2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

# RGB-D Sensor Data Correction and Enhancement by Introduction of an Additional RGB View

Artashes Mkhitaryan and Darius Burschka

***Abstract*— RGB-D sensors are becoming more and more vital to robotics. Sensors such as the Microsoft Kinect and time of flight cameras provide 3D colored point-clouds in real time can play a crucial role in Robot Vision. However these sensors suffer from precision deficiencies, and often the density of the point-clouds they provide is insufficient. In this paper, we present a multi-camera system for correction and enhancement of the data acquired from an RGB-D sensor. Our system consists of two sensors, the RGB-D sensor (main sensor) and a regular RGB camera (auxiliary sensor). We perform the correction and the enhancement of the data acquired from the RGB-D sensor by placing the auxiliary sensor in a close proximity to the target object and taking advantage of the established epipolar geometry. We have managed to reduce the relative error of the raw point-cloud from a Microsoft Kinect RGB-D sensor by $74.5\,\%$ and increase its density up to $2.5$ times.**

## I. Introduction

The invention of RGB-D sensors such as the Microsoft Kinectand time of flight (ToF) cameras has resulted in a tremendous leap ahead in many areas of computer/robot vision. The ability to acquire dense 3D colored point-clouds in real time allowed significant advancements in dense stereo reconstruction, object registration, grasping, indoor navigation and motion capture among others.

For example, Holz et. al. [1] introduce an approach for mobile robots that allows them to create the obstacle map of the environment and classify the graspable objects by analyzing the data acquired from an RGB-D sensor. In [2], a technique is presented in which a hierarchical, multi-view dataset of objects is created based on the RGB-D data. This is then used to identify and register objects. Schwarz et. al. [3], [4] introduce a system for human motion capture. The authors use dense 3D point-clouds in combination with the corresponding RGB images to create a graph-based representation of the target person and determine her or his pose.

Although RGB-D sensors have many advantages, it is not always possible to achieve the desirable or necessary performance without significant post-processing. Early ToF sensors had very small image resolution ($204 \times 204$) and where introducing distortions in depth values around the corners of large objects due to double reflections. Modern RGB-D sensors, like the Microsoft Kinect, don't suffer from the resolution issue, however the average error of the depth value estimation can amount up to $4cm$ [5] depending on the distance to the target object.

This work has been supported by an internal grant "Real-Time Perception and Exploration with Collaborating Agents" of the German Aerospace Center (DLR)

A. Mkhitaryan and D. Burschka are with the Machine Vision and Perception Group, Technical University of Munich, Munich, Germany 80339 {mkhitary|burschka}@.in.tum.de

Fig. 1. Illustration of the experimental setup. The RGB-D sensor is on the bottom of the image while the RGB camera is on the left side.

In this paper, we introduce a system for enhancing and correcting the data acquired from an RGB-D sensor. The system consists two sensors, the RGB-D sensor or main sensor and a regular RGB camera or the auxiliary sensor. The auxiliary sensor is placed closer to the target object than the main sensor and the two sensors are calibrated against each other based on the data acquired from them. In the second step we leverage the estimated epipolar geometry to perform a correction of the 3D points acquired from the RGB-D sensor. As a final step, we enhance the 3D point-cloud by taking advantage of the fact that the auxiliary sensor is located closer to the target object, thus it has a more detailed view of the target than the main sensor. Using this approach, for a best case scenario we succeeded in correcting the positions of the points in the point-cloud by an average of $74.5\,\%$, and increase its density up to 2.5 times. In contrast to many temporal approaches such as Kinect fusion [6], [7] our system provides the results based on only one time step.

## II. Related Work

In the early RGB-D sensors (before the Microsoft Kinect) correction or enhancement of the depth information was done with the help of a stereo system. Zhu et. al. [8] introduced a technique where a ToF range sensor is used in combination with a stereo pair to compute the depth probability distribution, which in turn can be used to improve the accuracy of the point-cloud. Another approach for ToF and stereo pair fusion is introduced in [9]. Here the authors are acquiring an initial, "rough" point cloud from the ToF and enhancing it by optimizing an energy function through a probabilistic multi-view framework. In both of the above presented approaches, the authors use a ToF camera as

an RGB-D sensor, which provides sparse inaccurate point-clouds. In the former case authors cope with the correction of the point-cloud, and in the latter they enhance the density, however neither of the approaches deal with both problems simultaneously. Gould et. al. [10] present a technique for object detection, by enhancing the 2D images. Where they use a a laser range scanner as an RGB-D sensor to acquire a sparse point cloud, which is later used with the RGB data from a video camera to achieve a high resolution, however the question of the accuracy of the point cloud is not addressed by the authors.

