

Report WP2-A1:

The survey of current Artificial Intelligence courses and the ILOs they follow



Result

Report on current AI and ML-related courses

Related to

WP2-A1: The survey of current Artificial Intelligence courses and the ILOs they follow

Statement of originality

This report contains original unpublished work, except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation, or both.

Disclaimer

This report contains material, which is the copyright of TAI Consortium Parties. All TAI Consortium Parties have agreed that the content of the report is licensed under a Creative Commons Attribution Non-Commercial Share Alike 4.0 International License. TAI Consortium Parties does not warrant that the information contained in the Report is capable of use, or that use of the information is free from risk and accept no liability for loss or damage suffered by any person or any entity using the information.

Copyright notice

© 2024-2027 TAI Consortium Parties

Note

For anyone interested in having more information about the project, please see the website at: <https://tai-erasmus.eu/>



This publication is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International Public License](https://creativecommons.org/licenses/by-nc/4.0/) (CC BY-NC 4.0).

Document version history

Version	Date	Author	Comment
1.0	1.3.2024	Tomaž Požrl, UNILJ	Initial version
2.0	20.8.2025	Tomaž Požrl, UNILJ	Final version

Table of contents

1. Introduction	4
2. Courses selection and syllabi collection.....	5
2.1 Relevant topics in artificial intelligence and machine learning for education.....	5
2.2 Collection protocol.....	8
2.3 Collection outcome	9
3. Course syllabi analysis.....	13
3.1 Course AI/ML-related topics	13
3.1.1 Core concepts and learning paradigms.....	14
3.1.2 Key algorithms and techniques.....	16
3.1.3 Model development and evaluation techniques.....	18
3.1.4 Essential Python tools and libraries	20
3.1.5 Engineering-oriented applications of AI/ML.....	22
3.1.6 Ethical, legal, and societal aspects.....	24
3.2 Course content analysis	25
3.2.1 Institutional reports	26
3.2.2 Cross-institution analysis	39
3.3 Course educational goals analysis (Bloom taxonomy).....	43
3.3.1 Institutional reports	45
3.3.2 Cross-institution analysis	62
3.4 Courses' Intended Learning Outcomes analysis	68
3.4.1 ILOs: naming, format, and structure.....	68
3.4.2 Number of ILOs per course	70
3.4.3 Length of ILOs	72
3.4.4 Bloom verbs in ILOs.....	73
4. Conclusion.....	76

1. Introduction

The Teaching Artificial Intelligence (TAI) project brings together five partner universities to strengthen AI education by jointly designing a coherent set of AI MOOCs. To ensure these MOOCs complement existing provision rather than duplicate it, Work Package 2 (WP2) establishes the curricular baseline. WP2 surveys current AI/ML courses at UNILJ, UNSA, UNBG, UoM, and UNIRI and analyses how their syllabi articulate Intended Learning Outcomes (ILOs) with respect to clarity, measurability, and cognitive depth.

This report presents the WP2 analysis. It describes the course corpus and the method used to examine ILO structure (structured vs. unstructured), quantity (ILOs per course), wording (length in words), and cognitive targeting (number and type of Bloom verbs). It then compares results across partners to identify consistent practices, misalignments, and coverage gaps that matter for teaching and assessment. The findings provide the evidence base for defining a standard ILO template and for prioritizing content areas in the forthcoming AI MOOCs, so that each module sets precise, assessable outcomes aligned with the skills learners actually need.

The University of Ljubljana (UNILJ) is the leading organization and coordinator for this work package.

2. Courses selection and syllabi collection

To underpin the courses and the intended learning outcome (ILO) analysis in Work Package 2, we first curated a corpus of university syllabi that explicitly address Artificial Intelligence (AI) and/or Machine Learning (ML) topics.

Project coordinator UNILJ provided a list of AI and ML-related topics that they deemed important and necessary for the latter parts of the TAI project that aims at developing AI courses for higher education.

2.1 Relevant topics in artificial intelligence and machine learning for education

Core concepts & learning paradigms

- Supervised learning (regression, classification)
- Unsupervised learning (clustering, pattern discovery)
- Reinforcement learning (interaction-based learning)
- Semi-supervised & self-supervised learning
- Transfer learning (domain adaptation)
- Online & incremental learning
- Hybrid AI (integration of symbolic and statistical methods)

Key algorithms and techniques

Supervised learning

- Linear regression & logistic regression
- Decision trees & random forests
- Support vector machines (SVMs)
- K-nearest neighbours (KNN)
- Gradient boosting (XGBoost, LightGBM, CatBoost)
- Neural networks (MLPs, CNNs, etc.)

Unsupervised learning

- K-means, DBSCAN, hierarchical clustering
- Principal component analysis (PCA)
- t-SNE (for dimensionality reduction and visualization)

Deep learning & sequential modelling

- Convolutional neural networks (CNNs) – for image-based tasks

- Recurrent neural networks (RNNs), LSTMs, GRUs – for time-series/sequential data
- Transformers – for advanced sequence modelling (optional for more advanced programs)

Time-series & control models

- Hidden Markov models (HMMs)
- Kalman filters / particle filters – for dynamic systems and estimation
- Markov decision processes – for modelling decision-making over time

Optimization & evolutionary techniques

- Genetic algorithms – for design optimization
- Swarm intelligence (PSO)
- Simulated annealing
- Constraint satisfaction & metaheuristics

Fuzzy logic & expert systems

- Fuzzy inference systems – control, automation, imprecise data
- Expert systems – rule-based reasoning in diagnostics and decision support
- Hybrid systems – e.g., neuro-fuzzy approaches

Model development & evaluation techniques

Data preparation

- Feature engineering & selection
- Data cleaning & imputation
- Outlier detection & removal
- Normalization and scaling
- Data augmentation (especially in vision tasks)

Model evaluation

- Train/test split
- K-fold cross-validation
- Confusion matrix, ROC/AUC, F1 score
- Sensitivity and uncertainty analysis

Model optimization

- Hyperparameter tuning (grid, random, Bayesian search)
- Regularization (L1, L2)
- Early stopping
- Ensembling (bagging, boosting, stacking)

Explainability & robustness

- SHAP & LIME – model interpretability
- Model explainability and trust
- Handling class imbalance (e.g., smote)
- Fairness and bias mitigation

Essential Python tools and libraries

Programming and data tools

- core Python (data structures, control flow)
- NumPy – numerical operations
- Pandas – data wrangling
- SciPy – scientific computing
- HDF5 / SQLite / MongoDB – data storage

ML/DL frameworks

- Scikit-learn – classical ML
- TensorFlow & keras – deep learning
- Pytorch – deep learning and research
- Stable Baselines – reinforcement learning
- XGBoost / lightGBM / CatBoost – boosting frameworks

Visualization

- matplotlib
- seaborn
- plotly (interactive dashboards)

Domain-specific libraries

- OpenCV – computer vision
- NLTK / SpaCy – NLP
- SimPy – system simulation
- PyFuzzy – fuzzy logic systems

Tools for deployment & workflow

- Flask / FastAPI – model serving via APIs
- Docker – containerization
- MLflow / DVC – experiment and model tracking
- AutoML tools (auto-sklearn, tpot) – automation of model design

Engineering-oriented applications of AI/ML

- predictive maintenance
- fault detection and diagnosis
- robotics and autonomous systems

- digital twins & simulation-based optimization
- intelligent control systems (fuzzy + AI integration)
- process optimization and quality control
- signal processing using ML
- human-machine interaction (adaptive interfaces)
- embedded and real-time AI (TinyML, Edge AI, FPGA-based inference)

Ethical, legal, and societal aspects

- AI ethics & fairness
- bias in data and models
- data privacy & security (e.g., GDPR)
- explainability in safety-critical applications
- sustainable AI (energy use, model size)

2.2 Collection protocol

Syllabi were collected at all five partner organisations and followed the collection protocols that were discussed and agreed at several project-coordination meetings.

1. Convenience sampling of a minimum of 10 courses per partner

- a. **Rationale:** Balance statistical breadth with feasibility; recognize that AI teaching density varies across disciplines.
- b. **Operational rule:** Mechanical-, electrical- and computer-engineering programmes plus mathematics/physics departments were flagged as the primary targets. Final inclusion decisions rested with partners' coordinators.

2. Institution-centred sampling

- a. **Rationale:** Keep results interpretable at partner level; avoid interinstitutional hurdles
- b. **Operational rule:** Each partner harvested syllabi from its own faculties first. If < 10 suitable courses could be found, the search could expand to other faculties within the same university (but not beyond).

3. Topic- relevance filter

- a. **Rationale:** Exclude tangential courses while respecting curricular idiosyncrasies
- b. **Operational rule:** A course qualified if any part of its content dealt with AI/ML concepts (topic list supplied by UNILJ). Local coordinators made the final call.

4. Minimum syllabus metadata

- a. **Rationale:** Enable cross-course comparison without over-burdening lecturers

- b. **Operational rule:** Every syllabus had to contain course content and ILOs; objectives, competencies and teaching methods were recommended.

5. Format and language constraints

- a. **Rationale:** Facilitate automated parsing and ensure archivability
- b. **Operational rule:** Documents were deposited as PDFs in a SharePoint repository. HTML pages were accepted only after export to PDF. English or local-language versions were both acceptable; no translation requirement.

6. Anonymity safeguards

- a. **Rationale:** Avoid disputes about perceived misinterpretations of individual courses.
- b. **Operational rule:** Public outputs report only aggregated statistics; individual courses, lecturers, and assessment details remain confidential.

2.3 Collection outcome

In total 94 AI/ML-related courses were captured following protocols explained above and the convenience sampling approach. Despite inevitable heterogeneity in structure - ranging from Bologna-style templates with Bloom-mapped ILOs to free-text outlines - the dataset offers a sufficiently rich, multi-disciplinary basis for the subsequent gap-analysis of AI competencies across partner institutions.

Table 1 presents the number of collected courses by partners' institutions. The **University of Ljubljana** supplies the largest tranche with thirty-five courses spread across its mechanical-, electrical- and computer-engineering curricula. **University of Belgrade** contributes twenty-one courses - most at bachelor and master level, with three drawn from its PhD programme - thereby adding breadth without letting research-specialist modules dominate the sample. Sixteen syllabi come from the **University of Rijeka**; all are home-grown offerings rather than material borrowed from other Croatian campuses. **The universities of Montenegro and Sarajevo** each add eleven courses, ensuring that the smaller national systems remain visible in the analysis. Every partner has therefore met the agreed minimum of ten courses as per *Collection criterion 1*, giving us a well-rounded foundation for the forthcoming examination of intended learning outcomes and topic coverage.

Institution	UNILJ	UNSA	UNBG	UoM	UNIRI	Total
No. of courses	35	11	21	11	16	94

Table 1: Collected courses by partner institutions

Regarding institutional departments, partners primarily drew on their own faculties unless too few courses were available (as per *Collection criterion 2*). Table 2 presents the faculties at the partner institutions where the collected courses are offered. Because academic disciplines are organised differently at each institution, faculty names and their associated domains vary substantially. Some faculties are mono-disciplinary, concentrating on a single scientific field, whereas others combine several major - or subsidiary - fields within one administrative unit. Faculty affiliation mirrors how each university clusters its engineering and science disciplines. The largest share of syllabi - 42 of 94 - comes from Faculties of Electrical Engineering. Every partner maintains such a unit except the University of Rijeka, where all engineering programmes are consolidated in a single, multi-disciplinary faculty. The next-largest contributor is the Faculty of Computer and Information Science (18 syllabi), represented exclusively by the University of Ljubljana. Faculties of Mechanical Engineering supply a further 11 syllabi drawn from Ljubljana, Sarajevo and Podgorica. The University of Rijeka's dedicated Faculty of Informatics and Digital Technologies adds 9 courses, while Faculties of Sciences and Mathematics (or stand-alone Mathematics faculties) account for the remaining 7 syllabi. This distribution underscores the partners' divergent organisational models: some segregate disciplines into specialised faculties, whereas others aggregate multiple engineering and science fields within broader administrative units.

Faculty	UNILJ	UNSA	UNBG	UoM	UNIRI	Total
Electrical Engineering	14	3	21	4		42
Computer and Information Science	18					18
Mechanical Engineering	3	4		4		11
Information and Digital Technologies					9	9
Sciences and Mathematics		4		3		7
Engineering					5	5
Mathematics					2	2

Table 2: Collected courses by faculties

Programme affiliation is even more heterogeneous than faculty affiliation. The 94 syllabi gathered from the five partner universities are distributed across twelve distinct study programmes (Table 3). Although the majority of courses remain embedded in a single programme, a growing share belongs to formally interdisciplinary curricula that allow

students from multiple faculties to enrol in the same AI/ML offering. To maintain consistency, our analysis cites only the primary programmes in which each course is listed. As some courses have more than one primary programmes, the sum of number in Table 3 is higher than the actual number of courses.

As with the faculty breakdown, Electrical-Engineering programmes dominate, accounting for 31 of the 94 courses. They are followed by the Computer- and Information-Science programme with 21 courses. Three further programmes—Mechanical Engineering, general Computer Science/Engineering, and Informatics—each contribute ten or more courses, underscoring the breadth of AI integration across engineering and computing domains. Additional diversity is introduced by isolated courses housed in Multimedia, Software Engineering, and Mathematics & Computer Science curricula, as well as single offerings from Mechatronics, Road-Traffic Engineering, and Mechatronics & Robotics programmes.

This wide programme spread ensures that the subsequent ILO and topic-coverage analyses are based on a dataset that captures both traditional disciplinary pathways and newer interdisciplinary trajectories in AI and machine learning education.

Study Programme	UNIJ	UNSA	UNBG	UoM	UNIRI	Total
Electrical Engineering	12		16	3		31
Computer and Informatics Science	18	3				21
Informatics			3		9	12
Computer Science / Computing		4		4	3	11
Mechanical Engineering	3	4		2	1	10
Software Engineering			4			4
Multimedia	4					4
Mathematics and Computer Science					2	2
Useful Statistics	2					2
Road Traffic				1		1
Mechatronics				1		1
Mechatronics and Robotics					1	1

Table 3: Collected courses by study programmes

A final stratification was performed by study cycle. Pedagogical expectations would place introductory AI/ML material in first-cycle (BSc) programmes, with progressively more specialised analysis, design and development topics reserved for second-cycle (MSc) and third-cycle (PhD) curricula. Table 4 confirms this pattern. Second-cycle programmes account for 56 of the 94 courses, reflecting the widespread adoption of advanced AI/ML electives at the master’s level. First-cycle programmes contribute 24 courses, typically covering foundational concepts. The remaining 14 courses are delivered within third-cycle doctoral programmes. The comparatively small PhD share is consistent with our convenience-sampling protocol: once the minimum quotas had been met with bachelor- and master-level syllabi, many interdisciplinary doctoral modules were excluded to keep the dataset tractable for subsequent analysis.

Study level / cycle	UNILJ	UNSA	UNBG	UoM	UNIRI	Total
1 st , BSc	8	4	8	1	3	24
2 nd , MSc	21	7	10	9	9	56
3 rd , PhD	6	0	3	1	4	14

Table 4: Collected courses by study levels (cycles)

3. Course syllabi analysis

In accordance with *Collection criterion 3*, every syllabus included in the dataset addresses artificial-intelligence and/or machine-learning topics to some extent. *Collection criterion 4* required that each document contain a minimum set of course metadata in order to be analysable. Despite variations in formatting and level of detail, all syllabi feature a common core of information:

- Course title
- Study programme
- Study cycle
- Lecturer
- Prerequisites
- Content (lectures or description)
- Literature / readings
- Objectives / goals
- Intended learning outcomes
- Learning and teaching methods
- Grading

For Work Package 2, the analysis focuses on two elements: course content and intended learning outcomes. The content descriptions - whether structured schedules or free-text summaries - enable us to catalogue the AI/ML topics currently taught. The ILOs reveal the competencies that students are expected to achieve. These two data streams define the baseline against which the consortium will design its forthcoming AI-MOOC modules, ensuring that the new materials are both comprehensive and complementary to existing offerings.

