



ENTERPRISE
ARTIFICIAL INTELLIGENCE
ENGINEER



CADS

FUTURE PROOFING THE WORLD!

Enterprise Artificial Intelligence Engineer

Turn data into effective
business insights.

APPLY NOW!

Instructor-led workshops
+ online learning



**Pre-
Requisite/Requirement for
Registration:**

Python Programming,
Descriptive Statistics

**Minimum Training Qualification for
Audiences (e.g., undergraduate
degree, postgraduate, etc.)**

Undergraduate Degree
Skills



Learning Outcome



Collect, summarize, visualize and explore raw data to transform it into actionable business insights.



Tame your data with Python - write a workable program to solve data related business problem using Python's powerful data processing libraries.



Learn how to think like a programmer by breaking complicated business problems into smaller ones and then solving these byte-sized problems using coding and automation.



Speak the language of Structured Query Language (SQL) fluently and unlock the capability to participate in database application development teams.



Uncover interconnected relations between data sets and generate comprehensive reports for business insights through drill-down analysis.



Explore how programming and automation can be used to chunk data faster and save you time.



Super-charge yourself with big data analytics skills to analyze huge volumes of data effectively and effortlessly by building smart programs.

Programs



Prepare to embark on the next frontier of AI/ML Engineer.

This comprehensive program is designed to equip participants with the foundational and advanced skills needed to excel as an AI/ML Engineer.

The curriculum spans key areas of Machine Learning, Deep Learning, Reinforcement Learning, Natural Language Processing (NLP), and Large Language Models (LLMs). Participants will gain hands-on experience in implementing and fine-tuning models using Python, explore real-world applications, and address challenges such as bias, efficiency, and ethical considerations.

By the end of the course, attendees will be proficient in developing cutting-edge AI/ML solutions for various industries, including healthcare, finance, and customer service.

Course Summary

Machine Learning (Supervised and Unsupervised Learning)

- Covering regression, classification, clustering, dimensionality reduction, and recommender systems.
- Hands-on techniques like hyperparameter tuning, cross-validation, and model evaluation ensure robust understanding.

Deep Learning

- Exploring Artificial Neural Networks (ANNs), Convolutional Neural Networks (CNNs), and Recurrent Neural Networks (RNNs).
- Case studies include digit classification, image recognition, and text analysis.

Reinforcement Learning


- Introducing foundational principles like Markov Decision Processes, key algorithms (Q-Learning, SARSA), and advanced methods like Deep Q-Networks.
- Real-world applications such as robotics, gaming, and autonomous decision-making systems.

Natural Language Processing (NLP)

- Focusing on text pre-processing, feature extraction, sentiment analysis, and topic modeling.
- Advanced tasks include using neural embeddings, CNNs, LSTMs, and transformers for NLP challenges.

Large Language Models (LLMs)

- Exploring the architecture and applications of state-of-the-art models like GPT and BERT.
- Hands-on projects in fine-tuning, deploying, and optimizing LLM-powered applications like chatbots and summarizers.



Throughout the course, participants will engage in practical projects to reinforce learning, gain experience with cutting-edge tools, and prepare for real-world challenges in AI and ML.

Summary Of Information On Each Course/Module

Introduction to Machine Learning

- Machine Learning is a broad field. This lecture gives an overview of all the different topics covered by Machine Learning

Supervised Machine Learning: Regression

- Regression describes the process of predicting a continuous, numerical value from features in data. Models covered include Linear regression & Regularization

Supervised machine Learning: Classification

- Classification describes the process of predicting a discrete category from features in data. Models covered include
- K-nearest neighbours
- Decision trees
- Logistic regression
- Naïve Bayes
- Ensemble Methods
- Support Vector Machines

Best Training and Evaluation Practices

- Correctly training and evaluating a model is important to ensure that results are robust and reliable. In this module, participants will learn best practices for training models, the consequences of not doing so, how to fine-tune hyperparameters of their models, and how to assess a model's performance. Topics covered include
- Validation metrics
- Cross validation
- Hyper-parameter tuning

Unsupervised Machine Learning: Clustering

- Clustering describes the process of grouping data points based only on their features without the help of labeled data to train a model on. Models covered include
- Hierarchical clustering
- K-Means clustering
- DBSCAN

