https://doi.org/10.2352/ISSN.2470-1173.2020.2.SDA-138

Objective and Subjective Evaluation of a Multi-Stereo 3D Reconstruction System

Christian Kapeller [1 2] ***; Braulio Sespede*** [2] ***; Matej Nezveda*** [3] ***; Matthias Labschütz*** [4] ***; Simon Flöry*** [4] ***; Florian Seitner*** [3] ***; Margrit Gelautz*** [2]
[1] ***AIT Austrian Institute of Technology GmbH, High Performance Vision, 1210 Vienna, Austria***
[2] ***Vienna University of Technology, Institute of Visual Computing & Human-Centered Technology, 1040 Vienna, Austria***
[3] ***emotion3D GmbH, 1040 Vienna, Austria***
[4] ***Rechenraum GmbH, 1050 Vienna, Austria***

Abstract

In recent years, 3D reconstruction systems comprising multiple depth sensors have received increasing interest for dynamic scene reconstruction and related applications. Publicly available ground truth data are of limited usefulness when dealing with quality assessment of self-recorded data delivered by customized stereo configurations. In this paper, we propose a framework that incorporates versatile strategies for quantitative and qualitative evaluation of a multi-stereo reconstruction system and its intermediate products. Besides the design of suitable calibration objects for quantitative measurements, the framework exploits multi-view data redundancy and generated novel views for objective quality assessment and to obtain subjective ratings from users. We demonstrate the applicability of our evaluation system in experiments with several stereo matching algorithms and view fusion approaches along with a pair-comparison based user study. We believe that our proposed evaluation framework is beneficial for the assessment of 3D products derived from self-recorded dynamic data of comparable set-ups, for example, in the context of subsequent augmented reality applications.

Introduction and Related Work

In recent years, 3D model acquisition systems for dynamic scenes - for example, for performance capturing of actors - have received increasing attention. A common practice in 3D reconstruction is an evaluation based on ground truth data (e.g. [1, 23]).

The utility of publicly available data sets, however, is limited when applied to systems using divergent sensor configurations, or when characteristics of self-recorded data need to be assessed. Moreover, ground truth data sets typically provide a means for quantitative evaluation, while human judgment may constitute the true benchmark for 3D models acquired in media production. Since objective error measures may not correlate with user perception, quantitative results need to be supplemented with subjective evaluations.

In this work, we present a framework suitable for evaluating self-recorded dynamic 3D models acquired by multiple depth sensors oriented toward a joint, central, spot of a scene. This setup is often used in media production where 3D models of an artist's performance are to be used as an element in a movie or a virtual environment (e.g. [18]). Our work relies on [3], where a more detailed discussion of the presented system can be found.

A major goal of our project was the incorporation of versatile evaluation strategies based on system-independent calibration objects, exploitation of multi-view data redundancy, and the computation of novel views for objective and subjective quality assessment. A key technical challenge was to design evaluation strategies that (i) can characterize the system's geometric accuracy, (ii) enable generic quantitative evaluation of intermediate and final products, and (iii) provide meaningful subjective results without pre-existing ground truth.

We apply the presented framework to evaluate the accuracy of 3D models that were acquired and generated with our 3D reconstruction pipeline, described in detail in [2, 3]. The tested processing pipeline is illustrated in Fig. 1. We use multiple synchronized stereo sensors to acquire pairs of image sequences individually per sensor. Subsequently, we generate depth maps using stereo matching algorithms. After semi-automatic segmentation of the filmed actors [17], we project the depth maps into 3D space as point clouds. Then, we fuse the point clouds into a volumetric grid representation with an APSS-based algorithm [21]. This approach is similar to [19]. Finally we use the marching cubes algorithm [22] for surface triangulation.

Dynamic point clouds are often acquired with RGB-D sensors [20] employing active stereo principles, such as structured-light, or time-of-flight. While convenient, available devices have fixed optical properties. Further, the algorithms which generate 3D data are embedded in these devices, which limits their applicability in general studio settings. In these cases, custom sensors [18] are suitable. Our 3D reconstruction pipeline uses passive stereo sensors to acquire stereo image sequences, from which we generate point clouds by applying stereo matching [9, 12] algorithms and projecting the resulting depth maps into 3D space.

