Mathematical Tools for Image Collections

Paulo Carvalho
IMPA

Outline

- Problems to solve
- Mathematical models
- Probability and Statistics
- Graph Cuts

Problems to solve

- Registration
- Segmentation
- Model Estimation
- Fusion
- Re-Synthesis

Problems to solve

- Basic tasks:
  - Classify data according to similarities and dissimilarities
  - Generate new data according to observed patterns

Example: matting

- Chuang et al. *A Bayesian Approach to Digital Matting*

classify

generate

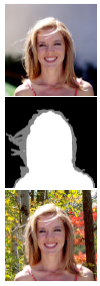

Mathematical Tools

- Probability
  - Mathematical models for uncertainty
- Statistics
  - Inference on probabilistic models
- Optimization
  - Formulating goals as functions to be minimized (and doing it...)

Probability

- Random vs deterministic phenomena
- Uncertainty expressed by means of probability distributions.
- Describe relative likelihood.

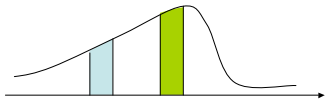

Statistics

- Basic problem: **infer** attributes of probability models from **samples**

(Statistics: probability models + samples)

Statistics

- Parametric vs non-parametric
  - Parametric models: distributions known, except for a finite number of parameters (e.g., Normal(μ, σ^2))
  - Otherwise, non-parametric (e.g., distribution symmetric about μ).

Statistics

- Classical vs Bayesian
  - Classical statistics: parameters are unknown constants (*states of Nature*) and inference is purely objective: based only on observed data.
  - Bayesian statistics: parameters are random themselves and inference is based on observed data and on prior (subjective) belief (probability distribution).

Bayesian Inference

- Appropriate for learning models
- Mimics learning process
  - Start with previous (imprecise) knowledge
  - Observe new data
  - Revise knowledge (still imprecise, but better informed)

Types of inference problems

- Classification
- Regression
- Density estimation
- Dimension reduction
- Clustering
- Model selection

Bayesian Inference

- Good starting point:
  - A. Hertzmann. *Machine Learning for Computer Graphics: A Manifesto and a Tutorial*
  - Provides basic references on machine learning
  - Relates techniques to papers in Computer Graphics.

Bayesian Inference

- Given
  - prior distribution $\pi(\theta)$
  - conditional distribution $p(\mathbf{x}|\theta)$
  - data $\mathbf{x}$
- Compute
  - posterior distribution

$$p(\theta \mid \mathbf{x}) = \frac{p(\theta, \mathbf{x})}{p(\mathbf{x})} = \frac{\overbrace{\pi(\theta)}^{\text{prior}}\,\overbrace{p(\mathbf{x} \mid \theta)}^{\text{likelihood}}}{\underbrace{\int \pi(\theta)\, p(\mathbf{x} \mid \theta)\, d\theta}_{\text{observation}}}$$

(Bayes' theorem)

Bayesian Inference

- Use the posterior distribution $p(\theta|\mathbf{x})$ to make inferences on θ
- Classification: MAP estimate (maximum a posteriori)

$$\hat{\theta} = \arg\max\, p(\theta \mid \mathbf{x})$$

- Estimation (minimize expected quadratic error):

$$\hat{\theta} = E(\theta \mid \mathbf{x}) = \int p(\theta \mid \mathbf{x})\, d\theta$$

Segmentation/classification

- Easy to use (Bayesian or other) inference methods for each pixel.
  - Local solution
  - Bad overvall results
- How to produce spatially-integrated results?

Example [Rabih et al]

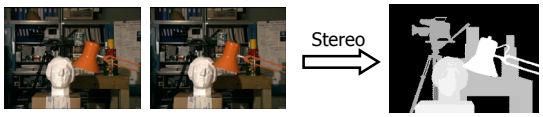

- Classify pixels according to depths (disparity)
- Neighboring pixels should be at similar depths
  - Except at the borders of objects!

Depth segmentation

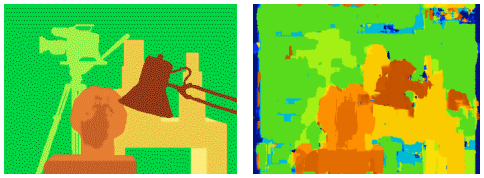

Ground truth | Local method (maximum correlation)

Pixel labeling problem

- Global formulation
  - Attribute labelings $f_p \in \{1, 2, ..., m\}$ to pixels
  - Assignment cost $D_p(f_p)$ for assigning label f_p to pixel p.
  - Separation cost $V(f_p, f_q)$ for assigning labels f_p, f_q to neighboring pixels p, q
  - Minimize total cost (energy function):

$$\min \sum_{p \in I} D_p(f_p) + \sum_{\text{neigh.}\, p,q} V(f_p, f_q)$$

Example: segmentation

- Labels are 0-1 (background-foreground)
- $D_p(f_p)$ measures the individual likelihood to belong to the background or the foreground (e.g, a posteriory probability)
- If $f_p \neq f_q$, $V(f_p, f_q) = \begin{cases} 2K, \text{if } |I_p - I_q| < C \\ K, \text{otherwise} \end{cases}$

