



## Dr. B. C. Roy Engineering College, Durgapur

### Department of CSE(AIML)

Field	Details
Course Name	AI in Finance and Economics
Course Code	AM-802
Semester	8
Course Category	Professional Elective Courses
Credits	3
Hours per Week	3L:0T:0P

### 1. Prerequisites

- Proficiency in Python programming (including data-analysis libraries such as pandas, NumPy, and scikit-learn)
- Fundamental knowledge of probability, statistics, and linear algebra (including concepts like distributions, hypothesis testing, matrix operations, and calculus basics)
- Introductory understanding of finance/economics fundamentals (e.g., financial statements, basic time-series concepts, and key financial ratios)

### 2. Course Learning Objectives

- Equip students with a solid conceptual and practical foundation in AI/ML techniques specifically tailored to financial and economic data, enabling them to design, implement, and evaluate predictive models across a variety of finance-related problems.
- Develop students' ability to acquire, clean, and engineer high-quality financial and macro-economic datasets using modern programming tools and AI-assisted workflows, emphasizing reproducibility and scalability.

- Cultivate critical understanding of how core AI models (regression, classification, clustering, ensemble methods, and basic time-series approaches) can be applied to key financial functions such as valuation, risk assessment, forecasting, and algorithmic trading.
- Foster awareness of the ethical, regulatory, and governance challenges associated with deploying AI in finance and economic policy, preparing students to propose responsible, transparent, and compliant AI solutions.
- Prepare students to synthesize AI-driven insights for strategic decision-making in finance and macro-economics, including portfolio optimization, credit scoring, policy simulation, and emerging fintech innovations.

### 3. Teaching Methodology

- Lectures and Presentations
- Interactive Discussions and Case Studies
- Lab Sessions
- Guest Lectures

### 4. Evaluation System

Activities	Class Test Full Marks	Assignment Full Marks	Attendance Full Marks	Total Marks
CIA-1	25	10	05	40
CIA-2	25	10	05	40
End Semester Examination (ESE)	-	-	-	60
Total				100 Marks

### 5. Course Modules

Module	Topics	Hours
1	<b>Foundations of AI &amp; Machine Learning for Finance &amp; Economics</b> - Training and testing workflow - Feature scaling and data normalization - Cross-validation pipelines - Gradient descent basics (batch, mini-batch,	5

	<p>stochastic)</p> <ul style="list-style-type: none"> <li>- Learning paradigms: supervised vs unsupervised</li> <li>- Linear vs non-linear models (conceptual)</li> <li>- Regression vs classification</li> <li>- Cross-sectional vs sequential (time-series) models (conceptual)</li> <li>- Overview of AI applications in finance</li> <li>- Basic model evaluation metrics (accuracy, RMSE, MAE)</li> </ul>	
2	<p><b>Financial Data Foundations &amp; Feature Engineering</b></p> <ul style="list-style-type: none"> <li>- Types of financial data (cross-sectional, time-series)</li> <li>- Fundamentals of probability &amp; statistics for finance</li> <li>- Data extraction (APIs, web-scraping) and automated cleaning with AI tools</li> <li>- Handling missing values, outliers, and data imputation</li> <li>- Feature engineering techniques for financial datasets</li> <li>- Brownian motion (random walk) and mean-reverting processes - conceptual view</li> <li>- Role and significance of big data in economics</li> <li>- Practical data-preprocessing workflow in Python/R</li> </ul>	6
3	<p><b>Core AI Models for Finance</b></p> <ul style="list-style-type: none"> <li>- Linear regression and its assumptions</li> <li>- Logistic regression for binary outcomes</li> <li>- Decision trees and tree-based learning</li> <li>- Random forests (ensemble of trees)</li> <li>- Support Vector Machines (basic linear kernel)</li> <li>- k-means clustering for unsupervised grouping</li> <li>- Principal Component Analysis (introductory dimensionality reduction)</li> <li>- Model evaluation &amp; validation (confusion matrix, ROC, precision/recall)</li> <li>- Regularization basics (L1/L2) and bias-variance trade-off</li> </ul>	7
4	<p><b>Financial Applications &amp; Metrics</b></p> <ul style="list-style-type: none"> <li>- Core financial metrics (profitability, liquidity, risk ratios)</li> <li>- AI-enhanced financial statement analysis</li> <li>- Cash-flow forecasting using regression models</li> <li>- Working-capital management with predictive analytics</li> <li>- Capital budgeting decisions aided by AI (NPV, IRR forecasts)</li> </ul>	8

