ARTIFICIAL INTELLIGENCE FRAMEWORKS · FIELD GUIDE

# Go from reading about AI to being certified in it.

**28 PAGES**

The full syllabus, the exam blueprint and the complete Learn-by-Doing list — the field guide that takes you from a working understanding of AI to the Certified AI Frameworks Professional credential.

- **01** Types-of-AI field guide
- **02** Framework infographics
- **03** Full official syllabus
- **04** Exam blueprint & weightings
- **05** Complete Learn-by-Doing list
- **06** Governance & responsible-AI maps

Issued by the Global Skill Development Council — companion to the Certified AI Frameworks Professional (CAIF) program.

ORIENTATION

# How to use this guide

This is a study companion, not a brochure. It walks the same path as the certification: understand the landscape, learn the frameworks, see how they are examined, then practise. Keep it beside you as you prepare.

## What is inside



## How to read each framework entry

**WHAT** A plain-language definition.

**WHY** The problem it solves for an organisation.

**PARTS** Its main components or stages.

**WATCH** Where teams most often get it wrong.

SECTION A

# Types-of-AI Field Guide

Before frameworks come categories. Knowing how AI is classified — by capability, by technique, and by autonomy — is the vocabulary the rest of the certification builds on.

### By capability 1

Narrow AI (one task), General AI (human-level breadth), Super AI (beyond human) — today everything in production is Narrow.

### By function 2

Reactive machines, limited-memory systems, theory-of-mind, and self-aware AI — a ladder of increasing context.

### By technique 3

Machine learning, deep learning, NLP, computer vision, and generative AI — the methods under the hood.

### By autonomy 4

Assisted, augmented, and autonomous systems — how much a human stays in the loop.

**Why it matters.** Most governance and risk decisions hinge on which category a system falls into. Misclassifying a system is the first step toward mis-governing it.

TYPES OF AI

# AI by Capability

**NARROW** Performs a single, well-defined task — recommendation, translation, detection. Every system in real use today is narrow.

**GENERAL** Hypothetical AI matching human flexibility across any task. An active research goal, not a product.

**SUPER** A speculative stage where AI exceeds human ability across the board — central to long-term safety debates.

THE CAPABILITY LADDER

```
NARROW   -> one task, superhuman in a slice   [TODAY]
GENERAL  -> any task, human-level breadth     [RESEARCH]
SUPER    -> beyond human across all domains   [THEORETICAL]
```

**Anchor your language.** When a vendor says "AI", they almost always mean narrow AI. Precise category language keeps strategy conversations honest.

TYPES OF AI

# AI by Function

**REACTIVE** No memory; responds to the current input only. Classic game-playing engines are the textbook case.

**LIMITED** Limited-memory systems learn from recent data — the category nearly all modern ML belongs to.

**THEORY** Theory-of-mind AI would model others' beliefs and intentions. Emerging, not solved.

**SELF** Self-aware AI with its own consciousness — purely theoretical.

FUNCTION TIERS

```
Reactive        -> stateless responses
Limited-memory  -> learns from recent history   [MOST ML]
Theory-of-mind  -> models other agents
Self-aware      -> hypothetical
```

**Place your systems.** Almost everything you will govern is limited-memory AI. Naming the tier clarifies what it can and cannot do.

TYPES OF AI

# AI by Technique

| | |
|---|---|
| ML | Machine learning: systems that improve from data rather than explicit rules. |
| DL | Deep learning: layered neural networks behind modern vision and language. |
| NLP | Natural-language processing: understanding and generating human language. |
| CV | Computer vision: interpreting images and video. |
| GEN | Generative AI: producing new text, images, audio or code. |

**HOW THE TECHNIQUES NEST**

```
Machine Learning
  +- Deep Learning
       +- NLP / Computer Vision
            +- Generative AI (LLMs, diffusion)
```

**They are layers, not rivals.** Generative AI sits on deep learning, which sits on machine learning. Knowing the stack tells you where risk enters.

TYPES OF AI

# AI by Autonomy

**ASSISTED** The system informs; the human decides and acts. Lowest risk, highest oversight.

**AUGMENTED** Human and system decide together; the human can override.

**AUTONOMOUS** The system decides and acts within set bounds; humans monitor by exception.

**THE HUMAN-IN-THE-LOOP DIAL**

```
ASSISTED    [human acts]        oversight: HIGH
AUGMENTED   [human + system]    oversight: MEDIUM
AUTONOMOUS  [system acts]       oversight: BY EXCEPTION
```

**Autonomy sets the controls.** The more autonomy you grant, the more your governance must shift from approving actions to auditing them.

