

A Semi-Local Method for Iterative Depth-Map Refinement

David M^cKinnon, Ryan N. Smith and Ben Upcroft

Abstract—Building a photorealistic, 3D model of an object or a complete scene from image-based methods is a fundamental problem in computer vision, and has many applications in robotic perception, navigation, exploration and mapping. In this paper, we extend current state-of-the-art in the computation of depth maps by presenting an accurate and computationally efficient iterative hierarchical algorithm for multi-view stereo. The algorithm is designed to utilise all available contextual information to compute highly-accurate and robust depth maps by iteratively examining different image resolutions in an image-pyramid. The novelty in our approach is that we are able to incrementally improve the depth fidelity as the algorithm progresses through the image pyramid by utilising a local method. This is achieved in a computationally efficient manner by simultaneously enforcing the consistency of the depth-map by continual comparison with neighbouring depth-maps. We present a detailed description of the algorithm, and describe how each step is carried out. The proposed technique is used to analyse multi-view stereo data from two well-known, standard datasets, and presented results show a significant decrease in computation time, as well as an increase in overall accuracy of the computed depth maps.

I. INTRODUCTION

Multiple robotic platforms are currently able to gather large-scale datasets at unprecedented spatiotemporal resolutions in both indoor and outdoor settings with the purpose to explore, map, reconstruct and ultimately understand the surrounding environment. Existing methods for such tasks rely on laser scanners and/or vision sensors, i.e., cameras to generate dense data sets from which to extract information. Laser scanners provide three-dimensional (3D) measurements in real-time to the robot, but are expensive and bulky. Coupling this with an increasing availability of low-cost, high-resolution vision sensors motivates our investigation into image-based 3D reconstruction.

Constructing a complete and photorealistic 3D model of an object or scene has been a topic of much interest for almost two decades, and is still an area of active research in computer vision, image-based modelling and computer graphics. In recent years, multiple, high-quality algorithms have been developed, and the state-of-the-art is improving rapidly. Additionally, with significant increases in computational power on smaller and smaller devices, image-based, 3D reconstruction is becoming accessible to the robotics community for use in a wide variety of applications.

Here, we present an iterative hierarchical method for multi-view stereo which provides rapid and accurate depth-




Fig. 1. Example image (left) and associated depth-map (right) for the fountain data set [4]. This depth-map was generated using our iterative refinement method.

maps for image-based 3D reconstruction. There are two novel aspects to our approach: 1) incrementally improved depth fidelity as the algorithm progresses through an image pyramid; and 2) simultaneous enforcement of the consistency of the depth-map by continual comparison with neighbouring depth-maps. The proposed method demonstrates significant improvements to the accuracy of the depth-maps, memory usage, and computation time of existing techniques, and is conceivable to implement onto robotic platforms to perform image-based 3D reconstruction in near-real time, e.g., see [1].

The method presented here is an alternative to patch-based methods such as [2], but can still be utilized to provide a dense set of oriented points. Furthermore, our method does not employ any form of global optimisation, but instead integrates neighbourhood pixel information (similar to [3]) to guide the solutions. We demonstrate that uncertainty associated with assigning depths to pixels can be circumvented by an iterative search, and show that this is better suited to identifying the local minima of the photo-consistency function. The proposed approach results in highly-accurate depth-maps with low error rates and superior performance to existing state-of-the-art schemes at a much lower computational cost. An example depth-map calculated using this scheme is shown in Fig. 1.

II. BACKGROUND

Existing stereo methods can be classed in two categories. First, the local method which calculates the depth for each pixel of the image independently, not taking into account depths allocated to surrounding pixels. Local methods are fast and memory efficient, but are prone to incorrectly allocating depths. Second, the global method which attempts to relate the choice of depth for a given pixel with the depth choices for those pixels immediately surrounding it. In this sense, the local continuity of a surface is taken into

D. M^cKinnon is with IPPLEX, Suite 4 2043 Colorado Ave., Santa Monica, CA 90404 USA david@ipplex.com

R.N. Smith and B. Upcroft are with the School of Engineering Systems at the Queensland University of Technology, Brisbane, QLD, 4000 Australia. ryan.smith,ben.upcroft@qut.edu.au

consideration when determining each and every depth and consequently the algorithm has a global perspective. Global methods are known to provide superior estimates of depth but at the expense of being orders of magnitude times slower and more memory intensive to compute than local methods.

