Probabilistic Graphical Models in Computer Vision

Lecture: Dr. Tao Wu
Exercises: Yuesong Shen, Zhenzhang Ye
Summer Semester 2019

Computer Vision Group
Institut für Informatik
Technische Universität München

Weekly Exercises 9

Room: 02.09.023

Wednesday, 17.07.2019, 12:15 - 14:00

MRF/CRF learn (Due: 15.07) (6+6 Points)

NOTE: There will be another programming assignment next week on EM algorithm and graph cut (due 29.07).

Exercise 1 (6 Points)**.** (Hessian of NLL) Prove the following result from the lecture (P9 of learning):

$$\nabla_\theta^2 l(\theta) = \mathbb{E}_{x\sim p(\cdot;\theta)}[\psi(x)\psi(x)^\top] - \mathbb{E}_{x\sim p(\cdot;\theta)}[\psi(x)]\,\mathbb{E}_{x\sim p(\cdot;\theta)}[\psi(x)]^\top = Cov_{x\sim p(\cdot;\theta)}[\psi(x)] \tag{1}$$

Exercise 2 (6 Points)**.** (partially observed CRF) Denoting x_i, y_i, z_i as the conditioned / observed / hidden part of the i-th sample respectively, $r(y,x)$ the empirical data distribution and $r_x(x) = \sum_y r(y,x)$. Similar to the cases with fully observed CRF and partially observed MRF, the partially observed CRF has the following distribution:

$$p(y,z|x;\theta) = \frac{1}{Z(\theta;x)}\exp(\theta^\top\psi(x,y,z)), \tag{2}$$

and we learn the parameters θ by minimizing the negative conditional data likelihood

$$l(\theta) = -\frac{1}{|\mathcal{S}|}\sum_{1\le i\le|\mathcal{S}|}\log p(y_i|x_i;\theta). \tag{3}$$

Prove that its gradient is

$$\nabla_\theta l(\theta) = -\mathbb{E}_{(x,y)\sim r}[\mathbb{E}_{z\sim p(\cdot|x,y;\theta)}[\psi(x,y,z)]] + \mathbb{E}_{x\sim r_x}[\mathbb{E}_{(y,z)\sim p(\cdot,\cdot|x;\theta)}[\psi(x,y,z)]] \tag{4}$$

Programming (Due:22.07) (12 Points)

In this programming exercise, you are asked to classify MNIST hand-written digits with Naive Bayes model and learn the parameters. See the ipython file for more details.