Resolution Enhancement of PMD Range Maps

A.N. Rajagopalan[1], Arnav Bhavsar[1], Frank Wallhoff[2], and Gerhard Rigoll[2]

[1] Department of Electrical Engineering
Indian Institute of Technology Madras, Chennai 600 036, India
[2] Lehrstuhl für Mensch-Maschine-Kommunikation
Technische Universität München, 80333 München, Germany
{raju,ee04s036}@ee.iitm.ac.in, Wallhoff@ei.tum.de, Rigoll@tum.de

Abstract. Photonic mixer device (PMD) range cameras are becoming popular as an alternative to algorithmic 3D reconstruction but their main drawbacks are low-resolution (LR) and noise. Recently, some interesting works have stressed on resolution enhancement of PMD range data. These works use high-resolution (HR) CCD images or stereo pairs. But such a system requires complex setup and camera calibration. In contrast, we propose a super-resolution method through induced camera motion to create a HR range image from multiple LR range images. We follow a Bayesian framework by modeling the original HR range as a Markov random field (MRF). To handle discontinuities, we propose the use of an edge-adaptive MRF prior. Since such a prior renders the energy function non-convex, we minimize it by graduated non-convexity.

1 Introduction

In contemporary computer vision and multimedia, numerous applications harness the information in the 3D shape of a scene. Apart from the conventional *shape from X* techniques, in recent years, direct acquisition of 3D shape from range sensors has gained importance. Laser range scanners produce high quality range maps but their use is currently limited by high cost and long acquisition times [1]. Photonic mixer device (PMD) scanners are less accurate and have a depth resolution of approximately 6 mm. However, they are attractive due to their cost, speed and simplicity. They are being used in many applications in pattern recognition, computer vision and multimedia [2,3,4].

PMD range scanners work on the basis of time-of-flight principle. The complete scene is illuminated using modulated infra-red waves and the reflected light is received by the PMD sensor. The output signal of the sensor is related to the phase delay between the reflected optical signal and the modulation signals. The electro-optic mixing process occurs at every pixel and hence the phase delay measurement is carried out at every pixel in the PMD pixel array [4,5,6]. Applications involving 3D visualization, image based rendering, augmented reality etc. require an overall scene representation that differentiates major objects. Similarly, an approximate range map is usually sufficient in applications related to scene segmentation, object recognition, video matting, surveillance and robot navigation [7].

G. Rigoll (Ed.): DAGM 2008, LNCS 5096, pp. 304–313, 2008.

A major drawback of PMD scanners is their low *spatial* range resolution which is 64×48 or at the most 160×120 [3,6] pixels. The resultant poor localization in 3D space inherently restricts the eventual utility of the PMD scanner. Recently, many interesting works have addressed the issue of enhancing the spatial resolution of PMD range data so that its scope can be extended to a multitude of applications. Prasad et al. [8] interpolate the range map and register it with a HR CCD image. However, this smooths the range data and causes loss of high frequencies especially at object discontinuities. An MRF-based energy minimization framework that uses a range map and a HR CCD image claims better performance at discontinuities [9]. Based on the assumption that discontinuities in range and image tend to coincide, this approach weights the smoothness terms in the energy function by a gradient measure of the HR intensity image. A fusion of LR PMD range images with a stereo algorithm to produce HR depth maps is proposed in [3], where the data term of the energy function comes from the HR stereo pairs as well as the LR range. The above mentioned works typically use a registered CCD HR image or a CCD HR stereo pair. However, this necessitates the requirement of an elaborate setup and an involved calibration procedure for both the CCD and the PMD cameras.

In this work, we propose to principally exploit translational camera motion (relative to the object) for super-resolution of PMD range data. Importantly, our approach uses only the PMD scanner. Camera motion results in multiple LR images with relative sub-pixel shifts thus effectively yielding a higher sampling rate. These LR range images can be modeled to have been formed by down-sampling sub-pixel shifted versions of the HR range image that is to be estimated. The use of multiple images also enhances the ability to reduce noise. Our technique requires only image-to-image registration which is much simpler than the complicated calibration involving CCD cameras followed in other works. Shifted range images of a 3D scene will ideally tend to produce motion parallax. However, since we capture videos by simple camera translation, the assumption of global shifts is valid for consecutive or near consecutive LR frames. These global shifts can be determined by any good sub-pixel registration technique [10].

