Statistical Modeling and Calibration of Triangulation Lidars

Anas Alhashimi, Damiano Varagnolo and Thomas Gustafsson

Control Engineering Group, Department of Computer Science, Electrical and Space Engineering,
Luleå University of Technology, Luleå 97187, Sweden

Keywords: Maximum Likelihood, Least Squares, Statistical Inference, Distance Mapping Sensors, Lidar, Nonlinear System, AIC.

Abstract: We aim at developing statistical tools that improve the accuracy and precision of the measurements returned by triangulation Light Detection and Rangings (Lidars). To this aim we: *i)* propose and validate a novel model that describes the statistics of the measurements of these Lidars, and that is built starting from mechanical considerations on the geometry and properties of their pinhole lens - CCD camera systems; *ii)* build, starting from this novel statistical model, a Maximum Likelihood (ML) / Akaike Information Criterion (AIC) - based sensor calibration algorithm that exploits training information collected in a controlled environment; *iii)* develop ML and Least Squares (LS) strategies that use the calibration results to statistically process the raw sensor measurements in non controlled environments. The overall technique allowed us to obtain empirical improvements of the normalized Mean Squared Error (MSE) from 0.0789 to 0.0046.

1 INTRODUCTION

Lidars are ubiquitously used for mapping purposes. Different types of Lidar technologies, such as Time of Flight (ToF) and triangulation, have different statistical performance. For example, ToF Lidars have generically lower bias and measurement noise variances than triangulation ones. At the same time, triangulation Lidars are generally cheaper than ToF ones. The market pull is then to increase the performance of cheaper Lidars in a cost-effective way.

Improving the accuracy and precision of sensors can then be done in different ways, e.g., by improving their mechanical properties. *In this paper we have a precise target: improve the performance indexes of triangulation Lidars by removing their biases and artifacts through opportune statistical manipulations of the raw information coming from the sensor.*

The following literature review analyzes a set of algorithms that are related to our aim.

Literature Review.   It is convenient to categorize the algorithms in the existing and relevant literature as:

- procedures for the characterization or calibration of the devices. Here characterization means a thorough quantification of the measurement noisiness of the device, while calibration means an

algorithm that aims at diminishing this noisiness level;

- when dealing with calibration issues, procedures for the intrinsic or extrinsic calibration. Here intrinsic means that the focus is on estimating the parameters of the Lidar itself, while extrinsic means that the focus is on estimating the parameters resulted from sensor positioning and installation.

Characterization issues: several papers discuss Lidar characterization issues for both ToF (Kneip et al., 2009; Reina and Gonzales, 1997; Lee and Ehsani, 2008; Sanz-Cortiella et al., 2011; Tang et al., 2009; Tuley et al., 2005; Ye and Borenstein, 2002; Anderson et al., 2005; Alhashimi et al., 2015) and triangulation Lidars (Lima et al., 2015; Campos et al., 2016). Notice that, at the best of our knowledge, for triangulation Lidars there exist only two manuscripts: (Lima et al., 2015), that discusses the nonlinearity of Neato Lidars, and (Campos et al., 2016), that analyzes the effect of the color of the target on the measured distance. Importantly, (Lima et al., 2015) models nonlinear effects on the measurements and the variance of additive measurement noises as two independent effects that can be modeled with a second order polynomials on the actual distance. From a statistical perspectives the authors, therefore, decouple the learning process into two separate parts.

Calibration issues: as for the calibration issues there is a relatively large number of papers describing how to calibrate extrinsic parameters either using additional sensors (such as cameras) (Zhang and Pless, 2004b; Mei and Rives, 2006; Jokinen, 1999; Tiddeman et al., 1998), or just requiring knowledge on the motion of the Lidar itself (Andreasson et al., 2005; Wei and Hirzinger, 1998; McIvor, 1999; Zhang and Pless, 2004a).

Still considering calibration issues, there has been also a big effort on how to perform intrinsic calibration for multi-beam Lidar systems, where the results from one beam is used to calibrate the intrinsic parameters of other beams (Chen and Chien, 2012; Muhammad and Lacroix, 2010; Atanacio-Jiménez et al., 2011; Glennie and Lichti, 2010; Glennie and Lichti, 2011; Glennie, 2012; Gordon and Meidow, 2013; Mirzaei et al., 2012; Gong et al., 2013; Park et al., 2014). As for single-beam Lidar systems, instead, (Mirzaei et al., 2012) proposes a method for the intrinsic calibration of a revolving-head 3D Lidar and the extrinsic calibration of the parameters with respect to a camera. The technique involves an analytical method for computing an initial estimate for both the Lidar's intrinsic parameters and the Lidar-camera transformation, that is then used to initialize an iterative nonlinear least-squares refinement of all of the calibration parameters.

We also mention the topic of on-line calibration of sensor parameters for mobile robots when doing Simultaneous localization and mapping (SLAM), very useful in navigation tasks. In this category, (Kümmerle et al., 2011) proposes an approach to simultaneously estimate a map of the environment, the position of the on-board sensors of the robot, and its kinematic parameters. These parameters are subject to variations due to wear of the devices or mechanical effects like loading. An other similar methodology for the intrinsic calibration of depth sensor during SLAM is presented in (Teichman et al., 2013).

Statement of Contributions. We focus specifically on triangulation Lidars for robotic applications, and aim to increase their performance in a cost-effective way through statistical processing techniques. Our long term vision is to arrive at a on-line automatic calibration procedure for triangulation Lidars like in (Kümmerle et al., 2011; Teichman et al., 2013); before reaching this above long-term goal, we must nonetheless solve satisfactorily the problem of calibrating triangulation Lidars off-line.

