A Comparative Error Analysis of Current Time-of-Flight Sensors

Peter Fürsattel, Simon Placht, Michael Balda, Christian Schaller, Hannes Hofmann, Andreas Maier, and Christian Riess

***Abstract*—Time-of-flight (ToF) cameras suffer from systematic errors, which can be an issue in many application scenarios. In this paper, we investigate the error characteristics of eight different ToF cameras. Our survey covers both well established and recent cameras including the Microsoft Kinect V2. We present up to six experiments for each camera to quantify different types of errors. For each experiment, we outline the basic setup, present comparable data for each camera, and discuss the respective results. The results discussed in this paper enable the community to make appropriate decisions in choosing the best matching camera for a certain application. This work also lays the foundation for a framework to benchmark future ToF cameras. Furthermore, our results demonstrate the necessity for correcting characteristic measurement errors. We believe that the presented findings will allow 1) the development of novel correction methods for specific errors and 2) the development of general data processing algorithms that are able to robustly operate on a wider range of cameras and scenes.**

***Index Terms*—Computer vision, time-of-flight cameras, range imaging.**

I. Introduction

In the last years multiple time-of-flight depth cameras have established themselves in academic, industrial and entertainment environments. Especially the PMD CamCube 3.0[1] and the Mesa SwissRanger cameras[2] have been studied extensively. More recently new ToF camera models have been presented to the industrial and consumer market, for example the Microsoft Kinect V2 [3], Fotonic[4] cameras or Bluetechnix cameras[5], which did not receive much attention in state-of-the-art publications. Even though ToF cameras are already used in commercial applications they still suffer from various systematic and non-systematic errors. In previous works, multiple groups presented correction methods which allow to compensate for some error sources. However, most methods have only been evaluated only on a single camera model or on two camera models with similar error characteristics. In this work, we present an extensive evaluation of multiple state-of-the-art ToF cameras with respect to several effects which have been observed in the past and point out the differences between the evaluated camera models. To the best of our knowledge, there is no comparable review of established, well-known and recent time-of-flight sensors.

The target audience of this work includes researchers, product- and application developers who seek to choose the right sensor for their application. Furthermore, the information presented in this study is valuable as a source of information for two types of research projects. First, it supports the development of correction methods for particular cameras. Second, it supports the development of data processing algorithms that are robust across one or more cameras to one or more of the various influences examined here.

Beyond that, we aimed to carefully assemble a benchmark protocol that is reusable, such that comparison results for new cameras can be added by evaluating them in the same manner. On the algorithmic side, this benchmark can help to quantitatively assess the effectiveness of new correction algorithms by following the protocol for a camera with the proposed correction in place.

II. Related Work

In the past years multiple publications on ToF errors and error correction methods have been presented. However, ToF technology is a quickly moving field. Some errors which have been observed in the past are now less easy to identify or cannot be found anymore with more recent ToF cameras. This is due to improved camera hardware or due to correction methods which are either implemented in hardware or in manufacturer-provided software. As a consequence, some of the former state-of-the art correction methods became meanwhile partly or completely ineffective.

Many correction methods from the literature rely on a physical model which represents potential error sources like, for example, the so-called wiggling error. Sinusoidal functions are commonly used for compensating this error (for example [1]).

Manuscript received March 24, 2015; revised October 20, 2015; accepted December 04, 2015. Date of publication December 18, 2015; date of current version March 01, 2016. This work was supported in part by the German Federal Ministry of Education and Research as part of the Spitzencluster Medical Valley program 13GW0029A, in part by the Research Training Group 1773 "Heterogeneous Image Systems", funded by the German Research Foundation (DFG), and in part by the Erlangen Graduate School in Advanced Optical Technologies (SAOT) by the German National Science Foundation (DFG) in the framework of the excellence initiative. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Neel Joshi.

P. Fürsattel, A. Maier, and C. Riess are with the Pattern Recognition Laboratory, University of Erlangen-Nuremberg, Erlangen 91058, Germany (e-mail: peter.fuersattel@fau.de; andreas.maier@fau.de; christian.riess@fau.de).

S. Placht, M. Balda, C. Schaller, and H. Hofmann are with Metrilus GmbH, Erlangen, Germany.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCI.2015.2510506

[1] pmdtechnologies GmbH, Siegen, Germany

[2] Mesa Imaging, Zuerich, Switzerland

[3] Microsoft, http://www.microsoft.com/en-us/kinectforwindows

[4] Fotonic, Stockholm, Sweden

[5] Bluetechnix Group GmbH, Wien, Austria

In this work we will show measurement data which is not affected by this error any more. Chow *et al.* [2] present a correction method that the authors evaluate with the PMD CamBoard nano. However, this work only presents an evaluation of the overall error with no detailed evaluation of all the known error sources.

Multiple authors point out that a significant amount of the measurement errors is introduced by internal scattering of light within the camera [3]–[6]. Jamtsho *et al.* further investigated this effect and claim that the scattering is independent of the camera's integration time [7]. Many of the proposed scattering compensation methods have been evaluated with the Mesa SR3000, of which already three successors exist up to date, the SR3100, SR4000 and SR4500. Chiabrando *et al.* [8] claim that the SR4000 does not suffer from internal scattering anymore. In this work, we demonstrate that internal scattering is still present with this camera. Furthermore, we show that scattering can be observed with all state-of-the-art ToF sensors.

In [9] a comparison of the Mesa SR4000 and the PMD CamCube 3.0 is presented. The authors provide an evaluation with respect to warm-up time, integration time and measurement accuracy.

Several groups identified error sources which are related to the intensity or reflectivity of the scene, for instance Pfeifer *et al.* [3] and Lindner *et al.* [10]. These errors are also known as amplitude-related errors and investigated in one of our experiments.

Finally, there also exist publications which simply analyze the measurement accuracy of different depth sensors like for example Stoyanov *et al.* [11]. The authors of this work present a comparison between the well-known Mesa SR4000, a Fotonic B70 sensor and the first Kinect sensor, which is a structured light device. In this work, we present comparisons between established ToF sensors and the more recent Kinect V2 sensor which, in contrast to its predecessor, is based on time-of-flight technology. Instead of analyzing the absolute measurement accuracy, which is influenced by multiple error sources, we evaluate systematic errors separately.

Similar work has been presented by Rauscher *et al.* [12]. The authors compare multiple structured-light and ToF sensors for robotics applications, focusing on accuracy and noise of distance measurements. In our work, we set the focus on a thorough study of ToF sensors with an in-depth analysis of their characteristic errors.

Recently, evaluations of the Kinect V2 sensor [13], [14] and comparisons to its structured light-based predecessor have been presented [15]. While both works provide a thorough study of this camera, they do not compare the performance of the Kinect V2-ToF sensor with other established cameras.

III. Investigated Error Sources

We investigate six dominant error sources of ToF cameras which we consider typical for this imaging modality. The types of errors depend on the camera model, and not primarily on the scene. We extend the list of systematic errors of Foix *et al.* [16] by internal scattering as done by Karel *et al.* [4]. This is based on the insight that the occurrence of internal scattering largely depends on camera characteristics.

Scene-dependent errors like multi-path interference are not addressed in this study. A standardized, fair, and meaningful evaluation of such effects is a large work in its own right. For more information, the reader is referred to, for example, [17]–[19]. Beyond that, our rationale for omitting scene-dependent errors is that the distortion of scene-dependent errors are caused by (potentially varying) external objects. As such, these errors occur in addition to the systematic errors.

A. Temperature Related Errors

A well-known error is introduced by the internal temperature of the camera itself. As the camera captures images, its illumination unit and its sensor heat up. With the increase in temperature the characteristics of the components change, leading to a temperature-dependent drift of distance measurements. The length of the warm-up period differs between camera models. Some cameras do compensate for internal warm-up effects, while others do not.

B. Temporal Variation

The measurement of a static pixel can vary significantly over time, caused by measurement noise. This error source becomes most noticeable if the observed surface has low reflectivity. Therefore, little light is reflected back to the camera, which results in a low signal-to-noise ratio.

C. Integration Time

The time span in which the ToF sensors detect photons for calculating a single depth image is called integration time and is similar to the exposure time of regular cameras. During this time span multiple phase measurements are performed for each pixel. In a second step, the final distance measurement is computed. It has been observed that changing the integration time has influence on the absolute measured distance, but this error also exhibits a temporal component (e.g. [4]). In this work, we demonstrate and measure these effects for cameras with adjustable integration time.

D. Internal Scattering

Several groups reported internal scattering of light between the sensor and the lens of the camera [4]–[6]. This effect induces a strong influence of high-reflectivity areas on low-reflectivity areas, especially when the respective areas are at different distances. The magnitude of this error can range from a couple of centimeters to more than a meter.

E. Amplitude Related Errors

Another effect which occurs with time-of-flight cameras is the so-called amplitude related distance error which is directly related to the amount of incident light that reaches a pixel.

When capturing low-reflectivity, poorly illuminated or distant objects, only few photons hit the respective pixels. This can deteriorate the accuracy of the distance measurements. This error can be visualized best when looking at a plane with strong variations in reflectivity, for example a checkerboard. Pixels which represent dark quads typically deviate from the true checkerboard plane by up to multiple centimeters. While the true origin of this effect is unknown, it is believed that it originates from non-linearities of the sensor which are introduced during the read-out phase [20].

F. Wiggling

Some cameras suffer from a nonlinear distance error or wiggling error. It occurs due to imperfect generation of the modulated light which is sent out into the scene. Typically, a sinusoidal shape of the emitted light is assumed when computing the distance image. The deviations from the ideal sinusoidal shape lead to a periodically oscillating distance error [21].

IV. Evaluation

The evaluation is performed in a dark room, excluding any daylight or sources of infrared light. There were no persons present during the experiments. In all experiments, the camera is mounted on a tripod and oriented towards a low-reflectivity molleton curtain which covers the whole field of view. All unnecessary objects close to the camera are removed in order to minimize multi-path effects. Relevant experiment-specific distances and dimensions are listed in the respective sections. Except for the warm-up time experiment all experiments are performed after a warm-up time of two hours in order to eliminate warm-up effects of the camera. Integration times are set to fixed values and auto-exposure functionality is deactivated before the camera warms up. Whenever possible internal filters of the cameras are deactivated during the experiments.

Note that we treat cameras as black boxes and do not use any custom external light sources. Employing an external illumination unit would also make a comparison between the Kinect V2 and the other cameras impossible, as the Kinect V2 uses multiple modulation frequencies. While using a single external light source would allow a better comparison of the sensors itself, such a setup would not be suitable to compare cameras as they can be purchased. Furthermore, not all cameras which are studied in this evaluation allow the usage of external light.