Unsupervised machine Learning: Dimensionality Reduction

- Dimensionality reduction describes the process of transforming features to be more informative and less redundant. Models covered include
- Principal Component Analysis (PCA)
- t-Distributed Stochastic Neighbour Embedding (t-SNE)

Recommender Systems

- Recommender systems can make use of prior user-item interactions, e.g. movie ratings, to predict future interactions and make recommendations. They make use of both supervised and unsupervised methods to achieve this. Topics covered include
- Collaborative filtering
- Content-based filtering
- Matrix factorization

Supervised and Unsupervised
Machine Learning: Days 1 to 5

Summary Of Information On Each Course/Module

Introduction to Artificial Neural Networks

- It is important to know why Artificial Neural Networks exists, how it came to be, and what it has achieved so far. Several pursuits in Artificial Intelligence, including Computer Vision, have collectively contributed to this area of knowledge. Microscopic to cosmic processes have already benefited from Artificial Neural Networks. Remarkable applications of Artificial Neural Networks will be covered.

Fully Connected (Densely Connected) Neural Networks

- The basic form of Artificial Neural Network is the Fully Connected Network. Trainees will learn its intuition and why it is also called a Densely Connected Network while covering its components: Neurons and Layers. Expertise in Deep Learning requires a thorough understanding of Fully Connected Neutral Networks. The process of Training the Neural Network is otherwise called Deep Learning. Activation Functions are set to capture the complexity unique to the problem. Trainees will also learn how to initialize parameters and tune their neural networks.
- Case Study - Trainees will build a handwritten digit classifier by building a Fully Connected Neural Network and train it with churn dataset. Hyperparameter tuning will be done as part of optimization. Number of layers and neurons will be varied, and accuracy will be compared.

Convolutional Neural Networks

- Some datasets such as large multi-channel (color) imageries are often impossible to train. Convolutional Neural Networks (CNN) employs Convolving Layers (Kernels) that takes the most salient and unique features of the image data. Computational efficiency is increased by employing layers of Dimensionality Reduction (Max-pooling) and Sub-sampling. Drop-out Layers are also introduced as a way to combat overfitting. Trainees will learn how to construct a CNN and discern its difference with Fully Connected Networks.
- Case Study - Trainees will build an image classifier by building a Convolutional Neural Network and train it with Fashion-MNIST dataset. Hyperparameter tuning will be done as part of optimization. Number of layers and neurons will be varied, and accuracy will be compared.

Recurrent Neural Networks

- Fully Connected and Convolutional Neural Network lack memory. Recurrent Neural Networks (RNN) mimics how biological memory is built on previous experience (data) and constantly updated when new data is introduced. RNN is a powerful tool in solving problems involving sequence and timeseries data. Trainees will learn the steps in building an RNN and its usability in classifying textual data.
- Case Study I - Trainees will build Simple Recurrent Neural Networks, Long-Short Term Memory, and Gated Recurrent Units in for Text Classification.
- Case Study II - Trainees will be given either image or sequence/timeseries datasets to work on. They will choose, build, and train the appropriate Neural Network architectures (ANN, CNN, or RNN). Most neural network involve a combination of two or more architectures.

Deep Learning: Days 6 to 8

Summary Of Information On Each Course/Module

Reinforcement Learning

Participants will gain a deep understanding of the theoretical foundations of RL, learn about key algorithms, and explore real-world use cases in gaming, robotics, and decision-making systems.

Introduction to Reinforcement Learning

- What is RL? How is it different from supervised and unsupervised learning?
- Core concepts: agents, environments, states, actions, rewards.
- Real-world applications of RL in gaming, robotics, and recommendation systems.

Mathematical Foundations

- Markov Decision Processes (MDPs): states, actions, rewards, transitions.
- The Bellman equation and optimality principles.

Key Algorithms in RL

- Dynamic Programming: Policy Iteration and Value Iteration.

Monte Carlo methods.

- Temporal Difference (TD) Learning: Q-Learning and SARSA.

Deep Reinforcement Learning

- Overview of Deep Q-Networks (DQN).
- Policy Gradient methods and Actor-Critic architectures.
- Advanced concepts: Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO).

