

Handling Occlusions in Dense Multi-view Stereo

Sing Bing Kang
Microsoft Research
Microsoft Corporation

Richard Szeliski
Microsoft Research
Microsoft Corporation

Jinxiang Chai
The Robotics Institute
Carnegie Mellon University

Abstract

While stereo matching was originally formulated as the recovery of 3D shape from a pair of images, it is now generally recognized that using more than two images can dramatically improve the quality of the reconstruction. Unfortunately, as more images are added, the prevalence of semi-occluded regions (pixels visible in some but not all images) also increases. In this paper, we propose some novel techniques to deal with this problem. Our first idea is to use a combination of shiftable windows and a dynamically selected subset of the neighboring images to do the matches. Our second idea is to explicitly label occluded pixels within a global energy minimization framework, and to reason about visibility within this framework so that only truly visible pixels are matched. Experimental results show a dramatic improvement using the first idea over conventional multibaseline stereo, especially when used in conjunction with a global energy minimization technique. These results also show that explicit occlusion labeling and visibility reasoning do help, but not significantly, if the spatial and temporal selection is applied first.

1 Introduction

One of the classic research problems in computer vision is that of stereo, i.e., the reconstruction of three-dimensional shape from two or more intensity images. Such reconstruction has many practical applications, including robot navigation, object recognition, and more recently, realistic scene visualization (image-based rendering).

Why is stereo so difficult? Even if we disregard non-rigid effects such as specularities, reflections, and transparencies, we still have to deal with depth discontinuities, which cause occlusions and disocclusions, and with lack of texture in images. Depth discontinuities cause objects to appear and disappear at different viewpoints, thus confounding the matching process at or near object boundaries. The lack of texture, meanwhile, results in ambiguities in depth assignments caused by the presence of multiple good matches.

1.1 Previous work

A substantial amount of work has been done on stereo; a survey on stereo methods can be found in [10]. Stereo can generally be described in terms of the following components: matching criterion, aggregation method, and winner

selection [21, 26]. (Due to space limitations, we only cite the most recent or relevant references.)

The matching criterion is used as a means of measuring the similarity of pixels or regions across different images. A typical error measure is the RGB or intensity difference between images (these differences can be squared, or robust measures can be used). Some methods compute subpixel disparities by computing the analytic minimum of the local error surface or using gradient-based techniques [15]. Birchfield and Tomasi [3] measure pixel dissimilarity by taking the minimum difference between a pixel in one image and the interpolated intensity function in the other image.

The aggregation method refers to the manner in which the error function over the search space is computed or accumulated. The most direct way is to apply search windows of a fixed size over a prescribed disparity space for multiple cameras [19]. Others use adaptive windows [18], shiftable windows [6, 27], or multiple masks [17]. Another set of methods accumulates votes in 3D space, e.g., the space sweep approach [9] and voxel coloring and its variants [22, 25, 14].

Once the initial or aggregated matching costs have been computed, a decision must be made as to the correct disparity assignment for each pixel $d(x, y)$. Local methods do this at each pixel independently, typically by picking the disparity with the minimum aggregated value. Cooperative/competitive algorithms can be used to iteratively decide on the best assignments [29].

Dynamic programming can be used for computing depths associated with edge features or general intensity similarity matches. These approaches can take advantage of one-dimensional ordering constraints along the epipolar line to handle depth discontinuities and unmatched regions [2, 6]. However, these techniques are limited to two frames.

Fully global methods attempt to find a disparity surface $d(x, y)$ that minimizes some smoothness or regularity property in addition to producing good matches. Such approaches include Markov Random Field optimization with simulated annealing [1], nonlinear diffusion of support at different disparity hypotheses [21], graph cut methods [20, 11, 7], and the use of graph cuts in conjunction with planar surface fitting [4].

While occlusions are usually only explicitly handled in the dynamic programming approaches (where semi-




Figure 1: *Overview of our stereo approach.*

occluded regions are labelled explicitly), some techniques have been developed for reasoning about occlusions in a multiple-image setting. These approaches include using multiple matching templates [17], voxel coloring and its variants [22, 25, 14], estimating a depth map per image [24], and graph cuts with the enforcement of unique correspondences [13].

