



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 10-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.
- 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 Reflection, Planning, and Tool Use.
- 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



Transformers & Attention Mechanism

Develop a clear, hands-on understanding of LLMs—from transformers and attention to embeddings, fine-tuning, and RAG.



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.



Vector Databases and Agentic RAG

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



Context Engineering

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



Agentic Design Patterns

Learn agentic design patterns that improve LLM performance by structuring reasoning, control, and coordination.



Agentic AI Protocols

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



Model Context Protocol

Learn how the Model Context Protocol (MCP) standardizes and secures LLM access to tools, data, and context.



Evaluation of Agents

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



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 | 10 Weeks

Transformers & Attention Mechanism

Key Topics:

- Introduction to LLMs: Strengths and weaknesses of large language models.
- Discriminative versus Generative AI: Predictive models contrasted with generative models.
- Transformer Architecture: Tokenization, embeddings, positional encoding, and attention.
- Embeddings and Similarity: Representing words as vectors and measuring closeness.
- Attention Mechanism: Keys, queries, and values in self-attention.
- Softmax and Probabilities: Converting scores to probabilities for next-word prediction.
- Training and Fine-Tuning: Adapting models with curated data for new tasks.
- Search and Retrieval: Building a semantic search engine with embeddings.
- Retrieval Augmented Generation (RAG): Combining retrieval with generation for grounded answers.
- Hands-On Exercises: Sentence Transformers, semantic search, attention scoring, and attention mechanisms.

Introduction to Agentic AI

Key Topics:

- Foundations of Agentic AI: From next-token prediction to reasoning models; limitations of classic LLMs vs reasoning LLMs; the core pillars of agentic systems—reasoning, context, and autonomy.
- Understanding LLMs: Context windows, session memory, and long-term memory (vector databases, knowledge graphs, summaries); data sources including pre-training, fine-tuning, and in-context learning.
- Retrieval-Augmented Generation (RAG): Naïve RAG workflows and common challenges; RAG as a context enhancement strategy; preparing and structuring data for effective RAG pipelines.
- Agentic AI Components: Cognition (reasoning, planning, self-reflection), knowledge representation, and autonomy through tool use, action execution, and monitoring.

- Agentic Design Patterns: Planning, tool use, and reflection loops; Agentic RAG, routers, and iterative loops; sequential, parallel, and hierarchical workflows.
- Architectures for Agents: Single-agent vs multi-agent systems; human-in-the-loop strategies; hybrid reasoning pipelines and decision graphs.
- Advanced Context Techniques: Session summaries, hybrid memory systems, Model Context Protocol (MCP), and scalable context management.
- Observability, Safety & Governance: Guardrails, explainability, monitoring and evaluation, ethical alignment, and compliance strategies.
- Hands-On Exercises: Practical implementations of reasoning workflows, memory systems, RAG pipelines, and safe agent design.

Mastering LangChain

Key Topics:

- Introduction to LangChain: Purpose and scope of LangChain; building LLM-powered applications; common challenges in implementing Retrieval-Augmented Generation (RAG).
- Core Components: LLMs and chat models; prompt templates and example selectors; document loaders and transformers for preprocessing data.
- Output Parsers: Extracting structured data; enforcing consistent output formats; handling parsing errors and validation failures.
- Retrieval: Embedding and vectorization strategies; retrievers and metadata filtering; parent document retrieval for contextual completeness.
- Vector Stores: Storing embeddings efficiently; performing scalable similarity search; optimizing retrieval for large datasets.
- Chains: Sequential prompt logic; pre- and post-LLM processing steps; integrating retrieval and tool use into end-to-end chains.
- Tool Use: Connecting APIs and external systems; feeding tool outputs back into workflows; managing retries and error handling.
- LangChain Expression Language (LCEL): Building modular workflows using runnable components; piping operations; parallel branches and composable pipelines.
- Hands-On Exercises: Constructing retrieval chains; parsing structured outputs; combining LangChain modules into coherent, production-ready workflows.

