

Shape from Silhouettes I

CV book Szelisky 11.6.2



Guido Gerig

CS 6320, Spring 2012

(slides modified from Marc Pollefeys

UNC Chapel Hill, some of the figures and slides
are adapted from M. Pollefeys, J.S. Franco, J.
Matusik's presentations, and referenced papers)

Shape from silhouettes



Slides from
Lazebnik,
Matusik
Yerex
and others

Automatic 3D Model Construction for Turn-Table Sequences,
A.W. Fitzgibbon, G. Cross, and A. Zisserman, SMILE 1998

Big Picture



- Multi-camera environments

- Dynamic scene




Outdoor data capturing with 9 video cameras behind the Ackland Museum, UNC-Chapel Hill, 2006/8/24. Pictured by Jae Hak Kim.

- N cameras observe the scene and produce N video streams

- What can we do with this data?




Motivation: Movies



Sinha Sudipta, UNC PhD 2008




Motivation: Movies



Sinha Sudipta, UNC PhD 2008



Motivation: Movies




Sinha Sudipta, UNC PhD 2008


What can we do with this data?




- Reconstruct scene objects:
 - shape from silhouettes
 - photo-consistency



- Calibrate cameras
 - recover epipolar geometry



- Fit specific models (articulated models)



Outline



- Silhouettes
 - basic concepts
 - extract silhouettes
 - fundamentals about using silhouettes
 - reconstruct shapes from silhouettes
 - use uncertain silhouettes
 - calibrate from silhouettes
- Perspectives and interesting ideas

Silhouettes of objects of interest



- Silhouettes are the regions where objects of interest project in images
- Silhouettes can generally be obtained using low level information (fast)
- They give information about the global shape of scene objects

How to extract silhouettes?



- Sometimes done manually (for offline applications, ground truth and verifications)
- Region based-extraction (automatic)
 - silhouette extraction is a 2-region image segmentation problem, w/ specific solutions:
 - chroma keying (blue, green background)
 - background subtraction (pre-observed static or dynamic background)

(refer to segmentation course)

How to extract silhouettes?

- Contour-based extraction
- focus on silhouette *outline* instead of region itself
 - snakes, active contours: fitting of a curve to high gradients in image, local optimization






How to extract silhouettes? (cont.)

- Background subtraction
 - Simple thresholding
 - Train an appearance model for each pixel, from a set of background images
 - RGB 3D-Gaussian model
 - HSV model
 - GMM model
 - Non-parametric model (histogram/kernel density function)
 - Apply the pixel color to the model, then classify it to be foreground/background
- We will talk about this in more detail later



Why use a Visual Hull?

- Good shape representation
- Can be computed efficiently
- No photo-consistency required
- As bootstrap of many fancy refinement ...



-



=

foreground



Outline




- Silhouettes
 - basic concepts
 - extract silhouettes
 - fundamentals about using silhouettes
 - reconstruct shapes from silhouettes
 - use uncertain silhouettes
 - calibrate from silhouettes
- Perspectives and cool ideas




What is shape from silhouette?

- The silhouette, or occluding contour of an object in an image contains some information about the 3D shape of the object.
- Given a single silhouette image of an object, we know that the 3D object lies inside the volume generated by back-projecting the silhouette area using the camera parameters.






What is shape from silhouette?




- With multiple views of the same object, we can intersect the *generalized cones* generated by each image, to build a volume which is guaranteed to contain the object.
- The limiting smallest volume obtainable in this way is known as the *visual hull* of the object.

One-view silhouette geometry






Multi-view silhouette geometry: the Visual Hull



- **Maximal volume consistent with silhouettes**
[Laurentini94] [Baumgart74]
- Can be seen as the intersection of viewing cones
- Properties:
 - Containment property: contains real scene objects
 - Converges towards the shape of scene objects **minus concavities** as N increases
 - Projective structure: simple management of visibility problems

Visual Hull: A 3D Example



Outline




- Silhouettes
 - basic concepts
 - extract silhouettes
 - fundamentals about using silhouettes
 - reconstruct shapes from silhouettes
 - use uncertain silhouettes
 - calibrate from silhouettes
- Perspectives and cool ideas


