Homework 1: Winter 2012. Stat 238. Due Wednesday. 22 February 2012.

A.L. Yuille:

Abstract

The homework assigning is to read the course notes for the first seven lectures, to put comments on the pdf files (sticky notes) and email to yuille@stat.ucla.edu

The comments can address clarity, errors (e.g., typos), request more examples, more motivation, more figures, more references to background material. The questions below are intended to motivate you to read the notes. These questions should have short answers. A few sentences and not long mathematical derivations.

1. Questions

1. What is linear filtering? What are the main purposes of filtering images.?

2. What is segmentation? What is the weak smoothness model of images?

3. What is statistical edge detection? What are conditional probability distributions? What ere the priors on edges and why are they needed? What are the advantages of this approach? What are the limitations?

4. What are the properties of non-linear filters that make them insensitive to image transformations?

5. What are factorizable graphical models? How do they make learning and inference easier?

6. What is the Rudin, Osher, Fatemi model? How does it exploit spatial context? Why is it called a functional? What calculus is used to take derivatives of functionals?

7. What is convexity? And what performance guarantees does it give for inference? What are steepest decent algorithms? What are variational bounding and CCCP algorithms?

8. What is an MRF model? What are the differences between a directed and undirected graphical model? How do MRFs impose constraints like spatial smoothness and impose spatial context? What is a Gibbs distribution? How can functionals (sometimes) be converted to MRFs?

9. What is Gibbs sampling for MRFs? What is it guaranteed to converge to?

10. What are mean field theory and variational methods? What are their advantages/disadvantages compared to Gibbs sampling?

11. What are exponential distributions? What are sufficient statistics? What is the maximum entropy principle?

12. What is maximum likelihood for exponential distributions? Name two algorithms for performing maximum likelihood learning? What is the main difficulty? How do the expected statistics of the data relate to the observed statistics?

13. How do the statistics of images justify the weak membrane model? How do statistics on the image (or other data) yield MRF models? How do the statistics of derivative filters on images relate to exponential models and MRFs?

14. What is the purpose of spectral clustering? What are affinities and the graph laplacians? How can spectral clustering be used to decompose a graph into its connected components? What is a "natural affinity" for images?

15. Briefly describe the Berkeley edge detector. How are cues combined? How are spectral methods used to add global (non-local) information?