

CELLO-3D: Estimating the Covariance of ICP in the Real World

David Landry François Pomerleau Philippe Giguère

Abstract—The fusion of Iterative Closest Point (ICP) registrations in existing state estimation frameworks relies on an accurate estimation of their uncertainty. In this paper, we study the estimation of this uncertainty in the form of a covariance. First, we scrutinize the limitations of existing closed-form covariance estimation algorithms over 3D datasets. Then, we set out to estimate the covariance of ICP registrations through a data-driven approach, with over 5 100 000 registrations on 1020 pairs from real 3D point clouds. We assess our solution upon a wide spectrum of environments, ranging from structured to unstructured and indoor to outdoor. The capacity of our algorithm to predict covariances is accurately assessed, as well as the usefulness of these estimations for uncertainty estimation over trajectories. The proposed method estimates covariances better than existing closed-form solutions, and makes predictions that are consistent with observed trajectories.

I. INTRODUCTION

The ICP algorithm [1], [2] is ubiquitous in mobile robotics for the tasks of localization and mapping. It estimates the rigid transformation between the reference frames of two point clouds, by iteratively pairing closest points in both point clouds and minimizing a distance between those pairs. This is equivalent to optimizing an objective function that maps rigid transformations to a scalar optimization score for a pair of point clouds. There is an abundance of ICP variants [3], each of which yields slightly different transformations due to their different objective functions. One notable variation is the choice of error metric between each pair of points, where common choices of metric are point-to-point [1] and point-to-plane [2]. The registration process is subject to a number of sources of uncertainty and error, because of a bad adequation between the objective function and the desired result. Chief among them is the presence of local minima in the objective function. Other causes of uncertainty comprise noise from the range sensor, and underconstrained environments such as featureless hallways [4].

The fusion of ICP measurements in existing state estimation frameworks (e.g. SLAM) relies on an appropriate estimation of the uncertainty of ICP, expressed as a covariance [5]. In this context, ICP is modeled as a function of input point clouds and an initial estimate which yields a registration transformation that is normally distributed. Optimistic covariance estimates can lead to inconsistency and navigation failures, whereas pessimistic ones inhibit efficient state estimation. Figure 1 illustrates the process of estimating the covariance of a registration. The registration process shown takes place in a hallway, and is consequently loosely

The authors are with the Department of Computer Science and Software Engineering, Université Laval, Quebec City, QC, G1V 0A6, Canada (emails davidlandry93@gmail.com, {philippe.giguere, francois.pomerleau}@ift.ulaval.ca).




Fig. 1: A reading point cloud (red) is being registered against a reference point cloud (blue). This paper studies the estimation of covariances such as the green ellipsoid to allow integration of ICP in state estimation toolchains. The result of ICP was sampled, and the black balls indicate the density of registration transformations at a particular location. There are three larger clusters which correspond to local minima in the objective function caused by the regularly spaced pillars. We project the translation part of a covariance in \mathbb{R}^3 for illustration purposes, but in general the covariances of a 6 degrees of freedom phenomenon is studied here.

constrained in one axis. In this figure, optimistic covariance estimation frameworks might miss this underconstrainedness and encompass only the central samples.

This paper focuses on the problem of estimating the covariance of ICP in such real 3D environments. To this effect, we first provide an experimental explanation as to why current covariance estimation algorithms may perform poorly in that context. Furthermore, we present CELLO-3D, a data-driven approach to estimating the uncertainty of 3D ICP that works with any error metric.

II. RELATED WORKS

There are many approaches to estimating the covariance of the ICP algorithm, each of which must balance quality of prediction and computation time. On one end of the spectrum, Monte-Carlo (also called brute force) algorithms such as [6], [7] provide an accurate estimate ICP's covariance. They consist in sampling a large number of ICP registration transformations, and using the covariance of the sampled results as the covariance estimation. If a model of the environment is available, brute-force algorithms can take sensor noise into account by simulating many scans of the environment given some noise model. However, Monte-Carlo algorithms cannot be used online due to their high computational cost. This limits their practicality in mobile robotics applications.

Another important category of covariance estimation algorithms rely on the objective function's Hessian [4], [6], [8]–[11]. These closed-form methods are motivated by the need for a covariance estimation that can be used online. Their underlying assumption is that the objective function




Fig. 2: The objective function of point-to-plane ICP around the ground truth when registering a simple cube. The black dots are sampled registration transformations, with their associated translation projected on the xy plane. The green circle is a 3σ covariance ellipse of the distribution of registration transformations. The white dashed circle is a similar representation, but of Censi’s covariance estimate. *Rightmost frame*: the transitions from one plateau to another are explained by point reassociations. At a scale smaller than the scale of the covariance of the registration transformations, the error landscape is dominated by point reassociations.