3.1 Course AI/ML-related topics

We analysed course content and mapped it to the AI and ML topics presented in the previous section. The selected topics covered AI- and ML-related fields grouped into six categories:

Topic category	Reason for inclusion in the analysis
Core concepts and learning paradigms	Establishes the conceptual foundation of AI/ML—e.g., supervised, unsupervised, reinforcement, transfer and self-supervised learning. Coverage here indicates whether

	students receive the theoretical framework required to understand later specialisations.
Key algorithms and techniques	Captures exposure to the classical and modern methods—decision trees, SVMs, neural networks, clustering, evolutionary optimisation, etc.—that form the practitioner’s toolbox. The presence of these algorithms signals the depth of methodological training.
Model development and evaluation techniques	Encompasses data preparation, feature engineering, hyper-parameter tuning, cross-validation, performance metrics and explainability tools. Courses that address this category equip students to build models that are reliable, reproducible and interpretable.
Essential Python tools and libraries	Python is the de-facto language of contemporary AI/ML. Mapping syllabi against libraries such as NumPy, pandas, scikit-learn, TensorFlow or PyTorch reveals how far curricula align with industry-standard workflows.
Engineering-oriented applications of AI/ML	Assesses whether programmes translate theory into domain-specific practice—e.g., robotics, computer vision, predictive maintenance, digital twins, embedded AI. This category is central to the TAI project’s goal of producing application-ready graduates.
Ethical, legal, and societal aspects	Reflects the growing requirement for responsible AI. Inclusion signals that students are taught to consider bias, fairness, privacy, sustainability and relevant regulatory frameworks—competencies essential for compliance with emerging EU and international guidelines.

Table 5: AI and ML-related topics

3.1.1 Core concepts and learning paradigms

Core concepts and learning paradigms category included supervised learning, unsupervised learning, reinforcement learning, semi-supervised and self-supervised learning, transfer learning, online and incremental learning, and hybrid AI.

Topic	Brief elaboration
-------	-------------------

Supervised learning	A paradigm in which a model is trained from pairs (x, y) , where x represents input data and y the ground-truth label. Typical tasks are regression (predicting continuous values) and classification (assigning discrete classes). Fundamental algorithms include linear and logistic regression, decision trees, k-nearest neighbours and neural networks.
Unsupervised learning	Learning structure from unlabelled data. The two canonical goals are clustering —grouping observations by similarity—and pattern discovery / density estimation . Representative techniques are k-means, hierarchical clustering, Gaussian mixture models, principal-component analysis and auto-encoders.
Reinforcement learning	An interaction-centric framework where an <i>agent</i> learns a policy π that maximises cumulative reward by acting in an environment. Core concepts include Markov-decision processes, value functions, policy gradients and actor–critic architectures. RL underpins robotics, game-playing agents and sequential decision systems.
Semi-supervised and self-supervised learning	<i>Semi-supervised</i> methods exploit small labelled and large unlabelled sets simultaneously (e.g., consistency regularisation, pseudo-labelling). <i>Self-supervised</i> learning formulates pretext tasks (e.g., masked-token prediction, contrastive objectives) that create supervisory signals from the data itself, enabling label-efficient pre-training of large models.
Transfer learning	Models trained on a <i>source</i> domain are adapted to a <i>target</i> domain with limited data. Strategies include fine-tuning deep networks, feature-based transfer, domain-adversarial training and model-agnostic meta-learning. Transfer learning is central to modern vision and language systems, reducing computational cost and data requirements.
Online and incremental learning	Algorithms that update their parameters sequentially as new data arrive, rather than retraining from scratch. Examples are stochastic-gradient descent, adaptive filters and streaming decision trees. Online learning is essential for non-stationary environments and real-time analytics.

<p>Hybrid AI</p>	<p>Architectures that integrate symbolic methods (logic, knowledge graphs, constraint solvers) with statistical learners (e.g., neural networks). The synergy aims to combine data-driven pattern recognition with explicit reasoning, yielding systems that are both performant and explainable. Representative approaches include neuro-symbolic networks, differentiable logic layers and probabilistic programming.</p>
-------------------------	---

Table 6: AI and ML-related core concepts and learning paradigms

3.1.2 Key algorithms and techniques

This category specifies the concrete methods that operationalise the learning paradigms outlined above. Mastery of these algorithms equips students to select, implement and critique AI/ML solutions across a wide range of tasks.

Topic	Algorithms and use-cases
<p>Supervised-learning methods</p>	<p>Algorithms:</p> <ul style="list-style-type: none"> • Linear & logistic regression – baseline models for continuous and binary outcomes; interpretable coefficients. • Decision trees & random forests – non-linear decision boundaries, built-in feature-importance metrics. • Support-vector machines (SVMs) – maximise margin; effective in high-dimensional spaces with kernel tricks. • K-nearest neighbours (KNN) – instance-based learner; no explicit training phase; sensitive to distance metric. • Gradient-boosting ensembles (XGBoost, LightGBM, CatBoost) – stage-wise additive trees; state-of-the-art on tabular data. • Neural networks (MLPs, CNNs etc.) – universal function approximators; scalable to large, unstructured data sets. <p>Typical use-cases: Prediction, classification, fault detection, recommendation, image recognition, structured-data analytics.</p>
<p>Unsupervised-learning methods</p>	<p>Algorithms:</p> <ul style="list-style-type: none"> • k-means, DBSCAN, hierarchical clustering – discover latent groupings; DBSCAN handles arbitrary shapes & noise.

	<ul style="list-style-type: none"> • Principal-component analysis (PCA) – linear dimensionality reduction; preserves maximal variance. • t-SNE / UMAP – non-linear embedding for visual inspection of high-dimensional spaces. <p>Typical use-cases: Market segmentation, anomaly detection, exploratory data analysis, feature compression.</p>
<p>Deep-learning and sequential modelling</p>	<p>Algorithms:</p> <ul style="list-style-type: none"> • Convolutional neural networks (CNNs) – weight sharing & local connectivity for images, speech, spatial data. • Recurrent nets (RNN, LSTM, GRU) – gated memory for time-series, language and sensor streams. • Transformers – self-attention architecture for long-range dependencies in text, audio and vision (advanced option). <p>Typical use-cases: Computer vision, natural language processing, speech recognition, multivariate time-series forecasting.</p>
<p>Time-series & control models</p>	<p>Algorithms:</p> <ul style="list-style-type: none"> • Hidden Markov models (HMMs) – generative models for discrete hidden states. • Kalman / particle filters – recursive Bayesian estimation in linear or non-linear dynamical systems. • Markov decision processes (MDPs) – formalism for sequential decision-making under uncertainty (foundation for RL). <p>Typical use-cases: Target tracking, sensor fusion, speech tagging, predictive maintenance, adaptive control.</p>
<p>Optimisation & evolutionary techniques</p>	<p>Algorithms:</p> <ul style="list-style-type: none"> • Genetic algorithms (GA) – population-based search via crossover/mutation. • Swarm intelligence (PSO, ant colony) – collective behaviour heuristics for continuous or combinatorial spaces. • Simulated annealing – probabilistic hill-climber with temperature schedule.

	<ul style="list-style-type: none"> • Constraint satisfaction & meta-heuristics – tabu search, GRASP, variable-neighbourhood search. <p>Typical use-cases: Engineering design, hyper-parameter tuning, scheduling, network layout, feature selection.</p>
Fuzzy logic & expert systems	<p>Algorithms:</p> <ul style="list-style-type: none"> • Fuzzy inference systems (Mamdani, Sugeno) – reason with linguistic variables and imprecise data. • Rule-based expert systems – forward/backward chaining for diagnostic or advisory tasks. • Hybrid (neuro-fuzzy) systems – combine fuzzy rules with trainable neural parameters. <p>Typical use-cases: Industrial control, HVAC, medical diagnosis decision support, automotive systems, knowledge-based controllers.</p>

Table 7: AI and ML-related key algorithms and techniques

Together these algorithms constitute the operational core of the TAI competency framework: they turn learning paradigms into deployable models, support both data-driven and knowledge-driven reasoning, and provide the basis for advanced application domains addressed in subsequent categories.

3.1.3 Model development and evaluation techniques

Robust AI systems require more than algorithm selection; they depend on disciplined data preparation, rigorous validation, systematic optimisation and transparent interpretation. This category captures the tool-chain that transforms raw data and base models into dependable, reproducible solutions.

Topic	Practices, tools, and purpose
Data preparation	<p>Practices and tools:</p> <ul style="list-style-type: none"> • Feature engineering & selection – deriving informative variables (one-hot encoding, interaction terms, embeddings) and pruning redundant ones (filter / wrapper methods). • Data cleaning & imputation – resolving missing, duplicate or inconsistent records with statistical and model-based imputers.

	<ul style="list-style-type: none"> • Outlier detection & removal – distance-, density- or model-based flagging of anomalous samples. • Normalisation & scaling – z-score, min–max, log or Box-Cox transforms to stabilise optimisation. • Data augmentation – synthetic transformations (flip, crop, noise, mix up) to enlarge training sets, mainly in vision and audio. <p><u>Purpose and typical outcomes:</u> Ensures that downstream models receive informative, well-conditioned input and limits overfitting to idiosyncrasies of the training data.</p>
<p>Model evaluation</p>	<p><u>Practices and tools:</u></p> <ul style="list-style-type: none"> • Train/test split – hold-out validation for first-pass performance estimation. • k-fold cross-validation – variance-reduced metrics and model selection on limited data. • Confusion matrix, ROC/AUC, precision-recall, F1 – task-appropriate performance indicators. • Sensitivity / uncertainty analysis – Monte-Carlo dropout, perturbation tests, prediction intervals. <p><u>Purpose and typical outcomes:</u> Provides unbiased estimates of generalisation performance and quantifies uncertainty—prerequisites for deployment in risk-sensitive domains.</p>
<p>Model optimisation</p>	<p><u>Practices and tools:</u></p> <ul style="list-style-type: none"> • Hyper-parameter tuning – grid, random, Bayesian (e.g., Hyperopt, Optuna) or evolutionary searches. • Regularisation – L1/L2 penalties, dropout, weight decay to control model complexity. • Early stopping – halting training on a validation plateau to prevent over-fit. • Ensembling – bagging, boosting, stacking for variance and bias reduction. <p><u>Purpose and typical outcomes:</u> Improves predictive accuracy and robustness while mitigating over-fitting and computational cost.</p>

<p>Explainability and robustness</p>	<p><u>Practices and tools:</u></p> <ul style="list-style-type: none"> • SHAP, LIME, counterfactual explanations – local and global interpretation of feature influence. • Class-imbalance handling – resampling, SMOTE, cost-sensitive loss. • Fairness & bias mitigation – disparate-impact metrics, reweighing, adversarial debiasing. • Adversarial / stress testing – FGSM, PGD, noise injections to probe model stability. <p><u>Purpose and typical outcomes:</u></p> <p>Facilitates regulatory compliance, stakeholder trust and ethical deployment, addressing EU AI Act and similar governance frameworks.</p>
---	---

Table 8: AI and ML-related model development and evaluation techniques

Courses that cover these techniques provide learners with the full life-cycle skill-set required to build, assess and harden AI/ML models in real-world settings.

3.1.4 Essential Python tools and libraries

A contemporary AI/ML curriculum must equip students with the software stack that dominates both research and industrial practice. The tools listed below cover the full pipeline - from exploratory coding and data storage to large-scale model training, interactive visualisation and deployment.

Topic	Representative packages and their roles in the end-to-end workflow
<p>Programming & data-handling tools</p>	<p><u>Python packages:</u></p> <p>Python 3 language features (data structures, control flow, comprehensions)</p> <p>NumPy for vectorised computation</p> <p>pandas for tabular data wrangling</p> <p>SciPy for numerical linear algebra, optimisation, signal processing</p> <p>HDF5, SQLite, MongoDB connectors for structured and semi-structured storage</p>

	<p><u>Role in the workflow:</u> Provide a high-performance yet accessible environment in which to ingest, clean and transform data prior to modelling.</p>
<p>Machine-learning / deep-learning frameworks</p>	<p><u>Python packages:</u> <code>scikit-learn</code> for classical algorithms and model-selection utilities <code>TensorFlow + Keras</code> and <code>PyTorch</code> for GPU-accelerated deep networks <code>Stable-Baselines3</code> for off-the-shelf reinforcement-learning agents <code>XGBoost</code>, <code>LightGBM</code>, <code>CatBoost</code> for gradient-boosted decision-tree ensembles</p> <p><u>Role in the workflow:</u> Offer optimised, production-ready implementations of core algorithms, enabling rapid prototyping and reproducible experiments.</p>
<p>Visualisation and exploratory analysis</p>	<p><u>Python packages:</u> <code>matplotlib</code> as the low-level plotting backbone <code>seaborn</code> for statistical graphics <code>plotly / Dash</code> for interactive dashboards</p> <p><u>Role in the workflow:</u> Support diagnostic plotting, hyper-parameter exploration and stakeholder communication.</p>
<p>Domain-specific libraries</p>	<p><u>Python packages:</u> <code>OpenCV</code> for image/video I/O, filtering and feature extraction <code>NLTK</code>, <code>spaCy</code>, <code>transformers</code> for natural-language processing <code>SimPy</code> for discrete-event simulation of industrial systems <code>PyFuzzy</code>, <code>scikit-fuzzy</code> for fuzzy-logic controllers</p> <p><u>Role in the workflow:</u> Allow application projects—vision, NLP, simulation, control—to be tackled without re-implementing domain primitives.</p>
<p>Deployment & workflow tools</p>	<p><u>Python packages:</u> <code>Flask</code>, <code>FastAPI</code> for REST/GRPC model serving <code>Docker</code> (and <code>docker-compose</code>) for environment isolation and reproducibility <code>MLflow</code>, <code>DVC</code> for experiment tracking, model registry and data</p>

	<p>versioning</p> <p>AutoML frameworks (<code>auto-sklearn</code>, <code>TPOT</code>, <code>FLAML</code>) for automated pipeline search</p> <p>Role in the workflow:</p> <p>Bridge the gap from research code to maintainable, monitored services in production environments.</p>
--	--

Table 9: Essential Python tools and libraries

3.1.5 Engineering-oriented applications of AI/ML

Topic	Typical objectives and representative AI/ML techniques
Predictive maintenance	<p>Objectives:</p> <p>Forecast remaining useful life (RUL) of equipment; schedule maintenance before failure.</p> <p>AI/ML techniques:</p> <p>Degradation modelling, survival analysis, LSTM/transformer RUL predictors, anomaly detection in sensor streams.</p>
Fault detection & diagnosis	<p>Objectives:</p> <p>Identify abnormal operating states and isolate root causes.</p> <p>AI/ML techniques:</p> <p>Auto-encoders, one-class SVMs, Bayesian networks, rule-based expert systems, hybrid neuro-fuzzy classifiers.</p>
Robotics & autonomous systems	<p>Objectives:</p> <p>Perception, localisation, motion planning and control of mobile or articulated robots.</p> <p>AI/ML techniques:</p> <p>CNN-based vision, SLAM, reinforcement learning for policy optimisation, model-predictive control with learned dynamics.</p>
Digital twins & simulation-based optimisation	<p>Objectives:</p> <p>Maintain a real-time virtual replica of physical assets; optimise design or operation via what-if scenarios.</p> <p>AI/ML techniques:</p> <p>Physics-informed neural networks, surrogate models for CFD/FEM, evolutionary or Bayesian optimisation loops</p>
Intelligent control systems	<p>Objectives:</p>

	<p>Combine classical control with AI to handle non-linear, uncertain or poorly modelled plants.</p> <p><u>AI/ML techniques:</u> Fuzzy inference, adaptive neuro-fuzzy controllers (ANFIS), reinforcement-learning controllers, gain-scheduling with ML predictors.</p>
Process optimisation & quality control	<p><u>Objectives:</u> Maximise throughput, minimise energy use, ensure conformance to quality specs.</p> <p><u>AI/ML techniques:</u> Multi-objective evolutionary algorithms, gradient-boosted regressors for soft-sensing, SPC with ML-driven anomaly flags.</p>
Signal processing using ML	<p><u>Objectives:</u> Extract information from audio, vibration, radar or biomedical signals.</p> <p><u>AI/ML techniques:</u> 1-D CNNs, wavelet-domain features, RNN/transformer sequence models, blind-source separation with non-negative matrix factorisation.</p>
Human-machine interaction (HMI)	<p><u>Objectives:</u> Adapt interfaces to user intent, context or emotion; improve usability and safety.</p> <p><u>AI/ML techniques:</u> Eye-gaze and gesture recognition, affective computing (emotion classification), adaptive dialogue managers, reinforcement-learning UI personalisation.</p>
Embedded & real-time AI	<p><u>Objectives:</u> Deploy inference on resource-constrained or latency-critical platforms (MCUs, FPGAs).</p> <p><u>AI/ML techniques:</u> Model compression (pruning, quantisation), TinyML tool-chains (TensorFlow Lite Micro, Edge Impulse), hardware-aware NAS, high-level synthesis for FPGA accelerators.</p>

Table 10: Engineering-oriented applications of AI/ML

Including these domains in the mapping allows the TAI project to assess not only theoretical coverage but also how effectively current curricula translate AI/ML methods into engineering practice.

