




**datasciencedojo**  
— data science for everyone —

# Large Language Models Bootcamp

Get started with building LLM Applications



# Course Overview

---

Pre-trained Large Language Models like ChatGPT offer impressive capabilities but they cannot be used in scenarios where the underlying data is proprietary and requires industry-specific knowledge. Businesses are rushing to build custom LLM applications that offer enhanced performance, control, customization and most importantly, competitive advantage.

This bootcamp offers a comprehensive introduction to get started with building a ChatGPT on your own data. By the end of the bootcamp, you will be capable of building LLM-powered applications on any dataset of your choice. It is an instructor-led course taught by industry experts and covers:

- Generative AI and LLM Fundamentals
- Canonical Architectures of LLM Applications
- Embeddings and Vector Databases
- Prompt Engineering
- Orchestration Frameworks: LangChain and Llama Index
- Deployment of LLM Applications
- Customizing Large Language Models
- Building An End-to-End Custom LLM Application

The course encompasses a series of lectures supplemented with hands-on exercises. It is designed for beginners, data professionals and product leaders. For beginners, it offers a headstart with hands-on experience in building LLM applications. For data professionals, their data skills will be supercharged through cutting-edge generative AI tools and techniques. For product leaders in enterprises or startups, it provides insights to leverage LLMs for enhancing products, processes, and services.



# Learning Outcomes

- Understand the foundational concepts and architecture of large language models.
- Develop skills in designing and engineering effective prompts for various LLM applications
- Acquire expertise in text embedding techniques and efficient vector database management for storage and retrieval
- Acquire practical understanding of transformer models and attention mechanisms.
- Hands-on experience in fine-tuning large language models, understanding their limitations and advanced tuning methods.
- Understand the use of orchestration frameworks in LLM applications and gain practical experience with LangChain and LlamaIndex.
- Knowledge in evaluating LLMs with metrics like BLEU, ROUGE and understanding LLMOps principles.
- Develop the ability to create and deploy RAG based custom LLM applications for industry-specific use cases.

# What you will learn

## The program will enable you to:

- LLM fundamentals, Transformers, attention mechanism
- Vector databases, semantic and hybrid search
- Retrieval-augmented generation (RAG)
- Langchain fundamentals
- Building multi-agent applications
- Observability and monitoring
- Evaluation datasets, tasks and metrics
- Guardrails and responsible AI
- Knowledge graphs and graph RAG
- Fine tuning and deploying a large language model



In today's fast-changing landscape, leaders who understand Large Language Models (LLMs) are better equipped to unlock business value, drive innovation, and stay competitive. By mastering LLMs, executives and professionals can confidently lead AI transformation across teams and workflows.

The LLM Bootcamp provides a solid foundation in LLM and generative AI concepts—covering prompting, fine-tuning, integration, and real-world use cases—while also helping professionals think strategically about implementation, ethics, and scalability in enterprise environments.

**Raja Iqbal**

*Founder at* **Ejento AI & Data Science Dojo**

# Program Highlights



## LLM Application Architectures

Understanding the components of a large-scale enterprise LLM application



## Challenges and Risks

Key challenges and risks in enterprise adoption of large language models



## Transformers and Attention

A comprehensive introduction to attention mechanism and transformer architecture



## Vector Databases

A comprehensive introduction to vector databases



## Prompt Engineering

An introduction to prompt engineering fundamentals



## Fine Tuning

A practical introduction to fine tuning



## Model Context

Introduction to prompt templates, retrievals, document loaders, and memory for context



## Agentic Workflows

Introduction to tools and multi-agent workflows



## Retrieval Augmented Generation (RAG)

Challenges in building an enterprise RAG pipeline



## LLM Evaluation

Evaluating RAG-based LLM applications



## Capstone Project

Building and deploying an LLM application

# Learning Journey

The LLM Bootcamp by Data Science Dojo is an immersive, hands-on training program designed to help professionals master large language models and build real-world generative AI applications. Covering everything from prompt engineering to fine-tuning and deployment, this bootcamp equips participants with the practical skills needed to harness the power of LLMs in production settings – all without requiring a deep background in AI or machine learning.

