

Matching Line Segment Scans with Mutual Compatibility Constraints

Mladen Mazuran¹

Francesco Amigoni²

Abstract—Over the years, proposals have been made to employ line segments to build 2D maps of indoor environments. One of the basic steps of these approaches is the matching between scans (or, more generally, sets) of line segments, which is usually addressed using variants of the Iterative Closest Line (ICL) paradigm. ICL is based on the idea of associating closest line segments belonging to the two scans and of reducing the distance between them. In this paper, we propose two algorithms that go beyond this approach by exploiting the *mutual compatibility* between associations of line segments. Experimental results show that our algorithms significantly outperform, in terms of matching accuracy, traditional algorithms based on ICL, at the cost of a slightly longer execution time.

I. INTRODUCTION

Building 2D maps of indoor environments is often performed by acquiring environmental data with laser range scanners and integrating these scans with a SLAM framework [18]. While scans composed of raw points are usually considered, an alternative approach exploits higher-level geometric features like *line segments* [5, 7, 15], which are extracted from point data [13]. The advantages consist in more compact representations for maps (which also mean faster methods) and in better capturing and leveraging on the regular structures of indoor environments [5].

An enabling step in building line segment maps is the matching between two scans composed of line segments. The scan matching problem can be described as that of finding a roto-translation that maximizes the overlap between the line segments of the two scans, according to some distance measure. The dominant paradigm for scan matching is Iterative Closest Line (ICL) [10], which iteratively finds the roto-translation that minimizes the distance between pairs of nearest (associated) line segments belonging to the two scans. Neira and Tardós [12] provide a sound argument that nearest neighbor associations do not guarantee that correlations between the underlying features will be preserved. This is especially true for ICL: due to the small number of line segments which usually compose a scan, even a single failed association may have devastating effects on the alignment quality. The effect is most evident when a line segment in a scan has no corresponding line segment in the other scan, as usually happens in incremental map building. An example is shown in Fig. 1, where trying to minimize the distance between nearest line segments leads to a wrong roto-translation estimate.

¹Mladen Mazuran is with the Institute of Computer Science, University of Freiburg, Germany mazuran@informatik.uni-freiburg.de

²Francesco Amigoni is with the Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy francesco.amigoni@polimi.it

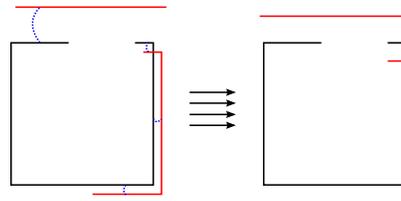


Fig. 1. On the left, the red scan is matched against the black scan by using the associations represented by the dashed blue arcs. The absence of mutual compatibility between associations results in a failed roto-translation estimate, on the right.

In this paper, we propose an approach for matching scans composed of line segments that considers the *mutual compatibility* of associations. Informally, two associations between line segments are mutually compatible if they “induce” the same roto-translation. The same idea can be found in the work of Tardos et al. [17].

We formulate a constrained optimization problem that captures the idea of mutual compatibility and we introduce two algorithms for solving this problem. Experimental results show that our algorithms significantly outperform, in terms of matching accuracy, state-of-the-art algorithms based on ICL, at the cost of slightly longer execution times.

II. RELATED WORK

The current state of the art in line segment based mapping efforts encompasses various subjects but, considering SLAM applications, three major topics can be identified: line extraction, scan matching, and global alignment.

Line extraction deals with determining line segments from point scans and sports a plethora of different approaches such as split and merge and Hough transform. See [13] for a survey.

