Pattern Recognition and Machine Learning: An Overview

- Roadmap: classical methods (k-NN, SVM, trees) to foundation models
- 2026 impact: vision, NLP, diagnostics, fraud detection, quality control
- ML as a subset of AI: supervised, unsupervised, semi-supervised, reinforcement
- Pipeline: collection → features → training → evaluation → deployment → monitoring
- 2026 toolkit: Python, NumPy, scikit-learn, PyTorch, TensorFlow, JAX

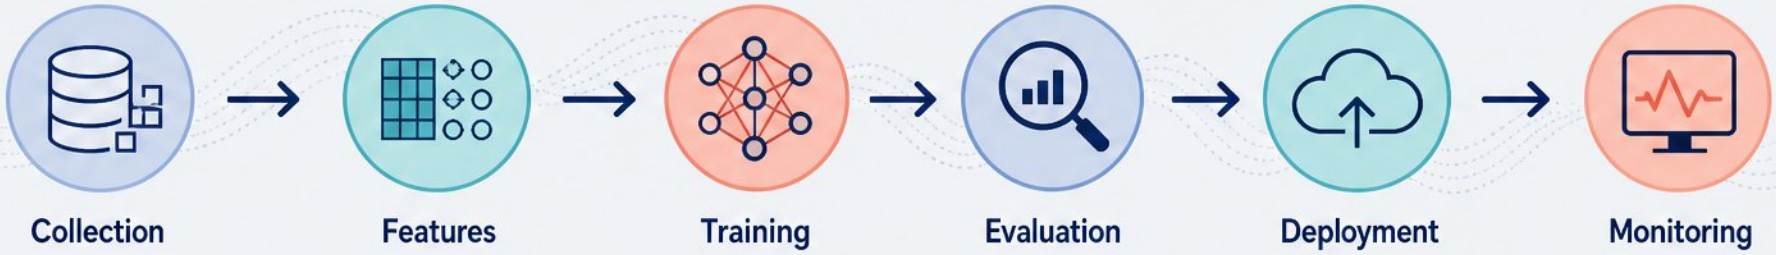

Mathematical and Historical Foundations

- **Probability:** random variables, distributions, Bayes' theorem — bedrock of classification
- **Linear algebra:** vectors, matrices, dot products, eigenvalues — how data flows through models
- **Optimization:** convex vs. non-convex loss surfaces; gradient descent, SGD, Adam
- The curse of dimensionality; PCA and t-SNE remain critical tools
- **Milestones:** perceptron (1958), backprop (1986), SVM (1995), deep learning (2012+), transformers (2017)