In this work we investigate the error characteristics of eight ToF cameras. Some of the evaluated cameras have already been studied in previous applications and will serve as a baseline for comparisons with more rarely used cameras. We also present error analyses for cameras which, to the best of our knowledge, have not yet been published. The complete list of cameras, together with their release date and their controllable parameter are given in Table I. In the remainder of this work we will use t_i to denote the integration time and f_m to describe the modulation frequency.

All experiments are performed with each camera, except for the analysis of integration-time-related errors, as some of the cameras do not offer the possibility of adjusting this parameter.

TABLE I
Evaluated Cameras and Their Controllable Parameters

Model	Year	used f_M	f_M	t_i	Filters
PMD CamCube 3.0	2010	20 MHz	✓	✓	✗
Mesa SR3100	2005	20 MHz	✓	✓	✓
Mesa SR4000	2008	20 MHz	✓	✓	✓
Creative Senz3D	2013	30 MHz	✗	✗	✗
Microsoft Kinect V2	2013	multiple	✗	✗	✗
Fotonic E70	2012	15 MHz	✗	✓	✓
Bluetechnix Argos P100	2013	30 MHz	✓	✓	✓
PMD CamBoard nano	2012	30 MHz	✓	✓	✓

Upon acceptance of the paper, all data of the experiments will be publicly available for download[6].

In the following sections we will outline the setup for each single experiment and discuss the respective results.

A. Warm-Up Time

Multiple publications report that the distance measurements of ToF sensors change over time, even though neither the scene nor the camera perspective is altered (e.g. [3]). This variation must be avoided when measuring camera characteristics like wiggling, internal scattering or other effects. Thus, we evaluate in the first experiment the variance of depth measurements with respect to time. We capture distance data for 120 minutes, under the assumption that the sensor heats up when continuously capturing images. To avoid temporal noise effects, the distance of a 9×9 pixels window at the center of the image is averaged over one second. The integration time is set for each camera separately with respect to two criteria: the integration time should be as long as possible to ensure accurate distance measurements, and second, no foreground pixels should be saturated, which serves as a upper bound for the integration time. The camera is oriented towards a white board perpendicular to the optical axis. Depending on the opening angle of the lens the distance to the foreground is chosen such that the board covers the whole field of view. The results are shown in Fig. 6. For plotting the results the measured distances of the last 10 minutes are averaged and considered as the measurement in the heated or steady state of the camera. All distance measurements are normalized by this final, averaged distance to highlight the variation over time.

We begin with the evaluation of the CamCube 3.0 to illustrate the influence of the integration time t_i on the warmup behavior. Fig. 1 shows two warmup curves with two different integration times, 100 µs (dashed) and 800 µs (solid). Clearly the total drift depends on the integration time which would have to be considered when correcting this offset. After approximately 30 minutes the difference between the drifts becomes less than 2 millimeters. The influence of the integration time on the length of the warm-up period is difficult to quantify due to the small differences between the measurement errors at timestamps larger than 30 minutes. This plot is shown exemplarily to demonstrate this correlation which is present for other sensors as well. For the sake of completeness the plot

[6]http://www.metrilus.de

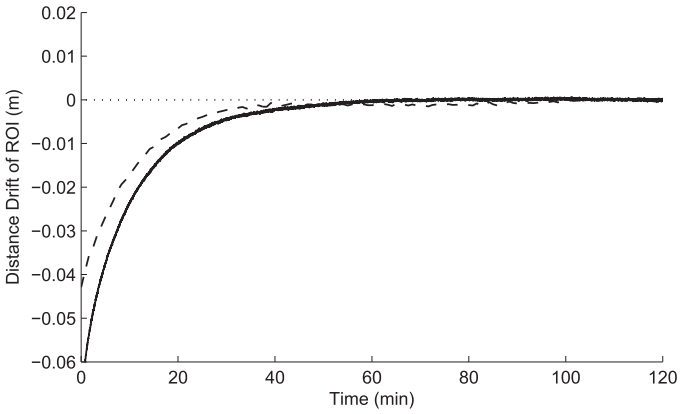

Fig. 1. Averaged distance drifts for the CamCube 3.0 sensor for two integration times (dashed: 100 µs, solid 800 µs). The dotted line represents the steady state distance by which all previous measurements have been normalized.

for $t_i = 800$ µs is also shown in Fig. 6a next to the results of the other cameras. In a later experiment we will focus on integration-time-related errors. Note that this effect might also be related to the amplitude of the incident light as shown in a later experiment.

The data captured with the SR3100 is illustrated in Fig. 6b. The sensor's auto-exposure feature suggests an integration time of 5400 µs, however we have set t_i to 1000 µs as lowering the integration time reduces the impact of fixed pattern noise drastically. Again there exists a drift in the measured distance values which, in comparison to the CamCube 3.0, already begins to stabilize after a warm-up period of approximately 5 minutes. Furthermore there exists a noticeable oscillation of the average distance due to camera's internal fan turning on and off.

The average distance plot for the successor of the SR3100, the SR4000, is shown in Fig. 6c. Again we choose a t_i which is lower than the value the automatic exposure mechanism would suggest. In this experiment t_i is set to 1000 µs. This camera does not come with active cooling which could reduce the overall warm-up time. The sensor can be considered to be warmed up after approximately 40 minutes. The total temperature-related drift is less significant than with the SR3100.

Due to the jagged pattern of the distance curve of the Senz3D (see Fig. 6d) one may assume that this camera performs an internal measurement correction based on a lookup table with correction terms for integer temperature values. Under the assumption that the warm-up process of the camera slows down with increasing time, the increasing length of the intervals between different correction values appears plausible. After approximately 30 minutes the distance measurements do not differ from the final distance measurement significantly anymore. Nonetheless, the last adaptation of the correction term happens after 45 minutes.

The Kinect V2 sensor shows only little variation over time (see Fig. 6e). Even right after powering up the sensor the difference to the steady-state measurement is less than two millimeters. After approximately 15 minutes of warming up the distance to the foreground does not change significantly anymore.

The E70 sensor shows a significant drift which does not stabilize before warming up for approximately 100 minutes. Furthermore, the distance measurements drift by more than 4 cm during this period. However, the true warm-up time is difficult to estimate with the current data as the plot shows a similar jagged shape as with the Senz3D sensor. Note the different scale of the y-axis of the plot when comparing it to the other camera models. For evaluating the warm-up time, the integration time of the sensor has been set to 400 µs.

The Argos P100 does not contain a fan for cooling. Nonetheless, the camera warms up comparably fast or internal correction routines ensure that measured distances do not drift after an initial warm-up phase of approximately 4 minutes, as can been seen in Fig. 6g. The warm-up time has been determined with t_i set to 2000 µs.

The last plot shown in this experiment illustrates the distance measurements of the CamBoard nano. Compared to the other cameras this camera shows significant measurement noise with a variation of up to one centimeter. The reason for these high noise levels might be poor illumination due to the distance between the camera and the foreground, as this camera has been designed for gesture-applications and not for long distance range measurements. With an integration time of 1000 µs the warm-up period totals to approximately 18 minutes.

B. Temporal Noise

Once the cameras have warmed up further experiments can be performed. The goal of this experiment is to determine the influence of scene reflectivity on the measurement noise. Due to a standardized setup a comparison between the cameras is possible.

We assume the measurement noise to be a combination of multiple noise effects, including Poisson or shot noise, dark current and others. In this experiment we evaluate the sum of all these effects as this is the noise that the user of a ToF camera is able to observe. In the following we demonstrate that the evaluated cameras differ with respect to their noise characteristics and show that internal filters can be used to reduce the influence of noise.

To measure the temporal noise we use a similar setup as in the previous experiment. The distance between camera and foreground equals approximately 80 cm, measured with the camera itself when looking at a white paper foreground. The integration time is set by the same criteria as in the previous experiment. The noise is measured by the mean value and the standard deviation of 1000 consecutive measurements at the central pixel.

Additionally the same measurement is performed with a gray and a black paper foreground without altering the integration time or distance to the camera to evaluate the dependency on the foreground's reflectivity. With some cameras it is possible to control internal temporal or adaptive noise filtering. In Table II we present results for all cameras without filtering or with internal filtering disabled. Table III shows the temporal noise evaluation with noise filtering enabled and for cameras which do not allow to control the filters.

TABLE II
TEMPORAL NOISE EXPERIMENT WITH INTERNAL FILTERS DISABLED

Model	White (m)	Gray (m)	Black (m)
CamCube 3.0	0.800 ± 0.007	0.787 ± 0.008	0.761 ± 0.044
SR3100	0.806 ± 0.019	0.801 ± 0.023	0.807 ± 0.165
SR4000	0.800 ± 0.005	0.790 ± 0.006	0.781 ± 0.034
Senz3D	0.798 ± 0.009	0.789 ± 0.010	0.792 ± 0.051
E70	0.800 ± 0.008	0.780 ± 0.008	0.771 ± 0.029
Argos P100	0.801 ± 0.008	0.791 ± 0.009	0.806 ± 0.050
CamBoard nano	0.800 ± 0.002	0.784 ± 0.003	0.760 ± 0.025

TABLE III
TEMPORAL NOISE EXPERIMENT WITH INTERNAL FILTERS ENABLED

Model	White (m)	Gray (m)	Black (m)
SR3100	0.800 ± 0.008	0.793 ± 0.009	0.790 ± 0.054
SR4000	0.800 ± 0.002	0.792 ± 0.002	0.788 ± 0.009
Senz3D	0.797 ± 0.004	0.792 ± 0.004	0.795 ± 0.020
Kinect V2	0.799 ± 0.001	0.792 ± 0.001	0.797 ± 0.002
E70	0.798 ± 0.001	0.784 ± 0.002	0.777 ± 0.001
Argos P100	0.800 ± 0.002	0.791 ± 0.002	0.780 ± 0.014
CamBoard nano	0.794 ± 0.002	0.776 ± 0.003	0.767 ± 0.031

As expected the standard deviation increases as the foreground gets darker, resulting in a smaller signal to noise ratio (SNR). This characteristic can be observed with all evaluated cameras and is independent on whether any filters are active or not. However, the magnitude of temporal noise is clearly influenced by these filters as can be seen when comparing the standard deviation for cameras which allow the user to control the filters.

Note that there exists a relation between modulation frequency and temporal noise [22] which is not investigated in this work. By increasing the modulation frequency the SNR can be increased. While this may hold for a single camera with an adjustable modulation frequency one cannot assume that a higher f_m will generally lead to a better SNR for different camera models, as can be seen with the SR4000 and the Senz3D.

