

Deep Learning for Social Sciences - Assignment: Facial Expression Classification

Giordano De Marzo

2 June 2024

To be handed in by 24.6.2024 23:59 at <https://github.com/orgs/DLSS-24/>.

Overview

In this assignment, you will use a Convolutional Neural Network (CNN) to classify human emotions. Follow the steps outlined below to complete this assignment. This is a challenging problem, where also humans can make mistakes. The dataset contains a very small class and is thus unbalanced. This assignment is closer to a real life application of deep learning so don't expect a perfect accuracy.

Complete tasks for receiving points, the maximum number of points you can get is 30. You can get bonus points completing the optional tasks. Discuss the bonus tasks in the report, you have up to 1.5 additional pages for the bonus tasks. Bonus points will not increase your mark over 30, but they can increase it up to 30 if you made some errors in the mandatory tasks.

Tasks

1. Data Import, Visualization and Preparation (10 points)

- Register on Kaggle and download the dataset `icml_face_data.csv` from https://www.kaggle.com/competitions/challenges-in-representation-learning-facial-expression-recognition-challenge/data?select=icml_face_data.csv

- Preprocess the data and visualize some images with the corresponding labels to be sure everything is working
- Use the public test as validation and the private test as test set. This means that you don't have to split the data in train test and validation, this has already been done by the creators of the dataset

2. CNN Model Building, Training and Testing (10 points)

- Build a CNN to classify the images. An accuracy in the range 50% to 60% is a good result. See the leaderboard at <https://www.kaggle.com/competitions/challenges-in-representation-learning-facial-expression-recognition-challenge/leaderboard>
- Train and test the model using all the tricks you know for getting good performances (for instance data augmentation, early stopping, regularization, batch normalization, dropout ...)

3. Report Writing (10 points)

- Write a report (max 3 pages) detailing the procedure you followed and the architecture you choose. Report all the challenges you encountered and the steps you followed to solve them.
- The report must contain the learning curves of your best model both on the training and validation set.
- Clearly report the performances of the model on the test set (accuracy, F1 Score, classification report).
- A person must be able to replicate all your analysis by only reading your report.

Bonus Tasks

Bonus Task 1: MLP (3 bonus points)

- Train a MLP to classify the images and compare its performances with the CNN.

Bonus Task 2: CNN Features visualization and discussion (3 bonus points)

- Visualize the activations of the last convolutional layer and discuss them. Show some relevant images in the report.

Bonus Task 3: Comparison with human performances (3 bonus points)

- Select 50/100 random images from the test set and try to label them by hand without looking at the ground truth.
- Run the CNN on the same small subset of images and look at its performances. Do you perform better than the CNN?
- Upload the hand labeled dataset to github.