

From Chaos to Calibration: A Geometric Mutual Information Approach to Target-Free Camera LiDAR Extrinsic Calibration

Jack Borer Jeremy Tschirner Florian Ölsner Stefan Milz
 Spleenlab GmbH

Abstract

Sensor fusion is vital for the safe and robust operation of autonomous vehicles. Accurate extrinsic sensor to sensor calibration is necessary to accurately fuse multiple sensor's data in a common spatial reference frame. In this paper, we propose a target free extrinsic calibration algorithm that requires no ground truth training data, artificially constrained motion trajectories, hand engineered features or offline optimization and that is accurate, precise and extremely robust to initialization error.

Most current research on online camera-LiDAR extrinsic calibration requires ground truth training data which is impossible to capture at scale. We revisit analytical mutual information based methods first proposed in 2012 and demonstrate that geometric features provide a robust information metric for camera-LiDAR extrinsic calibration. We demonstrate our proposed improvement using the KITTI and KITTI-360 fisheye data set.

1. Introduction

Modern autonomous vehicles and robots are powered by algorithms that fuse data from multiple sources in order to increase confidence and reduce decision making uncertainty. This algorithmic requirement for multiple independent observations of the environment requires that sensors be accurately co-registered in 3D space. Extrinsic calibration is the process of acquiring these co-registration parameters and is a fundamental building block of an autonomous system.

Camera and LiDAR is a commonly used multi-sensor suite used in high level autonomy algorithms for self-driving cars and trucks. Accurately mapping image texture onto the LiDAR pointcloud or pointcloud texture (depth, speed or intensity) onto the camera image are two common functions that require the extrinsic calibration between camera and LiDAR. Even small errors in the rotational component of the extrinsic rigid body transform can introduce fatal errors into this data fusion process. For a review of image-pointcloud fusion applications, including extrinsic calibration

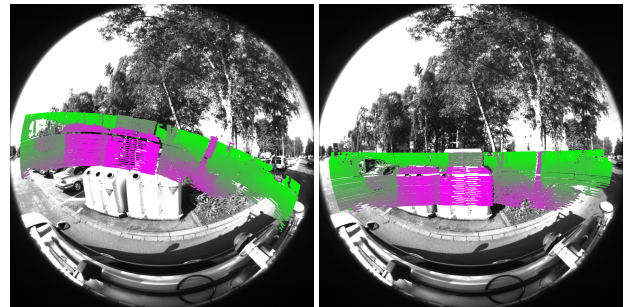


Figure 1. An accurate extrinsic calibration is required to fuse camera and LiDAR data together. A failed calibration results in a misaligned projection (left). Our target free calibration automatically provides the correct extrinsic calibration (right).

tion see the work of Cui *et al.* [5].

Many extrinsic calibration algorithms, discussed in Section 2, are challenging to automate. Therefore, it is common that extrinsic calibration is performed only once when a vehicle is factory commissioned. Naturally, this introduces the unrealistic constraint that the sensors will never move during the life of the vehicle. Online calibration methods, capable of calibrating sensors using only sensor data during the vehicles lifetime without requiring ground truth training data or artificial "calibration trajectories" are very relevant.

Our contribution is an improvement to Mutual Information (MI) camera-LiDAR extrinsic calibration that dramatically increases robustness to initialization error, greatly increases precision and is accurate. We propose using a depth-to-depth (D2D) feature instead of the legacy intensity-to-intensity (I2I) feature (Fig. 3). To support our claim we provide a detailed justification and a comprehensive set of experiments on the KITTI and KITTI-360 data sets.

2. Related Work

2.1. Feature Extraction Based

Feature extraction based extrinsic calibration algorithms can be broken down into three steps:

1. Feature extraction

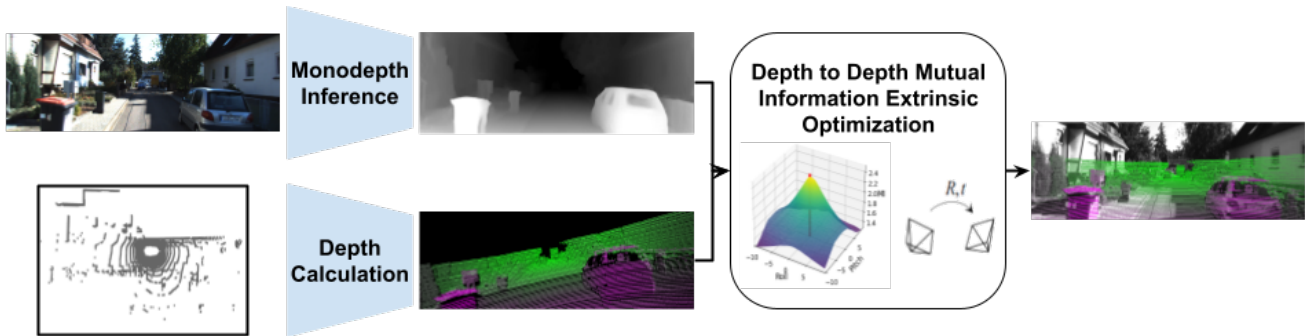


Figure 2. Our proposed depth-to-depth mutual information extrinsic calibration workflow. A monocular depth estimation network generates depth predictions from the camera image and the LiDAR pointcloud is projected to an image with intensity proportional to depth. Finally, the extrinsic calibration is optimized by maximizing the mutual information between the two intramodal measurements.

2. Feature correspondence
3. Reprojection error optimization

Target-based and target-free feature based methods follow the same three steps. Target-based methods are provided strong correspondent features by a user configured fiducial target, a checkerboard for example. This is in contrast to target-free methods which are not provided strong correspondent features and must generate features which can be easily extracted and matched from inter-modal camera-LiDAR data.

