Sparse Stereo Disparity Map Densification using Hierarchical Image Segmentation

Sébastien Drouyer, Serge Beucher, Michel Bilodeau, Maxime Moreaud, Loïc Sorbier

► To cite this version:

Sébastien Drouyer, Serge Beucher, Michel Bilodeau, Maxime Moreaud, Loïc Sorbier. Sparse Stereo Disparity Map Densification using Hierarchical Image Segmentation. 13th International Symposium, ISMM 2017, May 2017, Fontainebleau, France. pp.172-184. hal-01484143

HAL Id: hal-01484143

<https://hal-mines-paristech.archives-ouvertes.fr/hal-01484143>

Submitted on 29 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike| 4.0 International License

Sparse Stereo Disparity Map Densification using Hierarchical Image Segmentation

Sébastien Drouyer^{1,2} ✉, Serge Beucher¹, Michel Bilodeau¹, Maxime Moreaud^{2,1}, and Loïc Sorbier²

¹ Mines ParisTech, PSL Research University, CMM - Centre de morphologie mathématique, France

² IFP Energies nouvelles, Rond-point de l'échangeur de Solaize, BP 3, 69360 Solaize, France

{sebastien.drouyer, serge.beucher, michel.bilodeau@mines-paristech.fr
maxime.moreaud, loic.sorbier@ifpen.fr}

Abstract. We describe a novel method for propagating disparity values using hierarchical segmentation by waterfall and robust regression models. High confidence disparity values obtained by state of the art stereo matching algorithms are interpolated using a coarse to fine approach. We start from a coarse segmentation of the image and try to fit each region's disparities using robust regression models. If the fit is not satisfying, the process is repeated on a finer region's segmentation. Erroneous values in the initial sparse disparity maps are generally excluded, as we use robust regressions algorithms and left-right consistency checks. Final disparity maps are therefore not only denser but can also be more accurate. The proposed method is general and independent from the sparse disparity map generation: it can therefore be used as a post-processing step for any stereo-matching algorithm.

Keywords: Stereo, hierarchical segmentation, robust regression model, waterfall, disparity map, densification

1 Introduction

One of the main research interest in computer vision is the matching of stereoscopic images, as it allows to obtain a 3d reconstruction of the observed scene. A common practice is to first rectify the pair of images [2], that is to say slightly distort them so that a point of the scene is projected on the same ordinate on both images. This rectification allows to simplify the matching process to a 1D search problem. The difference of abscissa of a point's projection between the rectified pair of images is called disparity. A disparity map estimates the disparity for each pixel of the rectified pair.

Many different algorithms evaluating these disparity maps already exist in the literature. They generally use local [14], semi-global [10] [?] or global optimization methods [13] [24] [8] that minimize matching cost of image features between two images. While some methods, especially the global ones, can produce dense disparity maps, other methods produce sparse disparity maps - where

there are some pixels with undefined disparity values. This can happen for several reasons. Disparity values can be removed when a confidence measure, such as texture or uniqueness, is under a certain threshold. They can also be removed when a part of the scene is occluded, either by global and semi-global methods, or by using Left-Right Consistency (LRC) checks[?].

Global disparity evaluation methods often feature disparity propagation in order to generate a complete map [1] [21], but these propagation techniques are directly linked to the method and are not generalizable.

Apart from diffusion [23] and interpolation [18], a few general propagation methods exist in the literature. From a disparity map where only edges are matched, VMP [17] diffuses disparity values by using predefined masks and a voting process. FGS [15] has been used by the OpenCV library [11] for propagating LRC checked disparity values using an edge-preserving diffusion model. RBS [3] proposes a similar diffusion model, but using a modified version of the bilateral filter.

This work proposes a new densification method that is able to produce a denser and more accurate disparity map from an initial sparse disparity map. We called our method TDSR for Top Down Segmented Regression. From a coarse segmentation of the image, we try to fit each region's disparities using robust regression models. If the fit is not satisfying, the process is repeated on a finer region's segmentation.

