AWS Certified AI Practitioner Cheat Sheet

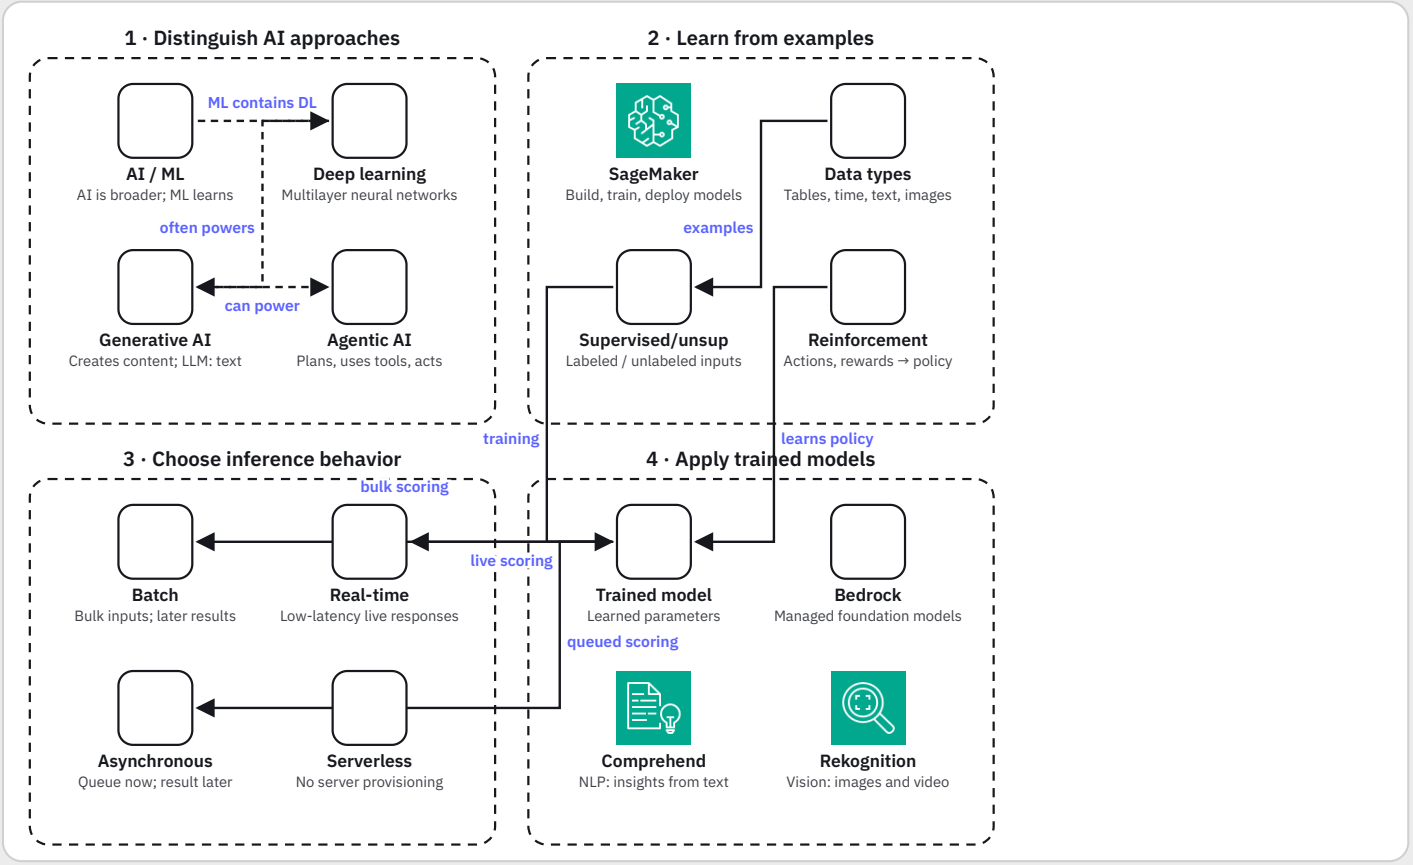
AIF-C01 · Domain 1 of 5

BlowStack

DOMAIN 1 · FUNDAMENTALS OF AI AND ML · 20% OF THE EXAM · TASK STATEMENT 1.1

Explain basic AI concepts and terminologies

Distinguish, learn, infer, apply



1 Distinguish AI approaches

Agentic AI coordinates steps, tools, and actions; GenAI creates content. Deep learning is ML using multilayer neural networks with learned weights. An LLM is a large language model, not a synonym for all AI.