Connectivity and continuity constraints for depth estimates across an image were first exploited by Zitnick and Kanade [5] and Zhang and Shan [6]. Hierarchical (or coarse-to-fine) algorithms have also been a part of some of the very early approaches to stereo [7]. The algorithm described by Cornelis and van Gool [3] appears to be the first approach to employ the idea of connectivity constraints within the hierarchical refinement process.

The authors in [3] showed that by swapping depth estimates with neighbouring pixel values with improved photo-consistency scores, highly accurate local methods could be achieved without adversely affecting the computation time. We term this sort of approach a *semi-local method*, since it involves some introspection from surrounding pixels, without a computationally-intensive global minimisation. Cornelis and van Gool also show that aliasing effects, that are commonly observed with hierarchical implementations, can be minimised through appropriate filtering and by employing connectivity constraints throughout the hierarchical levels. We will draw most heavily from this approach and adapt it to refine the resolution of the depth estimates at each level.

We also exploit the consistency of depth estimates for a scene element between two or more depth-maps. This information is commonly used in two-view stereo algorithms in the form of the left-right consistency check and more recently in multi-view stereo algorithms such as [8], [9], [10]. To avoid ambiguity with photo-consistency, we term this *depth-consistency*. We show that this information can be integrated progressively into the hierarchical depth-map iterations (to make best use of the redundancy of the scene measurements) to filter erroneous depth estimates as they emerge.

Campbell et al. showed that often the location of the global minimum of the photo-consistency score does not correlate with the true value for the depth of a pixel [11]. We show that by simply increasing the number of locations that the photo-consistency score is evaluated at, we can obtain more reliable and highly accurate measures of a pixel's depth. Figure 2 illustrates this concept, where the true location of the minima of the photo-consistency score do not coincide with the location of the discrete depth evaluations.

Clearly, increasing the number of depth evaluations for a stereo problem will result in a considerably higher cost for the algorithm. To counter this, we propose an iterative hierarchical scheme, which proves to be extremely efficient, particularly for large image sizes. The iterative scheme we employ attempts to integrate local assumptions of depth continuity (similar to a global method) in addition to the mutual consistency of depth estimates between images. The resulting depth-map has (effectively) a continuous depth resolution as opposed to the discrete resolution of most of the prior state-of-the-art, plane-sweeping based methods [3],




Fig. 2. Example of the photo-consistency cost function for a pixel. The locations of the discrete depth evaluations are indicated (bold vertical lines), as well as the global minimum, refined (parabolic) update and the initial (discrete) depth evaluation.

[12], [13].

III. ALGORITHM

We begin by establishing some of the notation, and defining the variables that will be used throughout the paper. The goal of the proposed algorithm is to produce a depth-map \mathbf{D}_b for the base image \mathbf{I}_b using M comparison images and their associated camera projection matrices $\mathbf{Q}_q = [\mathbf{N}_q | \mathbf{n}_q]$ (for $q = 1 \dots M$).