Scan matching between sets of line segments has been mainly tackled by algorithms belonging to the ICL class. ICL is heavily inspired by Iterative Closest Point (ICP) [6], with point-to-line ICP [4] in particular sharing many dualities. Given two point scans, a query and a reference, along with an estimate of their relative pose, first, a set of associations between points and surface lines belonging to different scans is determined on a nearest neighbor basis. Then, a roto-translation is computed by minimizing point-to-line distances, according to these associations. This transformation is applied to the query set of line segments and the process is iteratively repeated until convergence. ICL adopts the same four-step optimization process, the only differences being that points are replaced with line segments and that the notions of distance and roto-translation estimation change.

Numerous efforts based on this paradigm can be found in the literature, all of which differ either on the distance

function adopted or on the procedure for the computation of the roto-translation. For example, Li and Griffiths [10] measure the distance between line segments according to a weighted sum of their differences in position, direction, and length. At each iteration of their ICL algorithm, they find the translation in closed form and the rotation using an eigen- or a single-value decomposition algorithm. Amigoni et al. [2] use a distance between line segments inspired by the Hausdorff distance and calculate the roto-translation with an algorithmic approach that minimizes a function defined *ad hoc*. Alshawa [1] does not specify explicitly a distance function but provides a means of computing, at each iteration, the rotation and the translation in closed form. With our work, we both extend those algorithms by considering mutual compatibility, and move to a new paradigm, which goes beyond their iterative nature altogether.

Global alignment of line segments has been dealt with both probabilistic and *ad hoc* formulations. The basic probabilistic form has been proposed by Vandorpe et al. [19], although, due to absence of correlation between line segments, their work is unable to deal with incremental errors. This probabilistic formulation is extended by Castellanos et al. [3], Garulli et al. [7], and Pfister [15], who correlate line segments with a Kalman filter, although they provide evaluation either on simplistic test scenarios or synthetic data altogether. Nguyen et al. [14] provide, perhaps the first, line segment based approach consistently evaluated on a real dataset, although in doing so they introduce orthogonality constraints. A further work worth of notice is that of Elseberg et al. [5], who provide an iterative method for aligning line segment scans, based on a powerful distance function that can be also adopted in ICL applications.

The work by Tardos et al. [17] first introduced the concept of mutual, or joint, compatibility in the context of line segment approaches. They do so by using the *joint compatibility branch and bound* algorithm [12], which, contrary to our work, requires an underlying probabilistic representation and may result in computational requirements exponential in the number of line segments.

III. THE PROPOSED APPROACH

In this section we illustrate our method for matching pairs of line segment scans while ensuring mutual compatibility between associations. First, we introduce a general optimization problem aimed at computing the optimal set of associations, followed by the definition of possible objective and constraint functions for such problem. Then, we introduce two algorithms that ensure mutual compatibility. One adds a filtering step to ICL, while the other one deals away with the ICL paradigm altogether. Finally, a method for computing the roto-translation from a set of associations is introduced. The method ensures favorable properties when dealing with singular environments, although the proposed scan matchers are theoretically independent of such method. Note also that both algorithms are independent of the particular interpolation procedure and distance measure adopted.

When matching a pair of scans, we assume their line segments to be expressed in a global reference frame. We further assume one scan to be fixed, while the other to be observed from a pose in a locality of an *initial guess*. We call *assignment* any set of associations between line segments.

A. An optimization problem for mutual compatibility

In order to account for mutual compatibility between associations, we define a constrained optimization problem that describes how a mutually compatible assignment should behave. Under such a model we will seek an assignment where each association is constrained to provide a roto-translation close to all others, while at the same time maximizing the quality of the assignment.

Suppose that two collections of line segments (two scans, or a query scan and a map), indexed by the integer sets R and Q ¹, are given; then an optimal assignment S^* with respect to a threshold $\eta \in \mathbb{R}^k$ is defined as the one that maximizes the following:

$$\begin{aligned} \arg \max_{S \subseteq R \times Q} \sum_{(i,j) \in S} w_{ij}(S) \\ \text{s.t. } \mathbf{disp}(S) \leq \eta \end{aligned} \quad (1)$$

Where $\mathbf{disp} : \mathcal{P}(R \times Q) \rightarrow \mathbb{R}^k$ is a vector function that evaluates the “dispersion” of the assignment over k dimensions ($\mathcal{P}(\cdot)$ is the power set) and $w_{ij}(S)$ is a particular importance weight for the association of line segments i and j , given the full assignment set S . Note that, although the problem is defined for a generic k , in this work we will consider only the case $k = 2$. Dispersion and weight functions will be detailed in the next section. Note also that, in general, (1) is an NP-hard problem.

