

Segment-based Stereo Matching Using Graph Cuts

Li Hong

George Chen

Advanced System Technology San Diego Lab, STMicroelectronics, Inc.

li.hong@st.com

george-qian.chen@st.com

Abstract

In this paper, we present a new segment-based stereo matching algorithm using graph cuts. In our approach, the reference image is divided into non-overlapping homogeneous segments and the scene structure is represented as a set of planes in the disparity space. The stereo matching problem is formulated as an energy minimization problem in the segment domain instead of the traditional pixel domain. Graph cuts technique is used to fast approximate the optimal solution, which assigns the corresponding disparity plane to each segment. Experiments demonstrate that the performance of our algorithm is comparable to the state-of-the-art stereo algorithms on various data sets. Furthermore, strong performance is achieved in the conventionally difficult areas such as: textureless regions, disparity discontinuous boundaries and occluded portions.

1. Introduction

Stereo matching is one of the most active research areas in computer vision and it serves as an important step in many applications (e.g., view synthesis, image based rendering, etc). The goal of stereo matching is to determine the disparity map between an image pair taken from the same scene. Disparity describes the difference in location of the corresponding pixels and it is often considered as a synonym for inverse depth. Due to the ill-posed nature of the stereo matching problem, the recovery of accurate disparity still remains challenging, especially in textureless regions, disparity discontinuous boundaries and occluded areas.

An excellent review of stereo work can be found in [8]. In general, stereo algorithms can be categorized into two major classes. The first class is local (window-based) algorithms, where the disparity at a given pixel depends only on intensity values within a finite neighboring window. Local methods can easily capture accurate disparity in highly textured regions, however they often tend to produce noisy disparities in textureless regions, blur the disparity discontinu-

ous boundaries and fail at occluded areas. The second class is global algorithms, which make explicit smoothness assumptions of the disparity map and solve it through various minimization techniques. Recently, global methods such as graph cuts [2, 11, 12, 13, 14] have attracted much attention due to their excellent experimental results.

In this paper, we present a new segment-based stereo method that achieves strong performance in the commonly challenging image regions. The proposed algorithm is a global stereo algorithm inspired upon the recent work of Tao *et al.* [10] and the graph cuts-based algorithms of Boykov *et al.* [2] and Kolmogorov and Zabih [11].

Tao *et al.* [10] describes a color segment-based stereo framework, which is based on the assumption that there are no large disparity discontinuities inside homogeneous color segments. The main idea is that if a disparity hypothesis is correct, warping the reference image to the other view according to its disparity will render an image that matches the real view. Therefore, the stereo matching problem is solved through minimizing this global image similarity energy. Color segment representation is used to reduce the high solution space and enforce disparity smoothness in homogeneous color regions. A greedy local search mechanism by neighboring disparity hypothesis is proposed to further reduce the expensive warping cost and infer reasonable disparities for unmatched regions.

Boykov *et al.* [2] and Kolmogorov and Zabih [11] present efficient graph cuts-based stereo algorithms to find a smooth disparity map that is consistent with the observed data. In their approaches, stereo matching problem is formulated as an energy minimization problem, which mainly¹ includes: (i) a smoothness energy (E_s) that measures the disparity smoothness between neighboring pixel pairs; (ii) a data energy (E_d) that measures the disagreement between corresponding pixels based on the assumed disparities. A weighted graph is then constructed in which graph nodes represent image pixels,

¹ In [11], an additional occlusion energy E_{occ} is included that explicitly addresses occlusions.

graph label set (or terminals) relate to all possible disparities (or all discrete values in the disparity range interval) and graph edge weights correspond to the defined energy terms. Graph cuts technique is then used to approximate the optimal solution, which assigns the corresponding disparity (graph label) to each pixel (graph node).

We take advantages of both the color segment-based stereo framework and the graph cuts technique. In our approach, the reference image is divided into non-overlapping homogeneous color segments and the scene structure is approximated as a set of planes in disparity space (not necessarily the fronto-parallel planes). The stereo matching problem becomes assigning the corresponding disparity plane to each segment, which can be easily formalized as an energy minimization problem in the segment domain. Specifically, the energy function contains two parts: (i) a smoothness energy (E_{smooth}) that measures the disparity smoothness between neighboring segment pairs; (ii) a data energy (E_{data}) that measures the disagreement of segments and their matching regions based on the assumed disparity planes.

