



datasciencedojo
— data science for everyone —

Agentic AI Bootcamp

Learn to build enterprise-grade agentic ai applications



Course Overview

The Agentic AI Bootcamp by Data Science Dojo is an 9-week, hands-on program designed to teach professionals how to build intelligent, autonomous AI agents capable of reasoning, planning, and acting with minimal human input. Through a blend of live sessions and practical projects, participants will gain deep expertise in context-aware LLM application development, vector databases, multi-agent system design with LangGraph, and agentic workflows. The curriculum also covers agent communication protocols, interoperability, observability, and modern design patterns like ReAct and Reflection. Taught by leading experts from top AI organizations, the bootcamp equips developers, data scientists, and engineers with the tools and knowledge to create production-ready agentic AI applications.



Learning Outcomes

- Understand the foundations of agentic AI and how it extends traditional LLM applications with reasoning, memory, and safe design.
- Design and implement LLM-powered agents capable of planning, tool use, and multi-step reasoning.
- Apply prompt and workflow design to structure decision-making, reflection, and context-driven actions.
- Integrate external tools, APIs, Tavily search, and enterprise data via function calling, routing, and MCP.
- Use orchestration frameworks like LangChain and LangGraph to build modular, reliable agent pipelines.
- Build agents with memory and state awareness using Qdrant, vector stores, and hybrid retrieval loops.
- Evaluate, monitor, and debug agent behavior using tracing, observability tools, and metrics like RAGAs or G-Eval.
- Deploy and maintain production-ready multi-agent applications on platforms such as Streamlit or Docker.
- Mitigate risks in autonomous systems by applying safety guardrails, alignment strategies, and hallucination control.
- Translate business workflows into intelligent, interoperable agentic systems for automation and augmentation.

What you will learn

The program will enable you to:

- Build context-aware LLM applications that reason, retrieve, and act safely.
- Use vector databases like Qdrant to store embeddings and power RAG pipelines.
- Design multi-agent systems and collaborative flows with LangGraph.
- Create agentic workflows using node-based execution, planning, and memory layers.
- Implement agent communication and interoperability via MCP, A2A, and ACP protocols.
- Apply agentic design patterns such as ReAct, Reflection, Tree of Thoughts, and CodeAct.
- Establish observability and monitoring to track, debug, and optimise agentic workflows.



In today's AI-driven world, leaders who understand how to harness agentic workflows are better positioned to drive transformation, enhance productivity, and maintain a competitive edge. By mastering the principles of agentic AI, professionals can unlock new levels of automation and decision-making across business functions.

The Agentic AI Bootcamp offers a practical and strategic foundation in building AI agents using tools like LangChain and OpenAI. It covers planning and designing agent workflows, integrating APIs, deploying autonomous agents, and exploring real-world use cases—all while addressing key considerations such as evaluation, reliability, and scalability for enterprise environments.

Raja Iqbal

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

Program Highlights



Introduction to Agentic AI

Learn how reasoning-enabled agents extend LLMs with memory, RAG, and safe design patterns.



LangChain Fundamentals

Build modular LLM apps using prompts, retrieval, parsers, and LCEL-powered workflows.



Context Engineering

Design reliable, context-rich systems with LangGraph, memory layers, tools, and human oversight.



Vector Databases and Agentic RAG

Implement Qdrant and hybrid search to power RAG pipelines and scalable memory.



Agentic AI Protocols and Interoperability

Connect agents with MCP, A2A, and ACP for secure, cross-system collaboration.



Evaluation of Agents

Measure reliability, faithfulness, and text quality using RAGAs, G-Eval, and benchmark suites.



AgentOps: Operational Excellence for Managing Agents

Apply best practices for observability, tracing, and prompt governance to keep agents robust in production.



Final Project: Build a Multi-Agent LLM Application

Build and deploy a multi-agent LLM app—chatbot, RAG assistant, or MCP-powered agent—ready for real use.

Learning Journey

The Agentic AI Bootcamp offers a comprehensive, hands-on introduction to building such LLM-powered autonomous agents. By the end of this bootcamp, you'll be equipped to design, build, and deploy advanced agentic systems tailored to real-world tasks and business needs.

