Course Title	***Artificial Intelligence Fundamentals***						
Course Code	***MAI 611***						
Course Type	Compulsory						
Level	Graduate						
Year / Semester	Fall						
Teacher's Name	Elpida Keravnou-Papailiou						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	0 hours
Course Purpose and Objectives	The purpose of the course is to introduce students to the fundamental principles and techniques that underline software systems that exhibit "intelligent" behavior.						
Learning Outcomes	Upon completion of this course, students will have acquired a good understanding of symbolic Artificial Intelligence, the problems that it addresses and the fundamental solution methods that it uses. More specifically the students will know the main knowledge representation techniques and reasoning methods that underlie artificial intelligence problem solving and be able to develop simple solvers for artificial intelligence systems.						
Prerequisites	None			Required		Knowledge of a high-level programming language, object-based data concepts and structures.	

Course Content	The course introduces the fundamental principles and methods used in Artificial Intelligence to solve problems, with a special focus on the search in the state space, action planning, knowledge representation and reasoning, constraint satisfaction, intelligent agents and on the methods for dealing with uncertain knowledge. The course content includes the following thematic units: • Introduction to Artificial Intelligence: Turing test for machine intelligence, historical perspective, symbol processing, algorithms and heuristics, main application fields, introduction to knowledge-based systems and architectural organization. • Problem Solving in AI: Representation problem, navigation mechanisms (blind search, heuristic guidance), classification and synthesis problems, frame problem. Games. Constraint satisfaction problems. Search methods, meta heuristics, solving through decomposition, constraint relaxation, branch-and-bound techniques. Linear planning. • Intelligent Agents and Multiagent Systems: Autonomous/rational agents, abstract and concrete architectures, communication, interaction, and communication between agents. • Knowledge Representation and reasoning: Distinction between data, information and knowledge, expertise (types of knowledge), desirable practical and theoretical properties, descriptive/declarative versus procedural representation. Predicate logic. Conjunctive Normal Form and resolution-refutation. Horn clauses and negation as failure. Representing Terminological Knowledge: semantic networks, frames and inheritance. Uncertainty and probabilistic reasoning. • Rule-based systems: Production rules, control structure and rule chaining, forward chaining and RETE, backward chaining, goal-driven reasoning and explanations, meta-rules, strategic explanations. Limitations of rule-based systems. • Expert Systems Technology: Basic forms of reasoning (deduction, abduction, induction). First and second generation of expert systems (deep knowledge-based systems, multi-modelling, strategic explanations, collaborating specialists, generic tasks architecture). • Knowledge engineering methodologies: Developing and modelling expertise, total task investigation methods, basic principles, and multiple models of CommonKADS methodology. • Case-based reasoning as an alternative paradigm to model-based reasoning supporting incremental learning. • Intelligent Data Analysis: Data mining and data abstraction. • AI and Explanations: The significance of explanations in AI, some theories of explanation that have influenced AI, tracing the history of explanations in symbolic AI (rule-based, abductive, causal, case-based, strategic, explanations), the resurgence of interest in explanation in connectionist AI (XAI) – opening the 'black box'.
Teaching Methodology	Lectures, discussions of practical examples and (unsupervised) lab activities where the active learning element is encouraged and supported. Students would be strongly guided to view all topics presented and discussed with a critical eye, identifying the limits of AI both in its foundational years and the current state-of-affairs characterized by an explosion of multimedia data of varying degrees of usability, quality, and ethical considerations.

Bibliography	Main text: 1. S. Russel and P. Norvig, Artificial Intelligence: A Modern Approach, 4th Edition, Pearson, 2021. Other reading: 1. R. J. Brachman, H. J. Levesque, Knowledge Representation and Reasoning, Elsevier, 2004. 2. N. J. Nilsson: The Quest for Artificial Intelligence: A history of ideas and achievements, Cambridge University Press, 2010. 3. M. Ginsberg: Essentials of Artificial Intelligence, Morgan Kaufman,1993. 4. P. H. Winston: Artificial Intelligence, 3rd Edition, Addison-Wesley, 1992. 5. E. Keravnou, Artificial Intelligence and Expert Systems (in Greek), Greek Open University, 2000. 6. G.F. Luger and W.A. Stubblefield, Artificial Intelligence: Structures and Strategies for Complex Problem Solving, 5th edition, Addison-Wesley, 2005. 7. G. Weiss (editor), Multiagent Systems: a modern approach to distributed AI, The MIT Press, 2nd edition, 2013. 8. P. Jackson, Introduction to Expert Systems, 3rd edition, Addison-Wesley, 1999. 9. J. Giarratano και G. Riley, Expert Systems: Principles and Programming, 4th edition, International Thomson Publishing, 2004. 10. J-M. David, J-P. Krivine and R. Simmons (editors), Second Generation Expert Systems, Springer-Verlag, 2011. 11. J. Breuker και W. Van de Velde (editors), CommonKADS Library for Expertise Modelling: Reusable Problem Solving Components, IOS Press, 1994.
Assessment	Final exam (50%), midterm exam (30%) and homework (theoretical and/or programming assignments) (20%).
Language	English

Course Title	***Research Methodologies and Professional Practices in AI***						
Course Code	***MAI 613***						
Course Type	Compulsory						
Level	Graduate						
Year / Semester	Fall						
Teacher's Name	Stelios Timotheou						
ECTS	6	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	0 hours
Course Purpose and Objectives	The purpose of the mandatory course on research methodologies and professional practices in AI is to introduce students to the methods and tools of Artificial Intelligence Research, professional practices, and associated technological culture, bearing in mind European Commission's regulatory framework. Furthermore, the course will introduce the mathematical foundation and associated software tools for understanding AI techniques. Moreover, the course objectives encompass familiarization with reading, reviewing and presenting of relevant literature, technical writing and literature surveying.						
Learning Outcomes	Upon completion of the course the students will be sufficiently conversant with the key methodological steps involved in carrying out research in Artificial Intelligence and the safeguards for mitigating risks in potentially high-risk AI research and applications. Consequently, they will be familiar with the obligatory requirements for professional practices in AI to be characterized as secure, trustworthy and ethical. Furthermore, the students will gain good knowledge of the mathematical foundations of AI and of the utilization of software tools for the practical solutions of AI problems. In addition, they will acquire experience in surveying some topic, writing a technical report on it and presenting it.						
Prerequisites	None	Required	Experience in orally presenting some topic to an audience and in writing a structured report (objectives, analysis, conclusions).				
Course Content	The course comprises the following component units: • AI Research Methodologies, including the coverage of modern Library and other technological tools for aiding research, by professionals from UCY's Library. • Research or technical literature reviewing in AI. Presentation of technical study in an AI related topic. • European Commission's regulatory framework for the development of secure, trustworthy and ethical AI. • Mathematical foundation of AI techniques including Linear Algebra, Analytic Geometry, Matrix Decompositions, Vector Calculus, Probability and Statistics, and Optimization.						
Teaching Methodology	Lectures, in-class discussion on good and bad practices for undertaking literature reviews, in-class programming assignments to learn the usage of software tools on mathematical foundations of AI, and a group study of an AI-related research subject.						
Bibliography	1. Selected research articles from international literature. 2. Course Presentation Slides. 3. Regulation of the European Parliament and of the European Council in Laying down harmonized rules for Artificial Intelligence (COM (2021) 206 final). 4. P.R. Cohen, Empirical Methodology for Artificial Intelligence, MIT Press, 1995. 5. Deisenroth, M. P., Faisal A. A. and Ong C. S. Mathematics for machine learning. Cambridge University Press, 2020.						
Assessment	• Group study of an AI-related research subject and technical presentation of the group study (50%) • Final exam on the mathematical foundation of AI (50%)						
Language	English						