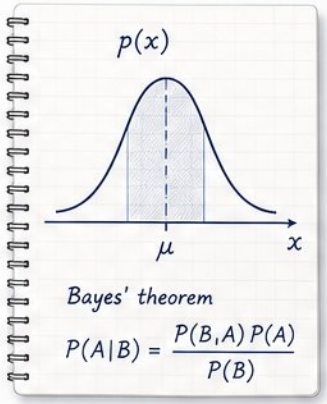

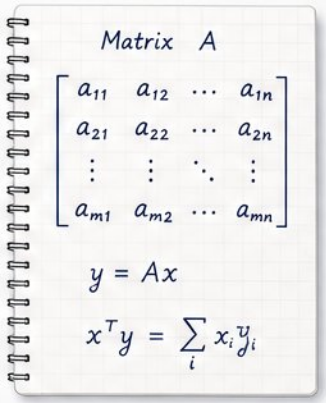

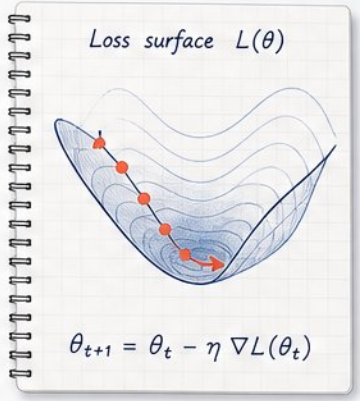

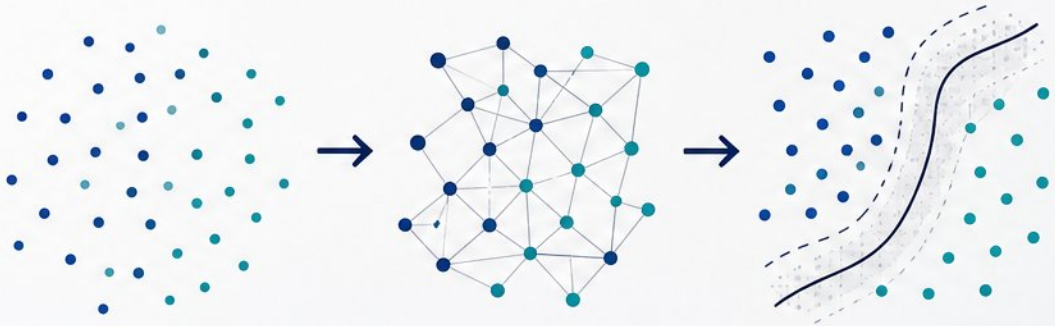

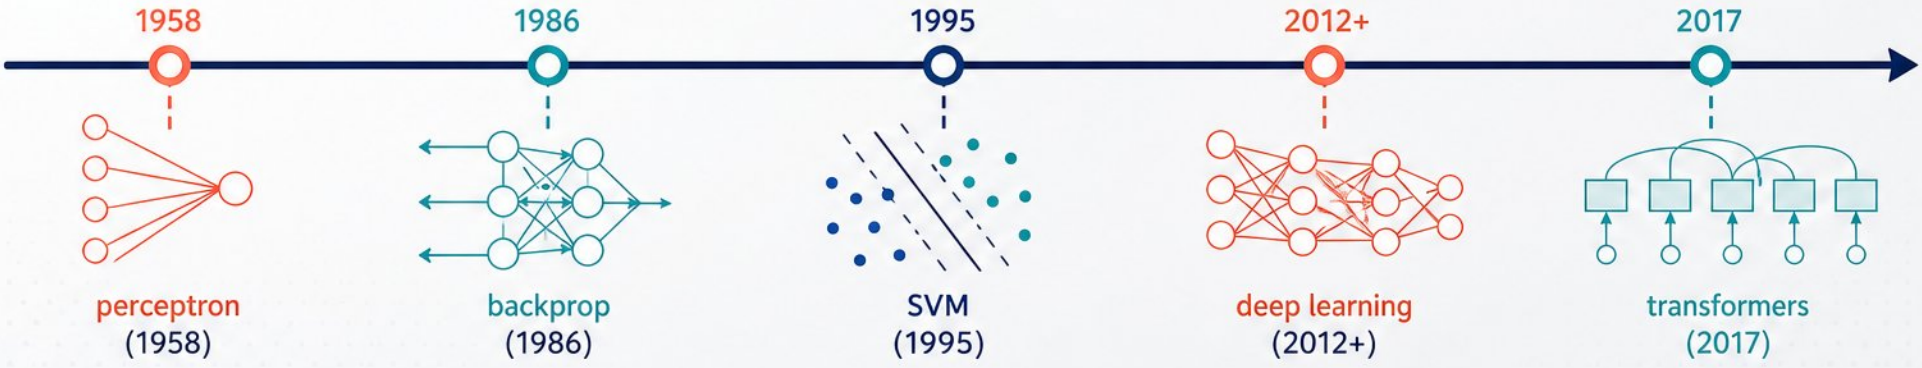

Core Learning Paradigms and Concepts

Supervised Learning

Classification and regression;
losses like cross-entropy and MSE

Unsupervised Learning

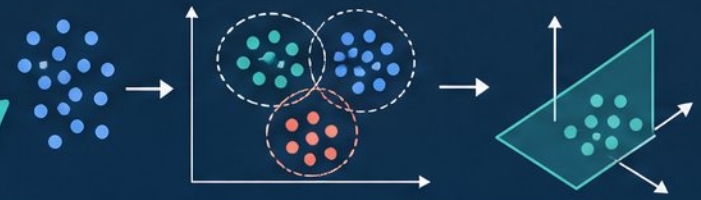

k-means clustering, PCA reduction,
anomaly detection

Semi-Supervised Learning

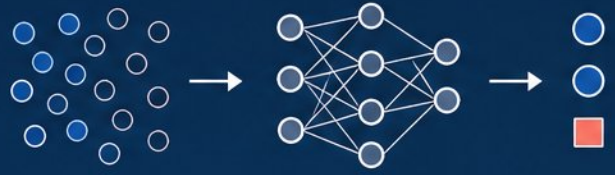

Leverages both labeled and
unlabeled data to improve learning

Reinforcement Learning

Agents, environments, rewards,
and policy learning

- Feature extraction vs. learning: handcrafted features versus learned representations
- Supervised: classification and regression; losses like cross-entropy and MSE
- Unsupervised: k-means clustering, PCA reduction, anomaly detection
- Reinforcement: agents, environments, rewards, and policy learning
- Splits, overfitting, underfitting, and the bias-variance tradeoff

Classical Pattern Recognition Algorithms

Linear Classifiers

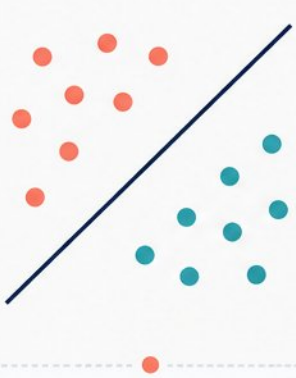

- Linear classifiers: logistic regression and perceptron — simple, interpretable, still competitive on linearly separable data

SVMs

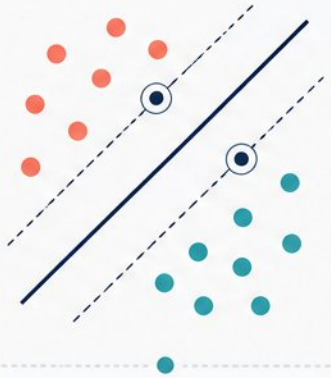

- SVMs: maximum margin plus kernel trick for non-linear boundaries; scikit-learn ready

Trees and Forests

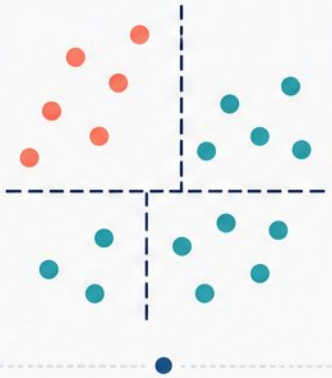

- Trees and forests: axis-aligned splits; bagging and boosting add ensemble robustness

k-NN

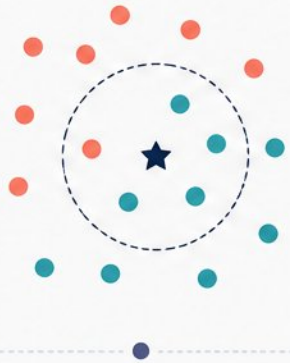

- k-NN and Naive Bayes: lazy learning and probabilistic classification for small-to-medium datasets

Naive Bayes

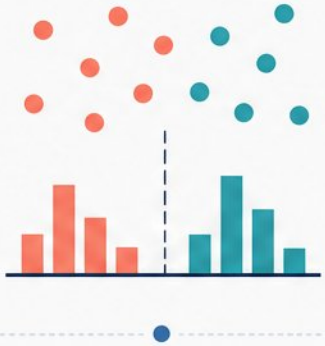

- Classical algorithms still win on small tabular data, interpretability, and low-latency edge inference

Neural Networks and Deep Learning Foundations

- From perceptron to MLPs; ReLU, GELU, and SiLU activation trade-offs
- Backpropagation intuition: gradients flow backward to tune weights
- Regularization: dropout, weight decay, batch norm, early stopping
- CNNs for images: convolution, pooling; ResNet and ConvNeXt
- RNNs and LSTMs gave way to transformers and self-attention
- 2026 frontier: looped transformers, latent-space models, MoE routing

