DAYANANDA SAGAR UNIVERSITY

SCHOOL OF ENGINEERING

SCHOOL OF ENGINEERING
Department of AI & ROBOTICS

SCHEME AND SYLLABUS

M.Tech. PROGRAMME

ARTIFICIAL INTELLIGENCE & ROBOTICS

Academic Year 2026-27

Devarakaggalahalli,Harohalli,
Kanakapura Main Road,
Bangalore South– 562112

Definitions / Descriptions

Definition of Credit	
1 Hour lecture (L) Per Week	01 Credit
1 Hour Tutorial (T) Per Week	01 Credit
1 Hour Practical (P) Per Week	0.5 Credit
1 Hour Project (J) Per Week	0.5 Credit

Course code and Definition

BSC	Basic science courses
ESC	Engineering Science Courses
HSMC	Humanities and Social Sciences including Management Courses
IPCC	Integrated Professional Core Course
PCC	Professional Core Course
PEC	Professional Elective Courses
OEC	Open Elective Courses
SEC	Skill Enhancement Courses

UHV	Universal Human Value Course
PROJ	Project Work
INT	Internship

DAYANANDA SAGAR UNIVERSITY

SCHOOL OF ENGINEERING

Implementation of National Education Policy (NEP) 2020 for the M. Tech students of Batch 2026-2028

The implementation of the curriculums follows NEP 2020 and addresses the following features and categories of courses:

1. Student-centric flexible curriculum,
2. Interdisciplinary Courses,
3. Multi-disciplinary Courses,
4. Ability Enhancement Courses,
5. Skill Enhancement Courses,
6. Value Added Courses,
7. Product Design and Development,
8. Internship (Rural Internship, Industry Internship, Research / Development Internship), and
9. Multiple Exit and multiple Entry
    - Certificate in Engineering after completion of first year.
    - Diploma in Engineering after completion of second year.
    - Advanced Diploma in Engineering after completion of third year.
    - Degree in Engineering after completion of fourth year.

SCHEME 2026 – 2028 Batch

I-SEMESTER

S.N	Program Code	Course Code	Course Title	Teaching Department	Teaching Hours /Week				Examination				Credits
					Lecture	Tutorial	Practical	Project	Contact Hours	CIE Marks	SEE Marks	Total Marks	
					L	T	P	J					
1			Higher Engineering Mathematics	Math	3	1	0	0	4	60	40	100	**4**
2			Artificial Intelligence and Machine Learning	AIR	3	0	2	0	5	60	40	100	4
3			Robot Operating System (ROS 2)	AIR	3	0	2	0	5	60	40	100	4
4			Responsible AI and Ethics	AIR	3	1	0	0	5	60	40	100	4
5			Adaptive and Robust Control Systems	AIR	3	1	0	0	5	60	40	100	4
6			Seminar	AIR	0	0	0	2	2	100	--	100	1
			TOTAL		**15**	**3**	**4**	**2**	**26**	**400**	**200**	**600**	**21**

	II-SEMESTER												
S.N	Program Code	Course Code	Course Title	Teaching Department	Teaching Hours / Week				Examination				Credits
					Lecture	Tutorial	Practical	Project	Contact Hours	CIE Marks	SEE Marks	Total Marks	
					L	T	P	J					
1			Deep Learning and Computer Vision for Robotics	AIR	3	0	4	0	7	60	40	100	5
2			Generative AI and Large Language Models	AIR	3	1	2	0	6	60	40	100	5
3			Explainable AI (XAI)	AIR	3	2	0	0	5	60	40	100	5
4			Professional Elective I	AIR	3	0	0	0	3	60	40	100	3
5			Professional Elective II	AIR	3	0	0	0	3	60	40	100	3
6			Seminar II	AIR	0	0	0	2	2	100	--	100	1
		TOTAL			**15**	**3**	**6**	**2**	**26**	**400**	**200**	**600**	**22**

Professional Elective-I		
S.N	**Course Code**	**Course Name**
1		Advanced Robotics
2		Perception and Sensor Fusion for Autonomous Robotic Systems

Professional Elective-II		
S.N	**Course Code**	**Course Name**
1		Intelligent Agents & Multi-Agent Systems
2		Introduction to Transformer Based NLP

	III-SEMESTER													
S.N	Program Code	Course Code	Course Title	Teaching Department	Teaching Hours /Week				Examination				Credits
					Lecture	Tutorial	Practical	Project	Contact Hours	CIE Marks	SEE Marks	Total Marks	
					L	T	P	J					
1			Agentic AI Architectures and LLM Systems	AIR	3	0	0	0	3	60	40	100	3
2			Professional Elective III	AIR	4	0	0	0	4	60	40	100	4
3			Professional Elective IV	AIR	4	0	0	0	4	60	40	100	4
4			Internship	AIR	0	0	0	24	24	100	-	100	12
			TOTAL		**11**	**0**	**0**	**24**	**35**	**280**	**120**	**400**	**23**

Professional Elective-III		
S.N	Course Code	Course Name
1		Autonomous Vehicles and SLAM
2		Space and Swam Robotics

Professional Elective-IV		
S.N	Course Code	Course Name
1		Cyber Security in Robotics
2		Augmenting LLMs Using Retrieval Augmented Generation

	IV-SEMESTER												
S.N	Program Code	Course Code	Course Title	Teaching Department	Teaching Hours / Week				Examination				Credits
					Lecture	Tutorial	Practical	Project	Contact Hours	CIE Marks	SEE Marks	Total Marks	
					L	T	P	J					
1			Capstone Project	AIR	0	0	0	24	24	100	-	100	**12**
2			Industry/Research Internship	AIR	0	0	0	8	8	100	-	100	4
			TOTAL		**0**	**0**	**0**	**32**	**32**	**200**	**-**	**200**	**16**

SEMESTER	CREDITS
I	21
II	22
III	23
IV	16
TOTAL CREDITS	**82**

SEMESTER - I

Higher Engineering Mathematics

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - I

Course Code	**:**		**Credits**	**:**	04
Hours / Week	**:**	04	**Total Hours**	**:**	60
L–T–P–J	**:**	3–1–0–0	**CIE+SEE**	**:**	60 + 40

Course Vision:
To develop strong mathematical foundations required for Artificial Intelligence, Machine Learning, Robotics, Control Systems and Intelligent Engineering by equipping students with analytical, computational and optimization techniques for modelling and solving complex engineering problems.

Course Objectives:
This course will enable students to:

1. Apply advanced mathematical concepts to formulate and solve engineering problems in AI and Robotics.
2. Develop proficiency in linear algebra, multivariable calculus, probability, graph theory and optimization for intelligent systems.
3. Interpret mathematical models underlying machine learning, deep learning and autonomous systems.
4. Apply computational tools to analyse mathematical models and engineering datasets.
5. Develop analytical thinking for research and emerging AI technologies.

Teaching-Learning Process (General Instructions)
These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.

1. Interactive lectures with mathematical derivations and problem solving.
2. Visualization using MATLAB, Python, Jupyter Notebook and GeoGebra.
3. Case studies from AI, Robotics and Intelligent Systems.
4. Tutorials emphasizing analytical reasoning and computational implementation.
5. Research-oriented assignments using recent AI applications

MODULE – I	**12 Hours**

Calculus and Optimization for Artificial Intelligence
Topics: Review of functions, limits, continuity, differentiation, partial differentiation, multivariable calculus, gradients, Jacobian, Hessian, Taylor series, convexity, unconstrained and constrained optimization, Gradient Descent, SGD, Mini-batch GD, convergence.
Applications: Parameter optimization in ML, robotic motion planning, intelligent control.
AI Integration: Python, NumPy, SciPy, TensorFlow, PyTorch, Jupyter Notebook, Google Colab.
Authentic Intelligence: Understanding mathematical foundations of optimization, convergence, stability and computational efficiency.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3
2. **Building a Brain in 10 Minutes**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+T-FX-01+V1
3. **Getting Started with Deep Learning**
   https://www.nvidia.com/en-us/training/online/

MODULE – II	**12 Hours**

Linear Algebra for Intelligent Systems
Topics: Scalars, vectors, matrices, tensors, vector spaces, rank, orthogonality, projections, eigenvalues, eigenvectors, SVD, PCA, least squares.
Applications: Feature extraction, robotics transformations, computer vision.
AI Integration: NumPy, MATLAB, Scikit-learn.
Authentic Intelligence: Numerical stability, efficient representations and interpretability

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3
2. **Building a Brain in 10 Minutes**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+T-FX-01+V1

MODULE – III	**12 Hours**

Probability, Statistics and Information Theory
Topics: Probability, random variables, Bayes theorem, distributions, expectation, covariance, MLE, Bayesian estimation, entropy, KL divergence, mutual information, bias-variance.
Applications: Probabilistic reasoning, sensor fusion, pattern recognition.
AI Integration: Python, SciPy, Pandas, Scikit-learn.
Authentic Intelligence: Reliable and explainable AI under uncertainty.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3

MODULE – IV	**12 Hours**

Discrete Mathematics and Graph Theory
Topics: Sets, logic, graphs, trees, connectivity, BFS, DFS, shortest paths, MST, graph Laplacian, spectral graph concepts.
Applications: Knowledge graphs, navigation, recommendation systems.
AI Integration: NetworkX, Neo4j.
Authentic Intelligence: Graph-based reasoning in intelligent systems.

Key Online Resources (Relevant to this Module):

1. **Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM**
   https://www.nvidia.com/en-us/on-demand/session/gtc25-dlit71491/

MODULE – V	**12 Hours**

Mathematical Foundations for Emerging AI Technologies
Topics: Convex optimization, Lagrange multipliers, numerical stability, reinforcement learning mathematics, deep learning foundations, mathematical analysis of research papers.
Applications: AI optimization, robotics, edge AI, generative AI.
AI Integration: Python, MATLAB, TensorFlow, PyTorch, Google Scholar, Semantic Scholar, Connected Papers.
Authentic Intelligence: Critical evaluation of mathematical models and research.

Key Online Resources (Relevant to this Module):

1.Generative AI Explained
https://www.nvidia.com/en-in/training/online/

2.Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM
https://www.nvidia.com/en-us/training/

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Apply multivariable calculus and optimization techniques to formulate and solve AI and Robotics problems.	L3
2	Apply linear algebra concepts for data representation, dimensionality reduction and intelligent systems.	L3
3	Analyse probabilistic models and information-theoretic measures for decision making under uncertainty.	L4
4	Model and solve graph-based engineering problems using discrete mathematical concepts.	L3
5	Evaluate mathematical models used in emerging AI technologies and interpret research outcomes.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	-	-	-	-	1	-	-	3	3	3
CO-2	3	2	2		-	-	-	-	1	-	-	3	3	3
CO-3	3	2	2	1	-	-	-	-	1	-	-	3	3	3
CO-4	3	2	2	1	-	-	-	-	1	-	-	3	3	3
CO-5		2	2	2	-	-	-	-	1	2	-	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

➢**PSO-1:** Design and build robotic systems for different applications using analytical, logical, and

problem-solving skills.

- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. Marc Peter Deisenroth, A. Aldo Faisal and Cheng Soon Ong, Mathematics for Machine Learning, Cambridge University Press, 2020.
2. Stephen Boyd and Lieven Vandenberghe, Introduction to Applied Linear Algebra, Cambridge University Press, 2018.
3. Gilbert Strang, Linear Algebra and Learning from Data, Wellesley-Cambridge Press, 2019.

REFERENCE BOOKS:

1. Erwin Kreyszig, Advanced Engineering Mathematics, 10th Ed., Wiley.
2. Stephen Boyd and Lieven Vandenberghe, Convex Optimization, CUP.
3. Kevin P. Murphy, Probabilistic Machine Learning: An Introduction, MIT Press, 2022.
4. Christopher M. Bishop and Hugh Bishop, Deep Learning: Foundations and Concepts, Springer, 2024.

E-RESOURCES:

1. NPTEL: Mathematical Foundations for Machine Learning (IISc Bangalore)
2. NPTEL: Linear Algebra (IIT Madras)
3. NPTEL: Probability and Statistics
4. MIT OpenCourseWare: Linear Algebra (Gilbert Strang)
5. Stanford Online: Convex Optimization (Stephen Boyd)
6. NVIDIA Deep Learning Institute
7. **Distributed Optimization and Machine Learning – IIT Bombay**
   NPTEL – Distributed Optimization and Machine Learning – Optimization and distributed ML techniques.
8. **Edge Computing – IIT Kanpur**
   NPTEL – Edge Computing – Emerging trends from cloud to edge systems.

RECOMMENDED TOOLS

Python, NumPy, SciPy, Pandas, Jupyter Notebook, Google Colab, MATLAB, TensorFlow, PyTorch, Scikit-learn, NetworkX, Neo4j, Mathematica, GeoGebra.

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - I

Course Code	:		**Credits**	:	04
Hours / Week	:	05	**Total Hours**	:	45 (Theory) + 30 (Lab)
L–T–P–J	:	3–0–2–0	**CIE+SEE**	:	60+40

Course Vision:

To develop a strong foundation in Artificial Intelligence and Machine Learning by enabling students to design, implement, evaluate and deploy intelligent systems for robotics, automation and real-world engineering applications using modern AI techniques and responsible AI practices.

Course Objectives:

This course will enable students to:

1. Understand the theoretical foundations of Artificial Intelligence and Machine Learning.
2. Design, implement and evaluate supervised, unsupervised and reinforcement learning algorithms.
3. Develop deep learning models for perception, prediction and intelligent decision-making.
4. Apply modern AI tools and deployment frameworks for robotics and edge AI applications.
5. Develop trustworthy, explainable and ethically responsible AI solutions.

Teaching-Learning Process (General Instructions)

These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.

1. Interactive lectures supported with mathematical derivations and coding demonstrations.
2. Hands-on implementation using Python, Scikit-learn, PyTorch and TensorFlow.
3. Case studies in robotics, healthcare, manufacturing and autonomous systems.
4. Research-paper discussions and mini-project based learning.
5. Model evaluation, explainability and responsible AI practices.

MODULE – I	**09 Hours**

Foundations of Artificial Intelligence and Machine Learning

Topics: Intelligent agents; Problem formulation; Supervised, Unsupervised and Reinforcement Learning; Bias-Variance trade-off; Regularization (L1, L2, Elastic Net); Cross-validation; Performance metrics; Maximum Likelihood Estimation; Bayesian Inference; MAP Estimation.
Applications: Predictive maintenance, medical diagnosis, fault detection and intelligent decision support.
AI Integration: Python, Scikit-learn, AutoML (H2O, Auto-sklearn), Jupyter Notebook.
Authentic Intelligence: Fairness, bias mitigation, reproducibility and responsible data-driven AI.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3
2. **Accelerating End-to-End Data Science Workflows**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-DS-01+V2
3. **Enhancing Data Science Outcomes With Efficient Workflow**
   https://learn.nvidia.com/courss/course-detail?course_id=course-v1:DLI+C-DS-04+V1

MODULE – II	09 Hours

Supervised Learning
Topics: Linear and Logistic Regression; Decision Trees; Random Forests; Gradient Boosting (XGBoost, LightGBM); Support Vector Machines; Kernel Methods; Gaussian Processes; Multi-class and Multi-label classification.
Applications: Robot perception, predictive analytics, industrial quality inspection.
AI Integration: SHAP, LIME, Optuna, MLflow.
Authentic Intelligence: Interpretable and trustworthy predictive modelling.

Key Online Resources (Relevant to this Module):

1. **Best Practices in Feature Engineering for Tabular Data with GPU Acceleration**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-DS-06+V1
2. **Accelerated Machine Learning for Time Series Forecasting**
   https://developer.nvidia.com/blog/accelerating-time-series-forecasting-with-rapids-cuml/
3. **Applications of AI for Predictive Maintenance**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-MF-02+V1

MODULE – III	09 Hours

Unsupervised Learning and Representation Learning
Topics: K-Means, DBSCAN, Hierarchical Clustering; PCA, Kernel PCA; t-SNE, UMAP; Autoencoders; Variational Autoencoders; Semi-supervised and Self-supervised Learning; Anomaly Detection.
Applications: Sensor analytics, feature extraction, industrial monitoring.
AI Integration: SimCLR, BYOL, Hugging Face Datasets.
Authentic Intelligence: Transparent pattern discovery and robust feature learning.

Key Online Resources (Relevant to this Module):

1. **Introduction to Transformer-Based Natural Language Processing**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-08+V1
2. **Synthetic Tabular Data Generation Using Transformers**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-09+V1
3. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1

MODULE – IV	09 Hours

Deep Learning for Intelligent Systems
Topics: MLP, Backpropagation; CNNs; ResNet; EfficientNet; RNN, LSTM, GRU; Attention; Transformers; Transfer Learning; Fine-tuning; Batch Normalization; Dropout.
Applications: Computer vision, NLP, speech, autonomous robots.
AI Integration: PyTorch, TensorFlow, Hugging Face, Vision Transformers, LoRA.
Authentic Intelligence: Explainable deep learning and ethical deployment.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3
2. **Building Transformer-Based Natural Language Processing Applications**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-03+V3
3. **Building Real-Time Video AI Applications**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-IV-01+V1

MODULE – V	09 Hours

Reinforcement Learning, Edge AI and Deployment
Topics: Markov Decision Processes; Q-Learning; Policy Gradient Methods; Imitation Learning; Edge AI; ONNX; TensorRT; MLOps; Model Monitoring; Responsible AI.
Applications: Autonomous navigation, robot learning, edge intelligence.
AI Integration: Jetson, MLflow, DVC, FastAPI, RT-2, Isaac Lab.
Authentic Intelligence: Safe, reliable and human-centred AI deployment.

Key Online Resources (Relevant to this Module):

1. **Getting Started With AI on NVIDIA Jetson Nano**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-02+V2
2. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
3. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain the principles of Artificial Intelligence and formulate machine learning problems.	L2
2	Develop and evaluate supervised and unsupervised learning models using appropriate algorithms.	L3
3	Build deep learning models for intelligent perception and prediction.	L3
4	Design complete AI pipelines incorporating model optimization, explainability and deployment.	L6
5	Apply AI techniques to robotics and edge AI applications while adhering to responsible AI principles.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	-	1	1	-	3	3	3
CO-2	3	3	3	2	2	2	1	3	-	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	-	2	3	3	-	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, MIT Press, 2016.
2. Christopher M. Bishop and Hugh Bishop, Deep Learning: Foundations and Concepts, Springer, 2024.
3. Aurélien Géron, Hands-On Machine Learning with Scikit-Learn, Keras and TensorFlow, 3rd Edition, O'Reilly, 2022.

REFERENCE BOOKS:

1. Christopher M. Bishop, Pattern Recognition and Machine Learning, Springer.
2. Sebastian Raschka and Vahid Mirjalili, Python Machine Learning, Packt.
3. Kevin P. Murphy, Probabilistic Machine Learning: An Introduction, MIT Press, 2022.

E-RESOURCES:

- NVIDIA Deep Learning Institute
- NPTEL – Introduction to Machine Learning
- Hugging Face Course
- Google ML Crash Course
- OpenAI Cookbook
- PyTorch Tutorials
- TensorFlow Tutorials

RECOMMENDED TOOLS

Python, NumPy, Pandas, Scikit-learn, PyTorch, TensorFlow, Keras, OpenCV, Hugging Face Transformers, MLflow, DVC, FastAPI, ONNX Runtime, TensorRT, Jupyter Notebook, Google Colab, Kaggle.

LAB EXERCISES

S.No	Exercise Title	Aim / Objective	Tools & Platform	Hrs	CO
MODULE I · Exercises 1–2 : Supervised Learning Algorithms (04 Hours)					
1	Regression & Classification Pipeline	Train linear regression and random forest models on a tabular dataset; compare MAE / RMSE and visualize feature importance using SHAP.	Python, Scikit-learn, SHAP, Matplotlib. Dataset: California Housing / UCI Adult	2 h	CO1

2	SVM & Kernel Methods	Train SVM with RBF kernel for robot fault classification; tune C and gamma using grid search; plot decision boundary and support vectors.	Python, Scikit-learn. Dataset: Robot Fault Dataset / UCI Wine	2 h	CO1
MODULE II · Exercises 3–5 : Deep Learning for Robotics (06 Hours)					
3	CNN for Image Classification	Fine-tune ResNet-18 on a robotics image dataset; track training / validation accuracy, apply data augmentation, and compare with VGG-16.	PyTorch, torchvision. Dataset: CIFAR-10 / Custom Robot Images	2 h	CO2
4	LSTM for Time-Series Prediction	Predict robot joint velocity using LSTM; compare with GRU and 1D-CNN, and analyze temporal dependencies.	PyTorch, Pandas. Dataset: Robot Arm Sensor Data	2 h	CO2
5	VAE for Anomaly Detection	Train a Variational Autoencoder (VAE) on normal robot sensor data; use reconstruction error as anomaly score and plot latent space using t-SNE.	PyTorch, Matplotlib. Dataset: KUKA Robot Sensor Data / Custom	2 h	CO2
MODULE III · Exercise 6 : Reinforcement Learning for Navigation (02 Hours)					
6	Q-Learning for Grid Navigation	Implement Q-learning for a grid-world robot navigation task; visualize Q-table convergence and compare with SARSA.	Python, NumPy, Matplotlib. Environment: Custom GridWorld	2 h	CO3
MODULE IV · Exercise 7 : Model Deployment & Optimization (02 Hours)					
7	Model Deployment with ONNX	Export trained PyTorch model to ONNX; benchmark inference latency; deploy with ONNX Runtime and compare with TorchScript.	PyTorch, ONNX Runtime. Dataset: Previously Trained Model	2 h	CO4
MODULE V · Exercise 8 : End-to-End MLOps Pipeline (04 Hours)					
8	End-to-End ML Pipeline	Build a complete MLOps pipeline including data versioning (DVC), experiment tracking (MLflow), model registry, and deployment endpoint.	Python, MLflow, DVC, FastAPI. Dataset: Any Classification Dataset	4 h	CO5

Robot Operating System (ROS 2)
[As per Choice Based Credit System (CBCS) scheme]
SEMESTER - I

Course Code	:		**Credits**	:	04
Hours / Week	:	05	**Total Hours**	:	45 (Theory) + 30 (Lab)
L–T–P–J	:	3–0–2–0	**CIE+SEE**	:	60+40

Course Vision:
To provide comprehensive theoretical knowledge and hands-on expertise in ROS 2, enabling students to design, develop, simulate and deploy intelligent robotic systems using modern middleware, autonomous navigation, motion planning and AI-enabled robotic software frameworks.

Course Objectives:
This course will enable students to:
1. Understand ROS 2 architecture, DDS middleware and distributed robotic communication.
2. Develop ROS 2 nodes, packages, services, actions and launch systems using Python and C++.
3. Model, simulate and visualize robotic systems using URDF, RViz2 and Gazebo/Isaac Sim.
4. Implement autonomous navigation, SLAM, localization and motion planning using Nav2 and MoveIt 2.
5. Deploy secure, scalable and AI-enabled robotic applications on embedded and real robotic platforms.

Teaching-Learning Process (General Instructions)
These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.
1. Lecture method along with hands-on ROS 2 programming, simulation-driven learning, and robotics middleware demonstrations to explain distributed robotic software systems.
2. Interactive teaching involving brainstorming, discussions, collaborative learning, analytical problem-solving, and group activities.
3. Demonstration of robotic communication, navigation, simulation, and manipulation systems using RVIZ2, Gazebo, Isaac Sim, and industrial robotics case studies.
4. Encouraging project-based and research-oriented learning through ROS 2 package development, autonomous navigation, motion planning, and robot deployment tasks.
5. Developing critical thinking through real-time communication analysis, robotic system integration, motion planning optimisation, and deployment debugging activities.
6. Demonstrating multiple approaches for solving robotic software integration problems and encouraging students to design scalable, reliable, and intelligent robotic systems.

MODULE – I	**09 Hours**

ROS 2 Architecture and Communication
Topics: DDS, QoS, Nodes, Topics, Services, Actions, Parameters, Lifecycle Nodes, Components, Launch System, rosbag2, tf2, colcon, ament.
Applications: Distributed robotic systems, multi-robot communication, industrial automation.
AI Integration: ROS2 + TensorRT, Isaac ROS, FastDDS, Python/C++.

Authentic Intelligence: Reliable, secure and scalable middleware for trustworthy robotics.

Key Online Resources (Relevant to this Module):

1. **Accelerating ROS 2 with NVIDIA GPU-Powered Libraries and AI Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1
2. **Getting Started With Isaac ROS**
   https://nvidia-isaac-ros.github.io/getting_started/index.html
3. **An Introduction to AI-Based Robot Development With Isaac ROS**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-ros/latest/an-introduction-to-ai-based-robot-development-with-isaac-ros/index.html

MODULE – II	09 Hours

Robot Modelling, Simulation and Perception
Topics: URDF, XACRO, robot_state_publisher, tf2, RViz2, Gazebo Harmonic, Isaac Sim, LiDAR, IMU, RGB-D cameras, ros2_control basics.
Applications: Digital twins, industrial manipulators, AMRs, UAV simulation.
AI Integration: Isaac Sim, Gazebo, OpenCV, perception pipelines.
Authentic Intelligence: Simulation-driven validation and safe deployment.

Key Online Resources (Relevant to this Module):

1. **Getting Started: Simulating Your First Robot in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/building-your-first-robot-in-isaac-sim/index.html
2. **Ingesting Robot Assets and Simulating Your Robot in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/ingesting-robot-assets-and-simulating-your-robot-in-isaac-sim/index.html
3. **Developing Robots with Software-in-the-Loop (SIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/developing-robots-with-sil-in-isaac-sim/index.html

MODULE – III	09 Hours

Localization, Mapping and Autonomous Navigation
Topics: slam_toolbox, Cartographer, AMCL, Nav2 architecture, planners, controllers, behavior trees, multi-robot navigation.
Applications: Warehouse robots, service robots, AGVs.
AI Integration: Learning-based planners, RL-assisted navigation.
Authentic Intelligence: Safe, explainable and robust autonomous navigation.

