



Date: March 20, 2014

Concerns: Exercises ImageMiner for ASCI course

Tutors: Ork de Rooij and Harro Stokman

This document contains a series of exercises for which you will need to hand in answers. These exercises can only be performed during the course. You can pair up in teams of 2.

Getting started

- Download the ImageMiner client from <http://z.ork.be/EuVisionImageMiner.zip> and extract this to C:\Temp
- The ImageMiner server is currently set up in the Amazon cloud. As such, there are a limited number of instances available, and we will distribute these evenly amongst you.

Ork will have a list of available IP addresses, ask him which address your team can use.
- Open C:\Temp\EuVision ImageMiner\ascicourse.ImageMinerProject (in notepad) and change the IP address to the one assigned to you.
- Start C:\Temp\EuVision ImageMiner\imageminer.exe and proceed with the questions below.
- Note: The ImageMiner server will only be available during the course. Do answer all questions during the course, there will be no possibility to run the application afterwards.

Finding your way

- On the left hand side of the user interface, browse through different subsets of images, to get a feel of the type of images you'll work on.

- On the right hand side of the user interface, browse through the series of builtin concepts, to get an idea of the images incorporated in the dataset.

Exercises and questions

While answering questions, do motivate your answer.

Create your first concept, e.g. “Black and White”. Annotate a few (1-3) of black and white images as positives by pressing the G key (Good) and annotate a few (1-3) of colorful images as negatives (by pressing False). If you make a mistake, press Delete. Finally, press the Find button.

Note: Pressing the Find button sends your annotations to your server in the AWS cloud, and processing takes place there. This will take a couple of minutes, so use that time to write down answers to questions, such as this one:

1. The ImageMiner is based on Bag of Words / Support Vector Machine algorithms. As such, when pressing the Find button, it executes 8 steps. Describe per step briefly what these algorithms do.

Inspect the results. Find the first false positives and first false negatives.

2. Describe in words what the first false positives and first false negatives look like.

Annotate these. Run the ImageMiner again.

3. Did the results improve?

Create your second concept, and call it “Fireworks. However, give your concept a unique name like “Fireworks-by-Stokman”, to prevent confusion with the other team working on the same Amazon instance. Annotate a series of images to represent the model (positive and negative examples). Train the model. Inspect the classification results.

4. When annotating negative examples for the next round of training, what would your criteria be for the most effective **negative** set? What are the effects if you do? Write this as a clear instruction to someone you do not know.
5. Compare your results after a few rounds of training with the predefined concept Fireworks (under Man-Made Objects). Which results are better? Why?

Create your third concept, and call it “More than one person”. Train the model again and again and again.

6. Does it work? Why (not)?

Download the Impala app (which uses deep learning) on your mobile phone, and let it categorize your own images. Here, the “More than one person” is called “Friends.”

Note: you can get Impala for Android from

<https://play.google.com/store/apps/details?id=eu.euvt.impalamobile>

and Impala for iOS from

<https://itunes.apple.com/us/app/impala/id736620048>

7. Which algorithm (deep learning or BoW/SVM) produces the best results for this concept?

Create your own fourth concept. It can be your favorite product (e.g. sunglasses or mobile phone) of your own choice. The object should be professionally photographed against a white background. Train the concept in a series of rounds.

8. Why did you choose it? What is difficult detecting the concept?

Create your fifth concept: “Car back seat.” To start, use the three images in the carinterior-20130304/crimescene directory for positives. The images depict the same backseat of the same car, in different angles. The backseat is empty. Train the model in a series of rounds to find other images of the same car. Answer the following question:

9. What object(s) are on the backseat of the car?
10. How sure are you of your answer? What do you expect is the recall percentage of your answer?
11. How would you have answered question 10 if you’d used deep learning?

Summarizing

12. What have you learned during the training rounds?
13. From now on, what would be your strategy to acquire your training data?

