

A hierarchical stereo algorithm using dynamic programming

G. Van Meerbergen, M. Vergauwen, M. Pollefeys, L. Van Gool

Abstract

In this paper, a new hierarchical stereo algorithm is presented. The algorithm matches individual pixels in corresponding scanlines by minimizing a cost function. Several cost functions are compared. The algorithm achieves a tremendous gain in speed and memory requirements by implementing it hierarchically. The images are downsampled an optimal number of times and the disparity map of a lower level is used as 'offset' disparity map at a higher level. An important contribution consists of the complexity analysis of the algorithm. It is shown that this complexity is independent of the disparity range. This result is also used to determine the optimal number of downsample levels. This speed gain results in the ability to use more complex (compute intensive) cost functions that deliver high quality disparity maps. Another advantage of this algorithm is that cost functions can be chosen independent of the optimization algorithm. Finally, the algorithm was carefully implemented so that a minimal amount of memory is used. It has proven its efficiency on large images with a high disparity range as well as its quality. Examples are given in this paper.

1. Introduction

The goal of this work is to compute fast accurate disparity maps, even for big images with large disparity ranges. It is assumed that these images are rectified first. The importance of the disparity maps for good 3D reconstruction is well known. The results of research on this topic include a wide variety of algorithms [1]. Most of these algorithms are minimizing a cost function. The cost has 2 aspects: A first component decides how well a pixel of the left image corresponds to a pixel at the right image, i.e. the dissimilarity between 2 pixels. A second, respectively third component rely on the continuity in the horizontal, respectively vertical dimension of the disparity map. The latter is only dominant if no strong image data are available.

The dimension of the optimization problem is one way to classify these algorithms. The cost function decides how well a pixel of the left image corresponds to a pixel at the right image (zero dimensional cost function) [12]. More robust is to optimize the sum of all costs of all matches in one scanline (one dimensional). The best is to optimize the sum of all costs over the whole image (two dimen-

sional) [8]. This algorithm, like many other algorithms [5] [7], has chosen the second method, because it seems the best tradeoff between robustness and computational complexity. However, we'll take into account information of the other scanlines by propagating information between them.

Our stereo algorithm has a number of specific properties. The algorithm is based on the hierarchical solution of the correspondence problem. Using this approach, we'll show that the complexity decreases from $O(W.H.\delta_{range}^2)$ to $O(W.H)$, where W and H represent respectively the width and height of the pictures. Note that the complexity becomes independent of the disparity search range δ_{range} . The algorithm offers a significant advantage over other algorithms, especially when this search range is large. [12] [8]

Another advantage of this hierarchical approach is that the memory usage can be greatly reduced. The acceleration discussed in the first topic does not lead to an explosion of memory usage. As will be explained in following paragraphs, rather the contrary is true. Considering today's arithmetic calculation power and memory bottleneck, this is very important in implementing the algorithm. Even for large images (1500x1500), it uses only a 100kbyte, so that it fits easily into Pentium memory caches.

This algorithm is a candidate for real time implementation, as it offers a good tradeoff between speed and quality. Because of the gain in speed, it was possible to use more complex and computation intensive cost functions. Secondly, because of decoupling between algorithm and cost function, several cost functions could be tested.

The paper is structured as follows. First a formal problem definition is given. In section three, a tree search algorithm that has been used as a starting point for this work is given. In the next section, an efficient hierarchical implementation is developed, followed by its complexity analysis in section five. In section six, a comparison is made with known solutions from literature. Finally, results and conclusions are presented.

2. Problem Definition

In this section, we first introduce some important definitions that are used throughout this paper. Then, we discuss the cost function and go more in detail about one of its components, the dissimilarity function. We end up with a formal problem

definition.

Match sequence The correspondence problem, that we are trying to solve, searches for matches between left and right image. A match is noted $M = (M_{xL}, M_{xR})$, M_{xL} is the index of the pixel in the left image and M_{xR} is the pixel index in the right image. The disparity δ of a match must lie between 2 predefined constants, as postulated by constraint C1:

$$C1 : \delta_{min} \leq \delta = xL - xR \leq \delta_{max} \quad (1)$$

We define $\delta_{range} = \delta_{max} - \delta_{min} + 1$.