In this paper we thus:

- propose and assess a model for the measurement process of triangulation Lidars (see Section 3 and

model (1)). Our model not only generalizes the model proposed in (Lima et al., 2015; Campos et al., 2016), but also motivates it starting from mechanical and physical interpretations;

- on top of this model, propose and assess a ML calibration procedure that uses data from a Motion Capture (MoCap) system. Importantly, our calibration procedure extends the one proposed in (Lima et al., 2015): there authors decoupled the learning process into two separate stages (corresponding to estimate two different sets of parameters), while here the calibration is performed simultaneously on both sets of parameters;

- propose and assess novel ML and LS strategies for correcting the measurements from the sensor with the model inferred during the calibration stage.

As reported in (31) and (32), the overall strategy is then shown to be capable to improve the normalized MSE of the raw information from the sensor from 0.0789 to 0.0046.

1.1 Organization of the Manuscript

Section 2 describes the working principles of triangulation Lidars. Based on these working principles, Section 3 proposes a statistical model of the measurement process of the device. Section 4 then validates this statistical model using data acquired through a MoCap system. Section 5 then presents a calibration algorithm for sensors deployed in a test environment. Section 7 eventually concludes the paper with the description of future research issues.

2 The Triangulation Lidar Range Sensor

We now describe the functioning principle of the triangulation scanners; this discussion will be useful for explaining why the moments of the measurement noise depend on the actual measured distance. More details about the constructive details of triangulation Lidars can be found in (Blais, 2004; Konolige et al., 2008).

A prototypical triangulation Lidar is the one in Figure 1. Its working principles are then explained with the diagram in Figure 2 and its caption.

This simple working principle helps keeping the cost of the sensor low[1], and making it commercially

---

[1]Incidentally, the sensor was costing $135.00 as of February 2016 in Ebay. Nonetheless, the original industrial goal was to reach an end user price of $30.00.
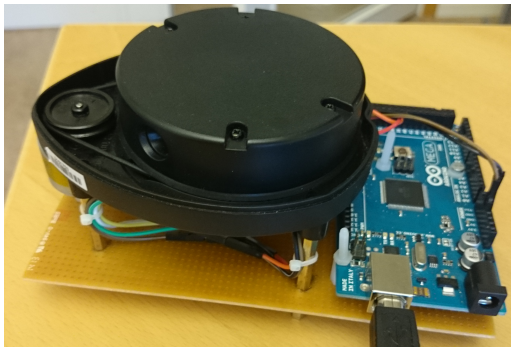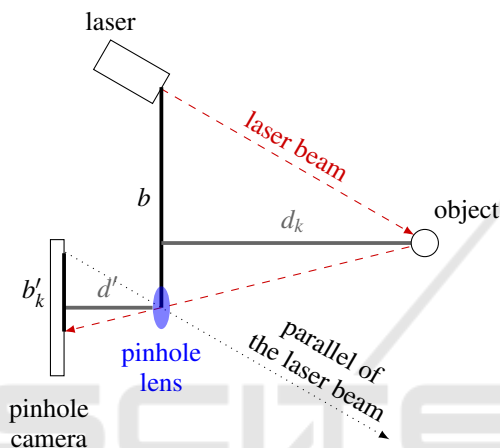
Figure 1: Photo of a triangulation Lidar.

Figure 2: Diagram exemplifying the working principle of a triangulation Lidar . The laser emits an infra-red laser signal that is then reflected by the object to be detected. The beam passes through a pinhole lens and hits a CCD camera sensor. By construction, thus, the triangles defined by (b, d_k) and by (b'_k, d') are similar: this means that the distance to the object is nonlinearly proportional to the angle of the reflected light, and as soon as the camera measures the distance b'_k one can estimate the actual distance d_k using triangles similarities concepts.

usable in low-cost devices like robotic vacuum cleaners. The low cost of the sensor comes nonetheless with some well-defined mechanical problems (Konolige et al., 2008):

- low-cost lens, that generate nonlinear distortion effects;

- imprecise pointing accuracy, that is known of at best 6 degrees;

- not rigid physical linkages among lens elements, camera, laser, and laser optics, that may suffer from distortion effects during the life of the device.

As it can be seen in Figure 3, all these problems induce measurement errors; more precisely, triangulation Lidars suffer from strong nonlinearities in both

the bias and the standard deviation of the measurement noise. This pushes towards finding some signal processing tools that can alleviate these problems, and keep the sensor cheap while improving its performance.
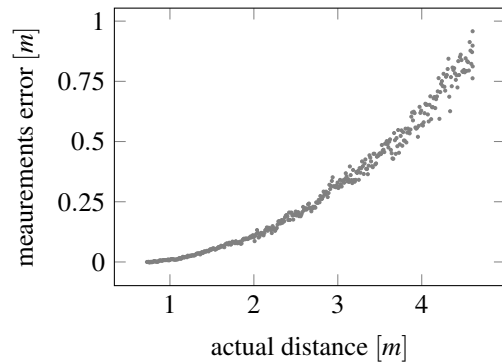

Figure 3: Dataset.

3 A NOVEL STATISTICAL MODEL FOR THE Lidar MEASUREMENTS

Let y_k be the k-th measurement returned by the Lidar when the true distance is d_k. Physically, y_k is computed by the logic of the sensor through a static transformation of b'_k in Figure 2; we assume here that this static transformation is unknown, that b'_k is not available, and that we want to improve the estimation for d_k from just y_k.