3.1.6 Ethical, legal, and societal aspects

Topic	Questions addressed, illustrative methods, standards, or regulations
<p>AI ethics & fairness</p>	<p><u>Question addressed:</u> Do autonomous systems act in accordance with accepted moral principles and distributive justice? Are outcomes equitable across demographic groups?</p> <p><u>Methods, standards, regulations:</u> Ethical-impact assessment, value-sensitive design, equality of opportunity / demographic-parity metrics, IEEE 7000 and ISO/IEC 24028 guidance.</p>
<p>Bias in data and models</p>	<p><u>Question addressed:</u> Are training sets or algorithms producing systemic error patterns that disadvantage protected groups or distort scientific inference?</p> <p><u>Methods, standards, regulations:</u> Bias audits, dataset documentation (“datasheets for datasets”), disparate-impact testing, counterfactual fairness, model debiasing via adversarial training or re-weighting.</p>
<p>Data privacy & security</p>	<p><u>Question addressed:</u> Does model development and deployment comply with GDPR and related data-protection statutes? Are models resilient to malicious manipulation?</p> <p><u>Methods, standards, regulations:</u> Privacy-by-design, differential privacy, federated learning, secure multi-party computation, model watermarking; threat modelling against data-poisoning and adversarial attacks.</p>
<p>Explainability in safety-critical applications</p>	<p><u>Question addressed:</u> Can model decisions be interpreted by domain experts and regulators when human life, health or large financial stakes are involved?</p> <p><u>Methods, standards, regulations:</u></p>

	Post-hoc tools (SHAP, LIME, counterfactuals), inherently interpretable models (GAMs, monotonic networks), model-card reporting, EU AI Act risk-classification compliance.
Sustainable AI	<p><u>Question addressed:</u> What is the environmental footprint of training, deploying and maintaining models, and how can it be reduced?</p> <p><u>Methods, standards, regulations:</u> Energy-aware architecture search, carbon-tracking of experiments, model compression, hardware-efficient inference (Edge AI), alignment with ISO 14064 and corporate carbon-accounting frameworks.</p>

Table 11: Ethical, legal, and societal aspects of AI/ML

Integrating these themes into curriculum mapping ensures that technical competence is accompanied by an understanding of the broader responsibilities and constraints governing modern AI practice.

3.2 Course content analysis

We mapped the course content against the selected AI/ML-related topics to determine how well the current courses cover these topics. This approach let us create a detailed, institution-by-institution report and later compare the current state of AI/ML education across institutions.

The final categorization and structure of the topics considered 6 categories and 33 topics as follows:

Category	Topics
Core concepts and learning paradigms	<ul style="list-style-type: none"> • Supervised Learning • Unsupervised Learning • Reinforcement Learning • Semi-supervised & Self-supervised Learning • Transfer Learning • Online & Incremental Learning • Hybrid AI

Key algorithms and techniques	<ul style="list-style-type: none"> • Supervised learning algorithms and techniques • Unsupervised learning algorithms and techniques • Deep learning & sequential modelling • Time-series & control models • Optimization & evolutionary techniques • Fuzzy logic & expert systems
Model development & evaluation techniques	<ul style="list-style-type: none"> • Data preparation • Model evaluation • Model optimization • Explainability and robustness
Essential python tools and libraries	<ul style="list-style-type: none"> • Programming and data tools • Machine learning / Deep learning frameworks • Visualization • Domain-specific libraries • Tools for development and workflow
Applications of AI/ML	<ul style="list-style-type: none"> • Predictive maintenance • Fault detection and diagnosis • Robotics and autonomous systems • Digital twins and simulation-based optimization • Intelligent control systems • Process optimization and quality control • Signal processing using ML • Human-machine interaction • Embedded and real-time AI
Ethical, legal, and societal aspects	<ul style="list-style-type: none"> • Ethical aspects • Legal and societal aspects

Table 12: AI/ML related categories and topics for course content analysis

3.2.1 Institutional reports

University of Ljubljana (UNILJ)

The synthesis covers 35 undergraduate, master's and doctoral courses offered across the Faculty of Electrical Engineering, Faculty of Mechanical Engineering and Faculty of Computer and Information Science. Table 13 presents the number and percentage of UNILJ courses that cover a specific AI/ML-related topic.

AI topics	UNILJ courses that cover the topic	Share of all UNILJ courses
Core concepts & learning paradigms		
Supervised Learning	33	0.94
Unsupervised Learning	18	0.51
Reinforcement Learning	9	0.26
Semi- & Self-supervised Learning	2	0.06
Transfer Learning	1	0.03
Online & Incremental Learning	2	0.06
Hybrid AI	9	0.26
Key AI algorithms and techniques		
Supervised learning algorithms & techniques	33	0.94
Unsupervised learning algorithms & techniques	18	0.51
Deep learning & sequential modelling	16	0.46
Time-series & control models	14	0.40
Optimization & evolutionary techniques	7	0.20
Fuzzy logic & expert systems	8	0.23
Model development & evaluation techniques		
Data preparation	23	0.66
Model evaluation	24	0.69
Model optimization	14	0.40
Explainability & robustness	10	0.29
Essential Python tools and libraries		
Programming and data tools	23	0.66
Machine learning / Deep-learning frameworks	8	0.23
Visualization	5	0.14
Domain-specific libraries	13	0.37
Tools for deployment & workflow	9	0.26
Applications of AI in engineering		
Predictive Maintenance	0	0.00
Fault Detection & Diagnosis	2	0.06
Robotics & Autonomous Systems	7	0.20
Digital Twins & Simulation-Based Optimization	1	0.03
Intelligent Control Systems	6	0.17
Process Optimization & Quality Control	3	0.09

Signal Processing using ML	5	0.14
Human-Machine Interaction	3	0.09
Embedded & Real-Time AI	8	0.23
Ethical, legal and societal aspects		
Ethical aspects	7	0.20
Legal and societal aspects	4	0.11

Table 13: UNILJ courses mapping against selected AI/ML topics

The topic mapping reveals a pronounced hierarchy of coverage. Supervised-learning fundamentals—regression, classification and the standard algorithmic toolkit—are effectively universal, appearing in 94 % of courses. Second-tier competencies in data preparation and formal model evaluation are also well represented (66–69 %), indicating that most programmes emphasise sound experimental practice. Coverage then drops to 40–51 % for advanced material such as deep-learning architectures, time-series and control models, optimisation strategies and unsupervised learning; many curricula therefore still rely on “shallow” ML methods. Only about a quarter of courses (20–29 %) address reinforcement learning, hybrid or evolutionary techniques, model explainability, robustness or MLOps tooling. Topics that are increasingly critical for deployment—transfer and self-supervised learning, online or incremental updating, predictive-maintenance and digital-twin applications, as well as ethical, legal and societal considerations—are found in ≤ 11 % of syllabi.

Areas of excellence

- **Classical ML foundations.** Students receive a thorough grounding in supervised methods, evaluation metrics and optimisation mathematics.
- **Data-handling skills.** Approximately two-thirds of courses include data cleaning, feature engineering and cross-validation.
- **Early exposure.** Introductory AI/ML modules in first-year Computer-Science programmes provide baseline literacy at the BSc level.

Concentration of teaching effort

- **Supervised learning predominates** (33 of 35 courses).
- **Python tooling is widespread** but generally limited to NumPy, pandas and scikit-learn; only eight courses (~23 %) employ modern deep-learning frameworks such as PyTorch or TensorFlow.

- **Application coverage is uneven;** robotics is reasonably well served, whereas predictive maintenance, digital-twin optimisation and signal-processing applications are seldom addressed.

Conclusion

In summary, the University of Ljubljana offers a dense, technically rigorous core in classical machine-learning: supervised learning, data preparation and model evaluation are embedded in nearly every AI-related course. To align with industry needs and emerging regulations, the curriculum should expand toward advanced learning paradigms, deployment engineering, responsible-AI principles and sector-specific applications. Closing these gaps will better equip graduates to design, deploy and govern AI systems responsibly and at scale.

University of Sarajevo (UNSA)

The University of Sarajevo contributes eleven AI- and ML-related syllabi to the consortium dataset; their coverage of the selected AI/ML topics is summarised in Table 14.

AI topics	UNSA courses that cover the topic	Share of all UNSA courses
Core concepts & learning paradigms		
Supervised Learning	7	0.64
Unsupervised Learning	4	0.36
Reinforcement Learning	0	0.00
Semi- & Self-supervised Learning	0	0.00
Transfer Learning	0	0.00
Online & Incremental Learning	0	0.00
Hybrid AI	4	0.36
Key AI algorithms and techniques		
Supervised learning algorithms & techniques	7	0.64
Unsupervised learning algorithms & techniques	4	0.36
Deep learning & sequential modelling	5	0.45
Time-series & control models	0	0.00
Optimization & evolutionary techniques	3	0.27
Fuzzy logic & expert systems	5	0.45
Model development & evaluation techniques		

Data preparation	8	0.73
Model evaluation	7	0.64
Model optimization	2	0.18
Explainability & robustness	0	0.00
Essential Python tools and libraries		
Programming and data tools	7	0.63
Machine learning / Deep-learning frameworks	6	0.55
Visualization	2	0.18
Domain-specific libraries	4	0.36
Tools for deployment & workflow	0	0.00
Applications of AI in engineering		
Predictive Maintenance	0	0.00
Fault Detection & Diagnosis	0	0.00
Robotics & Autonomous Systems	0	0.00
Digital Twins & Simulation-Based Optimization	0	0.00
Intelligent Control Systems	2	0.18
Process Optimization & Quality Control	0	0.00
Signal Processing using ML	0	0.00
Human-Machine Interaction	0	0.00
Embedded & Real-Time AI	0	0.00
Ethical, legal and societal aspects		
Ethical aspects	0	0.00
Legal and societal aspects	0	0.00

Table 14: UNSA courses mapping against selected AI/ML topics

At the foundational level, **supervised learning—together with its classical algorithms—appears in 7 of the 11 courses (64 %)**, and unsupervised learning is present in four courses (36 %). Hybrid AI notions (typically fuzzy–neural combinations) are covered to the same 36 % extent. Deep-learning and other sequential-modelling techniques reach five courses (45 %), whereas reinforcement learning and transfer-, semi- or self-supervised paradigms are absent.

Development practice is reasonably well represented. **Data-preparation procedures are included in 8 syllabi (73 %)**, and formal model-evaluation routines in seven (64 %). However, systematic model optimisation appears in only two courses (18 %), and no course addresses explainability or robustness.

Tool support shows a similar pattern: core Python, NumPy/pandas and related data-handling libraries feature in seven syllabi (64 %); six courses (55 %) adopt mainstream ML/DL frameworks such as scikit-learn or TensorFlow. Visualisation and domain-specific libraries are touched upon in two to four courses, yet no syllabus introduces deployment-workflow tools (e.g., Docker, MLflow).

Application coverage is limited. Apart from two offerings (18 %) on intelligent control systems, none of the engineering-oriented verticals—predictive maintenance, fault diagnosis, robotics, digital twins, signal processing or embedded AI—are represented. Likewise, **ethical, legal and societal dimensions of AI are absent from all eleven syllabi.**

Areas of excellence

- Classical ML foundations – supervised learning and core algorithms covered in 64 % of courses.
- Strong data-handling practice – data preparation appears in 73 % of syllabi; model-evaluation routines in 64 %.
- Solid Python tooling – NumPy/pandas programming in 64 % of courses; mainstream ML/DL frameworks in 55 %.

Concentration of teaching effort

- Supervised learning dominates (7 / 11 courses).
- Deep learning present but not dominant (5 courses, 45 %).
- Minimal attention to model optimisation (18 %), visualisation (18 %) and zero coverage of deployment workflows.
- Application scope narrow: only intelligent-control systems (18 %); no courses on robotics, predictive maintenance, digital twins, signal processing or embedded AI.
- Responsible-AI themes absent: no syllabus covers ethics, fairness, privacy or legal aspects.

Conclusion

In summary, UNSA provides solid exposure to classical supervised learning, data preparation and standard Python tooling, with a growing but still partial inclusion of deep-learning content. The portfolio would benefit from expansion into modern learning paradigms, deployment workflows, industrial application scenarios and responsible-AI principles to align more closely with emerging professional requirements.

University of Belgrade (UNBG)

The University of Belgrade’s 21 AI- and ML-related syllabi display a clear hierarchy of topic coverage (Table 15). Virtually every course (95 %) addresses supervised-learning fundamentals together with their standard algorithms, and over four-fifths (81 %) include data-preparation procedures. Formal model-evaluation techniques—train/test splits, cross-validation and metric reporting—appear in two-thirds of syllabi, signalling solid experimental practice. Coverage then falls to mid-range levels for more advanced material: unsupervised learning is present in 57 % of courses, deep-learning and sequential models in 33 %, and model-optimisation strategies in 38 %. Reinforcement learning, hybrid AI and evolutionary optimisation remain niche (19–29 %), and only a quarter of courses expose students to deployment workflows such as Docker or MLflow. Industrial application domains are lightly represented: signal-processing use-cases surface in four courses, while predictive maintenance, digital twins and quality control are absent. No syllabus tackles the ethical, legal or societal dimensions of AI.