Although $R \times Q$ represents the set of all possible associations, the solution can be sought only across a subset of $R \times Q$, as line segments which are exceedingly dissimilar will not be considered as plausible associations. This can be achieved by a pre-processing step, using a distance function or a statistic test, such as the one proposed in [5] or in [15].

B. Dispersion and weight functions

In order to apply the problem (1) to actual data, it is necessary to define both dispersion and weight functions, which are non-unique and subjective, at least to some extent.

In order to determine a dispersion function that accounts for mutual compatibility, we first consider a single association (i, j) : enforcing these two line segments to overlap introduces two constraints, which allow to uniquely determine a rotation $\theta(i, j)$ around the origin and a translation $\mathbf{t}(i, j)$ along the axis normal to one of the lines. We call $\mathbf{H}(i, j)$ the roto-translation obtained composing $\theta(i, j)$ and $\mathbf{t}(i, j)$. Thus, in the ideal case, the associations (i, j) in an assignment S are *exactly mutually compatible* if all the rotations $\theta(i, j)$ are equal and each $\mathbf{t}(i, j)$ is the projection (on the axis along which $\mathbf{t}(i, j)$ is defined) of the same overall translation.

¹With a slight abuse of notation, in the following we will denote a line segment both using its integer index (e.g., i) and as a set of points S_i , when we would like to emphasize its geometrical nature.

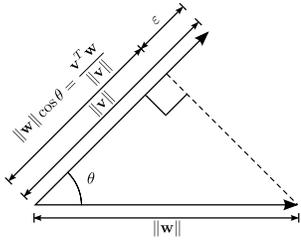


Fig. 2. Our deviation measure from a geometrical point of view.

For the rotation, a deviation from this ideal model can be easily quantified by means of a subtraction, the same is not true for the translation. For the latter, we define a deviation measure ε , of a translation vector $\mathbf{v} \in \mathbb{R}^2$ with respect to another $\mathbf{w} \in \mathbb{R}^2$:

$$\varepsilon(\mathbf{v}, \mathbf{w}) = \left\| \|\mathbf{v}\| - \frac{\mathbf{v}^T \mathbf{w}}{\|\mathbf{v}\|} \right\|$$

This deviation can be geometrically interpreted as the amount of length by which \mathbf{v} should be changed in order for the projection of \mathbf{w} on \mathbf{v} to be equal to \mathbf{v} . A sample of such deviation is shown in Fig. 2.

Suppose now that, given an assignment S , we know the “best” rotation $\theta(S)$ and translation $\mathbf{t}(S)$ (we address their actual definition and computation in Section III-E). It is then reasonable to consider as dispersion of an assignment S the maximum deviation in rotation and translation of each association (i, j) of S . We thus define the *maximum deviation* function:

$$\text{disp}_m(S) = \begin{bmatrix} \max_{(i,j) \in S} |\theta(i, j) - \theta(S)| \\ \max_{(i,j) \in S} \varepsilon(\mathbf{t}(i, j), \mathbf{t}(S)) \end{bmatrix} \quad (2)$$