Key Online Resources (Relevant to this Module):

1. **Leveraging ROS 2 and Hardware-in-the-Loop (HIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/leveraging-ros-2-and-hil-in-isaac-sim/index.html
2. **Scalable Multi-Robot Scene Workflows Using ROS Simulation Interfaces Standard in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/scalable-multi-robot-scene-workflows-using-ros-simulation-interfaces-standard-in-isaac-sim/index.html
3. **A Beginner's Guide to Autonomous Robots**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-35+V1

MODULE – IV	09 Hours

Motion Planning and Manipulation
Topics: MoveIt 2, SRDF, Planning Scene, OMPL, collision checking, inverse kinematics, grasp planning, pick-and-place, TAMP.
Applications: Industrial and service robot manipulation.
AI Integration: Deep-learning based grasping, Isaac Manipulator.
Authentic Intelligence: Reliable manipulation and collision-free motion.

Key Online Resources (Relevant to this Module):

1. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
2. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
3. **Best Practices for Robotics and OpenUSD**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/best-practices-for-robotics-and-openusd/index.html

MODULE – V	03 Hours

Advanced ROS 2 Deployment
Topics: micro-ROS, SROS2, Docker, Kubernetes, CI/CD, monitoring, logging, embedded deployment, edge robotics, LLM integration.
Applications: Embedded robotics, edge AI, cloud robotics.
AI Integration: LLM agents, Foundation Models, Isaac ROS.
Authentic Intelligence: Secure deployment and ethical AI-enabled robotic systems.

Key Online Resources (Relevant to this Module):

1. **Accelerating ROS 2 with NVIDIA GPU-Powered Libraries and AI Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1
2. **Getting Started With Isaac ROS**
   https://nvidia-isaac-ros.github.io/getting_started/index.html
3. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain ROS 2 architecture, middleware and communication mechanisms.	L2
2	Develop ROS 2 software packages and robotic simulations.	L3
3	Implement localization, SLAM and autonomous navigation solutions.	L3
4	Apply MoveIt 2 for robotic manipulation and motion planning.	L3
5	Design and deploy secure AI-enabled robotic applications on	L5

	embedded and real robotic platforms.	

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	-	1	1	-	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	-	3	2	1	2	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	-	3	3	2	3	3	3	3
CO-5	3	2	3	3	-	3	2	3	3	-	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. A.Koubaa (Ed.), Robot Operating System (ROS): The Complete Reference, Springer, Latest Edition.
2. Morgan Quigley, Brian Gerkey and William D. Smart, Programming Robots with ROS, O'Reilly.
3. The Construct, ROS 2 in 5 Days, Latest Edition.

REFERENCE BOOKS:

1. Official ROS 2 Documentation
2. Nav2 Documentation
3. MoveIt 2 Documentation
4. NVIDIA Isaac ROS Documentation

E-RESOURCES:

1. Official ROS 2 Documentation
2. NVIDIA Isaac ROS Documentation
3. The Construct ROS 2 Courses
4. Gazebo Simulator Documentation
5. MoveIt 2 Documentation

RECOMMENDED TOOLS

ROS 2 Jazzy/Humble, Ubuntu Linux, Python, C++, colcon, RViz2, Gazebo Harmonic, Isaac Sim, MoveIt 2, Nav2, slam_toolbox, OpenCV, NumPy, Docker, micro-ROS, ros2_control, Git.

LAB EXERCISES

S.No	Exercise Title	Aim / Objective	Tools & Platform	Hrs	CO
MODULE I · Exercises 1–2 : ROS 2 Communication & Robot Modelling (04 Hours)					
1	ROS 2 Publisher-Subscriber & Custom Messages	Create a ROS 2 package with publisher and subscriber nodes; define a custom message and test with QoS settings.	ROS 2 Humble, Python / C++, Ubuntu 22.04	2 h	CO1, CO2
2	URDF Robot Modelling in RVIZ2	Build a 6-DOF robotic arm URDF; visualize in RVIZ2 and animate joints using joint_state_publisher_gui.	ROS 2, RVIZ2, XACRO	2 h	CO2
MODULE II · Exercises 3–5 : Robot Simulation, SLAM & Navigation (06 Hours)					
3	Gazebo Simulation with Sensors	Spawn a differential drive robot in Gazebo with LiDAR and IMU; publish sensor data to ROS 2 topics and visualize in RVIZ2.	ROS 2 Humble, Gazebo Ignition	2 h	CO2, CO3
4	SLAM & 2D Map Building	Run slam_toolbox in a Gazebo environment; drive the robot using teleoperation and save/load generated maps.	ROS 2, slam_toolbox, Gazebo	2 h	CO3
5	Autonomous Navigation with Nav2	Configure Nav2 for the mapped environment; send goal poses, monitor navigation, and tune the DWB local planner.	ROS 2, Nav2, RVIZ2	2 h	CO3
MODULE III · Exercise 6 : Motion Planning with MoveIt! 2 (04 Hours)					
6	MoveIt! 2 Pick-and-Place	Configure MoveIt! 2 for a simulated robotic arm and execute a pick-and-place task with collision avoidance.	ROS 2, MoveIt! 2, Gazebo	4 h	CO4
MODULE IV · Exercise 7 : Robot Hardware Interface (02 Hours)					
7	ros2_control Hardware Interface	Implement a mock ros2_control hardware interface and control a simulated joint using a ROS 2 controller.	ROS 2, ros2_control	2 h	CO4
MODULE V · Exercise 8 : Embedded ROS Integration (04 Hours)					
8	micro-ROS on Embedded System	Flash micro-ROS on an Arduino / ESP32 and publish IMU data to a ROS 2 host over serial transport.	micro-ROS, Arduino / ESP32, ROS 2	4 h	CO5

RESPONSIBLE AI AND ETHICS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - I

Course Code	:		**Credits**	:	04
Hours / Week	:	03	**Total Hours**	:	60
L–T–P–J	:	3–1–0–0	**CIE+SEE**	:	60+40

Course Vision:

To develop ethically responsible AI professionals capable of designing, deploying and governing trustworthy, transparent, fair, secure and human-centric AI systems in accordance with global regulatory frameworks and societal values.

Course Objectives:

This course will enable students to:

1. Understand ethical, legal and societal dimensions of Artificial Intelligence.
2. Evaluate fairness, transparency, accountability and explainability in AI systems.
3. Apply privacy-preserving AI techniques and responsible data governance principles.
4. Assess AI systems against emerging international regulations and governance frameworks.
5. Design trustworthy AI solutions for robotics and intelligent systems.

Teaching-Learning Process (General Instructions)

These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.

1. Case-study driven learning using real-world AI failures and successes.
2. Policy discussions and ethical debates.
3. Hands-on demonstrations using fairness, explainability and privacy tools.
4. Research paper review and AI audit exercises.
5. Industry-oriented governance and compliance case studies.

MODULE – I	**09 Hours**

Foundations of Responsible AI

Topics: History of AI Ethics; Ethical theories; IEEE, OECD, UNESCO and NIST AI principles; EU AI Act overview; AI autonomy; value alignment; human-centered AI.
Applications: Healthcare, criminal justice, autonomous vehicles, social media.
AI Integration: ChatGPT, Constitutional AI, RLHF, OpenAI Evals.
Authentic Intelligence: Ethical reasoning, value alignment and trustworthy AI.

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **Building LLM Applications With Prompt Engineering**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-11+V1

MODULE – II	09 Hours

Fairness, Bias and Explainability
Topics: Data, algorithmic and societal bias; fairness metrics; bias detection; mitigation; SHAP; LIME; counterfactual explanations; intersectionality.
Applications: Hiring, finance, facial recognition, healthcare.
AI Integration: IBM AIF360, Fairlearn, SHAP, LIME.
Authentic Intelligence: Inclusive, transparent and explainable AI.

Key Online Resources (Relevant to this Module):

1. **Evaluating RAG and Semantic Search Systems**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-32+V1
2. **Enhancing Data Science Outcomes With Efficient Workflow**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-DS-04+V1
3. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1

MODULE – III	0 Hours

Privacy, Security and Responsible Data Governance
Topics: Privacy by design; differential privacy; federated learning; secure multiparty computation; GDPR; India DPDP Act; synthetic data; consent frameworks.
Applications: Healthcare AI, IoT, smart cities.
AI Integration: TensorFlow Privacy, PySyft, OpenDP.
Authentic Intelligence: Privacy-preserving and secure AI.

Key Online Resources (Relevant to this Module):

1. **Introduction to Federated Learning with NVIDIA FLARE**
   https://www.nvidia.com/en-us/on-demand/session/nvidiaflareday25-nvfd01/
2. **Decentralized AI at Scale with NVIDIA FLARE**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-29+V1
3. **Introduction to Multi-Modal Data Curation**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:NVIDIA+S-FX-27+V1

MODULE – IV	09 Hours

AI Governance, Accountability and Regulation
Topics: Model cards; datasheets; AI audits; risk assessment; algorithmic impact assessment; NIST AI RMF; ISO/IEC 42001; EU AI Act; monitoring.
Applications: Public sector AI, medical AI, finance.
AI Integration: Responsible AI Toolbox, Azure Responsible AI Dashboard.
Authentic Intelligence: Accountable AI lifecycle management.

Key Online Resources (Relevant to this Module):

1. **Evaluating RAG and Semantic Search Systems**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-32+V1
2. **Building Observable and Scalable Multi-Agent Workflows for Asset Lifecycle Management**
   https://www.nvidia.com/gtc/session-catalog/sessions/gtc26-dliw82268/

MODULE – V	09 Hours

Societal Impact and Sustainable AI
Topics: AI and employment; misinformation; deepfakes; Green AI; human-AI collaboration; participatory AI; responsible GenAI; AI for SDGs.
Applications: Education, governance, manufacturing, climate.
AI Integration: Deepfake detection, watermarking, Green AI tools.
Authentic Intelligence: Socially beneficial, sustainable and responsible AI

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **Techniques for Improving the Effectiveness of RAG Systems**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-20+V1
3. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain ethical principles, legal frameworks and societal implications of AI.	L2
2	Analyse and mitigate fairness, bias and explainability issues in AI systems.	L4
3	Apply privacy-preserving and secure AI techniques for responsible data governance.	L3
4	Evaluate AI systems using governance, accountability and regulatory frameworks.	L5
5	Design trustworthy, transparent and sustainable AI solutions for real-world applications.	L6

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	-	3	1	-	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	-	2	2	3	3	3
CO-4	3	2	3	3	3	-	2	-	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	-	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. Mark Coeckelbergh, AI Ethics, MIT Press, 2024.
2. Michael Kearns and Aaron Roth, The Ethical Algorithm, Oxford University Press, 2019.
3. Virginia Dignum, Responsible Artificial Intelligence, Springer, 2019.

REFERENCE BOOKS:

1. Cathy O'Neil, Weapons of Math Destruction, Crown, 2016.
2. Virginia Eubanks, Automating Inequality, St. Martin's Press, 2018.
3. EU AI Act (Latest Version), NIST AI Risk Management Framework 1.0, OECD AI Principles.

E-RESOURCES:

- NPTEL – AI: Ethics and Society
- NVIDIA DLI – Generative AI Explained
- OECD AI Principles
- NIST AI Risk Management Framework
- IBM AI Fairness 360 Documentation
- Hugging Face Responsible AI Course

RECOMMENDED TOOLS:

IBM AI Fairness 360 (AIF360), Fairlearn, SHAP, LIME, TensorFlow Privacy, PySyft, OpenDP, Hugging Face, OpenAI Evals, Azure Responsible AI Dashboard, Python, Scikit-learn.

ADAPTIVE AND ROBUST CONTROL SYSTEMS
[As per Choice Based Credit System (CBCS) scheme]
SEMESTER - I

Course Code	:		**Credits**	:	04
Hours / Week	:	04	**Total Hours**	:	60
L–T–P–J	:	3–1–0–0	**CIE+SEE**	:	50 Marks (CIE only)

Course Vision:
To equip students with advanced adaptive and robust control techniques for the analysis, design and implementation of intelligent, autonomous and safety-critical control systems under uncertainty, enabling reliable control of modern robotic and cyber-physical systems.

Course Objectives:
This course will enable students to:
1. Analyze nonlinear dynamic systems using Lyapunov stability theory.
2. Design adaptive controllers for uncertain dynamic systems.
3. Develop robust controllers using H∞, μ-synthesis and sliding mode control.
4. Apply intelligent and learning-based control strategies for robotic systems.
5. Implement adaptive and robust controllers using modern simulation tools.

Teaching-Learning Process (General Instructions)
These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.
1. Mathematical derivation supported by simulation-driven learning.
2. MATLAB/Simulink and Python-based controller design.
3. Industrial case studies in robotics, UAVs and autonomous systems.
4. Research paper discussions on intelligent control.
5. Mini-projects on adaptive and robust control implementation.

MODULE – I	**09 Hours**

Stability Theory and Nonlinear Systems
Topics: State-space representation; equilibrium points; phase portraits; Lyapunov stability; LaSalle invariance principle; input-output stability; passivity; describing functions.
Applications: Robot manipulators, UAVs, balancing robots.
AI Integration: Neural Lyapunov Functions, Isaac Lab, reinforcement learning stability verification.
Authentic Intelligence: Safe and trustworthy autonomous system stability.

Key Online Resources (Relevant to this Module):
1. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **A Beginner's Guide to Autonomous Robots**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-35+V1

MODULE – II	09 Hours

Adaptive Control Systems
Topics: MRAC; MIT rule; Lyapunov adaptation; Self-Tuning Regulators; Recursive Least Squares; persistent excitation; adaptive backstepping; adaptive observers.
Applications: Industrial drives, robot joints, UAV payload adaptation.
AI Integration: Meta-learning, model-based RL, adaptive digital twins.
Authentic Intelligence: Reliable adaptation under uncertainty.

Key Online Resources (Relevant to this Module):

1. **Train Your First Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-first-robot-in-isaac-lab/index.html
2. **Train Your Second Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-second-robot-in-isaac-lab/index.html
3. **Transferring Robot Learning Policies from Simulation to Reality**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/transferring-robot-learning-policies-from-simulation-to-reality/index.html

MODULE – III	09 Hours

Robust Control Design
Topics: Structured/unstructured uncertainty; small gain theorem; H∞ control; mixed sensitivity; μ-synthesis; DK iteration; loop shaping; LMIs.
Applications: Aerospace, precision motion systems, flexible manipulators.
AI Integration: Data-driven robust control, robust RL, ROS2-enabled controllers.
Authentic Intelligence: Reliable control under disturbances.

Key Online Resources (Relevant to this Module):

1. **Accelerating ROS 2 with NVIDIA GPU-Powered Libraries and AI Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1/
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **Leveraging ROS 2 and Hardware-in-the-Loop (HIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/leveraging-ros-2-and-hil-in-isaac-sim/index.html

MODULE – IV	09 Hours

Sliding Mode and Intelligent Nonlinear Control
Topics: Sliding surfaces; equivalent control; chattering reduction; higher-order SMC; super-twisting; disturbance observers; terminal SMC.
Applications: Electric drives, mobile robots, hydraulic systems.
AI Integration: Neural SMC, deep-learning disturbance observers.
Authentic Intelligence: Safe nonlinear control and disturbance rejection.

Key Online Resources (Relevant to this Module):

1. **Developing Robots with Software-in-the-Loop (SIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/developing-robots-with-sil-in-isaac-sim/index.html

2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **Train Your Second Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-second-robot-with-isaac-lab/index.html

MODULE – V	**09 Hours**

Advanced Intelligent Control
Topics: Gain scheduling; LPV systems; switched systems; Model Predictive Control; fuzzy control; neural control; adaptive MPC; safe reinforcement learning.
Applications: Autonomous vehicles, smart manufacturing, humanoid robots.
AI Integration: Gaussian Process MPC, Isaac Sim, reinforcement learning controllers.
Authentic Intelligence: Human-centred, explainable and reliable intelligent control.

Key Online Resources (Relevant to this Module):

1. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Analyze nonlinear dynamic systems using Lyapunov stability concepts.	L4
2	Design adaptive controllers for uncertain engineering systems.	L6
3	Develop robust controllers using H∞ and μ-synthesis techniques.	L6
4	Apply sliding mode and intelligent control strategies for nonlinear robotic systems.	L3
5	Evaluate advanced intelligent control architectures for autonomous and safety-critical applications.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	-	1	-	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	2	2	2	2	3	3	3
CO-4	3	2	3	3	3	-	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. K. J. Åström and B. Wittenmark, Adaptive Control, 2nd Edition, Dover, 2008.
2. S. Skogestad and I. Postlethwaite, Multivariable Feedback Control, 2nd Edition, Wiley, 2005.
3. J.-J. E. Slotine and W. Li, Applied Nonlinear Control, Prentice Hall, 1991.

REFERENCE BOOKS:

1. V. I. Utkin, Sliding Modes in Control and Optimization, Springer.
2. Hassan K. Khalil, Nonlinear Systems, 3rd Edition, Pearson.
3. Frank L. Lewis et al., Optimal and Robust Estimation, CRC Press.

E-RESOURCES:

- NPTEL – Control Engineering
- MIT OpenCourseWare – Feedback Control Systems
- MathWorks Control Tutorials
- NVIDIA Isaac Lab Learning Path
- ROS2 Control Documentation

RECOMMENDED TOOLS

MATLAB/Simulink, Control System Toolbox, Robust Control Toolbox, Model Predictive Control Toolbox, Python, NumPy, SciPy, Python-Control, CVXPY, ROS2, Isaac Sim, PyTorch, TensorFlow, GNU Octave.

SEMINAR – I

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - I

Course Code	:		**Credits**	:	01
Hours / Week	:	02	**Total Hours**	:	15 (Theory) + 30 (Lab)
L–T–P–J	:	0–0–0–2	**CIE**	:	100

Course Vision:

To develop students' ability to independently investigate emerging research problems in Artificial Intelligence and Robotics, critically evaluate scientific literature, communicate technical ideas effectively, and demonstrate ethical and responsible research practices through structured seminar presentations and academic writing.

Course Objectives:

This course will enable students to:

1. Identify emerging research problems in Artificial Intelligence, Robotics, and allied engineering disciplines.
2. Conduct systematic literature surveys using reputed scholarly databases and digital libraries.
3. Critically analyse, synthesise, and compare research contributions to identify knowledge gaps and future research opportunities.
4. Prepare technically sound reports using accepted academic writing and referencing standards.
5. Deliver professional seminar presentations and effectively defend technical ideas through discussion and question-answer sessions.

Teaching-Learning Process (General Instructions)

These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.

1. Lecture method includes traditional and modern teaching approaches.
2. Interactive Teaching: brainstorming, discussion, group work.
3. Use video/animation for conceptual clarity.
4. Encourage collaborative learning.
5. Problem-based learning.
6. Real-world applications discussion.

MODULE – I	**10 Hours**

Research Discovery, Literature Analysis and Research Ethics

Topics

- Identifying emerging research problems
- Formulating research questions and defining research objectives
- Conducting systematic literature surveys
- Searching scholarly databases
- Identifying research gaps and future research directions
- Annotated bibliography preparation

- Citation managemen
- Research quality assessment
- Ethical literature review and plagiarism awareness

Applications
Artificial Intelligence, Robotics, Autonomous Systems, Intelligent Manufacturing, Healthcare Robotics, Computer Vision, Edge AI, Generative AI, Multi-Agent Systems.

AI Integration
Students may utilize AI-assisted research tools such as:

- Google Scholar
- IEEE Xplore
- ACM Digital Library
- Scopus
- Semantic Scholar
- Connected Papers
- ResearchRabbit
- Consensus AI
- Elicit AI

for literature discovery, citation analysis, and research mapping, while ensuring independent verification of all AI-generated outputs.

MODULE – II	**10 Hours**

Scientific Writing, Technical Communication and Seminar Presentation

Topics

- Structure of scientific reports
- Writing Abstract, Introduction, Literature Review, Methodology, Results and Discussion, Conclusion
- IEEE manuscript preparation
- LaTeX document preparation
- Scientific presentation design
- Visual communication techniques
- Technical discussion and question handling
- Peer review methodology
- Responsible use of AI in academic writing

Applications
Conference paper preparation, seminar presentation, technical documentation, research proposal writing, thesis preparation.

AI Integration
Students may utilize

- Overleaf
- LaTeX
- IEEE Templates
- Grammarly

- LanguageTool
- ChatGPT
- Microsoft Copilot
- Perplexity AI
- Canva Presentation
- Gamma AI

for drafting, proofreading, formatting, visualization and presentation support while maintaining academic integrity and proper attribution.

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Identify and select a relevant research topic with well-defined scope and objectives.	L3
2	Conduct a systematic literature survey using academic databases and reference tools.	L3
3	Synthesise and critically analyse findings from multiple research papers.	L4
4	Prepare a structured technical seminar report following academic standards.	L5
5	Deliver an effective oral presentation and respond to technical queries.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.

- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. Wayne C. Booth, Gregory G. Colomb, Joseph M. Williams, Joseph Bizup and William T. FitzGerald, **The Craft of Research** 4th Edition, University of Chicago Press
2. Justin Zobel, **Writing for Computer Science** 4th Edition Springer

REFERENCE BOOKS:

1. N. J. Higham, *Handbook of Writing for the Mathematical Sciences*, SIAM, Second Edition, 2020.
2. W. Strunk Jr. and E. B. White, *The Elements of Style*, 4th Edition, Pearson.
3. Michael Alley, *The Craft of Scientific Presentations*, Springer, 2013.

E-RESOURCES:

Scholarly Databases

- IEEE Xplore Digital Library
- ACM Digital Library
- SpringerLink
- ScienceDirect
- Scopus
- Web of Science
- Google Scholar

Research Discovery Tools

- Semantic Scholar
- Connected Papers
- ResearchRabbit
- OpenAlex

Reference Management

- Zotero
- Mendeley
- EndNote

Academic Writing

- Overleaf
- IEEE Author Center
- Grammarly
- LanguageTool

AI-assisted Research

- ChatGPT
- Microsoft Copilot
- Perplexity AI
- Elicit AI
- Consensus AI

LAB EXERCISES

The current lab exercises are already good. I would only make them progressively aligned with the research workflow:

1. **Research Topic Identification & Literature Search**
    - IEEE Xplore, Scopus, Google Scholar, Semantic Scholar, Zotero
2. **Critical Literature Review & Research Gap Identification**
    - Connected Papers, ResearchRabbit, Elicit AI
3. **Technical Report Preparation**
    - IEEE Template, Overleaf, LaTeX, Grammarly
4. **Seminar Presentation**
    - PowerPoint, Canva, Gamma AI
5. **Peer Review and Technical Defence**
    - Structured evaluation rubric, plagiarism checking, technical discussion

SEMESTER – II

DEEP LEARNING AND COMPUTER VISION FOR ROBOTICS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - II

Course Code	:		**Credits**	:	5
Hours / Week	:	7	**Total Hours**	:	45 (Theory) + 60 (Lab)
L–T–P–J	:	3–0–4–0	**CIE+SEE**	:	60+40

Course Vision:

To develop technically proficient and ethically responsible professionals capable of designing, analysing, and deploying deep learning and computer vision systems for intelligent robotics and autonomous applications.

Course Objectives:

1. Understand deep learning fundamentals, neural network architectures, and optimisation techniques for AI and robotics applications.
2. Apply convolutional neural networks and computer vision techniques for image classification, detection, and segmentation tasks.
3. Analyse and implement 3D vision techniques, including stereo vision, depth estimation, and point cloud processing for robot perception.
4. Develop visual odometry and SLAM-based localisation systems for autonomous robotic applications.
5. Deploy optimised deep learning and computer vision models on real-time robotic and edge AI platforms.

Teaching-Learning Process (General Instructions)

These are sample innovative pedagogical methods that may be adopted to accelerate the attainment of the various course outcomes:

1. Lecture method along with modern AI-enabled teaching approaches to explain deep learning and computer vision concepts.
2. Interactive teaching involving brainstorming, discussions, focused learning, collaborative problem-solving, and group activities.
3. Demonstration of robotics and computer vision applications using videos, simulations, animations, and real-world case studies.
4. Encouraging collaborative and project-based learning through the implementation of deep learning and robotic vision systems.
5. Developing critical thinking through quizzes, research-oriented discussions, model evaluation, and implementation of advanced AI algorithms.
6. Demonstrating multiple approaches for solving robotics perception and deployment problems and encouraging students to design innovative AI-based solutions.

MODULE – I	**09 Hours**

Deep Learning Foundations and Neural Networks

Topics:

Introduction to Artificial Intelligence, Machine Learning and Deep Learning, Biological and Artificial Neurons, Perceptron and Multilayer Perceptron, Feedforward Neural Networks, Activation Functions Including Sigmoid, Tanh, ReLU and Softmax, Loss Functions and Gradient Descent, Backpropagation Algorithm, Optimisation Techniques Including SGD, Adam and RMSProp, Overfitting, Underfitting and Regularisation, Hyperparameter Tuning and Model Evaluation.

Applications:
Handwritten Digit Recognition, Medical Diagnosis Systems, Recommendation Systems, Speech Recognition and Intelligent Assistants.
AI Integration:
Transfer Learning Using Pretrained CNN Models, Deep Learning Frameworks Including TensorFlow and PyTorch.
Authentic Intelligence:
Understanding Bias in Training Data, Reliability of AI Predictions, Ethical Use of Neural Network Models in Decision-Making Systems.