Later approaches are mainly using the Microsoft Kinect as the RGB-D sensors, and are trying to correct or enhance the provided depth information, by either using an additional stereo reconstruction for the IR and RGB images, or by fusing the data acquired from the same sensor over time. Chiu et. al. [11] introduce a technique for object registration based on RGB-D data. Here the authors combine the depth map from the range sensor and the disparity map from the IR-RGB pair to increase the reliability of the depth estimate. This approach mainly focuses on reconstruction of the areas of the scene that have shiny or glossy surfaces, since the IR camera is not reliable in these cases. However the density and precision of the general point-cloud is not addressed. In [12] an approach is presented where the authors use a similar technique to increase the accuracy of the point-cloud. Henry et. al. [13] introduce a method for combining the visual features from the RGB camera and the shape-based alignment from the range sensor to construct reliable 3D maps of the environment. In [14], the authors propose the use of multiple Kinect cameras to increase the reliability and robustness of image processing systems, by fusing the data acquired from each individual Kinect. And last, Matyunin et. al. [15] use temporal filtering of the RGB-D sensor data to cope with occlusions and improve the temporal stability.

All of the above described approaches cope with the issue of enhancing or correcting the data from RGB-D sensors. However these approaches either deal with the correction of the point-cloud or its enhancement, and do not address both issues combined. In this paper we propose an approach that allows both the enhancement and the correction of the data acquired from an RGB-D sensor by introducing an additional RGB sensor. Our approach allows an automatic estimation of the epipolar geometry between the two sensors. It then performs the correction followed by an enhancement of the raw point-cloud data by taking an advantage of the estimated geometry.

The rest of the paper is organized as follows. We start with description of our approach for calibration (Section III-A), correction (Section III-B) and enhancement (Section III-C). In the following section, we present the evaluation of our approach. Specifically in Section IV-A we describe the experimental setup, and give a quantitative (Section IV-B) and qualitative (Section IV-C) assessment. Finally, conclusions are presented in Section V.

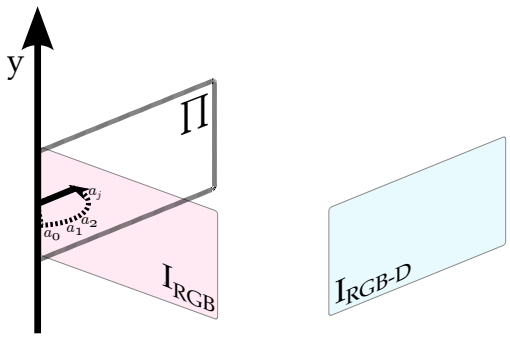


Fig. 2. Illustration of the calibration process.

## III. Approach

We propose an algorithm for correction and enhancement of an RGB-D sensor (hereinafter main sensor) data, by an introduction of an auxiliary RGB camera (hereinafter auxiliary sensor). Since most of the known RGB-D sensors consist of two parts, an RGB camera, and a depth sensor, we model the main sensor similarly. We presume that the mentioned components of the main sensor are calibrated to each other, and that both intrinsic and extrinsic parameters are known. For simplicity, the colored point-cloud $\{p_i\}$ acquired from the main sensors will further on be referred to as *point-cloud* $\{p_i\}$.

The proposed algorithm can be divided into three parts: calibration, correction and enhancement. Initially, we compute the pose of the auxiliary sensor with respect to the main one, based on the images acquired from their RGB cameras and point-cloud $\{p_i\}$ (Section III-A). Further, we use the established epipolar geometry (Fig. 4, Section III-B) to correct the poses of the points in $\{p_i\}$. And finally we use the corrected point-cloud $\{p'_i\}$ and its projection to the image plane of the auxiliary sensor to increase the density of $\{p'_i\}$ (Section III-C).

### A. Calibration

We establish the geometry between the sensors by solving a perspective n point problem [16]. This requires point correspondences between the point-cloud $\{p_i\}$ and the image of the auxiliary sensor. We obtain the correspondences by matching SURF [17] features from the RGB images acquired by both sensors. Since the SURF features are sensitive to large viewpoint foreshortening effects, to insure maximum number of matches we inherit an iterative binary search like algorithm which establishes a rough rotation between the sensors along y axes. For stationary sensors the calibration step is only performed once.

Given the images $I_{RGB}$ and $I_{RGB-D}$ acquired from the auxiliary and main sensors respectively as well as the point-cloud $\{p_i\}$ we proceed as follows. First, we project the *image* $I_{RGB}$ onto a virtual 3D plane $\Pi$ (Fig. 2). Following, we rotate the plain $\Pi$ for a range of angles $\{\alpha_j\}$ along its $y$ axis. Here the range $\{\alpha_j\}$ covers a broad spectrum, however the step $||\alpha_j - \alpha_{j+1}||$ is large. For each $\alpha_j$, we project the rotated plane $\Pi$ to the image plane of the auxiliary sensor,