<table>
<tr><td>Course Title</td><td colspan="7">AI Ethics</td></tr>
<tr><td>Course Code</td><td colspan="7">MAI 621</td></tr>
<tr><td>Course Type</td><td colspan="7">Compulsory</td></tr>
<tr><td>Level</td><td colspan="7">Graduate</td></tr>
<tr><td>Year / Semester</td><td colspan="7">Spring</td></tr>
<tr><td>Teacher’s Name</td><td colspan="7">Michalis Agathocleous</td></tr>
<tr><td>ECTS</td><td>6</td><td>Lectures / week</td><td>3 hours</td><td>Recitation / week</td><td>0 hous</td><td>Laboratories / week</td><td>0 hours</td></tr>
<tr><td>Course Purpose and Objectives</td><td colspan="7">This course aims to raise awareness of the dangers that can arise from the development, deployment, and usage of intelligent autonomous systems and to introduce the students to socio-technical solutions for mitigating the risk of exhibiting unwanted non-ethical behaviour. Students will understand the basics of implementing systems that are not only high performing, but also adhere to our ethical socio-legal cultural values.</td></tr>
<tr><td>Learning Outcomes</td><td colspan="7">The key learning outcomes of the course are:
Reflect upon the socio-ethical issues that arise upon the development, deployment, and usage of intelligent systems.
Critically discuss commonly occurring narratives and perspectives related to the use of AI.
Reason about the decisions made during a system’s lifecycle and how the decision makers have a moral responsibility for their choices.
Learn how to develop systems that exhibit a desired ethical behaviour, understand the main research challenges for this and design AI solutions aligned with responsible innovation principles.
Understand how to judge and evaluate AI systems for their “ethicacy”.
Appreciate the socio-technical mechanisms for the governance of AI systems.
By completing the above outcomes, the student will have a fundamental understanding of how intelligent systems influence—and are influenced by—our societies and of the socio-ethical responsibilities they have as developers and users of such tools.</td></tr>
<tr><td>Prerequisites</td><td colspan="3">AI Fundamentals</td><td>Required</td><td colspan="3">Basic technical knowledge of AI</td></tr>
</table>

Course Content	The course introduces the fundamental aspects of AI ethics by providing a holistic multidisciplinary view of the discipline. The course structure is such as to first help students understand the potential impact AI systems can have on societies and individuals (LO1, LO3). Then, they proceed by being introduced to ongoing discussions related to AI (LO2) and gaining the basic background in moral philosophy needed. Next, students are asked to put into knowledge gained into practice by learning how to develop systems that exhibit a desired ethical behaviour (LO4). The course finishes by examining the means of verifying and evaluating the ethical compliance or "ethicacy" of AI systems (LO5). The course is divided into following thematic units: • Introduction to AI Ethics: Establishes the motivation behind the field of AI ethics by using real-world use cases related to algorithmic biases, generation of disinformation, and attempts to escape accountability. It disproves the myth of 'responsibility damages innovation.' • Machine Ethics from a Philosophy Perspective: Provides an introduction into philosophical questions related to AI ethics: moral agency and patiency, anthropomorphising intelligence systems, mimicry human behaviour, • Ethics By Design: Delivers theoretical and practical understandings of how to do ethics in design; i.e. how can we develop autonomous AI systems that follow well-established ethical frameworks. • Ethics In Design: Provides a complete overview of procedures and methods that ensure compliance of systems to core values and governance mechanisms introduce in the previous module. Technical mechanisms for ensuring the ethical compliance of AI systems are introduced: software engineering practices, technical solutions for AI assessment, explainable and fairness AI toolkits. • Ethics For Design: Introduces fundamental concepts related to policymaking and governance to the students. Provides an up-to-date overview of governance initiatives related to AI: from the European Commission's AI Regulatory framework and GDPR to the UNESCO guidelines to the IEEE and ISO standardisation work. The unit provides the students working definitions and basic understanding of core AI ethics concepts accountability, responsibility, transparency, explainability, privacy, and fairness.
Teaching Methodology	A variety of teacher-led and student-led activities. Weekly lectures will introduce and provide overview of topics. Students will conduct self-study of the weekly material. Students will be given the opportunity to participate in problem-based solving group exercises, where they will conduct critical analysis and debate timely issues related to AI ethics. In unsupervised technical labs, students will be given the opportunity to test technical solutions for compliance checking and implement machine ethics, i.e. agents with moral reasoning.
Bibliography	Main Text: 1. V. Dignum. Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way, Springer, 2019. Other reading: 1. M. Coeckelbergh. AI Ethics, MIT-Press, 2020. 2. D. Gunkel, An Introduction to Communication and Artificial Intelligence, Willie, 2020. 3. C. O'Neil, Weapons of Math Destruction, Crown Books, 2016. 4. F. Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information, Belknap Press, 2014. 5. Papers, as reading material, will be made available to students on a weekly basis.
Assessment	2 major assignments (one group, one individual) and final exam.
Language	English

Course Title	***AI Entrepreneurship***						
Course Code	***MAI 622***						
Course Type	Compulsory						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	Marios Dikaiakos						
ECTS	8	Lectures / week	3 hours	Recitation / week	0 hours	Laboratories / week	1 hour
Course Purpose and Objectives	Progress in modern economies is shaped by scientific inventions and technological innovations that provide improved products or solutions to traditional needs, disrupt traditional business models or satisfy unanticipated consumer needs opening up new markets. Entrepreneurship is the primary mechanism through which inventions and innovations are brought to global markets. In recent years, the success of global Internet platforms, the rapid advance of digitalization across numerous sectors, and the ensuing abundance of digital data available on Internet have turned Artificial Intelligence (AI) and Machine Learning (ML) into major driving forces in innovative entrepreneurship, leading to unprecedented economic, social and cultural opportunities, challenges, and global impact. This course seeks to help students explore and master key concepts and challenges of relevance to AI and Data-driven entrepreneurship. The course introduces students to the world of AI entrepreneurship through case studies that demonstrate successes, failures and challenges. The course provides also an overview of and an introduction to key steps to develop a company, design a business model, explore product-market fit, manage intellectual property, and attract investment. Students will explore acknowledged innovation-driven entrepreneurship methodologies and experiment with them and associated tools to pursue the translation of their ideas into entrepreneurial endeavors. The course examines issues faced by Startup Founders and Chief Technology Officers who need to innovate at the boundaries of AI, Information Technology and Business by understanding all perspectives.						
Learning Outcomes	After taking this course, students should be able to: • Understand and explain the interplay between Big Data, Machine Learning and various application domains. • Evaluate technological ideas and apply the key stages of turning an idea or invention into a commercial product. • Apply the Business Model Canvas methodologies in Information Technology and Scientific application contexts. • Recognize and undertake the steps of the Disciplined Entrepreneurship methodology, and manage the key activities required to bring an innovative product or service to the market: product definition and market segmentation; value proposition analysis and high-level product specification; market and competition analysis; business model definition and revenue models; customer and user acquisition; minimum viable product definition and product implementation planning. • Understand the basics of fundraising and financing options for a startup. • Understand the basics of incorporation and company structure. • Understand the key challenges for attracting talent, establishing and managing a startup team. • Apply tools for project and team management, collaboration, ideation, rapid prototyping: Trello, Slack, SimpleMind, Proto.io, Github, Google AdService, Google Cloud, Heroku, etc. Prepare pitch decks, and pitch in front of potential investors, an AI-related business idea/product/service.						
Prerequisites	None		Required		None		