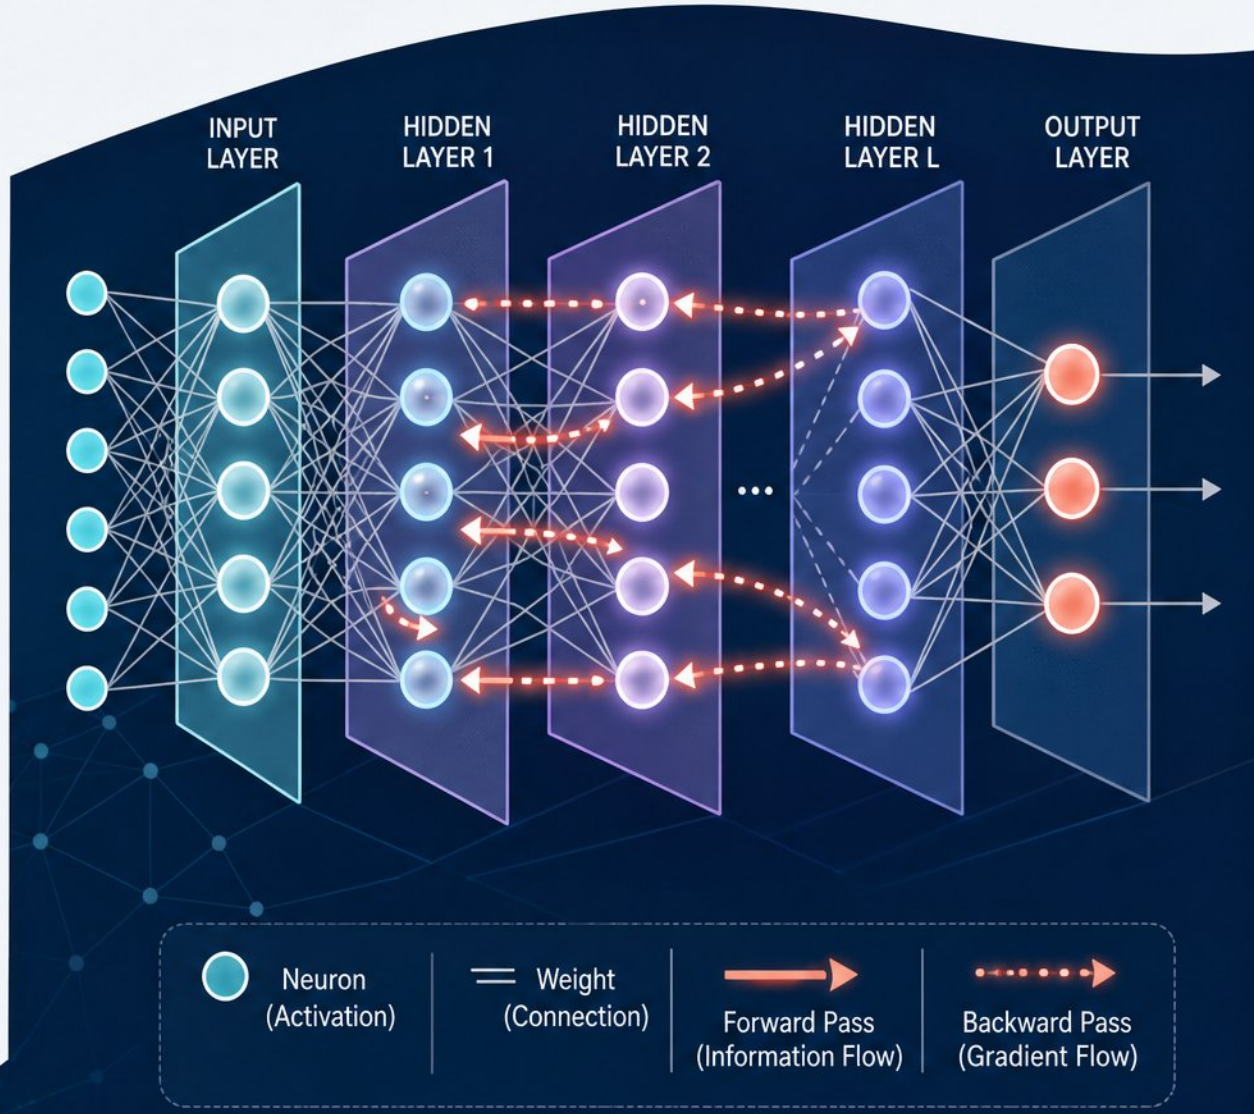

Model Evaluation and Validation

- Accuracy misleads on imbalanced data; use precision, recall, F1, ROC-AUC
- Read the confusion matrix to separate false positives from false negatives
- Validate with k-fold, stratified k-fold, and nested cross-validation
- Tune hyperparameters via grid search, random search, or Bayesian optimization (Optuna)
- Prevent data leakage; track runs with MLflow and Weights & Biases