	<ul style="list-style-type: none"> <li>- Credit scoring and default risk modeling</li> <li>- Time-series forecasting basics (ARIMA, simple recurrent approaches)</li> <li>- Portfolio optimization fundamentals (mean-variance, risk budgeting)</li> <li>- Introduction to rule-based algorithmic trading systems</li> <li>- Fraud detection techniques using classification models</li> <li>- Risk assessment and stress-testing with machine-learning tools</li> </ul>	
5	<p><b>AI for Economic Policy, Development &amp; Macro-analysis</b></p> <ul style="list-style-type: none"> <li>- Predictive modelling of macro-indicators (GDP, inflation, unemployment)</li> <li>- AI-assisted macro-economic forecasting frameworks</li> <li>- Automated acquisition and processing of economic data for policy analysis</li> <li>- Scenario planning and policy simulation using ML models</li> <li>- Analysis of AI's impact on labour markets and employment trends</li> <li>- AI-driven trade-policy assessment and balance-of-payments forecasts</li> <li>- Measuring poverty reduction and SDG progress with AI tools</li> <li>- Ethical considerations in macro-AI applications</li> </ul>	8
6	<p><b>Governance, Ethics, and Future Directions in AI Finance</b></p> <ul style="list-style-type: none"> <li>- AI-first finance and market-efficiency concepts</li> <li>- FinTech innovation landscape and learning components</li> <li>- AI in wealth and asset management</li> <li>- AI for risk management, compliance and regulatory reporting</li> <li>- Consumer-finance AI applications (credit scoring, robo-advisors)</li> <li>- AI-driven investment-fund strategies</li> <li>- Financial-stability and systemic-risk monitoring with AI</li> <li>- Governance frameworks, regulatory considerations and oversight</li> <li>- Data ethics, privacy, and responsible AI practices</li> <li>- Future trends: explainable AI, model interpretability, and emerging technologies</li> </ul>	8

## 6. References

### Textbooks:

1. Yves Hilpisch - "Artificial Intelligence in Finance: A Python-Based Guide" (O'Reilly, 2020)

### Reference Books:

1. Yves Hilpisch - "Artificial Intelligence in Finance & Algorithmic Trading"
2. Stephen Hansen - "Machine Learning for Economics and Policy"

## 7. Course Outcomes

ID	Statement	Action Verb	Knowledge Level
AM-802.1	Course Outcome 1	Understand	Understand
AM-802.2	Course Outcome 2	Understand	Understand
AM-802.3	Course Outcome 3	Understand	Understand
AM-802.4	Course Outcome 4	Understand	Understand
AM-802.5	Course Outcome 5	Understand	Understand
AM-802.6	Course Outcome 6	Understand	Understand

## 8. CO-PO Mapping

CO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1	3	2	1	1	1	1	1	1	1	1	1	2
CO2	3	2	1	1	2	1	1	1	1	1	1	2
CO3	2	2	3	2	2	1	2	1	1	1	1	2
CO4	2	2	2	3	2	1	2	1	1	1	1	2
CO5	2	2	2	2	3	1	2	1	1	1	1	2
CO6	2	2	1	1	1	1	1	1	2	1	1	3

## 9. CO-PSO Mapping

CO	PSO1	PSO2	PSO3
CO1	3	2	1
CO2	2	3	1
CO3	1	2	3
CO4	2	2	2
CO5	3	1	2
CO6	1	3	2



## Dr. B. C. Roy Engineering College, Durgapur

### Department of CSE(AIML)

Field	Details
Course Name	AI for Aquaculture
Course Code	AM-803
Semester	8
Course Category	Professional Elective Courses
Credits	3
Hours per Week	3L:0T:0P

### 1. Prerequisites

- Proficiency in Python programming (including libraries such as NumPy, pandas, and matplotlib)
- Fundamental knowledge of statistics and linear algebra (e.g., probability, descriptive statistics, matrix operations)
- Introductory understanding of machine learning concepts and basic data-science workflow (data cleaning, model training, evaluation)