Vector Databases and Agentic RAG

Key Topics:

- Vector Database Fundamentals: Embeddings and vector storage; approximate nearest neighbor (ANN) vs k-nearest neighbor (kNN) search; modern vector database architectures and data models.
- Hybrid Retrieval Design: Combining dense and sparse vectors; applying metadata filters and payload indexing; full-text tokenization for mixed semantic and keyword queries.
- Advanced Techniques: Maximal Marginal Relevance (MMR) for improving result diversity; Discovery APIs for broader coverage; monitoring and maintaining HNSW index health.
- Agentic RAG Concepts: Using AI-native vector databases as long-term memory for agents; multi-step retrieval with reasoning loops; context selection strategies and hallucination mitigation.

- Semantic Caching: Caching semantically similar queries using vector similarity; time-to-live (TTL) and invalidation policies; optimizing cost and latency.
- Hands-On Exercises: Exploring AI-native vector database fundamentals; implementing hybrid search; applying re-ranking techniques such as MMR; monitoring HNSW index health; building a RAG pipeline with agentic orchestration.

Context Engineering

Key Topics:

- Complex Agentic Workflows: Designing workflows with system and user prompts; integrating retrieval, memory layers, web search, and vector databases; implementing critique and refinement loops.
- Deterministic Chains and Control Flows: Building sequential pipelines with pre- and post-LLM steps; enforcing structured task execution and predictable control logic.
- Agent Reliability and Dynamic Decisions: Using router agents and conditional flows; balancing autonomy with control for robust task execution.
- LangGraph Fundamentals: Understanding nodes, edges, and state management; condition-based execution for reliable and auditable workflows.
- Tool Integration: Connecting APIs, databases, and external systems through node-based tool calls; updating shared state after execution.
- Agentic Design Patterns: Reflection for self-critique; tool use for external actions; planning for task decomposition and structured reasoning.
- Multi-Agent Collaboration: Implementing parallel, sequential, loop, and router flows; incorporating error handling and human supervision strategies.
- Multi-Agent Architectures: Designing hierarchical delegation systems; approval nodes; shared memory and resource coordination.
- Hands-On Experience: Practical implementation of advanced context engineering concepts across agentic workflows and multi-agent systems.

Agentic Design Patterns

Key Topics:

- Why Agentic Patterns Matter: Transforming single-pass prompting into iterative, goal-oriented reasoning loops for more adaptive and reliable systems.
- Reflection Pattern: Enabling agents to evaluate, critique, and refine their own outputs through structured feedback and revision cycles.
- Planning Pattern: Designing stepwise reasoning flows that decompose complex goals, manage task dependencies, and adapt dynamically to new information.
- Tool Use Pattern: Connecting models with external systems to retrieve data, execute actions, and extend problem-solving capabilities beyond the model's internal knowledge.
- Multi-Agent Collaboration Pattern: Coordinating specialized agents with defined roles; enabling structured communication and collaborative problem-solving.
- Pattern Trade-Offs: Balancing autonomy with control, flexibility with stability, and creativity with reliability in agent design.
- Pattern Composition: Integrating reflection, planning, and tool use into hybrid workflows for building more capable and adaptive agents.
- Hands-On Labs: Implementing individual patterns and combining them into complete, production-ready agent workflows for real-world tasks.

Agentic AI Protocols

Key Topics:

- Multi-Agent Coordination: Collaboration challenges in multi-agent systems; message routing and task orchestration; scaling cooperation while maintaining stability and control.
- Need for Agentic Protocols: Establishing discovery and negotiation mechanisms; structured task and state management; enabling secure and reliable agent cooperation.
- Model Context Protocol (MCP): Client-server architecture for LLM tool integration; standardized access to data and prompts; exposing tools, resources, and templates in a structured way.
- MCP Architecture: Roles of hosts, clients, and servers; message exchange formats and artifacts; connecting applications, IDEs, and assistants within unified workflows.
- Agent-to-Agent Protocol (A2A): Task-oriented communication flows; capability discovery using Agent Cards; structured message parts and artifact formats for coordination.
- Agent Communication Protocol (ACP): Open ecosystem for cross-agent interaction; routing, discovery, and dynamic updates; interoperability across frameworks and platforms.
- MCP vs ACP vs A2A: Comparing scope, architectural complexity, and message types; selecting appropriate protocols for different workflow requirements.
- Hands-On Exercise – MCP Client with Streamlit: Setting up the development environment; installing dependencies and configuring API access; connecting an MCP client to servers; discovering tools; automating workflows through tool invocation, data retrieval, and output validation.

