GUVI | HCL

Skill Up. Level Up

intel

IITM PRAVARTAK

Become an Industry-Ready AI Engineer with the

Intel & IITM Pravartak Certified Artificial Intelligence & Machine Learning Program

A 6-month learning journey where you'll gain hands-on experience in high-impact areas like Generative AI, Agentic AI, Deep Learning, and MLOps.

Duration of the course: 6 Month

200+ Hours

Real-World Projects

About the Program

HCL-GUVI’s Intel & IIT-M Pravartak-Certified Artificial Intelligence & Machine Learning Program is a comprehensive 6-month career-launching journey designed to transform learners into industry-ready AI professionals. It begins with a strong foundation in Python, statistics, and machine learning, and progresses into advanced concepts like Deep Learning, LLMs, GenAI, and Agentic AI systems.

Participants gain hands-on experience through Intel®-curated labs, real-world capstone projects, and cutting-edge tools such as PyTorch, TensorFlow, LangChain, and FAISS. You’ll learn to build and deploy end-to-end solutions—from classical ML models to modern LLM-powered APIs—using MLOps, containerization (Docker), and cloud platforms like AWS, GCP, and Azure.

Beyond technical mastery, the program offers live weekend sessions led by industry experts, mentorship for career planning, portfolio reviews, and placement preparation through mock interviews, coding challenges, and career acceleration via GUVI’s HyreNet.

From beginners to professionals looking to upskill for high-growth AI/ML roles, this program ensures you're equipped to build scalable, intelligent systems that drive impact in top-tier tech teams.

The Right Program to Code, Build & Scale

This 6-month program is designed for aspiring AI professionals who want to master machine learning, deep learning, and Generative AI, build real-world projects, and gain hands-on experience with tools like LLMs, MLOps, and Cloud Deployment. It prepares you for high-impact roles by combining technical depth, industry-ready skills, and career-focused mentorship.

What You'll Learn to Build

A rock-solid foundation in Machine Learning, Deep Learning, Generative AI & Agentic AI

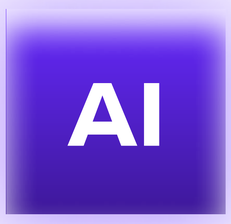

Mastery in AI Concepts for coding interviews

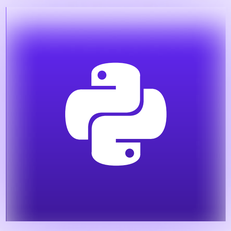

Full-cycle development skills—from Python to MLOps & LLMOps

Build end-to-end projects & perform hands-on lab exercises on Intel AI PC

Experience in building LLM-based Projects, Vector Search & RAG systems

A professional GitHub portfolio with real-world projects & open-source contributions

The Lasting Impact You'll See

Bridges the Skill Gap between theory and real-world application - certification from Academia (IITM-Pravartak) & Industry (Intel, AWS) both.

Master the AI/ML Stack from End-to-End.

Builds Confidence through practical, project-based learning.

Prepares you for AI & ML Interviews with mock rounds and curated prep kits.

Develop Problem-Solving Agility Across Tools & Platforms - from Python to cloud deployment to ML Ops workflows.

Delivers Job-Ready Outcomes with Machine learning, Generative AI & MLOps.

Strengthens Your Personal Brand through career mentorship and portfolio reviews.

Connects You to Top Tech Recruiters through GUVI's HyreNet platform.

Complete Your AIML Journey With One Future-Ready Course

Certifications That Set You Apart

Co-branded certificate with Intel, AWS, IITM-Pravartak & HCL GUVI validating your expertise in Artificial Intelligence & Machine Learning.

Learn from Certified Mentors

Get guided by Industry-certified instructors with deep expertise in scalable AI architectures, edge-to-cloud deployments, and GenAI-integrated system design.

Special Sessions by Intel Experts

Learn from those building next-gen LLM applications, AI accelerators, and real-time inference APIs using cutting-edge Intel toolchains.

AI-Powered Mock Interviews

Sharpen your interview skills with real-time AI-assisted mock sessions and targeted feedback from top industry mentors.

12-Month HyreNet Access

Gain exclusive job support and placement opportunities through HyreNet, our official hiring platform with partner companies across India and beyond.

Chance to get Offline Immersion

Opportunity to join an in-person engineering immersion to understand real-world software development workflows, collaborate on group projects, and receive mentorship directly.

Intel Powered Lab Exercises & Projects

Work on cutting-edge AI scenarios with hands-on labs and real-world projects crafted by Intel experts—focused on optimizing performance for Intel AI PCs. These labs not only simulate industry challenges but also equip you with the skills to deploy scalable GenAI solutions across edge and cloud environments.

Who Should Enroll in this Program

College Students	Fresh Graduates
Build job-ready skills & a standout portfolio before you graduate.	Crack AI ML interviews with in-demand stack & hands-on projects.
Career Switchers	**Self-taught Programmers**
Transition smoothly into tech with expert mentorship & AI tools.	Get certified, gain structure, and scale with expert guidance.

Program Overview

From programming fundamentals to Cloud Deployment for LLM-Based Applications, unlock the potential of Artificial Intelligence and Machine Learning with this comprehensive, industry-aligned program powered by Intel.

Section 1: Programming & Data Analytics

Module 1: Introduction to AI and Current Trends (During Orientation)

This module offers a powerful introduction to the evolution and current landscape of Artificial Intelligence. You'll gain a clear understanding of what AI is, where it came from, and how it's rapidly transforming the world today.

