EDWARTENS IN

ONLINE · SELF-PACED

SYLLABUS · REVIEWED 23 OCTOBER 2026

AIP

Artificial Intelligence Program

From your **first line of Python** to **AI on the factory floor**: machine learning, deep learning, computer vision, predictive maintenance and AI agents, with **5 graded industrial projects**. Then **placement support until you hold an offer letter**. No coding experience needed.

JOB-ASSURED · CONDITIONS APPLY · ASSURED job support until your offer letter

8	88	154h	5
phases	modules	of video	graded projects

UK REGISTERED LEARNING PROVIDER

Wartens Ltd, UK · UKPRN 10096641

UKRLP UK Register of Learning Providers · UKPRN : 10096641

ISO 9001:2015 certified

AWARD-WINNING

7 awards, 2025-2026

UKStartUp Awards 2025 National Winner · UKStartUp Awards 2025 Midlands Winner · The Sunday Times Great British Entrepreneur Awards 2025 FINALIST · radio LEMON Live 2025 WINNER · Business Awards UK 2025 FINALIST · 2025 WINNER · **Karnataka Excellence Award** Winner 2026

Job-assured, conditions apply: the conditions are in the course overview brochure and at edwartens.co.in/aip-terms.

GLOBAL PRESENCE India · United Kingdom · Germany · UAE · United States

edwartens.co.in/courses/aip · +91 96068 68686 · WhatsApp +91 81979 78653

HOW IT IS ORGANISED

Eight phases, at your own pace

88 modules of recorded lessons with notes and quizzes; a gate test closes each phase; the last phase is 5 graded projects, the final assessment and your career sessions. 12 months of access from the day you enrol.

PHASE 0 — Orientation and entry assessment
3 modules · 2 h | Entry assessment: 20 questions, 50% to pass

PHASE 1 — Python and programming fundamentals
9 modules · 11 h | Gate test: 25 questions, 70% to pass

PHASE 2 — Data science and analytics
12 modules · 20 h | Gate test: 25 questions, 70% to pass

PHASE 3 — Machine learning and deep learning
15 modules · 30 h | Gate test: 25 questions, 70% to pass

PHASE 4 — Computer vision and industrial AI
13 modules · 26 h | Gate test: 25 questions, 70% to pass

PHASE 5 — Generative AI, LLMs and agents
14 modules · 29 h | Gate test: 25 questions, 70% to pass

PHASE 6 — MLOps and deployment
13 modules · 24 h | Gate test: 25 questions, 70% to pass

PHASE 7 — Industrial AI projects and career
9 modules · 13 h | Final assessment: 50 questions, 70% to pass

Job-assured, conditions apply. The AIP is a job-assured programme. Once you pass the final assessment, placement support has no end date: it continues until an employer gives you an offer letter. Full terms: edwartens.co.in/aip-terms.

PHASE 0

Orientation and entry assessment

How the AIP works, what AI engineers do in a plant, and a short aptitude test to start you on the right footing. No coding needed.

0.1	**Welcome to the AIP: how a self-paced programme runs** The eight phases and what unlocks what; Gate tests: pass marks, attempts and waiting times; Projects and how they are graded; Booking your resume review and mock interviews; A study routine of 10 to 12 hours a week	25 min
0.2	**What an industrial AI engineer actually does** The industrial AI use-case map; Quality, maintenance, energy and process optimisation; The data a plant already has; AI roles and what each one does all day; Reading a job posting: must-have versus nice-to-have	39 min
0.3	**Thinking with numbers: the maths and logic AI is built on** Percentages, ratios and proportions; Mean, median and spread; Reading graphs and tables; Functions and rates of change; Basic probability	66 min

Entry assessment: 20 questions, 50% to pass

PHASE 1

Python and programming fundamentals

Python from the first line to classes, the Linux command line, Git and GitHub, and using an AI coding assistant without handing it your judgement.