2.1.1 Target

Zhang and Pless [35] is the seminal work for target based offline camera-LiDAR calibration. In the image the checkerboard corners are detected and used to calculate the plane-normal and in the LiDAR pointcloud the points on the plane are segmented. This plane-normal and set of points on the plane is a geometric constraint which is used by a iterative nonlinear least-squares optimizer to solve for the camera-LiDAR extrinsic transformation. In follow-up work, the idea is expanded to new sensor models, initialization techniques and optimization constraints [6, 15, 17, 31].

2.1.2 Target Free

Scaramuzza *et al.* [23] removed the need for a calibration target by hand selecting features in the camera image and LiDAR pointcloud and solved a perspective-from-n-point problem followed by a non-linear refinement. Hand selected features are however tedious to collect and require an expert user.

A natural feature that exists in both camera and LiDAR data and can be semi-automatically extracted are edges. It is common that these inter-modal edges align and therefore metrics which measure edge magnitude and direction similarity can be used as a metric for camera-LiDAR extrinsic

calibration. Multiple works use edge based metrics as an optimization constraint [2, 11, 16, 29, 30]. Edge based methods are very promising for autonomous vehicle calibration because they require neither an operator or fiducial target but are highly sensitive to initialization error.

Learning based methods starting with the work of Schneider *et al.* [24] cast all three steps: feature extraction, matching and optimization into a neural network learning problem. These methods require minimal feature engineering and no specific calibration targets. Work to improve the loss function and network architecture has made progress but has not yet solved the fundamental domain gap between camera and LiDAR data [9, 10, 12, 26, 36, 37]. Notably, these methods require ground truth training data, in the form of a ground truth extrinsic calibration or semantic labels. This data is impossible to collect at scale for real world applications. Our proposed D2D optimization algorithm does not require any ground truth labeled data and is therefore easily used, in contrast to the above listed black-box over fit neural networks.

2.2. Odometry Based

Robot hand-eye calibration using structure from motion was used for the extrinsic calibration of a camera and a mobile platform with known ego-motion in the work of Andreff *et al.* [1]. With the advent of LiDAR based ego-motion algorithms, the natural extension of this was to calculate the extrinsic calibration between a camera and a LiDAR with calculated ego-motion [28]. The value of the extrinsic calibration is regressed from the camera and LiDAR trajectory just as in the hand-eye calibration process. Multiple works present examples of hand-eye camera-LiDAR extrinsic calibration algorithms [3, 8, 38]. These methods are promising for autonomous vehicle calibration because they do not require an operator, fiducial target or even overlapping sensor field of view. A flaw of odometry based methods is that most vehicle motion rotates around only a single axis (yaw) and therefore the roll and pitch are often now observable

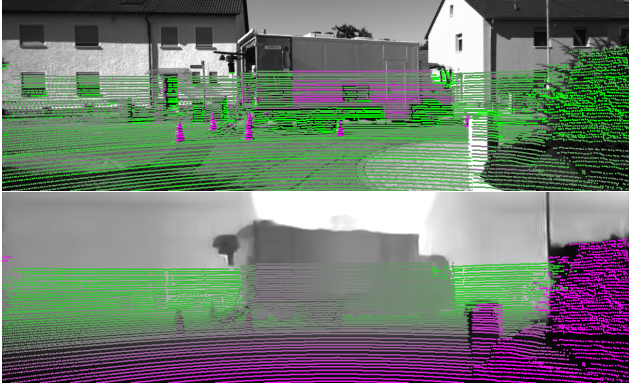


Figure 3. Comparison of intensity-to-intensity features (top) and depth-to-depth features (bottom). Depth features share significantly more information than intensity features. The depth estimate of the image (bottom) is provided by a self-supervised monocular depth estimation network that does not use any ground truth training data.

without specially designed calibration trajectories.

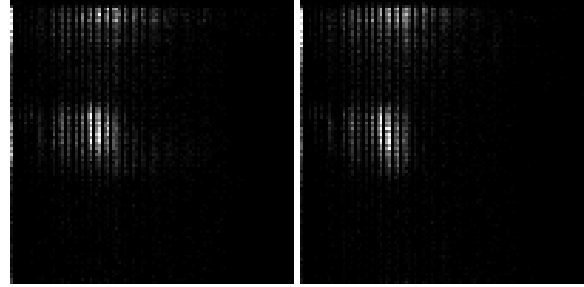
2.3. Mutual Information Based

Mutual information, also called relative entropy, is an established generic measure of the statistical dependence between two random variables [4, 32]. Maximizing the amount of mutual information, which can also be interpreted as “sharpening” the joint histogram of the distributions (Fig. 4), is the criterion for MI based extrinsic calibration algorithms. This approach is based on the assumption that at the correct extrinsic calibration the MI of the distributions is maximized. Mutual information, which makes no assumption about the distribution’s shape, form or dependency has been found to be a highly effective metric for extrinsic calibration of inter-modal medical images [13, 33, 34].

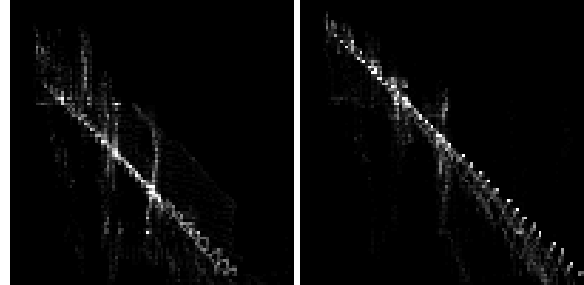
2.3.1 Robotics

Mutual information extrinsic calibration of a camera-LiDAR pair was first applied to aerial survey data [14]. Instead of maximizing the MI however, the entropy of the image and LiDAR features is assumed constant for small perturbations, and the joint entropy $H(X, Y)$ is minimized. The camera image’s intensity and the LiDAR’s elevation and probability of detection was used to construct two joint feature entropies that were minimized. Ground based surveys also require accurate extrinsic camera-LiDAR calibration. Taylor and Nieto [27] calibrated a stationary hyperspectral camera and 3D LiDAR using the MI between image intensity and LiDAR pointcloud point normal features.

