Course Syllabus
Artificial Intelligence Workshop

A. General Background

1. **Academic Unit**	Faculty of Engineering							
2. **Degree Program**	Civil Engineering in Computer Science and Technological Innovation							
3. **Course Code**	IIT414W							
4. **Placement in the Curriculum**	4th Year, 13th Trimester							
5. **Credits**	UDD	10		SCT		6		
6. **Course Type**	Mandatory	X	Elective		Optional			
7. **Duration**	Bimonthly		Semester		Annual		Other	X
8. **Weekly Modules**	Theoretical Classes	2	Practical Classes	2	Teaching Assistant Sessions		--	
9. **Academic Hours**	Classes	48	Teaching Assistant Sessions	48	Other hours per full academic period			
10. **Prerequisite**	Numerical Methods, Business Intelligence, B1+ Level English							

B. Contribution to the Graduate Profile

The course **Artificial Intelligence Workshop** (**IIT414W**) belongs to the advanced cycle of the Civil Engineering in Computer Science and Technological Innovation program (ICIIT) at Universidad del Desarrollo. It is offered in the 13th trimester (4th year) and directly addresses the core competencies of the graduation profile, which describes a professional with:

- Solid knowledge of basic sciences, engineering fundamentals, and emerging technologies including artificial intelligence.
- A highly experiential and practical training approach, developed through hands-on project work with industry partners.
- The capacity to solve complex problems using state-of-the-art technologies to improve quality of life, generate business value, and enhance organizational productivity.
- The ability to lead creative, collaborative, interdisciplinary, and agile processes with concrete and measurable impact.
- A global, innovative, entrepreneurial, and ethical perspective in the design and management of technology projects.

IIT414W operationalizes these graduate profile attributes through a project-based, studio-mode pedagogy in which students design, build, evaluate, and defend an end-to-end ML system applied to a real-world context. The course fosters scientific discipline, responsible AI use, and clear technical communication — all essential capabilities for the ICIIT graduate profile.

C. Competencies and General Learning Outcomes Developed by the Course

Generic Competencies	General Learning Outcomes
Critical Thinking Digital Transformation Communication Ethical Commitment	**GLO1 — Reproducible ML Systems**: Design and implement reproducible ML pipelines that include documented data splits, fixed random seeds, explicit baselines, and metric justification; apply leakage-control techniques and produce minimal runbooks to reproduce results.
Specific Competencies	
Problem Solving under a Systemic Approach Information Systems Design Continuous Learning	**GLO2 — Experimental Evaluation**: Select and train supervised learning models (regression and classification), conduct ablation studies, perform structured error analysis (top failure modes), and compare models under identical validation conditions. **GLO3 — Interpretability and Communication**: Apply feature-importance and SHAP-based interpretability techniques; communicate model behavior and limitations to both technical and non-technical audiences through structured technical memos and oral defenses. **GLO4 — Critical AI Use**: Evaluate LLM-generated outputs critically by documenting the prompts used, modifications made, verification steps performed, and limitations identified, attaching an AI Usage Appendix to every deliverable. **GLO5 — Responsible AI in Practice**: Integrate AI tools ethically into an end-to-end project by addressing potential biases, privacy considerations, and documentation standards; reflect on societal implications of deployed systems and practice transparent, honest scientific reporting.