Key Online Resources (Relevant to this Module):

1. Fundamentals of Deep Learning
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3

MODULE – II	09 Hours

Convolutional Neural Networks and Visual Recognition
Topics:
Introduction to Computer Vision, Image Representation and Preprocessing, Convolution Operation and Feature Extraction, CNN Architecture, Pooling Layers, Padding and Stride, Popular CNN Architectures Including LeNet, AlexNet, VGGNet and ResNet, Feature Extraction Using CNNs, Transfer Learning Using ResNet, Data Augmentation Techniques, Batch Normalisation and Dropout, Visual Recognition and Classification.
Applications:
Face Recognition, Industrial Defect Detection, Autonomous Driving and Medical Image Analysis.
AI Integration:
Vision Transformers (ViT), CLIP-Based Image Understanding and Segment Anything Model (SAM).
Authentic Intelligence:
Explainability in Image Recognition Systems, Responsible Use of Surveillance Technologies and Bias in Visual Datasets.

Key Online Resources (Relevant to this Module):

1. Building Real-Time Video AI Applications
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-IV-01+V1
2. Getting Started With AI on NVIDIA Jetson Nano
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-02+V2

MODULE – III	09 Hours

Object Detection and Segmentation
Topics:
Object Detection Fundamentals, Region-Based CNN (R-CNN), Fast R-CNN and Faster R-CNN, YOLO Architecture and SSD Detector, Semantic Segmentation and Instance Segmentation, U-Net and Mask R-CNN, Stereo Vision and Depth Estimation, RGB-D Vision, Point Cloud Processing, PointNet and PointNet++, 3D Object Detection, Bounding Box Regression, Evaluation Metrics Including IoU, Precision, Recall and Mean Average Precision (mAP).
Applications:

Robot Navigation, Autonomous Vehicles, Medical Image Segmentation, Industrial Automation and 3D Robot Perception.

AI Integration:

Grounding DINO, Open-Vocabulary Detection, SAM-Based Segmentation Pipelines and Open3D-Based Point Cloud Processing.

Authentic Intelligence:

Fair and Transparent Object Detection, Safety in Autonomous Perception Systems and Ethical Concerns in Public Monitoring Systems.

Key Online Resources (Relevant to this Module):

1. Synthetic Data Generation for Perception Model Training in Isaac Sim https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/synthetic-data-generation-for-perception-model-training-in-isaac-sim/index.html
2. Getting Started: Simulating Your First Robot in Isaac Sim https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/building-your-first-robot-in-isaac-sim/index.html

MODULE – IV	**09 Hours**

Sequence Models and Deep Learning for Robotics

Topics:

Sequential Data Processing, Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), Encoder-Decoder Models, Attention Mechanism and Transformers, Visual Odometry Fundamentals, ORB-SLAM and Visual SLAM Concepts, DeepVO and Deep Learning-Based SLAM, Robot Localisation and Mapping, ROS-Based Vision Pipelines.

Applications:

Autonomous Drone Navigation, Robot Localisation, Indoor Navigation, Human Activity Recognition and Intelligent Robotics Systems.

AI Integration:

Transformer-Based Robotics Models, Vision-Language Models, AI-Assisted Robot Planning and Deep Learning-Based SLAM Systems.

Authentic Intelligence:

Safe Decision-Making in Autonomous Systems, Accountability in Robot Actions and Human-Centred AI Design for Robotics.

Key Online Resources (Relevant to this Module):

1. An Introduction to AI-Based Robot Development With Isaac ROS https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-ros/latest/an-introduction-to-ai-based-robot-development-with-isaac-ros/index.html
2. Leveraging ROS 2 and Hardware-in-the-Loop (HIL) in Isaac Sim https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/leveraging-ros-2-and-hil-in-isaac-sim/index.html
3. Developing Robots with Software-in-the-Loop (SIL) in Isaac Sim https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/developing-robots-with-sil-in-isaac-sim/index.html

MODULE – V	**09 Hours**

Edge AI and Real-Time Deep Learning Deployment

Topics:

Model Optimisation Techniques, Including Quantisation, Pruning, and Model Compression, TensorRT and ONNX Frameworks, Edge AI and Embedded Deployment, NVIDIA Jetson Platform, Real-Time Inference and Benchmarking, Multi-Camera Systems and Sensor Fusion, Explainable AI (XAI), Safety, Reliability, and Ethical AI Deployment, Cloud and Edge Integration for AI Systems.

Applications:

Real-Time Robotics, Smart Surveillance Systems, Industrial Quality Inspection, Healthcare Monitoring and Edge Computing Systems.

AI Integration:

Generative AI on Edge Devices, Multimodal AI Systems and AI Accelerators Including NPUs.

Authentic Intelligence:

Responsible AI Deployment, Safety Assurance in Edge Devices, Privacy Preservation and Trustworthy AI Systems.

Key Online Resources (Relevant to this Module):

1. Build High-Performance and AI-Enabled Sensor Processing Applications
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-30+V1
2. Accelerating ROS 2 with NVIDIA GPU-Powered Libraries and AI Models
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain deep learning fundamentals, neural networks, and optimisation techniques for AI systems.	L2
2	Apply CNN architectures and computer vision techniques for image classification and recognition tasks.	L3
3	Analyse and implement object detection, segmentation, and 3D vision techniques for robotics applications.	L4
4	Develop sequence models, visual odometry, and SLAM techniques for robotics and autonomous systems.	L4
5	Evaluate and deploy optimised deep learning models on edge and real-time robotic platforms.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	–	2	1	–	2	1	1	1	3	3	3
CO-2	3	3	2	1	3	2	–	3	1	1	1	3	3	3
CO-3	3	3	3	2	3	2	1	3	2	1	1	3	3	3
CO-4	3	2	3	3	3	2	1	3	2	2	2	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. R. Szeliski – Computer Vision: Algorithms and Applications, Springer, 2022.
2. I. Goodfellow – Deep Learning, MIT Press, 2016.

REFERENCE BOOKS:

1. D. Forsyth & J. Ponce – Computer Vision: A Modern Approach, 2nd Ed., Pearson, 2011.
2. OpenCV Documentation – docs.opencv.org
3. Papers With Code – Vision benchmarks: paperswithcode.com

E-RESOURCES:

1. https://pytorch.org/tutorials/
2. https://docs.opencv.org/
3. https://paperswithcode.com
4. https://huggingface.co/docs
5. https://docs.ultralytics.com/
6. https://www.tensorflow.org/tutorials

RECOMMENDED TOOLS:

1. Python 3.10+, NumPy, Pandas, Matplotlib, Seaborn
2. PyTorch, TensorFlow, Hugging Face Transformers
3. OpenCV, Open3D, Detectron2
4. Ultralytics YOLOv8 Framework
5. ROS 2, RVIZ2, ORB-SLAM3
6. NVIDIA Jetson, TensorRT, ONNX Runtime
7. Isaac Sim and NVIDIA Omniverse
8. Jupyter Notebook, Google Colab, VS Code

LAB EXERCISES

S.No	Exercise Title	Aim / Objective	Tools & Platform	Hrs	CO
MODULE I · Exercises 1–5 : Deep Learning Foundations and Neural Networks (12 Hours)					
1	**Python Environment Setup for Deep Learning**	Install and configure Python, Jupyter Notebook, TensorFlow and PyTorch environments for deep learning experiments.	Python, Anaconda, Jupyter Notebook	2 h	CO1
2	**Perceptron and Feedforward Neural Network**	Implement single-layer perceptron and multilayer feedforward neural networks for classification tasks.	Python, NumPy, TensorFlow	2 h	CO1
3	**Activation Functions and Loss Analysis**	Analyse Sigmoid, Tanh, ReLU and Softmax activation functions and compare different loss functions.	PyTorch, Matplotlib	2 h	CO1
4	**Optimisation Techniques in Neural Networks**	Implement SGD, Adam and RMSProp optimisers and compare convergence performance.	TensorFlow, Keras	3 h	CO1
5	**Handwritten Digit Recognition using CNN**	Train a CNN model on MNIST dataset and evaluate classification accuracy.	PyTorch, TensorFlow. Dataset: MNIST	3 h	CO1
MODULE II · Exercises 6–10 : CNN and Visual Recognition (12 Hours)					
6	**Image Preprocessing and Augmentation**	Perform image preprocessing, augmentation, resizing and normalization techniques for computer vision tasks.	OpenCV, Albumentations	2 h	CO2
7	**CNN-Based Image Classification**	Implement CNN architecture for image classification and evaluate model accuracy.	TensorFlow, Keras. Dataset: CIFAR-10	2 h	CO2
8	**Transfer Learning using ResNet50**	Apply transfer learning using pretrained ResNet50 for custom image classification.	PyTorch, HuggingFace	3 h	CO2
9	**Fine-Tuning ViT for Robot Object Recognition**	Fine-tune ViT-B/16 model on robot object datasets and compare with CNN models.	PyTorch, Transformers. Dataset: YCB Object Dataset	3 h	CO2
10	**Face Recognition using CNN**	Develop a face recognition pipeline using CNN feature extraction techniques.	OpenCV, TensorFlow. Dataset: LFW	2 h	CO2

MODULE III · Exercises 11–15 : Object Detection and 3D Vision (12 Hours)					
11	**YOLOv8 Real-Time Object Detection**	Train YOLOv8 model for object detection and evaluate mAP metrics.	Ultralytics YOLOv8, OpenCV	3 h	CO3
12	**Instance Segmentation using Mask R-CNN**	Train Mask R-CNN model for instance segmentation and visualise segmentation masks.	Detectron2, PyTorch	3 h	CO3
13	**Stereo Depth Estimation**	Implement stereo disparity estimation and generate depth maps using stereo images.	OpenCV, PyTorch. Dataset: KITTI Stereo	2 h	CO3
14	**Point Cloud Processing with Open3D**	Process and visualise 3D point cloud data for robotics perception tasks.	Open3D, Python	2 h	CO3
15	**PointNet for 3D Object Classification**	Train PointNet model for 3D object classification and compare results with PointNet++.	PyTorch, Open3D. Dataset: ModelNet40	2 h	CO3
MODULE IV · Exercises 16–20 : Sequence Models and Robotics Vision (12 Hours)					
16	**Sequence Prediction using RNN**	Implement Recurrent Neural Networks for sequence prediction problems.	TensorFlow, Keras	2 h	CO4
17	**Human Activity Recognition using LSTM**	Develop LSTM-based human activity recognition models using sequential sensor data.	PyTorch, TensorFlow	2 h	CO4
18	**Transformer-Based Vision Models**	Implement transformer-based models for image understanding and robotics applications.	HuggingFace Transformers, PyTorch	2 h	CO4
19	**Visual Odometry with ORB-SLAM3**	Execute ORB-SLAM3 on monocular video sequences and visualise robot trajectory.	ORB-SLAM3, ROS 2, RVIZ2	3 h	CO4
20	**ROS 2 Vision Pipeline Integration**	Integrate object detection and visualisation pipelines into ROS 2 robotics framework.	ROS 2, YOLOv8, RVIZ2	3 h	CO4
MODULE V · Exercises 21–25 : Edge AI and Deployment (12 Hours)					
21	**ONNX Model Export and Deployment**	Convert deep learning models to ONNX format and perform inference testing.	ONNX Runtime, PyTorch	2 h	CO5

22	**TensorRT Optimisation for Inference**	Optimise trained models using TensorRT and benchmark inference performance.	TensorRT, NVIDIA Jetson	3 h	CO5
23	**Edge AI Deployment on Jetson Nano**	Deploy computer vision models on NVIDIA Jetson Nano for real-time inference.	NVIDIA Jetson Nano, OpenCV	3 h	CO5
24	**Multi-Camera Sensor Fusion**	Implement sensor fusion techniques using multiple camera inputs for robotics applications.	ROS 2, OpenCV	2 h	CO5
25	**Explainable AI for Vision Models**	Analyse model predictions using Grad-CAM and explainability techniques for trustworthy AI systems.	PyTorch, Grad-CAM	2 h	CO5

GENERATIVE AI AND LARGE LANGUAGE MODELS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - II

Course Code	:		**Credits**	:	5
Hours / Week	:	6	**Total Hours**	:	45(Theory) + 30(lab)+15(Tutorial)
L–T–P–J	:	3–1–2–0	**CIE+SEE**	:	60+40 Marks

Course Vision:

To develop comprehensive expertise in generative AI architectures and large language models, enabling students to build, fine-tune, evaluate, and deploy state-of-the-art generative systems for robotics, engineering, and intelligent applications.

Course Objectives:

This course will enable students to:

1. Understand the theory and architecture of Transformers, diffusion models, and GANs.
2. Fine-tune and align large language models using RLHF, PEFT, and instruction tuning.
3. Build multimodal systems combining vision, language, and action for robotics applications.
4. Develop retrieval-augmented generation (RAG) pipelines and knowledge-grounded intelligent systems.
5. Evaluate, align, and safely deploy generative AI systems with guardrails and trustworthy AI practices.

Teaching-Learning Process (General Instructions)

1. Lecture method along with modern AI-enabled teaching approaches to explain large language models, generative AI, and multimodal learning concepts.
2. Interactive teaching involving brainstorming, discussions, collaborative learning, focused problem-solving, and group activities.
3. Demonstration of transformer architectures, RAG systems, diffusion models, and multimodal AI applications using videos, simulations, and case studies.
4. Encouraging collaborative and project-based learning through implementation of LLM agents, generative AI pipelines, and robotics AI applications.
5. Developing critical thinking through model evaluation, hallucination analysis, alignment strategies, and responsible AI discussions.
6. Demonstrating multiple approaches for solving intelligent automation and generative AI problems and encouraging students to design safe, reliable, and innovative AI systems.

MODULE – I	**09 Hours**

Transformer Architecture and Pretraining

Topics:

Self-Attention and Multi-Head Attention Mechanisms, Positional Encoding Techniques Including Sinusoidal Encoding, Rotary Positional Embedding (RoPE) and ALiBi, Transformer Variants Including Encoder-Only Models (BERT), Decoder-Only Models (GPT) and Encoder-Decoder Models (T5), Scaling Laws and Emergent Abilities in Large Language Models, Pretraining Objectives Including Masked Language Modelling (MLM), Causal Language

Modelling (CLM) and Span Corruption, Efficient Transformer Architectures Including Flash Attention and Sparse Attention.

Applications:

Code Generation Systems, Intelligent Document Summarisation, Conversational AI Systems and Robotic Task Instruction Following.

AI Integration:

LLaMA 3, Mistral and Phi-3 Architectures, Mixture of Experts (MoE) in Mixtral and GPT-4, Large-Scale Foundation Models and Efficient Transformer Deployment.

Authentic Intelligence:

Transparency in attention-based decision making, ethical implications of large language models, bias in transformer pretraining data, and responsible foundation model development.

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-01+V1

MODULE – II	09 Hours

Fine-Tuning and Alignment

Topics:

Supervised Fine-Tuning (SFT) and Instruction Tuning, Parameter-Efficient Fine-Tuning Methods Including LoRA, QLoRA, Adapters and Prefix Tuning, Reinforcement Learning from Human Feedback (RLHF), Proximal Policy Optimisation (PPO), Direct Preference Optimisation (DPO), Constitutional AI and Reinforcement Learning from AI Feedback (RLAIF), Hallucination in Large Language Models Including Causes, Detection and Mitigation Strategies.

Applications:

Domain-Specific Chatbots, AI-Based Code Assistants, Medical Question Answering Systems and Intelligent Enterprise Assistants.

AI Integration:

OpenAI RLHF Pipeline, Anthropic Constitutional AI Framework, Meta LLaMA Alignment Techniques and Human Preference Optimisation Systems.

Authentic Intelligence:

Alignment of AI systems with human values, safe and trustworthy fine-tuning, mitigation of hallucination risks, and ethical human feedback integration.

Key Online Resources (Relevant to this Module):

1. Building AI Agents with Multimodal Models
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

MODULE – III	09 Hours

Retrieval-Augmented Generation and Agents

Topics: Retrieval-Augmented Generation (RAG) Pipeline Including Retrieval, Chunking, Embedding, Reranking and Generation, Vector Databases Including Chroma, FAISS, Pinecone and Weaviate, Sparse and Dense Retrieval

Techniques Including BM25 and Bi-Encoder Architectures, LLM Agents Including ReAct, Tool Use and Function Calling, Memory Systems Including Episodic, Semantic and Procedural Memory, Multi-Agent Frameworks Including AutoGen, CrewAI and LangGraph.

Applications:

Knowledge Base Question Answering Systems, Robot Task Planning Using LLM Agents, Intelligent Research Assistants and Autonomous Code Generation Agents.

AI Integration:

LangChain, LlamaIndex and DSPy Frameworks for LLM Systems, OpenAI Assistants API, Tool-Augmented LLMs and Autonomous Agent Workflows.

Authentic Intelligence:

Reliable knowledge grounding, reducing hallucinations in retrieval systems, ethical use of autonomous agents, and trustworthy tool-integrated AI systems.

Key Online Resources (Relevant to this Module):

1. **Augment Your LLM Using Retrieval Augmented Generation**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-11+V1
2. **Agentic AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-12+V1

MODULE – IV	09 Hours

Generative Models: Diffusion and GANs

Topics:

Generative Adversarial Networks (GANs) Including GAN Training, Mode Collapse and Wasserstein GAN, Variational Autoencoders (VAE), Denoising Diffusion Probabilistic Models (DDPM and DDIM), Latent Diffusion and Stable Diffusion Models, Score-Based Generative Models, Consistency Models and Flow Matching, Evaluation Metrics for Generative Models Including FID, Inception Score (IS) and CLIP Score.

Applications:

Synthetic Data Generation, Robot Simulation Data Augmentation, Image Editing and Image-Conditioned Manipulation Systems.

AI Integration:

Stable Diffusion XL, ControlNet, InstructPix2Pix, AI-Based Video Generation Systems Including Sora and CogVideo.

Authentic Intelligence:

Responsible synthetic content generation, mitigation of deepfake misuse, ethical image generation, and safe deployment of generative models.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-01+V1
2. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

MODULE – V	**09 Hours**

Multimodal LLMs and Robotics Applications

Topics:

Multimodal Architectures Including CLIP, LLaVA, Flamingo and GPT-4V, Vision-Language-Action Models for Robotics Including RT-2 and OpenVLA, Language-Conditioned Robot Learning, LLM-Based Task and Motion Planning Including SayCan and Inner Monologue, Safety and Evaluation of Generative AI Systems, EU AI Act Implications for Generative AI and Robotics Systems.

Applications:

Natural Language Robot Control, Embodied AI Assistants, Visual Question Answering for Robotics Inspection and Autonomous Human-Robot Interaction Systems.

AI Integration:

OpenVLA, RoboFlamingo, pi0 Generalist Robot Policy, Gemini Robotics and Multimodal Foundation Models for Robotics.

Authentic Intelligence:

Safe human-robot interaction, ethical multimodal decision making, grounding language to real-world actions, and responsible deployment of embodied AI systems.

Key Online Resources (Relevant to this Module):

1. **An Introduction to AI-Based Robot Development With Isaac ROS**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-05+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain Transformer architecture, pretraining strategies, and scaling properties of large language models.	L2
2	Fine-tune and align large language models using LoRA, instruction tuning, RLHF, and DPO.	L3
3	Build RAG pipelines and LLM agent systems with tool use and memory integration.	L3
4	Design and train generative models including diffusion models and GANs.	L4
5	Develop multimodal LLM applications for vision-language and robotics tasks.	L3

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. J. Alammar and M. Grootendorst, Hands-On Large Language Models, O'Reilly, 2024.
2. I. Goodfellow, Y. Bengio and A. Courville, Deep Learning, MIT Press, 2016.

REFERENCE BOOKS:

1. S. Bubeck et al., *Sparks of Artificial General Intelligence: Early Experiments with GPT-4*, arXiv, 2023.
2. Hugging Face NLP Course. Available at: https://huggingface.co/learn/nlp-course
3. fast.ai Practical Deep Learning. Available at: https://course.fast.ai/

E-RESOURCES:

1. https://huggingface.co/learn
2. https://huggingface.co/docs
3. https://python.langchain.com/docs/introduction/
4. https://docs.llamaindex.ai/en/stable/
5. https://platform.openai.com/docs/overview
6. https://pytorch.org/tutorials/
7. https://www.tensorflow.org/tutorials
8. https://huggingface.co/docs/diffusers/index
9. https://docs.ragas.io/en/latest/
10. https://course.fast.ai/

RECOMMENDED TOOLS:

1. Python 3.10+, NumPy, Pandas and Matplotlib
2. PyTorch, TensorFlow and Hugging Face Transformers
3. LangChain, LlamaIndex and DSPy

4. Ollama, OpenAI API and Gradio
5. Diffusers, Accelerate, PEFT and bitsandbytes
6. FAISS, ChromaDB and Pinecone
7. ROS 2, Gazebo and Isaac Sim
8. Jupyter Notebook, Google Colab and VS Code

LAB EXERCISES

S.No	Exercise Title	Aim / Objective	Tools & Platform	Hrs	CO
MODULE I · Exercises 1–3 : Transformer Architecture and Pretraining (6 Hours)					
1	Transformer Attention Visualisation	Implement self-attention and multi-head attention mechanisms and visualise attention maps for text sequences.	PyTorch, HuggingFace Transformers, Matplotlib	2 h	CO1
2	Positional Encoding and Transformer Variants	Implement sinusoidal positional encoding and compare encoder-only, decoder-only and encoder-decoder transformer architectures.	PyTorch, TensorFlow	2 h	CO1
3	Language Modelling using GPT-style Architecture	Train a small causal language model for next-word prediction and analyse generated text quality.	HuggingFace Transformers, Python	2 h	CO1
MODULE II · Exercises 4–6 : Fine-Tuning and Alignment (6 Hours)					
4	Fine-Tuning using LoRA and PEFT	Fine-tune a pretrained language model using LoRA and parameter-efficient fine-tuning techniques.	HuggingFace PEFT, Transformers	2 h	CO2
5	Instruction Tuning for Domain-Specific Chatbot	Develop and fine-tune an instruction-following chatbot for domain-specific question answering.	PyTorch, HuggingFace Datasets	2 h	CO2
6	Hallucination Detection and Response Evaluation	Analyse hallucination in LLM outputs and evaluate responses using alignment and safety metrics.	OpenAI API, Python, RAGAS	2 h	CO2
MODULE III · Exercises 7–9 : Retrieval-Augmented Generation and Agents (6 Hours)					
7	Building a Basic RAG Pipeline	Implement a Retrieval-Augmented Generation pipeline using embeddings, chunking and vector databases.	LangChain, ChromaDB, HuggingFace Embeddings	2 h	CO3
8	Document Question Answering using FAISS	Develop a document-based question answering system using FAISS vector database and LLM retrieval.	FAISS, LangChain, Python	2 h	CO3
9	LLM Agent with Tool Use	Create an LLM agent capable of tool usage and function calling for intelligent task execution.	LangChain Agents, OpenAI API	2 h	CO3

MODULE IV · Exercises 10–12 : Diffusion Models and GANs (6 Hours)					
10	Image Generation using GANs	Implement and train a simple GAN model for synthetic image generation tasks.	PyTorch, TensorFlow	2 h	CO4
11	Stable Diffusion Text-to-Image Generation	Generate images using Stable Diffusion models with prompt engineering techniques.	Diffusers, HuggingFace, Python	2 h	CO4
12	Image Editing using ControlNet	Perform image-conditioned editing and generation using ControlNet and diffusion pipelines.	Stable Diffusion XL, ControlNet	2 h	CO4
MODULE V · Exercises 13–15 : Multimodal LLMs and Robotics Applications (6 Hours)					
13	Vision-Language Understanding using CLIP	Implement CLIP-based image-text similarity and multimodal retrieval tasks.	OpenAI CLIP, PyTorch	2 h	CO5
14	Multimodal Chatbot using LLaVA	Develop a multimodal chatbot capable of processing image and text inputs.	LLaVA, HuggingFace Transformers	2 h	CO5
15	Language-Controlled Robotics Simulation	Implement language-guided robot task execution in a simulated robotics environment.	ROS 2, Gazebo, OpenVLA	2 h	CO5

EXPLAINABLE AI (XAI)

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - II

Course Code	:		**Credits**	:	5
Hours / Week	:	5	**Total Hours**	:	45(Theory) + 30(Lab)
L–T–P–J	:	3–2–0–0	**CIE+SEE**	:	60+40

Course Vision:

To develop technically proficient and ethically responsible AI professionals capable of designing, analysing, and deploying transparent, interpretable, and trustworthy artificial intelligence systems for real-world applications.

Course Objectives:

This course will enable students to:

1. Understand the fundamentals of Explainable AI (XAI) and interpretable machine learning models.
2. Analyse transparency, fairness, accountability, and ethical considerations in AI systems.
3. Apply feature attribution, model interpretation, and visual explanation techniques for machine learning and deep learning models.
4. Implement interpreability techniques for NLP, Computer Vision, and tabular data applications.
5. Develop trustworthy and explainable AI systems for real-world deployment.

Teaching-Learning Process (General Instructions)

1. Lecture method along with modern AI-enabled teaching approaches to explain interpretable machine learning and Explainable AI concepts.
2. Interactive teaching involving brainstorming, discussions, focused learning, collaborative problem-solving, and group activities.
3. Demonstration of explainability techniques using real-world AI systems, case studies, visualisations, videos, and simulation tools.
4. Encouraging collaborative and project-based learning through implementation of SHAP, LIME, Grad-CAM, fairness auditing, and interpretability frameworks.
5. Developing critical thinking through ethical AI discussions, fairness evaluation, bias analysis, and comparison of interpretable and black-box models.
6. Demonstrating multiple approaches for explaining AI systems and encouraging students to design trustworthy, transparent, and human-centred AI solutions.

MODULE – I	**09 Hours**

Introduction to Explainable AI and Interpretability

Topics:

Introduction to Artificial Intelligence and Machine Learning, Need for Explainable AI (XAI), Black Box vs White Box Models, Interpretability vs Explainability, Transparency, Trustworthiness and Accountability, Ethical and Responsible AI, Types of Interpretability, Global Interpretability, Local Interpretability, Human-Centred AI, Explainability Challenges in Deep Learning, Overview of XAI Frameworks and Tools.

