SYLLABUS
Academic year 2025-2026

Dean,
Prof. dr. eng. Vasile-Ion Manta

1. Program data

1.1 Higher education institution	"Gheorghe Asachi" Technical University of Iași
1.2 Faculty	Automatic Control and Computer Engineering
1.3 Department	Computers
1.4 Field of studies	Computers and Information Technology
1.5 The cycle of studies[1]	Master
1.6 Study program	Artificial Intelligence

2. Subject data

2.1 Name of the subject / Code				**Natural Language Processing** (*Procesarea Limbajului Natural*) / **AI.105**			
2.2 Course coordinator				Lect. dr. eng. Stefan-Daniel Achirei / Dr. eng. Matei Stefan Neagu			
2.3 Application instructor				Lect. dr. eng. Stefan-Daniel Achirei / Dr. eng. Matei Stefan Neagu / Assist. drd. eng. Codruț-Georgian Artene			
2.4 Year of study[2]	1	2.5 Semester[3]	2	2.6 Type of assessment[4]	exam	2.7 Type of subject[5]	DS

3. Estimated total time of daily activities (hours per semester)

3.1 Number of hours per week	4	3.2 lectures	2	3.3a sem.		3.3b laboratory	2	3.3c project	
3.4 Total hours in curriculum[6]	56	3.5 lectures	28	3.6a sem.		3.6b laboratory	28	3.6c project	

Distribution of the time fund[7]	No. hours
Study by textbook, course support, bibliography and notes	35
Additional documentation in the library, on specialist electronic platforms and in the field	25
Preparation of seminars/labs/projects, assignments, reports and portfolios	20
Tutorial[8]	10
Examinations[9]	4
Other activities:	

3.7 Total hours of individual study[10]	94
3.8 Total hours per semester[11]	150
3.9 Number of credits	6

4. Prerequisites (where applicable)

4.1 curriculum[12]	
4.2 competences	

5. Conditions (where applicable)

5.1 conducting the lectures[13]	• Blackboard, video projector
5.2 conducting the seminar / laboratory / project[14]	• Laboratory room with computers and Internet access • The Visual Studio (academic license) and PyCharm programming environments; jupyter notebooks

[1] *Bachelor / Master*
[2] *1-4 for Bachelor's, 1-2 for Master's*
[3] *1-8 for Bachelors, 1-3 for Masters*
[4] *Exam, colloquium or VP A/R – from the curriculum*
[5] *DF - fundamental subject, DID - subject in the field, DS - specialized subject or DC - complementary subject - from the education plan*
[6] *It is equal to 14 weeksx number of hours from point 3.1 (similar for 3.5, 3.6abc)*
[7] *The lines below refer to the individual study; the total is completed at point 3.7.*
[8] *Between 7 and 14 hours*
[9] *Between 2 and 6 hours*
[10] *The sum of the values on the previous lines, which refer to the individual study.*
[11] *The sum of the number of hours of direct teaching activity (3.4) and the number of hours of individual study (3.7); must be equal to the number of credits allocated to the subject (point 3.9)x 25 hours per credit.*
[12] *Mention the subjects that must be passed previously or equivalent*
[13] *Blackboard, video projector, flipchart, specific teaching materials, etc.*
[14] *Computing technique, software packages, experimental stands, etc.*

6. Specific competences accumulated[15]

Number of credits assigned to the subject[16]:		**6**	Distribution of credits per competences[17]
Professional competences	CP1	Knowledge of advanced concepts of computer science and information technology and the ability to work with these concepts.	1.0
	CP2	Scientific and practical research in the field of artificial intelligence.	1.0
	CP3	Problem solving using artificial intelligence methods and techniques.	2.0
	CP4	Design and development of artificial intelligence systems.	1.0
	CP5	Utilization of artificial intelligence tools and technologies.	0.7
	CP6		
	CPS1		
	CPS2		
Transversal competences	CT1	Legislation compliant application of the intellectual property rights and of the principles, norms and values of the professional ethics code within their own strategies for rigorous, effective and responsible work.	0.2
	CT2	Application of communication techniques and effective group work; developing empathic interpersonal communication skills and assuming leadership roles/functions in a multi-specialized team.	
	CT3	Creating opportunities for continuous training and the effective utilization of learning resources and techniques for personal development.	0.1
	CTS		