With some cameras it is possible to identify a relation between measured distance and reflectivity of the foreground, for example with the CamCube 3.0, the CamBoard nano or the E70. As the foreground gets darker the measured distance decreases, even though the camera is not moved. In case of the CamCube 3.0 this offset totals to more than three centimeters. This is a magnitude greater than any positioning errors, and can therefore not be attributed to changing the foreground object. This effect can also be observed with internal filters being active. In Section IV-E reflectivity-related and amplitude-related errors will be investigated in depth.

C. Effects of Changing the Integration Time

Some cameras allow the user to change the integration time t_i such that the measurements for the region or distance of interest become more reliable. The Mesa SR3100 and SR4000 even support auto exposure. However, it has been shown that changing t_i can affect the measurement accuracy both, absolute and over a certain amount of time [16].

In this experiment we use a similar setup as Karel *et al.* [4]. Each camera is positioned in front of a diffuse, white plane at a distance of 1 meter. During the experiment we measure the distance to the foreground object within a 9×9 pixels window at the image center. Additionally the measurements are averaged for 1 second. Every 10 minutes t_i is changed to a different value programmatically. The values are selected to cover the whole range of possible values for this setting while avoiding too noisy measurements and oversaturated pixels at the same time.

The resulting data for each camera is plotted in Fig. 2. The unit of the integration time is microseconds with the respective values given as additional numbers in the plot.

In this experiment all cameras show similar characteristics: first, the distance measurements of all cameras change whenever t_i is altered, and second, some time is required until the depth measurements stabilize at a certain distance.

The CamCube 3.0 exactly shows this behavior (Fig. 2a). During the warm-up phase, the integration time has been set to 50 µs. During the experiment, the integration time is increased to 400 µs, resulting in a drift of the distance measurements of more than 5 cm. The time which is required for the measurements to stabilize is difficult to estimate as the total required time might be larger than the ten minute intervals.

The largest drift of the depth measurements can be observed with the SR3100 (see Fig. 2b). During the warm-up phase the integration time has been set to 400 µs. During the experiment the t_i is increased to 6000 µs. The measured distance drifts by more than 9 cm in this setup. Whenever the integration time is altered roughly 3 minutes are required until the depth values become stable again.

A similar behavior can be observed with the more recent SR4000. However, the absolute variation is more than a magnitude smaller as shown in Fig. 2c. The time required for stabilization cannot be read from the plot as easily as with the SR3100. When considering only the first three intervals one can hardly see a temporal variation, leading to the assumption that internal corrections are able to compensate for this effect very well. On the contrary the forth and fifth interval show a drift after changing the integration time. The time required for stabilization might even exceed the ten minute intervals chosen in this experiment.

The distance measurements of the Fotonic E70 drift by multiple centimeters when changing the integration time (see Fig. 2d). Furthermore, the plot shows jags which can also be found in the respective warm-up plot (see Fig. 6f). This may indicate that there exists a correlation between integration time and internal temperature. The time which is required for the distance measurements to stabilize after altering t_i may be longer than the ten minute period chosen in this experiment.

The Argos P100 (Fig. 2e) also shows the typical behavior. Again the distance measurements drift by more than 3 cm. The period required for stabilization after changing t_i is comparably short.

The results of the CamBoard nano show a different shape than those of the Argos P100, even though their ToF-chips are technically identical. In comparison to the Argos P100, the CamBoard nano requires more time until the camera gets into a

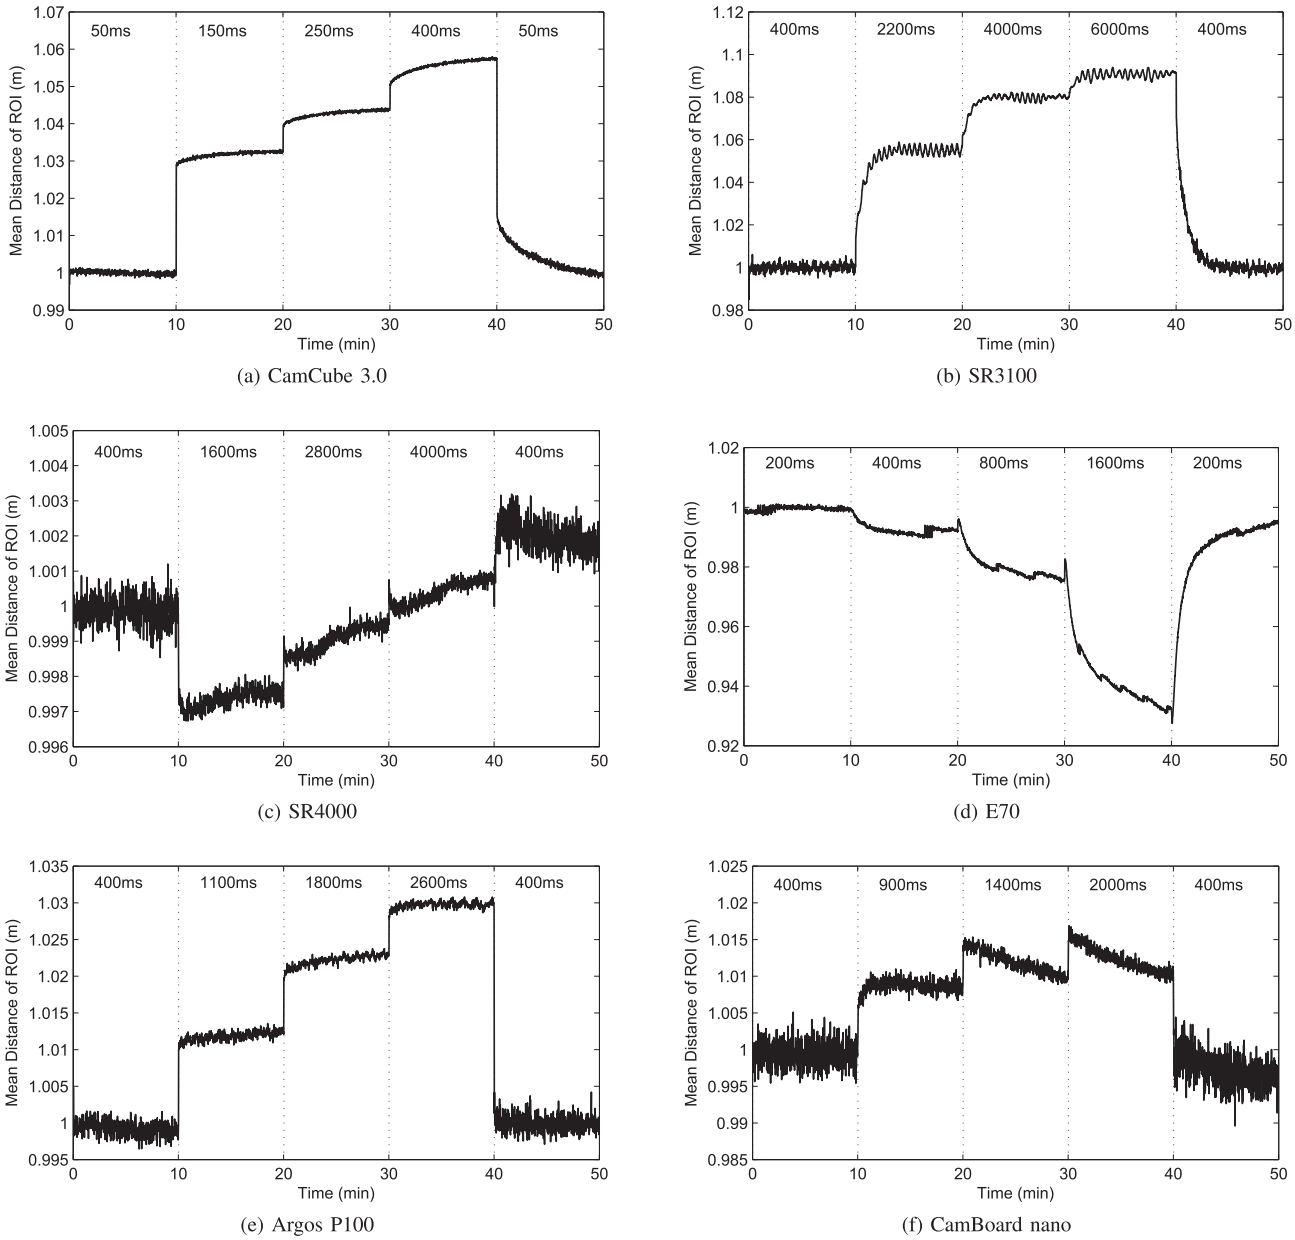

Fig. 2. Mean distance measurement to the ROI measured for different integration times. Each integration time setting lasts for 10 minutes and is delimited by dotted lines. The settings for each interval are given in their respective plot regions. Note the different scalings of the y-axis.

steady state again. A possible explanation might be that this camera uses an illumination unit with different temperature-dependent characteristics or a camera casing with different warm-up characteristics.

Multiple conclusions can be drawn from the results presented in this section. First, the distance measurements of all cameras are influenced by the integration time. Second, whenever t_i is changed, a certain amount of time is required for the distance measurements to converge again. The length of this period might correlate to the length of the warm-up phase and also on the absolute difference between the previous and new integration time. This results in the insight that a large part of the integration-time-related error is actually a temperature-dependent effect. Nevertheless, the total error is presumably a combination of multiple effects, for example a higher signal to noise ratio due to shorter illumination cycles and potential amplitude-related errors.

D. Internal Scattering

As reported by several groups, light can be scattered within the camera, which can negatively affect the measurement accuracy. This scattering occurs when incident light is reflected at the sensor and backscattered onto the sensor by the lens. Consequently, incident light from an object can influence phase measurements at other pixels. This effect can be observed best when the distance between foreground and background is large, and if the foreground has a higher reflectivity than the background. An illustration of this effect is given in Fig. 3.

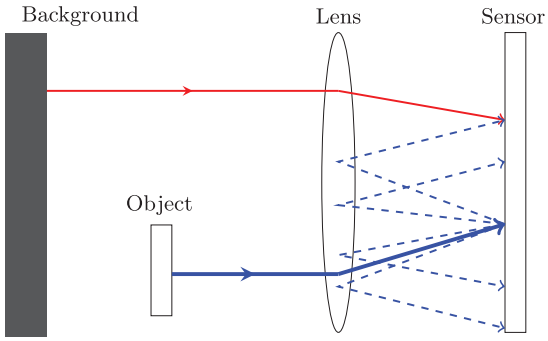

Fig. 3. Illustration of internal scattering. Incident light of the bright foreground object is scattered between lens and sensor and consequently influences other regions on the sensor. Typically low reflectivity regions in the scene, like a dark background, are affected most.