Our ansatz for the whole transformation from d_k to y_k is then

$$y_k = f(d_k) + f(d_k)^2 e_k \qquad (1)$$

where

- $f(\cdot)$ is an unknown non-linear function;

- $e_k \sim \mathcal{N}\left(0, \sigma_e^2\right)$ is a Gaussian and white additive measurement noise.

In the following Section 3.1 we motivate the presence of $f(\cdot)$ from mechanical considerations, while in the following Section 3.2 we motivate the presence of the $f(\cdot)^2$ multiplying the noise e_k starting from physical considerations.

3.1 Explaining the Presence of the Nonlinear Term $f(\cdot)$ in Model (1)

The nonlinear term $f(\cdot)$ in (1) is related to what is called the *radial distortion* in camera calibration literature (Zhang, 2000; Weng et al., 1992; Brown, 1964;

Duane, 1971). Indeed camera lenses are notoriously nonlinear at their borders, with this nonlinearity increasing as the light beam passes closer to the lens edges. In our settings this thus happens when targets are very close or very far.

Radial distortions are usually modeled in the camera calibration literature as a series of odd powers, i.e., as

$$f(d_k) = \sum_{i=0}^{n} \alpha_i d_k^{2i+1} \qquad (2)$$

where the α_i's are the model parameters.

As numerically shown during the validation of (1) in Section 4, model (2) does not describe well the evidence collected in our experiments. Indeed the specific case of triangulation Lidars lacks of the symmetries encountered in computer vision settings (see (4) and the discussion on that identity), and thus in our settings there is no need for odd symmetries in the model (in other words, doubling d does not lead to doubling b'). We thus propose to remove this constraint and use a potentially non-symmetric polynomial, i.e.,

$$f(d_k) = \sum_{i=0}^{n} \alpha_i d_k^{i}. \qquad (3)$$

The numerical validations of model (3) shown in Section 4 confirm then our physical intuition.

3.2 Explaining the Presence of the Multiplicative Term $f(d_k)^2$ in Model (1)

Assume for now that there are no lens-distortion effects. The similarity between the triangles in Figure 2 then implies

$$\frac{d_k}{b} = \frac{d'}{b'_k}. \qquad (4)$$

In (4) d_k and b'_k are generally time-varying quantities, while b and d' are constants from the geometry of the Lidar. Assume now that the quantity measured by the CCD at time k is corrupted by a Gaussian noise, so that $z_k = b'_k + w_k$ with $w_k \sim \mathcal{N}\left(0, \sigma_{\mathrm{CCD}}^2\right)$ and σ_{CCD}^2 constant and independent of d_k. Thus $z_k \sim \mathcal{N}\left(b'_k, \sigma_{\mathrm{CCD}}^2\right)$; since

$$y_k = \frac{bd'}{z_k}, \qquad (5)$$

assuming a Gaussian measurement noise on the CCD implies that y_k is a reciprocal Gaussian r.v. This kind of variable is notoriously difficult to treat (e.g., their statistical moments cannot be derived from closed form expressions starting from the original Gaussian variables). For this reason we perform a first order

Taylor approximation of the nonlinear map (5) above. In general, if

$$\begin{cases} z_k \sim \mathcal{N}\left(b, \sigma^2\right) \\ y_k = \phi(z_k) \end{cases} \qquad (6)$$

then the first order Taylor approximation of the distribution of y_k is (Gustafsson, 2010, (A.16))

$$y_k \sim \mathcal{N}\left(\phi(b), \phi'(b)^2 \sigma^2\right) \qquad (7)$$

where $\phi'(\cdot)$ is the first derivative of $\phi(\cdot)$ w.r.t. z_k. Substituting the values of our specific problem into formula (7) leads then to the novel approximated model

$$y_k \sim \mathcal{N}\left(\frac{bd'}{b'_k}, \left(\frac{-bd'}{b'^2_k}\right)^2 \sigma_{\mathrm{CCD}}^2\right), \qquad (8)$$

or, equivalently,

$$y_k = d_k + d_k^2 e_k \qquad e_k \sim \mathcal{N}\left(0, \sigma_e^2\right) \qquad (9)$$

where $\sigma_e^2 = \dfrac{\sigma_{\mathrm{CCD}}^2}{b^2 d'^2}$ is a scaled version of σ_{CCD}^2 independent of d_k and to be estimated from the data.

Consider now that actually there are some lens distortion effects that imply the presence of the nonlinear term $f(d_k)$. We can then repeat the very same discussion above, and obtain model (1) by substituting d_k with $f(d_k)$ in (9).

4 VALIDATION OF THE APPROXIMATION (8)

The approximation introduced by the first order Taylor expansion in (8) can be seen as arbitrary. Nonetheless we show in this section that on the collected datasets it actually corresponds to the most powerful approximation in a statistical sense.

To this aim we perform this two-step validation:

1. *(check if the noises are independent and identically distributed (iid) and normal)* perform a normality test on the y_k's assuming that measurements are collected at a fixed distance (i.e., d_k is constant): indeed e_k is approximately Gaussian as much as y_k is;

2. *(check the order of the term multiplying e_k)* compare the following alternative statistical models for the measurements y_k:

$$\begin{aligned} H^0: & \quad y_k = f(d_k) + e_k \\ H^1: & \quad y_k = f(d_k) + f(d_k)e_k \\ H^2: & \quad y_k = f(d_k) + f(d_k)^2 e_k \\ H^3: & \quad y_k = f(d_k) + f(d_k)^3 e_k \end{aligned} \qquad (10)$$

and check which one describes better the collected information.

