generative models.
- Explore practical applications of generative AI in various domains, such as text generation, image synthesis, music composition, and code generation.
- Work on projects that apply generative AI techniques to real-world problems.
- Enhance communication skills to effectively present and document AI projects and findings.
- Develop critical thinking by evaluating the strengths and limitations of different generative models.
- How to Build Chatbots.

CORE TRAINING CONTENT

Week 01: PYTHON BASICS

- Basics of Data Science Flow
- Anaconda Installation
- Intro to Jupyter Notebook
- Intro to Python
- Python Objects & Data Structure
- Subsetting (Strings, Lists, Dictionaries)
- Python Comparison Operators
- Python Statements
- Methods & Functions
- Importing Data in Python
- NumPy & Pandas Basics in Python
- Subsetting Dataframes in Pandas
- Data Aggregation using Group By, Pivot and Melt

Week 02: Introduction to Large Language Models (LLMs) and

1. **GEN AI Overview**
 - Introduction to the course structure and objectives.
 - Brief overview of Large Language Models (LLMs).
2. **Understanding LLMs**
 - Overview of LLM applications and importance.
 - Key concepts: Tokens, context window, token limits.
3. **Introduction to Embeddings**
 - Basics of text embeddings.
 - Classical text encoding techniques: One-hot encoding, TF-IDF.
4. **Hands-on Exercise: TF-IDF and One-Hot Encoding**
 - Practical session on creating and comparing TF-IDF and one-hot vectors.
 - Group activity to explore the impact of encoding on text data.
5. **Introduction to Word Embeddings**
 - Word2Vec and its importance.
 - Use cases of Word2Vec in text analytics.
6. **Hands-on Exercise: Building Word2Vec Models**
 - Practical session on building and exploring Word2Vec models.
7. **Text Similarity Measures**
 - Explanation of cosine similarity, dot product.
 - Practical examples of measuring text similarity.
8. **Review and Q&A**
 - Summary of key learnings.
 - Open floor for questions and discussion.

Week 03: Advanced Embedding Techniques and Vector Databases

1. **Review of Week 2 and Introduction to Advanced Embeddings**
 - Recap of key concepts from Week 2.
 - Overview of advanced embedding techniques: GloVe, FastText.
2. **Semantic Encoding Techniques**
 - In-depth discussion of semantic encoding.
 - Practical applications of GloVe and FastText.
3. **Hands-on Exercise: Semantic Embeddings**
 - Creating and comparing semantic embeddings.
 - Group activity: Analyze the effectiveness of different embeddings.
4. **Introduction to Vector Databases**
 - Importance of vector databases in LLMs.
 - Overview of popular vector databases.
5. **Indexing Techniques**
 - Introduction to indexing techniques: Product Quantization, LSH, HNSW.
 - Hands-on: Creating an index using HNSW.
6. **Retrieval Techniques**
 - Understanding retrieval methods: Cosine similarity, nearest neighbor search.
 - Practical examples of retrieval techniques.
7. **Building a Vector Store**
 - Practical session on creating and querying a vector store.
 - Group activity: Implement a vector database for a given dataset.
8. **Review and Q&A**
 - Summary of key learnings.
 - Open floor for questions and discussion.

Week 04: Attention Mechanisms and Transformer Models

- **Review of Week 3 and Introduction to Attention Mechanisms**
 - Recap of key concepts from Week 3.
 - Introduction to attention mechanisms in LLMs.
- **Understanding Self-Attention**
 - Explanation of self-attention and its role in transformers.
 - Practical examples of self-attention.
- **Hands-on Exercise: Self-Attention**
 - Implementing self-attention for contextual word analysis.
 - Group activity: Explore the impact of self-attention on text data.
- **Introduction to Transformer Models**
 - Overview of transformer architecture.
 - Key components: Tokenization, embedding, positional encoding.

- **Multi-Head Attention and Transformer Blocks**
 - Explanation of multi-head attention.
 - Exploring transformer blocks and their functionality.
- **Hands-on Exercise: Building a Transformer Model**
 - Practical session on constructing a simple transformer model.
 - Group activity: Analyze the performance of the transformer model.
- **Application of Transformer Models**
 - Real-world applications of transformers in NLP tasks.
 - Case studies of transformer models in action.
- **Review and Q&A**
 - Summary of key learnings.
 - Open floor for questions and discussion.

