EDVANTA

Ignite your growth, inspire your mind

12-WEEK PROGRAM

AI & ML Mastery: From Foundations to Agentic AI

A comprehensive 3-month online course designed for early professionals, graduate students, and career-transitioners ready to break into the world of Artificial Intelligence, Machine Learning, and Generative AI — including the cutting-edge frontier of Agentic AI systems.

Foundation Phase	Intermediate Phase	Advanced Phase
Weeks 1–4: Core concepts, math, and ML fundamentals	Weeks 5–8: Deep Learning, NLP, and Gen AI	Weeks 9–12: Agentic AI, deployment, and capstone

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

Why This Course Matters Today

AI and Generative AI are no longer niche technologies — they are the operating system of the modern economy. Every major industry is actively recruiting professionals who can build, fine-tune, and deploy intelligent systems.

The Market Reality

- AI job postings grew **over 74%** in the last 2 years
- Gen AI roles command 20–40% salary premiums
- Agentic AI is the #1 emerging skill per LinkedIn 2024
- Every Fortune 500 has an active AI transformation agenda

What Sets This Program Apart

- Industry-aligned curriculum updated for 2024–25
- Hands-on tools: Python, LangChain, OpenAI API, HuggingFace
- Live mentoring + recorded weekend sessions
- Real capstone project simulating industry problems
- Career mapping from Day 1 through job readiness

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

Program Structure at a Glance

Each of the 12 weeks follows a consistent rhythm designed for working learners and students — combining self-paced learning with live mentoring support.

Weekend Recorded Session

Concept-focused video lectures covering theory, tools, and demonstrations. Available on-demand so you can revisit anytime.

Live Mentoring Session

Weekly 1:1 or small-group live sessions for doubt-clearing, case discussions, code reviews, and personalized feedback.

Weekly Assignment

Practical exercises and mini-projects aligned to weekly outcomes. Graded with instructor feedback to reinforce concepts.

Learning Outcomes

Each week ends with clear, measurable outcomes so you always know what you've mastered and what comes next.

EDVANTA

EDVANTA

FOUNDATION PHASE

Week 1: Introduction to AI & the ML Landscape

This first week sets the stage — students build a mental model of the entire AI/ML ecosystem and establish their Python environment.

🎯 Learning Objectives

- Understand the AI/ML/DL/Gen AI hierarchy
- Distinguish supervised, unsupervised, and reinforcement learning
- Set up a complete Python ML environment
- Navigate the AI job market landscape
- Understand ethical considerations in AI
- Map tools to use cases across industries

📋 Key Topics

- AI vs. ML vs. Deep Learning vs. Gen AI — clear distinctions
- Types of ML: Supervised, Unsupervised, Reinforcement
- Python for ML: NumPy, Pandas, Matplotlib basics
- Jupyter Notebooks and Google Colab setup
- Industry verticals adopting AI (healthcare, finance, retail)
- AI ethics, bias, and responsible deployment

🛠 Tools

Python 3.x, Anaconda, Google Colab, Jupyter Notebook

📝 Assignment

Write a 1-page AI use case analysis for an industry of your choice and set up your working Colab environment.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

FOUNDATION PHASE

Week 2: Mathematics for Machine Learning

No shortcuts here — week 2 builds the mathematical intuition that makes ML algorithms understandable, not just usable. Taught with visual examples and practical code.

🎯 Learning Objectives

- Apply linear algebra concepts to ML problems
- Understand calculus fundamentals behind gradient descent
- Use probability and statistics for data interpretation
- Implement key math operations using NumPy
- Understand cost functions and optimization intuitively
- Distinguish underfitting and overfitting mathematically

📋 Key Topics

- Vectors, matrices, dot products, and matrix multiplication
- Derivatives, chain rule, and partial derivatives
- Gradient descent: intuition and implementation
- Probability distributions, Bayes' theorem
- Mean, variance, standard deviation, and covariance
- Loss functions: MSE, Cross-Entropy explained

🛠️ Tools

NumPy, Matplotlib, SymPy (symbolic math), Khan Academy references

📝 Assignment

Implement gradient descent from scratch to minimize a quadratic cost function. Visualize the loss curve.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

FOUNDATION PHASE

Week 3: Data Preparation & Exploratory Data Analysis

Real-world data is messy. This week teaches learners to clean, explore, and engineer features — the most time-consuming part of any ML project (often 60–70% of total effort).

