Quantification of gaseous structures with volumetric reconstruction from visual hulls

S. Seipel[1,2] and P. Jenke[1]

[1]Faculty of Engineering and Sustainable Technology, University of Gävle, Sweden
[2]Centre of Image Analysis, Uppsala University, Sweden

Abstract

3D reconstruction from visual hulls is a well established technique for camera based reconstruction of 3D objects in computer graphics. We propose in this paper to employ visual hull techniques to quantify the volume of diffusely defined gaseous structures. In our evaluation, visual quality of the 3D reconstructions is secondary. Instead, using synthetic ground truth data, we determine the number of independent silhouette images needed to achieve a stable volume estimate. We also estimate the influence of different segmentation results of the silhouette images on final volume estimates. Our results show that comparably few camera views yield to convergent volume estimates. For the type of 3D data studied, visual hull reconstructions overestimate actual volumes with about 50%. This proportion seems to be consistent for different data sets tested and may serve for re-calibration of volume estimation of gaseous structures.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Digitizing and scanning I.4.1 [Computer Graphics]: Digitization and Image Capture—Imaging geometry I.4.9 [Computer Graphics]: Applications—I.4.10 [Computer Graphics]: Image Representation—

1. Introduction

In context of environmental analysis, methane (CH_4) is one of the most critical greenhouse gases, as it remains in the atmosphere for approximately twelve years. According to Wuebbles [Wue02] methane is more than 20 times more effective in trapping heat in the atmosphere than carbon dioxide (CO_2). Methane gas is released from a variety of natural and man-made sources. Emissions of methane from open dumping and improper land filling of municipal solid waste are estimated to contribute to 3 - 19% of the anthropogenic emissions word-wide [Tal07]. Identification, quantification, and reduction of CH_4 emissions from landfills can make a significant contribution to decrease the greenhouse gas stock. Among the largely available and low cost methods to detect methane gas are pellistor based sensors. They are used for fixed point measurements of gas concentrations at specific positions. Recently, infrared (IR) based measuring principles have been used to develop new sensors for fixed point measurements. They exploit the fact that many gases absorb energy within a narrowly defined band in the IR spectrum. While traditional fixed point measurements can reveal concentrations of gases at specific locations, they are of limited value for quantifying the actual volume and its spatial distribution. Very recent research aims at using novel IR imaging devices to capture gas distributions [SM10]. In previous work segmentation methods based on optical flow analysis have been proposed to identify the presence of gas in the camera image [hS11].

The objective of the work presented here is to quantify and describe the 3D spatial distribution of gaseous structures using IR camera images from multiple views. In particular we aim to determine the number of camera views needed in order to obtain volume estimates of sufficient precision. Another objective is to evaluate the potential effect of different segmentation results of the 2D silhouette images on the final 3D volume estimate.

2. Related work

Reconstruction of 3D geometric objects is the purpose of several different multi-view geometry algorithms in computer vision. One class of these algorithms identifies feature points in different views [HZ03]. Using a RANSAC procedure, a robust estimate of the fundamental matrix is obtained by applying the eight-point algorithm based on best estimates for corresponding points [Har97]. The fundamental matrix describes the geometric relationships between two camera views and it is subsequently used to triangulate the positions in 3D space for sets of homologous feature points [LF95]. One advantage of these image based algorithms is their capability to describe surfaces of objects through dense point clouds, which form the basis of highly detailed polygonal surfaces.

Images of methane gas (and other gases) generated by IR imaging devices are, however, not well suited for this type of multi-view 3D reconstruction approach. As is shown in Figure 1, gas becomes even in the infrared spectral range apparent only as a transparent shadow, since gas only absorbs (attenuates) the IR radiation from the background. Gas clouds do therefore not expose any clearly structured surface towards the observer. Instead, they can be identified only as diffuse areas with rather fuzzy boundaries.