As for point 1 we can use standard iid tests (like the Wald-Wolfowitz runs (Croarkin and Tobias, 2006)) and standard normality tests (like the Shapiro-Wilk normality test). These tests performed on our registered data showed p-values of 0.56 and 0.42, so we can safely consider the measurement noises to be iid and Gaussian.

As for point 2, we instead consider the following strategy: for every model above, assuming that measurements are collected at a fixed distance (i.e., d_k is constant), we can perform a simple algebraic manipulation of (1) to obtain

$$\frac{y_k - y_{k-1}}{f(d_k)^\star} = e_k - e_{k-1} \qquad (11)$$

where $\star$ indicates the order of the model (that means $\star \in \{0, \ldots, 3\}$). (11) in its turn indicates that, since e_k and e_{k-1} are assumed iid,

$$\frac{y_k - y_{k-1}}{f(d_k)^\star} \sim \mathcal{N}\left(0, 2\sigma_e^2\right), \quad \star \in \{0, \ldots, 3\}. \qquad (12)$$

Assume now that the dataset is composed by different batches each corresponding to d_k's that are constant in the batch, but different among batches. Moreover assume that each batch is sufficiently rich to make it is possible to estimate with good confidence the unknown $f(d_k)$ through the empirical mean of the y_k relative to that batch. By combining the information from different batches it is then possible to check which model $\star$ describes better the measured information.

Indicate then with B the number of batches in the dataset, with $b = 1, \ldots, B$ the index of each batch, and with $\mathcal{B}_b$ the set of k's that are relative to that specific batch b. In formulas, we thus:

1. estimate, for each model batch $b = 1, \ldots, B$, the distance

$$\widehat{f_b} = \frac{1}{|\mathcal{B}_b|} \sum_{k \in \mathcal{B}_b} y_k; \qquad (13)$$

2. estimate, for each model $\star = 0, \ldots, 3$, the variance of e_k as

$$\widehat{\sigma_e^2} := \frac{1}{B} \sum_{b=1}^{B} \left(\frac{1}{2|\mathcal{B}_b|} \sum_{k, k-1 \in \mathcal{B}_b} \left(\frac{y_k - y_{k-1}}{\widehat{f_b}^\star} \right)^2 \right). \qquad (14)$$

3. compute, for each model $\star = 0, \ldots, 3$, the log-likelihood of the data as

$$-\log \mathbb{P}\left[\boldsymbol{y} ; \boldsymbol{d}, \widehat{\sigma_e^2} \right] = \sum_{b=1}^{B} \left(|\mathcal{B}_b| \log \left(\widehat{f_b}^{2\star} \widehat{\sigma_e^2} \right) + \sum_{k=1}^{\mathcal{B}_b} \frac{\left(y_k - \widehat{f_b} \right)^2}{\widehat{f_b}^{2\star} \widehat{\sigma_e^2}} \right) \qquad (15)$$

where $\boldsymbol{y} := [y_1, \ldots, y_N]^T$ and $\boldsymbol{d} := [d_1, \ldots, d_N]^T$.

In Figure 4 we then show the log-likelihoods for the different models. *As it can be seen, hypothesis H^2 is the one that best describes the collected evidence.*
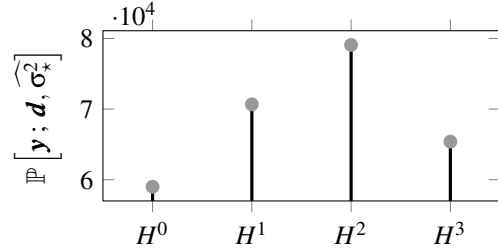

Figure 4: Evaluation of (15) on the collected datasets.

A non rigorous (but graphical and intuitive) argument supporting H^2 as the hypothesis best describing the evidence is then the one offered in Figure 5. The argument goes as follows: for the exact $\star \in \{0, \ldots, 3\}$ the quantities

$$\frac{y_k - y_{k-1}}{f(d_k)^\star} \qquad \star \in \{0, \ldots, 3\}. \qquad (16)$$

should be iid independently of d_k. This iid-ness is indeed a necessary condition for iid-ness of the measurement noises (one of our assumptions).

Since $f(\cdot)$ is actually unknown, this iid-ness test must be performed by means of some estimate of $f(\cdot)$. In the following we use the estimator defined in Section 5 over an experiment where we manually increase the true distance d_k. As it can be seen, the hypothesis H^2 is the unique one for which the quantities $\frac{y_k - y_{k-1}}{\widehat{f}(d_k)^\star}$ are homoscedastic. Thus the normalizing factor $\star = 2$ is the unique one guaranteeing iid-ness for the measurement noises. Notice that this argument is a non rigorous wishful thinking, since we use some estimates as the ground truth; nonetheless the heteroscedasticity of the noises for $\star = 0, 1, 3$ indicates that these hypotheses are non-descriptive.

5 CALIBRATING THE Lidar

Our overall goal is not just to propose the statistical model (1) describing the measurement process of the Lidar but also to find a calibration procedure for estimating the unknowns $f(\cdot)$ and σ_e^2 from some collected information.