Regarding subjective evaluation of 3D models, we consider the work of [6] the closest to ours. The authors assess the subjective impact of geometric mesh complexity and texture resolution. In contrast, our work suggests an evaluation framework for self-acquired 3D data. Subjective evaluation of point clouds is often performed on synthetic mesh models, such as [1]. Herein, we focus on self-acquired data.

While our tests have been carried out with stereo cameras, the proposed methodology can also be applied to setups with other sensors, for example RGB-D cameras.

Method

We design, implement, and test a framework to assess the performance of a 3D reconstruction system comprising multiple stereo cameras in terms of geometrical accuracy and perceived quality of its resulting 3D models. We determine the accuracy using two complementary evaluation strategies. (i) We compare acquired views of spherical and cuboid validation objects against their corresponding ideal shapes. (ii) We perform a subjective evaluation of dynamic 3D models. In these experiments, we assess user acceptance by carrying out a Pair Comparison-based [4] user study that considers the 3D nature of the shown material.

Objective Evaluation

Ideally, 3D reconstruction recovers the true metric properties of captured objects. Lengths and angles measured on real objects correspond to their reconstructed counterparts. The acquisition process and subsequent processing, however, introduce inaccuracies in resulting 3D models. Our objective evaluation aims to quantify the system's capability to reconstruct geometrically simple object surfaces. In this paper, we consider spherical bodies specifically. Another interesting type of validation object are cuboids, as they can be used to determine planar accuracy as well as the angular reconstruction fidelity. A detailed evaluation employing cuboid objects can be found in [3].

From a geometric perspective, spheres are simple and defined by their position in 3D space and radius. Filtering, for example, which is often employed in surface reconstruction algorithms, can lead to spherical reconstructions that appear "bumpy" instead of being round. A useful measure is the deviation of generated points from the true sphere surface.

We measure spherical reconstruction quality in terms of outliers, as they are a major obstacle for faithful reconstruction. In this context, outliers are 3D points which really belong to an object's surface, but whose reconstructed location is far away from the ideal sphere surface. Even a small number of outliers can cause unnatural deformations in reconstructed models.

We created a colorful spherical validation object of known radius. The texture was chosen to provide well-posed features for reconstruction algorithms. We then obtain reconstructions of the sphere placed at multiple positions within the system's working volume using a stereo matching algorithm [9]. Next, we fit ideal spheres of a radius corresponding to the validation object into the reconstructions with the Super4PCS algorithm. We measure the distance of each reconstructed point to the fitted sphere center, and measure the signed error to the known radius. Points that deviate from the sphere radius by more than a threshold value t in millimeter units are considered outliers. The ratio of outliers with respect to all 3D points is then used as a measure for geometric reconstruction accuracy.

Subjective Evaluation

The perceived quality of 3D models can be assessed using subjective quality evaluation in user studies [24]. For our subjective evaluation we adopt the Pair Comparison (PC) methodology [4]. In PC, pairs of stimuli (e.g. 3D models, videos, images, etc.) are shown to the participants. They express their preference for one item of a pair with an "A is better than B" or "B is better than A" choice. Due to the qualitative similarity of the compared approaches, participants are also offered a "no preference" option [8].

Compared Approaches

Our subjective evaluation aims to determine the influence of the choice of particular depth-generation strategies on the quality of the reconstructed mesh models. Further, we want to determine the qualitative impact of how the three separate single-view point clouds are fused into a single model. The compared approaches are summarized in Table 1 and described below.

Depth reconstruction algorithms. We compare three different stereo matching methods used to generate depth maps that are later projected into 3D space as point clouds. Two of them use cost volume filtering (CVF) [9]. The third method is a PatchMatch (PM) algorithm [12]. In Table 2 we summarize the algorithm parameters we use.

The *integer disparities (ID)* approach constitutes a base-line algorithm. Its output are integer-valued disparity values. Apart from enforcing total left/right consistency of the disparity maps, no further refinement is performed. The discrete nature of the ID disparity values causes the projected point cloud distances to lie on discrete slices in space. The *depth refined (DR)* approach also uses CVF, but in this case we additionally refine depth data. In addition to the left/right consistency check, we perform floating-point disparity refinement using parabolic fitting, which spreads reconstructed points more evenly in 3D space than the ID approach. The third method is a *PatchMatch (PM)* algorithm [12], which generates smoother disparity maps than the other approaches, due to the use of an energy function over slanted planes.

