

AI/ML DECODED | FREE STRATEGY SESSION

Don't Fear AI/ML

Complete syllabus breakdown · Exactly how deep to study · How Oliveboard will cover it

Phase I: 15 October 2026 · PFRDA Grade A Complete Course

14

**syllabus
heads**

5

**conceptual
buckets**

6

**Oliveboard
classes**

Download FREE PDF

Why is everyone panicking?

1

Completely new territory

AI/ML has never appeared in the General Stream syllabus before. There is no previous-year reference to fall back on.

2

The vocabulary sounds technical

LLM, Agentic AI, Responsible AI, Cross-Validation, Reinforcement Learning — unfamiliar words, so they feel advanced.

3

Nobody knows the expected depth

Aspirants cannot tell whether PFRDA wants computer-science depth or only conceptual understanding.

The reaction we keep hearing

"Do I need coding?"

"Do I need Python?"

"Is this a CS degree now?"

"Will they ask algorithms?"

The honest answer: none of the above. The syllabus is new — not difficult.

The one message of this session

AI/ML looks intimidating because the terminology is unfamiliar — not because the syllabus is unmanageably deep.

By the end of the next 40 minutes you will know

- 1 Exactly what is in the syllabus
- 2 Exactly how deep to go
- 3 What you should NOT study
- 4 How Oliveboard covers it

Don't panic. Just follow the syllabus — and the right depth.

The official syllabus — all 14 heads

Paper 2, General Stream · common to Phase I and Phase II · exactly as printed in the notification

Machine Learning	Supervised Learning	Unsupervised Learning	Data Preprocessing	Model Evaluation
Introduction to ML Model	NLP	Sentiment Analysis	Overfitting & Underfitting	LLM
Agentic AI	Responsible AI	Cross-Validation	Reinforcement Learning	

Fourteen topics. That is the entire section.

Do not expand this list yourself — the biggest mistake you can make is turning a 14-topic syllabus into a computer-science degree.

Download FREE PDF

What PFRDA has NOT said

None of these words appear anywhere in the AI/ML syllabus — so do not build your preparation around them

Python

TensorFlow / PyTorch

Neural Networks

Deep Learning architectures

Coding / programming

Data Structures

Algorithm design

Mathematical derivations

An honest caveat: the notification does not prescribe question depth, so nobody can guarantee the exact difficulty PFRDA will choose. Our strategy is to comprehensively cover the conceptual and application layer suggested by these syllabus terms — without going into programming or engineering depth.

Master what is named. Ignore what is not.

14 topics → just 5 conceptual buckets

Group the syllabus this way and the anxiety drops immediately

BUCKET 1 ML Foundations	BUCKET 2 Data & Evaluation	BUCKET 3 Language AI	BUCKET 4 New-Age AI	BUCKET 5 Another Paradigm
• Machine Learning • Introduction to ML Model • Supervised Learning • Unsupervised Learning What is ML, and what are the different ways a machine learns?	• Data Preprocessing • Model Evaluation • Overfitting • Underfitting • Cross-Validation How do we prepare data and judge whether a model actually works?	• NLP • Sentiment Analysis • LLM How do machines understand and generate human language?	• Agentic AI • Responsible AI How are AI systems becoming autonomous — and what safeguards apply?	• Reinforcement Learning How can a machine learn through rewards and penalties?

Stop looking at fourteen scary terms. Look at five ideas.

Exactly how deep should you study?

One formula. Apply it to every single topic in the syllabus — and stop there.



WORKED EXAMPLE • Supervised Learning

What it means:
learning from
labelled data.

Labelled dataset,
input and output,
classification,
regression.

How it differs
from
Unsupervised
Learning.

Fraud vs Not
Fraud; predicting
house price.

Loan default
prediction from
borrower history.

That is the depth. No mathematical derivations, no code.

Machine Learning & Intro to ML Model

Depth required: conceptual + application. Two syllabus heads, one continuous story.

