

Sharp Performance Bounds for PASTA

Matteo Marchi, Jonathan Bunton, Yskandar Gas, Bahman Ghahesifard, and Paulo Tabuada

Abstract—LiDAR is a standard sensor choice for self-localization and SLAM on indoor autonomous robots. While there are many methods to estimate a robot’s location using LiDAR measurements, most rely on algorithms that solve a generic LiDAR scan matching problem. When safety is a concern, these algorithms must provide a bound on the localization error to enable safety enforcing controllers, such as those based on Control Barrier Functions. Unfortunately, most existing scan matching algorithms offer no formal guarantees, and are tailored to structured, high-resolution 3D point clouds.

In this paper, we present an improved theoretical analysis for a low-cost alternative to these methods named PASTA (Provably Accurate Simple Transformation Alignment), originally introduced in [8]. We provide a formal worst-case guarantee on the localization error and show, experimentally, that it is tight. This characterization of the localization error simplifies the interface between high-dimensional perception data and safety-critical control algorithms.

I. INTRODUCTION

Light Detection And Ranging (LiDAR) scanners are popular sensors for autonomous robot and vehicle applications [7], [14]. Their accuracy, insensitivity to lighting conditions, low cost, and increasing availability in recent years make LiDAR sensors an indispensable alternative and companion to other sensors, such as camera vision. LiDAR is also widely used in self-localization, mapping, and SLAM algorithms, with a large body of literature devoted to its various properties [3].

In this paper, we consider an indoor localization problem, wherein a small robot must estimate its current position and orientation with respect to a fixed reference frame purely from two LiDAR measurements. This task is equivalent to estimating the rigid transformation (the rotation and translation) that relate the robot’s current pose to the fixed reference frame from the LiDAR measurements. We may also frame localization with LiDAR data as the general point cloud registration problem, where we search for a transformation that aligns two LiDAR data point clouds, for which a variety of algorithms have been developed.

In the ideal setting with known point-to-point correspondences between both LiDAR measurements, there is an analytic solution [6]. This solution, however, cannot be used in most scenarios and faces additional difficulties when working with point clouds from LiDAR measurements. For example, LiDAR measurements are non-uniformly spaced and thus points measured at different poses are not, in general, related by a pure rigid transformation.

Other point cloud registration algorithms we might use for localization fall into one of three categories: local (iterative), global, or learning-based. While we refrain from listing all existing methods in detail (see [9] for an excellent review), we provide some remarks to contextualize our work.

Arguably the most well-known local method for point cloud registration is Iterative Closest Point (ICP) [1]. The algorithm follows a simple loop of choosing correspondences (using nearest neighbors under the estimated transformation) then using them to estimate a new associated transformation. Since its inception, countless variants of ICP have been developed to handle more structured environments [13]. All of these ICP variants, however, rely on the same alternating optimization approach, meaning they converge to an incorrect transformation when initialized poorly. Moreover, when the correspondences (point-wise, planar, or otherwise) fail to exist, such as in point clouds from LiDAR data, ICP-based methods may fail entirely.

In contrast, global algorithms require no initialization to estimate a transformation. These methods are typically based on branch-and-bound techniques to solve the ICP problem [16], or on random-sampling consensus [12]. While global methods can have performance that is superior even to well-initialized local algorithms, they often suffer from prohibitively long runtimes or rely on carefully crafted features persisting between measurements.

Various machine learning algorithms have also been developed to solve the point cloud registration problem. In particular, machine learning methods are well-suited for extracting features from raw point cloud data [10], or even for finding correspondences between features in multiple point clouds [5]. The identified features or correspondences can then be directly handed to any of the aforementioned local or global algorithms [15]. The major downfall of these methods is their need for large amounts of labeled data coming from the deployment environment and the inherent brittleness of learned features.

Despite their practical success, these algorithms *cannot be used in safety-critical applications*, as guaranteed safe control requires the state to be either known exactly, or formal bounds on its uncertainty. [2]. To the authors’ knowledge, only one other point cloud registration algorithm, TEASER++ [15], provides formal guarantees, but relies on the presence of at least a few point-to-point correspondences that may be absent in LiDAR point clouds. While this problem can be circumvented by extracting point-features from the data, most scenarios require significant feature engineering.

The main contribution of this work is a strong theoretical guarantee for PASTA (Provably Accurate Simple Transformation Alignment), an algorithm first introduced in [8].

Given two LiDAR measurements, PASTA constructs sets (i.e., regions of the 2D plane for 2D LiDAR) and compares their first and second moments in an alignment step, the result of which can immediately be used for localization.

Comparing these sets, rather than the point clouds themselves, is a defining feature of PASTA that makes the method robust to the variable point density of LiDAR measurements. Moreover, this approach does not rely on persistent local features, like other fast global methods. We note that PASTA bears minor similarities to principal component analysis (PCA) methods [4], [11], but instead relies on moments—the mean and covariance—of a *set*, rather than those of a point cloud or a grid of pixels.

While PASTA was introduced in [8] with some guarantees, we significantly tighten the bounds on the error of the produced location estimates. This improved theoretical guarantee allows PASTA to seamlessly connect with existing safety-critical control frameworks. After proving these improved guarantees, we validate their tightness with real LiDAR data collected for an indoor localization task. These experiments highlight the theoretical limits of our bound, but also its good empirical performance.