View fusion methods The effectiveness of the model generation step of our pipeline is assessed by comparing two strategies of view fusion. We either *fuse views after model generation (FA)* as a union of individual views, or *fuse views before model generation (FB)*. In the first case, surface reconstruction is performed on point clouds of individual depth sensors, which are merged into combined models afterwards. In the second case, individual point clouds are combined prior to surface reconstruction.

Table 1. Overview of compared approaches.

Abbr.	Description
ID	CVF [9] integer disparities, no subpixel refinement, no depth refinement
DR	CVF [9] floating-point disparities, subpixel refinement [10] and depth refinement [11]
PM	Patch Match [12] with floating point disparities
FA	Point clouds of individual views are fused after model generation
FB	Point clouds of individual views are fused before model generation

Table 2. Parameters used for point cloud generation approaches.

Algorithm Stage	Method	Parameter	Value
Cost computation	Census [13]	Window size	5×5
Cost aggregation	Permeability Filter [14] [9]	Permeability	$\sigma = 12$
Hier. Matching		Hierarchy levels	$0-2$
Post-processing	L/R consistency	Consistency threshold	0
Temporal filtering	Permeability Filter [14]	Permeability, Temporal consistency	$\sigma = 12$, $\sigma^t = 1$

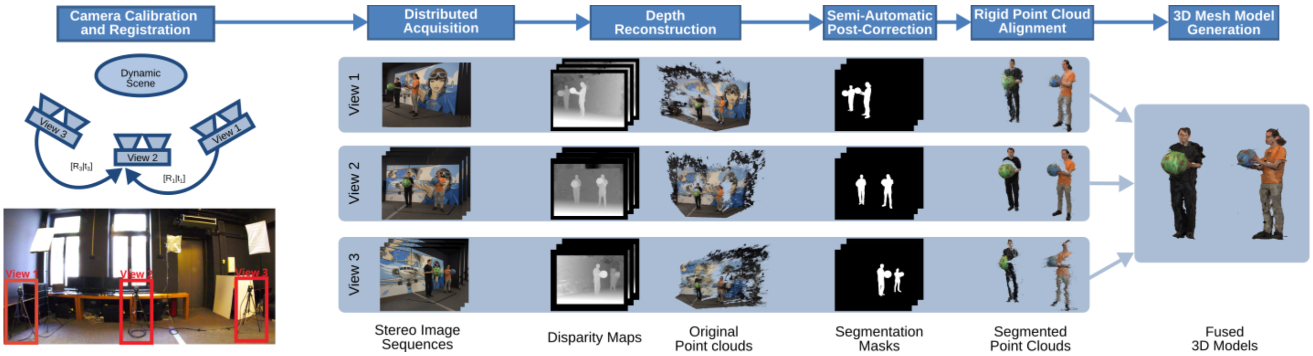

Figure 1. Illustration of the 3D reconstruction pipeline examined with our presented evaluation framework. Data is acquired using 3 passive stereo sensors.

Figure 2. Illustration of the mesh models used fur the subjective evaluation

Test Material Preparation

For our subjective evaluation we acquired five different dynamic scenes. Examples of our data set are shown in Fig. 2. Acquired raw data consisting of a sequence of stereo image pairs that we process into colored dynamic mesh models with our processing pipeline [2, 17]. Actors in the captured scenes are segmented from their background [17]. Next, we render them in front of an artificial background, showing them from a camera perspective that moves from the position of Sensor 1 to that of Sensor 3. We show pre-rendered videos because they guarantee the same visual impression for all participants and relieve them from having to navigate the shown material themselves. The 5 data sets are processed in 6 different variants consisting of 2 view fusion methods and 3 point cloud reconstruction algorithms, as previously described. In total we create 30 videos with a length of 15 seconds. In Pair Comparison, any approach is compared against all others. Thus, a comparison set for a single data set consists of $\binom{6}{2} = 15$ comparisons, resulting in 75 comparisons for all our 5 data sets. The net material show time is $75 \times 15 = 1125$ seconds or 18.75 minutes.

Test Environment and Participants

We have set up a lab environment according to the requirements of [5]. A view of the lab environment is shown in Fig. 3. We use an uncalibrated 27" display (ASUS VG278), set to a color

Figure 3. Example of the lab environment. Source: [3]

Table 3. Demographic information on the user study participants.