Core Modules | 9 Weeks

Introduction to Agentic AI

Learn the foundations of agentic systems and how reasoning-enabled agents extend standard LLMs. This module explores the principles, components, and memory strategies that make AI applications context-aware and reliable.

Key Topics:

- Foundations of Agentic AI: from next-token prediction to reasoning models, limits of classic LLMs vs reasoning LLMs, and the core pillars—reasoning, context, autonomy
- Understanding LLMs: context windows, session memory, long-term memory (vector DBs, graphs, summaries), and data sources (pre-training, fine-tuning, in-context learning)
- Retrieval-Augmented Generation (RAG): naïve RAG workflow and challenges, RAG as a context booster, preparing data for RAG pipelines
- Agentic AI Components: cognition (reasoning, planning, self-reflection), knowledge representation, autonomy (actions, tools, monitoring)
- Agentic Design Patterns: planning, tool use, reflection; Agentic RAG, router, loop; sequential, parallel, hierarchical flows
- Architectures for Agents: single-agent vs multi-agent systems, human-in-the-loop strategies, hybrid reasoning and decision graphs
- Advanced Context Techniques: session summaries, hybrid memory, Model Context Protocol (MCP), scalable context management
- Observability, Safety & Governance: guardrails, explainability, monitoring and evaluation, ethical alignment and compliance
- Hands-on exercises on reasoning, memory, RAG workflows, and safe agent design

LangChain Fundamentals

Learn to build robust, modular applications powered by large language models. This module introduces LangChain's key components and workflows, helping you assemble prompts, parsers, retrievers, and chains into maintainable, production-ready systems.

Key Topics:

- Introduction to LangChain: purpose and scope, building LLM-powered applications, challenges with RAG
- Core Components: LLMs and chat models, prompt templates and example selectors, document loaders and transformers

- Output Parsers: structured data extraction, consistent formatting, error handling
- Retrieval: embedding and vectorization, retrievers and metadata filtering, parent document retrieval
- Vector Stores: storing embeddings, efficient similarity search, optimizing for large datasets
- Chains: sequential prompt logic, pre/post-LLM steps, integrating tools and retrieval into chains
- Tool Use: integrating APIs and external actions, passing results back into workflows, handling errors and retries
- Domain specific language like LCEL : piping components with runnable, parallel branches, modular workflows
- Hands-on exercises on building retrieval chains, parsing structured outputs, and combining LangChain modules into coherent workflows

Context Engineering

Master the art of orchestrating context-aware, dependable agentic systems. This module teaches you how to structure prompts, manage memory, integrate tools, and coordinate agents using LangGraph and proven design strategies.

Key Topics:

- Complex Agentic Workflows: system prompts, user prompts, retrieved context; memory, web search, vector DBs; critique loops and reviewer models
- Chains: sequential control flows, pre/post-LLM steps, deterministic task execution
- Agents and Reliability: dynamic decision flows, router agents vs fully autonomous agents, balancing autonomy and reliability
- LangGraph Fundamentals: nodes, edges, state; condition-based execution; reliability for agent workflows
- Tool Integration: node-based tool calls, API connectors and database lookups, updating state after tool use
- Agentic Design Patterns: reflection for self-critique, tool use for external actions, planning for task decomposition
- Multi-Agent Collaboration: parallel, sequential, loop, and router flows; error checking and discussion; human-in-the-loop supervision
- Multi-Agent Architectures: hierarchical delegation, human approval nodes, shared resource tools
- Agentic Design Patterns – Hands-On Exercises:
- Reflection: Self-Discover framework for complex reasoning; iterative code refinement with LangGraph; stateful chatbot for prompt guidance
- Planning: Tree of Thoughts (ToT) for structured idea generation; DAG-based task compilation with LangGraph; efficient reasoning via parallel branches
- Tool Use: ReWOO (Reasoning Without Observation) for structured planning; function-calling for robust data extraction; handling unstructured inputs precisely
- Multi-Agent Collaboration: hierarchical agent teams for layered tasks; multi-agent supervisor with LangGraph; collaborative agent workflows

Vector Databases and Agentic RAG

Discover how to store, search, and optimize embeddings for scalable, reasoning-driven applications. This module explains vector database fundamentals, hybrid retrieval techniques, caching, and how to integrate retrieval-augmented generation (RAG) with agents.