The problem of estimating high-resolution range data from low-resolution range observations is basically ill-posed and prior constraints must be imposed to enforce regularization. We model the original HR range image by a discontinuity-adaptive MRF (DAMRF) prior. This model not only regularizes the solution but also enables depth discontinuities to be preserved effectively. The solution that we seek is the maximum a posteriori (MAP) estimate of the HR range image given the LR range images. The DAMRF prior renders the resultant cost function non-convex. To avoid local minima problems, we use graduated non-convexity (GNC) to arrive at the MAP estimate of the super-resolved range image.

Section 2 explains the relationship between HR and LR range data. In section 3, we propose a MAP-MRF framework to solve for the HR range data. Section 4 discusses optimization using the GNC algorithm tailored to our problem. This is followed by results and conclusions in sections 5 and 6, respectively.

2 LR-HR Relationship of Range Data

We move the camera relative to the object and capture several frames of low-resolution range data. Suppose we have N relatively shifted low-resolution data $[y_1, y_2, ..., y_N]$ of size $N_1 \times N_2$ from the PMD range scanner. These LR observations can be modeled to have been generated from a high-resolution range image x of size $L_1 \times L_2$ by warping followed by down-sampling (by a factor of $\frac{L_1}{N_1} \times \frac{L_2}{N_2}$). Down-sampling is caused by averaging of $\frac{L_1}{N_1} \times \frac{L_2}{N_2}$ pixels. In Fig. 1, we show an example of down-sampling by 2 along both the spatial dimensions. Each pixel of LR1 (the first range image) is formed by averaging of four pixels

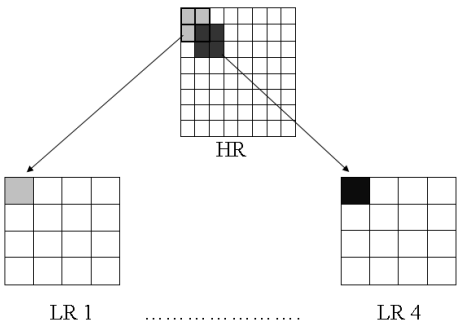

Fig. 1. The formation of LR observations from the HR range image

in the reference HR range image. Similarly, each pixel in LR2 to LR4 (second to fourth range image) is formed by averaging four pixels in the HR image when the reference HR image is shifted by one pixel in each direction. Thus, the pixels in the LR images carry unique information about different regions in the HR image effectively resulting in a higher sampling rate to enable super-resolution [11]. The use of multiple range images also facilitates to average noise effects.

The above process can be expressed mathematically as

$$\mathbf{y}_i = DW_i\mathbf{x} + \eta_i \tag{1}$$

Here, $\mathbf{y}_i$ is the lexicographically arranged i^{th} LR observation and D and W_i are down-sampling and warping matrices, respectively, that produce $\mathbf{y}_i$ from the HR range image $\mathbf{x}$. The term η_i represents noise. Equation (1) can be expressed in scalar form as

$$y_i(n_1, n_2) = \sum_{l_1, l_2 = 1}^{L_1, L_2} d(n_1, n_2, l_1, l_2) \cdot x(\theta_{1i}, \theta_{2i}) + \eta_i(n_1, n_2) \tag{2}$$

where $d(n_1, n_2, l_1, l_2)$ is the element of the D matrix that maps the $(l_1, l_2)^{th}$ pixel in HR image $x(\theta_{1i}, \theta_{2i})$ to the $(n_1, n_2)^{th}$ pixel in the ith LR image. The

transformations θ_{1i} and θ_{2i} are the warping transformations that are encoded in the matrix W_i. For a translating camera, equation (2) simplifies to

$$y_i(n_1, n_2) = \sum_{l_1, l_2=1}^{L_1, L_2} d(n_1, n_2, l_1, l_2) \cdot x(l_1 - \delta_{1i}, l_2 - \delta_{2i}) + \eta_i(n_1, n_2) \tag{3}$$

where δ_{1i} and δ_{2i} are the shifts in the x and y directions, respectively.

