Advanced Methods in ML 2017 - Exercise 2

1. In this exercise you will implement the *Loopy Belief Propagation* (LBP), for a 2D image completion task. In this task, we are given an image with some unobserved pixels, that need to be filled in. To do so, we define a *Markov Network* over a 2D grid graph (corresponding to the image pixels) as follows. The variables $x_1, \ldots, x_n$ will correspond to the greyscale values of the pixels (integers between 0 and 255). Define the following pairwise function:

$$\phi(x_i, x_j) = e^{-\min(|x_i - x_j|, v_{max})}$$

(i.e the pairwise factor is an exponent of a thresholded quadratic, where v_{max} is the threshold). The overall MRF will be defined as:[1]

$$p(x_1, \ldots, x_n) = \prod_{ij \in E} \phi(x_i, x_j) \tag{1}$$

where E are the edges of the grid graph. Convince yourself that this is a reasonable approach to modeling smoothness in the image.

We consider the MAP problem (namely, finding the assignment with maximum probability). The LBP algorithm (known as max-product in this context) is then defined as follows.
The algorithm updates as set of messages $m_{ij}(x_j)$ defined for every edge i and j that are neighbors in the graph, and every value of the variable x_j (so for each edge we will have messages $m_{ij}(x_j)$ and $m_{ji}(x_i)$. The messages are updated as follows. First calculate:

$$m'_{ij}(x_j) = \max_{x_i} \phi_{ij}(x_i, x_j) \prod_{k \in N(i)/\{j\}} m_{ki}(x_i) \tag{2}$$

And then update the new message to be a normalized version of m'_{ij} (note that for tree graphs we do not need this normalization, but in the loopy case we need it to prevent the messages from exploding or decaying to zero. Convince yourself that this is a "legal" thing to do in the tree case).

$$m_{ij}(x_j) = \frac{m'_{ij}(x_j)}{\sum_{\bar{x}_j} m'_{ij}(\bar{x}_j)} \tag{3}$$

These updates are repeated until convergence (although they generally not guaranteed to converge, in which case you just stop after some number of updates). There are many options for choosing the sequence of messages. Here's a simple one you can try: choose a certain order on the variables (say row by row in the part of the image you want to complete), and then for each node update all messages from it to its neighbors).

The approximate MAP assignment is then obtained from:

$$x_i^* = \arg\max_{x_i} \prod_{k \in N(i)} m_{ki}(x_i) \tag{4}$$

[1]Note we not using ϕ_{ij} since all the ϕ functions are the same, just applied to different edges.

To use this algorithm for an image completion task, run multiple iterations of it over a image segment containing the unobserved pixels, while fixing the values of the pixels that are observed.
Submission Guidelines:

- The image to be completed in this exercise is downloadable from the course website under the name 'pinguin-img.png'.
- The output of your code should be a reasonable completion of the input image. You can try $v_{max} = 50$, although other values might work as well.
- Add your IDs and the path to your code's folder in the following Google Form: https://tinyurl.com/kqsd9uh
- Make sure your code's folder has the proper read permissions.
- Your code should be run by the command 'python mrf_completion.py -outfile', where 'outfile' is the name of your output image, which should be in '.png' format.
- Your code should be *readable* and *well-documented*.

Hints and Tips:

- For numerical stability, use the logs of the factors written above.
- For building the graph "infrastructure" etc., you are free to use the skeleton code uploaded in the course website.

2. In this question you will show that LBP generates pairwise and singleton marginals that satisfy the property from Question 7 of Exercise 1. The question refers to sum-proudct LBP. Namely, the algorithm which on trees calculates exact marginals. Consider the pairwise MRF:

$$p(x_1, \ldots, x_n) \propto \prod_{ij \in E} \phi_{ij}(x_i, x_j) \tag{5}$$

where E are edges of a give graph.

Define the following *pseudo-marginals* $\mu_{ij}(x_i, x_j)$ and $\mu_i(x_i)$:

$$\mu_i(x_i) \quad \propto \quad \prod_{k \in N(i)} m_{ki}(x_i) \tag{6}$$

$$\mu_{ij}(x_i, x_j) \quad \propto \quad \phi_{ij}(x_i, x_j) \prod_{k \in N(i) \setminus j} m_{ki}(x_i) \prod_{k \in N(j) \setminus i} m_{kj}(x_j) \tag{7}$$

(a) Show that for tree graphs the above pairwise pseud-marginals $\mu_{ij}(x_i, x_j)$ are indeed the correct marginals of the MRF (you can use the fact that the messages are in that case the correct sum of assignments in the corresponding sub-tree).

(b) Show that for any graph E (possibly non-tree) it holds that:

$$p(x_1, \ldots, x_n) \propto \prod_i \mu_i(x_i) \prod_{ij} \frac{\mu_{ij}(x_i, x_j)}{\mu_i(x_i)\mu_j(x_j)} \tag{8}$$

Note this is exactly the property shown in Q7 in Ex1, but here we are considering non-tree graphs as well. And μ_{ij} and μ_i are not the marginals of $p(x_1, \ldots, x_n)$

(c) Note that generally the pairwise and singleton marginals might not be consistent. Namely it isn't true that $\sum_{x_j} \mu_{ij}(x_i, x_j) = \mu_i(x_i)$. Show that at a fixed point of LBP this property does hold.