Module Curriculum

- Brief history of AI: From rule-based systems to deep learning
- Key milestones in AI development
- AI vs. Machine Learning vs. Deep Learning
- Current trends in AI (2024–2025)
- Generative AI and foundation models (e.g., GPT, DALL·E, Claude)
- AI in everyday life and industries
- Introduction to LLMs (Large Language Models)
- Rise of multimodal AI and agentic systems
- AI in India and global innovation landscape
- Future possibilities and emerging AI technologies

Module 2: Setting up AI Development Environment on PC - Maximize Your Intel AI PC

Quickly configure your development environment with WSL, VS Code, Git, and database tools. Build and deploy full-stack AI apps using FastAPI and Prisma—optimized for seamless GenAI development.

Module Curriculum

- Introduction to Windows Subsystem for Linux (WSL)
- Setting Up Your AI Development Environment
- Setting Up Your Favorite IDE - Visual Studio Code
- Integrating Git for Source Code Management
- Running a Full-Stack Application on Your Device

Module 3: Introduction to Python

Kickstart your coding journey with a beginner-friendly, hands-on Python module—perfect for mastering data handling, clean coding, and app building from scratch.

Module Curriculum

- Getting started with Python: Installation and setup
- Python syntax and writing your first program
- Variables and data types
- Type casting and input/output operations
- Operators: Arithmetic, comparison, logical, bitwise
- Control structures: if-else, nested conditions
- Loops: for, while, break, continue
- Functions: definition, arguments, return values, scope
- Working with strings: slicing, formatting, built-in functions
- Lists, tuples, and dictionaries: creation and manipulation
- Using built-in Python functions and methods
- Writing modular and reusable code
- Introduction to error handling (try-except blocks)
- Writing and executing Python scripts from terminal/IDEs

Module 4: Advanced Python Concepts

Level Up Your Python Skills- Advance to writing scalable, modular, and high-performance code with OOP, advanced data structures, and essential modules—laying the groundwork for real-world AI development.

Module Curriculum

- Object-Oriented Programming (OOP) in Python
- Advanced functions and functional programming- Lambda functions, map(), filter(), and reduce()& Decorators and closures
- Working with modules and packages
- Exception handling
- File handling
- List comprehensions and generator expressions
- Understanding Python iterators and generators
- Working with regular expressions (re module)
- Introduction to working with JSON and APIs

Module 5: Data Handling with Pandas

Master Data with Pandas- Learn to clean, transform, and analyze real-world datasets using Python's most powerful data library. This hands-on module builds essential data wrangling skills for AI, ML, and data science.

Module Curriculum

- Introduction to Pandas and its role in data science
- Installing and importing Pandas
- Understanding Series and DataFrames
- Reading and writing data (CSV, Excel, JSON, etc.)
- Exploring and summarizing datasets
- Data selection and filtering techniques
- Indexing and slicing DataFrames
- Handling missing data (NaNs)
- Data cleaning and transformation
- Merging, joining, and concatenating datasets
- Grouping and aggregation operations
- Sorting and ranking data
- Applying functions with apply() and lambda
- Exporting clean datasets for modeling

Module 6: Introduction to SQL

Gain hands-on skills in querying, filtering, and managing relational data. This module empowers you to extract insights and structure data efficiently—an essential foundation for any AI, ML, or analytics workflow.

Module Curriculum

- What is SQL and why it matters in AI & data workflows
- Understanding relational databases and schemas
- Setting up a SQL environment (e.g., MySQL, SQLite, PostgreSQL)
- Writing basic queries with SELECT

- Filtering data using WHERE, BETWEEN, LIKE, IN
- Sorting and limiting results (ORDER BY, LIMIT)
- Combining data using JOIN (INNER, LEFT, RIGHT, FULL)
- Aggregation with GROUP BY, HAVING, and built-in functions (SUM, COUNT, AVG, etc.)
- Working with subqueries and nested statements
- Data types and constraints (NULL, PRIMARY KEY, FOREIGN KEY)
- Creating, updating, and deleting records (INSERT, UPDATE, DELETE)
- Best practices for structuring and querying databases

Module 7: Exploratory Data Analysis (EDA)

Master the art of uncovering insights through Exploratory Data Analysis (EDA). Learn to visualize patterns, detect anomalies, and prepare data for ML using tools like Pandas, Matplotlib, and Seaborn.

Module Curriculum

- What is EDA and why it’s important
- Workflow of an EDA process
- Descriptive statistics and summary metrics
- Data profiling and quality assessment
- Detecting and handling missing values
- Identifying outliers and anomalies
- Visualizing distributions (histograms, box plots, KDE)
- Exploring relationships (scatter plots, heatmaps, pair plots)
- Correlation analysis and multicollinearity checks
- Categorical vs numerical data exploration
- Using Seaborn and Matplotlib for custom visualizations
- Creating data stories and dashboards with insights
- Documenting EDA findings for modeling phases

Module 8: Data Visualization in Python

Visualize Data with Impact- Turn raw numbers into compelling visuals. Learn to

build charts, graphs, and interactive dashboards using Matplotlib, Seaborn, and Plotly—perfect for data storytelling and AI-driven insights.

Module Curriculum

- Importance of data visualization in AI and data science
- Introduction to Python visualization libraries: Matplotlib, Seaborn, Plotly
- Creating line, bar, and scatter plots with Matplotlib
- Styling and customizing plots (labels, legends, titles, colors)
- Building advanced visualizations: histograms, box plots, pie charts
- Heatmaps and correlation matrices
- Plotting time series and trends
- Using Seaborn for statistical visualizations (violin plots, pair plots, swarm plots)
- Introduction to interactive visualizations with Plotly
- Creating subplots and multi-panel charts
- Saving and exporting high-quality visualizations
- Best practices for visual storytelling and design clarity

Module 9: Essential Mathematics for Data Science

Grasp core concepts in Linear Algebra, Calculus, Probability, and Statistics—tailored for AI and ML. From vectors, derivatives, and matrix operations to Bayes' Theorem and gradient descent, this module builds the mathematical intuition needed to train and optimize intelligent models.