2 Learn from examples

Supervised learning predicts labeled targets; unsupervised learning finds patterns without labels. Tables are structured; text/images are unstructured; time-series is time-ordered. Labels supply target values.

3 Choose inference behavior

Serverless describes infrastructure management, not response timing. Real-time serves live requests; asynchronous queues requests; batch handles datasets. SageMaker Serverless Inference can have cold starts.

4 Apply trained models

Inference applies a trained model to new inputs; training learns parameters. An algorithm is a learning procedure. Bedrock provides foundation models without customers operating the underlying model-serving infrastructure.

Exam traps

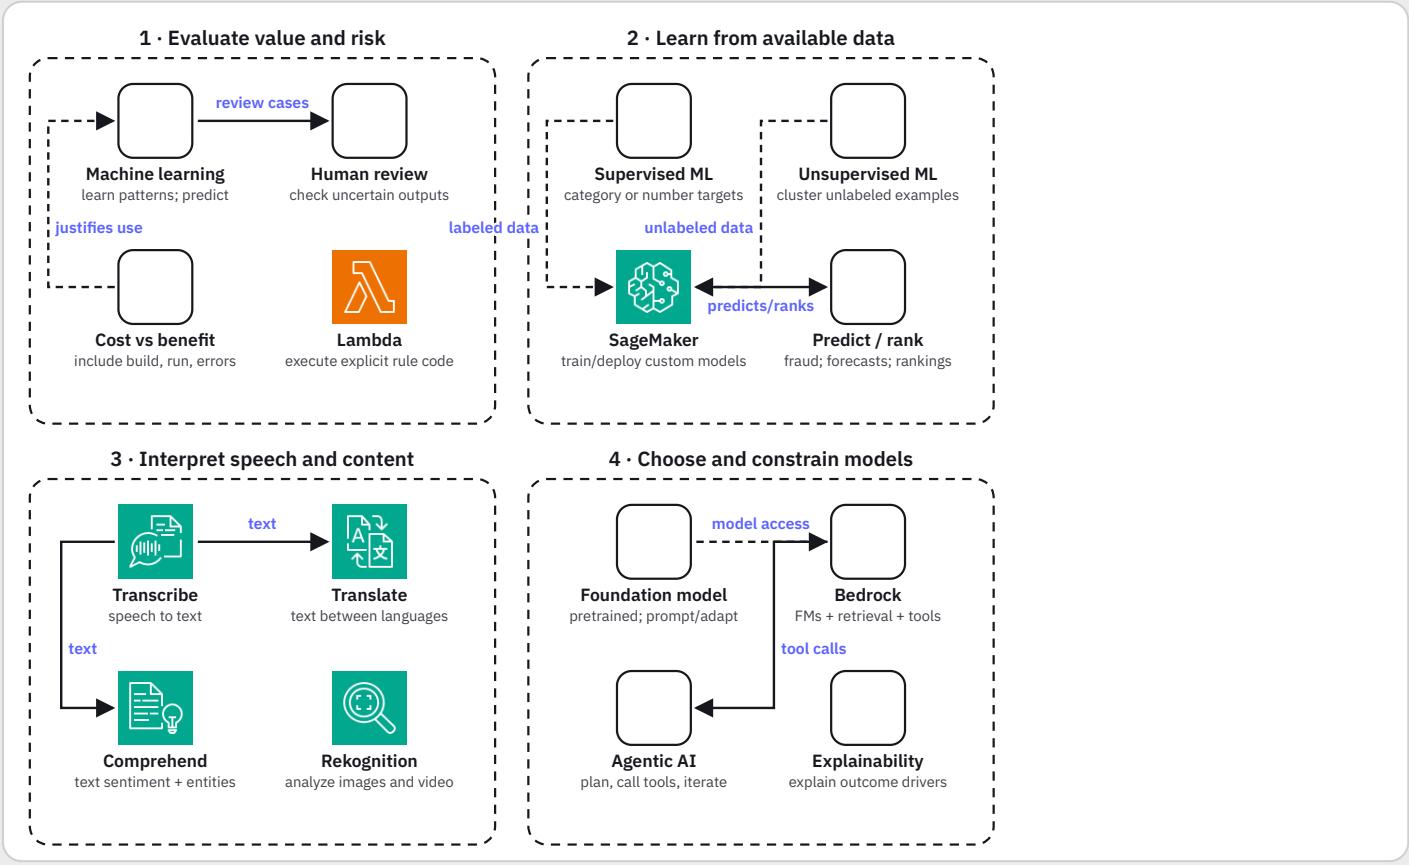
- **Bias is not fairness**: data, labels, design, or deployment can create unequal outcomes. Fairness requires context-specific evaluation; removing sensitive fields alone does not guarantee it.
- **Overfitting** means strong training performance but poor generalization. Underfitting misses patterns even in training data. Training accuracy alone cannot establish a good fit.
- **Agentic AI is not reinforcement learning**: agents can use LLMs and tools without reward-based training. Reinforcement learning learns a policy from action-and-reward feedback.

