



National University of Engineering (UNI)
School of Computer Science
Syllabus 2027-I

1. COURSE

AI268. Computational Vision (Elective)

2. GENERAL INFORMATION

2.1 Course	: AI268. Computational Vision
2.2 Semester	: 8 th Semester.
2.3 Credits	: 3
2.4 Hours	: 2 HT; 2 HP;
2.5 Duration of the period	: 16 weeks
2.6 Type of course	: Elective
2.7 Learning modality	: Face to face
2.8 Prerequisites	: AI264. Deep Learning. (7 th Sem)

3. PROFESSORS

Meetings after coordination with the professor

4. INTRODUCTION TO THE COURSE

Computer Vision is a key area of Artificial Intelligence that enables machines to interpret and understand the visual world. This course covers fundamental techniques for image processing, feature extraction, and object recognition. With the advent of deep learning, computer vision has become essential for autonomous vehicles, medical imaging, and augmented reality. Students will learn the mathematical foundations and algorithmic approaches to solve real-world visual perception problems.

5. GOALS

- Understand the fundamentals of image formation and digital representation.
- Master image processing techniques such as filtering and edge detection.
- Apply feature extraction methods to identify key patterns in images.
- Implement object detection and recognition using deep learning models.
- Analyze 3D vision and motion estimation techniques.

6. COMPETENCES

AG-C08) Problem Analysis: Identifies, formulates, and analyzes complex computing problems. (Usage)

AG-C09) Design and Development of Solutions: Designs, implements, and evaluates solutions for complex computing problems. (Assessment)

AG-C12) Uses systemic approaches to select, develop, apply, integrate, and manage secure computing technologies to achieve user objectives. (Usage)

7. TOPICS

Unit 1: Introduction to AI: Concepts and Intelligent Behavior (5 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Overview of AI problems, Examples of successful recent AI applications • Definitions of agents with examples (e.g., reactive, deliberative) • What is intelligent behavior?: – The Turing test and its flaws – Multimodal input and output – Simulation of intelligent behavior – Rational versus non-rational reasoning • Overview of AI Applications, growth, and impact (economic, societal, ethics) • Philosophical issues • History of AI	• Describe the Turing test and the "Chinese Room" thought experiment [Explain] • Differentiate between optimal reasoning/behavior and human-like reasoning/behavior [Evaluate] • Differentiate the terms: AI, machine learning, and deep learning [Explain]
Readings : [Szeliski2022], [Gonzales2018]	

Unit 2: Problem Characteristics and Agent Nature (5 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Problem characteristics: – Fully versus partially observable – Single versus multi-agent – Deterministic versus stochastic – Static versus dynamic – Discrete versus continuous • Nature of agents: – Autonomous, semi-autonomous, mixed-initiative autonomy – Reflexive, goal-based, and utility-based – Decision making under uncertainty and with incomplete information – The importance of perception and environmental interactions – Learning-based agents – Embodied agents: * sensors, dynamics, effectors • Practice identifying problem characteristics in example environments • Additional depth on nature of agents with examples • Additional depth on AI Applications, Growth, and Impact (economic, societal, ethics, security)	• Enumerate the characteristics of a specific problem [Apply]
Readings : [Szeliski2022], [Gonzales2018]	

Unit 3: Perception and Computer Vision (12 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Computer vision: – Image acquisition, representation, processing, and properties – Shape representation, object recognition, and segmentation – Motion analysis – Generative models • Audio and speech recognition • Touch and proprioception • Other modalities (e.g., olfaction) • Modularity in recognition • Approaches to pattern recognition: – Classification algorithms and measures of classification quality – Statistical techniques – Deep learning techniques	• Summarize the importance of image and object recognition in AI and indicate several significant applications of this technology [Summarize] • List at least three image-segmentation approaches, such as thresholding, edge-based and region-based algorithms, along with their defining characteristics, strengths, and weaknesses [List] • Implement 2d object recognition based on contour-based and/or region-based shape representations [Apply] • Distinguish the goals of sound-recognition, speech-recognition, and speaker-recognition and identify how the raw audio signal will be handled differently in each of these cases [Distinguish] • Provide at least two examples of a transformation of a data source from one sensory domain to another, e.g., tactile data interpreted as single-band 2d images [Generate] • Implement a feature-extraction algorithm on real data, e.g., an edge or corner detector for images or vectors of Fourier coefficients describing a short slice of audio signal [Apply] • Implement an algorithm combining features into higher-level percepts, e.g., a contour or polygon from visual primitives or phoneme hypotheses from an audio signal [Apply] • Implement a classification algorithm that segments input percepts into output categories and quantitatively evaluates the resulting classification [Apply] • Evaluate the performance of the underlying feature-extraction, relative to at least one alternative possible approach (whether implemented or not) in its contribution to the classification task (8), above [Evaluate] • Describe at least three classification approaches, their prerequisites for applicability, their strengths, and their shortcomings [Describe] • Implement and evaluate a deep learning solution to problems in computer vision, such as object or scene recognition [Apply]
Readings : [Szeliski2022], [Prince2012]	