### 2. Course Learning Objectives

- This course introduces students to fundamental concepts and applications of the subject
- Students will learn theoretical foundations and practical skills relevant to the field

### 3. Teaching Methodology

- Lectures and Presentations

- Interactive Discussions and Case Studies
- Lab Sessions
- Guest Lectures

#### 4. Evaluation System

Activities	Class Test Full Marks	Assignment Full Marks	Attendance Full Marks	Total Marks
CIA-1	25	10	05	40
CIA-2	25	10	05	40
End Semester Examination (ESE)	-	-	-	60
Total				100 Marks

#### 5. Course Modules

Module	Topics	Hours
1	<p><b>Foundations of AI in Aquaculture &amp; Fisheries Management</b></p> <ul style="list-style-type: none"> <li>- Introduction to AI in Aquaculture and Fisheries Management <ul style="list-style-type: none"> <li>* Definition, significance and brief evolution of AI</li> <li>* AI lifecycle basics (data -&gt; model -&gt; deployment -&gt; monitoring)</li> </ul> </li> <li>- Overview of Aquaculture Production Systems <ul style="list-style-type: none"> <li>* Fresh-water, brackish, marine; ponds, cages, recirculating systems</li> </ul> </li> <li>- Core Challenges in Aquaculture that AI can address <ul style="list-style-type: none"> <li>* Disease detection, feed optimisation, yield forecasting, environmental monitoring</li> </ul> </li> <li>- Types of AI Solutions <ul style="list-style-type: none"> <li>* Rule-based systems, classic Machine Learning, introductory Deep Learning</li> </ul> </li> <li>- Foundations of AI Ethics &amp; Sustainability <ul style="list-style-type: none"> <li>* Data privacy, bias awareness, basic sustainability metrics (resource use, CO<sub>2</sub> footprint)</li> </ul> </li> <li>- Importance of Data Quality &amp; Documentation <ul style="list-style-type: none"> <li>* Data provenance, versioning, simple data-sheet standards</li> </ul> </li> </ul>	5
2	<b>Sensor Technologies, IoT &amp; Data Acquisition for</b>	6

	<p><b>Smart Aquaculture</b></p> <ul style="list-style-type: none"> <li>- Core Sensor Suite for Aquaculture <ul style="list-style-type: none"> <li>* Temperature, pH, dissolved oxygen, ammonia, salinity, turbidity</li> </ul> </li> <li>- Underwater Imaging &amp; Acoustic Sensing <ul style="list-style-type: none"> <li>* Low-cost cameras, sonar basics, UAV/ASV platforms</li> </ul> </li> <li>- IoT Architecture for Aquaculture <ul style="list-style-type: none"> <li>* Edge node concepts, cloud-based storage, simple data-stream pipelines</li> </ul> </li> <li>- Data Modalities &amp; Pre-processing Fundamentals <ul style="list-style-type: none"> <li>* Time-series sensor streams, image/video frames, basic cleaning &amp; labeling</li> </ul> </li> <li>- Introductory Networking &amp; Security for IoT <ul style="list-style-type: none"> <li>* Wi-Fi/LTE basics, authentication, data-in-transit protection</li> </ul> </li> <li>- Practical Activity (1 hr) <ul style="list-style-type: none"> <li>* Simulate a water-quality data stream in Python, visualise trends and perform basic cleaning</li> </ul> </li> </ul>	
3	<p><b>Machine Learning &amp; Deep Learning Fundamentals for Aquaculture Data</b></p> <ul style="list-style-type: none"> <li>- Machine-Learning Foundations <ul style="list-style-type: none"> <li>* Supervised vs unsupervised learning, regression, classification, clustering</li> </ul> </li> <li>- Feature Engineering for Sensor &amp; Image Data <ul style="list-style-type: none"> <li>* Statistical descriptors, colour/texture features, data augmentation basics</li> </ul> </li> <li>- Introductory Neural Networks <ul style="list-style-type: none"> <li>* Simple feed-forward network, basic Convolutional Neural Network (CNN) concepts</li> </ul> </li> <li>- Model Evaluation &amp; Validation Basics <ul style="list-style-type: none"> <li>* Accuracy, precision/recall, F1-score, confusion matrix, k-fold cross-validation, over-fitting checks</li> </ul> </li> <li>- Transfer Learning for Small Datasets <ul style="list-style-type: none"> <li>* Re-using pretrained CNN weights, fine-tuning strategy</li> </ul> </li> <li>- Hands-On Labs (2 hrs) <ul style="list-style-type: none"> <li>* Build a fish-growth regression model (ML)</li> <li>* Apply clustering for disease-anomaly detection on sensor data</li> </ul> </li> </ul>	7
4	<p><b>Computer Vision for Fish Identification, Segmentation &amp; Behaviour Analysis</b></p> <ul style="list-style-type: none"> <li>- Fish Detection &amp; Localization <ul style="list-style-type: none"> <li>* YOLOv8 basics, handling turbidity, real-time inference considerations</li> </ul> </li> <li>- Multi-Object Tracking &amp; Motion Analysis <ul style="list-style-type: none"> <li>* Simple tracking-by-detection, optical-flow overview</li> </ul> </li> </ul>	8