$J(\mathbf{T})$ used in ICP can be linearized around the point of convergence. This allows the use of linear regression theory to derive a covariance from the Hessian of J . If J is analytically differentiable, then the Hessian can be computed directly [6], [10]. Otherwise, it can also be approximated numerically by sampling [8], [9]. This approach accounts for errors that are due to the environment structure, but not for errors due to sensor noise. However, modelling the effect of sensor noise on the objective function J proves to be crucial for covariance estimation algorithm. Censi [4] addresses this by using the implicit function theorem. His model of the covariance contains the Hessian of J , but also the effect of the sensor noise on it. It is successfully evaluated on 2D datasets. The equations for the 3D case are derived in [11], but Censi’s algorithm is considered to be largely optimistic in that situation [12]. More elaborate noise models alleviate this difficulty, but only for specific sensors [13].

Closed-form approaches have the shortcoming of not taking point reassociation into account. Therefore, they must assume that 1) ICP converged to a loosely defined region of attraction of the “true” solution [4], and 2) the reassociation of points that occur in that region have a negligible influence on the objective function. Bonnabel *et al.* [14] show that for point-to-point ICP variants, the second assumption is broken. However, they present a proof that Censi’s method is accurate using the point-to-plane ICP variant in a noiseless context. They do so by demonstrating that changes in J due to point reassociations are small enough for the covariance estimation methods to remain valid under certain conditions. In spite of this, Mendes *et al.* [12] indicates that this covariance is still optimistic in a noisy experimental context.

A data-driven alternative to covariance estimation emerged from the Covariance Estimation and Learning through Likelihood Optimization (CELLO) framework [15]. It is a general covariance estimation strategy that projects point clouds in a descriptor space, then estimates the covariance within this space. It uses a machine learning algorithm that first estimates a distance metric between the predictors, and then uses this metric to weigh the learning examples during online inference. This procedure can be done with ground-truth data [15], but also without it [16] exploiting expectation-maximization. In the presence of ground-truth data, expectation-maximization should be avoided to simplify the machine learning process. The general strategy

of CELLO was successfully applied to the estimation of the reliability of visual features for visual-inertial navigation [17], [18]. Peretroukhin *et al.* [18] used the prediction space to generate a noise model for visual landmarks, which in turn was used to predict the ego-motion of the sensor. For an application to 3D ICP, these data-driven approaches are challenging in that extracting relevant features from 3D point clouds is still an open problem. Hand-crafted feature designers must tread carefully between the expressiveness and the generality of the descriptor for this approach to be viable. Consequently, this method was never assessed in a 3D ICP context to the best of our knowledge.

Finally, Liu *et al.* [19] circumvent the explicit design of features by training a deep neural network directly on the sensor data. Their method successfully estimates the covariance of 2D visual odometry. Unfortunately, the use of deep neural networks on 3D point clouds is still an active research area, and further work is required to apply this method to such data.

III. SHORTCOMINGS OF CLOSED-FORM COVARIANCE ESTIMATION ALGORITHM

As discussed earlier, Bonnabel *et al.* [14] point out that closed-form covariance estimation methods are potentially ill-founded if point reassociations occur at a scale that is smaller than that of the covariance to be estimated. Our own analysis on 3D simulated data shows that it is likely that point reassociations happen at a scale this small. For example, Figure 2 shows the objective function $J(\mathbf{T})$ observed when registering a pair of $1 \times 1 \times 1$ m cube shaped point clouds. The points lie on the surface of the cube and have a $\sigma = 0.01$ m noise applied on them on every axis. At a larger scale, this objective function $J(\mathbf{T})$ corresponds to our intuition, with a seemingly-smooth slope towards large global minimum. At a smaller scale, however, it is composed of a large number of “plateaus”, each of them corresponding to one fixed association of points between the reading and the reference. This litters the objective function with local minima to which ICP is sensitive. In turn, a larger covariance of ICP results is observed.

There is a mathematical explanation for the optimism of Censi’s algorithm for the point-to-point variant of ICP in Bonnabel *et al.* [14]. We do not know of such a proof for the point-to-plane case. Furthermore, the results in Bonnabel *et al.* [14] were encouraging about the validity of Censi’s al-




Fig. 3: Trace of covariance computations of ICP against sensor noise. Censi’s covariance estimate increases slowly as the sensor noise model grows. The sampled covariance of ICP increases dramatically with sensor noise, even on simulated datasets. This is attributed to the point reassociation provoked by sensor noise.

gorithm in the point-to-plane case. They show that the change in the objective function provoked by point reassociations is bounded under certain conditions, proving the correctness of Censi’s estimate. On the contrary, our own experiments show that caution is required even in the point-to-plane case. Indeed, Figure 3 shows that sensor noise significantly impacts the empirically sampled covariance of ICP as it grows. On the other hand, Censi’s covariance estimate grows slowly as the estimate of the sensor noise grows. To use Censi’s estimate in that experimental context, we would need to inflate our estimation of the sensor noise to values that are beyond a meaningful range.

