Image Segmentation by Total Variation Methods

Dennis Sun
Stanford University
Dept. of Statistics
dlsun@stanford.edu

Matthew Ho
Stanford University
Dept. of Electrical Engineering
matthew.ho@stanford.edu

1. Introduction

Total-variation denoising (TVD) is a robust technique for reconstructing noisy images. Because it imposes an L1 penalty on differences between adjacent pixels, it tends to result in images with piecewise constant regions. The goal of image segmentation is to assign pixels to clusters such that pixels within each cluster are similar. Our goal is to see if TVD can be used to perform segmentation, by taking the constant regions of the denoised image as the segments. We hope that this project can yield insight into image segmentation, which is a central problem in computer vision and is necessary for many higher-order tasks, such as object recognition.

2. Data

We would like to develop our algorithm for grayscale and color images, and if time permits, RGB-Z images. For grayscale and color images, we will use the Berkeley Segmentation Dataset. For RGB-Z images, we will use images that we will collect around campus using Kinect. (We would appreciate it if we could use images collected from past student projects.)

3. Methods

In the original paper by Rudin, Olsher, and Fatemi, TVD was applied to grayscale images. We intend to make use of existing TVD scripts in Matlab and apply them first to grayscale images, where "piecewise constant" is most clearly defined. Although TVD has also been applied to RGB images, it is less clear how we will search for "piecewise constant" regions, since it is unlikely that the reconstructed image will be piecewise constant in all three coordinates. We will implement and test a variety of methods, such as majority voting (if two of three coordinates are piecewise constant) and thresholding (if the L1 norm of the channel differences is less than some ϵ). There are also techniques to scalarize RGB data, as decribed in Ivanovska. Finally, if time permits, we would like to try TVD on RGB-Z images, combining the channels in a similar fashion.

4. Background

L. Rudin, S. Osher, E. Fatemi. "Nonlinear Total Variation based noise removal algorithms". *Physica D*. 60: 259-268. 1992.

P. Blomgren, T. Chan. "Color TV:Total Variation Methods for Restoration of Vector Valued Images (1996)". IEEE Trans. Image Processing (1996)

T. Chan, S. Kang, J. Shen, Total Variation Denoising and Enhancement of Color Images Based on the CB and HSV Color Models, Journal of Visual Communication and Image Representation, Volume 12, Issue 4, December 2001,

T. Ivanovska, L. Linson. "Converting RGB Volume Data to Scalar Fields for Segmentation Purposes".

Berkeley Vision Group. "The Berkeley Segmentation Dataset and Benchmark".

5. Evaluation

For the images in the Berkeley database, we assume as "true" the human-produced segmentation. To come up with numerical benchworks, we will train our algorithm on a set of provided training images, and then test on the set of test images. The ouput will be a soft boundary map of likelihood values. Then for different thresholds, we will calculate precision and recall and subsequently an F-value, a weighted harmoinc mean of both. (The Berkeley Computer Vision Group provides a harness to do this). We can also print out figures of the segmentation results so that readers can judge for themselves whether the proposed segmentation is reasonable.