Get practice exams and lessons for AWS Certified AI Practitioner

DOMAIN 1 · FUNDAMENTALS OF AI AND ML · TASK STATEMENT 1.2

Identify practical use cases for AI

Evaluate, learn, interpret, act



1 Evaluate value and risk

ML adds value through prediction, scalable automation, and assisted decisions. **Exact specifications** favor explicit rules over predictions. Benefits must justify data, operating, oversight, and error costs.

2 Learn from available data

Known targets favor supervised learning: classification predicts categories; regression predicts numbers. Clustering discovers unlabeled groups. Amazon SageMaker supports custom fraud, forecasting, and recommendation models.

3 Interpret speech and content

Choose a **managed task API** for a supported capability without training a custom model. Transcribe recognizes speech; Translate changes text language; Comprehend extracts text insights; Rekognition analyzes visual content.

4 Choose and constrain models

Use **foundation models** for open-ended content and knowledge-base retrieval. Task-specific ML may better meet explainability, regulatory, cost, or latency constraints. Agents use tools; consequential actions need oversight.

Exam traps

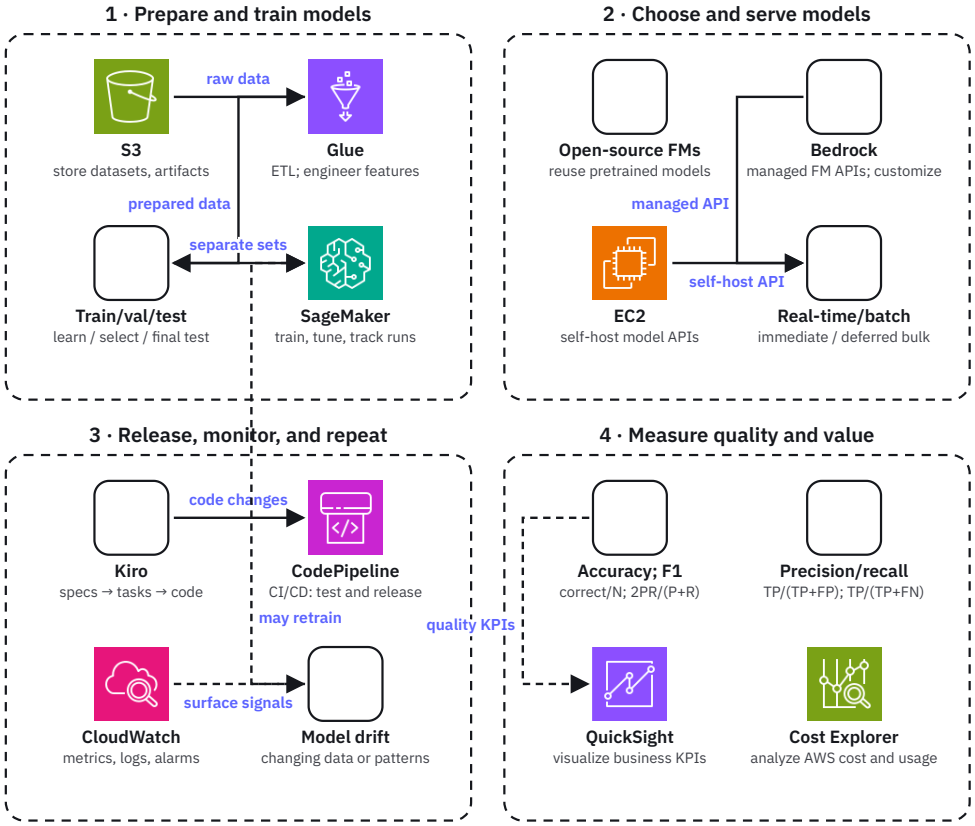
- **Deterministic does not mean correct**: a model can repeat the same wrong prediction. Exact compliance with a specification favors explicit rules, not merely repeatable outputs.
- **Classification needs known categories** during supervised training; clustering discovers groups without target labels. Fraud detection can use supervised labels or unlabeled anomaly detection.
- **Retrieval is not retraining**: a knowledge base supplies context to a foundation model without changing its weights. Grounding can improve relevance but does not guarantee factual answers.