Consequently, we argue that a viable covariance estimation algorithm for ICP should take into account the effect of noise. More precisely, it should model both the direct effect of the sensor noise on the objective function of ICP, *and also* the point reassociations that it provokes around ground truth. Monte-Carlo based approaches circumvent those difficulties by incorporating the effect of noise directly. The complexity we observe in the 3D registration process motivates our shift from analytical to data-driven solutions. Our general approach is to implement the CELLO framework for 3D ICP and work towards learning covariance models from training data generated through a sampling process. We aim at getting the best of both worlds: the accuracy of brute-force methods and the rapid inference of machine learning approaches.

IV. 3D COVARIANCE ESTIMATION OF ICP

Casting CELLO onto a 3D registration problem requires a primer on notations. A rigid transformation ${}^a_b\mathbf{T} \in \text{SE}(3)$ allows a point cloud to be expressed ${}^b\mathbf{P} \in \mathbb{R}^{3 \times n}$ in the coordinate system b in a second coordinate system a . Using the Lie algebra, we can express the matrix \mathbf{T} as a vector $\boldsymbol{\xi} \in \mathfrak{se}(3)$ using $\log(\mathbf{T})$ and reverse the process using $\exp(\boldsymbol{\xi})$. The vector $\boldsymbol{\xi}$ is split into a translation $\mathbf{u} \in \mathbb{R}^3$ and an angle-axis rotation $\boldsymbol{\omega} \in \mathbb{R}^3$. This allows us to express the uncertainty on a rigid transformation as a covariance matrix $\mathbf{Y} \in \mathbb{R}^{6 \times 6}$, such that

$$\boldsymbol{\xi} = \begin{bmatrix} \mathbf{u} \\ \boldsymbol{\omega} \end{bmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{bmatrix} \mathbf{Y}_{uu} & \mathbf{Y}_{u\omega} \\ \mathbf{Y}_{\omega u} & \mathbf{Y}_{\omega\omega} \end{bmatrix}. \quad (1)$$

Generally speaking, we need to rely on prior information, for which we use the notation $\widetilde{(\cdot)}$, to produce an estimate $\widehat{(\cdot)}$ of a true quantity (\cdot) . For example, having access to a

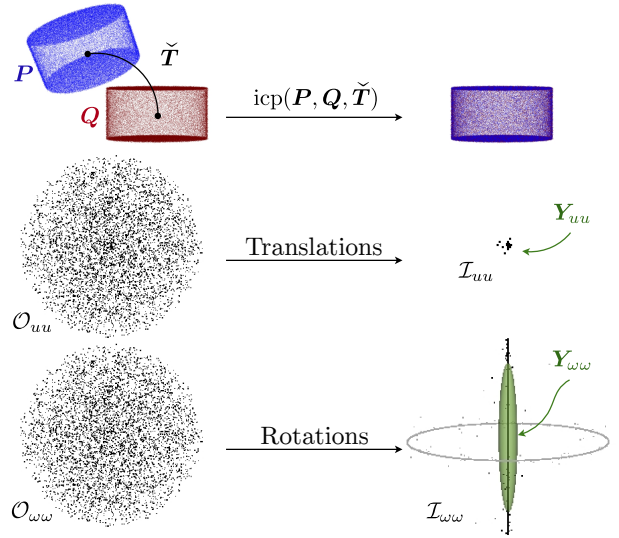


Fig. 4: Overview of different variables used to estimate the covariance of ICP. *Top row*: A single example of two cylinders being registered together using ICP and converging to well aligned point clouds \mathbf{P} (blue) and \mathbf{Q} (red). *Middle row*: A view of the translation components \mathbf{u} of a set of transformations \mathcal{O} before ICP (left) and after ICP (right). *Bottom row*: Same view, but for the rotational components $\boldsymbol{\omega}$. For the rotations, the spread along the vertical axis is explained by the cylinder being unconstrained around one axis. The presence of points in the outer ring is explained by \mathbf{P} first turning upside down and *then* being unconstrained around the same axis.

reference point cloud $\mathbf{Q} \in \mathbb{R}^{3 \times m}$, it is possible to produce a transformation estimate $\widehat{\mathbf{T}}$ that reduces the alignment error between \mathbf{P} and \mathbf{Q} by relying on a prior transformation $\widetilde{\mathbf{T}}$. A typical solution to this registration problem is the ICP algorithm:

$$\widehat{\mathbf{T}} = \text{icp}(\mathbf{P}, \mathbf{Q}, \widetilde{\mathbf{T}}). \quad (2)$$

