

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 1: Mathematical Foundations for AI

Duration: Weeks 1–13 | Format: Fascinating, Logical, Hands-on + Python Coding Integrated

Module Objective

To build a foundational yet intuitive understanding of the core mathematical concepts that power artificial intelligence, enabling learners from 10th standard onwards to think structurally, model abstractly, reason algorithmically, and code practically.

Week 1–2: Logic, Sets, and Relations

Key Question: What is the foundation of machine reasoning?

- Propositional and Predicate Logic (truth tables, logical connectives)
- Set Theory (unions, intersections, Venn diagrams)
- Relations (reflexive, symmetric, transitive), Functions (injective, surjective)

Activities:

- Python coding: Build a digital truth table generator
- Set operations using Python sets
- Logic puzzles solver using boolean expressions

Example: Basic logic gate simulation

A = True

B = False

AND_gate = A and B

OR_gate = A or B

NOT_A = not A

Outcome: Understanding the logic of programming and decision rules in AI systems

Week 3–4: Functions, Algebra, and Probability Intuition

Key Question: How do we use functions and probability to model real-world processes?

- Linear, quadratic, exponential functions
- Function transformations, inverses
- Introduction to probability theory and axioms
- Conditional probability and basic Bayes' Rule

Activities:

- Python coding: Graphing functions using matplotlib
- Probability simulation: coin flips, dice rolls

```
import matplotlib.pyplot as plt
```

```
import numpy as np
```

```
x = np.linspace(-10, 10, 100)
```

```
y = 2*x + 3
```

```
plt.plot(x, y)
```

```
plt.title("Linear Function  $y = 2x + 3$ ")
```

```
plt.grid(True)
```

```
plt.show()
```

Outcome: Ability to reason about uncertainty and create mathematical models

 Week 5–7: Linear Algebra Intuition

Key Question: How do we represent and operate on multi-dimensional data?

- Scalars, vectors, matrices, dot product, norms
- Matrix multiplication and linear transformations
- Geometric intuition: projection, span, orthogonality

Activities:

- Python coding: Matrix operations with NumPy
- Compute cosine similarity between document vectors

```
import numpy as np
```

```
A = np.array([[1, 2], [3, 4]])
```

```
B = np.array([[2, 0], [1, 2]])
```

```
product = np.dot(A, B)
```

```
print("Matrix Product:\n", product)
```

Outcome: Understanding of how data is structured, manipulated, and transformed in AI

 Week 8–10: Calculus and Optimization

Key Question: How do machines learn from feedback?

- Derivatives, gradients, and cost functions

- Chain rule and backpropagation intuition
- Gradient descent and convergence
- Convex vs non-convex landscapes

Activities:

- Python coding: Implement basic gradient descent for a quadratic loss

Simple gradient descent

x = 10 # initial value

learning_rate = 0.1

for i in range(20):

 grad = 2 * x

 x = x - learning_rate * grad

 print(f"Iteration {i+1}: x = {x:.4f}")

Outcome: Appreciation for the learning process and optimization at the heart of AI

Week 11–13: Graph Theory, Trees, and Search Paths

Key Question: How do we represent networks, relationships, and solve path problems?

- Graphs, adjacency matrices, directed vs undirected
- Trees, DFS, BFS, Dijkstra's algorithm
- Centrality, shortest paths, and spanning trees

Activities:

- Python coding: Build and traverse a graph using dictionaries and queues

Simple BFS implementation

from collections import deque

graph = {

 'A': ['B', 'C'],

 'B': ['D', 'E'],

 'C': ['F'],

 'D': [], 'E': [], 'F': []

}

def bfs(start):

 visited = []

```
queue = deque([start])

while queue:
    node = queue.popleft()
    if node not in visited:
        visited.append(node)
        queue.extend(graph[node])

return visited

print(bfs('A'))
```

Outcome: Ability to visualize and compute with network structures used in search, recommendation, and navigation systems

🌟 Capstone Week 13: "Mathematics as the Language of Machines"

Project: Create an illustrated Jupyter notebook portfolio that demonstrates one concept from each of the 5 mathematical domains above, with code, visualizations, and real-world application examples.

Module Completion Outcome

Learners complete this module with a rigorous, visual, and hands-on understanding of the mathematical foundations that enable logic, structure, optimization, and learning in AI systems—preparing them for statistical learning and model building in the next module.