Assume the base image has a camera projection matrix $\mathbf{P}_b = [\mathbf{M} | \mathbf{m}]$ that projects a 3D homogeneous point \mathbf{X} into its 2D homogeneous location in the base image \mathbf{x}_b (i.e., $\mathbf{x}_b = \mathbf{P}_b \mathbf{X}$) and the camera is located at \mathbf{C} where $\mathbf{P}_b \mathbf{C} = \mathbf{0}$. The reprojection of a point in a comparison image \mathbf{x}_q from a point (\mathbf{x}_b) in the base image with a depth d is

$$\mathbf{x}_q = d \mathbf{N}_q \mathbf{M}^{-1} \mathbf{x}_b + \mathbf{N}_q \mathbf{C}. \quad (1)$$

We assume that there exists a photo-consistency metric that is used to evaluate the cost $\mathcal{R}_q^d \in [0, 1]$ between a patch of pixels \mathcal{P}_b from a location in the base image \mathbf{x}_b to the corresponding patch \mathcal{Q}_q^d in a comparison image at a given depth d . This patch is determined by a direct application of (1). The photo-consistency measure is normalised so that a value $\mathcal{R}_q^d = 0$ indicates similarity and $\mathcal{R}_q^d = 1$ dissimilarity. We use the Normalised Cross-Correlation (NCC) measure exclusively (with a window size $[2N + 1 \times 2N + 1]$), although there are numerous other alternatives to this that could also be used.

A. Depth Refinement

The presented algorithm operates over an image-pyramid of k_l levels. Each subsequent level of the image-pyramid (k) is a power of 2 reduction from the previous level. The first level of the pyramid is chosen as an image size with dimension no smaller than 128 pixels. The resolution is increased from this point to the desired extent.

The first stage in the hierarchical depth estimation process is determining a suitable initial depth-map \mathbf{D}_b^0 for the smallest level of the image pyramid. Recall that the first level is orders of magnitude smaller than the original image, and consequentially the depth-map corresponding to this can be calculated rapidly. For this depth-map initialisation, we evaluated the use of global methods; the Semi-Global Matching (SGM) method [13], and the Winner Take All

(WTA) plane-sweeping approach. It was found that WTA produced sufficiently accurate initial estimates, without the need to introduce the additional complexity (and superior accuracy) of SGM.

The initial depth-map for the first level can now be used as a guide in determining the depth maps for images of increasing pixel and depth resolution higher in the image pyramid. As discussed above, the iterative algorithm we present attempts to search the depths of nearby pixels in the image to employ a depth connectivity constraint. We assume that neighbouring pixels will not have significant differences in depth, and these are good candidates as starting points for a search. Due to the erratic nature of the photo-consistency function (over the depths of a pixel) the chances of a discrete plane-sweep locating a minimum are small. For this reason, we search depths at a finer increment (to that of the plane-sweep) to better ascertain the local nature of the photo-consistency function around a given depth (see Fig. 2).

At each iteration (n), a select group of candidate depths (C_p) are sampled from the surrounding pixels, note that the candidate depth set does not include the current depth estimate for the pixel. The patterns used to isolate the candidate depths are the same as those given in [3]. The implementation alternates between these candidate depth patterns at each iteration. For each of the candidate depths the surrounding depths C_p^+ and C_p^- (generated by adding and subtracting the depth increment $\frac{1}{d_{inc}^k}$ from the original value C_p) are evaluated and the corresponding photo-consistency scores \mathcal{R}_p^+ , \mathcal{R}_p and \mathcal{R}_p^- determined.

If the smallest of the resulting photo-consistency scores is \mathcal{R}_p , then a parabolic interpolation using the neighbouring scores and depth values is used to accurately estimate the corresponding depth minimum. If either depth corresponds to the minimum, then it is selected as the updated value for the pixel depth. This iterative process is repeated n_l times on each level of the pyramid. We have observed experimentally that this process of searching neighbouring pixels for improved depth estimates results in the ability of the algorithm to jump-out of local minima of the photo-consistency score, and dramatic increases in the depth fidelity for the pixels (via the refinement).

The value for d_{inc}^k can be set at each pyramid level, and governs the step size for each new depth candidate test. Larger values of the d_{inc}^k increase the rate of convergence to a solution, but may over shoot the true depth value. Smaller values increase computation time and ensure that search proceeds more cautiously, thus increasing the likelihood that a precise local minimum is found. It has been demonstrated experimentally that larger values of d_{inc}^k perform better on more oblique surfaces where the depth gradient is high. This implies that the choice in tuning this parameter may be data driven.