7. Objectives of the subject (resulting from the grid of specific competences accumulated)

7.1 General objective of the subject	Understanding and becoming familiar with the theoretical principles and the development of practical skills in the field of natural language processing.
7.2 Specific objectives	The goal of this course is to present the general issue of natural language processing while learning theoretical concepts and acquiring practical skills. Review of current tools and technologies in the field of NLP. Stimulating research skills in this field. Exploring how NLP can be applied in different domains (text classification, similarity analysis, text generation, machine translation).

8. Contents

8.1 Course[18]	Teaching methods[19]	Remarks
Week 1: Introduction to Natural Language Processing (NLP) ● *Overview of NLP and its applications* ● *NLP, NLU, NLG* ● *Introduction to deep learning for NLP* ● *Basic NLP tasks and datasets* ***Week 2: Text Preprocessing and Text Representation*** ● *Text normalization and cleaning* ● *Tokenization and stemming/lemmatization* ● *N-grams and bag-of-words representation* ***Week 3: Word embeddings*** ● *Word embedding techniques: Word2Vec, GloVe, ELMo* ● *Contextualized Word Embeddings* ● *Multi-lingual Aligned Word Vectors* ***Week 4: Attention Mechanism and Transformer Networks*** ● *The attention mechanism and its significance in NLP* ● *Transformer architecture and its applications* ● *Self-attention and multi-head attention* ● *Transformer-based models: BERT, GPT-3, RoBERTa*	Lectures with Powerpoint presentations, explanations and answers to questions	

[15] *Competencies from the G1 and G1bis Grids of the study program, adapted to the specifics of the subject, for which credits are allocated (www.rncis.ro or the faculty website)*

[16] *From the education plan*

[17] *The credits allocated to the subject are distributed on professional and transversal competences according to the specifics of the subject*

[18] *Chapter and paragraph headings*

[19] *Exposition, lecture, blackboard presentation of the studied issue, use of video projector, discussions with students (for each chapter, if applicable)*

Week 5: Text Classification ● *Binary and multi-class text classification* ● *Logistic regression and support vector machines for text classification* ● *Neural network models for text classification: CNNs, RNNs, Transformers* ● *Evaluation metrics for text classification: accuracy, precision, recall, F1 score* ***Week 6: Text Similarity*** ● *Measuring text similarity: cosine similarity, Jaccard similarity, word embedding similarity* ● *Applications of text similarity: plagiarism detection, sentiment analysis, product recommendation* ● *Deep learning models for text similarity: siamese networks, sentence-BERT* ***Week 7: Text Generation*** ● *Statistical language modeling: n-grams, recurrent neural networks* ● *Generative adversarial networks (GANs) for text generation* ● *Transformer-based models for text generation: GPT-2, GPT-3* ● *Decoding strategies* ● *Creative text generation: poetry, scripts, musical pieces* ***Week 8: Question Answering & Machine Translation (MT)*** ● *Extractive question answering (QA)* ● *End-to-end question answering: BERT, RoBERTa, XLNet* ● *Question answering evaluation metrics: SQuAD, MRQA* ● *Statistical machine translation (SMT)* ● *Neural machine translation (NMT)* ● *Transformer-based models for NMT: Transformer, ConvS2S* ● *MT evaluation metrics: BLEU, ROUGE* ***Week 9: Part-of-Speech (POS) Tagging and Named Entity Recognition (NER)*** ● *POS tagging: Hidden Markov Models (HMMs), conditional random fields (CRFs)* ● *NER: Bidirectional LSTMs, CRF-based models* ● *Applications of POS tagging and NER: information extraction, natural language understanding* ***Week 10: Evaluating NLP Models*** ● *Common evaluation metrics for NLP tasks* ● *Bias and fairness in NLP models* ● *Human-centered evaluation of NLP systems* ● *Explainability and interpretability of NLP models* ● *Ethics in Natural Language Processing* ***Week 11: Advanced Topics in NLP*** ● *Dialogue systems and chatbots* ● *Natural language generation (NLG)* ● *Sentiment analysis and opinion mining* ● *Summarization and abstractive text generation* ● *Multimodal NLP: combining text, audio, and video data* ***Week 12: NLP Applications*** ● *NLP in cybersecurity* ● *NLP in healthcare and medicine* ● *NLP in finance and economics* ● *NLP in social media and marketing* ● *NLP in education and learning* ● *NLP in law and government* ***Week 13: Course Review and Project Presentations*** ● *Review of key concepts and techniques covered in the course* ● *Student presentations on their final projects* ● *Discussion of future directions in NLP research* ***Week 14: Text-to-Speech (TTS) and Speech-to-Text (STT)*** ● *Text-to-speech synthesis: waveform generation, vocoders* ● *Speech recognition: hidden Markov models (HMMs), deep neural networks (DNNs)* ● *Applications of TTS and STT: assistive technologies, smart speakers, voice assistants* ***TOTAL: 28h***		
Course references:		

1. Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit 1st Edition, Steven Bird, Ewan Klein, Edward Loper, O'Reilly Media, 2009 (https://tjzhifei.github.io/resources/NLTK.pdf)
2. Natural Language Processing, Jacob Eisenstein, 2018 (https://github.com/jacobeisenstein/gt-nlp-class/blob/master/notes/eisenstein-nlp-notes.pdf)
3. Thumbs up? Sentiment Classification using Machine Learning Techniques, Bo Pang et al., 2002
4. Baselines and Bigrams: Simple, Good Sentiment and Topic Classification, Sida Wang and Christopher Manning, 2012
5. Convolutional Neural Networks for Sentence Classification, Yoon Kim, 2014
6. [GitHub] NLP Progress on Sentiment Analysis (https://github.com/sebastianruder/NLP-progress/blob/master/english/sentiment_analysis.md)
7. A large annotated corpus for learning natural language inference, Sam Bowman et al., 2015
8. Authorship Attribution of Micro-Messages, Roy Schwartz et al., 2013
9. 50 Years of Test (Un)fairness: Lessons for Machine Learning, Ben Hutchinson and Margaret Mitchell, 2018
10. [Article] Amazon scraps secret AI recruiting tool that showed bias against women
11. [Blog] Neural Networks, Manifolds, and Topology, Chris Olah
12. Dropout: a simple way to prevent neural networks from overfitting, Nitish Srivastava et al., 2014
13. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, Sergey Ioffe and Christian Szegedy, 2015
14. Adam: A Method for Stochastic Optimization, Durk Kingma and Jimmy Ba, 2015
15. The Marginal Value of Adaptive Gradient Methods in Machine Learning, Ashia Wilson et al., 2017

8.2a Seminar	Teaching methods[20]	Remarks

8.2b Laboratory	Teaching methods[21]	Remarks
Weeks 1-3: Linear Sentiment Classification (6h) *Goals: The main goal of this project is for the students to get experience extracting features and training classifiers on text. They will get a sense of what the standard machine learning workflow looks like (reading in data, training, and testing), how standard learning algorithms work, and how the feature design process goes.* ● *Framework code setup; understanding the dataset; implementing a perceptron classifier with a bag-of-words unigram featurization (2h)* ● *Implementing a logistic regression classifier with unigram bag-of-words feature (2h)* ● *Implementing two additional feature extractors (BigramFeatureExtractor and BetterFeatureExtractor) (2h)* ***Weeks 4-6: Feedforward Neural Networks, Word Embeddings, and Generalization (6h)*** *Goals: The main goal of this project is for the students to get experience training neural networks over text. They will play around with feedforward neural networks in PyTorch and see the impact of different sets of word vectors on the sentiment classification problem from Project 1.* ● *Familiarize with function optimization; implement the gradient (2h)* ● *Implement a Deep Averaging Network (4h)* ***Weeks 7-9: Transformer Language Modeling (8h)*** *Goals: The primary goal of this assignment is to give students hands-on experience implementing a Transformer language model. Understanding how these neural models work and building one from scratch will help the students understand not just language modeling but also systems for many other applications, such as machine translation.* ● *Building a "Transformer" Encoder (4h)* ● *Transformer for Language Modeling (4h)* ***Weeks 10-14: Factuality and ChatGPT (8h)*** *Goals: The primary goal of this project is to give the students hands-on experience analyzing outputs from large language models (LLMs) like ChatGPT. The students will understand what sorts of non-factual outputs LLMs can generate and what is involved in the process of verifying those outputs against Wikipedia. In addition, they will conduct some error analysis of the mistakes the fact-checking model makes.* ● *Bag-of-words Overlap (4h)* ● *Textual Entailment (2h)* ● *Error Analysis (2h)*	General and individual explanations, individual computer work	