(discontinuity preserving)

Example: stereo

- Labels are disparity between corresponding pixels

- $D_p(f_p)$ is difference in intensity for each depth (disparity)
- If $f_p \neq f_q$, $V(f_p, f_q) = \begin{cases} 2K, \text{if } |I_p - I_q| < C \\ K, \text{otherwise} \end{cases}$

Pixel labeling program

- NP-Hard
- General purpose optimization methods
  - Simulated annealing, or some such
  - Bad answers, slowly
- Local methods
  - Each pixel chooses a label independently
  - Bad answers, fast

Efficient solution: min cuts

- Maximum flow problem

  Given a directed graph, with distinguished nodes s, t and edge capacities c_a, find the maximum flow from s to t

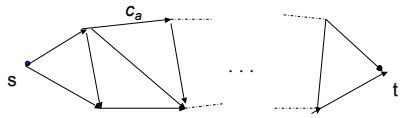

Max flow – min cut theorem

- [Ford & Fulkerson] The maximum s-t flow in a network equals the capacity of the minimum s-t cut

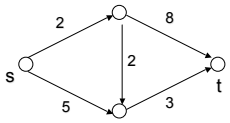

Max flow – min cut theorem

- [Ford & Fulkerson] The maximum s-t flow in a network equals the capacity of the minimum s-t cut

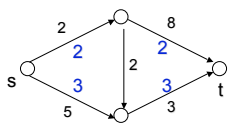

max flow = 5

Max flow – min cut theorem

- [Ford & Fulkerson] The maximum s-t flow in a network equals the capacity of the minimum s-t cut

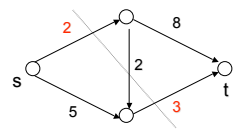

max flow = 5

min cut = 5

Max flow – min cut theorem

- Inequality easy.
- Equality shown by the flow augmenting path algorithm.
- Provides the basis for building efficient algorithms for max flow (min cut)
- State of the art: low-degree polynomial complexity, with small constants (i.e., fast)

Pixel labeling and min cuts

- Pixel labeling problems can be efficiently solved by a sequence of min cut problems.
- Each min-cut problem represents a labeling problem with only two labels (label expansion).
- Cuts are labelings, cut costs are energy functions.
- Local minima, but global for certain energy functions.

Pixel labeling and min cuts

p q r s

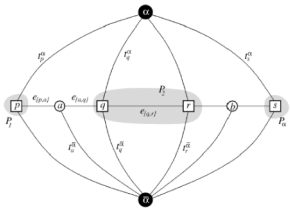

edge	weight	for
t_p^α	∞	$p \in \mathcal{P}_\alpha$
t_p^α	$D_p(f_p)$	$p \notin \mathcal{P}_\alpha$
t_p^α	$D_p(\alpha)$	$p \in \mathcal{P}$
$e_{\{p,a\}}$	$V(f_p, \alpha)$	
$e_{\{a,q\}}$	$V(\alpha, f_q)$	$\{p,q\} \in \mathcal{N},\ f_p \neq f_q$
t_a^α	$V(f_p, f_q)$	
$e_{\{p,q\}}$	$V(f_p, \alpha)$	$\{p,q\} \in \mathcal{N},\ f_p = f_q$

Pixel labeling and min cuts

p q r s

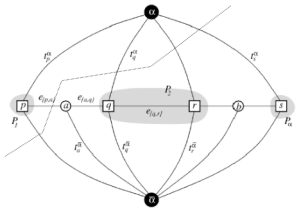

edge	weight	for
t_p^α	∞	$p \in \mathcal{P}_\alpha$
t_p^α	$D_p(f_p)$	$p \notin \mathcal{P}_\alpha$
t_p^α	$D_p(\alpha)$	$p \in \mathcal{P}$
$e_{\{p,a\}}$	$V(f_p, \alpha)$	
$e_{\{a,q\}}$	$V(\alpha, f_q)$	$\{p,q\} \in \mathcal{N},\ f_p \neq f_q$
t_a^α	$V(f_p, f_q)$	
$e_{\{p,q\}}$	$V(f_p, \alpha)$	$\{p,q\} \in \mathcal{N},\ f_p = f_q$

Example - Stereo [Zabih]

For each move we choose expansion that gives the largest decrease in the energy: **binary energy minimization subproblem**

initial solution

-expansion

-expansion

-expansion

-expansion

-expansion

-expansion

-expansion

Energy minimization via min cut

- Zabih's web page:
  http://www.cs.cornell.edu/~rdz/graphcuts.html