SURFACE STEREO WITH SOFT SEGMENTATION

Bleyer M.[1*], Rother C.[2], Kohli P.[2]

[1]Institute for Software Technology and Interactive Systems
Vienna University of Technology, Austria

[2]Microsoft Research Cambridge
Cambridge, UK

Abstract

We propose a new stereo model that employs a surface-based representation, where each pixel is assigned to a 3D surface (planes or B-splines). This representation enables two important contributions: (1) We formulate a higher-order prior which incorporates the very popular color segmentation constraint in a soft and principled way. (2) We use a global MDL prior to penalize the number of surfaces.

Further details are found in the corresponding conference paper: M. Bleyer, C. Rother, P. Kohli, "Surface Stereo with Soft Segmentation", CVPR 2010.

Surface-Based Representation

- Our method assigns pixels to 3D surfaces (**planes** and **B-splines**)
- Our contributions (soft segmentation term and MDL prior) would most likely not work if a disparity-based representation was chosen.

Mapping of pixels to surfaces

Surfaces implicitly define disparities

Reference image

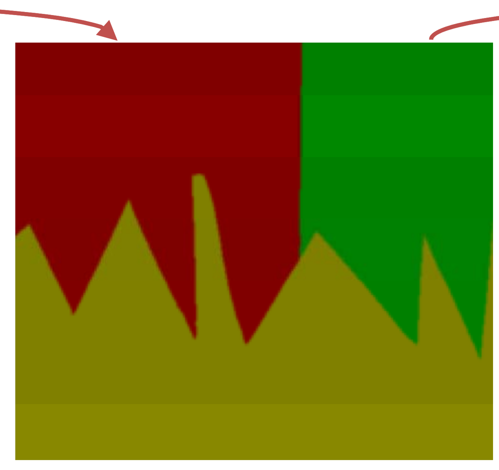

Surface map (Identical colors mean identical 3D surface)

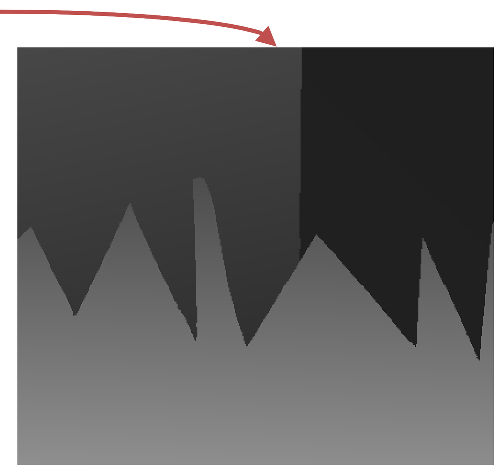

Disparity map

Energy Model

- We search an assignment of pixels to surfaces that minimizes an energy function:

Data Term: Measures pixel dissimilarity via Mutual Information – Imposes a penalty for occluded pixels

Smoothness Term: Penalizes neighboring pixels (4-connectivity) that are assigned to different surfaces

Soft Segmentation Term: Penalizes violations of segmentation assumption (described bellow)

MDL Term: Puts a penalty on the number of surfaces (described bellow)

Curvature Term: Penalizes surface curvature at each pixel

Soft Segmentation Term

- Most top-performing methods assign whole segments to a single surface. (Typically planar models are used.)
- They cannot recover from cases where a segment overlaps a disparity discontinuity. (Segmentation is a **hard constraint**.)
- We **only prefer** solutions that are consistent with a given segmentation by assigning a lower energy. (Segmentation is a **soft constraint**.)

Map reference image

Ground truth disparities

Result of using segmentation as a hard constraint

Our soft segmentation result

Construction:

- We are given a color segmentation of the left image.
- For each pixel *p*, we generate a **sub-segment** *Lp* by intersecting a square window centered on *p* with *p*'s color segment (see figure bellow).
- We give 0 penalty if all pixels within the sub-segment *Lp* are assigned to the same surface and a constant penalty, otherwise.