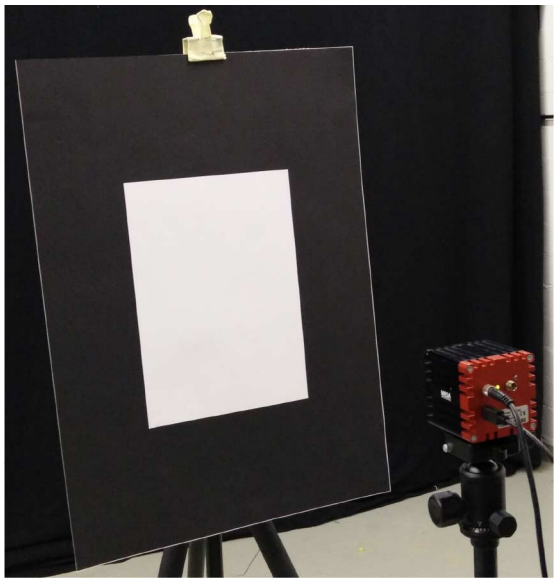

Fig. 4. The setup for measuring amplitude-related errors.

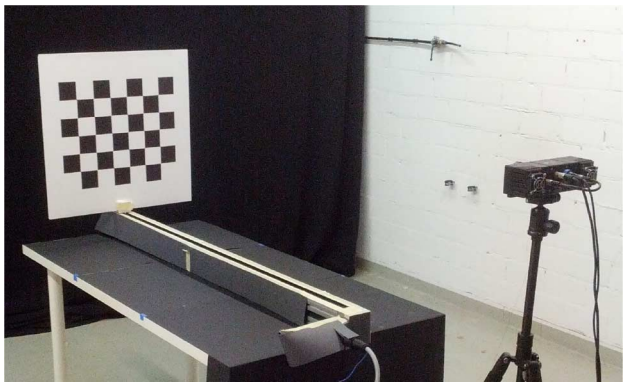

Fig. 5. The setup for the wiggling experiments.

We evaluate how strong this effect is present with the evaluated cameras by comparing two scenes with each other. The first scene consists of a white paper board which covers the right half of the camera's field of view. The foreground is positioned at approximately 80 cm. The distance to the low-reflectivity background is roughly 250 cm. The second scene simply shows the background with no foreground object.

This layout aims towards simulating the worst case setup for internal scattering, maximizing the ratio between scattered foreground light and incident background light. This setup is ideal to demonstrate the presence of internal scattering, but may also include amplitude-related errors due to different reflectivities of foreground and background. In the next section we evaluate amplitude-related errors separately and demonstrate how to distinguish them from internal scattering.

For both scenes, mean distance images of 1000 consecutive frames are computed. Internal scattering is visualized by plotting the horizontal profile of both distance images along the center pixel. Similar setups have been used by Karel *et al.* [4] and Kavli *et al.* [6] to show the influence of scattering as well as to demonstrate the performance of scattering compensation algorithms. In our experiments the integration time is fixed, as Karel *et al.* showed that this setting has no influence on the amount of scattering.

The results of the scattering experiments are shown in Fig. 7. The solid line represents the horizontal profile with the foreground object visible, the dashed line the same scene with the white board removed. Ideally the solid line would match the dashed line on the left half of the plot. For the Argos P100 and the E70 an amplitude-based confidence filter can be enabled. We plotted for these two cameras a third, dotted profile of distance measurements indicating that this filter is active.

The first plot (Fig. 7a) presents the CamCube 3.0 results computed from images captured with an integration time of 300 µs. The measurements in the background-only region are clearly influenced by the measurements of the foreground. Furthermore the whole sensor is affected, not only pixels in the area of the border of the foreground.

Similar results can be seen in the plots of the Mesa cameras (Fig. 7b and 7c). With both cameras the background measurements are significantly influenced by scattering.

The Senz3D suffers from scattering as well, however, its negative influence is not as severe as with the previous cameras. The global influence of the backscattered light is comparably small. The magnitude of the measurement error falls off faster for pixels further away from the foreground as with the other cameras. Next to scattering effects one can also observe that the distance measurements to the background are wrong. An explanation for this error is the decreased amount of reflected light as the distance to the background increases radially. Therefore, this error might be related to the amplitude.

The horizontal profile of the Kinect V2 shows the strongest global influence of all evaluated cameras. Most of the measurements within the background area lie below 80 cm, which would be even closer to the camera than the actual foreground. This effect might be caused by the algorithm that computes the distance measurements from multiple modulation frequencies. However, this algorithm is not publicly available, therefore the true reason for this effect is unknown. For additional information about the modulation frequencies and how they can be used to derive distance data can be found in [19].

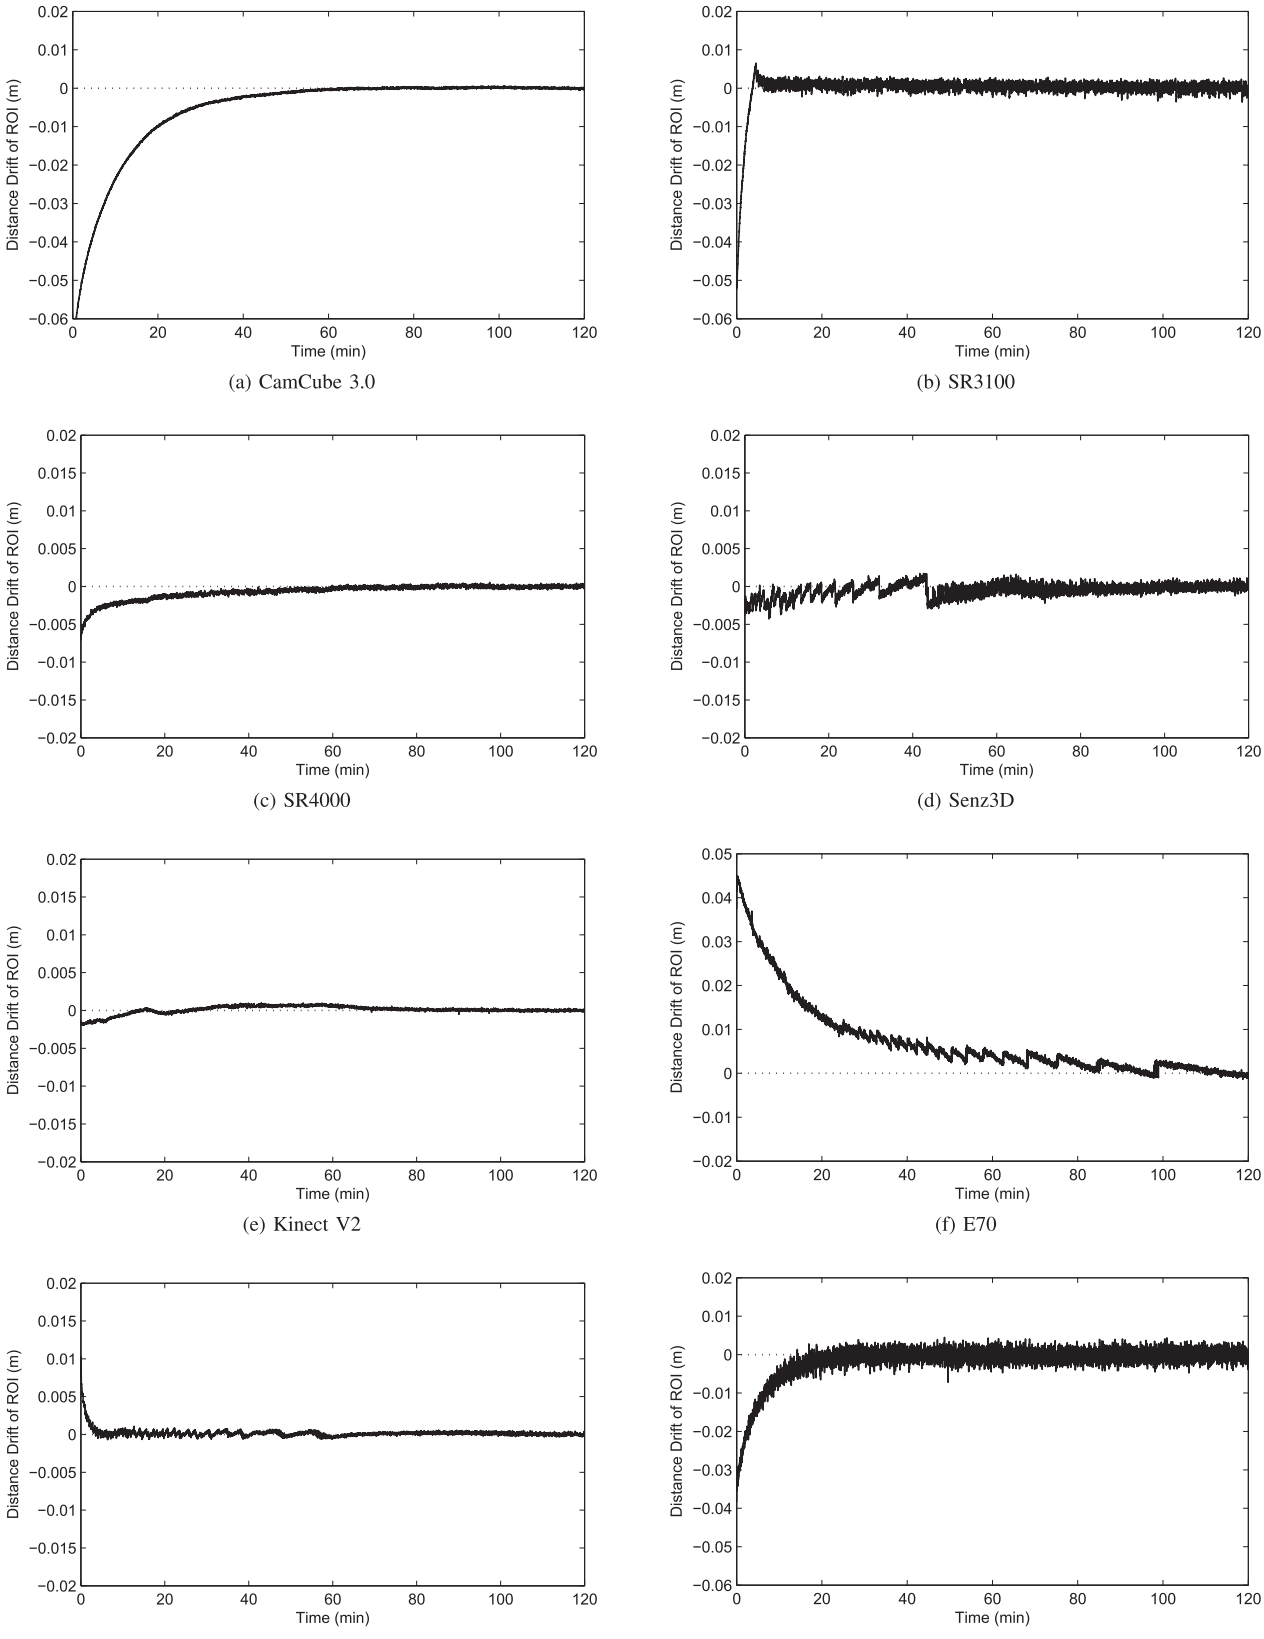

(a) CamCube 3.0 (b) SR3100

(c) SR4000 (d) Senz3D

(e) Kinect V2 (f) E70

(g) Argos P100 (h) CamBoard nano

Fig. 6. Warm-up experiment. Mean distance to the ROI captured for 120 minutes for all evaluated cameras plotted as a black line. The dotted line represents the mean distance computed from the measurements of the last 10 minutes and is considered as the steady state measurement. All measurements have been normalized by the steady state. Note the different scalings of the y-axis with some figures.