Course Content	The course will comprise weekly live and recorded lectures by the professor and by invited speakers on various aspects of entrepreneurship and innovation. The students will be required to establish teams and work on an idea, producing a business plan and a prototype of an MVP, and several writeups for class readings and invited lectures. Lectures will cover Case Studies in AI Entrepreneurship, Basic concepts in Entrepreneurship and Innovation, and elements of Preparatory Analysis for establishing a startup company, Setting up a company, Value Proposition, Market Analysis and Competition, Business Modeling for AI Products and Services, Customer acquisition and Sales. Module 1: Innovation, Entrepreneurship and AI • Invention, Entrepreneurship, Innovation, Research, Start-ups, Ecosystems, Risk, Venture Capital • Intellectual Property Elements • Steps involved to turn an Invention to a Start-up • Explore success stories and failures of AI entrepreneurship; discuss visions for the future • AI-first companies and the Lean AI methodology. Module 2: Business Modeling • Business Model Canvas • Drafting a Mission Statement Module 3: Disciplined Entrepreneurship • Introduction to the DE Methology and DE Canvas • Customer & Market Exploration • Market segmentation – DE Step 1 o Beachhead market selection – DE Step 2 o End-user Profile Definition – DE Step 3 o Total Addressable Market Size (TAM) of Beachhead – DE Step 4 o Profile Persona development for the Beachhead – DE Step 5 o Identify your Next 10 Customers – DE Step 9 • Product & Competition o Full Life Cycle Use Case – DE Step 6 o High-level Product Specification – DE Step 7 o Value Proposition: Definition and Quantification – DE Step 8 o Define your Core – DE Step 10 o Charting your Competitive Position – DE Step 11 • Business Modeling o Design a Business Model – DE Step 15 o Introduction to Platform Economy, Network effects, Platform-based services o Business Model Generation – Business Model Canvas o Set Your Pricing Framework – DE Step 16 o Calculate Lifetime Value of Acquired Customer – DE Step 16 o Cost of Customer Acquisition (COCA) Analysis – DE Step 18 • Product Design/Prototyping o Design and test key assumptions – DE Steps 20, 21 o Minimum Viable Business Product – DE Step 22 o Product demonstration and customer-satisfaction assessment – DE Step 23 o Lean Product Methodology Overview • Customer acquisition/Sales o Customer's Decision-Making Unit Definition – DE Step 12 o Map Process to Acquire Paying Customer – DE Step 13 o Map the Process to Acquire a Customer – DE Step 18 • Scaling Up o Calculate TAM Size for Follow-on Markets – DE Step 14 o Develop a Product Plan – DE Step 24 Module 4: Pitching and Fundraising • Introduction to Start-up Financing and Fundraising • Pitching

Teaching Methodology	• Course Lectures: 3 hours every week (13 x 3hrs) • Invited Seminars (Precept): 1 hour every week (13 x 1hr). This refers to participation to invited lectures with lecturers brought to campus by the Centre for Entrepreneurship, talking about experiences with entrepreneurship and start-ups. • Team Project (all semester).
Bibliography	1. Bill Aulet, Disciplined Entrepreneurship, Wiley, 2024. 2. Bill Aulet, Disciplined Entrepreneurship Workbook, Wiley, 2017. 3. Alexander Osterwalder et al, Business Model Generation, Wiley, 2010. 4. Ash Fontana, The AI-First Company: How to Compete and Win with Artificial Intelligence, Penguin, 2021. 5. Peter Thiel και Blake Masters, Zero to One: Notes on Startups, or How to Build the Future, Virgin Books, 2015. 6. Cade Metz (2021). "The Genius Makers: The Mavericks Who Brought A.I. to Google, Facebook, and the World." Random House Business. 7. Lee, Kai-Fu (2018). "AI Superpowers: China, Silicon Valley, And The New World Order." Houghton Mifflin Harcourt Company. 8. O'Neil, C. (2016). "Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy." Crown. 9. Alexander Osterwalder et al, "Value Proposition Design: How to Create Products and Services Customers Want." Wiley, 2014. 10. Ben Horowitz, "The Hard Thing about Hard Things." Harper Business, 2014. 11. Clayton Christensen, "The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail (Management of Innovation and Change)." Harvard Business Review Press, 2016. 12. Jeff Bezos, "The Everything Store: Jeff Bezos and the Age of Amazon." Corgi, 2014. 13. Geoffrey G. Parker, Marshall W. Van Alstyne and Angeet Paul Choudary, "Platform Revolution." W.W. Norton and Co., 2016. 14. European Patent Office. Inventors' Handbook. https://www.epo.org/learning-events/materials/inventors-handbook.html 15. Y Combinator's Resources, https://www.ycombinator.com/resources/ 16. Steve Blank, "How to build a startup?" Udacity, https://classroom.udacity.com/courses/ep245 17. Sam Altman, "How to start a startup?" http://startupclass.samaltman.com/
Assessment	Student progress is evaluated continuously through class participation and the assessment of in-class participation and presentations, writing assignments, group project deliverables, and final exam. The final grade is based on the following formula: • Individual Class Participation and Attendance: 15% • Individual Assignments, Deliverables and Peer Reviews: 10% • Term Project Group Deliverables (based on slides prepared for each of 24 steps of DE): 30% • Term Project Final Deliverables: 20% • Final Exam: 25%
Language	English

Course Title	***Natural Language Processing***						
Course Code	***MAI 623***						
Course Type	Compulsory						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	Demetris Paschalides						
ECTS	8	Lectures / week	3 hours	Recitation / week	0 hours	Laboratories / week	1.5 hours
Course Purpose and Objectives	Natural language processing (NLP) seeks to provide computers with the ability to intelligently process human language, extracting meaning, information, and structure from text, speech, web pages, and social networks. The goal of this course is to provide the fundamental aspects of NLP systems, as well as introduce recent advancements in the field of NLP and Deep Learning. The course is organized into two parts: Part 1: Fundamental knowledge, concepts, and techniques of NLP. Part 2: Introduction to Deep Learning methods for NLP.						
Learning Outcomes	The students who complete this course successfully will be able to: • Comprehend various fundamental concepts of NLP: Text processing (normalization, lemmatization, stemming, etc.), language models (N-Grams), word representation (word embeddings), and text classification with Machine Learning. • Familiarize with known NLP tasks: Named Entity Recognition (NER), Part-of-Speech tagging (PoS), Dependency and Syntax parsing. • Employ Machine Learning (ML) techniques for text classification (e.g. Naive Bayes) and be able to properly apply the NLP feature engineering process. • Extend their knowledge with advanced methods in NLP and Deep Learning: Word Vectors, Word2vec algorithm, Transformers (e.g. BERT), and Large Language Models (LLMs). • Apply their knowledge on real-world research applications of NLP and recognize the societal impact in cases of misinformation and hate-speech identification. • Design efficient and effective NLP solutions to a variety of problems, using state-of-the-art tools.						

Prerequisites	• Machine Learning Course: Throughout the course we will refer to fundamental aspects of machine learning and text classification, as well as deep learning model structures. If the students have basic machine learning and/or deep learning knowledge, this course will be straightforward. • Familiarity with Python: Any programming code and libraries referred to or used during the course will be in Python. It is not required for the students to know Python, and there is flexibility in the programming languages that can be used, but a familiarity with Python will significantly improve the student resources in the topics of both Machine Learning and Natural Language Processing.	Required	None

Course Content	The content of the course is organized in the following groups: Part 1: Fundamental Knowledge, Concepts, and Techniques of NLP. • Introduction: A brief introduction of Natural Language Processing (NLP) in the matters of extracting information from language using methods such as Information Retrieval (IR), text classification with Machine Learning (ML) and sentiment analysis. • Basic Text Processing: Fundamental applications of text processing using regular expressions, sentence segmentation, word tokenization and word normalization using lemmatization and stemming. Also introduction to string similarity techniques such as minimum edit distance (MED). • Language Modeling: Introduction to probabilistic language models using N-gram models (e.g. Unigrams, Bigrams etc). Methods for language model evaluation (intrinsic and extrinsic) and generalization, such as perplexity. • Text Classification: Use of ML algorithms (e.g. Naive Bayes) to classify text into categories. Identify distinguishable textual features using feature extraction and feature selection methods. Classification task evaluation measures of Precision, Recall and F-measure. Introduction to sentiment analysis task. • Natural Language Processing Tasks: Introduction to widely known NLP tasks such as Part-of-Speech (PoS) tagging, Named Entity Recognition (NER), and Dependency and Constituency parsing. Part 2: Introduction to Deep Learning methods for NLP. • Word Vector Semantics: Introduction to vector space models and word distributional similarity. Analysis of the word2vec algorithm and word embeddings usage in NLP tasks. • NLP with Deep Learning: Introduction to Deep Learning (DL) for NLP. Take a look into Neural Network structures used in NLP, such as Multi-Layer Perceptrons (MLP), Recurrent Neural Networks (RNN) and Convolutional Neural Networks (CNN). Introduction to recent DL advancements in NLP with Seq2Seq, followed by a detailed overview of how Transformers operate (e.g. encoder-decoder, self-attention). We will then introduce Large Language Mmodels (LLMs), covering key concepts such as in-context learning , fine-tuning, and specific metrics. • Future of NLP: Discuss the future of NLP using examples from recent advancements in both industry and research. Also discuss the ethical issues that arise along with the advancement in NLP. Tools & Libraries: Jupyter Notebooks, Scipy, NumPy, SpaCy, Gensim, NLTK, PyTorch, Transformers

Teaching Methodology	Students will meet the expected learning outcomes through participation in lectures, active participation and in class discussions, and actual practice with programming assignments and the final project. The lectures will be hybrid, with the possibility of both physical and virtual presence of the students.
	In summary, the teaching and learning methods are the following:
	• Participation to Lectures:, where the main concepts and methodologies covered by the course are presented and critically appraised. Students are required to review and study the materials assigned for each lecture and actively participate in class. Participation in precepts, where the class will expand on topics covered in the lectures, through presentation of tools, open discussion, and viewing and discussing of relevant online material.
	• Coding Assignments: There are 3 coding assignments, which will improve both theoretical understanding and practical skills of the students. All assignments contain both written questions and programming parts.
	• Final project, where each team of students is expected to work on a project that will be assigned by the professor. Students can also suggest their own project ideas, which must first be evaluated. The project comprises three deliverables:
	  o Project Description and Plan report.
	  o Oral Presentation.
	  o Developed Code and Written Final report.