	<ul style="list-style-type: none"> <li>- Image Segmentation for Morphometrics <ul style="list-style-type: none"> <li>* U-Net architecture, extracting shape, fin-ratio and body-index metrics</li> </ul> </li> <li>- Behaviour Monitoring &amp; Pose Estimation <ul style="list-style-type: none"> <li>* Intro to DeepLabCut-style pose extraction, stress-pattern identification</li> </ul> </li> <li>- Model Optimisation for Edge Deployment <ul style="list-style-type: none"> <li>* Quantisation, pruning, inference speed-up tips</li> </ul> </li> <li>- Practical Labs (2 hrs) <ul style="list-style-type: none"> <li>* Train a detection model on a small fish dataset</li> <li>* Apply segmentation to compute morphological indices</li> <li>* Implement a basic behaviour-tracking pipeline</li> </ul> </li> </ul>	
5	<p><b>Multimodal Fusion, Predictive Modelling &amp; Decision Support</b></p> <ul style="list-style-type: none"> <li>- Multimodal Data Fusion Strategies <ul style="list-style-type: none"> <li>* Early-fusion (feature-level) and late-fusion (decision-level) concepts for vision + acoustic + sensor streams</li> </ul> </li> <li>- Simple Time-Series Forecasting for Water-Quality <ul style="list-style-type: none"> <li>* Moving-average, exponential smoothing, introductory ARIMA (parameter-free demo)</li> </ul> </li> <li>- Feed Intake &amp; Biomass Prediction <ul style="list-style-type: none"> <li>* LSTM basics for growth trends, comparison with linear regression</li> </ul> </li> <li>- Rule-Based Feed-Scheduling Optimisation (instead of advanced RL) <ul style="list-style-type: none"> <li>* Decision tables, threshold-based control loops</li> </ul> </li> <li>- Decision-Support System Design <ul style="list-style-type: none"> <li>* Dashboard fundamentals, alert generation, what-if analysis</li> </ul> </li> <li>- Ensemble Techniques for Robustness &amp; Uncertainty <ul style="list-style-type: none"> <li>* Bagging, simple model averaging, basic confidence intervals</li> </ul> </li> <li>- Hands-On Projects (3 hrs) <ul style="list-style-type: none"> <li>* Feature-level fusion of image and sensor data for health classification</li> <li>* Build an LSTM-based growth forecast model</li> <li>* Implement a rule-based feed-scheduling prototype with a dashboard</li> </ul> </li> </ul>	8
6	<p><b>AI-Driven Automation, Edge Deployment, Explainability, Sustainability &amp; Capstone</b></p> <ul style="list-style-type: none"> <li>- AI-Driven Automation &amp; Robotics in Aquaculture <ul style="list-style-type: none"> <li>* Feeding robots, basic ROV/AUV concepts, vision-guided path planning</li> </ul> </li> <li>- Edge AI Deployment &amp; Cyber-Physical Integration <ul style="list-style-type: none"> <li>* Containerisation (Docker), model serving on</li> </ul> </li> </ul>	8