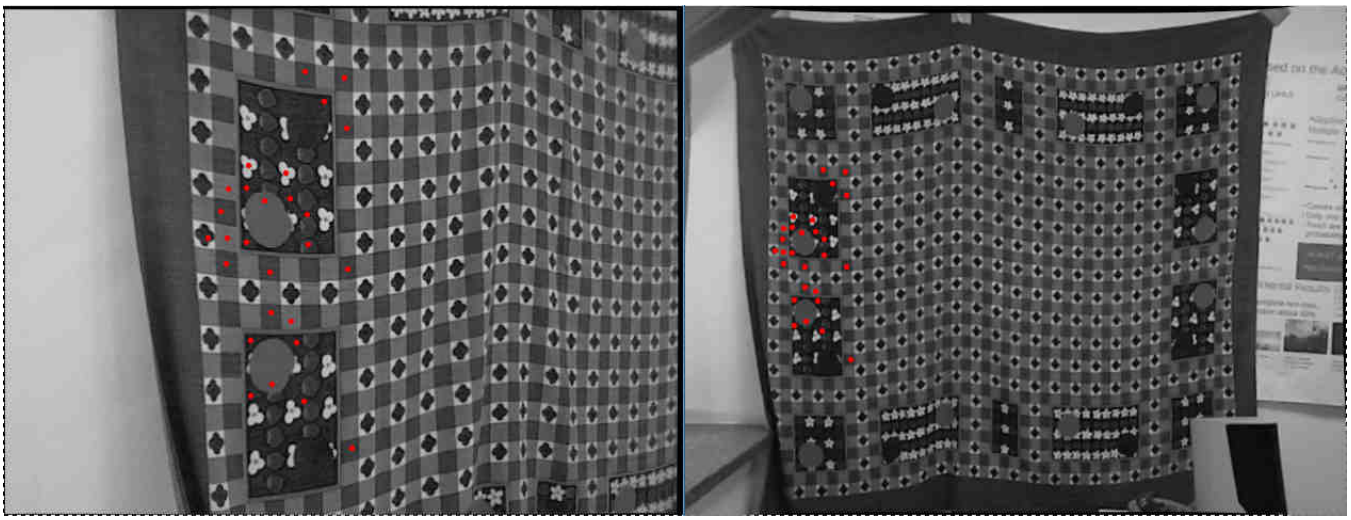

Fig. 3. Illustration of the calibration process. On the left is the image registered by the auxiliary sensor, and on the right the image acquired by the RGB camera of the main sensor. The red dots are the matching features on both images.

which results in set of images $\{I^j_{RGB}\}$ and homographies $\{H^j\}$, such that:

$$H^j I^j_{RGB}(u,v) = I_{RGB}(u,v) \tag{1}$$

Next, for each $I^j_{RGB}$ we extract SURF features and match them with SURF features from $I_{RGB-D}$ (Fig. 3). Further, we define a new range of angles $\{\alpha'_j\}$ with a finer resolution (i.e. $||\alpha'_j - \alpha'_{j+i}|| < ||\alpha_j - \alpha_{j+1}||$) around the angle $\alpha_j$ that corresponds to the pair $\{I^j_{RGB}, I_{RGB-D}\}$ with most feature matches. We repeat this process iteratively until the difference between the largest amount of feature matches from previous and current iterations is smaller than a certain threshold.

Then we establish the rotation and translation $\{R,T\}$ between the two sensors by using RANSAC combined with Perspective-N-Point problem for the $H^j\{f(u,v)\}_j$ and $\{\hat{p_j}\} \in \{p_i\}$, where $\{f(u,v)\}_j$ are the matched features from $I^j_{RGB}$ and $\{\hat{p_j}\}$ are the 3D points corresponding to the matched features from $I_{RGB-D}$.

$$H^j\{f(u,v)\}_j = K[R|T]\{\hat{p_j}\} \tag{2}$$

Where $K$ is the camera matrix.

### B. Correction

The main reason for inaccuracies in point cloud $\{p_i\}$ is the faulty depth value estimation by the RGB-D sensor [5]. We cope with this issue by taking advantage of the fact that a different view, i.e. the image plane of the RGB sensor, provides a better observability of the target object Fig. 4. Therefore, due to the depth inaccuracies the corresponding pixels to the projected points $p_i$ would be shifted along the epipolar line Fig. 4. Note, that here the size of the shift depends on the angle $\phi$ between the sensors, the closer the angle to $\pi/2$ the better the observability.