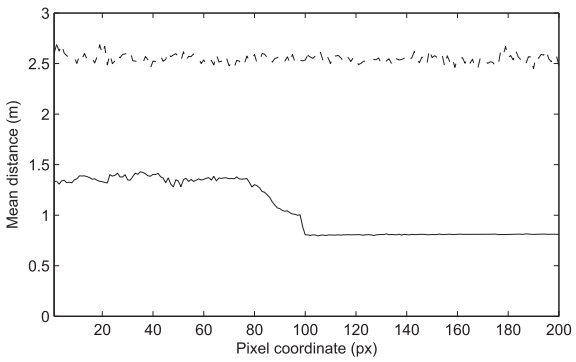

(a) CamCube 3.0

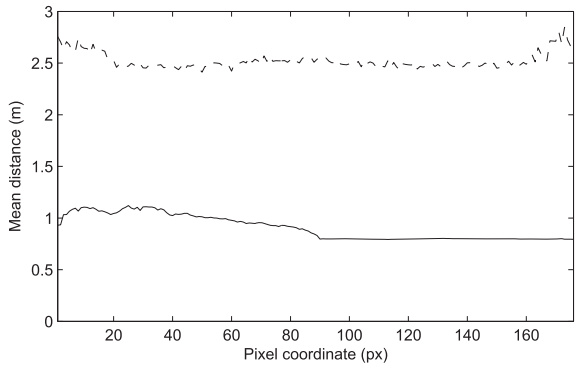

(b) SR3100

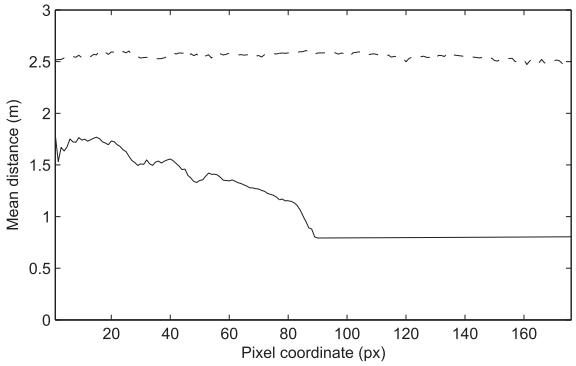

(c) SR4000

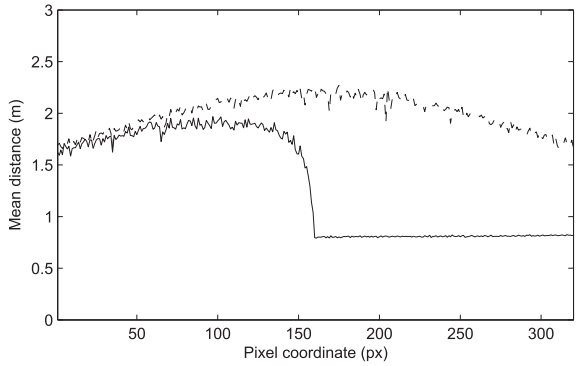

(d) Senz3D

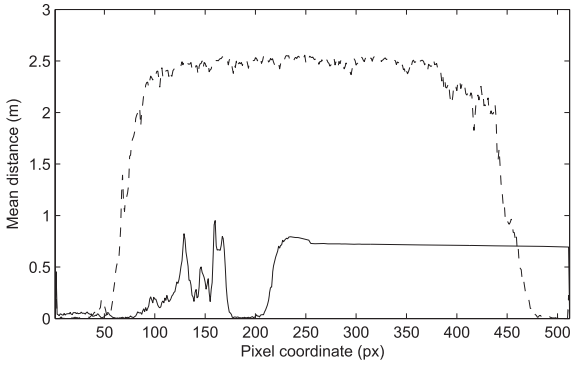

(e) Kinect V2

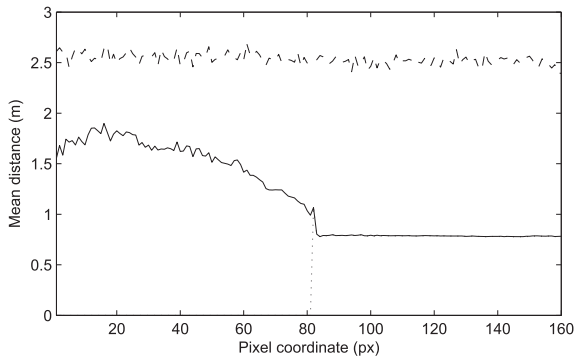

(f) E70

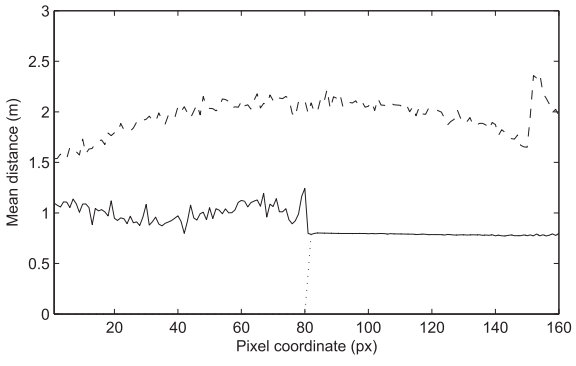

(g) Argos P100

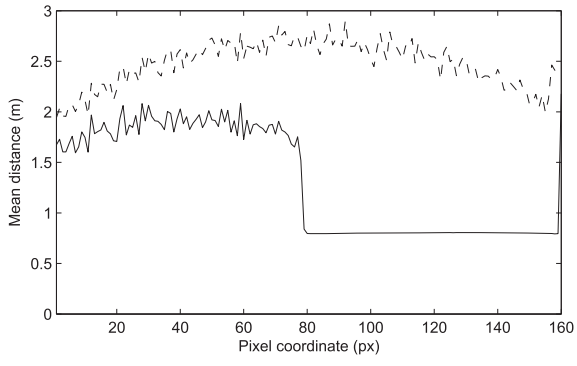

(h) CamBoard nano

Fig. 7. Internal scattering experiment. The solid line represents the averaged horizontal profile at the center of the Z-image with the foreground object present. The dashed line shows the profile without the foreground object. The E70 and Argos P100 have amplitude-based confidence filters to discard unreliable measurements. The profiles with these filters being active are shown as dotted lines.

One solution to cope with scattering is to perform confidence-based filtering for every pixel to invalidate unreliable measurements. This confidence value may be computed based on the amplitude value of a certain pixel, meaning that at least a certain amount of light must be received for a valid depth measurement. The E70 offers such a filter which is enabled by default. With this camera, the default, conservative threshold basically invalidates all background measurements as shown in Fig. 7f as a dotted line. Nonetheless, scattering can be observed when disabling this filter. For some applications it might be desirable to enable the filter as low amplitude values typically indicate noisy distance measurements.

By default the Argos P100 also performs amplitude based filtering. However, when disabling this behavior scattering effects become visible (see Fig. 7g). Similar as with the E70 all background pixels are invalidated if the amplitude based thresholding is active (dotted line). Two additional observations can be made in the plot. First, due to the wide opening angle of the camera, one can see the side wall of the room on the far right side of the plot, starting at a pixel coordinate of 151. Second, the low-amplitude distance measurements of the background differ from the true distance by at least 50 cm. In contrast to the false dark distance measurements the distances to pixels which present parts of the white side wall are correct with approximately 250 cm. The reason for the wrong measurements may be due to the low amount of incident light for the dark background pixels and hence, similar as with the Senz3D, be related to the amplitude. The magnitude of this effect will be investigated in detail in Section IV-E.

The CamBoard nano also shows severe scattering effects, even though the difference of the distance measurements at the border of the foreground objects is comparably large. Nonetheless, distance measurements throughout the background area are heavily influenced by the foreground plane, resulting in measurement errors of almost one meter.

E. Amplitude-Related Distance Error

This experiment aims to quantify the so-called amplitude-related distance error and also to show that this effect is not related to scattering. This effect can be observed when looking at a planar surface with high reflectivity variations. With some sensors the distance measurements for pixels with different amplitudes do not lie on the same plane, even though they should.

To the best of our knowledge no evaluation setup has been presented for this error source so far. In the past this error has been typically observed with images of checkerboards or other high contrast patterns. However, the analysis of single images allows no differentiation between amplitude-related errors and internal scattering.

In this work we measure the average distance to two different planar, high-contrast patterns. The first pattern consists of a black background with a white sheet of paper in its center as shown in Fig. 4. The second pattern is simply the inverse of the first pattern. Both patterns can be exchanged without altering the distance or orientation to the camera. In all experiments the patterns are positioned orthogonally to the optical axis and cover the whole field of view in horizontal direction. Due to different opening angles of the camera the pattern has to be positioned at a different distance for each sensor. Due to the low image resolution we chose to rather vary the distance between camera and foreground instead of using a large, unified pattern for all cameras. The downside of such a large pattern would be that the size of the central area of the pattern would vary and may become to small. To avoid temporal noise effects, average distance images (1000 frames) are captured of both patterns.

For this experiment three different outcomes can be imagined. Ideally one would not see any distance changes at the transition from high reflectivity to low reflectivity for both patterns, which results in a smooth distance image. In this case there would be neither any amplitude-related error nor any scattering effects. If significant distance changes can be observed in the distance images, one has to distinguish the directions towards which the low-reflectivity areas are pulled. These directions indicate whether the wrong measurements are caused by low reflectivity or by scattering.