This coarse to fine approach allows the algorithm to rely on general purpose parameters. Since each region is processed separately, the algorithm can be easily parallelized to improve efficiency. As we use linear regression models, it is possible to get not only the disparity values but also normals in the disparity space, which can be convenient for multi-view stereo algorithms and object recognition. Finally, the proposed method is general and independent from the sparse disparity map generation: it can therefore be used as a post-processing step for any stereo-matching algorithm.

2 Top Down Segmented Regression Method

2.1 General Concept

We have two rectified stereo images (left and right) depicting the same scene (see example in Figure 1a). We also have a sparse disparity map between the left and the right images obtained using an existing stereo matching algorithm (see Figure 1b). Large areas of the disparity map can be undefined (gray pixels), and defined areas can contain errors. Furthermore, we know the camera matrix: if a point is defined in the sparse disparity map we can know its 3D relative position in the scene, and vice versa. From these information, we want to obtain a complete and accurate disparity map (see Figure 1c).

We will suppose that the scene is a collection of different **objects**. Each object's surface is composed of one or multiple **simple shapes**. We will consider here that those shapes can be represented by a plane: all the 3D points inside

Fig. 1: From sparse disparity map to complete one, illustrated by the Vintage pair of the Middlebury dataset[20].

a shape can be modeled using a bivariate linear polynomial, that we will call **model**.

Our method will therefore try to separate all the simple shapes in the reference (left) image. If each region is a simple shape, we can obtain a 3D point cloud estimate of the corresponding region from the sparse disparity map (by converting each disparity value to a 3d point) and modelize it by a bivariate linear polynomial. The model can then be used to correct existing disparity values and fill undefined values in the region, by converting back each 3d point of the model to a disparity value. However, this is not automatically feasible by a segmentation algorithm, so we will instead rely on hierarchical segmentation.

The hierarchical segmentation of an image can be represented by a **partition tree**. The concept is very similar to the binary partition tree[19] presented by Salembier and Garrido, except that a node can have more than two children. The top node represents the whole image, its children represent the most important regions of the image according to the algorithm used, their children represent secondary regions, and so on until a node can't be separated into more sub regions.

An adapted segmentation can't be obtained by simply choosing a specific level of the partition tree: nodes at the lower levels will produce a very coarse segmentation of the image, nodes at the higher levels will produce an over-segmentation. At the middle, some simple shapes are over-segmented and others are merged with neighboring shapes.

We will therefore proceed the following way. We start from the top node, and try to modelize its associated 3D point cloud. If the modelization is satisfying or if the node doesn't have any children, we stop here, associate the model to the node, and delete all the node's children. If it is not satisfying, we retrieve the node's children, and apply the same process individually to each node.

If we combine all the leafs of the resulting tree, we can obtain a **modeled segmentation** of the image, where each region is associated to a model. Some post-processing are done on this modeled segmentation that we will describe in Section 2.6. It can then be converted to a complete disparity map.

In order to concretize this general concept, we need to define the following: the sparse disparity map S , the partition tree H , the region modelization and the conditions defining a satisfying modelization.

We will now present those different components. For doing that, we will need to define some parameters. They have been chosen using a grid search to minimize the average error of the disparity maps produced for the middlebury dataset[20].

2.2 The Sparse Disparity Map S

In order to compute a complete disparity map, we first need an initial sparse disparity map. The quality of our complete disparity map is highly correlated to the quality of the initial disparity map, so choosing an accurate algorithm at this stage is crucial.

We chose to compute this disparity map between the left and right image using the MC-CNN algorithm[?] algorithm as it evaluates high quality disparity maps. The algorithm use a convolutional neural network for computing the matching cost of two 11x11 blocks (one from the left image and one from the right image). The results are then refined using a series of post-processing steps: cross-based cost aggregation, semiglobal matching, a left-right consistency check [?], subpixel enhancement, a median filter, and a bilateral filter. The authors provide two sets of architectures: one optimized for speed, and one for accuracy. We chose the one for accuracy.