Age			Sex		Experienced	
Min	**Max**	**Mean**	**Male**	**Female**	**Yes**	**No**
19	59	36.6	9	13	11	11
Total participants			22			

temperature of 5000K, to present the material in full-HD resolution approximately 0.9 away from the participants. The room is subjected to controlled, dark illumination conditions while the trial is in progress. The material is shown with an application developed for this purpose.

Table 3 presents demographic information on the participants. In total, 22 persons (9 male, 13 female) participated in the main study. Half of them are trained in the evaluation of images or 3D models, while the other half considered themselves untrained. The participants' age ranges from 19 to 59 years with an average of 36.6 years. One participant failed the screening procedure due to limited colour perception, and was excluded from subsequent statistical processing.

Trials

A trial session is conducted with one participant at a time, and lasts for approximately 40 minutes. A session comprises five stages: (1) introduction, (2) screening, (3) practice, (4) experimental trial and (5) interview. Participants are briefly introduced orally (1) and are given written instructions explaining their task. In the following user screening (2), participants perform a visual acuity and a color vision test. After a short practice session (3), the experimental trial (4), illustrated in Fig. 4, is carried out. While a pair of videos (A, B) is shown, the participant can switch between video A and video B by pressing a mouse button. After 15 seconds, a voting screen appears, and the user is asked to make their judgement. The PC method calls for the presentation of both stimuli in both orders, that is AB and BA for a pair of stimuli (A, B) [4]. We modify this approach by showing one model at a time in full-screen, while allowing the participants to switch between stimuli A and B as they choose. After finishing the evaluation procedure, participants are interviewed (5) and fill out a questionnaire.

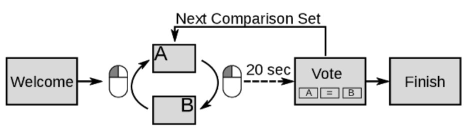

Figure 4. Schematic illustration of the pair comparison-based trial process. Source: [3]

Processing

The trial data is screened for biased opinions using circular triad detection [15]. Circular triads are triples of contradictory stimulus preferences. No such votes were detected at all, thus the opinions of all participants are used for further processing. Next, the combined comparison data is separated into depth reconstruction and view-fusion approaches. Opinion data is grouped according to the type of respective compared approach. The separated data is processed into continuous opinion scores with the Bradley-Terry model [16], that estimates the maximum likelihood of a log-likelihood probability distribution according to the given opinions. The opinion scores are grouped by data set and compared approach and constitute the evaluation result.

Results of the Objective Evaluation

In our experiments, we use a validation sphere with 150 mm radius. We obtain point clouds with our processing pipeline (c.f. Fig. 1) with the DR depth-reconstruction approach described above and in Table 1. The sphere test data set comprises 29 reconstructions of blue and green spheres. Out of these, 26 samples are used. Three samples were excluded due to larger clusters of erroneous depth values biasing the fitting procedure. Sphere point cloud samples range from 0.4 to 2.1 million points, with an average point count of approximately 1 million points. The distance from depth sensor to sphere ranges from 2.3 to 2.9 meters.

An example of a sphere fitted into the point cloud is shown in Fig. 5. The reconstruction result is shown from a sensor's perspective as well as from the side. A sphere fitting result is shown overlaid. The right column shows outliers visualized at an outlier threshold of $t = \pm$ 12.5 mm.

In our quantitative evaluation of the geometrical system accuracy, shown in Fig. 6, we vary the outlier threshold t from 5 to 100 mm. At $t = \pm$ 2.5 mm, only 31.3 percent of the points are counted as inliers. The inlier ratio goes up to almost 90 percent at a $\pm$ 12.5 mm threshold, and reaches 99 percent at $\pm$ 55 mm. Table 4 summarizes the sensor accuracy of the individual depth sensor. Sensor 2 is the most accurate with 97.41 percent, the least accurate is Sensor 1 with 93.76 percent, while Sensor 3 has an inlier ratio of 96.09 percent. Note that the accuracy of individual sensors is slightly higher than that of combined point clouds.