An ordered set of matches $\{M_i\}$ that obeys constraints C2..C4, is called a match sequence MS.

$$\begin{aligned} C2 : & M_{i,xL} < M_{j,xL} \& M_{i,xR} < M_{j,xR} \\ & \text{where } 0 \leq i < j < N_m - 1 \\ C3 : & M_{i+1,xL} = M_{i,xL} + 1 \vee M_{i+1,xR} < M_{i,xR} + 1 \\ & \text{where } i = 0..N_m - 2 \\ C4 : & M_{0,xL} = 0 \vee M_{0,xR} = 0 \end{aligned} \quad (2)$$

It is important to have a closer look at these constraints. The first constraint C2 expresses that no pixel can ever belong to 2 matches in the same match sequence and an ordering constraint is imposed. Constraint C3 imposes that the match sequence may not contain gaps. If there is a match M, then either the next pixel in the left, or the next pixel in the right scanline must be matched to the match following M. We say that the match sequence is complete if also the last pixel of left and/or right scanline appears in the (last) match. In Figure 1, an example is given to clarify these definitions. The match sequence used in the example is:

$$MS = \{(2, 0), (3, 1), (4, 2), (5, 3), (6, 4), (7, 6), (8, 7), (9, 8), (12, 9), (13, 10), (14, 11), (15, 12)\} \quad (3)$$

Instead of coding a match sequence as given in the example above, one can as well just code the disparities in what we call a δ - sequence. The disparity map consists of




Figure 1: Graphical representation of a match sequence with a left and a right occlusion. The left scanline is on top. The corresponding disparities are also indicated.

nothing else than all δ - sequences. The δ - sequence that corresponds with the example above is¹:

$$DS = \{Undef, Undef, 2, 2, 2, 2, 2, 1, 1, 1, \underline{2}, \underline{2}, 3, 3, 3, 3\} \quad (4)$$

Cost function Once we defined a match sequence, a cost can be associated to it. As said in the introduction, the cost function consists of 3 components (The intensity of the images is represented by I):

$$\begin{aligned} \kappa(DS) = & \sum_{xL=0}^{W-1} (D(xL, xR) |_{match}) + \\ & \sum_{xL=0}^{W-1} (\beta |_{occl}) + \sum_{xR=0}^{W-1} (\beta |_{occl}) + \\ & \sum_{xL=0}^{W-1} (\alpha | DS(xL) - DS_{-1}(xL) | IG^Y(xL)) \end{aligned} \quad (5)$$

- For each unoccluded pixel, a cost that equals the dissimilarity between 2 pixels is taken into account. This dissimilarity function is discussed in detail further on.
- If the pixel is occluded, a cost equal to β is counted. In contrast to other algorithms in the literature [4], larger occlusions are more severely penalized. Left occlusions are taken into account in the first summation, while right occlusions are in the second summation.
- The last term looks at the vertical dimension in the image. If this δ - sequence differs a lot from the scanline above, its penalty will be higher. This term is also scaled by an intensity gradient factor IG^Y that decreases this penalty if there is a large vertical intensity gradient near the pixel under consideration. The intensity gradients $\frac{\partial I(xL)}{\partial y}$ are easily calculated using the Sobel operator. The actual function used is:

$$IG^Y(xL) = \max\left(1, \frac{a^2}{64 + S_Y(xL)}\right) \quad (6)$$

where $S_Y(xL) = Sobel(xL)$

The parameter a is scaled such that a sobel output of 70 gives $IG^Y = 1$. Outputs above 1 are clipped.

Dissimilarity Function The dissimilarity function indicates how well (the neighbourhood of) the pixel at position xL matches with pixel xR. It is the most important term in the cost function. Three dissimilarity functions are discussed.

¹The underlined disparities are interpolated. This is necessary for the hierarchical implementation.