From the above model, we observe that we require the W_i matrices that denote warping at high resolution. Since we consider only translational motion, we can compute the HR shifts by simply multiplying the LR shifts by the resolution factor. We compute the LR shifts using the well-known sub-pixel motion estimation algorithm proposed in [10].

3 Regularization Using MRF

Having discussed the formation of shifted LR images and estimation of their spatial shifts, we now address the problem of deriving the super-resolved (SR) range data $\mathbf{x}$ given observations $\mathbf{y_1}$, $\mathbf{y_2}$, ... $\mathbf{y_N}$. We propose to solve for the maximum a posteriori (MAP) estimate of x within a Bayesian framework. Let $\mathbf{Y_1}$, $\mathbf{Y_2}$, ... $\mathbf{Y_n}$ be the random fields associated with the observations $\mathbf{y_1}$, $\mathbf{y_2}$, ... $\mathbf{y_n}$ and let $\mathbf{X}$ be the random field associated with the SR range map. We wish to estimate $\widehat{\mathbf{x}}$ such that

$$\widehat{\mathbf{x}} = \max_{\mathbf{x}} P(\mathbf{X} = \mathbf{x} | \mathbf{Y_1} = \mathbf{y_1} ... \mathbf{Y_n} = \mathbf{y_n}) \tag{4}$$

Using Bayes rule, the above equation can be written as

$$\widehat{\mathbf{x}} = \max_{\mathbf{x}} P(\mathbf{Y_1} = \mathbf{y_1} ... \mathbf{Y_n} = \mathbf{y_n} | \mathbf{X} = \mathbf{x}) \cdot P(\mathbf{X} = \mathbf{x}) \tag{5}$$

Solving for $\mathbf{x}$ is clearly an ill-posed problem due to the down-sampling and warping operators, and due to the presence of noise [12,13]. We need to incorporate constraints on the solution through a suitably chosen prior.

The first term in the product on the right-hand side of the equation (5) is the likelihood term that arises from the image formation model. From equation (1), assuming a pin hole camera model and considering the noise to be additive white Gaussian with variance σ^2, we have

$$P(\mathbf{Y_1} = \mathbf{y_1} ... \mathbf{Y_n} = \mathbf{y_n} | \mathbf{X} = \mathbf{x}) = \frac{1}{(2\pi\sigma^2)^{N_1 N_2}} \exp\left(- \sum_{i=1}^{N} \frac{\|\mathbf{y}_i - DW_i\mathbf{x}\|^2}{2\sigma^2} \right) \tag{6}$$

We model the prior probability $P(\mathbf{X} = \mathbf{x})$ for the SR range image by a Markov random field. MRF modeling provides a natural way to embed constraints on the solution. The Markovianity property implies that the label at a pixel depends only on its neighborhood i.e., only neighboring labels have interactions with one another. This property is quite natural in the sense that the range value at a

particular pixel does not depend on the range values of pixels that are located far away from it. Due to the MRF - Gibbs equivalence [14], the prior probability of $\mathbf{X}$ can be expressed in analytical form as

$$P(\mathbf{X} = \mathbf{x}) = K \exp\left(-\sum_{c \in C_{\mathbf{x}}} V_c^{\mathbf{x}}(\mathbf{x})\right) \tag{7}$$

Here, $V_c^x(x)$ is the potential function and c is called a clique which is a subset of the MRF neighborhood. The potential function captures the manner in which neighboring pixels interact. For details on MRF, refer to [15]. From equations (6) and (7), we can rewrite equation (5) as

$$\widehat{\mathbf{x}} = \min_{\mathbf{x}} \left(\sum_{i=1}^{n} \frac{\|\mathbf{y}_i - DW_i\mathbf{x}\|^2}{2\sigma^2} + \sum_{c \in C_{\mathbf{x}}} V_c^{\mathbf{x}}(\mathbf{x}) \right) \tag{8}$$

The MAP - MRF framework results in an energy minimization formulation to estimate the SR range data $\widehat{\mathbf{x}}$, where the cost function is the bracketed term in equation (8). The first term in the cost is the data term. It measures how closely the transformed (warped and down-sampled) $\mathbf{x}$ compares with the observations. The form of the second term is crucial for a good solution. It is usually known as the smoothness term in image super-resolution works [16]. This is because the potential function usually has the form $V_c^{\mathbf{x}}(\mathbf{x}) = (x(i,j) - x(p,q))^2$ where pixels (p,q) belong to the neighborhood of (i,j). But this form of the potential function tends to select solutions that are smooth and results in loss of discontinuities and high frequencies which one would typically wish to preserve in the HR image.

