hugomcp@di.ubi.pt, 2024/25

Machine Learning

Practical Project 2
Convolutional Neural Networks

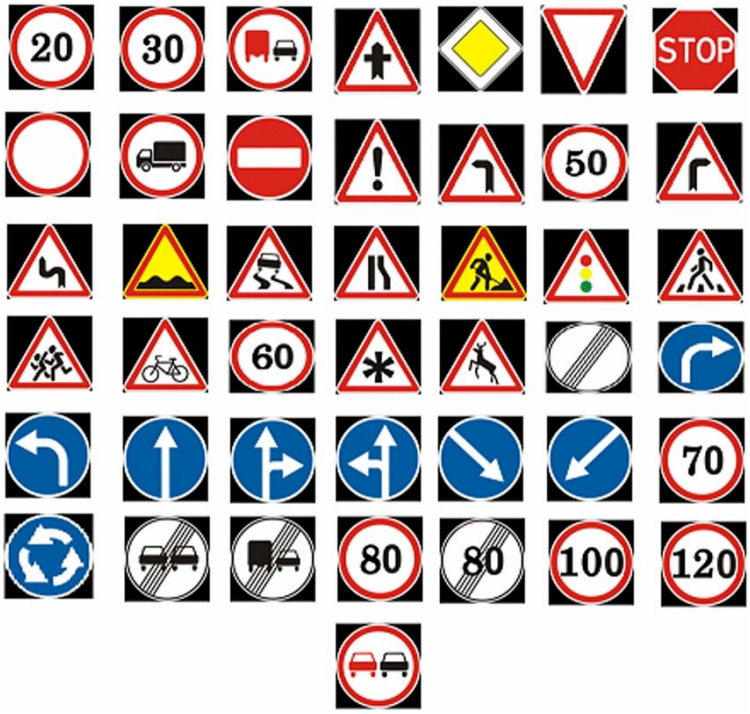

Consider the "GTSRB.zip" dataset, available at the course web page. This is a widely used image dataset designed for traffic sign classification and recognition tasks in computer vision and machine learning.
It was created to support the development and evaluation of algorithms for automatic traffic sign recognition, a key component in intelligent transportation systems and driver assistance technologies.

- Total images: More than 50,000 color images.
- Number of classes: 43 different types of traffic signs.
- Image resolution: Varies between 15×15 and 250×250 pixels — images are not all the same size.
- Format: RGB color images (usually in PNG or PPM format).
- Conditions: Captured under real-world conditions, with variations in:
    - Lighting
    - Weather
    - Viewing angle
    - Partial occlusion
    - Motion blur

The dataset is divided into two main folders:

- **Training**. It is divided into 43 sub-folders, each one with all instances of the corresponding class;

- **Test**. It contains 12,629 images for test purposes. The gorund truth is provided in the "GT-final_test.csv" file.

1. Create a Keras + Python CNN neytwork from scratch, that distinguishes between the 43 classes (traffic signs).

2. Compare the effectiveness attained by your neural neteork, against a fine tuned version of different well known architectures (e.g., ResNet, MobileNet, Inception,…)