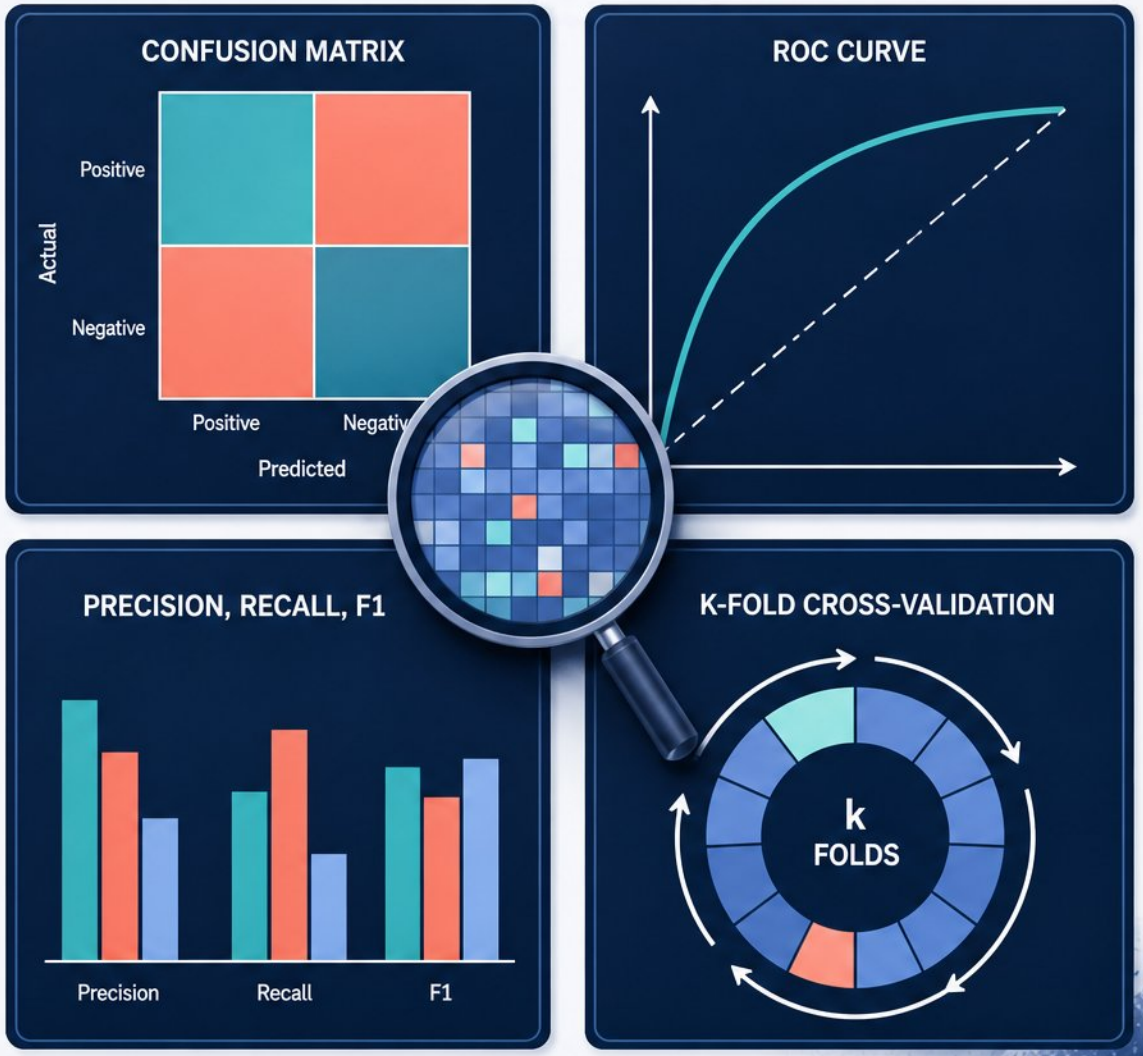

Practical Workflow: Data Preparation and Feature Engineering

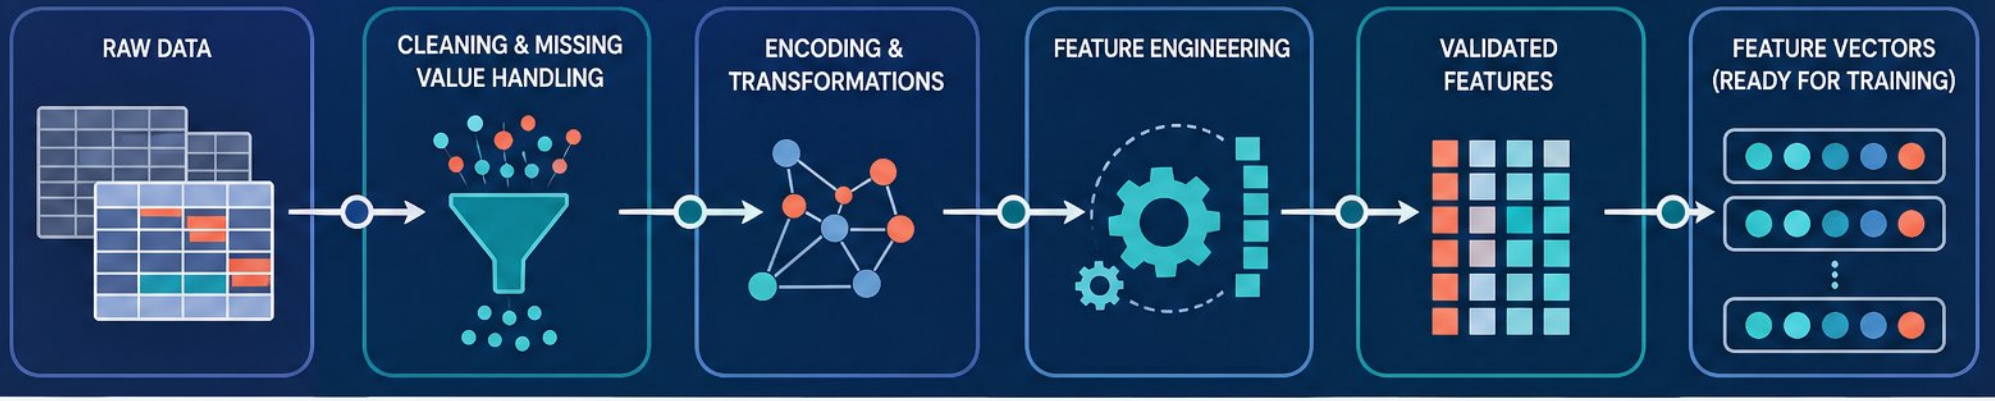

- Clean, encode, and handle missing values before anything else matters
- Great Expectations validates schema, null rates, and feature drift
- Feature stores (Feast, Tecton) keep training and serving consistent
- Version data with DVC so every pipeline run is fully reproducible
- Data validation gates prevent silent failures before training begins

