masai

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

FROM VISHLESAN i-HUB, IIT PATNA

Program Overview

The foundations in AI and Machine Learning Program is designed to accommodate learners from diverse academic and professional backgrounds. Through structured coursework and hands-onprojects, students will build a solid understanding of Artificial Intelligence (AI) and Machine Learning (ML) techniques—starting from foundational mathematics and basic programming, progressingthrough core machine learning models, and culminating in cutting-edge AI applications. By the endof the program, participants will be able to confidently implement AI/ML solutions, interpret results,and address real-world problems across various domains.

Skills you will gain:

- Python Programming
- Data Cleaning and Preprocessing
- Exploratory Data Analysis (EDA)
- Model Evaluation and Validation
- Regression and Classification Techniques
- Agentic AI Development
- Neural Networks and Deep Learning
- Natural Language Processing (NLP) and Transformers
- Computer Vision & Image Processing
- Unsupervised Learning (Clustering, Dimensionality Reduction)
- Gen AI Engineering

Why Choose This Course?

- No coding background required a beginner-friendly entry into AI/ML
- Taught by IIT faculty – learn from top academic experts
- Hands-on projects – apply concepts to real-world problems
- Industry-relevant curriculum – covers latest tools, trends, and techniques
- Career-focused outcomes – prepares you for high-demand AI/ML roles

What Will You Learn?

Begin your AI/ML journey with a beginner-friendly, industry-aligned curriculum from Vishlesan i-Hub, IIT Patna. Covering everything from Python basics to advanced topics like deep learning and generative AI, the program builds real-world skills through hands-on projects and expert-led learning.

Toolkit

python™ docker PyTorch NumPy

pandas GitHub OpenAI

& more

Course Summary

Courses
3

Course Duration
36 Weeks (12 + 12 + 12)

Evaluations (Per Course)
2 Quizzes | 3 Trimesters

Break Duration
4 Weeks (2 + 2)

Weekly Commitment
8-10 Hours

Course Details

Code	Course
AIM101	Introduction to AI and ML for Beginners
AIM201	Core Machine Learning Techniques and Practices
AIM301	Advanced AI: Deep Learning, NLP, and Emerging Trends

Course Descriptions

AIM101

Introduction to AI and ML for Beginners

Geared toward students with minimal or no coding background, this course lays the groundwork for artificial intelligence and machine learning. The curriculum covers essential mathematical concepts, basic programming in Python, introductory data manipulation, and an overview of AI/ML applications. By the end, students will be comfortable writing simple Python scripts, performing basic data analysis, and understanding how AI/ML fits into various real-world scenarios.

AIM201

Core Machine Learning Techniques and Practices

After establishing the basics, this course dives deeper into machine learning algorithms. Students will work with both traditional and modern approaches—such as regression, classification, and clustering—while honing their skills in model evaluation, data preprocessing, and ethical considerations. By the end of AIM201, learners will be able to build, tune, and interpret effective ML models for diverse datasets, preparing them for more advanced AI techniques.

AIM301

Advanced AI: Deep Learning, NLP, and Emerging Trends

Building on the foundations from AIM101 and AIM201, this course explores cutting-edge AI topics like deep learning, natural language processing, and computer vision. Students will learn how to leverage advanced models (e.g., convolutional neural networks, transformers) and deploy them at scale. Throughout AIM301, participants will also stay current with emerging AI trends, addressing real-world complexities such as big data, reinforcement learning, and AI ethics in depth.

