Group project: hand gesture recognition

Introduction to Machine Learning, KU Leuven, academic year 2026-2027

Task

In pairs, you build classifiers that recognise 18 static hand gestures from photos. At least one of your models works on features that you design yourselves, and at least one learns its own features from the pixels. The project asks which features carry the gesture, and how close designed features come to learned ones.

Data

Each image is a photo cropped around one hand, in colour and at its original resolution. The training set holds 25,416 crops of 6,399 people, each with its gesture and the person who made it. The test set holds 6,356 crops of 1,569 other people, without labels. Crops whose shorter side is below 24 pixels were left out.

The function `load_gestures` of the course package loads both sets, and the starter notebook shows how. Several crops come from the same person, so a fair validation keeps each person on one side of a split, for instance with `GroupKFold` and the `groups` returned by `load_gestures`.

The crops are cut from the HaGRID sample of 30k images at 384p by C. Mills, a sample of HaGRID by A. Kapitanov, K. Kvanchiani, A. Nagaev, R. Kraynov and A. Makhliarchuk. They are published under the licence CC BY-SA 4.0, and your report credits both sources.

Rules

- Train at least one random forest or logistic regression on features that you design and compute from the pixels.
- Train at least one neural network, an MLP or a convolutional network, on the pixels.
- A third model from the course is optional.
- Pretrained networks, external data and hand landmark detectors are not allowed, since each would replace the features that the project asks you to design.
- Estimate the accuracy of each model by validation on the training set, split by person. The test set is used once, for the predictions you hand in.

What you hand in

One zip file per group on Toledo, by 16 December 2026 at 23:59. It holds:

- the report as a PDF of at most 6 pages plus appendix, written from the LaTeX or Word template below, with the sections data and split, features, models and tuning, results, and discussion;
- the notebooks or scripts that produce your numbers;

- exactly one file named `predictions.csv`, at the top level of the zip, holding the predictions of your final model, written by `write_submission` and passed by `check_submission`.

You receive the test accuracy of your predictions on the morning of 17 December 2026. Your slides show it next to your own validation estimate.

Presentation

Each group presents on 17 December 2026, for 7 minutes, followed by 3 minutes of questions.

Grading

Criterion	Weight
Features: designed with reasons, and tested for what each adds	30%
Evaluation: validation split by person, and the test score compared with the validation estimate	25%
Models and tuning: choices justified, and models compared on equal terms	20%
Discussion: the errors analysed, such as which gestures are confused and why	15%
Presentation: clear, and within time	10%

The test score itself does not enter the grade.

Material

- Starter notebook, which loads the data and takes a baseline to a checked predictions file
- Report template: LaTeX, Word
- This brief as a PDF