II. PROBLEM STATEMENT

Given a point cloud $\{\mathbf{r}_1^{(i)}\}_{i=1}^{m_1}, \mathbf{r}_1^{(i)} \in \mathbb{R}^n$ from a LiDAR located¹ at the pose $(\mathbf{p}_1, \mathbf{R}_1)$, and a second point cloud $\{\mathbf{r}_2^{(i)}\}_{i=1}^{m_2}$ from a pose $(\mathbf{p}_2, \mathbf{R}_2)$ (both in some global coordinate frame) estimate the relative translation vector $\hat{\mathbf{p}}$ and rotation matrix $\hat{\mathbf{R}}$ that transforms the first pose into the second, i.e., such that $\hat{\mathbf{R}}\mathbf{R}_1 = \mathbf{R}_2$ and $\hat{\mathbf{p}} = \mathbf{R}_1^T(\mathbf{p}_2 - \mathbf{p}_1)$.

III. A BRIEF REVIEW OF PASTA

In [8], the authors introduced a LiDAR point cloud registration algorithm named PASTA (Provably Accurate Simple Transformation Alignment) inspired by classical 2D image alignment techniques, and proved that the algorithm enjoys worst case bounds on the computed transformation error. Our main contribution is a simpler theoretical analysis and a substantially improved error bound. We first review the concepts necessary to understand the theory to follow, and defer to [8] for other details.

A. Moments of a Set

PASTA relies on the notion of *moments* of a compact set $H \subset \mathbb{R}^n$. Given a measure μ on \mathbb{R}^n and the indicator function² on the set H , $\mathbf{1}_H$, the zeroth moment of H (e.g., the area in 2D or volume in 3D) is defined as:

$$|H| = \int \mathbf{1}_H(\mathbf{x}) \, d\mu. \quad (1)$$

Similarly, the first moment (centroid) and second moment (covariance) of H are:

$$\mathbf{c} = \frac{1}{|H|} \int \mathbf{x} \cdot \mathbf{1}_H(\mathbf{x}) \, d\mu \quad (2)$$

$$\Sigma = \frac{1}{|H|} \int (\mathbf{x} - \mathbf{c})(\mathbf{x} - \mathbf{c})^T \mathbf{1}_H(\mathbf{x}) \, d\mu. \quad (3)$$

¹A pose is a pair $(\mathbf{p}, \mathbf{R}) \in \mathbb{R}^n \times \mathcal{SO}(n)$, where \mathbf{p} denotes a translation vector and \mathbf{R} denotes a rotation matrix between a fixed reference frame and a frame fixed to the LiDAR. The location of the LiDAR frame is the origin of the rays used by the sensor to perform the distance measurements.

²The indicator function of the set H , $\mathbf{1}_H : \mathbb{R}^n \rightarrow \{0, 1\}$, is equal to one for all points in H and zero elsewhere.

Algorithm 1 PASTA

Input: Point clouds $\{\mathbf{r}_1^{(i)}\}_{i=1}^{m_1}, \{\mathbf{r}_2^{(i)}\}_{i=1}^{m_2}$

Output: Transformation $\hat{\mathbf{R}}, \hat{\mathbf{p}}$

for each point cloud i **do**

$H_i \leftarrow$ convex hull of $\{\mathbf{r}_i^{(j)}\}_{j=1}^{m_i}$

$\mathbf{c}_i, \Sigma_i \leftarrow$ first and second moments of H_i

end for

$\hat{\mathbf{R}} \leftarrow$ closed-form solution of (4)

$\hat{\mathbf{p}} \leftarrow$ closed-form solution of (5)

PASTA algorithm details. Note that the rotation equation (4) may have multiple solutions, and we must choose one properly. Details on selecting which one to use can be found in [8].

For a convex hull H constructed from a (non-empty) point cloud, the moments are well-defined. Moreover, by partitioning H into a set of disjoint simplices (triangles in \mathbb{R}^2 , tetrahedrons in \mathbb{R}^3), we can compute the moments of H as a weighted sum of the moments of the simplices, using just the coordinates of their vertices (see [8]).

B. PASTA Overview

The main intuition behind PASTA (and other PCA-inspired methods) is as follows: if two compact sets H_1 and H_2 are related by a rigid transformation (\mathbf{p}, \mathbf{R}) , then their first $(\mathbf{c}_1, \mathbf{c}_2)$ and second (Σ_1, Σ_2) moments are related by the following algebraic expressions that we can sequentially solve:

$$\Sigma_2 = \mathbf{R}\Sigma_1\mathbf{R}^T, \quad (4)$$

$$\mathbf{c}_2 = \mathbf{R}\mathbf{c}_1 + \mathbf{p}. \quad (5)$$

This idea does not directly apply to LiDAR data point clouds since, even without occlusions, points and sets are never related by a pure rotation and translation. LiDAR measurements are noisy and are not uniformly distributed over the environment when the sensor changes pose, as seen in Fig. 1, so (4) and (5) are not satisfied by the true rotation and translation. However, we show that if the sets H_1 and H_2 (e.g., the *convex hulls* of both point clouds) are still nearly related by a rotation and translation, PASTA’s estimated $\hat{\mathbf{R}}$ and $\hat{\mathbf{p}}$ are also nearly correct. For clarity, Algorithm 1 describes the operation of PASTA in pseudo-code.