The initial transformation $\widetilde{\mathbf{T}}$ can be seen as an element randomly selected from a distribution of transformations \mathcal{O} , for which the shape typically depend on odometry computation. Similarly, the estimated value $\widehat{\mathbf{T}}$ comes from a distribution of transformations \mathcal{I} which has a complex shape that depends on both point clouds, \mathcal{O} and the configuration of $\text{icp}(\cdot)$. Estimating the covariance \mathbf{Y} of ICP corresponds to making the (oversimplifying but tractable) assumption that the \mathcal{I} is normally distributed

$$\mathcal{I} \approx \mathcal{N}(\bar{\mathbf{T}}, \mathbf{Y}), \quad (3)$$

where the right hand side is shorthand for

$$\widetilde{\mathbf{T}} = \exp(\boldsymbol{\xi})\bar{\mathbf{T}} \quad \text{with} \quad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{Y}). \quad (4)$$

Note the use of $\bar{\mathbf{T}}$ in Equation 3, which supposes that ICP is unbiased. Figure 4 shows an example of those symbols, with two cylinders being registered starting with a wrongful initial alignment $\widetilde{\mathbf{T}}$. The distribution \mathcal{O} was manually configured and sampled to feed multiple $\widetilde{\mathbf{T}}$ to ICP. The 5000 resulting transformations $\widehat{\mathbf{T}}$ form a distribution with covariance \mathbf{Y} that approximates \mathcal{I} . The covariance approximates the translations well, but some points in rotations were considered divergent and filtered away.

A. Covariance prediction

Vega-Brown *et al.* [15] propose a data-driven framework for covariance estimation, in which the estimation function $\widehat{\mathbf{F}}$

is posed as a weighted average of training examples. The first step in such an approach is to collect a training dataset $\mathcal{D} = \{(\mathbf{d}_0, \mathbf{Y}_0), \dots, (\mathbf{d}_n, \mathbf{Y}_n)\}$ composed of point cloud descriptors \mathbf{d}_k and sampled covariances \mathbf{Y}_k . Our formulation differs from the original CELLO framework, which has error vectors in the place of the covariance matrices \mathbf{Y}_k . This was rendered necessary due to the limited availability of 3D point cloud pairs with associated ground truth: it allows us to extract more knowledge from existing point cloud pairs.

The descriptors \mathbf{d}_k are computed by a function $\mathbf{g}({}^a\mathbf{P}, {}^b\mathbf{T}^b\mathbf{Q})$ that extracts features from a registered point cloud pair. One can then predict a covariance $\hat{\mathbf{F}}(\mathbf{d})$ for an unseen example \mathbf{d} using

$$\hat{\mathbf{F}}(\mathbf{d}) = \frac{1}{\sum_k s(\rho(\mathbf{d}, \mathbf{d}_k))} \sum_k s(\rho(\mathbf{d}, \mathbf{d}_k)) \mathbf{Y}_k. \quad (5)$$

The function ρ is a distance between a pair of point cloud descriptors which is defined as

$$\rho(\mathbf{d}, \mathbf{d}') = (\mathbf{d} - \mathbf{d}')^\top \Theta^\top \Theta (\mathbf{d} - \mathbf{d}') \quad (6)$$

where Θ is an upper triangular matrix. The weighing function $s(x) = e^{-x}$ is chosen here, but any decreasing positive function is appropriate [15]. Large datasets could motivate a choice of s that completely ignores examples with large distances to make runtime predictions more efficient [18]. For the training of Θ , the loss for an individual covariance prediction $\hat{\mathbf{F}}(\mathbf{d}_k)$ is

$$\mathcal{L}(\hat{\mathbf{F}}(\mathbf{d}_k) | \Theta) = \det(\hat{\mathbf{F}}(\mathbf{d}_k)) + \text{tr}(\hat{\mathbf{F}}(\mathbf{d}_k)^{-1} \mathbf{Y}_k), \quad (7)$$

along with a regularization term [15].

B. 3D point cloud descriptor

It is important for a point cloud descriptor \mathbf{d}_k to contain relevant information to the prediction of the covariance for \mathbf{P} , \mathbf{Q} , while being small enough to be amenable to machine learning algorithms. Thus, the descriptor extraction function $\mathbf{g}(\cdot)$ should capture the geometry of the scene, in a way that translates our assumption that this geometry is an important factor of the covariance. One consequence of this is that the extracted descriptors should *not* be rotation invariant, so as to capture the correct orientation of the covariance. For instance, the covariance of ICP in a featureless hallway should be aligned with its walls, as depicted in Figure 1. There is a wide variety of descriptors that could be used for 3D point clouds [20], but a more thorough evaluation of the existing descriptors is a question that is left for future work.