For this type of boundary based 2D object representations, visual hull algorithms seem to be a more feasible approach for 3D reconstruction. Reconstruction of the visual hull of an object assumes that an object is captured with many different camera views of known positions and directions [Lau94, LW10]. After determining the silhouette of the object in the camera images, silhouette edges are projected back into 3D space forming a set of general cones. The intersection of all cones represents the visual hull of the object, and it is an estimate for the geometry of the reconstructed object. Laurentini has shown that for non-concave geometries an object can be reconstructed faithfully with $O(n^5)$ silhouettes where the number of faces of this object is n [Lau97]. It is obvious that the accuracy of the reconstructed geometry increases as the number of camera views is increased. Nevertheless, concavities in an object cannot generally be reconstructed accurately [Lau95]. Most algorithms published for reconstruction of the visual hull [Sze93, LMS03, LBN08] use discrete silhouette representations from real or synthetic camera images. Efficient memory techniques have to be employed to maintain adequate resolutions for the reconstructed volume. Szeliski proposed an octree for backprojection [Sze93]. Recent implementations use 3D texture mapping [LMS03, LMS04, LBN08] and GPU accelerated algorithms [LMS04, KSK*08, LBN08] for the back-projection into 3D space. These algorithms have generally limited precision depending on texture sizes and render buffer limits.

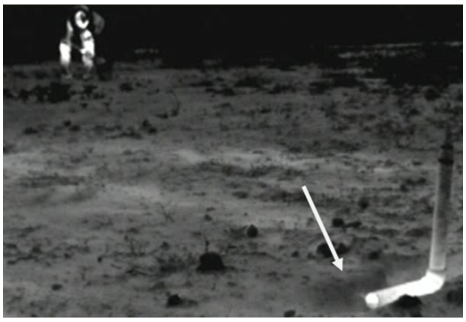

Figure 1: *Infra-red (IR) image showing methane gas leaking from a deduction pipe.*

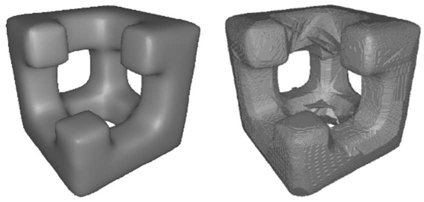

Figure 2: *An example of a polygonal model (left) and its volumetric reconstruction using visual hull methods (right).*

3. Method

3.1. Evaluation approach

Figure 2 shows an example of an original polygonal object rendered with simple Phong shading (left) and its visual hull retrieved from a set of silhouette images (right). The reconstructed object is initially a discrete voxel model, which can be transformed into an explicit polygon representation using iso-surface reconstruction. Alternatively, as is the case here, it can be rendered directly using a volume ray-caster. The picture of the visual hull in Figure 2 clearly shows that visual artifacts are inevitable as fine detail is not preserved in the reconstruction process. For most visual applications, the reconstructed objects are again texture mapped with the original camera images which makes artifacts less salient. In our specific application we aim at using visual hulls to estimate the volume of gas and smoke clouds. For this purpose, we have some different requirements and limitations:

- As the visual hull is primarily intended for volume quantification, visual artifacts are not critical.
- Monitoring of gas leakage will require expensive IR imaging, which calls for a low number of different camera positions.
- Detection of gas clouds in IR imagery is a very delicate image processing task, which at present state implies segmentation results with some uncertainty.

In our experiments we want to find out how many camera images are sufficient to reconstruct gas clouds with an acceptable quantification error. Another issue is to identify, to what extend different segmentation results in the camera images affect the quantification of the reconstructed volume. Figure 3 illustrates the different steps in our evaluation approach; they are described in the following subsections.

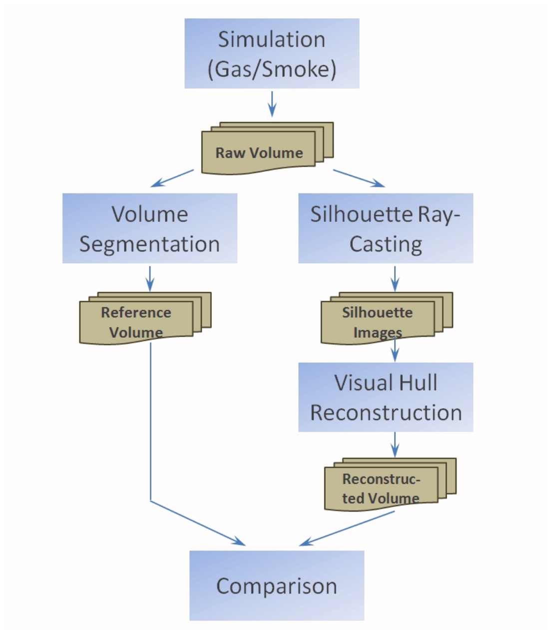

Figure 3: *Flowchart illustrating the evaluation approach.*

3.2. Data acquisition by simulation

For evaluation of the volumetric accuracy of visual hull reconstructions, a known ground truth data set is needed. Smoke and non-visible gases are, as stated before, very diffusely defined and no practical means exist to capture a reliable reference volume of gas. We therefore decided to use a synthetic dataset as reference standard. Various methods exist in computer graphics to simulate fluids and smoke, among others wavelet-based approaches for volumetric simulation [KTJG08]. For our experiment we used the reference implementation by Kim [KT] and simulated volumetric smoke on a 512^3 voxel grid. From each simulation sequence we captured 300 snapshots (volumes) and we sporadically chose a few volumes out of these sequences as reference volumes. We parametrized the smoke/gas simulation so as to accomplish plausible appearances of smoke plums.