Unit 4: Machine Learning Fundamentals (5 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Definition and examples of a broad variety of machine learning tasks: – Supervised learning: * Classification * Regression – Reinforcement learning – Unsupervised learning: * Clustering • Fundamental ideas: – No free lunch theorem: no one learner can solve all problems; representational design decisions have consequences. – Sources of error and undecidability in machine learning • A simple statistical-based supervised learning such as linear regression or decision trees: – Focus on how they work without going into mathematical or optimization details; enough to understand and use existing implementations correctly • The overfitting problem/controlling solution complexity (regularization, pruning - intuition only): – The bias (underfitting) - variance (overfitting) tradeoff	• Describe the differences among the three main styles of learning (supervised, reinforcement, and unsupervised) and determine which is appropriate to a particular problem domain [Explain] • Differentiate the terms of AI, machine learning, and deep learning [Evaluate] • Frame an application as a classification problem, including the available input features and output to be predicted (e.g., identifying alphabetic characters from pixel grid input) [Apply] • Identify overfitting in the context of a problem and learning curves and describe solutions to overfitting [Evaluate] • Explain how machine learning works as an optimization/search process [Explain]
Readings : [Goodfellow2016], [Russell2020]	

Unit 5: Machine Learning Data, Representations and Evaluation (5 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Working with Data: – Data preprocessing: * Importance and pitfalls of preprocessing choices – Handling missing values (imputing, flag-as-missing): * Implications of imputing vs flag-as-missing – Encoding categorical variables, encoding real-valued data – Normalization/standardization – Emphasis on real data, not textbook examples • Representations: – Hypothesis spaces and complexity – Simple basis feature expansion, such as squaring univariate features – Learned feature representations • Machine learning evaluation: – Separation of train, validation, and test sets – Performance metrics for classifiers – Estimation of test performance on held-out data – Tuning the parameters of a machine learning model with a validation set – Importance of understanding what a model is doing, where its pitfalls/shortcomings are, and the implications of its decisions	• Apply two or more simple statistical learning algorithms to a classification task and measure the classifiers' accuracy [Apply] • Implement a statistical learning algorithm and the corresponding optimization process to train the classifier and obtain a prediction on new data [Apply] • Explain proper ML evaluation procedures, including the differences between training and testing performance, and what can go wrong with the evaluation process leading to inaccurate reporting of ML performance [Explain] • Compare two machine learning algorithms on a dataset, implementing the data preprocessing and evaluation methodology (e.g., metrics and handling of train/test splits) from scratch [Apply]
Readings : [Goodfellow2016], [Russell2020]	

Unit 6: Neural Networks and Ethics (4 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Basic neural networks: – Fundamentals of understanding how neural networks work and their training process, without details of the calculations – Basic introduction to generative neural networks (e.g., large language models) • Ethics for Machine Learning: – Focus on real data, real scenarios, and case studies – Dataset/algorithmic/evaluation bias and unintended consequences	• Describe the neural network training process and resulting learned representations [Explain]
Readings : [Goodfellow2016], [Russell2020]	

Unit 7: Machine Learning Formulation and Evaluation (3 hours)	
Competences Expected: AG-C09	
Topics	Learning Outcomes
• Formulation of simple machine learning as an optimization problem, such as least squares linear regression or logistic regression: – Objective function – Gradient descent – Regularization to avoid overfitting (mathematical formulation) • Ensembles of models: – Simple weighted majority combination • Performance evaluation: – Other metrics for classification (e.g., error, precision, recall) – Performance metrics for regressors – Confusion matrix – Cross-validation: * Parameter tuning (grid/random search, via cross-validation)	• Compare and contrast several learning techniques (e.g., decision trees, logistic regression, naive Bayes, neural networks, and belief networks), providing examples of when each strategy is superior [Compare] • Evaluate the performance of a simple learning system on a real-world dataset [Evaluate] • Characterize the state of the art in learning theory, including its achievements and shortcomings [Explain] • Explain the problem of overfitting, along with techniques for detecting and managing the problem [Explain] • Explain the triple tradeoff among the size of a hypothesis space, the size of the training set, and performance accuracy [Explain]
Readings : [Goodfellow2016]	