TYPES OF AI

# Spotlight: Generative AI

**WHAT** Models that generate new content — text, image, audio, code — from patterns learned in training data.

**WHY** It collapses the cost of drafting, summarising and prototyping, which is why it dominates enterprise AI roadmaps.

**PARTS** Foundation model, prompt interface, retrieval/grounding layer, and a guardrail/eval wrapper.

**WATCH** Hallucination, data leakage, and IP exposure — the three risks every governance review must cover.

GEN-AI SYSTEM SHAPE

```
Input -> [grounding/retrieval] -> Foundation model
      -> [guardrails + evals] -> Output -> Human review
```

**New power, familiar discipline.** Generative AI is novel; the governance it needs — data control, testing, oversight — is not.

TEAR-OUT

# Types of AI — Reference Card

The one-glance map. Classify any system across all four axes before you govern it.

| Axis | Categories | Where today's systems sit |
|---|---|---|
| **Capability** | Narrow / General / Super | Narrow |
| **Function** | Reactive / Limited / Theory / Self | Limited-memory |
| **Technique** | ML / DL / NLP / CV / GenAI | All, layered |
| **Autonomy** | Assisted / Augmented / Autonomous | Assisted → Augmented |

## Use it as a one-line profile

EXAMPLE CLASSIFICATION

```
A fraud-detection model:
  Capability: Narrow | Function: Limited-memory
  Technique:  ML/DL  | Autonomy: Augmented
```

SECTION B

# The Core AI Frameworks

The certification is built on six frameworks that an enterprise needs to adopt AI responsibly — from strategy down to day-to-day risk control. They interlock: strategy sets direction, governance sets rules, the lifecycle does the work.

### AI Strategy 1

Aligning AI initiatives to business value and readiness.

### AI Governance 2

The policies, roles and controls that keep AI accountable.

### AI Risk Management 3

Identifying, assessing and mitigating AI-specific risk.

### Responsible AI 4

Fairness, transparency, privacy and human oversight in practice.

### AI Lifecycle (MLOps) 5

From data and modelling to deployment and monitoring.

### Enterprise AI Operating Model 6

Structuring people, process and platform to scale AI.

**They form a system.** Strategy without governance is reckless; governance without a lifecycle is theatre. The certification teaches them together.

CORE FRAMEWORKS

# Framework 1 — AI Strategy

**WHAT** A structured way to choose where AI creates value and whether the organisation is ready to deliver it.

**WHY** Without it, AI spend scatters across pilots that never reach production.

**PARTS** Value mapping, use-case prioritisation, data & talent readiness, build-vs-buy, and a roadmap.

**WATCH** Starting from the technology instead of the business problem.

**PRIORITISATION GRID**

```
             LOW effort        HIGH effort
HIGH value   QUICK WINS         BIG BETS
LOW value    FILL-INS           AVOID
```

**Value first.** Score every candidate use case on value and feasibility before a single model is trained.

CORE FRAMEWORKS

# Framework 2 — AI Governance

**WHAT** The system of policies, roles, and decision rights that keeps AI lawful, ethical and accountable.

**WHY** Regulators, customers and boards now expect demonstrable control over how AI is built and used.

**PARTS** Principles, policies, an oversight body, model inventory, and audit trails.

**WATCH** Policies on paper with no owner, no inventory, and no enforcement.

**GOVERNANCE STACK**

```
PRINCIPLES   -> what we believe
POLICIES     -> the rules that follow
CONTROLS     -> how rules are enforced
OVERSIGHT    -> who is accountable + audit
```

**Map to the regulation.** Align your governance to recognised references such as the NIST AI RMF and ISO/IEC 42001 so audits have something to anchor on.

CORE FRAMEWORKS

# Framework 3 — AI Risk Management

**WHAT** A repeatable cycle for finding, rating and reducing the risks an AI system introduces.

**WHY** AI fails in ways traditional software does not — bias, drift, opacity, and novel attack surfaces.

**PARTS** Identify → assess → mitigate → monitor, mirroring the NIST AI RMF functions.

**WATCH** Treating risk as a one-time sign-off instead of a continuous loop.

**NIST AI RMF, IN BRIEF**

```
GOVERN   -> culture and accountability (wraps all)
MAP      -> identify context and risk
MEASURE  -> analyse and track
MANAGE   -> prioritise and act
```

**Risk is continuous.** Models drift after deployment. Risk management that stops at launch is the most common and costly mistake.

CORE FRAMEWORKS

# Framework 4 — Responsible AI

**WHAT** Putting ethical principles — fairness, transparency, privacy, accountability — into operational practice.

**WHY** Trust is the licence to operate; one visible failure can end a program.

**PARTS** Fairness testing, explainability, privacy-by-design, human oversight, and documentation.

**WATCH** Principles that never translate into a test, a metric, or a gate in the pipeline.

**PRINCIPLES -> PRACTICE**

```
Fairness        -> bias tests across groups
Transparency    -> model cards + explanations
Privacy         -> minimisation + consent
Accountability  -> named owner + human override
```

**Make it testable.** A responsible-AI principle you cannot measure is a slogan. Each one should map to a check in the lifecycle.

