

Shape Packing for Object Recognition

Jianbo Shi, Ph.D.

University of Pennsylvania

Abstract

We introduce a method for 'packing' salient image contours/segments into recognizable object shapes. Most object recognition methods rely on one-to-one matching of contours/segments to a model. However, bottom-up image contours/segments often fragment unpredictably. We resolve this difficulty by using many-to-one matching of image contours to a model.

In operation, our system achieves three goals: it locates an object, identifies its part configuration, and segments out its contours. To learn a descriptive object shape model, we combine contours from a few representative images. The goal is to construct a model that can be many-to-one matched to most of the contours in the training images.

For detection, our challenges are inferring the object contours and part locations, in addition to object location. Because the locations of object parts and matches of contours are not annotated, they appear as latent variables during training. We use the latent SVM learning formulation to discriminatively tune the many-to-one matching score using the max-margin criterion. There are several computational implementations, using Linear Programming (LP) or Semi-Definite Programming (SDP). We evaluate on the challenging ETHZ shape categories dataset and outperform all existing methods.

This is joint work with Praveen Srinivasan, Qihui Zhu

About Jianbo Shi:

He studied Computer Science and Mathematics as an undergraduate at Cornell University where he received his B.A. in 1994. He received his Ph.D. degree in Computer Science from University of California at Berkeley in 1998, for his thesis on Normalize Cuts image segmentation algorithm. He joined The Robotics Institute at Carnegie Mellon University in 1999 as a research faculty, where he lead the Human Identification at Distance(HumanID) project, developing vision techniques for human identification and activity inference. In January 2003, he joined the Department of Computer & Information Science at University of Pennsylvania where he is currently an Associate Professor.

His current research focus on human behavior analysis and image recognition-segmentation. His other research interests include image/video retrieval, and vision based desktop computing. His long-term interests center around a broader area of machine intelligence, he wishes to develop a "visual thinking" module that allows computers not only to understand the environment around us, but also to achieve higher level cognitive abilities such as machine memory and learning.