3.3. Volume segmentation

The simulation of gas and smoke results in raw data that represents smoke through density values in the range between 0.0 and 1.0. As in a real scenario, the appearance of a gas cloud is not well defined by clear borders or contours in the density field. The identification of gas will rather depend on the employed segmentation methods and parameters; this will imply some uncertainty in regard to the actual shape of smoke. We mimic this uncertainty in the segmentation by applying different thresholds to the raw data, in order to retrieve *reference volumes* of varying conformance with the actual gas shape. The segmented reference volumes serve as ground truth for the volumes to be estimated using the visual hull reconstruction.

3.4. Silhouette image rendering

The reconstruction of the 3D geometry of the reference volume requires a set of silhouette images of the same raw dataset. We implemented a ray-caster with a simple ray evaluation function that evaluates to one as soon as a voxel above the threshold along the ray has been encountered. The resulting images are binary silhouettes of the raw volume from an arbitrary number of randomly chosen camera positions on a sphere around the raw data volume.

3.5. Visual hull reconstruction

The reconstruction of the volume using the visual hull approach is accomplished following an algorithm similar to the one published in [KSK*08]. The algorithm operates in a slice-order; that is, the volume is reconstructed by rendering subsequent slices of the volume using screen-filling quads. For each slice, every binary silhouette image is rasterized into the frame-buffer. Hereby the corresponding camera parameters are used to configure the texture projection matrix. The stencil-buffer is used to increment buffer elements where the silhouette image is not empty. The overlapping region in each slice is hence determined by all buffer elements that are equal to the number of camera views. The resulting frame-buffer for each slice is copied back into the volume memory. To reconstruct a volume of size 512^3, 512 render passes are required, each rasterizing n silhouette images into the stencil buffer. The volume resulting from this stage is in the following referred to as the *reconstructed volume*.

3.6. Comparison

The final step of our evaluation is a simple volumetric comparison, whereby classified voxels in the *reference volume* are enumerated and compared with the number of voxels in the *reconstructed volume*. Here, we take the number of voxel as a discrete measure for the volume of smoke in each volume, which allows for a fairly basic overall quantitative comparison. In order to get a more qualitative picture of the differences between *reference volume* and *reconstructed volume*, a voxel wise comparison can be achieved by calculating the subtraction volume.
Variables to be controlled in this evaluation are the number of silhouette images for visual hull reconstruction, and the chosen segmentation threshold.

4. Results

This section presents some results of our initial simulation based evaluation. Figure 4 shows in the upper row an example of one gas volume rendered with surface shading. The