B. Depth-Consistency Checking

Campbell *et al.* have shown that depth-consistency checking can be successfully used with multi-view stereo pipelines to reduce the occurrences of bad depth estimates and also

Algorithm 1: Iterative Depth-Map Refinement

Input : Given a base image I_b , M accompanying images Q_q and their camera matrices P_b and P_i .

Output: Refined estimate of D_b^0 .

```

for  $k = 0$  to  $k_l - 1$  do
  for  $n = 1$  to  $n_l$  do
    foreach Pixel  $x_b$  of  $D_b^{k,n}$  do
      if  $X \cap V = \emptyset$  then continue
      Sample the patch  $P_b$  at this position
      Sample a set of surrounding depth
      candidates  $C_p$ 
      foreach Depth candidate  $C_p$  do
        for  $q = 1$  to  $M$  do
          Sample the patches  $Q_q^+$ ,  $Q_q$  and  $Q_q^-$ 
          Accumulate scores  $\mathcal{R}_p^+$ ,  $\mathcal{R}_p$  and  $\mathcal{R}_p^-$ 
          if  $\mathcal{R}_p < \mathcal{R}_p^+$  and  $\mathcal{R}_p < \mathcal{R}_p^-$  then
             $\bar{C}_p = \text{interp}(C_p^+, C_p, C_p^-)$ 
          else if  $\mathcal{R}_p^+ < \mathcal{R}_p$  and  $\mathcal{R}_p^+ < \mathcal{R}_p^-$  then
             $\bar{C}_p = C_p^+$ 
          else  $\bar{C}_p = C_p^-$ 
        Update  $D_b^{k,n}$  with the value of  $\bar{C}_p$ 
        corresponding to the lowest score
      Perform consistency-checking on  $D_b^{k,n}$  by
      comparing to neighbouring views  $D_q^{k,n}$ 

```

as a means to regularise good depth estimates (filtering estimates from neighbouring views) [11]. They also noted that regularising during the calculation of depth has a negative effect on the resulting depth-map accuracy. Instead of attempting to alter the value of a depth estimate, we employ a process of consistency checking strictly to remove the instance of spurious depth estimates, and those that are not well supported by the surrounding views. This technique is similarly to the confidence-based fusion approach of Merrell *et al.* [8].

Previous authors have employed depth-consistency checking as a type of post-process to clean and enhance depth-maps upon completion of the stereo process. By utilising a neighbouring view's depth-maps, we have access to important information that can be used during the depth refinement process. We are able to eliminate spurious and unsubstantiated depth estimates propagating through to the final depth-map. The depth-consistency checking is inserted as a final stage at each pyramid level of the depth-map refinement. For this to be achieved, all neighbouring view's depth-maps must be refined simultaneously and then compared to each other at each level.

IV. IMPLEMENTATION

To take full advantage of the parallel nature of the algorithms proposed in the previous section, they were implemented on a Graphics Processing Unit (GPU). The



Fig. 3. (left) Example image of the rendered ground-truth depth-map for the *fountain* sequence. (right) An example of the error statistics gathered during the experiments, green denotes no ground-truth measure, red a gross error ($> 10\sigma$), white no error increasing towards black (10σ).

OpenGL Shading Language (GLSL) was chosen to allow the programs to run on all major GPU hardware types (including OpenGL ES platforms). Graphics hardware is particularly well suited for stereo algorithms, due to its built in capacity for texture sampling and parallelism. The refinement process is implemented as a set of fragment shaders similarly to [12], [3].

V. RESULTS

A. Ground-Truth Analysis

To evaluate the effectiveness of the presented scheme we performed a series of experiments with a publicly available ground-truth dataset. The experiments were designed to determine the performance of the presented algorithm over a range of different parameter values and settings.