Machine Learning — must know

- Meaning of ML · AI vs ML · why ML is used
- Data → learning → prediction
- The basic ML workflow
- Model · training · prediction / inference

Introduction to ML Model — must know

- A model learns patterns from training data and uses them on new data
- Features · target · training data · test data
- Prediction · generalisation
- No model architecture needed

EXAMPLE Loan default prediction: historical borrower data → model learns patterns → new borrower data → predicted probability of default.

PFRDA has listed “Introduction to ML Model” separately — so treat it as its own topic, not a footnote to ML.

Supervised vs Unsupervised Learning

The single most important comparison in this bucket — you should identify the type from a scenario in five seconds.

Supervised Learning

MUST KNOW

- **Labelled dataset** · input and output
- **Classification** · regression
- **The machine learns from known answers**

EXAMPLE **Classification:** Fraud / Not Fraud
Regression: predicting house price or revenue

Unsupervised Learning

MUST KNOW

- **Unlabelled data** · pattern detection
- **Clustering** · segmentation
- **The machine finds structure on its own**

EXAMPLE **A bank has millions of customers and no categories. ML groups them by behavioural similarity.**

LIKELY QUESTION A model is trained on historical transactions already labelled fraudulent or genuine. What type of learning is this?

Supervised Learning.

Data Preprocessing

It sounds technical. It is actually the most intuitive topic in the syllabus — cleaning data before a model sees it.

Why raw data is unusable	Missing values	Duplicate records	Incorrect / inconsistent entries
Outliers	Categorical data	Numerical data	Scaling / normalisation

THE INTUITION Imagine a customer income column that looks like this

₹50,000	₹85,000	(blank)	₹7,00,00,000
---------	---------	-----------	--------------

A blank value and an implausible outlier sit in the same column. Before this dataset reaches a model, it has to be cleaned. That is preprocessing.

Depth needed: recognise the problem and name the fix. No code, no formulas.

Model Evaluation

The core question: how do we know whether a model is actually good? Keep the four metrics intuitive.

Accuracy

Of all predictions made, how many were correct?

Precision

Of the cases the model flagged, how many were genuinely right?

Recall

Of all the real cases that existed, how many did the model catch?

F1 Score

One balanced number that combines precision and recall.

THE FRAUD TRAP

1,000 transactions · 10 are fraud · the model labels everything “genuine” → Accuracy = 99%, frauds caught = 0.

Accuracy alone can mislead. That one example explains why precision, recall and F1 exist.

Overfitting, Underfitting & Cross-Validation

The most MCQ-friendly area of the entire section. Learn this table.

	Underfitting	Good Model	Overfitting
Model complexity	Too low	Appropriate	Too high
Training performance	Poor	Good	Excellent
Test / unseen performance	Poor	Good	Poor

CROSS-VALIDATION One train-test split can mislead. Split the data repeatedly, test repeatedly, average the result. Know: purpose · basic idea · K-fold · link to generalisation.

Overfitting = memorising. Underfitting = not learning. Cross-validation = checking honestly.

NLP & Sentiment Analysis

How machines work with human language — the most familiar and most scoring part of this section.

NLP — Natural Language Processing

Technology that lets computers process and work with human language.

Applications: chatbots · translation · text summarisation · question answering · information extraction · text classification

Sentiment Analysis

Judging whether a piece of text is positive, negative or neutral — an application of language processing.

Used on: customer feedback · news articles · social-media posts · investment sentiment

REGULATORY ANGLE PFRDA receives thousands of grievances. NLP can categorise complaints, detect repeated issues, summarise text and route cases.

LIVE QUESTION “I have been waiting for my withdrawal request for 18 days and no one has responded.”

Sentiment: Negative.

LLM — Large Language Models

Start by decoding the name itself — students remember the topic far better that way.

