6.881 Representation and Modeling for Image Analysis

Instructor: Polina Golland

Goals

- Learn useful mathematical techniques
- Read vision papers that use them
  - Old and new “classics”
- Practice
  - Presenting work
  - Writing/reviewing papers
- Understand the interconnectedness of everything

Syllabus

- Subspace (Manifold) learning
  - Theory: PCA
  - Applications: Eigen faces, Active Shape & Active Appearance Models.
  - Additional topics: kernel PCA, LLE
- Boundary Detection
  - Theory: Calculus of variations
  - Applications: Mumford-Shah functional, active contours: snakes and level sets
- EM:
  - Theory: EM algorithm
  - Applications: segmentation, tracking

Syllabus cont'd

- Graph algorithms
  - Theory: Graph cut algorithms
  - Applications: segmentation, stereo
- Clustering
  - Theory: hierarchical, k-means, spectral
  - Applications: grouping in images
- Graphical Models
  - Theory: MRFs, inference in graphical models
  - Applications: regularization, part/layer models

Syllabus (Approximate) cont'd

- Shape descriptors
- Transformations and manipulating them
- Information Theoretic Methods
- Classification

Format

- For each topical cluster
  - Tutorial on the method/theory
  - Paper presentations
  - Discussion

- We will expect that you
  - Have read the assigned reading material
  - Have questions and are ready to discuss

When you present

- Meet with us the week before to go over the presentation

- Present the paper/method in class and set the stage for discussion
  - What is the problem?
  - What methods are used?
  - What are the main assumptions?
  - How robust are the results?
  - Possible extensions: point to others' work or suggest your own ideas

Grades

- 40% method/paper presentations
- 20% participation in the discussions
- 30% final paper & presentation
  - Project or analysis paper
- 10% paper review

Final Project/Paper

The goal is to write a conference paper and to get feedback on it.

We will also have final presentations on the projects in the last three classes.

Proposals due: March 4.
Final papers due: May 1.
Paper reviews due: May 12.

Course participation

- Please register for the course
  - It increases the likelihood of your actually reading the papers
  - It increases the likelihood I will be able to offer this seminar again in the future

- Sitting in on this course is ok, but you still have to read the papers and participate in the discussions. You might also be asked to present.
  - so you might as well get the credit for it

First Topic Cluster: PCA

- Today:
  - tutorial on Principal Component Analysis

- Next week
  - Eigenfaces
    - » PCA on intensity values of the faces images
    - » More recent work: kernel PCA
    - » Presenter: Polina Golland

  - Active Shape & Active Appearance Models
    - » PCA on point locations and intensity variation
    - » Looking for a presenter

Looking Ahead

- Tutorial on calculus of variations
- Mumford-Shah functional, snakes
- Level Set curve evolution

- Tutorial on EM
- EM segmentation
- EM for tracking in layer models
- EM for generative model training

End of logistics

Principal Component Analysis

- Samples $\{\mathbf{x}_1,\ldots,\mathbf{x}_n\}$
- Find the best low dimensional approximation of the samples
  - Restrict to linear approximations
  - Minimize sum squares error
- Blackboard derivation of PCA

Properties

- **m** is the sample mean
- $\{\mathbf{v}_1, \ldots \mathbf{v}_k\}$ form an orthonormal basis

- $\mathbf{v}_i$'s are the *k* largest eigenvectors of the sample covariance

- they can also be found sequentially!

- Mahalanobis distance

$$(\mathbf{x}-\mathbf{y})^T S^{-1} (\mathbf{x}-\mathbf{y})$$

Geometric Interpretation

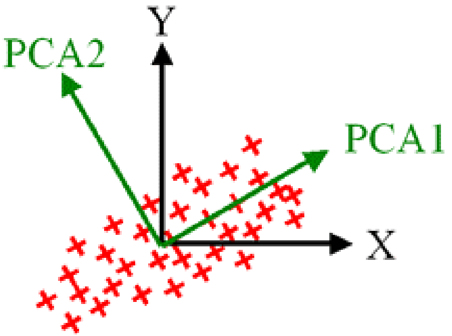

- Linear subspace that explains most variability in the sample.
- "Prefix optimal"
  - Error measure: sum of squares.
- K-dimensional Gaussian distribution that best approximates the sample.
- In physics: principal axes of inertia.

Extensions

- Data lies in a non-linear space
  - Kernel PCA
  - Locally linear embedding (LLE)
  - Other manifold learning methods