AI PRODUCT MANAGEMENT 2 · DETAILED COURSE

Artificial Intelligence and Machine Learning Foundations

How AI and machine learning work, the main problem types and how models are trained and judged.

Chapter 11 of 26 · 32 topics · Edition 2026 · skillnavora.in

CHAPTER 11

Artificial Intelligence and Machine Learning Foundations

How AI and machine learning work, the main problem types and how models are trained and judged.

Introduction to Artificial Intelligence

Systems that perform tasks associated with human intelligence, such as perception, language and decision making.

AI refers to systems that perform tasks associated with human intelligence, such as perceiving images, understanding language and making decisions. Most AI in products today is machine learning, which learns patterns from data.

History and Evolution of AI

Rules and symbolic AI, expert systems, statistical ML, deep learning from 2012, transformers from 2017, and generative AI.

The field moved from rule-based and symbolic systems to expert systems, then statistical machine learning, deep learning from around 2012, transformer models from 2017 and today's generative models. Each shift changed what products could do.

Traditional AI vs Generative AI

Traditional models classify or predict from data; generative models create new text, images, audio or code.

Traditional models classify or predict from data, such as flagging fraud. Generative models create new text, images, audio or code. Traditional AI outputs a label or number; generative AI outputs content that needs review.

Narrow AI vs General AI

Narrow AI does specific tasks; general AI would match human breadth, and does not yet exist.

Narrow AI performs specific tasks well, such as translation or recommendation. General AI would match human breadth across tasks. Only narrow AI exists today, so products should be designed around specific tasks.

AI Applications Across Industries

Fraud detection, recommendations, diagnostics, document processing, support automation, forecasting and more.

Examples include fraud detection in banking, recommendations in e-commerce, diagnostics in healthcare, document processing in insurance, support automation and demand forecasting in retail.

Machine Learning Fundamentals

Algorithms that learn patterns from data to predict or decide without explicit rules.

Machine learning algorithms learn patterns from data to predict or decide without explicit rules for every case. Quality depends mostly on data quality and on defining the target correctly.

Supervised Learning

Learning from labelled examples to predict labels for new data.

Supervised learning trains on labelled examples to predict labels for new data. Example: historical loan applications labelled repaid or defaulted, used to predict risk for a new applicant.

Unsupervised Learning

Finding structure in unlabelled data, such as clusters or anomalies.

Unsupervised learning finds structure in unlabelled data, such as clusters of similar customers or unusual transactions. Results need interpretation, because no correct answer is given.

Semi-Supervised Learning

Using a small labelled set plus a large unlabelled set.

It uses a small labelled dataset together with a large unlabelled one. This helps when labelling is expensive, such as medical images labelled only by specialists.

Reinforcement Learning

An agent learns by trial and error to maximise cumulative reward.

An agent learns by trial and error to maximise cumulative reward. It suits sequential decisions such as pricing or robot control, but needs a safe simulation or careful rollout.

Classification

Predicting a category, such as spam or not spam.

Classification predicts a category, such as spam or not spam, or churn or stay. Its output is usually a probability, which the product converts into a decision using a threshold.

Regression

Predicting a continuous number, such as price or demand.

Regression predicts a continuous number, such as delivery time or demand next week. Errors are measured as average distance from the true value.

Clustering

Grouping similar items without labels, used for segmentation.

Clustering groups similar items without labels. It is used for customer segmentation and for finding natural groups in support tickets. Clusters must be named and checked by people.

Recommendation Systems

Suggesting items using collaborative filtering, content features or hybrid models.

Recommenders suggest items using collaborative filtering (what similar users chose), content features, or hybrid models. They need careful evaluation for diversity and for avoiding narrow feedback loops.

Predictive Analytics

Using historical data to forecast outcomes such as churn or demand.

Predictive analytics uses historical data to forecast outcomes such as churn, demand or default. A forecast is only useful if someone acts on it, so define the intervention before building the model.

Feature Engineering

Creating informative inputs from raw data to help models learn.

Feature engineering creates informative inputs from raw data, such as days since last purchase. Good features often matter more than a more complex algorithm.

Training, Validation and Testing

Train to fit, validate to tune, and test once on unseen data for an honest estimate.

The model is trained on one set, tuned on a validation set and evaluated once on an unseen test set. Keeping the test set untouched gives an honest estimate of performance.

Model Training and Inference

Training learns parameters from data; inference uses the trained model to make predictions.

Training learns model parameters from data. Inference uses the trained model to make predictions on new inputs. Inference cost and latency matter in production even when training is done once.

Overfitting and Underfitting

Overfitting memorises training noise and fails on new data; underfitting is too simple to capture the pattern.

Overfitting memorises training noise and performs poorly on new data. Underfitting is too simple to capture the pattern. Compare training and validation performance to detect either.

