TU München
Fakultät für Informatik
PD Dr. Rudolph Triebel
John Chiotellis, Maximilian Denninger

Machine Learning for Computer Vision

May 5, 2018
Topic: Graphical Models 2

Exercise 1: Iterated Conditional Modes

Consider the use of iterated conditional modes (ICM) to minimize the energy function given by (1), where y_i are the observed values and x_i are the true values.

$$E(\boldsymbol{x}, \boldsymbol{y}) = h \sum_i x_i - \beta \sum_{j \in N(i)} x_i x_j - \eta \sum_i x_i y_i \tag{1}$$

a) Write down an expression for the difference in the values of the energy associated with the two states of a particular variable x_j , with all other variables held fixed, and show that it depends only on quantities that are local to x_j in the graph.

b) Write the joint probability distribution corresponding to this energy function.

c) Consider a particular case of the energy function given by (1) in which the coefficients $\beta = h = 0$. Show that the most probable configuration of the latent variables is given by $x_i = y_i,\ \forall i$.

Exercise 2: Programming ICM

a) Download the images.zip file from the website. The file contains a binary image and the same image with added random noise. Use the ICM algorithm to denoise the image. Begin by setting $h = 0, \beta = 1.0, \eta = 2.1$ as hyperparameters. How fast does the algorithm converge?

b) Perform a grid search over the hyperparameters and pick the setting that gives the solution with lowest energy. How close is the solution to the groundtruth?

c) Try applying different levels of noise (20%, 40%, 60%) to the groundtruth image. Does the algorithm still work? If yes, why? If not, why not?