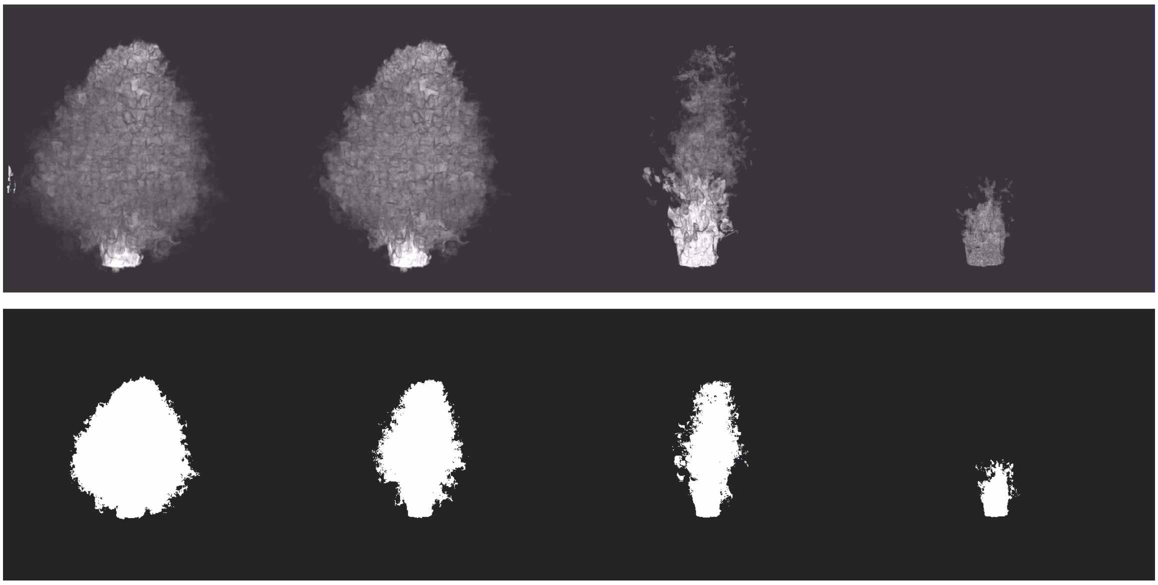

Figure 4: *Results of the simulated smoke dataset. The upper row shows surface renderings of the smoke volume for different volume segmentation thresholds. The bottom row shows silhuette images rendered with a ray-evaluation function that binarizes voxels with corresponding thresholds. Thresholds from left to right are: 0,01; 0,1; 0,2; and 0,4.*

four different pictures show the same raw data volume segmented with four different threshold values. Apparently the size and structure of the gas cloud change dramatically as the threshold is increased. In the bottom row, the same raw data volume has been rendered as a silhouette image using the corresponding threshold values and from one camera position. Note that the camera position for silhouette rendering is not exactly the same as the one used to render the pictures in the upper row. The pictures of the silhouettes illustrate that a sharp definition of the borders of a smoke/gas cloud is problematic, as smoke is very diffusely defined without sharp borders. Over- or undersegmented silhouette images are therefore likely to occur and they hamper reliable volume quantification.

Figure 5, upper row, shows surface rendered images of the reconstructed volumes. In this figure, reconstructions were accomplished with increasing numbers (3, 7, 14, and 35 images) of silhouette images and for the reference volume with a fixed threshold of 0,01 (same as the leftmost volume in Figure 4). The image series demonstrates that the overall volume for a simple convex object is reconstructed fairly well with only a few reconstruction views. Results are nevertheless not acceptable from a visual quality point of view, even in the reconstruction with 35 silhouettes. That is, visual appearance of the reconstruction still lacks visual detail that is present in the ground truth data volume. The bottom row of Figure 5 shows a semitransparent rendering of the subtraction volume. It illustrates where and how the *reference volume* and the *reconstructed volume* differ. Yellow areas indicate voxels that have been reconstructed without actually being part of the original volume. Blue voxels (on top) indicate small structures that have been lost in the reconstruction process. At large, the reconstruction is conservative and yields to volumes, which are bigger than the original volume.

A quantitative comparison of the reconstructed volume with the original volume is illustrated in Figure 6. It shows the ratio R/O of the volume sizes of the reconstructed volume with the original (reference) volume. It is clearly visible that this ratio decreases for an increasing number of silhouettes used for reconstruction, which suggests that the reconstruction of the volume becomes increasingly refined for more reconstruction views and hence better approximates the original object. Figure 6 contains observations for the same dataset and ten different segmentation thresholds.

5. Conclusion and Discussion

In this work we have investigated, if visual hull reconstructions can be used to quantify gas volumes recorded with IR imaging devices. Based on the result of our experiments we can so far conclude the following: For the purpose of volume quantification, visual hulls appear as a promising estimation technique, even if only comparably few camera views are used for reconstruction. Our initial results show that reconstructed volume sizes converge comparably soon to some steady level. With approximately 16-20 camera views the estimated volumes reached a fairly constant size, which was not significantly improved even with more camera views. It seems as if the inclusion of even more camera views is im-

