CERTIFICATION SYLLABUS

Certified Artificial Intelligence Practitioner

Exam Code: CAIP™



A complete preparation guide covering exam objectives, core competencies, domain weightage, and topics to prepare.

©2026 Certaining, Inc. All rights reserved.

Table of Contents

1. **Exam Name & Code** 1
2. **Table of Contents** 2
3. **Objective of the Certification** 3
4. **Core Competencies** 4-5
   - *Machine Learning Workflow & Data Preparation*
   - *Supervised Learning: Regression & Classification*
   - *Unsupervised Learning & Feature Engineering*
   - *Model Evaluation, Tuning & Deployment*
   - *Deep Learning, NLP & Computer Vision*
   - *Generative AI & Prompt Engineering in Practice*
   - *Responsible AI, Security & MLOps*
5. **Who Should Take This Certification** 6
6. **Exam Information** 6
7. **Examination Domain & Weightage** 7
8. **Topics to Prepare** 8-9
   - *Machine Learning Workflow & Data Preparation*
   - *Supervised Learning: Regression & Classification*
   - *Unsupervised Learning & Feature Engineering*
   - *Model Evaluation, Tuning & Deployment*
   - *Deep Learning, NLP & Computer Vision*
   - *Generative AI & Prompt Engineering in Practice*
   - *Responsible AI, Security & MLOps*

©2026 Certaining, Inc. All rights reserved.

Objective of the Certification

- **Build a practical understanding of the machine learning workflow:** from problem framing and data collection through training, evaluation, deployment, and monitoring.
- **Develop skill in data preparation:** cleaning, transforming, splitting, and validating data so that models are trained on reliable inputs.
- **Apply supervised learning techniques:** choose and apply regression and classification algorithms such as linear and logistic regression, decision trees, random forests, and support vector machines.
- **Apply unsupervised learning and feature engineering:** use clustering and dimensionality reduction, and create or select features that improve model performance.
- **Evaluate and improve models:** select suitable metrics, use cross-validation, diagnose overfitting and underfitting, and tune hyperparameters.
- **Understand deep learning fundamentals:** neural network architectures, training concepts such as backpropagation, and when deep learning is the right choice.
- **Work with natural language processing and computer vision solutions:** apply common techniques and pre-trained models to text and image problems.
- **Use generative AI effectively in practice:** apply prompt engineering, retrieval-augmented generation, and fine-tuning concepts to business tasks while managing their limitations.
- **Apply responsible AI, security, and MLOps practices:** address bias, privacy, explainability, adversarial risks, versioning, and monitoring across the model lifecycle.
- **Create a foundation for progression to advanced credentials:** building on CAIF™ and preparing for advanced AI engineering and leadership certifications.

©2026 Certaining, Inc. All rights reserved.

Core Competencies — CAIP™ (Artificial Intelligence Practitioner)

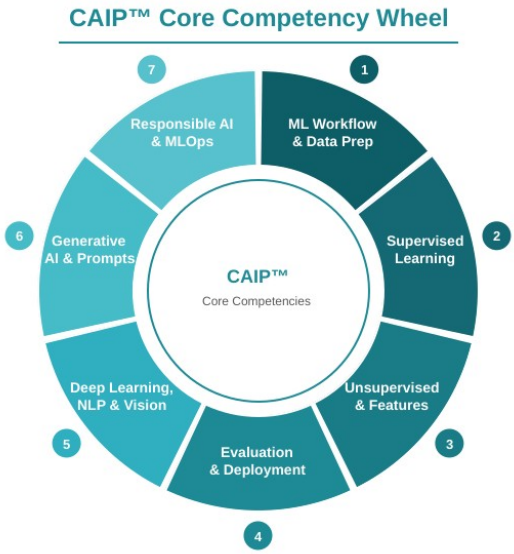

1. Machine Learning Workflow & Data Preparation

- Describe each stage of the machine learning lifecycle, from problem definition to monitoring.
- Clean, transform, and encode data, and handle missing values, outliers, and imbalanced classes.
- Split data into training, validation, and test sets while avoiding data leakage.

2. Supervised Learning: Regression & Classification

- Choose between regression and classification for a given business problem.
- Explain how linear regression, logistic regression, decision trees, random forests, k-nearest neighbors, and support vector machines work at a practical level.
- Compare algorithms on interpretability, accuracy, training cost, and data requirements.

3. Unsupervised Learning & Feature Engineering

- Apply clustering methods such as k-means and hierarchical clustering and interpret the results.
- Explain dimensionality reduction, including principal component analysis, and when to use it.
- Create, scale, and select features that improve model performance.

©2026 Certaining, Inc. All rights reserved.

4. Model Evaluation, Tuning & Deployment

- Select metrics such as accuracy, precision, recall, F1 score, ROC-AUC, and RMSE that match the business goal.
- Use cross-validation, regularization, and hyperparameter tuning to improve generalization.
- Describe deployment options, including batch, real-time, and edge, and how deployed models are monitored.

5. Deep Learning, NLP & Computer Vision

- Explain neural network components: layers, activation functions, loss functions, and backpropagation.
- Distinguish convolutional, recurrent, and transformer architectures and their typical uses.
- Apply pre-trained models and transfer learning to text and image tasks.

6. Generative AI & Prompt Engineering in Practice

- Write effective prompts using roles, context, examples, constraints, and output formats.
- Explain retrieval-augmented generation and when to use it instead of fine-tuning.
- Evaluate generative AI output for accuracy, hallucination, and fit for purpose.