DATA VERSIONING WITH DVC

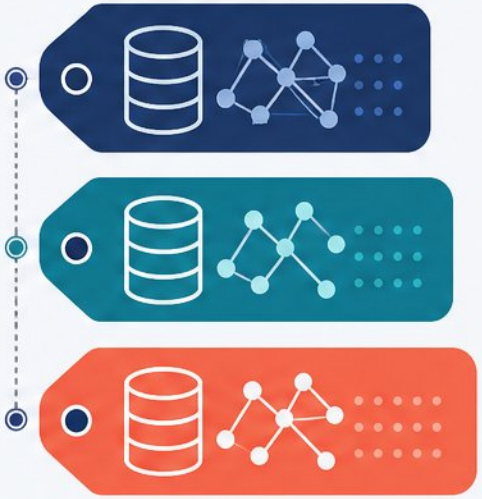

Practical Workflow: Training, Tuning, and Deployment

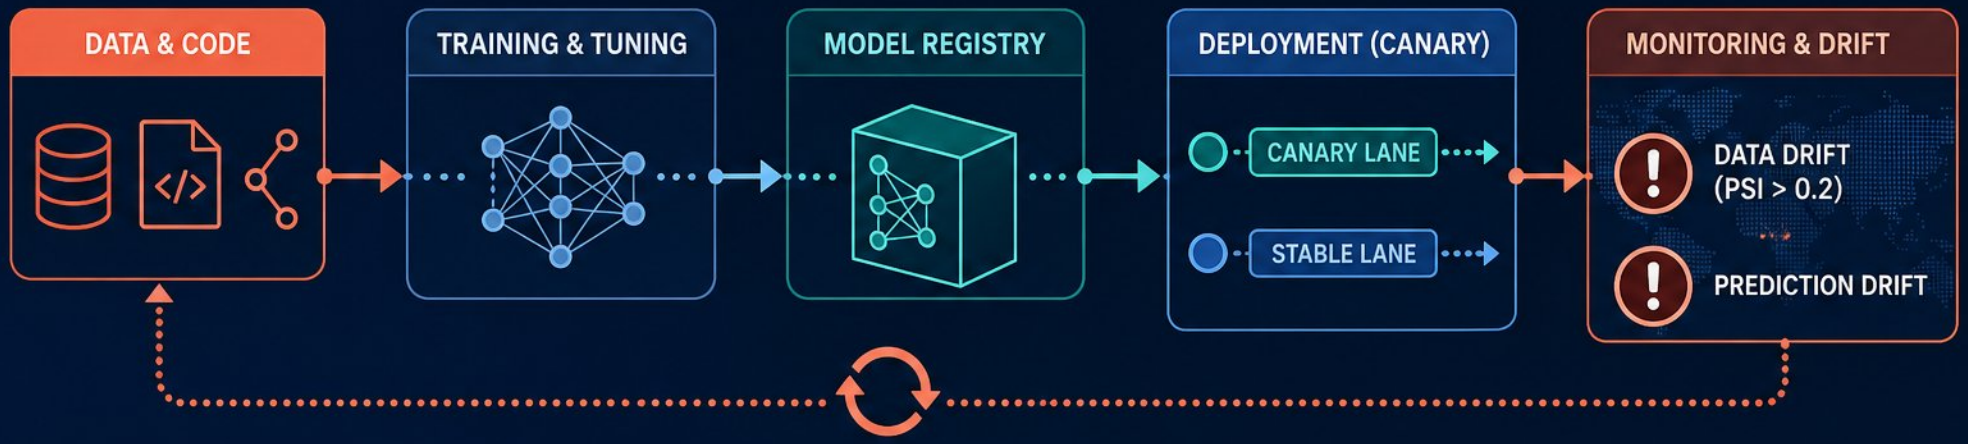

- Version everything together: data, code, environments, and hyperparameters
- Orchestrate with Airflow, Kubeflow, or Prefect; register models in MLflow
- Serve with FastAPI, KServe, TorchServe, or vLLM; quantize to FP8/FP4
- Monitor data drift (PSI > 0.2) and prediction drift in production
- Automate retraining with canary and blue-green rollback strategies

Ethical AI, Fairness, and Explainability

- Bias sources: historical data, sampling, labeling, proxy variables, and feedback loops
- Fairness metrics conflict: demographic parity, equalized odds, disparate impact cannot all hold
- Toolkits: IBM AI Fairness 360, Fairlearn, vfairness, EthicLens, relationship-aware counterfactual testing
- Explainability: SHAP, LIME, and model cards for EU AI Act and GDPR compliance
- Model risk management runs parallel independent fairness gates before deployment approval

Case Studies and Hands-On Examples

- **Image classification:** CIFAR-10/ImageNet, ResNet transfer learning, ConvNeXt backbones

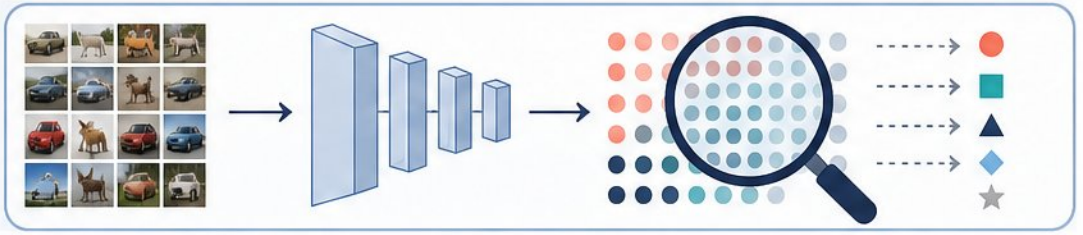

- **Sentiment analysis:** bag-of-words → BERT/GPT, fine-tuning with Hugging Face

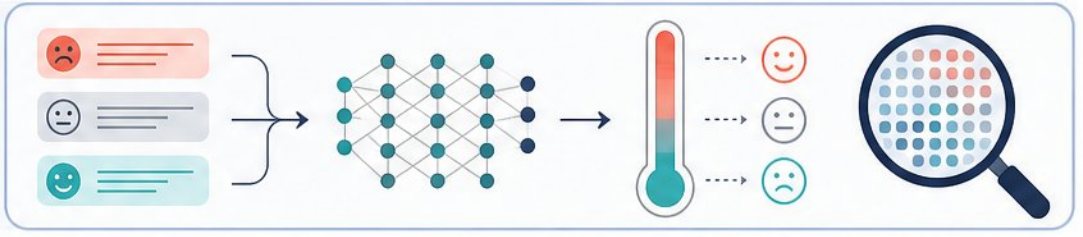

- **Customer segmentation with k-means:** choose k, interpret clusters, apply to business