1.1	**Python in Colab: variables, numbers and strings** Google Colab and notebooks; Variables and assignment; int, float, bool and str; String methods and f-strings; Type conversion	36 min
1.2	**Lists, tuples, sets and dictionaries** Lists, indexing and slicing; Tuples and immutability; Sets and membership; Dictionaries and key lookup; Nested data structures	94 min
1.3	**Conditions, loops and comprehensions** Comparison and boolean operators; if, elif, else; for loops and range; while loops; break, continue and else on loops	52 min
1.4	**Functions, modules and the standard library** def, parameters and return; Default and keyword arguments; Scope; Modules and imports; Standard library tour	72 min
1.5	**Files, CSV and JSON, errors and debugging** open() and with; csv and json modules; pathlib; Exceptions and try/except/finally; Reading a traceback	86 min
1.6	**Object-oriented Python: classes, objects and methods** Classes and instances; __init__ and self; Methods and attributes; Inheritance; Dunder methods	47 min
1.7	**The Linux command line for AI engineers** The terminal and shell; Filesystem and paths; Creating and editing files; Pipes and redirection; grep and searching	139 min
1.8	**Git and GitHub: version control and your portfolio** Why version control; init, add, commit, log; Branches and merges; Remotes, push and pull; Undoing changes	52 min
1.9	**Working with AI coding assistants, spec first** Agentic coding tools; Spec-first prompting; Reviewing generated code; Tests as guardrails; Secrets and confidentiality	109 min

Gate test: 25 questions, 70% to pass

PHASE 2

Data science and analytics

NumPy, Pandas and Matplotlib on real engineering data, SQL from SELECT to joins, the statistics a model rests on, where plant data comes from, and a Power BI report a plant head can read.

2.1	**NumPy: arrays, vectorisation and signals** ndarray, shape and dtype; Indexing, slicing and masks; Vectorised arithmetic; Broadcasting; Aggregations along an axis	95 min
2.2	**Pandas: DataFrames from real files** Series and DataFrame; read_csv and read_excel; loc and iloc; Filtering; describe and value_counts	65 min
2.3	**Cleaning, joining and grouping data** Missing values; Duplicates and outliers; Type conversion; merge and concat; groupby and agg	52 min
2.4	**Time series in Pandas: resample, rolling and shift** Datetime parsing; DatetimeIndex; resample; rolling windows; shift and diff	54 min
2.5	**Visualisation with Matplotlib and the reproducible report** Figure and axes; Line, bar, histogram, scatter; Subplots; Labels, units and legends; Choosing a chart	90 min
2.6	**SQL I: tables, keys and SELECT** Tables and data types; Primary and foreign keys; SELECT and WHERE; BETWEEN, IN, LIKE and IS NULL; ORDER BY and LIMIT	104 min
2.7	**SQL II: aggregates, joins and subqueries** Aggregate functions; GROUP BY and HAVING; INNER and LEFT JOIN; Subqueries and EXISTS; Views	139 min
2.8	**Statistics I: distributions, spread and capability** Data types and measurement scales; Central tendency and spread; The normal distribution; Sampling; Normality checks	63 min
2.9	**Statistics II: hypothesis tests, ANOVA and regression** Null and alternative hypotheses; p-values and significance; t-tests; ANOVA; Correlation	192 min
2.10	**Industrial data: where plant data lives, OPC UA and MQTT** The automation pyramid and data sources; OPC UA address space; OPC UA from Python; MQTT brokers and topics; Publish/subscribe	109 min
2.11	**Exploratory analysis of process data** Time-series storage; Cleaning process data; Resampling; Gaps and flat-lines; Outliers	168 min
2.12	**Dashboards a plant head reads: Power BI** Power BI Desktop and Power Query; Measures with DAX; Choosing visuals; Report layout and interaction; OEE and downtime pages	76 min

Gate test: 25 questions, 70% to pass

PHASE 3

Machine learning and deep learning

scikit-learn from regression to boosting and clustering, honest evaluation, then neural networks in PyTorch and Keras: CNNs, transfer learning and sequence models, and forecasting process and demand series.