IV. MAIN RESULTS

In this section we present our main theoretical result, a new and tighter bound on the error in PASTA’s estimates $(\hat{\mathbf{p}}, \hat{\mathbf{R}})$. This error is a function of how “closely” a compact set H and a perturbed version of H —here denoted by H' —are related by some true rotation and translation, as measured in variations of their second moments. In the following, we use $\|\cdot\|$ to refer to the ℓ_2 norm on \mathbb{R}^n and its induced norm on $\mathbb{R}^{n \times n}$ and use $\lambda_i(\mathbf{A})$ to denote the i th eigenvalue of a symmetric matrix \mathbf{A} . For compactness, we also define:

$$\begin{aligned} \bar{\Delta}_\lambda(\mathbf{A}) &= \max_{i \neq j} |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})| \\ \underline{\Delta}_\lambda(\mathbf{A}) &= \min_{i \neq j} |\lambda_i(\mathbf{A}) - \lambda_j(\mathbf{A})|. \end{aligned} \quad (6)$$

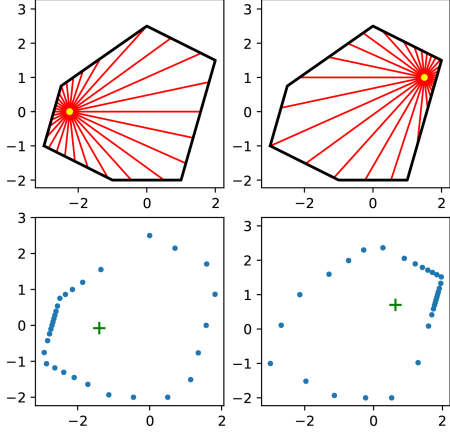


Fig. 1. Top row: LiDAR rays from different positions in a 2D environment. Bottom row: corresponding distance measurements converted into a point cloud. The centroid, or first moment, of the points (green cross) differs despite the shapes being aligned in the same coordinate frame, thus related by an identity rotation and no translation in (5).

A. Eigenvector Perturbation Bounds

The second moment (3) corresponds to a symmetric and positive definite covariance matrix, and assuming distinct eigenvalues, this matrix has a set of orthogonal *eigenvectors*. If we additively perturb any symmetric matrix \mathbf{A} with another symmetric matrix \mathbf{B} , the eigenvectors of \mathbf{A} will be perturbed by some angle. The following lemma bounds the magnitude of this angle as a function of the *norm* of \mathbf{B} , and is a stepping stone for a stronger result to follow.

Theorem 1 (Symmetric Matrix Eigenvector Perturbation): Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be symmetric matrices, where \mathbf{A} has simple eigenvalues. Let λ_i and γ_i , $i = 1, \dots, n$, be the ordered eigenvalues of \mathbf{A} and $\mathbf{A} + \mathbf{B}$ respectively, with corresponding unit eigenvectors \mathbf{r}_i and \mathbf{s}_i . If $\|\mathbf{B}\| < \frac{1}{2} \min_{i \neq j} |\lambda_i - \lambda_j|$, the angle θ_k between \mathbf{r}_k and \mathbf{s}_k satisfies:

$$|\theta_k| \leq \sin^{-1} \left(\frac{2\|\mathbf{B}\|}{\min_{i \neq j} |\lambda_i - \lambda_j|} \right). \quad (7)$$

Proof: \mathbf{A} and \mathbf{B} are symmetric, thus their sum $\mathbf{C} = \mathbf{A} + \mathbf{B}$ is also symmetric and we can write the eigenvalue decompositions $\mathbf{A} = \mathbf{R}\mathbf{\Lambda}\mathbf{R}^T$ and $\mathbf{C} = \mathbf{S}\mathbf{\Gamma}\mathbf{S}^T$. Note that $\mathbf{\Lambda}$ and $\mathbf{\Gamma}$ are diagonal matrices whose i th entries are the i th eigenvalues of \mathbf{A} and \mathbf{C} , respectively, in ascending order. Similarly, \mathbf{R} and \mathbf{S} are orthonormal matrices whose i th columns \mathbf{r}_i and \mathbf{s}_i are the eigenvectors associated to λ_i and γ_i , respectively.

Consider the following for any unit-eigenvector \mathbf{s}_k of \mathbf{C} :

$$\begin{aligned} & \mathbf{s}_k^T (\mathbf{A} - \mathbf{C})^T (\mathbf{A} - \mathbf{C}) \mathbf{s}_k \\ &= \mathbf{s}_k^T (\mathbf{R}\mathbf{A}\mathbf{R}^T - \mathbf{S}\mathbf{\Gamma}\mathbf{S}^T)^T (\mathbf{R}\mathbf{A}\mathbf{R}^T - \mathbf{S}\mathbf{\Gamma}\mathbf{S}^T) \mathbf{s}_k \\ &= \left(\mathbf{s}_k^T \mathbf{R}\mathbf{A}\mathbf{R}^T - \gamma_k \mathbf{s}_k^T \right) \left(\mathbf{R}\mathbf{A}\mathbf{R}^T \mathbf{s}_k - \gamma_k \mathbf{s}_k \right) \\ &= \mathbf{s}_k^T \left(\sum_{i=1}^n \lambda_i^2 \mathbf{r}_i \mathbf{r}_i^T \right) \mathbf{s}_k - 2\gamma_k \mathbf{s}_k^T \left(\sum_{i=1}^n \lambda_i \mathbf{r}_i \mathbf{r}_i^T \right) \mathbf{s}_k + \gamma_k^2 \\ &= \sum_{i=1}^n (\lambda_i^2 - 2\lambda_i \gamma_k) \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k + \gamma_k^2 \sum_{i=1}^n \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k \\ &= \sum_{i=1}^n (\lambda_i - \gamma_k)^2 \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k, \end{aligned}$$