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 2: Statistical Learning

Duration: Weeks 14–21 | Format: Conceptual + Applied + Python Integrated

Module Objective

To develop deep statistical reasoning and modeling intuition in learners. This module empowers students to use data for inference, testing, prediction, and model evaluation, laying the groundwork for classical and modern machine learning techniques.

Week 14–15: Descriptive and Inferential Statistics

Key Question: What is data telling us? What conclusions can we trust?

Topics:

- Descriptive statistics: Mean, median, mode, variance, standard deviation
- Data visualization: Histograms, boxplots, KDE plots
- Population vs sample, sampling techniques
- Confidence intervals and margin of error

Python Coding Activities:

```
import pandas as pd
```

```
import seaborn as sns
```

```
import matplotlib.pyplot as plt
```

```
data = pd.read_csv("https://raw.githubusercontent.com/mwaskom/seaborn-data/master/tips.csv")
```

```
print(data.describe())
```

```
sns.boxplot(x=data["total_bill"])
```

```
plt.title("Boxplot of Total Bill")
```

```
plt.show()
```

Week 16: Random Variables, Distributions, and Sampling

Key Question: How does randomness behave at scale?

Topics:

- Discrete vs continuous random variables
- Probability distributions: Normal, Binomial, Poisson

- Central Limit Theorem, Law of Large Numbers
- Bias, variance, and sample design

Python Coding Activities:

```
import numpy as np
```

```
import matplotlib.pyplot as plt
```

```
samples = [np.mean(np.random.normal(100, 15, 50)) for _ in range(1000)]
```

```
plt.hist(samples, bins=30, edgecolor='black')
```

```
plt.title("Sampling Distribution of the Mean")
```

```
plt.xlabel("Sample Mean")
```

```
plt.ylabel("Frequency")
```

```
plt.show()
```

Week 17: Hypothesis Testing & Confidence Intervals

Key Question: Can we test assumptions using data?

Topics:

- Null and alternative hypothesis
- p-values, significance levels
- Type I & II errors
- Z-test, T-test

Python Coding Activities:

```
from scipy import stats
```

```
sample1 = np.random.normal(60, 10, 30)
```

```
sample2 = np.random.normal(65, 10, 30)
```

```
t_stat, p_value = stats.ttest_ind(sample1, sample2)
```

```
print(f"T-statistic: {t_stat}, P-value: {p_value}")
```

Week 18: Correlation vs Causation

Key Question: Are patterns always meaningful?

Topics:

- Covariance and correlation coefficients
- Spurious correlations and confounding
- Introduction to causal inference

Python Coding Activities:

```
correlation = data[["total_bill", "tip"]].corr()
print("Correlation Matrix:\n", correlation)
sns.heatmap(correlation, annot=True)
plt.title("Correlation Heatmap")
plt.show()
```

Week 19–20: Linear and Logistic Regression

Key Question: How do we predict and classify using statistical models?

Topics:

- Linear regression (simple and multiple)
- Logistic regression for binary classification
- Model diagnostics and interpretability
- R-squared, confusion matrix, ROC curve (intro)

Python Coding Activities:

```
from sklearn.linear_model import LinearRegression
from sklearn.model_selection import train_test_split

X = data[["total_bill"]]
y = data["tip"]
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2)
model = LinearRegression()
model.fit(X_train, y_train)
print(f"Model Coefficients: {model.coef_}, Intercept: {model.intercept_}")
```

Week 21: Bias-Variance Tradeoff

Key Question: What makes models overfit or underfit?

Topics:

- Underfitting vs Overfitting
- Bias-variance decomposition
- Introduction to regularization (L1, L2 overview)

Python Coding Activities:

```
from sklearn.preprocessing import PolynomialFeatures  
from sklearn.metrics import mean_squared_error
```

```
poly = PolynomialFeatures(degree=5)  
X_poly = poly.fit_transform(X)  
model.fit(X_poly, y)  
y_pred = model.predict(X_poly)  
print("MSE:", mean_squared_error(y, y_pred))
```

◆ Module Outcome

By the end of this module, learners will:

- Build and interpret statistical models
 - Use Python to simulate randomness, test hypotheses, and visualize uncertainty
 - Understand and evaluate basic predictive models
 - Be ready to transition into formal machine learning algorithms in Module 3
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 3: Classical Machine Learning

Duration: Weeks 22–29 | Format: Num-Sense Driven, Intuitive, and Applied for the Layman

Module Objective

To demystify machine learning by grounding every concept in intuitive numerical sense and real-world metaphors. Learners will be guided through the core ML algorithms with minimal math intimidation and maximum intuition, enabling them to build and evaluate models confidently using Python.