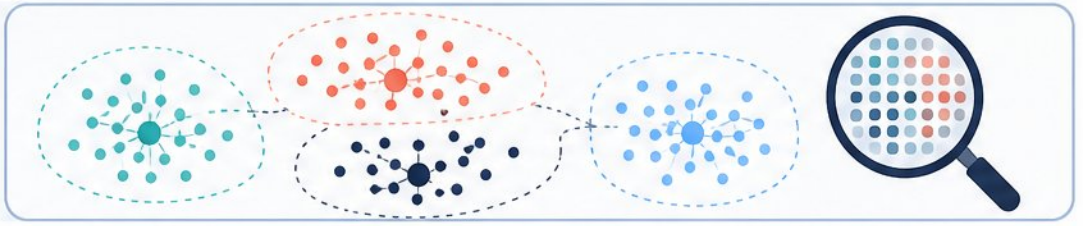

- **Anomaly detection:** fraud detection, predictive maintenance, healthcare monitoring

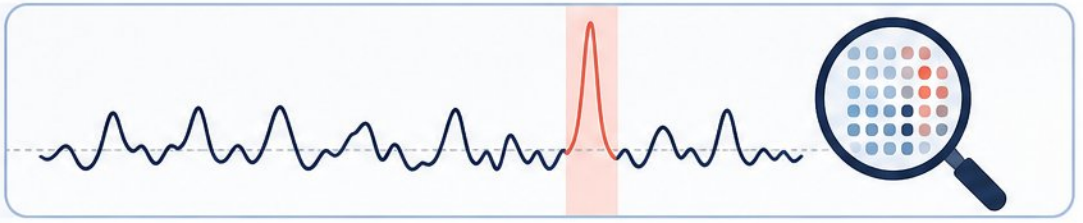

- **End-to-end project:** scikit-learn classifier, proper evaluation, FastAPI deployment

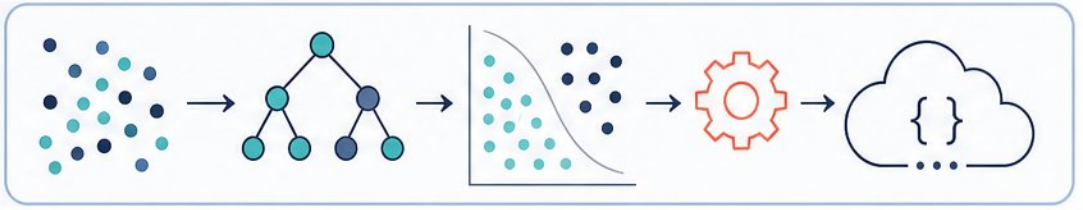

Current Trends and Advanced Topics

- Foundation models: GPT-6 Astra, Kimi K3, Llama 4, Gemini 2.5; open-weight vs closed frontier
- 2026 architectures: latent-space NCP models, looped/recursive transformers, MoE, sparse attention (MoVA, KDA)
- Self-supervised and contrastive pretraining without labels powers vision transformers
- Generative models: diffusion (latent vs pixel-space, accelerated sampling), GANs, VAEs, flow matching
- AutoML and neural architecture search; graph neural networks and geometric deep learning

