UCLA CS Dept TR 070006 April 2007

Proximity Distribution Kernels for Geometric Context in Category Recognition

Haibin Ling and Stefano Soatto

April 17, 2007

Abstract:
We propose using the proximity distribution of vector-quantized local feature descriptors for object and category recognition. To this end, we introduce a novel ``proximity distribution kernel'' that naturally combines local geometric as well as photometric information from images. It satisfies Mercer's condition and can therefore be readily combined with a support vector machine to perform visual categorization in a way that is insensitive to photometric and geometric variations, while retaining significant discriminative power. In particular, it improves on the results obtained both with geometrically un-constrained ``bags of features'' approaches, as well as with over-constrained ``affine procrustes.'' Indeed, we test this approach on several challenging data sets, including Graz-01, Graz-02, and the PASCAL challenge, and in all tests it outperforms all previously proposed approaches. We registered the average performance at 91.5% on Graz-01, 82.7% on Graz-02, and 74.5% on PASCAL. Our approach is designed to enforce and exploit geometric consistency among objects in the same category; therefore, it does not improve the performance of existing algorithms on datasets where the data is already roughly aligned and scaled. Our method has the potential to be extended to more complex geometric relationships among local features, as we illustrate in the experiments.