Vision Algorithms (CSE/EE 577, CSE590CV)

Staff

- Prof: Steve Seitz (seitz@cs)
- TA (part time): Colin Zheng (kzheng@cs)

Web Page

- http://www.cs.washington.edu/education/courses/cse577/04sp/

Handouts

- these notes
- class web page
- project/partner request form

Course Overview

Motivation—tools you should know to do research in computer vision

Administrative

- Get on class mailing list!
  - http://mailman.cs.washington.edu/mailman/listinfo/cse577
  - (if you got the TEST message today, you're already on)
- room, video taping
- format:
  - 577: Tue/Fri 1:30-2:50: invited tutorial talks, lectures, project presentation, discussion
  - 590CV: Mon (mostly) 10:30 at MSR (113/1021). Van leaves promptly at 9:45am from parking lot N. of HUB.
- By Thursday noon: complete project/partner form
  - choose partner for first two projects
  - algorithm preferences
- Bus to MSR (590CV)—space is limited to registered students

Assignments

Implement three algorithms/applications

- one from each category
  - cat1: belief propagation, expectation maximization, graph cuts
  - cat2: level sets, nonlinear least squares/sparse matrix, discriminative methods
  - cat3: dimensionality reduction, distance transforms/matching, monte carlo sampling (MCMC)
- work with a partner
- each group implements both an algorithm and an application
- present results in class

Topics/speakers

- Belief propagation (Yair Weiss, Hebrew U)
  - intro to graphical models
- Expectation maximization (Nebojsa Jojic, MSR)
  - variational methods
- Graph cuts (Ramin Zabih, Cornell)
- Level set evolution (Guillermo Sapiro, Minnesota)
- Nonlinear least squares (Rick Szeliski, MSR)
  - sparse matrix methods
- Discriminative methods (Paul Viola, MSR)
  - Adaboost
- Dimensionality reduction (Sam Roweis, Toronto)
- Distance transforms, matching (Dan Huttenlocher, Cornell)
- Markov Chain Monte Carlo (Frank Dellaert, Georgia Tech)

(Loopy) belief propagation (Yair Weiss)

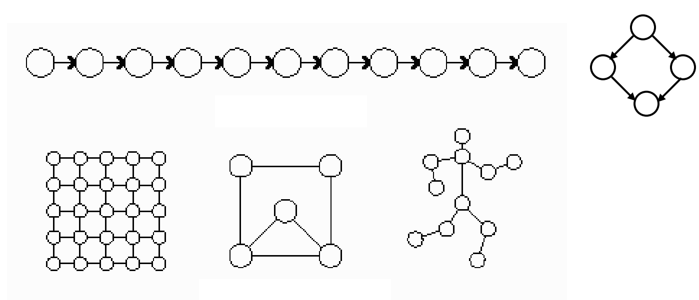

Inference on *graphical models*

- nodes are random variables X
- edges are probabilistic dependences P(X | Y)
- goal (e.g.): find most likely value assignment

Applications

- stereo, detection/recognition, tracking, ...

Expectation Maximization (Nebojsa Jojic)

EM is an iterative optimization method to estimate some unknown parameters Θ, given measurement data $\mathbf{U}$. However, we are not given some "hidden" nuisance variables $\mathbf{J}$, which need to be integrated out. In particular, we want to maximize the posterior probability of the parameters Θ given the data $\mathbf{U}$, marginalizing over $\mathbf{J}$:

$$\Theta^* = \underset{\Theta}{\operatorname{argmax}} \sum_{\mathbf{J} \in \mathcal{J}^n} P(\Theta, \mathbf{J} | \mathbf{U}) \qquad (1)$$

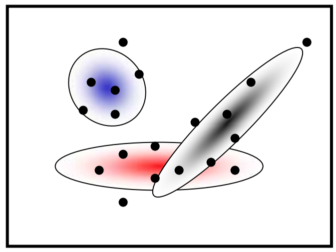

Applications

- segmentation, reconstruction, recognition, ...