Model Context Protocol

Key Topics:

- Origins & Motivation: Addressing fragmented integrations and brittle bespoke adapters; introducing a unified, interoperable interface – the “USB-C for AI.”
- Protocol Structure: Client-server handshake model; defining resources, tools, and prompts; JSON-RPC transport with structured, schema-driven messages.
- Context Exposure: How MCP surfaces tools, data, and metadata through a consistent schema to enable discoverability, governance, and controlled access.
- Agentic Integration: Connecting MCP endpoints to reflection, planning, tool-use, and multi-agent coordination patterns for modular and scalable systems.
- Hands-On Labs: Setting up an MCP client in Streamlit; discovering and registering tools; automating workflows through data retrieval and validation; logging traces for monitoring and review.

Evaluation of Agents

Key Topics:

- Need for Evaluation: Reliability, accuracy, and safety; business and ethical alignment; transparency and user trust.
- Challenges in Evaluation: Hallucinations and prompt sensitivity; weak context grounding; subjectivity and trade-offs between accuracy, fluency, and creativity.
- Benchmarking Approaches: MMLU for multitask accuracy; HELM for robustness and fairness; BBH and HotpotQA for reasoning and multi-hop QA.
- Text Quality Metrics: BLEU for precision; ROUGE for recall; BERTScore for semantic similarity.
- RAG Evaluation (RAGAs): RAGAs for faithfulness and answer relevance; context precision and recall; joint retrieval-generation scoring.
- G-Eval: G-Eval for fluency, faithfulness, relevance; claim-level scoring of open-ended outputs.
- Additional Benchmarks: GLUE for NLU; TriviaQA for QA; RealToxicityPrompts for safety; Blended Skill Talk for dialogue quality.
- Other Metrics: Perplexity for confidence; METEOR for alignment; MRR and MAP for ranking; ROSCOE for reasoning quality.
- Hands-On Exercises: Apply RAGAs to RAG pipelines; compare BLEU, ROUGE, BERTScore; evaluate agent outputs using G-Eval.

Final Project: Build a Multi-Agent LLM Application

Key Topics:

Project Tracks:

- Conversational Workflow Orchestration: Design a multi-turn assistant coordinating tasks across specialized agents.
- Knowledge-Enhanced Agent: Integrate search and APIs for grounding, fact-checking, and real-time data access.
- Document-Aware Action Agent: Retrieve and reason over documents; trigger external tools or services based on insights.
- Orchestrated Collaboration (MCP): Build coordinated multi-agent systems using the Model Context Protocol for seamless tool and enterprise integration.

Attendees Will Receive:

- Comprehensive Datasets: Industry-spanning document collections for robust development and testing.
- Step-by-Step Implementation Guides: Clear instructions from environment setup to deployment.
- Ready-to-Use Code Templates: Prebuilt templates within Data Science Dojo's sandbox for accelerated development.

Learners Can Choose to Implement:

- Virtual Assistant
- Content Generation (Marketing Co-pilot)
- Conversational Agent (Legal & Compliance Assistant)
- Content Personalizer
- MCP Chatbot – AI agent with calendar, CRM, and API integrations

Outcome:

A production-ready multi-agent application demonstrating 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



Shub Argha

Head of Solutions
Engineering Arcade.dev



Thierry Damiba

Developer Advocate,
Qdrant



**Kartik
Talamadupula**

Head of AI, Wand 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.”

