

Module 10: Capstone & Ethics

Data Science with Python – BSc Course

Why Ethics in ML?

Apple Card was investigated for gender-biased credit decisions in 2019. Amazon scrapped an AI recruiting tool that penalized women. Goldman Sachs faced scrutiny when their algorithm offered women lower credit limits. With great power comes great responsibility — and in finance, irresponsible AI has legal consequences.

The EU AI Act classifies credit scoring as high-risk. Regulators require explainability for model decisions. Bias scandals destroy reputations. The message is clear: technical excellence without ethical awareness is not enough.

Responsible AI is not optional – it is a legal and ethical requirement

- **Legal risk:** EU AI Act classifies credit scoring as high-risk; non-compliance has penalties
- **Reputational damage:** Bias scandals destroy trust and hurt business
- **Fairness requirements:** Equal treatment regardless of gender, race, or age is legally mandated
- **Explainability demands:** Regulators require you to explain model decisions to affected parties

Ethics is not a nice-to-have – it is a core competency in modern ML

By the end of this module, you will be able to:

- Execute a complete data science project from EDA to deployment
- Present technical findings to a non-technical audience effectively
- Identify and mitigate algorithmic bias in ML models
- Understand regulatory requirements for AI in finance (EU AI Act, SR 11-7)
- Implement fairness metrics and explainability techniques

From technical skills to professional competence and ethical awareness

Lesson Roadmap

Lesson	Topic	Focus
L45	Project Work Session 1	EDA, feature engineering
L46	Project Work Session 2	Modeling, evaluation
L47	ML Ethics in Finance	Bias, fairness, explainability
L48	Final Presentations	Communication, peer review

Your group project is where everything comes together

- **End-to-End Project Workflow** – From raw data to deployed model and presentation
- **Technical Presentation Skills** – Communicate complex findings to non-technical stakeholders
- **Algorithmic Bias & Fairness** – Detect and mitigate unfair treatment across groups
- **Explainability & Regulatory Compliance** – Meet legal requirements for model transparency

Professional data scientists balance technical excellence with ethical responsibility

Scenario: Full ML Project

Your group will:

- Choose a finance dataset (credit risk, stock prediction, portfolio optimization)
- Perform exploratory data analysis and feature engineering
- Build and evaluate multiple models
- Check for bias and implement fairness metrics
- Present findings to the class

This is exactly what data science teams do in industry: end-to-end projects with stakeholder communication and ethical review.

Your group project is your portfolio piece – make it count

Who Uses This?

- **Regulators:** EU AI Act and Fed SR 11-7 model risk guidance mandate ethical AI practices
- **Legal Cases:** Apple Card gender bias investigation shows real-world consequences
- **Industry Standards:** IEEE Ethics in AI and ACM Code of Ethics guide responsible development
- **Best Practice:** Google's Model Cards and Microsoft's Responsible AI framework are industry benchmarks

Ethics is embedded in the regulatory and professional landscape

The Journey Ahead

You've learned Python, statistics, ML, deep learning, NLP, and deployment.

What to do next:

- **Build your portfolio** – Put your group project on GitHub
- **Keep learning** – Kaggle competitions, open-source contributions
- **Specialize** – Quant finance, fintech, data engineering, research
- **Stay current** – Follow AI/ML developments in finance

This course gave you the foundation – your career builds on it

Let's Finish Strong!

Your final project begins now.

Show us what you've learned.