Algorithms




- Standard voxel based method
 Marching Intersections



- Exact polyhedral methods



- Image-based visual hulls





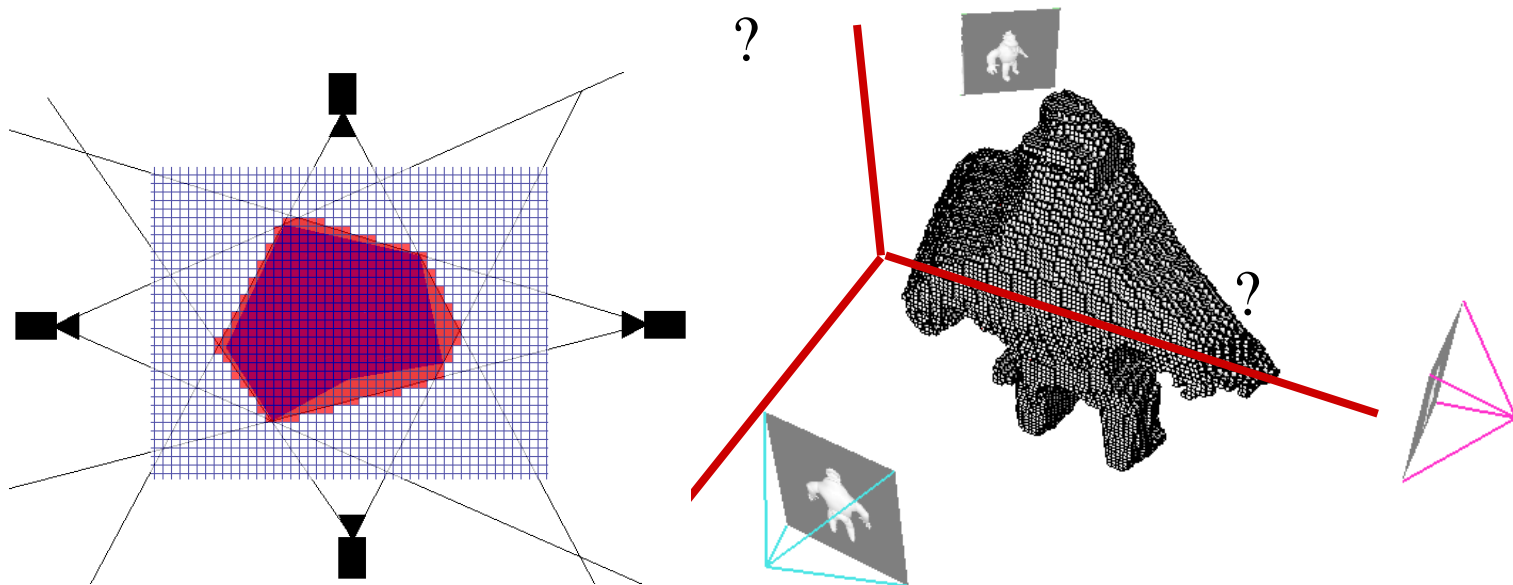
Voxel based

- First the object space is split up into a 3D grid of voxels.
- Each voxel is intersected with each silhouette volume.
- Only voxels that lie **inside all silhouette** volumes remain part of the final shape.

Visual hull as voxel grid






- Identify *3D region* using voxel carving
 - does a given voxel project inside all silhouettes?



- pros: simplicity
- cons: bad precision/computation time tradeoff

Classical voxel grid improvement: octrees

- Same principle, but refinement through space subdivision



[Szeliski TR 90']



Marching intersections

- The object space is again split up into a 3D grid.
- The grid used is made of 3 sets of rays, rather than voxels.
- Rays are aligned with the 3 axes, and store points of entry/exit into the volume
- Each silhouette cone can be converted to the marching intersections data structure.
- Then merging them is reduced to 1D intersections along each ray.

M. Tarini et al, *Marching intersections, An efficient Approach to Shape from Silhouette*

Marching intersections - example





Figure 1: The *MI* data structure and conversion algorithm in a 2D example: (top-left) the curve to be processed and the *MI* data structures collecting the intersections between the input curve and the vertical (top-center) and horizontal (top-right) lines of the user selected grid; (bottom-left) based on the horizontal and vertical intersections, an *MC* look up table entry code is located for each not empty (virtual) cell; (bottom-right) the reconstructed curve. The vertical intersections inside the small dotted box are collected and then deleted from the algorithm because belonging to the same virtual cell. This removal is a prerequisite for the correct operation of the algorithm and it leads to the removal of high frequency details.

Marching intersections -Concept



M. Tarini et al,
Marching intersections, An efficient Approach to Shape from Silhouette

- Given a curve
- Select reference grid
- Intersections between curve and horizontal and vertical lines: MI
- Create look-up-table for each non-empty box

Marching intersections - Silhouettes




Figure 5: Surface of a conoid extracted from a MI structure obtained with algorithm of Fig. 4.

```
1 for each ray-set (X,Y,Z)
2   for each ray (i,j) in it
3     compute the projection on I of ray(i,j)
4     scan-convert the corresponding 2D line:
5     for each intersection k with silhouette found
6       remove perspective distortion, obtaining k'
7       add k' to ray(i,j) of the current ray-set
```

Figure 4: A pseudo code for the conversion into a MI structure of a truncated conoid implicitly defined by a 2D silhouette on an image I and by a camera position.