Once again the long term goal is to calibrate (1) on-line and continuously using information from other sensors like odometry, ultrasonic sensors, etc. Instrumental to this future direction we now solve the first step, that is to estimate $f(\cdot)$ and σ_e^2 from a dataset
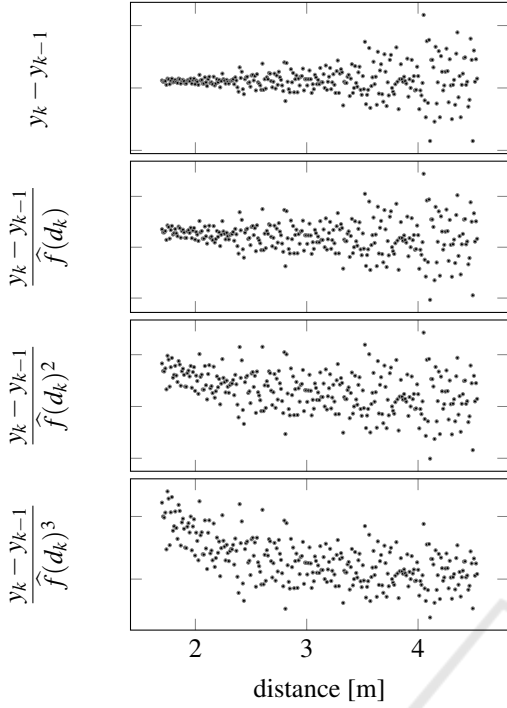
Figure 5: Plots of the quantities $\dfrac{y_k - y_{k-1}}{\widehat{f}(d_k)^\star}$ for $\star = 0, \ldots, 3$ and for increasing d_k and for $\widehat{f}(\cdot)$ computed as in Section 5. The results graphically suggest that $\widehat{f}(d_k)^2$ is the unique normalizing factor for which we obtain homoscedastic samples.

$\mathcal{D} = \{y_k, d_k\}$ in which we know d_k (e.g., thanks to a MoCap system).

Given our Fisherian setting, we seek for the ML estimate for both $f(\cdot)$ and σ_e^2, where we recall that (due to the radial distortion hypothesis as the source of $f(\cdot)$, see Section 3.1) $f(\cdot)$ is modeled as a non-symmetric polynomial, i.e., as $f(d_k) = \sum_{i=0}^{n} \alpha_i d_k^i$ as in (3). Since now model (1) implies

$$y_k - f(d_k) \sim \mathcal{N}\left(0, f(d_k)^4 \sigma_e^2\right), \qquad (17)$$

it follows immediately that the corresponding negative log-likelihood is proportional to

$$\mathcal{L} := \log\left(\det\Sigma\right) + \left(\boldsymbol{y} - \boldsymbol{f}(\boldsymbol{d})\right)^T \Sigma^{-1}\left(\boldsymbol{y} - \boldsymbol{f}(\boldsymbol{d})\right) \quad (18)$$

where

- $\boldsymbol{y} := [y_1, \ldots, y_N]^T$;
- $\boldsymbol{d} := [d_1, \ldots, d_N]^T$;
- $\boldsymbol{f}(\boldsymbol{d}) := [f(d_1), \ldots, f(d_N)]$;
- $\Sigma := \text{diag}\left(f(d_1)^4 \sigma_e^2, \ldots, f(d_N)^4 \sigma_e^2\right)$.

Finding the ML estimates in our settings thus means:

1. solving

$$\arg\min_{\theta \in \Theta} \mathcal{L}(\theta) \qquad (19)$$

for several different n, with

$$\theta := \left[\alpha_0, \ldots, \alpha_n, \sigma_e^2\right] \qquad (20)$$

and Θ the set of $\theta \in \mathbb{R}^{n+1}$ for which $\sigma_e^2 > 0$;

2. deciding which n is the best one using some model order selection criterion, e.g., AIC.

Unfortunately problem (19) is not convex, so it neither admits a closed form solution nor it can be easily computed using numerical procedures. Solving problem (19) is thus numerically difficult. Keeping in mind that our long-term goal is the development of on-line calibration procedures, where numerical problems will be even more complex, we strive for some alternative calibration procedure.

5.1 An Approximate Calibration Procedure

We here propose an alternative estimator that trades off statistical performance for solvability in a closed form. We indeed propose to seek an estimate for θ in (20) by using the alternative model

$$y_k = f(d_k) + d_k^2 e_k, \qquad (21)$$

that differs from (1) only for the fact that the noise is multiplied by d_k^2 instead of $f(d_k)^2$. This approximation is intuitively meaningful, since $f(d_k)$ represents a distortion term induced by the pinhole lens: ideally, indeed, $f(d_k)$ should be equal to d_k.

Assuming model (21) it is now possible do derive a ML estimator of θ. Indeed dividing both sides of (21) by d_k^2 we get

$$\frac{y_k}{d_k^2} = g(d_k) + e_k \qquad (22)$$

where (cf. (3))

$$g(d_k) = \sum_{i=0}^{n} \alpha_i d_k^{i-2}. \qquad (23)$$

This means that the estimation problem can be cast as the problem of estimating the parameters $\boldsymbol{\alpha} := [\alpha_0, \ldots, \alpha_n]^T$ and the noise variance σ_e^2 describing the linear system

$$\frac{y_k}{d_k^2} = \begin{bmatrix} d_k^{-2} & \cdots & d_k^{n-2} \end{bmatrix} \begin{bmatrix} \alpha_0 \\ \vdots \\ \alpha_n \end{bmatrix} + e_k, \qquad (24)$$

for which the ML solution is directly

$$\begin{aligned} \widehat{\boldsymbol{\alpha}} &= \left(H^T H\right)^{-1} H^T \widetilde{\boldsymbol{y}} \\ \widehat{\sigma_e^2} &= \frac{1}{N}\left(\widetilde{\boldsymbol{y}} - H\widehat{\boldsymbol{\alpha}}\right)^T \left(\widetilde{\boldsymbol{y}} - H\widehat{\boldsymbol{\alpha}}\right) \end{aligned} \qquad (25)$$

with

$$H := \begin{bmatrix} d_1^{-2} & \cdots & d_1^{n-2} \\ \vdots & & \vdots \\ d_N^{-2} & \cdots & d_N^{n-2} \end{bmatrix} \quad \widetilde{\mathbf{y}} := \begin{bmatrix} \dfrac{y_1}{d_1^2} \\ \vdots \\ \dfrac{y_N}{d_N^2} \end{bmatrix}. \quad (26)$$

Notice that the procedure above does not determine the model complexity n. For inferring this parameter we then propose to rely on classical model order selection criteria such as AIC.

