CPSC 535
Assignment 6: Pattern Recognition II:
Counting Beans

The goal of this assignment is to develop a simple pattern recognition system that can count the number of different objects in an image. In this case, the objects are a mix of beans, lentils, and grains.

Section 9.2 of Sonka, Hlavac and Boyle covers material on statistical pattern recognition.

1 The Data

For this assignment you get two sets of images.

The first set is the training data. It contains images with only one type of bean/grain, two images of each type. The bean/grain types are:

- barley,
- green lentils,
- green split peas,
- kidney beans,
- lima beans,
- navy beans,
- peas,
- red lentils, and
- yellow split peas.

Use these training images to *teach* your algorithm to recognize the grains and beans.

The second set of data contains two images for testing your algorithm. These images show a mix of the beans and grains that need to be counted.

2 Your Program

Your program should work as follows:

myprog *image-file-name*

The output of your program should be a table indicating the counts of each grain type, e.g.,

```
barley:            25
green lentils:      0
green split peas:  32
kidney beans:       4
lima beans:        39
```

```
navy beans:            10
peas:                   0
red lentils:            1
yellow split peas:      7
total:                118
```

I have no ground truth available.

To test your code, I have two other images like those in the test set for use in evaluating your program.

3   Hand in

Hand in the following:

1. your code,

2. a written description of how your algorithm recognizes the objects, and

3. a description and explanation of any salient observations you make during the development of the algorithm.

You will be graded on the quality of your code, the accuracy of counts, and your written observations.