Pandey *et al.* [18] applied the concept to an automotive camera-LiDAR data set, calculating the MI between



(a) De-calibrated (left, MI=0.39) and calibrated (right, MI=0.47) intensity-to-intensity histograms.



(b) De-calibrated (left, MI=2.40) and calibrated (right, MI=3.05) depth-to-depth histograms.

Figure 4. Maximizing mutual information is equivalent to sharpening the joint feature histogram. The depth feature histogram (bottom row) is more structured and sharpens significantly more than the intensity feature histogram (top row) for the same recalibration. Our proposed depth features are directly correlated therefore we see a linear relationship between the camera and LiDAR feature.

grayscale image intensity and LiDAR intensity. Interestingly, using pointcloud point depth as the feature, instead of intensity, was proposed but not evaluated. This was followed up by a 2015 paper [19] which expanded the idea to multiple new sensing modalities and environments. In both papers it was noted that environments with features close to the sensors and with high correlated image intensity and LiDAR intensity (i.e. metal signs, road paint or cars) result in better calibrations. Our contribution made in this paper is a direct successor of and largely inspired by this work.

2.4. Our Contribution

We propose maximizing geometric feature mutual information as the optimization metric for camera-LiDAR extrinsic calibration. Advances in monocular depth estimation [7, 21] have enabled the extraction of geometric information from camera images without the need for ground truth labeled training data. Our method takes advantage of this fact and calculates the MI between the inferred camera depth image and the LiDAR pointcloud’s inherent geometric depth feature.

The result is a target free camera-LiDAR calibration algorithm that requires no ground-truth training data or hand-

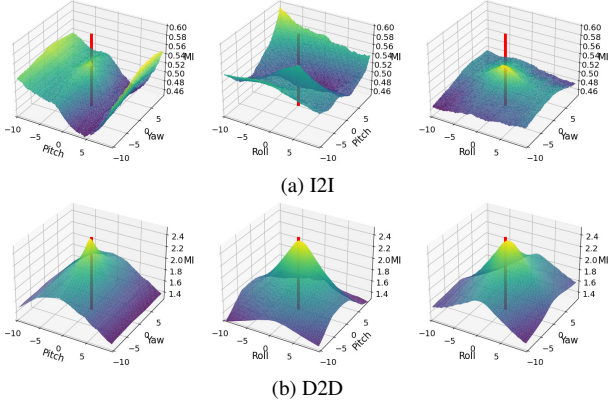


Figure 5. Mutual information surface plots of (i) classic I2I, and (ii) our proposed D2D feature from 25 frames. The D2D MI is clearly more suitable for optimization over a larger range of errors than I2I.

crafted features and is explainable. Our contribution extends naturally to other dense depth sensors, such as structured light sensors, time-of-flight sensors and stereo camera rigs, where image-LiDAR correspondences can easily be calculated with a camera model.

3. Theory

3.1. Statistical Theory

Statistical dependence refers to the relationship or association between two or more variables. If two variables are statistically independent, the distribution of one remains unchanged regardless of the knowledge of the other. Conversely, if two variables are dependent, knowledge of the distribution of one provides total information about the distribution of the other. These two extremes represent the limits of association, where independent variables share no information and dependent variables share complete information.

Using the definition of entropy $H = -\sum_i p_i \log p_i$ the mutual information of two variables can be written in terms of the entropy of the marginal distributions and the joint distributions with

$$I(X, Y) = H(X) + H(Y) - H(X, Y). \quad (1)$$

3.2. Measurement Theory

Camera and LiDAR both sense the surrounding environment. A camera passively captures the visible light texture and the LiDAR actively senses non-visible light texture which is then processed into a 3D position and intensity. Modeling how camera and LiDAR sensors measure intensity is notoriously challenging. For image sensors the amount of light measured is highly dependent on external

lighting conditions, object color and reflectance properties. Three variables which in commonly encountered environments are difficult to efficiently and accurately constrain. Measured LiDAR intensity is similarly difficult to model and is highly dependent on surface finish, color, and angle of incidence. Three variables which, like modeling camera intensity, are also nearly impossible to efficiently and accurately constrain.

Therefore the assumption that camera and LiDAR intensities are correlated is problematic. Practical examples of this problem are commonly found in autonomous driving data. Where for example two cars, one dark red and one black can have similar camera intensity values but completely different LiDAR intensity values. Or two road signs, one dark blue and one white that have nearly the same LiDAR intensity values but vastly different camera intensity values.

With the advent of monocular depth estimation algorithms we can now consider a third virtual depth sensor that directly maps image intensity to depth. Such algorithms are robust and can produce estimates of scene geometry which are accurate enough for our geometric MI calibration metric [7, 21]. LiDAR sensors have also experienced large improvements in their ability to measure scene geometry and are able to reliably estimate 3D position to within several centimeters at ranges over 100m.

The ability of both a virtual depth sensor and LiDAR to measure 3D geometry is the foundation of the extrinsic calibration algorithm presented in this paper.

4. Method

The foundation of our method is the mutual information based extrinsic calibration of a camera and LiDAR. We improve on the previous MI-based calibration approaches of Mastin *et al.* [14] and Pandey *et al.* [19] by using highly correlated depth features instead of poorly correlated intensity features.

4.1. Mathematical Formulation

In the following, we posit the algorithm in terms of the LiDAR coordinate frame O_L and the camera coordinate frame O_C . We denote the rigid transformation that transforms a point p_L in coordinate frame O_L to a corresponding point p_C in coordinate frame O_C by

$${}^C T_L = ({}^C R_L, {}^C t_L) \in \mathbb{R}^{4 \times 4} \quad (2)$$

where ${}^C R_L \in \text{SO}(3)$ and ${}^C t_L \in \mathbb{R}^3$ are the rotational and translational part of the transformation. Using homogeneous coordinates, a 3D point p_L in the LiDAR coordinate frame O_L can be transformed to a point p_C in the camera coordinate frame O_C with $p_C = {}^C T_L \cdot p_L$.

