Main information on the course		
Course name	Software Engineering for AI-Enabled Systems	
Degree	Computer Science (second-level degree in Computer Science)	
Academic year	2025/26	
European Credit Transfer and Accumulation System (ECTS), in Italian Crediti Formativi Universitari (CFU)		6 CFU (of which 4 T1 and 2 T2 CFU) each CFU corresponds to 25 hours (h) of student's time); CFU are of type T1, T2 or T3 T1 = 8 h lecture + 17 h individual study T2 = 15 h practice + 10 h individual study T3 = 25 h individual study
Settore Scientifico Disciplinare	INFO-01/A - Informatica (previously INF/01)	
Course language	English	
Course year	Second	
Course period	First Semester - exact dates can be found in the didactic regulations	
Course attendance requirement	None, but it is highly recommended to attend classes	
Website of the Degree	https://www.uniba.it/it/corsi/cdl-computer-science/corso-di-laurea-in-computer-science	

Teacher(s)	
Name and Surname	**Filippo Lanubile**
email	filippo.lanubile@uniba.it
phone	via MS Teams
office	6th floor
e-learning platform	https://elearning.uniba.it/
Teacher's homepage	https://collab.di.uniba.it/lanubile/
Office hours	on demand via MS Teams
Name and Surname	**Luigi Quaranta**
email	luigi.quaranta@uniba.it
phone	via MS Teams
office	6th floor
e-learning platform	https://elearning.uniba.it/
Teacher's homepage	https://collab.di.uniba.it/luigi-quaranta/
Office hours	on demand via MS Teams

Syllabus	
Course goals	The main goal of the course is to provide the students with theoretical and practical knowledge on building high-quality, production-grade machine learning (ML) components, ready to be integrated into AI-enabled systems. Thus, MLOps practices and tools are taught and applied into iterative lab sessions.
Prerequisites/requirements	- Fundamentals of software engineering - Fundamentals of data-driven AI / machine learning

Course program	Lectures - Introduction and key concepts - Requirements Engineering for ML - Code and data versioning - Experiment tracking - Quality assurance for ML - ML system design - APIs for ML - Containerization - CI/CD for ML - Monitoring Tutorials and Lab - MLOps tools
Books of reference	Main book - Christian Kastner, *Machine Learning in Production* — https://mlip-cmu.github.io/book/ Other suggested books - Emmanuel Ameisen. *Building Machine Learning Powered Applications*, O'Reilly Media, January 2020, ISBN: 9781492045113 - Mark Treveil and the Dataiku Team. *Introducing MLOps*. O'Reilly Media, November 2020, ISBN: 9781492083290 - Valliappa Lakshmanan, Sara Robinson, Michael Munn. *Machine Learning Design Patterns*. O'Reilly Media, October 2020, ISBN: 9781098115784 Suggested research papers and grey literature: - F. Lanubile, S. Martinez-Fernandez and L. Quaranta, "Training Future Machine Learning Engineers: A Project-Based Course on MLOps" in IEEE Software, vol. 41, no. 02, pp. 60-67, 2024. doi: 10.1109/MS.2023.3310768 https://doi.ieeecomputersociety.org/10.1109/MS.2023.3310768 - F. Lanubile, S. Martínez-Fernández and L. Quaranta, "Teaching MLOps in Higher Education through Project-Based Learning," ICSE 2023, *Software Engineering Education and Training Track,* pp. 95-100, doi: 10.1109/ICSE-SEET58685.2023.00015 - https://arxiv.org/abs/2302.01048v1 - Sato, Wider, Windheuser. "Continuous Delivery for Machine Learning". https://martinfowler.com/articles/cd4ml.html - D. Sculley et al., "Hidden technical debt in machine learning systems" NIPS'15 - A. Vogelsang and M. Borg, "Requirements Engineering for Machine Learning: Perspectives from Data Scientists," REW 2019, https://arxiv.org/abs/1908.04674 - Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19). https://doi.org/10.1145/3287560.3287596 https://arxiv.org/abs/1810.03993 - Luigi Quaranta, Fabio Calefato, and Filippo Lanubile. 2022. Eliciting Best Practices for Collaboration with Computational Notebooks. Proc. ACM Hum.-Comput. Interact. 6, CSCW1, Article 87 (April 2022), 41 pages. https://doi.org/10.1145/3512934 - M.T. Ribeiro, T. Wu, C. Guestrin, and S. Singh. "Beyond Accuracy: Behavioral Testing of NLP Models with CheckList." In Proceedings ACL, p. 4902–4912. (2020). https://homes.cs.washington.edu/~wtshuang/static/papers/2020-acl-checklist.pdf - J. Zhang, M. Harman, L. Ma and Y. Liu, "Machine Learning Testing: Survey, Landscapes and Horizons" in IEEE Transactions on Software Engineering, vol.

	, no. 01, pp. 1-1, 5555. doi: 10.1109/TSE.2019.2962027. https://arxiv.org/pdf/1906.10742.pdf - Jeremy Jordan. A simple solution for monitoring ML systems. https://www.jeremyjordan.me/ml-monitoring/ Collections of online resources: - MLOps Primer. https://elvissaravia.substack.com/p/mlops-primer-2021 - Made With ML: https://madewithml.com/
Notes to the books	Bibliography will be integrated with the slides available on the e-learning platform.

Organization of the didactic activities

Hours				
Total	Lectures	Practice sessions	Project work	Individual study
150 hours	32 hours	30 hours		88 hours
CFU/ETCS				
6 CFU	4 CFU	2 CFU		

Teaching methods

Lectures and tutorials supported by slides and demos.

Expected learning outcomes

Knowledge and understanding	The course adopts a software-engineering perspective on building AI-enabled systems, focusing on understanding how AI components are parts of larger systems and how to turn a machine learning (ML) idea into a scalable and reliable product. It assumes a working relationship between software engineers and data scientists, and focuses on issues of design, implementation, operation, and quality assurance.
Applying knowledge and understanding	The students will be able to turn a data-driven AI/ML model into a scalable and reliable product.
Other skills	*Making judgements* The students will be informed of the main challenges in engineering an AI-enabled system beyond accuracy. *Communication* The students will learn how to communicate in teamwork through individual and collaborative exercises. *Learning skills* The students will be able to autonomously learn theoretical concepts and empirical evidence by reading research papers and grey literature.

Assessment	
Assessment methods	Lab assessment: - A set of tasks consisting in building a reproducible ML pipeline and transforming a ML model from prototype to a production-level web API using containerization in a cloud-based environment. Oral assessment: - oral test, including questions about the course contents
Evaluation criteria	The students should know the concepts presented and discussed during classes and be familiar with the tools introduced in the tutorials and lab sessions.
Measurements and final grade	The lab assessment is measured in thirtieths and is shared by the entire work group. The individual evaluation can be corrected with bonus or malus points depending on the contribution in the workgroup and review meetings. The oral assessment is measured in thirtieths on an individual basis. The final grade is the arithmetic mean of the two assessments.
Further information	