Statistical Approaches to Texture Classification

Manik Varma

Abstract

submitted for the degree of
Doctor of Philosophy
Trinity Term 2004

This thesis investigates the problem of classifying textures from their imaged appearance without imposing any constraints on, or requiring any *a priori* knowledge of, the viewing or illumination conditions under which the images were obtained. Weak classification algorithms based on the statistical distribution of texton primitives are developed to categorise single, uncalibrated images into a set of pre-learnt material classes.

The thesis starts by introducing a filter bank based approach to the problem of texture classification. We design low dimensional, rotation and scale invariant filter sets which are nevertheless capable of extracting rich features at multiple orientations and scales. Textures are modelled by the frequency distribution of exemplar filter response features. Characterising a texture by multiple models allows the classification of single images without requiring any knowledge of the imaging conditions. Using this framework, it is demonstrated that the new filter sets achieve superior performance as compared to their traditional counterparts when benchmarked on real world databases conatining many classes with significant imaging variations.

There are two major approaches when it comes to building texture classifiers based on the statistical distribution of filter responses. One approach is to first determine the texton primitives of a material, next model the texture by their frequency distribution and finally classify novel images by nearest neighbour matching. An alternative is offered by the Bayesian paradigm which recommends learning the joint probability distribution of filter responses followed by MAP classification. We show that both approaches are essentially the same and that they can actually be made equivalent under suitable choices of PDF representation and distance measure.

The issue of whether filter banks are necessary for material classification is addressed next. A novel texture representation is developed based on the joint probability distribution of pixel intensities in compact neighbourhoods. Using this representation within the standard classification framework leads to two astonishing results: (a) very small neighbourhoods can yield superior performance as compared to multi-scale, multi-orientation filter banks with large support and (b) the performance of filter banks is always inferior to the new representation with equivalent neighbourhood size. Theoretical arguments are presented as to why these two results might hold.

Finally, the related problem of determining the illuminant's direction from textured images is explored. A theory for estimating the illuminant's azimuthal angle from images of Lambertian, rough surfaces with spatially varying albedo is formulated. In certain cases, the theory is able to accommodate the effects of non-Lambertian factors such as shadows, specularities, inter-reflections, etc. This is evidenced by the good results achieved on numerous real world images which deviate strongly from the ideal assumptions.