	Raspberry-Pi/Jetson, digital-twin overview (high-level) - Explainable AI for Aquaculture Decisions * Grad-CAM visual explanations, SHAP feature importance, bias-check checklist - Sustainability & Socio-Economic Impact * CO <sub>2</sub> -footprint estimation, biodiversity monitoring, community adoption considerations - Global Industry Case Studies * AquaByte (Norway), eFishery (Indonesia), Fisheyebox (India), AI4Fish (EU) - lessons learned - Capstone Project Framework * Problem definition, dataset collection, model development, edge-deployment plan, technical report & presentation guidelines - Assessment Overview * Theory (MCQs, short essays), practical labs, project reports, oral viva	
--	---	--

## 6. References

### Textbooks:

1. Balchen, J.G. (Ed), 1986. Automation and Data Processing in Aquaculture: Proceedings of the IFAC Symposium, Trondheim, Norway, 18-21 August 1986 (Ifac Symposia Series) 1stEdition
2. Artificial Intelligence in Fisheries Transformative Potentials and Challenges, Sylvester Chibueze Izah, Matthew Chidozie Ogwu, Adams Ovie Iyiola

## 7. Course Outcomes

ID	Statement	Action Verb	Knowledge Level
AM-803.1	Define the core AI concepts, lifecycle stages (data -> model -> deployment -> monitoring), major challenges in aquaculture, and basic ethical and sustainability considerations relevant to fisheries management.	Define	Remember
AM-803.2	Explain the architecture of IoT-based sensor networks for aquaculture, and demonstrate how to acquire, clean, and document multi-modal data streams (e.g.,	Explain	Understand

	temperature, pH, video) using Python tools.		
AM-803.3	Apply supervised and unsupervised machine-learning techniques--including regression, clustering, and transfer-learning CNNs--to build and evaluate predictive models for fish growth and disease-anomaly detection.	Apply	Apply
AM-803.4	Analyze the performance of computer-vision pipelines (YOLOv8 detection, U-Net segmentation, pose estimation) by comparing accuracy, inference speed, and robustness under varying water-clarity conditions.	Analyze	Analyze
AM-803.5	Evaluate multimodal fusion strategies and rule-based decision-support models by constructing early-fusion and late-fusion prototypes for health classification and feed-schedule optimisation, and reporting quantitative uncertainty metrics.	Evaluate	Evaluate
AM-803.6	Create a complete AI-driven automation solution for a smart-aquaculture scenario, including edge deployment (Docker on Raspberry Pi/Jetson), explainable-AI visualisations (Grad-CAM, SHAP), sustainability impact assessment, and a professional capstone report and presentation.	Create	Create

## 8. CO-PO Mapping

CO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1	3	2	1	1	1	2	3	3	1	1	1	2
CO2	3	2	2	2	3	1	1	1	1	2	1	2
CO3	3	3	3	3	3	2	2	1	1	2	2	3
CO4	3	3	2	3	3	2	2	1	1	2	2	3
CO5	3	3	3	3	3	2	2	1	1	2	2	3
CO6	3	3	3	3	3	3	3	3	3	3	3	3

## 9. CO-PSO Mapping

<b>CO</b>	<b>PSO1</b>	<b>PSO2</b>	<b>PSO3</b>
CO1	3	1	3
CO2	2	3	1
CO3	3	3	1
CO4	2	3	1
CO5	2	3	1
CO6	3	3	3



## Dr. B. C. Roy Engineering College, Durgapur

### Department of CSE(AIML)

Field	Details
Course Name	AI in Healthcare
Course Code	AM-804
Semester	8
Course Category	Professional Elective Courses
Credits	3
Hours per Week	3L:0T:0P

### 1. Prerequisites

- Proficiency in Python programming and common data-science libraries (e.g., NumPy, pandas, scikit-learn)
- Fundamental knowledge of probability, statistics, and linear algebra
- Introductory understanding of healthcare/clinical concepts or health-informatics basics

### 2. Course Learning Objectives

- Provide students with a comprehensive understanding of how AI technologies are applied across clinical, operational, and public-health domains, including the historical evolution, regulatory environment, and ethical considerations specific to healthcare.
- Equip learners with the ability to acquire, preprocess, and govern heterogeneous health-care data sets, applying appropriate standards and governance practices to ensure data quality, privacy, and interoperability.
- Develop students' proficiency in designing, training, evaluating, and interpreting machine-learning and deep-learning models for health-care problems, emphasizing rigorous validation, performance metrics, and fairness techniques.