1. The easiest dissimilarity function is to take the absolute value of difference in intensities.

$$D(xL, xR) = |I(xL) - I(xR)| \quad (7)$$

The advantages are its simplicity and speed. Its practical use is very limited because it is not robust. To make it more robust, this value can be calculated over a window. It is named Sum of Absolute Differences (SAD). The Sum of the Square Differences can also be calculated (SSD).

2. A second alternative is a more complex dissimilarity function that implements a cross correlation between the neighbourhoods of the pixels under consideration [10]. It is implemented according to following formula:

$$D(xL, xR) = \frac{\sum_{i=-by}^{by} \sum_{j=-bx}^{bx} L_{ij} R_{ij}}{\sqrt{\sum_{i=-by}^{by} \sum_{j=-bx}^{bx} L_{ij}^2 \sum_{i=-by}^{by} \sum_{j=-bx}^{bx} R_{ij}^2}}$$

where $L_{ij} = (I_L(y+i, xL+j) - \bar{I}_L)$
 $R_{ij} = (I_L(y+i, xR+j) - \bar{I}_R)$

$$\bar{I}_L = \frac{\sum_{i=-by}^{by} \sum_{j=-bx}^{bx} I_L(y+i, xL+j)}{(2.bx+1)(2.by+1)}$$

$$\bar{I}_R = \frac{\sum_{i=-by}^{by} \sum_{j=-bx}^{bx} I_R(y+i, xL+j)}{(2.bx+1)(2.by+1)} \quad (8)$$

This is the most robust dissimilarity function discussed in this paper. Because of the normalization in the denominator, it can deal with zone differences between left and right images, such as global intensity variations. The biggest disadvantage is its computational cost. Attempts to minimize the number of operations, result in an explosion of memory usage. With nowadays's computers, we should definitely avoid this.

3. The last dissimilarity function presented here is the one described by Birchfield [6]. This dissimilarity function is insensitive to image sampling. The latter phenomenon can significantly change the intensity value of a pixel where the intensity function is changing rapidly. Most stereo algorithms [7] [13] just work at pixel resolution, so it is important to use a dissimilarity function that eliminates errors due to sampling. It is illustrated in Figure 2.




Figure 2: Birchfield's dissimilarity measure: $I_{Rmin} \leq I_L^0 \leq I_{Rmax}$, thus D_1 equals 0.

First we calculate I_R^- and I_R^+ :

$$I_R^- \equiv \frac{1}{2} (I_R(xR) + I_R(xR-1))$$

$$I_R^0 \equiv I_R(xR)$$

$$I_R^+ \equiv \frac{1}{2} (I_R(xR) + I_R(xR+1)) \quad (9)$$

These are the linearly interpolated intensities halfway between xR and its neighbours. Let $I_{Rmin} \equiv \min(I_R^-, I_R^+, I_R^0)$ and $I_{Rmax} \equiv \max(I_R^-, I_R^+, I_R^0)$, then define

$$D_1 \equiv \max(0, I_L^0 - I_{Rmax}, I_{Rmin} - I_L^0)$$

$$D_2 \equiv \max(0, I_R^0 - I_{Lmax}, I_{Lmin} - I_R^0)$$

$$D(xL, xR) = \min(D_1, D_2) \quad (10)$$

Actually, Birchfield's dissimilarity measure D_1 is looking for the minimum distance between I_L^0 and the linearly interpolated intensity curve in an interval $[-\frac{1}{2}, \frac{1}{2}]$ around xR . Indeed, there are 3 possibilities for $I_L(xL) = I_L^0$:

- (a) $I_{Rmin} \leq I_L^0 \leq I_{Rmax}$: this is the case in Figure 2. $I_L^0 - I_{Rmax} \leq 0$ and $I_{Rmin} - I_L^0 \leq 0$, thus D_1 equals 0.
- (b) $I_L^0 \geq I_{Rmax}$: D_1 will be positive, and the dissimilarity equals the distance between the maximum and I_L^0
- (c) $I_L^0 \leq I_{Rmin}$: analogous as 2.