If both, the black background of the first pattern and the black center of the second pattern, are pulled towards the same direction, e.g. towards the camera, the main reason for this error is the amplitude or low reflectivity of the patch that the pixel represents. Whenever the black regions of the two patterns are pulled in different directions, scattering is the dominant error source. In this case the measurements in the low-reflectivity areas of the patterns are influenced by the white regions.

The results of the experiments are plotted in Fig. 8. Each plot shows two horizontal profiles along the center pixel of the Z-image. The solid line represents the measurements performed with the pattern shown in Fig. 4. The dashed line represents the profile measured with the inverse pattern. Additionally we show a reference distance (shaded gray) which we estimated with a checkerboard.

The CamCube 3.0 is a good example to illustrate the effects of amplitude-related distance errors. The variation between black and white areas is approximately 3.5 cm, with the black areas appearing to be closer to the camera. These results also match the mean distance measurements performed in the temporal noise experiment. Therefore, this camera shows a significant amplitude-dependent error.

In contrast to the results of the CamCube 3.0, the data of the SR3100 appears noisy. This makes it difficult to draw conclusions with respect to amplitude-related errors. Several different integration times have been evaluated, including the values suggested by the auto exposure feature without getting smoother distance measurements. Similarly, the results of the temporal noise experiment do not exhibit a clear trend. Reasons for the bad measurements may include fixed pattern noise and poor illumination of the border areas of the field of view. The latter may cause noisy measurements towards the image border, especially with the first pattern.

Significant improvements have been made with the more recent SR4000. The data shown in Fig. 8c shows less variation than its predecessor. However, measurement errors at the transition between black and white can still be observed. Noise levels are noticeably higher in black areas, but a definite amplitude-related error cannot be identified.

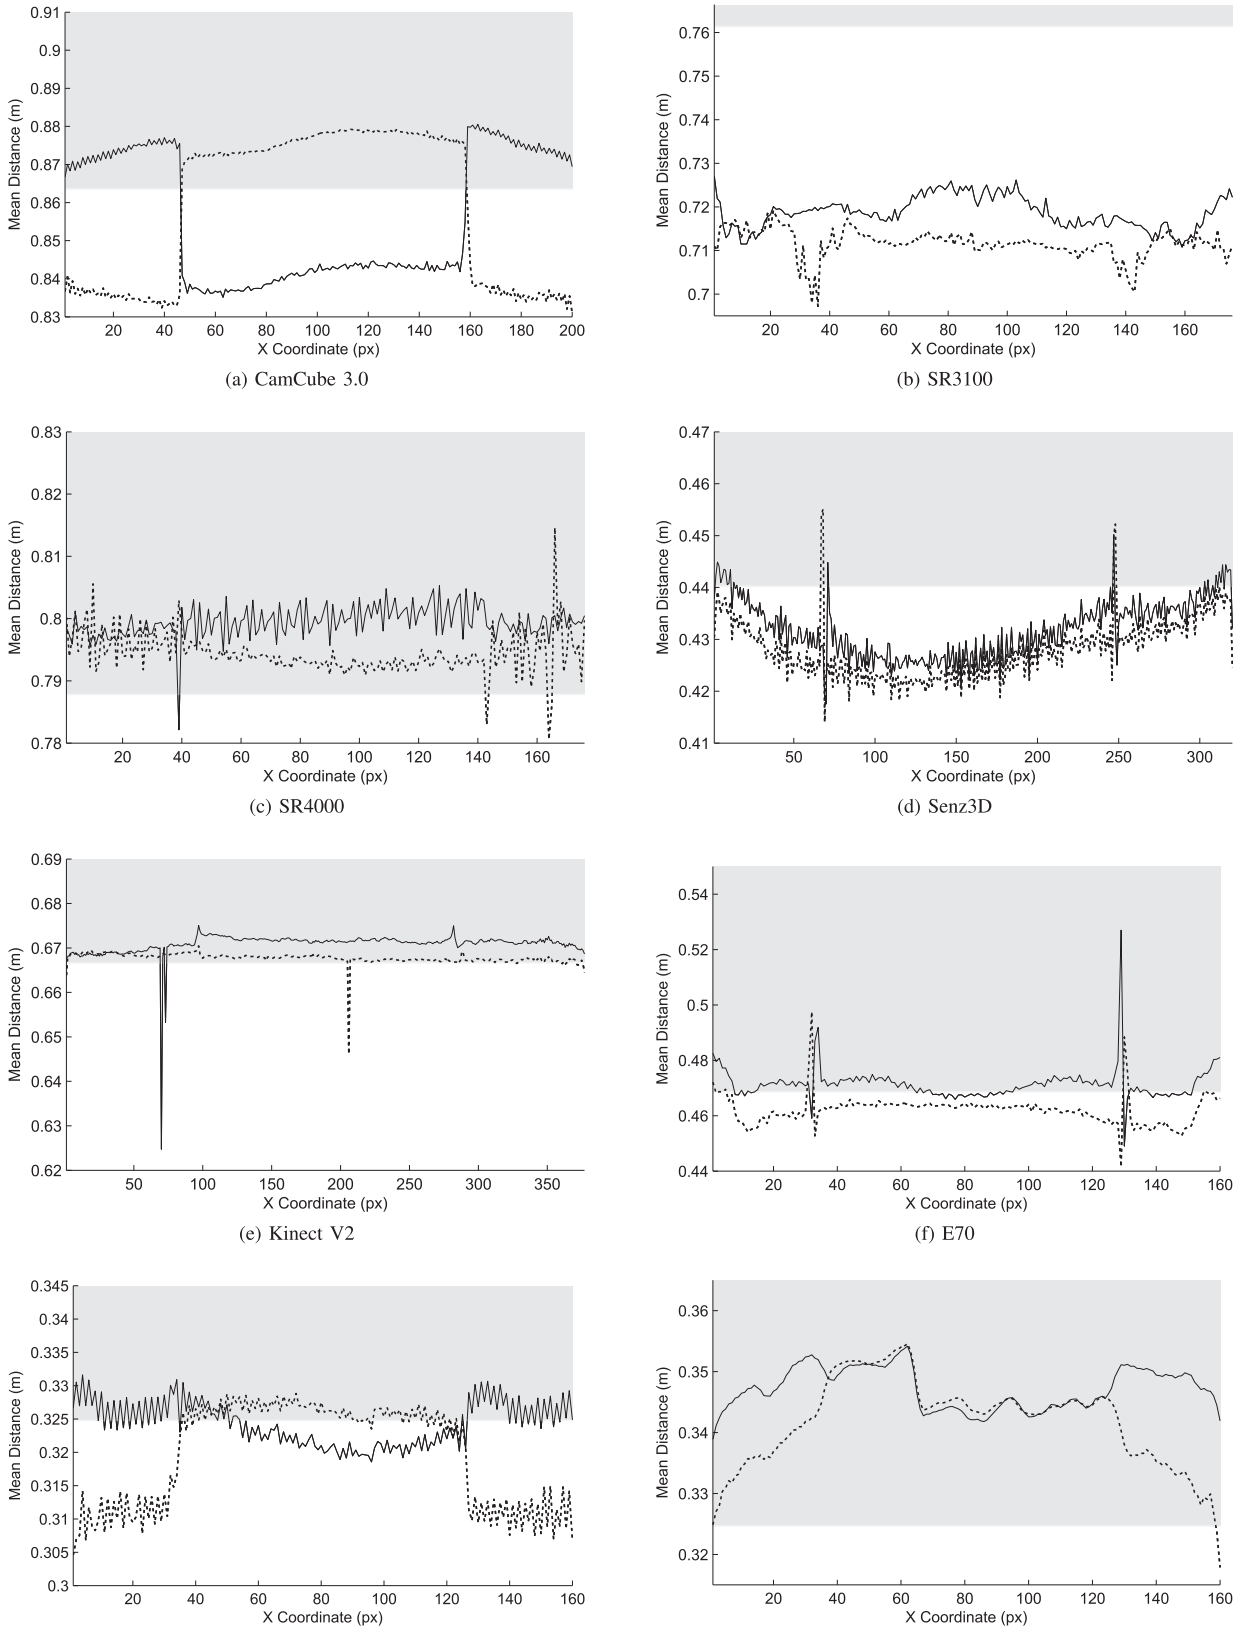

(a) CamCube 3.0

(b) SR3100

(c) SR4000

(d) Senz3D

(e) Kinect V2

(f) E70

(g) Argos P100

(h) CamBoard nano

Fig. 8. Amplitude-related error experiment. Each plot shows two horizontal profiles computed from two separate measurements. The first measurement (solid line) has been captured with the pattern shown in Fig. 4. The dotted line represents distance measurements with the inverse pattern. A checkerboard was used to estimate a reference distance which is shown in gray.

A similar amount of noise can be seen in the plot of the Senz3D (see Fig. 8d). Fixed pattern noise seems to be present, especially in the central region of the image. Still, a clear amplitude-related error cannot be identifier without doubt.

Due to camera-internal filtering, very little noise can be observed with the Kinect V2 as presented in Fig. 8e. Note that the horizontal axis does not cover the whole image width. The pattern used in our experiments is too small to cover the complete field of view of the camera at the minimum distance required by the sensor (approx. 70 cm). With the first pattern almost no transition nor amplitude-related effects can be observed. However, for the pattern with the white background one can see that the dark center region is pulled away from the camera. The reason for this behavior might be the internal scattering of white background measurements, as the high-reflectivity region is further away from the camera than the low-reflectivity center of the pattern. Please note that the spikes which can be seen in the plot originate from invalid distance measures at single pixels in single frames which distorted the average distance measurement.

Both temporal filtering and amplitude thresholding have been deactivated for evaluating the E70. As a consequence the noise levels in the black background region of the first pattern increase, which makes an interpretation of the data difficult. However, when assuming that white regions are measured accurately, one can argue that the black areas of the patterns are influenced by scattered light from bright regions. This is due to the observation that the distance drops with the second pattern at the transition from white to black. With the first pattern, no drop can be observed. Thus, in this example, scattering effects are more severe than amplitude-related errors.

Analogously to the E70, it is also required to disable the amplitude-based filter in the Argos P100. The results for this camera are shown in Fig. 8g. Similar as with the Kinect V2, it is only possible to see a clear offset offset in the profile with one of the patterns. Arguing analogously like for the Kinect V2 experiment, the reason for these results is that effects of internal scattering dominate the amplitude or reflectivity-related effects.

The CamBoard nano exhibits strong fixed pattern noise, which makes it difficult to decide whether the CamBoard nano suffers from an amplitude-related error. However, it is still possible to identify an offset at transition from white to black with the first pattern (solid line in Fig. 8h). This matches the results of the Argos P100, which uses the same sensor.