Week 22: Supervised vs Unsupervised Learning

Key Question: When does a machine need labels to learn?

Topics:

- Supervised learning: Regression and classification
- Unsupervised learning: Clustering, dimensionality reduction
- Types of ML tasks: Predicting vs grouping
- Feature-label structure explained through real-life analogies (e.g., restaurant ratings, clothing sizes)

Activities:

- Sort fruits using labeled vs unlabeled features
 - Python walkthrough with sklearn's datasets (Iris, Breast Cancer)
-

Week 23: k-Nearest Neighbors (k-NN)

Key Question: What does it mean to learn by similarity?

Topics:

- Instance-based learning and lazy learning
- Distance metrics (Euclidean, Manhattan)
- Choosing 'k' and the bias-variance impact
- Visual intuition with 2D decision boundaries

Python Code Sample:

```
from sklearn.neighbors import KNeighborsClassifier
from sklearn.datasets import load_iris
from sklearn.model_selection import train_test_split
```

```
iris = load_iris()
X_train, X_test, y_train, y_test = train_test_split(iris.data, iris.target, test_size=0.3)
model = KNeighborsClassifier(n_neighbors=3)
model.fit(X_train, y_train)
print("Accuracy:", model.score(X_test, y_test))
```

Week 24: Decision Trees & Rule-Based Learning

Key Question: How can machines learn human-like decision rules?

Topics:

- Decision trees as flowcharts
- Gini impurity and entropy (explained through coin flips and card sorting)
- Overfitting control: pruning and depth limits
- Visualizing and interpreting tree logic

Python Code Sample:

```
from sklearn.tree import DecisionTreeClassifier, plot_tree
import matplotlib.pyplot as plt

model = DecisionTreeClassifier(max_depth=3)
model.fit(X_train, y_train)
plot_tree(model, feature_names=iris.feature_names, class_names=iris.target_names, filled=True)
plt.show()
```

Week 25: Support Vector Machines (SVM)

Key Question: What's the best dividing line between two groups?

Topics:

- SVM intuition using margins and support vectors
- Kernel trick explained via shadows and curved spaces
- When to use linear vs non-linear SVMs

Python Code Sample:

```
from sklearn.svm import SVC
```

```
model = SVC(kernel='linear')
model.fit(X_train, y_train)
print("SVM Accuracy:", model.score(X_test, y_test))
```

Week 26: Clustering Algorithms (K-Means, Hierarchical)

Key Question: How does a machine find natural groups without guidance?

Topics:

- K-means clustering: centroids, inertia, elbow method
- Hierarchical clustering: dendrograms and linkage types
- Use cases: customer segmentation, topic grouping

Python Code Sample:

```
from sklearn.cluster import KMeans
import seaborn as sns

kmeans = KMeans(n_clusters=3)
kmeans.fit(iris.data)
sns.scatterplot(x=iris.data[:, 0], y=iris.data[:, 1], hue=kmeans.labels_)
plt.title("K-Means Clustering")
plt.show()
```

Week 27: Ensemble Models (Bagging, Random Forest, Boosting)

Key Question: How do many weak learners become a strong one?

Topics:

- Bagging with bootstraps (Random Forest intuition)
- Boosting (Gradient Boosting, AdaBoost concepts)
- Bias-variance improvement via ensembles
- Real-world metaphor: committee decisions

Python Code Sample:

```
from sklearn.ensemble import RandomForestClassifier
```

```
rf = RandomForestClassifier(n_estimators=100)

rf.fit(X_train, y_train)

print("Random Forest Accuracy:", rf.score(X_test, y_test))
```

Week 28: Cross-validation and Evaluation Metrics

Key Question: How do we trust our models?

Topics:

- Train-test split vs K-fold cross-validation
- Evaluation metrics: accuracy, precision, recall, F1-score
- ROC curve and AUC

Python Code Sample:

```
from sklearn.model_selection import cross_val_score

scores = cross_val_score(model, iris.data, iris.target, cv=5)

print("Cross-validated Scores:", scores)
```

Week 29: Fairness and Explainability in ML

Key Question: Can models be trusted and explained?