We project the point-cloud $\{p_i\}$ on the image plane of the RGB sensor:

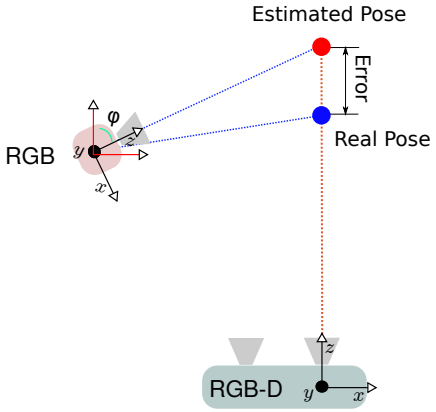


Fig. 4. Illustration of the faulty depth value estimation. Here the blue dot represents the real pose of the point, and the red one the estimated pose by the depth sensor. Note that from a different view, i.e. on the image plane of the auxiliary sensor (here depicted in red) the faulty estimate would be shifted along the epipolar line.

$$I_P = K(R\{p_i\} + T) \tag{3}$$

where $K$ is the intrinsic matrix of the auxiliary sensor and $I_P$ is the resulting image. Following, we compute the matches between all the pixels of $I_P$ and $I_{RGB}$ using Farneback's [18] algorithm. Next, for each matching pair of pixels $\{I_p(u_j,v_j), I_{RGB}(u'_j,v'_j)\}$ we compute the displacement values $(du_j, dv_j)$:

$$\begin{aligned} du_j &= u_j - u'_j \\ dv_j &= v_j - v'_j \end{aligned} \tag{4}$$

and for each $I_P(u_j + du_j, v_j + dv_j)$ pixel we compute $\hat{v}^j_P$ and $\hat{v}^j_d$ unit vectors:

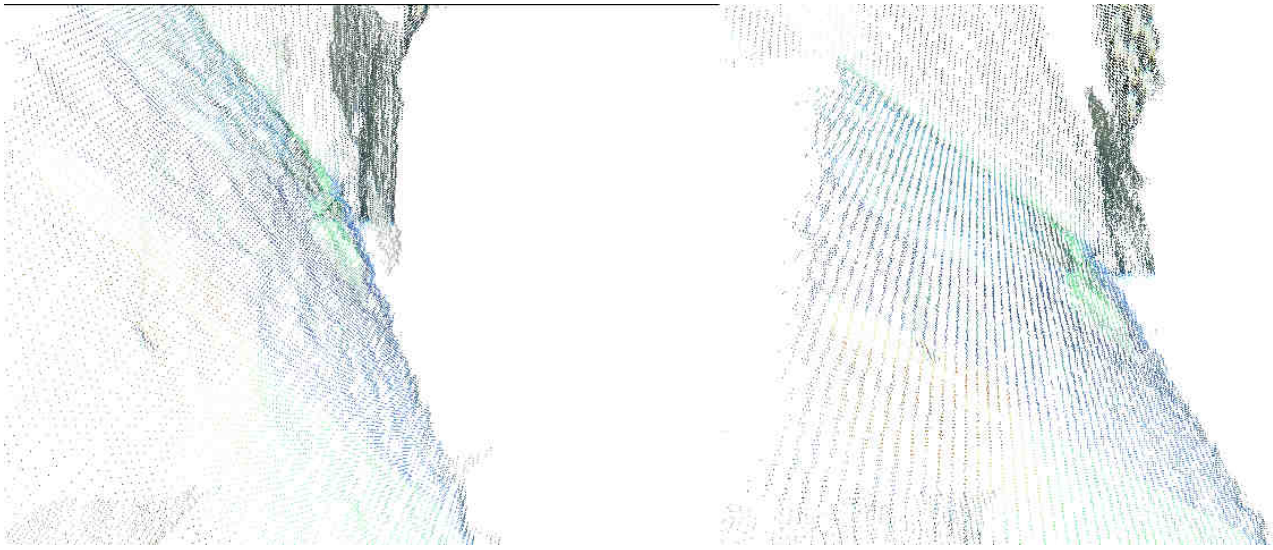

Fig. 5. Illustration of the point-cloud correction. Here the right side represents the raw point-cloud acquired from the RGB-d sensor, and the left side the same point-cloud after correction step. Note, that for the visualization purposes the example is brought on a flat surface, and that the placement of the virtual camera has been done by hand, thus viewing angle is not exactly the same in both of the images.

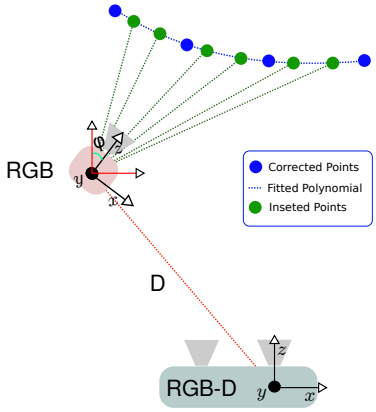


Fig. 6. Illustration of the point-cloud enhancement, here the blue dots represent the corrected points of the point-cloud, the green dashed line is the fitted polynomial, and the green dots are the new interpolated points of the enhanced point-cloud

$$\hat{v}_P^j = \frac{K\begin{pmatrix} u_j + du_j \\ v_j + dv_j \\ 1 \end{pmatrix}}{\|K\begin{pmatrix} u_j + du_j \\ v_j + dv_j \\ 1 \end{pmatrix}\|} \qquad (5)$$

$$\hat{v}_D^j = \frac{p_j}{\|p_j\|}$$

where $p_j$ is the 3D point that corresponds to $I_P(u_j, v_j)$ pixel, i.e. $I_P(u_j, v_j) = K(Rp_j + T)$. The new corrected point-cloud $\{p'_j\}$ (see Fig. 5) is computed by triangulating the $\{\hat{v}_P^j\}, \{\hat{v}_D^j\}$ sets. For triangulation we use the algorithm proposed by Hartley et. al. in [19] (Chapter 12.5).