We can define E_{smooth} and E_{data} to enable them as simplifications to energies E_s and E_d . Under the assumption that each homogeneous color region contains smooth continuous disparities, energy E_s only needs to measure the disparity smoothness between neighboring pixels that lie in different segments. Therefore, the disparity smoothness between neighboring segments only needs to be associated with the common segment boundaries, which leads the derivation of E_{smooth} from E_s . It is trivial to define the disagreement of each segment and its corresponding region as the accumulation of the disagreement from all pixels inside the segment. Therefore data energy E_{data} measures the disagreement of all corresponding pixels of the whole image, which is essentially data energy E_d .

We apply graph cuts technique in a similar manner as in [2] *et al.* to approximate the optimal solution of our energy function. Note that the graph nodes here represent segments instead of pixels. In general, the number of segments is much less than pixels, which leads to a simple graph structure and fast computation. In addition, note that the graph label set here only contains the approximate disparity planes of the scene instead of all possible discrete values in the disparity range. Birchfield and Tomasi [1] pointed out that by searching over discrete disparities, what is preserved is actually piecewise-constancy rather than piecewise-continuity. In our approach, piecewise-continuity is enforced inside each segment through the plane representation. Furthermore, disparity smoothness is also enforced between segments through the energy term E_{smooth} . This is especially important for textured regions with smooth disparities, as very likely many small segments will be present

in these regions. By enforcing smooth constraints between segments, the disparity smoothness of these regions is preserved.

Our algorithm also shares certain similarities with the algorithm of Birchfield and Tomasi [1]. Both algorithms use the smooth disparity function to represent disparity continuous regions, which enforces piecewise-continuity. However, in [1] the image is segmented into regions with smooth disparities through graph cuts, while we use color segmentation. The accuracy of their regions is influenced by the estimated disparity functions, while color segmentation only relies on the original image intensity. In general, good color segmentation can capture much sharper intensity boundaries without being confined to certain contour shapes, while in [1], their region boundaries often ensure biases against curved object contours. In addition, occlusions are incorporated in our approach (details see Section 3), but they are not addressed in [1]. Furthermore, it is worth noticing that in [1] graph cuts is still built in the pixel domain, while ours is in the segment domain, which leads a simpler computation in the stage of graph cuts.

The rest of the paper is organized as follows: First we briefly describe the color segmentation process in Section 2. Then we present in details the proposed algorithm, mainly focus on how to estimate disparity planes inside the scene (Section 3), and how to apply graph cuts to assign the corresponding disparity plane to each segment (Section 4). We provide various experimental results in Section 5 to demonstrate the algorithm's strong performance for traditional challenging image regions. Finally, we conclude and discuss future investigation areas in Section 6.

2. Color segmentation

Our approach is built upon the assumption that large disparity discontinuities only occur on the boundaries of homogeneous color segments. Therefore any color segmentation algorithm that decomposes an image into homogeneous color regions will work for us. In this paper, we strictly enforce disparity continuity inside each color segment, therefore under-segmentation is not preferred. However, over-segmentation is largely tolerated due to the smoothness disparity constraints between segments defined in our E_{smooth} energy term (details given in Section 4). In our current implementation, mean-shift color segmentation algorithm [3] is used.

3. Disparity plane estimation

We represent the scene structure as a collection of disparity continuous surfaces, and approximate each surface by a plane (could be extended to more sophisticated models as well). Note that we may lose some disparity accuracy due

to the limitation of the plane approximation, but such a simple model is sufficient for many applications such as view synthesis, image based rendering, *etc.* We estimate all the possible disparity planes inside the scene through the following processes. First we obtain the initial pixel disparity through local matching analysis. We then compute the initial plane parameters from each color segment (while skip very small segments). Finally, we refine the plane parameters through fitting them to grouped neighboring segments.