F. Wiggling

For evaluating the presence of wiggling, ground truth distance information is required. We calculate the true distance by setting up a stereo camera system. This system consists of the ToF camera to be evaluated and a high resolution monochrome camera (IDS UI-1241LE[7]) which we call the reference camera. To ensure accurate measurements with the stereo setup, a wide baseline of approximately 2 m is chosen. Both cameras are oriented towards a checkerboard. The board is mounted on a linear slider which allows to increase the distance between cameras and reference object automatically and reproducibly (see Fig. 5). The cameras are calibrated with Zhang's algorithm [23] with point correspondences computed with ROCHADE [24]. Ground truth is calculated by intersecting the rays of all ToF camera pixels with the 3D plane of the checkerboard. For higher accuracy, we compute this plane from corners detected in the reference image and transform the plane into the coordinate system of the ToF camera.

In this experiment, only measurements of white checkerboard quads are considered to avoid amplitude-related effects. For each step of the linear slider the average of 25 distance images is computed. For each pixel that represents a white quad both, the ground truth distance of the checkerboard and the measurement error is stored. Integration times are set such that there exist no saturation effects in the checkerboard area at the closest position of the linear slider.

The results of the wiggling experiments are shown in Fig. 9. Single measurement errors for a certain distance are plotted as gray dots. For visualization purposes the mean measurement errors for 1 cm intervals are plotted as a solid black line.

We begin with the discussion of the SR3100 as it is a good example to visualize wiggling (see Fig. 9b). Depth measurement errors of up to 4 cm can be observed. The typical sinusoidal shape of the wiggling error can be identified easily. In this experiment, t_i has been set to 1000 µs.

Deciding whether the CamCube 3.0 suffers from wiggling or not is difficult with the current evaluation, as the sinusoidal shape of the error cannot be identified with confidence (see Fig. 9a). If there are errors due to imperfectly modulated light, then the wavelength of the error signal is obviously larger than for the SR3100. To determine wiggling effects with this camera, one would need a different setup with a larger evaluation range. In this experiment, t_i has been set to 200 µs.

In comparison to its predecessor the SR4000 compensates wiggling well, as shown in Fig. 9c. The variation of the average measurement error is less than 1 mm. Auto-exposure is turned off and the integration time is fixed to a value which satisfies the aforementioned criteria.

The error shape of the Senz3D (Fig. 9d) may be approximated with a sinusoidal function, although the model will not fit as well as with the SR3100. Potentially some internal correction algorithms compensate some parts of this error and ignore others. The results for this camera can be considered reliable up to a distance of approximately 2.70 m. Beyond this distance the illumination of the checkerboard is too dark for accurate measurements.

A different error shape can be observed with the Kinect V2 in Fig. 9e. As this camera uses multiple modulation frequencies it can be expected that identifying a single sinusoidal error function will be difficult. The measurement errors vary w.r.t. distance and reach values of up to 2 cm.

Wiggling errors seem to be compensated well with the E70 (Fig. 9f) except for two distance ranges at approximately 0.75 meters and 1.5 meters. Reasons for these outliers might either be incorrect wiggling correction terms or slightly inaccurate checkerboard detections at these positions. With an integration time of 400 µs range measurements of the checkerboard can

[7]IDS Imaging Development Systems GmbH, Obersulm, http://www.ids-imaging.com

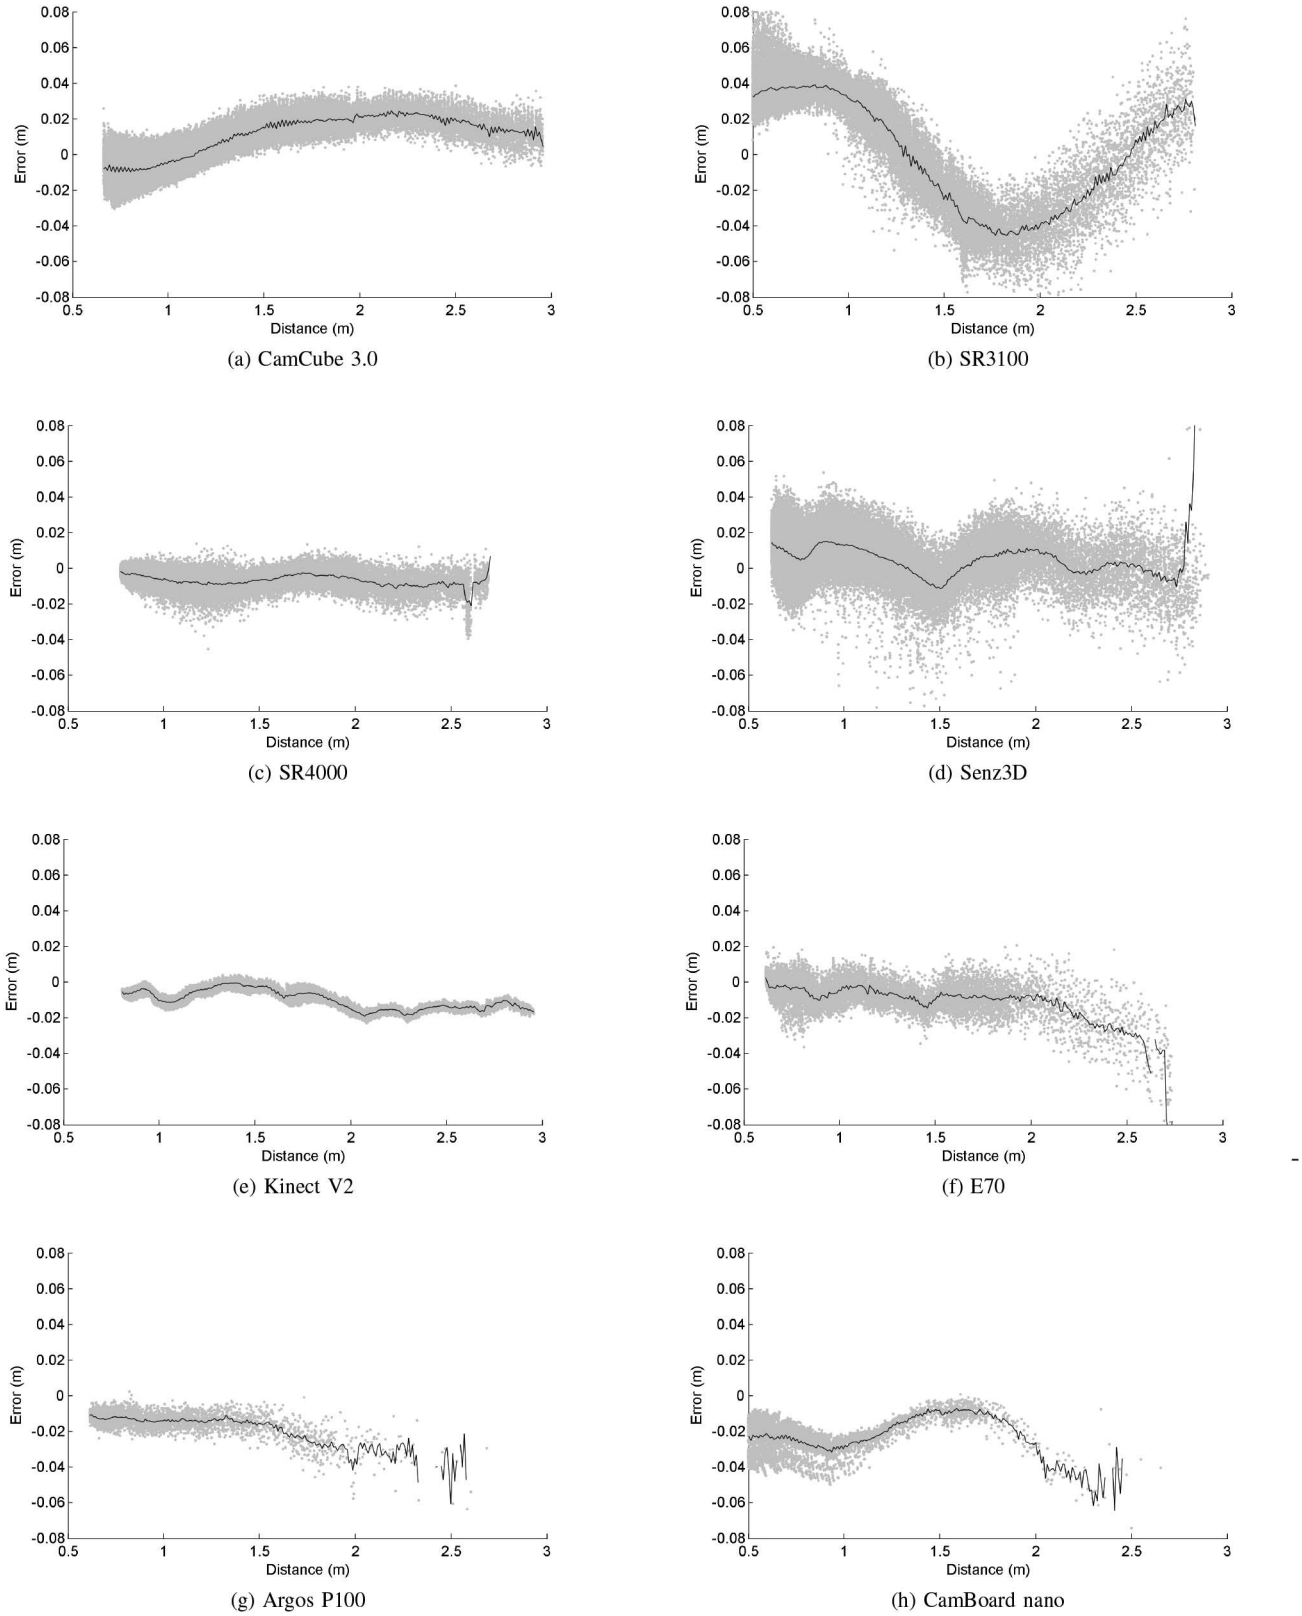

(a) CamCube 3.0

(b) SR3100

(c) SR4000

(d) Senz3D

(e) Kinect V2

(f) E70

(g) Argos P100

(h) CamBoard nano

Fig. 9. Mean measurement error for different distances to investigate wiggling. Gray dots represent single measurements. The black line is the average measurement error in a one centimeter window.

be considered reliable up to a distance of two meters. Beyond this distance noise levels increase significantly and an increased integration time is recommended.

The Argos P100 does not show any error characteristics which can be associated with typical wiggling (see Fig. 9g). However, the working range of the camera is comparably

limited as distance measurements beyond 1.5 m are increasingly noisy with the chosen integration time of 2000 µs.