Week 05: Semantic Search and Advanced Retrieval Techniques

- **Review of Week 4 and Introduction to Semantic Search**
 - Recap of key concepts from Week 4.
 - Introduction to semantic search and its importance.
- **Understanding Lexical vs. Semantic Search**
 - Comparison of lexical and semantic search techniques.
 - Practical examples of semantic search.
- **Hands-on Exercise: Building a Semantic Search Engine**
 - Implementing a basic semantic search engine.
 - Group activity: Enhance the search engine with additional features.
- **Multilingual Search Techniques**
 - Exploring multilingual search capabilities.
 - Practical applications of multilingual search.
- **Limitations and Improvements in Semantic Search**
 - Challenges in semantic search.
 - Techniques for improving search accuracy beyond embeddings.
- **Advanced Retrieval Techniques**
 - Overview of advanced retrieval methods.
 - Practical examples of applying retrieval techniques in semantic search.
- **Hands-on Exercise: Enhancing Semantic Search**
 - Practical session on optimizing a semantic search engine.
 - Group activity: Implement advanced retrieval techniques.
- **Review and Q&A**
 - Summary of key learnings.
- Open floor for questions and discussion

Week 06: Prompt Engineering and Model Control

1. **Review of Week 5 and Introduction to Prompt Engineering**
 - Recap of key concepts from Week 5.
 - Introduction to prompt engineering.
2. **Designing Effective Prompts**
 - Techniques for designing effective prompts.
 - Prompting by instruction and example.
3. **Hands-on Exercise: Creating Prompts**
 - Practical session on crafting prompts for various tasks.
 - Group activity: Analyze the effectiveness of different prompts.
4. **Controlling Model Output**
 - Techniques for controlling the output of LLMs.
 - Practical examples of controlling model behavior.
5. **Advanced Prompting Techniques**
 - Exploring advanced prompt engineering methods.
 - Case studies of successful prompt engineering.
6. **Hands-on Exercise: Advanced Prompt Engineering**
 - Practical session on advanced prompt design.
 - Group activity: Implement complex prompt engineering techniques.
7. **Real-World Use Cases of Prompt Engineering**
 - Applications of prompt engineering in various domains.
 - Examples of tasks like summarization, sentiment analysis, and translation.
8. **Review and Q&A**
 - Summary of key learnings.
 - Open floor for questions and discussion.

Week 07: Fine-Tuning Foundation Models

- **Review of Week 5 and Introduction to Fine-Tuning**
 - Recap of key concepts from Week 6.
 - Introduction to fine-tuning large language models.
- **Understanding Fine-Tuning Techniques**
 - Overview of fine-tuning methods: RLHF, transfer learning.
 - Practical applications of fine-tuning.
- **Hands-on Exercise: Fine-Tuning Basics**
 - Practical session on fine-tuning a pre-trained model.
 - Group activity: Evaluate the fine-tuned model on specific tasks.
- **Parameter-Efficient Fine-Tuning**
 - Introduction to PEFT methods like quantization, LoRA.
 - Practical examples of parameter-efficient fine-tuning.

- **Hands-on Exercise: Parameter-Efficient Fine-Tuning**
 - Implementing PEFT on a given model.
 - Group activity: Compare the performance of different fine-tuning approaches.
- **Fine-Tuning vs. Retrieval-Augmented Generation (RAG)**
 - Comparison of fine-tuning and RAG techniques.
 - Practical use cases for each approach.
- **Deploying Fine-Tuned Models**
 - Techniques for deploying fine-tuned models in real-world applications.
 - Hands-on: Deploying a fine-tuned model on a cloud platform.
- **Review and Q&A**
 - Summary of key learnings.
 - Open floor for questions and discussion.

Week 08: Orchestration Frameworks and Autonomous

1. **Review of Week 7 and Introduction to Orchestration Frameworks**
 - Recap of key concepts from Week 7.
 - Introduction to orchestration frameworks for LLMs.
2. **Understanding LangChain and Orchestration Needs**
 - Overview of LangChain and its components.
 - Practical applications of orchestration frameworks in LLMs.
3. ****Hands-on Exercise**

TOOLS:

- **OLLAMA**
- **WEAVITE**
- **OPENAI**
- **LANGCHAIN**

TRAINER PROFILE

Ali Raza Anjum

Co-Founder | DICE Analytics | Dicecamp | Viewlytics | KNIME Certified Trainer | Data Science | Big Data | Business Analytics | Artificial Intelligence | Generative AI | Consultant

14+ Years of Experience working in the Domain of Data Analytics, Data Science, Big Data, Artificial Intelligence and Machine Learning