We propose to use a discontinuity adaptive MRF (DAMRF) prior model for $\mathbf{x}$ in which the degree of interaction between pixels can be adjusted adaptively in order to preserve discontinuities. Li [15] suggests some models for DAMRF clique potentials. In this work, we use the potential function

$$V_c^{\mathbf{x}}(\mathbf{x}) = \gamma - \gamma e^{-(x(i,j)-x(p,q))^2/\gamma} \tag{9}$$

which is shown in Fig. 2(a). It is convex in the band $B_\gamma = (-\sqrt{\gamma/2}, \sqrt{\gamma/2})$ and the value of γ controls the shape of the function. Thus, choosing a large value of γ makes the function convex. Beyond B_γ, the cost of the prior tends to saturate as the difference between the pixel values increases. Hence, unlike the quadratic prior, the cost for a sudden change is not excessively high which allows discontinuities to be preserved in the solution. For the first-order MRF that we make use of in this work, the exact expression for the prior term in equation (7) is given by

$$\begin{aligned} \sum_{c \in \mathcal{C}} V_c^{\mathbf{x}}(\mathbf{x}) = \sum_{i=1}^{L_1} \sum_{j=1}^{L_2} 4 * \gamma - \gamma \exp\{-[x(i,j) - x(i,j-1)]^2/\gamma\} \\ -\gamma \exp\{-[x(i,j) - x(i,j+1)]^2/\gamma\} - \gamma \exp\{-[x(i,j) - x(i-1,j)]^2/\gamma\} \\ -\gamma \exp\{-[x(i,j) - x(i+1,j)]^2/\gamma\} \end{aligned} \tag{10}$$

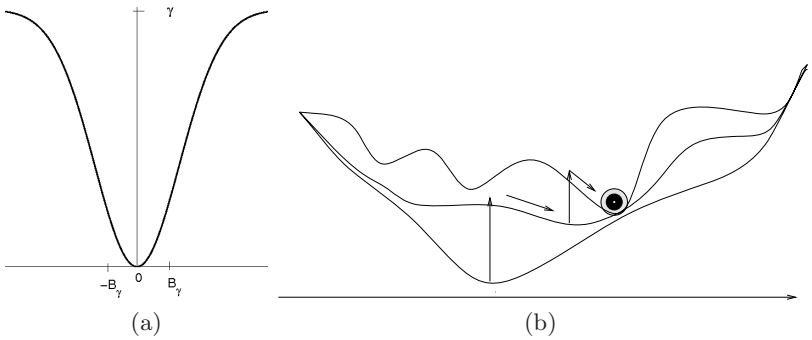

Fig. 2. (a) DAMRF clique potential function. (b) Graphical representation of GNC.

4 Energy Minimization: Graduated Non-convexity

Due to the non-convex nature of the DAMRF prior, the overall energy function in equation (8) becomes non-convex. Hence, minimization to obtain the MAP estimate of $\mathbf{x}$ becomes non-trivial in the sense that traditional local gradient-based techniques cannot be used as they can get trapped in local minima.

We minimize the energy using a deterministic annealing technique known as graduated non-convexity (GNC) [15,17]. The idea behind GNC is graphically illustrated in Fig. 2(b). It (initially) starts with a convex cost function by choosing a large value of γ and finds the minimum using simple gradient-descent. This value is then used as the initial estimate for the next iteration but now with a smaller γ. This step is iterated, reducing γ in each iteration. As shown in the figure, the lower curve symbolizes the energy function at an earlier iteration and the successive upper curves denote the energy functions at successive iterations. Note that the function non-convexity increases as the iterations progress. The vertical arrows denote that the located minimum of the previous iteration is considered as the initial state for the next iteration. The slant arrows depict convergence to the nearest minimum in each iteration.