Module Curriculum

- Linear Algebra: Scalars, Vectors, and Matrices, Basic Vector Operations, Dot Product
- Intro to Calculus: Matrix Transposition, Basic Matrix Operations, Matrix Multiplication
- Calculus Intuition: What is a Function?, The Concept of a Derivative, Finding Minima and Maxima

- Calculus for Optimization & Introduction to Probability: Partial Derivatives, The Gradient, Gradient Descent Algorithm
- Fundamentals of Probability: Basic Terminology, Probability Axioms & Basic Rules, Conditional Probability, Independence of Events
- Essential Probability & Statistics: Bayes' Theorem, Descriptive Statistics, Random Variables & Probability Distributions, Tying it All Together & Next Steps

Section 2: Machine Learning

Module 10: Machine Learning Fundamentals

This module introduces core ML concepts—data prep, model building, algorithm selection, and performance evaluation—laying a strong base for AI, data science, and automation careers.

Module Curriculum

- What is Machine Learning and how it differs from traditional programming
- Categories of ML: Supervised, Unsupervised, and Reinforcement Learning
- Understanding datasets: features, labels, and data splitting
- Overview of the ML pipeline: preprocessing to deployment
- Introduction to Scikit-learn and its ecosystem
- Data preprocessing techniques- Scaling, normalization, encoding categorical variables, missing value handling
- Model selection and training- Linear Regression, Logistic Regression, Decision Trees
- Evaluating models- Accuracy, precision, recall, F1-score, confusion matrix
- Cross-validation and overfitting vs underfitting
- Hyperparameter tuning basics

- Real-world use cases of ML in business and tech
- Ethical considerations and fairness in machine learning

Module 11: Advanced Machine Learning

Take your ML skills to the next level with advanced algorithms, classification, clustering, dimensionality reduction, and model optimization—designed to solve complex, real-world AI challenges with precision.

Module Curriculum

- Advanced Classification Techniques - Beyond simple logistic regression: enhancements with decision trees, ensemble methods (Random Forest, Gradient Boosting)
- Sophisticated Clustering Methods
- Dimensionality Reduction- Principal Component Analysis (PCA) and its variants, Linear Discriminant Analysis (LDA) for feature extraction
- Support Vector Machines (SVM)
- Model Evaluation & Improvement
  - Cross-validation techniques and model diagnostics
  - Hyperparameter tuning for optimal performance
  - Handling bias-variance trade-offs
- Ensemble Techniques
  - Combining models for improved accuracy
  - Introduction to bagging, boosting, and stacking
- Reinforcement Learning
  - Key concepts: agent, environment, action, reward, state
  - Types of RL: model-free vs model-based
  - The RL loop: exploration vs exploitation
  - Markov Decision Processes (MDPs)
  - Policy, Value Function, and Q-Function

Module 12: Introduction to MLOps: Bridge the gap between Machine Learning and real-world deployment

Learn to operationalize machine learning with industry-grade tools and workflows. This module covers model versioning, pipeline automation, testing, deployment, and monitoring—preparing you to build scalable, production-ready AI systems.

Module Curriculum

- What is MLOps? Why is it important?
- ML lifecycle: from experimentation to deployment
- Key differences between DevOps and MLOps
- Components of the MLOps ecosystem
  - Data engineering
  - Model development
  - Continuous integration and deployment (CI/CD)
  - Monitoring and governance
- Tools and platforms overview:
  - MLflow, DVC, Kubeflow, TFX
  - Weights & Biases, Neptune.ai, ClearML
- Model versioning and tracking
- Creating reproducible ML workflows
- Data version control and pipeline automation
- Introduction to containers (Docker) for ML reproducibility
- Deployment strategies:
  - Batch vs real-time inference
  - REST APIs, FastAPI, Flask for model serving
- Infrastructure options: cloud, on-prem, hybrid
- Monitoring models in production: drift, bias, accuracy
- Collaboration across teams (Dev, Data Science, Ops)
- CI/CD pipeline basics for ML projects
- Security and ethical considerations in ML deployments
- Hands-on demo:
  - Track experiments with MLflow
  - Containerize a model and deploy using FastAPI

Section 3: Deep Learning

Module 13: Introduction to Neural Networks

Step into deep learning by building your first neural network. Understand how layers, weights, and training processes power modern AI—from image recognition to language models.

Module Curriculum

- What are neural networks and how they work
- Biological inspiration vs artificial implementation
- Structure of a neural network: neurons, layers, weights, and biases
- Activation functions: ReLU, sigmoid, tanh, softmax
- Forward propagation and loss calculation
- Backpropagation and the learning process
- Gradient descent and optimization techniques
- Introduction to deep learning frameworks: TensorFlow and PyTorch
- Building a simple neural network from scratch
- Visualizing the training process and performance metrics
- Overfitting and regularization techniques
- Real-world applications of neural networks
- Key terminology: epochs, batches, learning rate, accuracy

Module 14: Deep Neural Networks

Advance your deep learning skills by building and training multi-layered neural networks. Learn to capture complex patterns using TensorFlow and Keras, and master techniques to tackle overfitting, vanishing gradients, and more.