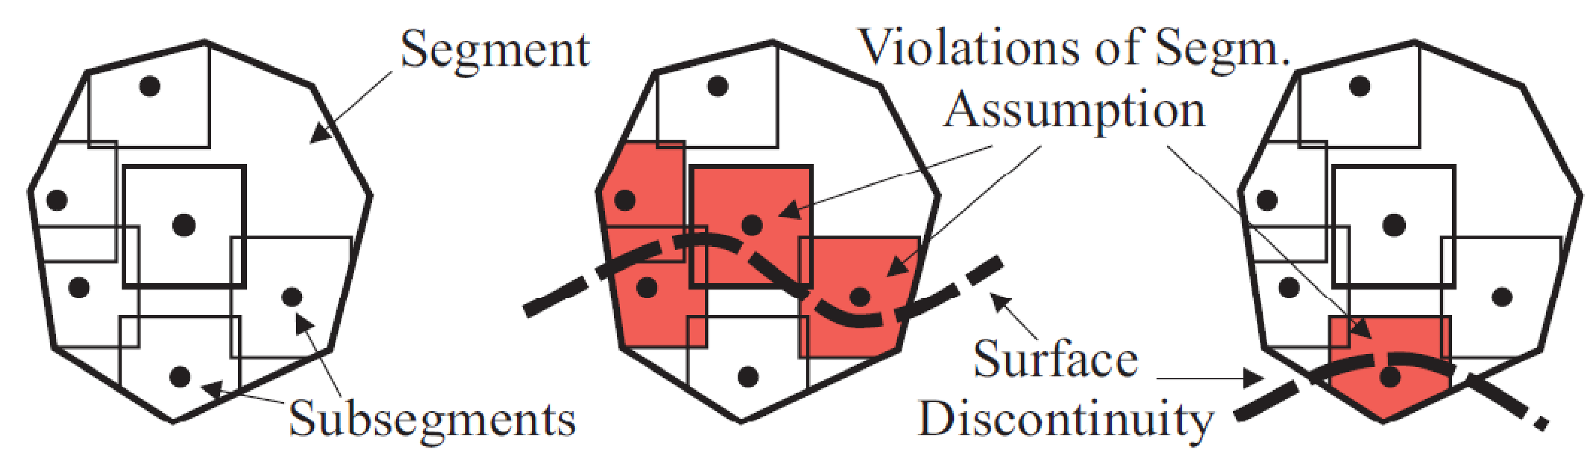

Construction of overlapping sub-segments

Soft segmentation term at work. Red sub-segments contain more than one surface. A penalty is imposed for these sub-segments.

MDL Term

- A simple explanation of the scene (small number of surfaces) is often better than an unnecessarily complex one (large number of surfaces).
- We impose a penalty on the number of surfaces. (A solution consisting of 5 surfaces is cheaper than a solution containing 100 surfaces.)

Crop of the Cones image

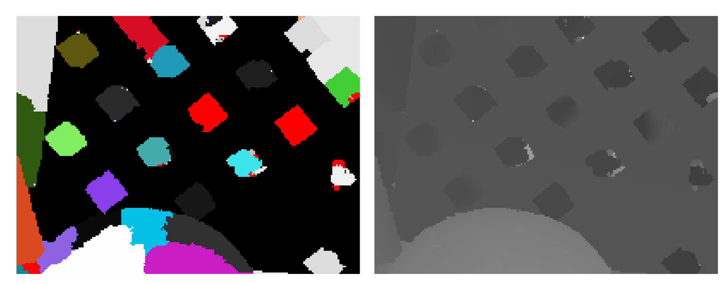

Solution without using our MDL term. The background is erroneously modeled by many slightly different surfaces.

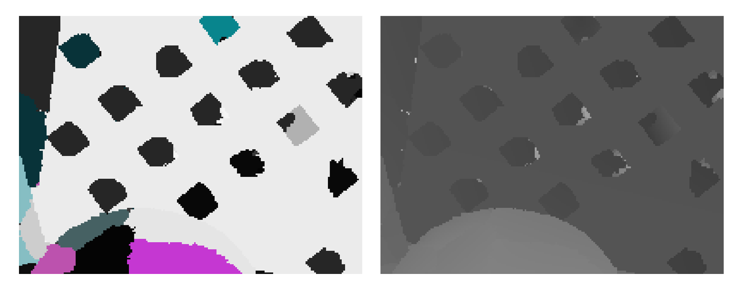

Our MDL term. The background is largely modeled by a single surface.

Energy Optimization

- Fusion move approach [Lempitsky, ICCV09] enables powerful graph cut-based optimization, despite the infinite number of possible labelings (surfaces).
- Computing the "optimal" fusion move:
  - We apply recent work on sparse higher-order cliques ([Kohli, CVPR07] and [Rother, CVPR09]) for implementing the soft segmentation term.
  - The non-submodular pseudo-boolean energy is optimized via QPBOI.
- We describe 6 different ways for computing good proposals in the paper.

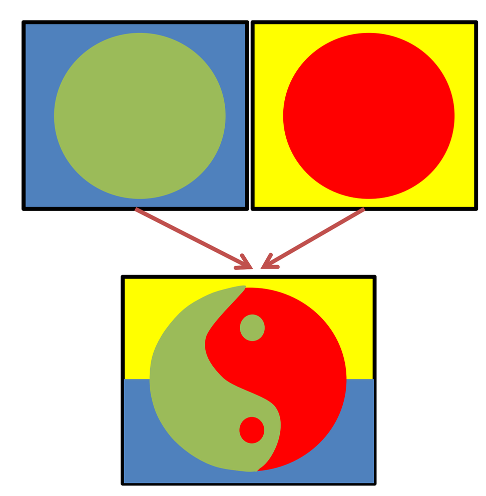

The fusion move

Results

- 6th rank out of over 80 submissions in the Middlebury online table
- 1th rank for the complex Teddy set on all error measures
- Lowest average error percentage among all methods