D. Content Units and Learning Outcomes

Content Units	Competency	Learning Outcomes
Unit I — ML Fundamentals, Reproducibility & Ethics Wk 1: Course orientation; AI-as-amplifier philosophy; reproducible environment setup (Python 3.10+, Git, Jupyter); Project data ecosystem. Wk 2: Decision-oriented EDA; data quality diagnosis; train/validation/test rationale; team formation. Lab 0 due (reproducible setup). Wk 3: Temporal validation; leakage taxonomy; explicit baselines; metric-choice justification. Lab 1 due (EDA + baseline).	*Critical Thinking* *Ethical Commitment* *Continuous Learning*	**LO1**: Configure a reproducible machine learning environment, implementing version control, dependency management, and fixed random seeds. **LO2**: Evaluate data quality and design appropriate train/validation/test splits to prevent data leakage and ensure robust temporal validation.
Unit II — Supervised Models, Pipelines & Feature Engineering Wk 4: sklearn pipelines; feature engineering for tabular/temporal data; Certamen 1 (90 min, in-class). Lab 2 due (feature engineering + improved baseline). Wk 5: Regression and classification models; probability calibration; threshold decision-making. Mini-Challenge 1 checkpoint. Wk 6: Ensemble methods (RF, GBM); hyperparameter tuning; Control Aplicado (60 min, in-class — documented prompt engineering case).	*Digital Transformation* *Problem Solving under a Systemic Approach* *Information Systems Design*	**LO1**: Build end-to-end machine learning pipelines using supervised algorithms to solve complex regression and classification tasks. **LO2**: Apply advanced feature engineering techniques to tabular and temporal data to improve baseline model performance and decision-making thresholds.
Unit III — Ensembles, Interpretability & Robustness Wk 7: Temporal feature engineering; contextual robustness testing; Mini-Challenge 2 checkpoint. Wk 8: SHAP-based interpretability; feature importance; technical writing for non-technical	*Critical Thinking* *Communication* *Information Systems Design*	**LO1**: Optimize ensemble machine learning methods (such as Random Forests and Gradient Boosting) through rigorous hyperparameter tuning and ablation studies to maximize predictive performance. **LO2**: Evaluate model robustness by applying temporal feature engineering and contextual testing,

audiences. Lab 3 due (model comparison + technical memo). Wk 9: Consolidation; ablation studies; Certamen 2 (90 min, in-class). Mini-Challenge 3 checkpoint.		ensuring the algorithm's reliability across different real-world scenarios. LO3: Interpret complex model predictions and internal behaviors by calculating and analyzing feature importance and SHAP (SHapley Additive exPlanations) values. **LO4**: Compose structured technical memos that effectively translate complex algorithmic results and limitations into accessible, actionable insights for non-technical stakeholders.
Unit IV — Capstone Project Wk 10: Capstone kickoff; experimental design; problem framing; Hito 1 due (framing + baseline + experiment plan, frozen dataset). Wk 11 (Last week of classes): Iteration; ablation studies; experiment logging; Hito 2 due (mid-point evaluation + error analysis). In-class demo prep and peer feedback. Final report submission deadline: Sun, 23:59.	*Critical Thinking* *Communication* *Problem Solving under a Systemic Approach* *Ethical Commitment*	**LO1**: Frame an open-ended real-world problem as a structured machine learning task, establishing rigorous experimental designs, frozen datasets, and explicit baselines. **LO2**: Diagnose model weaknesses by conducting systematic error analysis on the top failure modes, using these insights to iteratively improve the system's performance. **LO3**: Author a comprehensive and fully reproducible technical report that justifies all architectural decisions, metric choices, and the software engineering practices applied. **LO4**: Defend the final applied ML system in an oral pitch, critically addressing its ethical implications, potential biases, and the documented use of generative AI tools during its development.
Exam Week — Final Examination & Demo Day Demo Day — Team oral defense (7-min pitch + 5-min individual Q&A).	*Critical Thinking* *Digital Transformation*	**LO1**: Synthesize end-to-end machine learning concepts—from data validation to model evaluation—to independently architect and

Capstone Final Report already submitted. Final Examination (4 hours, individual, with AI access — integrative applied case on a frozen dataset).	*Communication* *Ethical Commitment* *Problem Solving under a Systemic Approach* *Information Systems Design* *Continuous Learning*	implement a solution for an unseen applied case under time constraints. **LO2**: Defend the technical, methodological, and ethical decisions of the capstone project in a formal oral presentation, demonstrating individual accountability and effective professional communication before an evaluation panel. **LO3**: Justify the critical use of generative AI tools and reproducible engineering practices during high-stakes problem-solving scenarios, ensuring transparency and methodological rigor in the final deliverables.