The ground-truth dataset is the *fountain* dataset from [4]. This dataset provides access to both the calibrated camera locations and images (11 images of resolution 3072×2048 which we downsampled by a factor of 2 to fit the image textures in the graphics card memory), in addition to LIDAR data (from a Zoller+Fröhlich IMAGER 5003) aligned to the same coordinate frame. The experiments use the ground-truth LIDAR data to create the corresponding ground-truth depth-maps, which in turn are used to evaluate the algorithm.

As we are primarily interested in the accuracy of the algorithm, we only gather statistics for pixels where the depth-refinement process and the LIDAR provide an estimate. Generally, the proportion of pixels that result in a depth estimate is quite high ($> 90\%$), with the discarded pixels often occluded in the comparison image or associated with homogeneous regions where reliable depth determination is not possible using a small correlation window. An example of the ground-truth depth-map and associated error image for the *fountain* dataset is shown in Fig. 3.

The experiments conducted consist of a set of trials adjusting the number of levels of hierarchical refinement, with and without depth-consistency checking (Fig. 4a.), trials sweeping the value of the parameter d_{inc} (Fig. 4b.), and trials adjusting the number of comparison views M (Fig. 4c.). A further set of trials was performed comparing the results of the algorithm to WTA, SGM and Cornelis' algorithms (Fig. 4d.). Unless otherwise noted, the parameters for each of the trials were set to $M = 2$, photo-consistency threshold 0.5, $d_{inc} = 2048$, $N = 1$, $n_l = 3$, $k_l = 4$ and consistency




Fig. 6. Renderings of the templeSR (left) and herzjesu (right) reconstructions.

checking activated. The variable σ in the trials is equivalent to 3mm.

The results show that our proposed approach achieves extremely accurate depth-maps when compared to the other methods (Fig. 4d.). Furthermore, depth-consistency checking and higher levels of refinement progressively improve the accuracy (and resolution of the depths) (Fig. 4a.). The addition of more comparison views (Fig. 4c.) and higher depth levels (Fig. 4b.) also lead to improved results.

B. Middlebury Analysis

To further test the performance of the algorithm, we selected several of the image sets from the Middlebury 2006 Stereo Evaluation page [14]. As the approach is a plane-sweeping algorithm designed to calculate depth-maps from multiple camera locations and not disparities from stereo pairs, we were not able to produce the disparity fields suitable for submission to the evaluation page. Instead, we rely on standard Structure From Motion methods [15] to auto-calibrate a selection of image sets and provide the approximate camera locations and depth-maps. Once calibrated, the image sets are used for a relative comparison between our (semi-local) method and the (global) SGM method. The results of this comparison are given in Fig. 5.

The results indicate that the algorithm consistently produces visually accurate depth-maps, consistent with depth-maps produced by SGM but at a much higher depth resolution (note the discretisation for SGM on *aloe*), in agreement with the results in Fig. 4d.

C. 3-D Reconstruction

To give an impression of how the resulting depth-maps affect the accuracy of a image-based 3D reconstruction, we present our algorithm used in conjunction with Poisson Surface estimation [16] to estimate a polygon model of objects from the *templeSR* and *Herzjesu* datasets [4] (Fig. 6). A simple approach to extract oriented points from the depth-maps resulting from our algorithm has been used.

D. Computation Time

All experiments were performed on an Intel Quad Core (2.66GHz) machine with a NVidia GTX 480 GPU. Table I gives the average computation time for this system to process depth-maps from the *fountain* dataset under different

parameter settings (a considerable portion ~ 1000 ms of this time is for image loading and transfer to the GPU). Cornelis, WTA and SGM were all timed using 128 depth levels.

k_l	M				
	w x h	1	2	3	4
0	192x128	1134	1127	1159	1249
1	384x256	1147	1267	1292	1421
2	768x512	1355	1637	1822	1995
3	1536x1024	2242	3064	3676	4143
WTA	1536x1024	4472	5053	6206	7065
SGM	1536x1024	64443	60724	62229	74575
Cor	1536x1024	1337	1448	1519	1617

