

Introduction to Machine Learning



- What this training covers and why it matters for your career and organization




- Who this is for: professionals who work with, manage, or decide around ML




- How machine learning is reshaping business and technology in 2026



- Training agenda and what you'll be able to do by the end



The State of ML in the Enterprise



- 78% of Global 2000 companies have AI in production, yet only 33% have scaled it enterprise-wide



- The ML Deployment Deficit: 88% adoption vs. 33% enterprise-scale deployment




- Enterprise AI spending reached \$247B globally, with a median 2.4x ROI



- The gap between pilot and production is the defining challenge for business leaders

What Is Machine Learning?

- Algorithms that learn patterns from data without explicit programming
- Distinct from traditional software, statistics, and broader AI
- Three paradigms: supervised, unsupervised, and reinforcement learning
- Powers recommendation engines, fraud detection, and predictive maintenance



The Machine Learning Workflow



- End-to-end steps: problem framing, data, modeling, evaluation, deployment




- The workflow is iterative; revisiting steps is a sign of rigor



- Roles: data engineer, data scientist, ML engineer, product manager



- Common pitfalls: wrong problem, data leakage, metric obsession, deployment neglect





Data: The Foundation of ML



Data quality and quantity matter more than algorithm choice




Features are inputs, labels are the target you predict



Always split data into training, validation, and test sets




Beware of sampling, historical, and labeling bias in data



Key Algorithm Families


- Regression, classification, clustering, and dimensionality reduction: each solves a distinct problem type
- Decision trees and ensembles: the workhorses for tabular business data
- Neural networks and deep learning: powerful but often overkill for structured data
- Choose by balancing interpretability, speed, accuracy, and data requirements

DECISION TREE




BEST FOR
Interpretable rules
on tabular data

ENSEMBLE OF TREES



BEST FOR
High accuracy on
tabular business data

NEURAL NETWORK (DEEP LEARNING)



BEST FOR
Complex patterns in
large or unstructured data



Interpretability | Speed | Accuracy | Data Requirements

Model Evaluation in Plain Terms



- Accuracy, precision, recall, F1: what they mean and when each matters



- Overfitting vs. underfitting: causes, symptoms, and detection



- Use the confusion matrix to communicate with non-technical stakeholders



- Align technical metrics with business outcomes the organization cares about



BUSINESS
STAKEHOLDER



MODEL METRICS

ACCURACY
92.4%

PRECISION
91.1%

RECALL
89.3%

F1 SCORE
90.2%

DATA QUALITY

96%
GOOD



TECHNICAL
PRACTITIONER



Clear metrics. Shared understanding. Better decisions. **Better outcomes.**

From Model to Business Impact



- Choose deployment: batch, real-time, or embedded inference




- MLOps monitors drift; production is the real test of model quality



- Governance meets 2026 regulations: EU AI Act, ISO 42001




- Link model outputs directly to business KPIs and ROI



The MLOps Control Plane

- MLOps is a control plane for traceability, accountability, and governance
- Core components: model registry, feature store, experiment tracking, pipelines
- Data lineage is a runtime requirement for auditors and regulators
- Separate risk assessment from model performance to prevent unsafe deployments



Production Monitoring and Observability



- Monitor four layers: infrastructure, data quality, model performance, and business impact




- Track feature validity, decision volatility, policy violations, and human feedback



- Alert only on actionable conditions to prevent alert fatigue



- Retraining is slow; design for fast rollback in production systems



ALERT FATIGUE
Too many noisy alerts
= missed critical issues



Design for fast rollback to minimize impact

FEATURE VALIDITY



DECISION VOLATILITY



POLICY VIOLATIONS



HUMAN FEEDBACK



Working with ML Teams



- Typical roles: data engineer, scientist, ML engineer, product manager, risk approver




- Ask good questions and set realistic expectations as a business stakeholder



- Agile adapts to ML's experimental nature with iterative, hypothesis-driven cycles



- Decide when to build custom models versus use existing solutions or AutoML



Model Accountability and Governance



- Every production model needs an owner, technical steward, and risk approver



- Build governance in: model cards, approvals, and audit trails from the start



- EU AI Act, ISO/IEC 42001 require provable traceability and explainability



- Fast-track low-risk models; multi-stage sign-off for high-risk deployments



ALERT
High-risk model detected
Additional review required

MODEL CARD

Model Name
Credit Risk Scoring Model v2.1

Purpose
Assess credit risk for loan applications

Model Type
Supervised Machine Learning

Training Data
Jan 2021 – Dec 2023
Internal transactional data

Performance
AUC-ROC: 0.87
Accuracy: 0.81

Limitations
May underperform for
new customer segments

Ethical Considerations
Fairness monitoring in place
No sensitive attributes used

Monitoring & Maintenance
Performance monitored monthly
Retraining cadence: Quarterly

Owner
Jane Smith

Technical Steward
Alex Kim

Risk Approver
Michael Brown



OWNER
Jane Smith
Approved



TECHNICAL STEWARD
Alex Kim
Approved



RISK APPROVER
Michael Brown
Approved



HIGH RISK

DEPLOYMENT RISK LEVEL
HIGH RISK



Agentic AI and LLMs in the Enterprise



- **23%** of organizations are scaling agentic AI, with **39%** experimenting



- LLMs need new ops: prompt versioning, RAG monitoring, hallucination detection



- Manage token usage, inference cost, and runaway expenses proactively



- **"Agent washing"** is widespread; verify genuine agentic capabilities, not marketing

GENUINE AGENTIC AI

AUTONOMOUS. ADAPTIVE. OUTCOME-DRIVEN.



INFERENCE COST



\$7,842
TODAY




TICKING UP RAPIDLY

VS

AGENT WASHING

MARKETING FIRST. SUBSTANCE LAST.



LIMITED CONTROL. LIMITED VALUE.

- Rigid workflows
- Minimal reasoning
- Human-dependent
- High risk of disappointment

Identifying High-Potential ML Opportunities



- Start with business problems, not technology; look for repeatable, high-volume decisions



- Assess data readiness: is the right data available, accessible, and of sufficient quality?



- Build stakeholder buy-in by communicating value, managing expectations, and addressing resistance



- Redesign workflows around AI, don't just bolt it onto existing processes



Building Data Literacy and AI Fluency

78%

- 78% of senior leaders expect AI fluency to become more important for employee roles




- Understand data types, distributions, and basic statistics



- Foster a data-driven culture: ask the right questions and interpret results critically

71%

- 71% of organizations report progress toward an integrated AI-human workforce



Your Next Steps



- ML is a team sport; data is the foundation; production is the real test



- Curated resources: courses, books, and communities for continued learning



- Action plan: identify one ML opportunity, assess data readiness, start a conversation



- Durable skills: validation discipline, data hygiene, and clear communication



PRODUCTION



DATA FOUNDATION



TEAM SPORT

PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/ebd765ef/public