



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

1. COURSE

CS261. Artificial Intelligence (Mandatory)

2. GENERAL INFORMATION

2.1 Course	: CS261. Artificial Intelligence
2.2 Semester	: 5 th Semester.
2.3 Credits	: 4
2.4 Hours	: 2 HT; 4 HP;
2.5 Duration of the period	: 16 weeks
2.6 Type of course	: Mandatory
2.7 Learning modality	: Face to face
2.8 Prerequisites	: STA251. Probability Calculation. (3 rd Sem)

3. PROFESSORS

Meetings after coordination with the professor

4. INTRODUCTION TO THE COURSE

Artificial Intelligence (AI) is a fundamental field of computer science that seeks to create systems capable of performing tasks that typically require human intelligence. This course provides an introduction to the techniques used to build such systems, including search algorithms, probabilistic reasoning, and machine learning. Students will learn to model complex problems and implement intelligent agents, while considering the social and ethical implications of these technologies.

5. GOALS

- Understand the fundamental concepts and history of Artificial Intelligence.
- Implement and evaluate search algorithms for problem-solving.
- Master the basics of machine learning, including supervised and unsupervised learning.
- Apply probabilistic reasoning to handle uncertainty in intelligent systems.
- Analyze the ethical challenges and social impact of AI.

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. (Usage)

7. TOPICS

Unit 1: Introduction to AI: Concepts and Intelligent Behavior (4 hours)	
Competences Expected: AG-C08,AG-C09	
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 : [Russell2020], [Luger2008]	

Unit 2: Problem Characteristics and Agent Nature (4 hours)	
Competences Expected: AG-C08,AG-C09	
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 : [Russell2020], [Luger2008]	

Unit 3: State and Uninformed Search (5 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• State space representation of a problem: – Specifying states, goals, and operators – Factoring states into representations (hypothesis spaces) – Problem solving by graph search: * e.g., Graphs as a space, and tree traversals as exploration of that space * Dynamic construction of the graph (not given upfront) • Uninformed graph search for problem solving: – Breadth-first search – Depth-first search: * With iterative deepening – Uniform cost search • Space and time complexities of graph search algorithms • Bidirectional search • Understanding the search space: – Constructing search trees – Dynamic search spaces – Combinatorial explosion of search space – Search space topology (e.g., ridges, saddle points, local minima)	• Design the state space representation for a puzzle (e.g., N-queens or 3-jug problem) [Apply] • Select and implement an appropriate uninformed search algorithm for a problem (e.g., tic-tac-toe), and characterize its time and space complexities [Apply]
Readings : [Russell2020], [Luger2008]	

Unit 4: Heuristic Search and CSP (6 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• Heuristic graph search for problem solving: – Heuristic construction and admissibility – Hill-climbing – Local minima and the search landscape: * Local vs global solutions – Greedy best-first search – A* search • Beam search • Implementation of A* search • Constraint satisfaction • Variations on A* (IDA*, SMA*, RBFS)	• Select and implement an appropriate informed search algorithm for a problem after designing a helpful heuristic function (e.g., a robot navigating a 2D gridworld) [Apply] • Evaluate whether a heuristic for a given problem is admissible/can guarantee an optimal solution [Evaluate] • Design and implement A*/beam search to solve a problem, and compare it against other search algorithms in terms of the solution cost, number of nodes expanded, etc [Apply] • Compare and contrast various heuristic searches vis-a-vis applicability to a given problem [Evaluate] • Model a logic or Sudoku puzzle as a constraint satisfaction problem, solve it with backtrack search, and determine how much arc consistency can reduce the search space [Apply]
Readings : [Russell2020], [Luger2008]	

Unit 5: Adversarial Search and Stochastic Optimization (5 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• Two-player adversarial games: – Minimax search – Alpha-beta pruning: * Ply cutoff • Local search • Tabu search • Two-player adversarial games: – The horizon effect – Opening playbooks/endgame solutions – What it means to "solve" a game (e.g., checkers) • Implementation of minimax search, beam search • Expectimax search (MDP-solving) and chance nodes • Stochastic search: – Simulated annealing – Genetic algorithms – Monte-Carlo tree search	• Apply minimax search in a two-player adversarial game (e.g., connect four), using heuristic evaluation at a particular depth to compute the scores to back up [Apply] • Apply minimax search with alpha-beta pruning to prune search space in a two-player adversarial game (e.g., connect four) [Apply] • Compare and contrast genetic algorithms with classic search techniques, explaining when it is most appropriate to use a genetic algorithm to learn a model versus other forms of optimization (e.g., gradient descent) [Evaluate] • Design and implement a genetic algorithm solution to a problem [Apply] • Design and implement a simulated annealing schedule to avoid local minima in a problem [Apply]
Readings : [Russell2020], [Luger2008]	