TABLE I

AVERAGE TIMINGS (IN MILLISECONDS) FOR OUR SEMI-LOCAL METHOD FOR DIFFERENT NUMBERS OF PYRAMIDAL LEVELS (k_l), AND FOR VARYING NUMBER OF COMPARISON IMAGES (M). COMPARISONS ARE ALSO GIVEN FOR THE LOCAL METHOD WTA, THE GLOBAL METHOD SGM, AND CORNELIS' (COR) ALGORITHM ALL WITH $k_l = 3$.

VI. DISCUSSION

We have presented a local stereo algorithm capable of producing very accurate depth estimates at a relatively low computational cost. The algorithm consistently makes the best use of all available information to speed-up and guide the creation of depth-maps.

The novelties in this approach include: 1) computing all depth-maps in parallel allowing for the integration of depth-consistency checking with neighbouring views at each level of the refinement process; 2) the use of depth searches at progressively finer increments seeded by neighbouring depth pixels to jump out of and avoid local minima in the photo-consistency score; and 3) the real-time implementation for highly accurate depth map computation suitable for mobile robotics applications in outdoor environments.

These approaches are simple to implement and were shown to provide significant improvements to the accuracy of the depth-maps, memory usage, and computation time of the algorithm (due to their parallel/local nature).

Future work for this study is focused on the generation of high-quality depth maps for the purpose of 3D reconstruction of objects and scenes. We will consider a two-tiered approach to this effort. First, we have investigated the use of the proposed algorithm to reconstruct objects in near-real time from robotic platforms [1]. This is an initial study into the ability and accuracy of our methods in aquatic environments. Secondly, we plan to exploit the parallelised nature of the algorithm and further demonstrate speed and accuracy benchmarks for the algorithm in the generation of 3D models of objects and scenes.

REFERENCES

[1] D. McKinnon, B. Upcroft, and R. N. Smith, "Towards automated and in-situ, near-real time 3-d reconstruction of coral reef environments," in *Proceedings of the MTS/IEEE Oceans '11 - Kona*, (Kona, Hawaii), September 2011. Accepted, to appear.

[2] Y. Furukawa and J. Ponce, "Accurate, dense, and robust multiview stereopsis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, pp. 1362–1376, 2010.

[3] N. Cornelis and L. V. Gool, "Real-time connectivity constrained depth map computation using programmable graphics hardware," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 1, pp. 1099–1104, 2005.

[4] C. Strecha, W. von Hansen, L. Van Gool, P. Fua, and U. Thoennessen, "On benchmarking camera calibration and multi-view stereo for high resolution imagery," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1 –8, 2008.

[5] C. Zitnick and T. Kanade, "A cooperative algorithm for stereo matching and occlusion detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 675 – 684, July 2000.

[6] Z. Zhang and Y. Shan, "A progressive scheme for stereo matching," *Lecture Notes in Computer Science*, vol. 2018, pp. 68 – 85, 2001.

[7] L. Falkenhagen, "Hierarchical block-based disparity estimation considering neighbourhood constraints," in *International Workshop on Synthetic-Natural Hybrid Coding and Three-Dimensional Imaging*, 1997.

[8] P. Merrell, A. Akbarzadeh, L. Wang, P. Mordohai, J.-M. Frahm, R. Yang, D. Nister, and M. Pollefeys, "Real-time visibility-based fusion of depth maps," in *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pp. 1 –8, 2007.

[9] C. Unger, E. Wahl, P. Sturm, and S. Ilic, "Probabilistic disparity fusion for real-time motion-stereo," in *Technical Report TU Munich*, 2010.

[10] G. Zhang, J. Jia, T.-T. Wong, and H. Bao, "Consistent depth maps recovery from a video sequence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, pp. 974–988, 2009.