3.1. Local matching in pixel domain

In a standard rectified stereo setup, the correspondence between a pixel (x, y) in the reference image I and a pixel (x', y') in the matching image J is given by: $x' = x + d(x, y)$, $y' = y$, where the disparity $d(x, y)$ can take any discrete value from the displacement interval $[d_{min}, d_{max}]$. Let $c(x, y, d)$ denote the matching cost (or dissimilarity) for the pixel (x, y) in image I with disparity d . We compute c as the average pixel intensity differences in a 3×3 surrounding window, i.e.,

$$c(x, y, d) = \frac{1}{9} \sum_{i=-1}^1 \sum_{j=-1}^1 |I(x+i, y+j) - J(x+i+d, y+j)|.$$

Among all possible disparities for the pixel (x, y) , the one that gives the minimum matching cost is selected as the initial estimated disparity $\hat{d}(x, y)$.

3.2. Initial plane fitting from single segment

Plane fitting from the initial disparities in a segment is discussed in details by Tao *et al.* [10]. We briefly review their approach here. A plane is used to model the continuous disparity of each segment, i.e., $d = c_1x + c_2y + c_3$, where c_1, c_2, c_3 are the plane parameters and d is the corresponding disparity of the image pixel (x, y) . (c_1, c_2, c_3) is the least square solution of a linear system

$$A[c_1, c_2, c_3]^T = B \quad (1)$$

where the i -th row of A is $[x_i \ y_i \ 1]$, the i -th element in B is $d(x_i, y_i)$. After the initial plane fitting, an iterative process is adopted to update the plane. In each iteration, the pixel disparity is changed within a given range of the fitted plane and the plane parameters are updated based on the modified disparities accordingly.

In our work, we introduce several schemes to enhance the robustness of the plane fitting algorithm.

- The solution of linear system (1) can be biased by pixels with unreliable initial disparities (we call them outliers), therefore we only include pixels with reliable initial disparity estimation to form system (1). We detect outliers through a simple crosscheck method. Let

pixel (x', y') (in the matching image) be the correspondence of pixel (x, y) (in the reference image) based on the initial disparity $\hat{d}(x, y)$, and let $\hat{d}'(x', y')$ be the initial disparity of pixel (x', y') in the matching image. If $\hat{d}(x, y) \neq \hat{d}'(x', y')$, we consider pixel (x, y) as an outlier.

- Weighted least square scheme is adopted in the iteration process. The plane parameters are computed initially through (1), while in the following iterations, the weight of each equation in (1) is adjusted based on the closeness of the pixel initial disparity to the estimated plane, e.g., $w(\beta_i) = e^{-2*\beta_i}$, where $\beta_i = |c_1x_i + c_2y_i + c_3 - \hat{d}(x_i, y_i)|$. The plane parameters are updated accordingly based on the modified weights.
- Very small segments are skipped as they lack sufficient data to provide reliable plane estimations.

A plane computed from a segment is added into the plane set only if there is no similar plane already inside the set. Generally, our plane set contains much less number of planes than the number of segments in the image.

3.3. Refined plane fitting from grouped segments

Our purpose is not to find the best plane for each segment but rather extract all possible planes for the image. Therefore, it is crucial to extract a set of disparity planes that accurately represent the scene structure. This is achieved through refining plane fitting on grouped segments. The reason behind our approach is that small fragmented regions should be grouped together to provide more reliable pixels to form the linear system (1). We proceed with the following steps:

1. Measure segment matching cost for each plane in the disparity plane set.
2. Assign each segment the plane ID that gives the minimum matching cost.
3. Group neighboring segments with the same plane ID.
4. Apply the plane fitting process mentioned in subsection 3.2 to each grouped segment.

Steps 2 - 4 are trivial. We focus on how to measure segment matching cost for each disparity plane. It is natural to compute it as the sum of the matching cost from each single pixel inside the segment, i.e.,

$$C(S, P) = \sum_{(x,y) \in S} c(x, y, d), \quad (2)$$

where S is a segment, P is a disparity plane, and $d = c_1^P x + c_2^P y + c_3^P$ (c_1^P, c_2^P and c_3^P are the parameters of P).

