

10907 Pattern Recognition

Lecturers

Prof. Dr. Thomas Vetter (thomas.vetter@unibas.ch)

Assistants

Dr. Adam Kortylewski (adam.kortylewski@unibas.ch)

Dennis Madsen (dennis.madsen@unibas.ch)

Dana Rahbani (dana.rahbani@unibas.ch)

Exercise 2 — Skin Detection

Introduction 08.10

Deadline **16.10** Upload code to Courses.**15.10+16.10** Group presentations, U1.001, see schedule online

In this exercise you will write a skin detector, which is able to classify pixels of an image as *skin* or *non-skin*. For the following experiments we provide two data files `skin.mat` and `nonskin.mat` containing RGB colour data (format: $3 \times \#samples$). This data represents pixel values from several photographs which were manually labelled as belonging to either *skin* or *non-skin* regions.

Grading: We will ask questions about the exercises listed here. The grading will be 50% coding and 50% theory questions. The questions will be about the required conceptual understanding. The coding part will focus on concrete figures and results. Please appear prepared with the required calculation results readily available (error rates etc.). The total time of discussion has to fit within 15 minutes and each student of the group should contribute.

You can download the data needed for this exercise and a code skeleton from the following repository: <https://bitbucket.org/gravis-unibas/pattern-recognition-2018> . A number of helper functions are provided - your task is to implement the *TODO* parts in the python (.py) files.

1 Maximum Likelihood

Estimate a Gaussian distribution for each dataset (skin + non-skin). You should use a full covariance matrix. Use the maximum likelihood estimators known from the lecture. The *main* function for this exercise is found in the file `ex2-MVND.py`.

1.1 Detection

- Multivariate Normal Distribution** Implement a MVND class. The general structure of the class is found in the file `myMVND.py`.
Hint: The MVND class takes as input a list of Gaussian distributions. In this exercise, only one list item will be given to the function. In Section 2, when you implement the GMM, you shall provide multiple list items.
- Prior Probabilities** Estimate the prior probabilities for each classification-class (skin & non-skin). The prior probabilities are estimated from the image `image.png`. The skin label for the image is provided in `mask.png`.
- Likelihood Classifier** Implement a Bayes classifier to distinguish between skin and non-skin. Provide the output as a binary image. In the file `classifyHelper.py` the *likelihood* function need to be finalized. This function will compute the likelihood for each pixel in an image, given a list of distributions.
Finalize the *TODO* parts in the `classifyHelper.py` file - Compute false positives/negatives (FP, FN) and the total error - both with and without the skin-prior. Finally, compute the false positive/negative images.

2 Gaussian Mixture Models

The task is now to extend the model the skin/non-skin classes with a Gaussian Mixture Model instead of modelling with a single Gaussian distribution. The *main* function for this exercise is found in the file `ex2-GMM.py`.

2.1 EM-Algorithm

The function `gmm_em()` needs to be implemented. The EM-Algorithm for Gaussian Mixture Models can be found in the lecture slides "*Density_Estimation_2_GMM*".

You can visualize the progress of the EM-Algorithm by using the function `gmm_draw` (provided in the `ex2-GMM.py` file). After each iteration plot, the data points are coloured according to their current cluster assignment for a visual debugging.

Develop and test your EM algorithm using the two-dimensional toy data provided in the file `gmmdata.mat`, work with $K=3$ components (see function `gmmToyExample()`). To initialize, the algorithm uses a random cluster assignment for each data point. Run the algorithm long enough to converge (check at least graphically) and report the obtained cluster means covariance matrices.

Make sure that your `gmm_em()` function works according to specifications:

- K - Number of GMM components, integer (≥ 1)
- $iter$ - Number of iterations, integer (≥ 0)

2.2 Skin Detection

Train a GMM with two $K=3$ components for each dataset (`skin` and `non-skin`) and use a Bayes classifier for classifying the pixels into skin and non-skin (see function `gmmSkinDetection()`). Report the error rate on the training and the test set. What are the error rates if you use 1, 2, 3 or 4 components?