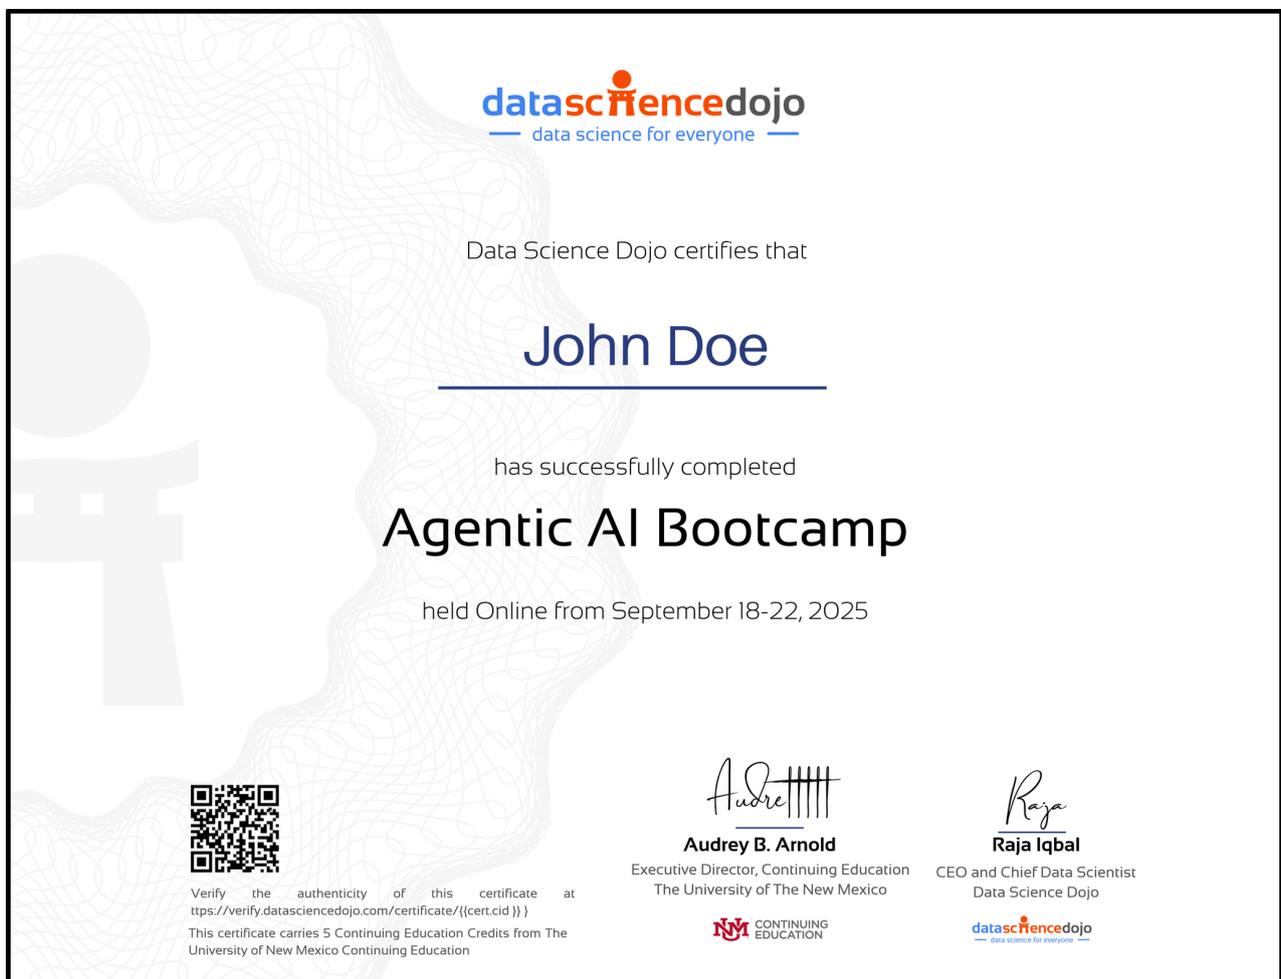
- AISHWARIYA RAMAN

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

Email: help@datasciencedojo.com

Phone: +1 (877) 360-3442

[Apply Now](#)

In collaboration with