The outlier ratio can give a notion of how small an object can be for a reconstruction of meaningful quality. In Fig. 5, we can see that reconstructions reach a 99 percent inlier ratio at a $\pm$ 50 mm threshold. This is acceptable for large objects without fine surface details. To reconstruct smaller objects like a human nose, a higher accuracy is needed. An improvement could be achieved by increasing the disparity range and, hence, depth resolution of acquired disparity maps by either increasing the depth sensor's stereo baseline or decreasing the object-to-camera distance.

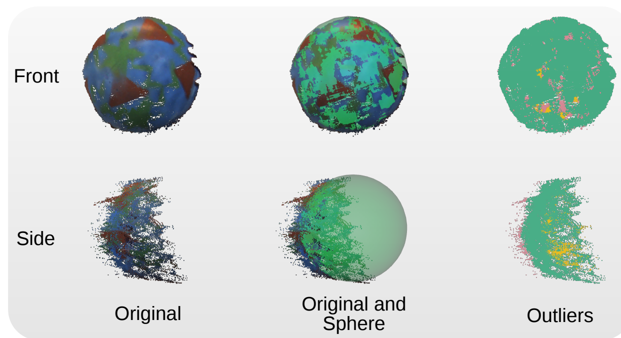

Figure 5. Illustration of the reconstructed spherical validation object overlaid with the fitted ideal sphere. Source: [3]

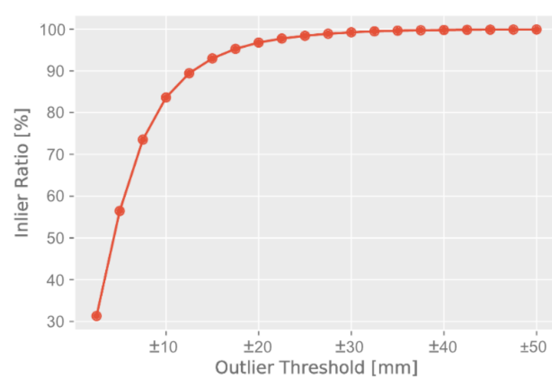

Figure 6. Quantitative results of the evaluation on spherical objects. Graph is given in terms of absolute threshold for both sides of the sphere surface. Source: [3]

Results of the Subjective Evaluation

In our subjective evaluation, we assess the the impact of different depth generation algorithms, and view fusion methods on the resulting dynamic 3D models, as described above. In this section, we present the results of our subjective user study that compares these approaches.

Table 4. Results of the objective evaluation on spherical objects for individual sensors at an outlier threshold $t = \pm 12.5$ mm.

Sensor ID	Sensor 1	Sensor 2	Sensor 3
Inliers (%)	93.76	97.41	96.09
Outliers (%)	6.24	2.59	3.91

Depth reconstruction algorithms. The results shown in Fig. 7 indicate a clear preference for PM, with DR second, and ID last. The ratings are consistent among the majority of models, namely *ball*, *people*, *standing*, and *face*. Only for the *box* data set ID was ranked before DR. Models generated with ID exhibit visible seams where views are fused, contrary to models generated with DR and PM. The comparatively low spatial resolution of ID point clouds leads to model fusion not properly aligning separate views. Participant's comments indicate that especially the appearance of faces had a significant impact on given opinions. Models generated with the PM method were presented as static models due to the high computational time needed. This fact needs to be taken into consideration in regard to the interpretation of the high scores for PM models. For those models, temporal artifacts, such as flickering and global color changes occurring over time could not be evaluated.

View fusion methods. Results of the view fusion evaluation are shown in Fig. 8. View fusion before model generation (FB) is preferred over view fusion after model generation (FA). The result can be explained by the fact that view fusion after model generation (FA) normalizes normal vectors among individual views, causing unnatural model texture appearance. View fusion before model generation (FB) on the other hand, performs normal estimation jointly for the fused point cloud that serves as basis for the final mesh. Textures are more blurred in FA than in FB variants, since in FB the model generation algorithm drops points from spatially crowded regions, leading to lower texture resolution. In FA both views constituting the fused model are processed into mesh models separately, hence more 3D points remain in the final fused model, due to the inherent point sub-sampling of the model generation algorithm.