- Enable learners to translate AI concepts into concrete, domain-specific solutions such as medical imaging analysis, clinical NLP, predictive risk modeling, and operational optimization, while critically assessing feasibility, safety, and clinical impact.
- Prepare students to plan and execute the end-to-end deployment of AI systems in real-world health-care settings, covering integration into clinical workflows, continuous monitoring, governance, and emerging trends toward value-based, personalized medicine.

### 3. Teaching Methodology

- Lectures and Presentations
- Interactive Discussions and Case Studies
- Lab Sessions
- Guest Lectures

### 4. Evaluation System

Activities	Class Test Full Marks	Assignment Full Marks	Attendance Full Marks	Total Marks
CIA-1	25	10	05	40
CIA-2	25	10	05	40
End Semester Examination (ESE)	-	-	-	60
Total				100 Marks

### 5. Course Modules

Module	Topics	Hours
1	<p><b>Foundations of AI in Healthcare</b></p> <ul style="list-style-type: none"> <li>- Role and types of AI in healthcare (diagnostic, predictive, operational)</li> <li>- Brief history and evolution of AI; distinction between AI, Machine Learning, and Deep Learning</li> <li>- Core AI paradigms relevant to health: reactive agents, limited-memory agents, rational agents</li> <li>- Ethical foundations: data ethics, privacy, informed consent, bias, fairness</li> <li>- Safety risks and practical guidelines (short-,</li> </ul>	5

	<p>medium-, long-term) for AI systems in clinical settings</p> <ul style="list-style-type: none"> <li>- Overview of regulatory landscape (FDA, CE, EMA) and compliance basics</li> </ul>	
2	<p><b>Healthcare Data Landscape &amp; Governance</b></p> <ul style="list-style-type: none"> <li>- Clinical data sources and formats (structured, unstructured, imaging, sensor, genomics)</li> <li>- Key health-informatics standards (HL7, FHIR, DICOM, LOINC, SNOMED CT)</li> <li>- Data governance principles: stewardship, security, privacy, and fragmentation challenges</li> <li>- Data quality, de-identification, and consent management</li> <li>- End-to-end data preparation pipeline: cleaning, normalization, handling missing values, feature engineering</li> <li>- Practical considerations for integrating fragmented datasets across institutions</li> </ul>	6
3	<p><b>Machine Learning &amp; Deep Learning Fundamentals</b></p> <ul style="list-style-type: none"> <li>- Fundamental concepts of Machine Learning</li> <li>- Overview of supervised, unsupervised and reinforcement learning (high-level)</li> <li>- Core algorithms with intuitive explanation: Linear Regression, Logistic Regression, Decision Trees, Support Vector Machines, k-Nearest Neighbours</li> <li>- Introduction to Neural Networks: perceptron, activation functions, simple feed-forward architecture</li> <li>- Deep Learning basics for health: Convolutional Neural Networks (CNN) for images, Recurrent Neural Networks (RNN) for sequences</li> <li>- Model training essentials: loss functions, optimization basics, over-fitting &amp; regularization</li> <li>- Real-world healthcare examples for each technique</li> </ul>	8
4	<p><b>Model Development, Evaluation &amp; Interpretability</b></p> <ul style="list-style-type: none"> <li>- End-to-end model building pipeline (data split, feature selection, model training)</li> <li>- Validation strategies: hold-out, k-fold cross-validation, temporal split</li> <li>- Evaluation metrics for classification and regression (accuracy, precision, recall, F1, ROC-AUC, RMSE, MAE)</li> <li>- Basics of statistical inference and significance testing for model results</li> <li>- Model interpretability techniques: feature</li> </ul>	8