This dissimilarity measure performs very well, and will be used indeed in our algorithm, unless extra robustness is needed. It needs only 3 additions per pixel in an efficient implementation. Further details can be found in Birchfield's paper. [6]

Problem Definition Now, all the material is ready for a formal problem definition: Find between all match sequences that obey constraints C1..C4, the sequence with the lowest cost, according to the cost function defined above. The cost of a scanline can be calculated incrementally, i.e.




Figure 3: Stereo tree with the successors of match (0,0) and match (1,0).

one can reuse the result if one adds a match to the match sequence. In the next section, it is shown how this sequence is found, making use of dynamic programming.

3. Tree search algorithm

In this section, we describe a first solution, using a tree representation in which the best path is found making use of dynamic programming [2]. The matrix structure that is used to solve this optimization problem is also covered in this section.

Stereo Tree To describe the stereo tree, we define the nodes and edges. Each node is just a valid match, obeying constraint C1. An edge is defined by the constraints of a match sequence: between every 2 nodes that comply with constraints C2 and C3 exists an edge. It is not difficult to see that a kind of network is build between all possible matches that include the first pixel (beginnodes) and all matches that have the last pixel of a scanline (endnode).

The construction of the tree is best explained with an example (The 2 scanlines used are the same as in the introduction). Figure 3 shows such a tree, with the matches as nodes in the shaded bullets. The beginnodes are found in the upper row and the left column. Also shown are the successors of match (0,0) and match (1,0). Also remark that the nodes (xL, xR) are placed in a special configuration, namely the matrix with horizontal xL and vertical xR axis. This way, a kind of bandmatrix from upper left to down right corner is formed. This is subject of the next paragraph.

Matrix The implementation uses 4 matrices. First it is explained using indices xL, xR, as in Figure 3. Afterwards, it will be explained how a transformation of xL, xR is applied yielding Figure 4.

- Every node has its own matchcost, calculated by the dissimilarity function. Because this function can be




Figure 4: The 4 compressed matrices Π_{xL} , Π_{xR} , γ and Γ . The filling has just started.

quite expensive from computational point of view, it is calculated before and stored in matrix γ .

- Every node has one and only one predecessor, by the definition of the tree structure. This predecessor can be unambiguously characterized by its (xL, xR) numbers. That's why we need 2 other matrices Π_{xL} and Π_{xR} that store xL, respectively xR of each node.
- Because each node has only one predecessor, this node also characterizes the path from one of the beginnodes to that node. The path cost of that path is stored in a 4th matrix Γ , thus $\Gamma(xL, xR) = \kappa(\text{MS})$.

The matrices are filled top-down, the endnode with the lowest cost (Γ) is selected and trivial backtracking is performed to find the best path.

To be more memory efficient, only the diagonal band of the matrices is stored. In fact, this comes down to a transformation by which the xR axis is transformed into the δ axis. This way, implementation is far more efficient and the δ - sequence can directly be read out from the Γ matrix. The 4 compressed matrices are shown in Figure 4, when the filling has just started. Note the special pattern by which the successors of node (0,0) can be found.

4. Hierarchical implementation

The complexity of this algorithm is $O(W.H.\delta_{range}^2)$, as will be shown in next section. In this section, we develop an hierarchical implementation of the above algorithm, which will be shown to have a complexity independent of δ_{range} . We first need to introduce the concept of an offset δ - sequence.

Offset stereo tree An offset stereo tree is a normal stereo tree, but the number of allowed nodes is limited by an extra constraint:




Figure 5: Overview of the hierarchical stereo calculation.

$$C5: \delta_{smin} \leq \delta_s \leq \delta_{smax} \text{ where } \delta_s = \delta - DS_{\Delta} \quad (11)$$

This means that the nodes that are allowed must have a disparity in the interval $[\delta_{smin} + DS_{\Delta}, \delta_{smax} + DS_{\Delta}]$, around a given offset disparity. This offset disparity sequence DS_{Δ} is given. How it is calculated is subject of next paragraph.