## Core Modules | 40 hrs

### LLM Application Architectures

In this module, we will understand the common use cases of large language models and the fundamental building blocks of such applications. Learners will be introduced to the following topics:

- Large language models and foundation models
- Prompts and prompt engineering
- Context window and token limits
- Embeddings and vector databases
- Build custom LLM applications by:
  - Training a new model from scratch
  - Fine-tuning foundation LLMs
  - In-context learning
- Canonical architecture for an end-to-end LLM application

### Challenges and Risks

In this module, we will explore the primary challenges and risks associated with adopting generative AI technologies. Learners will be introduced to the following topics at a very high level without going into the technical details:

- Misaligned behavior of AI systems
- Handling complex datasets
- Limitations due to context length
- Managing cost and latency
- Addressing prompt brittleness
- Ensuring security in AI applications
- Achieving reproducibility
- Evaluating AI performance and outcomes

### Transformers and Attention

Dive into the world of large language models, discovering the potent mix of text embeddings, attention mechanisms, and the game-changing transformer model architecture.

- Review of neural networks, deep learning and other fundamentals
- Encoder/decoder
- Transformer networks: tokenization, embedding, positional encoding and transformers block
- Attention mechanism
  - Self-Attention
  - Multi-head Attention
  - Transformer models

## Vector Databases

Learn about efficient vector storage and retrieval with vector database, indexing techniques, retrieval methods, and hands-on exercises.

- Overview
  - Rationale for vector databases
  - Importance of vector databases in LLMs
  - Popular vector databases
- Different types of search
  - Vector search, text search, hybrid search
- Indexing techniques
  - Product Quantization (PQ), Locality Sensitive Hashing (LSH) and Hierarchical Navigable Small World (HNSW)
- Retrieval techniques
  - Cosine Similarity, Nearest Neighbor Search
- Advanced Retrieval Augmented Generation techniques
  - Limitations of embeddings and similarity in semantic search
  - Query transformation for better retrieval
  - Relevance scoring in hybrid search using Reciprocal Rank Fusion (RRF)
  - Using auto-cut feature to remove irrelevant results dynamically
  - Improving search relevance by using language understanding to re-rank search results
- Challenges using vector databases in production
  - Scaling optimization
  - Reliability optimization
  - Cost optimization
- Hands-on Exercise
  - Learn how to perform similarity searches with vectors as input.
  - Learn how to perform queries using vector similarity searches with embedding models and vectors.
  - Learn how to combine the results of a vector search and a keyword (BM25F) search using hybrid search approach.
  - Learn how to use multi-tenancy features for the efficient and secure management of data across multiple users or tenants.
  - Learn how to compress vectors using product quantization to reduce memory footprint.

## Prompt Engineering

Unleash your creativity and efficiency with prompt engineering. Seamlessly prompt models, control outputs, and generate captivating content across various domains and tasks.

- Prompt Design and Engineering
  - Crafting Instructions for Effective Prompting
  - Utilizing Examples to Guide Model Behavior
- Innovative Use Case Development
  - Tailoring Prompts to Goals, Tasks, and Domains
  - Practical Examples:
    - Summarizing Complex Reports
    - Extracting Sentiment and Key Topics from Texts
- Understanding and Mitigating Prompt Engineering Risks
  - Identifying Common Risks: Prompt Injection, Prompt Leaking, Jailbreaking
  - Best Practices for Secure Prompt Engineering
- Advanced Prompting Techniques
  - Enhancing Performance with Few-Shot and Chain-of-Thought (CoT) Prompting
  - Exploring Program-aided Language Models (PAL) and ReAct Methods

## Fine Tuning

In-depth discussion on fine-tuning of large language models through theoretical discussions, exploring rationale, limitations, and Parameter Efficient Fine Tuning.

- Fine Tuning Foundation LLMs
  - Transfer learning, knowledge distillation and Fine-tuning
  - Different fine-tuning techniques
  - Limitations for fine-tuning
  - Parameter-efficient fine-tuning in depth.
  - \* Quantization of LLMs
  - \* Low-Rank Adaptation (LoRA) and QLoRA
  - Fine-tuning vs. RAG: When to use one or the other. Risks and limitations.
- Hands-on Exercise:
  - In-Class: Instruction fine-tuning, deploying, and evaluating a LLaMA2-7B 4-bit quantized model

## Model Context

Build LLM Apps using LangChain. Learn about LangChain's key components such as models, prompts, parsers, memory, chains, and Question-Answering.