Course Syllabus

AIM101: Introduction to AI and ML for Beginners

1. **Python & Jupyter Fundamentals**
   - Python install, virtualenv, Jupyter launch, CLI basics, first script, variables, data types, I/O
   - Setting Up Your Dev Environment: OS-specific installs, venv / conda, IDE vs Jupyter, shell shortcuts, debugging tips
   - Coding Fluency Boot-camp: Keyboard workflows, code snippets, refactoring habits, dot-files
2. **Essential Mathematics for AI/ML**
   - if/else, loops, functions, modules, Git intro, commit-push loop
   - Version Control in Teams: Branching, pull/merge requests, conflict fixes, code reviews
   - Packaging & Env Management: setup.py, requirements.txt, virtual-env patterns
3. **Vector Algebra with NumPy**
   - Vector arithmetic, dot product, array creation, broadcasting, visualising vectors
   - NumPy Power Tools: Fancy indexing, ufuncs, memory views, profiling
   - Optimising Numerical Code: Vectorised patterns, Cython/Numba primer
4. **Matrices & Images**
   - Matrix mult., transpose, determinants, eigen intuition, RGB image as 3-D matrix
   - Linear Algebra in Practice: Eigendecomp, SVD demo, PCA teaser, SciPy LA
   - Efficient Image Ops: OpenCV matrix tricks, batch transforms, GPU offload
5. **Calculus & Stats Essentials**
   - Derivatives, gradients, SymPy, descriptive stats, probability dists, CLT
   - Auto-Diff for ML: PyTorch autograd, gradient checks, numerical vs analytic
   - Statistical Testing Toolbox: SciPy stats, bootstrap demo, power analysis
6. **Pandas & Basic Plots**
   - DataFrame create/index, cleaning, joins, line/bar/scatter, design principles
   - Data Wrangling Clinic: Messy CSVs, merges, tidy-data rules, time-series tricks
   - Matplotlib Mastery: Custom styles, subplots, annotations, export quality

7. Data Storytelling
   - Chart selection, bias spotting, redesign bad charts
   - Fixing Bad Visuals: Real-world bad chart makeover, color-blind palettes, decluttering
   - Dashboard Quick-Start: Plotly, Streamlit, layout grids, narrative flow
8. History & Foundations of AI
   - Symbolic AI, expert systems, ML rise, Turing Test, milestones
   - AI Timeline Walk-through: Key industry eras, failed projects, inflection points
   - Modern AI Stacks: Typical production pipeline, build-vs-buy choices
9. Intro ML & Metrics
   - Supervised vs unsupervised, K-NN, simple linear regression, train/test split, accuracy, RMSE
   - Training Your First Model: scikit-learn fit/predict, metric trade-offs, baseline report
   - Model Debugging 101: Error analysis, confusion matrix, bias-variance checks
10. Data Collection & Cleaning
    - Open data portals, API pulls, scraping, missing values, encoding, outliers
    - Web-Scraping at Scale: Requests/BS4, pagination, rate-limits, proxies
    - Data Quality Pipelines: Great Expectations, imputation libs, anomaly rules
11. EDA & Statistical Tests
    - Distributions, heatmaps, chi-square, t-tests, sector case studies
    - Story-Driven EDA: Interactive dashboards, segmentation cuts, hypothesis logs
    - Case-Study Deep-Dive: Healthcare & finance EDA playbooks, insights write-up
12. Mini-Project & Ethics
    - Dataset › model pipeline, bias audit, privacy, showcase
    - Project Code-Clinic: Pair debugging, repo hygiene, doc strings, tests
    - Responsible AI in Action: Fairness metrics, privacy tools, audit checklist

AIM201 – Core Machine Learning Techniques and Practices

1. 1. ML Workflow & Pipelines
   - Pre-processing, feature eng., train/eval/deploy, sklearn Pipeline, DVC
   - Building Reproducible Pipelines: MLflow, DVC remotes, config files, CI hooks
   - Continuous Delivery for ML: GitHub Actions, Docker basics, staging vs prod
2. Multiple & Polynomial Regression
   - Multi-linear, polynomial fits, coef interpretation, residual diagnostics
   - Regression in Business: Statsmodels, KPI forecasting, diagnostic plots
   - Regularisation Preview: Ridge/Lasso intro, when & why
3. Regularisation & Hyper-search
   - Ridge, Lasso, ElasticNet, overfit guard, GridSearchCV, RandomisedSearchCV
   - Tuning at Scale: Cross-val strategies, Optuna sampler, early-stop rules
   - Experiment Tracking: Weights & Biases, comparative dashboards, report prep
4. Logistic Regression & Metrics
   - Binary logistic, confusion matrix, precision, recall, F1, ROC, AUC
   - Classification Metrics Clinic: Threshold tuning, imbalanced hacks, cost curves
   - Exec-Ready Visuals: Lift/gain charts, business translation, report tips
5. Decision Trees & Overfitting
   - CART, Gini/entropy, pruning, decision boundaries, CV, early stopping
   - Tree Visualisation: sklearn plot_tree, dtreeviz, feature importance
   - Serving Tree Models: Pickle vs ONNX, REST stub, latency tests
6. Ensembles – RF & GBM
   - Random Forest, Gradient Boosting, feature importance, early stop, mini-comp
   - Winning with XGBoost: GPU mode, regularisation, CV tuning, SHAP plots
   - Kaggle-Style Stacking: Stacker blueprints, blend scripts, leaderboard tricks
7. K-Means & DBSCAN
   - Distance metrics, elbow, silhouette, density clustering, geo-segmentation
   - Market Segmentation Lab: sklearn cluster, scaler choice, persona crafting
   - Visualising Clusters: UMAP, GeoPandas maps, interactive hover