Key Topics:

- Vector Database Fundamentals: embeddings and vector storage; ANN vs kNN search; modern vector DB architecture and data model
- Hybrid Retrieval Design: combining dense and sparse vectors; metadata filters and payload indexes; full-text tokenization for mixed queries
- Advanced Techniques: Maximal Marginal Relevance (MMR) for diversity; Discovery API for broader coverage; HNSW health monitoring and healing
- Agentic RAG Concepts: agents using AI native vector DBs as long-term memory; multi-step retrieval with reasoning loops; context selection and hallucination control
- Memory Tiers & Named Vectors: short-term, episodic, and long-term memory; multi-modal embeddings in one collection; modality-specific retrieval
- Semantic Caching: caching similar queries via vector similarity; TTL and invalidation policies; optimizing cost and latency
- APIs & Operations: search, scroll, recommendation, discovery; index optimization and quantization; deployment (OSS, Docker, Cloud)
- Monitoring & Observability: latency, recall, and throughput metrics; resource usage and drift detection; Web UI for search playground and ops
- Hands-on Exercises: explore AI-native vector database fundamentals, implement hybrid search, apply re-ranking techniques such as MMR, monitor HNSW index health, and build a RAG pipeline with agentic orchestration.

Agentic AI Protocols and Interoperability

Understand how protocols enable agents to communicate, discover capabilities, and coordinate tasks at scale. This module explains MCP, A2A, and ACP, highlighting their architectures and how to use them to connect tools, apps, and agents securely.

Key Topics:

- Multi-Agent Coordination: collaboration challenges with multiple agents; message routing and task management; scaling cooperation without chaos
- Need for Agentic Protocols: discovery and negotiation rules; task and state management; enabling secure cooperation
- Model Context Protocol (MCP): client-server architecture for LLM tools; standardized data and prompt access; tool, resource, and template exposure
- MCP Architecture: hosts, clients, and servers; message exchange and artifacts; connecting apps, IDEs, and assistants
- Agent-to-Agent Protocol (A2A): task-oriented communication flows; capability discovery with Agent Cards; message parts and artifact formats
- Agent Communication Protocol (ACP): open ecosystem for cross-agent exchange; routing, discovery, and dynamic updates; interoperability across frameworks
- MCP vs ACP vs A2A: scope and complexity comparison; message types and architecture; choosing protocols for different workflows
- Hands-on Exercise – MCP Client Using Streamlit: set up environment, install dependencies, and load your OpenAI API key; connect an MCP client to servers and discover available tools; automate workflows by invoking tools, retrieving data, and validating outputs

Evaluation of Agents

Learn how to assess the reliability, safety, and effectiveness of LLM-based agents. Explore benchmarking methods, text-quality metrics, and RAG-specific evaluation frameworks to ensure business, ethical, and technical alignment.

Key Topics:

- Need for Evaluation: reliability, accuracy, safety; business and ethical alignment; transparency and user trust
- Challenges in Evaluation: hallucinations, prompt sensitivity, weak context; subjectivity and multiple valid answers; trade-offs between accuracy, fluency, and creativity
- Benchmarking Approaches: MMLU for multitask accuracy; HELM for holistic metrics (accuracy, robustness, fairness); BBH & HotpotQA for reasoning and multi-hop QA
- Text Quality Metrics: BLEU for n-gram precision; ROUGE (ROUGE-N, ROUGE-L) for recall; BERTScore for semantic similarity
- RAG-Specific Evaluation (RAGAs): faithfulness and answer relevance; context precision and recall; joint retrieval + generation scoring
- G-Eval for Open-Ended Outputs: fluency and faithfulness; answer relevance; claim-level scoring
- Additional Benchmarks: GLUE for NLU tasks; TriviaQA for multi-hop QA; RealToxicityPrompts for safety and toxicity; BST for blended dialogue quality
- Other Metrics: perplexity for prediction confidence; METEOR for synonym and stem alignment; MRR and MAP for ranking performance; ROSCOE for reasoning quality (SA, SS, LI, LC)
- Hands-on Exercises: apply RAGAs to evaluate retrieval + generation pipelines, compare BLEU/ROUGE/BERTScore on generated text, and use G-Eval to assess open-ended agent responses