Bias-Variance Trade-Off

Simple models have high bias; complex models have high variance; aim for the balance with best generalisation.

Simple models have high bias and miss patterns. Complex models have high variance and react to noise. The aim is the level of complexity that generalises best to new data.

Model Performance Metrics

Accuracy, precision, recall, F1, AUC for classification; MAE and RMSE for regression, picked to match business cost of errors.

Classification uses accuracy, precision, recall, F1 and AUC. Regression uses mean absolute error and root mean squared error. Pick the metric that matches the business cost of each kind of error.

Precision, Recall and F1 Score

Precision is correctness of positive predictions; recall is coverage of actual positives; F1 balances both.

Precision is the share of positive predictions that are correct. Recall is the share of actual positives found. F1 combines both. In fraud detection, recall often matters more; in automatic account blocking, precision matters more.

Confusion Matrix

A table of true/false positives and negatives that shows the types of errors.

A confusion matrix tabulates true positives, false positives, true negatives and false negatives. It shows what kinds of mistakes a model makes, which a single accuracy number hides.

ROC-AUC

Area under the curve of true-positive vs false-positive rate across thresholds; 0.5 is chance, 1.0 is perfect.

The ROC curve plots true-positive rate against false-positive rate across thresholds. AUC is the area under it: 0.5 means chance and 1.0 means perfect ranking.

Deep Learning Fundamentals

Multi-layer neural networks that learn features automatically from large data.

Deep learning uses multi-layer neural networks that learn features directly from large data, such as images or text. It needs more data and compute than classical methods.

Neural Networks

Layers of connected units with weights, trained by backpropagation and gradient descent.

Neural networks connect layers of units with weights. Training adjusts the weights with backpropagation and gradient descent to reduce error.

CNN, RNN and LSTM

CNNs process images, RNNs process sequences, LSTMs handle longer dependencies in sequences.

CNNs process images using local patterns. RNNs process sequences step by step. LSTMs handle longer dependencies in sequences. Transformers have largely replaced RNNs for language tasks.

Transformers

An architecture based on attention that processes sequences in parallel and underlies modern language models.

Transformers process sequences in parallel using attention. They underpin modern language and multimodal models, and they scale well with data and compute.

Attention Mechanisms

A way for a model to weigh which parts of the input matter for each output.

Attention lets the model weigh which parts of the input matter for each output. It is the reason models can relate a word to a distant word in the same document.

Natural Language Processing

Techniques for understanding and generating human language, from tokenisation to LLMs.

NLP covers understanding and generating human language, from tokenisation and classification to question answering. Large language models now handle most NLP tasks in products.

Computer Vision Fundamentals

Techniques for classifying, detecting and segmenting content in images and video.

Computer vision classifies, detects and segments content in images and video. Applications include document scanning, quality inspection and product tagging.

AI Model Lifecycle

Problem framing, data, training, evaluation, deployment, monitoring and retraining.

The lifecycle covers problem framing, data collection, training, evaluation, deployment, monitoring and retraining. The PM owns framing, the success definition and the monitoring plan.

Summary of terms

Term	Meaning
Introduction to Artificial Intelligence	Systems that perform tasks associated with human intelligence, such as perception, language and decision making.
History and Evolution of AI	Rules and symbolic AI, expert systems, statistical ML, deep learning from 2012, transformers from 2017, and generative AI.
Traditional AI vs Generative AI	Traditional models classify or predict from data; generative models create new text, images, audio or code.
Narrow AI vs General AI	Narrow AI does specific tasks; general AI would match human breadth, and does not yet exist.
AI Applications Across Industries	Fraud detection, recommendations, diagnostics, document processing, support automation, forecasting and more.
Machine Learning Fundamentals	Algorithms that learn patterns from data to predict or decide without explicit rules.
Supervised Learning	Learning from labelled examples to predict labels for new data.