In this paper, we present two complementary approaches to better deal with occlusions in multi-view stereo matching. The first approach (Section 3) uses not only spatially adaptive windows, but also selects a temporal subset of the frames to match at each pixel. The second approach (Section 4) uses a global (MRF) minimization approach based on graph cuts that explicitly models occluded regions with a special label. It also reasons about occlusions by selectively freezing good matching points, and erasing these from the set of pixels that must be matched at depths further back. Both approaches can be combined into a single system, as shown in Figure 1.

2 Problem formulation

In a multi-view stereo problem, we are given a collection of images $\{I_k(x, y), k = 0 \dots K\}$ and associated camera matrices $\{\mathbf{P}_k, k = 0 \dots K\}$. $I_0(x, y)$ is the *reference image* for which we wish to compute a *disparity map* $d(x, y)$ such that pixels in $I_0(x, y)$ project to their corresponding locations in the other images when the correct disparity is selected.

In the classic forward-facing multi-baseline stereo configuration [19], the camera matrices are such that disparity (inverse depth) varies linearly with horizontal pixel motion,

$$\hat{I}_k(x, y, d) = I_k(x + b_k d(x, y), y), \quad (1)$$

where $\hat{I}_k(x, y, d)$ is image I_k warped by the disparity map $d(x, y)$. In a more general (plane sweep) multi-view setting [9, 25], each disparity corresponds to some plane equation in 3-D. Hence, the warping necessary to bring pixels at some disparity d into registration with the reference image can be represented by a homography $H_k(d)$,

$$\hat{I}_k(x, y, d) = H_k(d) \circ I_k(x, y), \quad (2)$$

where the homography can be computed directly from the camera matrices \mathbf{P}_0 and \mathbf{P}_k and the value of d [25]. In this

paper, we assume the latter generalized multi-view configuration, since it allows us to reconstruct depth maps from arbitrary collections of images.

Given the collection images warped at all candidate disparities, we can compute an initial *raw* (unaggregated) matching cost

$$E_{\text{raw}}(x, y, d, k) = \rho \left(I_0(x, y) - \hat{I}_k(x, y, d) \right), \quad (3)$$

where $\rho(\cdot)$ is some (potentially) robust measure of the color or intensity difference between the reference and warped image (see, e.g., [21, 26] for some comparative results with different robust metrics). In this paper, we use a simple squared color difference in our experiments.

The task of stereo reconstruction is then to compute a disparity function $d(x, y)$ such that the raw matching costs are low for all images (or at least the subset where a given pixel is visible), while also producing a “reasonable” (e.g., piecewise smooth) surface. Since the raw matching costs are very noisy, some kind of spatial aggregation or optimization is necessary. The two main approaches used today are local methods, which only look in a neighborhood of a pixel before making a decision, and global optimization methods.

3 Local techniques

The simplest aggregation method is the classic sum of sum of squared distances (SSSD) formula, which simply aggregates the raw matching score over all frames

$$E_{\text{SSSD}}(x, y, d) = \sum_{k \neq 0} \sum_{(u, v) \in \mathcal{W}(x, y)} E_{\text{raw}}(u, v, d, k), \quad (4)$$

where $\mathcal{W}(x, y)$ is an $n \times n$ square window centered at (x, y) . This can readily be seen as equivalent to a convolution with a 3-dimensional box filter.

After the aggregated errors have been computed, local techniques choose the disparity with the minimum SSSD error, which measures the degree of photoconsistency at a hypothesized depth. The best match can also be assigned a local confidence computed using the variance (across disparity) of the SSSD error function within the vicinity of the best match [16].

While window-based techniques work well in textured regions and away from depth discontinuities or occlusions, they run into problems in other cases. Figure 2 shows how a symmetric (centered) window may lead to erroneous matching in such regions. Two ways of dealing with this problem are spatially shiftable windows and temporal selection.

3.1 Spatially shiftable windows

The idea of spatially shiftable windows is an old one that has recently had a resurgence in popularity [17, 6, 27]. The basic idea is to try several windows that include the pixel we are trying to match, not just the window centered at that pixel (Figure 3).¹ This approach can improve the matching

¹When using square windows, finding the best matching shifted window can be computed by passing a min-filter over the original SSD scores.