- Introduction to LangChain:
  - Why do we need an orchestration tool for LLM application development?
  - What is LangChain?
  - Different components of LangChain
- Why are orchestration frameworks needed?
  - Eliminate the need for foundation model retraining
  - Overcoming token limits
  - Connectors for data sources
- Interface with any LLM using model I/O
  - Model I/O overview
  - Components of model I/O: Language models, chat models, prompts, example selectors, and output parsers
  - Overview of prompts, prompt templates, and example selectors
  - Different types of models: language, chat, and embedding models
  - Structuring language model responses using various types of output parsers
- Connecting external data with LLM application with retrieval
  - Retrieval overview
  - The rationale for the requirement of retrieval and how does it work with LangChain
  - Components of retrieval: Document loaders, text splitters, vector stores, and retrievers
  - Loading public, private, structured, and unstructured data with document loaders
  - Transforming documents to fewer chunks and extracting metadata using document transformers
  - Embedding and vector stores for converting documents into vectors and for efficient storage and retrieval
  - Optimizing retrieval using different retrieval techniques available in LangChain
- Creating complex LLM workflows with chains
  - Chains overview
  - Various foundational chain types: LLM, router, sequential, and transformation
  - Summarizing large documents using different document chains like stuff, refine, and map-reduce
- Retain context and refer to past interactions with the memory component
  - How memory can empower AI applications
  - Different types of memories: simple buffer memory, conversation summarization, vector-store-backed-memory
  - Overcoming token limit by using memory based on summarization of past conversations
  - Utilize vector stores for memory

## Agentic Workflows

- Dynamic decision-making with LLMs using agents
  - Agents overview
  - Components of agents: Tools, toolkits, prompt, and memory
  - Different types of agents: Self-ask with search, ReAct, JSON chat, structured chat
  - Working with agents using LangGraph
- Monitoring and logging using callbacks
  - Monitoring LLM application using callbacks
  - Understanding how callbacks work with different events
- Dynamic decision-making with LLMs using agents
  - Agents overview
  - Components of agents: Tools, toolkits, prompt, and memory
  - Different types of agents: Self-ask with search, ReAct, JSON chat, structured chat
  - Working with agents using LangGraph
- Monitoring and logging using callbacks
  - Monitoring LLM application using callbacks
  - Understanding how callbacks work with different events
- A Practical Guide to Coordinated LLM Agents Using LangGraph
  - Nodes (functions or agents)
  - Edges (data/control flow)
  - Cycles (iteration, self-correction)
  - State management (memory, messages, tools)
  - Create 2–3 simple agents (e.g., Researcher, Critic, Summarizer)
  - Use `Runnable` interfaces or `ToolRunnable` for integration
  - Add memory or context passing between agents

## Retrieval Augmented Generation (RAG)

In this module, we'll explore the challenges in developing RAG-based enterprise-level Large Language Model (LLM) applications. We will discuss the following:

- Basic RAG pipeline: Limitations of naïve approach
- Indexing: Chunking size optimization. Embedding Models
- Querying Challenges: Large Document Slices. Query Ambiguity
- Query Optimizations: Multi-Query Retrieval. Multi-Step Retrieval. Step-Back Prompting. Query Transformations
- Retrieval Challenges: Inefficient Retrieval of Large Documents. Lack of Conversation Context.
- Complex Retrieval from Multiple Sources.: Hybrid Search and Meta-data integration. Sentence window retrieval. Parent-child chunk retrieval. Hierarchical Index Retrieval. Hypothetical Document embeddings (HyDE).
- Generation Challenges: Information Overload. Insufficient Context Window. Chaotic Contexts. Hallucination. Inaccurate Responses.
- Generation Optimization: Information Compression. Thread of Thought (ThoT). Generator Fine-tuning. Adapter methods. Chain of Note (CoN). Expert Prompting
- Access control and governance

## LLM Evaluation

Dive into large language model (LLM) evaluation, examining its importance, common issues, benchmark datasets, and key metrics such as BLEU, ROUGE, and RAGAs, and apply these insights through a hands-on summarization exercise.

### Introduction to LLM evaluation

- What is evaluation and why is it important for LLMs?
- Overview of common mistakes made by LLMs
- A brief introduction to benchmark datasets and metrics
- Common LLM evaluation tasks

### Benchmark datasets

- Explore datasets for different tasks including natural language understanding, reasoning, knowledge retrieval, etc.
- Learn about different datasets such as MMLU, HELM, and BBH.