🎯 Learning Objectives

- Perform end-to-end EDA on a real dataset
- Handle missing values, outliers, and duplicates
- Engineer meaningful features from raw data
- Apply encoding for categorical variables
- Normalize and scale numerical features correctly
- Build visual dashboards for data storytelling

📋 Key Topics

- Pandas: data loading, slicing, groupby, merge operations
- Missing data strategies: imputation, deletion, flagging
- Outlier detection: IQR, Z-score methods
- Feature encoding: Label, One-Hot, Target encoding
- Feature scaling: StandardScaler, MinMaxScaler
- Seaborn and Plotly for EDA visualizations

🛠 Tools

Pandas, Seaborn, Matplotlib, Plotly, Scikit-learn preprocessing

📝 Assignment

Perform complete EDA on the Titanic or Housing Prices dataset. Submit a visual report with at least 8 charts and a feature engineering notebook.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

FOUNDATION PHASE

Week 4: Supervised Learning — Regression & Classification

Week 4 is where learners build and evaluate their first real ML models. Emphasis is on understanding the why behind algorithms, not just running library calls.

🎯 Learning Objectives

- Build and interpret Linear and Logistic Regression models
- Implement Decision Trees and Random Forests
- Use K-Nearest Neighbors for classification
- Evaluate models using appropriate metrics
- Apply cross-validation and hyperparameter tuning
- Understand the bias-variance tradeoff in practice

📋 Key Topics

- Linear Regression: assumptions, coefficients, R^2 score
- Logistic Regression: sigmoid, log-loss, decision boundary
- Decision Trees: Gini impurity, information gain, pruning
- Random Forests and ensemble learning basics
- Metrics: Accuracy, Precision, Recall, F1, ROC-AUC
- Cross-validation: k-fold, stratified k-fold

🛠 Tools

Scikit-learn, GridSearchCV, Yellowbrick (model visualization)

📝 Mini Project

Predict customer churn using a telecom dataset. Compare at least 3 algorithms. Present a model evaluation report.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

CHAPTER 2

Intermediate Phase: Weeks 5–8

With foundations locked in, the intermediate phase elevates learners into deep learning, natural language processing, and their first exposure to Generative AI. Industry tools and real-world workflows take center stage.

1

Week 5

Unsupervised Learning & Advanced Supervised Techniques

2

Week 6

Neural Networks & Deep Learning Foundations

3

Week 7

NLP & Transformers

4

Week 8

Introduction to Generative AI

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

INTERMEDIATE PHASE

Week 5: Unsupervised Learning & Advanced Supervised Techniques

🎯 Learning Objectives

- Apply K-Means and DBSCAN clustering algorithms
- Reduce dimensionality using PCA and t-SNE
- Implement SVM for classification and regression
- Use XGBoost and Gradient Boosting effectively
- Understand anomaly detection techniques
- Select the right algorithm for business problems

📋 Key Topics

- K-Means clustering: elbow method, silhouette score
- DBSCAN for density-based clustering
- PCA: variance explained, scree plots, dimensionality reduction
- t-SNE for high-dimensional data visualization
- SVM: kernel trick, margin maximization, C parameter
- XGBoost and LightGBM: boosting, feature importance

🛠️ Tools

Scikit-learn, XGBoost, LightGBM, UMAP

📝 Assignment

Perform customer segmentation on an e-commerce dataset using K-Means. Interpret and name each customer segment based on behavior patterns.

EDVANTA

Week 6: Neural Networks & Deep Learning Foundations

🎯 Learning Objectives

- Explain the architecture of a neural network
- Implement forward and backward propagation
- Build multi-layer perceptrons using TensorFlow/Keras
- Apply activation functions strategically
- Use regularization to prevent overfitting
- Train CNNs for image classification tasks

📋 Key Topics

- Perceptron to MLP: layers, weights, biases
- Activation functions: ReLU, Sigmoid, Softmax, GELU
- Backpropagation and optimizer comparison (SGD, Adam, RMSProp)
- Batch normalization and Dropout regularization
- Convolutional Neural Networks: filters, pooling, feature maps
- Transfer learning with pre-trained models (VGG16, ResNet)

🛠 Tools

TensorFlow 2.x, Keras, TensorBoard, Google Colab GPU

📝 Assignment

Build an image classifier using CNN on CIFAR-10 dataset. Apply transfer learning and compare accuracy vs. training from scratch.