The CamBoard nano shows, unlike the Argos P100, clear wiggling errors. However, the dynamic range of the camera is slightly higher, with a maximum distance of approximately 1.8 m with an integration time of 2000 µs. The total variation of the distance-dependent error is larger than two centimeters. As both cameras are based on the same sensor, possible explanations for the different wiggling characteristic include either the lack of a wiggling correction or the usage of a different illumination unit.

This last experiment concludes the wiggling evaluation. Some cameras compensate wiggling errors reliably. Some cameras suffer from high wiggling errors, either due to the lack of correction mechanisms or due to an insufficient distance calibration. The distance-dependent measurement errors of the evaluated cameras do not follow a common pattern. Instead, they exhibit very different characteristics which requires flexible models when developing new, generally applicable wiggling-correction methods.

V. Discussion

Each camera has its own strengths and weaknesses. In initial experiments, we empirically observed across a number of Kinect V2 and Mesa SR4000 devices that the presence and quality of the evaluated errors are characteristic for a camera model. The magnitude of the error can vary for different devices within a camera model. With respect to the evaluated systematic errors, there is no camera that consistently outperforms all the others. Hence, it depends on the type of application for the camera which errors are considered more critical than others. Nonetheless, it is possible to highlight cameras which only marginally suffer from certain error sources, or well compensate a particular type of error. In the remainder of this section we order these cameras in the order of the experiments.

The Kinect V2 and Argos P100 show only a small drift during the warm-up time. Measurements can be considered reliable after a short warm-up period of less than 10minutes.

Internal filters allow the Kinect V2 and E70 cameras to reduce the measurement noise drastically, even if the scene reflectivity is low. The downside of internal filtering is that these filters typically smooth the measurements either temporally or spatially. If a high temporal resolution is required, it might be reasonable to deactivate these filters. With all internal filters disabled the CamBoard nano or SR4000 perform best.

In our experiments we demonstrated that changing the integration time will cause the distance measurements to drift. In terms of total drift the SR4000 performs best.

Internal scattering is a problem that all evaluated cameras share. Solutions may include amplitude-based filtering as possible with the E70 and the Argos P100, especially if the region of interest has a high reflectivity or is illuminated well.

In Section IV-E we evaluated the presence and severity of amplitude-related errors. The least sensitivity with respect to this error source can be observed with the Kinect V2 and the Argos P100.

Wiggling or non-linear distance-dependent measurement errors can be observed with most cameras. Still, some compensate the so-called wiggling error effectively, e.g. the Argos P100, the E70 or the SR4000.

VI. Conclusion

We have presented an evaluation setup for six systematic error sources of ToF cameras and investigated the performance of eight cameras.

Due to the standardized setup we were able to show that each camera model shows different error characteristics.

By presenting results for the cameras which have been studied in the past we have also proved the applicability of the suggested setup. A few cameras compensate some error sources well, whereas other effects like internal scattering remain an unsolved problem with all evaluated cameras. Cameras which successfully compensate certain errors are, for example, the Microsoft Kinect V2 or the Fotonic E70, which remove the so-called wiggling error, or the Bluetechnix Argos P100 which, in our experiments, shows no amplitude-related error. Other causes like the sensor warmup can also be corrected or reduced as can be seen with the Creative Senz3D or the Mesa SR3100. Additionally, we show that the integration time has significant influence on the measured distance. Based on the results in this work, researchers and practitioners will gain valuable insight for matching their application with a proper ToF camera. Furthermore, the presented error characteristics deepen the understanding of the behavior of a wide range of ToF cameras. As a consequence, this work supports the development of more efficient, generally applicable correction methods.

Acknowledgment

We would like to thank the reviewers for their constructive comments that helped us to improve the quality of this work.

References

[1] T. Pattinson, "Quantification and description of distance measurement errors of a time-of-flight camera," M.S. thesis, Institute for Photogrammetry, University of Stuttgart, Stuttgart, Germany, 2010.

[2] J. C. K. Chow, D. D. Lichti, F. Remondino, M. R. Shortis, J. Beyerer, and F. Puente Len, "A study of systematic errors in the PMD CamBoard nano," in *Proc. SPIE Opt. Metrol.*, 2013, p. 87910X.

[3] N. Pfeifer, D. Lichti, J. Böhm, and W. Karel, "3D cameras: Errors, calibration and orientation," in *TOF Range-Imaging Cameras*, F. Remondino and D. Stoppa, Eds. New York, NY, USA: Springer, 2013, pp. 117–138.

[4] W. Karel, S. Ghuffar, and N. Pfeifer, "Modelling and compensating internal light scattering in time of flight range cameras," *Photogramm. Rec.*, vol. 27, no. 138, pp. 155–174, 2012.

[5] J. Mure-Dubois and H. Hügli, "Real-time scattering compensation for time-of-flight camera," in *Proc. Int. Conf. Comput. Vis. Syst.*, 2007.

[6] T. Kavli, T. Kirkhus, J. T. Thielemann, and B. Jagielski, "Modelling and compensating measurement errors caused by scattering in time-of-flight cameras," in *Proc. SPIE*, vol 7066, pp. 706604-1–706604-10, 2008, doi: 10.1117/12.791019.

[7] S. Jamtsho and D. D. Lichti, "Modelling scattering distortion in 3D range camera," *Int. Arch. Photogram. Remote Sens. Spatial Inf. Sci.*, vol. 38, no. 5, pp. 299–304, 2010.

[8] F. Chiabrando, D. Piatti, and F. Rinaudo, "SR-4000 TOF camera: Further experimental tests and first applications to metric surveys," *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, vol. 38, no. 5, pp. 149–154, 2010.

[9] D. Piatti and F. Rinaudo, "SR-4000 and CamCube3.0 time of flight (TOF) cameras: Tests and comparison," *Remote Sens.*, vol. 4, no. 12, pp. 1069–1089, 2012.
[10] M. Lindner, I. Schiller, A. Kolb, and R. Koch, "Time-of-flight sensor calibration for accurate range sensing," *Comput. Vis. Image Understand.*, vol. 114, no. 12, pp. 1318–1328, 2010.
[11] T. Stoyanov, A. Louloudi, H. Andreasson, and A. J. Lilienthal, "Comparative evaluation of range sensor accuracy in indoor environments," in *Proc. 5th Eur. Conf. Mobile Robots (ECMR'11)*, Sep. 7–9, 2011, Örebro, Sweden, 2011, pp. 19–24.
[12] G. Rauscher, D. Dube, and A. Zell, "A comparison of 3D sensors for wheeled mobile robots," in *Proc. Int. Conf. Intell. Auton. Syst. (IAS'13)*, Padova, Italy, 2014, pp. 29–41.
[13] E. Lachat, H. Macher, T. Landes, and P. Grussenmeyer, "Assessment and calibration of a RGP-D camera (Kinect v2 sensor) towards a potential use for close-range 3D modeling," *Remote Sens.*, vol. 7, no. 10, pp. 13070–13097, 2015.
[14] L. Yang, L. Zhang, H. Dong, A. Alelaiwi, and A. El Saddik, "Evaluating and improving the depth accuracy of Kinect for windows v2," *IEEE Sens. J.*, vol. 15, no. 8, pp. 4275–4285, Aug. 2015.
[15] H. Sarbolandi, D. Lefloch, and A. Kolb, "Kinect range sensing: Structured-light versus time-of-flight Kinect," *CoRR*, vol. abs/1505.05459, 2015.
[16] S. Foix, G. Alenya, and C. Torras, "Lock-in time-of-flight (TOF) cameras: A survey," *IEEE Sens. J.*, vol. 11, no. 9, pp. 1917–1926, Sep. 2011.
[17] A. A. Dorrington, J. P. Godbaz, M. J. Cree, A. D. Payne, and L. V. Streeter, "Separating true range measurements from multi-path and scattering interference in commercial range cameras," *Proc. SPIE*, vol. 7864, pp. 786404-1–786404-10, 2011, doi: 10.1117/12.876586.
[18] A. Bhandari *et al.*, "Resolving multipath interference in time-of-flight imaging via modulation frequency diversity and sparse regularization," *Opt. Lett.*, vol. 39, no. 6, pp. 1705–1708, 2014.
[19] D. Freedman, Y. Smolin, E. Krupka, I. Leichter, and M. Schmidt, "SRA: Fast removal of general multipath for TOF sensors," in *Computer Vision—ECCV 2014*. Cham, Switzerland: Springer International Publishing, 2014, vol. 8689, pp. 234–249 [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-319-10590-1_16
[20] M. Lindner, "Calibration and real-time processing of time-of-flight range data," Ph.D. dissertation, Comput. Graphics Multimedia Syst., Univ. Siegen, Siegen, Germany, 2010.
[21] D. Lefloch *et al.*, "Technical foundation and calibration methods for time-of-flight cameras," in *Time-of-Flight and Depth Imaging Sensors, Algorithms, and Applications*. Berlin, Germany: Springer, 2013, vol. 8200, pp. 3–24 [Online]. Available: http://link.springer.com/chapter/10.1007%2F978-3-642-44964-2_1
[22] M. Hansard, S. Lee, O. Choi, and R. Horaud, *Time-of-Flight Cameras: Principles, Methods and Applications*. New York, NY, USA: Springer, 2013.
[23] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
[24] S. Placht *et al.*, "ROCHADE: Robust checkerboard advanced detection for camera calibration," in *Computer Vision—ECCV 2014*. Cham, Switzerland: Springer International Publishing, 2014, vol. 8692, pp. 766–779 [Online]. Available: http://link.springer.com/chapter/10.1007%2F978-3-319-10593-2_50

Peter Fürsattel received the Dipl.-Ing. degree in information and communication technology from the University of Friedrich-Alexander Erlangen-Nuremberg, Erlangen, Germany, in 2012. He is currently pursuing the Ph.D. degree at the Pattern Recognition Laboratory, University Erlangen-Nuremberg. His research interests include the distance calibration of time-of-flight cameras, multicamera setups, and applications with depth cameras.

Christian Riess received the Ph.D. degree in computer science from the Friedrich-Alexander University of Erlangen-Nuremberg, Erlangen, Germany, in 2012. He did his Postdoc with the Radiological Sciences Laboratory, Stanford University, Stanford, CA, USA, from 2013 to 2015. Currently, he works as a Researcher and the Head of the Phase-Contrast X-ray Group, Pattern Recognition Laboratory, University of Erlangen-Nuremberg. His research interests include all aspects of image processing and imaging, particularly with applications in image forensics, X-ray phase contrast, color image processing, and computer vision.