[20] *Discussions, debates, presentation and/or analysis of papers, solving exercises and problems*

[21] *Practical demonstration, exercise, experiment*

TOTAL: 28h		

Applications (laboratory) references:

1. Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP).
2. Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daume III. 2015. Deep Unordered Composition Rivals Syntactic Methods for Text Classification. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL).
3. Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP).
4. Tomas Mikolov, Ilya Sutskever, Anoop Deoras, Hai-Son Le, Stefan Kombrink, and Jan Cernocky. 2012. Subword Language Modeling with Neural Networks. In Online preprint.
5. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In arXiv.
6. Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. DeBERTa: Decoding-enhanced BERT with Disentangled Attention. ArXiv, abs/2006.03654.
7. Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. SummaC: Re-visiting NLI-based models for inconsistency detection in summarization. Transactions of the Association for Computational Linguistics, 10:163–177.
8. Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. arXiv 2305.14251.
9. Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. Adversarial NLI: A new benchmark for natural language understanding. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 4885–4901, Online, July. Association for Computational Linguistics.
10. James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 809–819, New Orleans, Louisiana, June. Association for Computational Linguistics.
11. Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers).

9. Corroboration of the contents of the subject with the expectations of representatives of the epistemic community, professional associations and representative employers in the field related to the program[22]

An "Natural Language Processing" course directly addresses the IT industry's demand for NLP systems, as the industry seeks Large Language Models solutions due to the large addresabilit..

10. Evaluation

Type of activity	10.1 Evaluation criteria	10.2 Evaluation methods		10.3 Weight in the final grade
10.4a Exam	Acquired theoretical and practical knowledge (quantity, correctness, accuracy)	Periodic tests[23]:		50% (minimum 5)
		Homework:		
		Other activities[24]:		
		Final evaluation:	100% Test from the topics presented in the course	
10.4b Seminar	Frequency/relevance of interventions or responses			
10.4c Laboratory	Knowledge of equipment, how to use specific tools; evaluation of tools or achievements, processing and interpretation of results	• Practical demonstrations • Oral answers • Assessment of progress during laboratory hours, through discussions with students, questions, and checks of practical results obtained		50% (minimum 5)
10.4d Project	The quality of the completed project, the correctness of the project documentation, the justification of the chosen solutions			
10.5 Minimum performance standard[25]: grade 5 in the exam and laboratory work				

[22] *The connection with other subjects, the usefulness of the subject on the labor market*

[23] *The number of tests and the weeks in which they will be held will be specified.*

[24] *Scientific circles, professional competitions, etc.*

Date of completion, 4 December 2023	Signature of course coordinator, Lect. dr. eng. Stefan-Daniel Achirei	Signature of application instructor, Lect. dr. eng. Stefan-Daniel Achirei
	Signature of course coordinator, Dr. eng. Matei Stefan Neagu	Signature of application instructor, Dr. eng. Matei Stefan Neagu
		Signature of application instructor, Assist. dr. eng. Codruț-Georgian Artene
Date of approval in the department, 7 December 2023		Director of department, Assoc. prof. dr. eng. Andrei Stan

[25]*The minimum performance standard from the competences grid of the study program is customized to the specifics of the subject, if applicable.*