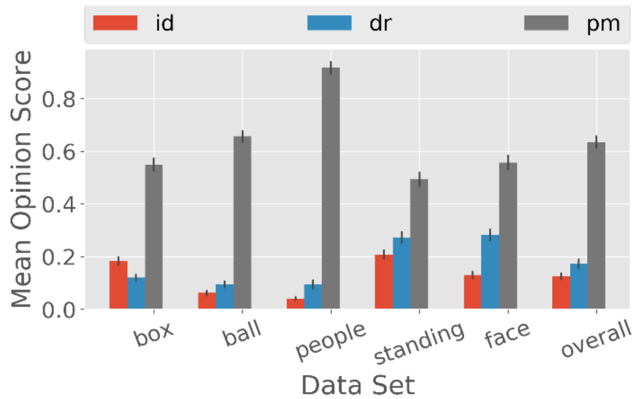

Figure 7. Subjective evaluation results grouped by different point cloud generation approaches: 3D models created with cost-volume filtering with integer disparities (id), cost-volume filtering with floating point disparity refinement (dr) and PatchMatch (pm).

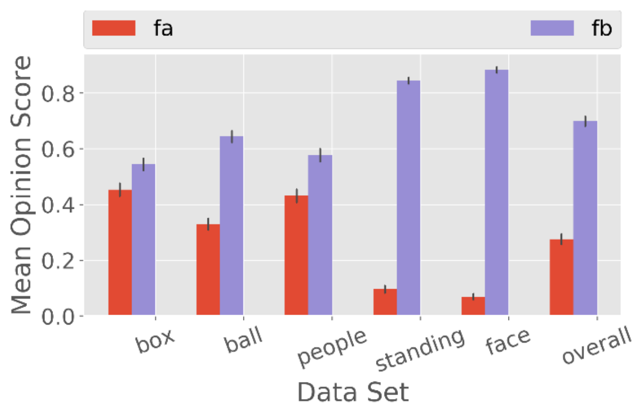

Figure 8. Results for two different view fusion methods: fusion of views before surface reconstruction (fb) and after surface reconstruction (fa)

Conclusion

In this paper, we have presented and tested a framework for objective and subjective evaluation of self-acquired dynamic 3D point clouds. Our objective evaluation aims to quantify the geometric reconstruction accuracy by leveraging the geometrically simple properties of spherical calibration objects to assess reconstructions in terms of the inlier ratio relative to a threshold distance around the ideal sphere radius. We have found that the tested system achieves an inlier ratio of 90 percent at a threshold of $t = \pm 12.5$ mm.

Our subjective evaluation, which comprises a Paired Comparison user study on dynamic colored mesh models, assesses the qualitative impact of several 3D model generation strategies on user acceptance. When testing 3 different depth-generation and 2 view-fusion strategies in a user study with 22 participants, we found that the PatchMatch (PM) strategy is the most well received. Further, view fusion of the point clouds before model generation (FB) was considered better than view fusion after model generation (FA), due to the reconstructed surface normals being consistent in regions belonging to fused views.

We believe that the presented framework for evaluation of self-acquired 3D models is an adequate approach to characterize geometric reconstruction accuracy and the model quality in terms of user acceptance. Our approach is not only applicable for systems employing stereo sensors, but is also useful for setups with other sensing modalities.

Acknowledgements

This work has been funded by the Austrian Research Promotion Agency (FFG) and the Austrian Ministry BMVIT under the program ICT of the Future (project "Precise3D", grant no. 855442)

References

[1] S. Seitz, B. Curless, J. Diebel, D. S., and R. Szeliski, A comparison and evaluation of multi view stereo reconstruction algorithms, IEEE CVPR, pp. 519-528, (2006).

[2] C. Kapeller, B. Sespede, M. Nezveda, M. Labschütz, F. Seitner, and M. Gelautz, A work-flow for 3D model reconstruction from multi-