The proposed algorithm for range super-resolution is summarized on the next page. We note that a basic requirement for GNC is computation of the gradient of the cost. From equations (8) and (10), the gradient at the kth iteration is

$$\text{grad}^{(k)} = \frac{1}{\sigma^2}\sum_{i=1}^{n} W_i^T D^T (DW_i\mathbf{x} - \mathbf{y_i}) + \lambda G^{(k)} \tag{11}$$

where λ is a smoothness parameter and the gradient $G^{(n)}$ at (p,q) is given by

$$\begin{aligned} G^{(k)}(p,q) = {} & 2[x(p,q) - x(p,q-1)]\exp\{-[x(p,q) - x(p,q-1)]^2/\gamma\} + \\ & 2[x(p,q) - x(p,q+1)]\exp\{-[x(p,q) - x(p,q+1)]^2/\gamma\} + \\ & 2[x(p,q) - x(p-1,q)]\exp\{-[x(p,q) - x(p-1,q)]^2/\gamma\} + \\ & 2[x(p,q) - x(p+1,q)]\exp\{-[x(p,q) - x(p+1,q)]^2/\gamma\} \end{aligned} \tag{12}$$

We would like to mention that all matrix computations in equation (11) can be incorporated easily through local image operations involving very few pixels. An important point in favor of the DAMRF prior from a computational perspective is that it is a continuous and differentiable function unlike other quasi-discontinuity handling priors such as line fields [14]. Hence, the computationally efficient GNC algorithm can be used as opposed to the much slower simulated annealing.

Algorithm: Range super-resolution using GNC.

Require: : Observations $\{y_i\}$ and motion parameters.

Calculate $\mathbf{x}^{(0)}$ as the average of the bi-linearly up-sampled and aligned images y_is

Choose a convex $\gamma^{(0)} = 2v$ where v is the maximum value of the gradient along the x and y directions of the initial estimate $\mathbf{x}^{(0)}$

$n = 0$

Do a. Update $\mathbf{x}^{(n)}$ as $\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \alpha\,\mathrm{grad}^{(n)}$

b. Set $n = n + 1$

c. If $(\mathrm{norm}(\mathbf{x}^{(n)} - \mathbf{x}^{(n-1)}) < \epsilon)$

set $\gamma^{(n)}$=max $[\gamma_{\mathrm{target}}, k\gamma^{(n-1)}]$

UNTIL $(\mathrm{norm}(\mathbf{x}^{(n)} - \mathbf{x}^{(n-1)}) < \epsilon)$ and $(\gamma^{(n)} = \gamma_{\mathrm{target}})$

Set $\widehat{\mathbf{x}} = \mathbf{x}^{(n)}$

Here, α is the step size, ϵ is a constant for testing convergence, and k is a factor that takes $\gamma^{(n)}$ slowly towards γ_{target}.

5 Experimental Results

In this section, we present results obtained using the proposed algorithm. The range data was captured with the PMD 3k-S range scanner from PMD Technologies. It operates at a wavelength of 870 nm with a modulating RF of 20 MHz. It captures range data at a frame-rate of 25 fps. The detector dimensions are 6.4 mm $\times$ 4.8 mm while the resolution of the range image is 64×48 pixels.

We captured range videos of objects by translating the object in front of the PMD scanner (which was held static). We ensured that the motion of the object deviated only negligibly from pure translation. We compute only frame-to-frame registration thus precluding the need for any controlled motion or calibration procedures. Also, we select successive or near-successive frames as observations for super-resolution to avoid any parallax effects. The captured video was converted into a readable ASCII format by the software CamVis Pro provided by PMD Technologies [6]. We attempted to super-resolve the range image by a factor of 4. Due to space constraints, we give only representative results here.

In the first experiment, we took range images of a wooden hand-model. It was translated fronto-parallel to the camera to ensure that the motion was translational. In fact, we translated only along the horizontal direction. (In general, the motion can be 2D. In fact, we require 2D motion for true 2D super-resolution).

