hugomcp@di.ubi.pt, 2024/25

# Machine Learning

## Practical Sheet 6: Self Supervised Learning

Consider the "AR.zip" dataset, available at the course web page. It contains 3,315 images, of 136 subjects, each one represented in the RGB color space and having dimensions 576 (rows) x 768 (columns).

1. Implement a Python function that receives one input image and returns a partially occluded version of that image (by a black rectangle)

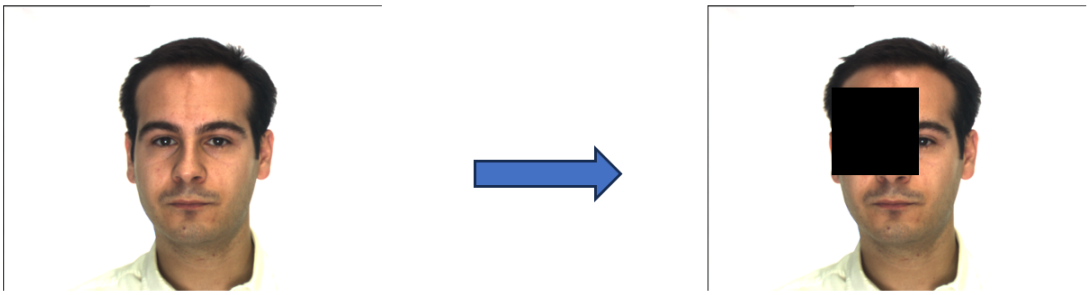

2. Learn a deep learning-based model, based in an auto-encoder architecture, that receives partially occluded images and returns a prediction of the corresponding un-occluded versions.

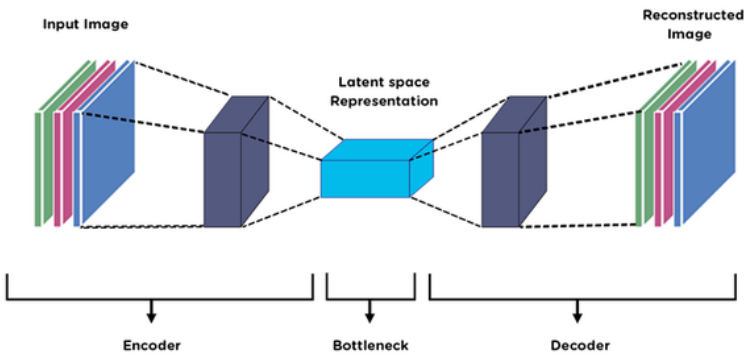

3. For each image in the learning/test set, obtain the corresponding feature representations provided by the “Bottleneck” layer of the model learned in 2).

4. Use t-SNE (t-distributed Stochastic Neighbor Embedding (t-SNE) to obtain 2D representations of the features previously found.

5. Analyze how these features/instances spread in the 2D feature space, according to:

    a. Gender;

    b. ID;

    c. Other features (e.g., Glasses, facial expressions,…)