view depth acquisitions of dynamic scenes, OAGM Workshop, pp. 116–119, (2018).
[3] C. Kapeller, Evaluation of a 3D reconstruction system comprising multiple stereo cameras, Dipl. Thesis, Vienna University of Technology, Austria, (2018).
[4] ITU, Subjective video quality assessment methods for multimedia applications, Rec. P.910, (2008).
[5] ITU, Methodology for the subjective assessment of the quality of television pictures, Rec. BT.500, (2012).
[6] A. Doumanoglou, N. Zioulis, E. Christakis, D. Zarpalas, and P. Daras, Subjective quality assessment of textured human full-body 3D-reconstructions, Int. Conf. QoMEX, pp. 1–6, (2018).
[7] M. Nezveda, N. Brosch, F. Seitner, and M. Gelautz, Depth map post-processing for depth-image-based rendering: a user study, SPIE 9011, Stereoscopic Displays and Applications XXV, pp. 170–178, (2014).
[8] J. S. Lee, L. Goldmann, and T. Ebrahimi, Paired comparison-based subjective quality assessment of stereoscopic images, Multimed. Tools Appl., 67, 1, pp. 31–48, (2013).
[9] F. Seitner, M. Nezveda, M. Gelautz, G. Braun, C. Kapeller, W. Zellinger, and B. Moser, Trifocal system for high-quality inter-camera mapping and virtual view synthesis, IC3D, pp. 1–8 (2015).
[10] Q. Yang, R. Yang, J. Davis, and D. Nister, Spatial-depth super resolution for range images, IEEE CVPR, pp. 1–8, (2007).
[11] Z. Ma, K. He, Y. Wei, J. Sun, and E. Wu, Constant time weighted median filtering for stereo matching and beyond, IEEE ICCV, pp. 49–56, (2013).
[12] L. Li, S. Zhang, X. Yu, and L. Zhang, PMSC: PatchMatch-based superpixel cut for accurate stereo matching, IEEE T-CSVT, 28(3), pp. 679–692, (2018).
[13] R. Zabih and J. Woodfill, Non-parametric local transforms for computing visual correspondence, ECCV, pp. 151–158, (1994).
[14] C. Çığla and A. Alatan, Information permeability for stereo matching, Signal Process. Image., 28, 9, pp. 1072–1088, (2013).
[15] J. S. Lee, L. Goldmann, and T. Ebrahimi, Paired comparison-based subjective quality assessment of stereoscopic images, Multimed. Tools Appl., 67, 1, pp. 31–48, (2013).
[16] R. A. Bradley and M. Terry, Rank analysis of incomplete block designs: I. The method of paired comparisons, Biometrika, 39, 3/4, pp. 324–345, (1952).
[17] B. Sespede, F. Seitner, and M. Gelautz, Semi-automatic post-processing of multi-view 2D-plus-depth video, Electronic Imaging, 2019, 3, pp. 657-1–657-5, (2019).
[18] D. Mingsong, J. Taylor, P. Kohli, V. Tankovich, et al., Fusion4D: Real-time performance capture of challenging scenes, ACM Trans. Graph., 35, 4, pp. 1–13, (2016).
[19] B. Curless, M. Levoy, A volumetric method for building complex models for range images, SIGGRAPH, pp. 303–312, (1996).
[20] H. Sarbolandi, D. Lefloch, and A. Kolb, Kinect range sensing: Structured-light versus Time-of-Flight Kinect, Computer Vision and Image Understanding, 130, pp. 1–20, (2015).
[21] G. Guennebaud, and M. Gross, Algebraic point set surfaces, ACM Trans. Graph., 26, 3, pp. 37–49, (2007).
[22] W.E. Lorensen, H.E. Cline, Marching cubes: A high resolution 3D surface construction algorithm, Computer graphics and interactive techniques, pp. 163–169, (1987).
[23] K. Fliegel, F. Battisti, M. Carli, M. Gelautz, L. Krasula, P. Le Callet and V. Zlokolica, 3D Visual Content Datasets, 3D Visual Content Creation, Coding and Delivery, Springer, Cham, pp. 299–325, (2018).
[24] M. Berreda-Ángeles, F. Battisti, G. Boato, M. Carli, E. Dumic, M. Gelautz, C. Hewage, D. Kukolj, P. Le Callet, A. Liotta, C. Pasquini, A. Pereda-Baños, C. Politis, D. Sandic, M. Tekalp, María Torres-Vega and V. Zlokolica, Quality of Experience and Quality of Service Metrics for 3D Content, 3D Visual Content Creation, Coding and Delivery, Springer, Cham, pp. 267–297, (2018).