Graph Cuts (Ramin Zabih)

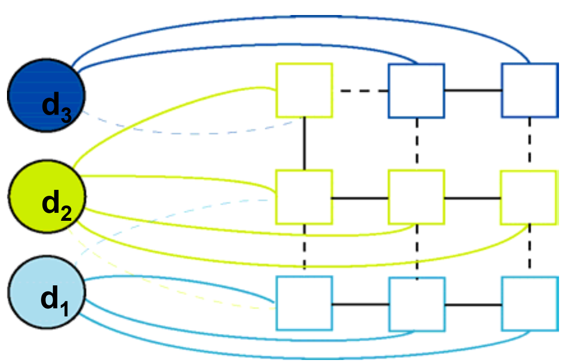

Some graphical models can be solved efficiently

- graph cuts formulations can give global solutions (exact or approx)

Applications

- stereo, shape from silhouettes, image blending, textures

Level Set Evolution (Guillermo Sapiro)

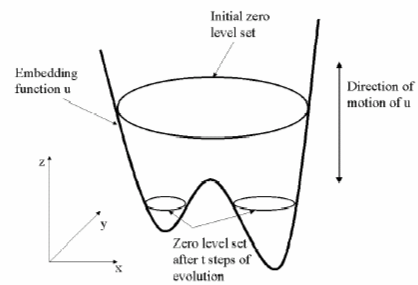

figure 2.2. *Embedding the curve in a higher dimensional function automatically solves topological problems. While the curve can be changing topology, the surface moves up and down, on a fixed coordinate system, without altering its topology. The topological change is just 'discovered' when computing the corresponding level-set of the function.*

Propagating a curve or surface

- level set formulation avoids parameterization, can change topology

Applications

- multi-view stereo, in-painting, noise-removal

Nonlinear least squares (Rick Szeliski)

We all know how to solve systems of linear equations

How about systems of nonlinear equations?

Rick will also present sparse matrix techniques

Applications

- bundle adjustment (structure from motion, ego-motion/calibration)
- alignment, registration, tracking

Discriminative Methods (Paul Viola)

Techniques:

- Support vector machines
- Adaboost
- Memory based classification and locally-weighted regression
- Building fast classifiers

Applications:

- Face and object detection
- Recognition: Characters/faces/objects
- Body pose estimation

Dimensionality Reduction (Sam Roweis)

Given a data set, find a low-dimensional embedding/parameterization

Linear methods are well-known (PCA)

Nonlinear methods are newer, more general

Applications

- Recognition, compression, learning control knobs

Distance Transforms + Matching (Dan Huttenlocher)

Fast methods for matching points and other features

Definition

- Distance at each location to nearest point
- Relationship to Voronoi diagrams and Delaunay triangulations
- Different metrics (Chamfer, Hausdorff)

Applications

- Model fitting, recognition, registration/alignment...

Markov Chain Monte Carlo Sampling (Frank Dellaert)

Generate sequence of samples of an arbitrary PDF

- How can we use this? Design a PDF where high probability points correspond to solutions of your problem
- In many cases it is possible to evaluate such a PDF, but not directly maximize it.

Applications

- matching, recognition, animation, rendering

Algorithms by area in computer vision

Reconstruction

- Level sets, graph cuts, nonlinear LS, MCMC

Recognition

- Discrim. methods, distance transforms, dim. reduct., EM, BP, MCMC

Matching/motion

- Distance transforms, nonlin LS, BP, graph cuts

Segmentation

- EM, BP, MCMC

Image processing

- BP, graph cuts

Not meant to be comprehensive...

This week

Yair Weiss (Hebrew University)

- Wed MSR talk (bus leaves 9:45am, N. Hub parking lot)
  - "Learning to Perceive from Image Statistics--A Computational Challenge"
- Fri UW talk (CSE 305)
  - "Approximate inference in graphical models using loopy belief propagation"

Reading (for Friday)

- Jordan/Weiss: "Probabilistic Inference in Graphical Models"
  - http://www.cs.washington.edu/education/courses/577/04sp/readings/jordan.pdf