Bibliography	1. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, Dan Jurafsky and James H. Martin, Prentice Hall PTR, Upper Saddle River, NJ, United States, 2025 2. Introduction to Natural Language Processing, Jacob Eisenstein, 2019, MIT Press
Assessment	Student progress is evaluated continuously through class participation and the assessment of at-home assignments, group project deliverables, and two quizzes (midterm and final). The final grade is based on the following formula (the percentages are indicative): • Assignments and Participation: 25% • Final Project: 35% • Quizzes: 40%
Language	English

Course Title	***Artificial Intelligence on the Edge Webinars***						
Course Code	***MAI 632***						
Course Type	Mandatory						
Level	Graduate						
Year / Semester	Fall						
Teacher's Name	George Pallis						
ECTS	2	Lectures / week	0 hours	Recitation / week	0 hours	Laboratories / week	0 hours
Course Purpose and Objectives	This course aims to expose students to the latest developments in Artificial Intelligence, with a focus on emerging trends and the breadth of AI research worldwide. It fosters an appreciation of the diverse approaches and global research efforts that drive innovation in AI.						
Learning Outcomes	Upon successful completion of the course, students will be able to: • Demonstrate awareness of current advancements and challenges in AI. • Critically analyze key topics presented in AI-focused webinars and/or colloquiums. • Identify and contextualize research initiatives in AI.						
Prerequisites	None		Required		None		
Course Content	• A list of AI-related webinars is updated weekly. • Students select 12 webinars to attend throughout the semester.						
Teaching Methodology	For each webinar attended, each student must submit a structured report that includes: • A concise summary of the webinar/colloquium. • A critical analysis of the topic's relevance and approach. • A list of 3–5 key questions addressed in the webinar/colloquiums • A concluding section with 2–4 forward-looking questions to stimulate further exploration of the topic. Reports are evaluated for clarity, depth of analysis, and engagement with the topic. Webinars are made available to students on a weekly basis.						
Bibliography	List of AI-focused Webinars						
Assessment	• A report is submitted for each webinar that the student has attended • Grading: PASS / FAIL, based on the quality and completeness of reports for attended webinars/colloquiums.						
Language	English						

Course Title	***Responsible AI***						
Course Code	***MAI 633***						
Course Type	Mandatory						
Level	Graduate						
Year / Semester	Fall						
Teacher's Name							
ECTS	6	Lectures / week	3 hours	Recitation / week	0 hours	Laboratories / week	0 hours
Course Purpose and Objectives	This course introduces students to the foundational principles and practical approaches of responsible Artificial Intelligence (AI). It focuses on equipping students with the critical skills and ethical frameworks required to assess and address challenges related to fairness, transparency, accountability, and societal impact in AI systems. Through real-world case studies and regulatory contexts, students will gain the ability to evaluate ethical implications, design AI systems aligned with responsible innovation, and critically engage with emerging legal and policy frameworks. The course aims to develop not only technical awareness but also policy literacy and ethical reflexivity in the development and deployment of AI systems.						
Learning Outcomes	By the end of the course, students will be able to: • Identify and evaluate ethical challenges in AI development and use. • Analyze the societal and legal implications of AI technologies. • Apply frameworks for fairness, accountability, transparency, and ethics in AI systems. • Design AI solutions aligned with responsible innovation principles. • Critically assess AI policies, regulations (e.g., EU AI Act), and ethical guidelines.						
Prerequisites	MAI621: AI Ethics		Required		Basic technical knowledge of AI.		
Course Content	• Responsible AI Pillars Introduce Responsible AI pillars • Bias, Fairness, and Discrimination in Algorithms Techniques for identifying and mitigating algorithmic and data bias. • Explainability and Transparency in Machine Learning Understanding AI decisions through explainable AI (XAI) techniques and model interpretability. • Sustainability in Algorithms Understanding carbon footprint, environmental costs, techniques for making AI sustainable • Safety, Privacy, Reliability in Algorithms Techniques for safety, privacy, and reliability • Accountability, AI Governance, Policies, and the EU AI Act Frameworks for developing and implementing responsible AI governance, accountability and policy. • Case Studies and Real-World Applications Practical engagement with real-world examples from various domains and industries. Case studies of how Responsible AI went well and gone wrong, showing both successes and failures.						
Teaching Methodology	• Lectures • Practical exercises • Student presentations						

Bibliography	Main Text: 1. Ben Shneiderman's book on Human-Centered AI: https://global.oup.com/academic/product/human-centered-ai-9780192845290?cc=cy&lang=en& 2. Human-Centered AI: An Illustrated Scientific Quest: https://link.springer.com/book/10.1007/978-3-031-61375-3 3. D. Quercia and M. Constantinides, "Operationalizing Responsible AI: Opportunities and Challenges", Cambridge University Press with expected pub date Oct-Dec 2025. Papers, as reading material, will be made available to students on a weekly basis.
Assessment	2 assignments (one group, one individual) and a final exam.
Language	English

Course Title	***Master Thesis***						
Course Code	***MAI 641***						
Course Type	Elective						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	Elpida Keravnou-Papailiou						
ECTS	16	Lectures / week	0 hours	Recitation / week	0 hours	Laboratories / week	0 hours
Course Purpose and Objectives	The main objective of this course is to enable the students to develop deeper knowledge, understanding, capabilities and attitudes in the context of the programme of study. The thesis should be written at the end of the programme and offers the opportunity to delve more deeply into and synthesise knowledge acquired. The thesis will place emphasis on the technical and/or scientific aspects of the subject matter.						
Learning Outcomes	• Considerably more in-depth knowledge of the major subject/field of study, including deeper insight into current research and development work. • Deeper knowledge of methods in the major subject/field of study. • A capability to contribute to research and development work. • The capability to use a holistic view to critically, independently and creatively identify, formulate and deal with complex issues. • The capability to plan and use adequate methods to conduct qualified tasks in given frameworks and to evaluate this work. • The capability to create, analyze and critically evaluate different technical solutions. • The capability to critically and systematically integrate knowledge. • The capability to clearly present and discuss the conclusions as well as the knowledge and arguments that form the basis for these findings in written and spoken English. • A consciousness of the ethical aspects of research and development work.						
Prerequisites	A student must complete successfully courses, totaling at least 45 ECTS credits from the Graduate Programme	Required	None				
Course Content	-						
Teaching Methodology	For a student to undertake a Master Thesis a Research Advisor, based on the rules of the University Senate, is assigned to the student before s/he submits the Thesis Proposal. The Thesis must deal with a research topic or a technical issue. It must be of some original contribution or show a thorough and clear understanding of some special topic. A student participating in the AI Camp and/or completing successfully an industrial internship may discuss the possibility of doing his/her Master Thesis in collaboration with an industrial partner.						
Bibliography	The bibliography of this course will be determined by the Research Advisor.						