Author Biography

Christian Kapeller *received his MSc in Computational Intelligence from the Vienna University of Technology (2018). Since then he has worked at the High Performance Vision group of AIT Austrian Institute of Technology, Vienna as junior scientist, as well as the Institute of Visual Computing & Human-Centered Technology at the Vienna University of Technology. His main research areas are computational imaging and 3D reconstruction.*

Braulio Sespede *received his MSc in Software Engineering from the Buenos Aires Institute of Technology (2018). His research interests lie in the domain of visual computing, more specifically 3D imaging and interactive computer graphics.*

Matej Nezveda *received his MSc in computer science from the Vienna University of Technology (2014). Since then he has worked at emotion3D at the intersection of computer vision, AI, machine learning, and graphics with particular emphasis on interior monitoring systems.*

Matthias Labschütz *received his MSc in Computer Science from the Vienna University of Technology (2016). He has spent time at the King Abdullah University of Science and Technology (2015) and at the University of Bergen (2016). In 2017 he started working at Rechenraum GmbH. His current work focuses on registration, mesh reconstruction, geometry processing, animation and rendering.*

Simon Flöry *received his PhD in Applied Mathematics and Geometry from Vienna University of Technology (2010). In 2012 he founded the Austrian high-tech company Rechenraum GmbH, providing efficient yet robust cutting edge 3D software to process, analyze and model 3D data. The company's 3D/CAD/CAM solutions are trusted by customers in a wide field of applications, ranging from dental technology and industrial production to 3D sensors.*

Florian Seitner *received his Dipl.Ing. and Dr.techn. degrees in Computer Science from the Vienna University of Technology. After working as a system designer on embedded video processing solutions at On Demand Microelectronics, he co-founded emotion3D, a company specialized on AI-based camera solutions for automotive in-cabin monitoring. Initially acting as CTO, Florian became CEO of emotion3D in 2014 and is responsible for the company's corporate and business strategy.*

Margrit Gelautz *received her PhD in Telematics from Graz University of Technology (1997). She is an associate professor at Vienna University of Technology with a focus on image and video analysis & synthesis techniques. Her current research interests include 3D reconstruction for film, human-robot interaction and autonomous driving applications.*

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

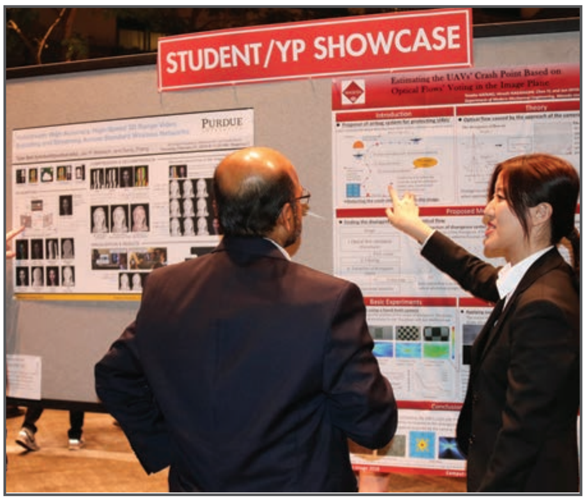

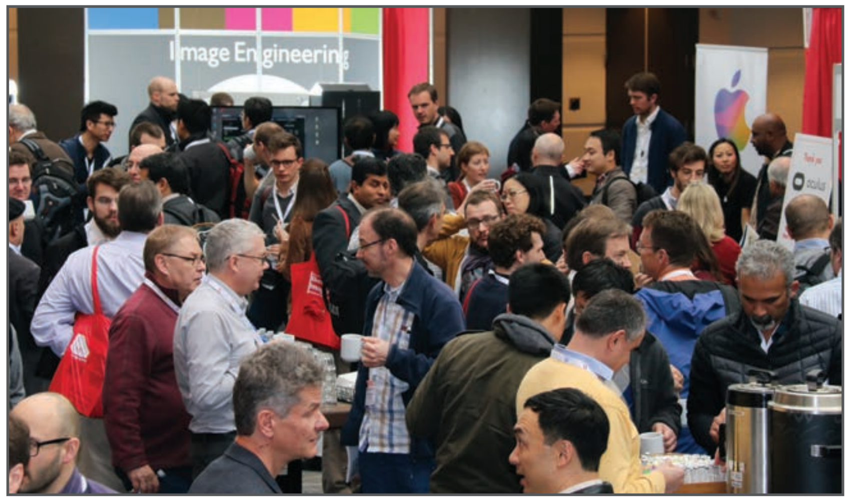

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org