Unit 6: Machine Learning Fundamentals (7 hours)	
Competences Expected: AG-C08,AG-C09	
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 : [Russell2020], [Alpaydin2020]	

Unit 7: Machine Learning Data, Representations and Evaluation (7 hours)	
Competences Expected: AG-C08,AG-C09	
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 : [Russell2020], [Alpaydin2020]	

Unit 8: Neural Networks and Ethics (6 hours)	
Competences Expected: AG-C08,AG-C09	
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 : [Russell2020], [Alpaydin2020]	

Unit 9: 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 10: 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 11: Knowledge Representation and Probabilistic Reasoning Foundations (6 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
- Types of representations: - Symbolic, logical: * Creating a representation from a natural language problem statement - Learned subsymbolic representations - Graphical models (e.g., naive Bayes, Bayesian network) - Review of probabilistic reasoning, Bayes theorem - Bayesian reasoning: - Bayesian inference - Random variables and probability distributions: - Axioms of probability - Probabilistic inference - Bayes' Rule (derivation) - Bayesian inference (more complex examples)	- Given a natural language problem statement, encode it as a symbolic or logical representation [Apply] - Compute a probabilistic inference in a real-world problem using Bayes' theorem to determine the probability of a hypothesis given evidence [Apply] - Apply Bayes' rule to determine the probability of a hypothesis given evidence [Apply] - Compute the probability of outcomes and test whether outcomes are independent [Apply]
Readings : [Russell2020]	

Unit 12: Independence, Markov Models and Decision Making (6 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
- Independence - Conditional Independence - Markov chains and Markov models - Utility and decision making	Explain how we can make decisions under uncertainty, using concepts such as Bayes theorem and utility [Explain]
Readings : [Russell2020]	

Unit 13: AI Applications and Societal Impact (4 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
• At least one application of AI to a specific problem and field, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose at least one for the CS Core): – Formulating and evaluating a specific application as an AI problem: * How to deal with underspecified or ill-posed problems – Data availability/scarcity and cleanliness: * Basic data cleaning and preprocessing * Data set bias – Algorithmic bias – Evaluation bias – Assessment of societal implications of the application • Deployed deep generative models: – High-level overview of deep image generative models (e.g., as of 2023, DALL-E, Midjourney, Stable Diffusion, etc.), their uses, and their shortcomings/pitfalls. – High-level overview of large language models (e.g., as of 2023, ChatGPT, Bard, etc.), their uses, and their shortcomings/pitfalls. • Overview of societal impact of AI: – Ethics – Fairness – Trust/explainability – Privacy and usage of training data – Human autonomy and oversight/regulations/legal requirements – Sustainability	• Given a real-world application domain and problem, formulate an AI solution to it, identifying proper data/input, preprocessing, representations, AI techniques, and evaluation metrics/methodology [Apply]
Readings : [Russell2020], [Alpaydin2020]	

Unit 14: Additional AI Applications and Societal Impact (4 hours)	
Competences Expected: AG-C08,AG-C09	
Topics	Learning Outcomes
- One or more additional applications of AI to a broad set of problems and diverse fields, such as medicine, health, sustainability, social media, economics, education, robotics, etc. (choose a different area from that chosen for the CS Core): - Formulating and evaluating a specific application as an AI problem: - How to deal with underspecified or ill-posed problems - Data availability/scarcity and cleanliness: - Basic data cleaning and preprocessing - Data set bias - Algorithmic bias - Evaluation bias - Assessment of societal implications of the application - Additional depth on deployed deep generative models: - Introduction to how deep image generative models work, (e.g., as of 2023, DALL-E, Midjourney, Stable Diffusion) including discussion of attention - Introduction to how large language models work, (e.g., as of 2023, ChatGPT, Bard) including discussion of attention - Idea of foundational models, how to use them, and the benefits/issues with training them from big data - Analysis and discussion of the societal impact of AI: - Ethics - Fairness - Trust/explainability - Privacy and usage of training data - Human autonomy and oversight/regulations/legal requirements - Sustainability	- Analyze the societal impact of one or more specific real-world AI applications, identifying issues regarding ethics, fairness, bias, trust, and explainability [Evaluate] - Describe some of the failure modes of current deep generative models for language or images, and how this could affect their use in an application [Describe]
Readings : [Russell2020], [Alpaydin2020]	

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