


Multi-View 3D-Reconstruction

Alexander Ladikos

Computer Aided Medical Procedures (CAMP)
Technische Universität München, Germany

Problem statement

- Given a set of images of an object we want to recover its shape and appearance
- Only the information contained in the images is used
- No active techniques like laser scanners or structured light are used
- The intrinsic and extrinsic camera parameters of the images may or may not be known
- The output is usually a textured mesh



shape



shape and appearance

Applications: Real-Time 3D-Reconstruction

Submission Id: 0063

VIRTUALIZATION GATE

**INRIA / Grenoble Universities
4D View Solutions**

Applications: Large-Scale Reconstruction




Applications: Large-Scale Reconstruction




Shape-from-Silhouette

- Shape-from-Silhouette is a technique for recovering the shape of an object called the visual hull from a set of calibrated silhouette images
- It is purely based on geometric information
- Often used in real-time systems due to its simplicity and speed
- Biggest drawback is that it cannot recover concavities




Shape-from-Silhouette: Visual Hull

The visual hull is the shape maximally consistent with the silhouettes



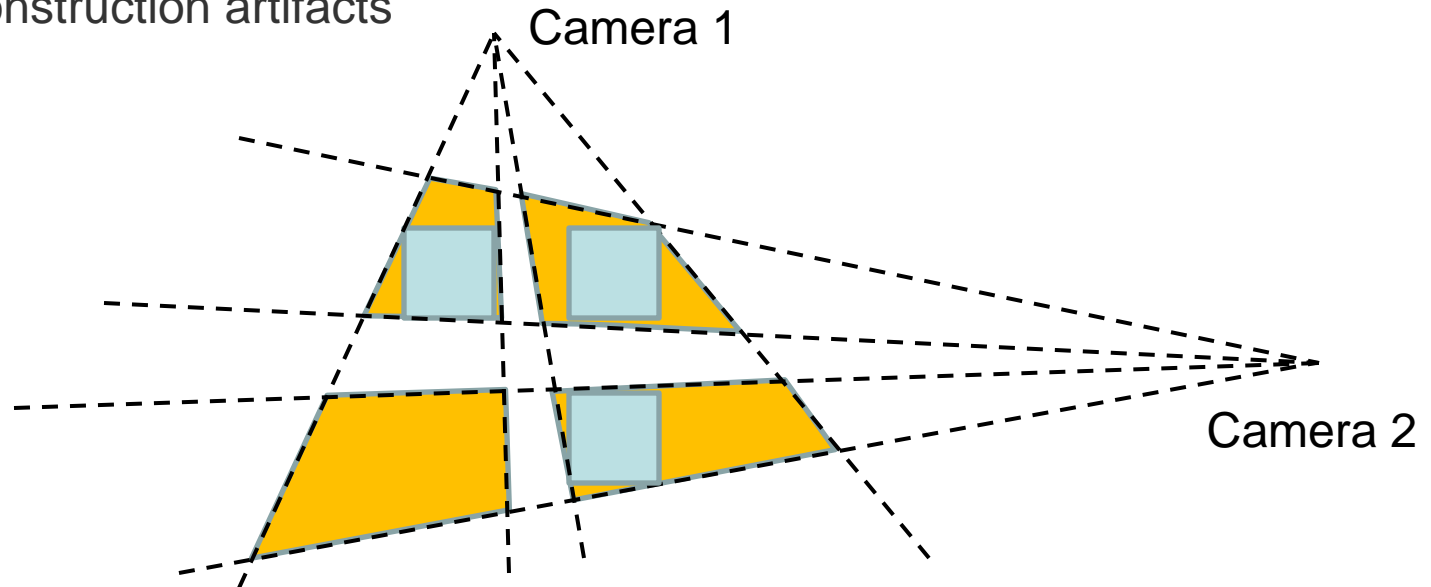
Shape-from-Silhouette: Limitations

The visual hull cannot capture concavities not visible in the silhouettes



Shape-from-Silhouette: Properties

- The accuracy of the reconstruction is limited by the number and the position of the cameras
- As the number of cameras approaches infinity the reconstruction will reach the true shape (except for the concavities)
- In the case of multiple objects in the scene a too small number of cameras can lead to reconstruction artifacts