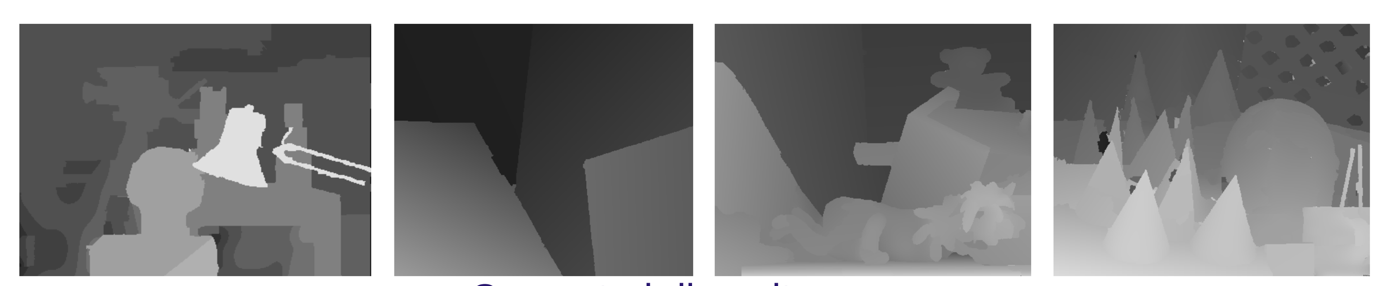

Computed disparity maps

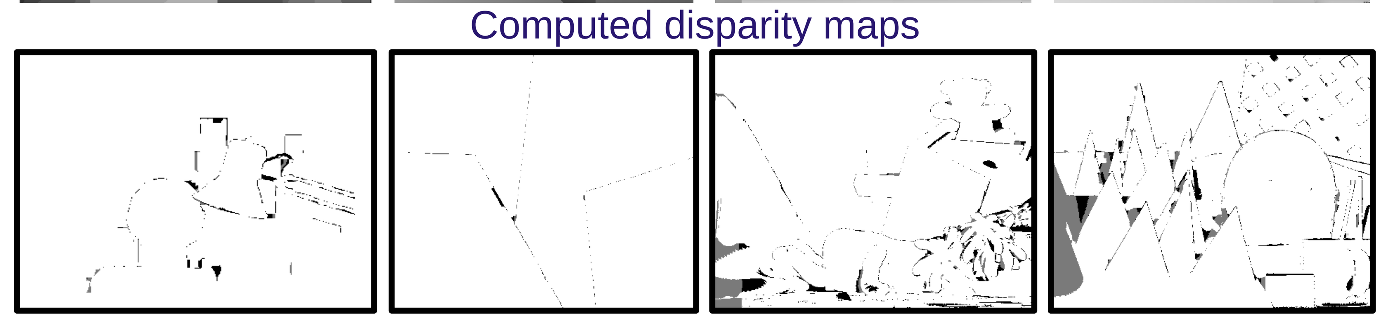

Disparity errors > 1 pixel

Assignment of pixels to surfaces

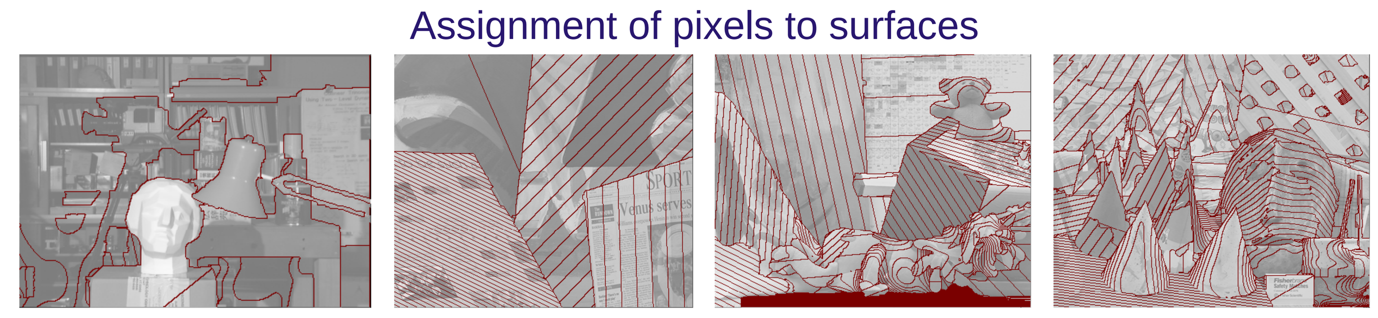

Left images with contour lines overlaid

* Michael Bleyer received financial support from the Austrian Science Fund (FWF) under project P19797 and the Vienna Science and Technology Fund (WWTF) under project ICT08-019.