Several frames of range data were captured and one of the low resolution depth maps is shown in Fig. 3(a). In order to perform super-resolution by a factor of 4, we selected a total of 8 frames from the range data set which were near-consecutive. Ideally, one would need 16 observations to upsample by a factor of 4. But since the motion is 1D in our example, we considered only 8 frames. These frames were then fed to the algorithm in [10] and the motion parameters for each frame was estimated. (As expected, the motion module gave negligible displacement along the vertical direction for this data set). The estimated motion parameters were then fed to the proposed GNC-based super-resolution algorithm. The initial value of γ was chosen as 1000 while γ_{target} was taken as 1 with $k = 0.9$. The initial estimate of the super-resolved image was obtained by averaging the bilinearly up-sampled and aligned low resolution observations. The values of λ and α were chosen as 0.01 and 6, respectively.

The GNC algorithm was run with the above parameters and the corresponding output (when super-resolved by a factor of 4) is given in Fig. 3(b). When compared with Fig. 3(a) in which the spatial localization of the range map of PMD is not good, we note that our output is significantly better. The contours of the hand are clearly discernible after performing super-resolution. In particular, note how well the thumb has been reconstructed in Fig. 3(b). We also tried bicubic interpolation of the low-resolution depth map but it suffers from the effect of range data bleeding into the background due to smoothing. This is to be expected since the low-resolution range map of PMD itself is not well-localized in space.

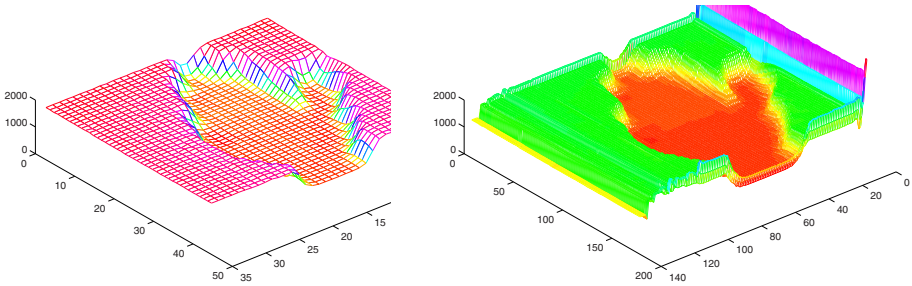

(a) (b)

Fig. 3. Wooden hand-model. (a) Low-resolution range data. (b) Super-resolved range map using the proposed method.

In the next experiment, we took images of the Alpha Rex robot. The robot was translated horizontally and the range data was recorded over many frames. We again chose 8 frames from this data set and attempted super-resolution by a factor of 4 using the proposed method. One of the low resolution frames is shown in Fig. 4(a). Even though one can make out the shape of the robot, the spatial extent of the robot is difficult to gauge due to poor spatial resolution of the range

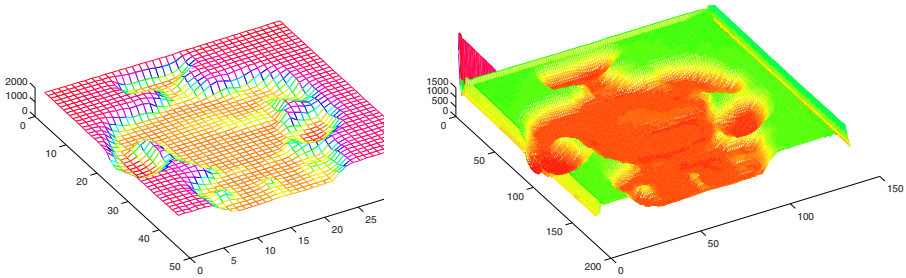

(a) (b)

Fig. 4. Alpha Rex robot. (a) Low-resolution range map. (b) Super-resolved range output of our method.

map from PMD. The bleeding effect is very much evident in the figure. (Please see the pdf file which has a color-coded version of the plot). In the output of the proposed method shown in Fig. 4(b) the edges of the robot are well-preserved and the localization is also very good due to improved spatial resolution. The noise level is comparatively much lower. The arms and elbow of the robot come out clearly, especially the right palm region. Note the wedge-shaped thighs and the holes in the legs which are difficult to infer from Fig. 4(a). The shape of the head emerges in its true shape. Also, the proposed algorithm correctly brings out the outline of the (slender) neck of the robot which appears smeared in Fig. 4(a).

