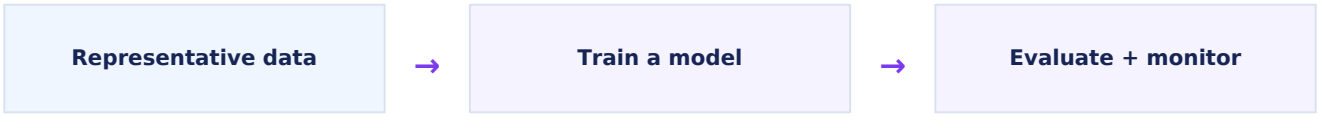
LEARNHUB / TECHNICAL COURSE LIBRARY

Artificial Intelligence and Machine Learning

Course introduction

Explore how machine-learning systems learn patterns from examples, evaluate performance on unseen data, and support practical decisions responsibly.



Level: Beginner · Category: AI

Original, beginner-friendly study content designed to accompany the course syllabus and its linked video lesson.

Contents

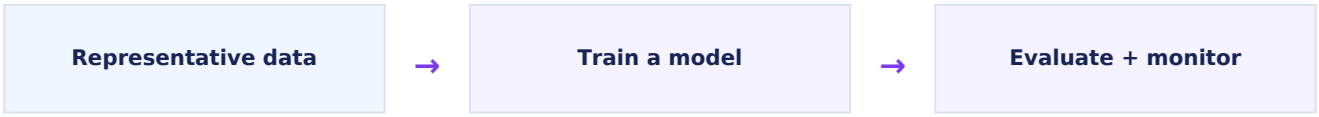
- How to use this resource
- Learning map
- Chapter 01 · AI, ML, and the learning workflow
- Chapter 02 · Supervised learning and data
- Chapter 03 · Training, validation, and overfitting
- Chapter 04 · Metrics and model evaluation
- Chapter 05 · Practical models and responsible use
- Summary and next steps

Course purpose: Explore how machine-learning systems learn patterns from examples, evaluate performance on unseen data, and support practical decisions responsibly.

How to use this resource

Read the chapter titles first, then study the explanations and examples. Pause at each practice prompt before reading its suggested answer. Write down questions that remain and revisit them during revision.

Learning map



Chapter 01 · AI, ML, and the learning workflow

The central idea

Artificial intelligence is a broad field of systems performing tasks associated with intelligent behavior. Machine learning is a set of methods that infer patterns from data to make predictions or decisions.

How to think about it

A practical workflow defines the task, gathers representative data, separates training and evaluation data, chooses a baseline, trains, evaluates, and documents limitations. A model's output is evidence, not automatically a correct decision.

First example

A support team estimates ticket priority from historical messages, then checks whether labels and examples represent current requests.

Try this

Define a prediction task for forecasting energy use and name one input feature, one target, and one risk.

A way to check your answer

> An input may be recent hourly temperature; the target may be next-hour electricity use. A risk is that training data excludes unusual weather and the model fails during a heat wave.

Chapter 02 · Supervised learning and data

The central idea

Supervised learning uses examples paired with labels. Classification predicts a category; regression predicts a numeric value. A feature is an input variable used by a model.

How to think about it

Split data before fitting preprocessing steps to avoid leakage. Check missing values, units, class balance, and the meaning of every label. Establish a simple baseline so complexity is justified by measured improvement.

First example

A model classifies loan applications into predefined risk bands using historical application features, subject to fairness and legal review.

Try this

Classify whether a product review is positive or negative. Identify a feature, label, and a possible source of leakage.

A way to check your answer

Text tokens or review length can be features; sentiment is the label. A post-review moderation outcome unavailable at prediction time would be leakage.

Chapter 03 · Training, validation, and overfitting

The central idea

Training adjusts model parameters to fit examples. Overfitting occurs when a model captures training-specific noise and performs poorly on new examples.

How to think about it

Use training data to fit, validation data to choose settings, and a final test set for an unbiased estimate. Regularization, simpler models, more representative data, and careful feature selection can reduce overfitting.

First example

A decision tree with a branch for every training record may memorize outcomes; limiting tree depth can improve generalization.

Try this

A model is 99 percent accurate on training data and 68 percent on held-out data. Name a likely issue and a useful next check.

A way to check your answer

> Overfitting is likely. Check the split and leakage, compare performance across simpler models, and verify the held-out data reflects the intended use population.

Chapter 04 · Metrics and model evaluation

The central idea

A metric summarizes a model's behavior for a specific objective. A confusion matrix counts true/false positives and negatives; precision and recall measure different error trade-offs.

How to think about it

Do not choose accuracy by habit. In an imbalanced screening task, a model can be accurate by predicting the majority class. Select metrics based on the cost of missed positives, false alarms, calibration, and downstream workflow.

First example

A fraud-screening system can use recall to measure the share of fraud cases caught while also tracking false alerts that burden investigators.

Try this

For a safety screening task where missed cases are costly, compare recall and precision and say what other measure matters.

A way to check your answer

> Recall captures the share of real positive cases found; precision captures how many flagged cases are truly positive. Review false-positive costs, subgroup performance, and calibration.

Chapter 05 · Practical models and responsible use

The central idea

A model is a parameterized function learned from data. Responsible deployment monitors distribution changes, access, explanations, privacy, and the consequences of automated decisions.

How to think about it

Start with interpretable baselines such as linear/logistic regression or a small tree. Compare them with more complex methods only when the gain is real. Document intended use, evaluation slices, known limits, and a human escalation path.

First example

A city transit team predicts passenger volume but lets planners review events and service changes before acting on the forecast.

Try this

Write two safeguards for using an image classifier in a workflow that affects people.

A way to check your answer

> Measure errors across representative groups and provide human review for high-impact cases. Obtain lawful data, minimize retention, and make uncertain predictions visible.

Summary and next steps

Artificial Intelligence and Machine Learning brings together 5 connected topics. Review each definition, explain one example without looking, and then try the practice prompt again.

Key ideas to carry forward

AI, ML, and the learning workflow: Artificial intelligence is a broad field of systems performing tasks associated with intelligent behavior. Machine learning is a set of methods that infer patterns from data to make predictions or decisions.

Supervised learning and data: Supervised learning uses examples paired with labels. Classification predicts a category; regression predicts a numeric value. A feature is an input variable used by a model.

Training, validation, and overfitting: Training adjusts model parameters to fit examples. Overfitting occurs when a model captures training-specific noise and performs poorly on new examples.

Metrics and model evaluation: A metric summarizes a model's behavior for a specific objective. A confusion matrix counts true/false positives and negatives; precision and recall measure different error trade-offs.

Practical models and responsible use: A model is a parameterized function learned from data. Responsible deployment monitors distribution changes, access, explanations, privacy, and the consequences of automated decisions.

A short review routine

1. Recall the main idea from each chapter before rereading. 2. Work one example by hand. 3. Explain a common mistake and how you would catch it. 4. Revisit a difficult topic after a short break.

Course outcomes

- Distinguish AI, machine learning, and common learning settings
- Prepare data and define features and labels
- Train a supervised model and evaluate it with suitable metrics
- Recognize overfitting, leakage, and responsible-use concerns

This course introduction is an original LearnHub learning aid. Use the referenced ideas for study, then consult the course video and the rest of the library for additional practice.