Hands-On Coding Session

- Setting up the RL environment using Python.
- Implementing Q-Learning with OpenAI Gym.
- Training a Deep Q-Network using Stable-Baselines.

Exploration vs. Exploitation

- Understanding the trade-off.
- Strategies for efficient exploration: ϵ -greedy, Boltzmann exploration, and more.

Challenges and Solutions in RL

- Sample efficiency and the curse of dimensionality.
- Reward shaping and designing appropriate reward functions.
- Dealing with instability in training.

Course Project

- Build and train an RL agent to play a simple game (e.g., CartPole or LunarLander).
- Evaluate performance and tune hyperparameters.

•

Reinforcement Learning: Days 9 to 10

Summary Of Information On Each Course/Module

Introduction to NLP

Participants will learn the basics of NLP, and the applications and tasks that can be performed using NLP. They will be introduced with the top Python libraries for NLP.

Feature Extraction

- Participants will learn to extract feature from text data. They will also use regular expressions to perform phrase matching. Topics covered include:
- Tokenization ·Phrase Matching ·Number of Words & Sentences ·Extract Stop-words ·Extract the Numbers
- Extract Specific Words ·Extract non ASCII Characters ·Extract the Date ·Frequency Distributions of Words

Text Cleaning & Pre-processing

- Participants will learn to clean and pre-process text data by removing noises. Topics covered include:
- Noise Removal ·Lexicon Normalization

Common NLP Tasks

- Participants will learn to perform some NLP tasks. Topics covered include Part of Speech Tagging, Phrase Detection and Named Entity Recognition.

Sentiment Analysis

- Participants will learn to perform sentiment analysis with a case study on reviews of phones from Amazon. They will learn to represent words as feature vectors. Topics covered include bag-of-words (BOW), term frequency inverse document frequency (tf-idf) and N-grams. They will also learn to build classification model using feature representations of words and visualize the top N-grams in a text data.

Topic Modeling

- Participants will learn to perform topic modelling with a case study on reviews of foods from Amazon. They will perform LDA on text data using two different Python modules, scikit-learn and gensim , and they also will visualize topic modelling using pyLDAvis.

Using Neural Embeddings in Text Classification

- Participants will discover a more powerful word representation in word embeddings, how to develop them as standalone models, how to learn them as part of neural network models, and how to visualize them. Topics covered include: ·Word Embedding ·Document Embedding ·FastText Embedding ·Visualizing Embeddings

Deep Learning for Text Classification

- How to leverage word embeddings, convolutional neural network (CNN) and Long short-term memory (LSTM) to learn deep learning models (in Keras) of text for classification and sentiment analysis, a successor to the bag-of-words model. They will learn to interpret text classification models with Lime and Shap. They will also learn to build deep learning models for classification with no or less training data. Topics covered include: ·CNNs for Text Classification ·CNNs for sentiment analysis ·LSTMs for Text Classification ·Text Classification with Large, Pre-Trained Language Models ·Interpreting Text Classification Models ·Learning with No or Less Data and Adapting to New Domains
-

Natural Language Processing: Days
11 to 13

Summary Of Information On Each Course/Module

This course will dive into the world of Large Language Models (LLMs), such as GPT and BERT, which are revolutionizing AI applications in natural language understanding, generation, and interaction.

Participants will learn the foundational principles behind LLMs, gain hands-on experience fine-tuning and deploying these models, and explore their transformative potential across industries.

By the end of the course, attendees will have the skills to harness the power of LLMs in real-world applications, from chatbots and content generation to sentiment analysis and personalized recommendations.

Introduction to Large Language Models

- What are LLMs, and how do they differ from traditional NLP models?
- Historical evolution of LLMs: from RNNs to Transformers.
- Applications across industries: healthcare, finance, customer service, and more.
- Understanding Transformer Architecture

Core components: self-attention, positional encoding, and feed-forward layers.

- Key breakthroughs: the attention mechanism and scaling laws.
- Differences between encoder-based models (e.g., BERT) and decoder-based models (e.g., GPT).
- Pretraining and Fine-Tuning

Pretraining on large corpora: unsupervised learning for language modeling

- Fine-tuning for specific tasks: classification, summarization, translation, and Q&A.
- Transfer learning and prompt engineering for task-specific use cases.
- Popular LLMs in Practice

Overview of state-of-the-art models: GPT, BERT, T5, LLaMA, and others.