where the fifth equality holds because \mathbf{s}_k has unit norm, implying $\sum_i \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k = \sum_i (\mathbf{r}_i^T \mathbf{s}_k)^2 = 1$. Further, note:

$$\begin{aligned} \min_{i \neq k} (\lambda_i - \gamma_k)^2 \sum_{i \neq k} \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k &= \sum_{i \neq k} \min_{i \neq k} (\lambda_i - \gamma_k)^2 \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k \\ &\leq \sum_{i \neq k} (\lambda_i - \gamma_k)^2 \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k \leq \mathbf{s}_k^T (\mathbf{A} - \mathbf{C})^T (\mathbf{A} - \mathbf{C}) \mathbf{s}_k \\ &= \mathbf{s}_k^T \mathbf{B}^T \mathbf{B} \mathbf{s}_k \leq \|\mathbf{B}\|^2, \end{aligned}$$

implying that as long as $\lambda_i - \gamma_k \neq 0$, we have that

$$\sum_{i \neq k} \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k \leq \frac{\|\mathbf{B}\|^2}{\min_{i \neq k} (\lambda_i - \gamma_k)^2}. \quad (8)$$

Because $\sum_i \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k = 1$, we can rewrite (8) as:

$$\begin{aligned} \sum_{i \neq k} \mathbf{s}_k^T \mathbf{r}_i \mathbf{r}_i^T \mathbf{s}_k &= 1 - \mathbf{s}_k^T \mathbf{r}_k \mathbf{r}_k^T \mathbf{s}_k = 1 - \cos^2 \theta_k \\ &= \sin^2 \theta_k \leq \frac{\|\mathbf{B}\|^2}{\min_{i \neq k} (\lambda_i - \gamma_k)^2}. \end{aligned} \quad (9)$$

Finally, by the triangle inequality,

$$|\lambda_i - \gamma_k| \geq |\lambda_i - \lambda_k| - |\lambda_k - \gamma_k|$$

and since $|\lambda_k - \gamma_k| \leq \|\mathbf{B}\|$ and $\|\mathbf{B}\| \leq \frac{1}{2} \min_{i \neq j} |\lambda_i - \lambda_j|$ by assumption, we conclude that:

$$|\lambda_i - \gamma_k| \geq \frac{1}{2} \min_{i \neq j} |\lambda_i - \lambda_j|.$$

Then, we can rewrite the bound (9) and find:

$$\begin{aligned} \sin^2 \theta_k &\leq \frac{4\|\mathbf{B}\|^2}{\min_{i \neq j} (\lambda_i - \lambda_j)^2}, \\ \Rightarrow |\theta_k| &\leq \sin^{-1} \left(\frac{2\|\mathbf{B}\|}{\min_{i \neq j} |\lambda_i - \lambda_j|} \right). \quad \blacksquare \end{aligned}$$

The bound above is expressed in terms of the norm of the perturbation \mathbf{B} and the difference in the eigenvalues of \mathbf{A} . We can exploit this result to obtain a separate and strictly tighter bound expressed in terms of the maximum *eigenvalue separation* of \mathbf{B} .

Corollary 1: Under the same assumptions of Theorem 1, the angle θ_k between the k th eigenvector of \mathbf{A} and $\mathbf{A} + \mathbf{B}$ is bounded by:

$$|\theta_k| \leq \sin^{-1} \left(\frac{\max_{i \neq j} |\zeta_i - \zeta_j|}{\min_{i \neq j} |\lambda_i - \lambda_j|} \right), \quad (10)$$

where ζ_i is the i th eigenvalue of \mathbf{B} .

Proof: First, note that adding a multiple of the identity $k\mathbf{I}$ to a matrix does not alter its eigenvectors, and shifts all its eigenvalues by k . Thus, the eigenvectors of $\mathbf{A} + \mathbf{B}$ are the same as those of the matrix

$$\mathbf{A} + \mathbf{B} - \frac{1}{2} (\bar{\lambda}(\mathbf{B}) + \underline{\lambda}(\mathbf{B})) \mathbf{I}$$

where $\bar{\lambda}(\mathbf{B})$ and $\underline{\lambda}(\mathbf{B})$, respectively, denote the maximum and minimum eigenvalues of \mathbf{B} .

Accordingly, we may consider the equivalent perturbation $\widehat{\mathbf{B}} = \mathbf{B} - \frac{1}{2}(\overline{\lambda}(\mathbf{B}) + \underline{\lambda}(\mathbf{B}))\mathbf{I}$. By construction, then, $\overline{\lambda}(\widehat{\mathbf{B}})$ and $\underline{\lambda}(\widehat{\mathbf{B}})$ are:

$$\begin{aligned}\overline{\lambda}(\widehat{\mathbf{B}}) &= \frac{1}{2}(\overline{\lambda}(\mathbf{B}) - \underline{\lambda}(\mathbf{B})) = \frac{1}{2} \max_{i,j} |\zeta_i - \zeta_j|, \\ \underline{\lambda}(\widehat{\mathbf{B}}) &= -\frac{1}{2}(\overline{\lambda}(\mathbf{B}) - \underline{\lambda}(\mathbf{B})) = -\frac{1}{2} \max_{i,j} |\zeta_i - \zeta_j|.\end{aligned}$$