Week 7: Natural Language Processing & Transformers

🎯 Learning Objectives

- Apply classical NLP pipelines for text processing
- Understand word embeddings and semantic similarity
- Explain the Transformer architecture intuitively
- Fine-tune BERT for text classification
- Build a text summarization or sentiment pipeline
- Understand tokenization and attention mechanisms

📋 Key Topics

- Text preprocessing: tokenization, stemming, lemmatization, stopwords
- TF-IDF and Bag of Words representations
- Word2Vec, GloVe, FastText word embeddings
- Attention mechanism: self-attention, multi-head attention
- Transformer architecture: encoder-decoder explained
- BERT, RoBERTa, DistilBERT: pre-training vs. fine-tuning

🛠 Tools

HuggingFace Transformers, SpaCy, NLTK, Sentence-Transformers

📝 Assignment

Fine-tune DistilBERT on a sentiment analysis dataset (IMDB or SST-2). Report F1 score. Compare with a baseline TF-IDF + Logistic Regression model.

EDVANTA

Ignite your growth, inspire your mind

INTERMEDIATE PHASE

Week 8: Introduction to Generative AI

🎯 Learning Objectives

- Distinguish generative from discriminative models
- Understand how LLMs are trained at a conceptual level
- Apply prompt engineering to real business scenarios
- Use OpenAI and open-source LLM APIs effectively
- Explore image generation with diffusion models
- Identify Gen AI limitations, hallucinations, and risks

📋 Key Topics

- GANs, VAEs, and Diffusion Models — how they generate content
- Large Language Models: GPT architecture and training overview
- Prompt engineering: zero-shot, few-shot, chain-of-thought
- System prompts, temperature, and token limits
- OpenAI API: completions, chat, embeddings endpoints
- Stable Diffusion, DALL-E: text-to-image generation

🛠️ Tools

OpenAI API, HuggingFace, Stability AI, Ollama (local LLMs)

📝 Mini Project

Build a prompt-engineered AI assistant for a specific domain (HR, legal, finance). Demonstrate zero-shot, few-shot, and CoT prompting. Evaluate output quality.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

CHAPTER 3

Advanced & Application Phase: Weeks 9–12

The final phase is where everything comes together. Learners move from understanding Gen AI to building production-grade systems — including RAG pipelines, Agentic AI frameworks, and full deployment workflows.

1

Week 9

RAG & Vector Databases

2

Week 10

Agentic AI & Multi-Agent Systems

3

Week 11

MLOps & Model Deployment

4

Week 12

Capstone Project & Career Launch

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

ADVANCED PHASE

Week 9: RAG Pipelines & Vector Databases

🎯 Learning Objectives

- Explain why RAG solves LLM hallucination problems
- Build an end-to-end RAG pipeline from scratch
- Index and query documents using vector databases
- Use embedding models for semantic search
- Implement chunking strategies for document processing
- Evaluate RAG output quality with relevance metrics

📋 Key Topics

- Retrieval-Augmented Generation: architecture and flow
- Embeddings: how they encode semantic meaning
- Vector databases: Pinecone, ChromaDB, Weaviate, FAISS
- Document ingestion: chunking, overlapping, and indexing
- LangChain RetrievalQA chain construction
- RAG evaluation: faithfulness, relevance, context precision

🛠 Tools

LangChain, ChromaDB, Pinecone, OpenAI Embeddings, RAGAS

📝 Assignment

Build a RAG-based Q&A system over a 50-page PDF document (company policy, research paper, or product manual). Evaluate with at least 20 test queries.

EDVANTA

Ignite your growth, inspire your mind

ADVANCED PHASE

Week 10: Agentic AI & Multi-Agent Systems

Agentic AI is the frontier — systems that can plan, reason, use tools, and collaborate with other agents to complete complex tasks autonomously.

🎯 Learning Objectives

- Understand the architecture of AI agents and tool use
- Build a ReAct-style reasoning agent with LangChain
- Design and orchestrate multi-agent workflows
- Implement memory types: short-term, long-term, episodic
- Use AutoGen or CrewAI for agent collaboration
- Define guardrails and safety boundaries for agents

📋 Key Topics

- AI Agent components: Perception, Memory, Planning, Action
- ReAct framework: Reasoning + Acting loop
- Tool use: web search, code execution, API calls, databases
- LangChain Agents: AgentExecutor, custom tools, callbacks
- Multi-agent orchestration with CrewAI and AutoGen
- Agent safety: hallucination guards, tool validation, human-in-the-loop

🛠 Tools

LangChain Agents, CrewAI, AutoGen, OpenAI Function Calling, Tavily Search

📝 Assignment

Build a 3-agent crew: Researcher, Analyst, and Writer. Task them to generate a market intelligence report on a company. Demonstrate tool use and inter-agent communication.