Figure 2: A simple three-image sequence (the middle image is the reference image), with a frontal gray square (marked F), and a stationary background. Regions B, C, D, and E are partially occluded. A regular SSD algorithm will make mistakes when matching pixels in these regions (e.g. the window centered on the black pixel in region B), and also in windows straddling depth discontinuities (the window centered on the white pixel in region F).




Figure 3: Shiftable windows help mitigate the problems in partially occluded regions and near depth discontinuities. The shifted window centered on the white pixel in region F now matches correctly in all frames. The shifted window centered on the black pixel in region B now matches correctly in the left image. Temporal selection is required to disable matching this window in the right image.

of foreground objects near depth discontinuities (so long as the object is not too thin), and also handle background regions that are being disoccluded rather than occluded (the black pixel in the middle and left image of Figure 3).

To illustrate the effect of shiftable windows, consider the flower garden sequence shown in Figure 6. The effect of using spatially shiftable windows over all 11 frames is shown in Figure 7(b) for a 5×5 window size. As can be seen, there are differences, but they are not dramatic. The errors seen can be attributed to ignoring the effects of occlusions and disocclusions.

3.2 Temporal selection

Rather than summing the match errors over all the frames, a better approach would be to pick only the frames where the pixels are visible. Of course, this is not possible in general without resorting to the kind of visibility reasoning present in volumetric [22, 25, 14] or multiple depth map [24] approaches, and also in the multiple mask approach of [17]. However, often a semi-occluded region in the reference image will only be occluded in the predecessor or successor frames, i.e., for a camera moving along a continuous path, objects that are occluded along the path in one




Figure 4: The spatio-temporal diagram (epipolar plane image) corresponding to the previous figure. The three images (middle, left, right) are slices through this EPI volume. The spatially and temporally shifted window around the black pixel is indicated by the rectangle, showing the the right image is not being used in matching.

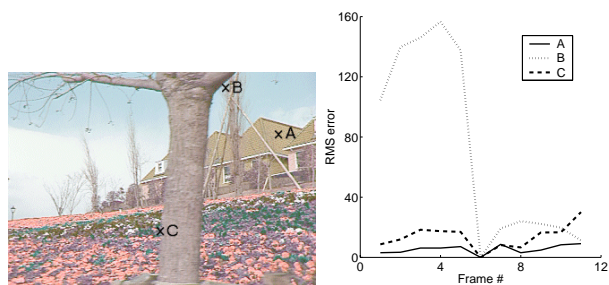


Figure 5: Error profiles for three points in reference image. A: point seen all the time, B: point occluded about half the time, C: occluding point. Left: Reference image, Right: Error graph at respective optimal depths with respect to the frame number (frame #6 is the reference).

direction tend to be seen along the reverse direction. (A similar idea has recently been applied to optic flow computation [23].) Figure 3 shows this behavior. The black pixel in region B and its surrounding (shifted) square region can be matched in the left image but not the right image. Figure 4 show this same phenomenon in a spatio-temporal slice (epipolar plane image). It can readily be seen that temporal selection is equivalent to shifting the window in time as well as in space.

Temporal selection as a means of handling occlusions and disocclusions can be illustrated by considering selected error profiles depicted in Figure 5. Points such as A, which can be observed at all viewpoints, work without shiftable windows and temporal selection. Points such as C, which is an occluding point, work better with shiftable windows but do not require temporal selection. Points such as B, however, which is occluded in a fraction of the viewpoints, work best with both shiftable windows and temporal selection.

Rather than just picking the preceding or succeeding frames (one-sided matching), a more general variant would be to pick the best 50% of all images available. In this case, we compute the local SSD error for each frame separately, and then sum up the lowest values. This kind of approach can better deal with objects that are intermittently visible,



Figure 6: Middle (reference) images from the three test sequences: 11-image flower garden sequence; 5-image Symposium sequence, courtesy of Dayton Taylor; 5-image sequence, courtesy of the University of Tsukuba.




Figure 7: Local (5×5 window-based) results for flower garden sequence: (a) non-spatially perturbed (centered) window, (b) spatially perturbed window, (c) using best 5 of 10 neighboring frames, (d) using better half sequence.