Shape-from-Silhouette: Algorithms

- There are two approaches for computing the visual hull
 - Volumetric approach
 - Polyhedral approach
- Due to its simplicity and efficient implementations the volumetric approach is often used
- It is based on a subdivision of the reconstruction volume into basic computational units (voxels)
- An independent computation determines the properties of each voxel (i.e. inside or outside the visual hull)
- There are many methods for speeding this method up
 - Octrees
 - Parallel implementation
 - GPU-based implementation


Shape-from-Silhouette: Volumetric Approach

- Discretize the volume into voxels
- Check occupancy for each voxel
 - Project voxel into each silhouette image
 - If the projection is outside at least one silhouette the voxel is empty
 - If the projection is inside all silhouettes the voxel is occupied

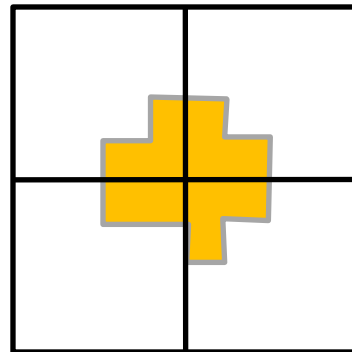


Shape-from-Silhouette: Volumetric Approach


- Octrees allow to speed up the computation by using an adaptive discretization
- If an octree cell is partially occupied it is subdivided further
- If it is empty or full it is not subdivided




Level 1



Level 2



Level 3




Level 4

Shape-from-Silhouette: Volumetric Approach

- The volumetric approach can be implemented very efficiently
- Since the occupancy test for each voxel is independent of the other voxels, the method can be parallelized easily
- Parallel octree implementations can achieve real-time performance
- Accuracy of reconstruction is limited by voxel resolution
- When using a too low spatial resolution discretization artifacts appear in the output
- If a mesh is desired as output the reconstructed volume has to be converted using the Marching Cubes algorithm

Shape-from-Silhouette: Polyhedral Approach

- In the polyhedral approach the silhouettes are backprojected and intersected in 3D
- It does not suffer from discretization artifacts
- Precision is only limited by silhouette resolution
- Problematic when silhouettes contain errors
- Does not scale very well with the number of images
- Implementation is non-trivial
- Good performance in real-time systems