AI topics	UNBG courses that cover the topic	Share of all Unable BG courses
Core concepts & learning paradigms		
Supervised Learning	20	0.95
Unsupervised Learning	12	0.57
Reinforcement Learning	4	0.19
Semi- & Self-supervised Learning	2	0.10
Transfer Learning	1	0.05
Online & Incremental Learning	0	0.00
Hybrid AI	6	0.29
Key AI algorithms and techniques		
Supervised learning algorithms & techniques	20	0.95
Unsupervised learning algorithms & techniques	11	0.52
Deep learning & sequential modelling	7	0.33
Time-series & control models	2	0.10
Optimization & evolutionary techniques	3	0.14
Fuzzy logic & expert systems	2	0.10
Model development & evaluation techniques		
Data preparation	17	0.81
Model evaluation	14	0.67
Model optimization	8	0.38
Explainability & robustness	6	0.29
Essential Python tools and libraries		
Programming and data tools	16	0.76

Machine learning / Deep-learning frameworks	9	0.43
Visualization	6	0.29
Domain-specific libraries	8	0.38
Tools for deployment & workflow	5	0.24
Applications of AI in engineering		
Predictive Maintenance	0	0.00
Fault Detection & Diagnosis	1	0.05
Robotics & Autonomous Systems	1	0.05
Digital Twins & Simulation-Based Optimization	0	0.00
Intelligent Control Systems	1	0.05
Process Optimization & Quality Control	0	0.00
Signal Processing using ML	4	0.19
Human-Machine Interaction	1	0.05
Embedded & Real-Time AI	3	0.14
Ethical, legal and societal aspects		
Ethical aspects	0	0.00
Legal and societal aspects	0	0.00

Table 15: UNBG courses mapping against selected AI/ML topics

Areas of excellence

- Classical ML foundations - supervised learning and core algorithms in 95 % of courses.
- Strong data-handling practice - data preparation (81 %) and model evaluation (67 %) are widely covered.
- Broad Python fluency - NumPy/pandas programming in 76 % of syllabi; major ML/DL frameworks in 43 %.

Concentration of teaching effort

- Supervised learning dominates (20 / 21 courses); unsupervised learning is secondary, and modern paradigms such as reinforcement or transfer learning are marginal.
- Only one course in three includes deep-learning architectures; few discuss time-series or control models.
- Deployment engineering and MLOps appear sporadically, and responsible-AI themes are entirely absent.
- Application focus is narrow—signal processing and a handful of embedded-AI or intelligent-control topics; manufacturing-oriented areas such as predictive maintenance or digital twins receive no attention.

Conclusion

UNBG offers a robust core in classical machine learning and data-centric practice, but coverage of modern paradigms, deployment engineering, industrial applications and responsible-AI considerations is limited. Addressing these gaps will better align the curriculum with current professional standards and regulatory expectations.

University of Montenegro (UoM)

The analysis covers eleven AI- and ML-related syllabi drawn from three faculties at the University of Montenegro: Electrical Engineering, Mechanical Engineering and Science and Maths Studies. Table 16 summarizes the mapping to select topics.

AI topics	UoM courses that cover the topic	Share of all UoM courses
Core concepts & learning paradigms		
Supervised Learning	8	0.73
Unsupervised Learning	2	0.18
Reinforcement Learning	1	0.09
Semi- & Self-supervised Learning	0	0.00
Transfer Learning	0	0.00
Online & Incremental Learning	1	0.09
Hybrid AI	3	0.27
Key AI algorithms and techniques		
Supervised learning algorithms & techniques	8	0.73
Unsupervised learning algorithms & techniques	2	0.18
Deep learning & sequential modelling	3	0.27
Time-series & control models	3	0.27
Optimization & evolutionary techniques	2	0.18
Fuzzy logic & expert systems	2	0.18
Model development & evaluation techniques		
Data preparation	5	0.45
Model evaluation	4	0.36
Model optimization	4	0.36
Explainability & robustness	1	0.09
Essential Python tools and libraries		
Programming and data tools	7	0.64
Machine learning / Deep-learning frameworks	3	0.27
Visualization	2	0.18
Domain-specific libraries	1	0.09
Tools for deployment & workflow	0	0.00
Applications of AI in engineering		

Predictive Maintenance	0	0.00
Fault Detection & Diagnosis	0	0.00
Robotics & Autonomous Systems	1	0.09
Digital Twins & Simulation-Based Optimization	0	0.00
Intelligent Control Systems	4	0.36
Process Optimization & Quality Control	0	0.00
Signal Processing using ML	2	0.18
Human-Machine Interaction	2	0.18
Embedded & Real-Time AI	6	0.55
Ethical, legal and societal aspects		
Ethical aspects	1	0.09
Legal and societal aspects	0	0.00

Table 16: UoM courses mapping against selected AI/ML topics

Foundational material is well represented: **supervised-learning concepts and their standard algorithms appear in 73 % of courses (8 / 11)**, whereas unsupervised learning is covered in two courses (18 %). Reinforcement learning and online/incremental learning surface only once each (9 %), and no syllabus addresses transfer or self-supervised paradigms. Hybrid AI—primarily fuzzy–neural combinations—appears in three courses (27 %).

Among key techniques, deep-learning and other sequential models feature in three syllabi (27 %), closely mirrored by coverage of time-series/control models. Evolutionary optimisation and fuzzy-logic systems are each present in roughly one course in five.

Core development practice is uneven: data-preparation routines appear in five syllabi (45 %) and formal model-evaluation procedures in four (36 %), with an identical share treating model optimisation. Explainability and robustness are touched on in only one course (9 %).

Python programming with NumPy/pandas is common (64 %), yet only three courses (27 %) use mainstream ML/DL frameworks, and no syllabus introduces deployment-workflow tools such as Docker or MLflow.

Application coverage favours **embedded and real-time AI**, present in six courses (55 %), and **intelligent control systems** (36 %). Robotics, signal processing, and HMI topics each appear in one or two syllabi, while predictive maintenance, digital twins and process-quality control are absent. Ethical aspects are acknowledged in a single course; legal and societal dimensions are not addressed.

Areas of excellence

- **Classical supervised learning** – covered in 73 % of syllabi, ensuring a solid theoretical base.
- **Embedded & real-time AI** – more than half the courses tackle deployment on constrained hardware, reflecting institutional strengths in embedded systems.
- **Intelligent control** – 36 % of offerings integrate AI with control-engineering principles.

Concentration of teaching effort

- Instruction centres on supervised methods and embedded/real-time deployment.
- Python data-handling skills are widespread, but exposure to modern ML/DL frameworks and MLOps tooling is limited.
- Advanced paradigms (transfer/self-supervised learning, reinforcement learning, explainability) receive minimal attention.
- Industry-oriented applications such as predictive maintenance and digital-twin optimisation are not yet represented.
- Responsible-AI coverage is embryonic; ethical considerations appear in only one syllabus, and legal/societal aspects are absent.

Conclusion

The University of Montenegro delivers a focused portfolio that blends classical machine learning with embedded and control-oriented applications—an alignment well suited to its engineering context. To meet evolving professional standards, the curriculum would benefit from deeper coverage of modern learning paradigms, deployment workflows, industrial verticals (e.g., predictive maintenance) and comprehensive responsible-AI training.

University of Rijeka (UNIRI)

The sample comprises sixteen AI- and ML-related syllabi drawn from the Faculties of Informatics and Digital Technologies, Engineering, and Mathematics. Each syllabus was mapped to the project’s 33 reference topics (Table 17).

AI topics	UNIRI courses that cover the topic	Share of all UNIRI courses
Core concepts & learning paradigms		
Supervised Learning	15	0.94
Unsupervised Learning	10	0.63
Reinforcement Learning	5	0.31

Semi- & Self-supervised Learning	3	0.19
Transfer Learning	3	0.19
Online & Incremental Learning	2	0.13
Hybrid AI	2	0.13
Key AI algorithms and techniques		
Supervised learning algorithms & techniques	15	0.94
Unsupervised learning algorithms & techniques	10	0.63
Deep learning & sequential modelling	8	0.50
Time-series & control models	4	0.25
Optimization & evolutionary techniques	1	0.06
Fuzzy logic & expert systems	2	0.13
Model development & evaluation techniques		
Data preparation	12	0.75
Model evaluation	11	0.69
Model optimization	8	0.50
Explainability & robustness	4	0.25
Essential Python tools and libraries		
Programming and data tools	12	0.75
Machine learning / Deep-learning frameworks	5	0.31
Visualization	3	0.19
Domain-specific libraries	7	0.44
Tools for deployment & workflow	2	0.13
Applications of AI in engineering		
Predictive Maintenance	0	0.00
Fault Detection & Diagnosis	0	0.00
Robotics & Autonomous Systems	0	0.00
Digital Twins & Simulation-Based Optimization	0	0.00
Intelligent Control Systems	1	0.06
Process Optimization & Quality Control	0	0.00
Signal Processing using ML	0	0.00
Human-Machine Interaction	0	0.00
Embedded & Real-Time AI	2	0.13
Ethical, legal and societal aspects		
Ethical aspects	5	0.31
Legal and societal aspects	5	0.31

Table 17: UNIRI courses mapping against selected AI/ML topics

Foundational material is extensive: **supervised-learning concepts and their standard algorithms occur in 94 % of courses (15 / 16)**, while unsupervised learning appears in 63 %. Reinforcement learning is treated in five syllabi (31 %), and both transfer and semi/self-

supervised learning in three (19 %); online or incremental learning is mentioned twice (13 %). Hybrid AI—typically fuzzy–neural systems—surfaces in two syllabi.

Among key techniques, **deep-learning and sequential models are covered in half of the syllabi (50 %)**, and time-series/control models in 25 %. Evolutionary optimisation and fuzzy-logic systems are niche, each present in ≤ 13 %.

Core development practice is strong: data preparation is taught in 75 % of courses, model evaluation in 69 % and model optimisation in 50 %. Explainability and robustness appear in four syllabi (25 %).

Tool coverage is mixed. **Python programming with NumPy/pandas features in 75 % of courses**, yet only five syllabi (31 %) employ major ML/DL frameworks. Visualisation libraries and domain-specific packages (e.g., OpenCV, NLP tool-kits) appear in 19 % and 44 % of courses respectively. Deployment/MLOps tools are introduced in just two syllabi (13 %).

Application-oriented content is limited. Intelligent-control systems appear once (6 %), embedded/real-time AI in two courses (13 %); other industrial verticals—predictive maintenance, robotics, digital twins, quality control—are absent. Ethical, legal and societal aspects are addressed in five courses (31 %), the highest share among partner institutions.

Areas of excellence

- **Classical ML foundation** – supervised learning and its algorithmic toolbox covered in 94 % of syllabi.
- **Data-centric practice** – data preparation (75 %) and model evaluation (69 %) are consistently included.
- **Attention to responsible AI** – ethical and legal considerations embedded in nearly one-third of courses.

Concentration of teaching effort

- Emphasis on supervised and unsupervised methods; reinforcement-learning and transfer/self-supervised paradigms are emerging but not yet mainstream.
- Half of the courses incorporate deep-learning architectures, yet only one in three uses contemporary ML/DL frameworks and only one in eight addresses deployment workflows.
- Industrial applications are sparse: intelligent control and embedded AI receive modest coverage; other engineering verticals are absent.

Conclusion

AI/ML teaching at the University of Rijeka is anchored in strong classical-ML and data-handling foundations and displays a comparatively high awareness of ethical and legal issues. Curriculum gaps remain in modern learning paradigms, MLOps engineering and industry-specific applications such as predictive maintenance or digital-twin optimisation. Bridging these gaps would move the programme portfolio closer to current professional practice.

3.2.2 Cross-institution analysis

The consolidated matrix comprises 94 syllabi: 35 from UNILJ, 21 from UNBG, 16 from UNIRI, and 11 each from UoM and UNSA. Percentages below refer to the share of courses *within an institution* that address a given topic.

Cross-institution findings

Table 18 summarizes cross institution findings and the subsequent tables present results of specific analysis.

Topic category	Consortium average	Highest coverage	Observations
Core concepts & paradigms	Supervised learning is almost universal (88 %); unsupervised learning reaches 49 %; reinforcement learning 20 %	UNBG & UNILJ (95 % supervised)	Advanced paradigms—transfer, semi/self-supervised, online learning—remain marginal (≤ 7 %).
Key algorithms & techniques	88 % of courses teach the classic supervised toolbox; deep learning appears in 42 %	UNILJ & UNBG (95 % supervised)	Optimisation, fuzzy logic and time-series models each fall below 20 % consortium-wide.
Model development & evaluation	Data preparation (69 %) and model evaluation (64 %) are strong; optimisation 38 %; explainability 22 %	UNBG (81 % data preparation)	Good experimental discipline overall, but interpretability is still emerging.

Python tools & libraries	Programming/data wrangling tools feature in 69 % of courses; ML/DL frameworks in 33 %; MLOps workflows in 17 %	UNBG (76 % core tools)	Practical deployment skills lag well behind modelling skills.
Engineering applications	Largely absent—only embedded/real-time AI (20 %) and intelligent control (15 %) reach double-digit coverage	UoM (55 % embedded AI)	Vertical use-cases in manufacturing, maintenance and quality control are effectively missing.
Ethical, legal, societal aspects	14 % ethical; 10 % legal/societal	UNIRI (31 % each)	Responsible-AI content is inconsistent and generally limited.

Table 18: Cross-institution analysis

Per-institution profile across six thematic categories

Core concepts & learning paradigms

Institution	Coverage profile	Notable observations
UNILJ	Very broad: supervised 94 %, unsupervised 51 %, reinforcement 26 %. Transfer, semi/self-supervised and online learning \leq 6 %.	Ljubljana is the only partner with material on every paradigm but depth drops sharply outside the classical set.
UNSA	Solid supervised (64 %) and hybrid (36 %) coverage; other paradigms absent.	Reliance on fuzzy–neural hybrids rather than newer transfer/self-supervision.
UNBG	Supervised 95 %, unsupervised 57 %, reinforcement 19 %; some semi/self-supervision (10 %) and a single transfer-learning course.	Best balance among partners, but advanced paradigms still niche.
UoM	Supervised 73 %; all other paradigms \leq 27 %.	Focus on embedded/control contexts; little exposure to data-efficient paradigms.
UNIRI	Supervised 94 %, unsupervised 63 %, reinforcement 31 %; modest semi-/self-supervision and transfer (19 %).	Only partner with \geq 30 % reinforcement learning; online learning at 13 %.

Key AI algorithms & techniques

Institution	Coverage profile	Notable observations
UNILJ	Classic supervised algorithms 94 %; deep learning 46 %; time-series/control 40 %.	Broadest technique portfolio; evolutionary methods at 20 %.
UNSA	Supervised algorithms 64 %; deep learning 45 %; optimisation & fuzzy logic 27–46 %.	Strong fuzzy-logic presence; no time-series models.
UNBG	Supervised algorithms 95 %; deep learning 33 %; time-series 10 %.	Technique mix skewed toward tabular ML; limited optimisation/fuzzy (≤ 14 %).
UoM	Supervised techniques 73 %; deep learning & time-series each 27 %.	Technique set aligned with embedded applications; low optimisation/fuzzy (18 %).
UNIRI	Supervised algorithms 94 %; deep learning 50 %; time-series 25 %.	Minimal evolutionary optimisation (6 %) despite broad paradigm coverage.

Model development & evaluation techniques

Institution	Strengths	Gaps
UNILJ	Data preparation 66 %, evaluation 69 %.	Optimisation 40 %; explainability 29 %.
UNSA	Data preparation 73 %.	Optimisation 18 %; no explainability.
UNBG	Data preparation 81 %, evaluation 67 %.	Explainability 29 %, but optimisation only 38 %.
UoM	Balanced preparation/evaluation/optimisation (45–36 %).	Explainability only 9 %.
UNIRI	Data preparation 75 %, evaluation 69 %, optimisation 50 %.	Explainability 25 %.

Engineering-oriented applications

Institution	Chief application strengths	Missing verticals
UNILJ	Robotics (20 %), embedded AI (23 %), intelligent control (17 %).	Predictive maintenance, digital twins, quality control.
UNSA	Intelligent control (18 %).	All other verticals absent.
UNBG	Signal processing (19 %), embedded AI (14 %).	Maintenance, twins, robotics, QA largely missing.