Figure 8: Local (5×5 window-based) results for the other two sequences: (a) and (c) non-spatially perturbed (centered) window; (b) and (d) using better half sequence.



Figure 9: Final results of incremental window size algorithm run on all three sequences. See the supplementary materials for intermediate (partial) results.

i.e., a ‘‘picket fence’’ phenomenon.

We have experimented with both variants, and found that they have comparable performance. Figure 7(c–d) shows the results on the flower garden sequence, while Figure 8 shows results on the other two sequences. (See our supplementary materials for higher quality versions of these images.) As you can see, using temporal selection yields a dramatic improvement in results, especially near depth discontinuities (occlusion boundaries) such as the edges of the tree.

3.3 Adaptive window size

If a purely window-based technique is to be used, a reasonable way to handle untextured areas would be to use variable window sizes. We have implemented a novel variable window size approach that works as follows.

Instead of simply selecting the best depth at each pixel for a fixed (initial) window size, only a fraction (currently 15%) of the depths computed are committed based on their reliability. The reliability (or local confidence) assigned to each depth is the local variance of the error function around that depth. The higher the variance, the higher the perceived reliability. At every new iteration, the process is repeated with a larger window size over the uncommitted pixels. After 12 iterations, any undecided pixels are forced to commit. By using the error variance as a measure of depth reliability, we ensure that larger regions of textureless surfaces get to be handled by larger windows.

Our approach bears some resemblance to the recent proposal by Zhang and Shan [28], which starts with point matches and grows matching regions around these points. In our approach, however, there is no requirement to grow existing regions; instead, the most confident pixels are simply selected at each iteration. Our idea of variable window sizes is also similar to [18]. However, we adopt a high-est confidence first approach [8] to choosing a window size rather than testing at each pixel location all the windows sizes in order to select an optimal size.

Results of using the incremental selection approach can be seen in Figure 9. While it generally interpolates across textureless regions reasonably well, determining the correct fraction of pixels to commit at each iteration requires a heuristic decision (i.e., it may be scene dependent).

4 Global techniques

The second general approach to dealing with ambiguity in stereo correspondence is to optimize a global energy function. Typically, such a function consists of two terms,

$$E_{\text{global}}(d(x, y)) = E_{\text{data}} + E_{\text{smooth}}. \quad (5)$$

The value of the disparity field $d(x, y)$ that minimizes this global energy is chosen as the desired solution.²

²Because of the tight connection between this kind of global energy and the log-likelihood of a Bayesian model using Markov Random Fields, these methods are also often called Bayesian or MRF methods [2, 7].

The *data* term E_{data} is just a summation of the local (aggregated or unaggregated) matching costs, e.g.,

$$E_{\text{data}} = \sum_{(x,y)} E_{SSSD}(x, y, d(x, y)). \quad (6)$$

Because a smoothness term is used, spatial aggregation is usually not used, i.e., the window $W(x, y)$ in the SSSD term is a single pixel (but see, e.g., [6] for a global method that starts with a window-based cost measure).

The smoothness term E_{smooth} measures the piecewise-smoothness in the disparity field,

$$E_{\text{smooth}} = \sum_{(x,y)} s_{x,y}^h \phi(d(x, y) - d(x + 1, y)) \quad (7) \\ + s_{x,y}^v \phi(d(x, y) - d(x, y + 1)).$$

The smoothness potential $\phi(\cdot)$ can be a simple quadratic, a delta function, a truncated quadratic, or some other robust function of the disparity differences [5, 7]. The smoothness strengths $s_{x,y}^h$ and $s_{x,y}^v$ can be spatially varying. The MRF formulation used by [7] makes $s_{x,y}^h$ and $s_{x,y}^v$ monotonic functions of the local intensity gradient, which greatly helps in forcing disparity discontinuities to be coincident with intensity discontinuities.

If the vertical smoothness term is ignored, the global minimization can be decomposed into an independent set of 1-D optimizations, for which efficient dynamic programming algorithms exist [2, 6]. Many different algorithms have also been developed for minimizing the full 2-D global energy function, e.g., [1, 21, 7].

In this section, we propose two extensions to the graph cut formulation introduced by [7] in order to better handle the partial occlusions that occur in multi-view stereo: explicit occluded pixel labeling, and visibility computation.

