Statistical image analysis, 7.5hp | **Solutions to written exam**
TMS016/MSA300 | **2018-08-29, 14:00 - 18:00**

---

1. (a) True. Covariance function A has a longer practical range than covariance function B, and Realisation 1 clearly is more slowly varying.
   (b) False.
   (c) False. It is Po(Λ)-distributed, where
   $$\Lambda = \int_0^1 \int_0^1 \lambda(s_1, s_2) ds_1 ds_2 = 3/4.$$
   (d) False. The Markov property does not imply that things far apart are uncorrelated.
   (e) True. The kriging predictor is the best linear unbiased predictor.
   (f) True. It is the special case when the smoothness parameter ν equals 0.5.

2. (a) Let A denote the set of pixels with non-zero values in I. The erosion of A is defined as
   $$A \ominus S = \{(i, j) : S_{(i,j)} \subseteq A\},$$
   where $S_{(i,j)}$ denotes the structure element centered at pixel (i, j). The dialation of A is defined as $A \oplus S = (A^c \ominus S)^c$ where A^c denotes the complement of the set A. Finally, the opening of A is defined as $(A \ominus S) \oplus S'$, where S' is S rotated 180 degrees.
   (b) A common way to remove noise is to compute the image opening with respect to a circular structure element with a radius that is larger than one pixel but smaller than the features that should be kept in the image.

3. (a) For supervised learning, we are given a set of images $X_1, \ldots, X_n$ and a set of labels $z_1, \ldots, z_n$ that decides which class each image belongs to. We use all this information to train the model. For unsupervised learning, we do not have access to the labels $z_1, \ldots, z_n$ when training the model.
   (b) The top-left value (77) shows the number of images which had class 1 and where the classifier correctly labelled the images as class 1. The top-mid value (39) shows the number of images where the correct class was 2 but where the classifier labelled the images as class 1. The top-right value shows the percentage of correctly labelled images among those that were labelled as class 1 by the classifier.

   The middle row shows the same results but when the output class of the classifier was 2: The mid-left value shows that 39 images which were labelled as class 2 actually had the correct class 1, and the mid-mid value shows that 60 images were correctly labelled as class 2. The mid-right value show the percentage of correctly labelled images among those that were labelled as class 2 by the classifier.

   The bottom row shows the percentage of correctly labelled images among those that had class 1 (left) and 2 (mid), as well as the total percentage of correctly classified images to the right. Thus in total 63.7% of the images were correctly classified. This is a quite low number, which is not much better than random guessing, so the classifier does not seem to be doing such a good job.

4. (a) The K-means algorithm classifies the data into K classes by iterating the following steps until the locations of the cluster centers stop changing places between iterations:
   1. Select K observations as cluster centers.
   2. Assign all observations to their nearest cluster center.

3. Compute the mean within each cluster and assign these values as new cluster centers.
4. Repeat from Step 2.

(b) Using the K-means algorithm corresponds to assuming a Gaussian mixture model for the data where the prior probabilities for each class is $1/K$ and the covariance matrix for each Gaussian distribution is $\sigma^2\mathbf{I}$. The algorithm will not work well for this example because of two reasons: First of all we have a lot more observations in one of the two classes, so the assumption that the prior probabilities are $1/2$ will likely force a lot of the observations into the wrong class. Secondly, the shape of the clusters will always be circular when using the K-means algorithm, but we see that we need elliptical clusters for this data.