Using the pinhole camera model for simplicity, a homogeneous point p_C can be further transformed to a point p_{CO} in the 2D camera optical coordinate frame O_{CO} with

$$p_{CO} = P \cdot p_C \quad (3)$$

where $P = K[I|0] \in \mathbb{R}^{3 \times 4}$ is the projective camera matrix parameterized by the calibration matrix K . It should be emphasized that the algorithm is generic to the selected camera model, and in this paper we use both the pinhole model for the KITTI experiment and double sphere model for the KITTI-360 experiment.

Using both transformations, ${}^C T_L$ and P , a 3D point p_L measured in the LiDAR coordinate frame can be projected to a 2D point p' in the camera optical coordinate frame. The projected point can then be represented by its pixel coordinates (u', v') in the image frame. Fig. 6 illustrates these transformations.

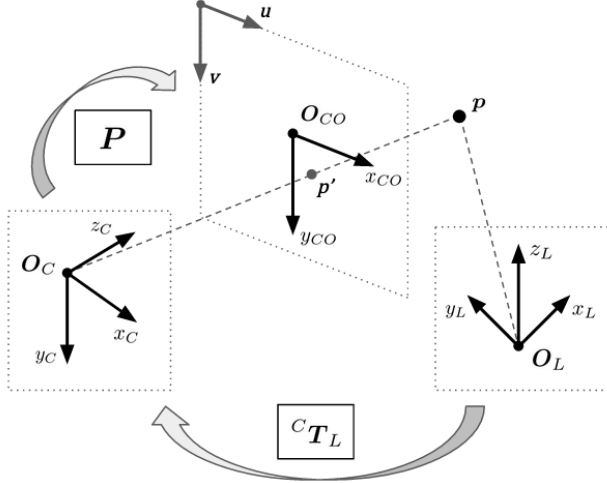


Figure 6. Using the rigid transformation ${}^C T_L$ and the projective transformation P , a 3D point p in the LiDAR frame can be transformed to a 2D pixel in the camera image.

4.2. Algorithm

We define $\mathcal{F} = \{(I_1, P_1), (I_2, P_2), \dots, (I_N, P_N)\}$ as a set of N time synchronized image pointcloud pairs. Each image I_i is a grid of pixels i_0, i_1, \dots, i_j measuring intensity. Each pointcloud P_i is a set of K 3D points p_1, p_2, \dots, p_K which can include additional measurement attributes such as intensity. In the first step a depth map D_i is inferred for each image by a monocular depth estimation network to generate a set of depth maps $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$. An initial guess of the extrinsic parameters $\Theta_0 = [\theta_x, \theta_y, \theta_z, t_x, t_y, t_z]^T$ is used to generate the transformation

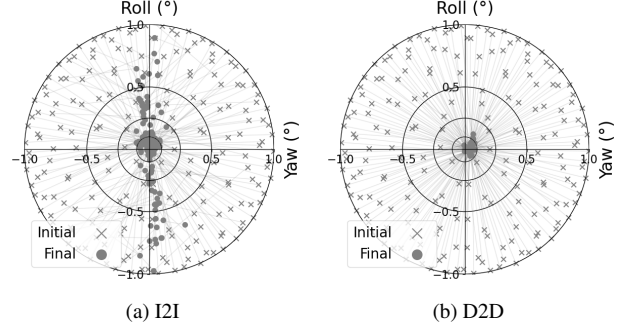


Figure 7. We introduce the Bull's Eye plot. It is designed such that the ground truth extrinsic parameter value is at the center of the plot. A single execution of the algorithm is represented by two connected points, the initial perturbed extrinsic value marked with an "x" and the final optimized value marked with a dot. Plotted here are the 200 executions of the 1° 3 DoF experiment. The improvement in accuracy and precision of D2D compared to I2I is clear.

$${}^C T_L^0 = [R|t] \quad (4)$$

where $R = R_x(\theta_x) \cdot R_y(\theta_y) \cdot R_z(\theta_z) \in \mathbb{R}^{3 \times 3}$ represents a chained rotation around the three canonical axes and $t \in \mathbb{R}^3$ the translation vector. Using this transformation together with the camera projection, each pointcloud P_i is projected into the image frame as explained in Sec. 4.1. Points that are outside the camera's field of view are omitted.

The resulting set $U_i^P = \{(u, v)_1, (u, v)_2, \dots, (u, v)_M\}$ with $M \leq K$ contains the pixel coordinates of all points from P_i that have been successfully projected into the image. From the corresponding 3D points we directly get a set of depth features $f^P = \{d_1^P, d_2^P, \dots, d_M^P\}$ for pointcloud P_i using the euclidean distance $d_i^P = \|p_{L_i}\|_2$ of each point to the LiDAR frame origin. Likewise, we extract a set of matching depth features $f^D = \{d_1^D, d_2^D, \dots, d_M^D\}$ from depth map D_i using the depth values at the pixel coordinates U_i^P . This step is the point in the algorithm where correspondence is directly established between camera and LiDAR features.

To calculate the mutual information between the feature sets, we approximate their marginal and joint distributions $p(f^P)$, $p(f^D)$ and $p(f^P, f^D)$ using normalized histograms (Fig. 4). We express the MI value as a function of the image pointcloud pair (P_i, D_i) and the extrinsic parameters Θ with

$$MI(P_i, D_i, \Theta) = H(f_\Theta^P) + H(f_\Theta^D) - H(f_\Theta^P, f_\Theta^D). \quad (5)$$

The above expression is based on Eq. (1) introduced in Sec. 3.1.

To get an accurate estimate of the MI value including all image pointcloud pairs we average over the whole set \mathcal{F} yielding the final value

$$MI(P_i, D_i, \Theta) = \frac{1}{N} \sum_{i=1}^N MI(P_i, D_i, \Theta). \quad (6)$$

Using Eq. (6) as an objective function, we can apply a suitable iterative algorithm to solve for a set of parameters $\hat{\Theta}$ that maximize the average MI value given a set of image pointcloud pairs. Due to the perspective projection of the 3D points it is not possible to analytically calculate the derivative of the objective function with respect to the parameters Θ . In contrast to Pandey *et al.* [19] we therefore apply Powell’s BOBYQA algorithm [20] which efficiently solves the bounded optimization problem without derivatives using quadratic approximation. Algorithm 1 shows pseudocode for the procedure.