### C. Enhancement

In the final stage of our algorithm we enhance point-cloud $\{p'_i\}$, i.e. we increase its density. Fig. 7 shows the difference between the image registered by the auxiliary sensor and the one obtained by projecting $\{p'_i\}$ on to it's image plane. Here the black dots represent perfect color matches between the two images. Note, that they are distributed sparsely in the image plane. This is due to the fact that the auxiliary sensor is located closer to the target object than the main sensor Fig. 9. Here, for each of the non-black pixels the only information missing for the 3D reconstruction of its pose is the depth value. Latter we interpolate based on the pose of the point $p'_i$ that correspond to the neighboring black pixels. Note, that this kind of interpolation has an advantage over simple sub-sampling techniques of range images such as interpolation of range values, in that it provides a better preservation of the original shape, due to the fact that here the interpolation is done only for one degree of freedom (Fig. 6).

The interpolation is performed as follows. Around each of the non-black pixels we define a window. Further, we analyze the depth discontinuities within that window. This is done by comparing the absolute distances between points $p'_i$ corresponding to nearest black-pixel neighbors, to a certain threshold. If there are no depth discontinues within the window, we interpolate the depth value of the non-black pixel by fitting a quadratic polynomial within the window. The latter is based on depth values of the 3D points corresponding to neighboring black pixels located within the window Fig. 6.

$$D(u, v) = q_1u^2 + q_2v^2 + q_3uv + q_4u + q_5v + p_6 \qquad (6)$$

Where $D(u, v)$ is the depth value of the colored pixel, $(u, v)$ are its pixel coordinates, and $q_i$ are the coefficients of the polynomial. An example of the final result is depicted in Fig. 8. Here, the borders of the target object are left

Fig. 7. Projection of the colored 3D point-cloud to the image plane of the auxiliary sensor. The difference between the registered image on the RGB sensor and the projected, corrected colored point cloud is evident. The black dots represent one-to-one color matches. Note, that since the auxiliary sensor is located closer to the target object than the main sensor, the black dots are sparse.

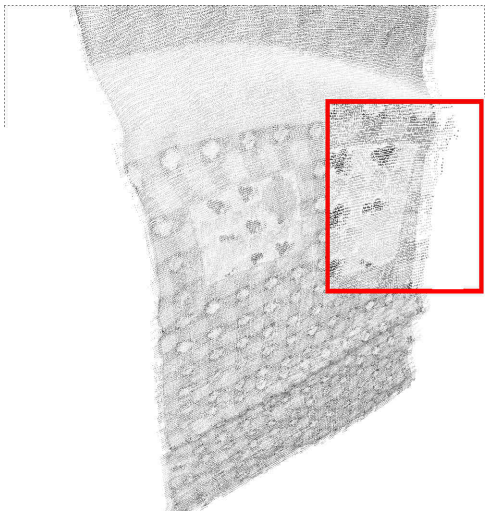

Fig. 8. Final result after the correction and the enhancement steps have been performed on Fig. 7. for the visualization purposes the content of the red rectangle has been zoomed and the borders of the object has not been processed to show the difference between the raw and processed data. Note, the difference between the densities.

unprocessed to illustrate the difference between the raw and processed data. To insure better visibility the content of the red rectangle has been zoomed.

## IV. Experiments and Results

In this Section, we first give an overview of the experimental setup used to evaluate our approach, followed by a description of the conducted experiments. Finally, we present results for qualitative and quantitative assessment.

### A. Experimental Setup

In the conducted experiments we used a Microsoft kinects as the main sensor, and a color web-cam with $640 \times 480px$ resolution as the auxiliary sensor. The sensors were placed as depicted in Figs. 1 and 9. We have conducted the experiments on a variety of target objects, such as a planar wall and a carton box with flat, planar sides to evaluate the performance

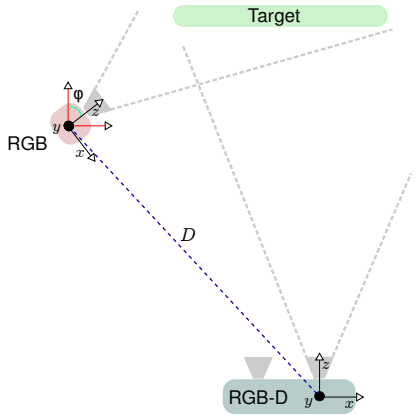


Fig. 9. System setup

of the correction stage, and a large carpet with a curved surface to evaluate the performance of the enhancement step.