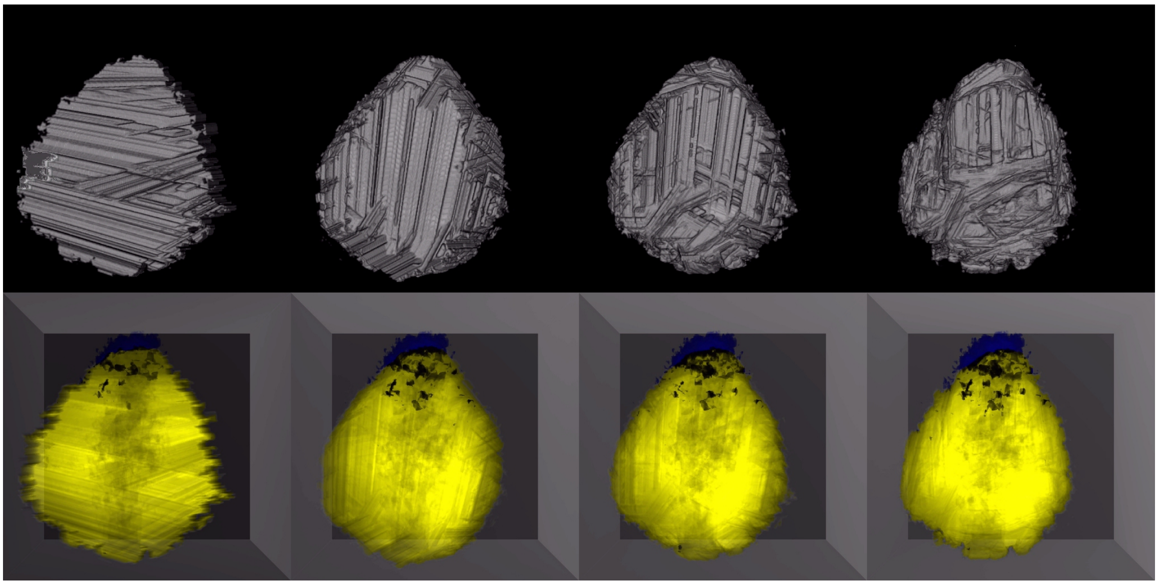

Figure 5: *The upper row shows surface renderings of the the visual hull reconstructions for increasing numbers of silhuette images (from left to right: 3 images, 7 images, 14 images, and 35 images). The bottom row is a visualization of the subtraction volumes for increasing numbers of silhuette images. In all cases, reconstructions were applied for a volume segmented with a threshold of 0,01.*

portant to reveal visual detail; for quantification of the volume, however, these details are not significant. For the type of data volumes tested in our experiment, estimated volume sizes reached a steady value of about 150% of the original volume size for many of the datasets we tested. This value might therefore be instrumental for calibration of the volume estimation procedure in this type of application.
Another interesting finding is that estimation of volume sizes follows the same trend regardless of the chosen segmentation threshold. However, absolute levels do vary significantly depending on the segmentation results. While for most segmentation thresholds (between 0,01 and 0,1) volume estimates are confined within narrow ranges, for higher and lower thresholds the volume estimates can differ by approximately 100% throughout. A comparison with Figure 4 may give one explanation for this. Depending on the segmentation criterion (here threshold) over- or undersegmentation can yield structurally very different results, which are partly reflected by hugely varying volume estimates in the reconstruction. For a final practical application this implies that reliable segmentation of the gaseous structures in the original IR image material is as least as important as maintaining a sufficient number of different camera views.

6. Acknowledgments

The authors wish to acknowledge Linus Fredriksson for provision of the images in Figure 2, and Julia Åhlén for the picture in Figure 1. We also wish to acknowledge the DaGIS project at Future PositionX, Gävle, through which this research has partly been funded.

References

[Har97] HARTLEY R.: In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence 19*, 6 (1997), 580–593. 2

[hS11] ÅHLÉN J., SEIPEL S.: Indication of methane gas in irimagery. In *Proceedings of IADIS International Conference Computer Graphics, Visualization, Computer Vision and Image Processing 2011 (CGVCVIP 2011)* (2011), pp. 187–192. 1

[HZ03] HARTLEY R., ZISSERMAN A.: *Multiple View Geometry in Computer Vision, 2nd ed.*, 2nd ed. ed. Cambridge University Press, 2003. 2

[KSK*08] KIM H., SAKAMOTO R., KITHARA I., TORIYAMA T., KOGURE K.: Compensated visual hull with gpu-based optimization. *Advances in Multimedia Information Processing Ű PCM 2008 5353* (2008), 573–582. 2, 3

[KT] KIM T., THÜREY N.: Wavelet turbulence source code. http://www.cs.cornell.edu/ tedkim/WTURB/source.html, visited 10/16/2011. 3

[KTJG08] KIM T., THÜREY N., JAMES D., GROSS M.: Wavelet turbulence for fluid simulation. *ACM Trans. Graph. 27* (August 2008), 50:1–50:6. 3

[Lau94] LAURENTINI A.: The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence 16*, 2 (1994), 150–162. 2