Topics:

- Introduction to fairness in ML (bias detection)
- Black box vs white box models
- Tools for explainability: SHAP, LIME (overview only)

Python Conceptual Example: Discuss how to measure fairness via subgroup accuracy and perform feature importance analysis using decision trees.

Module Outcome

By the end of this module, learners will:

- Understand core classical ML algorithms using visual, numeric, and intuitive tools
 - Implement and compare ML models using real datasets in Python
 - Grasp key issues of model trustworthiness, validation, and interpretability
 - Be prepared to tackle real-world AI problems using foundational ML knowledge
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 4: Deep Learning

Duration: Weeks 30–36 | Format: Intuitive + Visual + Code-Oriented + Layered Understanding (TensorFlow, PyTorch, and Keras)

Module Objective

To build a clear and layered understanding of Deep Learning — starting from the intuition of neurons and moving to practical neural networks for vision, sequences, and feature learning. Learners gain hands-on experience using **TensorFlow**, **PyTorch**, and **Keras**, while learning the principles behind today's AI systems.

Week 30: Neural Networks 101 (Perceptron to MLP)

Key Question: How does a machine learn through layers of neurons?

Topics:

- Biological vs artificial neurons
- Perceptron and multi-layer perceptrons (MLP)
- Activation functions: ReLU, Sigmoid, Tanh
- Forward pass and architecture design

TensorFlow/Keras Example:

```
from tensorflow.keras.models import Sequential
```

```
from tensorflow.keras.layers import Dense
```

```
model = Sequential([
    Dense(10, activation='relu', input_shape=(4,)),
    Dense(3, activation='softmax')
])
model.summary()
```

PyTorch Example:

```
import torch.nn as nn
```

```
class MLP(nn.Module):
    def __init__(self):
        super(MLP, self).__init__()
```

```
self.fc1 = nn.Linear(4, 10)

self.relu = nn.ReLU()

self.fc2 = nn.Linear(10, 3)

self.softmax = nn.Softmax(dim=1)

def forward(self, x):

    return self.softmax(self.fc2(self.relu(self.fc1(x))))
```

Week 31: Backpropagation and Optimization

Key Question: How does a network learn from its mistakes?

Topics:

- Loss functions (MSE, cross-entropy)
- Gradient descent recap
- Backpropagation: the chain of corrections
- Optimizers: SGD, Adam

TensorFlow/Keras Compilation:

```
model.compile(optimizer='adam',

              loss='sparse_categorical_crossentropy',

              metrics=['accuracy'])
```

PyTorch Training Loop (Simplified):

```
import torch.optim as optim

model = MLP()

optimizer = optim.Adam(model.parameters(), lr=0.001)

loss_fn = nn.CrossEntropyLoss()
```

Week 32: Convolutional Neural Networks (CNNs)

Key Question: How do machines see patterns in images?

Topics:

- Convolutions and filters
- Pooling layers and feature maps
- CNN architecture for image recognition

TensorFlow/Keras CNN:

```
from tensorflow.keras.layers import Conv2D, MaxPooling2D, Flatten
```

```
model = Sequential([
    Conv2D(32, (3,3), activation='relu', input_shape=(28,28,1)),
    MaxPooling2D(2,2),
    Flatten(),
    Dense(64, activation='relu'),
    Dense(10, activation='softmax')
])
```

PyTorch CNN (Simplified):

```
class CNN(nn.Module):
    def __init__(self):
        super(CNN, self).__init__()
        self.conv1 = nn.Conv2d(1, 32, kernel_size=3)
        self.pool = nn.MaxPool2d(2, 2)
        self.fc1 = nn.Linear(5408, 64)
        self.fc2 = nn.Linear(64, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = x.view(-1, 5408)
        x = F.relu(self.fc1(x))
        return self.fc2(x)
```

Week 33: Recurrent Neural Networks (RNNs) and LSTMs

Key Question: How do AI systems handle sequences and memory?