Algorithm 1 D2D MI Extrinsic Calibration

Input: Initial guess Θ_0 , set of N image pointcloud pairs
 $\mathcal{F} = \{(I_1, P_1), (I_2, P_2), \dots, (I_N, P_N)\}$
Output: Optimized parameter $\hat{\Theta}$
 $D_{1\dots N} \leftarrow \text{EXTRACTDEPTHFEATURES}(I_{1\dots N})$
 $\hat{\Theta} \leftarrow \text{BOBYQA}(\text{CALCMI}(\Theta_0, D_{1\dots N}, P_{1\dots N}))$
return $\hat{\Theta}$
procedure CALCMI($\Theta, D_{1\dots N}, P_{1\dots N}$)
 $MI = 0$
 ${}^C T_L = [R|t] \leftarrow \Theta$
for $i \leftarrow 1$ **to** N **do**
 $f^P, f^D \leftarrow \text{GETMATCHES}({}^C T_L, D_{1\dots N}, P_{1\dots N})$
 $MI += H(f^P) + H(f^D) - H(f^P, f^D)$
end for
return $\frac{MI}{N}$
end procedure

5. Experiments

Two experiments were performed using two different camera models in two data sets. The experiments are a 3 Degrees of Freedom (DoF) rotation optimization only experiment and a 6 DoF rotation and translation optimization experiment. The two data sequences used are from KITTI where we use a pinhole camera model and KITTI-360 where we use a double sphere camera model and use the raw unrectified 180° FoV images directly. For both the KITTI and KITTI-360 sequences only 25 frames out of the sequence, sampled uniformly from the entire length of the sequence were used.

In both experiments, the ground truth value of the extrinsic calibration is perturbed by applying a transformation to the LiDAR pointcloud. For each error range (1°, 2°,

Data	KITTI	KITTI-360
Date	09/30/2011	05/28/2013
Drive	0018	0000
Frames	2.762	11,518
Camera	image_02	image_03
Depth	FeatDepth	MiDaS

Table 1. Description of the two sequences used in the experiments.

etc.) 200 points are generated on a unit sphere using the Fibonacci sphere algorithm. The components of each of the 200 unit vectors are multiplied by the selected error range. The same is done for the translation perturbation in the 6 DoF experiment.

In the interest of reproducibility, and because monocular depth estimation itself is not a focus of this work, we use two publicly available networks. For the KITTI sequence we use the self-supervised trained FeatDepth network [25] and for KITTI-360 we use the partially-supervised MiDaS network [22].

A "hit" metric was defined to remove executions that are obviously catastrophic failures from the statistic calculations. A optimization is considered to have "hit" (been successful) if the magnitude of the optimized rotation parameters is less than 0.5° and of the translation parameters less than 20cm.

5.1. Rotation Only Experiment

In this experiment we model the error as a rotation of the LiDAR, however camera rotation can also be considered. This experiment reflects a real world use case where there is strong prior on the translation component of the extrinsic calibration. Such as in autonomous vehicles or robotic systems in which the translation is known from design drawings or a pre-existing factory calibration.

The results of this experiment are shown in Table 2. The superior hit rate of D2D is clear. At a 10° error, I2I hits none of the KITTI executions and only 0.5% of the KITTI-360 executions, D2D however hits 96.5% and 100% of executions respectively. Even up to an error of 20° D2D hits 50.5% and 89% of KITTI and KITTI-360 executions. The increased hit rate of KITTI-360 is due to the fact the LiDAR pointcloud never goes of out the FoV because of the wide FoV fisheye lens.

The precision of D2D is superior to I2I. At all four error levels, for both KITTI and KITTI-360, the mean value of each respective rotation parameter varies by less than 0.02°. There is no comparable consistency for the I2I optimization.

The accuracy of I2I is marginally better, when it hits, for the KITTI sequence but the same for the KITTI-360 sequences. For example, for the 1° KITTI experiment the magnitude of the I2I and D2D error is 0.16° and 0.29° re-