[11] N. D. Campbell, G. Vogiatzis, C. Hernández, and R. Cipolla, "Using multiple hypotheses to improve depth-maps for multi-view stereo," in *Proceedings of the 10th European Conference on Computer Vision: Part I, ECCV '08*, (Berlin, Heidelberg), pp. 766–779, Springer-Verlag, 2008.

[12] R. Yang and M. Pollefeys, "Multi-resolution real-time stereo on commodity graphics hardware," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 1, p. 211, 2003.

[13] H. Hirschmuller, "Accurate and efficient stereo processing by semi-global matching and mutual information," in *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2 - Volume 02*, CVPR '05, (Washington, DC, USA), pp. 807–814, IEEE Computer Society, 2005.

[14] H. Hirschmuller and D. Scharstein, "Evaluation of cost functions for stereo matching," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 0, pp. 1–8, 2007.

[15] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

[16] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Eurographics Symposium on Geometry Processing*, pp. 61–70, 2006.

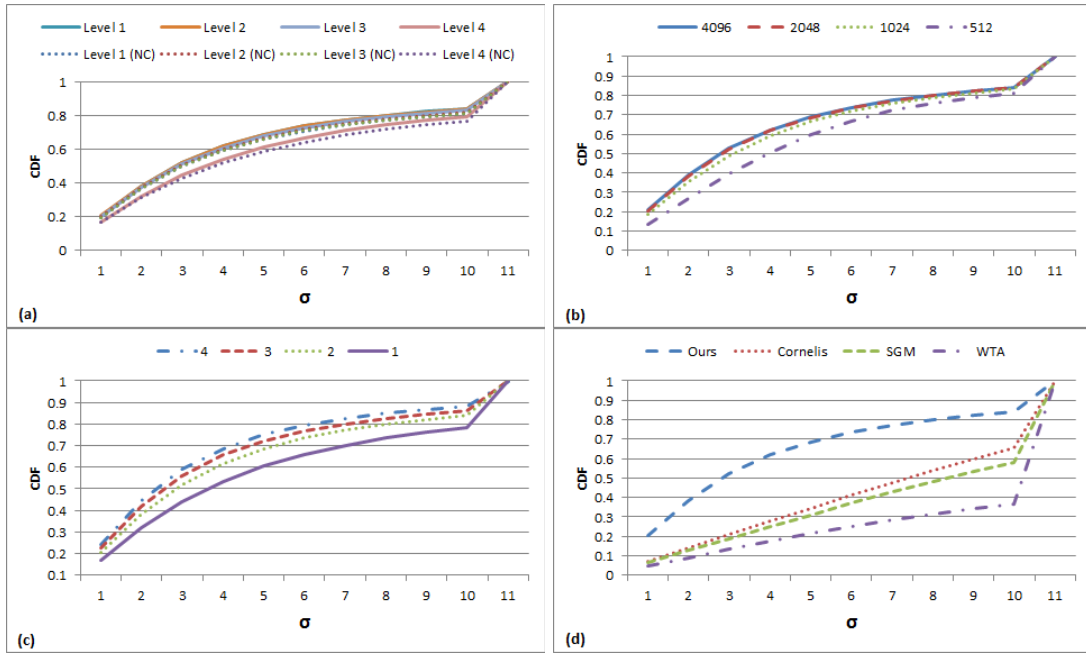


Fig. 4. Cumulative Distribution Function of estimates where the error is $< n\sigma$. (a) Adjusting the number of hierarchical levels (k_l) with and without (NC) depth-consistency (dashed). (b) Sweeping d_{inc} . (c) Adjusting the number of comparison views M . (d) Comparison with WTA, SGM and Cornellis.




Fig. 5. Results for the Middlebury 2006 Stereo Datasets. Top Row - original images for the *cloth4*, *wood2*, *rocks1* and *aloe* datasets. Middle Row - results from our algorithm (red pixels represent no measurement). Bottom Row - Results from Semi-Global Matching.