UoM	Embedded & real-time AI (55 %), intelligent control (36 %).	Predictive maintenance, digital twins, QA.
UNIRI	Embedded AI (13 %) and modest control coverage (6 %).	Predictive maintenance, robotics, signal processing.

Ethical, legal & societal aspects

Institution	Coverage	Comment
UNILJ	Ethics 20 %, legal 11 %.	Fragmented; no dedicated module.
UNSA / UNBG	0%	Absent from curricula.
UoM	Ethics 9 %.	Single elective; no legal coverage.
UNIRI	Ethics & legal each 31 %.	Only partner with systematic responsible-AI content.

Institution-specific highlights

- **UNILJ** – Broadest offering: high coverage of supervised/unsupervised methods and the largest share of robotics content; modest presence of explainability (29 %) and deployment workflows (26 %).
- **UNBG** – Strong classical pipeline (95 % supervised, 81 % data preparation) and best coverage of practical Python tooling (76 %); responsible-AI content absent.
- **UNIRI** – Matches UNILJ in foundational coverage and leads on ethics/legal (31 %) but has very limited optimisation (6 %) and industrial applications.
- **UoM** – Niche strength in embedded/real-time AI (55 %) and intelligent control (36 %); otherwise the narrowest topical spread, with limited deep learning and no MLOps.
- **UNSA** – Solid supervised foundation (64 %) and notable fuzzy-logic presence (46 %); weakest on optimisation (18 %), explainability (0 %) and engineering applications (nil).

Overall gaps

1. **Advanced learning paradigms** – Transfer, self-supervised and online learning receive minimal attention across all institutions.
2. **Deployment engineering (MLOps)** – Only one course in six introduces tools such as Docker or MLflow; UNSA and UoM have no coverage at all.
3. **Industrial verticals** – Predictive maintenance, digital twins, quality-control analytics and domain-specific robotics are essentially absent.

4. **Responsible AI** – Ethical, legal and societal issues are inconsistently addressed and almost entirely missing outside UNIRI.
5. **Explainability & robustness** – Covered in just 22 % of syllabi, leaving a compliance gap for safety-critical and regulated applications.

Conclusions

Across the consortium, AI/ML education provides a **solid classical core**—supervised learning, basic unsupervised methods, data preparation and evaluation practices appear in a majority of courses. **Deep learning and reinforcement learning are entering the mainstream** but remain far from universal, while cutting-edge paradigms (transfer, self-supervised, online learning) are still rare. **Practical software tooling is common at the data-handling level but weak in production-grade deployment workflows**, and **application-specific content is limited almost exclusively to embedded AI and intelligent control**. Finally, **responsible-AI training is patchy**, raising concerns about future compliance with EU and international governance frameworks.

Bridging these gaps will require targeted curriculum development: dedicated modules on advanced paradigms and MLOps, cross-faculty electives on industrial applications, and uniform integration of ethical and legal considerations. Coordinated action in these areas would align the partners' programmes with contemporary professional standards and better prepare graduates for AI deployment at scale.

3.3 Course educational goals analysis (Bloom taxonomy)

The next phase of our study analyses the educational goals articulated in the collected syllabi. While some documents label these goals as “Objectives,” “Goals,” or “Competencies,” most provide a dedicated **Intended Learning Outcomes (ILO)** section.

ILOs specify, in student-centred language, **what a learner will know or be able to do after successfully completing a course**. Well-formulated outcomes are measurable, achievable and assessable. The recommended phrasing begins with “*By the end of the course a successful student will be able to ...*” followed by a precise action verb that indicates the required cognitive level, typically chosen from Bloom’s taxonomy. Importantly, an ILO states the minimum performance that constitutes a pass.

This outcomes-based viewpoint differs from traditional syllabus design, which lists topics to be covered. Whereas a syllabus describes content, ILOs describe the **knowledge, skills, and attitudes that students must demonstrate**. Adopting an outcome-based approach

therefore shifts curriculum planning from “What will I teach?” to “What will students be able to do?” - a perspective central to the remainder of our analysis.

Formulating ILOs hinges on selecting appropriate action verbs that make the expected performance explicit. These verbs are commonly drawn from **Bloom’s (1956) taxonomy of educational objectives**, which organised cognitive learning into hierarchical levels—remember, understand, apply, analyse, evaluate and create. More recently, this taxonomy was revised and labels were partially renamed and the 2001 edition by Anderson and Krathwohl defines the taxonomy as follows:

Level	Typical action verbs	Brief description
1 Remember	list, recall, identify	Retrieve relevant knowledge from long-term memory (facts, terms, basic concepts).
2 Understand	explain, summarize, classify	Construct meaning from instructional messages—interpreting, exemplifying, translating.
3 Apply	use, execute, demonstrate	Carry out or implement a procedure in a given situation; solve routine problems.
4 Analyse	differentiate, organise, compare	Break material into parts, determine how they relate to one another and to an overall structure.
5 Evaluate	judge, critique, justify	Make decisions or assessments based on criteria and standards; check consistency or effectiveness.
6 Create	design, formulate, construct	Put elements together to form a novel, coherent whole or make an original product.

Table 19: Bloom taxonomy

By anchoring an outcome to a Bloom-level verb, instructors clarify both the **depth of learning** and the **assessment standard**. This verb-based convention also facilitates alignment: course activities and assessments can be checked to ensure they genuinely give students the opportunity to reach - and demonstrate - each stated outcome. Consequently, verb choice is not a stylistic detail but a critical design decision that links teaching methods, learning activities and evaluation criteria in a coherent, student-centred framework.

In operational terms, analysing a course’s ILOs entails assigning each outcome to one of Bloom’s six cognitive levels. The procedure is straightforward in principle: identify the action verb and match it to the appropriate level in the revised taxonomy. In practice, however, the syllabi reveal considerable variation in how ILOs are phrased. Some statements are concise—containing a single, well-chosen verb—whereas others combine multiple verbs or include vague language such as “understand” or “be familiar with.” This heterogeneity

makes reliable classification challenging; a single outcome may reference several cognitive operations simultaneously.

To maintain consistency in the subsequent analysis, we therefore adopted a **dominant-verb rule**: when an ILO contains multiple actions, we classify it by the verb that represents the highest cognitive demand that is both explicit and assessable. This approach provides a workable compromise between strict taxonomic precision and the realities of non-standardised syllabus writing.

3.3.1 Institutional reports

Each partner has manually checked syllabi for all of their institution courses and based on the defined ILOS they determined the **dominant Bloom level**. Furthermore, they extracted up to three action verbs that support the selected dominant Bloom level.

University of Ljubljana (UNILJ)

Thirty-five accredited syllabi spanning all three study cycles were examined. For each course we identified the dominant Bloom level - the highest-order cognitive category that appears most frequently in its set of Intended Learning Outcomes. Other ILOs at different levels may be present; the figures below therefore signal emphases rather than the full cognitive range of every course.

Distribution of dominant Bloom levels

Bloom level	Courses	Share
Remember	0	0%
Understand	2	6%
Apply	11	31%
Analyse	1	3%
Evaluate	7	20%
Create	14	40%

Table 20: UNILJ distribution of dominant Bloom levels

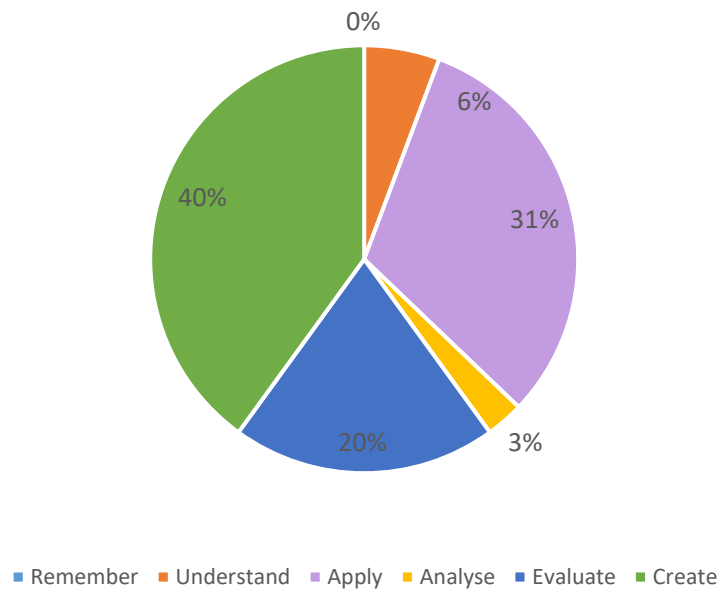


Figure 1: UNILJ ILOs dominant Bloom level distribution

The curriculum places strong emphasis on higher-order cognitive activities—design, implementation and critical appraisal - whereas pure recall seldom appears as the primary outcome.

Cycle-by-cycle overview

Cycle	Dominant levels	Comment
BSc (7 courses)	57 % Apply, 29 % Create, 14 % Analyse	First-cycle courses already target hands-on problem-solving and some design work.
MSc (22 courses)	41 % Create, 27 % Evaluate, 23 % Apply	Master’s syllabi move decisively toward design and critique of AI systems.
PhD (5 courses)	60 % Create, 40 % Apply	Doctoral offerings focus on original development, complemented by practical implementation skills.

Table 21: UNILJ dominant Bloom level by study cycle

Frequency of action verbs in ILOs

Verb	Frequency	Typical Bloom level
develop	11	Create / Apply
evaluate	10	Evaluate
use	9	Apply
implement	7	Apply / Create
program	6	Apply / Create
apply	6	Apply
design	6	Create
develop	4	Create
use	4	Apply
understand	4	Understand
choose	3	Apply / Evaluate
compare	3	Analyse
evaluate	3	Evaluate
model	3	Apply / Create
select	3	Apply / Analyse
adapt	2	Create
combine	2	Create
formulate	2	Create
identify	2	Apply / Analyse
interpret	2	Understand
analyse	2	Analyse
argument, assess, critique, differentiate, enhance, explain, gather, process, propose, search, visualize	1 each	levels vary

Table 22: UNIJ frequency of Bloom action verbs in ILOs



Figure 2: Word-cloud of UNILJ ILOs dominant action verbs

Observations on the current state

- **Design orientation.** 40% of courses list *Create* as the main cognitive target, indicating that many programmes expect students to produce novel artefacts or solutions.
- **Procedural competence.** Nearly one-third of courses focus on *Apply*, reinforcing hands-on modelling skills.
- **Critical evaluation.** One in five syllabi foregrounds *Evaluate*, confirming that judgement and assessment of AI systems receive explicit attention.
- **Analytical wording is rare.** Only a single course cites *Analyse* as dominant, although analytical verbs may still appear in subordinate outcomes.
- **Foundational recall implicit.** No course identifies *Remember* as its leading level; basic knowledge acquisition is likely embedded within broader, higher-level ILOs.

Reflections for interpretation

- The dominance of upper-level verbs (*Create*, *Evaluate*) suggests that **higher-order thinking is well embedded** in the accredited programmes.
- The absence of *Remember* as a leading outcome does not imply neglect of factual knowledge; it simply reflects how staff articulate threshold expectations.

- Given that each course contains multiple ILOs, the **overall cognitive spectrum is undoubtedly wider** than the dominant-level metric indicates. The present analysis should therefore be read as a snapshot of prevailing emphases.

UNILJ's syllabi form a steep cognitive curve that already begins at Apply in first-year modules and peaks at Create in most MSc and PhD courses. Students therefore move rapidly from hands-on modelling to genuine design work while also learning to evaluate. The only notable gap is that explicit Analyse outcomes are rare, suggesting that diagnostic reasoning - though probably practised - is under-articulated.

University of Sarajevo (UNSA)

Eleven accredited syllabi were reviewed: four from the Faculty of Electrical Engineering (cycle 1), four from the Faculty of Mechanical Engineering (cycles 1 & 2) and three from the Faculty of Natural Sciences and Mathematics (cycle 2). For each course the most frequent high-order verb in its Intended Learning Outcomes (ILOs) was used to assign a dominant Bloom level.

Distribution of dominant Bloom levels

Bloom level	Courses	Share
Remember	0	0%
Understand	0	0%
Apply	4	36%
Analyse	1	9%
Evaluate	5	45%
Create	1	9%

Table 23: UNSA distribution of dominant Bloom levels

UNSA dominant Bloom levels

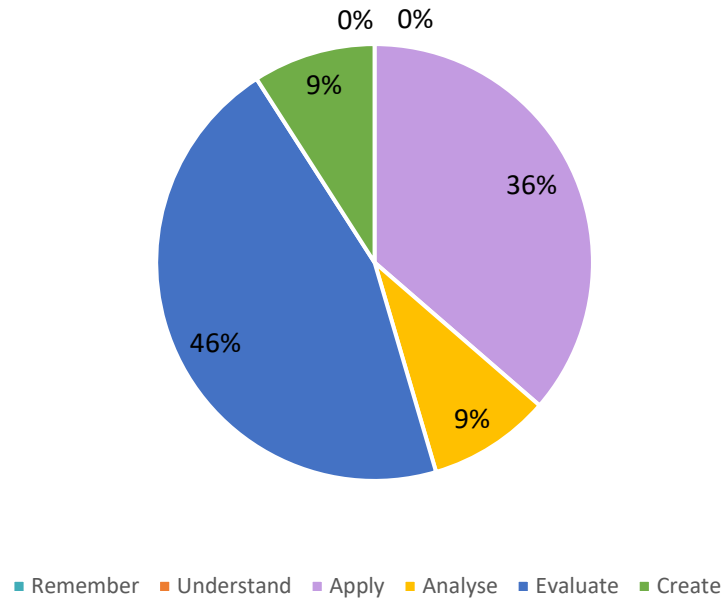


Figure 3: UNSA ILOs dominant Bloom level distribution

Nearly half the syllabi position students at the **Evaluate** level - critiquing or judging AI solutions - while a further third emphasise **Apply**, signalling hands-on procedural competence. Design-level work (**Create**) appears only once and pure recall (**Remember**) never constitutes the dominant target.

Cycle-by-cycle overview

Cycle	Dominant levels	Comment
BSc (4 courses)	Apply 50 %, Analyse 25 %, Evaluate 25 %	First-cycle modules balance practical execution with introductory critical thinking.
MSc (7 courses)	Evaluate 57 %, Apply 29 %, Create 14 %	Master's courses shift toward appraisal and occasional design tasks; no course lists Analyse as dominant.
PhD	–	No doctoral AI/ML-related syllabi were supplied.

Table 24: UNSA dominant Bloom level by study cycle

Frequency of action verbs in ILOs

Verb	Frequency	Typical Bloom level
apply	6	Apply

conclude	4	Evaluate
evaluate	4	Evaluate
analyse	3	Analyse
implement	2	Apply / Create
solve	2	Apply
create, explain, identify, model, select, simulate, understand, predict, develop, interpret, project, use	1 each	levels vary

Table 25: UNSA frequency of Bloom action verbs in ILOs

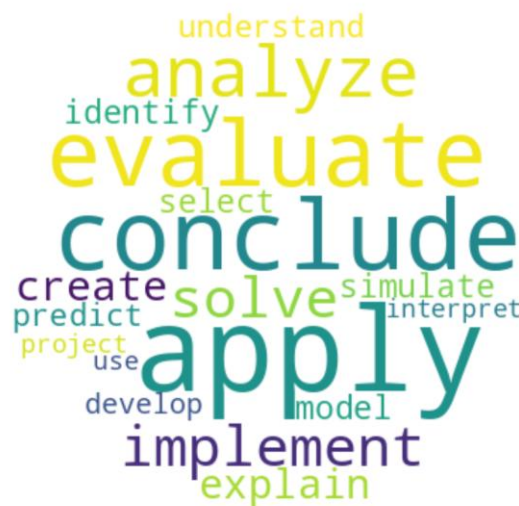


Figure 4: Word-cloud of UNSA ILOs dominant action verbs

The verb distribution underscores the evaluative orientation (evaluate, interpret, conclude) and a secondary focus on direct application (apply, use, solve).