Get practice exams and lessons for AWS Certified AI Practitioner

DOMAIN 1 · FUNDAMENTALS OF AI AND ML · TASK STATEMENT 1.3

Describe the AI/ML development lifecycle

Prepare, adapt, deploy, improve



1 Prepare and train models

Use **validation data** to select hyperparameters; reserve test data for final evaluation. Fit feature transformations on training data only. Experiment tracking links results to data, code, settings, and environment.

2 Choose and serve models

Choose **managed APIs** with Amazon Bedrock to avoid serving infrastructure; self-host on Amazon EC2 for control. Real-time inference serves interactive requests; batch inference prioritizes bulk throughput over immediacy.

3 Release, monitor, and repeat

Use **MLOps** to version data/code/models, automate repeatable releases, test scalability and production readiness, and limit technical debt. Monitor model quality as well as operations; drift requires investigation.

4 Measure quality and value

Favor **precision** when false positives are costly; recall when misses are costly. Accuracy can hide class imbalance; F1 balances precision and recall. Business metrics include cost/user, development costs, feedback, and ROI.

Exam traps

- **Pre-training** builds general capabilities from large datasets; fine-tuning adapts an existing model. Open-source availability does not guarantee unrestricted licensing or available training data.
- **Drift does not automatically trigger retraining.** Monitoring detects signals; teams investigate quality and causes, then configure retraining, evaluation, and release controls.
- **ROI** = (benefit – cost)/cost; high accuracy alone does not prove business value. AWS Cost Explorer reports AWS spending, not every development cost.

Get practice exams and lessons for AWS Certified AI Practitioner

DOMAIN 1 · FUNDAMENTALS OF AI AND ML · WHAT TO REMEMBER

Concepts and services of Fundamentals of AI and ML

Concepts to know

Machine learning

ML learns patterns from data; deep learning is ML using multilayer neural networks. **Training learns model parameters; inference applies them**. Choose ML for predictions from patterns, not when explicit rules must guarantee a specified outcome.

Supervised Learning

Supervised learning uses labeled examples: **classification predicts categories; regression predicts numbers**. Unsupervised learning finds structure without target labels; reinforcement learning improves action choices through reward feedback.

Classification Metrics

Accuracy measures overall correctness; precision measures correctness among predicted positives; recall measures positives found. **Prioritize recall for costly misses, precision for costly false alarms**. F1 balances precision and recall; imbalance can hide behind high accuracy.

Core services



Amazon SageMaker

Choose **custom ML development and lifecycle management**: prepare data, train and tune models, deploy batch or live inference, and monitor models. It offers more model control than a task-specific AI API.

Amazon Bedrock

Choose **managed foundation model access** for generation and content analysis without operating model-serving infrastructure. Retrieval, customization, and tool use support knowledge bases and agents.



Amazon Transcribe

Choose **speech-to-text** for recordings, captions, and call transcripts. Speaker labels and timestamps enrich transcripts; translating text and building conversational interfaces are different tasks.

Get practice exams and lessons for AWS Certified AI Practitioner