However, there are several problems associated with this approach. First, occluded pixels would easily bias this segment matching cost, especially in situations such as occlusion counts nontrivial percentage of a segment, or occlusion

exists in a smooth textureless segment. Second, although accumulating pixel matching cost in a segment helps to resolve the matching ambiguities, the problem still remains, especially in textureless regions.

We propose two remedies. First we exclude all possible occluded pixels in computing the segment matching cost. Second we augment the sum of pixel matching cost by the percentage of non-supporting pixels to the disparity plane.

In more details, we consider only textured outliers as possibly occluded pixels. Recall that outliers are pixels with unreliable initial disparity estimation and are detected through the crosscheck method. It is easy to see that occluded pixels are outliers under our outlier definition. We do not treat outliers in smooth textureless regions as occluded pixels for two reasons: (i) smooth textureless pixels tend to fail the crosscheck as a result of matching ambiguities rather than occlusion; (ii) occlusions normally occur around object boundaries, which are highly textured.

Among all non-occluded pixels, we consider a pixel support a disparity plane if its initial estimated disparity is within a small vicinity of the plane; otherwise, we consider it as a non-supporting pixel.

Let n be the number of non-occluded pixels in a segment S , and let s be the number of supporting pixels to a disparity plane P in segment S . We define the segment matching cost as follows:

$$C(S, P) = \sum_{(x,y) \in S-O} c(x, y, d) e^{1-\frac{s}{n}}, \quad (3)$$

where O represents the occluded portion in S .

Our matching cost function is composed of two parts: (i) an accumulated sum of non-occluded pixel matching cost, and (ii) an exponential function that increases along the non-supporting pixel percentage. Note that the defined matching cost function will favor a disparity plane with low accumulated sum and large supporting size. This is especially useful for textureless segments. As in those segments, the first part will be comparable for various planes, thus favoring the disparity plane with large supporting size helps to reduce the ambiguities. For textured segment, part (i) will vary significantly between the correct disparity plane and others, the role of part (ii) will enlarge even more the cost difference, as the correct disparity plane will normally has larger supports.

4. Disparity plane labeling by graph cuts

We have presented two major steps of the proposed algorithm, namely, color segmentation and the disparity plane estimation. In this section, we describe in details the formalization of the stereo matching as an energy minimization problem in the segment domain and its solution, i.e., label-

ing each segment with its corresponding disparity plane by graph cuts.

Let \mathcal{R} be the color segments of the reference image, \mathcal{D} be the estimated disparity plane set. The goal is to find a labeling f that assigns each segment $S \in \mathcal{R}$ a corresponding plane $f(S) \in \mathcal{D}$, where f is both piecewise smooth and consistent with the observed data. We formulate it as an energy minimization problem in the segment domain, i.e., we want to obtain a labeling f^* that minimizes:

$$E(f) = E_{data}(f) + E_{smooth}(f) \quad (4)$$

where,

- E_{data} is a data-dependent energy term containing the cost of assigning plane labels to the segments:

$$E_{data}(f) = \sum_{S \in \mathcal{R}} C(S, f(S)). \quad (5)$$

where $C(S, f(S))$ is defined in Eq. (3).

- $E_{smooth}(f)$ enforces smoothness by penalizing discontinuities, i.e., imposing penalties if different disparity planes are assigned to neighboring segments.

$$E_{smooth}(f) = \sum_{(S,S')} u_{S,S'} \cdot \delta(f(S) \neq f(S')). \quad (6)$$

where S and S' are neighboring segments, $u_{S,S'}$ is proportional to the common border length between segment S and S' , and $\delta(f(S) \neq f(S'))$ has value 1 if $f(S) \neq f(S')$, otherwise 0.

In general, minimizing the energy function given in Eq. (4) is very difficult [14]. However, several new algorithms based on graph cuts have been developed recently to efficiently solve certain energy minimization problems [2, 11, 12]. It is trivial to see our discontinuous function in E_{smooth} is a Potts model, i.e., discontinuities between any pair of labels are penalized irrelevant to the amount of label difference. Therefore we can use the methods of [11] to construct a graph \mathcal{G} and apply graph cuts to approximate the global minimum of our energy function. Notice that in our work, the graph nodes represent segments instead of pixels, and the label set is composed of all the estimated disparity planes instead of all possible discrete values in the disparity range. We proceed the following steps:

1. Start with an initial labeling f .
2. Set success := 0.
3. Select in a random (or fixed) order a disparity plane $P \in \mathcal{D}$.
 - 3.1 Find $\hat{f} = \arg \min E(f')$ among f' within one α -expansion of f .
 - 3.2 If $E(\hat{f}) \leq E(f)$, set $f := \hat{f}$ and success := 1
4. If success == 1 goto 2.
5. Return f .