Topics:

- Sequential data and time dependencies
- RNN architecture and limitations (vanishing gradient)
- LSTM and GRU networks

- Applications in text, finance, speech

Keras LSTM Example:

```
from tensorflow.keras.layers import LSTM, Embedding
```

```
model = Sequential([
    Embedding(input_dim=10000, output_dim=32),
    LSTM(64),
    Dense(1, activation='sigmoid')
])
```

PyTorch LSTM (Simplified):

```
class LSTMModel(nn.Module):
    def __init__(self):
        super(LSTMModel, self).__init__()
        self.embedding = nn.Embedding(10000, 32)
        self.lstm = nn.LSTM(32, 64)
        self.fc = nn.Linear(64, 1)

    def forward(self, x):
        x = self.embedding(x)
        output, (h_n, c_n) = self.lstm(x)
        return torch.sigmoid(self.fc(h_n[-1]))
```

Week 34: Transfer Learning

Key Question: How can models learn from prior knowledge?

Topics:

- Pretrained models (ResNet, VGG, MobileNet)
- Feature extraction vs fine-tuning
- Benefits for small data problems

Keras Transfer Learning:

```
from tensorflow.keras.applications import MobileNetV2
```

```
base_model = MobileNetV2(weights='imagenet', include_top=False, input_shape=(224, 224, 3))
```

```
base_model.trainable = False
```

PyTorch Transfer Learning:

```
from torchvision import models
```

```
model = models.resnet18(pretrained=True)
```

```
for param in model.parameters():
```

```
    param.requires_grad = False
```

```
model.fc = nn.Linear(model.fc.in_features, 10)
```

Week 35: TensorFlow vs PyTorch Primer

Key Question: What are the tools behind modern Deep Learning?

Topics:

- Overview of TensorFlow, Keras, and PyTorch
- Graph vs eager execution
- Syntax and architecture comparison
- When to use what and why

Activity: Build the same MLP classifier using all three libraries

Week 36: Capstone Project — “Build Your First Deep Model”

Challenge: Each learner builds a small deep learning application using an open dataset.

- Example choices: digit classification, fake news detection, emotion recognition
- Must include: data preprocessing, model architecture, training, and evaluation

Submission: GitHub link + demo video + short project report

Module Outcome

By the end of this module, learners will:

- Understand how deep networks operate and optimize
- Build CNN and LSTM models from scratch
- Apply transfer learning with pre-trained networks
- Be confident using both **TensorFlow/Keras** and **PyTorch** frameworks

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 5: Natural Language Processing (NLP)

Duration: Weeks 37–40 | Format: Intuitive + Application-Driven + Python + Transformer-Aware + Reinforcement Learning for NLP (RL)

Module Objective

To provide a structured and intuitive understanding of how machines process human language. This module covers traditional, neural, and reinforcement learning approaches to NLP — from tokenization to Transformers — with practical coding experience using spaCy, NLTK, Hugging Face Transformers, and RL-based methods for conversational systems.

Week 37: Text Preprocessing and Tokenization

Key Question: How do machines begin to understand raw language?

Topics:

- Text normalization: lowercasing, stemming, lemmatization
- Tokenization: word, sentence, subword
- Removing stopwords, punctuation, and noise
- Word frequency and Bag-of-Words models

Python (NLTK & spaCy):

```
import nltk

from nltk.tokenize import word_tokenize

from nltk.corpus import stopwords

nltk.download('punkt'); nltk.download('stopwords')

text = "AI is transforming the future."

tokens = word_tokenize(text)

filtered = [w for w in tokens if w.lower() not in stopwords.words('english')]

print(filtered)
```

spaCy Version:

```
import spacy

nlp = spacy.load("en_core_web_sm")
```

```
doc = nlp("AI is transforming the future.")  
print([token.lemma_ for token in doc if not token.is_stop])
```

Week 38: Word Embeddings and Semantic Representation

Key Question: How can we represent meaning mathematically?

Topics:

- Vector space models
- Word2Vec, GloVe: context-based meaning
- Cosine similarity and analogies
- Limitations of one-vector-per-word methods

Python (Gensim Example):

```
from gensim.models import Word2Vec  
  
sentences = [["ai", "is", "changing", "world"], ["machine", "learning", "is", "future"]]  
model = Word2Vec(sentences, vector_size=50, window=2, min_count=1, workers=2)  
print(model.wv.most_similar("ai"))
```

Application: Find semantic similarity between user queries

Week 39: Transformers and Modern NLP

Key Question: How does modern AI understand language contextually?

Topics:

- Sequence modeling limitations in RNNs
- Introduction to Attention and Transformers
- BERT, GPT (overview)
- Hugging Face Transformers for text classification

Python (Hugging Face Example):

```
from transformers import pipeline  
  
classifier = pipeline("sentiment-analysis")  
print(classifier("I love learning AI!"))
```

Mini Project: Sentiment analysis or topic classification using Transformers

Week 40: RL in NLP + Applications + Ethics

Key Question: How do conversational agents improve through feedback and exploration?