Unsupervised Learning	Finding structure in unlabelled data, such as clusters or anomalies.
Semi-Supervised Learning	Using a small labelled set plus a large unlabelled set.
Reinforcement Learning	An agent learns by trial and error to maximise cumulative reward.
Classification	Predicting a category, such as spam or not spam.
Regression	Predicting a continuous number, such as price or demand.
Clustering	Grouping similar items without labels, used for segmentation.
Recommendation Systems	Suggesting items using collaborative filtering, content features or hybrid models.
Predictive Analytics	Using historical data to forecast outcomes such as churn or demand.
Feature Engineering	Creating informative inputs from raw data to help models learn.
Training, Validation and Testing	Train to fit, validate to tune, and test once on unseen data for an honest estimate.
Model Training and Inference	Training learns parameters from data; inference uses the trained model to make predictions.
Overfitting and Underfitting	Overfitting memorises training noise and fails on new data; underfitting is too simple to capture the pattern.
Bias-Variance Trade-Off	Simple models have high bias; complex models have high variance; aim for the balance with best generalisation.
Model Performance Metrics	Accuracy, precision, recall, F1, AUC for classification; MAE and RMSE for regression, picked to match business cost of errors.
Precision, Recall and F1 Score	Precision is correctness of positive predictions; recall is coverage of actual positives; F1 balances both.
Confusion Matrix	A table of true/false positives and negatives that shows the types of errors.
ROC-AUC	Area under the curve of true-positive vs false-positive rate across thresholds; 0.5 is chance, 1.0 is perfect.
Deep Learning Fundamentals	Multi-layer neural networks that learn features automatically from large data.
Neural Networks	Layers of connected units with weights, trained by backpropagation and gradient descent.
CNN, RNN and LSTM	CNNs process images, RNNs process sequences, LSTMs handle longer dependencies in sequences.
Transformers	An architecture based on attention that processes sequences in parallel and underlies modern language models.
Attention Mechanisms	A way for a model to weigh which parts of the input matter for each output.
Natural Language Processing	Techniques for understanding and generating human language, from tokenisation to LLMs.
Computer Vision Fundamentals	Techniques for classifying, detecting and segmenting content in images and video.
AI Model Lifecycle	Problem framing, data, training, evaluation, deployment, monitoring and retraining.

Practice exercise: Train a simple classifier on a public dataset and report precision, recall and F1 at two thresholds. Say which one the business should pick.

Self-check

1. Explain Introduction to Artificial Intelligence in two sentences to a non-technical stakeholder, then say what you would measure for it.
2. Explain Training, Validation and Testing in two sentences to a non-technical stakeholder, then say what you would measure for it.
3. Explain AI Model Lifecycle in two sentences to a non-technical stakeholder, then say what you would measure for it.

Related SkillNavora material

- Machine Learning Products for PMs : ML concepts, metrics and lifecycle.
- AI Tutor : Ask questions on any topic.
- Study Library : Every lesson, class and PDF by programme.

© SkillNavora · skillnavora.in · Facts and tools change; check official sites before relying on a figure.

ORIGINAL CONTENT

Original AI Product Management content

The AI Product Management programme already on SkillNavora, included unchanged so the full course is in one place.

Quick links Study Library (all material) Learning sites for PM AI builder tools (how-to) Webhooks & Playwright MCP Webhook URLs and keys (start here) AI Skills Academy AI Skills Map Context engineering (PDF) AI skills & templates (PDF) Orchestration (PDF) PRD library (120 ideas) AI PRD Writer AI PM Toolkit AI Evals (PDF) Browser automation (PDF) AI second brain (PDF) PM handbook Top 100 PM questions AI product sense Q&A (86) PM practice test Practice Lab Tool tutorials AI Agents & MCP for PMs

PRD documents

PRD library: template, framing guide and 120 PRD ideas

Learn how top PMs frame a PRD, download the free template, and turn any of 120 ideas across 15 industries into a full PRD with AI.

Open the PRD library ✨ AI PRD Writer PRD template (Word) AI Evals for PMs (PDF)

Context window

Context window: what a model can see at once

The context window is the most tokens a model can consider in one request. It covers your prompt, any documents you paste, the conversation so far and the reply. Anything outside it is invisible to the model. For PMs: a longer window fits more material but costs more and can slow responses; put the most important facts near the instructions; use retrieval (RAG) rather than pasting whole libraries; and test answers on long inputs, where quality can drop.

Context engineering course Context windows in the AI PM 2 course

Tokens and context in the Grokbot course

AI-First Product Management

Resources for Product Managers

The PM Master Handbook, AI and product tools with their official websites, interview practice, a frameworks cheat sheet, a PRD template and free courses.

Open the PM handbook library: all 68 PDFs ↗ PM Glossary Handbook Tokens, Pricing and Tokenization Vibe Coding addendum

PM Master Handbook by SkillNavora
23 modules and 620+ pages: 12 Product Core modules, an 11-class AI Tools Lab, frameworks, case studies, projects and interview prep.