Assessment	The Master Thesis is submitted at the Department and defended within the time period decided by the Departmental Council. • When a Master Thesis is submitted the Chair of the Department appoints a Thesis Examination Committee of three members. The head of the Committee is the student's Research Advisor. It is possible a member of this Committee to be a faculty member of another department or to not be a faculty member but to hold a PhD Degree or have a reputation in the field of study. The Committee can have also external members who do not hold a faculty position. However, the membership must be approved by the Departmental Council and the external member must hold a PhD Degree or must have a reputation in the field of study. • The Master Thesis is defended in a presentation before the Thesis Examination Committee. The head of the Thesis Examination Committee is responsible for the procedures that must be followed during the defense. • The Thesis Examination Committee can accept (even with conditions) or reject a Master Thesis. The Thesis Examination Committee writes its decision in a Master Thesis Evaluation Form and submits it to the Departmental Council for approval. When the Thesis is accepted the Department informs the Student Affairs Office for the graduation procedure of the student. When the Thesis is rejected, the student can follow the suggestions of the Committee and resubmit it for the first and the last time. The submission and Thesis defense procedure is repeated from the beginning.
Language	English

Course Title	***Deep Learning***						
Course Code	***MAI 642***						
Course Type	Mandatory						
Level	Graduate						
Year / Semester	Fall						
Teacher's Name	Constantinos Pattichis, Theocharis Theocharides						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	1.5 hours
Course Purpose and Objectives	The objective of this course is to provide a complete introduction to deep machine learning. How to design a neural network, how to train it, and what are the modern techniques that specifically handle very large networks.						
Learning Outcomes	• Design a deep neural network [A1] [B2] • Design a training procedure [A1] [B1] [B2] • Design a sound evaluation protocol [B1] [B4]						
Prerequisites	None		Required		None		
Course Content	The course aims at teaching the required skills to use deep learning methods on applied problems. It will show how to design and train a deep neural network for a given task, and the sufficient theoretical basis to go beyond the topics directly seen in the course. The planned content of the course: - What is deep learning, introduction to tensors. - Basic machine-learning, empirical risk minimization, simple embeddings. - Linear separability, multi-layer perceptrons, back-prop. - Generalized networks, autograd, batch processing, convolutional networks. - Initialization, optimization, and regularization. Drop-out, activation normalization, skip connections. - Deep models for Computer Vision. - Analysis of deep models. - Auto-encoders, embeddings, and generative models. - Deep learning for sequences - Recurrent neural networks (RNNs); vanishing and exploding gradients; Long Short-Term Memory (LSTM); deep RNNs; bidirectional RNNs; combination of CNNs with RNNs - pytorch tensors, deep learning modules, and internals.						
Teaching Methodology	Lectures (3 hours weekly), Recitation (1 hour weekly) and Laboratory (1,5 hours weekly).						
Bibliography	1. Deep Learning - An MIT Press book, Ian Goodfellow and Yoshua Bengio and Aaron Courville - http://www.deeplearningbook.org						
Assessment	Final exam, midterm exam, homework (programming and theoretical assignments) and semester project.						
Language	English						

Course Title	***Artificial Intelligence in Medicine***						
Course Code	***MAI 643***						
Course Type	Elective						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	Elpida Keravnou-Papailiou Pavlos Antoniou (Project Instructor)						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	0 hours
Course Purpose and Objectives	The medical domain has presented key challenges to the AI community from the early days of AI research. It is not an exaggeration to say that this pioneering work, particularly in medical expert systems, and its undisputable successes, some in real-life settings, has helped both in restoring confidence in the promise of AI, that at some point was disturbed after its failure to deliver fully on the very ambitious initial goals that it had set, and in paving the way towards more viable paths harnessing the mechanization of knowledge and human expertise. AI in Medicine (AIM) is as old as AI itself, initially focusing on modelling human expertise, researching at the same time the cognitive processes involved in developing from a novice to an expert problem solver, as well as on intelligent tutoring systems for medical students and automated support for various clinical tasks. Over the years the initial focus on knowledge engineering has expanded to include ontologies and terminologies, natural language processing and text mining, guidelines and protocols, temporal information management, distributed and cooperative systems, uncertainty management, machine learning, image and signal processing and others. Recent interest focuses on medical analytics for healthcare intelligence. New challenges continuously arise, triggered from and/or triggering, important technological advances. The aim of this elective course is to familiarize students with the past, present and future of Artificial Intelligence in Medicine, illustrating the discussion with a number of case studies, and pinning down the human-centric and ethical aspects underlying the given applications.						
Learning Outcomes	Upon completion of the course the students will have a good understanding, from a critical perspective, of the span of applications of AI methods and techniques in the medical domain, and the methodologies used in developing such applications. More specifically the students will understand the importance of time in medical information systems and how time can be modelled, be conversant with data-driven clinical decision-making, and grasp the regulatory, social, ethical and legal issues of Artificial Intelligence in Medicine.						
Prerequisites	Artificial Intelligence Fundamentals	Required	None				
Course Content	The question that is often asked is whether medicine is an art or a science. The support of the various medical tasks (diagnosis, prognosis, treatment, patient monitoring) through computer systems concerns a number of scientific communities (artificial intelligence, databases, medical informatics, biomedical engineering). Moreover, in recent years there has been a shift in approach, from knowledge-intensive to data-intensive applications, and from consulting systems to information systems. The main challenge now is the intelligent use of data and not necessarily the mechanization of knowledge. Data analytics, explainability, time representation and reasoning and healthcare intelligence are central components of the current AIM landscape. The course aims to cover the following thematic units: • Introducing AIM and tracing its history • AIM: from knowledge-intensive to data-intensive applications • Probabilistic Graphical Models in medicine • Biomedical data • Public health applications and ethics • Clinical Cognition in AI • Time representation and temporal reasoning in medicine • Temporal clinical diagnosis • Case-based systems in medicine • Protocols and Guidelines • Explainability in Medical AI • The future of AIM						

Teaching Methodology	Lectures and discussions particularly around the presented applications and case studies. Students would be strongly guided to view all topics presented and discussed with a critical eye.
Bibliography	Main texts: 1. T.A. Cohen, V.L. Patel and E.H. Shortliffe (editors), Intelligent Systems in Medicine and Health: The Role of AI, Springer, 2022. 2. A.C. Chang, Intelligence-Based Medicine: Artificial Intelligence and Human Cognition in Clinical Medicine and HealthCare, Academic Press, 2020. 3. C. Combi, E. Keravnou-Papailiou and Y. Shahar, Temporal Information Systems in Medicine, Springer, 2010. 4. L. Xing, M.L. Giger and J.K. Min (editors), Artificial Intelligence in Medicine: Technical Basis and Clinical Applications, Academic Press, 2021. Other reading: 1. W. Hersh (ed.), Health Informatics: Practical Guide, 8th edition, Lulu.com, 2022. 2. E.H. Shortliffe and J.C. Cimino (editors) and M.F. Chiang (co-editor), Biomedical Informatics: Computer Applications in Health Care and Biomedicine, Springer, 2021. 3. A. Panesar, Machine Learning and AI for Healthcare: Big Data for Improved Health Outcomes, Apress, 2021. 4. T. Lawry, AI in Health: A Leader's Guide to Winning the New Age of Intelligent Health Systems, CRC Press, 2020. 5. N. Lavrac, E.T. Keravnou and B. Zupan (editors), Intelligent Data Analysis in Medicine and Pharmacology, Kluwer Academic Publishers, 1997. 6. E.T. Keravnou (editor), Deep Models for Medical Knowledge Engineering, Elsevier Science Publishers, 1992. 7. Scientific papers from thematic and/or standard issues of relevant journals, primarily the journal Artificial Intelligence in Medicine (AIME) published by Elsevier.
Assessment	Assignments (50%) [(i) Project on a clinical decision support system (35%) and (ii) a critical opinion article (15%)] and final exam (50%).
Language	English