To give a sense on how critical the quality of the sparse disparity map is for the performance of our algorithm, we previously used the semi-global MGM [?] algorithm for producing our initial maps. The final disparity maps' average error on the Middlebury dataset[20] was more than 50% higher compared to the current algorithm.

2.3 The Partition Tree H

Obtaining an adapted partition tree for our reference images (left images) is not a straightforward task.

First, the assumption of having simple shapes separated at least at the leafs of the partition tree can be difficult to attain on some images. Segmentation, whether it provides hierarchical information or not, relies on the image gradient. An object occluding another one with similar color can, in the image, have indistinguishable borders with very low or even null gradient.

Secondly, the partition tree must divide the image in a progressive manner. If it does not - for instance a simple shape is merged with another one at a level and divided into twenty parts at the next level of the tree - it can cause our approach to poorly perform in many ways. Since the input disparity map is sparse, it is possible that some regions at the over-segmented level would contain no disparity values making them impossible to modelize. Since regions are much smaller and contain fewer disparity points, the modelization would be more sensitive to errors in the sparse disparity map.

We took these two problems into account when constructing our partition tree. This partition tree will be obtained from the image gradient and defined markers, so we address these issues when computing them.

A first problem that can occur is to have a border with a very progressive change of color between two regions. If we compute a morphological gradient [4] of size 1 in those areas, we will have a very low gradient around those borders, preventing an accurate segmentation. There is a change of color that is occurring, but we need to look at a larger scale. However, we don't want to simply use a morphological gradient with a larger size since we might lose out on small details. That is why we used the multi-scale gradient presented in [7]. We computed this gradient from scale 1 to scale 6.

For markers, we initially used the h-minima transformation of the gradient, with $h = 5$. A second problem is that, when two adjacent regions have similar colors, the gradient can be less than h at some parts of the border between them. This effect is known as gradient leakage [22]. A way to detect these leakage is that markers will get thinner at these locations. Therefore, we want to encourage a split when the markers get thinner. That is why we applied on these markers the adaptive erosion presented in [6] with $\alpha = 0.25$.

We then compute the valued watershed segmentation from the gradient and markers, and apply on it the enhanced waterfalls operator discussed in [5], as it better preserves the important contours than the standard waterfall algorithm [4]. The result is a grayscale image S_h . If $S_h(x, y) = n$ and $n \neq 0$, then the pixel is a border between two or more regions at the level n or higher of the hierarchical segmentation. The higher the level, the coarser the segmentation. See figures 2a 2b 2c.

Algorithm 1 shows how we transform S_h to a partition tree:

Algorithm 1: From S_h to a partition tree

```

Function getPartitionTreeFromSh( $S_h$ )
     $N \leftarrow \max(S_h)$  ; // Set N as the highest level of segmentation
     $T \leftarrow \text{newPartitionTree}()$  ;
    foreach  $n$  in  $[N, N - 1, \dots, 2, 1]$  do
         $S_h^n \leftarrow S_h \geq n$  ; // set  $S_h^n$  as the segmentation at level n
         $L_h^n \leftarrow \text{label}(S_h^n)$  ; // Label  $S_h^n$  and save it in  $L_h^n$ 
         $L_h^n \leftarrow \text{unique}(L_h^n)$  ; /* Make each label unique in  $L_h^n$ : no similar
        labels in  $L_h^n$  should exist in previously processed levels */
        foreach label  $l$  in  $L_h^n$  do
             $R \leftarrow L_h^n == l$  ; // Set  $R$  as the region in  $L_h^n$  where label =  $l$ 
            // Set  $p$  as parent id of label  $l$ 
            if  $n == N$  then
                 $p \leftarrow T.\text{topNode}$  ; // The parent is the top node
            else
                 $p \leftarrow \text{Get label in } L_h^{n+1} \text{ inside the region } R$  ; /* Note: no
                region in  $L_h^n$  can belong to two regions in  $L_h^{n+1}$  */
             $T.\text{addNode}(\text{id}=l, \text{parentId}=p, \text{region}=R)$  ;