Module Curriculum

- Understanding the architecture of Deep Neural Networks (DNNs)
- Benefits of depth: capturing complex patterns and hierarchies
- Hidden layers: how many, how deep, and why it matters

- Training challenges in deep networks
  - Vanishing/exploding gradients
  - Overfitting and underfitting
- Optimization strategies
  - Batch normalization
  - Dropout regularization
  - Learning rate schedules
- Building and training DNNs with TensorFlow/Keras
- Activation functions in deep networks
- Monitoring performance: loss curves, accuracy tracking
- Model evaluation and validation techniques
- Using callbacks and checkpoints
- Real-world DNN applications and case studies
- Best practices for scaling deep models

Module 15: Applied Deep Learning with PyTorch

Master PyTorch by building and training deep neural networks from the ground up. This module also shows you how to harness Intel® Arc™ GPU acceleration for faster, more efficient AI model development using Intel's open-source tools.

Module Curriculum

- Introduction to PyTorch: tensors, automatic differentiation, and the computation graph
- PyTorch vs other frameworks: why researchers and developers prefer it
- Building neural networks using torch.nn and nn.Module
- Implementing forward passes and activation functions
- Loss functions and optimizers using torch.optim
- Writing custom training and validation loops
- Efficient data loading using torch.utils.data.Dataset and DataLoader
- Handling overfitting with dropout, regularization, and early stopping
- Visualizing metrics and loss curves with TensorBoard
- Saving and loading models using state_dict and checkpoints

- Training models on Intel Arc GPU
  - Installing oneAPI and enabling Arc support
  - Switching training backend from CUDA to Intel
  - Performance profiling and optimization tips
- Case study: Training an image classifier on Intel Arc GPU
- Best practices for deploying models trained with PyTorch

Module 16: Introduction to Computer Vision with CNNs

Explore how AI sees the world! Learn to process and analyze images using Convolutional Neural Networks. This module equips you to build image classifiers and solve real-world vision tasks using deep learning tools.

Module Curriculum

- Introduction to Computer Vision and real-world applications
- Challenges in interpreting images through AI
- Image representation: pixels, channels, resolution, and color spaces
- Why CNNs? Understanding the limitations of traditional neural networks for vision tasks
- Fundamentals of Convolutional Neural Networks (CNNs)
  - Convolution layers, filters, and feature maps
  - Pooling layers (max and average pooling)
  - Activation functions and non-linearity
- Building a simple CNN architecture from scratch
- Image classification workflow using TensorFlow/Keras or PyTorch
- Working with image datasets (CIFAR-10, MNIST, etc.)
- Data preprocessing and augmentation techniques
- Model training, validation, and evaluation
- Visualization of learned filters and feature maps
- Transfer learning and using pre-trained CNNs (VGG, ResNet, MobileNet)
- Performance tuning and avoiding overfitting in CNNs
- Hands-on: Building a digit/image classifier using CNN
- Introduction to edge deployment on AI-enabled devices

Module 17: Natural Language Processing (NLP)

Learn how machines understand and generate human language. This hands-on module covers text preprocessing, sentiment analysis, and basic NLP models—equipping you to build AI-driven language applications using leading Python libraries.

Module Curriculum

- Introduction to NLP and its real-world applications
- Understanding structured vs unstructured text data
- Text preprocessing techniques
  - Tokenization, stop-word removal, stemming, lemmatization
- Working with libraries like NLTK and spaCy
- Text representation methods
  - Bag of Words (BoW), TF-IDF
- Word embeddings
  - Introduction to Word2Vec, GloVe, and FastText
- Text classification basics
  - Spam detection, sentiment analysis
- Sequence models overview
  - Introduction to RNNs, LSTMs, and attention mechanisms
- Hands-on: Building a sentiment analysis classifier
- Evaluating NLP models: accuracy, precision, recall, F1-score
- Introduction to transfer learning in NLP
  - Pre-trained models (BERT, RoBERTa, GPT)
- Ethical challenges in NLP: bias and misinformation

Section 4: Generative AI

Module 18: Introduction to Generative AI

Discover how machines create human-like text, images, and more. This module introduces core generative models and their real-world impact across content creation, design, healthcare, and customer experience—empowering you to build the future with AI.

Module Curriculum

- What is Generative AI?
- Real-world applications of generative models
  - Text (ChatGPT, Jasper)
  - Images (DALL·E, Midjourney)
  - Audio/Music (ElevenLabs, Suno)
  - Video (Runway, Synthesia)
- Generative vs traditional AI: key differences
- Overview of foundational models
  - Language models (GPT, BERT, T5)
  - Vision models (Stable Diffusion, GANs)
- Understanding autoencoders and latent spaces
- Introduction to Generative Adversarial Networks (GANs)
- Introduction to Diffusion Models
- Introduction to Large Language Models (LLMs)
- Creative tools powered by Gen AI
- Ethical considerations in content generation
- Understanding hallucinations and content validation
- The role of prompt engineering in content control
- Hands-on demo:
  - Generate text with ChatGPT or Hugging Face
  - Generate art using text-to-image models
- Introduction to open-source generative AI tools
- The future of Generative AI and emerging trends

Module 19: Large Language Models (LLMs) and Prompt Engineering

Uncover how powerful LLMs like GPT, Claude, and LLaMA work. This module teaches prompt engineering, real-world use cases like summarization and Q&A, and equips you to build ethical, language-driven AI applications with generative intelligence.