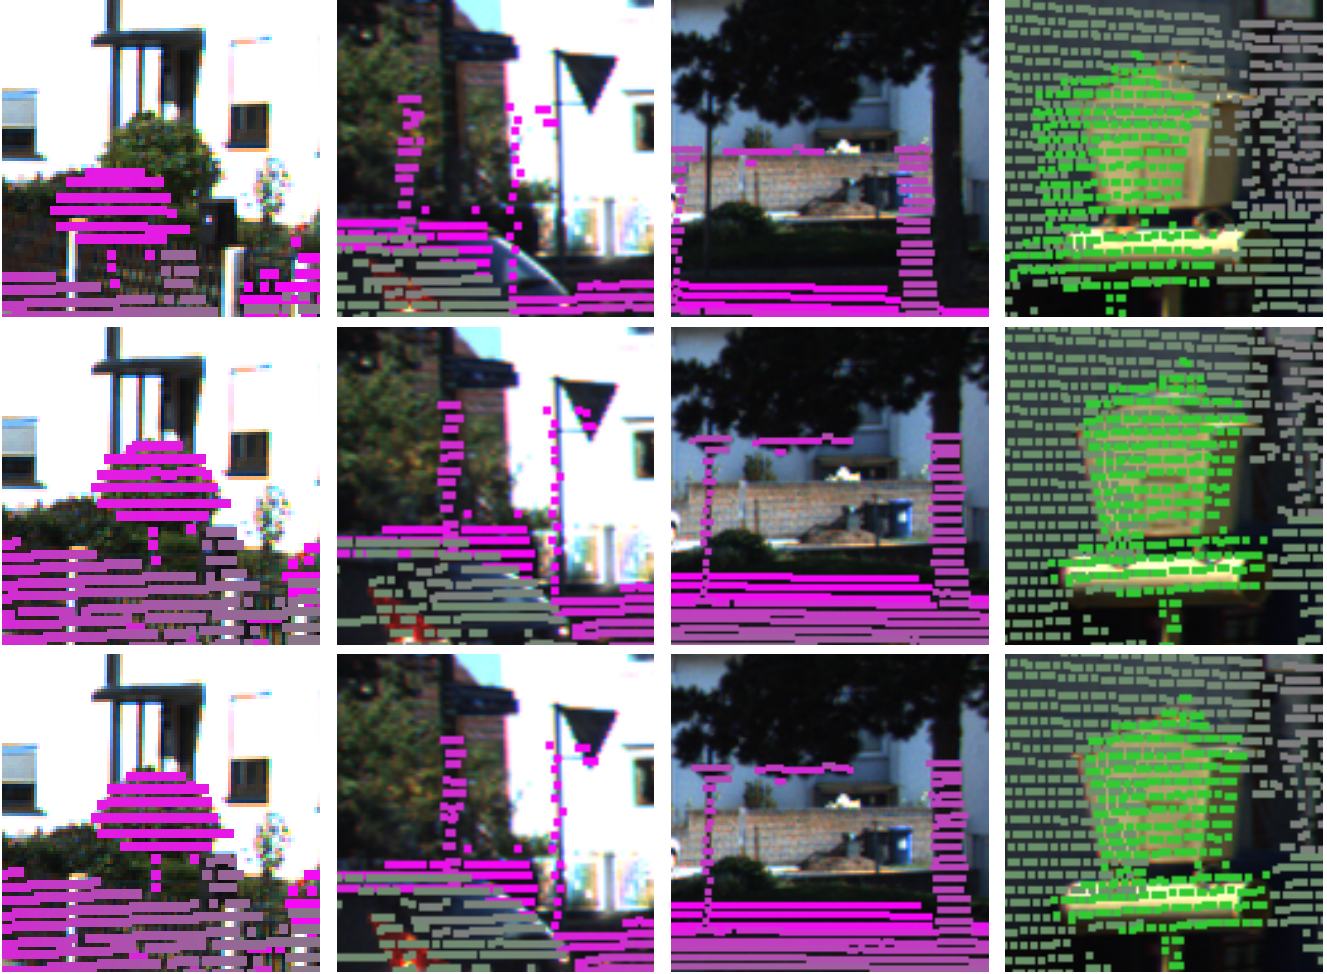


Figure 8. Qualitative results: projection before optimization (top), projection after proposed D2D optimization (middle), and projection with ground truth calibration (bottom).

spectively. For the KITTI-360 sequence the magnitude of the I2I and D2D error is 0.38° and 0.37° respectively.

We propose using Bull’s Eye plots, shown in Fig. 7 to easily characterize and communicate the accuracy of 3 DoF extrinsic calibration optimization algorithms.

5.2. Rotation and Translation Experiment

In this experiment we model the error as a rotation and translation of the LiDAR. This is the general case where no constraint is made on the motion of the camera or LiDAR. Here only a weak prior on the translation component is available. We model this with two translation error ranges of 25cm and 50cm.

The results of this experiment are shown in Tab. 3. Clearly the 6 DoF case is more challenging than the 3 DoF case. However, the increased hit rate of D2D with respect to I2I is again clear. For example at 0.5° and 50cm error only 9.5% and 22% of I2I executions hit whereas for D2D

88% and 76.5% of executions hit for KITTI and KITTI-360 respectively.

Again the precision of D2D is higher than I2I. The mean values of the D2D translation parameters across the four different error initialization vary by less than a magnitude of 2cm. For example, for the D2D KITTI t_x parameter, the mean value is 9.0, 9.4, 9.0 and 9.4cm across the four error ranges. In contrast to the I2I values of -8.0 , 3.7 , -5.7 and -2.0 cm. This consistency pattern is also seen in the value of the rotation parameters.

The low hit rate of I2I makes it challenging to compare the accuracy of I2I to D2D at any error range. Furthermore, the large deviation in the mean value makes direct comparison of the mean values less meaningful. However, D2D can recover the value of the ground truth calibration to within 0.18° and 9.9cm for KITTI and 0.08° and 5.4cm for KITTI-360.

Error	1°				2°				10°				20°			
Data	KITTI		KITTI-360		KITTI		KITTI-360		KITTI		KITTI-360		KITTI		KITTI-360	
Method	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D
Hit (%)	61	100	46.5	100	12.5	99.5	27.5	100	0	96.5	0.5	100	0	50.5	0	89
θ_R (°)	-0.07 ± 0.19	0.26 ± 0.05	-0.35 ± 0.04	0.30 ± 0.02	-0.10 ± 0.11	0.26 ± 0.05	-0.35 ± 0.05	0.30 ± 0.02	-	0.27 ± 0.05	-0.34 ± -	0.30 ± 0.02	-	0.27 ± 0.05	-	0.30 ± 0.02
θ_P (°)	-0.13 ± 0.08	-0.12 ± .07	0.11 ± 0.10	-0.03 ± 0.02	-0.10 ± 0.04	-0.11 ± 0.07	0.07 ± 0.10	-0.03 ± 0.02	-	-0.11 ± 0.08	0.08 ± -	-0.03 ± 0.02	-	-0.13 ± 0.06	-	-0.03 ± 0.02
θ_Y (°)	0.07 ± 0.07	0.04 ± 0.02	-0.08 ± 0.08	0.21 ± .04	0.06 ± 0.04	0.04 ± 0.02	-0.09 ± 0.08	0.21 ± .03	-	0.04 ± 0.02	-0.11 ± -	0.21 ± .04	-	0.03 ± 0.02	-	0.21 ± .03

Table 2. Rotation only extrinsic calibration optimization results using I2I and D2D mutual information at four different error levels. Our proposed D2D method can converge to the true value of the extrinsic calibration from over a 20° error. D2D optimization has significantly more "hits" and is more precise than I2I optimization. Statistics are calculated only for optimizations that hit.