### B. Quantitative Assessment

As mentioned in Section III, our approach can be divided into three logical stages, calibration, correction, and enhancement. Further we provide the qualitative assessment for each of the mentioned logical stages.

*1) Calibration of Sensors:* To evaluate the calibration of our system we ran the calibration stage several times with static placement of the sensors. For each iteration we have computed the distance $\|R' * T\|$ between the two sensors. The results are presented in Fig. 10(a). We have used a checker board calibration pattern to obtain the true distance of $0.524\,78\,\mathrm{m}$ between the sensors. On average the calibration stage performed with an absolute error of $\pm 0.0022\,\mathrm{m}$ which is $0.5\%$ of the absolute distance.

*2) Evaluation of point-cloud correction:* As is mentioned in Section III-B, the average distance between the raw and corrected corresponding points depends on the displacement of the pixel corresponding to point $p_i$ along the epipolar line in the image plane of the auxiliary sensor. Since the observability of the error by the auxiliary camera mainly depends on the angle $\phi$ (Fig. 9), the evaluation of the correction stage is done in reference to it. As a target object, we chose an object that has a planar flat surface Fig. 5. Since the mathematical description of a planar surfaces is well known, it is possible to compute the ground truth analytically. The cameras where placed as is depicted in Fig. 9. Where the distance between the target object and the RGB-D sensor was found to be approximately $1.35\,\mathrm{m}$.

We perform the evaluation of the correction stage of our approach in reference to the angle $\phi$ and average correction magnitude. The latter is defined as the difference between the average displacement values $Q$ and $Q'$ of the processed and raw point-clouds from the planar surface of the target object. We have processed the point-cloud corresponding to the target object for a range of angles $\phi$, and for each iteration we have computed the average correction magnitude. Each measurement was done $20$ times, and for each time the average was computed over $3000$ points. The results of our

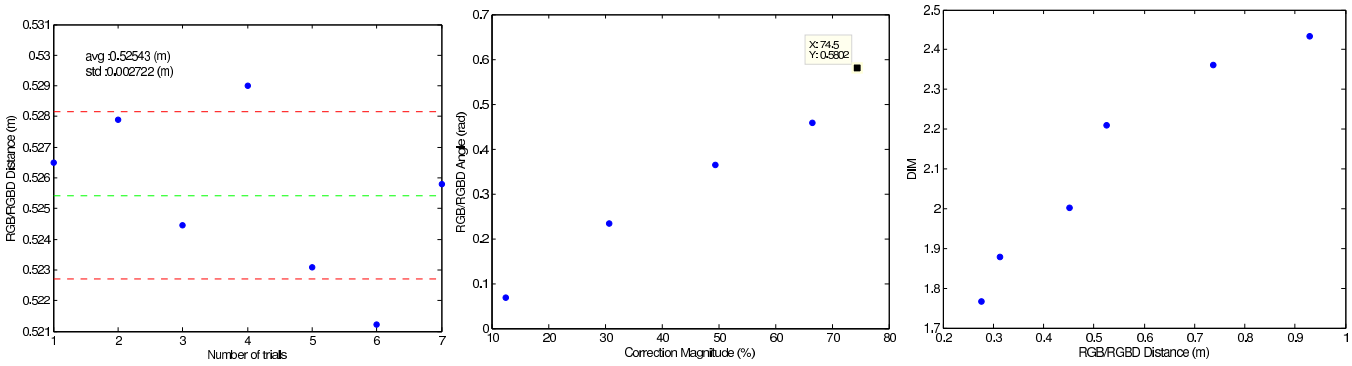


(a) Assessment of the accuracy of the calibration stage. Here, the absolute distance between the cameras is $0.524\,78\,\mathrm{m}$, the standard deviation of the estimated values is around $2.7\,\mathrm{mm}$ and worst case scenario error is around $3.5\,\mathrm{mm}$

(b) Relative correction magnitude relation to the angle $\phi$ (Fig. 9) between the cameras.

(c) Relation of the density increase multiplier (DIM) to the distance $D$ (Fig. 9) between the cameras.

Fig. 10.