The solution converges usually within 2-3 iterations. In addition, it is extremely insensitive to the initial labeling.

Table 1: Results of Different Stereo Algorithms.

Algorithms	Tsukuba			Sawtooth			Venus			Map	
	B_{all}	B_{untex}	B_{disc}	B_{Pall}	B_{untex}	B_{disc}	B_{all}	B_{untex}	B_{disc}	B_{all}	B_{disc}
Proposed	1.23	0.29	6.94	<u>0.30</u>	<u>0.00</u>	<u>3.24</u>	<u>0.08</u>	<u>0.01</u>	<u>1.39</u>	1.49	15.46
Layered [5]	1.58	1.06	8.82	0.34	<u>0.00</u>	3.35	1.52	2.96	2.62	0.37	5.24
Belief [9]	<u>1.15</u>	0.42	<u>6.31</u>	0.98	0.30	4.83	1.00	0.76	9.13	0.84	5.27
MultiCam [13]	1.85	1.94	6.99	0.62	<u>0.00</u>	6.86	1.21	1.96	5.71	0.31	4.34
GC+occl [11]	1.19	<u>0.23</u>	6.71	0.73	0.11	5.71	1.64	2.75	5.41	0.61	6.05
Improved Coop [6]	1.67	0.77	9.67	1.21	0.17	6.90	1.04	1.07	13.68	0.29	3.65
Symbiotic [4]	2.87	1.71	11.90	1.04	0.13	7.32	0.51	0.23	7.88	0.50	6.54
Var.win. [15]	2.35	1.65	12.17	1.28	0.23	7.09	1.23	1.16	13.35	0.24	2.98
Graph Cuts [2]	1.86	1.00	9.35	0.42	0.14	3.76	1.69	2.30	5.40	2.39	9.35
Multiw Cut [1]	8.08	6.53	25.33	0.61	0.46	4.60	0.53	0.31	8.06	0.26	3.27
Cooperative [16]	3.49	3.65	14.77	2.03	2.29	13.41	2.57	3.52	26.38	<u>0.22</u>	<u>2.37</u>
Bay. diff [7]	6.49	11.62	12.29	1.45	0.72	9.29	4.00	7.21	18.39	<u>0.20</u>	2.49

Note:The underlined bold number is the best in its category.

5. Experiments

The proposed algorithm has been implemented on a 2.4GHz Pentium IV PC platform. Typically, the disparity is computed within 3 seconds after the color segmentation for a 384x288 reference image. In this section, we present experimental results of the proposed algorithm on six image pairs. One set of thresholds is used for all the experiments.

The first four image pairs along with their ground truth data are taken from the Middlebury database. We evaluate the performance using the quality measures proposed in [8]. A "bad" pixel is defined as the one whose absolute disparity error is greater than one pixel. Percentages of "bad" pixels are measured among: all pixels (B_{all}), pixels in untextured areas (B_{untex}), and pixels near disparity discontinuities (B_{disc}). Only non-occluded pixels are considered in all three cases.

Table I and Fig. 1 present the overall performance of our algorithm, where Table I summarizes the quantitative evaluation results and Fig. 1 shows the extracted disparity maps. The proposed algorithm performs well, and it reaches the first place in the Middlebury database ranking² (at the submission time). Strong performance is achieved especially in the conventionally challenging areas such as, textureless regions, disparity discontinuous boundaries and occluded portions (see the extracted disparity maps). Note that for the *Map* stereo pair, our unsatisfying result is mainly due to the initial color segmentation errors. The low contrast grayscale texture causes many regions (especially foreground borders) mistakenly segmented. Since our algorithm assumes disparity continuity inside each color segment, we cannot separate background and foreground if they were grouped

² Although the ranking just gives a rough idea of the overall performance of an algorithm.

as one segment in the first place.