```

Fig. 2: Waterfall and hierarchy tree. (a) Extract of the left image of the Adirondack pair in the Middlebury dataset. (b) Top level of the enhanced waterfall segmentation of (a). (c) Intermediate level. (d) Example of a top level of waterfall segmentation, where each region has been labeled. (e) Example of a lower level of waterfall segmentation, labeled. (f) Hierarchy tree constructed from (d) and (e).

2.4 Region Modelization

When we process each node of the partition tree, we want to obtain a model that best reflects the real geometry of the region. For doing that, we have an estimated 3D point cloud extracted and converted from the sparse disparity map which contain errors and bias. We need to address two points. First, we have to define whether we take all the 3D points into account or we select a specific subset. Then, we have to specify how the model is computed from this subset.

Points Selection

A common issue is the well-known fattening effect[12]. For computing each point's disparity, stereo algorithms generally match blocks of pixels [14] [10] [?] instead a single pixels for more robustness. However, around strong edges, the center of the block inherits the disparity of the more contrasted pixels in the block, hence the fattening effect. The worst consequences of this effect take place around occlusions: near the border separating two objects, pixels on the occluded object will inherit the disparity of pixels on the occluding object. Such an issue can seriously impede the modelization process.

Since disparities around borders are not reliable, we decided not to take them into account when computing the model. For doing that, we first separate the region A into two separate regions: A_b being the border of A and A_i being the in-

side of A . More specifically $A_i = \varepsilon(A)$, $A_b = A \setminus A_i$, $\varepsilon(A)$ being the morphological erosion of A .

We want here to select all the pixels in A_i whose matching blocks don't intersect with the border. Therefore the retained pixels $\hat{A}_i$ belong to $\hat{A}_i = \varepsilon_{\frac{B}{2}}(A)$, where $\varepsilon_{\frac{B}{2}}$ is the morphological erosion of size $\lceil \frac{B}{2} \rceil$ (on a square grid) and B is the matching block size.

We will keep all pixels in A_b in the subset where the models will be fitted, as $\hat{A}_i$ might contain an insufficient number of disparity values for an accurate fitting. For instance, a region might have a completely textureless inside resulting in an absence of disparity values. Though there might be some outliers in the borders, due notably to occlusions, we are relying on the model computation method we will describe later to filter them out.

As a result, we will modelize only the 3D points from the sparse disparity map in the region $\hat{A} = (\hat{A}_i \cup A_b)$.

Points Modelization

Now that we have selected the 3D points, we can compute our model. There are however some issues that need to be addressed.

First, we have seen previously that some points from A_b can contain outliers which disparity have been contaminated by a nearby object. An other issue is the presence of small areas with disparity values very different from the ground truth [16]. The problem generally appears on areas where there is very little or repetitive texture, as the stereo matching algorithms generally fail on these areas.

We will therefore proceed the following way. We first modelize the points using a linear regression: we try to fit $z = A + Bx + Cy$. If the model is satisfying (see Section 2.5), we stop here and the linear equation is affected as the model of the region. If it isn't, we use RANSAC[9] to perform a robust regression on these 3D points: if the proportion of outliers isn't too large, they should be automatically excluded from the model evaluation process thanks to RANSAC. The obtained linear equation is then affected as the model of the region.

If there are no points in the sparse disparity map, we can't compute any model: we will affect to the region an undefined model.