Therefore, the induced ℓ_2 norm of $\widehat{\mathbf{B}}$ is:

$$\|\widehat{\mathbf{B}}\| = \frac{1}{2} \max_{i,j} |\zeta_i - \zeta_j|, \quad (11)$$

and by Theorem 1, since perturbing by $\widehat{\mathbf{B}}$ and \mathbf{B} produces the same eigenvalue perturbation to \mathbf{A} , we conclude that for the k th eigenvector angle:

$$|\theta_k| \leq \sin^{-1} \left(\frac{\max_{i \neq j} |\zeta_i - \zeta_j|}{\min_{i \neq j} |\lambda_i - \lambda_j|} \right). \quad \blacksquare$$

For a symmetric matrix \mathbf{B} it always holds that $\max_{i \neq j} |\zeta_i - \zeta_j| \leq 2\|\mathbf{B}\|$, making this a tighter bound than that in Theorem 1. Intuitively, Corollary 1 *minimizes* the bound (7) over all matrices $\widehat{\mathbf{B}}$ that produce the same eigenvectors as $\mathbf{A} + \mathbf{B}$.

B. From Overlap to Eigenvalue Separation

In [8], the authors prove that their notion of compact set overlap δ can be used to bound the *norm* of the difference between the second moments of each set. However, Corollary 1 requires a bound on the *eigenvalue separation*, rather than norm. In this section, we show how a similar analysis can also yield an equally tight bound on eigenvalue separation.

First, we need a measure of how “close” two sets H and H' are. To characterize this, we recall the following notions of size ρ and overlap δ from [8]:

$$\rho(H, H') = \min_{\mathbf{q} \in \mathbb{R}^n} \max_{\mathbf{x} \in H \cup H'} \|\mathbf{x} - \mathbf{q}\|, \quad (12)$$

$$\delta(H, H') = \frac{|H \cap H'|}{\max\{|H|, |H'|\}}. \quad (13)$$

Note that $\rho(H, H')$ describes the radius of the smallest ℓ_2 -ball containing both H and H' , while $\delta(H, H')$ is a measure of geometric similarity.

Remark 1: When using PASTA, the definitions of δ and ρ do not apply directly to the sets H_1 and $H_2 = \mathbf{R}H_1 + \mathbf{p}$, which will not overlap in general. Instead, we are interested in how similar the (noisy) *observed* set H_2 is to the reference set H_1 under the *true* rotation \mathbf{R} and translation \mathbf{p} . This is equivalent to using δ and ρ to compare a “perturbed” version of the set H_1 , which we call H'_1 , which is *defined* by $H_2 = \mathbf{R}H'_1 + \mathbf{p}$. Note that $H_1 \neq H'_1$ in general because of occlusions, LiDAR resolution, and noise. In the following, δ quantifies the overlap between H_1 and H'_1 .

We need the following result, which we recall from [8].

Theorem 2 (Theorem 2, [8]): Let $H, H' \subset \mathbb{R}^n$ be compact sets of non-zero measure, and let their first moments be $\mathbf{c}, \mathbf{c}' \in \mathbb{R}^n$ respectively. Further, let their size be $\rho(H, H') \in \mathbb{R}_{\geq 0}$ and their overlap be $\delta(H, H') \in [0, 1]$. Then:

$$\|\mathbf{c} - \mathbf{c}'\| \leq 2(1 - \delta(H, H'))\rho(H, H'). \quad (14)$$

We now present our new result.

Theorem 3: Let $H, H' \subset \mathbb{R}^n$ be compact sets of non-zero measure, and let their second moments be $\Sigma, \Sigma' \in \mathbb{R}^{n \times n}$ respectively. If their size is $\rho(H, H')$ and their overlap is $\delta(H, H') \in [0, 1]$, then:

$$\overline{\lambda}(\Sigma' - \Sigma) - \underline{\lambda}(\Sigma' - \Sigma) \leq (2(1 - \delta) + 4(1 - \delta)^2)\rho^2, \quad (15)$$

where $\overline{\lambda}(\cdot)$ and $\underline{\lambda}(\cdot)$ again denote the maximum and minimum eigenvalues of a symmetric matrix.

Proof: For space reasons, we omit a detailed derivation of the following claim, but its proof can be found in [8, Theorem 2]. Let \mathbf{c} and \mathbf{c}' be the first moments of H and H' respectively, and let us define $\Delta\mathbf{c} = \mathbf{c}' - \mathbf{c}$. Then, the following holds:

$$\begin{aligned}\Sigma' - \Sigma &= \int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_+(\mathbf{x}) d\mu \\ &\quad - \int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_-(\mathbf{x}) d\mu + \Delta\mathbf{c}\Delta\mathbf{c}^T,\end{aligned} \quad (16)$$

where f_+ and f_- are non-negative and satisfy:

$$\int f_+(\mathbf{x}) d\mu = \int f_-(\mathbf{x}) d\mu \leq 1 - \delta. \quad (17)$$

Note that $\Sigma' - \Sigma$ is the sum of three terms such that:

- 1) $\int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_+(\mathbf{x}) d\mu$ is positive definite with minimum eigenvalue zero and maximum eigenvalue:

$$\begin{aligned}\overline{\lambda} \left(\int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_+(\mathbf{x}) d\mu \right) \\ &= \left\| \int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_+(\mathbf{x}) d\mu \right\| \\ &\leq \int \|(\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T\| f_+(\mathbf{x}) d\mu \leq (1 - \delta)\rho^2.\end{aligned}$$

- 2) Analogously, $-\int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_-(\mathbf{x}) d\mu$ is negative definite with maximum eigenvalue zero and minimum eigenvalue:

$$\underline{\lambda} \left(-\int (\mathbf{x} - \mathbf{c}')(\mathbf{x} - \mathbf{c}')^T f_-(\mathbf{x}) d\mu \right) \geq -(1 - \delta)\rho^2.$$

- 3) The eigenvalues of $\Delta\mathbf{c}\Delta\mathbf{c}^T$ are all zero except for $\overline{\lambda}(\Delta\mathbf{c}\Delta\mathbf{c}^T) = \|\Delta\mathbf{c}\|^2 \leq 4(1 - \delta)^2\rho^2$. Where the inequality holds by Theorem 2.

Consequently, the eigenvalues of $\Sigma' - \Sigma$ are bounded as:

$$\begin{aligned}\overline{\lambda}(\Sigma' - \Sigma) &\leq (1 - \delta)\rho^2 + 0 + 4(1 - \delta)^2\rho^2, \\ \underline{\lambda}(\Sigma' - \Sigma) &\geq 0 - (1 - \delta)\rho^2 + 0.\end{aligned}$$

Taking the difference, we recover:

$$\begin{aligned}\overline{\lambda}(\Sigma' - \Sigma) - \underline{\lambda}(\Sigma' - \Sigma) &\leq \\ &\leq (1 - \delta)\rho^2 + 4(1 - \delta)^2\rho^2 - (-(1 - \delta)\rho^2) \\ &= (2(1 - \delta) + 4(1 - \delta)^2)\rho^2. \quad \blacksquare\end{aligned}$$

With Theorem 3 in hand, we can express our new, tighter error bounds as a function of $\delta(H, H')$ and $\rho(H, H')$.

Theorem 4 (PASTA error bound): Let $H, H' \subset \mathbb{R}^n$ be non-empty compact sets of non-zero measure with an overlap

of $\delta(H, H') \in [0, 1]$. Let \mathbf{c}, \mathbf{c}' and Σ, Σ' be the first and second moments of H, H' , and define the constants:

$$e_c = 2(1 - \delta)\rho(H, H'),$$

$$e_\Sigma = (2(1 - \delta(H, H')) + 4(1 - \delta(H, H'))^2) \rho^2(H, H').$$

Let (\mathbf{R}, \mathbf{p}) be the true transformation relating H to $\mathbf{R}H' + \mathbf{p}$, and $(\hat{\mathbf{R}}, \hat{\mathbf{p}}) = \text{PASTA}(H, \mathbf{R}H' + \mathbf{p})$ be the transformation estimated by PASTA. Then, if $\min_{i,j} |\lambda_i - \lambda_j| > 2e_\Sigma$, where λ_i is the i th eigenvalue of the second moment of H , the following holds:

$$\|\hat{\mathbf{R}} - \mathbf{R}\| \leq \sqrt{n} \frac{e_\Sigma}{\underline{\Delta}_\lambda(\Sigma)} \quad (18)$$

$$\|\hat{\mathbf{p}} - \mathbf{p}\| \leq \sqrt{n} \frac{e_\Sigma}{\underline{\Delta}_\lambda(\Sigma)} \|\mathbf{c}\| + e_c.$$

Proof: By the analysis in [8, Section IV], it holds that:

$$\|\hat{\mathbf{R}} - \mathbf{R}\| = \|\mathbf{V}'\mathbf{V}^T - I\| = \|\mathbf{V}' - \mathbf{V}\|, \quad (19)$$

where \mathbf{V} and \mathbf{V}' are the eigenvector matrices of the covariance matrices of H and H' respectively. Then, note that by simple trigonometry if a vector $\mathbf{v} \in \mathbb{R}^n$ is rotated by an angle θ into the vector \mathbf{v}' , their distance is $\|\mathbf{v}' - \mathbf{v}\| = 2\|\mathbf{v}\| \left| \sin \frac{\theta}{2} \right|$. Therefore, by Theorem 1 and Theorem 3:

$$\begin{aligned} \|\hat{\mathbf{R}} - \mathbf{R}\|^2 &= \|\mathbf{V}' - \mathbf{V}\|^2 \leq \|\mathbf{V}' - \mathbf{V}\|_F^2 \leq \sum_{i=1}^n \|\mathbf{v}'_i - \mathbf{v}_i\|^2 \\ &\leq \sum_{i=1}^n \left(2 \sin \left(\frac{1}{2} \sin^{-1} \left(\frac{\overline{\Delta}_\lambda(\Sigma' - \Sigma)}{\underline{\Delta}_\lambda(\Sigma)} \right) \right) \right)^2 \\ &\leq \sum_{i=1}^n \left(2 \sin \left(\sin^{-1} \frac{1}{2} \left(\frac{\overline{\Delta}_\lambda(\Sigma' - \Sigma)}{\underline{\Delta}_\lambda(\Sigma)} \right) \right) \right)^2 \\ &\leq n \left(\frac{\overline{\Delta}_\lambda(\Sigma' - \Sigma)}{\underline{\Delta}_\lambda(\Sigma)} \right)^2 \leq n \left(\frac{e_\Sigma}{\underline{\Delta}_\lambda(\Sigma)} \right)^2. \end{aligned}$$

