Chapter 11
Convex Relaxation Methods II: Multiview Reconstruction

Computer Vision I: Variational Methods

Winter 2018/19

Prof. Daniel Cremers
Chair of Computer Vision and Artificial Intelligence
Departments of Informatics & Mathematics
Technical University of Munich

Overview

1. **Convex Relaxation Methods: Recent Developments**
2. **Variational Multi-view Reconstruction**
3. **Silhouette-Consistency**
4. **Convex Multi-view Reconstruction**

Convex Relaxation Methods: Recent Developments

In Chapter 9, we saw that certain optimization problems arising in computer vision can be addressed by means of convex relaxation methods. In particular, Chan, Esedoglu, Nikolova (2006) showed that the two-region segmentation problem (for fixed region models) can be solved optimally by solving a convex problem and subsequent thresholding.

In the wake of the above work, a number of papers have aimed at generalizing this result to other classes of optimization problems arising in computer vision. This is a very active and ongoing effort.

Among various approaches, we will in the following discuss approaches to multiple-view reconstruction from calibrated images or videos. These were developed in Kolev et al., IJCV 2009, Cremers, Kolev, PAMI 2011 and Oswald et al., ECCV 2014.

Multiple-view Reconstruction

Multiple-view reconstruction deals with the reconstruction of geometrical structure observed in multiple images.

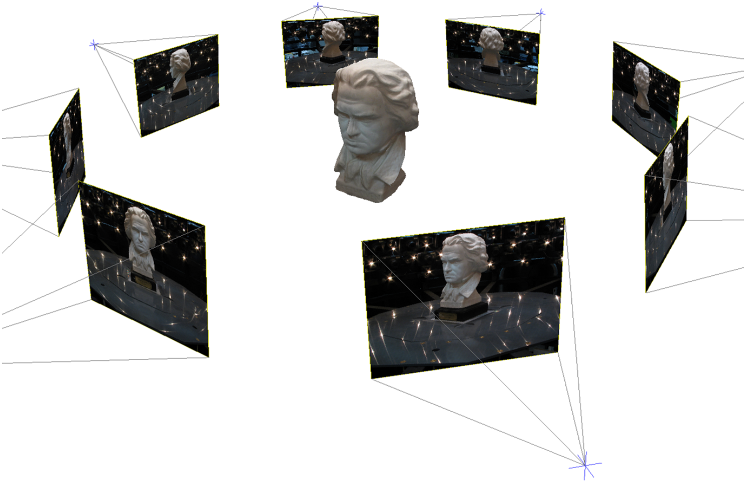

Variational Multiple-view Reconstruction

There are many ways to obtain 3D geometry from calibrated images. Among the most successful techniques is the concept of stereo photoconsistency.

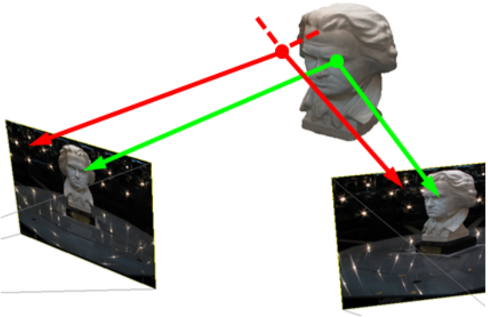

The key idea is to associate each voxel (volume element) in a specified volume $V \subset \mathbb{R}^3$ a photoconsistency function $\rho : V \to [0, 1]$ which takes on values near 0 if the projection into pairs of images gives rise to similar colors (or local patch texture) and values near 1 otherwise.

Variational Multiple-view Reconstruction

Faugeras and Keriven (1998) proposed an approach to reconstruct dense surfaces from multiple calibrated images. Essentially they aim at finding optimally photoconsistent reconstructions by minimizing the following type of functional on the space of surfaces S:

$$\min_S \int_S \rho(s)\, dA(s),$$

where $dA(s)$ denotes the area element at location $s \in S$. The key idea is that a surface is optimally photoconsistent if for all surface elements the value of ρ is small.

In the absence of photoconsistency information, i.e. in constant-color regions, ρ is spatially constant and the functional will boil down to a Euclidean minimal surface.

Yet, a key drawback of such approaches is their so-called shrinking bias: The globally optimal solution of the above problem is the empty set $S = \emptyset$ which gives rise to a cost of 0.

Variational Multiple-view Reconstruction

Faugeras and Keriven employ level set methods in order to minimize the functional locally.

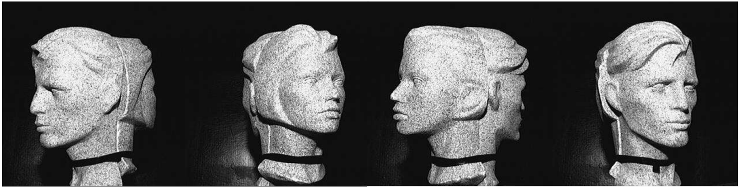

4 out of 18 images of a double-head

reconstruction process

Source: Faugeras, Keriven, TIP 1998

Silhouette-Consistent Multiple View Reconstruction

In Cremers, Kolev PAMI '11 it is proposed to impose silhouette-consistency in the photoconsistency-weighted minimal surface formulation.

The resulting optimization problem is given by:

$$\min_S \int_S \rho(s)\, dA(s),$$

$$\text{s.t. } \pi_i(S) = S_i \quad \forall i = 1, \dots, n.$$

Here π_i denotes the projection into the image i and S_i denotes the silhouettes of the object observed in this image.

