

Machine Learning

Practical Project 3 Convolutional Neural Networks



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).

Male images are stored as: “[M-xx-yy.raw.jpg](#)”, while Females as: “[F-xx-yy.raw.jpg](#)”

Here, 'xx' is a unique person identifier (from "00" to "76" for males and from "00" to "60" for females) and 'yy' specifies the features of each image, that were collected in two different sessions (1..13 and 14..26):

- 1 : Neutral expression
- 2 : Smile
- 3 : Anger
- 4 : Scream
- 5 : left light on
- 6 : right light on
- 7 : all side lights on
- 8 : wearing sun glasses
- 9 : wearing sun glasses and left light on
- 10 : wearing sun glasses and right light on
- 11 : wearing scarf
- 12 : wearing scarf and left light on
- 13 : wearing scarf and right light on
- 14 to 26 : regard second session (same conditions as 1 to 13)

Starting from the “[scr_deep_learning_keras.py](#)” script, adapt it to:

1. Learn from large datasets, using the `train_in_batch()` function, or feeding the data from data generators.
2. Use different network architectures, either created by you or well-known structures (e.g., VGG, ResNet, Inception or others)



3. Divide the available datasets into learn/validation and test subsets, in order to create models that predict:
 - a. ID;
 - b. Facial expression;
 - c. Gender;
 - d. Glasses.

4. Compare the effectiveness attained by each architecture in each task, when using the “learning from scratch” and “fine tuning” paradigms.