Shape-from-Silhouette: Real-time Systems

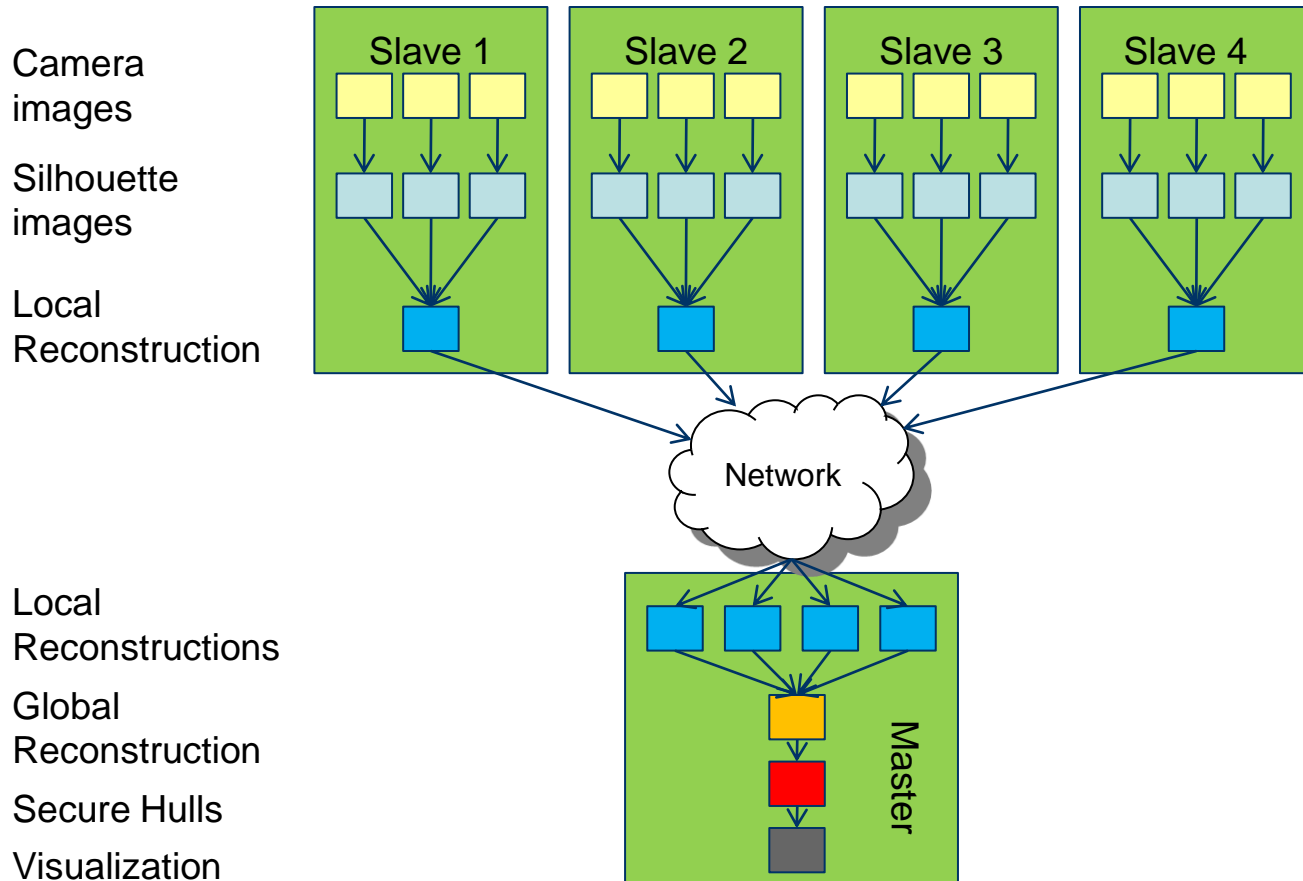
- There are many real-time systems in existence today
- Nowadays it is comparatively easy to build such a system due to the progress in computational power at reasonable prices
- Most systems use the volumetric approach due to its simplicity
- One notable exception is the GrImage platform at INRIA Rhône-Alpes

Shape-from-Silhouette: CAMP system


- 16 synchronized cameras mounted on the ceiling
- Working volume 3.5m x 3.5m x 2.5m
- Runs at 30 Hz




Shape-from-Silhouette: CAMP system



Shape-from-Silhouette: CAMP system



Shape-from-Silhouette: CAMP system



Shape-from-Silhouette: Summary

- Advantages
 - Only silhouette images are required
 - No need for correspondences or texture
 - Robust
 - Efficient and easy to implement
- Disadvantages
 - Cannot recover concavities not seen in the silhouette images
 - Artifacts for complex scenes and low number of cameras
 - Needs calibrated input images
 - Silhouettes have to be available

Space Carving

- Up to now only the silhouette images were used
- Due to this we could for instance not recover concave regions of an object which are not visible in the silhouette images
- By using the photometric information in the images (i.e. the actual images) more details can be reconstructed
- This makes the reconstruction process more complicated
- Many methods for photometric reconstruction for photometric reconstruction exist
- Space carving is one of the earliest methods
- It reconstructs the shape by removing all parts of the scene which are not consistent with the input images