[Lau95] LAURENTINI A.: How far 3d shapes can be understood from 2d silhouettes. *IEEE Transactions on Pattern Analysis and Machine Intelligence 17*, 2 (1995), 188–195. 2

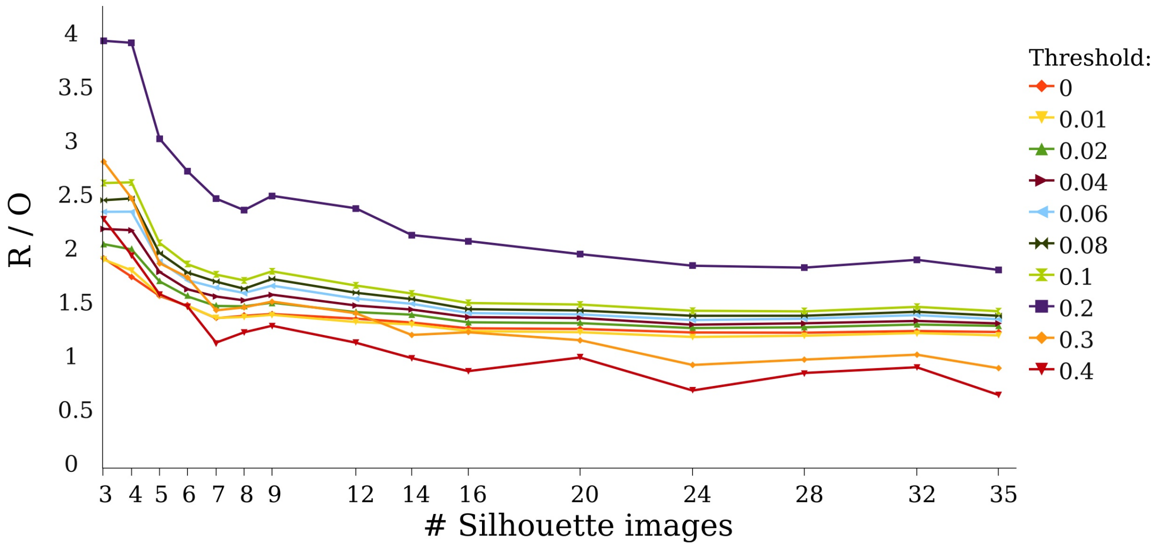

Figure 6: *Diagram showing the ratio of the volume sizes R/O of the reconstructed volume and the original (reference) volume for increasing numbers of silhouette images. Data series are plotted for the same raw data volume but with varying segmentation thresholds.*

[Lau97] LAURENTINI A.: How many 2d silhouettes does it take to reconstruct a 3d object? *Computer Vision and Image Understanding 67*, 2 (1997), 81–87. 2

[LBN08] LADIKOS A., BENHIMANE S., NAVAB N.: Efficient visual hull computation for real-time 3d reconstruction using cuda. In *Proceedings of Computer Vision and Pattern Recognition Workshops - CVPRW '08* (2008), pp. 1–8. 2

[LF95] LUONG Q.-T., FAUGERAS O.: The fundamental matrix: theory, algorithms, and stability analysis. *International Journal of Computer Vision 17* (1995), 43–75. 2

[LMS03] LI M., MAGNOR M., SEIDEL H. P.: Improved hardware-accelerated visual hull rendering. *Vision, Modeling, and Visualization*, Nov. 2003 (2003), 151–158. 2

[LMS04] LI M., MAGNOR M., SEIDEL H.: A hybrid hardware-accelerated algorithm for high quality rendering of visual hulls. In *Proceedings of Graphics Interface 2004* (2004), pp. 41–48. 2

[LW10] LIANG C., WONG K. K.: 3d reconstruction using silhouettes from unordered viewpoints. *Image and Vision Computing 28* (2010), 579–589. 2

[SM10] SAFITRI A., MANNAN M. S.: Methane gas visualization using infrared imaging system and evaluation of temperature dependence of methane gas emissivity. *Ind. Eng. Chem. Res. 49*, 8 (2010), 3926–3935. 1

[Sze93] SZELISKI R.: Rapid octree construction from image sequences. *CVGIP-Image Understanding 58*, 1 (1993), 23–32. 2

[Tal07] TALYAN V.: Quantification of methane emission from municipal solid waste disposal in delhi. *Resources, Conservation and Recycling 50*, 3 (2007), 240–259. 1

[Wue02] WUEBBLES D.: Atmospheric methane and global change. *Earth-Science Reviews 57* (2002), 177–210. 1