Unit 8: Deep Learning, RL and Ethics (3 hours)	
Competences Expected: AG-C09	
Topics	Learning Outcomes
• Deep learning: – Deep feed-forward networks (intuition only, no mathematics) – Convolutional neural networks (intuition only, no mathematics) – Visualization of learned feature representations from deep nets – Other architectures (generative NN, recurrent NN, transformers, etc.) • Overview of reinforcement learning methods • Two or more applications of machine learning algorithms: – E.g., medicine and health, economics, vision, natural language, robotics, game play • Ethics for Machine Learning: – Continued focus on real data, real scenarios, and case studies – Privacy – Fairness – Intellectual property – Explainability	• Visualize the training progress of a neural network through learning curves in a well-established toolkit (e.g., TensorBoard) and visualize the learned features of the network [Apply] • Given a real-world application of machine learning, describe ethical issues regarding the choices of data, preprocessing steps, algorithm selection, and visualization/presentation of results [Apply]
Readings : [Goodfellow2016]	

Unit 9: Perception and Computer Vision (12 hours)	
Competences Expected: AG-C08	
Topics	Learning Outcomes
• Computer vision: – Image acquisition, representation, processing, and properties – Shape representation, object recognition, and segmentation – Motion analysis – Generative models • Audio and speech recognition • Touch and proprioception • Other modalities (e.g., olfaction) • Modularity in recognition • Approaches to pattern recognition: – Classification algorithms and measures of classification quality – Statistical techniques – Deep learning techniques	• Summarize the importance of image and object recognition in AI and indicate several significant applications of this technology [Summarize] • List at least three image-segmentation approaches, such as thresholding, edge-based and region-based algorithms, along with their defining characteristics, strengths, and weaknesses [List] • Implement 2d object recognition based on contour-based and/or region-based shape representations [Apply] • Distinguish the goals of sound-recognition, speech-recognition, and speaker-recognition and identify how the raw audio signal will be handled differently in each of these cases [Distinguish] • Provide at least two examples of a transformation of a data source from one sensory domain to another, e.g., tactile data interpreted as single-band 2d images [Generate] • Implement a feature-extraction algorithm on real data, e.g., an edge or corner detector for images or vectors of Fourier coefficients describing a short slice of audio signal [Apply] • Implement an algorithm combining features into higher-level percepts, e.g., a contour or polygon from visual primitives or phoneme hypotheses from an audio signal [Apply] • Implement a classification algorithm that segments input percepts into output categories and quantitatively evaluates the resulting classification [Apply] • Evaluate the performance of the underlying feature-extraction, relative to at least one alternative possible approach (whether implemented or not) in its contribution to the classification task (8), above [Evaluate] • Describe at least three classification approaches, their prerequisites for applicability, their strengths, and their shortcomings [Describe] • Implement and evaluate a deep learning solution to problems in computer vision, such as object or scene recognition [Apply]
Readings : [Szeliski2022]	

8. WORKPLAN

8.1 Methodology

Individual and team participation is encouraged to present their ideas, motivating them with additional points in the different stages of the course evaluation.

8.2 Theory Sessions

The theory sessions are held in master classes with activities including active learning and roleplay to allow students to internalize the concepts.

8.3 Practical Sessions

The practical sessions are held in class where a series of exercises and/or practical concepts are developed through problem solving, problem solving, specific exercises and/or in application contexts.

9. EVALUATION SYSTEM

Theory Sessions:

The theory sessions are held in master classes with activities including active learning and roleplay to allow students to internalize the concepts.

Lab Sessions:

In order to verify their competences, several activities including active learning and roleplay will be developed during lab sessions.

Oral Presentations:

Individual and team participation is encouraged to present their ideas, motivating them with additional points in the different stages of the course evaluation.

Reading:

Throughout the course different readings are provided, which are evaluated. The average of the notes in the readings is considered as the mark of a qualified practice. The use of our virtual campus allows each student to access the course information, and interact outside the classroom with the teacher and with the other students.

The final grade NF depends on the following grades:

- PE : Partial Exam.
- FE : Final Exam.
- P : Permanent Evaluation (Average of practical work and graded assignments).

The final grade NF is calculated as follows:

$$NF = \frac{PE + FE + P}{3}$$

10. BASIC BIBLIOGRAPHY