3.1	**What machine learning is, for an engineer** Types of learning; Features and labels; Training versus inference; The scikit-learn API; Problem framing	241 min
3.2	**Linear regression and gradient descent** Linear regression; Multiple regression; Loss and MSE; Gradient descent; Learning rate	66 min
3.3	**Classification, splitting and leakage** Categorical encoding; Train/test split; Time-based splits; Logistic regression; Confusion matrix	63 min
3.4	**Decision trees, random forests and boosting** Decision trees; Overfitting in trees; Random forests; Gradient boosting; XGBoost and LightGBM	67 min
3.5	**SVM, KNN, Naive Bayes and feature scaling** Support vector machines; Kernels; K-nearest neighbours; Naive Bayes; StandardScaler and MinMaxScaler	53 min
3.6	**Evaluating models: cross-validation, bias, variance and metrics** Cross-validation; Bias-variance trade-off; Precision, recall and F1; ROC and AUC; Regression metrics	53 min
3.7	**Feature engineering, pipelines and tuning** Feature engineering; Outliers; Pipeline and ColumnTransformer; GridSearchCV and RandomizedSearchCV; L1 and L2 regularisation	86 min
3.8	**Unsupervised learning: clustering, PCA and anomalies** k-means; Choosing k; PCA; Explained variance; Isolation forest	49 min
3.9	**Neural networks: neurons, activations, loss and backpropagation** Neurons and layers; Activation functions; Loss functions; Gradient descent variants; Backpropagation	176 min
3.10	**PyTorch: tensors, autograd and the training loop** Tensors and devices; nn.Module; Loss functions and optimisers; The training loop; Evaluation mode	177 min
3.11	**Training well: data loaders, regularisation and early stopping** Dataset and DataLoader; Transforms; Dropout; Weight decay; Early stopping	119 min
3.12	**Convolutional neural networks** Convolution; Kernels, stride and padding; Pooling; Feature maps; CNN in PyTorch	101 min
3.13	**Transfer learning and augmentation** Pretrained models; Feature extraction versus fine-tuning; Freezing layers; Data augmentation; torchvision models	230 min
3.14	**Sequence models: 1-D CNNs, RNNs, LSTMs and autoencoders** Windowing; RNNs; LSTMs and GRUs; 1-D CNNs; Autoencoders	170 min
3.15	**Forecasting process and demand series** Trend, seasonality and stationarity; Naive and seasonal-naive baselines; ARIMA and exponential smoothing; Lag features with gradient boosting; Rolling-origin backtesting	137 min

Gate test: 25 questions, 70% to pass

PHASE 4

Computer vision and industrial AI

OpenCV and classical vision, cameras and lighting, YOLO detection for quality inspection, predictive maintenance and anomaly detection, and AI at the edge.

4.1	**OpenCV: images as arrays, colour and masks** Images as arrays; Colour spaces; Pixel arithmetic; Masks; Drawing	108 min
4.2	**Thresholds, morphology, edges and contours** Thresholding; Morphology; Smoothing and gradients; Canny edges; Contours	115 min
4.3	**Video, motion, tracking and counting** Video capture; Frame differencing; Background subtraction; Tracking; Line-crossing counters	76 min
4.4	**Cameras, lenses, lighting and calibration** Sensors and resolution; Lenses and field of view; Lighting techniques; Backlight, dome and dark-field; Camera calibration	180 min
4.5	**The inspection station: triggering, the PLC and classical tools** Inspection station anatomy; Triggering and timing; PLC handshake; Reject mechanisms; Classical inspection tools	64 min
4.6	**Datasets, labelling and training a YOLO detector** Dataset design; Labelling tools; Label quality; YOLO architecture; Training on a custom dataset	145 min
4.7	**Evaluating for the line and deploying vision at the edge** Escape rate and false reject rate; Cycle-time budgets; Edge devices for vision; YOLO on Raspberry Pi and Jetson; Monitoring a vision system	95 min
4.8	**Predictive maintenance: failure modes, signals and condition indicators** Failure modes; Sensors for condition monitoring; Sampling and aliasing; FFT and spectra; Condition indicators	76 min
4.9	**Anomaly detection on normal data** Learning normal; Statistical limits; Isolation forest; Autoencoders; Thresholds and false alarms	135 min
4.10	**Fault classification and remaining useful life** Fault classification; Labelling run-to-failure data; Remaining useful life; NASA C-MAPSS; RUL metrics	83 min
4.11	**Deep learning on sensor windows and the end-to-end workflow** Deep learning on windows; Feature-based versus end-to-end; The PdM workflow; Data pipelines; Validation on unseen machines	198 min
4.12	**Edge AI: hardware, latency and quantisation** Edge versus cloud; Microcontrollers, Raspberry Pi, Jetson; TinyML; Quantisation; Memory and latency budgets	162 min
4.13	**Alerts that get acted on, and measuring value** Alert design; Alarm fatigue; Hysteresis and confirmation; Measuring value; Deployment checklists	93 min