Hierarchical calculation of the disparity map As already suggested in the introduction, this offset disparity sequence is the result of the disparity calculation at a higher (downsampled) level. We assume in this paper that downsampling is always done by a factor 2 in 2 dimensions. An example is given in Figure 5, where the same scanlines are used again. Both of them are downsampled by which the disparity search range halves. Then, our algorithm is executed with an initial offset disparity map, which is chosen all-zero.² The result (in our example, all ones, but this is coincidence) is upsampled afterwards and multiplied by 2. This δ -sequence is then used as an offset δ -sequence in our next run at the original level in a δ -range $[-1, 1]$. From next section, it will be clear that executing the algorithm twice with a small δ_{range} is cheaper than one 'big' search in the full δ_{range} .

Next, the memory requirements for the 4 matrices can be reduced further by using the relative δ_s axis instead of the absolute disparity δ . The algorithm is programmed in such way that the offset DS_{Δ} sequence is read from the disparity map, used to calculate an updated (refined) map using our 4 matrices and rewritten in the same disparity map, after being interpolated. The 4 matrices are shown in Figure 6.

One question remains unanswered: What's the largest refinement that's allowed, i.e. what's the choice of δ_{smin} ,

²If a raw disparity map is available, it can be used here as offset disparity map.




Figure 6: The 4 reduced matrices used in the hierarchical stereo calculation.

δ_{smax} ? As already suggested in [9], this depends on the disparity estimation error in each level. Experiments have shown that this error is below 1.5 pixel. In the previous upsampled level, this error corresponds to 3 pixels. That's why we choose the constants $\delta_{smin} = -3$ and $\delta_{smax} = 3$. Disparity errors up to 3 pixels will be corrected in a lower level. As a consequence $\delta_{srange} = 7$.

5. Complexity Analysis

Reduction of complexity As already discussed in Birchfield's thesis [4], the complexity Θ of the non-hierarchical algorithm is proportional to the square of the disparity search range. L represents the number of times downsampling is applied. ($L = 0$ represents the non-hierarchical algorithm.)

$$\Theta(L = 0) = O(WH\delta_{range}^2) \quad (12)$$

Downsampling once ($L = 1$) reduces the complexity to³:

$$\Theta(L = 1) = O\left(WH\delta_{srange}^2 + \frac{W}{2}\frac{H}{2}\delta_{srange}^{(1)2}\right) \quad (13)$$

where $\delta_{srange}^{(1)} = \frac{\delta_{range} - \delta_{srange}}{2} + 1$

The complexity Θ , function of L , reaches a unique global minimum for $L = L^{opt}$. This optimal number of levels is only a function of the initial disparity search range δ_{range} . The complexity at this optimal level is

$$\Theta(L = L^{opt}) = O\left(WH\left(65 - \frac{4071}{(\delta_{range} + 5)^2}\right)\right) \quad (14)$$

So, if δ_{range} is large enough, we claim that the complexity is independent of the disparity search range. This is illustrated in Figure 7 which shows a logarithmic plot of the complexity Θ versus L . The tests are performed on large images: 1404x1092. The figure compares the calculated complexity with experimentally measured runtimes for 4 values

³It suffices to search in a $\delta_{srange}^{(1)}$ that is $\frac{\delta_{srange}}{2}$ smaller than $\frac{\delta_{range}}{2}$. This disparity range will be corrected (and enlarged) while upsampling and refining the disparity map.




Figure 7: Logarithmic plot of the complexity Θ versus L .

of the disparity range $\delta_{range} = 51, 107, 219, 443$. As we can see, the theoretical curves (solid without bullets) decrease exponentially towards an almost constant complexity, where they reach a unique minimum. Practical results (dashed with bullets) show the same tendency. All these minima for a varying δ_{range} parameter, as described by Equation 14, are connected by the solid line, that reaches asymptotically $\log_{10}(65)$.