Course Title	***Computer Vision***						
Course Code	***MAI 644***						
Course Type	***Elective***						
Level	***Graduate***						
Year / Semester	***Spring***						
Teacher's Name	***C. Pattichis / G. Chrysanthou***						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	1.5 hours
Course Purpose and Objectives	The objective of this course is to understand the basic issues in mechanical vision and the major approaches that address them. Through the duration of the course, vision is treated as a process of inference from noisy and uncertain data in order to answer the question of how computers can understand the visual world of humans.						
Learning Outcomes	Upon the course's completion, students should: • Understand the principles of vision. • Appreciate the commercial potential of computer vision, but also understand the limitations of current methods. • Identify computer vision's near-future goals. • Apply mathematical and statistical methods as a means to solve vision-related problems. • Be able to extract features of interest from images. • Understand the process of perspective projection, be able to calibrate cameras, and manipulate 3D and 2D transformations. • Be able to recover 3D shape information based on multiple viewpoints. • Understand the principle of inference techniques such as the Kalman filter and the particle filter, and their use in 3D tracking. • Be able to develop simple programs utilizing computer vision algorithms.						
Prerequisites	MAS 029: Linear Algebra MAI 231: Data Structures and Algorithms	Required	None				
Course Content	Βασικές έννοιες and μεθοδολογίες που αφορούν το αντικείμενο της Μηχανικής Όρασης. Σχηματισμός Εικόνας, επεξεργασία εικόνας, ανίχνευση χαρακτηριστικών, κατάτμηση εικόνων and ομαδοποίηση χαρακτηριστικών, επεξεργασία πολλαπλών εικόνων, μελέτη εφαρμογών.						
Teaching Methodology	Lectures (3 hours weekly), Recitation (1 hour weekly) and Laboratory sessions (1.5 hours weekly).						
Bibliography	1. D. Forsyth and J. Ponce, Computer Vision: A Modern Approach, Prentice-Hall, 2003. 2. R. Hartley and A. Zeisserman, Multiple View Geometry, Cambridge University Press, 2003. 3. C. Bishop, Pattern Recognition and Machine Learning, Springer-Verlag, 2007. 4. O. Faugeras and Q. T. Luong, Geometry of Multiple Images, MIT Press, 2001. 5. B. Horn, Robot Vision, MIT Press, Cambridge, Massachusetts, 1986.						
Assessment	Final exam, midterm exam and homework (programming assignments).						
Language	Greek						

<table>
<tr><td>Course Title</td><td colspan="7">Machine Learning for Graphics and Computer Vision</td></tr>
<tr><td>Course Code</td><td colspan="7">MAI 645</td></tr>
<tr><td>Course Type</td><td colspan="7">Elective</td></tr>
<tr><td>Level</td><td colspan="7">Graduate</td></tr>
<tr><td>Year / Semester</td><td colspan="7">Spring</td></tr>
<tr><td>Teacher’s Name</td><td colspan="7">Andreas Aristidou</td></tr>
<tr><td>ECTS</td><td>8</td><td>Lectures / week</td><td>3 hours</td><td>Recitation / week</td><td>1 hour</td><td>Laboratories / week</td><td>1.5 hours</td></tr>
<tr><td>Course Purpose and Objectives</td><td colspan="7">This course will offer an introduction to machine learning algorithms, the use of deep learning and its applications in computer vision and graphics. The course will also operate as a graduate-level seminar with weekly readings (1 hour per week), summarizations, and discussions of recent papers.
Machine Learning Topics:
• Classification, Regression
• Random Forests
• Deep Neural Networks
• Recurrent Neural Networks
• Generative Models
• Generative Adversarial Networks
• Transformers
Vision and Graphics Applications:
• Image Recognition, Object Detection
• Semantic Segmentation
• Stereo & Multi-view Reconstruction
• Inpainting, Rendering Faces,
• Composite Image generation, Style transfer
• Motion capture and synthesis,
• Denoising, virtual reality</td></tr>
<tr><td>Learning Outcomes</td><td colspan="7">Participants will explore the latest developments in neural network research and deep learning models that are enabling highly accurate and intelligent computer vision and graphics systems.

By the end, participants will:
• Be familiar with fundamental concepts and applications in computer vision and graphics.
• Grasp the principles of state-of-the art deep neural networks.
• Gain knowledge of high-level vision tasks, such as object recognition, scene recognition, face detection and human motion categorization.
• Gain knowledge of high-level graphics tasks, such as composite image generation, style transfer, motion reconstruction, and motion synthesis.
• Develop practical skills necessary to build highly-accurate, advanced computer vision and graphics applications</td></tr>
<tr><td>Prerequisites</td><td colspan="3">Knowledge of a high-level programing language, and experience in programming with Python.
Experience with linear algebra, calculus, statistics, and probability.</td><td>Required</td><td colspan="3">Previous knowledge on Computer graphics and Computer Vision is encouraged.</td></tr>
</table>

Course Content	• L01 Introduction Basic regression, understanding of linearity and non-linearity. DL for Computer Vision • L02 Learning from Images Deep learning for image classification and object detection. • L03 Learning from Videos Deep learning for video classification. • L04 Feature Extraction Deep learning for feature extraction. • L05 Semantic Understanding Deep learning for semantic segmentation, Deep learning: visualizing networks, impainting, saliency detection (GAN). • L06: Creative applications Photo collections: style and enhancement, Ambiguity and style, style transfer. • L07: Vision->Graphics Computer vision as inverse computer graphics, Novel image synthesis – compositional image generation. DL for Computer Graphics • L08 From 2D to 3D What is 3D Vision, 3D shape representations, 3D shape datasets, 3D Deep Learning architectures. 3D meshes and point clouds. • L09 Inverse graphics in practice: Generation Load and stored images/3D data. 3D labeling/classification • L10 Creative Applications Generative networks, generating faces, landscapes, portraits, Sketches, denoising. • L11 Motion Motion capture, character animation, and synthesis (style transfer, retargeting, control), deep reinforcement learning for physics-based animation and authoring, pose representation, popular deep character animation networks. • L12 (neural) rendering, physics, materials, virtual reality. • L13 Advanced Topics in Deep Learning Levels of supervision; Adversarial training, open problems.
	Weekly lectures, discussions of practical examples and (unsupervised) lab activities where the active learning element is encouraged and supported. Weekly or bi-weekly recitation and discussion on recent papers, with students' presentations. Students would be strongly guided to view all topics presented and discussed with a critical eye.
Bibliography	1. Deep Learning, by Ian Goodfellow, Yoshua Bengio, Aaron Courville, MIT Press, 2016 2. Computer Vision: Advanced Techniques and Applications, by Steve Holden, Clanrye International, 2019 3. Pattern Recognition and Machine Learning, Christopher Bishop, Springer, 2016
Assessment	Student paper presentations (15%); Programming assignments (35%); Final course project (50%)
Language	English

Course Title	***Computational Neuroscience***						
Course Code	***MAI 647***						
Course Type	Elective						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	Chr. Christodoulou						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	1.5 hours
Course Purpose and Objectives	Computational Neuroscience is an emerging and dynamically developing field aiming to elucidate the principles of information processing by the nervous system. This course aims to develop and apply computational methods for studying brain and behaviour as well as understanding the dynamics of the conscious mind.						
Learning Outcomes	The learning outcomes for the students are the following: • understand and be able to explain the fundamental principles of information processing by neural systems • appreciate the importance of computational neuronal models in the quest of understanding the brain and the fact that many aspects of neuroscience cannot be understood without appropriate computational modeling framework • understand the most important biophysical neuronal models and the different levels of description and complexity in computational neuronal modelling from the level of the single neuron to that of neural networks • understand neuronal dynamics and learn how high dimensional neuronal models can be reduced to low dimensional neural models • understand how experimentally recorded physiological signals enable us to understand the functionality of neurons/systems in the brain and how statistical approaches help in the analysis of such data • be able to implement/simulate basic computational neuronal models through programming • become familiar and be able to use various computational neuroscience simulation software packages for modelling complex biophysical models and experimentally observed phenomena • be able to grasp the importance of high level modelling abstraction from the underlying neuronal principles for understanding brain behaviours • critical reading and discussion of recently published scientific papers						
Prerequisites	Linear Algebra, Differential Equations	Required	None				
Course Content	Introduction to Computational Neuroscience; basic neurobiology: from the brain to single neurons; biophysics of single neurons; synapses; dendrites and axons. Conductance-based neuron models: the generation of action potentials and the Hodgkin and Huxley equations. Spiking neuron models and response variability: leaky integrator and leaky integrate-and-fire (LIF) type neuron models; spike time variability. Two dimensional (2D) neuron models: reduction of the four dimensional (4D) HH model to a 2D model; phase plane analysis of 2D models/nullclines; FitzHugh-Nagumo model; neuronal dynamics. Modelling synapses/inputs to neurons. Neuron models beyond HH – more ion channels and their functions. Cable Theory: neuronal structure; passive/active membranes; modelling axons and dendrites; action potential propagation. Compartmental models. Neural coding: firing rate; rate code; temporal code; neural operational modes – temporal integration/coincidence detection. Synaptic plasticity: Hebbian learning; Spike-Timing Dependent Plasticity. Bottom-up/top-down modeling of the brain: modeling of self-control behaviour as an example of top-down modelling. Bottom-up/top-down modeling of the brain: modeling of self-control behaviour as an example of top-down modelling. Modelling consciousness.						
Teaching Methodology	Lectures (3 hours weekly), Recitation (1 hour weekly) and Laboratory sessions (1.5 hours weekly).						