LARGE = scale of training data, parameters and computation, and capability
language

LANGUAGE MODEL = it works with human

Must know

- What an LLM is
- The broad training idea
- Prompt · token · context
- Generated output
- Connection to Generative AI

Capabilities

- Drafting
- Summarisation
- Translation
- Question answering
- Information extraction
- Classification

Risks & limitations

- Hallucination
- Bias
- Privacy concerns
- Incorrect output
- Heavy dependence on context

Traditional ML ≠ Generative AI ≠ LLM. Know the distinction — but don't over-engineer the taxonomy.

Agentic AI

The term students will search the most — and the one that sounds far harder than it is.

AI that works towards a goal by planning and taking a sequence of actions — rather than merely generating one response.

Goal	Planning	Memory / context	Tools	Actions	Feedback
LLM			Agentic AI		
- Generates an answer - Prompt → response - Primarily produces content - Lower autonomy			- Pursues an objective - Multi-step process - Can take actions - Greater autonomy		

EXAMPLE Normal AI: “Summarise this report.” Agentic AI: find the key risks → check relevant information → summarise → prepare recommendations → create follow-up tasks.

Responsible AI

Position this as easy marks — seven principles, all of them plain English.

Fairness

Transparency

Explainability

Accountability

Privacy

Safety

Human oversight

Why it matters →

Imagine AI decides whether somebody receives a loan.

What if the historical training data is biased? What if the model discriminates against a group? What if nobody can explain why the loan was rejected? What if personal financial data leaks?

HIGH-PROBABILITY MCQ Which principle relates most closely to understanding why an AI took a particular decision?
Explainability.

Reinforcement Learning

Teach it through games, never through mathematics. Learning by rewards and penalties.



Takes an action → receives feedback → changes behaviour → improves

	Supervised	Unsupervised	Reinforcement
Signal used	Labels	No labels	Reward / penalty
What it learns	Correct outputs	Hidden patterns	Which actions to take
Typical tasks	Classification, regression	Clustering	Sequential decisions

If a student can read a scenario and name the paradigm in five seconds, this bucket is done.

The trap list — what NOT to study

Screenshot this slide. Every hour spent here is an hour taken away from the syllabus that is actually printed.

Python programming

TensorFlow

PyTorch

Coding ML algorithms

Mathematical derivations

Advanced probability & statistics

Neural-network architecture

CNNs

RNNs

Transformer architecture in
engineering depth

Gradient descent mathematics

Backpropagation

The notification itself does not list these areas. Unless later PFRDA communication indicates otherwise, master the concepts explicitly named in the syllabus first.

You are preparing for a regulatory exam — not a data-science interview.

So, what could PFRDA actually ask?

Eight realistic question patterns. Notice how little technical depth each one needs.

Q1

A system learns from examples where every transaction is labelled "fraud" or "not fraud."

Supervised Learning

Q2

Customers are automatically divided into groups by spending behaviour, with no existing labels.

Unsupervised Learning

Q3

A model performs extremely well on training data but poorly on unseen data.

Overfitting

Q4

Which technique estimates performance by repeatedly dividing data into training and validation sets?

Cross-Validation

Q5

Analysing whether a customer complaint expresses positive or negative emotion is called:

Sentiment Analysis

Q6

An AI system independently plans several steps and uses tools to complete a goal.

Agentic AI

Q7

Ensuring AI does not discriminate against certain customer groups relates most closely to:

Fairness / Responsible AI

Q8

An agent improves its behaviour based on rewards and penalties.

Reinforcement Learning

Understand the concepts properly and none of these feel technical.

How Oliveboard will cover AI/ML

A tight, bounded teaching block — built exactly around the 14 syllabus heads.

4

Concept classes

Every syllabus head taught with definition, comparison, example and application.

1

Complete revision class

The whole module compressed — comparison tables, flowcharts, one-liners, common traps.

1

MCQ marathon

100+ exam-oriented questions: direct, application, scenario, statement and comparison based.