8. Dim-Red & Anomaly Detect
   - PCA, t-SNE preview, Isolation Forest, stat outlier methods, fraud heatmap
   - Fraud Detection in Finance: Feature engineering, IF tuning, scoring rules
   - Streaming Anomaly Pipelines: Kafka intro, window stats, alert thresholds
9. Neural Nets Fundamentals
   - Perceptron, activations, back-prop math, PyTorch net, dropout, augment
   - PyTorch Training Loop: Tensor ops, SGD vs Adam, LR schedulers, checkpoints
   - Profiling & Debugging: GPU utilisation, weight histograms, grad anomalies
10. LLM Basics & Tokenisation
    - Transformer overview, BPE, embeddings, 7-B demo, token/logit inspect
    - Hugging Face Quick-Start: transformers API, custom tokenisers, pipeline API
    - Memory-Efficient Inference: Bits-and-bytes, 8-bit load, GPU vs CPU
11. Vector DBs & RAG
    - FAISS/Chroma, text chunking, retrieval, FAQ bot, grounding eval
    - Building RAG Pipelines: Embedding batch gen, hybrid search, eval harness
    - RAG Evaluation & Tuning: Retrieval scores, answer-ground truth, hallucination fix
12. Mid-Project & MLOps Intro
    - Team ML+RAG build, CI, Git, DVC, API stub, fairness reflection
    - Containerising Models: Dockerfile patterns, dependency pinning, slim images
    - Monitoring in Prod: Prometheus, Grafana, model drift alerts, retrain triggers

AIM301 – Advanced AI: Deep Learning, NLP, and Emerging Trends

1. 1. CNN Fundamentals
   - Convolution, kernels, stride, pooling, cats-vs-dogs CNN
   - Image Augmentation Lab: Albumentations, mixup, CutMix, dataset loaders
   - TensorBoard Vision: Filters, activations, embeddings, PR curves
2. Adv. CNN & Detection
   - Batch norm, ResNet, fine-tune, SSD, YOLO, segmentation
   - Transfer Learning Primer: Pre-trained weights, freezing vs finetune, LR finder
   - Model Optimisation: Quantisation, pruning, TensorRT, edge deploy
3. Sequence Models & Sentiment
   - RNN, LSTM, GRU, Word2Vec/GloVe, IMDb LSTM
   - Text Pipeline Engineering: torchtext, token buckets, batching, padding tricks
   - FastAPI Deployment: Inference endpoint, async queue, caching
4. Transformers & BERT Tune
   - Self-attention, positional enc., mini-transformer, BERT finetune
   - PEFT Hands-On: LoRA, QLoRA, adapters, PEFT library
   - Experiment Tracking at Scale: wandb sweeps, dataset versions, hyper dashboards
5. LLM Internals & Scaling
   - Depth/width, heads, scaling laws, interpretability probes
   - Interpreting Transformers: Attention roll-out, probing layers, causal tracing
   - Cost & Performance Trade-offs: Token vs context length, GPU hours, batching
6. Fine-Tuning LLMs & AutoML
   - LoRA/QLoRA, Optuna, FAQ model tune
   - AutoML for LLMs: Bayesian opt, early stop, search space design
   - LLM Evaluation Suite: BLEU, ROUGE, perplexity, human eval loop
7. Prompt Eng. & Security
   - Roles, token budget, temp, top-p, function calls, injections
   - Red-Teaming Prompts: Jailbreak attempts, toxic prompt sets, safe completions
   - Guardrails & Monitoring: OpenAI moderators, regex filters, logging, alerts