In addition, Fig. 2 and Table II report the detailed intermediate results on the *Venus* data. These results demonstrate visually and quantitatively the effectiveness of our plane fitting and energy minimization methods.

The last two image pairs are taken from the *Garden* and the *Studio* sequences, respectively. Since there is no ground truth available, we evaluate the performance by synthesizing novel views through forward warping the reference image based on the computed disparity map. Fig. 3 shows the extracted disparity maps and the synthesized novel views.

Besides the strong numerical and visual performance, an interesting by-product of the proposed algorithm is disparity segmentation. Each disparity segment is a connected region of color segments assigned the same plane index. To overcome occasional disparity over-segmentation, we fit a new plane to each neighboring disparity segment pair and add it to the plane set. Disparity segmentation is then obtained by applying graph cuts based on the new plane set. The reason that this process can eliminate over-segmentation is straightforward: if the added new plane is correct, the smoothness energy term dominates, which favors joining the over-segmented disparity segments together; otherwise, the data energy term dominates, which favors keeping the original disparity segments. We demonstrate the disparity segmentation results for all testing image pairs in Fig. 4.

6. Conclusion

We have demonstrated a new algorithm for stereo matching. Our segment-based approach works well for images with sharp color discontinuities and slanted disparity surfaces. Especially, strong performance is obtained in the conventionally difficult areas such as textureless regions, disparity discontinuous boundaries and occluded portions. In

