



LAB 3

LAB 3 — MACHINE LEARNING ON DATA REPRESENTATIONS

How much performance comes from the model —
and how much from the representation?

PRESS **←** OR SPACE. **→**

The Task

- Classify land cover from geospatial data assembled in Lab 2
- Work with point-based samples from raster/geospatial layers
- Use one dataset consistently across all experiments
- Evaluate predictions on held-out data

Implementation is yours: Python in **Google Colab**, **VS Code**, or local Jupyter. CPU is enough.

Experimental Questions

Q1

Does feature engineering help?

Q2

Does complementary information help?

Q3

When do we need a nonlinear model?

Treat this as an experiment: hold conditions fixed and change one conceptual factor at a time.

Build the Dataset Once

- Choose/extract labeled sample locations
- Build one table: **rows = locations**, **columns = features**
- Target column: land-cover class
- Create one fixed train / validation / test split

Important: Do the split once and reuse it for every experiment.

Baseline: Logistic Regression

- Use `LogisticRegression` as the linear baseline
- Input: feature vector per sample
- Learns a linear decision boundary
- Fast, simple, and interpretable reference model

Key question: How good is your representation if linear separation is all you allow?

Change the Representation (Model Fixed)

Retrain the same `LogisticRegression` while changing features:

1. Small raw feature set (e.g., Red + NIR)
2. Raw features + engineered feature (e.g., NDVI)
3. Several Sentinel-2 bands
4. Bands + complementary Lab 2 variables (if available)

Prompt: What changes when the model stays fixed but the representation changes?

Compare with Random Forest

- Run `RandomForestClassifier` on the same representations
- Keep the exact same train/validation/test split
- Compare against logistic regression fairly

Logistic Regression

Linear decision boundary

Random Forest

Nonlinear decision boundaries

Ablation Table

Representation	Logistic Regression	Random Forest
Raw baseline	...	...
+ NDVI	...	...
+ additional bands	...	...
+ complementary features	...	...

Use accuracy (and optionally macro-F1). Focus on interpreting changes, not maximizing one score.

Interpret the Experiment

- Which representation change helped most?
- When did logistic regression approach random-forest performance?
- What does the model gap suggest about feature-space geometry?

Caution: More dimensions are not automatically better. Useful, complementary information matters.

Takeaway → Lecture 4

1. Feature engineering changes the representation.
2. Complementary information can make classification easier.
3. A strong representation can make simple linear models surprisingly powerful.

NEXT

**In Lab 3 you designed representations manually.
In Lecture 4: what if the model learns the representation itself?**