Observations on the current state

- **Critical judgement emphasised.** The prominence of *Evaluate* verbs suggests that existing courses require students to assess or justify AI solutions rather than merely reproduce them.
- **Procedural skills remain central.** Over one-third of courses still list *Apply* as their main outcome, reflecting the need for practical modelling competence.
- **Limited design and analysis articulation.** Only a single course each is dominated by *Create* and *Analyse*, indicating that innovative design and structured de-composition are less frequently foregrounded in stated outcomes—even if practised in class.

- **Foundational recall implicit.** Absence of *Remember* and *Understand* as leading levels implies that factual knowledge and basic comprehension are embedded within broader, higher-order outcomes rather than singled out for assessment.

Reflections for interpretation

- **High-order emphasis is deliberate but selective.** Nearly half of the UNSA syllabi foreground *Evaluate* as their dominant cognitive target, signalling a strong commitment to critical appraisal of AI methods. Yet only one course is centred on *Create* and one on *Analyse*, so innovative design and systematic decomposition appear less frequently as headline goals.
- **Procedural mastery is still valued.** With four courses led by *Apply*, the curriculum ensures that students can implement models and techniques—an essential bridge between theory and practice.
- **Foundational levels are implicit, not absent.** No syllabus lists *Remember* or *Understand* as the dominant level, but this does not mean factual recall and basic comprehension are ignored; rather, they are embedded within broader, higher-order statements. Teaching staff will still need to scaffold these fundamentals in classroom activities and assessments.
- **Faculty-level variation exists.** Mechanical-engineering modules are more evaluation-oriented, whereas electrical-engineering modules lean toward application; natural-science courses combine evaluation with predictive or analytical verbs. This diversity suggests opportunities for cross-faculty collaboration to balance design, analysis and evaluation skills.

UNSA's portfolio emphasises Evaluate and Apply, signalling a dual focus on critical appraisal and operational skill. Design-level targets appear only once and analytical decomposition is limited, so the cognitive ladder flattens in the upper tiers. Making a few outcomes explicitly Create- or Analyse-oriented would balance the strong evaluative culture and provide clearer scaffolding for innovation.

University of Belgrade (UNBG)

Twenty accredited syllabi from the School of Electrical Engineering were examined (seven BSc, ten MSc, three PhD). For each course the most frequent higher-order verb in its ILO set was taken as the dominant Bloom level.

Distribution of dominant Bloom levels

Bloom level	Courses	Share
Remember	0	0%

Understand	1	5%
Apply	9	43%
Analyse	4	19%
Evaluate	1	5%
Create	6	29%

Table 26: UNBG distribution of dominant Bloom levels

Almost half of the programmes set hands-on implementation (*Apply*) as the primary target, while a further 30 % expect students to design or originate artefacts (*Create*). Pure recall never appears as the leading outcome.

UNBG dominant Bloom levels

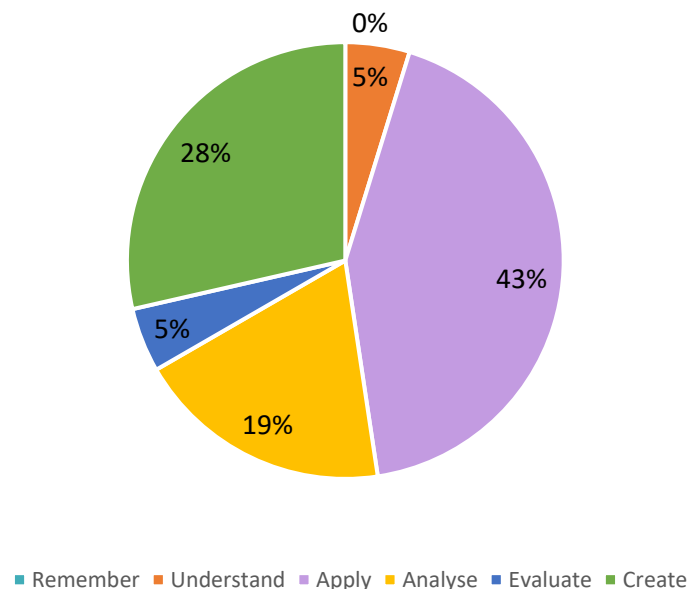


Figure 5: UNBG ILOs dominant Bloom level distribution

Cycle-by-cycle overview

Cycle	Dominant levels	Comment
BSc (7 courses)	86 % Apply, 14 % Create	Introductory courses emphasise operational competence; one course already targets design.
MSc (11 courses)	36 % Analyse, 27 % Apply, 18 % Create, 9 % Evaluate and Understand	Master's syllabi broaden the spectrum, adding structured reasoning (<i>Analyse</i>) and limited critical appraisal.

PhD (3 courses)	100 % Create	Doctoral offerings focus exclusively on original design and innovation.
------------------------	--------------	---

Table 27: UNBG dominant Bloom level by study cycle

Frequency of action verbs in ILOs

Verb	Frequency	Typical Bloom level
develop	8	Create / Apply
model	7	Analyse / Apply
apply	5	Apply
solve	4	Apply
identify	4	Analyse / Apply
compare	4	Analyse
create	4	Create
construct	3	Create / Apply
explain	2	Understand / Apply
examine	2	Analyse
discover	2	Analyse
classify	2	Analyse
design	2	Create
utilize	2	Apply
evaluate, improve, imagine, summarize, predict, illustrate, simulate, make use of, adapt, analyse, demonstrate, invent	1 each	levels vary

Table 28: UNBG frequency of Bloom action verbs in ILOs

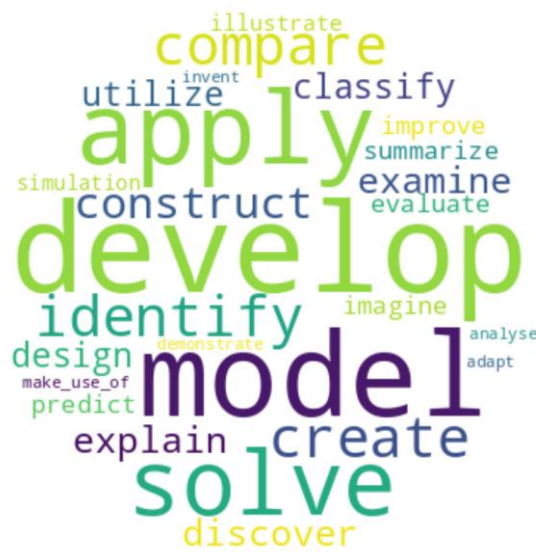


Figure 6: Word-cloud of UNBG ILOs dominant action verbs

Observations on the current state

- **Strong implementation focus.** Nearly half of all courses concentrate on *Apply*, underscoring the programme’s commitment to practical, hands-on mastery of AI/ML methods.
- **Clear design trajectory.** The proportion of *Create* outcomes rises from 14 % in BSc to 100 % in PhD courses, demonstrating a deliberate progression toward innovation and original research contributions.
- **Emerging analytical depth.** One-fifth of courses target *Analyse* as their dominant level, but these are confined to the MSc cycle; explicit analytical reasoning is absent in BSc and PhD syllabi.
- **Limited emphasis on critical appraisal.** Only one course lists *Evaluate* as its primary goal, suggesting that systematic evaluation of AI systems—vital for robust, ethical deployment—remains under-articulated.
- **Verb diversity reflects engineering culture.** High-frequency verbs such as *develop*, *model*, *apply*, and *solve* mirror an engineering mindset aimed at building functional artefacts, whereas verbs linked to critique or reflection occur less often.

Reflections for interpretation

- **Procedural dominance with design progression.** *Apply* remains the leading emphasis, especially at BSc level, ensuring that students can implement core

techniques. The step-up to *Create* becomes prominent in MSc and universal in PhD syllabi, signalling a clear developmental pathway toward innovation.

- **Analytical reasoning gains traction at MSc.** Four of ten master’s courses list *Analyse* as dominant, indicating deliberate cultivation of decomposition and comparison skills, yet only a single course prioritises *Evaluate*.
- **Foundational levels are implicit.** No course foregrounds *Remember* or *Understand*; basic knowledge acquisition is assumed rather than explicitly assessed.
- **Verb palette confirms technical orientation.** High counts for *develop*, *model*, *apply*, *solve* and *identify* highlight a curriculum geared toward building and refining working AI systems.

UNBG presents an engineering-shaped profile: almost half the courses concentrate on implementation (*Apply*) and another third on original design (*Create*), while *Analyse* adds structured reasoning in several MSc subjects. *Evaluate* appears only once, leaving systematic critique as the missing rung in an otherwise well-graded ladder from practice to invention.

University of Montenegro (UoM)

Eleven accredited AI/ML or AI-related syllabi were reviewed: one BSc course, nine MSc courses and one PhD course. For each syllabus the most frequent high-order verb in its Intended Learning Outcomes (ILOs) was taken as the dominant Bloom level. Other outcomes at different levels certainly exist; the statistics below capture the prevailing emphasis only.

Distribution of dominant Bloom levels

Bloom level	Courses	Share
Remember	0	0%
Understand	1	9%
Apply	3	27%
Analyse	3	27%
Evaluate	1	9%
Create	3	27%

Table 29: UoM distribution of dominant Bloom levels

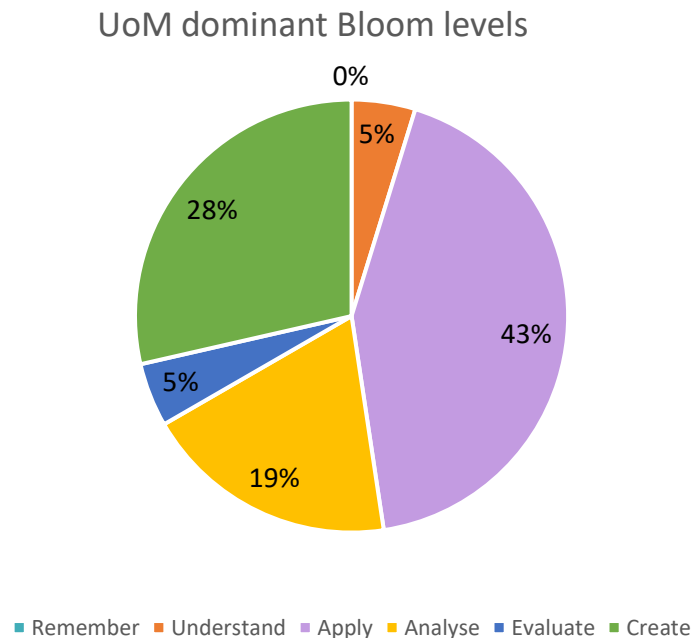


Figure 7: UoM ILOs dominant Bloom level distribution

Cognitive demand is evenly split between practical application, analytical reasoning and creative design. Pure recall is never the main target.

Cycle-by-cycle overview

Cycle	Dominant levels	Comment
BSc (1 course)	100% Evaluate	Introductory AI course centres on critical appraisal.
MSc (9 courses)	Apply 33 %, Analyse 33 %, Create 22 %, Understand 11 %	Master’s curriculum balances implementation, analysis and innovation; one course emphasises conceptual grasp.
PhD (1 courses)	100 % Create	Doctoral module focuses on original design and integration.

Table 30: UoM dominant Bloom level by study cycle

Frequency of action verbs in ILOs

Verb	Frequency	Typical Bloom level
apply	5	Apply

analyse	4	Analyse
explain	3	Understand
use	3	Apply
define	2	Understand / Apply
compare	2	Analyse
understand	2	Understand
evaluate, create, recognise, model, distinguish, identify, design, simulate, perform, classify, review, integrate	1 each	levels vary

Table 31: UoM frequency of Bloom action verbs in ILOs

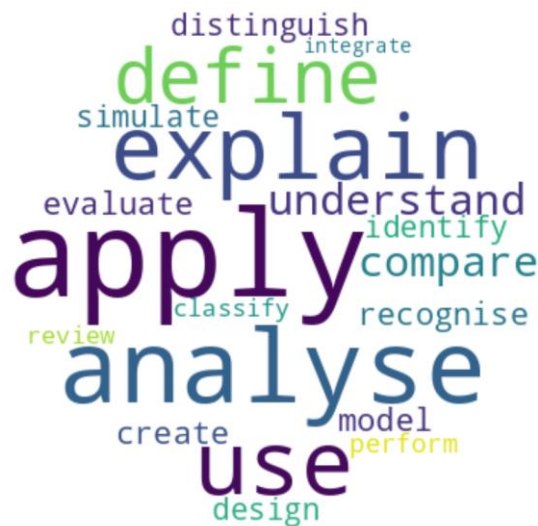


Figure 8: Word-cloud of UoM ILOs dominant action verbs

Observations on the current state

- **Balanced cognitive spread.** The trio *Apply–Analyse–Create* each dominates roughly one-quarter of courses, showing a curriculum that moves students from implementation through reasoning to design.
- **Conceptual and evaluative elements present but modest.** One course each leads with *Understand* or *Evaluate*; these competencies may be embedded more broadly but are seldom foregrounded.
- **Foundational recall implicit.** No dominant *Remember* outcomes appear—basic factual knowledge is assumed within higher-order statements.

- **Verb palette reflects engineering practice.** High-frequency verbs (*apply, use, analyse, compare, create*) point to pragmatic modelling, analytical scrutiny and design work - aligned with UoM's embedded-systems and mechatronics focus.

Reflections for interpretation

- **Strength in application & design.** Students are expected to implement and originate AI solutions, matching the university's engineering orientation.
- **Analytical reasoning well-articulated.** A third of syllabi highlight *Analyse*, indicating deliberate cultivation of diagnostic and comparative skills.
- **Limited explicit emphasis on critical appraisal.** With only one course led by *Evaluate*, systematic judgement of AI systems could be more visible in programme documentation - even if practiced implicitly.
- **Conceptual scaffolding may be hidden.** A single *Understand*-dominant course suggests that foundational conceptual goals are nested within broader outcomes rather than assessed separately.
- **PhD focus on creation aligns with research objectives.** The sole doctoral module targets *Create*, signalling expectation of novel contributions.

UoM offers one of the most even spreads: *Apply, Analyse* and *Create* each headline roughly one-quarter of courses. This symmetry nurtures implementation skill, diagnostic reasoning and innovation in parallel, though only a single course leads with *Evaluate* and foundational levels are implicit.

University of Rijeka (UNIRI)

Sixteen accredited syllabi were analysed: three BSc, nine MSc and four PhD. Each course was assigned a dominant Bloom level based on the action verb that appears most frequently in its Intended Learning Outcomes (ILOs).

Distribution of dominant Bloom levels

Bloom level	Courses	Share
Remember	0	0%
Understand	0	0%
Apply	7	43%
Analyse	1	6%
Evaluate	6	38%
Create	2	13%

Table 32: UNIRI distribution of dominant Bloom levels

UNIRI dominant Bloom levels

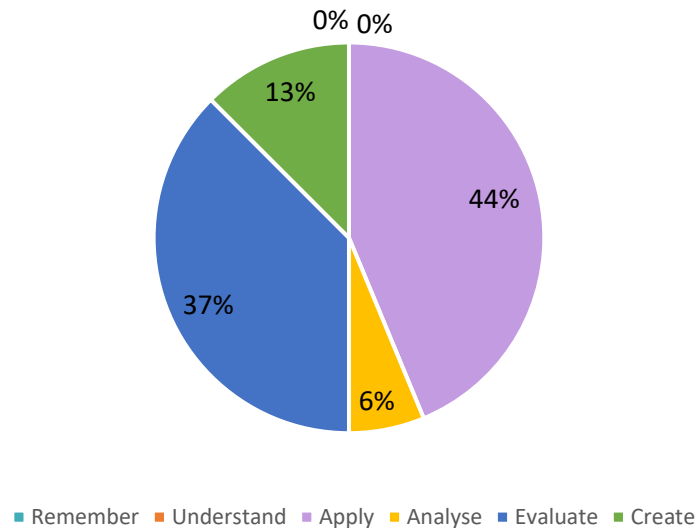


Figure 9: UNIRI ILOs dominant Bloom level distribution

Almost half of all courses prioritise procedural competence (Apply), while a substantial 38 % centre on critical appraisal (Evaluate). Design (Create) and analytical decomposition (Analyse) appear, but less frequently; foundational recall and basic comprehension are never dominant.

