Computer Vision
James L. Crowley and Nachwa Aboubakr

Face Detection with a sliding window detector at multiple scales.

The objective for this exercise is to use your best MLP face detector constructed last week to detect faces at multiple scales using windows from a Gaussian Pyramid.   You will first construct a sliding window face detector using your best MLP, and then optimize this detector using a full octave Gaussian Pyramid. Do this with the FDDB and WIDER data sets

This exercise is composed of three parts.
1) Write a program to construct a scale-invariant Gaussian pyramid, using the algorithm shown in section 3.2 of the course notes. Demonstrate the impulse response of your pyramid by creating a 512 x 512 image with a single non-zero pixel at the center position (256x256). Display the contents of central 13 columns (cols 250 to 262) from row 256 from each channel of each level of your pyramid. Do this for $\sigma_0=1$ and $\sigma_0=\sqrt{2}$ and compare the results.
2) Use the best MLP from project 2 to detect faces at multiple scales from a full octave pyramid with $\sigma_0=1$. Write a program to extract and flatten a sliding window from an image over a range of sizes from 16x16 to 100x100 using a scale factor of 1.2. Each window should be transformed to the standard size of input vector for your MLP face detector.   Report precision, recall, F1 and computing time for evaluation with the images in folds 9 and 10 of FDDB and WIDER.
3) Use the same MLP to detect faces from each level of your pyramid over a range of sizes from 16x16 to 33x33 using a scale factor of 1.2. Each window should be transformed to the standard size of input vector for your MLP face detector.   Report, precision, recall, F1 and computing time for evaluation with the images in folds 9 and 10 of FDDB and WIDER.

Document your work in the Jupyter Notebook by commenting it and send the .ipynb file to: James.Crowley@inria.fr,   Nachwa.Aboubakr@inria.fr, Yangtao.Wang@inria.fr. Results are due before class on Thursday 12 nov.