2.5 Conditions Defining a Satisfying Modelization

We have constructed a model from a set of 3D points, and now we want to know if the model is satisfying. By satisfying, we mean that we want a large proportion of points that are close to the model. First, we convert the model to disparity values thanks to the camera matrix. For each point i of the set, we know therefore its value in the sparse disparity map d_i and we also know its disparity value according to the model $\tilde{d}_i$. A point is an outlier if $|d_i - \tilde{d}_i| > t_d$, t_d being a custom threshold (we chose $t_d = 2$). We define the fitting score s as

the percentage of points in the set that are not outliers. A model is satisfying if $s > t_s$, with $t_s = 70\%$, and if the number of outliers $n_o < 100$.

2.6 Post-processing

We have at this stage a modeled segmentation S_m of our image. To further improve the results obtained, some post-processing are done on S_m . We will note D the disparity map obtained from S_m .

The first problem is that probably some regions have an undefined model because of the absence of points in the sparse disparity map. We have to affect a valid model to these regions, and to do that we will use the modeled segmentation's concept to our advantage. Indeed, these regions are surrounded by regions associated with a model: the main idea is to fill the undefined region with the most probable model in the adjacent regions.

However, since each undefined region might cover multiple simple shapes, we will cut these regions the following way. First, we create a label map ℓ_1 labeling all regions with an undefined model. We create a second label map ℓ_2 labeling all the regions of a similar watershed segmentation than the one computed in Section 2.3 but with $h = 12$. We then compute ℓ_f such as $\ell_f(x) = \ell_f(y) \iff ((\ell_1(x) = \ell_1(y)) \wedge (\ell_2(x) = \ell_2(y)))$. Each region labeled in ℓ_f will then be processed independently.

For each label l and associated area A_l , we define its external border as $B_l = (\delta A_l) - A_l$, δA_l being the morphological dilation of A_l of size 1. We list all different models in $S_m(B_l)$. We want to choose the model that creates the largest consensus across adjacent regions separated to the current region by a low gradient. For doing that, we create a list l_p of all positions where $g(B_l) < \min(g(B_l)) + t_g$, g being the gradient of the left image obtained following the process explained in Section 2.3. t_g is a threshold to be defined (we chose 10). We affect to $S_m(A_l)$ the model M that maximize the number of points p in l_p such as $(M(p) - D(p)) < t_d$, t_d being the threshold used in Section 2.5.

The labels are processed from the one that contain the lowest proportion of undefined models in B_l to the one that contain the highest.

Once S_m and D are filled, we compute S'_m and D' using the same process than for S_m and D but by inverting the left and right images. We remove values in D that are inconsistent according to the LRC check with threshold set to 1. We also remove values at the same positions in S_m . We then fill the values using the same process as described earlier.

3 Results

We benchmarked our TDSR method using the well-known Middlebury dataset[20]. The dataset contains several stereo-images pairs with their associated disparity map ground truth. An executable is also provided allowing to quantitatively compare disparity maps to their ground truth according to various criteria. We chose the following two commonly used criteria in the Middlebury comparative

Table: “Bad 2.0”, which computes the percentage of disparity values farther than 2.0 from their associated ground truth, and “Avg. error”, that computes the average absolute difference between the disparity values and their associated ground truth. These criteria have been separately computed on all values and only on values where no occlusion is occurring. They were computed on full resolution (F).

We compared the results of our TDSR method with three interpolations techniques directly applied on the sparse disparity map: nearest, linear, and cubic.

We also compared our results with the Weighted Least Squares disparity map post filtering method of the OpenCV library [11]. This method uses the Fast Global Smoothing [15] algorithm for diffusing high confidence disparity values using an edge preserving scheme. However, this method is parametric, bringing additional challenges to the comparison. There are three main parameters: “depthDiscontinuityRadius” (or “depthDR”), defining the size of low-confidence regions around depth discontinuities, “lambda”, defining the amount of regularization, and “sigma”, defining how sensitive the filtering process is to the source image edges. We computed three complete maps: one with the default parameters, one where the parameters have been optimized to minimize “Bad 2.0”, and one where the parameters have been optimized to minimize “Avg. error”. The parameters have been optimized using a grid search. Tested parameters are shown in Table 1.