Gate test: 25 questions, 70% to pass

PHASE 5

Generative AI, LLMs and agents

How large language models work, prompting and structured output, retrieval-augmented generation with LangChain, local models with Ollama, tool-calling agents, LangGraph workflows and MCP, transformers and LoRA fine-tuning, and how to keep it all safe.

5.1	**How large language models work** What generative AI is; Tokens and probabilities; Transformers and attention; Pretraining and fine-tuning; Hallucination	111 min
5.2	**Choosing a model and calling it from Python** Model families; Hosted versus local; API keys and secrets; Chat messages and roles; Parameters	62 min
5.3	**Prompt engineering and structured output** Prompts as specifications; Zero-shot and few-shot; Step-by-step reasoning; System prompts; JSON and schemas	88 min
5.4	**Embeddings, vector search and the idea of RAG** Embeddings; Similarity search; Vector databases; What RAG fixes; Hybrid search	56 min
5.5	**Building a RAG chain with LangChain** LangChain models, prompts and chains; Document loaders; Chunking; Retrievers; The basic RAG chain	57 min
5.6	**Better retrieval: query rewriting, routing and re-ranking** Multi-query and RAG fusion; Decomposition and HyDE; Routing; Structured data; Re-ranking	65 min
5.7	**Evaluating RAG systems and prompts** Evaluation sets; Retrieval metrics; Faithfulness and relevance; LLM-as-judge and its limits; Prompt regression tests	178 min
5.8	**Local LLMs with Ollama and offline RAG** Why local; Ollama and the Modelfile; Quantisation formats; Hardware sizing; The local API	146 min
5.9	**Tool calling and AI agents** What an agent is; Function and tool calling; Tool schemas; The agent loop; Memory	114 min
5.10	**The Model Context Protocol and industrial data** MCP architecture; Tools, resources and prompts; Building an MCP server; MCP for OPC UA and Modbus; Read-only by design	141 min
5.11	**Guardrails, agent evaluation and LLM security** Human approval; Allow-lists and rate limits; Audit logs; Agent evaluation; Prompt injection	165 min
5.12	**Prompting, RAG or fine-tuning, and Hugging Face models** Prompting versus RAG versus fine-tuning; Hugging Face Hub and transformers; Pipelines; LoRA and parameter-efficient tuning; Images and multimodal models	100 min
5.13	**Agent workflows with LangGraph** Workflows against autonomous agents; LangGraph state, nodes and edges; Conditional edges and loops; Checkpoints, memory and human-in-the-loop; Corrective and agentic RAG	312 min
5.14	**Transformers hands-on, and fine-tuning with LoRA** Tokens, embeddings and attention; Encoder, decoder and encoder-decoder models; Fine-tuning with the Trainer API; LoRA adapters with PEFT; QLoRA and 4-bit quantisation	120 min

Gate test: 25 questions, 70% to pass

PHASE 6

MLOps and deployment

Taking a model out of the notebook: MLflow tracking and registry, pipelines, Docker, serving as an API, ONNX for line PCs, monitoring and drift, tests and CI/CD, the plant data path, FastAPI and demo apps, and the cloud.