Open the interactive handbook ↗ Complete handbook PDF (620+ pages, 23 MB) ↗

Product Core modules

1Fundamentals of Product Management*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
2Product Strategy and Strategic Thinking*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
3Problem Discovery and Market Research*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
4Problem Framing and Prioritization*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
5UX and Design Collaboration*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
6Effective Communication and Stakeholder Management*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
7Product Analytics and Metrics*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
8Product Launch, Growth, and Monetization*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
9Tech 101 and System Design for PMs*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
10LLMs, RAGs, and AI Foundations*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
11Vibe Coding for Product Managers*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides
12Agentic AI and AI Systems*Handbook chapter, detailed PDF, handbook and class slides*PDF Detailed PDF Handbook Class Slides

AI Tools Lab

13AI Tools Landscape and Your Starter StackPDF ↗ 14Claude for Daily PM WorkPDF ↗ 15Claude Code: Basic to AdvancedPDF ↗ 16Lovable: Basic to AdvancedPDF ↗ 17Replit and Other AI App BuildersPDF ↗ 18Figma and AI Design ToolsPDF ↗ 19Notion, PRD Tools and AI DocumentationPDF ↗ 20Automation with Zapier, Make and n8nPDF ↗ 21AI Chatbots for BusinessPDF ↗ 22Voice Agents for BusinessesPDF ↗ 23Learning Paths, Courses and CommunitiesPDF ↗

Business AI Stack (chapters 24 to 31)

24The business AI stack: tools and where to learn themPDF ↗ 25Next.js: customer-facing dashboardsPDF ↗ 26Zoho CRM: managing customers and leadsPDF ↗ 27WhatsApp Business Platform: customer messagingPDF ↗ 28Deployment with Vercel and AWSPDF ↗ 29Jira: tracking product and development workPDF ↗ 30Meta Business Suite, social media marketing and Google AdsPDF ↗ 31A one-page website that generates leadsPDF ↗

Complete AI tool guides (15)

1n8n: workflows, Gmail and an AI reply agentPDF ↗ 2Make: scenarios, routers and lead routingPDF ↗ 3Zapier: invoices from Gmail to Drive and a sheetPDF ↗ 4Webhooks: how apps talk in real timePDF ↗ 5Replit: one app for the web and a phonePDF ↗ 6Lovable: a job application trackerPDF ↗ 7Figma: components, auto layout and handoffPDF ↗ 8Claude Code: an AI coding agent in your projectPDF ↗ 9MCP: connecting AI assistants to your toolsPDF ↗ 10LLMs: how they work and how to call themPDF ↗ 11RAG: answers from your own documentsPDF ↗ 12AI agents: a job-search agent with approvalsPDF

↗ 13Agentic AI: planning, guardrails and teamsPDF ↗ 14Generative AI: uses, limits and a weekly reportPDF ↗ 15Other AI tools: choosing one per jobPDF ↗

Complete AI tool guides (one PDF, 30 pages) ↗

Libraries

★Framework Library*29 frameworks with worked examples and templates*PDF ↗ ★Case Study Library*12 product case studies*PDF ↗ ★Practical Projects and Portfolio*5 end-to-end portfolio projects*PDF ↗ ★Interview Handbook*100+ PM interview questions with model answers*PDF ↗ ★AI Tools Lab*Every AI tools class with step-by-step labs*PDF ↗

PM interview practice (free, on the original sites)

Hand-picked public question banks and guides from other PM communities. Each link opens the original page in a new tab.

PM frameworks cheat sheet

Framework	Use it to	How it works
RICE	Prioritise a backlog	Reach × Impact × Confidence ÷ Effort. Score each idea, then sanity-check the ranking.
Kano	Decide which features delight	Classify features as Must-be, Performance, Attractive, Indifferent or Reverse using a short survey.
MoSCoW	Scope a release	Must have, Should have, Could have, Won't have (this time).
Jobs-to-be-done	Understand motivation	When [situation], I want to [motivation], so I can [outcome].
AARRR (Pirate metrics)	Map the growth funnel	Acquisition, Activation, Retention, Referral, Revenue.
HEART	Measure user experience	Happiness, Engagement, Adoption, Retention, Task success, each with goals, signals and metrics.
North Star metric	Align the team	One metric that captures the value users get, supported by input metrics the team can move.
Five whys	Find the root cause	Ask 'why?' repeatedly until you reach a cause you can act on.
CIRCLES	Answer product design questions	Comprehend, Identify users, Report needs, Cut through prioritisation, List solutions, Evaluate trade-offs, Summarise.

PRD template

PRD: [Feature name] Owner: [Name] · Status: Draft · Last updated: [Date] 1. Problem Who has the problem, what it is, and the evidence (data, interviews, tickets). 2. Goal and success metrics Primary metric (with baseline and target), secondary metrics, guardrails. 3. Users and use cases Target segment, top 3 user stories: "As a [user], I want [action] so that [benefit]." 4. Solution What we'll build (link designs), what's out of scope, open questions. 5. AI behaviour (if applicable) Inputs, model, prompt/approach, what it must never do, human review, evaluation set and pass criteria, cost per request. 6. Launch plan Experiment design, rollout stages, support and comms, rollback criteria. 7. Risks and dependencies Technical, legal/privacy, operational; owners and mitigations.

Free courses