Bibliography	1. P. Dayan and L. Abbott, Theoretical Neuroscience: Computational and Mathematical Modelling of Neural Systems, MIT Press, 2001. 2. D. Sterratt, B. Graham, A. Gilles and D. Willshaw, Principles of Computational Modelling in Neuroscience, Cambridge University Press, 2011. 3. W. Gerstner, W. M. Kistler, R. Naud, L. Paninski, Neuronal Dynamics: From single neurons to networks and models of cognition, Cambridge: Cambridge University Press, 2014. 4. C. Koch, Biophysics of Computation: Information Processing in Single Neurons, Oxford University Press, 1998. 5. E. M. Izhikevich, Dynamical Systems in Neuroscience: the Geometry of Excitability and Bursting, MIT Press, 2007.
Assessment	Final exam, midterm exam and laboratory exercises/oral presentations of selected research papers.
Language	Greek

<table>
<tr><td>Course Title</td><td colspan="7">Human-centered Intelligent User Interfaces</td></tr>
<tr><td>Course Code</td><td colspan="7">MAI 648</td></tr>
<tr><td>Course Type</td><td colspan="7">Elective</td></tr>
<tr><td>Level</td><td colspan="7">Graduate</td></tr>
<tr><td>Year / Semester</td><td colspan="7">Fall or Spring</td></tr>
<tr><td>Teacher’s Name</td><td colspan="7">Argyris Constantinides</td></tr>
<tr><td>ECTS</td><td>8</td><td>Lectures / week</td><td>3 hours</td><td>Recitation / week</td><td>0 hours</td><td>Laboratories / week</td><td>1.5 hours</td></tr>
<tr><td>Course Purpose and Objectives</td><td colspan="7">The purpose of the course is to introduce students to fundamental principles and methods within the intersection of Artificial Intelligence (AI) and Human-Computer Interaction (HCI) aiming to design and develop more efficient and effective user interfaces through the use of intelligent computation methods.</td></tr>
<tr><td>Learning Outcomes</td><td colspan="7">Upon completion of this course, students will have acquired: i) an in-depth understanding of theoretical and practical aspects of intelligent user interfaces; ii) skills to design, develop and evaluate intelligent interactive systems by considering a variety of human factors, such as human cognitive and emotional characteristics for improving the efficiency, effectiveness and user experience in interactive systems; and iii) abilities to synthesize and evaluate the potential of this knowledge in relation to deploying intelligent user interfaces in real-life applications.</td></tr>
<tr><td>Prerequisites</td><td colspan="3">None</td><td colspan="2">Required</td><td colspan="2">Knowledge on AI and HCI fundamentals; general programming knowledge</td></tr>
<tr><td>Course Content</td><td colspan="7">This interdisciplinary course covers a variety of topics within the intersection of AI and HCI, with an emphasis on incorporating human factors in intelligent interactive systems’ design. Specifically, it covers: i) theoretical foundations of intelligent user interfaces and the importance of human factors in such contexts; ii) state-of-the-art processes and techniques for implementing intelligent interactive systems; and iii) research and practice on how intelligent user interfaces can be applied in various application domains. The course content includes the following topics:
• Introduction to Intelligent User Interfaces: Overview of intelligent interactive systems; history and evolution of intelligent user interfaces; fundamentals of AI; fundamentals of HCI; importance of incorporating intelligent capabilities in human-computer interaction.
• Conceptual and Architectural Overview of Intelligent Interactive Systems: Conceptual frameworks, architectural design and core components of intelligent interactive systems.
• Human Cognitive Factors in Intelligent Interactive Systems: The role of human factors in the design of intelligent interactive systems; theory on human cognition and information processing; cognitive elicitation techniques; implications of cognitive factors in intelligent interactive systems.
• Affective Computing: Theories of human emotion; emotion elicitation techniques based on physiological signals; conversational agents; assistive agents.
• User Modelling: User modeling factors; user data collection methods; user model elicitation and generation techniques; modeling human factors; clustering; classification.
• Adaptive User Interfaces and Recommender Systems: Personalization categories; adaptation and decision-making mechanisms; content-based filtering techniques, collaborative filtering techniques; adaptation effects; design guidelines.
• Voice User Interfaces: Speech recognition; state-of-the-art methods and techniques of voice-based user interfaces; text-to-speech.
• Natural Language Processing (NLP): NLP technologies and frameworks; NLP Application Programming Interfaces.
• Explainable Artificial Intelligence (XAI): Importance of XAI; state-of-the-art techniques; guidelines for designing XAI user experiences.
• Examples of Intelligent User Interface Systems: Voice user interfaces; emotion recognition; recommender systems; personalized and adaptive user interfaces; biometric technologies.</td></tr>
</table>

Teaching Methodology	Lectures covering the theoretical foundations of intelligent user interfaces, discussion of practical examples, and lab activities for designing and implementing intelligent user interfaces.
Bibliography	Relevant readings: 1. Germanakos, P., Belk, M. (2016). Human-Centred Web Adaptation and Personalization - From Theory to Practice. Human-Computer Interaction Series, Springer, DOI: 10.1007/978-3-319-28050-9 2. Brusilovski, P., Kobsa, A., Nejdl, W. (2007). The Adaptive Web: Methods and Strategies of Web Personalization, Springer, DOI: 10.1007/978-3-540-72079-9 3. Shneiderman, B., Plaisant, C., Cohen, M., Jacobs, S., Elmqvist, N., Diakopoulos, N. (2017). Designing the User Interface: Strategies for Effective Human-Computer Interaction (6th Edition), Pearson, ISBN: 9780134380384 4. Preece, J., Sharp, H., Rogers, Y. (2015). Interaction Design: Beyond Human-Computer Interaction (4th Edition), Wiley, ISBN: 9781119088790 5. Mourlas, C., Germanakos, P. (2008). Intelligent User Interfaces: Adaptation and Personalization Systems and Technologies, Information Science Reference, ISBN: 9781605660325 Relevant conferences and journals: 1. ACM Transactions on Interactive Intelligent Systems (TiiS), ACM Press 2. User Modeling and User-Adapted Interaction (UMUAI): The Journal of Personalization Research, Springer 3. Intelligent User Interfaces (IUI), ACM Press 4. Recommender Systems (RecSys), ACM Press 5. User Modeling, Adaptation and Personalization (UMAP), ACM Press 6. Human Factors in Computing Systems (CHI), ACM Press
Assessment	Final exam, midterm exam and homework (theoretical and programming assignments).
Language	English

Course Title	***Principles of Ontological Databases***						
Course Code	***MAI 649***						
Course Type	Elective						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	A. Pieris						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	0 hours
Course Purpose and Objectives	Nowadays we need to deal with data that is very large, heterogeneous, distributed in different sources, and incomplete. At the same time, we have very large amounts of knowledge about the application domain of the data in the form of ontologies that can be used to provide end users with flexible and integrated access to data . This gave rise to ontological databases, which lie at the intersection of traditional databases, and knowledge representation and reasoning. The purpose of the course is to introduce students to the principles of ontological databases and demonstrate the importance of studying data-intensive problems in a mathematically rigorous way, as well as the implications of such studies for real-life applications.						
Learning Outcomes	Upon completion of this course, the students will be able to: • Abstract relational data and relational queries from their physical implementation and formalize them in a rigorous way. • Analyze the complexity of querying relational data and isolate the source of complexity. • Explain the semantics of Datalog queries, analyze the complexity of evaluating Datalog queries, and model queries in a declarative way. • Abstract rule-based ontologies from their physical implementation and formalize them in a rigorous way. • Explain and use the main (forward- and backward-chaining) techniques underlying ontological query answering. • Analyze the complexity of ontological query answering and isolate the source of complexity.						
Prerequisites	None		Required		While there are no formal prerequisites, it is recommended that students have passed an introductory course in Databases (some familiarity with the relational model, and the main relational query languages). It is also recommended that students have some basic familiarity with computational logic (first-order logic), and complexity theory (standard complexity classes such as PTIME and NP).		