Applications:

Healthcare Diagnosis Systems, Financial Risk Prediction, Recommendation Systems, Fraud Detection and Intelligent Decision-Making Systems.

AI Integration:

Use of Explainable AI Frameworks, Integration of SHAP and LIME for Model Interpretation, Explainability Dashboards and Human-Centred AI Systems.

Authentic Intelligence:

Understanding Bias in AI Systems, Trustworthy Decision-Making, Ethical Use of AI Predictions and Responsible AI Practices.

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1

MODULE – II	**09 Hours**

Intrinsically Interpretable Models

Topics:

Interpretable Machine Learning Concepts, Linear Regression and Logistic Regression Interpretation, Decision Trees and Rule-Based Learning, Naive Bayes Interpretation, k-Nearest Neighbour Interpretation, Feature Importance Analysis, Partial Dependence Plots (PDP), Individual Conditional Expectation (ICE), Surrogate Models, Model-Agnostic Interpretability, SHAP (Shapley Additive Explanations), LIME (Local Interpretable Model-Agnostic Explanations).

Applications:

Customer Churn Prediction, Loan Approval Systems, Medical Data Analysis, Predictive Maintenance and Business Intelligence Applications.

AI Integration:

Model Interpretation using SHAP and LIME Libraries, Explainability Integration with Scikit-Learn Pipelines and Interactive Feature Visualisation Tools.

Authentic Intelligence:

Fairness in Predictive Models, Transparency in Decision Trees, Responsible Use of Interpretable Models and Bias Detection in Data-Driven Systems.

Key Online Resources (Relevant to this Module):

1. **Fundamentals of Deep Learning**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-01+V3
2. **Accelerating End-to-End Data Science Workflows**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-DS-01+V1

MODULE – III	**09 Hours**

Post-Hoc Local Explanation Methods

Topics:

Explainability in Neural Networks, Feature Visualisation Techniques, Saliency Maps, Gradient-Based Interpretability, CAM and Grad-CAM, Attention Mechanisms, Interpretable CNN Architectures, Explainability for Object Detection Models, Explainability in Vision Transformers (ViTs), Adversarial Attacks and Robustness, Bias Detection in Vision Models.

Applications:

Medical Image Interpretation, Autonomous Vehicle Vision Systems, Facial Recognition Systems, Smart Surveillance and Industrial Defect Detection.

AI Integration:
Grad-CAM Based Visualisation, Explainable Vision Transformers, OpenCV-Based Interpretation Systems and Deep Learning Explainability Frameworks.
Authentic Intelligence:
Responsible Computer Vision Systems, Explainability in Autonomous Decision-Making, Bias Reduction in Vision Datasets and Ethical AI-Based Surveillance.

Key Online Resources (Relevant to this Module):

1. **Introduction to Multi-Modal Data Curation**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:NVIDIA+S-FX-27+V1
2. **Building a Brain in 10 Minutes**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+T-FX-01+V1

MODULE – IV	09 Hours

Post-Hoc Global and Visual Methods
Topics:
Interpretability in Natural Language Processing, Attention and Transformer Explainability, Explainability in Large Language Models (LLMs), Token Attribution Methods, Counterfactual Explanations, Prompt Interpretability, Bias and Hallucination in Generative AI, Explainability in Chatbots and Conversational AI, Explainability Metrics for NLP Models, Ethical Challenges in Generative AI.
Applications:
Chatbots, Conversational AI Systems, Sentiment Analysis, Machine Translation and AI-Assisted Content Generation.
AI Integration:
Transformer-Based Explainability Tools, Hugging Face Interpretability Libraries, Prompt Analysis Frameworks and Explainable Generative AI Systems.
Authentic Intelligence:
Detection of Hallucination in AI Responses, Ethical Use of Generative AI, Bias Mitigation in NLP Models and Human-Centred Conversational Systems.

Key Online Resources (Relevant to this Module):

1. **Building Real-Time Video AI Applications**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-IV-01+V1
2. **Getting Started With AI on NVIDIA Jetson Nano**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-02+V2

MODULE – V	09 Hours

Evaluation, Deployment, and Human Factors
Topics:
Fairness in AI Systems, Bias Detection and Mitigation, AI Governance and Regulations, Robustness and Security in AI, Explainability in Healthcare Applications, Explainability in Finance and Banking, Explainability in Autonomous Systems, Responsible AI Frameworks, Case Studies in Explainable AI, Future Trends in Explainable and Trustworthy AI.
Applications:

Healthcare Diagnosis, Banking and Financial Analytics, Smart Governance Systems, Autonomous Vehicles and AI-Based Security Systems.

AI Integration:

AI Fairness Toolkits, Responsible AI Dashboards, Explainability Monitoring Systems and Real-Time Trustworthy AI Deployment.

Authentic Intelligence:

Transparency in Critical AI Applications, Ethical AI Governance, Privacy Preservation, Accountability in Autonomous Systems and Socially Responsible AI Deployment.

Key Online Resources (Relevant to this Module):

1. **Introduction to Federated Learning with NVIDIA FLARE**
   https://developer.download.nvidia.com/assets/Clara/flare/hroth_nvflare_webinar_Feb2025.pdf
2. **Decentralised AI at Scale with NVIDIA FLARE**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-29+V1
3. **Introduction to NVIDIA NIM Microservices**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-23+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain the fundamentals, principles, and importance of Explainable AI and interpretable machine learning models.	L2
2	Apply model-specific and model-agnostic interpretation techniques for machine learning applications.	L3
3	Analyse deep learning interpretability methods for computer vision and neural network models.	L4
4	Evaluate explainability techniques in NLP, Generative AI, fairness, robustness, and ethical AI systems.	L5
5	Design trustworthy and explainable AI solutions for real-world applications using suitable frameworks and tools.	L6

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	-	-	2	2	-	2	-	-	-	3	3	3
CO-2	3	3	2	2	3	-	-	1	-	-	-	3	3	3
CO-3	3	3	3	3	3	-	-	1	-	-	-	3	3	3
CO-4	2	3	2	3	2	3	2	3	-	-	-	3	3	3
CO-5	3	3	3	3	3	3	2	3	2	2	1	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. C. Molnar – Interpretable Machine Learning, 2nd Ed., Lulu Press, 2022. (Free: christophm.github.io/interpretable-ml-book)
2. W. Samek et al. (Eds.) – Explainable AI: Interpreting, Explaining and Visualizing Deep Learning, Springer LNCS 11700, 2019.

REFERENCE BOOKS:

1. A. Barredo Arrieta et al. – Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges. Information Fusion, 2020.
2. F. Doshi-Velez & B. Kim – Towards A Rigorous Science of Interpretable Machine Learning, arXiv:1702.08608, 2017.

E-RESOURCES:

1. https://shap.readthedocs.io/en/latest/
2. https://lime-ml.readthedocs.io/en/latest/
3. https://captum.ai/
4. https://interpret.ml/docs/
5. https://aif360.readthedocs.io/en/latest/
6. https://fairlearn.org/main/
7. https://huggingface.co/docs/transformers/index
8. https://pytorch.org/tutorials/
9. https://www.tensorflow.org/tutorials

RECOMMENDED TOOLS:

1. Python 3.10+, NumPy, Pandas, Matplotlib, Seaborn
2. Scikit-learn, XGBoost, InterpretML
3. SHAP, LIME, DiCE, Captum
4. PyTorch, TensorFlow, Hugging Face Transfomers
5. IBM AI Fairness 360 (AIF360), Fairlearn
6. Plotly Dash and Streamlit for Explainability Dashboards
7. Jupyter Notebook, Google Colab, VS Code

ADVANCED ROBOTICS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - II

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	45 (Theory)
L–T–P–J	:	3–0–0–0	**CIE+SEE**	:	60 + 40

Course Vision:

To develop deep expertise in modelling, kinematics, dynamics, and intelligent motion planning of robotic systems, enabling students to design, simulate, and deploy advanced robot platforms for real-world applications..

Course Objectives:

This course will enable students to:

1. Analyse forward and inverse kinematics of serial and parallel manipulators using DH convention.
2. Derive and apply Newton-Euler and Lagrangian dynamic models for robotic arms.
3. Design trajectory planning and motion control algorithms for manipulators.
4. Apply force control and compliance strategies for robot-environment interaction.
5. Evaluate advanced robot configurations including parallel robots and continuum robots.

Teaching-Learning Process (General Instructions)

These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes.

1. Lecture method along with simulation-driven teaching approaches to explain advanced robotics concepts including kinematics, dynamics, motion planning, and robot control.
2. Interactive teaching involving brainstorming, discussions, collaborative learning, analytical problem-solving, and group activities.
3. Demonstration of robotic systems, simulations, and autonomous navigation using videos, animations, digital twins, and industrial case studies.
4. Encouraging project-based and research-oriented learning through robot modelling, trajectory planning, and AI-enabled robotic system design.
5. Developing critical thinking through motion planning problems, dynamic analysis, controller design, and robot-environment interaction studies.
6. Demonstrating multiple approaches for solving robotics problems and encouraging students to design innovative, reliable, and human-centered robotic solutions.

MODULE – I	9 Hours

Robot Kinematics and Spatial Descriptions

Topics: Homogeneous Transformations and Denavit-Hartenberg (DH) Convention; Forward Kinematics of Serial Manipulators; Inverse Kinematics: Analytical and Numerical Methods; Jacobian Matrix and Singularity Analysis; Workspace Analysis and Reachability; Parallel Manipulator Kinematics (Stewart Platform).

Applications: Industrial Pick-and-Place Arms, Surgical Robots (da Vinci), and Delta Robots for High-Speed Packaging.

AI Integration: AI-Assisted Inverse Kinematics Solvers Using Neural Networks; Reinforcement Learning for Kinematic Calibration.

Authentic Intelligence: Understanding Singularities and Workspace Limitations, Reliable Robot Motion Planning, Safety in Human-Robot Interaction, and Trustworthy Robot Manipulation Systems.

Key Online Resources (Relevant to this Module):

1. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **A Beginner's Guide to Autonomous Robots**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-35+V1

MODULE – II	9 Hours

Robot Dynamics and Simulation

Topics: Newton-Euler Recursive Formulation; Lagrangian Dynamics and Equations of Motion; Inertia Tensor and Dynamic Parameters; Gravity, Coriolis, and Centrifugal Terms; Dynamic Simulation Using Gazebo and MuJoCo; Parameter Identification and Calibration.

Applications: High-Speed Industrial Manipulators, Space Robot Dynamics, and Legged Robot Systems.

AI Integration: Physics-Informed Neural Networks for Dynamic Modelling; Machine Learning-Based Parameter Identification.

Authentic Intelligence: Physically Interpretable Dynamic Models, Reliable Robot Simulation, Safety-Aware Dynamic Analysis, and Trustworthy AI-Assisted Robot Modelling.

Key Online Resources (Relevant to this Module):

1. **Developing Robots with Software-in-the-Loop (SIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/developing-robots-with-sil-in-isaac-sim/index.html
2. **Leveraging ROS 2 and Hardware-in-the-Loop (HIL) in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-sim/latest/leveraging-ros-2-and-hil-in-isaac-sim/index.html
3. **Train Your First Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-first-robot-with-isaac-lab/index.html

MODULE – III	9 Hours

Trajectory Planning and Motion Control

Topics: Joint-Space and Cartesian-Space Trajectory Generation; Polynomial and Spline Interpolation; Via-Point Trajectory Planning; Minimum-Jerk and Time-Optimal Trajectories; PID and Computed Torque Control; Impedance and Admittance Control Fundamentals.

Applications: CNC Machining, Robotic Welding Path Planning, and Rehabilitation Exoskeleton Control.

AI Integration: Learning from Demonstration for Trajectory Generation; Gaussian Process Regression for Trajectory Optimisation.

Authentic Intelligence: Safe Motion Generation, Reliable Human-Robot Interaction, Interpretable Motion Planning, and Ethical Deployment of Autonomous Robotic Systems.

Key Online Resources (Relevant to this Module):

1. **Transferring Robot Learning Policies from Simulation to Reality** https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/transferring-robot-learning-policies-from-simulation-to-reality/index.html
2. **Train Your Second Robot in Isaac Lab** https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-second-robot-with-isaac-lab/index.html

MODULE – IV	9 Hours

Advanced Motion Planning

Topics: Configuration Space and C-Space Obstacles; Sampling-Based Planners Including PRM, RRT, and RRT*; Potential Field Methods; Graph Search Algorithms Including A* and D* Lite; Dynamic Motion Primitives (DMPs); Multi-Robot Coordination and Collision Avoidance.

Applications: Autonomous Mobile Manipulation, Surgical Path Planning, and Warehouse Multi-Robot Systems.

AI Integration: Learning-Based Motion Planners; Neural RRT; Reinforcement Learning for Motion Planning.

Authentic Intelligence: Safe Autonomous Navigation, Transparent Path Planning Decisions, Reliable Multi-Robot Coordination, and Ethical Deployment of Autonomous Systems.

Key Online Resources (Relevant to this Module):

1. **Train Your Second Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-second-robot-with-isaac-lab/index.htmll
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **Scalable Multi-Robot Scene Workflows Using ROS Simulation Interfaces Standard in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/scalable-multi-robot-scene-workflows-using-ros-simulation-interfaces-standard-in-isaac-sim/index.html

MODULE – V	**12 Hours**

Advanced Robot Mechanisms and Soft Robotics

Topics: Parallel Robots: Kinematics and Workspace; Tendon-Driven and Cable-Driven Mechanisms; Continuum and Soft Robot Modelling; Compliant Mechanisms and Flexure Joints; Tactile Sensing and Force/Torque Sensors; Wearable and Exoskeleton Robotics.

Applications: Soft Grippers for Food Handling, Continuum Robots for Minimally Invasive Surgery, and Wearable Rehabilitation Devices.

AI Integration: Model-Free Control for Soft Robots; Deep Learning for Deformable Object Manipulation.

Authentic Intelligence: Safe Human-Centered Soft Robotics, Reliable Physical Interaction, Ethical Use of Assistive Exoskeletons, and Trustworthy Adaptive Robot Systems.

Key Online Resources (Relevant to this Module):

1. **Training Healthcare Robots From Scratch With Isaac for Healthcare**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-for-healthcare/latest/training-healthcare-robots-from-scratch/index.html
2. **Best Practices for Robotics and OpenUSD**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/best-practices-for-robotics-and-openusd/index.html
3. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Analyse forward/inverse kinematics and dynamic models of serial and parallel manipulators.	L4
2	Derive Newton-Euler and Lagrangian dynamic equations and simulate robotic systems.	L3
3	Design trajectory planning algorithms in joint and Cartesian space.	L6
4	Apply advanced motion planning algorithms (RRT*, A*) for collision-free path planning.	L3
5	Evaluate advanced robot mechanisms including parallel robots and soft robotic systems.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	-	1	-	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	-	2	1	3	2	2	2	3	3	3

CO-4	3	2	3	-	3	3	-	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	-	-	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. J.J. Craig – Introduction to Robotics: Mechanics and Control, 3rd Ed., Pearson, 2005.
2. B. Siciliano et al. – Robotics: Modelling, Planning and Control, Springer, 2010.

REFERENCE BOOKS:

1. K.M. Lynch & F.C. Park – Modern Robotics: Mechanics, Planning, and Control, Cambridge, 2017. (Free online: modernrobotics.org)
2. M.W. Spong, S. Hutchinson & M. Vidyasagar – Robot Modeling and Control, 2nd Ed., Wiley, 2020.

E-RESOURCES:

1. NPTEL – Robotics Course
2. MIT OpenCourseWare – Underactuated Robotics
3. ROS 2 Documentation/

RECOMMENDED TOOLS

1. MATLAB/Simulink Robotics System Toolbox
2. ROS2, MoveIt!, Gazebo, MuJoCo
3. NVIDIA Isaac Sim and Isaac Lab
4. Python 3.10+, NumPy, SciPy, PyTorch, OpenCV
5. RViz, Webots, CoppeliaSim (V-REP)

PERCEPTION AND SENSOR FUSION FOR AUTONOMOUS ROBOTIC SYSTEMS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER – II

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	45
L–T–P–J	:	3–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To develop engineers proficient in designing and implementing multi-sensor perception systems for autonomous robots — integrating camera, LiDAR, RADAR, and IMU sensing with deep learning-based detection, segmentation, and 3D scene understanding — and leveraging NVIDIA Isaac Sim, OpenUSD, and NVIDIA Cosmos for synthetic perception data generation and digital twin-based system validation, enabling robust and safe autonomous robot operation in unstructured real-world environments.

Course Objectives: This course will enable students to:

1. **Understand** the sensor modalities used in autonomous robotic perception — camera geometry, LiDAR point clouds, RADAR range-Doppler, and IMU dynamics — and their complementary characteristics.
2. **Apply** deep learning architectures for camera-based 2D/3D object detection, semantic segmentation, depth estimation, and instance tracking in robotic scenes.
3. **Implement** 3D LiDAR perception pipelines using point cloud processing networks and sensor fusion architectures for bird's-eye view (BEV) scene understanding.
4. **Design** multi-sensor fusion systems integrating camera, LiDAR, and IMU data using OpenUSD scene representation and NVIDIA Cosmos-generated synthetic training data.
5. **Evaluate** perception system performance, safety, and robustness using digital twins in NVIDIA Isaac Sim and responsible AI deployment principles for autonomous systems.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Lectures on perception theory with detailed mathematical derivations of camera geometry, point cloud processing, and fusion architectures supported by visual examples.
2. **Lab Sessions:** Hands-on PyTorch perception model implementation and training on nuScenes, KITTI, and Waymo Open Dataset benchmarks.
3. **Isaac Sim Labs:** NVIDIA Isaac Sim sessions for synthetic perception data generation and digital twin-based perception system testing.
4. **OpenUSD Workshops:** Structured OpenUSD scene composition exercises for building 3D robot environment representations used in perception pipelines.
5. **Paper Reading Club:** Weekly review of landmark perception papers — PointPillars, BEVFusion, DETR3D, DINOv2, SAM, and NVIDIA's DriveX perception stack.
6. **Project:** Teams build and evaluate a complete multi-sensor perception pipeline for a chosen autonomous robot scenario — indoor navigation or outdoor field robot.

MODULE – I	**08 Hours**
Sensor Modalities and Camera Geometry for Robot Perception	

Topics: Autonomous robot perception overview: the sense-plan-act loop; Sensor modalities comparison: camera (rich texture, no depth), LiDAR (accurate 3D, no colour), RADAR (all-weather, velocity), ultrasonic (low-cost proximity), event cameras (high dynamic range); Camera models: pinhole model, projection matrix, homogeneous coordinates; Intrinsic calibration: focal length, principal point, distortion coefficients — Zhang's calibration method; Extrinsic calibration: rotation-translation, camera-to-robot transform, multi-camera rig calibration; Stereo vision: epipolar geometry, fundamental and essential matrices, stereo rectification, disparity computation; Structured light depth sensing: Intel RealSense, Azure Kinect — IR pattern projection and depth decoding; Event cameras: dynamic vision sensor (DVS) — asynchronous event stream processing for high-speed robot perception; NVIDIA Isaac ROS camera node and camera calibration tools.

Applications: Zhang's checkerboard calibration of a stereo camera rig on a mobile robot — computing intrinsics and extrinsics using OpenCV; Stereo disparity map generation at 30 FPS using NVIDIA Isaac ROS stereo depth node; Event camera-based obstacle detection for a drone at 1000 FPS — outperforming frame-based camera at high speed.

AI Integration: Assembling Your First Digital Twin with Omniverse and OpenUSD (NVIDIA DLI) — setting up perception test environments in Isaac Sim; NVIDIA Isaac ROS camera perception pipeline with GPU-accelerated stereo depth; An Introduction to NVIDIA Cosmos for Physical AI (NVIDIA DLI) — world model context for perception system simulation.

Authentic Intelligence: Privacy implications of always-on robot cameras in shared human spaces; Sensor failure modes and graceful degradation in safety-critical autonomous systems; Calibration drift in deployed robots — responsibility for regular recalibration in production; Equitable perception performance across diverse lighting conditions, weather, and environments.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Assembling Your First Digital Twin with Omniverse and OpenUSD [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Learn OpenUSD: Creating Composition Arcs [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to NVIDIA Cosmos for Physical AI [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – II	**09 Hours**

Camera-Based 2D and 3D Object Perception

Topics: 2D object detection review: anchor-based (YOLOv8, YOLO-NAS), anchor-free (FCOS, CenterNet), transformer-based (DETR, RT-DETR); Tracking: SORT, DeepSORT, ByteTrack, GHOST — multi-object tracking in robot perception; Semantic segmentation: FCN, U-Net, DeepLabV3+, Mask2Former for robot scene understanding; Panoptic segmentation: combining semantic and instance segmentation; Open-vocabulary detection: Grounding DINO (GDINO), OWLv2 — detecting arbitrary object categories from text prompts; Foundation model-based segmentation: SAM (Segment Anything Model) in robotic perception; Monocular depth estimation: DepthAnything V2, Metric3D, ZoeDepth — scale-ambiguous vs. metric depth; 3D object detection from monocular camera: FCOS3D, PGD, GUP-Net — lifting 2D to 3D; Multi-camera 3D detection: DETR3D, PETR, BEVFormer — bird's-eye view from surround cameras; NVIDIA Isaac ROS perception modules: object detection, segmentation, and visual SLAM.

Applications: RT-DETR pedestrian and obstacle detection at 60 FPS on NVIDIA Jetson AGX Orin for a hospital service robot; ByteTrack multi-object tracking of 20 simultaneous targets in a warehouse robot camera stream; BEVFormer surround-view 3D object detection on nuScenes benchmark — 56.9 NDS score.

AI Integration: Learn OpenUSD: Asset Structure Principles and Content Aggregation (NVIDIA DLI) — building 3D perception test assets in OpenUSD; NVIDIA Isaac ROS detection and segmentation nodes for GPU-accelerated robot perception; Synthetic Data Generation for Perception Model Training in Isaac Sim — generating labelled training data for 2D/3D detection.

Authentic Intelligence: Open-vocabulary detection and the risk of robots acting on incorrect natural language object descriptions; Segmentation model failure in out-of-distribution scenes — industrial vs. domestic robot perception transfer; Privacy of individuals detected and tracked by autonomous service robots; Transparency of perception system confidence to human operators supervising autonomous robots.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Learn OpenUSD: Asset Structure Principles and Content Aggregation [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Assembling Your First Digital Twin with Omniverse and OpenUSD [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to NVIDIA Cosmos for Physical AI [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – III	**09 Hours**

LiDAR and RADAR Perception for 3D Scene Understanding

Topics: LiDAR fundamentals: time-of-flight vs. FMCW, spinning vs. solid-state, Velodyne HDL-64E, Ouster OS1, Livox HAP; Point cloud representation: unordered 3D points, range images, voxel grids, pillars, octrees; Point cloud pre-processing: voxel grid downsampling, statistical outlier removal, normal estimation (Open3D, PCL); 3D LiDAR object detection networks: PointNet, PointNet++, VoxelNet, SECOND, PointPillars, CenterPoint, VoxelNeXt; LiDAR semantic segmentation: RandLA-Net, SPVCNN, Cylinder3D for outdoor scene labelling; RADAR fundamentals: FMCW radar, range-Doppler-angle estimation, CFAR detection, radar cross-section; RADAR object detection: RadarNet, CRN (Camera-Radar Net), RCBEV for robust all-weather perception; Occupancy grid mapping: Bayesian occupancy from LiDAR and RADAR for navigation; LiDAR datasets: KITTI, nuScenes, Waymo Open Dataset, Argoverse 2.

Applications: PointPillars 3D car detection on KITTI dataset — 77.3 AP (moderate) at 62 FPS; CenterPoint pedestrian and cyclist detection on nuScenes with velocity estimation for trajectory prediction; FMCW radar + camera fusion for detecting occluded pedestrians in rain — outperforming camera-only by 34% recall.

AI Integration: Learn OpenUSD: Developing Data Exchange Pipelines (NVIDIA DLI) — structuring LiDAR and camera data in USD for perception pipelines; NVIDIA Cosmos-generated synthetic LiDAR point clouds for rare scenario augmentation (adverse weather, night, sensor failure); NVIDIA Isaac Sim LiDAR sensor simulation for point cloud perception model training.

Authentic Intelligence: LiDAR sensing in public spaces — privacy of 3D scanned individuals and unintended mapping of private property; RADAR reliability vs. camera — trade-offs in autonomous safety for vulnerable road users; Sensor degradation disclosure — knowing when autonomous perception is operating outside safe performance bounds; Responsible occupancy grid design — avoiding discrimination in robot navigation around people.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Learn OpenUSD: Developing Data Exchange Pipelines [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Assembling Your First Digital Twin with Omniverse and OpenUSD [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to NVIDIA Cosmos for Physical AI [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – IV	**10 Hours**

Multi-Sensor Fusion and Bird's-Eye View Perception

Topics: Sensor fusion motivation: complementary strengths — camera texture + LiDAR geometry + RADAR velocity + IMU dynamics; Fusion paradigms: early fusion (raw data), intermediate fusion (feature-level), late fusion (decision-level) — trade-offs; Camera-LiDAR calibration: extrinsic estimation, point-to-pixel projection, LiDAR-camera geometric alignment; Camera-LiDAR fusion architectures: PointFusion, F-PointNet (frustum-based), BEVFusion (parallel backbone), TransFusion; Bird's-Eye View (BEV) unified representation: perspective-to-BEV transformation using IPM, LSS (Lift-Splat-Shoot), BEVDet; IMU pre-integration for LiDAR motion compensation (LIO-SAM, FAST-LIO2); Occupancy networks: OccNet, OpenOccupancy — 3D volumetric scene understanding; 4D occupancy flow: predicting future occupancy for robot trajectory planning; Temporal fusion: aggregating multi-frame perception for improved detection of slow-moving objects; OpenUSD for multi-sensor scene composition: structuring camera, LiDAR, and RADAR data streams in USD.