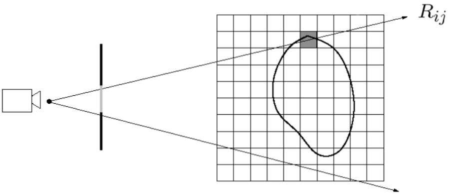

Convex Relaxation of Multiview Reconstruction

Let $u : V \to \{0,1\}$ denote the indicator function of the interior of S. Then Cremers, Kolev (PAMI 2011) propose to rewrite the silhouette-constrained photoconsistency-weighted minimal surface problem as follows:

$$\min_{u:V\to\{0,1\}} \int_V \rho(x)|\nabla u|\,dx$$

$$\text{s.t. } \forall i,j: \int_{R_{ij}} u(x)\,dx \geq 1, \quad \text{if } j \in S_i,$$

$$\int_{R_{ij}} u(x)\,dx = 0, \quad \text{else},$$

where R_{ij} denotes the visual ray for camera i through pixel j.

Intuition: If a pixel j in image i is inside the silhouette area then the ray R_{ij} must intersect the object somewhere, i.e. one of the u-values must be 1. Otherwise all u-values along this ray must be 0.

Convex Relaxation of Multiview Reconstruction

Dropping the integrality constraint leads to the following relaxed convex problem:

$$\min_{u:V\to[0,1]} \int_V \rho(x)|\nabla u|\,dx$$

$$\text{s.t. } \forall i,j: \int_{R_{ij}} u(x)\,dx \geq 1, \quad \text{if } j \in S_i,$$

$$\int_{R_{ij}} u(x)\,dx = 0, \quad \text{else},$$

Interestingly, in this implicit representation of geometry, the silhouette-consistency constraints are simple linear constraints.

One can now solve the above convex optimization problem and binarize the solution in order to obtain a provably silhouette-consistent reconstruction. Since there is no thresholding theorem, the binary solution is no longer provably optimal but merely within computable bounds of the optimum.

Reconstruction of Thin and Elongated Structures

3 out of 24 input images

Reconstructed geometry

Source: Cremers, Kolev, PAMI ’11

Reconstruction of Texture-less Objects

3 out of 27 input images

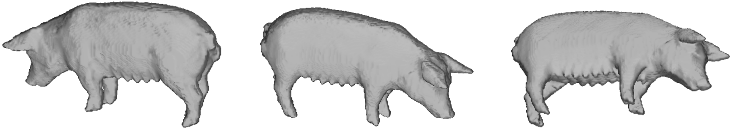

Reconstructed geometry

Source: Cremers, Kolev, PAMI ’11

Reconstruction of the Niobid Statues

3 out of 28 input images

Reconstructed geometry

Source: Cremers, Kolev, PAMI ’11

Convex Space-time Reconstruction from Multiview Video

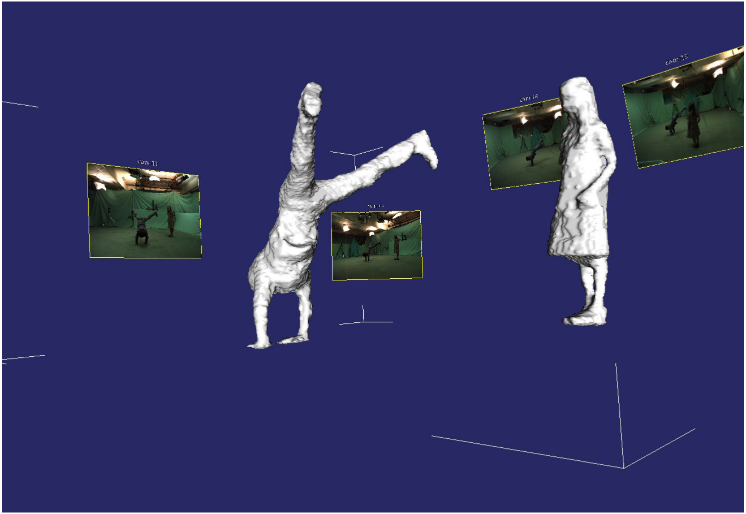

Source: Oswald, Cremers ’13

Convex Space-time Reconstruction from Multiview Video

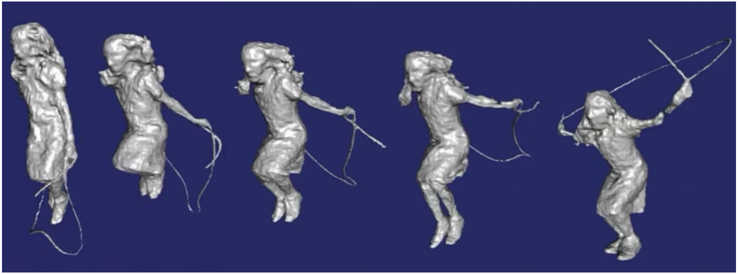

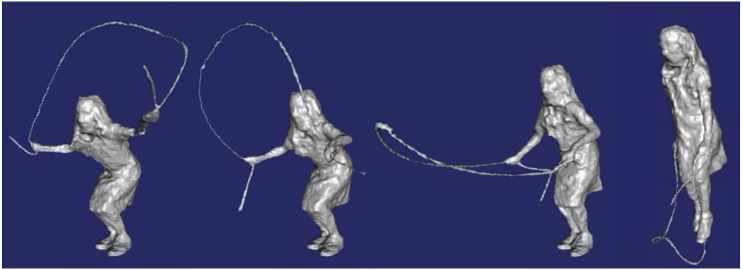

Oswald, Stühmer, Cremers, ECCV '14