Parameter	Tested	Default	“Bad 2.0” optimized	“Avg. error” optimized
depthDR	1, 3, 5, 7, 10	5	3	5
lambda ($\times 1000$)	0, 0.1, 0.3, 0.5, 1, 2, 4, 6, 8, 10, 12, 14, 16, 20	8	0.1	6
sigma (/10)	1, 2.5, 5, 7.5, 10, 12.5, 15, 17.5, 20, 30	10	10	15

Table 1: Tested parameters for FGS.

Finally, we compared our results with the RBS [3] method, which uses a variant of the bilateral algorithm to propagate disparity values by preserving edges.

Produced disparity maps can be qualitatively compared on the “Vintage” pair of the Middlebury dataset in Figure 3. They can be quantitatively compared on the whole training set in Table 2.

Compared to other stereo methods in the Middlebury benchmark, on the training set, using the “Avg. error” criterion on all pixels, our solution currently ranks 1st over 59. The results are therefore comparable with state of the art methods.

Interpolation methods are very sensitive to the noise in the sparse disparity map and they don’t work well on occlusions. Depending on the chosen parameters, the Fast Global Smoothing method can either produce a less noisy disparity

Fig. 3: Image and disparity map produced for the Vintage pair of the Middlebury dataset.

Method	Bad 2.0 (no occ)	Avg. error (no occ)	Bad 2.0 (all)	Avg. error (all)
Nearest interpolation	10.1	2.83	18.3	7.53
Linear interpolation	10.6	2.71	20.2	7.61
Cubic interpolation	10.6	2.75	20.1	7.70
FGS (Optimized on Bad 2.0)	10.4	2.62	17.4	8.20
FGS (Default)	16.5	2.60	23.1	6.80
FGS (Optimized on Avg. error)	21.3	2.85	27.7	6.61
MC-CNN+RBS[3]	10.8	2.60	19.3	6.66
MC-CNN+TDSR	10.2	2.28	15.6	4.56
→ Improvement to best interp.	-1%	16%	15%	40%
→ Improvement to best FGS	2%	12%	10%	31%

Table 2: Quantitative comparison between interpolation, FGS, RBS, and our proposed TDSR method. *all*: Taking all points into account. *no occ*: Taking only non-occluded points into account.

map but with the risk of having the disparity of some objects contaminated by neighboring ones (high “lambda”, optimized for “Avg. error”), or more noise but less contamination (low “lambda”, optimized for “Bad 2.0”). Default parameters give a good tradeoff between both effects.

Qualitatively, compared to the interpolation, FGS and RBS methods, our approach appears robust to noise and seems to globally manage occlusions. However, some occlusions, notably near the computer screen in the foreground (see Figure 3f), are not well evaluated due to lack of data. Some regions can also have their disparity contaminated by nearby objects with similar color.

Quantitatively, compared to the interpolation, FGS and RBS methods, whether we choose to only look at non occluded pixels or to look at all the pixels, our method produces at least similar results according to all chosen criteria. It is important to note that our algorithm’s parameters were fixed, though we optimized the parameters of the FGS algorithm for each criterion and compared our algorithm’s results to the best cases of the interpolation and FGS methods.

4 Conclusion

We have presented the Top Down Segmented Regression (TDSR) algorithm that allowed us to densify noisy sparse disparity maps. Our method generated promising results. TDSR is more robust to noise and better preserves the edges than interpolation or diffusion algorithms. The quality of the produced complete disparity map is comparable with state of the art methods.

Our approach can also be used in complement to any stereo matching algorithm as a post-processing step, though its performance will depend on the accuracy of the initial sparse disparity map.

