The University of Chicago
Department of Statistics
Ph.D. Seminar

ELLIOT JOEL BERNSTEIN
Department of Statistics
The University of Chicago
Adviser: Yali Amit, Professor of Statistics

Statistical Models for Object Classification and Detection

Monday, November 7, 2005 at 12:30 pm
110 Eckhart Hall, 5734 S. University Avenue

ABSTRACT

Classification of pre-segmented images and detection of objects in images are two fundamental problems in computer vision. Traditional approaches to these problems involve the use of non-parametric classifiers that use training data to explicitly learn boundaries between various object categories (for classification) or between the object and non-object populations (for detection). In recent years, attention has turned to the use of statistical models to represent object categories, although these are frequently complex Bayesian models that are carefully constructed for a particular type of object.

We present an alternative approach, using straightforward statistical mixture models that can be estimated using the EM algorithm and applied to a wide variety of object categories. The same model form is used both to learn a collection of generic local features and to learn a model for each object category of interest. In its simplest form, the local feature model identifies each non-background local image patch with one or more indices, indicating the features found at that location. By adding additional latent variables to the mixture, the feature set can be extended to encode additional information, such as feature orientation and polarity.

We demonstrate the utility of this framework through three visual tasks: classification of pre-segmented handwritten digits, detection of car profiles in photographs, and detection of faces. When performing classification, the likelihood of the observed data is calculated under the model for each category, and the category with the highest likelihood is selected. For detection, likelihood ratio tests are computed using a coarse-to-fine computational strategy that focuses attention on image regions likely to contain the object of interest. If orientation information has been incorporated into the local features, the object models can be rotated to allow detection of objects at a variety of poses. Of particular note is the fact that the object detector can be learned using only data from the category of interest; no non-object examples are required.

Information about building access for persons with disabilities may be obtained in advance by calling the department office at (773) 702-8333.