4.1 Explicit occluded pixel labeling

When using a global optimization framework, pixels that do not have good matches in other images will still be assigned some disparity. Such pixels are often associated with a high local matching cost, and can be detected in a post-processing phase. However, occluded pixels also tend to occur in contiguous regions, so it makes sense to include this information within the smoothness function (i.e., within the MRF formulation).

Our solution to this problem is to include an additional label that indicates pixels that are either outliers or potentially occluded. A fixed penalty is associated with adopting this label, as opposed to the local matching cost associated with some other disparity label. (In our current implementation, this penalty is set at 18 intensity levels.) The penalty should be set to be somewhat higher than the largest value observed for correctly matching pixels. The smoothness term for this label is a delta function, i.e., a fixed penalty is paid for every non-occluded pixel that borders an occluded one.



Figure 10: *Effect of using the undefined label for 11-frame flower garden sequence (64 depth levels, no visibility terms, using best frames): (a) Reference image is 1st image, (b) Reference image is 6th image, (c) Reference image is 11th image. The undefined label is black, while the intensities for the rest are bumped up for visual clarity.*

Examples of using such a label can be seen in Figure 10. Unfortunately, this approach sometimes fails to correctly label pixels in occluded textureless regions (since these pixels may still match correctly at the frontal depth). In addition, the optimal occluded label penalty setting depends on the amount of contrast in a given scene.

4.2 Visibility reasoning

An idea that has proven to be very effective in dealing with occlusions in volumetric [22, 25, 14] or multiple depth map [24] approaches is that of visibility reasoning. Once a pixel has been matched at one disparity level, it is possible to “erase” that pixel from consideration when considering possible matches at disparities further back. This is the most principled way to reason about visibility and partial occlusions in multi-view stereo. However, since the algorithms cited above make independent decisions between pixels or frames, their results may not be optimal.

To incorporate visibility into the global optimization framework, we compute a visibility function similar to the one presented in [25]. The visibility function $v(x, y, d, k)$ can be computed as a function of the disparity assignments at layers closer than d . Let $o(x, y, d') = \delta(d', d(x, y))$ be the *opacity* (or indicator) function, i.e., a binary image of those pixels assigned to level d' . The *shadow* $s(x, y, d', d, k)$ that this opacity casts relative to camera k onto another level d can be derived from the homographies that map between disparities d' and d

$$s(x, y, d', d, k) = (H_k(d)H_k^{-1}(d')) \circ o(x, y, d') \quad (8)$$

(we can, for instance, use bilinear resampling to get “soft” shadows, indicative of partial visibility). The visibility of a pixel (x, y) at disparity d relative to camera k can be computed as

$$v(x, y, d, k) = \prod_{d' < d} (1 - s(x, y, d', d, k)). \quad (9)$$

Finally, the raw matching cost (3) can then be replaced by

$$E_{\text{vis}}(x, y, d, k) = v(x, y, d, k) \rho \left(I_0(x, y) - \hat{I}_k(x, y, d) \right). \quad (10)$$

The above visibility-modulated matching score thus provides a principled way to compute the goodness of a particular disparity map $d(x, y)$ while explicitly taking into account occlusions and partial visibility. For any given labeling $d(x, y)$, we can compute the opacities, shadows, and visibilities, and then sum up the visibility-modulated matching scores (10) to obtain the final global energy (5). Unfortunately, it is not obvious how to minimize such a complicated energy function.

One possibility would be to start with all pixels visible, and to then run the usual graph-cut algorithm. From the initial $d(x, y)$ solution, we could recompute visibilities, and then re-optimize the modified energy function. Unfortunately, this process may not converge, since the energy function is being modified from iteration to iteration, and the visibilities assumed for one iteration may be undone by a re-assignment of labels in that iteration.