CORE FRAMEWORKS

# Framework 5 — AI Lifecycle & MLOps

| | |
|---|---|
| WHAT | The end-to-end path a model travels: data, build, validate, deploy, monitor, retire. |
| WHY | Most AI value is lost between a working prototype and a reliable production system. |
| PARTS | Data pipeline, training, evaluation, deployment, monitoring, and retraining. |
| WATCH | No monitoring. A model unwatched in production silently decays. |

THE LIFECYCLE LOOP

```
DATA -> TRAIN -> EVALUATE -> DEPLOY -> MONITOR
  ^                                      |
  +---------------- RETRAIN <------------+
```

**Close the loop.** Monitoring feeds retraining. A lifecycle that is a straight line, not a loop, will degrade.

CORE FRAMEWORKS

# Framework 6 — Enterprise AI Operating Model

**WHAT** How an organisation structures people, process and platform so AI scales beyond isolated pilots.

**WHY** Scaling AI is an operating-model problem long before it is a technology problem.

**PARTS** Operating pattern (centralised, federated, or hub-and-spoke), a centre of excellence, and shared platforms.

**WATCH** Every team rebuilding the same tooling because there is no shared platform.

**THREE OPERATING PATTERNS**

```
CENTRALISED   one team builds for all     [control]
FEDERATED     each unit builds its own     [speed]
HUB & SPOKE   shared core + local teams    [balance]
```

**Hub-and-spoke wins at scale.** A central platform and standards with embedded local teams is the pattern most mature enterprises converge on.

SECTION C

# Framework Infographics

The frameworks at a glance — how they connect, and the questions each one answers. These are the mental models to keep on a wall.

## How the six frameworks fit together

THE STACK, TOP TO BOTTOM

```
STRATEGY        -> where and why we use AI
  GOVERNANCE    -> the rules and accountability
    RISK MGMT   -> what could go wrong, and controls
    RESPONSIBLE -> fairness, transparency, oversight
  LIFECYCLE     -> how a model is actually built/run
OPERATING MODEL -> how it all scales across the org
```

**Read it vertically.** Strategy sets direction at the top; the operating model carries the weight at the bottom. Governance, risk and responsibility are the load-bearing middle.

AT A GLANCE

# Infographic — Govern & Manage Risk

**Govern**
Set culture, accountability and decision rights across the AI program.

**Map**
Establish context: purpose, stakeholders, and where risk can arise.

**Measure**
Analyse and track risk with metrics, tests and monitoring.

**Manage**
Prioritise risks and act — mitigate, transfer, or accept with sign-off.

**Model inventory**
A live register of every model, its owner, and its risk tier.

**Audit trail**
Evidence of decisions and tests, ready for internal or external review.

AT A GLANCE

# Infographic — Responsible AI in Practice

**Fairness**
Test outcomes across groups; correct disparate impact before launch.

**Transparency**
Publish model cards and provide explanations for decisions.

**Privacy**
Minimise data, secure it, and honour consent and retention rules.

**Accountability**
Name a human owner with authority to override the system.

**Human oversight**
Keep a person in or on the loop, sized to the autonomy level.

**Documentation**
Record data sources, limits, and intended use for every model.

AT A GLANCE

# Infographic — The AI Lifecycle

**Data**
Collect, clean, label and version the data that feeds the model.

**Train**
Select an approach and fit the model to the prepared data.

**Evaluate**
Test accuracy, fairness, robustness and fitness for purpose.

**Deploy**
Release behind controls, with rollback and access management.

**Monitor**
Watch for drift, performance decay and emerging harms.

**Retrain**
Feed monitoring back into the next version — close the loop.

SECTION D

# The Official Syllabus

What the Certified AI Frameworks Professional program covers, module by module. This is the spine of both the course and the exam.

| Module | Focus |
|---|---|
| **M1 · AI Foundations** | Types of AI, key techniques, and the vocabulary of the field. |
| **M2 · AI Strategy** | Value mapping, use-case prioritisation, readiness and roadmaps. |
| **M3 · AI Governance** | Principles, policies, oversight bodies, and model inventories. |
| **M4 · Risk Management** | The identify-assess-mitigate-monitor cycle and the NIST AI RMF. |