AgentOps: Operational Excellence for Managing Agents

Learn how to achieve operational excellence when managing agentic systems. Explore practices for prompt management, observability, tracing, and continuous improvement to ensure agents remain reliable, transparent, and aligned with organizational, ethical, and user goals.

Key Topics:

- Need for Operational Excellence: reliability, accuracy, safety; business and ethical alignment; transparency and user trust as ongoing goals of AgentOps
- Challenges in Operations: hallucinations and prompt sensitivity; brittle reasoning and weak context; subjectivity and multiple valid answers; trade-offs between accuracy, fluency, creativity, cost, and latency
- Prompt & Policy Management: version control, modular structures, and governance; managing updates to prompts, instructions, and policies to reduce drift and ensure consistency
- Observability & Tracing: logging agent decisions and actions; capturing inputs, outputs, and intermediate steps; tracing multi-step reasoning and tool use for debuggability and accountability
- Operational Benchmarks: tracking success rates, error recovery, efficiency, and user satisfaction; aligning operational metrics with business KPIs
- Lifecycle Practices: deployment pipelines with staging and production; continuous monitoring for drift and degradation; A/B testing and controlled rollouts; incorporating human feedback for sustainable improvement
- Hands-on Exercise: Learn how to manage and version prompts, and how to debug, monitor, and trace AI applications.

Final Project: Build a Multi-Agent LLM Application

Apply everything you've learned by designing and deploying a complete multi-agent LLM application. This capstone gives you structured guidance, datasets, and templates so you can deliver a production-grade solution that showcases reasoning, retrieval, tool use, and interoperability.

Key Topics:

- Project Tracks:
 - Conversational Workflow Orchestration: design a multi-turn assistant that coordinates tasks across multiple specialized agents
 - Knowledge-Enhanced Agent: integrate external search or APIs for fact-checking, grounding, and real-time information use
 - Document-Aware Action Agent: build an agent that retrieves and reasons over documents, then triggers external tools or services to act on insights
 - Orchestrated Collaboration (with MCP): build a system where agents seamlessly assume specialized roles and coordinate through the Model Context Protocol (MCP), enabling effortless integration with external tools, data sources, and enterprise systems.
- Attendees Will Receive:
 - Comprehensive Datasets: a broad collection of documents across industries to support data needs and ensure robust functionality
 - Step-by-Step Implementation Guides: detailed instructions for each phase, from setup to deployment
 - Ready-to-Use Code Templates: prebuilt templates in Data Science Dojo's sandbox environment for faster development
- Learners Can Choose to Implement:
 - Virtual Assistant
 - Content Generation (Marketing Co-pilot)
 - Conversational Agent (Legal & Compliance Assistant)
 - Q&A Bot (IRS Tax Advisor)
 - Content Personalizer
 - MCP Chatbot – AI agent with calendar, CRM, and API integrations

Outcome:

By the end of this project, you'll have a fully functional, production-ready multi-agent application that demonstrates your mastery of reasoning, retrieval, tool use, and protocol-driven interoperability.

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



Zain Hasan

Senior DevRel Engineer, Together AI



Kartik Talamadupula

Head of AI, Wand AI



John Gilhuly

Head of Developer Relations, Arize AI



Sage Elliot

AI Engineer, Union AI



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."