Course Content	The main purpose of the course is to introduce students to the principles of ontological databases. To this end, it is vital to first cover the principles of relational databases, without taking ontologies into account, on top of which the principles of ontological databases are built. In particular, the course will cover the following topics: • Relational Model: data model, relational algebra, relational calculus (first-order queries), first-order query evaluation, static analysis of first-order queries (satisfiability, containment, and equivalence). • Conjunctive Queries (CQs): syntax and semantics, CQ evaluation, static analysis of CQs (satisfiability, containment, and equivalence), minimization of CQs, acyclicity of CQs, evaluation of acyclic CQs (Yannakaki's algorithm), semantically acyclic CQs and their evaluation. • Adding Recursion – Datalog: inexpressibility of recursive queries, syntax and semantics of Datalog, Datalog query evaluation, static analysis of Datalog queries (satisfiability, containment, equivalence, and boundedness). • Ontological Databases: rule-based ontologies (syntax and semantics), combining relational databases with rule-based ontologies, ontological query answering (OQA), universal models, ontology-based data access. • Ontological Query Answering: forward-chaining (the chase procedure), backward-chaining (resolution-based query rewriting), linear rule-based ontologies (tractable data complexity, intractable combined complexity, fixed-parameter tractability). • Advanced Topics (time permitting): expressive rule-based ontology languages, chase termination (necessary and sufficient conditions), static analysis of ontological queries (containment and boundedness).
Teaching Methodology	Lectures, discuss solutions to non-trivial problems given in advance (during the weekly recitation hour), review of recent research papers.
Bibliography	1. S. Abiteboul, R. Hull, V. Vianu, Foundations of Databases, 1995 2. M. Arenas, P. Barcelo, L. Libkin, W. Martens, A. Pieris, Database Theory, υπό συγγραφή – προκαταρκτική έκδοση διαθέσιμη στο σύνδεσμο https://github.com/pdm-book/community 3. F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, 2017 4. L. Libkin, Elements of Finite Model Theory, 2012
Assessment	For a technical course of this type, which focuses on the mathematical side of (ontological) databases, exams do not allow us to properly evaluate the students' knowledge of the material. For a proper evaluation, students must be presented with non-trivial problems and tasks, rather than "toy" ones that can be solved in a limited time. Therefore, the assessment of the course consists of the following three components: • Engagement component: During the course, the students will be given 10 exercises that will cover the various topics described above. A serious attempt to solve an exercise will be awarded all the marks, no matter if the provided solution is correct. The solutions of the exercises will be discussed during the recitation hours. • Essay and in-class presentation on the principles of databases (without ontologies): Students will choose a research paper from a given list, and present (i) a summary of the paper and (ii) analysis and critical thoughts (criticism of the paper, discussion on follow-up papers that show how the ideas of the paper under review have influenced the field, ideas for future research directions). There will be also an in-class presentation based on the essay. • Essay and in-class presentation on the principles of ontological databases: As above.
Language	Greek

Course Title	***Internet of Things***						
Course Code	***MAI 650***						
Course Type	Elective						
Level	Graduate						
Year / Semester	Spring						
Teacher's Name	V. Vassiliou						
ECTS	8	Lectures / week	3 hours	Recitation / week	1 hour	Laboratories / week	1.5 hours
Course Purpose and Objectives	The main objective of the course will be to provide an overview of the building blocks of IoT such as sensors and smart devices, M2M communication, data collection and processing, and the role of people and applications.						
Learning Outcomes	Upon completion of this course, students should be able to: • Explain the definition and usage of the term "Internet of Things" in different contexts • Understand and describe the key components that make up an IoT system • Apply the knowledge and skills acquired during the course to build and test a complete, working IoT system involving prototyping, programming and data analysis • Independently research the technological trends which have led to IoT • Understand where the IoT concept fits within the broader ICT industry and recognize possible future trends • Valuate the impact of IoT on society by analysing IoT systems with regard to sustainability, safety, integrity and ethics. • Appreciate the role of big data, cloud computing and data analytics in a typical IoT system						
Prerequisites	MAI 606		Required		None		
Course Content	General Principles and Architecture of IoT Systems, Devices, Detection and Response, Communication Technologies, IoT Communication Protocols, IoT Architecture, IoT Applications						
Teaching Methodology	Lectures (3 hours weekly), Recitation (1 hour weekly) and Laboratory sessions (1.5 hours weekly).						
Bibliography	1. Internet of Things: A Hands-On Approach, by Arshdeep Bahga and Vijay Madisetti, 2016 2. Internet of Things (IoT): Architectures, Protocols and Standards, by Simone Cirani, Gianluigi Ferrari, Marco Picone, and Luca Veltri, Wiley 2018 3. IoT and Edge Computing for Architects – Second Edition" by Perry Lea, O'Reilly, 2020						
Assessment	Final exam, midterm exam and homework.						
Language	English						

<table>
<tr><td>Course Title</td><td colspan="7">Research/Industrial Internship</td></tr>
<tr><td>Course Code</td><td colspan="7">MAI 602</td></tr>
<tr><td>Course Type</td><td colspan="7">Elective</td></tr>
<tr><td>Level</td><td colspan="7">Graduate</td></tr>
<tr><td>Year / Semester</td><td colspan="7">Fall</td></tr>
<tr><td>Teacher’s Name</td><td colspan="7"></td></tr>
<tr><td>ECTS</td><td>8</td><td>Lectures / week</td><td>0 hours</td><td>Recitation / week</td><td>0 hours</td><td>Laboratories / week</td><td>0 hours</td></tr>
<tr><td>Course Purpose and Objectives</td><td colspan="7">This course provides graduate students with hands-on experience in applying advanced AI concepts and methods within a research or industrial setting. The primary objective is to bridge academic knowledge with real-world practice by immersing students in professional environments where they can contribute to active AI projects. Students will refine their technical skills, develop problem-solving abilities, and gain insight into the practical, ethical, and collaborative aspects of AI innovation. The internship also aims to foster professional development and prepare students for careers in research, industry, or entrepreneurship in the AI field.</td></tr>
<tr><td>Learning Outcomes</td><td colspan="7">Upon completion of the course, students will be able to:
Apply AI tools, and techniques to real-world research or industrial challenges.
Design, implement, and evaluate AI solutions in a practical context, under academic or industry supervision.
Demonstrate the ability to work independently and collaboratively within interdisciplinary teams.
Communicate technical ideas, progress, and outcomes effectively to both technical and non-technical audiences.
Critically reflect on ethical, societal, and professional implications of AI deployment in practical settings.
Integrate feedback and adapt to dynamic project requirements in a real-world environment.
Develop professional competencies and identify career pathways in AI research or industry sectors.</td></tr>
<tr><td>Prerequisites</td><td colspan="3">None</td><td>Required</td><td colspan="3">None</td></tr>
<tr><td>Course Content</td><td colspan="7">N/A</td></tr>
<tr><td>Teaching Methodology</td><td colspan="7">N/A</td></tr>
<tr><td>Bibliography</td><td colspan="7">N/A</td></tr>
<tr><td>Assessment</td><td colspan="7">The course is assessed on a Pass/Fail basis. To successfully complete the internship, students must fulfil all of the following requirements:
Internship Report: Submit a comprehensive written report detailing the objectives, methods, AI techniques used, outcomes, and reflections on the internship experience.
Supervisor Evaluation: Receive a satisfactory evaluation from the host organization or academic supervisor, confirming meaningful contribution, professional conduct, and engagement throughout the internship.</td></tr>
<tr><td>Language</td><td colspan="7">English</td></tr>
</table>