Topics:

- Overview of Reinforcement Learning in NLP
- RL for dialogue systems and conversational agents
- Reward-based learning for policy optimization
- Ethical concerns: hallucinations, bias amplification, misinformation

Python Preview (RL in NLP - Conceptual Simulation):

```
# Simulating reward-based improvement in chatbot responses (conceptual)
```

```
responses = {"hi": 1, "hello": 2, "howdy": 0.5}
```

```
feedback = {"hi": -1, "hello": 1, "howdy": 2}
```

```
updated_rewards = {k: responses[k] + 0.1 * feedback[k] for k in responses}
```

```
print(updated_rewards)
```

Activity: Group discussion: "Should RL-fine-tuned models speak without supervision?"

Project Challenge: Build a mini NLP pipeline: clean → vectorize → classify → improve using feedback signal

 Module Outcome

By the end of this module, learners will:

- Understand how machines represent and manipulate language
 - Apply traditional, transformer-based, and reinforcement learning models to NLP tasks
 - Use Python libraries such as NLTK, spaCy, Hugging Face, and RL simulations
 - Build ethical awareness around language technologies
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 6: Reinforcement Learning (RL)

Duration: Weeks 41–43 | Format: Interactive + Goal-Driven + Real-World Applications

Module Objective

To provide learners with a strong foundation in **Reinforcement Learning (RL)** as a paradigm of learning through interaction, reward, and feedback. This module expands from basic RL theory to applications in robotics, games, recommendation systems, and conversational AI.

Week 41: Foundations of Reinforcement Learning

Key Question: How do agents learn by interacting with their environment?

Topics:

- Core components: Agent, Environment, State, Action, Reward
- The RL loop and feedback cycle
- Markov Decision Processes (MDPs)
- Exploration vs Exploitation
- Return, Discount Factor, and Value Functions

Python Simulation (Q-table update):

```
Q = {"state1": {"a1": 1, "a2": 0.5}}  
reward = 1.0  
alpha = 0.1  
Q["state1"]["a1"] += alpha * (reward - Q["state1"]["a1"])  
print(Q)
```

Application Cases: Grid-world navigation, elevator scheduling

Week 42: RL Algorithms and Applications

Key Question: How do we implement policies and optimize them?

Topics:

- Policy-based vs Value-based methods
- Q-learning and SARSA
- Policy Gradient Methods
- Deep Q Networks (DQN)

- Multi-armed bandit problems

Hands-on Projects:

- Build a grid-based RL agent
- Reward-maximizing recommendation policy (simple bandit)

Python (Q-learning):

```
import numpy as np
```

```
Q = np.zeros((5, 2)) # 5 states, 2 actions
```

```
alpha, gamma = 0.1, 0.9
```

```
state, action, reward, next_state = 0, 1, 10, 2
```

```
Q[state, action] += alpha * (reward + gamma * np.max(Q[next_state]) - Q[state, action])
```

Application Examples:

- Game AI (e.g., Tic Tac Toe, CartPole)
- RL for energy management, traffic control

Week 43: Safe, Aligned, and Ethical RL

Key Question: What are the risks of letting machines learn on their own?

Topics:

- Reward Hacking and unintended behaviors
- Exploration risk and catastrophic forgetting
- Reinforcement Learning from Human Feedback (RLHF)
- Guardrails for safe deployment (especially in robotics and LLMs)

Group Discussion:

- “Designing RL agents that don’t cheat the reward system”

Ethical Case Studies:

- OpenAI RLHF in ChatGPT
- DeepMind’s AlphaGo and AlphaZero
- Sim2Real Transfer in Robotics

Module Outcome

By the end of this module, learners will:

- Understand foundational RL theory and practical training loops

- Implement tabular and simple neural RL agents in Python
 - Apply RL to gaming, recommendation, robotics, and conversational systems
 - Recognize safety, interpretability, and ethical issues in autonomous learning systems
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 7: Recommender Systems

Duration: Weeks 44–46 | Format: Intuition + System Thinking + Algorithmic Implementation



Module Objective

To equip learners with an understanding of how modern recommender systems personalize content across domains such as e-commerce, entertainment, education, and social media. This module blends mathematical intuition with implementation using Python.