- Convert conoid structures to MI datastructure
- Intersection tested in 2D image: purely 2D operation
- Intersection of conoids: AND operations on MI datastructures

Marching intersections - Silhouettes



Figure 6: One object (top left) and its silhouette with 2D lines traced over it to find intersections along rays in the X , Y and Z ray-set of the MI, respectively. The number of lines has been reduced for illustration purposes: in typical cases they would cover most of the area of the image.




Figure 7: Some slices of the conoid visible in Fig. 5, composed by intervals defined by some of the MI ray along X and Y . Notice that each slice is the replication of another one at a different scale. Cropping is applied on the sides of the volume to make the front and back faces.

Final step: Convert MI datastructure representing all intersections to triangular mesh

Marching intersections - Silhouettes




Figure 9: Examples of reconstruction results: the Lady, obtained by silhouettes like the one in Fig. 6, and the familiar Stanford bunny. In both cases a $256 \times 256 \times 256$ MI structure has been carved by 128 (synthetic) 1000×1000 silhouettes images taken from a (virtual) camera rotating around it (like when a rotating plate is used). The two meshes are composed by 257K and 451K faces respectively.


Size	Hits	Time0	Time1
Lady silhouettes			
64	12.3%	0.692	0.070
128	44.0%	1.805	0.310
256	80.2%	2.643	1.542
512	94.6%	3.190	8.021
Bunny silhouettes			
64	15.3%	1.350	0.091
128	52.8%	3.079	0.511
256	85.1%	4.049	2.604
512	96.0%	4.895	13.189

Figure 8: Some results of our application running on a normal PC (Athlon 1.4 GHz 512 MB). *Size* is the grid size, *Hits* is the percentage of cache hits, that is, how many out of the $3Size^2$ rays composing the structure were found in the cache described in Sec. 4.2 and were not scan-converted. *Time0* is the time in seconds required to convert a silhouette into a MI structure for the conoid, and to intersect it with the previous MI structure. *Time1* is the time in seconds to extract the final mesh from the MI structure. Total time to compute the models shown in Fig. 9 is $((Time0 \times \text{number of silhouettes}) + Time1)$, which is 340 sec. for the Lady and 577 sec. for the Bunny.






Example

- Compute visual hull with silhouette images from multiple calibrated cameras
- Compute Silhouette Image
- Volumetric visual hull computation
- Display the result



Algorithms

- Standard voxel based method
- Exact polyhedral methods
- Image-based visual hulls





Exact Polyhedral Methods

- First, silhouette images are converted to polygons. (convex or non-convex, with holes allowed)
- Each edge is back projected to form a 3d polygon.
- Then each polygon is projected onto each image, and intersected with each silhouette in 2D.
- The resulting polygons are assembled to form the polyhedral visual hull



Exact Polyhedral Methods






Fig. 1. A single silhouette cone face is shown, defined by the edge in the center silhouette. Its projection in two other silhouettes is also shown.



Exact Polyhedral - example



Metric Cameras and Visual-Hull Reconstruction from 4 views



Final calibration quality comparable to explicit calibration procedure

Image-based Visual Hulls

IBVH Results

- Approximately constant computation per pixel per camera
- Parallelizes
- Consistent with input silhouettes





IBVH Results



Figure 6 – Four segmented reference images from our system.




Figure 8 - Example IBVH images. The upper images show depth maps of the computed visual hulls. The lower images show shaded renderings from the same viewpoint. The hull segment connecting the two legs results from a segmentation error caused by a shadow.



Image-based Visual Hulls

<http://www.youtube.com/watch?v=Lw9aFaHobao>



What if my silhouettes are noisy?



- Idea: we wish to find the content of the scene from images, as a probability grid
- Modeling the forward problem - explaining image observations given the grid state - is easy. It can be accounted for in a *sensor model*.
- Bayesian inference enables the formulation of our initial inverse problem from the sensor model
- Simplification for tractability: independent analysis and processing of voxels



Idea

- Unreliable silhouettes: do not make decision about their location
- Do sensor fusion: use all image information simultaneously
- Fusion of “probability information”: Related to Radon Transform (used for Computer Tomography Reconstruction)