= 6 focused classes. No bloated computer-science course. No 20-hour detour.

Complete AI/ML module finished by 7 October — eight days before Phase I.

Download FREE PDF

The four concept classes

Every one of the 14 syllabus heads is mapped to a class — nothing left floating.

CLASS 1 ML Foundations	CLASS 2 Building & Evaluating Models	CLASS 3 NLP & Generative AI	CLASS 4 New-Age AI
AI vs ML · Machine Learning · ML workflow · Introduction to ML Model · features and target · Supervised Learning · classification · regression · Unsupervised Learning · clustering · Supervised vs Unsupervised	Data Preprocessing · data quality · missing values · outliers · training and test data · Model Evaluation · accuracy · precision · recall · F1 score · Overfitting · Underfitting · Cross-Validation	NLP · NLP applications · Sentiment Analysis · Generative AI basics · LLM · prompt · token and context · capabilities · hallucination · LLM limitations	Agentic AI · AI agents · LLM vs Agentic AI · Responsible AI · fairness · transparency · accountability · explainability · privacy · Reinforcement Learning · agent, environment, reward

Four classes. Complete conceptual coverage. No detours.

Revision, MCQ marathon & course assets

Teaching is only half the job — the other half is repetition and question practice.

CLASS 5 · Complete Revision

The whole module compressed: ML → Model → Evaluation → NLP → LLM → Agentic AI → Responsible AI → RL

Plus comparison tables, flowcharts, one-liners and the common traps.

CLASS 6 · MCQ Marathon

100+ exam-oriented questions in one sitting.

Direct concept · application · scenario · statement-based · comparison questions.

EVERY ENROLLED STUDENT ALSO GETS

Lecture notes

Revision notes

MCQ practice sets

One-liner revision sheet

AI/ML questions inside
Paper 2 mocks

Learn it once. Revise it three times. Practise it enough.

Download FREE PDF

How much time should you personally give?

Realistic numbers, so you can slot AI/ML into a schedule that already has five other subjects in it.

8–10

focused hours

First learning — classes plus notes. That is the whole first pass.

150–200

quality MCQs

Across course tests and mocks, built up over time — not in one day.

3

revision rounds

Immediately after the classes · 7–10 days later
· final round before Phase I.

Do NOT study AI/ML a little every day for a month. Finish it quickly — then revise it repeatedly.

Target state before Phase I: “I have revised this multiple times and practised enough questions.”

PFRDA AI/ML — your roadmap

One slide to screenshot and keep: the order to study in, what to skip, and the method to use.

STUDY — in this order

1 ML basics	6 LLM
2 Supervised vs Unsupervised	7 Agentic AI
3 Data + Model Evaluation	8 Responsible AI
4 Overfitting + Cross-Validation	9 Reinforcement Learning
5 NLP + Sentiment Analysis	

DON'T STUDY

Coding-heavy ML
Deep mathematics
Advanced architectures not mentioned in the syllabus

METHOD — for every single topic

Concept → Comparison → Example → Application → MCQ

The syllabus is bounded. Follow the order, respect the depth, and this section is finished.

An unfamiliar syllabus — not a difficult one

Yes, the topics are modern. No, you do not need to become a data scientist in five weeks.

You only need to understand what these technologies mean, how they differ, where they are used, and what their limitations and risks are.

AND THAT IS EXACTLY WHAT WE WILL DO

4

Concept classes

1

Complete revision

1

MCQ marathon

7 Oct

Coverage complete

Don't fear the new topic. Follow the syllabus, the right depth, and a system.

AI/ML classes begin 14 September · Oliveboard PFRDA Grade A Complete Course

Download FREE PDF



PFRDA Gr. A 2026 Complete Course

Phase 1 + 2 Prep!



Medium: English

- 120+ QRE & GA Video Lectures
- 100+ Live Paper 2 Classes
- Paper 2 Subject-Wise Notes
- Full Length & Descriptive Mocks

Enrol Now