As for the position error, also by the analysis in [8, Section IV] and Theorem 2, it holds that:

$$\begin{aligned} \|\hat{\mathbf{p}} - \mathbf{p}\| &\leq \|\hat{\mathbf{R}} - \mathbf{R}\| \|\mathbf{c}\| + \|\mathbf{c}' - \mathbf{c}\| \\ &\leq \sqrt{n} \frac{e_\Sigma}{\underline{\Delta}_\lambda(\Sigma)} \|\mathbf{c}\| + \|\mathbf{c}' - \mathbf{c}\| \leq \sqrt{n} \frac{e_\Sigma}{\underline{\Delta}_\lambda(\Sigma)} \|\mathbf{c}\| + e_c. \end{aligned}$$

We now substantiate the claim that this bound is a major improvement over the one in [8, Theorem 3]. In Fig. 2 we illustrate the difference between the bounds by plotting them as functions of δ and fixing their other parameters at some arbitrary value (the bound from our Theorem 4 is tighter than the one from [8, Theorem 3] regardless of choice). In particular, we plot the functions for the parameters: $n = 2$, $\rho = 5\text{m}$, $\underline{\Delta}_\lambda(\Sigma) = 1$, and $\|\mathbf{c}\| = 1\text{m}$. The rotation error is expressed as an angle estimation error for ease of interpretation.

Remark 2: PASTA's theoretical guarantees can be applied to other algorithms with a supervisory approach. First run PASTA, which produces an estimate with worst-case guarantees. Then, run any other algorithm to find a new estimate. The triangle inequality immediately provides a naive guarantee on the estimate provided by this new algorithm.




Fig. 2. Comparison of the bound in Theorem 4 and [8, Theorem 3].

Remark 3: Theorem 4 depends on the similarity, δ , between the two convex hulls H_1 and H_2 . Occlusions or non-convex geometry cause dramatic changes in these shapes, and the overlap δ naturally reduces and the estimation error increases according to (18). Note that when the shapes are sufficiently different and δ is very small, the measurements have so little in common that it is unrealistic to expect *any algorithm* to correctly solve the localization problem.

Remark 4: PASTA's error bound can be computed through either $\overline{\Delta}_\lambda(B)$ or the geometric parameters (δ, ρ) . This flexibility raises the question of which quantity to use in a practical setting. Either set of parameters can be estimated by performing calibration experiments in an environment with access to the ground truth poses by computing the true $\overline{\Delta}_\lambda(B)$ or δ at these poses, then estimating their worst-case values outside. Note that the bound computed directly from an estimate of $\overline{\Delta}_\lambda(B)$ provides a significantly tighter bound (see Fig. 4). δ and ρ , however, have a clearer *physical* interpretation and are easier to reason about during the calibration and subsequent worst-case estimation.

V. FROM THEORY TO PRACTICE

While the bounds from Corollary 1 and Theorems 2 and 4 hold in theory, we now establish that they are reasonably tight and behave as expected in extreme settings.

We equipped a small wheeled robot with a 360° 2D LiDAR with 0.5° of angular resolution (i.e., each LiDAR point cloud consists of 720 points in the plane). We placed the robot in a closed indoor environment with several obstacles, measuring both ground truth poses and LiDAR sensor measurements as the robot constantly moved around the environment. See Fig. 3 for a sample LiDAR measurement and the robot's trajectory in the environment.

Remark 5: We use a 2D LiDAR sensor in this regime since we are considering the indoor navigation setting, where 2D LiDAR is often sufficient. The *theory* developed for PASTA extends naturally to three dimensions.

For a given pair of scans from the experiment, we compute the error bounds in two ways: i) We use the ground truth measurements to align the convex hulls of LiDAR point clouds in the two compared scans, then we compute the size and overlap parameters ρ and δ , and the error bound via Theorem 4; ii) Since both scans are now aligned with the true transformation, we can also directly compute the perturbations to the second moments $\overline{\Delta}_\lambda(\Sigma' - \Sigma) = \overline{\Delta}_\lambda(B)$, which characterizes the bound in Corollary 1. The first



Fig. 3. Sample LiDAR scan and trajectory (left) and image (right) of the robot from the experimental setup.

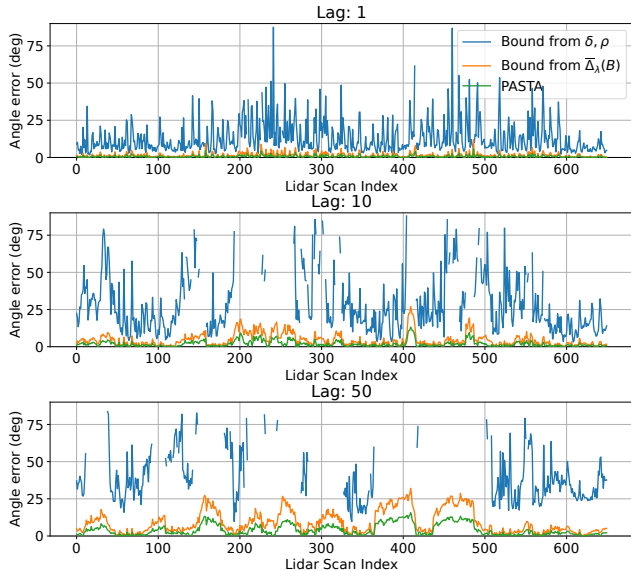


Fig. 4. Comparison of the error bound in Corollary 1 (orange) and Theorem 4 (blue) and the empirical error of PASTA (green) on real LiDAR data. Increased “lag” implies larger times between compared LiDAR measurements.

approach is equivalent to first computing an upper bound for $\bar{\Delta}_\lambda(\Sigma' - \Sigma)$ from δ and ρ , and then using this upper bound in Corollary 1.