ICP—at least in its robust version—acts mainly on the overlapping region of point clouds. Consequently, descriptors should only be extracted from this overlapping region. In our approach, \mathbf{g} first extracts a single point cloud \mathcal{S} containing the subset of points from both \mathbf{P} and \mathbf{Q} that are overlapping, after registration. Then, descriptors are extracted from \mathcal{S} . The extraction pipeline separates the space into a fixed-size grid, from which all local descriptors are extracted. First, two descriptors capture the overall “planarity” p and “cylindricity” c of the entire voxel, from the average values of these two

TABLE I: ICP pipeline used for the sampled covariance computation

Pipeline Element	Configuration
Point cloud filters	Maximum density, random subsampling
Point matcher	k -d tree, 3 nearest neighbors
Outlier filter	Trimmed distance (keep the closest 70% of associations)
Error minimizer	point-to-plane
Transformation checkers	Max. 80 iterations

metrics defined in [21] computed at each point within the voxel. Then, the orientation of the estimated surface normals vectors for every point in a voxel are summarized in a 9-histogram $\{h_1, \dots, h_9\}$ [22]. The local descriptor of voxel i , j , k is $\mathbf{v}_{ijk} = \{p, c, h_1, \dots, h_9\}$. The full descriptor \mathbf{d} for the overlapping point cloud \mathcal{S} is the concatenation of the local descriptors for every voxel. This way, our descriptor preserves a certain amount of global information, namely the spatial distribution of the local features.

C. Covariance sampling

We employ a brute-force approach similar to [7] to estimate the covariance of ICP for training. We sample the result of ICP for every point cloud pair in our dataset. Every sample uses an initial estimate $\bar{\mathbf{T}}$ drawn from \mathcal{O} . The distribution of results is then used to compute a sampled covariance \mathbf{Y}_k for the point cloud pairs through

$$\mathbf{Y}_k = \frac{1}{n-1} \sum_i \xi_i \xi_i^\top \quad (8)$$

with $\xi_i = \log(\bar{\mathbf{T}}_k^{-1} \hat{\mathbf{T}}_i)$ the n perturbations of the sampled transformations. In some cases, like in the lower row of Figure 4, ICP converges to many clusters. If we suppose that we estimate the covariance of ICP *when it converges*, it is ultimately up to the designer to decide what converged from what did not. We use the DBSCAN clustering algorithm [23] on the set of ξ_i , and keep only the points in the cluster which is closest to ground truth. This filtering method has the benefit of avoiding unrepresentative sampled covariances \mathbf{Y} in the training dataset. It does so *without* imposing an upper bound to the covariance of the samples. One more look at Figure 4 illustrates the results of this procedure. The central line of samples is explained by the fact that rotations around the z axis are not constrained on this cylinder. The outer ring corresponds to situations where ICP converged upside-down, and then spun freely around the z axis. Our filtering strategy is able to remove results that converged incorrectly (i.e. the outer ring) while capturing the information about an underconstrained axis (i.e. the central line).

V. EXPERIMENTS

A. Training datasets

To be realistic, we used datasets that are representative of a wide variety of environments that a mobile robot can encounter. Consequently, we used a subset of the *Challenging data sets for point cloud registration algorithms* [24]. It comprises point clouds taken in environments ranging from structured to unstructured, and indoor to outdoor. Every *Challenging* dataset contains a sequence of l point clouds \mathbf{P}_i as well as ground truth positions $\bar{\mathbf{T}}_i$ for them. To generate a

learning dataset $\mathcal{D} = \{(d_0, \mathbf{Y}_0), \dots, (d_n, \mathbf{Y}_n)\}$, we considered pairs of point clouds P_i, P_j for all i, j such that $i < j < l$ and $j - i \leq 4$.

The descriptors d_k were generated using $g(\cdot)$. We used a grid of $4 \times 4 \times 4$ spanning 25 m in the x and y axes (parallel to the ground plane) and 10 m in the z axis (perpendicular to the ground plane) [20]. This grid was chosen because it encapsulates the typical spatial extent of a point cloud from the *Challenging* datasets.

Each sampled covariance \mathbf{Y}_k was computed from 5000 registrations. Every registration had a \mathcal{O} that was centered at the ground truth and a covariance of $a\mathbf{I}$. We set $a = 0.05$ to simulate a reasonable odometry scenario. A typical ICP registration pipeline was used, featuring the point-to-plane error metric. Table I describes the full registration pipeline, in terms of the framework layed down in [25]. As discussed in Section IV-C, the outliers were filtered from the registration transformations to enforce the assumption that ICP converged. Point cloud pairs where ICP failed consistently were removed from the training dataset. To do so we identified the pairs where ICP converged on average more than 1m or 1rad away from ground truth. In total, this work uses the data from about 5 100 000 registrations on 1020 pairs. These registrations were performed on Compute Canada’s computing clusters using a total of approximately 5 CPU-Years.

Data augmentation was performed to extract more knowledge from the computationally expensive sampling of the \mathbf{Y}_k . By applying a transformation \mathbf{T} on the point clouds before sending them in $g(\cdot)$, we obtained new descriptors d_{aug} . The sampled covariances were transformed similarly using the adjoint representation of \mathbf{T} such that $\mathbf{Y}_{\text{aug}} = \text{Adj}_{\mathbf{T}} \cdot \mathbf{Y}_k \cdot \text{Adj}_{\mathbf{T}}^T$. For our application we chose to perform data augmentation only by rotating the frames of reference around the z axis of the reference point cloud. The z axis was chosen to preserve the 2.5D aspect of the dataset, while giving some rotation invariance to our algorithm. Other rotations were expected to create examples that are not close to the registration pairs encountered online, such as examples where the trees of a forest are sideways. A 6-DOF robot would warrant a more complete data augmentation.