5.2 Using the Calibration Results to Estimate d_k

Once the sensor has been calibrated, i.e., a $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\sigma_e^2}$ have been computed, it is possible to invert the process and use the learned information for testing purposes. This means that given some measurements y_k collected in an unknown environment we can, through $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\sigma_e^2}$, estimate d_k.

5.2.1 Computing the ML Estimate of d_k

Rewriting model (3) as

$$f(d_k) = \boldsymbol{d}_k^T \boldsymbol{\alpha} \qquad \boldsymbol{d}_k := \begin{bmatrix} d_k^0 \\ d_k^1 \\ \vdots \\ d_k^n \end{bmatrix} \quad (27)$$

and equating the score of y_k parametrized by $\boldsymbol{\alpha}$ and σ_e^2 to zero leads to the equation

$$\left(y_k - \boldsymbol{d}_k \boldsymbol{\alpha}\right)\left(y_k - \boldsymbol{d}_k\left(I - K\right)\boldsymbol{\alpha}\right) = \sigma_e^2 d_k^4 \quad (28)$$

with

$$K := \mathrm{diag}\left(0, \frac{1}{2}, \ldots, \frac{n}{2}\right). \quad (29)$$

This means that estimating d_k from y_k, $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\sigma_e^2}$ can be performed by solving (28) in d_k after substituting the real values $\boldsymbol{\alpha}$ and σ_e^2 with their estimates.

Since polynomial (28) is quartic for $n = 0, 1, 2$, and of order at least 6 for any other n, the ML estimate for d_k must then either rely on complex algebraic formulas or numerical roots finding methods.

5.2.2 Computing the LS Estimate of d_k

Given our assumption (3) on the structure of $f(\cdot)$, and given an estimate $\widehat{f}$ for f, the problem of estimating d_k from y_k is the one of minimizing the squared loss $\left(y_k - \widehat{f}(d_k)\right)^2$. Once again, the problem is of finding

the roots of a polynomial, since the solutions of the LS problem above are directly

$$\widehat{d}_k \in \left\{ \widetilde{d} \text{ s.t. } y_k - \widehat{f}\left(\widetilde{d}\right) = 0 \right\}. \quad (30)$$

Thus if the Lidar has heavy nonlinear radial distortions (that means that it requires high order polynomials $f(\cdot)$) then one is again required to compute polynomial roots.

Notice also that some of the roots above may not belong to the measurement range of the sensor (e.g., some roots may be negative); these ones can safely be discarded from the set of plausible solutions. The other ones, instead, are equally plausible.

This raises a question on how to decide which root should be selected among the equally plausible ones. This question is actually non-trivial, and cannot be solved by means of the frequentist approach used in this manuscript. We thus leave this question unanswered for now, and leave it as a future research question. Bayesian formulations will be explored to see if they solve the multiple plausible roots problem.

6 NUMERICAL EXPERIMENTS

Our experiments consist of a robot with the Lidar mounted on top moving with piecewise constant speeds towards a target. We recorded several datasets for training and testing purposes, consisting of the Lidar measurements and a ground truth information collected by a MoCap system (see Figure 6). Training

Figure 6: Experimental setup used for recording the dataset. The Lidar was mounted over a Pioneer 3AT robot facing an obstacle; the photo moreover shows some of the cameras of the MoCap system.

datasets were thus initially used to estimate $\boldsymbol{\alpha}$ and σ_e^2 as described in Equation (25). As for the model order selection, we empirically detected that $n = 2$ was always the best choice when using AIC measures. E.g.,
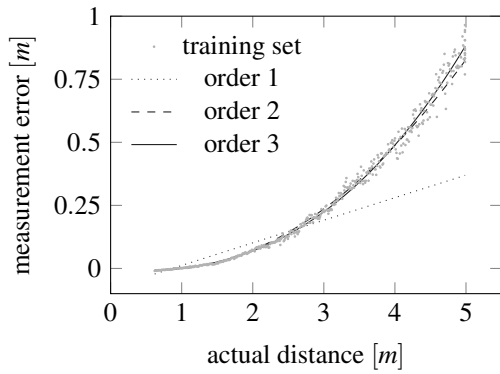
Figure 7: A typical training set collected in our experiments. The plotted quantities correspond to the measurement errors and to the polynomial models fitting these errors.

Table 1: AIC scores for the different models complexities involved in the training set of Figure 7.

polynomial order	AIC score
1	-5.774
2	-7.380
3	-5.824
4	-3.890

for the dataset shown in Figure 7 we obtained the AIC scores reported in Table 1.