Module Curriculum

- Introduction to LLMs: what they are and how they work
- Overview of popular models: GPT, LLaMA, Claude, Mistral, and Gemini
- LLM capabilities and limitations
- Tokenization and model context window
- Few-shot, zero-shot, and chain-of-thought prompting
- Prompt Engineering fundamentals
  - Anatomy of a good prompt
  - Instruction-tuned models and role prompting
- Prompt patterns for specific tasks
  - Summarization, rewriting, classification, and translation
- Using OpenAI's GPT and HuggingFace models via APIs
- Tools for prompt testing and refinement
- Introduction to prompt chaining and dynamic prompt templates
- Safety, alignment, and ethical concerns when working with LLMs
- Hands-on: Crafting and refining prompts for real-world tasks

Module 20: Building AI-Powered Applications with Flask and Streamlit

Learn to turn your AI models into interactive web applications using Flask and Streamlit. This module teaches you to build user-friendly interfaces, connect them to ML models, and deploy apps accessible via the browser—bridging data science with real-world usability.

Module Curriculum

- Introduction to web applications for AI and ML
- Overview: Flask vs Streamlit – use cases and strengths
- Basics of Flask:
  - Routes, templates, and form handling
  - Connecting Python logic with HTML

    - Returning model predictions via API endpoints
- Basics of Streamlit:
    - Creating interactive UI components
    - Data visualizations with charts and plots
    - Deploying models with drag-and-drop simplicity
- Structuring a basic AI application
- Integrating pre-trained ML/LLM models into apps
- Connecting front-end inputs to back-end predictions
- Uploading files and handling user input
- Hosting local vs cloud-based applications
- Debugging and testing application workflows
- Hands-on:
    - Build a sentiment analysis web app using Streamlit
    - Build an image classifier or text generator with Flask
- Preparing apps for deployment (Heroku, Render, or cloud)
- UX tips for building user-friendly AI interfaces

Module 21: Advanced Prompt Engineering and LLM Fine-Tuning

Go beyond basics with techniques like few-shot chaining, function calling, and embedding-based prompts. Learn to fine-tune open-source LLMs for domain-specific use—boosting accuracy, control, and performance in real-world AI applications.

Module Curriculum

- Beyond the basics: Prompt refinement strategies for precision
- System prompts and role-based persona engineering
- Few-shot learning with structured examples
- Chain-of-thought prompting and multi-step reasoning
- Embedding-based prompts and semantic similarity
- Prompt chaining for multi-turn, multi-step tasks
- Retrieval-augmented prompting vs vanilla prompting
- Controlling tone, length, style, and response format

- Function calling and tool use with LLMs (OpenAI, Claude)
- Limitations of prompting and when fine-tuning is needed
- Introduction to fine-tuning LLMs
  - Full fine-tuning vs parameter-efficient tuning (LoRA, PEFT)
- Datasets for fine-tuning (custom, open-source)
- Using Hugging Face tools to fine-tune a model (e.g., LLaMA, Mistral)
- Evaluating fine-tuned models: performance and generalization
- Best practices for safe and ethical deployment

Module 22: Reinforcement Learning: Fine-tune generative models through feedback-driven learning

Discover how Reinforcement Learning enhances LLMs through real-time feedback. This module covers RLHF and other cutting-edge techniques used to align models like ChatGPT and Claude for safer, smarter, and more adaptive AI outputs.

Module Curriculum

- Overview of RL in the context of Generative AI
- What is Reinforcement Learning from Human Feedback (RLHF)?
- Role of feedback loops in fine-tuning LLM behavior
- Architectures that support RLHF (e.g., PPO, DPO)
- Stages of RLHF:
  - Pretraining with supervised learning
  - Reward model training from human preferences
  - Policy optimization with RL algorithms
- Algorithms used in RLHF
  - Proximal Policy Optimization (PPO)
  - Direct Preference Optimization (DPO)
- Designing reward models for generative tasks
- Aligning language models with ethical guidelines and safety constraints
- Case study: How OpenAI uses RLHF in ChatGPT

- RLHF vs supervised fine-tuning – when to use what
- Challenges in scaling RL for large models
- Bias, interpretability, and feedback reliability
- Hands-on example: Use a small-scale transformer with simulated feedback
- Tools and libraries: TRL (Hugging Face), Accelerate, RLlib

Module 23: Retrieval-Augmented Generation (RAG) for AI Models

Enhance LLMs with real-time, dynamic knowledge retrieval. This module teaches you to build RAG pipelines using tools like FAISS, Pinecone, and LangChain—enabling AI systems to deliver accurate, context-rich responses powered by your own knowledge base.

Module Curriculum

- What is Retrieval-Augmented Generation (RAG) and why it matters
- Key components of a RAG pipeline
    - Retriever, Vector Store, Generator
- Vector databases: FAISS, Pinecone, Weaviate overview
- Document chunking strategies for optimal retrieval
- Text embeddings and semantic search
    - Using OpenAI, Hugging Face, or Intel-generated embeddings
- Creating knowledge bases from PDFs, HTML, Notion, etc.
- Connecting LLMs with retrieval layers using LangChain
- Prompting with contextual information from retrievers
- RAG vs traditional search vs fine-tuning: when to use what
- Deploying a basic RAG-powered Q&A app
- Evaluation metrics: relevancy, latency, hallucination reduction
- Best practices for scaling RAG systems
- Privacy and compliance in RAG-powered enterprise AI
- Case studies: RAG in customer support, internal documentation bots

Module 24: Building a Local Retrieval-Augmented Generation (RAG) System on Intel AI PC

Build and deploy RAG systems locally using tools like FAISS and LangChain. This module leverages Intel® Core™ Ultra processors' hybrid architecture (CPU + GPU + NPU) to accelerate real-time data retrieval and AI-powered responses—enabling efficient, context-rich LLM applications.