B. ICP trajectories computation

Once the models were trained, we evaluated our covariance estimation algorithm for state estimation on indoor and outdoor trajectories. We thus computed an ICP odometry from our trajectory datasets. Since the point cloud pairs are in a sequence of length l , we compounded their ICP registration transformations using $\hat{\mathbf{T}}_F = {}^0\hat{\mathbf{T}}_1 {}^1\hat{\mathbf{T}}_2 \dots {}^{l-1}\hat{\mathbf{T}}_l$ where $\hat{\mathbf{T}}_F$ is the final pose estimate of the odometry. The distribution of initial transformations was $\mathcal{O} \approx \mathcal{N}({}_{i+1}^i\hat{\mathbf{T}}, a\mathbf{I})$ with $a = 0.05$. In a second step, we computed covariance estimation $\hat{\mathbf{Y}}_k$ for each $\hat{\mathbf{T}}_k$ using a covariance estimation model $\hat{\mathbf{F}}$. Finally, we compounded the covariances with the 4th order approximation in Barfoot *et al.* [26] to obtain a final covariance estimate $\hat{\mathbf{Y}}_F$. This setup allowed us to assess the quality of the covariance estimation in context, over trajectories.

VI. RESULTS

We trained on one trajectory dataset, while testing on one or many others that had the same type of environment/structure. Testing on separate (but similar) datasets was done to obtain a fair evaluation, while avoiding an overly optimistic result due to overfitting. This correspondence is detailed in Table II. The weighting matrix Θ was trained by stochastic gradient descent using the loss from Equation 7. The batch size was 4 and the learning was $1e-5$.

A. Single-Pair Covariance Prediction

First, we validated the quality of our approach, for single pairs of point clouds. As a quality metric, we used the Kullback-Leibler (KL) divergence from our estimated distribution to the sampled distribution (covariance). This metric measures the amount of information lost when using the covariance $\hat{\mathbf{F}}(d_k)$ instead of \mathbf{Y}_k to express the distribution of ICP results \mathcal{I} . Results reported in Table II are the average KL divergence over all pairs within a testing group. For comparison, we also computed the average KL divergence from a baseline covariance $\mathbf{Y}_{\text{base}} = \frac{1}{n} \sum_k \mathbf{Y}_k$, using the \mathbf{Y}_k of the training dataset. We made similar computations with a Censi estimate $\mathbf{Y}_{\text{Censi}}$, for the same point cloud pairs. One should keep in mind that the mere prediction of the scale of the ICP covariance has been historically challenging. As hinted by Figure 3, Censi’s covariance estimates have large divergences, since they are orders of magnitude smaller than the sampled covariances. At worst, our approach makes predictions that have the correct order of magnitude, as seen in Table II.

Trajectory in self-similar environments will have point clouds (and covariances) that are similar to one another. For these self-similar environments, gains over the baseline are expected to be modest. In this situation, we observe that our predictor’s parameters Θ converges to nearly uniform weights after training. Consequently, CELLO-3D treats training examples nearly equally, and outputs (something close to) their mean. This phenomenon is clearly visible in Table II, where gains over the baseline are limited for *Wood Autumn* and *Wood Summer*. For the *Gazebo* environments, they also exhibit a certain degree of self-similarity, as measured by the low KL-divergence of the baseline. Again, gains for our approach are modest there. However, for the three indoor environments (apt, haupt, and stairs), we can see that they have the highest baseline KL-divergence. This indicates that the sampled covariance varies largely throughout the trajectory, in comparison to the average (baseline) one. As expected, this is where our algorithm makes the strongest gains. These gains for indoor environments are also explained by their structured nature, well-suited for our descriptors.

B. Consistency over Trajectories

We evaluated the consistency of the estimated covariances when computing ICP odometry trajectories in diverse locations. This represents the fundamental situation where our predictor is used within a state-estimation algorithm. To visualize the error on the covariance in all dimensions at

TABLE II: Loss of the CELLO algorithm in various training scenarios.

Dataset	Trained on	N. Pairs	Avg. KL Divergence		
			Baseline	Ours	Censi
Apartment	Haupt. & Stairs	1190	34.1	26.6	9.19e7
Hauptgebaude	Apt & Stairs	938	34.0	26.7	2.06e8
Stairs	Apt & Haupt.	798	33.7	27.0	7.65e7
Gazebo Summer	Gzb. Winter	826	20.8	19.8	2.55e6
Gazebo Winter	Gzb. Summer	798	20.3	18.9	2.25e6
Wood Autumn	Wd Summer	812	13.2	11.5	4.94e6
Wood Summer	Wd Autumn	966	13.6	11.3	3.52e7

TABLE III: Final odometry error and consistency of CELLO-3D.