The estimated α and σ_e^2 were also used for testing purposes to refine the estimate of the distances d_k in non-controlled environments. Notice that the selected model order was always 2, so it was always possible to solve the LS problem in a closed form and also discard one of the roots in (30), so that the set of roots was always a singleton. As shown in Figure 8, $\widehat{d_k}$ is much closer to d_k than y_k. For example, the empirical normalized MSEs for the test set in Figure 8 were

$$\frac{1}{N}\sum_{k=1}^{N}\frac{\left\|\widehat{d_k}-d_k\right\|^2}{\|d_k\|^2}=0.0046, \qquad (31)$$

$$\frac{1}{N}\sum_{k=1}^{N}\frac{\|y_k-d_k\|^2}{\|d_k\|^2}=0.0789. \qquad (32)$$

7 CONCLUSIONS

We derived, starting from a combination of physical and statistical considerations, a model that describes the statistical behavior of the measurements returned by triangulation Lidars. This statistical model, given in (1), is based on two assumptions:

1. the effects of radial distortions in the pinhole lens can be captured by means of a polynomial function;
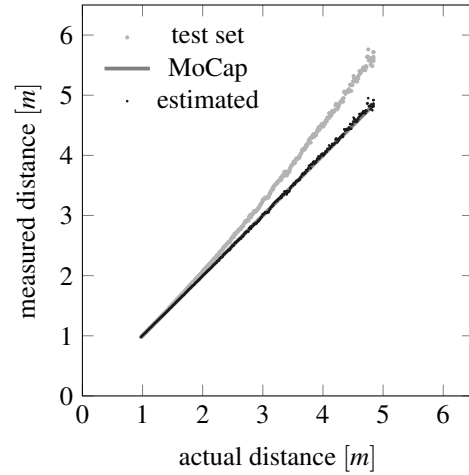

Figure 8: Effects of the estimation procedure on the original Lidar measurement. It can be noticed how the overall strategy removes the nonlinearities induced by the pinhole lens - CCD camera system.

2. the nonlinearities induced by the geometry of the laser-CCD system can be captured by means of a heteroscedastic noise which standard deviation depends in first approximation quadratically with the measured distance.

This model, validated through some experiments on real devices, allows to build tailored triangulation Lidars calibration strategies that follow the classical training-testing paradigm:

- in the training phase, collect information in a controlled environment and use it to estimate through ML paradigms the parameters defining the statistical behavior of the sensor;

- in the test phase, use this information and some statistical inference techniques such as ML or LS to correct the measurements from the sensor when this is in a non-controlled environment.

It turns then out that both the ML and LS estimation strategies may be numerically demanding, specially for sensors suffering from strong radial distortions in the pinhole camera. In this case, indeed, the estimators may require to use numerical root finding procedures and lead to some computational disadvantages.

Irrespectively of these issues, that can in any case be mitigated by limiting the complexity of the polynomials describing the radial distortions, the estimation strategies above have been proved to be effective in our tests. Real-life experiments indeed showed that the techniques allow to reduce the empirical MSE of the sensor of a factor 17.15.

Despite this promising result, the research associated to triangulation Lidars is not finished. First

of all, the techniques should be modified so to be implementable using recursive estimation schemes. Moreover, by following a classical training-testing approach, the techniques above present some limitations. Different sensors may in fact differ even if nominally being constructed in the same way. Moreover sensors may change their statistical behavior in time, due to aging or mechanical shocks. This means that techniques based on results from a controlled environment on just one sensor and just once are eventually not entirely meaningful.

A robust approach must indeed perform continuous learning for each sensor independently in a non-controlled environment by performing information fusion steps, e.g., combining also information from other sensors like odometry, ultrasonic and accelerometers.

This information-fusion continuous-learning algorithm nonetheless must be based on some preliminary results on what are the statistical models of triangulation Lidars and on how inference can be performed on them. This paper can thus be seen as the first step towards more evolved strategies.

REFERENCES

Alhashimi, A., Varagnolo, D., and Gustafsson, T. (2015). Joint temperature-lasing mode compensation for time-of-flight lidar sensors. *Sensors*, 15(12):31205–31223.

Anderson, D., Herman, H., and Kelly, A. (2005). Experimental characterization of commercial flash ladar devices. In *International Conference of Sensing and Technology*, volume 2.

Andreasson, H., Triebel, R., and Burgard, W. (2005). Improving plane extraction from 3d data by fusing laser data and vision. In *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 2656–2661. IEEE.

Atanacio-Jiménez, G., González-Barbosa, J.-J., Hurtado-Ramos, J. B., Ornelas-Rodríguez, F. J., Jiménez-Hernández, H., García-Ramirez, T., and González-Barbosa, R. (2011). Lidar velodyne hdl-64e calibration using pattern planes. *International Journal of Advanced Robotic Systems*, 8(5):70–82.

Blais, F. (2004). Review of 20 years of range sensor development. *Journal of Electronic Imaging*, 13(1).

Brown, D. C. (1964). An advanced reduction and calibration for photogrammetric cameras. Technical report, DTIC Document.

Campos, D., Santos, J., Gonçalves, J., and Costa, P. (2016). Modeling and simulation of a hacked neato xv-11 laser scanner. In *Robot 2015: Second Iberian Robotics Conference*, pages 425–436. Springer.

Chen, C.-Y. and Chien, H.-J. (2012). On-site sensor recalibration of a spinning multi-beam lidar system using automatically-detected planar targets. *Sensors*, 12(10):13736–13752.

Croarkin, C. and Tobias, P. (2006). Nist/sematech e-handbook of statistical methods. *NIST/SEMATECH, July. Available online: http://www.itl.nist.gov/div898/handbook.*

Duane, C. B. (1971). Close-range camera calibration. *Photogrammetric engineering*, 37(8):855–866.