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

ADVANCED PHASE

Week 11: MLOps, Fine-Tuning & Model Deployment

🎯 Learning Objectives

- Understand the full ML model lifecycle in production
- Fine-tune open-source LLMs using PEFT/LoRA
- Track experiments with MLflow and Weights & Biases
- Package models and deploy with FastAPI and Docker
- Monitor deployed models for drift and degradation
- Implement CI/CD pipelines for ML workflows

📋 Key Topics

- MLOps maturity model: from notebooks to production
- Fine-tuning with LoRA and QLoRA on HuggingFace
- Experiment tracking: MLflow, Weights & Biases
- Model serving: FastAPI REST endpoints, Streamlit demos
- Containerization with Docker for reproducibility
- Model monitoring: data drift, concept drift, alerting

🛠 Tools

HuggingFace PEFT, MLflow, FastAPI, Docker, Streamlit, Evidently AI

📝 Assignment

Fine-tune a small open-source LLM (Mistral 7B or Phi-2) on a domain-specific dataset. Deploy as a FastAPI endpoint. Create a Streamlit demo interface.

EDVANTA

Industry Example 1: AI-Powered Fraud Detection at a FinTech Firm

The Business Problem

A mid-sized FinTech company was losing **$2.3M annually** to fraudulent transactions. Their rules-based system had an 18% false positive rate, frustrating legitimate customers and increasing churn.

How This Course Applies

Weeks 3–5 skills: EDA revealed behavioral anomalies. XGBoost (Week 5) outperformed the rules system. An anomaly detection layer using Isolation Forest was added to flag edge cases in real time.

Tools & Technologies

- XGBoost + Isolation Forest for detection
- Pandas + Feature engineering for behavioral signals
- FastAPI for real-time scoring endpoint
- MLflow for experiment tracking and model registry

Measurable Business Impact

- False positive rate reduced from 18% → 4.2%
- Fraud losses reduced by 67% in Year 1
- Customer complaint tickets dropped by 31%

Role of Early Professionals

Junior ML engineers built the feature pipeline, ran model comparisons, and owned the FastAPI deployment. They reported directly to a senior data scientist.

Industry Example 2: RAG-Based Internal Knowledge Assistant at a Consulting Firm

The Business Problem

A 3,000-person consulting firm had knowledge scattered across 40,000+ documents, proposals, and case studies. Consultants spent an average of **4.5 hours/week** searching for relevant precedents — a massive productivity drain.

How This Course Applies

Week 9 RAG skills were directly applied: documents were chunked, embedded, and indexed into Pinecone. A GPT-4-powered query interface allowed natural language search across the entire knowledge base.

Tools & Technologies

- LangChain for the RAG pipeline orchestration
- Pinecone as the vector store
- OpenAI GPT-4 for generation
- Streamlit for the internal chat UI

Measurable Business Impact

- Search time reduced by 78% per consultant
- Proposal win rate improved by 12% (faster, better-informed bids)
- ROI achieved in under 5 months post-deployment

Role of Early Professionals

Two junior AI engineers built the ingestion pipeline and Streamlit UI. They also maintained the vector index and handled document update workflows.

Industry Example 3: Agentic AI for Automated Market Intelligence

The Business Problem

A product team at a SaaS company needed weekly competitive intelligence reports — covering competitor features, pricing changes, and news. Manual research took **12+ hours per analyst per week**. Insights were often stale by the time they reached leadership.

How This Course Applies

Week 10 Agentic AI skills: A CrewAI multi-agent system was designed with three agents: Researcher (web scraping + news APIs), Analyst (synthesizes trends), and Writer (formats executive briefing).

Tools & Technologies

- CrewAI for multi-agent orchestration
- Tavily Search + SerpAPI for real-time web data
- OpenAI GPT-4 Turbo as the reasoning backbone
- Notion API for automatic report publishing

Measurable Business Impact

- Report generation time: 12 hours → 35 minutes
- Reports delivered every Monday morning, automatically
- Product team cited 3 feature decisions directly influenced by agent output

Role of Early Professionals

A junior Gen AI engineer designed the agent roles, wrote custom tools, and built the Notion publishing integration. Full ownership from prompt to pipeline.