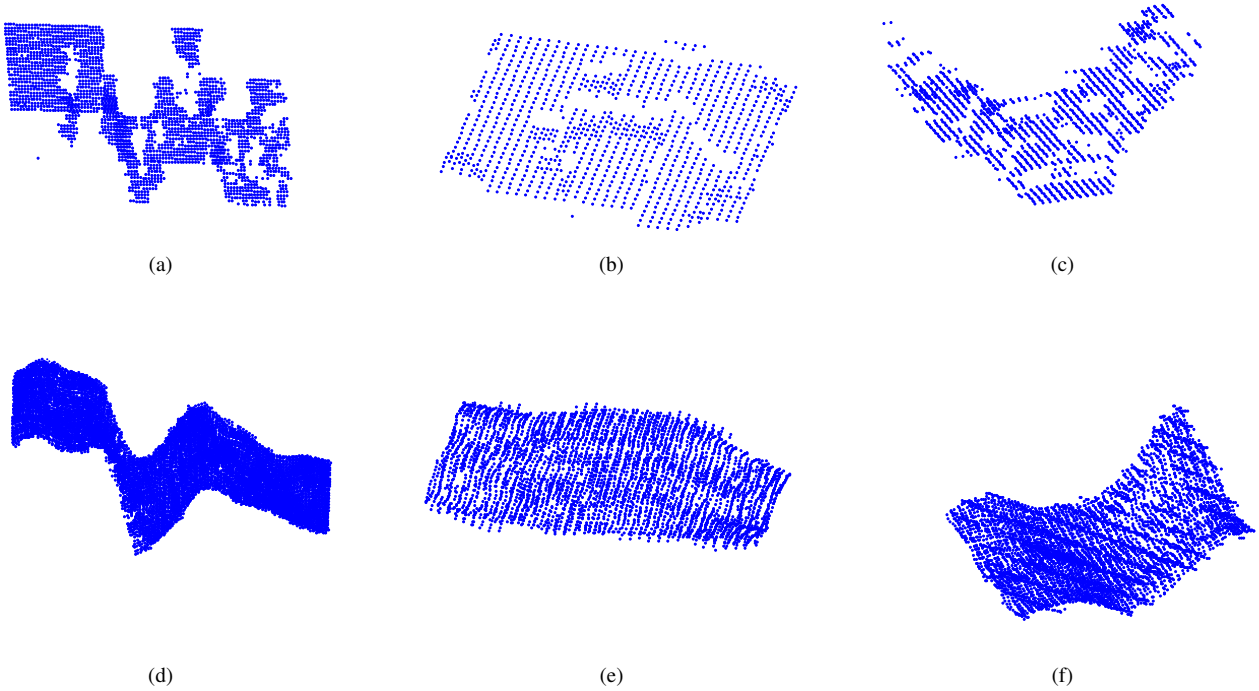

(a) (b) (c)

(d) (e) (f)

Fig. 11. Depiction of the final result for qualitative assessment. Here Figs. 11(a) to 11(c) illustrate the raw data as acquired by the RGB-D sensor while Figs. 11(d) to 11(f) are the same patches after the correction and enhancement steps have been applied. The patches represent parts of the carpet illustrated in Fig. 8. Note, that as a result of processing some of the gaps from the raw data have been filled and the patches obtained smoother and more realistic curvatures.

experiments are depicted in Fig. 10(b). We have achieved a relative average correction magnitude of around $74.5\,\%$. Here, the absolute average displacement $Q$ of the raw point-cloud was around $4\,\mathrm{cm}$, and the absolute value of $Q'$ was around $1\,\mathrm{cm}$. Note that the average correction magnitude also depends on the distance of the auxiliary sensor to the target object, however the corresponding experiments has shown that this dependency is week and negligible in comparison to dependency to the angle $\phi$, thus the results of those experiments are omitted.

*3) Evaluation of Enhancement:* To evaluate the enhancement step, we first introduce a metric called density increase multiplier (DIM). The latter is defined as the ratio between the amount of points after enhancement to the amount of points before. As was mentioned the increase in density after the enhancement stage is due to the fact that the auxiliary sensor is located closer to the target object than the main sensor. Thus, we evaluate the enhancement stage in reference

to the distance $D$ between the sensors and the DIM. This is realized by keeping the main sensor static and moving the auxiliary sensor towards the target object. For each distance we performed 20 iterations to estimate the average DIM. The results of the experiments are illustrated in the Fig. 10(c). The original density of the patch was found to be around 3600 points. In the best case scenario, where the cameras were 1 m apart from each other, we achieved up to 2.5 times increase in the density of the point-cloud.

### C. Qualitative Assessment

*1) Correction Step:* For qualitative assessment of the correction step we use a carton box with planar flat sides. This is done so it is easier for a naked eye to notice the changes. The results are illustrated in Fig. 5. The left side of the image depicts the raw point-cloud as acquired by the RGB-D sensor. The right side is the same point-cloud after a correction stage has been performed (the colors of the illustration are inverted, to insure a better visibility on printed formats). Note, that in both cases the viewing angle was adjusted by hand, thus a reasonable displacement of the scene is to be assumed.