Traj.	Length (m)	Translation		Rotation		D_M
		$\ u\ $ (m)	D_M	$\ \omega\ $ (rad)	D_M	
Apartment	22	0.115	0.274	0.0331	0.160	0.540
Haupt.	24	0.168	0.467	0.00910	0.346	1.15
Stairs	12	0.0664	0.0998	0.0127	0.307	0.592
Gzb. Smmr	14	0.0396	0.278	0.0165	0.278	0.491
Gzb. Wntr	15	0.0311	2.000	0.0144	2.90	2.50
Wd Atmm	18	0.217	0.205	0.0178	0.394	0.405
Wd Smmr	21	0.332	0.208	0.0299	0.533	0.762

once, we resort to computing the Mahalanobis distance D_M between \hat{T}_F and the ground truth 0T

$$D_M = \sqrt{\xi^\top \hat{Y}_F^{-1} \xi}, \quad (9)$$

in which $\xi = \log({}^0T^{-1} \hat{T}_F)$. The distance D_M can be thought of as the number of standard deviations between a sample and the mean, with a value zero meaning that \hat{T}_F was exactly on the ground truth. Table III shows the average D_M of 100 trajectories for every covariance model. It also lists the average Mahalanobis distances of the translation u against Y_{uu} and rotation ω against $Y_{\omega\omega}$ for the trajectories. The average values of D_M indicate that our algorithm is consistent overall. We consider that an average D_M above 3 indicates an optimistic covariance estimate, while an average below 1.5 indicates a pessimistic estimate. In that sense, the covariance estimation algorithm is pessimistic for the *Stairs* or *Apartment* datasets. However, not extreme values were found demonstrating a functional solution.

In Figure 5, we take a closer look at the behaviour of CELLO-3D over short trajectories, in the *Wood Summer* (unstructured) and *Gazebo Winter* (semi-structured) environments. The figure compares the compounded sampled covariances Y_k with our estimated covariances $\hat{F}(d_k)$. In there, we sampled 20 ICP odometry trajectories to compare against the covariance predictions at each step (the grey ellipses). The green dots represent the final \hat{T}_F for each trajectory, with the green ellipses being \hat{Y}_F . This figure shows that our covariance estimates are consistent with the compounded trajectories, within 2 sigma. Note that the CELLO-3D uncertainty estimate grows more steadily and uniformly than that of the sampled covariances. This indicates that, while our algorithm does not seem to model the input descriptors in very sharp detail, it is able to extract a consistent knowledge from the training dataset.




Fig. 5: Comparison of covariance estimations for *Wood Summer* (top) and *Gazebo Winter* (bottom). Every frame as 20 ICP odometry trajectories, although some are not visible because they are overlapping. The ground truth trajectory is shown as a thick black line. The green dots represent the final poses \hat{T}_F , while the green ellipses represent \hat{Y}_F . The data is projected on the ground plane for the sake of visualization.

VII. CONCLUSION

In this work, we presented CELLO-3D, an online covariance estimation algorithm for ICP that works well in 3D. CELLO-3D uses the covariances of a learning dataset to predict the covariance of similar point cloud pairs at runtime. It was successfully validated on individual pairs of point clouds and over trajectories, on challenging datasets. It provides also a better uncertainty estimate when compared to existing solutions. Our predicted covariances are neither too optimistic nor too pessimistic, and represent well sampled particles over trajectories of several meters.

Throughout the course of this work, some challenges became visible with the transition from 2D to 3D. Due to the curse of dimensionality, generating a dataset for covariance in 3D requires more samples than in 2D. With this larger number of samples, we noticed that approximating \mathcal{I} as a normal distribution is prone to larger estimation errors in $SE(3)$, as we observe mainly multimodal distributions (see Figure 1 and Figure 4). Moreover, the original CELLO framework proposed the use of weak descriptors to alleviate the difficulties of descriptor design. However, the quality of the input features is found to be critical to the success of covariance predictions. Inspired by [19], we intend to meet this challenge using Deep Neural Networks (DNNs) directly on 3D point clouds.

ACKNOWLEDGMENT

This work was partially funded by the *Fonds de Recherche du Québec - Nature et technologies* (FRQNT) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

