NATIONAL UNIVERSITY OF SAN MARCOS

FACULTY OF MATHEMATICAL SCIENCES

SYLLABUS

1. GENERAL INFORMATION

1.1 Course name	:	Topics in Artificial Intelligence
1.2 Code	:	AI367
1.3 Academic Program	:	Scientific Computing
1.4 Credits	:	5
1.5 Prerrequisites	:	AI264. Deep Learning. (9^{th} Sem)
1.6 Type of Course	:	Mandatory
1.7 Area of Studies	:	Estudios Generales
1.8 Number of Weeks	:	16
1.9 Hours per week	:	3 HT; 4 HP;
1.10 Cycle	:	I
1.11 Learning modality	:	Face to face
1.12 Academic Period	:	2026-I
1.13 Semester Start Date	:	
1.14 Semester End Date	:	

2. PROFESSORS

Meetings after coordination with the professor

3. INTRODUCTION TO THE COURSE

Artificial Intelligence is a rapidly evolving field. This course explores advanced and contemporary topics that are currently shaping the industry and research. Focus is placed on Generative AI, Large Language Models (LLMs), and Advanced Reinforcement Learning. Students will analyze the architectures behind modern AI systems, such as Transformers, and discuss the technical and societal challenges of deploying autonomous intelligent agents at scale.

4. GOALS

- Understand the architecture and training process of Large Language Models.
- Master the principles of Generative AI for text and image synthesis.
- Implement advanced Reinforcement Learning algorithms for decision making.
- Apply prompt engineering and fine-tuning techniques to existing models.
- Critically evaluate the latest research papers and trends in AI.

5. COMPETENCES

ABET-1) An ability to identify, formulate, and solve broadly defined technical or scientific problems by applying knowledge of mathematics and science and/or technical topics to areas relevant to the discipline. (Usage)

6. TOPICS

Unit 1: Machine Learning Fundamentals (6 hours)

Competences Expected: ABET-1

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**], [**Foster2022**]

Unit 2: Machine Learning Data, Representations and Evaluation (5 hours)

Competences Expected: ABET-1

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**], [**Foster2022**]

Unit 3: Neural Networks and Ethics (5 hours)	
Competences Expected: ABET-1	
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**], [**Foster2022**]	

Unit 4: Machine Learning Formulation and Evaluation (3 hours)	
Competences Expected:	
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 5: Deep Learning, RL and Ethics (3 hours)	
Competences Expected:	
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 6: Natural Language Processing (20 hours)

Competences Expected: ABET-1

Topics	**Learning Outcomes**

• Deterministic and stochastic grammars
• Parsing algorithms:
  – CFGs and chart parsers (e.g., CYK)
  – Probabilistic CFGs and weighted CYK
• Representing meaning/Semantics:
  – Logic-based knowledge representations
  – Semantic roles
  – Temporal representations
  – Beliefs, desires, and intentions
• Corpus-based methods
• N-grams and HMMs
• Smoothing and backoff
• Examples of use: POS tagging and morphology
• Information retrieval:
  – Vector space model:
  ∗ TF & IDF
  – Precision and recall
• Information extraction
• Language translation
• Text classification, categorization:
  – Bag of words model
• Deep learning for NLP:
  – RNNs
  – Transformers
  – Multi-modal embeddings (e.g., images + text)
  – Generative language models
	• Define and contrast deterministic and stochastic grammars, providing examples to show the adequacy of each [Define]
	• Simulate, apply, or implement classic and stochastic algorithms for parsing natural language [Apply]
	• Identify the challenges of representing meaning [Analyze]
	• List the advantages of using standard corpora. Identify examples of current corpora for a variety of NLP tasks [List]
	• Identify techniques for information retrieval, language translation, and text classification [Analyze]
	• Implement a TF/IDF transform, use it to extract features from a corpus, and train an off-the-shelf machine learning algorithm using those features to do text classification [Apply]

Readings : [**Tunstall2022**], [**Russell2020**]

Unit 7: Machine Learning Fundamentals (6 hours)

Competences Expected: ABET-1

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 : [**Sutton2018**]

Unit 8: Machine Learning Data, Representations and Evaluation (5 hours)	
Competences Expected: ABET-1	
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 : [**Sutton2018**]	

Unit 9: Neural Networks and Ethics (5 hours)	
Competences Expected: ABET-1	
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 : [**Sutton2018**]	

Unit 10: AI Applications and Societal Impact (6 hours)

Competences Expected: ABET-1

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**], [**Bostrom2014**]

Unit 11: Additional AI Applications and Societal Impact (6 hours)	
Competences Expected: ABET-1	
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**], [**Bostrom2014**]	

7. WORKPLAN

7.1 Methodology

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

7.2 Theory Sessions

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

7.3 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.

8. EVALUATION SYSTEM

********* EVALUATION MISSING ********

9. BASIC BIBLIOGRAPHY