In fact, TDSR could also be adapted to densify other sparse spatial data or to refine dense but noisy data. Dense disparity maps refinement, depth map super-resolution or semantic segmentation post-processing are potential applications that would require very little changes on the proposed approach.

The source code and executable of our implementation are available at:

<https://hal-mines-paristech.archives-ouvertes.fr/hal-01484143>

References

1. L. Alvarez, R. Deriche, J. Sánchez, and J. Weickert. Dense disparity map estimation respecting image discontinuities: A PDE and scale-space based approach. *JVCIR*, 13(1-2):3–21, mar 2002.
2. N. Ayache and C. Hansen. Rectification of images for binocular and trinocular stereovision. In *ICPR 88*.
3. J. T. Barron and B. Poole. The fast bilateral solver. In *Computer Vision – ECCV 2016*, pages 617–632. Springer Nature, 2016.
4. S. Beucher. *Image segmentation and mathematical morphology*. Theses, École Nationale Supérieure des Mines de Paris, June 1990.
5. S. Beucher. Towards a unification of waterfalls, standard and P algorithms. working paper or preprint, Jan. 2013.

6. J.-C. Bricola, M. Bilodeau, and S. Beucher. A top-down methodology to depth map estimation controlled by morphological segmentation. Technical report, 2014.
7. J.-C. Bricola, M. Bilodeau, and S. Beucher. A multi-scale and morphological gradient preserving contrast. In *14th International Congress for Stereology and Image Analysis*, Liège, Belgium, July 2015. Eric Pirard.
8. J.-C. Bricola, M. Bilodeau, and S. Beucher. Morphological processing of stereoscopic image superimpositions for disparity map estimation. working paper or preprint, Mar. 2016.
9. M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, jun 1981.
10. H. Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):328–341, feb 2008.
11. Itseez. Open source computer vision library. <https://github.com/itseez/opencv>, 2015.
12. T. Kanade and M. Okutomi. A stereo matching algorithm with an adaptive window: theory and experiment. In *Proceedings. IEEE ICRA.*, 1991.
13. V. Kolmogorov. *Graph Based Algorithms for Scene Reconstruction from Two or More Views*. PhD thesis, Ithaca, NY, USA, 2004. AAI3114475.
14. K. Konolige. Small vision systems: Hardware and implementation. In *Robotics Research*, pages 203–212. Springer Nature, 1998.
15. D. Min, S. Choi, J. Lu, B. Ham, K. Sohn, and M. N. Do. Fast global image smoothing based on weighted least squares. *IEEE Transactions on Image Processing*, 23(12):5638–5653, dec 2014.
16. K. Moravec, R. Harvey, and J. A. Bangham. Improving stereo performance in regions of low texture. In *BMVC*, 1998.
17. J. Ralli, J. Díaz, and E. Ros. A method for sparse disparity densification using voting mask propagation. *Journal of Visual Communication and Image Representation*, 21(1):67–74, jan 2010.
18. J. Ralli, F. Pelayo, and J. Diaz. Increasing efficiency in disparity calculation. In *Lecture Notes in Computer Science*, pages 298–307. Springer Nature.
19. P. Salembier and L. Garrido. Binary partition tree as an efficient representation for filtering, segmentation and information retrieval. In *ICIP 98*.
20. D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Lecture Notes in Computer Science*, pages 31–42. Springer Nature, 2014.
21. D. Scharstein and R. Szeliski. Stereo matching with non-linear diffusion. In *CVPR*, 1996.
22. C. Vachier and F. Meyer. The viscous watershed transform. *Journal of Mathematical Imaging and Vision*, 22(2-3):251–267, may 2005.
23. J. Weickert. *Anisotropic diffusion in image processing*. PhD thesis, 1998.
24. Q. Yang, L. Wang, R. Yang, H. Stewenius, and D. Nister. Stereo matching with color-weighted correlation, hierarchical belief propagation and occlusion handling. In *CVPR*, 2006.