Applications: BEVFusion camera-LiDAR fusion on nuScenes — 72.9 mAP, outperforming camera-only (56.0) and LiDAR-only (65.3); Lift-Splat-Shoot BEV representation from 6 surround cameras enabling parking space occupancy detection; IMU-LiDAR motion compensation reducing point cloud smearing during fast robot turns — 23% improvement in detection at 2m/s.

AI Integration: Assembling Your First Digital Twin with Omniverse and OpenUSD (NVIDIA DLI) — multi-sensor digital twin for fusion system testing; An Introduction to NVIDIA Cosmos for Physical AI (NVIDIA DLI) — world model for multi-sensor scene generation; NVIDIA DriveX perception stack for GPU-accelerated multi-sensor fusion in autonomous systems.

Authentic Intelligence: Fusion system failure modes: what happens when one sensor fails in a multi-sensor fusion pipeline; Safety certification of fusion systems — how to validate that fusion improves not just average but worst-case safety; Thermal imaging for night perception — ethical implications of continuous nighttime surveillance capability; Responsible synthetic data generation — avoiding bias in simulated training environments.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Assembling Your First Digital Twin with Omniverse and OpenUSD [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **An Introduction to NVIDIA Cosmos for Physical AI [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Learn OpenUSD: Creating Composition Arcs [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – V	**09 Hours**

Digital Twins, Synthetic Data, and Responsible Perception Deployment

Topics: Digital twin for perception testing: NVIDIA Isaac Sim environments — photorealistic rendering, physics-accurate sensors; Synthetic data generation for perception: domain randomisation — lighting, weather, object placement, texture variation; NVIDIA Cosmos world foundation model: generating diverse and physically plausible synthetic scenes for perception training; OpenUSD for asset standardisation: reusable sensor configurations, scene composition, and perception data pipelines; Sim-to-real transfer for perception: domain gap mitigation using photorealistic rendering and reality augmentation (RA); Perception evaluation metrics: AP, mAP, nuScenes metrics (NDS), KITTI benchmark; Perception safety testing: edge case generation, adversarial patch attacks on LiDAR and camera, sensor spoofing; Uncertainty quantification for perception: ensemble detectors, conformal prediction for robot-safe bounding boxes; Foundation model perception: DINOv2, Grounding DINO as zero-shot perception backbone for new robot environments.

Applications: NVIDIA Isaac Sim domain randomisation: training a warehouse robot perception model on 100k synthetic scenes — achieving 89% real-world AP with zero real labels; Adversarial patch attack on a YOLOv8 pedestrian detector — physical printout causing complete detection failure — and mitigation; Conformal prediction bounding box calibration achieving guaranteed 95% coverage with tighter intervals than standard YOLO output.

AI Integration: An Introduction to NVIDIA Cosmos for Physical AI (NVIDIA DLI) — world-model-based synthetic scene generation for perception; Assembling Your First Digital Twin with Omniverse and OpenUSD (NVIDIA DLI) — perception system digital twin validation; Learn OpenUSD: Developing Data Exchange Pipelines (NVIDIA DLI) — perception data pipeline with USD.

Authentic Intelligence: Synthetic data bias: if the simulator is biased, all models trained on it inherit the bias; Adversarial perception attacks as safety risks in deployed autonomous systems — responsible disclosure; Transparency about simulation-trained vs. real-world-validated perception systems to end users; Long-term safety monitoring of deployed robot perception systems — drift, environment change, and sensor aging.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **An Introduction to NVIDIA Cosmos for Physical AI [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Assembling Your First Digital Twin with Omniverse and OpenUSD [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Learn OpenUSD: Developing Data Exchange Pipelines [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Describe sensor modalities for robot perception — camera geometry, LiDAR point clouds, RADAR range-Doppler, and IMU dynamics — and their complementary characteristics.	**L2**

2	Apply CNN, DETR, SAM, and BEVFormer architectures for 2D/3D object detection, semantic segmentation, and monocular depth estimation in robotic scenes.	**L3**
3	Implement 3D LiDAR perception pipelines using PointPillars, CenterPoint, and camera-LiDAR BEVFusion for bird's-eye view autonomous robot scene understanding.	**L3**
4	Design multi-sensor fusion systems and OpenUSD-based 3D scene representations using NVIDIA Cosmos synthetic data for robust autonomous robot perception.	**L6**
5	Evaluate perception systems using digital twins in NVIDIA Isaac Sim, adversarial robustness testing, and responsible AI deployment principles for safe autonomy.	**L5**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**1**	**2**	**3**
CO-1	3	2	2	1	2	2	–	3	–	–	2	3	3	3
CO-2	3	3	3	2	3	2	–	3	2	–	2	3	3	3
CO-3	3	3	3	2	3	3	–	3	2	1	2	3	3	3
CO-4	3	3	3	3	3	3	–	2	2	1	2	3	3	3
CO-5	2	2	2	2	3	3	3	3	2	2	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build multi-sensor perception systems integrating camera, LiDAR, RADAR, and IMU sensing with deep learning for robust autonomous robot operation.

PSO-2: Develop GPU-accelerated perception pipelines using NVIDIA Isaac Sim, OpenUSD, Cosmos, and Isaac ROS for synthetic data generation and digital twin perception validation.

PSO-3: Apply perception and sensor fusion expertise to enable safe, robust, and responsible autonomous robot operation in industrial, agricultural, healthcare, and public environments.

TEXT BOOKS:

1. *Roland Siegwart, Illah R. Nourbakhsh & Davide Scaramuzza, Introduction to Autonomous Mobile Robots, 2nd Edition, MIT Press, 2011.*

2. *Simon J. D. Prince, Computer Vision: Models, Learning, and Inference, Cambridge University Press, 2012. (Free at computervisionmodels.com)*

3. *Charles R. Qi, Hao Su et al., PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation, CVPR 2017.*

REFERENCE BOOKS:

1. NVIDIA Isaac Sim Perception Documentation: https://developer.nvidia.com/isaac/sim
2. OpenUSD Documentation: https://openusd.org/docs/
3. nuScenes Benchmark: https://www.nuscenes.org/

E-RESOURCES:

1. NVIDIA DLI — Industrial Digital Twins: Assembling OpenUSD Data for Digital Twins: https://learn.nvidia.com/en-us/training/learning-paths
2. NVIDIA Cosmos: https://www.nvidia.com/en-us/ai/cosmos/
3. KITTI Vision Benchmark: https://www.cvlibs.net/datasets/kitti/
4. Waymo Open Dataset: https://waymo.com/open/
5. Open3D — 3D Data Processing: https://www.open3d.org/

RECOMMENDED TOOLS:

1. Perception Frameworks: PyTorch, MMDetection3D, OpenPCDet, Detectron2, Ultralytics
2. 3D Processing: Open3D, PCL (Point Cloud Library), PyVista, k3d
3. Simulation: NVIDIA Isaac Sim, Gazebo, CARLA, AirSim
4. Digital Twin: NVIDIA Omniverse, OpenUSD, NVIDIA Cosmos
5. ROS 2 Perception: NVIDIA Isaac ROS (detection, segmentation, stereo depth)
6. Benchmarks: nuScenes devkit, KITTI, Waymo Open Dataset, Argoverse 2 API

INTELLIGENT AGENTS & MULTI-AGENT SYSTEMS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - II

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	45
L–T–P–J	:	3–0–0–0	**CIE+SEE**	:	60+40

Course Vision:
To develop technically proficient and ethically responsible professionals capable of designing, analysing, and deploying intelligent agent and multi-agent systems for robotics, distributed AI, and autonomous decision-making applications.

Course Objectives:
This course will enable students to:

1. Understand the theoretical foundations of intelligent agents, environment models, and agent architectures.
2. Analyse game-theoretic models and coordination mechanisms in multi-agent systems.
3. Apply reinforcement learning techniques for cooperative and competitive multi-agent environments.
4. Design communication, coordination, and distributed problem-solving mechanisms in autonomous agent systems.
5. Develop and deploy multi-agent frameworks for robotics, swarm intelligence, and distributed AI applications.

Teaching-Learning Process (General Instructions)
These are sample innovative pedagogical methods that may be adopted to accelerate the attainment of the various course outcomes:

1. Lecture method along with modern AI-enabled teaching approaches to explain intelligent agents, game theory, and multi-agent coordination concepts.
2. Interactive teaching involving brainstorming, discussions, collaborative learning, focused problem-solving, and group activities.
3. Demonstration of autonomous agents, reinforcement learning, and robotic coordination systems using videos, simulations, and case studies.
4. Encouraging collaborative and research-oriented learning through analysis of swarm robotics, MARL systems, and distributed AI frameworks.
5. Developing critical thinking through strategic decision-making problems, coordination challenges, and multi-agent optimisation tasks.
6. Demonstrating multiple approaches for solving distributed AI and robotics coordination problems and encouraging students to design innovative multi-agent solutions.

MODULE – I	**09 Hours**

Agent Foundations and Architectures
Topics:
Introduction to Intelligent Agents, Rational Agents and Performance Measures, PEAS Description, Environment Types Including Observable, Deterministic, Episodic and Static Environments, Agent Types Including Reactive,

Deliberative, Hybrid and BDI Agents, Reactive Architectures Including Subsumption and Behaviour-Based Systems, Deliberative Planning Using STRIPS and PDDL, Beliefs, Desires and Intentions in BDI Models, Agent Rationality and Decision Making.

Applications:

Autonomous Robot Navigation, Warehouse Automation Systems, Intelligent Virtual Assistants and Game-Playing Agents.

AI Integration:

LLM-Based BDI Agents, Neuro-Symbolic Agent Architectures, AI-Based Planning Systems and Autonomous Decision-Making Agents.

Authentic Intelligence:

Ethical Constraints in Autonomous Agent Decision Making, Trustworthy Goal-Oriented Behaviour, Rationality vs Optimality and Responsible Autonomous Planning.

Key Online Resources (Relevant to this Module):

1. **Agentic AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-39+V1

MODULE – II	**09 Hours**

Game Theory and Multi-Agent Interaction

Topics:

Game Theory Foundations, Utility and Payoff Representation, Normal and Extensive Form Games, Dominant Strategies, Nash Equilibrium, Cooperative and Non-Cooperative Games, Mechanism Design and Auctions, VCG Mechanisms, Cooperative Game Theory and Shapley Values, Coalition Formation, Voting Mechanisms and Social Choice Theory, Argumentation Frameworks and Agent Negotiation Models.

Applications:

Robot Task Allocation, Autonomous Trading Agents, Intelligent Traffic Management Systems and Distributed Resource Allocation.

AI Integration:

Multi-Agent Reinforcement Learning Foundations, LLM-Based Negotiation Agents, AI-Based Strategic Decision Making and Autonomous Economic Agents.

Authentic Intelligence:

Fair Resource Allocation, Ethical Negotiation Strategies, Transparency in Autonomous Economic Decisions and Accountability in Multi-Agent Coordination.

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

MODULE – III	**09 Hours**

Multi-Agent Reinforcement Learning

Topics:

Reinforcement Learning Fundamentals, Stochastic Games and Markov Games, Independent Q-Learning, Cooperative and Competitive Multi-Agent Reinforcement Learning, Centralised Training and Decentralised

Execution (CTDE), MADDPG and QMIX Architectures, Credit Assignment Problem in MARL, Communication Models Including DIAL and CommNet, Emergent Communication and Coordination in Multi-Agent Systems.

Applications:

Swarm Robotics, Multi-Robot Warehouse Logistics, Competitive Game Agents Including AlphaStar and OpenAI Five and Autonomous Drone Coordination.

AI Integration:

OpenAI Multi-Agent Particle Environment, PettingZoo Multi-Agent Reinforcement Learning Framework, StarCraft Multi-Agent Challenge (SMAC) and AI-Based Cooperative Learning Systems.

Authentic Intelligence:

Responsible Coordination Among Autonomous Agents, Fairness in Cooperative Learning, Safe Exploration Strategies and Accountability in Autonomous Multi-Agent Decision Making.

Key Online Resources (Relevant to this Module):

1. **An Introduction to Robot Learning and Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/an-introduction-to-robot-learning-and-isaac-lab/index.html
2. **Train Your First Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-first-robot-with-isaac-lab/index.html
3. **Train Your Second Robot in Isaac Lab**
   https://docs.nvidia.com/learning/physical-ai/getting-started-with-isaac-lab/latest/train-your-second-robot-with-isaac-lab/index.html

MODULE – IV	**09 Hours**

Agent Communication and Coordination

Topics:

Agent Communication Languages Including FIPA ACL, Communication Performatives and Message Passing Mechanisms, Ontologies and Shared Knowledge Bases, Coordination Mechanisms Including Blackboard and Mediator Architectures, Contract Net Protocol, Consensus Algorithms, Distributed Constraint Optimisation Problems (DCOP), Max-Sum Algorithm and Byzantine Fault Tolerance in Distributed Systems.

Applications:

Multi-Robot SLAM Systems, Distributed Sensor Networks, Intelligent Supply Chain Automation and Smart Manufacturing Systems.

AI Integration:

LLM-Based Multi-Agent Communication, Structured JSON Message Protocols, Distributed AI Coordination Systems and Collaborative Autonomous Agents.

Authentic Intelligence:

Trustworthy Agent Communication, Fault-Tolerant Distributed Coordination, Ethical Information Sharing and Reliable Multi-Agent Collaboration.

Key Online Resources (Relevant to this Module):

1. **Introduction to Federated Learning with NVIDIA FLARE**
   https://developer.download.nvidia.com/assets/Clara/flare/hroth_nvflare_webinar_Feb2025.pdf
2. **Decentralised AI at Scale with NVIDIA FLARE**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-29+V1

MODULE – V	09 Hours

MAS Platforms and Robotics Applications

Topics:

JADE and SPADE Multi-Agent Platforms, Agent Lifecycle Management, ROS 2 Multi-Robot Systems, Swarm Robotics Concepts Including Stigmergy and Self-Organisation, Task Allocation Methods Including Auction-Based and Market-Based Allocation, Formation Control and Flocking Behaviour, Human-Agent Interaction Design, Multi-Robot Coordination and Distributed Robotic Intelligence.

Applications:

Drone Swarms, Multi-Robot Search and Rescue Operations, Collaborative Manufacturing Systems and Intelligent Autonomous Robotics.

AI Integration:

AutoGen and CrewAI for Multi-Agent Software Systems, ROS 2 Based Multi-Robot Navigation, AI-Based Swarm Coordination and Autonomous Collaborative Robotics.

Authentic Intelligence:

Human-Centred Multi-Agent Design, Ethical Swarm Coordination, Safety in Multi-Robot Deployment and Responsible Human-Agent Collaboration.

Key Online Resources (Relevant to this Module):

1. **A Beginner's Guide to Autonomous Robots**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-35+V1
2. **Develop, Simulate, and Deploy Robot Intelligence With General Robotics**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/develop-simulate-and-deploy-intelligence-with-general-robotics/index.html
3. **Scalable Multi-Robot Scene Workflows Using ROS Simulation Interfaces in Isaac Sim**
   https://docs.nvidia.com/learning/physical-ai/going-further-with-robotics/latest/scalable-multi-robot-scene-workflows-using-ros-simulation-interfaces-standard-in-isaac-sim/index.html

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Explain intelligent agent architectures including reactive, deliberative, hybrid, and BDI models.	L2
2	Apply planning, game-theoretic, and coordination techniques in multi-agent environments.	L3
3	Analyse reinforcement learning approaches for cooperative and competitive multi-agent systems.	L4
4	Design agent communication and distributed coordination mechanisms for autonomous systems.	L5

5	Develop and deploy multi-agent systems for robotics, swarm intelligence, and distributed AI applications.	L6

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**1**	**2**	**3**
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	3	3	3	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles.

TEXT BOOKS:

1. S. Russell & P. Norvig – Artificial Intelligence: A Modern Approach, 4th Ed., Pearson, 2020.
2. M. Wooldridge – An Introduction to MultiAgent Systems, 2nd Ed., Wiley, 2009.

REFERENCE BOOKS:

1. Y. Shoham & K. Leyton-Brown – Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge, 2009. (Free: masfoundations.org)
2. PettingZoo Multi-Agent RL Library – pettingzoo.farama.org

E-RESOURCES:

1. https://masfoundations.org/
2. https://pettingzoo.farama.org/
3. https://gymnasium.farama.org/
4. https://stable-baselines3.readthedocs.io/en/master/
5. https://pytorch.org/tutorials/
6. https://huggingface.co/docs
7. https://openai.com/research/
8. https://paperswithcode.com/

RECOMMENDED TOOLS:

1. Python 3.10+, NumPy, Pandas, Matplotlib
2. OpenAI Gymnasium, PettingZoo, Stable-Baselines3
3. PyTorch, TensorFlow, RLlib

4. ROS 2, RVIZ2, Gazebo
5. JADE and SPADE Multi-Agent Frameworks
6. Isaac Sim and NVIDIA Omniverse
7. AutoGen, CrewAI and LangChain
8. Jupyter Notebook, Google Colab, VS Code

INTRODUCTION TO TRANSFORMER BASED NLP

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER – II

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	45
L–T–P–J	:	3–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To build a strong foundation in natural language processing using transformer architectures - empowering students to understand, implement, and deploy BERT-based and GPT-based language models for real-world NLP applications including text classification, named entity recognition, question answering, text generation, and retrieval-augmented generation - using NVIDIA's transformer training and deployment ecosystem.

Course Objectives: This course will enable students to:

1. **Understand** the evolution of NLP from classical methods to transformer architectures including attention mechanisms and positional encoding.
2. **Implement** transformer models and fine-tune pre-trained BERT and GPT variants for downstream NLP tasks using HuggingFace and NVIDIA NeMo.
3. **Apply** encoder-based models (BERT, RoBERTa, DistilBERT) for classification, NER, and semantic similarity tasks.
4. **Develop** decoder-based language models (GPT, LLaMA) for text generation, summarization, and instruction following.
5. **Design** production-ready NLP applications using RAG pipelines, vector databases, and NVIDIA NIM microservices deployment.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Lectures on transformer theory supported with interactive attention visualization tools (BertViz, Transformer Explainer).
2. **Coding Labs:** Hands-on implementation of self-attention, multi-head attention, and full transformer encoder/decoder in PyTorch from scratch.
3. **HuggingFace Workshops:** Fine-tuning BERT and GPT variants on standard NLP benchmarks (GLUE, SQuAD, CoNLL-2003) using HuggingFace Trainer.
4. **NVIDIA NeMo Sessions:** Students train and deploy NLP models using NVIDIA NeMo framework and evaluate with NVIDIA NIM microservices.
5. **Paper Reading Club:** Weekly reading and discussion of landmark NLP papers: Attention is All You Need, BERT, GPT-3, LLaMA, RAG.
6. **Capstone Project:** Teams build and deploy a domain-specific NLP application (legal, medical, or multilingual) using transformers and RAG.

MODULE – I	**08 Hours**

Foundations of NLP and the Road to Transformers

Topics: NLP task taxonomy: tokenization, POS tagging, NER, parsing, classification, generation, Text preprocessing: cleaning, stopword removal, stemming, lemmatization, Classical representations: Bag of Words, TF-IDF, n-grams, Word embeddings: Word2Vec (CBOW, Skip-gram), GloVe, FastText - training and properties, Recurrent Neural Networks for NLP: LSTM, GRU, bidirectional RNNs, Sequence-to-sequence models and encoder-decoder architecture, Limitations of RNNs: vanishing gradient, inability to parallelize, Attention mechanism introduction: additive attention (Bahdanau), dot-product attention, The transformer motivation - parallelization and long-range dependencies.

Applications: Sentiment analysis using TF-IDF + Logistic Regression, Word2Vec embedding for semantic similarity in product recommendation, LSTM-based text classifier for spam detection.

AI Integration: NVIDIA GPU-accelerated Word2Vec training with RAPIDS cuML, NVIDIA NeMo tokenizer and vocabulary tools, HuggingFace Tokenizers library for fast BPE and WordPiece tokenization.

Authentic Intelligence: Bias in word embeddings - gender, racial, and cultural stereotypes encoded in Word2Vec, Responsible dataset curation for NLP model training, Language equity - NLP resource gaps for low-resource Indian languages, Data privacy in text datasets used for training (PII removal).

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building Transformer-Based NLP Applications [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Generative AI Explained [Self-paced, Free]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Introduction to Transformer-Based NLP [Self-paced + Workshop]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – II	**09 Hours**

Transformer Architecture: Self-Attention and Positional Encoding

Topics: Scaled dot-product attention: Query, Key, Value formulation, Multi-head attention: parallel attention heads, concatenation and projection, Positional encoding: sinusoidal encodings, learned positional embeddings, Transformer encoder block: LayerNorm, residual connections, Feed-Forward Network (FFN), Transformer decoder block: masked self-attention, cross-attention, Full encoder-decoder transformer for machine translation, Computational complexity of attention: $O(n^2d)$ and efficiency variants - Sparse Attention, Longformer, BigBird, Implementing a transformer from scratch in PyTorch, Visualizing attention heads with BertViz and Transformer Explainer.

Applications: Neural machine translation (NMT) with encoder-decoder transformer, Text summarization using encoder-decoder with copy mechanism, Question answering with cross-attention between passage and question.

AI Integration: NVIDIA Tensor Core acceleration for matrix multiplications in multi-head attention, Flash Attention for memory-efficient attention computation on NVIDIA GPUs, NVIDIA NeMo transformer building blocks for custom model construction.

Authentic Intelligence: Explainability of attention weights as a proxy for model reasoning, Limitations of attention visualization as a true explanation (Jain & Wallace, 2019), Carbon cost of large transformer training runs, Responsible reporting of compute in NLP research papers (model cards).
Key Online Resources / DLI Courses (Relevant to this Module): **1. Building Transformer-Based NLP Applications [Self-paced + Workshop ★]** https://learn.nvidia.com/en-us/training/learning-paths **2. Introduction to Transformer-Based NLP [Self-paced + Workshop]** https://learn.nvidia.com/en-us/training/learning-paths **3. Building LLM Applications With Prompt Engineering [Self-paced + Workshop ★]** https://learn.nvidia.com/en-us/training/learning-paths

MODULE – III	**09 Hours**
Encoder-Based Models: BERT and Variants **Topics:** BERT (Bidirectional Encoder Representations from Transformers): masked language modelling (MLM) and next sentence prediction (NSP), BERT input representation: [CLS], [SEP] tokens, segment embeddings, Fine-tuning BERT for downstream tasks: classification (GLUE), NER (CoNLL-2003), QA (SQuAD), BERT variants: RoBERTa (robust pretraining), DistilBERT (knowledge distillation), ALBERT (parameter sharing), DeBERTa, Domain-specific BERT: BioBERT (biomedical), LegalBERT, IndicBERT (Indian languages), Sentence transformers (SBERT) for semantic textual similarity, Cross-lingual models: mBERT, XLM-RoBERTa for multilingual NLP, Efficient fine-tuning: LoRA, Adapter Layers, Prompt Tuning (PEFT). **Applications:** BERT-based sentiment analysis for Indian language social media text, BioBERT fine-tuning for clinical NER on MIMIC-III discharge summaries, SBERT-based semantic search engine for legal document retrieval. **AI Integration:** NVIDIA NeMo BERT training with mixed-precision (FP16/BF16) on multi-GPU, NVIDIA TensorRT for BERT inference optimization (10× latency reduction), NVIDIA NIM microservices for instant BERT-based API deployment. **Authentic Intelligence:** Bias in BERT: gender and racial bias in masked language model predictions, De-biasing techniques: counterfactual data augmentation, bias probes, Ethical use of IndicBERT for regional language processing with community consent, Responsible fine-tuning - not amplifying existing corpus biases in specialized BERT variants.	
Key Online Resources / DLI Courses (Relevant to this Module): **1. Building Transformer-Based NLP Applications [Self-paced + Workshop ★]** https://learn.nvidia.com/en-us/training/learning-paths **2. Introduction to Transformer-Based NLP [Self-paced + Workshop]** https://learn.nvidia.com/en-us/training/learning-paths **3. Augment Your LLM Using Retrieval Augmented Generation [Self-paced]** https://learn.nvidia.com/en-us/training/learning-paths	

MODULE – IV	**09 Hours**

Decoder-Based Language Models: GPT and Large Language Models

Topics: GPT architecture: autoregressive language modelling, causal (masked) self-attention, GPT-1 → GPT-2 → GPT-3 → GPT-4 evolution: scale, emergent capabilities, Pretraining objectives: causal language modelling, next-token prediction, Decoding strategies: greedy decoding, beam search, top-k sampling, top-p (nucleus) sampling, temperature scaling, Open-source LLMs: LLaMA 2/3, Mistral, Gemma, Falcon - architecture differences, Instruction tuning and RLHF: SFT, reward model, PPO fine-tuning, Parameter-Efficient Fine-Tuning (PEFT): LoRA, QLoRA, DoRA, Quantization for LLMs: GPTQ, AWQ, GGUF for consumer-grade inference, Text summarization, translation, and code generation using GPT models.

Applications: LLaMA 3-based multilingual customer support chatbot with instruction tuning, GPT-based automatic code documentation generator, QLoRA fine-tuning of Mistral-7B for domain-specific medical Q&A.

AI Integration: NVIDIA NeMo for LLM pretraining and instruction tuning on multi-GPU clusters, TensorRT-LLM for optimized LLaMA and GPT inference on NVIDIA GPUs, NVIDIA NIM for one-click LLM API deployment with auto-scaling.