Error	0.5°, 25cm				1°, 25cm				0.5°, 50cm				1°, 50cm			
Data	KITTI		KITTI-360		KITTI		KITTI-360		KITTI		KITTI-360		KITTI		KITTI-360	
Method	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D	I2I	D2D
Hit (%)	4	84.5	26	71.5	0.5	51.5	9	49	9.5	88	22	76.5	2	40.5	3	40.5
θ_R (°)	-0.05 ± 0.14	0.15 ± 0.09	-0.11 ± 0.18	-0.06 ± 0.25	-0.06 ± -	0.18 ± 0.11	-0.04 ± 0.20	-0.02 ± 0.23	-0.08 ± 0.16	0.15 ± 0.10	-0.10 ± 0.21	0.07 ± 0.23	-0.04 ± 0.09	0.15 ± 0.11	0.12 ± 0.23	0.07 ± 0.26
θ_P (°)	-0.25 ± 0.09	-0.04 ± 0.21	0.08 ± 0.14	0.00 ± 0.07	-0.44 ± -	-0.03 ± 0.23	0.04 ± 0.21	0.03 ± 0.08	-0.13 ± 0.12	-0.05 ± 0.22	0.07 ± 0.14	0.00 ± 0.06	0.20 ± 0.15	-0.03 ± 0.23	-0.21 ± 0.11	0.02 ± 0.07
θ_Y (°)	0.13 ± 0.18	-0.06 ± 0.18	-0.13 ± 0.18	0.14 ± 0.18	0.1 ± -	-0.03 ± 0.20	-0.03 ± 0.19	0.03 ± 0.23	0.8 ± 0.20	-0.06 ± 0.20	-0.21 ± 0.16	0.12 ± 0.21	-0.11 ± 0.14	-0.10 ± 0.22	-0.14 ± 0.20	0.02 ± 0.22
t_x (cm)	-8.0 ± 4.6	9.0 ± 4.5	0.0 ± 3.6	-0.1 ± 2.8	3.7 ± -	9.4 ± 4.4	0.3 ± 3.8	1.5 ± 3.5	-5.7 ± 5.7	9.0 ± 4.9	1.1 ± 3.0	0.1 ± 3.0	-2.0 ± 1.8	9.4 ± 4.8	2.2 ± 2.7	1.6 ± 3.3
t_y (cm)	-0.3 ± 6.6	2.1 ± 4.2	3.0 ± 3.7	5.7 ± 1.5	3.0 ± -	1.3 ± 4.6	1.4 ± 2.4	5.1 ± 1.7	1.4 ± 6.8	2.1 ± 4.5	1.8 ± 2.0	5.3 ± 1.6	4.6 ± 3.3	2.9 ± 5.1	0.0 ± 3.1	4.9 ± 1.8
y_z (cm)	-3.7 ± 2.2	0.2 ± 4.7	1.4 ± 1.3	-2.4 ± 2.0	-8.2 ± -	0.6 ± 5.0	1.5 ± 1.0	-2.2 ± 1.9	-1.6 ± 2.6	0.0 ± 4.8	1.5 ± 1.5	-1.4 ± 2.0	5.5 ± 2.9	0.6 ± 5.0	2.0 ± 1.0	-1.5 ± 2.2

Table 3. Rotation and translation extrinsic calibration optimization results using I2I and D2D mutual information at four different error levels. D2D optimization has significantly more "hits" and is more precise than I2I optimization. However the optimization is noticeably more sensitive to rotation error than in the rotation optimization only case. Statistics are calculated only for optimizations that hit.

6. Conclusion

We have presented a robust and explainable camera-LiDAR extrinsic calibration algorithm using geometric mutual information. We introduced three contributions: (i) the use of monocular depth estimation networks as a virtual sensor for camera-LiDAR extrinsics calibration, (ii) the use of geometric mutual information for camera LiDAR extrinsic calibration, (iii) and the first extrinsic calibration experiments on the KITTI-360 fisheye data set. Our proposed algorithm requires no ground truth data, hand engineered features or offline optimization and is suitable for the continuous online extrinsic calibration (COEC) of camera-LiDAR sensors found in automated vehicles and robotics systems.

6.1. Future Work

Further work in this area should include:

- Experimental results exploring the temporal consistency and data dependency
- Understanding the impact of occlusion and large baseline camera-LiDAR systems
- Exploration of optimization strategies for the 6 DoF case

- Research on the effect of degraded monocular depth estimation quality
- Geometric camera-LiDAR mutual information for continuous online intrinsic camera calibration (COIC)

References

- [1] Nicolas Andreff, Radu Horaud, and Bernard Espiau. Robot hand-eye calibration using structure-from-motion. *The International Journal of Robotics Research*, 20:228 – 248, 2001. [2](#)
- [2] Stanley M. Bileschi. Fully automatic calibration of lidar and video streams from a vehicle. *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1457–1464, 2009. [2](#)
- [3] Hsiang-Jen Chien, Reinhard Klette, N. Schneider, and Uwe Franke. Visual odometry driven online calibration for monocular lidar-camera systems. *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2848–2853, 2016. [2](#)
- [4] Thomas M. Cover and Joy A. Thomas. Elements of information theory. 1991. [3](#)
- [5] Yaodong Cui, Ren-Hao Chen, Wenbo Chu, Long Chen, Daxin Tian, Ying Li, and Dongpu Cao. Deep learning for image and point cloud fusion in autonomous driving: A review. *IEEE Transactions on Intelligent Transportation Systems*, 23:722–739, 2020. [1](#)