We end this section with a practical implementation for calculating the optimal level of downsampling. The real formula, a result of the analytical optimization problem, is

$$L^{opt} = \log_2 \left(\frac{(\delta_{range}^2 + 5) (\sqrt{4924} - 45)}{242} \right) \quad (15)$$

A good approximation of this curve, illustrated in Figure 8, is given by:

$$L^{opt} \approx \log_2 \left(\frac{(\delta_{range} - 2 + \delta_{srange})}{2(\delta_{srange} - 1)} \right) \quad (16)$$

This is nothing else than the result of the solution of $\delta_{range}^{(L)} = \delta_{srange}$. So we keep on downsampling as long as δ_{range} is greater than $\delta_{srange} = 7$.

6. Comparison

It is instructive to compare our stereo algorithm with other stereo algorithms. First, it is compared with the algorithm of Koch [12] and Falkenhagen [10]. They developed an algorithm that searches for correspondences by matching blocks (for instance NCC). Local boundary conditions are taken into account, but no global cost function is minimized. Interpolation of disparities is non-existent that way. To increase the robustness of the local optimisation problem, and




Figure 8: Interpretation of the optimal level of downsampling: approximation of the theoretic curve.

to reduce the corona effect, one estimates the disparity in a large region (large robust block comparison), after which these disparities are refined with smaller blocks. Although this kind of hierarchy is comparable with the one used in our algorithm, it arises from a fundamental different reason, i.e. a lack of robustness of the local optimisation problem. It is important to see that the acceleration of our algorithm, under certain conditions proven by Birchfield [?], has no effect on the outcome! The complexity of the algorithm described in [10] is $O(W.H.\delta_{range}^2)$.

Secondly, our algorithm is compared with Birchfield's algorithm [4]. A modified version is implemented in Intel's OpenCV. The core of this algorithm is very similar to our algorithm, but there is a major difference in how the acceleration is performed. Birchfield's fast algorithm (complexity $O(W.H.\delta_{range} \log(\delta_{range}))$) developed a method to prune bad nodes in the search tree. The disadvantage is that the exact minimum (best path) is not longer guaranteed, and to achieve an acceptable minimum (path), extra constraints are put on the cost function. It can no longer be chosen arbitrarily, and Birchfield had to postpone several takes to a postprocessing step. For instance, in Birchfield's algorithm, propagation of disparities between scanlines, intensity gradient information, processing occlusions happens in the postprocessing. To make all this happen, a lot of postprocessing parameters are used without rules to tune them.

Finally, a comparison is made with the algorithm of S. Roy and I. Cox [8]. A complete 2D optimisation problem is solved, making use of mincut-maxflow problem. The algorithm offers very high quality disparity maps. The disadvantage is its complexity $O(W.H.\delta_{range}^2 \cdot \log(W\delta_{range}))$.

7. Results

A first testcase includes a set of images of the 'Arenbergkasteel'. These two images (792x635) can be found at the top of Figure 9. It is a good example to show the accuracy of the disparitymap. There are lots of depth discontinuities and lots of details. The hierarchical algorithm achieves sharp edges (roof, crenels, ...) and a good overall quality. The total runtime is 16 seconds.⁴




Figure 9: Images from the 'Arenbergkasteel' Original left and right image are shown at the top. The disparity map is presented at the bottom.

The second example is provided by ESA (European Space Agency). Figure 10 shows images of Mars' surface. This are large images (1404x1092) with large disparity ranges (500 and more). This was the actual reason to search for a new algorithm which resulted in the hierarchical algorithm, since most existing algorithms failed. This is a good example to illustrate the memory needs and speed of our algorithm. The images were downsampled 4 times, and the total processing time did not exceed 70 seconds. The amount of memory used was 202 kbyte. Because of imperfections of the cameras, yielding an intensity profile in the images, the more expensive NCC dissimilarity function is used.

The next image is provided by the University of Tsukuba

⁴The tests were performed on an intel PentiumIII workstation 700MHz




Figure 10: Mars images. Original left and right image are shown at the top. The disparity map is presented in the middle. At the bottom, a stone at Mars' surface and the corresponding part of the disparity map is shown.

as a dataset for testing accuracy. The depth map ground truth is also given in Figure 11. A lot of details, including the camera in the background, are visible in the disparitymap. The images (384x288) are not downsampled during the execution, that took 4 seconds of processing time.