Module Curriculum

- Setting Up the Intel AI PC Environment
  - Hardware Overview: Utilizing Intel AI Capabilities
  - Installing Necessary Software and Libraries
  - Configuring Development Tools
- Developing the RAG Pipeline
- Optimization Techniques for Intel Hardware
  - Leveraging Intel AI Accelerators for Enhanced Performance
  - Utilizing OpenVINO for Model Optimization
- Testing and Evaluation of the RAG System
- Deployment and Maintenance
  - Strategies for Local Deployment
  - Ensuring System Scalability and Reliability
  - Regular Maintenance and Updates
- Hands on: Building a local RAG based application

Module 25: Agentic AI – Autonomous AI Systems

Design intelligent agents that act, adapt, and make decisions independently. This module uses frameworks like LangChain, Auto-GPT, and CrewAI to build AI agents that automate tasks, interact with APIs, and operate with minimal human input in dynamic environments.

Module Curriculum

- What is Agentic AI and how it differs from traditional AI models
- Use cases of AI agents in research, productivity, business, and

automation

- Core components of an AI agent: memory, planning, reasoning, and tool use
- Frameworks for building agents
  - LangChain Agents, CrewAI, Auto-GPT, BabyAGI, SuperAGI
- Agent architecture and task decomposition
- Integrating tools and APIs into agent workflows
- ReAct (Reasoning + Acting) and other planning strategies
- Short-term vs long-term memory in agent systems
- Using vector databases and document loaders as knowledge tools
- Creating multi-agent systems with specialized roles
- Orchestration of concurrent agents for complex goals
- Logging, monitoring, and controlling autonomous behavior
- Security, constraints, and ethical boundaries in autonomous AI
- Hands-on: Build a research assistant or data analyst AI agent
- Future directions: self-improving and recursive agents

Section 5: MLOps & LLMOps

Module 26: Setting up Docker for AI Development

Learn to containerize AI apps with Docker for seamless development and deployment. This module covers building, running, and managing lightweight containers—making your AI workflows scalable, portable, and cloud-ready.

Module Curriculum

- Introduction to Docker
  - What is Docker and why it’s essential for AI development
  - Docker vs traditional virtual environments
- Installing and Configuring Docker
  - Installing Docker Desktop (Windows/Linux/Mac)
  - Understanding Docker Engine, CLI, and Docker Desktop interface

- Docker Fundamentals
  - Docker images vs containers
  - Writing your first Dockerfile
  - Using Docker Hub and pulling pre-built images
- Creating a Containerized AI Development Environment
  - Building a Python-based container with essential ML libraries (NumPy, pandas, scikit-learn, TensorFlow, etc.)
  - Persisting data with Docker volumes
  - Managing dependencies with requirements.txt
- Running and Managing Containers
  - Starting, stopping, and interacting with containers
  - Exposing ports and networking between containers
- Docker Compose for Multi-Service AI Projects
  - Setting up a docker-compose.yml
  - Running full-stack setups: UX + FastAPI + Database (Prisma + MySQL)
- Best Practices
  - Version control for Dockerfiles
  - Optimizing image size and build time
  - Sharing and collaborating on AI containers

Module 27: Cloud Deployment for LLM-Based Applications

Take your AI apps live with cloud deployment! This module covers the essentials of deploying LLM-powered solutions using AWS, Azure, or GCP. Learn containerization, serverless setups, cloud storage, and scalable model hosting to turn your prototypes into production-ready apps.

Module Curriculum

Introduction to cloud computing: key concepts and benefits
Overview of major platforms: AWS, Azure, GCP, and Intel Developer Cloud
Cloud infrastructure essentials: compute, storage, networking
Hosting and deploying containerized AI apps with Docker
Using cloud services to run FastAPI or Streamlit apps

- Deploying LLM APIs (e.g., OpenAI, Hugging Face) via cloud functions or REST endpoints
- Environment setup and configuration
- Introduction to serverless computing (Lambda, Azure Functions, Google Cloud Functions)
- CI/CD pipelines for AI app updates
- Working with secrets, environment variables, and model endpoints securely
- Logging, monitoring, and scaling applications in the cloud
- Using GPU instances to run inference with large models
- Cost management and optimization tips
- Hands-on: Deploy a simple LLM-powered chatbot or document Q&A system on the cloud
- Bonus: Introduction to Intel Developer Cloud for AI acceleration

Get Industry Ready with Top Tools & Technologies

Python

SQL

Computer Vision

Data Cleaning

Machine Learning

Deep Learning

Data Visualization

Artificial Intelligence

LLMs

Natural Language Processing

Generative AI

RAG

Agentic AI

MLOps

All-Round Career Support To Help You Get Hired

Our dedicated placement team and industry mentors guide you every step of the way, ensuring you're not just prepared, but truly job-ready.

Personalized Career Coaching

Tailored guidance to align your strengths with job market expectations.

AI-Powered Mock Interviews

Practice real-world technical interviews with AI-driven evaluation & expert feedback.

Company-Specific Prep Kits

Get curated interview prep resources tailored for top tech companies and roles.

1:1 Resume & LinkedIn Optimization

Build standout resumes and online profiles that get noticed.

Dedicated Referrals & Hiring Drives via HyreNet

Get connected to top recruiters through exclusive placement channels.

Earn Recognition that Employers Trust

Gain industry-recognized certifications from GUVI, Intel & AWS that enhance your credibility.