For each LiDAR scan at index i , we align it with the scan at index $i + k$, and plot the computed error bounds along the trajectory for three different values of “lag” $k \in \mathbb{N}$, displayed in Fig. 4. Generally, with higher k the scans we compare are taken from poses further apart along the trajectory, thus more likely to be different, with the expectation of a worse bound and estimation error. Other variations naturally occur depending on the obstacles and their relative location to the robot, seen as small spikes in Fig. 4. We only show angular error bounds, as translation error bounds are a simple affine function of this value. For a given pair of scans, PASTA provides two candidate estimated angles separated by 180 degrees. We plot the error for the correct choice, as there are multiple ways to reliably select the correct one (see [8]), and the bounds only apply to this “correct choice”. Missing values in the plot correspond to scan pairs where $\min_{i,j} |\lambda_i - \lambda_j| > 2e_\Sigma$ fails to hold, violating the assumptions for our results.

We observe that the perturbation bound in Corollary 1 is very tight, and most of the difference between the bound in Theorem 4 and the actual estimation error comes from upper bounding $\bar{\Delta}_\lambda(\Sigma' - \Sigma)$ via knowledge of δ and ρ . As expected, larger distances between poses (as measured by

EXPERIMENT SUMMARY

Lag k	Bound from δ, ρ	Bound from $\bar{\Delta}_\lambda(\Sigma' - \Sigma)$	PASTA error
1	13.0°, 99.1%	1.4°, 100%	0.6°, n/a
10	27.9°, 80.5%	4.4°, 100%	1.9°, n/a
50	38.8°, 50.0%	9.0°, 100%	4.1°, n/a

For each lag and error type, we list the mean error over the experiment and the fraction of data points that satisfy the necessary assumptions.

the “lag” k) lead to higher error bounds. To counteract these effects, one may need a weak correspondence method that identifies common measured regions of both scans, which is a direction of current research.

VI. CONCLUSIONS

In this paper, we presented a LiDAR localization algorithm called PASTA and derived new theoretical worst-case bounds on its estimation error. These worst-case bounds are crucial for interfacing robotic systems using LiDAR for localization with safety-critical control algorithms. We also provided experimental evidence highlighting the tightness of the bounds and where improvements could be made.

REFERENCES

- [1] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):239–256, feb 1992.
- [2] S. Dean, A. Taylor, R. Cosner, B. Recht, and A. Ames. Guaranteeing safety of learned perception modules via measurement-robust control barrier functions. In *Proc. of the 2020 Conference on Robot Learning*, volume 155 of *PMLR*, pages 654–670, 16–18 Nov 2021.
- [3] C. Debeunne and D. Vivet. A review of visual-LiDAR fusion based simultaneous localization and mapping. *Sensors*, 20(7):2068, 2020.
- [4] J. Flusser, T. Suk, and B. Zitová. *2D and 3D Image Analysis by Moments*. John Wiley & Sons, Ltd., 2017.
- [5] Z. Gojcic, C. Zhou, J.D. Wegner, and A. Wieser. The perfect match: 3d point cloud matching with smoothed densities. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [6] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *Josa a*, 4(4):629–642, 1987.
- [7] M. Labbé and F. Michaud. Rtab-map as an open-source LiDAR and visual simultaneous localization and mapping library for large-scale and long-term online operation. *Journal of Field Robotics*, 36(2):416–446, 2019.
- [8] M. Marchi, J. Bunton, B. Ghahesifard, and P. Tabuada. Lidar point cloud registration with formal guarantees. In *61st IEEE Conference on Decision and Control (CDC)*, pages 3462–3467, 2022.
- [9] F. Pomerleau, F. Colas, and R. Siegwart. A review of point cloud registration algorithms for mobile robotics. *Foundations and Trends® in Robotics*, 4(1):1–104, 2015.
- [10] C. Qi, H. Su, K. Mo, and L. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, July 2017.
- [11] H.Z.U. Rehman and S. Lee. Automatic image alignment using principal component analysis. *IEEE Access*, 6:72063–72072, 2018.
- [12] R. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (FPFH) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217, 2009.
- [13] A. Segal, D. Haehnel, and S. Thrun. Generalized ICP. In *Robotics: science and systems*, volume 2, page 435. Seattle, WA, 2009.
- [14] R.W. Wolcott and R.M. Eustice. Fast LiDAR localization using multiresolution gaussian mixture maps. In *2015 IEEE International Conference on Robotics and Automation*, pages 2814–2821. IEEE, 2015.
- [15] H. Yang, J. Shi, and L. Carlone. TEASER: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2021.
- [16] J. Yang, H. Li, and Y. Jia. Go-icp: Solving 3d registration efficiently and globally optimally. In *2013 IEEE International Conference on Computer Vision*, pages 1457–1464, 2013.