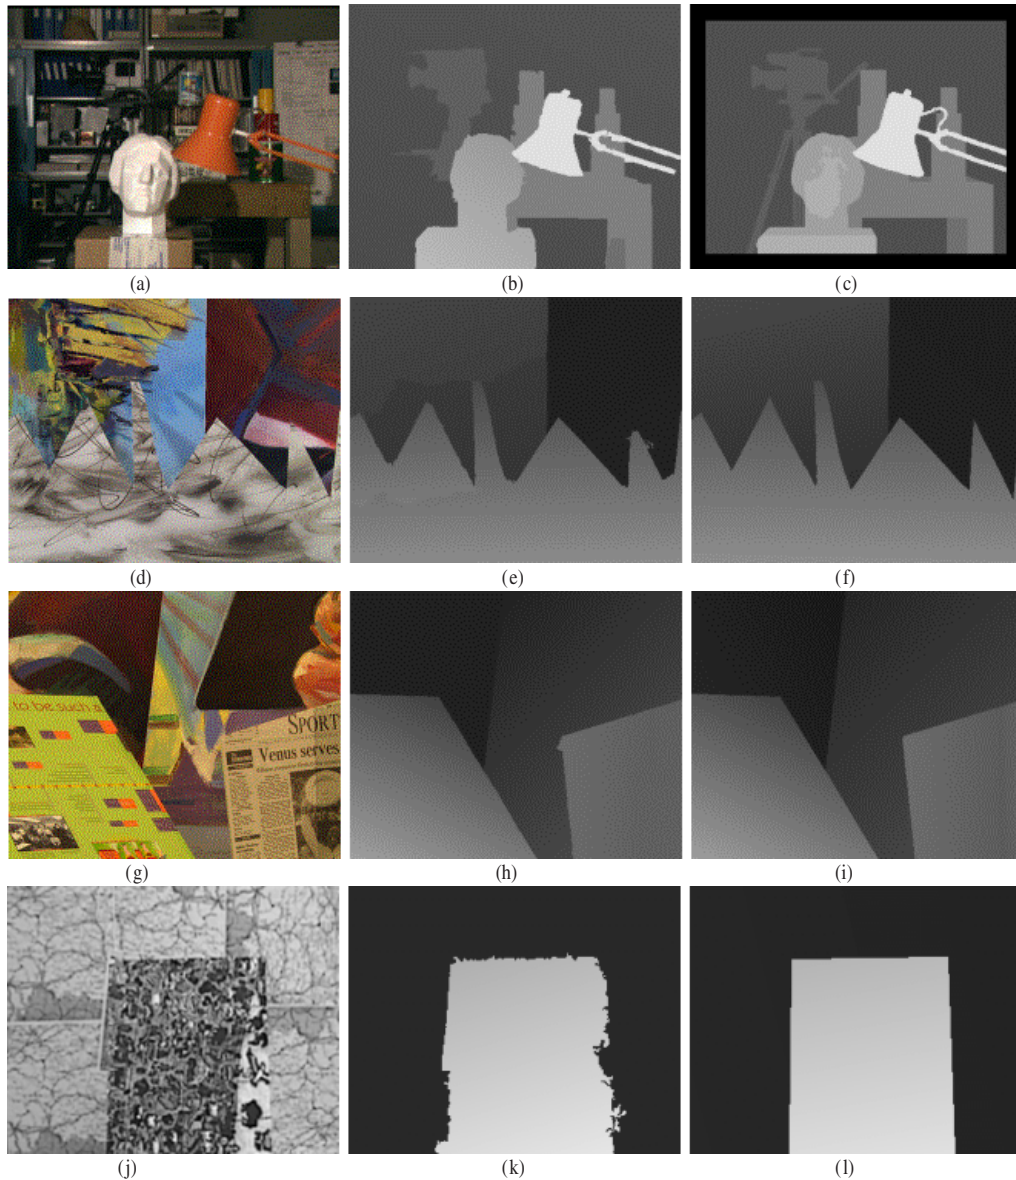


Figure 1: Results on Middlebury data sets. From top to down order: *Tsukuba*, *Sawtooth*, *Venus*, *Map*. From Left to right order: reference images, extracted disparity maps and the ground truth disparity maps (Disparities are scaled according to the provided scaling factor in Middlebury data set).

addition, representing the disparity map as segmentations and plane models is not only much compact than the conventional dense pixelwise disparity map, but also beneficial to other tasks such as image-based rendering, scene analysis, *etc.*

Apart from the advantages of our approach, the current version of our algorithm will not be able to handle the situation if there are disparity boundaries appearing inside the initial color segments. In the future, we plan to investigate more on how to split these problematic segments by combining both color and disparity information.

Acknowledgement

We would like to thank prof. Scharstein and Dr. Szeliski for providing the testing stereo images and evaluation results on their Middlebury website.

References

- [1] S. Birchfield and C. Tomasi, "Multiway Cut for Stereo and Motion with Slanted Surfaces," *Proc. Int'l Conf. Computer Vision*, pp. 489-495, Sept. 1999.




Figure 2: Detailed results on the *Venus* data. (a) Color segmentation. (b) Edge contours of color segmentation. (c), (d), (e), (f) are extracted disparity maps from each intermediate procedure: (c) local matching, (d) initial plane fitting, (e) refined plane fitting, (f) graph cuts labeling.

Table 2: Intermediate results on the *Venus* data

Procedures	B_{all}	B_{untex}	B_{disc}
1	35.91	65.96	30.87
2	6.35	6.25	12.30
3	2.61	4.22	8.91
4	0.08	0.01	1.39

Procedure 1: Local matching. Procedure 2: Initial plane fitting from single segment. Procedure 3: Refined plane fitting from grouped segments. Procedure 4: Plane labeling from graph cuts.

