Training Sessions Outline: IBM Watson AI

Duration: 5 days

Day 1: Introduction to IBM Watson & AI Fundamentals

A. What is AI and Cognitive Computing?

- Defining AI, Machine Learning, Deep Learning

- The role of cognitive computing in business

B. Introduction to IBM Watson

- Overview of the Watson ecosystem

- Key capabilities and use cases across industries

- The "build, run, manage" paradigm of Watson services

C. Setting up your IBM Cloud Account & Watson Services

- Navigating the IBM Cloud platform

- Provisioning Watson services (Discovery, NLP, Orchestrate)

- Understanding service credentials and API keys

D. Data Privacy and Ethical AI Considerations with Watson

- IBM's approach to data governance and security

- Ethical guidelines for AI development and deployment

Deep Dive into IBM Watson Discovery

A. Understanding Watson Discovery's Core Functionality

- What is Discovery? Information extraction and search

- Key components: Collections, Configurations, Queries

- Use cases: Enterprise search, document analysis, content intelligence

B. Data Ingestion and Management

- Supported data sources (web crawls, file uploads, Box, SharePoint, etc.)

- Configuring data connectors

- Data enrichment processes (Smart Document Understanding, optical character recognition)

- Document conversion and normalization

C. Hands-on Lab: Building a Custom Search Solution with Watson Discovery

- Ingesting sample data

- Creating a custom SDU model

- Running basic and advanced queries

- Basic integration demonstration

Day 2: Automating Workflows with IBM Watson Orchestrate & Advanced Discovery

A. Enriching Your Data with Smart Document Understanding (SDU)

- Introduction to SDU: Defining custom fields and entities

- Training SDU models: Labeling documents, iterating and refining

- Applying SDU to enhance search relevance

B. Querying Watson Discovery

- Discovery Query Language (DQL) basics

- Constructing complex queries: Filtering, aggregations, natural language queries

- Relevancy training: Improving search results

- Understanding query parameters and result formats

C. Introduction to Watson Orchestrate

- What is process automation and hyper automation?

- The role of Watson Orchestrate in business process automation

- Key features: Digital skills, conversational AI, intelligent automation

D. Understanding Digital Skills and Connectors

- What are digital skills? Reusable automation components

- Connecting to applications (Salesforce, Workday, Slack, etc.)

- Building custom digital skills for unique needs

E. Hands-on Lab: Automating a Simple Business Process with Orchestrate

- Connecting to a sample application

- Creating a digital skill

- Building and testing a multi-step workflow

- Interacting with the workflow via conversational AI

Day 3: Mastering IBM Watson Natural Language Processing (NLP)

A. Fundamentals of Natural Language Processing (NLP)

- What is NLP? Text analysis and understanding

- Key NLP tasks: Tokenization, stemming, lemmatization, part-of-speech tagging

- Challenges in NLP and how Watson addresses them

B. Introduction to Watson NLP Services

- Overview of available NLP models and capabilities (Emotion, Sentiment, Syntax, Entity extraction, etc.)

- Understanding pre-trained models vs. custom models

C. Text Analysis with Watson NLP

- Sentiment Analysis: Detecting positive, negative, and neutral sentiment

- Emotion Analysis: Identifying joy, sadness, anger, fear, disgust

- Syntax Analysis: Understanding grammatical structure

- Keyword Extraction: Identifying important terms and phrases

D. Entity Extraction and Customization

- Named Entity Recognition (NER): Identifying people, organizations, locations

- Custom Entity Extraction: Training custom models for domain-specific entities

- Using Watson Knowledge Studio for custom models

E. Integrating Watson NLP into Applications

- Utilizing the Watson NLP API for various programming languages

- Real-world examples of NLP in action

F. Hands-on Lab: Analyzing Text Data with Watson NLP

- Performing sentiment and emotion analysis on sample text

- Extracting entities and keywords

- Building a simple custom entity extraction model

Day 4: Introduction to watsonx.ai, Prompt Engineering, and watsonx.ai.studio

A. What is IBM watsonx.ai?

- Overview of the watsonx platform and its components.

- The role of Foundation Models and Generative AI.

- Key capabilities and use cases for watsonx.ai.

B. Deep Dive into watsonx.ai.studio

- Understanding the core functionality of watsonx.ai.studio.

- Key components: Prompt Lab, Foundation Models, and projects.

- Using Prompt Lab to interact with various foundation models.

C. Fundamentals of Prompt Engineering

- Principles of effective prompting.

- Prompt techniques: Zero-shot, few-shot, and Chain-of-Thought prompting.

- Strategies for controlling model output and ensuring desired results.

D. Hands-on Lab: Exploring watsonx.ai.studio and Prompting

- Creating a new project in watsonx.ai.studio.

- Experimenting with different foundation models in the Prompt Lab.

- Crafting and refining prompts for various tasks (e.g., text summarization, content generation).

- Evaluating the output of different foundation models.

Day 5: Building AI Assistants and Deploying Models with watsonx.ai

A. Introduction to watsonx.ai.assistant

- Overview of watsonx.ai.assistant and its role in building conversational AI.

- Key components: Actions, intents, entities, and dialog flow.

- Designing and building a simple conversational experience.

B. Integrating with Back-end Systems & External Data

- Connecting watsonx.ai.assistant to external APIs and applications.

- Handling dynamic responses and data retrieval.

- Combining watsonx.ai with Watson Discovery for enhanced knowledge retrieval.

C. Understanding watsonx.ai.runtime and Model Deployment

- The role of watsonx.ai.runtime for deploying and managing models.

- Key concepts: Deployments, endpoints, and inference.

- Accessing the deployed model via a REST API endpoint.

D. Hands-on Lab: Deploying and Testing a Model

- Creating a new watsonx.ai.assistant and defining an intent.

- Deploying a model from watsonx.ai.studio to watsonx.ai.runtime.

- Calling the deployed model's API endpoint using a tool like curl or a simple Python script.

- Presenting and testing the final assistant.