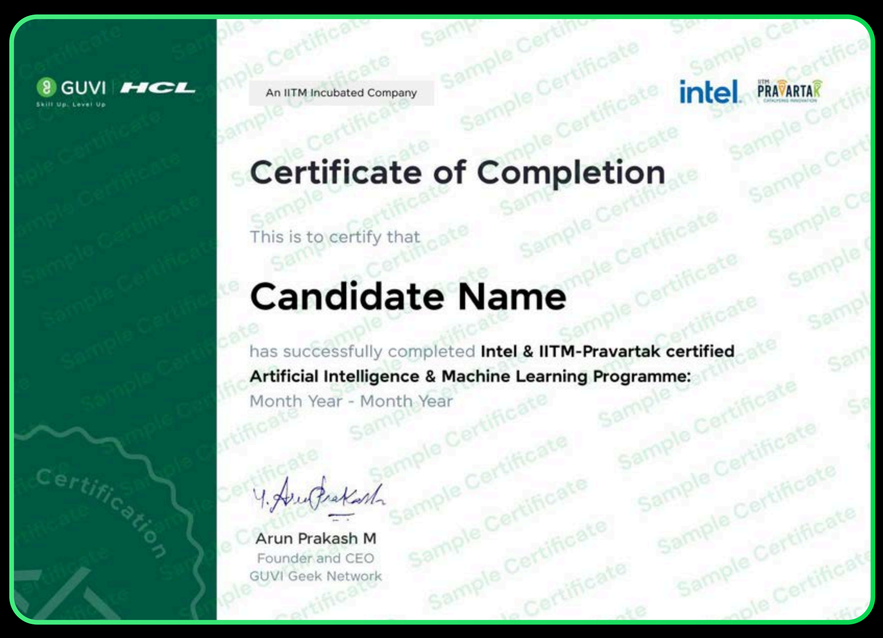

Intel Special Access

- Expert-led sessions conducted by Intel experts
- Labs & Real-world Projects designed for Intel AI PCs
- Joint certification from Intel, IITM Pravartak & HCL GUVI
- Exclusive learner discounts on Intel AI PCs

GUVI Career Acceleration

- Opportunity to join Campus Immersion Program at Intel
- Leadership and Management Masterclass by top tech leaders
- 12-month exclusive access to the HYRENET placement portal

NSDC Certification

- NSDC certificate adds government recognition to your qualification
- Improves your job prospects with validated, industry-relevant skills
- Enhance global employability through internationally aligned standards

Meet the Experts behind the Program

This program is thoughtfully curated by industry experts at HCL-GUVI, combining hands-on engineering experience with a deep understanding of emerging tech trends to prepare learners for real-world software development and AI integration.

Thillaikkarasan M

Lead Data Scientist,
Wells Fargo

Shyam Kumar

Machine Learning Solutions Lead,
Skit.AI

Kousik Krishnan

Quant Trader,
Independant Trading Desk

Jasmine Jeniffer J

Data Scientist,
Accenture

Get Hired with Our End-to-End Placement Support

- One-on-One Mock Interviews
- One-on-One Resume Evaluation
- One-on-One Mentor Guidance
- Interview Preparation Support
- Job Opportunities

Program Details

PROGRAM DURATION AND FORMAT

4 Month Weekday/6 Month Weekend

PROGRAM START DATES

Please contact our Intel & IITM Pravartak Certified AIML Program coordinator team : **+919736097320**

PROGRAM FEE

Total Course Fee **₹ 1,80,000**

Available EMI Options

Now become an AIML Professional at Affordable Instalments!

Master AIML at just **₹14,051*** /Month

Upto 12 Months

Note: Valid documents are required for the EMI Process. An additional processing fee will be applied. EMI Amount might vary with Vendors.

For Refunds, Registrations and Learning Policies Kindly **refer T&C**

About HCL-GUVI

GUVI is India’s first Vernacular EdTech platform of its kind. GUVI stands for ‘Grab Ur Vernacular Imprint’, dedicated to making technical education accessible and effective by breaking down language barriers. Our pioneering EdTech company is incubated by India's premier institutions, ensuring the highest standards of quality and innovation. We aim to make a significant impact in the field of tech upskilling, opening doors for learners across India to acquire valuable technical skills in their vernacular languages. By democratizing tech education online through prominent partnerships with Google-for-Education, UiPath, NASSCOM, & AICTE, GUVI has made it possible to impart job-ready tech skills to The ambitious aspirants.

