Exercises for Graphical Models in Computer Vision - No. 1

Due Date: 18.4.2016

Hand in before the exercise, either in person or via mail to `joel.janai@tue.mpg.de`.

1. Decision Theory (5 + 5 Points)

Remember that the Bayes optimal decision is to predict the value y that incurs the lowest expected loss

$$y^* = \operatorname*{argmin}_{y \in \mathcal{Y}} \mathbb{E}_{p(\cdot|x)}[\Delta(y, \cdot)]. \tag{1}$$

1. Show that for a random variable $y \in \mathcal{Y}$, the 0/1 loss $\Delta(y, y') = [y \neq y']$ (1 if $y \neq y'$, 0 otherwise, i.e. $[a] = 1$ if a is true, 0 otherwise), the Bayes optimal decision is

$$y^* = \operatorname*{argmax}_{y \in \mathcal{Y}} p(y|x). \tag{2}$$

   in other words the optimal decision for the 0/1 loss is to predict the maximum-a-posteriori value. This is true for general $\mathcal{Y}$, maybe start with $\mathcal{Y} = \{0, 1\}$.

2. For a vector valued $y \in \mathbb{R}^D$, a frequently used loss is the sum-of-squared error (SSE) loss

$$\Delta(y, y') = \|y - y'\|^2 = \sum_{d=1}^{D} |y_d - y'_d|^2.$$

   For this case, and given a conditional probablity distribution $p(y|x)$, what is the Bayes optimal decision, i.e. Eq.(1)?

$$y^* =? \tag{3}$$

2. Bayes Decision Theory?!

This question does not count, but I thought that maybe you wonder.

- Bayes decision theory makes sense. But, given a belief $p(\cdot|x)$ and a loss function Δ can you think of a case where it is *not optimal* to predict with Eq.(1)?