Authentic Intelligence: Hallucination in LLMs: causes, detection, and mitigation strategies, RLHF alignment: whose values are encoded - annotator selection bias, Misuse of LLMs for disinformation, deepfake text, and academic dishonesty, Responsible LLM deployment - content moderation, output filtering, and human oversight mechanisms.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building LLM Applications With Prompt Engineering [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Rapid Application Development with Large Language Models [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Building Transformer-Based NLP Applications [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – V	**10 Hours**

NLP Applications: RAG, Conversational AI, and Deployment

Topics: Retrieval-Augmented Generation (RAG): motivation, architecture - retriever + generator, Dense retrieval: DPR (Dense Passage Retrieval), bi-encoder vs. cross-encoder, Vector databases: FAISS, ChromaDB, Weaviate, NVIDIA cuVS - indexing and ANN search, GraphRAG: knowledge graph augmented generation with cuGraph and NVIDIA NIM, Conversational AI: dialogue state tracking, multi-turn context management, Named Entity Recognition (NER), relation extraction, and information extraction pipelines, NLP evaluation metrics: BLEU, ROUGE, BERTScore, perplexity, Production NLP deployment: NVIDIA NIM microservices, REST APIs, latency optimization, Multilingual and code-switching NLP for Indian language applications.

Applications: RAG-based technical documentation Q&A system using LLaMA + FAISS, Conversational AI agent for hospital appointment booking using NVIDIA NIM, Multilingual NER pipeline for Hindi-English code-switched social media data.

AI Integration: NVIDIA NIM for production NLP API deployment with GPU-accelerated inference, NVIDIA cuVS (cuGraph + FAISS) for GPU-accelerated vector similarity search, Structure From Chaos: GraphRAG with cuGraph and NVIDIA NIM for knowledge-rich Q&A.

Authentic Intelligence: Privacy in RAG systems - indexed document confidentiality and access control, Preventing prompt injection and jailbreaking in production conversational AI, Transparency of RAG citations - source attribution and hallucination prevention, Equitable NLP: investing in low-resource Indian language models for inclusive AI access.
Key Online Resources / DLI Courses (Relevant to this Module): **1. Augment Your LLM Using Retrieval Augmented Generation [Self-paced]** https://learn.nvidia.com/en-us/training/learning-paths **2. Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM [Self-paced]** https://learn.nvidia.com/en-us/training/learning-paths **3. Building Conversational AI Applications [Instructor-Led Workshop ★]** https://learn.nvidia.com/en-us/training/learning-paths

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Describe the evolution of NLP from classical representations to transformer architectures, explaining attention mechanisms and positional encoding.	**L2**
2	Implement multi-head self-attention, positional encoding, and complete transformer encoder-decoder blocks in PyTorch.	**L3**
3	Apply BERT-based encoder models (RoBERTa, DistilBERT, IndicBERT) for fine-tuning on classification, NER, and semantic similarity tasks.	**L3**
4	Develop decoder-based LLMs (GPT, LLaMA) using instruction tuning and PEFT (LoRA, QLoRA) for text generation and domain adaptation.	**L6**
5	Design production-ready NLP applications integrating RAG pipelines, vector databases, and NVIDIA NIM for scalable deployment.	**L6**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	3	1	2	2	–	3	–	–	2	3	3	3
CO-2	2	3	3	2	3	2	–	2	1	–	2	3	3	3
CO-3	3	3	3	2	3	3	–	3	2	–	2	3	3	3
CO-4	3	2	3	3	3	3	–	3	2	1	2	3	3	3

CO-5	3	3	3	3	3	3	2	3	2	2	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and implement transformer-based NLP systems for a variety of language understanding and generation applications.

PSO-2: Develop production-ready NLP software systems using HuggingFace, NVIDIA NeMo, NIM, and RAG frameworks.

PSO-3: Apply transformer-based NLP and LLMs to solve real-world problems in robotics, healthcare, agriculture, and multilingual AI systems.

TEXT BOOKS:

1. *Lewis Tunstall, Leandro von Werra & Thomas Wolf, Natural Language Processing with Transformers, O'Reilly Media, 2022.*

2. *Yoav Goldberg, Neural Network Methods for Natural Language Processing, Morgan & Claypool, 2017.*

3. *Jacob Devlin et al., BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, NAACL 2019.*

REFERENCE BOOKS:

1. NVIDIA NeMo Framework Documentation - LLM Training and NLP: https://docs.nvidia.com/nemo-framework/

2. Vaswani et al., Attention Is All You Need, NeurIPS 2017. https://arxiv.org/abs/1706.03762

3. HuggingFace Course - Free Transformers Course: https://huggingface.co/learn/nlp-course/

E-RESOURCES:

1. NVIDIA DLI - Gen AI & LLMs: Build Applications with LLMs: https://learn.nvidia.com/en-us/training/learning-paths

2. NVIDIA NeMo: https://developer.nvidia.com/nemo

3. NVIDIA NIM Microservices: https://www.nvidia.com/en-us/ai/

4. HuggingFace Hub: https://huggingface.co/models

5. IndicNLP and IndicBERT: https://indicnlp.ai4bharat.org/

RECOMMENDED TOOLS:

1. Model Development: HuggingFace Transformers, NVIDIA NeMo, PyTorch Lightning

2. Fine-Tuning: PEFT (LoRA, QLoRA), TRL (RLHF), Axolotl

3. Vector Databases: FAISS, ChromaDB, NVIDIA cuVS, Weaviate

4. Inference & Deployment: NVIDIA NIM, TensorRT-LLM, vLLM, Ollama

5. Evaluation: NLTK, spaCy, BERTScore, LM-Evaluation-Harness

6. Visualization: BertViz, Transformer Explainer, Weights & Biases

SEMESTER - III

AGENTIC AI ARCHITECTURES AND LLM SYSTEMS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER III

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	45
L–T–P–J	:	3–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To develop mastery in designing, building, and evaluating autonomous AI agent systems — where large language models act as planners, reasoners, and orchestrators capable of executing complex multi-step tasks in open-ended environments — equipping students to architect production-grade agentic pipelines integrating memory systems, tool use, and multi-agent collaboration frameworks such as LangChain, AutoGen, and CrewAI, while extending agentic intelligence to embodied robotic systems for real-world task execution.

Course Objectives: This course will enable students to:

1. Understand the architecture and capabilities of LLM-based autonomous agents — including reasoning loops, context management, and the role of foundation models as planners and decision-makers.

2. Design and implement memory, planning, and tool-use modules for AI agents — covering short-term context windows, long-term vector memory, ReAct-style reasoning, and API-based tool integration.

3. Build multi-agent systems and orchestrate agent collaboration for complex tasks using frameworks such as AutoGen, CrewAI, and LangGraph for role-based and hierarchical agent coordination.

4. Evaluate agent performance, reliability, and safety in deployment scenarios — assessing task completion rates, hallucination risks, tool misuse, and failure recovery across agentic pipelines.

5. Integrate agentic AI with robotic systems for embodied task execution — enabling LLM-driven task planning, environment perception, and action execution in physical and simulated robot platforms.

Teaching-Learning Process (General Instructions)

1. Lecture Method: Traditional lectures enriched with real-world agentic AI case studies — AutoGPT, GitHub Copilot Workspace, and LLM-driven robotics — to ground agent architecture concepts in current industry practice.

2. Interactive Teaching: Brainstorming on agent failure modes, group discussions on multi-agent role design, and role-playing scenarios where students act as orchestrators and tool providers to understand coordination dynamics.

3. Video and Animation-Based Learning: Curated videos demonstrating LLM reasoning traces, ReAct loops, memory retrieval pipelines, and multi-agent workflows to make abstract agentic architectures visually intuitive.

4. Collaborative Group Learning: Team-based projects where groups jointly design and evaluate multi-agent systems — assigning planner, executor, and critic agent roles — and present end-to-end agentic task demonstrations.

5. Critical Thinking and Higher-Order Assessment: Scenario-based challenges requiring students to diagnose hallucination failures, design tool-use safety guardrails, and architect reliable multi-agent pipelines for high-stakes environments.

6. Creative Problem Solving: Encouraging students to propose novel memory architectures, original tool-use workflows, and custom coordination protocols beyond standard LangChain and AutoGen implementations.

MODULE – I	**9 Hours**

Foundations of LLM Agents

Topics: Agent loop: Perceive-Think-Act; Chain-of-thought, Tree-of-thought, and ReAct reasoning; Prompt engineering for agents: zero-shot, few-shot, structured outputs; LLM tool use: function calling and JSON schemas; Structured output generation with Pydantic; Agent evaluation: AgentBench, WebArena, GAIA.

Applications: Autonomous web browsing agents performing multi-step information retrieval, form filling, and summarisation tasks without human intervention; Code generation agents autonomously writing, executing, debugging, and iterating on software using tool-augmented reasoning loops; Research agents decomposing complex queries into subgoals, retrieving papers, synthesising findings, and generating structured research reports.

AI Integration: Claude Computer Use for GUI-based autonomous agent interaction with desktop applications and browsers; GPT-4o and Gemini function calling for structured tool invocation, API chaining, and multi-step task execution within agent pipelines; OpenAI Assistants API for building stateful agents with persistent thread memory, code interpreter, and file retrieval tool integration; AI-assisted predictive analytics — agents autonomously querying data sources, running statistical models, and generating decision-support insights.

Authentic Intelligence: Students develop the ability to design autonomous reasoning workflows using LLM-based agents for real-world problem solving — building accountability-aware pipelines that balance agent autonomy with human oversight, transparency, and responsible deployment practices.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **AI for All: From Basics to Gen AI Practice**
   https://www.nvidia.com/en-us/learn/learning-path/generative-ai-llm/

MODULE – II	**9 Hours**

Memory and Knowledge Systems

Topics: Memory types: in-context, external (vector DB), episodic, semantic; Vector databases: FAISS, Chroma, Pinecone, Weaviate; Embedding models: sentence-transformers, text-embedding-3; RAG vs long-context LLMs; Knowledge graphs and graph RAG; Memory consolidation and forgetting strategies.

Applications: Personal AI assistants with persistent memory retaining user preferences, past interactions, and long-term context across sessions; Robot task history recall enabling agents to retrieve past execution experiences for improved replanning and error recovery; Long-term planning agents maintaining goal state, intermediate results, and evolving world knowledge across extended multi-session task execution.

AI Integration: MemGPT and Letta architecture for hierarchical memory management — enabling LLM agents to self-manage context windows through intelligent memory paging and retrieval; GraphRAG (Microsoft) for knowledge graph-augmented retrieval combining structured entity relationships with vector search for richer, more accurate agent reasoning; Mem0 memory layer for persistent, user-specific long-term memory injection into LLM agent pipelines across sessions.

Authentic Intelligence: Students design intelligent memory architectures enabling agents to learn continuously, retain context, and perform long-horizon reasoning — while critically evaluating data retention risks, user privacy obligations, and the ethical boundaries of persistent agent memory in personal and sensitive deployment contexts.

Key Online Resources / DLI Courses (Relevant to this Module):

1. Accelerate Data Science Workflows with Zero Code Changes
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+T-DS-03+V1

2. Introduction to Multi-Modal Data Curation
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:NVIDIA+S-FX-27+V1

MODULE – III	**9 Hours**

Planning and Reasoning

Topics: Task decomposition: hierarchical task networks (HTN); LLM-based planning: SayCan, Inner Monologue, DEPS; Symbolic planning with PDDL and LLM synthesis; Monte Carlo Tree Search with LLMs; Reflexion: verbal reinforcement learning; Self-consistency and majority voting.

Applications: Robot long-horizon task planning using LLM-based task decomposition and symbolic grounding for multi-step manipulation and navigation in unstructured environments; Scientific experiment design agents autonomously formulating hypotheses, planning experimental steps, and iterating based on result feedback; Software engineering agents performing end-to-end code planning, implementation, testing, and debugging across complex multi-file repositories.

AI Integration: OpenAI o1/o3 reasoning chains for extended chain-of-thought planning and self-verified multi-step problem solving in complex task environments; DeepSeek-R1 for open-source long-horizon reasoning with reinforcement learning-trained deliberative planning capabilities; AlphaCode 2 and SWE-agent for autonomous software engineering — planning, writing, and iteratively debugging code solutions to real-world programming tasks.

Authentic Intelligence: Students develop adaptive reasoning systems capable of multi-step planning, self-correction, and intelligent decision-making — while critically examining the risks of autonomous planning in high-stakes domains, accountability gaps when self-correcting agents make compounding errors, and the ethical boundaries of delegating complex real-world decisions to LLM-based planners.

Key Online Resources / DLI Courses (Relevant to this Module):

1. Building a Brain in 10 Minutes
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+T-FX-01+V1

MODULE – IV	**9 Hours**

Multi-Agent Systems

Topics: Multi-agent architectures: sequential, parallel, hierarchical; Agent roles and persona design; Inter-agent communication protocols; Debate and critique between agents; AutoGen, CrewAI, LangGraph, and Swarm frameworks; Agent monitoring, logging, and observability; Conflict resolution and consensus in multi-agent systems.

Applications: Software development pipelines using specialised agents for requirements analysis, code generation, code review, testing, and documentation in an automated end-to-end engineering workflow; Autonomous research teams comprising literature search, hypothesis generation, experiment design, and report writing agents collaborating to accelerate scientific discovery; Multi-robot task allocation using multi-agent LLM orchestration for dynamic role assignment, workload balancing, and collaborative execution in warehouse and field robotics scenarios.

AI Integration: Microsoft AutoGen 0.4 for flexible multi-agent conversation orchestration with dynamic agent creation, tool use, and human-in-the-loop intervention patterns; CrewAI Flows for role-based multi-agent pipeline design with structured crew coordination, task chaining, and agent specialisation; OpenAI Swarm and Anthropic multi-agent patterns for lightweight agent handoff, context passing, and scalable distributed agent execution architectures.

Authentic Intelligence: Students design collaborative AI ecosystems where multiple agents coordinate, negotiate, and solve complex distributed problems — while critically evaluating emergent failure modes in multi-agent pipelines, accountability diffusion when no single agent owns a decision, and the ethical risks of fully autonomous agent collectives operating without human oversight in consequential real-world tasks.

Key Online Resources / DLI Courses (Relevant to this Module):

1. Generative AI Explained
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1

MODULE – V	**9 Hours**

Safety, Evaluation, and Embodied Agents

Topics: Agent safety: prompt injection, goal misgeneralisation, reward hacking; Sandboxing and tool call validation; Capability control and minimal footprint principles; Embodied AI agents: vision-language-action models; LLMs for robot task planning: SayCan, Code-as-Policies; Agent observability: LangSmith, Langfuse; Deployment patterns for production agents.

Applications: Safe autonomous coding agents with sandboxed execution environments, tool call validation, and human-in-the-loop checkpoints for production software engineering pipelines; Robotic household assistants using vision-language-action models for open-vocabulary instruction following, object manipulation, and safe human-robot interaction in unstructured home environments; Industrial process automation agents executing multi-step manufacturing workflows with real-time observability, anomaly detection, and fail-safe intervention mechanisms.

AI Integration: Constitutional AI for agents — embedding principle-based self-critique and behavioural constraints into agent reasoning loops to enforce safe and aligned autonomous decision-making; OpenVLA and pi0 vision-language-action models for grounding natural language instructions into real-world robot motor actions in embodied agent deployments; Sparse reward shaping in embodied agents for training safe, goal-directed behaviours in simulation before physical robot deployment.

Authentic Intelligence: Students ensure responsible, safe, and trustworthy deployment of agentic AI systems in real-world robotic and industrial environments — critically examining prompt injection vulnerabilities, the risks of goal misgeneralisation in open-ended deployments, minimal footprint principles for limiting agent permissions, and the ethical imperative of maintaining meaningful human oversight in consequential autonomous systems.

1. Getting Started with AI on NVIDIA Jetson Nano

https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-02+V2

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Explain LLM agent architectures including reasoning, tool use, and structured output generation.	**L2**
2	Implement memory and knowledge systems for AI agents using vector databases and RAG.	**L3**
3	Design planning and reasoning pipelines for complex multi-step agentic tasks.	**L4**
4	Build multi-agent systems using orchestration frameworks with safety and observability.	**L4**
5	Integrate agentic AI with robotic systems for embodied and autonomous task execution.	**L3**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	2	2	3	2	3	3	3	3	2	2	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. *J. Alammar S M. Grootendorst – Hands-On Large Language Models, O'Reilly, 2024.*

2. *LangChain Documentation – python.langchain.com.*

REFERENCE BOOKS:

1. L. Wang et al. – A Survey on Large Language Model based Autonomous Agents, arXiv:2308.11432, 2023.

2. AutoGen Documentation – microsoft.github.io/autogen.

E-RESOURCES:

1. OpenVLA Repository **—** **https://github.com/openvla/openvla**

2. NVIDIA Deep Learning Institute **—** **https://learn.nvidia.com/courses/**

3. IEEE Xplore **—** **https://ieeexplore.ieee.org/**

4. LangChain Documentation **—** **https://python.langchain.com/**

5. Microsoft AutoGen Documentation **—** **https://microsoft.github.io/autogen/**

6. CrewAI Documentation **—** **https://docs.crewai.com/**

7. LlamaIndex Documentation **—** **https://docs.llamaindex.ai/**

8. Hugging Face Documentation **—** **https://huggingface.co/docs**

RECOMMENDED TOOLS:

1. Programming Languages — Python 3.x, C/C++

2. ML / AI Frameworks — TensorFlow, PyTorch, scikit-learn, Keras

3. Agentic AI S Orchestration Frameworks — LangChain, LangGraph, CrewAI, Microsoft AutoGen, OpenAI Swarm

4. Memory S Vector Database Tools — FAISS, ChromaDB, Pinecone, Weaviate

5. Multimodal S VLM Tools — Hugging Face Transformers (CLIP, LLaVA), OpenAI API, NVIDIA NIM, Google AI Studio (Gemini)

6. Observability S Evaluation Tools — LangSmith, Langfuse, Weights S Biases

7. Robotics S Simulation — NVIDIA Isaac Lab, Isaac Sim, ROS 2, ManiSkill2

Autonomous Vehicles and SLAM

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER III

Course Code	:		**Credits**	:	04
Hours / Week	:	04	**Total Hours**	:	60
L–T–P–J	:	4–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To develop comprehensive expertise in autonomous vehicle systems — spanning sensor fusion, perception, simultaneous localisation and mapping (SLAM), motion planning, and closed-loop control — enabling students to architect, implement, and evaluate full autonomous driving stacks using industry-standard tools such as ROS 2, NVIDIA Isaac Sim, CARLA, and OpenCV, while bridging theoretical foundations with real-world deployment challenges in structured and unstructured driving environments.

Course Objectives: This course will enable students to:

1. **Understand** the end-to-end architecture of autonomous vehicle systems, including sensor modalities such as LiDAR, cameras, radar, and IMU, and their roles in the perception-planning-control pipeline.
2. **Implement** SLAM algorithms — including EKF-SLAM, particle filter SLAM, graph-based SLAM, and ORB-SLAM — for accurate simultaneous localisation and mapping in dynamic environments.
3. **Design and evaluate** perception pipelines for 3D object detection, semantic segmentation, depth estimation, and lane detection using deep learning frameworks and point cloud processing.
4. **Apply** motion planning algorithms including A*, RRT, RRT*, Hybrid A*, and model predictive control (MPC) for safe navigation in both structured road environments and unstructured off-road scenarios.
5. **Evaluate** autonomous vehicle systems using industry-standard simulation platforms such as CARLA, Gazebo, and NVIDIA DRIVE Sim, assessing performance across safety, robustness, and real-time execution metrics.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Traditional lectures enriched with real-world autonomous driving case studies from Waymo, Tesla Autopilot, and ISRO rover missions to contextualise perception, SLAM, and planning concepts within deployed systems.
2. **Interactive Teaching:** Incorporating brainstorming sessions, group discussions, focused listening, question formulation, notetaking, annotation, and role-playing activities such as simulating sensor failure scenarios and debugging autonomous pipelines collaboratively.
3. **Video and Animation-Based Learning:** Showing curated video content and animations to explain the functioning of SLAM algorithms, sensor fusion pipelines, object detection architectures, and motion planning visualisations in simulated and real environments.
4. **Collaborative Group Learning:** Encouraging team-based lab projects where student groups jointly build and test perception modules, integrate SLAM outputs with planning stacks, and present autonomous mission demonstrations in simulation.
5. **Critical Thinking and Higher-Order Assessment:** Posing complex scenario-based quiz questions and programming challenges that require students to debug SLAM drift, resolve sensor conflicts, and design planning solutions for edge-case driving scenarios.

6. Creative Problem Solving: Encouraging students to explore novel sensor fusion strategies, propose original SLAM optimisations, and develop custom planning heuristics — fostering innovation beyond standard algorithmic implementations.

MODULE – I	**9 Hours**

Autonomous Vehicle Architecture and Sensors

Topics: AV system architecture: perception, prediction, planning, control stack; SAE automation levels; Sensor modalities: LiDAR, RADAR, camera, GPS/IMU, ultrasonic; Sensor calibration and synchronisation; HD maps: OpenDRIVE, Lanelet2; Data recording and annotation.

Applications: Self-driving passenger cars in urban environments using LiDAR, camera, and RADAR fusion with HD map localisation; Autonomous sidewalk delivery robots using ultrasonic and camera stacks for low-speed pedestrian navigation; Mining and agricultural autonomous vehicles using LiDAR terrain mapping and RADAR obstacle detection in GPS-degraded environments.

AI Integration: BEVFusion and PointPillars for multi-modal LiDAR-camera bird's-eye-view sensor fusion and 3D object detection; Foundation models for autonomous driving — GAIA-1 for world model scene generation and DriveLM for vision-language end-to-end driving reasoning; NVIDIA DRIVE Sim and Isaac Sim for physics-accurate synthetic sensor data generation for perception model training and validation.

Authentic Intelligence: Sensor failure accountability — ethical responsibility when deploying systems with known degradation in adverse weather; Data privacy obligations for continuous public-space LiDAR and camera data collection; Dataset bias — KITTI and nuScenes skewed toward Western urban environments, reducing performance in Global South contexts; Legal accountability frameworks when autonomous perception systems cause fatal errors.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building Robotics Applications Using ROS and NVIDIA Isaac**
   https://developer.nvidia.com/blog/building-collaborative-robotics-using-ros-and-isaac-sdk/

MODULE – II	**9 Hours**

SLAM and Localisation

Topics: EKF-SLAM and UKF-SLAM; Graph-based SLAM: g2o, GTSAM; LiDAR SLAM: LOAM, LeGO-LOAM, LIO-SAM; Visual-Inertial Odometry (VIO): VINS-Mono, MSCKF; Multi-session mapping; GPS-denied localisation; NDT (Normal Distribution Transform) scan matching.

Applications: Autonomous indoor robots navigating warehouse and hospital environments without GPS using LiDAR and visual SLAM; GPS-denied mining AV operations using LiDAR-based SLAM for underground tunnel mapping and localisation; Delivery drone navigation in urban canyons and indoor spaces using visual-inertial odometry and real-time mapping.

AI Integration: Deep learning SLAM: NICE-SLAM for neural implicit scene representation and NeRF-SLAM for photorealistic dense 3D mapping using neural radiance fields; Learning-based loop closure detection using CNN and transformer-based place recognition for robust long-term localisation; Deep odometry networks for IMU and camera-based ego-motion estimation in feature-sparse environments.

Authentic Intelligence: Reliability of GPS-denied localisation in safety-critical environments — accountability when SLAM drift causes navigation failure; Privacy concerns in dense indoor mapping — building layout capture and unauthorised spatial data collection; Robustness bias — SLAM systems trained in structured Western environments underperforming in cluttered or unstructured Global South deployment contexts; Ethical deployment of mapping drones in restricted or sensitive areas.

Key Online Resources / DLI Courses (Relevant to this Module):

1. Building Realtime Video AI Applications
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-IV-01+V1

MODULE – III	**9 Hours**

Perception for Autonomous Driving

Topics: detection: PointPillars, CenterPoint, BEVFusion; Lane detection: PolyLane, CLRNet; Traffic sign and light recognition; Occupancy grid prediction; Trajectory prediction: Social Force, Trajectron++; Bird's Eye View (BEV) perception.

Applications: Highway autopilot systems using camera-based lane detection, vehicle tracking, and adaptive cruise control in structured high-speed environments; Urban intersection navigation using multi-sensor fusion for simultaneous pedestrian, cyclist, and vehicle detection with intent prediction; Pedestrian safety systems in smart city infrastructure using real-time human pose estimation and trajectory forecasting to prevent collisions.

AI Integration: Tesla FSD bird's-eye-view (BEV) approach — multi-camera surround view fusion into a unified BEV feature space for lane, object, and free-space detection without LiDAR; BEVFusion for unified LiDAR-camera sensor fusion in BEV representation enabling robust 3D detection across adverse conditions; Unified perception transformers — DETR3D and UniAD for end-to-end multi-task perception covering detection, tracking, mapping, and motion prediction in a single architecture.

Authentic Intelligence: SPerception failure accountability in safety-critical scenarios — responsibility when object detection misses a pedestrian or misclassifies a cyclist; Demographic bias in pedestrian detection — models underperforming on darker skin tones or non-Western clothing due to skewed training datasets; Transparency obligations — right of affected road users to understand how autonomous perception systems make decisions affecting their safety; Ethical limits of predictive systems — risks of over-reliance on trajectory forecasting in genuinely unpredictable human behaviour.

Key Online Resources / DLI Courses (Relevant to this Module):

1. Getting Started with AI on Jetson Nano
https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-02+V2

MODULE – IV	**9 Hours**

Motion Planning for AVs

Topics: Frenet frame path planning; Lattice planner and Polynomial trajectory; Spline-based path generation; Behaviour prediction and decision making; Finite state machines for driving behaviour; Contingency planning; Real-time re-planning.

Applications: Highway lane change manoeuvres using velocity-aware path planning and gap acceptance algorithms for safe merging in high-speed traffic; Roundabout navigation using rule-based and learning-based planners for yield detection, entry timing, and circular path following in unstructured intersections; Emergency braking systems using real-time obstacle detection and minimum stopping distance computation for collision avoidance in safety-critical scenarios.