*2) Assessment of the Entire Algorithm:* Fig. 11 illustrates the improvements achieved in different patches of the point-cloud. These patches were obtained using the carpet depicted in Fig. 8 as a target object. The latter was chosen due to its curved structure, which makes it easier to observe the improvements of the entire approach. Figs. 11(a) to 11(c) illustrate the unprocessed patches, while Figs. 11(d) to 11(f) are the respective patches after the correction and enhancement. Note, that the processed patches have smoother and more realistic representation of the curves, as well as that some of the gaps notable in the raw patches has been filled.

## V. Conclusion

Currently, RGB-D sensors play a major role in such robotics applications as navigation, grasp planning, object recognition and perception. In most of these applications the precision of the point-clouds is a decisive factor in overall system performance. In this paper we presented a system for correcting and enhancing the data acquired from an RGB-D sensor by introducing an auxiliary RGB sensor. We have managed to correct the accuracy of the point-cloud from RGB-D sensor by an average magnitude of 74.5 %, and increase its density up to 2.5 times. Note that the Microsoft Kinect sensor, is one of the most used RGB-D sensors to our knowledge, and has an average absolute error of 4 cm in depth estimation [5].

## References


[1] D. Holz, S. Holzer, R. B. Rusu, and S. Behnke, "Real-time plane segmentation using rgb-d cameras," in *Proceedings of 15th RoboCup International Symposium*, 2011.

[2] K. Lai, L. Bo, X. Ren, and D. Fox, "A large-scale hierarchical multi-view rgb-d object dataset," in *IEEE ICRA*, 2011.

[3] L. Schwarz, A. Mkhitaryan, D. Mateus, and N. Navab, "Estimating human 3d pose from time-of-flight images based on geodesic distances and optical flow," in *IEEE Face and Gesture Recognition (FG)*, 2011.

[4] ——, "Human skeleton tracking from depth data using geodesic distances and optical flow," *Image and Vision Computing, Elsevier*, vol. 30, p. 217226, 2012.

[5] K. Khoshelham, "Accuracy analysis of kinect data," in *ISPRS workshop laser scanning*, 2011.

[6] D. Molyneaux, "Kinectfusion rapid 3d reconstruction and interaction with microsoft kinect," in *International Conference on the Foundations of Digital Games*, 2012.

[7] S. Izadi, D. Kim, O. Hilliges, D. Molyneaux, R. Newcombe, P. Kohli, J. Shotton, S. Hodges, D. Freeman, A. Davison, and A. Fitzgibbon, "Kinectfusion: Real-time 3d reconstruction and interaction using a moving depth camera," in *ACM Symposium on User Interface Software and Technology*, 2011.

[8] J. Zhu, L. Wang, R. Yang, and J. Davis, "Fusion of time-of-flight depth and stereo for high accuracy depth maps," in *Computer Vision and Pattern Recognition CVPR*, 2008.

[9] Y. M. Kim, C. Theobalt, J. Diebel, J. Kosecka, B. Miscusik, and S. Thrun, "Multi-view image and tof sensor fusion for dense 3d reconstruction," in *IEEE Computer Vision Workshops ICCV Workshop*, 2009.

[10] S. Gould, P. Baumstarck, M. Quigley, A. Y. Ng, and D. Koller, "Integrating visual and range data for robotic object detection," in *Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications M2SFA2*, 2008.

[11] W.-C. Chiu, U. Blanke, and M. Fritz, "Improving the kinect by cross-modal stereo," in *British Machine Vision Conference (BMVC)*, 2011.

[12] K. Xu, L. Qin, and L. Yang, "Rgb-d fusion toward accurate 3d mapping," in *IEEE International Conference on Robotics and Biomimetics*, 2011.

[13] P. Henry, M. Krainin, E. Herbst, X. Ren, and D. Fox, "Rgb-d mapping: Using depth cameras for dense 3d modeling of indoor environments," in *International Symposium on Experimental Robotics ISER*, 2010.

[14] L. Sumar and A. Bainbridge-Smith, "Feasibility of fast image processing using multiple kinect cameras on a portable platform," in *Feasibility of Fast Image Processing Using Multiple Kinect Cameras on a Portable Platform*, .

[15] S. Matyunin, D. Vatolin, Y. Berdnikov, and M. Smirnov, "Temporal filtering for depth maps generated by depth camera," in *The True Vision - Capture, Transmission and Display of 3D Video (3DTV-CON)*, 2011.

[16] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnp: An accurate o(n) solution to the pnp problem," *International Journal of Computer Vision*, vol. 81, pp. 155–166, 2009.

[17] H. Bay, T. Tuytelaars, and L. V. Gool, "Surf: Speeded up robust features," in *Lecture Notes in Computer Science ECCV*, 2006.

[18] G. Farneback, "Two-frame motion estimation based on polynomial expansion," in *Springer Berlin Heidelberg*, 2003, pp. 363–370.

[19] R. Hartley and A. Zisserman, *Multiple View Geometry in computer vision*, ., Ed. Cambridge University Press, 2008.