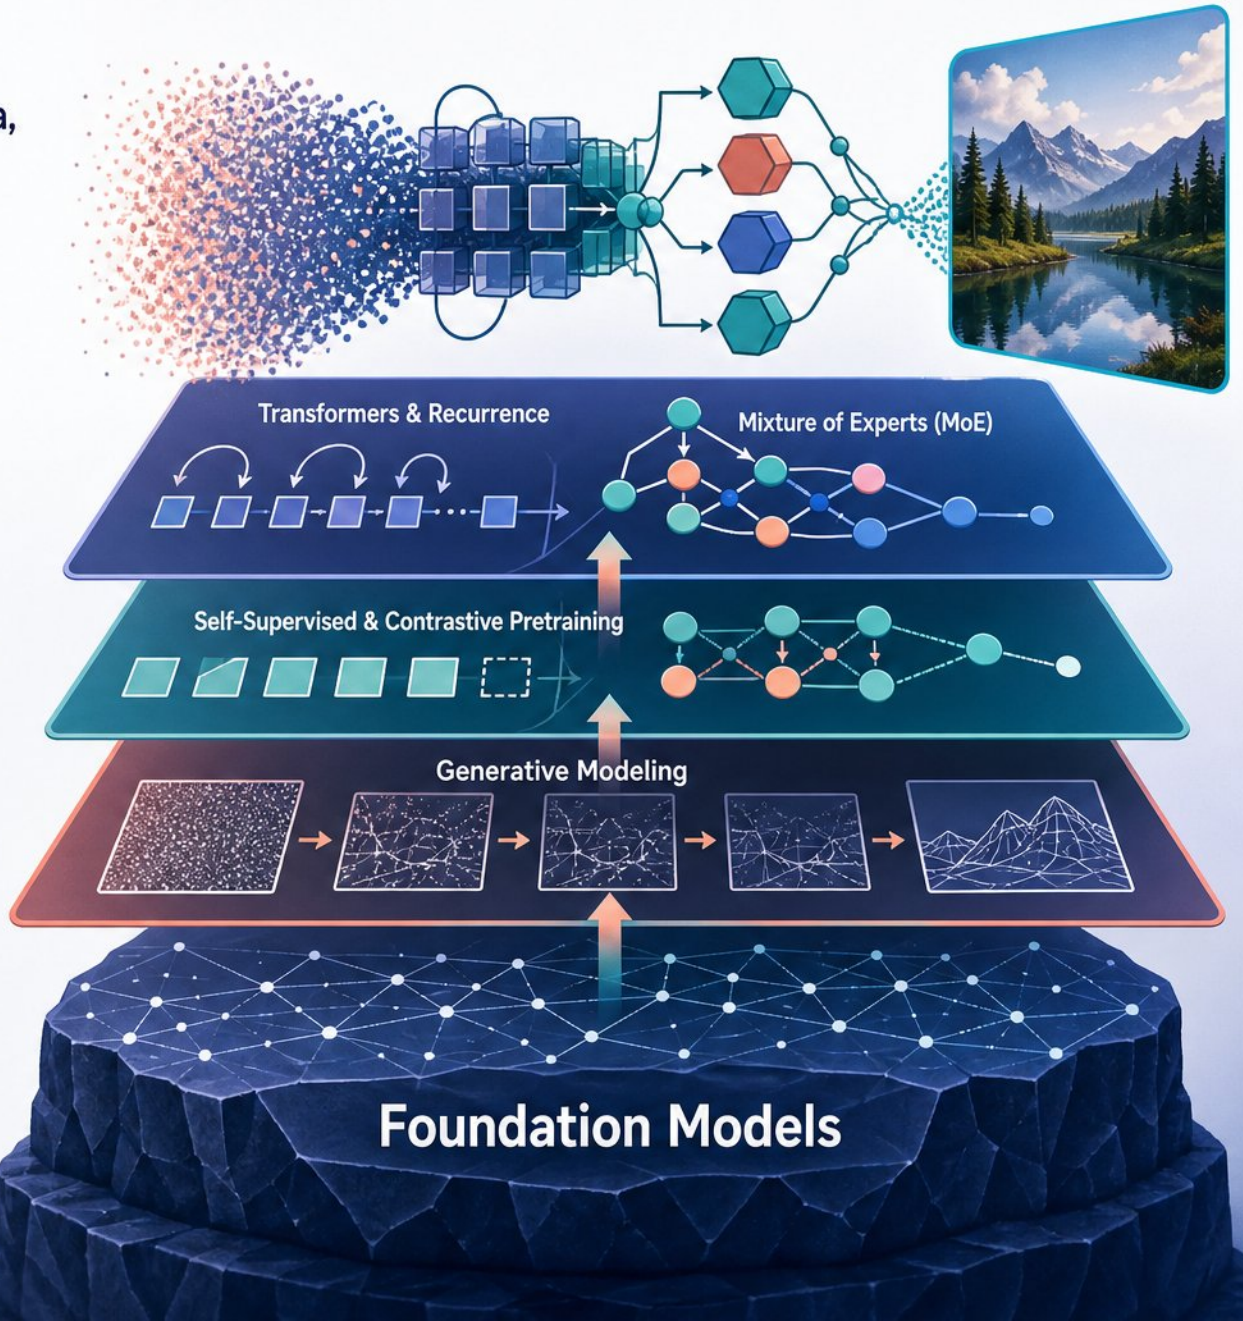

MLOps and Production Systems in 2026

- MLOps as a control plane: version data, code, environments, hyperparameters together
- Feature stores and data lineage: immutable snapshots, traceability, audit trails
- Monitoring beyond drift: feature validity, decision volatility, policy violations, human feedback
- LLMOps extensions: prompt versioning, RAG monitoring, hallucination scoring, token cost
- Governance: model cards, approval workflows, automated bias audits, EU AI Act, ISO/IEC 42001

Learning Resources, Tools, and Career Paths

- **Core libraries:**
  scikit-learn, PyTorch, TensorFlow/Keras, Hugging Face, vLLM, LangGraph

- **Structured paths:**
  Andrew Ng's ML Specialization, fast.ai, Karpathy's Zero to Hero, MIT 6.S191, Stanford CS336

- **Key books:**
  Hands-On ML (Géron), Dive into Deep Learning, Designing ML Systems, Understanding Deep Learning

- **Portfolio:**
  Kaggle competitions, open-source contributions, deployed end-to-end projects

- **Stay current:**
  NeurIPS 2026 (Sydney, Dec 6–12), ICML 2026 (Seoul, Jul 6–11), EMNLP 2026 (Budapest, Oct 24–29), CIARP 2026 (Mexico City, Nov 24–27), arXiv, HF Daily Papers

Next Steps: Building Your Capstone and Continuing to Grow

- **Capstone:** pick a real Kaggle or UCI dataset, define a business metric, build, evaluate, deploy
- **Portfolio:** public GitHub repo plus blog post or model card; present at a meetup
- **Specialize:** computer vision, NLP/LLMs, MLOps, responsible AI, or research
- **Weekly paper reading**; follow NeurIPS, ICML, CVPR, ICLR, ACL; contribute to open source
- **Community:** Kaggle, Hugging Face, local AI meetups, IAPR/ACM/IEEE, and mentorship