Cycle-by-cycle overview

Cycle	Dominant levels	Comment
BSc (3 course)	100 % Apply	Introductory modules focus squarely on hands-on implementation.
MSc (9 courses)	Apply 44 %, Evaluate 22 %, Create 22 %, Analyse 11 %	Master’s curriculum balances implementation, critique and some design work.
PhD (4 courses)	100 % Evaluate	Doctoral offerings concentrate on critical assessment of advanced AI systems.

Table 33: UNIRI dominant Bloom level by study cycle

Frequency of action verbs in ILOs

Verb	Frequency	Typical Bloom level
------	-----------	---------------------

apply	9	Apply
evaluate	8	Evaluate
analyse	8	Analyse
develop	4	Create / Apply
design	4	Create
understand	2	Understand
compare	2	Analyse
explain	2	Understand / Analyse
assess	2	Evaluate
solve, implement, select, define, acquire, investigate, interpret	1 each	levels vary

Table 34: UNIRI frequency of Bloom action verbs in ILOs



Figure 10: Word-cloud of UNIRI ILOs dominant action verbs

Observations on the current state

- **Implementation is a constant.** *Apply* dominates BSc syllabi and remains substantial at MSc level, ensuring strong practical skills.
- **Evaluation culture is robust.** A high share of master's courses and every PhD course foreground *Evaluate*, indicating systematic emphasis on appraisal and judgement of AI artefacts.
- **Design and innovation present but moderate.** Only two courses list *Create* as dominant; design verbs (develop, design) nonetheless appear frequently in supporting positions.

- **Limited explicit analysis.** Despite frequent use of “analyse” verbs within ILOs, only one MSc course has *Analyse* as its primary cognitive target.

Reflections for interpretation

- The curriculum **progresses from implementation (BSc) to critique (PhD)**, with MSc acting as a bridge that blends practice, evaluation and some design.
- **Analytical reasoning** is articulated verb-wise but seldom emerges as the headline outcome; making it explicit in a few additional courses could close the evaluation-analysis loop.
- **Design-level ambition** is visible yet could be expanded, especially in MSc offerings where two-fifths of courses already target *Apply*.
- Because this analysis uses **dominant verbs only**, lower-level verbs are still present in many ILO sets. Thus, factual knowledge and conceptual understanding are likely taught even though they do not headline the outcome lists.

The UNIRI's AI/ML syllabi trace a clear cognitive staircase: “apply” dominates at the bachelor tier, “evaluate” at the doctoral tier, with a mixed middle layer that introduces design (*Create*) and deeper analysis (*Analyse*). This trajectory ensures that graduates can both implement and rigorously appraise AI systems. Strengthening explicit analytical and creative outcomes in a few additional MSc modules—and confirming that foundational knowledge is adequately assessed—would round out an already well-structured, competence-oriented curriculum.

3.3.2 Cross-institution analysis

Having examined every syllabus in the consortium, we now shift from single-institution snapshots to a collective view of course educational goals. This cross-institution lens allows us to address questions that single-faculty analyses cannot:

- Where does the collective curriculum sit on the lower-order → higher-order spectrum?
- Is there a clear cognitive progression from bachelor to master to doctorate across the region?
- Which verbs - and therefore which kinds of learning - dominate, and which ones are rare or missing?

By quantifying these aspects, we provide a robust baseline for collaborative curriculum design, highlight convergences that can underpin joint courses, and expose gaps that future project work packages may wish to fill.

Distribution of dominant Bloom levels

More than two-thirds of all courses position students at Apply-Create-Evaluate, showing a strong orientation toward doing, designing and judging rather than basic recall.

Bloom level	Courses	Share
Remember	0	0%
Understand	3	3%
Apply	34	37%
Analyse	10	11%
Evaluate	20	22%
Create	26	28%

Table 35: Distribution of dominant Bloom levels across all institutions

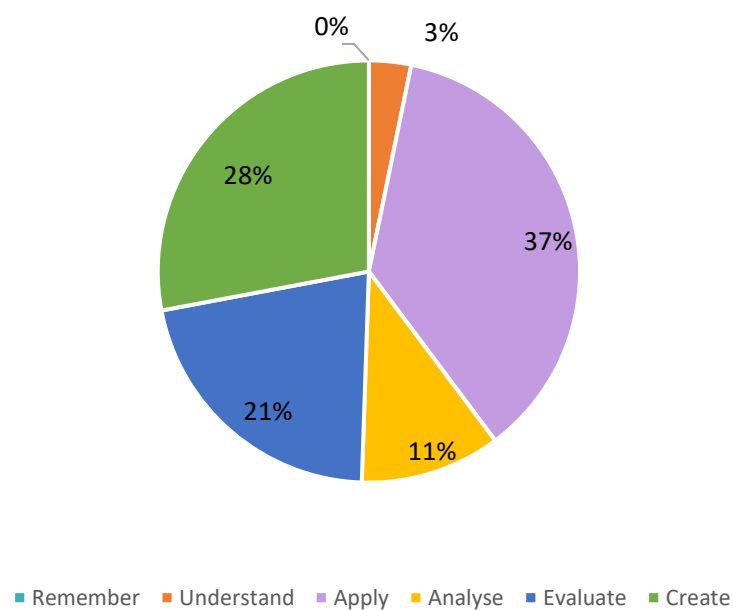


Figure 11: Overall ILOs dominant Bloom level distribution

The merged dataset confirms a decisive move away from lower-order cognition. No course lists *Remember* as its dominant outcome and only three (3 %) foreground *Understand*. Instead, almost four courses in ten (37 %) concentrate on *Apply*: students are expected to implement algorithms, configure tools or run experiments as a threshold competence. Another 28 % of syllabi push beyond execution to *Create*, requiring learners to design

architectures, invent solutions or integrate subsystems. A further 22 % are led by *Evaluate*, signalling systematic appraisal of model quality, trade-offs or ethical fitness. *Analyse* - decomposition and comparison - registers at 11 %, providing some but not uniform attention to structural reasoning.

Implication. The five universities collectively privilege “doing, designing and judging” over “remembering and describing”. This aligns with the practical ethos of engineering and computer-science education, but it also means conceptual scaffolding and analytical decomposition are often embedded rather than made explicit.

Cycle-by-cycle overview

- BSc programmes focus on procedural mastery (*Apply*).
- MSc curricula are mixed, blending design (*Create*), evaluation and some analysis.
- PhD syllabi pivot to innovation (*Create*) and critical appraisal (*Evaluate*).

	BSc	MSc	PhD
All courses	22	57	14
Apply	15	17	2
Analyse	2	8	0
Evaluate	2	13	4
Create	3	16	8
Understand	0	3	0
Remember	0	0	0

Table 36: Dominant Bloom level by study cycle across all institutions

Bachelor - Implementation rules: two-thirds of first-cycle courses list *Apply* as dominant. The absence of *Create* or *Evaluate* reflects a deliberate choice to establish hands-on fluency before students move to critique or innovation.

Master - Cognitive diversity expands. *Apply* drops to one-third, while *Create* and *Evaluate* each account for about a quarter. *Analyse* also appears (11 %), showing that MSc curricula start asking students how and why models work, not just whether they work.

Doctorate - The profile flips: more than one course in two (57 %) now specifies *Create* or *Evaluate* as the headline level - usually in tandem - with *Apply* becoming a minority concern. PhD students are therefore pushed toward original contribution and rigorous assessment rather than further procedural practice.

Implication. The cognitive staircase is clear: *Apply* → *Analyse* → *Create/Evaluate*. Where gaps appear (e.g., analysis seldom dominant at PhD), they may be an artefact of wording rather than intent, but they nevertheless obscure the developmental logic to external reviewers.

Frequency of action verbs in ILOs

Verbs that can straddle levels (e.g. implement, identify) were counted once but may map to different Bloom categories depending on context.

Verb	Frequency	Typical Bloom level
apply	30	Apply
develop	21	Create / Apply
evaluate	21	Evaluate
analyse	16	Analyse
model	12	Analyse / Apply
design	12	Create
compare	11	Analyse
use	9	Apply
implement	9	Apply / Create
solve	8	Apply
next tier: choose, identify, select, construct, explain, adapt, assess, formulate, predict, imagine, etc.	3 - 7	levels vary

Table 37: Frequency of Bloom action verbs in ILOs across all institutions

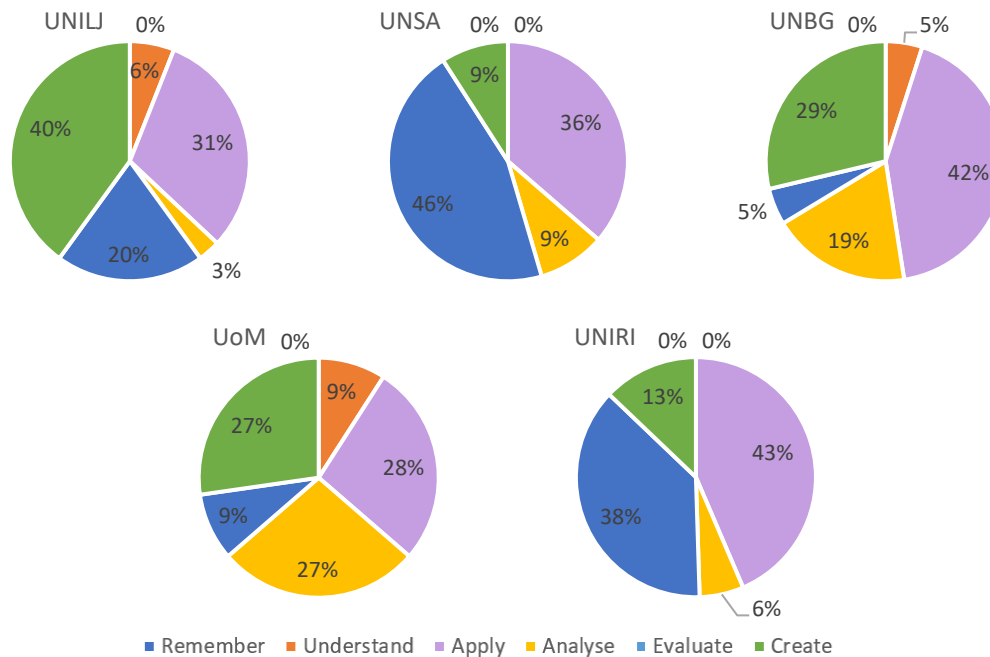


Figure 13: Comparison of ILOs dominant Bloom level distribution across institutions

- **UNILJ & UNBG** push furthest into *Create*, reflecting vibrant research groups and project-centred teaching; *Analyse* remains modest, hinting at an opportunity for more structured reasoning outcomes.
- **UNSA** stands out with an *Evaluate*-heavy profile, showing strong attention to statistical inference and fuzzy-logic appraisal, yet with limited design-level articulation.
- **UoM** offers the most even cognitive spread, but - with only one *Evaluate*-dominant course - systematic critique is under-represented.
- **UNIRI** excels in *Evaluate* at PhD level, but BSc/MSc stages could benefit from more explicit analysis and design verbs.

These nuances illustrate that while the overall patterns are similar, each institution has a distinctive cognitive “accent” that reflects local expertise, regulatory frameworks and historical curriculum choices.

Opportunities and recommendations

- **Make analysis explicit.** In several programmes analyse verbs appear within ILO sentences but rarely as the dominant target. Elevating a few of these to headline status would clarify assessment expectations and balance the strong design-and-evaluation culture.

- **Embed responsible-AI appraisal earlier.** The high share of Evaluate at PhD level shows commitment to critical thinking; introducing dedicated evaluation ILOs into late-Bachelor or early-Master courses would scaffold that capability sooner.
- **Articulate conceptual scaffolding.** Even a single Understand-level outcome per course can help teaching teams verify that prerequisite knowledge is explicitly taught and assessed.
- **Diversify verb repertoire.** Incorporating verbs such as justify, critique, reflect, or validate would broaden the evaluative lens beyond technical metrics to include ethical, social and sustainability considerations - aligning with EU AI Act expectations.
- **Share exemplar ILO phrasing.** Cross-institution workshops on outcome writing could harmonise terminology, reduce ambiguity (e.g., “implement” vs “develop”) and make competency progression clearer to external stakeholders.

Closing reflection

Collectively, the five universities deliver AI/ML education that is ambitious and practice-centred, guiding students from hands-on implementation toward innovation and critique. The absence of lower-level outcomes suggests confidence in students’ prerequisite knowledge, but it also places a premium on coherent scaffolding and formative assessment. By sharpening analytical articulation, broadening evaluative verbs, and foregrounding responsible-AI dimensions, the consortium can transform an already strong competence framework into a fully rounded, internationally distinctive educational offering.

3.4 Courses’ Intended Learning Outcomes analysis

In this section we analyse the syllabi’s Intended Learning Outcomes (ILO) blocks, focusing on the section titles used, the number of ILOs per course, the length of individual ILOs, and the number of Bloom-aligned action verbs employed in each ILO.

3.4.1 ILOs: naming, format, and structure

The collected syllabi differ in language and template, but all are structured documents and each includes a clearly identifiable section that states the course’s Intended Learning Outcomes (ILOs). In many cases this block is explicitly titled “**Intended Learning Outcomes**” or “**Learning outcomes**”; in others, institutions use locally standard labels (e.g., “**Expected study results**” or “**Course outcomes**”).

Using the English headings where available - and the bilingual equivalents where necessary - we consistently located the outcome blocks for analysis. Table 38 summarises the section titles observed across institutions and faculties.

Institution	naming
UNILJ	intended learning outcomes, expected study results
UNSA	ishodi učenja (Bosnian: learning outcomes), ishodi modula (Bosnian: module outcomes)
UNBG	the outcome, learning outcome of the course, ishod predmeta (Serbian: course outcome)
UoM	learning outcomes, ishodi učenja (Serbian: learning outcomes)
UNIRI	expected learning outcomes, expected outcomes for the course, expected course learning outcomes

Table 38: Intended Learning Outcomes naming in collected syllabi.

With respect to document layout - and, by extension, the layout of the ILO block - Table 39 summarises three observed formats: (i) standard institutional templates, (ii) non-standard faculty-specific templates, and (iii) web format (HTML captures):

ILO layout	UNILJ	UNSA	UNBG	UoM	UNIRI
standard	21	8	18	10	16
non-standard	0	3	3	0	0
web format	14	0	0	0	0

Table 39: Distribution of ILO layouts across partner institutions.

Except for the University of Ljubljana's Faculty of Electrical Engineering, whose syllabi were collected from the website (n = 14 in the "web format" category), all syllabi were provided in official template form. Most courses use the standard format (UNILJ 21; UNSA 8; UNBG 18; UoM 10; UNIRI 16). Non-standard formats occur only at UNSA (n = 3) and UNBG (n = 3). No web-format syllabi were present outside UNILJ.