E. Teaching Strategies

The course is delivered through a Studio Mode methodology — a lab-style learning environment in which each 2.5-hour session combines a short theory block (approximately 20–30 minutes), a guided practice activity, and a structured reflection. The following strategies are applied throughout the trimester:

- **Active Learning and Studio Mode:** Students engage in hands-on coding, experimentation, and peer discussion during every class. Theory is introduced in short blocks (20 min max), immediately followed by application. Socratic questioning and exit tickets promote metacognitive reflection.
- **Project-Based Learning (PBL):** The capstone project is the integrating thread of the course. Students work in teams of 2–3 across staged milestones (Hito 1, Hito 2, Final Report, Demo Day), accumulating a reproducible GitHub repository throughout the trimester.
- **Assessment for Learning:** Frequent formative feedback is embedded in the course through Lab rubrics with explicit criteria, in-class checkpoints (Mini-Challenges 1, 2, and 3), and written comments on all graded work within one week of submission. Students receive rubrics before each deliverable.
- **Scientific Discipline:** Every graded deliverable involving experiments must include fixed random seeds, explicit train/validation/test rationale, at least one baseline, metric-choice justification, structured error analysis (minimum: top-3 failure modes), compute and library version notes, and a reproducibility runbook.
- **Responsible and Transparent AI Use:** AI language models are treated as a professional tool, not a substitute for student thinking. Students are required to attach an AI Usage Appendix to every submission, documenting: prompts used, outputs accepted or rejected, modifications made and

why, verification steps taken, and identified limitations. Submitting AI-generated content as original thinking is an academic integrity violation.

- **Contextualized Learning:** A particular domain provides a consistent real-world context throughout the course, offering rich time-series data, domain constraints, and an authentic prediction task that mirrors industry ML pipelines.
- **Technical Communication:** Students practice writing for both technical and non-technical audiences through technical memos (Lab 3), a final written report (8–12 pages), and an oral defense with individual Q&A accountability (Demo Day).

F. Assessment Strategies

For each assessment, clear criteria are established and communicated to students before the due date via Canvas. All evaluations are individual unless explicitly marked as team assessments. AI tool access is permitted in all assessments but must be documented through the AI Usage Appendix. The course includes two midterm examinations, one applied quiz, four lab assignments, a team capstone project with four milestones, and one final examination.

F.1 Grading Scheme

Assessment Component	Weight in NP	Notes
Midterm Examination No. 1	18%	90 min · In-class · Individual · Open AI access
Midterm Examination No. 2	18%	90 min · In-class · Individual · Open AI access
Applied Quiz	15%	60 min · In-class · Individual · Documented prompt engineering
Lab Assignments (Lab 0 + Lab 1 + Lab 2 + Lab 3)	24%	Individual/pairs · Reproducible notebooks · AI Usage Appendix required
Capstone Project (Hito 1 + Hito 2 + Final Report + Demo Day)	25%	Teams of 2–3 · Frozen F1 dataset · Peer evaluation included
Coursework Grade (NP)	**70%**	NP = weighted sum of all components above
Final Examination (EF)	**30%**	Wed May 20 · 4 hours · Individual · Integrative case · Open AI access
Final Course Grade (NF)	**NF = 0.70×NP + 0.30×EF**	Minimum grade of 3.0 required on the Final Examination to pass.
Required Attendance	70%	Attendance tracked each session. Below 70% may result in course withdrawal.