**Maps to your work.** Each module is written to be applied to a system you already own — not learned in the abstract.

SECTION D

# The Official Syllabus — cont.

| Module | Focus |
|---|---|
| **M5 · Responsible AI** | Fairness, transparency, privacy, accountability and human oversight. |
| **M6 · AI Lifecycle & MLOps** | Data, training, evaluation, deployment, monitoring and retraining. |
| **M7 · Enterprise Operating Model** | Operating patterns, centres of excellence and shared platforms. |
| **M8 · Regulation & Standards** | NIST AI RMF, ISO/IEC 42001, the EU AI Act and sector rules. |

## Standards you will be able to reference

- **NIST** AI Risk Management Framework
- **ISO** ISO/IEC 42001 AI management systems
- **EU** EU AI Act risk tiers
- **OECD** AI principles

SECTION D

# Syllabus — Skills & Outcomes

What each module proves you can do — phrased for your résumé and your next project brief.

| Module | You can demonstrate… |
|---|---|
| **M1–M2** | Classifying AI systems and building a value-ranked AI roadmap. |
| **M3–M4** | Standing up a governance structure and a working risk-management cycle. |
| **M5** | Turning responsible-AI principles into pipeline checks and metrics. |
| **M6** | Designing a monitored, retrainable AI lifecycle. |
| **M7** | Selecting an operating model that scales AI across the organisation. |
| **M8** | Mapping a deployment to NIST, ISO and EU AI Act requirements. |

**Portfolio-ready.** Every row is a deliverable you can build during the program and defend in an interview.

SECTION E

# The Exam Blueprint

How the certification exam is weighted, so you can study to the test rather than guess at it.

| Domain | Weight | What it tests |
|---|---|---|
| **Foundations** | 15% | Types, techniques, terminology |
| **Strategy** | 15% | Prioritisation, readiness, roadmaps |
| **Governance** | 20% | Policies, oversight, inventories |
| **Risk & Responsible AI** | 25% | RMF cycle, fairness, oversight |
| **Lifecycle & Operating** | 15% | MLOps, scaling patterns |
| **Regulation** | 10% | NIST, ISO 42001, EU AI Act |

**Where to spend your time.** Risk, responsible AI and governance together are nearly two-thirds of the exam. Weight your study accordingly.

SECTION E

# Exam Format & Preparation

## Format at a glance

| Item | Detail |
|---|---|
| **Style** | Multiple-choice and scenario questions |
| **Basis** | The eight syllabus modules, by blueprint weight |
| **Mode** | Online, remotely accessible |
| **Result** | Digital certificate and shareable badge on passing |

## A four-step study plan

- ✓ Read the Types-of-AI guide until classification is automatic.
- ✓ Work each framework against a real system you know.
- ✓ Drill the high-weight domains: governance, risk, responsible AI.
- ✓ Complete the Learn-by-Doing list, then sit a practice run.

**Study to the blueprint.** Let the weightings, not your comfort zone, decide where your revision hours go.

SECTION F

# The Learn-by-Doing List

The certification is applied. These are the deliverables you build — each maps to a framework in this guide.

**1. AI use-case portfolio**
Rank candidate use cases on value and feasibility.

**2. AI strategy roadmap**
A phased plan tied to readiness and business value.

**3. Governance charter**
Principles, policies, roles and an oversight body.

**4. Model inventory + risk register**
A live record of models, owners and risk tiers.

**5. Risk assessment**
Run one system through map-measure-manage.

**6. Responsible-AI checklist**
Fairness, privacy and oversight, as pipeline gates.

**7. Lifecycle + monitoring plan**
Define drift metrics and a retraining trigger.

**8. Operating-model design**
Choose and justify a pattern to scale AI.

SECTION F

# Learn-by-Doing — How It Works

Each deliverable follows the same rhythm, so practice compounds across the program.

### Brief

A real-world scenario sets the goal and the constraints.

### Build

You produce the artefact using the framework from the guide.

### Review

It is assessed against a rubric mapped to the syllabus.

### Refine

Feedback drives one improvement pass before sign-off.

### Reuse

The artefact becomes part of your professional portfolio.

### Relate

Each artefact links to the next, building one coherent program.

**Evidence, not just answers.** You leave with eight artefacts you can show an employer — proof you can run AI frameworks, not just define them.

FROM READER TO CERTIFIED

# Your path to certification

Four AI frameworks. Eight modules. One credential.

| | |
|---|---|
| **Know the types** | Capability · function · technique · autonomy |
| **Learn the frameworks** | Strategy · governance · risk · responsible · lifecycle · operating model |
| **Study the syllabus** | Eight modules, mapped to the exam blueprint |
| **Know the exam** | Weighted by domain — study where it counts |
| **Learn by doing** | Eight portfolio-ready deliverables |
| **Get certified** | Digital certificate + shareable badge |