	<p>importance, partial dependence, SHAP, LIME</p> <ul style="list-style-type: none"> <li>- Algorithmic fairness concepts and simple bias-mitigation approaches</li> <li>- Clinical risk stratification and disease-progression modeling fundamentals</li> <li>- Predictive analytics use-cases: readmission risk, treatment outcome, outbreak forecasting</li> </ul>	
5	<p><b>Domain-Specific AI Applications in Healthcare</b></p> <ul style="list-style-type: none"> <li>- AI-driven diagnosis and treatment-recommendation systems</li> <li>- Medical image analysis fundamentals (radiology, pathology, dermatology): classification, segmentation, detection using CNNs</li> <li>- Computer-vision tasks specific to health (e.g., tumor detection, organ segmentation)</li> <li>- Natural Language Processing for clinical notes, discharge summaries, and tele-medicine interactions</li> <li>- AI for hospital operations: resource allocation, scheduling, patient flow optimization</li> <li>- Wearable devices, IoT health monitoring, and real-time analytics</li> <li>- Introductory AI for drug discovery and personalized medicine (target identification, virtual screening)</li> <li>- Public-health surveillance and epidemiological forecasting using AI</li> </ul>	8
6	<p><b>Implementation, Governance, Ethics &amp; Future Trends</b></p> <ul style="list-style-type: none"> <li>- Strengths, limitations, and realistic expectations of AI in healthcare</li> <li>- Integration of AI tools into clinical workflows and evidence-based practice</li> <li>- Safety analysis, design, and reporting guidelines for AI systems (e.g., IEC 82304, FDA Good Machine Learning Practice)</li> <li>- Governance frameworks, policy considerations, and post-market surveillance</li> <li>- Deployment essentials: cloud vs. edge, CI/CD pipelines, model monitoring and drift detection</li> <li>- Emerging trends: connected medicine, virtual assistants, remote monitoring, smart implantables, digital therapeutics</li> <li>- Transition from volume-based to value-based care and the role of AI in personalized medicine</li> </ul>	7

## 6. References

### Textbooks:

1. Adam Bohr, Kaveh Memarzadeh (eds.) - "Artificial Intelligence in Healthcare" (Academic Press / Elsevier, 2020)
2. Kayvan Najarian et al. (eds.) - "Artificial Intelligence in Healthcare and Medicine" (Routledge/CRC)

### Reference Books:

1. Rashmi Agrawal et al. - "Machine Learning for Healthcare: Handling and Managing Data" (CRC Press)
2. Eric Topol - "Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again" (Basic Books)

## 7. Course Outcomes

ID	Statement	Action Verb	Knowledge Level
AM-804.1	Identify and differentiate the major types of AI used in healthcare, the evolution of AI terminology (AI, ML, DL), core AI paradigms, and the key ethical and regulatory considerations governing clinical AI systems.	Identify	Remember
AM-804.2	Explain the structure and purpose of major health-informatics standards (HL7, FHIR, DICOM, LOINC, SNOMED CT) and construct a complete data-preparation pipeline--including cleaning, de-identification, and feature engineering--for fragmented clinical datasets.	Explain	Understand
AM-804.3	Implement supervised and deep-learning models (e.g., logistic regression, decision trees, CNNs, RNNs) on real-world healthcare data sets and tune hyper-parameters to achieve predefined performance thresholds.	Implement	Apply

AM-804.4	Analyze model outcomes using appropriate classification/regression metrics, conduct statistical significance testing, and apply interpretability techniques (SHAP, LIME, partial dependence) to assess fairness and clinical relevance.	Analyze	Analyze
AM-804.5	Evaluate alternative AI approaches for at least three domain-specific healthcare problems (medical image analysis, clinical NLP, and operational optimization) and recommend the most suitable technique based on accuracy, interpretability, and implementation constraints.	Evaluate	Evaluate
AM-804.6	Create a comprehensive end-to-end deployment plan for a healthcare AI solution that incorporates cloud/edge architecture, CI/CD pipelines, model monitoring, drift detection, and compliance with FDA/EMA governance frameworks.	Create	Create

## 8. CO-PO Mapping

CO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1	3	2	1	2	1	3	-	3	-	-	-	2
CO2	3	3	3	2	3	2	-	2	2	2	1	2
CO3	3	2	3	3	3	2	-	2	2	2	1	2
CO4	3	3	2	3	3	3	-	3	1	2	-	2
CO5	3	3	3	2	2	2	-	2	2	2	2	2
CO6	3	3	3	2	3	3	1	3	2	3	3	2

## 9. CO-PSO Mapping

CO	PSO1	PSO2	PSO3
CO1	3	1	3
CO2	3	3	2
CO3	3	3	1
CO4	3	2	2
CO5	3	2	2
CO6	3	3	3