AI Integration: Imitation learning for driving , behaviour cloning and inverse reinforcement learning from human driver demonstrations for naturalistic motion planning in complex urban scenarios; RL-based planners — deep reinforcement learning agents trained in CARLA and NVIDIA DRIVE Sim for adaptive decision-making in dynamic and adversarial traffic environments; Diffusion models for trajectory prediction — denoising diffusion probabilistic models generating multimodal future trajectory distributions for uncertainty-aware motion planning.

Authentic Intelligence: Accountability for planning failures, ethical responsibility when motion planners make incorrect right-of-way decisions causing accidents; Trolley problem in AV planning — moral implications of encoding collision trade-off decisions into autonomous systems; Algorithmic transparency — obligation to make planning decisions interpretable and auditable for regulators and affected road users; Bias in imitation learning, planners trained on data from specific demographics or geographies replicating unsafe or culturally inappropriate driving behaviours.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Applications of AI for Anomaly Detection**
   https://learn.nvidia.com/courses/course?course_id=course-v1:DLI+S-AC-06+V1

MODULE – V	**9Hours**

Simulation and Safety Evaluation

Topics: CARLA simulator; SUMO traffic simulation; Scenario-based testing: OpenSCENARIO; Safety metrics: Time-to-Collision, ETTC; Adversarial scenario generation; Regulatory frameworks: ISO 26262, SOTIF; AV safety case methodology.

Applications: Virtual validation of full AV stacks in CARLA and NVIDIA DRIVE Sim across thousands of simulated scenarios before physical road testing; Regulatory certification pipelines using simulation-based evidence generation to meet ISO 26262 functional safety and UNECE WP.29 autonomous vehicle approval standards; Edge case stress testing — simulating rare but critical scenarios such as sensor occlusion, adverse weather, sudden pedestrian crossings, and sensor spoofing attacks.

AI Integration: Adversarial scenario generation with GANs — generative adversarial networks creating rare, safety-critical traffic scenarios and sensor degradation conditions to stress-test perception and planning stacks beyond manually curated test suites; Simulation with neural rendering — NeRF and Gaussian splatting-based photorealistic scene synthesis for sensor-accurate camera and LiDAR simulation enabling high-fidelity closed-loop AV evaluation.

Authentic Intelligence: Simulation-to-real gap accountability, ethical risk of over-certifying AV systems based on simulation performance that does not transfer to real-world edge cases; Transparency in safety reporting — obligation to disclose simulation assumptions, coverage gaps, and failure rates to regulators and the public; Liability in certification — determining responsibility when a simulation-validated system fails in real-world deployment; Equitable safety standards — ensuring certification frameworks account for diverse global road conditions beyond Western highway-centric test scenarios.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **NVIDIA Isaac for Accelerated Robotics**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-OV-05+V1

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Explain AV system architecture, sensor modalities, and calibration processes.	**L2**
2	Implement LiDAR SLAM and visual-inertial odometry for robust localisation.	**L3**
3	Design perception pipelines for 3D detection, lane detection, and trajectory prediction.	**L4**
4	Apply motion planning algorithms for autonomous vehicles in structured environments.	**L3**
5	Evaluate AV systems in simulation using safety metrics and scenario-based testing.	**L5**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. *R. Siegwart, I.R. Nourbakhsh S D. Scaramuzza – Introduction to Autonomous Mobile Robots, 2nd Ed., MIT Press, 2011*

2. *T. Thrun, W. Burgard S D. Fox – Probabilistic Robotics, MIT Press, 2005.*

REFERENCE BOOKS:

1. CARLA Autonomous Driving Simulator Documentation – carla.org

2. Apollo Autonomous Driving Platform – github.com/ApolloAuto/apollo

3. Autoware Foundation – autoware.org

E-RESOURCES:

1. OpenAI Platform Documentation — https://platform.openai.com/docs

2. Anthropic Claude Documentation — https://docs.anthropic.com/

3. Google AI Studio (Gemini) — https://aistudio.google.com/

4. LangChain Documentation — https://python.langchain.com/

5. LangGraph Documentation — https://langchain-ai.github.io/langgraph/

6. Microsoft AutoGen Documentation — https://microsoft.github.io/autogen/

7. CrewAI Documentation — https://docs.crewai.com/

8. OpenAI Swarm Repository — https://github.com/openai/swarm

RECOMMENDED TOOLS:

1. Programming Languages — Python 3.x, SQL, JSON

2. Agentic AI S Orchestration Frameworks — LangChain, LangGraph, CrewAI, Microsoft AutoGen, OpenAI Swarm

3. LLM S API Platforms — OpenAI API, Anthropic Claude API, Google Gemini API, Hugging Face Transformers, Pydantic AI

4. Memory S Knowledge Systems — FAISS, ChromaDB, Pinecone, Weaviate, Mem0, Microsoft GraphRAG

Planning S Reasoning Tools — DeepSeek-R1, SWE-agent, AlphaCode 2

5. Observability S Evaluation Tools — LangSmith, Langfuse, Weights S Biases

6. Robotics S Embodied AI — ROS 2, OpenVLA Repository, NVIDIA Isaac Sim, NVIDIA Isaac Lab, NVIDIA Jetson Platform

SPACE AND SWARM ROBOTICS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER III

Course Code	**:**		**Credits**	**:**	04
Hours / Week	**:**	04	**Total Hours**	**:**	60
L–T–P–J	**:**	4–0–0–0	**CIE + SEE**	**:**	60 + 40 Marks

Course Vision: To develop engineers who can design, simulate, and deploy intelligent robotic systems for space exploration and swarm robotics applications — integrating orbital mechanics, space robot kinematics, multi-agent coordination, swarm intelligence algorithms, and NVIDIA Isaac Sim-based simulation — to build robust, autonomous, and collectively intelligent robot systems for the frontier domains of space exploration and large-scale distributed robotics.

Course Objectives: This course will enable students to:

1. **Understand** the unique operational constraints, kinematics, and control challenges of space robotic systems in microgravity environments.
2. **Apply** multi-robot coordination algorithms, consensus protocols, and distributed sensing for swarm robotic systems.
3. **Implement** swarm intelligence algorithms including PSO, ACO, artificial potential fields, and stigmergy-based collective behaviors.
4. **Design** space robot manipulation systems and swarm robotic platforms using NVIDIA Isaac Sim and ROS 2.
5. **Evaluate** multi-robot system performance using coverage, connectivity, scalability, and fault-tolerance metrics.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Lectures on space robotics and swarm theory supported with NASA/ESA mission case studies and multi-robot simulation demonstrations.
2. **Simulation Labs:** NVIDIA Isaac Sim and ARGoS sessions for simulating swarm behaviors and space robot manipulation in zero-gravity environments.
3. **Algorithm Implementation:** Coding PSO, ACO, artificial potential fields, and consensus protocols in Python for multi-robot coordination.
4. **Mission Design Project:** Teams design a swarm robotic mission — lunar surface mapping or disaster response — from task allocation to execution.
5. **Guest Lectures:** ISRO scientists, space robotics researchers, and swarm robotics experts share insights on real-world deployments.
6. **Paper Reading:** Discussion of landmark papers in swarm robotics (Dorigo, Brambilla) and space robotics (NASA Mars rovers, Robonaut 2).

MODULE – I	**9 Hours**

Introduction to Space Robotics and Orbital Mechanics

Topics: Overview of space robotics: planetary rovers, orbital servicing robots, in-space assembly; Historical missions: Sojourner, Spirit, Opportunity, Curiosity, Perseverance (NASA); Chandrayaan and Gaganyaan robotics context (ISRO); Orbital mechanics fundamentals: Keplerian orbits, orbital parameters, launch windows; Microgravity effects on robot dynamics: absence of gravity, thermal extremes, radiation; Free-floating robot dynamics: conservation of momentum, zero-G kinematics; Space environment challenges: vacuum, thermal cycling, cosmic radiation, electrostatic discharge; Space communication constraints: signal delay (light-time), autonomous operation necessity; Introduction to NVIDIA Isaac Sim for space robot simulation.

Applications: Autonomous terrain navigation for Mars rover in rocky terrain; On-orbit satellite servicing robot for fuel transfer; Chandrayaan-inspired lunar regolith sampling robot design.

AI Integration: NVIDIA Isaac Sim for space robot simulation with physics-accurate microgravity; Deep RL for autonomous rover navigation in planetary terrain using Isaac Lab; Foundation models for autonomous decision-making in communication-delayed space missions.

Authentic Intelligence: Space debris and responsible orbital operations; COSPAR planetary protection policies for preventing biological contamination of celestial bodies; Equitable access to space — space resource mining ethics; Safety of autonomous decisions in communication-delayed deep space missions.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Getting Started: Simulating Your First Robot in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Disaster Risk Monitoring Using Satellite Imagery [Self-paced, Certificate ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to AI-Based Robot Development With Isaac ROS [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – II	**9 Hours**

Space Robot Kinematics, Manipulation, and Control

Topics: Space manipulator kinematics: forward and inverse kinematics of free-floating robot arms; Reaction dynamics: base disturbance from arm motion, Generalized Jacobian Matrix (GJM); Reaction Null Space (RNS) control for disturbance-free manipulation; Space robot control strategies: computed torque control, impedance control; Grasping and docking in space: tumbling target capture, compliant grippers; Vision-based pose estimation for non-cooperative satellite capture (ORB-SLAM3, SPEED dataset); Space robot examples: Canadarm2 (ISS), ERA (European Robotic Arm), JAXA ETS-VII; Soft robotics for space: shape memory alloys, pneumatic actuators in vacuum; Simulation of space manipulators in NVIDIA Isaac Sim.

Applications: Canadarm2-style docking simulation for ISS resupply vehicle berthing; Vision-based tumbling satellite capture using pose estimation with SPEED dataset; Free-floating dual-arm manipulation for on-orbit assembly of large structures.

AI Integration: Deep RL for space robot arm trajectory planning in Isaac Lab; Computer vision-based non-cooperative target pose estimation using CNN on NVIDIA Jetson; NVIDIA Isaac Sim physics engine for accurate microgravity arm dynamics simulation.

Authentic Intelligence: Safety-critical failure modes in space robot manipulation — no repair once launched; Human-robot interaction in space suits — crew safety in proximity operations; Responsible use of dual-use space technologies (military vs. civilian); Debris generation risk from failed space robot operations.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Ingesting Robot Assets and Simulating Your Robot in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Synthetic Data Generation for Perception Model Training in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to AI-Based Robot Development With Isaac ROS [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – III	**9 Hours**

Foundations of Swarm Robotics and Collective Intelligence

Topics: Swarm robotics definition: scalability, flexibility, robustness, and decentralization; Biological inspiration: ant colonies (ACO), bee swarms (ABC), bird flocking (Reynolds Boids), fish schooling; Self-organization and emergence in swarm systems; Stigmergy: indirect communication through environment modification; Robot-to-robot communication: local broadcast, infrared, WiFi mesh, acoustic; Collective behaviors taxonomy: aggregation, dispersion, foraging, flocking, formation, task allocation; Swarm hardware platforms: Kilobot, Crazyflie 2.x, e-puck, Thymio, Lego Mindstorms; ARGoS simulator for large-scale swarm simulation (1000+ robots); Modular self-reconfigurable robots: SMORES, M-TRAN, SuperBot.

Applications: 1000-robot Kilobot swarm aggregation and shape formation experiment; Crazyflie drone swarm collective exploration of GPS-denied warehouse; ACO-inspired swarm foraging for disaster debris search and mapping.

AI Integration: Deep reinforcement learning for emergent swarm behavior (MADDPG, QMIX); NVIDIA Isaac Sim multi-agent environment for swarm RL training; Graph Neural Networks (GNN) for decentralized swarm coordination and communication.

Authentic Intelligence: Dual-use concerns of autonomous swarm systems — military drone swarms; Accountability in emergent swarm behavior — who is responsible when 1000 robots cause harm?; Privacy implications of distributed swarm sensing in public spaces; International regulation of autonomous weapon swarms (LAWS — Lethal Autonomous Weapon Systems).

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Getting Started: Simulating Your First Robot in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Developing Robots with Software-in-the-Loop (SIL) in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **An Introduction to Robot Learning and Isaac Lab [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – IV	**9 Hours**

Swarm Intelligence Algorithms and Multi-Robot Coordination

Topics: Particle Swarm Optimization (PSO): velocity update, inertia weight, cognitive and social components; Ant Colony Optimization (ACO): pheromone deposit and evaporation, ant decision rule; Artificial Bee Colony (ABC): employed, onlooker, and scout bees; Artificial Potential Fields (APF): attractive and repulsive potentials,

local minima problem; Flocking algorithms: Reynolds' rules — separation, alignment, cohesion; Formation control: leader-follower, virtual structure, behaviour-based; Consensus protocols: average consensus, max consensus, distributed estimation; Distributed task allocation: market-based (CBBA), threshold-based, auction mechanisms; Multi-robot SLAM and cooperative mapping; Fault tolerance in swarm systems: redundancy, role reassignment.

Applications: ACO-based path planning for swarm of ground robots in maze environment; APF-based multi-drone formation flight with obstacle avoidance; CBBA-based dynamic task allocation for disaster response robot swarm.

AI Integration: MADDPG (Multi-Agent Deep Deterministic Policy Gradient) for swarm formation control; Graph Attention Networks (GAT) for decentralized multi-robot communication; NVIDIA Isaac Lab multi-agent RL environment for swarm coordination training.

Authentic Intelligence: Emergent behavior validation — ensuring swarm behaves as intended in edge cases; Fail-safe mechanisms when consensus protocols fail or are jammed; Transparency of distributed decision-making to human supervisors; Ethical task allocation — avoiding unfair load distribution among robot agents.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Train Your First Robot in Isaac Lab [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Leveraging ROS 2 and Hardware-in-the-Loop (HIL) in Isaac Sim [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Transferring Robot Learning Policies from Simulation to Reality [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – V	**9 Hours**

Advanced Topics: Space Swarms, Future Missions, and System Integration

Topics: Intersection of space and swarm robotics: planetary swarm exploration, in-orbit swarm assembly; NASA SWAMP, ESA Daedalus, MIT AeroCube swarm concepts; In-space manufacturing and assembly using swarm robots; CubeSat swarms for distributed Earth observation and space weather monitoring; Applying AI weather and satellite monitoring to space robotics support (NVIDIA Earth-2); Multi-robot system integration: ROS 2 multi-robot infrastructure, DDS, RTPS; Testing and validation of swarm systems: simulation, hardware-in-the-loop, field trials; ISRO future robotic missions: Chandrayaan-4, Shukrayaan, Gaganyaan robotic precursors; Regulatory and policy framework for autonomous space robots: ITU, COSPAR, UN COPUOS.

Applications: CubeSat swarm constellation design for persistent lunar surface monitoring; Swarm of micro-rovers for distributed ice detection in permanently shadowed lunar craters; NVIDIA Earth-2-assisted orbital weather monitoring for space launch window selection.

AI Integration: NVIDIA Earth-2 for space weather prediction supporting satellite swarm operations; NVIDIA Isaac Sim for end-to-end space swarm mission simulation and validation; Agentic AI for autonomous long-duration swarm mission planning and replanning.

Authentic Intelligence: Space debris mitigation from failed swarm satellites — active debris removal ethics; Long-term sustainability of cislunar and orbital robotic swarm operations; Indigenous knowledge and equity in space resource access (Moon Treaty considerations); Environmental impact of space launch emissions from swarm constellation deployment missions.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Disaster Risk Monitoring Using Satellite Imagery [Self-paced, Certificate ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Applying AI Weather Models with NVIDIA Earth-2 [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Transferring Robot Learning Policies from Simulation to Reality [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Describe the operational constraints, kinematics, and control challenges unique to space robotic systems operating in microgravity.	**L2**
2	Apply space manipulator kinematics, reaction dynamics, and vision-based pose estimation for on-orbit robot manipulation and docking.	**L3**
3	Implement swarm intelligence algorithms (PSO, ACO, APF, flocking) and consensus protocols for multi-robot coordination.	**L3**
4	Design multi-robot task allocation, formation control, and cooperative SLAM systems with fault-tolerant swarm architectures.	**L6**
5	Evaluate space swarm mission concepts, system integration frameworks, and policy implications for autonomous space and swarm robotic systems.	**L5**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**1**	**2**	**3**
CO-1	3	2	3	1	2	2	1	3	–	–	2	3	3	3
CO-2	3	3	3	2	3	2	–	3	1	–	2	3	3	3
CO-3	3	3	3	2	3	3	–	2	2	–	2	3	3	3
CO-4	3	3	3	3	3	3	–	3	2	1	2	3	3	3
CO-5	2	2	3	2	3	3	3	3	2	2	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. *Marco Dorigo & Mauro Birattari, Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press, 2004.*
2. *Kazuya Yoshida & Brian Wilcox (Eds.), Space Robotics: Dynamics and Control, Kluwer Academic Publishers, 1993.*
3. *Maja Mataric, The Robotics Primer, MIT Press, 2007.*

REFERENCE BOOKS:

1. NVIDIA Isaac Sim Documentation — Multi-Robot and Space Simulation: https://docs.isaacsim.omniverse.nvidia.com/
2. NASA JPL Open Source Rover Documentation: https://opensourcerover.jpl.nasa.gov/
3. ARGoS Swarm Simulator Documentation: https://www.argos-sim.info/documentation.php

E-RESOURCES:

1. NVIDIA DLI — Robotics: Getting Started with Isaac Sim and Isaac Lab: https://learn.nvidia.com/en-us/training/learning-paths
2. NVIDIA Earth-2 for Space Weather: https://www.nvidia.com/en-us/high-performance-computing/earth-2/
3. NASA Mars Rover Simulation: https://mars.nasa.gov/technology/robotics/
4. ISRO Future Missions: https://www.isro.gov.in/
5. ARGoS Multi-Robot Simulator: https://www.argos-sim.info/

RECOMMENDED TOOLS:

1. Space Robot Simulation: NVIDIA Isaac Sim, NASA JPL OpenSPACE, MATLAB/Simulink Aerospace Toolbox
2. Swarm Simulation: ARGoS, Webots, NVIDIA Isaac Sim (Multi-agent), Gazebo
3. Swarm Hardware Platforms: Kilobot, Crazyflie 2.x, e-puck 2, Thymio II
4. Robot Programming: ROS 2, NVIDIA Isaac ROS, DDS/RTPS (FastDDS)
5. Multi-Agent RL: NVIDIA Isaac Lab, Ray RLlib (MADDPG, QMIX), PettingZoo
6. Satellite & Orbital Tools: GMAT (NASA), Orekit, NVIDIA Earth-2

CYBER SECURITY IN ROBOTICS

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER III

Course Code	:		**Credits**	:	04
Hours / Week	:	04	**Total Hours**	:	60
L–T–P–J	:	3–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To develop specialised knowledge in securing robotic and cyber-physical systems against evolving cyber threats — ensuring safety, integrity, and resilience of connected robotic platforms in industrial and critical infrastructure environments.

Course Objectives: This course will enable students to:

1. **Analyse** the cybersecurity threat landscape specific to robotic and cyber-physical system (CPS) environments.
2. **Apply** cryptographic techniques and secure communication protocols to robotic systems.
3. **Identify** and exploit vulnerabilities in ROS 2 and robotic middleware through ethical penetration testing and responsible disclosure practices.
4. **Design** intrusion detection and anomaly detection systems for robotic platforms.
5. **Develop** a comprehensive security framework and incident response plan for robotic systems.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Traditional lectures supplemented with real-world cybersecurity incident case studies involving robotic and CPS environments to ground theoretical concepts in practice.
2. **Simulation Labs:** Incorporating brainstorming sessions, group discussions, focused listening exercises, formulating higher-order questions, notetaking, annotation, and role-playing threat actor and defender scenarios.
3. **Algorithm Implementation:** Showing curated video and animation content to explain functioning of cryptographic protocols, attack vectors, intrusion detection mechanisms, and secure ROS 2 architectures.
4. **Mission Design Project:** Encouraging team-based lab activities such as group penetration testing exercises, joint security audits, and collaborative incident response simulations.
5. **Critical Thinking and Higher-Order Assessment:** Posing complex, scenario-based quiz questions and programming challenges that require students to think analytically about threat modelling, vulnerability exploitation, and security framework design.
6. **Creative Problem Solving:** Encouraging students to explore diverse approaches to securing robotic systems and to develop original security methodologies, tools, and mitigation strategies beyond prescribed solutions.

MODULE – I	**9 Hours**

Cyber-Physical Systems Security Fundamentals

Topics: CPS and IoT attack surfaces; CIA triad in robotics context; STRIDE threat modelling; Common Vulnerabilities and Exposures (CVE) for robotics; ICS/SCADA security; Robot-specific attack vectors: sensor spoofing, actuator hijacking, kinematic attacks.

Applications: Industrial robot security; Medical robot safety; UAV cyber defence.

AI Integration: AI-powered threat intelligence; Automated vulnerability scanning.

Authentic Intelligence: Students develop the ability to analyse and defend robotic systems against real-world cyber-physical attacks by integrating security principles with AI-driven threat detection.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building AI-Based Cybersecurity Pipelines**
   https://www.nvidia.com/en-gb/training/instructor-led-workshops/building-ai-based-cybersecurity-pipelines/

MODULE – II	9 Hours

Cryptography and Secure Communication

Topics: Symmetric and asymmetric encryption: AES, RSA, ECC; Digital signatures and PKI; TLS/DTLS for robotic communication; SROS2: ROS 2 security with DDS security; Secure key management; Hardware Security Modules (HSM) for robots.

Applications: Secure robot fleet communication; Authenticated robot commands; Drone telemetry encryption.

AI Integration: Post-quantum cryptography for robotic systems; AI-assisted key management.

Authentic Intelligence: Students design secure communication architectures ensuring confidentiality, integrity, and authentication in robotic networks.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Accelerated Computing**
   https://www.nvidia.com/en-us/learn/learning-path/accelerated-computing/

MODULE – III	9 Hours

ROS 2 Security and Vulnerability Analysis

Topics: ROS 2 DDS threat model; SROS2 configuration and certificate management; Network traffic analysis of ROS 2 topics; Common ROS 2 vulnerabilities: unauthenticated topic injection, replay attacks; Penetration testing methodology for robotic systems; Secure ROS 2 deployment patterns.

Applications: Securing hospital robots; Protecting industrial ROS 2 deployments; ROS 2 audit.

AI Integration: AI-assisted fuzzing for ROS 2 interfaces; ML-based anomaly detection in ROS topics.

Authentic Intelligence: Students gain hands-on expertise in identifying vulnerabilities and securing robotic middleware systems against cyber threats.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building Robotics Applications Using ROS and NVIDIA Isaac**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1

MODULE – IV	9 Hours

Intrusion and Anomaly Detection

Topics: Network intrusion detection systems (IDS) for CPS; ML-based anomaly detection: isolation forest, autoencoders, LSTM for time-series; Behavioural monitoring of robot actuators; Digital forensics for robotic incidents; Honeypot deployment in robotic networks; SIEM for robotic infrastructure.

Applications: Industrial robot anomaly detection; Drone hijacking detection; Fleet security monitoring. **AI Integration:** Deep learning IDS for CPS; Federated anomaly detection for robot fleets. **Authentic Intelligence:** Students build intelligent monitoring systems capable of detecting anomalies and cyber intrusions in robotic environments.
Key Online Resources / DLI Courses (Relevant to this Module): **1. Building Real-Time Cybersecurity Applications with AI** **https://learn.nvidia.com/courses/course?course_id=course-v1:DLI+S-AC-06+V1**

MODULE – V	**9 Hours**
Security Frameworks and Incident Response **Topics:** IEC 62443 for industrial cybersecurity; NIST Cybersecurity Framework applied to robotics; Security risk assessment and FMEA; Secure development lifecycle (SDL) for robotic software; Incident response plan for robotic systems; Security testing and red team exercises; Regulatory compliance: NIS2 Directive, EU Cyber Resilience Act. **Applications:** Security audit of a robot manufacturing line; Incident response drill for hospital robots. **AI Integration:** AI-assisted security posture management; Automated compliance checking. **Authentic Intelligence:** Students design complete security frameworks and incident response strategies ensuring resilience of robotic systems.	
Key Online Resources / DLI Courses (Relevant to this Module): **1. Building Robotics Applications Using ROS and NVIDIA Isaac** **https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-OV-41+V1** **2. Transform Cybersecurity With Agentic Blueprints** **https://www.nvidia.com/en-us/on-demand/session/gtc25-s72417/**	

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Analyse cyber threats and attack surfaces specific to robotic and CPS environments.	**L4**
2	Apply cryptographic techniques and secure communication protocols to robotic systems.	**L3**
3	Identify and assess vulnerabilities in ROS 2 and robotic middleware.	**L4**
4	Design and implement anomaly detection and intrusion detection systems for robots.	**L4**
5	Develop security frameworks and incident response plans compliant with IEC 62443 and NIST CSF.	**L5**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	**3**	**3**	**3**
CO-2	3	3	3	2	2	2	1	3	2	1	2	**3**	**3**	**3**
CO-3	3	2	3	2	3	2	1	3	2	2	2	**3**	**3**	**3**
CO-4	3	2	3	3	3	3	2	3	3	2	3	**3**	**3**	**3**
CO-5	3	2	3	3	3	3	2	3	3	3	3	**3**	**3**	**3**

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

PSO-2: Develop AI-powered software systems to control and operate robotic systems.