### Evaluation metrics

- Explain commonly used automatic metrics (BLEU, ROUGE, BERTScore)
- Compare strengths and weaknesses of different metrics
- Discuss the role of human evaluation and techniques (Likert scale)

### RAGAs

- Introduction and basic workflow
- Evaluation metrics: Faithfulness. Context precision. Answer relevancy. Context recall
- Detailed workflow stages
- Practical Applications: Summarization. Open-domain QA. Fact-checking

### Hands-on exercise

- Evaluating LLMs summarization using metrics like ROUGE, METEOR, and Bertscore
- Evaluation using G-Eval
- Evaluation of end-to-end RAG pipeline with RAGAs

## Capstone Project

On the last day of the LLM bootcamp, the learners will apply the concepts and techniques learned during the bootcamp to build an LLM application. Learners will choose to implement the following:

- Naive RAG assistant: A simple RAG assistant designed to answer general queries.
- Tool connection: An advanced agent that integrates with your data to provide more tailored responses.
- Connect web search tool: Extending your agent by adding web search and other sources.

Attendees will receive the following:

- Comprehensive Datasets: Access a vast collection of documents from a variety of industries to support your project's data needs and ensure robust functionality.
- Step-by-Step Implementation Guides: Detailed instructions that guide you through each phase of your project, from initial setup to final deployment.
- Ready-to-Use Code Templates: Utilize code templates available in Data Science Dojo's sandbox environments to streamline the development process and get your application up and running quickly.
- Cloud-Based Resources: Gain exclusive access to powerful cloud resources, including your own OpenAI key, facilitating the hassle-free deployment of your application on platforms like Streamlit.

At the culmination of the bootcamp, you will have a fully operational LLM application deployed on a public cloud platform, such as Streamlit. This deployment process includes setting up a continuous integration and continuous deployment (CI/CD) pipeline to ensure that your application can be updated and maintained effortlessly. By the end of the bootcamp, you'll be equipped not only with a finished project but also with the knowledge and skills to deploy and scale applications in real-world scenarios.

# Instructors and guest speakers

Learn from thought leaders at the forefront of building agentic AI applications



**Raja Iqbal**

Founder at **Ejento AI**, Founder at **Data Science Dojo**



**Luis Serrano**

Founder, Serrano Academy



**Sebastian Witalec**

Director of Developer Relations, Weaviate



**John Gilhuly**

Head of Developer Relations, Arize AI



**Kartik Talamadupula**

Head of AI, Wand AI



**Jerry Liu**

CEO/Co-founder, LlamaIndex



**Zain Hasan**

Senior DevRel Engineer, Together AI



**Sophie Daly**

Staff Data Scientist, Stripe



**Rehan Jalil**

Co-Founder | CEO, Securiti AI



**Sage Elliot**

AI Engineer, Union AI



**Adam Cowley**

Developer Advocate, Neo4j



**Hamza Farooq**

Founder, Travesaal AI

# Loved by customers and partners

More than **10,000 working professionals** have gone through our training program and recommend us.



"Partnering with Data Science Dojo aligns with our mission to make Data Science accessible. Their bootcamps contribute to safe AI deployment education."

- **ALESSYA VISNJIC**  
CEO & Co-founder, WhyLabs

"LLM Bootcamp provides hands-on experience with expert instructors, a comprehensive framework, and extensive resources for substantial upskilling."

- **BURNARD TUMANJONG**  
Chief, Operations Analysis, USA Airforce



"Collaborating with Data Science Dojo nurtures the next generation of LLM Developers commendable for fostering a dedicated creator community."

- **HARRISON CHASE**  
CEO & Co-founder, LangChains

"The bootcamp exceeded my expectations, offering comprehensive trainings for both beginners, and experienced data scientists. Highly recommended for all!"

- **SHAKEEB SYED**  
Data Scientist, Nationwide





"The bootcamp helped me fill the gap between academia and industry, offering insightful talks and practical problem-solving exercises."

**- YASHWANTH REDDY**

AI Leader, Fidelity Investments

"The bootcamp accelerated my text analysis skills by engaging in-person experience, industry insights, and valuable hands-on learning. Highly recommended!"

**- FLORIAN KLONEK**

Senior Lecturer, Deakin University



"Data Science Dojo's LLM bootcamp provides exceptional, hands-on initiation into LLMs & practical applications. A deep dive into the subject with theoretical knowledge, is highly recommended."

**- JEFF FONG**

Principal Product Manager, HOVER Inc.

"This bootcamp provided great content, equipping me with skills and confidence for efficient job execution. Enjoyable class and projects, a rewarding decision."

**- LUIS ARMANDO M**

Sales and Services Manager, RAISA



"The LLM Bootcamp offered an enlightening, well-structured learning experience, invaluable networking, & profound insights into real-world AI challenges. Highly recommend Raja's teaching approach."