6.1	**Why MLOps exists, and the working environment** The ML lifecycle; MLOps maturity; Virtual environments; Dependency pinning; Notebook to script	112 min
6.2	**Experiment tracking with MLflow** Why track experiments; MLflow tracking; Parameters, metrics and artefacts; Comparing runs; Reproducibility	132 min
6.3	**Model registry and model management** The model registry; Versions and aliases; Promotion and rollback; Loading from the registry; Model signatures	106 min
6.4	**From notebook to a re-runnable pipeline** Pipeline steps; Idempotence; Configuration; Scheduling; Orchestrators	82 min
6.5	**Docker for machine learning** Containers versus virtual machines; Images and registries; Dockerfiles; Docker Compose; Volumes	66 min
6.6	**Serving models: web services, batch and streaming** Deployment modes; REST APIs with Flask and FastAPI; Serving in Docker; Batch scoring; Streaming	93 min
6.7	**Exporting models: ONNX, TensorFlow Lite and line PCs** Saving models; torch.export and ONNX; ONNX Runtime; TensorFlow Lite; Numerical checks	97 min
6.8	**Monitoring models in production: drift and dashboards** What to monitor; Data drift; Concept drift; Monitoring services; Dashboards	87 min
6.9	**Tests, data quality checks and CI/CD** Unit tests; Data quality checks; Integration tests; Code quality; Makefiles	128 min
6.10	**The plant data path: MQTT, Unified Namespace and the historian** MQTT fundamentals; Unified Namespace design; Topic hierarchies and payloads; Telegraf, InfluxDB and Grafana; Publishing predictions	184 min
6.11	**Responsible AI and governance: risk, ISO 42001 and the EU AI Act** Trustworthy AI; NIST AI RMF; ISO/IEC 42001; EU AI Act risk tiers; Deployer duties	87 min
6.12	**Model APIs and demo apps: FastAPI, Streamlit and Gradio** FastAPI routes and Pydantic models; Validation, units and error responses; Loading the model once at start-up; Streamlit for data apps; Gradio Interface and Blocks	108 min
6.13	**Cloud for ML: AWS SageMaker and Azure Machine Learning** Regions and availability zones; Compute, storage and networking; Identity and access: IAM and managed identities; Managed endpoints on SageMaker and Azure Machine Learning; Microsoft Foundry and its model catalogue	128 min

Gate test: 25 questions, 70% to pass

PHASE 7

Industrial AI projects and career

Five graded industrial projects (quality inspection, predictive maintenance, SCADA analytics, an operator assistant and a deployed capstone), interview preparation, the final assessment and your booked resume and interview sessions.

7.1	**How the AIP projects are run and graded** The five projects; Rubrics and pass marks; Repository structure; The write-up; Resubmissions	63 min
7.P1	**Project 1: Visual quality inspection of surface defects** Defect datasets; Transfer learning or YOLO; Thresholds; Escape and false reject rates; Inspection report	170 min
7.P2	**Project 2: Predictive maintenance, from vibration to remaining life** Run-to-failure data; Features and windows; Validation on unseen units; Alert policy; Value estimate	63 min
7.P3	**Project 3: SCADA analytics, anomaly detection and an OEE dashboard** SCADA data; Anomaly detection; OEE; Downtime Pareto; Dashboard	88 min
7.P4	**Project 4: An operator assistant over equipment manuals (RAG)** Document ingestion; Retrieval; Citations and refusal; Evaluation set; Guardrails	199 min
7.P5	**Project 5 (capstone): Deploy and monitor an industrial model** Experiment tracking; Docker; Serving; Monitoring; Runbook	38 min
7.2	**Resume, LinkedIn and portfolio for AI roles** CV structure for AI roles; Keywords and ATS; Project bullets with results; GitHub profile; LinkedIn profile	31 min
7.3	**Technical interviews for ML and AI engineer roles** Interview rounds; Project walkthroughs; ML fundamentals questions; Python and SQL exercises; System design for ML, at fresher level	97 min
7.4	**HR interviews, offers and your first ninety days** Tell me about yourself; Common HR questions; Salary and relocation; Offer letters; The first ninety days	48 min

Final assessment: 50 questions, 70% to pass

Ready to start AIP?

Buy online and start straight away, or talk to admissions first.

Call
+91 96068 68686

WhatsApp
+91 81979 78653

Online
edwartens.co.in/courses/aip

Centre
NVR Complex, GS Palya Road, Electronic City, Bangalore 560100, Karnataka

Industry photographs are free-licence stock (Peter Xie (Pexels), Mohammad Rahmani (Unsplash), Noman Khan (Unsplash)) and illustrate the work, not EDWartens premises. Vendor and tool names and logos are trademarks of their owners and are used only to describe what the course teaches; EDWartens is not affiliated with or endorsed by them. EDWartens India is the training brand of Wartens Automation Private Limited. This syllabus follows the curriculum the portal runs on 23 October 2026; topics may be updated as software and standards change.