While English versions were sought, some partners could provide syllabi only in the local language. This did not affect our analysis: where ILOs were not in English, we applied a bilingual verb list to map local action verbs to Bloom levels. Table 40 shows the language distribution of ILOs. All UNSA outcomes are in Bosnian (n = 11). At UNBG, three courses list ILOs in Serbian (Cyrillic). UoM includes eight courses with English ILOs and two in Serbian/Montenegrin. UNILJ and UNIRI provide ILOs in English throughout.

ILO language	UNILJ	UNSA	UNBG	UoM	UNIRI
--------------	-------	------	------	-----	-------

English	35	0	18	8	16
Serbian	0	0	3	2	0
Bosnian	0	11	0	0	0

Table 40: Distribution of ILO languages across partner institutions.

For the sake of clarity and understandability of the syllabi ILOs, it is recommended that the outcomes are defined in a consistent format or structure, preferably following a standardized categorization. As shown in Table 41, the analysed syllabi ILOs are for the most part not structured. In most cases, ILOs are defined as unordered and unstructured lists (60 courses out of 94), while many courses present ILOs in completely free-text form, which makes the learning outcomes harder to identify and interpret (21 courses). Only in 12 courses were we able to extract a recognizable structure of ILOs, where the intended learning outcomes were explicitly linked to knowledge, skills, competences, or similar categories.

ILO structure	UNILJ	UNSA	UNBG	UoM	UNIRI
unstructured list	26	5	5	8	16
free text	1	2	16	2	0
structured	8	4	0	0	0

Table 41: Distribution of ILO structures across partner institutions.

3.4.2 Number of ILOs per course

In the next phase of our analysis, we examined how many ILOs are defined in course syllabi. A common rule of thumb at some institutions is to align one ILO per ECTS, but this approach is not universally accepted and can be misleading if applied rigidly. Since 1 ECTS typically corresponds to 25–30 hours of student work, the heuristic links each outcome to a substantial workload and assessment component, helping to avoid extremes (e.g., only 2 ILOs in a 10 ECTS course). However, an ILO does not represent a workload unit; rather, it specifies what a student is expected to achieve. It is generally preferable to define fewer, well-formulated outcomes that encompass knowledge, skills, and competences, rather than many fragmented ones. For complex, integrative courses (e.g., a 10 ECTS project course), 4–6 broad ILOs may be sufficient, such as Design a system..., Evaluate..., or Communicate results.... In contrast, technical lecture-based courses with the same credit load may require 8–10 more specific ILOs. Therefore we omitted the ECTS points and focused solely on the number of ILOs in a course.

Table 42 shows statistics on the number of ILOs defined in the collected course syllabi. The analysis of the number of ILOs per course across the partner institutions reveals notable differences in both the level of detail and the consistency of outcome definition. At the University of Ljubljana (UNILJ), syllabi typically contain around five ILOs per course, with an average of 5,43 and a median of 5. The distribution is relatively balanced, although some courses list as many as 13 ILOs, indicating that while most syllabi are concise, there are cases where outcomes are described in much greater detail.

	UNILJ	UNSA	UNBG	UoM	UNIRI
courses	35	11	21	11	16
ILOs	190	33	52	57	107
average	5,43	3,00	2,48	5,18	6,69
standard deviation	1,71	1,60	1,05	1,59	1,99
min	2	1	1	1	3
max	13	5	5	7	12
median	5	4	2	5	7

Table 42: Statistics on number of ILOs per course.

The University of Sarajevo (UNSA) shows a more minimal approach. With an average of 3 ILOs per course and a maximum of 5, their syllabi are consistent but define outcomes in a limited way. This contrasts with the University of Belgrade (UNBG), where the average is even lower at 2,48 ILOs per course and the median is only 2. The narrow range and low variability confirm a very concise and uniform style of outcome formulation, which may enhance clarity but at the expense of comprehensiveness.

The University of Montenegro (UoM) lies closer to the Ljubljana pattern, with an average of 5,18 and a median of 5 ILOs per course. The range of outcomes per syllabus is smaller, with a maximum of 7, which suggests that syllabi are consistently defined and avoid extremes in either direction.

The University of Rijeka (UNIRI), by contrast, stands out as the partner with the most detailed outcome definitions. On average, courses include 6,69 ILOs, with the median reaching 7. Some syllabi list as few as 3 outcomes, while others include up to 12, which explains the relatively high variability observed. This indicates that while many courses aim for comprehensiveness and detail, the degree of elaboration is not uniform across the institution.

Overall, the comparison highlights three distinct tendencies within the consortium. UNBG and UNSA adopt a minimalist approach with very few ILOs per course, UNILJ and UoM

maintain a balanced and consistent formulation with around five outcomes, and UNIRI places the strongest emphasis on detail, often defining a larger set of outcomes per syllabus. These differences suggest that while some institutions prioritize clarity and brevity, others aim for broader coverage and a more comprehensive description of learning achievements.

3.4.3 Length of ILOs

ILOs are generally expected to be short, precise statements that clearly describe what a learner will be able to do after completing a course. While there is no universal rule on their exact length, most recommendations suggest that an ILO should normally be between 10 and 25 words long, and rarely exceed 40 words. Anything much shorter often ends up being too vague, while overly long formulations tend to combine several different achievements into one sentence and become difficult to interpret or assess. A good practice is to keep each ILO to a single, well-structured sentence that begins with a measurable verb, states the object of learning, and, if needed, specifies the context or conditions in which the learning should be demonstrated.

In our courses, ILOs generally followed these recommendations (Table 43).

	UNILJ	UNSA	UNBG	UoM	UNIRI
courses	35	11	21	11	16
ILOs	190	33	52	57	107
words	1613	774	676	559	1655
average	8.49	23.45	13.00	10.75	15.47
standard deviation	2.84	12.11	5.18	4.64	8.55
min	3	4	5	3	3
max	18	56	28	25	38
median	8	20	12	10	13
ILOs with > 25 words	0	12	1	0	13
perc. of ILOs with > 25 words	0.00	36.36	1.92	0.00	12.15

Table 43: Statistics on lengths of ILOs

The analysis of the number of words used in ILOs across the five partner institutions reveals considerable differences in both length and consistency of formulation. At the University of Ljubljana (UNILJ), ILOs are very concise, with an average length of just 8.5 words and a median of 8. The shortest outcomes contain as few as three words, while the longest reach 18, but none exceed the recommended 25-word limit. This indicates that UNILJ syllabi generally formulate outcomes in a clear but very brief manner, which may risk vagueness if the wording is too minimal.

In contrast, the University of Sarajevo (UNSA) stands out with particularly long ILOs. The average outcome length is 23.5 words, with a median of 20 and some ILOs extending to over 50 words. More than one third of all UNSA outcomes exceed 25 words, making them the least concise of all partners. While this level of detail may enhance clarity and contextualization, it also risks overloading a single outcome with multiple objectives, which complicates assessment and reduces comparability.

The University of Belgrade (UNBG) falls between these two extremes. With an average of 13 words per outcome and a median of 12, most ILOs are within the recommended range. Only one outcome exceeds the 25-word threshold, accounting for less than 2% of their total. This suggests that UNBG generally balances clarity with brevity, although the relatively wide spread of lengths shows some inconsistency in how syllabi are prepared.

The University of Montenegro (UoM) is close to UNBG in its approach. Its average ILO length is 10.8 words, with a median of 10 and a minimum of 3. None of the outcomes exceed the 25-word guideline, which suggests a high degree of consistency and adherence to concise formulation. Their outcomes are on the shorter side of the recommended spectrum, but generally within an acceptable range.

The University of Rijeka (UNIRI), on the other hand, uses longer ILOs than most other partners, with an average of 15.5 words and a median of 13. While many of their outcomes are still concise, more than 12% exceed 25 words, with the longest stretching to 38. This points to a mixed practice in which some syllabi are well-structured and succinct, while others include overly detailed or compound formulations.

Taken together, these findings highlight three distinct tendencies among the partners. UNILJ and UoM favour very short and concise ILOs, sometimes at the risk of oversimplification. UNBG maintains a moderate and balanced style, staying well within the recommended range. UNIRI and especially UNSA, by contrast, formulate much longer outcomes, often exceeding recommended lengths and reflecting a tendency to embed multiple ideas within a single statement. These differences show that while some institutions prioritize brevity and measurability, others emphasize detail and contextualization, and aligning practices across the consortium will require balancing these two extremes.

3.4.4 Bloom verbs in ILOs

An ILO should normally contain a single Bloom's verb. The reason is simple: each outcome is meant to capture one distinct, measurable achievement, and the verb defines the level of cognitive engagement expected. When several verbs are combined in the same outcome,

the statement quickly becomes compound, mixing different actions and sometimes even different Bloom levels, which makes it difficult to interpret and assess. For instance, an outcome such as “analyse and design control systems” merges two quite different levels of learning, whereas “identify and describe the main components of a mechatronic system” is acceptable because both verbs belong to the same level and reinforce each other.

The comparison of Bloom’s verbs per ILO (Table 44) shows important differences across partner institutions in how outcomes are formulated. At the University of Ljubljana (UNILJ), the average number of verbs per ILO is 1,31, with most outcomes containing only one verb. About a quarter of ILOs (26 percent) contain more than one, which suggests that the majority of outcomes are written clearly and measurably, but there is still a significant share of compound formulations.

	UNILJ	UNSA	UNBG	UoM	UNIRI
courses	35	11	21	11	16
ILOs	190	33	52	57	107
verbs	248	80	75	70	155
average	1.31	2.42	1.44	1.35	1.45
stdev	0.55	1.30	0.72	0.55	0.73
min	1	1	1	1	1
max	4	6	4	3	4
median	1	2	1	1	1
ILOs with > 1 Bloom verb	50	24	17	16	35
perc. of ILOs with > 1 Bloom verb	26.32	72.73	32.69	30.77	32.71

Table 44: Statistics on Bloom verbs in ILOs

The University of Sarajevo (UNSA) is an outlier in this respect. Here, the average number of verbs per ILO is 2,42, the highest among all partners, and nearly three quarters of outcomes include more than one verb. The median is two verbs, showing that compound outcomes are the norm rather than the exception. This indicates a tendency to overload outcomes with multiple actions, which risks reducing their clarity and making assessment more difficult.

At the University of Belgrade (UNBG), the average is 1,44 verbs per ILO, with about one third of outcomes containing more than one verb. This places Belgrade in the middle of the group: while many ILOs are well defined with a single verb, a notable proportion still combine two or more.

The University of Montenegro (UoM) shows a similar pattern, with an average of 1,35 verbs per ILO. Roughly 31 percent of the outcomes contain more than one verb, which indicates that while most outcomes are concise, nearly one in three are compound.

The University of Rijeka (UNIRI) also falls into this middle group, with an average of 1,45 verbs per ILO. Here, about 33 percent of outcomes contain more than one verb, and while most outcomes remain well-structured, a significant share still includes multiple actions.

Overall, the results show three distinct tendencies. UNILJ, UoM, UNBG, and UNIRI generally adhere to the recommended practice of keeping ILOs focused on a single measurable action, but each still shows around one third of outcomes with multiple verbs. UNSA stands out as an exception, where the overwhelming majority of outcomes are compound, often including two or more Bloom verbs. From a quality perspective, this suggests that while most partners show a reasonable level of alignment with best practice, harmonization across the consortium would benefit from encouraging more consistent use of single-verb formulations, particularly in Sarajevo where the overuse of multiple verbs is most pronounced.

4. Conclusion

Across the consortium the picture is clear: partners deliver a solid classical core - supervised learning and its standard tool-box are near-universal; data preparation and model evaluation are widely taught - yet coverage thins as soon as we move to modern paradigms, deployment engineering, and real-world verticals. Deep learning and reinforcement learning are entering the mainstream but remain uneven; transfer, self-/semi-supervised and online learning are rare. Practical software fluency is strong at the NumPy/pandas level, while production-grade workflows (Docker, MLflow and similar) are scarcely present. Application content concentrates in a few niches (notably embedded/real-time AI), and responsible-AI themes are inconsistently embedded, with systematic treatment only at one partner. These findings set a clear agenda for our MOOCs.

The MOOCs should therefore be designed to complement, not duplicate, existing strengths while closing the most visible gaps. A coherent sequence that starts from “modernising the core” and progresses to “deployable and responsible AI” will add the greatest value: concise modules that introduce transfer and self-supervision, a hands-on MLOps strand that takes models from notebook to service, and application-driven units where learners solve problems in manufacturing and operations (predictive maintenance, digital twins, quality control). Parallel to this, explicit training in explainability, fairness, privacy and compliance should be woven through the learning path rather than parked at the end, so that evaluation becomes as much ethical and societal as it is technical.

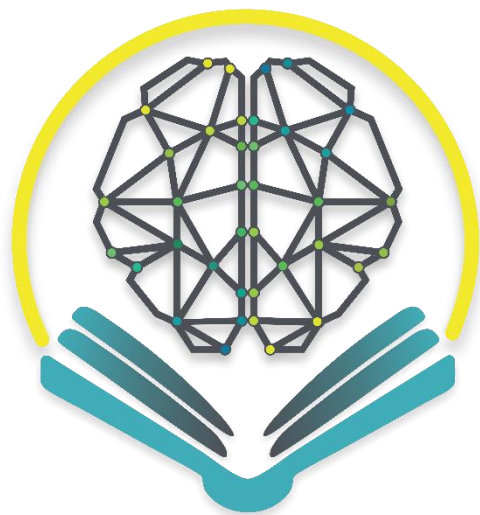
The ILO review points to a second, equally important task: tightening outcome design so that learning achievements are easy to teach, assess and compare across institutions. Syllabi presently use heterogeneous blocks - many in unstructured lists or free text; only a minority tie outcomes to knowledge, skills and competences - so the MOOCs should model better practice. Each course page should expose a single, clearly labelled “Intended Learning Outcomes” section that follows a standard pattern and is consistently categorised. Where outcomes mix several actions, they become hard to classify and assess; in our analysis we had to adopt a “dominant-verb” rule precisely because compound ILOs are common. The MOOCs should avoid this by design.

For the MOOCs themselves, outcomes should be written as short, measurable statements beginning with a single Bloom verb, followed by the object and, where useful, the condition or performance criterion: “By the end of the module, a successful learner will be able to ...”. One verb per outcome keeps the cognitive target and the assessment method aligned; additional verbs should signal additional outcomes rather than being chained inside the

same sentence. Across a module, verbs should map a deliberate progression—typically from Apply and Analyse toward Evaluate and, where project work is involved, Create—so that the assessment suite naturally ranges from implementation and diagnosis through to judgement and design.

In terms of count and granularity, a small set of well-formed outcomes is preferable to long enumerations. For compact, integrative MOOCs, four to six broad ILOs often suffice; for more technical, lecture-style modules, a few more specific outcomes can be justified provided each can be assessed distinctly. Either way, the emphasis should be on coverage of knowledge, skills and competences rather than on meeting a numeric quota, and on phrasing that is concise enough to be understood at a glance yet precise enough to anchor rubrics and evidence.

Finally, outcome content should mirror the curricular gaps we aim to close. At minimum, every MOOC should commit to one outcome on responsible AI (e.g., evaluating fairness, privacy or sustainability implications of a model), one on robust evaluation and explainability, and one on deployment or reproducibility (packaging, versioning, or monitoring a model). Application-facing MOOCs should promise demonstrable competence in an industrial scenario rather than generic familiarity. Taken together, these choices will move the consortium from a strong classical base to a contemporary, deployable and accountable AI offering—and, just as importantly, they will make our outcomes legible to students, teachers and external reviewers alike.



Teaching AI

Lead Partner



UNIVERZA V LJUBLJANI
University of Ljubljana

Partners



УНИВЕРЗИТЕТ У БЕОГРАДУ
UNIVERSITY OF BELGRADE



Univerzitet Crne Gore