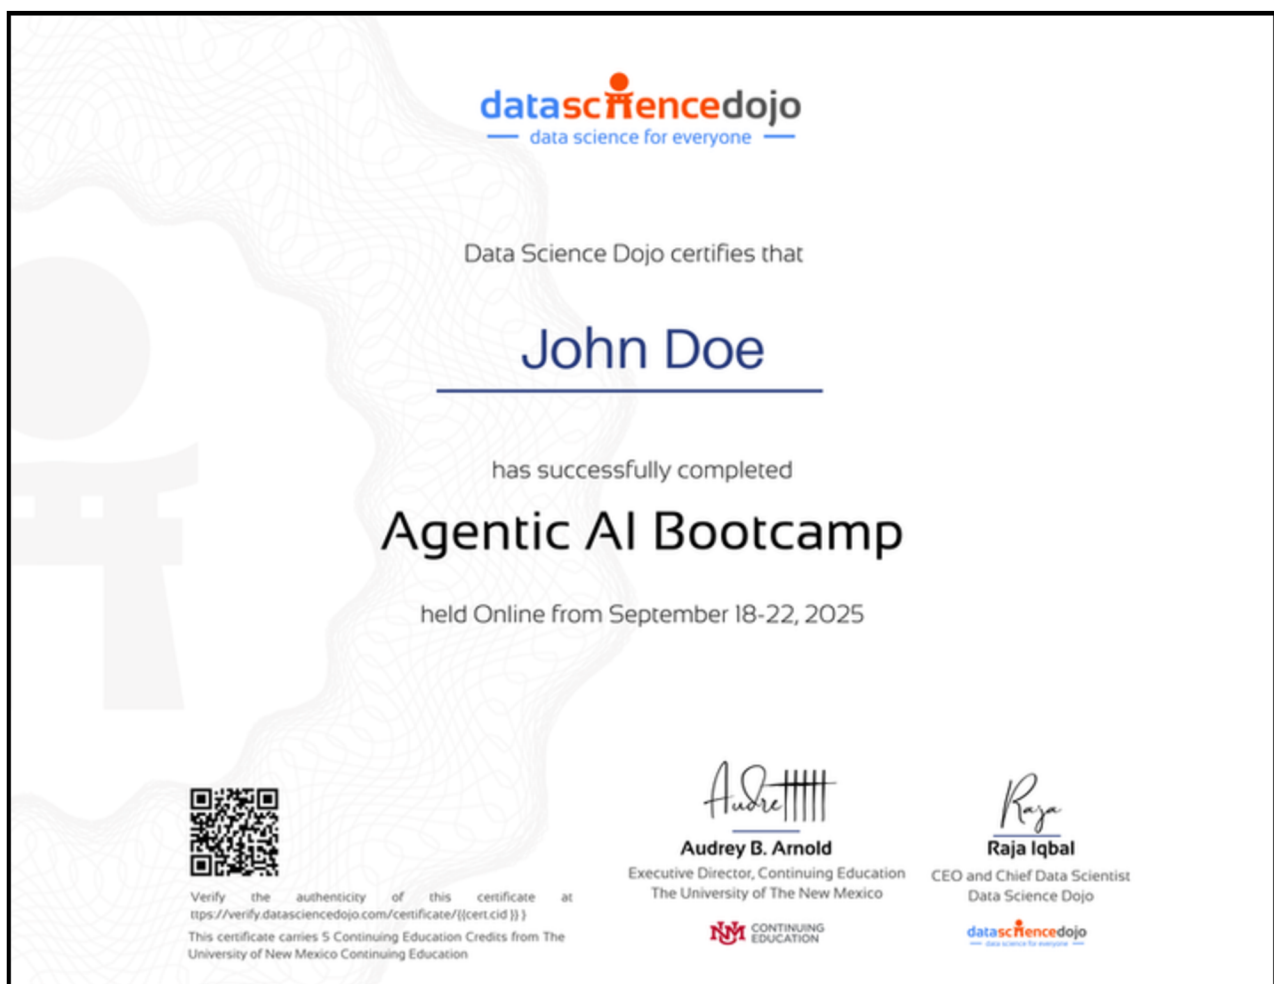
- AISHWARIYA RAMAN

Senior Software Engineer, Microsoft



Certificate

Upon successful completion of the 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

Apply Now

Email: help@datasciencedojo.com

Phone: +1 (877) 360-3442

In collaboration with