Space Carving: Idea

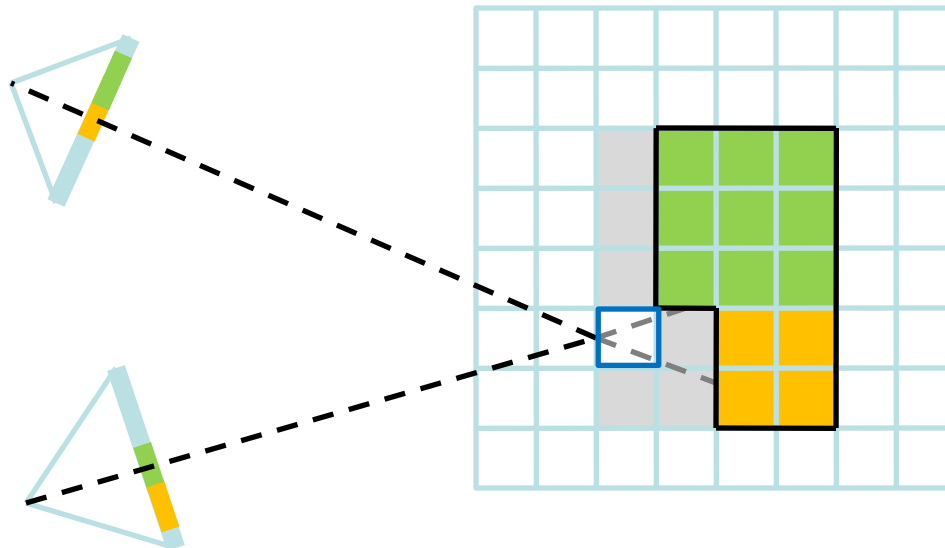
- Given an initial reconstruction of an object (bounding box, visual hull, ...) we want to test if it is consistent with the input images
- To test this we need to check whether the color of a surface point is consistent across all the input images
- If it is, it is called **photoconsistent**
- The measure used to compute the photoconsistency is called the **photoconsistency measure**
- For a perfect reconstruction every point on the object has to be photoconsistent
- Space carving computes the most photoconsistent surface starting from an initial surface estimate by eliminating the non-photoconsistent parts
- The photoconsistent surface is called the **photo hull** and is the tightest possible bound on the true scene

Space Carving: Photoconsistency Measures

- There are a number of photoconsistency measures used
 - Color variance
 - Color histograms
 - SSD (requires some assumption/knowledge about the orientation of the surface)
 - NCC (requires some assumption/knowledge about the orientation of the surface)
- Often the color variance is used
- Problems occur when the surface has highlights or non-uniform lighting
- This leads to bad photoconsistency scores and the removal of correct voxels



Space Carving: Algorithm

- Initialize a volume with a superset of the true scene
- Project a surface voxel into all images in which it is visible
- Remove if not photoconsistent
- Repeat until convergence




Space Carving: Visibility Ordering

- When evolving the surface the visibility has to be considered
- The photoconsistency is only evaluated in the views in which a voxel is visible
- When a voxel is deleted new voxels become visible and the visibility has to be updated
- This is efficiently done using a multi-pass plane-sweep algorithm
- Scene is swept with a plane in each of the six principle directions and only cameras on one side of the plane are considered



Space Carving: Limitations

- The photohull is only guaranteed to be the tightest superset of the true reconstruction
- If a voxel is wrongly removed it can lead to the removal of other correct parts of the object
- Photoconsistency measure is critical




Space Carving: Limitations

- Needs calibrated input images
- Problematic for non-lambertian objects
- Reconstructed photo hull is only a superset of the true shape
- Greedy approach
 - No regularization (e.g. smoothing)
 - Removed voxels cannot be readded to the reconstruction
- Accuracy limited by voxel resolution
 - Voxels should be small
- Discretization artifacts

Space Carving: Results



Space Carving: Results



(c)



(d)



(e)



(f)

Space Carving: Summary

- Reconstructs objects using photoconsistency
- An initial shape is updated until it is photoconsistent
- Comparatively easy to implement