7. Responsible AI, Security & MLOps

- Identify and reduce bias, and use explainability techniques to support transparent decisions.
- Recognize security threats to AI systems, such as adversarial inputs, data poisoning, prompt injection, and data leakage, and the controls that reduce them.
- Describe MLOps practices: version control for data and models, automated pipelines, monitoring for drift, and retraining.

©2026 Certaining, Inc. All rights reserved.

Who Should Take This Certification

- Data analysts and business analysts moving into machine learning and AI projects.
- Software developers and engineers who build or integrate AI-enabled applications.
- Data engineers and database professionals who prepare data for machine learning.
- Cloud and DevOps engineers who support AI and MLOps workflows.
- Project and product managers who deliver AI initiatives and need a hands-on understanding of the work.
- Professionals who hold CAIF™ or equivalent knowledge and want to demonstrate practical AI skills.

Exam Information

Exam Name	Artificial Intelligence Practitioner
Exam Code	CAIP™
Exam Format	Multiple Choice Questions (MCQ), including scenario-based questions
Duration	120 minutes
Number of Questions	75
Passing Score	70%
Mode of Exam	Test Center Based Testing / Online Testing
Language	English
Validity	Lifetime
Open Book	No

©2026 Certaining, Inc. All rights reserved.

Examination Domain & Weightage

Domain	Weightage
Machine Learning Workflow & Data Preparation	16%
Supervised Learning: Regression & Classification	16%
Unsupervised Learning & Feature Engineering	12%
Model Evaluation, Tuning & Deployment	16%
Deep Learning, NLP & Computer Vision	16%
Generative AI & Prompt Engineering in Practice	12%
Responsible AI, Security & MLOps	12%
Total	**100%**

©2026 Certaining, Inc. All rights reserved.

Topics to Prepare

The topics below represent the core areas candidates should focus on for each domain, reflecting the practical, intermediate depth expected at the CAIP™ level.

1. Machine Learning Workflow & Data Preparation (16%)

- **ML Lifecycle:** Problem framing, data collection, training, evaluation, deployment, and monitoring.
- **Data Collection and Quality:** Data sources, labeling, completeness, accuracy, and consistency.
- **Data Cleaning and Transformation:** Handling missing values and outliers, encoding categories, and scaling numeric data.
- **Data Splitting:** Training, validation, and test sets, and how to prevent data leakage.
- **Imbalanced Data:** Resampling and weighting approaches for rare outcomes.

2. Supervised Learning: Regression & Classification (16%)

- **Regression:** Linear and polynomial regression, and regularized variants.
- **Classification:** Logistic regression, k-nearest neighbors, and support vector machines.
- **Tree-Based Models:** Decision trees, random forests, and gradient boosting at a practical level.
- **Choosing an Algorithm:** Trade-offs among accuracy, interpretability, speed, and data size.

3. Unsupervised Learning & Feature Engineering (12%)

- **Clustering:** K-means, hierarchical clustering, and how to judge cluster quality.
- **Dimensionality Reduction:** Principal component analysis and its uses in visualization and noise reduction.
- **Anomaly Detection:** Identifying unusual records without labeled examples.
- **Feature Engineering:** Creating, scaling, and selecting features that improve results.

4. Model Evaluation, Tuning & Deployment (16%)

- **Evaluation Metrics:** Accuracy, precision, recall, F1 score, ROC-AUC, MAE, and RMSE, and when to use each.
- **Validation Techniques:** Cross-validation and hold-out testing.
- **Improving Generalization:** Detecting overfitting and underfitting, regularization, and hyperparameter tuning.
- **Deployment Patterns:** Batch, real-time, API-based, and edge deployment.
- **Monitoring in Production:** Tracking performance, data drift, and model drift.

©2026 Certaining, Inc. All rights reserved.

5. Deep Learning, NLP & Computer Vision (16%)

- **Neural Network Basics:** Layers, activation functions, loss functions, gradient descent, and backpropagation.
- **Architectures:** Convolutional networks, recurrent networks, and transformers.
- **Natural Language Processing:** Tokenization, embeddings, text classification, and named entity recognition.
- **Computer Vision:** Image classification, object detection, and segmentation.
- **Transfer Learning:** Reusing pre-trained models to save data and training time.

6. Generative AI & Prompt Engineering in Practice (12%)

- **Prompt Design:** Roles, context, examples, constraints, and output formats.
- **Prompting Techniques:** Zero-shot, few-shot, and step-by-step reasoning prompts.
- **Retrieval-Augmented Generation:** Grounding model output in trusted organizational data.
- **Fine-Tuning versus Prompting:** Choosing the right approach for cost, accuracy, and control.
- **Evaluating Output:** Checking for hallucinations, bias, and consistency.

7. Responsible AI, Security & MLOps (12%)

- **Bias and Fairness:** Detecting and reducing bias, and measuring fairness across groups.
- **Explainability:** Feature importance and other methods that make model decisions understandable.
- **AI Security Threats:** Adversarial inputs, data poisoning, model theft, and prompt injection.
- **Privacy and Governance:** Data minimization, consent, documentation, and regulatory awareness.
- **MLOps:** Versioning, automated pipelines, monitoring, and retraining.

©2026 Certaining, Inc. All rights reserved.