F.2 Capstone Project Component Breakdown

Capstone Sub-component	Weight in Capstone
Hito 1 — Problem Framing + Baseline + Experiment Plan	20%
Hito 2 — Mid-point Evaluation + Error Analysis	20%
Final Report — Reproducible Report + GitHub Repository	40%
Demo Day — Oral Defense (team pitch + individual Q&A)	20%

F.3 Quizzes and Midterm Examination Policy

- Two midterm examinations are administered during scheduled class time (90 min each). Format: applied case study — critical analysis, decision justification, and documented AI use.
- One applied quiz focuses on documented prompt engineering and iterative AI collaboration: students work through a methodological case and submit a maximum 2-page process log.
- No quiz elimination rule applies, as the course has fewer than five quiz-type assessments in total.
- In the event of justified absence from a certamen or quiz, the student must inform the Faculty of Engineering before the assessment. Make-up assessments are governed by Faculty of Engineering regulations.

F.4 Attendance Policy

- A minimum of 70% attendance is required. Students falling below this threshold may be subject to course withdrawal in accordance with Faculty regulations.
- Classes will not be suspended during certamen weeks, as all evaluations occur within the regular class session block.

F.5 Late Submission Policy

- Lab assignments and capstone milestones submitted late will be penalized 1.0 grade point per 24-hour period of delay, up to a maximum of 3 days. Submissions more than 3 days late receive a grade of 1.0.
- Extensions may be requested at least 48 hours before the deadline by emailing the professor, subject to approval.

F.6 Academic Integrity and AI Use Policy

- All submissions must be the student's own intellectual work. The AI Usage Appendix (attached to every deliverable) must include: (1) representative prompts used, (2) outputs accepted or rejected, (3) modifications made and rationale, (4) verification steps and tests, (5) identified limitations and failure modes.
- Prohibited: Submitting AI-generated content as original thinking; fabricating results, data, or citations; omitting the AI Usage Appendix; sharing assessment materials outside the course; plagiarism in any form.

- Violations will be handled in accordance with the University's academic integrity regulations and may result in a grade of 1.0 for the assessment and/or disciplinary action.

F.7 Reproducibility Standards

Every deliverable involving experiments must include all of the following:

- Fixed random seeds (numpy, sklearn, torch if applicable) stated explicitly in code.
- Train/validation/test separation rationale (especially temporal splits to prevent leakage).
- At least one explicit baseline for comparison.
- Metric-choice justification linked to the prediction task and business objective.
- Structured error analysis (minimum: top-3 failure modes with examples).
- Compute and environment notes: Python version, library versions (requirements.txt or environment.yml).
- A minimal runbook — step-by-step instructions to reproduce results from a clean environment.

Deliverables that do not meet reproducibility standards will receive a reduction of up to 20% on the technical quality dimension of the rubric.

G. Learning Resources

Required Bibliography:

- Burkov, A. (2019). The Hundred-Page Machine Learning Book. Andriy Burkov. Available at: http://themlbook.com/
- VanderPlas, J. (2016). Python Data Science Handbook. O'Reilly Media. Open access: https://jakevdp.github.io/PythonDataScienceHandbook/
- Molnar, C. (2022). Interpretable Machine Learning (2nd ed.). Open access: https://christophm.github.io/interpretable-ml-book/
- FastF1 Documentation (v3.x). https://docs.fastf1.dev/
- Jolpica F1 API Documentation. https://github.com/jolpica/jolpica-f1

Supplementary Bibliography:

- Géron, A. (2022). Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow (3rd ed.). O'Reilly Media.
- Scikit-learn Documentation — User Guide: Pipelines and Feature Engineering. https://scikit-learn.org/stable/modules/pipeline.html
- Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. NeurIPS 2017. SHAP library: https://shap.readthedocs.io/
- Kapoor, S., & Narayanan, A. (2023). Leakage and the Reproducibility Crisis in ML-based Science. Patterns, 4(9). https://doi.org/10.1016/j.patter.2023.100804
- Anthropic. (2025). Claude API Documentation — Prompt Engineering Guide. https://docs.anthropic.com/
- OpenAI. (2024). Best Practices for Prompt Engineering. https://platform.openai.com/docs/guides/prompt-engineering