The alternative we have come up with (inspired by Chou’s Highest Confidence First algorithm [8]) is to progressively commit the best-matching depths (i.e., freeze their labels) and apply graph cut on the remaining pixels. This approach is related to the voxel coloring work [22], where voxels are tagged from front to back. However, in our approach, the best 15% of the pixels (based on the current visibility-modulated matching score (10)) whose depths have been computed by the graph cut are frozen. The visibility function and matching costs are then recomputed, which may affect costs at more distal voxels. Within each iteration, graph-cut labeling effectively takes into account neighboring pixels’ preferences and tries to make the disparity function piecewise-smooth, whereas the voxel coloring approach only uses per-pixel photo-consistency. After 12 iterations, the remaining uncommitted pixels are frozen at their best value.

Figure 11 shows the results of adding visibility reasoning to the graph cut algorithm when starting with all frames as the data cost (no temporal selection). As can be seen by comparing the first two columns, the improvement is significant. (See our supplementary materials for higher quality versions of these images, as well as additional ex-




Figure 11: *Effect of applying incremental visibility-based graph cuts: (a,d,g) results using all frames; (b,e,h) results using all frames and visibility; (c,f,i) results using best half sequence and no visibility. (There is no significant improvement in adding visibility reasoning when starting with the best half sequence (or best frames) data term.)*

perimental results.) Surprisingly, the addition of visibility computation to the graph cut did not produce significant improvements to our algorithm when the original matching costs were computed using a shiftable window and temporal selection. (The third column of Figure 11 shows the results of regular graph cuts combined with temporal selection using the best half sequence. The results are almost identical when visibility reasoning is added.) This suggests that the shiftable window, and especially temporal selection, handled the occlusion problem well.

5 Discussion

In this section, we summarize our experimental results applied to the image sequences shown in Figure 6. Due to space limitations, we only present a brief sampling from the larger set of experiments we have run. Please see the

accompanying supplementary material Web pages and our longer technical report [12] for a more complete set of results.

Figures 7 and 8 show some results using spatially shiftable windows optionally combined with temporal selection, followed by a simple winner-take-all. The effects of temporal selection are more dramatic than spatial shifting, and yield their biggest improvements near depth discontinuities. Large textureless regions are still not recovered well.

Using incremental window sizes (Figure 9) helps fill in more reasonable disparity values in textureless regions, but still does not do that well in some areas such as the sky in the Symposium sequence and the upper right corner of the U. Tsukuba sequence. Global optimization techniques generally outperform this idea.

Adding an extra occluded pixel label to global optimization helps find regions that are visible in only one image, such as the start and end frames for a multi-view sequence (Figure 10). This should help the most when only a small number of frames is available (e.g., in classical two-frame matching).

Finally, visibility reasoning is a good way to obtain better results near depth discontinuities when the complete set of images is used as input to the data cost term (Figure 11). Reasoning about which pixels are occluded allows us to iteratively re-compute a better data term. However, to our surprise, this idea does not seem to help much if temporal selection has already been applied to heuristically reject possibly occluded pixels, at least not on the data sets we have currently tried.

6 Conclusions

In this paper, we have presented several new ideas for improving the results of multi-view stereo correspondence algorithms. Our particular emphasis has been on better dealing with pixels and regions that are occluded in some images but not in others. Some of our ideas, such as temporal selection, can be applied at the initial matching cost stage. Other ideas, such as outlier/invisible pixel labeling and visibility reasoning, can be used to enhance the performance of global optimization techniques such as graph cut algorithms.

Of all the ideas we have developed, using temporal selection (using only a subset of all frames for computing the matching cost), followed by a regular graph cut global optimization, seems to yield the best results for the least computational effort. It will be interesting to see how these ideas generalize to other versions of multi-view stereo reconstruction, such as the extraction of multiple layers and volumetric reconstruction techniques.

