Syllabus for Graduate Computer Vision

Erik G. Learned-Miller
Department of Computer Science
University of Massachusetts, Amherst
Amherst, MA 01003

January 4, 2010

Abstract

1 Introduction and Motivations

1. What is computer vision? What is the goal of computer vision? What does human vision do? What should computer vision do? When is human vision "wrong"? Is this bad? Illusions and their role in computer vision.
2. Intro to MATLAB. Image representation.

2 Uncertainty, probability, estimation, and decision making

1. Formalizing the decision making process. Minimizing errors. Maximizing utility. Quick review of probability theory. Sample spaces. Discrete probabilities. Continuous probabilities. Conditional probability. Joint probability. Marginal probabilities. Independence of random variables. Factorization of distributions.
2. Bayes rule. Bayes rule in a conditional setting. Bayes rule in a multi-dimensional setting. Bayes rule and independence.
3. Estimation. Models. Trade-off between form of models and speed of estimation. Relation to computer vision.
   Examples:
   (a) Binary random variables. Which ones are hard to estimate?
   (b) Multinomials. Do they have a "topology"?
   (c) Continuous random variables. Options for discretization. Soft binning.
   (d) Multi-dimensional distributions. How to get enough data?
4. Supervised learning.
5. Nearest neighbor. In one dimension. In two dimensions. In three dimensions. In 10,000 dimensions.
6. Generative models.
7. Discriminative models.
8. Semi-supervised learning.

3 Alignment

1. What is alignment? Criteria for alignment.
2. Exhaustive search. Gradient descent. Random restarts. Simulated annealing. Combination approaches.
3. Smoothing the space.

4. Joint alignment. Congealing.

5. Prokudin-Gorsky photos. Alignment by maximization of mutual information.

6. Classification with alignment.

   (a) Invariance. Penalty for alignment. Partial invariance. Revisiting congealing.

   (b) Classification with alignment.

4 Light and Optics

1. The electromagnetic spectrum. Visible light. Infrared, ultraviolet, and other wavelengths.

2. Pinhole cameras and simple lenses. Image formation.

3. Point sources of light and extended sources. Power, irradiance, radiance, brightness, lightness. Tracing the path of light from source to sensor.

4. Sensors. CCD and CMOS sensors. The human retina. Discretization. Number of levels. Discretization pattern. Sensor response. Gain control. Local gain control. Dynamic range.

5. Blurring. Blurring between pixels. Blurring between frames during motion. Deblurring. Deconvolution with kernel. Blind deconvolution. Estimates of distance from blur. Estimates of motion from blur.

5 Motion

1. Depth from motion. Pigeon and Praying mantis motion. Comparison to stereo vision for depth.

2. Backgrounding. Thresholding. Likelihood backgrounding. Bayesian backgrounding. Panoramic backgrounding.

3. Tracking.

4. Optical Flow. Image derivatives. Smoothness assumption. Horn-Schunk. Lucas-Kanade.

5. Bayesian optical flow?

6 Edges, Boundaries, Contours, and Regions

1. Edges, image derivatives, apparent edges. Texture edges, color edges, other edges.
2. The EM-algorithm. Splitting an image into regions.
3. Layers in optical flow.
4. Segmentation by combining cues. Conditional Random Fields. Markov Random Fields.