Week 44: Foundations of Recommendation

Key Question: How do platforms decide what you might like next?

Topics:

- Types of recommender systems: Content-based, Collaborative, Hybrid
- User-item matrix intuition
- Cold start and sparsity problems
- Applications: Netflix, Amazon, Spotify, LinkedIn

Python Concept (User-Item Matrix):

```
import pandas as pd

ratings = pd.DataFrame({
    'User1': [5, 3, 0],
    'User2': [4, 0, 2],
    'User3': [0, 4, 4]
}, index=['ItemA', 'ItemB', 'ItemC'])

print(ratings)
```

Discussion: What does “personalized” mean in different contexts?



Week 45: Collaborative & Content-Based Filtering

Key Question: How can machines learn from user preferences and item characteristics?

Topics:

- User-based vs Item-based collaborative filtering
- Cosine similarity and neighborhood methods
- TF-IDF for content profiling
- Hybrid recommendation strategies

Python (Cosine Similarity):

```
from sklearn.metrics.pairwise import cosine_similarity
similarity_matrix = cosine_similarity(ratings.fillna(0).T)
print(similarity_matrix)
```

Mini Project: Movie recommender using user ratings + genre tags

Week 46: Matrix Factorization and Beyond

Key Question: How do we learn latent preferences?

Topics:

- Matrix Factorization (SVD, ALS)
- Latent factors and dimensionality reduction
- Introduction to Neural Collaborative Filtering (NCF)
- Evaluation metrics: RMSE, Precision@k, Recall@k

Python (SVD using Surprise):

```
from surprise import SVD, Dataset, Reader
from surprise.model_selection import cross_validate

data = Dataset.load_builtin('ml-100k')
model = SVD()
cross_validate(model, data, measures=['RMSE', 'MAE'], cv=3, verbose=True)
```

Project Challenge: Build and evaluate a recommendation engine using real or synthetic data.

◆ **Module Outcome**

By the end of this module, learners will:

- Understand the core logic and types of recommendation systems

- Build collaborative and content-based recommenders using Python
 - Learn matrix factorization for performance optimization
 - Explore evaluation and interpretability of recommender systems
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 8: Sentiment and Behaviour Analytics

Duration: Weeks 47–49 | Format: Human-Centered + Computational + Predictive + Visual

Module Objective

To develop the ability to analyze and interpret user behavior and sentiment across digital touchpoints. This module blends emotional intelligence with computational models to extract meaning from human actions, preferences, and expressions.

Week 47: Understanding Sentiment — Signals and Systems

Key Question: What do users reveal through their words and expressions?

Topics:

- Introduction to sentiment: polarity, intensity, subjectivity
- Sentiment lexicons (VADER, TextBlob, SentiWordNet)
- Preprocessing for emotion capture
- Emotions vs opinions vs attitudes

Python Example (VADER):

```
from vaderSentiment.vaderSentiment import SentimentIntensityAnalyzer  
  
analyzer = SentimentIntensityAnalyzer()  
  
print(analyzer.polarity_scores("This product is absolutely fantastic!"))
```

Mini Activity: Compare sentiment across product reviews and social media tweets

Week 48: Behavioural Signals — What People Do

Key Question: How can we quantify user behavior?

Topics:

- Behavioral metrics: clickstream, dwell time, bounce rate, churn rate
- User journey mapping and funnel analysis
- Web/app log mining
- Feature engineering for behavioral traits

Python (Log Parsing Example):

```
import pandas as pd  
  
logs = pd.read_csv("user_logs.csv")
```

```
logs['duration'] = logs['logout_time'] - logs['login_time']  
engaged = logs[logs['duration'] > 5]  
print(engaged.head())
```

Case Study: Predicting drop-off in a learning app

Week 49: Integrating Emotion and Action — Predictive Models

Key Question: Can we predict outcomes based on how people feel and behave?