REFERENCES

- [1] P. Besl and N. McKay, "A Method for Registration of 3-D Shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [2] C. Yang and G. Medioni, "Object modelling by registration of multiple range images," in *Image and Vision Computing*, vol. 10, 1991, pp. 145–155.
- [3] F. Pomerleau, F. Colas, and R. Siegwart, "A Review of Point Cloud Registration Algorithms for Mobile Robotics," *Foundations and Trends in Robotics*, vol. 4, no. 1, pp. 1–104, 2015.
- [4] A. Censi, "An accurate closed-form estimate of ICP's covariance," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 3167–3172.
- [5] Z. Zhang, "Iterative Point Matching for Registration of Free-Form Curves and Surfaces," *International Journal of Computer Vision*, vol. 13, no. 2, pp. 119–152, 1994.
- [6] O. Bengtsson and A.-J. Baerveldt, "Robot localization based on scan-matching—estimating the covariance matrix for the IDC algorithm," *Robotics and Autonomous Systems*, vol. 44, no. 1, pp. 29–40, 2003.
- [7] T. M. Iversen, A. G. Buch, and D. Kraft, "Prediction of ICP pose uncertainties using Monte Carlo simulation with synthetic depth images," in *IEEE International Conference on Intelligent Robots and Systems*, 2017, pp. 4640–4647.
- [8] P. Biber and W. Strasser, "The normal distributions transform: a new approach to laser scan matching," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, vol. 3, 2003, pp. 2743–2748.
- [9] J. Nieto, T. Bailey, and E. Nebot, "Scan-SLAM: Combining EKF-SLAM and scan correlation," *Springer Tracts in Advanced Robotics*, vol. 25, pp. 167–178, 2006.
- [10] M. Bosse and R. Zlot, "Map matching and data association for large-scale two-dimensional laser scan-based SLAM," *International Journal of Robotics Research*, vol. 27, no. 6, pp. 667–691, 2008.
- [11] S. M. Prakhya, Liu Bingbing, Yan Rui, and Weisi Lin, "A closed-form estimate of 3D ICP covariance," in *Proceedings of the 14th IAPR International Conference on Machine Vision Applications, MVA 2015*, IEEE, 2015, pp. 526–529.
- [12] E. Mendes, P. Koch, and S. Lacroix, "ICP-based pose-graph SLAM," in *SSRR 2016 - International Symposium on Safety, Security and Rescue Robotics*, IEEE, 2016, pp. 195–200.
- [13] M. Barczyk and S. Bonnabel, "Towards Realistic Covariance Estimation of ICP-based Kinect V1 Scan Matching : the 1D Case," in *Proceedings of the American Control Conference*, 2017, pp. 4833–4838.
- [14] S. Bonnabel, M. Barczyk, and F. Goulette, "On the covariance of ICP-based scan-matching techniques," January 2015, IEEE, 2016, pp. 5498–5503.
- [15] W. Vega-Brown, A. Bachrach, A. Bry, J. Kelly, and N. Roy, "CELLO: A fast algorithm for Covariance Estimation," in *IEEE International Conference on Robotics and Automation*, May 2013, pp. 3160–3167.
- [16] W. Vega-Brown and N. Roy, "CELLO-EM: Adaptive sensor models without ground truth," in *IEEE International Conference on Intelligent Robots and Systems*, IEEE, Nov. 2013, pp. 1907–1914.
- [17] V. Peretroukhin, L. Clement, M. Giamou, and J. Kelly, "PROBE: Predictive robust estimation for visual-inertial navigation," in *IEEE International Conference on Intelligent Robots and Systems*, 2015, pp. 3668–3675.
- [18] V. Peretroukhin, W. Vega-Brown, N. Roy, and J. Kelly, "PROBE-GK: Predictive robust estimation using generalized kernels," in *IEEE International Conference on Robotics and Automation*, Jun. 2016, pp. 817–824.
- [19] K. Liu, K. Ok, W. Vega-Brown, and N. Roy, "Deep inference for covariance estimation: Learning gaussian noise models for state estimation," in *IEEE International Conference on Robotics and Automation*, May 2018, pp. 1436–1443.
- [20] M. Bosse and R. Zlot, "Place Recognition Using Regional Point Descriptors for 3D Mapping," in *Field and Service Robotics*, Springer Berlin Heidelberg, 2010, pp. 195–204.
- [21] J. Demantke, C. Mallet, N. David, and B. Vallet, "Dimensionality based scale selection in 3d lidar point clouds," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 38, no. Part 5, W12, 2011.
- [22] M. Magnusson, H. Andreasson, A. Nüchter, and A. J. Lilienthal, "Appearance-based loop detection from 3D laser data using the normal distributions transform," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, vol. 3, 2009, pp. 23–28.
- [23] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, "DBSCAN Revisited, Revisited," *ACM Transactions on Database Systems*, vol. 42, no. 3, pp. 1–21, 2017.
- [24] F. Pomerleau, M. Liu, F. Colas, and R. Siegwart, "Challenging data sets for point cloud registration algorithms," *The International Journal of Robotics Research*, vol. 31, no. 14, pp. 1705–1711, 2012.
- [25] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing ICP Variants on Real-World Data Sets," *Autonomous Robots*, vol. 34, no. 3, pp. 133–148, Feb. 2013.
- [26] T. D. Barfoot and P. T. Furgale, "Associating uncertainty with three-dimensional poses for use in estimation problems," *IEEE Transactions on Robotics*, vol. 30, no. 3, pp. 679–693, 2014.