**- AMITY FOX**

Data Scientist, 343 Industries

"Instructors simplified complex topics effectively. Hands-on learning enhanced my LLM understanding. Intensive but immeasurable rewards. A robust foundation gained in just five days."

**- FATEMEH SHAH-MUHAMMADI**

Research Assistant Professor, University of Utah School of Medicine





"A rewarding opportunity to help students of all backgrounds learn new techniques. Being part of a passionate and caring data science community was a real privilege."

**- SAGE ELLIOT**

AI Engineer | ML Ops & Robotics, WhyLabs

"The effort in providing resources was commendable. Teachers and assistants were helpful, making it one of the best courses. I learned a lot for future use."

**- FAHAD ALSWAINA**

Senior Data & Model Solutions Associate, KAPSARC



"Six months of course material access post-training is valuable. Instructors ensure success, making it a top-notch learning experience in AI."

**- ABRAR BHUIYAN**

Supervisor Software Application Engineer, ICF

"Top-notch speakers, hands-on workshops, and networking make it a wonderful tech experience. Raja's in-depth teaching focuses on learning concepts."

**- KEVIN SONG**

Software Engineer, ICF



"Outstanding collaboration on the LLM Bootcamp. Comprehensive curriculum in generative AI, prompt engineering, and data retrieval. Excited about practical training opportunities."

**- TAIMUR RASHID**

Advisor, Weaviate

"Seamless navigation, invaluable hands-on exercises, and well-structured technical aspects for optimal, cost-effective results. Highly recommended for practical LLM knowledge."

**- MARYAM BAGHERI**

Senior Data Scientist, Deloitte Business





“Comprehensive curriculum for in-depth understanding, seamless hands-on learning with cloud-based tools, and insightful talks by industry experts. Highly recommended.”

- **JARED MILLER**

Software Development Manager, Gurseey

“Outstanding boot camp, engaging discussions on coding and data security. A valuable and refreshing experience, recommended for advancing data science skills.”

- **KSHITIJ SINGH**

Machine Learning Lead, Trade Union



“Explored Data Science Dojo's LLM Bootcamp, gaining confidence with a comprehensive understanding. In-person experience, insightful instructors, and diverse insights.”

- **VICTOR GREEN**

Dev-Ops Engineer

“LLM Bootcamp surpassed expectations, bridging theory with practical examples. Engaging with cohorts, invaluable insights, and balanced intensity elevated understanding. Highly recommended for language models.”

- **KEN BUTLER**

Solutions Architect, ICF



“LLM Bootcamp enriched my understanding of the language-tech intersection. Collaborating with peers was wonderful. Sophie DA's insightful presentation on practical applications is recommended for foundational knowledge.”

- **ERIKA DAVIS**

Senior Business Reporting Analyst, Starbucks

“LLM Bootcamp shifted my problem-solving approach, sparking creativity with limitless applications. Highly recommended for anyone building their toolkit, fostering creativity in problem solvers, & providing invaluable insights for diverse applications.”

- **AISHWARIYA RAMAN**


Senior Software Engineer, Microsoft



# Certificate

Upon successful completion of the Large Language Model Bootcamp, you'll earn a verified certificate from The University of New Mexico Continuing Education, awarding 5 Continuing Education Units (CEUs). This certificate is recognized by employers for tuition reimbursement, valid for professional licensing renewal, and verifiable through The University of New Mexico Registrar's Office. You can also add it to your LinkedIn profile and share it with your professional network to showcase your achievement.

For the Agentic AI Bootcamp, participants receive a verified certificate with 3 CEUs, carrying the same benefits of recognition, validation, and shareability.



Note: Upon successful completion of the bootcamp, your verified digital certificate will be emailed to you using the name provided at registration. Certificate designs are for illustrative purposes only and may be subject to change at the discretion of The University of New Mexico Continuing Education and Data Science Dojo.

Data Science Dojo is collaborating with The University of New Mexico Continuing Education to offer a portfolio of high-impact bootcamps. These programs combine Data Science Dojo's expertise in practical AI and data science training with the academic excellence and credibility of a leading public university.

[Schedule a call](#) with an advisor to learn more.

You can apply to the program here

**Email:** [help@datasciencedojo.com](mailto:help@datasciencedojo.com)

**Phone:** +1 (877) 360-3442

**Apply Now**

In collaboration with