Tools & Technology Stack

This course uses an industry-standard toolkit that mirrors what top AI teams use today. Every tool is free or has a generous free tier for learners.

Core ML & DL

- Python 3.x
- Scikit-learn
- TensorFlow / Keras
- XGBoost / LightGBM
- NumPy, Pandas

NLP & Gen AI

- HuggingFace Transformers
- OpenAI API
- Ollama (local LLMs)
- SpaCy / NLTK
- Sentence-Transformers

RAG & Agents

- LangChain
- ChromaDB / Pinecone
- CrewAI / AutoGen
- Tavily Search API
- RAGAS (evaluation)

MLOps & Deployment

- MLflow
- FastAPI
- Docker
- Streamlit
- Weights & Biases

Career & Job Role Alignment

This course is engineered to open specific, high-demand job titles. Here's how your learning maps to real career tracks in the AI industry.

Job Role	Core Skills Required	Weeks Most Relevant
ML Engineer	Scikit-learn, XGBoost, MLOps, FastAPI, Docker	Weeks 3–5, 11
Data Scientist	EDA, Supervised/Unsupervised ML, Statistics, Visualization	Weeks 2–5
NLP Engineer	Transformers, BERT, HuggingFace, Fine-tuning	Weeks 7, 11
Gen AI Engineer	Prompt Engineering, OpenAI API, RAG, LangChain	Weeks 8–9
AI Agent Developer	LangChain Agents, CrewAI, AutoGen, Tool Design	Week 10
MLOps Engineer	MLflow, Docker, FastAPI, CI/CD, Model Monitoring	Week 11
AI Product Analyst	Use Case Mapping, Prompt Engineering, Business Framing	Weeks 1, 8, Capstone

EDVANTA

Ignite your growth, inspire your mind

Skills Progression Map

Track your growth from absolute beginner to job-ready AI professional across the three phases of this program.

Beginner (Weeks 1–4)

Python fluency, ML fundamentals, EDA mastery, model evaluation, first supervised learning models

Intermediate (Weeks 5–8)

Deep learning, NLP pipelines, Transformer fine-tuning, Gen AI APIs, prompt engineering

Advanced (Weeks 9–11)

RAG systems, Agentic AI, fine-tuning LLMs, MLOps deployment pipelines, production monitoring

Job Ready (Week 12+)

Capstone portfolio, GitHub presence, interview-ready skills, real business problem solved end-to-end

EDVANTA

EDVANTA

Ignite your growth, inspire your mind

Key Takeaways

Everything in this program is designed with one goal: to make you genuinely job-ready in AI and ML, not just certificate-ready. Here's what you walk away with.

Unshakeable Foundations

From linear algebra to deep learning — you understand the *why* behind the algorithms, making you a stronger engineer and better problem-solver than peers who only know the *how*.

Gen AI & Agentic Fluency

You can build RAG pipelines, design multi-agent systems, and fine-tune LLMs — skills that put you at the frontier of what the industry is hiring for right now.

A Deployable Portfolio

Every weekly project and your capstone become live, portfolio-ready artifacts. You arrive at interviews with working demos, not slide decks about "what you would build."

Career Clarity

You know which AI roles fit your strengths, how to map your skills to job descriptions, and how to present yourself as a capable, business-aware AI professional from Day 1.

EDVANTA

Start Your AI Journey Today

The AI industry is not waiting. Every week you delay is a week your peers are committing code, building agents, and landing interviews. This 12-week program is the most direct path from where you are now to where the market demands you to be.

> *"The best time to start learning AI was two years ago. The second best time is this weekend."*

🗓️ Week 1 Starts

Set up Python, Colab, and your GitHub. Read the Week 1 brief and watch your first recorded session.

🤝 Join the Community

Connect with your cohort on Discord. Introduce yourself, share your background, and state your 3-month goal publicly.

🚀 Build in Public

Post your first project update on LinkedIn by Week 2. Document your learning journey. The community that watches you grow becomes your professional network.

💡 **Commitment matters more than talent.** Students who submit every weekly assignment — even imperfect ones — are 3× more likely to complete the capstone and get hired within 6 months of graduation.