Glennie, C. (2012). Calibration and kinematic analysis of the velodyne hdl-64e s2 lidar sensor. *Photogrammetric Engineering & Remote Sensing*, 78(4):339–347.

Glennie, C. and Lichti, D. D. (2010). Static calibration and analysis of the velodyne hdl-64e s2 for high accuracy mobile scanning. *Remote Sensing*, 2(6):1610–1624.

Glennie, C. and Lichti, D. D. (2011). Temporal stability of the velodyne hdl-64e s2 scanner for high accuracy scanning applications. *Remote Sensing*, 3(3):539–553.

Gong, X., Lin, Y., and Liu, J. (2013). 3d lidar-camera extrinsic calibration using an arbitrary trihedron. *Sensors*, 13(2):1902–1918.

Gordon, M. and Meidow, J. (2013). Calibration of a multi-beam laser system by using a tls-generated reference. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences II-5 W*, 2:85–90.

Gustafsson, F. (2010). *Statistical sensor fusion*. Studentlitteratur,.

Jokinen, O. (1999). Self-calibration of a light striping system by matching multiple 3-d profile maps. In *3-D Digital Imaging and Modeling, 1999. Proceedings. Second International Conference on*, pages 180–190. IEEE.

Kneip, L., Tâche, F., Caprari, G., and Siegwart, R. (2009). Characterization of the compact hokuyo urg-04lx 2d laser range scanner. In *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, pages 1447–1454. IEEE.

Konolige, K., Augenbraun, J., Donaldson, N., Fiebig, C., and Shah, P. (2008). A low-cost laser distance sensor. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 3002–3008. IEEE.

Kümmerle, R., Grisetti, G., and Burgard, W. (2011). Simultaneous calibration, localization, and mapping. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 3716–3721. IEEE.

Lee, K.-H. and Ehsani, R. (2008). Comparison of two 2d laser scanners for sensing object distances, shapes, and surface patterns. *Computers and electronics in agriculture*, 60(2):250–262.

Lima, J., Gonçalves, J., and Costa, P. J. (2015). Modeling of a low cost laser scanner sensor. In *CONTROLO'2014– Proceedings of the 11th Portuguese Conference on Automatic Control*, pages 697–705. Springer.

McIvor, A. M. (1999). Calibration of a laser stripe profiler. In *3-D Digital Imaging and Modeling, 1999. Proceedings. Second International Conference on*, pages 92–98. IEEE.

Mei, C. and Rives, P. (2006). Calibration between a central catadioptric camera and a laser range finder for robotic applications. In *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, pages 532–537. IEEE.

Mirzaei, F. M., Kottas, D. G., and Roumeliotis, S. I. (2012). 3d lidar–camera intrinsic and extrinsic calibration: Identifiability and analytical least-squares-based initialization. *The International Journal of Robotics Research*, 31(4):452–467.

Muhammad, N. and Lacroix, S. (2010). Calibration of a rotating multi-beam lidar. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 5648–5653. IEEE.

Park, Y., Yun, S., Won, C. S., Cho, K., Um, K., and Sim, S. (2014). Calibration between color camera and 3d lidar instruments with a polygonal planar board. *Sensors*, 14(3):5333–5353.

Reina, A. and Gonzales, J. (1997). Characterization of a radial laser scanner for mobile robot navigation. In *Intelligent Robots and Systems, 1997. IROS'97., Proceedings of the 1997 IEEE/RSJ International Conference on*, volume 2, pages 579–585. IEEE.

Sanz-Cortiella, R., Llorens-Calveras, J., Rosell-Polo, J. R., Gregorio-Lopez, E., and Palacin-Roca, J. (2011). Characterisation of the lms200 laser beam under the influence of blockage surfaces. influence on 3d scanning of tree orchards. *Sensors*, 11(3):2751–2772.

Tang, P., Akinci, B., and Huber, D. (2009). Quantification of edge loss of laser scanned data at spatial discontinuities. *Automation in Construction*, 18(8):1070–1083.

Teichman, A., Miller, S., and Thrun, S. (2013). Unsupervised intrinsic calibration of depth sensors via slam. In *Robotics: Science and Systems*. Citeseer.

Tiddeman, B., Duffy, N., Rabey, G., and Lokier, J. (1998). Laser-video scanner calibration without the use of a frame store. In *Vision, Image and Signal Processing, IEE Proceedings-*, volume 145, pages 244–248. IET.

Tuley, J., Vandapel, N., and Hebert, M. (2005). Analysis and removal of artifacts in 3-d ladar data. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 2203–2210. IEEE.

Wei, G.-Q. and Hirzinger, G. (1998). Active self-calibration of hand-mounted laser range finders. *Robotics and Automation, IEEE Transactions on*, 14(3):493–497.

Weng, J., Cohen, P., and Herniou, M. (1992). Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 14(10):965–980.

Ye, C. and Borenstein, J. (2002). Characterization of a 2-d laser scanner for mobile robot obstacle negotiation. In *ICRA*, pages 2512–2518.

Zhang, Q. and Pless, R. (2004a). Constraints for heterogeneous sensor auto-calibration. In *Computer Vision and Pattern Recognition Workshop, 2004. CVPRW'04. Conference on*, pages 38–38. IEEE.

Zhang, Q. and Pless, R. (2004b). Extrinsic calibration of a camera and laser range finder (improves camera calibration). In *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 3, pages 2301–2306. IEEE.

Zhang, Z. (2000). A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334.