8. **LLM APIs & Streamlit Bot**
   - OpenAI & Anthropic calls, streaming, Streamlit chatbot, logging
   - Rapid Prototyping: Async API, error handling, cost calc, prompt versioning
   - Deploying Streamlit: Docker-compose, secrets, CI/CD to cloud run
9. **RAG at Scale w/ Spark**
   - Spark embed gen, hybrid search, re-rank, Docker, cloud deploy
   - Distributed Embeddings: PySpark UDFs, GPU scheduling, cluster tuning
   - Service Observability: Logs, traces, metrics, autoscale
10. **AI Agents with LangGraph**
    - Agent arch, tool calls, memory, retries, citations
    - Building Autonomous Agents: LangGraph flows, sub-agents, eval harness
    - Ethics & Safety for Agents: Goal mis-spec, sandboxing, human-in-loop
11. **Generative Models 101**
    - GANs, VAEs, diffusion, fashion image gen
    - Diffusion Fine-Tuning: Stable Diffusion LoRA, DreamBooth basics, prompt templates
    - Quality & Metrics: FID, CLIP score, human eval, prompt libraries
12. **Reinforcement Learning & Capstone**
    - MDP, Q-learning, policy gradients, RLHF, governance, demos
    - RL in Practice: Gymnasium, vector envs, PPO, reward shaping
    - Capstone Coaching: Demo polish, storytelling, viva prep, next-steps

Our Mentors

Dr. Sriparna Saha

Associate Professor at Department of Computer Science and Engineering, IIT Patna

Dr. Sriparna Saha received the M.Tech and Ph.D. degrees in computer science from Indian Statistical Institute Kolkata, Kolkata, India, in the years 2005 and 2011, respectively. She is currently a Faculty Member in the Department of Computer Science and Engineering, Indian Institute of Technology Patna, India. She is the author of a book published by Springer-Verlag. She has authored or coauthored more than 120 papers in reputed journals and conferences including IEEE/ACM transactions, core ranked conferences. Her current research interests include text mining pattern recognition, natural language processing, multi-objective optimization and biomedical information extraction. Her h-index is 19 and total citation count of her papers is 2200 (according to Google scholar). She is also a senior member of IEEE. She is the recipient of the Lt Rashi Roy Memorial Gold Medal from the Indian Statistical Institute for outstanding performance in MTech (computer science). She is the recipient of the Google India Women in Engineering Award,2008, NASI YOUNG SCIENTIST PLATINUM JUBILEE AWARD 2016, BIRD Award 2016 and IEI Young Engineers' Award 2016, Humboldt Research Fellowship.

Gaurav Kandel

Data Scientist, FinBox

Gaurav Kandel is a Data Scientist at FinBox with extensive experience across NIRA, KPMG India, and multiple AI mentorship roles. He frequently shares insights on practical machine learning topics—from deploying models in production to building scalable recommender systems and mastering prompt engineering in Vertex AI. Gaurav blends hands-on technical expertise with a passion for demystifying AI for learners, professionals, and businesses through teaching, writing, and real-world problem solving.

Sriram Desai

Software Engineer, ByteDance

Sriram Desai is a seasoned Software Engineer currently at ByteDance, where he leads development within TikTok LIVE's Money Platform. With prior roles at Amazon Prime Video and PayPal, Sriram has architected scalable systems for invoicing, content localization, and alternate payments. A Georgia Tech graduate, he brings deep expertise in backend systems, REST APIs, cloud services, and data-intensive platforms. Known for operational excellence and developer tooling, Sriram is a driving force in building resilient, compliant, and high-throughput financial and media infrastructures.

Admission Process

Clear Qualifier Test
You must pass the exam to confirm your seat for the program.

Complete Onboarding
Only shortlisted candidates go through the onboarding process.

Start Learning
Learn from India's top educators and stand out from the crowd.

Fees Structure

Qualifier Test Fee (Non-Refundable)	**₹99**

	Option 1	Option 2
	Upfront	**EMI** Through our NBFC partners
Secure Seat Fee (Non-Refundable)	₹4,000	₹4,000
Remaining Course Fee (Non-Refundable)	₹51,085	₹6,527 x 9 months
Total Program Fee	₹55,085*	₹62,743*

*GST at 18% extra, as applicable

For More Queries

iitp.programs@masaischool.com