Topics:

- Combining sentiment + behavior for churn prediction, engagement scoring
- Clustering users by behavioral and emotional similarity
- Real-time dashboards for behavioral insights
- Ethics and explainability in profiling and personalization

Capstone Project Ideas:

- Predicting customer churn using behavioral and emotional signals
- Building a dashboard for real-time user sentiment and behavior
- User clustering for personalization using hybrid features

Optional Tools: Python (scikit-learn, Streamlit, Plotly), RapidMiner, Power BI

Module Outcome

By the end of this module, learners will:

- Detect, classify, and visualize sentiment in text
 - Model and interpret user behavior from digital interactions
 - Build predictive pipelines using emotion + behavior features
 - Understand ethical boundaries in emotional and behavioral analytics
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 9: Time Series Analysis

Duration: Weeks 50–51 | Format: Temporal Thinking + Forecasting + Decomposition + Python Implementation

Module Objective

To equip learners with the tools and intuition required to analyze and forecast data that changes over time. From trend discovery to anomaly detection, this module reveals the temporal dimension behind data and empowers decision-making through time-aware models.

Week 50: Time Series Foundations and Exploration

Key Question: What makes time-based data special?

Topics:

- Characteristics of time series: trend, seasonality, noise
- Stationarity and differencing
- Lag plots and autocorrelation (ACF, PACF)
- Time-based feature engineering

Python (Exploration):

```
import pandas as pd
```

```
import matplotlib.pyplot as plt
```

```
from statsmodels.graphics.tsaplots import plot_acf
```

```
data = pd.read_csv('air_passengers.csv', parse_dates=['Month'], index_col='Month')
```

```
data.plot()
```

```
plot_acf(data['Passengers'])
```

```
plt.show()
```

Mini Task: Identify seasonal trends and lags in monthly sales or weather data

Week 51: Forecasting Techniques and Anomaly Detection

Key Question: Can we predict what happens next?

Topics:

- Moving averages and exponential smoothing
- ARIMA and SARIMA models

- Prophet for business forecasting
- Anomaly detection in time series (rolling statistics, IQR, Z-score)

Python (Forecasting with Prophet):

```
from prophet import Prophet
```

```
df = data.rename(columns={"Passengers": "y"})
```

```
df['ds'] = df.index
```

```
model = Prophet()
```

```
model.fit(df)
```

```
future = model.make_future_dataframe(periods=12, freq='M')
```

```
forecast = model.predict(future)
```

```
model.plot(forecast)
```

Project Challenge:

- Forecast future demand for a product or service
- Detect anomalies in server uptime or stock prices



Module Outcome

By the end of this module, learners will:

- Understand the unique features of time series data
 - Visualize and decompose temporal trends
 - Build and evaluate forecasting models
 - Apply time-aware anomaly detection in real-world scenarios
-

Certified Enterprise Transformation Analyst (CETA) Foundation Program

Module 10: Applied Data Science Projects & Portfolio

Duration: Week 52 | Format: Experiential + Integrative + Reflective + Showcase-Based

Module Objective

To consolidate the learning journey by engaging learners in real-world, end-to-end data science projects. This module enables learners to demonstrate their skills through applied case studies and a curated portfolio aligned with industry readiness.

Week 52: Capstone Project, Presentation, and Portfolio Development

Part A: Capstone Project Execution

Focus: Solve a realistic, complex problem using the full pipeline:

- **Problem Framing:** Business or research problem to ML objective
- **Data Collection:** Use open datasets or synthetically generate structured data
- **Preprocessing:** Cleaning, EDA, feature engineering
- **Modeling:** Choose ML/Deep learning techniques appropriately
- **Evaluation:** Metrics + interpretation
- **Visualization & Reporting:** Dashboards, notebooks, presentations

Capstone Ideas:

- Predicting student dropout in online learning using behavior + sentiment
 - Retail sales forecasting using time series + holidays
 - Chatbot sentiment improvement using RLHF simulation
 - Building a movie recommender from scratch using hybrid methods
-

Part B: Portfolio & GitHub Curation

Goal: Publish and document project outcomes with professionalism

Tasks:

- Document 2–3 notebooks with clean code, visuals, explanations
 - Host projects on GitHub with README.md, dependencies, and usage
 - Create an infographic résumé or portfolio site (optional: Streamlit app)
 - Record a 3-minute video explanation of one key project
-

Part C: Peer Review & Viva-Style Evaluation

- Present project to mentor panel or peer group
 - Respond to questions on design decisions, assumptions, and improvements
 - Receive structured feedback for real-world enhancement
-

Module Outcome

By the end of this final module, learners will:

- Demonstrate competency across data science, ML, DL, NLP, and deployment
 - Develop a showcase-ready project for internships or job applications
 - Reflect on their journey from fundamentals to enterprise transformation
 - Be certified as **CETA Foundation Graduates** with a complete public portfolio
-