- [6] Andreas Geiger, Frank Moosmann, Omer Car, and Bernhard Schuster. Automatic camera and range sensor calibration using a single shot. *2012 IEEE International Conference on Robotics and Automation*, pages 3936–3943, 2012. 2
- [7] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Digging into self-supervised monocular depth estimation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3827–3837, 2018. 3, 4
- [8] Ryoichi Ishikawa, Takeshi Oishi, and Katsushi Ikeuchi. Lidar and camera calibration using motions estimated by sensor fusion odometry. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7342–7349, 2018. 2
- [9] Ganesh Iyer, R. Karnik Ram, Krishna Murthy Jatavallabhula, and K. Madhava Krishna. Calibnet: Geometrically supervised extrinsic calibration using 3d spatial transformer networks. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1110–1117, 2018. 2
- [10] Yurim Jeon and Seung-Woo Seo. Efgnet: A versatile image-to-point cloud registration network for extreme outdoor environment. *IEEE Robotics and Automation Letters*, PP:1–8, 2022. 2
- [11] Jesse Levinson and Sebastian Thrun. Automatic online calibration of cameras and lasers. In *Robotics: Science and Systems*, 2013. 2
- [12] Xudong Lv, Boya Wang, Ziwen Dou, Dong Ye, and Shuo Wang. Lccnet: Lidar and camera self-calibration using cost volume network. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2888–2895, 2021. 2
- [13] Frederik Maes, André M. F. Collignon, Dirk Vandermeulen, Guy Marchal, and Paul Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16:187–198, 1997. 3
- [14] Andrew Mastin, Jeremy Kepner, and John W. Fisher III. Automatic registration of lidar and optical images of urban scenes. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2639–2646, 2009. 3, 4
- [15] Faraz M. Mirzaei, Dimitrios G. Kottas, and Stergios I. Roumeliotis. 3d lidar–camera intrinsic and extrinsic calibration: Identifiability and analytical least-squares-based initialization. *The International Journal of Robotics Research*, 31:452–467, 2012. 2
- [16] Miguel Ángel Muñoz-Bañón, Francisco A. Candelas, and Fernando Torres. Targetless camera-lidar calibration in unstructured environments. *IEEE Access*, 8:143692–143705, 2020. 2
- [17] Gaurav Pandey, James R. McBride, Silvio Savarese, and Ryan M. Eustice. Extrinsic calibration of a 3d laser scanner and an omnidirectional camera. *IFAC Proceedings Volumes*, 43:336–341, 2010. 2
- [18] Gaurav Pandey, James R. McBride, Silvio Savarese, and Ryan M. Eustice. Automatic targetless extrinsic calibration of a 3d lidar and camera by maximizing mutual information. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2012. 3
- [19] Gaurav Pandey, James R. McBride, Silvio Savarese, and Ryan M. Eustice. Automatic extrinsic calibration of vision and lidar by maximizing mutual information. *Journal of Field Robotics*, 32, 2015. 3, 4, 6
- [20] M. Powell. The bobyqa algorithm for bound constrained optimization without derivatives. *Technical Report, Department of Applied Mathematics and Theoretical Physics*, 01 2009. 6
- [21] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:1623–1637, 2019. 3, 4
- [22] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3), 2022. 6
- [23] Davide Scaramuzza, Ahad Harati, and Roland Y. Siegwart. Extrinsic self calibration of a camera and a 3d laser range finder from natural scenes. *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4164–4169, 2007. 2
- [24] N. Schneider, Florian Piewak, Christoph Stiller, and Uwe Franke. Regnet: Multimodal sensor registration using deep neural networks. *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 1803–1810, 2017. 2
- [25] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. Feature-metric loss for self-supervised learning of depth and egomotion. In *ECCV*, 2020. 6
- [26] Yi Sun, Jun Yu Li, Yuru Wang, Xinan Xu, Xiaohui Yang, and Zhenping Sun. Atop: An attention-to-optimization approach for automatic lidar-camera calibration via cross-modal object matching. *IEEE Transactions on Intelligent Vehicles*, 8:696–708, 2023. 2
- [27] Zachary Taylor and Juan I. Nieto. A mutual information approach to automatic calibration of camera and lidar in natural environments. *2012 Australasian Conference on Robotics and Automation (ACRA)*, 2012. 3
- [28] Zachary Taylor and Juan I. Nieto. Motion-based calibration of multimodal sensor arrays. *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4843–4850, 2015. 2
- [29] Zachary Taylor, Juan I. Nieto, and David Johnson. Automatic calibration of multi-modal sensor systems using a gradient orientation measure. *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1293–1300, 2013. 2
- [30] Zachary Taylor, Juan I. Nieto, and David Johnson. Multimodal sensor calibration using a gradient orientation measure. *Journal of Field Robotics*, 32, 2015. 2
- [31] Ranjith Unnikrishnan and Martial Hebert. Fast extrinsic calibration of a laser rangefinder to a camera. 2005. 2
- [32] Igor Vajda. Theory of statistical inference and information. 1989. 3
- [33] Paul A. Viola and William M. Wells. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24:137–154, 1995. 3

- [34] William M. Wells, Paul A. Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical image analysis*, 1 1:35–51, 1996. [3](#)
- [35] Qilong Zhang and Robert Pless. Extrinsic calibration of a camera and laser range finder (improves camera calibration). *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, 3:2301–2306 vol.3, 2004. [2](#)
- [36] Zhikang Zhang, Zifan Yu, Suya You, Raghuvver M. Rao, Sanjeev Agarwal, and Fengbo Ren. Enhanced low-resolution lidar-camera calibration via depth interpolation and supervised contrastive learning. *ArXiv*, abs/2211.03932, 2022. [2](#)
- [37] Ganning Zhao, Jiesi Hu, Suya You, and C. C. Jay Kuo. Calibdnn: multimodal sensor calibration for perception using deep neural networks. In *Defense + Commercial Sensing*, 2021. [2](#)
- [38] David Zuñiga-Noël, José-Raúl Ruiz-Sarmiento, Ruben Gomez-Ojeda, and Javier Gonzalez-Jimenez. Automatic multi-sensor extrinsic calibration for mobile robots. *IEEE Robotics and Automation Letters*, 4:2862–2869, 2019. [2](#)