- Comparative analysis of model strengths, limitations, and use cases.
- Hands-On Implementation

Setting up a Python environment for NLP tasks.

- Using pre-trained models with Hugging Face Transformers.
- Fine-tuning an LLM for a specific task (e.g., sentiment analysis or summarization).
- Prompt Engineering and Zero/Few-Shot Learning

Crafting effective prompts for language understanding and generation tasks.

- Strategies for zero-shot and few-shot learning using LLMs.
- Practical exercises to improve model responses.
- Challenges and Ethical Considerations

Computational costs and environmental concerns of training LLMs.

- Addressing biases and ensuring fairness in generated outputs.
- Risks of misuse: misinformation, plagiarism, and data privacy.
- Exploring Deployment

Deploying LLMs via APIs (e.g., OpenAI, Hugging Face Inference API).

- Custom deployment with frameworks like Flask and Streamlit.
- Optimizing for performance and scalability.

Project

- Build and deploy an LLM-powered application, such as:
- A chatbot for customer support.
- A content generation tool for blogs and articles.

Large Language Models:
Days 14 to 15



YAYASAN PENERAJU FINANCING SCHEME



SILVER PACKAGE WITH
FLEXIBLE FINANCING



MAXIMUM FINANCING
UP TO RM 15,000.00



INTEREST

Eligibility Criteria:

- Bumiputera status**
- Malaysian citizen**
- Minimum age requirement is 16 years old or older**
- Must have an active individual bank account in Malaysia**

Silver Flexible Highlights:

- Training / Tuition fee**
- Professional / Examination fee**
- Monthly Allowance***
- Repayment*** **0%**
*Silver Package - Talent may be eligible for zero repayment, subject to the Terms and Conditions. Please refer to Peneraju.org for further details.
- Scheme's limit for reapply*** **One-time only**



Participants Testimonials

I didn't know that Big Data itself is huge. At least I know how to optimize and utilize data better now!

Nur Amira Binti Zamri
Executive, Culture & Engagement, UMW

Love the graphics and animation, information is bite-sized, the content is relevant to the current industry trend. Thank you! 5 stars.

Huwaina Farizal
Student, Universiti Kebangsaan Malaysia (UKM)

Positive, impactful and helpful for talents to apply skills and techniques of better analyzing, preparing and presenting data to the relevant stake holders of the business.

Anonymous
Kian Joo Group

Learned about Cloud, Hadoop ecosystem HDFS, Managing Big Data Software helped me to enhance my knowledge & improve my skills in order to enroll myself into future projects.

Sri Thiba Nathan
RHB



42 CPD
HOURS

For the members of the Malaysian Board of Technologists (MBOT), 42 CPD hours will be awarded upon completion of the Enterprise Data Practitioner (EDP) certification course.

Industry-leading organizations are now driven by data. Because of this, organizations need a data-driven strategy and a plan to develop their **analytics & data science** talents to support this strategy.



PDU Approved

INFORMS is the leading international association for professionals in operations research, analytics, management science, economics, behavioral science, statistics, artificial intelligence, data science, applied mathematics, and other relevant fields. CADS is a Professional Education Partner (PEP) with INFORMS. PEP is an INFORMS program designed to provide a reliable list of organization and university courses that learners can use to earn professional development units (PDUs) and help learners remain Certified Analytics Professionals (CAPs). Holder of AEDA certification can earn PDUs recognized by INFORMS.

We tailor-make and develop specialized skills development training program for all levels of workforce. All of the programs are HRDF claimable

- Data-driven leadership
- Data-driven culture development
- Forming a data science team
- Data driven organization roadmap



ENTERPRISE
ARTIFICIAL INTELLIGENCE
ENGINEER



Per user certification (min 25 pax)

If you have any questions not answered on the website and would require additional information, please get in touch with us; we will be happy to assist.



Unit A10-01, 10th Floor, Blok A, Pusat Dagangan
Phileo Damansara 1, No.9, Jalan 16/11, Off Jalan
Damansara, 46350 Petaling Jaya, Selangor.

SCAN ME



+603 7931 2963



engage@cad.s.ai



<https://www.cads.ai>

**Futureproof
your career**