Our method takes a few minutes to run on a 1 GHz Athlon PC with 64 MB RAM (for a non-optimized Matlab code). There is enough scope to increase the speed multi-fold by resorting to an efficient implementation in C on a faster machine.

6 Conclusions

In this paper, we proposed a new method for super-resolution of range data captured from a low-resolution PMD camera that (unlike previous works) avoids the need for cumbersome camera calibration. An edge-preserving MRF prior was used to adaptively retain discontinuities. We used GNC for non-convex optimization and showed results on real data to demonstrate the effectiveness of our technique. We are currently investigating fusion of range and intensity data. We are also working on extending our method to more general motion of the range sensor.

Acknowledgments

The first author gratefully acknowledges support from the Humboldt Foundation, Germany. The work was supported in part within the DFG excellence

initiative research cluster *Cognition for Technical Systems - CoTeSys* (visit website `www.cotesys.org`).

References

1. Kil, Y., Mederos, B., Amenta, N.: Laser scanner super-resolution. In: Eurographics Symposium on Point-Based Graphics, pp. 9–16 (2006)
2. Ghobadi, S., Hartmann, K., Weihs, W., Netramai, C.: Detection and classification of moving objects - stereo or time-of-flight images. In: International Conference on Computational Intelligence and Security, pp. 11–16 (2006)
3. Hahne, U., Alexa, M.: Combining time-of-flight depth and stereo images without accurate extrinsic calibration. In: International Workshop on Dynamic 3D Imaging, pp. 1–8 (2007)
4. Beder, C., Bartczak, B., Koch, R.: A comparison of PMD-cameras and stereo-vision for the task of surface reconstruction using patchlets. In: IEEE International Conference Computer Vision and Pattern Recognition, pp. 1–8 (2007)
5. Reulke, R.: Combination of distance data with high resolution images. In: ISPRS Commission V Symposium Image Engineering and Vision Metrology, pp. 1-6 (2006)
6. `http://www.pmdtec.com`
7. Wallhoff, F., Ruß, M., Rigoll, G., Gobel, J., Diehl, H.: Improved image segmentation using photonic mixer devices. In: Proceedings IEEE Intl. Conf. on Image Processing, vol. VI, pp. 53–56 (2007)
8. Prasad, T., Hartmann, K., Weihs, W., Ghobadi, S., Sluiter, A.: First steps in enhancing 3D vision technique using 2D/3D sensors. In: Computer Vision Winter Workshop, pp. 82–86 (2006)
9. Huhle, B., Fleck, S., Schilling, A.: Integrating 3D time-of-flight camera data and high resolution images for 3DTV applications. In: 3DTV-Conference, pp. 1–4 (2008)
10. Irani, M., Peleg, S.: Improving resolution by image registration. Graphical Models and Image Processing 53(3), 231–239 (1991)
11. Park, S.C., Park, M.K., Kang, M.G.: Super-resolution image reconstruction: A technical overview. IEEE Signal Processing Magazine 16(3), 21–36 (2003)
12. Farsiu, S., Robinson, D., Elad, M., Milanfar, P.: Fast and robust super-resolution. In: IEEE International Conference on Image Processing, pp. 14–17 (2003)
13. Rav-Acha, A., Zomet, A., Peleg, S.: Robust super resolution. In: IEEE International Conference on Computer Vision and Pattern Recognition, pp. 645–650 (2001)
14. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distribution and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence 6(6), 721–741 (1984)
15. Li, S.Z.: Markov random field modeling in computer vision. Springer, Tokyo (1995)
16. Suresh, K., Rajagopalan, A.N.: Robust and computationally efficient super-resolution algorithm. Journal of the Optical Society of America - A 24(4), 984–992 (2007)
17. Blake, A., Zisserman, A.: Visual reconstruction. The MIT Press, Cambridge (1987)