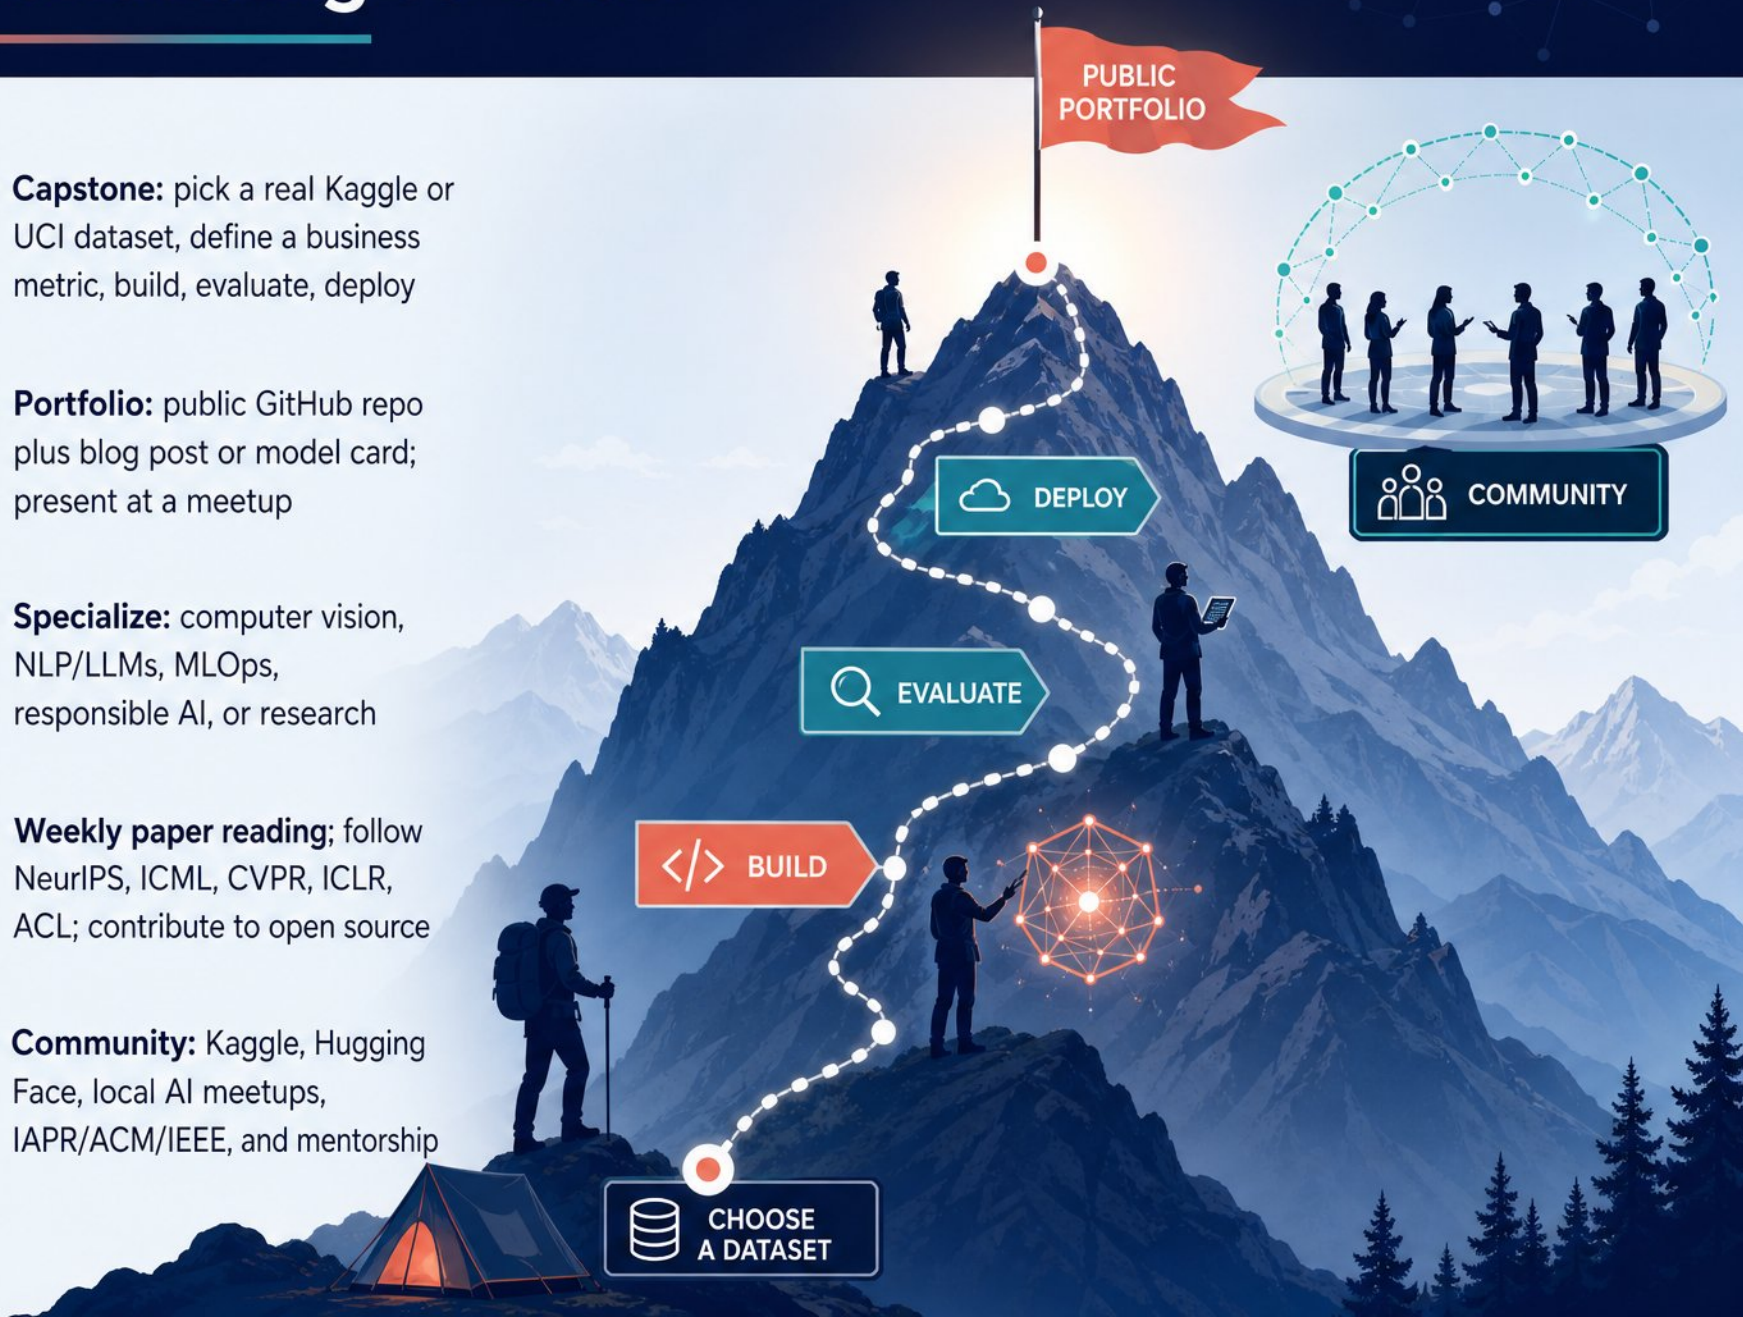

PersonWise.

PersonWise • AI digital-human course platform

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

personwise.ai/t/bb2d3ccd/public