PSO-3: Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. *W. Stallings – Cryptography and Network Security, 8th Ed., Pearson, 2019.*
2. *H. Benson & P. Campagna – Industrial Cybersecurity: Efficiently Secure Critical Infrastructure Systems, Packt, 2022.*

REFERENCE BOOKS:

1. SROS2 Documentation – github.com/ros2/sros2
2. IEC 62443 Standard – iec.ch/cybersecurity
3. NIST Cybersecurity Framework – nist.gov/cyberframework

E-RESOURCES:

1. SROS2 Documentation — https://github.com/ros2/sros2
2. ROS 2 Documentation — https://docs.ros.org/en/rolling/index.html
3. NIST Cybersecurity Framework — https://www.nist.gov/cyberframework
4. IEC 62443 Industrial Cybersecurity Standards — https://www.iec.ch/cybersecurity

RECOMMENDED TOOLS:

1. Programming Languages — Python 3.x, C/C++, Bash Scripting

2. Cybersecurity & Penetration Testing Tools — Wireshark, Nmap, Metasploit Framework, Burp Suite Community Edition, Kali Linux

3. Cryptography & Secure Communication — OpenSSL, GnuPG (GPG), HashiCorp Vault

4. ROS 2 Security & Robotics — ROS 2, SROS2, NVIDIA Isaac Sim, Gazebo Simulator

5. AI & Machine Learning Frameworks — TensorFlow, PyTorch, scikit-learn, Keras

6. Intrusion Detection & Monitoring — Snort IDS, Suricata IDS, Wazuh SIEM, Splunk, Elastic Security

7. AI-Assisted Security Tools — NVIDIA Morpheus, OpenAI API, Hugging Face Transformers

AUGMENTING LLMS USING RETRIEVAL AUGMENTED GENERATION

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER – III

Course Code	:		**Credits**	:	03
Hours / Week	:	03	**Total Hours**	:	60
L–T–P–J	:	3–0–0–0	**CIE + SEE**	:	60 + 40 Marks

Course Vision: To develop engineers proficient in designing, building, and deploying production-grade Retrieval Augmented Generation (RAG) systems that ground large language models in verified, domain-specific knowledge — mastering dense retrieval, vector databases, re-ranking, agentic RAG, GraphRAG, and evaluation frameworks — using NVIDIA NIM, cuGraph, and the Agentic AI Professional Certification path to build trustworthy, scalable, and responsible AI knowledge systems.

Course Objectives: This course will enable students to:

1. **Understand** the limitations of parametric LLM knowledge and the RAG paradigm — retrieval pipeline, chunking, embedding, and context injection.
2. **Implement** dense retrieval systems using embedding models, FAISS/ChromaDB/Milvus vector databases, and GPU-accelerated semantic search.
3. **Design** advanced RAG architectures — HyDE, FLARE, Self-RAG, Corrective RAG, and multi-hop reasoning pipelines for complex queries.
4. **Build** agentic RAG systems with tool-augmented LLMs, RAG agents, and scalable multi-agent workflows using NVIDIA NIM and LangChain.
5. **Evaluate** RAG pipeline quality using RAGAS, TruLens, and DeepEval metrics and apply responsible AI principles for grounded, transparent AI knowledge systems.

Teaching-Learning Process (General Instructions)

1. **Lecture Method:** Lectures on RAG theory with architecture diagrams and real-world deployment case studies from enterprise knowledge management and legal AI systems.
2. **Coding Labs:** Hands-on Python labs using LangChain, LlamaIndex, NVIDIA NIM, and ChromaDB to build progressively advanced RAG pipelines.
3. **Evaluation Workshops:** Students evaluate their RAG systems using RAGAS and TruLens frameworks — measuring faithfulness, answer relevancy, and context recall.
4. **Paper Reading Sessions:** Guided reading of landmark RAG papers — original RAG (Lewis 2020), Self-RAG, FLARE, HyDE, GraphRAG, and RAG-Fusion.
5. **Project-Based Learning:** Teams build a domain-specific RAG system — legal case retrieval, medical literature QA, or engineering standards search — deployed via NVIDIA NIM.
6. **Certification Preparation:** Structured NVIDIA Certified Professional Agentic AI exam preparation with practice questions, mock tests, and hands-on graded assignments.

MODULE – I	**9 Hours**

LLM Knowledge Limitations and RAG Foundations

Topics: Parametric vs. non-parametric knowledge in LLMs; Hallucination problem: causes, taxonomy, and measurement; Knowledge cutoff problem and context window constraints; RAG paradigm: retrieve → augment → generate; RAG pipeline components: document loader, text splitter, embedding model, vector store, retriever, LLM, and response synthesiser; Chunking strategies: fixed-size, recursive, semantic, sentence-window, document-aware chunking; Embedding models: sentence transformers (SBERT, E5, BGE, UAE-Large), OpenAI text-embedding-3; Sparse retrieval: TF-IDF, BM25 — Okapi BM25 formula; Dense retrieval: bi-encoder architecture, DPR (Dense Passage Retriever); Hybrid retrieval: combining BM25 + dense with RRF (Reciprocal Rank Fusion); Introduction to LangChain and LlamaIndex RAG frameworks.

Applications: Building a basic RAG QA system over NVIDIA documentation using LangChain and ChromaDB; BM25 vs. dense retrieval benchmark on a legal judgement dataset; Chunking strategy comparison — fixed vs. semantic chunking on a 500-page engineering textbook.

AI Integration: NVIDIA NIM embedding microservices for GPU-accelerated text embedding at scale; NVIDIA cuVS (RAPIDS) for GPU-accelerated approximate nearest neighbour search; Augment Your LLM Using Retrieval Augmented Generation (NVIDIA DLI).

Authentic Intelligence: Transparency obligations for RAG-generated content — source citation and attribution; Hallucination vs. retrieval failure — distinguishing model errors from retrieval gaps; Intellectual property and copyright of retrieved source documents in RAG responses; Privacy risks when RAG indexes proprietary or personal documents.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building RAG Agents with LLMs [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Introduction to Deploying RAG Pipelines for Production at Scale [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Evaluating RAG and Semantic Search Systems [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – II	**9 Hours**

Vector Databases and Semantic Search Infrastructure

Topics: Vector database fundamentals: index types — Flat, IVF, HNSW, PQ; Approximate Nearest Neighbour (ANN) algorithms: HNSW, FAISS (IndexIVFFlat, IndexIVFPQ); GPU-accelerated vector search with NVIDIA cuVS and RAPIDS; Vector database platforms: FAISS, ChromaDB, Milvus, Weaviate, Qdrant, Pinecone — feature comparison; Metadata filtering: combining semantic search with structured filters; Hybrid search architectures: dense + sparse with RRF re-ranking; Embedding model selection: domain-specific fine-tuning of embedding models; Multi-modal embeddings: CLIP for image-text RAG; Embedding dimensionality reduction: PCA, UMAP for visualisation; Scalability: vector database sharding, replication, and production deployment; NVIDIA cuGraph for GPU-accelerated graph-based vector indexing.

Applications: Deploying a production Milvus vector database cluster for a 10M-document enterprise knowledge base; GPU-accelerated FAISS index (IVFPQ) achieving <5ms latency on 1M embedding queries; Multi-modal image-text RAG for engineering diagram retrieval using CLIP embeddings.

AI Integration: NVIDIA cuVS (RAPIDS) for GPU-accelerated ANN search — 100x speedup over CPU FAISS; Use GPUs to Accelerate Vector Databases (NVIDIA DLI); NVIDIA NIM embedding microservices for high-throughput production embedding generation.

Authentic Intelligence: Data governance for vector databases: PII detection and redaction before indexing; Secure vector database access control — namespace isolation and RBAC; Bias in embedding spaces — demographic and linguistic bias in pretrained embedding models; Responsible choice of vector database vendor — data residency and vendor lock-in risks.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Building RAG Agents with LLMs [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Evaluating RAG and Semantic Search Systems [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – III	**9 Hours**

Advanced RAG Architectures and Multi-Hop Reasoning

Topics: Naive RAG limitations: low retrieval precision, lost-in-the-middle problem, context window saturation; Query transformation: HyDE (Hypothetical Document Embeddings), query expansion, step-back prompting, sub-question decomposition; Re-ranking: cross-encoder re-rankers (BGE-reranker, Cohere Rerank), ColBERT late interaction; Corrective RAG (CRAG): relevance scoring and web search fallback; Self-RAG: self-reflection tokens — retrieve, critique, and refine; FLARE: forward-looking active retrieval; Multi-hop reasoning: iterative retrieval for complex multi-step questions (IRCoT, BeamRetrieval); RAG-Fusion: parallel query generation and reciprocal rank fusion; Contextual compression: LLM-based document compression, metadata-aware filtering; Sentence-window and parent-document retrieval strategies; Long-context RAG vs. RAG for extended context models.

Applications: HyDE pipeline improving retrieval recall by 23% on TechQA benchmark; Cohere Rerank integration increasing answer faithfulness in a legal case RAG system; IRCoT multi-hop reasoning chain for complex 3-hop questions on MuSiQue dataset.

AI Integration: NVIDIA NIM reranker microservices for production-grade cross-encoder re-ranking; GPU-accelerated ColBERT late interaction retrieval with NVIDIA cuVS; LlamaIndex advanced RAG modules (HyDE, sub-question engine, recursive retrieval).

Authentic Intelligence: Retrieval faithfulness vs. fluency trade-off in advanced RAG; Transparency of multi-hop reasoning chains for end-user trust; Avoiding retrieval hallucination amplification in iterative RAG pipelines; Responsible web-search fallback in Corrective RAG — source credibility and misinformation risk.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building RAG Agents with LLMs [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Adding New Knowledge to LLMs [Instructor-Led Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Introduction to Deploying RAG Pipelines for Production at Scale [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – IV	**9 Hours**

Agentic RAG, GraphRAG, and Production Deployment

Topics: Agentic RAG: ReAct framework, tool-augmented LLMs, RAG as a tool; RAG agents with LangChain and LlamaIndex agent frameworks; Multi-agent RAG workflows: research agents, summarisation agents, verification agents; GraphRAG: Microsoft GraphRAG — community detection, global and local search modes; NVIDIA cuGraph-accelerated GraphRAG with NVIDIA NIM; Knowledge graph construction from documents: entity extraction, relation extraction, graph population; Production RAG deployment: FastAPI + LangServe, Triton + TensorRT-LLM for LLM, NVIDIA NIM; RAG pipeline optimisation: latency profiling, caching (semantic cache, prompt cache), async retrieval; Streaming RAG responses for real-time user experience; Document preprocessing pipeline: PDF extraction (PyMuPDF, Unstructured), OCR (Tesseract, NVIDIA OCR), table extraction.

Applications: Agentic RAG research assistant using LangChain ReAct agent with web search + vector DB tools; GraphRAG system over 50,000 research papers using NVIDIA cuGraph for community-aware global summarisation; Production RAG API with semantic caching reducing LLM calls by 40% on repeated queries.

AI Integration: Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM (NVIDIA DLI); NVIDIA NIM for production LLM and embedding microservice deployment in RAG; Introduction to Deploying RAG Pipelines for Production at Scale (NVIDIA DLI).

Authentic Intelligence: Agentic RAG autonomy limits — when should RAG agents require human confirmation; Security vulnerabilities in tool-augmented RAG agents — prompt injection via retrieved content; Responsible GraphRAG: privacy risks in entity graph construction from personal documents; Production RAG logging and audit trails for regulated industries (healthcare, finance, legal).

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Building RAG Agents with LLMs [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Structure From Chaos: Accelerate GraphRAG With cuGraph and NVIDIA NIM [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **Introduction to Deploying RAG Pipelines for Production at Scale [Self-paced + Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths

MODULE – V	**9 Hours**

RAG Evaluation, Responsible AI, and NVIDIA Agentic AI Certification

Topics: RAG evaluation framework: component-level vs. end-to-end evaluation; RAGAS metrics: faithfulness, answer relevancy, context recall, context precision, answer correctness; TruLens evaluation: RAG triad — context relevance, groundedness, answer relevance; DeepEval: LLM-as-judge evaluation, hallucination score, toxicity; Retrieval evaluation: MRR, NDCG, Recall@K on benchmark datasets (BEIR, MTEB); Human evaluation protocols for RAG quality; RAG failure mode taxonomy: retrieval failure, context window overflow, conflicting sources, hallucination; NVIDIA Certified Professional Agentic AI examination overview and preparation; Responsible RAG: bias in knowledge bases, outdated retrieved content, source diversity; RAG governance: content moderation, PII redaction, access-controlled retrieval; Future of RAG: multimodal RAG (video, audio, image), agentic retrieval with memory systems.

Applications: RAGAS evaluation dashboard for a 100-question QA benchmark on a company knowledge base RAG system; TruLens faithfulness monitoring for a medical literature RAG with automated alerts for low-

groundedness responses; NVIDIA Agentic AI Certification preparation: end-to-end graded RAG project submission.

AI Integration: NVIDIA Certified Professional Agentic AI Certification Exam preparation and submission; Evaluating RAG and Semantic Search Systems (NVIDIA DLI); NVIDIA NeMo Guardrails for responsible content filtering in deployed RAG systems.

Authentic Intelligence: Epistemic responsibility of RAG-generated answers in high-stakes domains; Auditability of RAG pipelines — complete retrieval and generation logging for accountability; Equity in RAG knowledge bases — underrepresentation of non-English and low-resource language content; Responsible decommissioning of RAG systems handling sensitive personal data.

Key Online Resources / DLI Courses (Relevant to this Module):

1. **Evaluating RAG and Semantic Search Systems [Self-paced]**
   https://learn.nvidia.com/en-us/training/learning-paths
2. **Adding New Knowledge to LLMs [Instructor-Led Workshop ★]**
   https://learn.nvidia.com/en-us/training/learning-paths
3. **NVIDIA Certified Professional Agentic AI [Certification Exam + Digital Badge ★]**
   https://learn.nvidia.com/en-us/training/learning-paths

Course Outcome

Course Outcome		
Course Outcomes (COs)	**Description**	**Bloom's Taxonomy Level**
At the end of the course the student will be able to:		
1	Explain the RAG paradigm, LLM knowledge limitations, chunking strategies, and dense vs. sparse retrieval architectures.	**L2**
2	Implement GPU-accelerated vector database systems using FAISS, Milvus, and NVIDIA cuVS for high-throughput semantic search.	**L3**
3	Design advanced RAG pipelines incorporating HyDE, self-RAG, re-ranking, multi-hop reasoning, and corrective retrieval strategies.	**L6**
4	Build agentic RAG systems and GraphRAG workflows deployed at production scale using NVIDIA NIM, LangChain, and cuGraph.	**L6**
5	Evaluate RAG pipeline quality using RAGAS and TruLens metrics and apply responsible AI governance for grounded and transparent knowledge systems.	**L5**

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs														
	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	–	3	–	–	2	3	3	3
CO-2	3	3	3	2	3	2	–	3	2	–	2	3	3	3
CO-3	3	3	3	3	3	3	–	2	2	1	2	3	3	3
CO-4	3	2	3	3	3	3	1	3	2	1	2	3	3	3
CO-5	2	2	2	2	3	3	3	3	2	2	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

PSO-1: Design and build production RAG and agentic AI knowledge systems integrating vector databases, LLMs, and NVIDIA NIM microservices.

PSO-2: Develop GPU-accelerated RAG pipelines using NVIDIA cuVS, cuGraph, and NIM for grounded, scalable, and low-latency AI knowledge retrieval systems.

PSO-3: Apply RAG, GraphRAG, and agentic AI to solve real-world problems in enterprise knowledge management, legal AI, clinical decision support, and engineering documentation systems.

TEXT BOOKS:

1. *Jerry Liu, LlamaIndex: Building Data-Augmented Applications with LLMs, O'Reilly Media, 2024.*
2. *Ben Auffarth, LangChain Masterclass: Build LLM-Powered Applications with Python, Packt Publishing, 2024.*
3. *Lewis et al., Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, NeurIPS 2020. arXiv:2005.11401.*

REFERENCE BOOKS:

1. NVIDIA NIM Documentation: https://www.nvidia.com/en-us/ai/
2. Microsoft GraphRAG Documentation: https://microsoft.github.io/graphrag/
3. RAGAS Evaluation Framework: https://docs.ragas.io/

E-RESOURCES:

1. NVIDIA DLI — Agentic AI: Professional Certification Path: https://learn.nvidia.com/en-us/training/learning-paths
2. NVIDIA Certified Professional Agentic AI Exam: https://learn.nvidia.com/en-us/training/learning-paths
3. LangChain Documentation: https://docs.langchain.com/
4. LlamaIndex Documentation: https://docs.llamaindex.ai/
5. BEIR Retrieval Benchmark: https://github.com/beir-cellar/beir

RECOMMENDED TOOLS:

1. RAG Frameworks: LangChain, LlamaIndex, Haystack, NVIDIA NIM
2. Vector Databases: FAISS, ChromaDB, Milvus, Weaviate, Qdrant, Pinecone
3. GPU Acceleration: NVIDIA cuVS, RAPIDS, NVIDIA NIM Embeddings
4. Graph RAG: NVIDIA cuGraph + NIM, Microsoft GraphRAG, Neo4j
5. Evaluation: RAGAS, TruLens, DeepEval, MTEB, BEIR Benchmark
6. Deployment: FastAPI + LangServe, NVIDIA Triton + TensorRT-LLM, Docker

SEMESTER - IV

CAPSTONE PROJECT

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - IV

Course Code	:	26RA23XX	**Credits**	:	12
Hours / Week	:	24	**Total Hours**	:	-
L–T–P–J	:	0–0–0–24	**CIE**	:	100 Marks (CIE only)

Course Vision:
To provide students with a faculty-supervised research and innovation experience that enables them to identify, investigate and solve a real-world problem in Artificial Intelligence and Robotics through systematic research, engineering design, prototype development, validation and professional dissemination.

Course Objectives:
This course will enable students to:

1. Identify a significant research or industry problem in AI and Robotics.
2. Conduct a systematic literature review using leading scholarly databases.
3. Design a technically sound solution with appropriate algorithms, datasets and software/hardware tools.
4. Develop, validate and evaluate a working prototype using reproducible experimental practices.
5. Communicate research outcomes through technical reports, presentations and publications.

Teaching-Learning Process (General Instructions)
These are sample innovative pedagogical methods that may be adopted to accelerate the attainment of the various course outcomes:

1. Faculty-guided project-based learning with periodic technical reviews.
2. Research mentoring through literature discussions and design reviews.
3. Simulation, prototyping and experimental validation using modern AI and robotics tools.
4. Independent learning supported by AI-assisted research and software engineering practices.
5. Emphasis on reproducibility, documentation, ethical AI and professional communication.

MODULE – I

Problem Definition, Literature Review and Project Planning
Topics: Problem formulation; systematic literature review using IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect and Scopus; research gap identification; feasibility analysis; SRS; benchmarking; project planning; proposal seminar.
Applications: Healthcare AI, autonomous robots, computer vision, NLP, industrial automation, edge AI.
AI Integration: PyTorch, ROS2, Hugging Face, LangChain, GitHub Copilot, ChatGPT, Claude, Isaac Sim.
Authentic Intelligence: Responsible literature review, originality, ethical AI, transparent planning and research integrity.

Key Online Resources (Relevant to this Module):

1. **Generative AI Explained**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-07+V1
2. **Building AI Agents with Multimodal Models**
   https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+C-FX-17+V1

MODULE – II
System Design, Development and Validation **Topics:** Architecture design; algorithm selection; dataset preparation; implementation; version control; experiment tracking; statistical validation; ablation studies; error analysis; baseline comparison; reproducibility; demonstration; thesis preparation. **Applications:** Project-specific AI and Robotics applications. **AI Integration:** PyTorch, TensorFlow, MLflow, Weights & Biases, DVC, Docker, ROS2, Isaac Lab. **Authentic Intelligence:** Responsible experimentation, reproducible research, transparent evaluation and trustworthy deployment.
Key Online Resources (Relevant to this Module): 1. **Introduction to NVIDIA NIM Microservices** https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-FX-23+V1 2. **An Introduction to Robot Learning and Isaac Lab** https://learn.nvidia.com/courses/course-detail?course_id=course-v1:DLI+S-RX-06+V1

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Define a well-scoped research problem and project plan.	L3
2	Conduct a systematic literature review and identify research gaps.	L4
3	Design and justify an AI & Robotics solution architecture.	L5
4	Develop, validate and benchmark a working prototype using reproducible practices.	L6
5	Communicate outcomes through reports, demonstrations, publications and professional presentations.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.
- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. Wohlin et al., Experimentation in Software Engineering, Springer.
2. Booth et al., The Craft of Research, 5th Edition, University of Chicago Press.
3. Zobel, Writing for Computer Science, Springer.

REFERENCE BOOKS:

1. Creswell & Creswell, Research Design.
2. IEEE Author Center and Conference Publication Guidelines.

E-RESOURCES:

- IEEE Xplore
- ACM Digital Library
- SpringerLink
- ScienceDirect
- Google Scholar
- Semantic Scholar
- Overleaf
- GitHub Docs
- PyTorch Docs
- Hugging Face Docs

RECOMMENDED TOOLS

- Python, NumPy, Pandas
- PyTorch, TensorFlow
- ROS2, Gazebo, RViz2, Isaac Sim, Isaac Lab
- Git, GitHub, GitHub Actions
- Docker
- MLflow, Weights & Biases, DVC
- LangChain, LlamaIndex
- VS Code, JupyterLab, Google Colab
- Overleaf and LaTeX

INDUSTRY / RESEARCH INTERNSHIP

[As per Choice Based Credit System (CBCS) scheme]

SEMESTER - IV

Course Code	:		**Credits**	:	04
Hours / Week	:	08	**Total Hours**	:	
L–T–P–J	:	0–0–0–8	**CIE+SEE**	:	100 Marks (CIE only)

Course Vision:
To provide immersive industry or research experience that enables students to apply advanced AI and Robotics knowledge to real-world engineering problems while developing technical competence, innovation, professionalism, ethical responsibility and lifelong learning skills.

Course Objectives:
This course will enable students to:

1. Apply advanced AI and Robotics concepts in industrial or research environments.
2. Gain hands-on experience with modern software, hardware and deployment workflows.
3. Collaborate effectively with multidisciplinary teams under professional mentorship.
4. Develop reproducible technical documentation and engineering reports.
5. Demonstrate professional ethics, communication and project management skills.

Teaching-Learning Process (General Instructions)
These are sample new pedagogical methods, where teacher can use to accelerate the attainment of the various course outcomes:

1. Faculty and industry mentor guided experiential learning.
2. Weekly technical reviews and reflective learning.
3. Hands-on implementation using industrial tools and workflows.
4. Technical seminars, demonstrations and peer discussions.
5. Continuous assessment through reports, presentations and mentor feedback.

MODULE – I

Internship Planning and Professional Practice
Topics: Problem definition; requirement analysis; work package allocation; technology selection; implementation planning; weekly progress reviews; interaction with mentors; engineering ethics; workplace safety.
Applications: Industrial automation, autonomous systems, embedded AI, healthcare, manufacturing and research laboratories.
AI Integration: PyTorch, TensorFlow, ROS2, Isaac Sim, GitHub, Docker, NVIDIA AI ecosystem.
Authentic Intelligence: Professional conduct, responsible AI deployment, teamwork, transparency and engineering integrity.

MODULE – II
Implementation, Reporting and Knowledge Dissemination **Topics:** System implementation; validation; benchmarking; documentation; code repository management; technical presentation; internship report; reflection; future work recommendations. **Applications:** Prototype development, product engineering, research validation and technology transfer. **AI Integration:** MLflow, Weights & Biases, GitHub Actions, AI-assisted documentation, data visualization tools. **Authentic Intelligence:** Reproducible engineering practice, ethical reporting, lifelong learning and professional growth.

Course Outcome

Course Outcomes (COs)	Description	Bloom's Taxonomy Level
At the end of the course the student will be able to:		
1	Apply advanced AI and Robotics knowledge to solve real-world engineering problems.	L3
2	Work effectively in multidisciplinary professional teams.	L3
3	Develop reproducible technical artefacts including code repositories and documentation.	L4
4	Prepare a professional internship report demonstrating technical competence and critical evaluation.	L5
5	Present outcomes professionally and reflect on technical and professional growth.	L5

Mapping Levels of COs to POs / PSOs

Table: Mapping Levels of COs to POs / PSOs

	Program Outcomes (POs)											PSOs		
COS	1	2	3	4	5	6	7	8	9	10	11	1	2	3
CO-1	3	2	2	1	2	2	1	3	1	1	2	3	3	3
CO-2	3	3	3	2	2	2	1	3	2	1	2	3	3	3
CO-3	3	2	3	2	3	2	1	3	2	2	2	3	3	3
CO-4	3	2	3	3	3	3	2	3	3	2	3	3	3	3
CO-5	3	2	3	3	3	3	2	3	3	3	3	3	3	3

3: Substantial (High) 2: Moderate (Medium) 1: Poor (Low)

- **PSO-1:** Design and build robotic systems for different applications using analytical, logical, and problem-solving skills.

- **PSO-2:** Develop AI-powered software systems to control and operate robotic systems.
- **PSO-3:** Apply robotics and AI to solve real-world problems in a variety of domains, including industrial robotics, service robots, exoskeletons, surgical robots, delivery vehicles, autonomous vehicles, and crewless micro aerial vehicles

TEXT BOOKS:

1. Booth et al., The Craft of Research, 5th Edition.
2. Zobel, Writing for Computer Science, 3rd Edition.
3. Industry/Research Organisation Technical Documentation.

REFERENCE BOOKS:

1. Creswell & Creswell, Research Design.
2. IEEE Author Center and Publication Guidelines.

E-RESOURCES:

- IEEE Xplore
- ACM Digital Library
- Google Scholar
- NVIDIA DLI
- GitHub Documentation
- Hugging Face Documentation
- Kaggle Learn
- Overleaf

Recommended Tools

- Python, JupyterLab, VS Code
- PyTorch, TensorFlow
- ROS2, Gazebo, RViz2, Isaac Sim
- Git, GitHub, GitHub Actions
- Docker
- MLflow, Weights & Biases
- Google Colab, Kaggle
- Overleaf and LaTeX