References

- [1] S. T. Barnard. Stochastic stereo matching over scale. *Intl. J. Comp. Vision*, 3(1):17–32, 1989.
- [2] P. N. Belhumeur. A Bayesian-approach to binocular stereopsis. *Intl. J. Comp. Vision*, 19(3):237–260, August 1996.
- [3] S. Birchfield and C. Tomasi. A pixel dissimilarity measure that is insensitive to image sampling. *IEEE Trans. PAMI*, 20(4):401–406, April 1998.
- [4] S. Birchfield and C. Tomasi. Multiway cut for stereo and motion with slanted surfaces. In *ICCV'99*, pp. 489–495, Kerkyra, Greece, September 1999.
- [5] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *Intl. J. Comp. Vision*, 19(1):57–91, 1996.
- [6] A. F. Bobick and S. S. Intille. Large occlusion stereo. *Intl. J. Comp. Vision*, 33(3):181–200, September 1999.
- [7] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. In *ICCV'99*, pp. 377–384, Kerkyra, Greece, September 1999.
- [8] P. B. Chou and C. M. Brown. The theory and practice of Bayesian image labeling. *Intl. J. Comp. Vision*, 4(3):185–210, June 1990.
- [9] R. T. Collins. A space-sweep approach to true multi-image matching. In *CVPR'96*, pp. 358–363, San Francisco, 1996.
- [10] U. R. Dhond and J. K. Aggarwal. Structure from stereo—a review. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(6):1489–1510, November/December 1989.
- [11] H. Ishikawa and D. Geiger. Occlusions, discontinuities, and epipolar lines in stereo. In *ECCV'98*, pp. 232–248, Freiburg, Germany, June 1998.
- [12] S. B. Kang, R. Szeliski, and J. Chai. Handling occlusions in dense multi-view stereo. Technical Report MSR-TR-2001-80, Microsoft Research, September 2001.
- [13] V. Kolmogorov and R. Zabih. Computing visual correspondence with occlusions via graph cuts. In *ICCV 2001*, vol. II, pp. 508–515, Vancouver, July 2001.
- [14] K. N. Kutulakos and S. M. Seitz. A theory of shape by space carving. In *ICCV'99*, pp. 307–314, Kerkyra, Greece, September 1999.
- [15] B. D. Lucas and T. Kanade. An iterative image registration technique with an application in stereo vision. In *IJCAI-81*, pp. 674–679, Vancouver, 1981.
- [16] L. H. Matthies, R. Szeliski, and T. Kanade. Kalman filter-based algorithms for estimating depth from image sequences. *Intl. J. Comp. Vision*, 3:209–236, 1989.
- [17] Y. Nakamura, T. Matsuura, K. Satoh, and Y. Ohta. Occlusion detectable stereo - occlusion patterns in camera matrix. In *CVPR'96*, pp. 371–378, San Francisco, June 1996.
- [18] M. Okutomi and T. Kanade. A locally adaptive window for signal matching. *Intl. J. Comp. Vision*, 7(2):143–162, 1992.
- [19] M. Okutomi and T. Kanade. A multiple baseline stereo. *IEEE Trans. PAMI*, 15(4):353–363, April 1993.
- [20] S. Roy and I. J. Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. In *ICCV'98*, pp. 492–499, Bombay, January 1998.
- [21] D. Scharstein and R. Szeliski. Stereo matching with non-linear diffusion. *Intl. J. Comp. Vision*, 28(2):155–174, July 1998.
- [22] S. M. Seitz and C. M. Dyer. Photorealistic scene reconstruction by voxel coloring. In *CVPR'97*, pp. 1067–1073, San Juan, Puerto Rico, June 1997.
- [23] S. Sun, D. Haynor, and Y. Kim. Motion estimation based on optical flow with adaptive gradients. In *ICIP-2000*, vol. I, pp. 852–855, Vancouver, September 2000.
- [24] R. Szeliski. A multi-view approach to motion and stereo. In *CVPR'99*, vol. 1, pp. 157–163, Fort Collins, June 1999.
- [25] R. Szeliski and P. Golland. Stereo matching with transparency and matting. *Intl. J. Comp. Vision*, 32(1):45–61, August 1999.
- [26] R. Szeliski and R. Zabih. An experimental comparison of stereo algorithms. In *Intl. Work. Vision Algorithms*, pp. 1–19, Kerkyra, Greece, September 1999. Springer.
- [27] H. Tao, H.S. Sawhney, and R. Kumar. A global matching framework for stereo computation. In *ICCV 2001*, vol. I, pp. 532–539, Vancouver, July 2001.
- [28] Z. Zhang and Y. Shan. A progressive scheme for stereo matching. *SMILE 2000*, pp. 68–85, Dublin, July 2000.
- [29] C. L. Zitnick and T. Kanade. A cooperative algorithm for stereo matching and occlusion detection. *IEEE Trans. PAMI*, 22(7):675–684, July 2000.