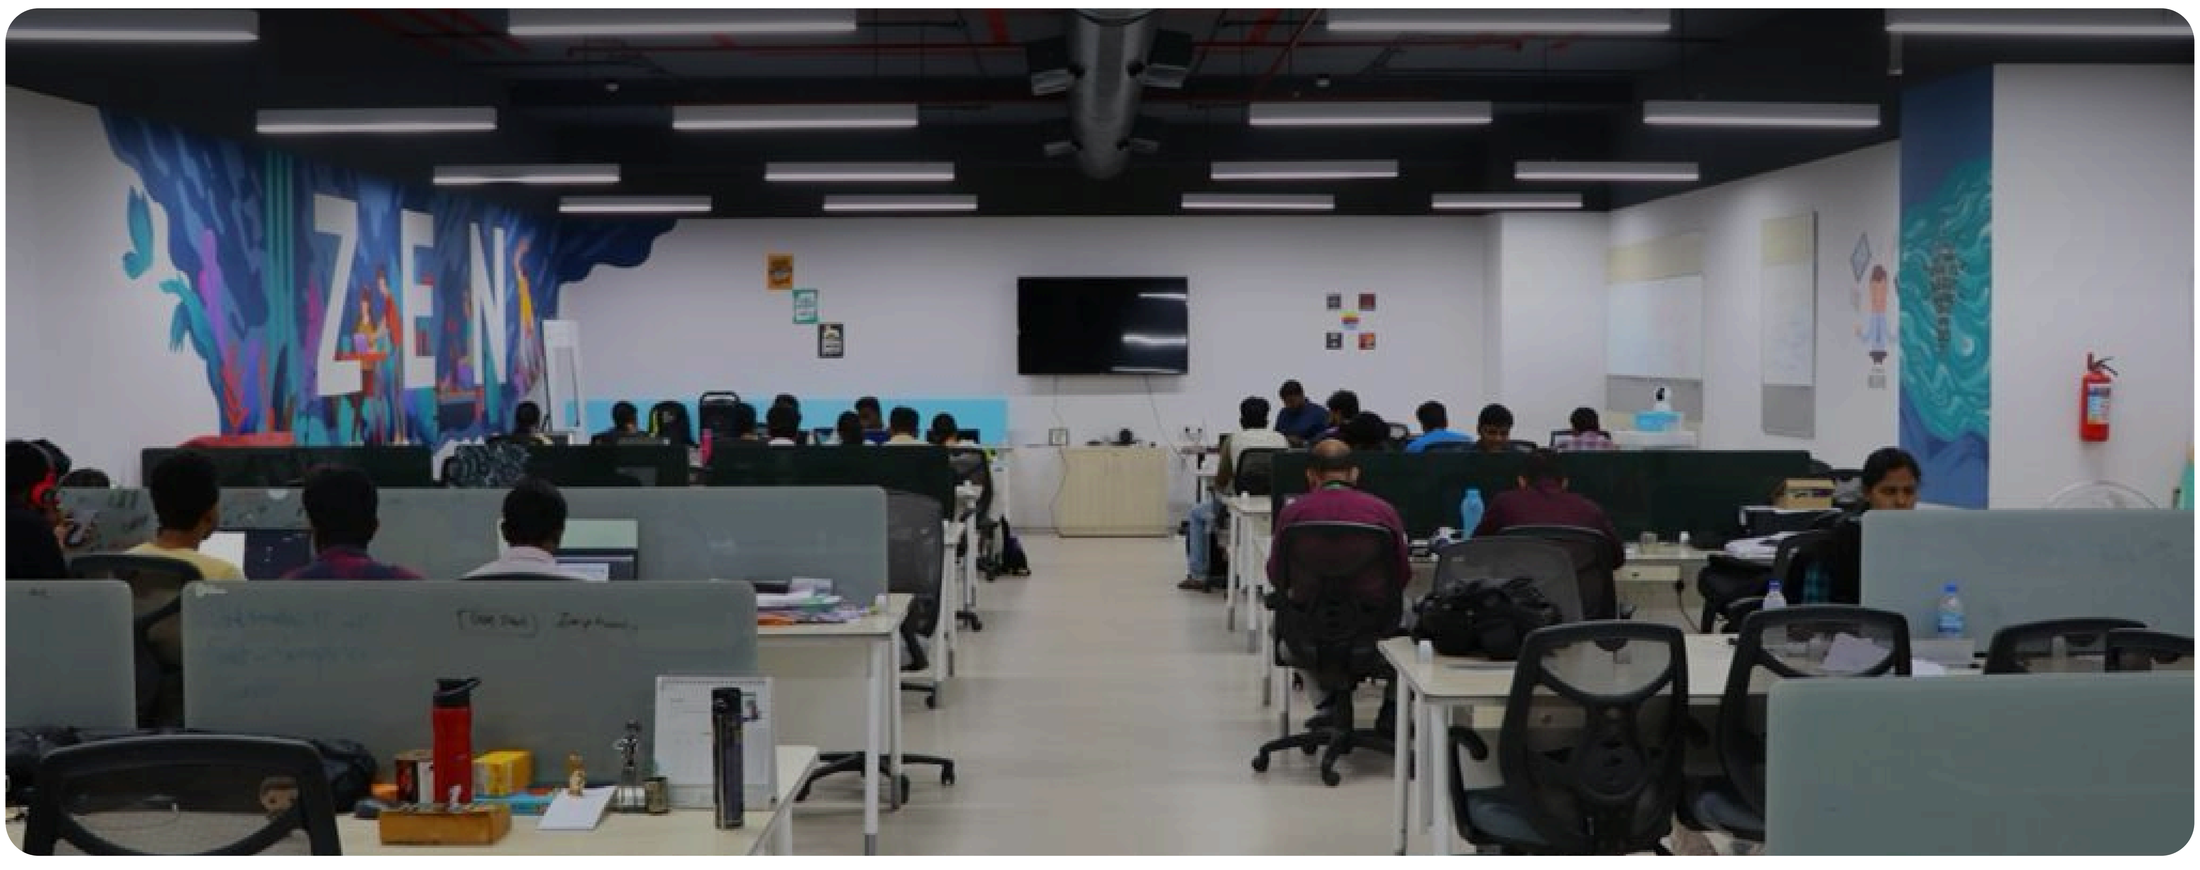

We are Accredited by

AICTE – All India Council for Technical Education

N·S·D·C National Skill Development Corporation – Transforming the skill landscape

NASSCOM®

ISO 27001-9001 CERTIFIED

Google for Education Partner

Built by Passionate Educators

Arun Prakash

CEO & Founder at HCL GUVI

20+ years of Technical Expertise

PayPal Honeywell Symantec

& more

Built 7 Products from Scratch Mentored

1000+ students Hosted 200+ sessions &

25+ webinars

Bala Murugan

Co - Founder at HCL GUVI

17+ years of experience with IT industry

PayPal

Technologist with 9+ years of Entrepreneurial experience & Member of the Syllabus Sub-Committee at Anna University

Get Ready to Launch Your AI-ML Career!

Join the **Intel & IITM Pravartak Certified Artificial Intelligence & Machine Learning Program** and equip yourself with the tools, mentorship, and industry-driven curriculum to code with confidence, build with purpose, and scale with impact.

Take the leap! Apply now ....

For more information, Reach out to us!

cs@guvi.in