- [2] Y. Boykov, O. Veksler, and R. Zabih, "Fast Approximate Energy Minimization via Graph Cuts," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 23(11), Nov. 2001.
- [3] D. Comaniciu and P. Meer, "Robust Analysis of Feature Spaces: Color Image Segmentation," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 750-755, 1997
- [4] J. Y. Goulermas and P. Liatsis, "A Collective-based Adaptive Symbiotic Model for Surface Reconstruction in Area-based Stereo," *IEEE Trans. Evolutionary Computation*, Vol. 7(5), pp. 482-502, 2003.
- [5] M. Lin and C. Tomasi, "Surfaces with Occlusions from Layered Stereo," Ph.D thesis, Stanford University 2002.
- [6] H. Mayer, "Analysis of Means to Improve Cooperative Disparity Estimation," *ISPRS Conf. Photogrammetric Image Analysis*, 2003.
- [7] D. Scharstein and R. Szeliski, "Stereo Matching with Non-linear Diffusion," *Int'l J. Computer Vision*, Vol. 28, no. 2, pp. 155-174, 1998.
- [8] D. Scharstein and R. Szeliski, "A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms," *Int'l J. Computer Vision*, Vol. 47, no. 1, pp.7-42, 2002.
- [9] J. Sun, N. N. Zheng and H. Y. Shum, "Stereo Matching Using Belief Propagation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 25, no. 7, July 2003.
- [10] H. Tao, H. S. Sawhney and R. Kumar, "A Global Matching Framework for Stereo Computation," *Proc. Int'l Conf. Computer Vision* 2001.
- [11] V. Kolmogorov and R. Zabih, "Computing Visual Correspondence with Occlusions using Graph Cuts," *Proc. Int'l Conf. Computer Vision* 2001.
- [12] V. Kolmogorov and R. Zabih, "What Energy Functions can be Minimized via Graph Cuts?" *Proc. European Conf. Computer Vision* 2002.
- [13] V. Kolmogorov and R. Zabih, "Multi-Camera Scene Reconstruction via Graph Cuts," *Proc. European Conf. Computer Vision* 2002.
- [14] O. Veksler, "Efficient Graph-based Energy Minimization Methods in Computer Vision", PhD thesis, Cornell University, July 1999.
- [15] O. Veksler, "Fast Variable Window for Stereo Correspondence Using Integral Images," *Proc. Int'l Conf. Computer Vision and Pattern Recognition*, 2003.
- [16] C. L. Zitnick and T. Kanade, "A Cooperative Algorithm for Stereo Matching and Occlusion Detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(7), pp. 675-684.

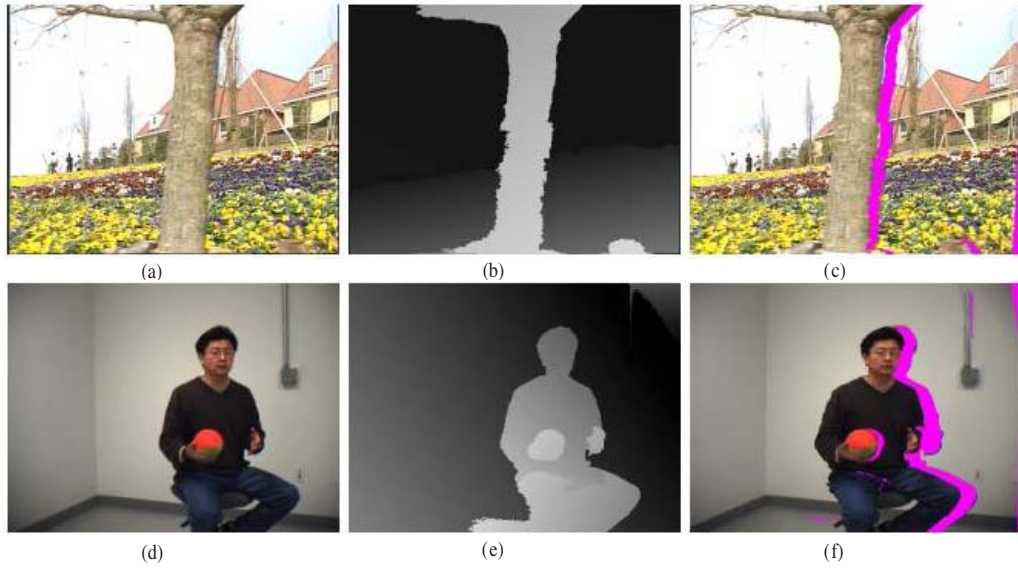


Figure 3: Results on the *Garden* and *Studio* data. (a), (d) Reference images from *Garden* and *Studio*. (b), (e) the corresponding extracted disparity maps. (c), (f) Novel views (Purple pixels are occluded regions).




Figure 4: Disparity segmentation results. Each plane index corresponds to a different color. (a) *Tsukuba*, (b) *Sawtooth*, (c) *Venus*, (d) *Map*, (e) *Garden*, (f) *Studio*.