The last image is provided by the Microsoft Research, again as a dataset for testing accuracy. The depth map ground truth is also given in Figure 12. The black spots



Figure 11: Tsukuba dataset for testing accuracy. Original left and right image are shown at the top. The disparity map is presented in the middle. The ground truth can be found at the bottom.

in the disparitymap represent occlusions. In previous images, they were interpolated. Because of the large occlusion present in the image, this is not done here. The images (284x216) are downsampled only once during execution that took only 2 seconds of processing time.




Figure 12: Slanted dataset for testing accuracy. Original left and right image are shown at the top. The disparity map is presented in the middle. The ground truth can be found at the bottom.

8. Conclusion

The algorithm presented in this paper is based on the hierarchical solution of the stereo correspondence problem, using subsampling. The complexity is $O(W.H)$, which is independent of the disparityrange and lower than other classical stereo algorithms which have a complexity of $O(W.H.\delta_{range}^2)$. A careful implementation yields minimum memory usage, so that the computation fits in the

cache of most modern processors. The disparity maps show a high quality, thanks to a new and more complex cost function. This cost function is completely decoupled from the optimisation problem, so it is easy to switch cost functions to make a tradeoff between quality and execution time.

References

- [1] U. Dhond and J. Aggarwal, Structure from Stereo - A Review, *IEEE Trans. Syst., Man and Cybern.* 19, 1489-1510, 1989.
- [2] Y. Ohta and T. Kanade, Stereo by Intra- and Inter-scanline Search Using Dynamic Programming, *IEEE Trans. on Pattern Analysis and Machine Intelligence* 7(2), 139-154, 1985.
- [3] M. Okutomi and T. Kanade, A Locally Adaptive Window for Signal Processing, *International Journal of Computer Vision*, 7, 143-162, 1992.
- [4] Stan Birchfield, Depth and Motion Discontinuities, PhD thesis, Dept. of Electrical Engineering, Stanford University, jun 1999.
- [5] Stan Birchfield and Carlo Tomasi, Depth and Motion Discontinuities, *International Journal of Computer Vision*, Vol. 35, No. 3, pg 269-293, dec 1999.
- [6] Stan Birchfield and Carlo Tomasi, A Pixel Dissimilarity Measure that is Insensitive to Image Sampling, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 20, No. 4, pg 401-406, apr 1998.
- [7] Cox, I., Hingorani, S., Rao, S., 1996, A Maximum Likelihood Stereo Algorithm, *Computer Vision and Image Understanding*, Vol. 63, No. 3.
- [8] S. Roy, I. Cox, 1998, A Maximum-Flow Formulation of the N-camera Stereo Correspondence Problem, *Proceedings of the international Conference on Computer Vision*, Bombay, India, 1998
- [9] Falkenhagen, L., 1997, Hierarchical Block-Based Disparity Estimation Considering Neighbourhood Constraints. *Proceedings International Workshop on SNHC and 3D Imaging*, Rhodes, Greece, pp.115-122.
- [10] L. Falkenhagen, 1997, Depth Estimation from Stereoscopic Image Pairs assuming Piecwise continuous surfaces in: Paker, Y. and Wilbur, S. (Ed.), *Image Processing for Broadcast and Video Production*, ISBN 3-540-19947-0, pp. 115-127, Springer Great Britain, 1994.
- [11] R. Hartley, R. Gupta, and T. Chang. "Stereo from uncalibrated cameras". *Proc. Conference Computer Vision and Pattern Recognition*, pp. 761-764, 1992.
- [12] Koch, R., 1996, Automatische Oberflächenmodellierung starrer dreidimensionaler Objekte aus stereoskopischen Rundum-Ansichten, PhD thesis, University of Hannover, Germany, also published as *Fortschritte-Berichte VDI*, Reihe 10, Nr.499, VDI Verlag, 1997.
- [13] S.S. Intille and A.F. Bobick. Disparity-Space Images and Large Occlusion Stereo. *Proc. Third European Conference Computer Vision*, pp. 179-186, 1994.