For the weights $w_{ij}(S)$ in the objective function of (1) we consider the amount of overlap between the pair of line segments (i, j) . The main motivating factor is that, often, for perfectly aligned scans, the mutual overlap is maximized. More formally, let \mathcal{L} denote the line obtained by fusing line segments \mathcal{S}_i and \mathcal{S}_j with a procedure such as those of [9] or [11]. Projecting \mathcal{S}_i and \mathcal{S}_j onto \mathcal{L} yields two collinear line segments, which, when intersected, provide a (possibly empty) line segment \mathcal{S}_o . Let $\mathcal{O}(\mathcal{S}_i, \mathcal{S}_j)$ be the length of \mathcal{S}_o (which represents the overlap between projections of \mathcal{S}_i and \mathcal{S}_j), we choose as weight function between two line segments \mathcal{S}_i and \mathcal{S}_j the value:

$$w_{ij}(S) = \mathcal{O}(\mathcal{S}_i, \mathbf{H}(S) \cdot \mathcal{S}_j) \quad (3)$$

Where $\mathbf{H}(S) \cdot \mathcal{S}_j$ denotes that the rototranslation induced by the assignment S (assumed, for now, to be provided by an oracle) is applied to \mathcal{S}_j .

C. Approximate solution by Filtered ICL

Solving problem (1) can be done directly by an appropriate algorithm (see next section) or by re-adapting existing ICL algorithms to at least ensure that the computed associations satisfy the dispersion requirement. We discuss this latter option in this section.

Starting from an initial assignment \bar{S} , determined in a nearest neighbor fashion (according to any distance function,

such as the one in [5] or in [10]), it is possible to provide a greedy algorithm which computes an assignment $S \subseteq \bar{S}$ such that $\text{disp}(S) \leq \eta$, while also approximately minimizing $|\bar{S}| - |S|$. Thus, we formulate a vanilla ICL algorithm with this aim, called *Filtered ICL*, by adding a filtering step before computing the rototranslation from the associations in the assignment.

Intuitively, this can be done by iteratively removing the “most deviant” association in \bar{S} and repeating the process until the resulting assignment satisfies the dispersion constraint. Such an approach is approximate and greedy, hence does not guarantee an optimal solution.

This kind of formulation requires the definition of an order relation over the associations, based on their deviation. Considering that (2) defines a two dimensional vector subspace, determining which association is the most deviant effectively results in a multi-criteria decision problem. As an order statistic for (2) we propose a simple linear combination of the deviations:

$$\xi((i, j), S) = \lambda^T \begin{bmatrix} |\theta(i, j) - \theta(S)| \\ \varepsilon(\mathbf{t}(i, j), \mathbf{t}(S)) \end{bmatrix} \quad (4)$$

Where $\lambda \in \mathbb{R}^2$ is a weighting vector chosen experimentally. Note that (4) is scale-dependent, since the translational and rotational part have different units, which need to be taken into account when selecting λ .

D. A polynomial-time approximation algorithm

While the Filtered ICL algorithm provides a simple approach to account for mutual compatibility of associations, in this section we introduce a feasible algorithm to solve the full optimization problem (1). We first introduce a theoretical property of the optimization problem, and then we exploit this aspect to provide an alternative non-iterative algorithm.

In absence of noise, scans acquired in strongly structured environments are characterized by the fact that a pair of associated corners (or, more generally, a pair of non-parallel line segments) allows to unambiguously identify a rototranslation. Thus, the search for an assignment can be reduced to the search for a pair of associations between line segments, and we can state the following proposition.

Proposition. *Let the dispersion function be disp_m , let the importance weights $w_{ij}(S)$ be such that the resulting objective function is monotone non-decreasing with respect to the addition of an association, and let S^* be the optimal solution to (1). Then, if there exists a pair of associations $a_1 = (i, j)$ and $a_2 = (i', j')$ such that $\theta(\{a_1, a_2\}) = \theta(S^*)$ and $\mathbf{t}(\{a_1, a_2\}) = \mathbf{t}(S^*)$, the optimization problem (1) can be solved optimally in polynomial time.*

By assumption, the addition of an association to an assignment cannot decrease the value of the objective function, therefore, if $\{a_1, a_2\}$ are known, the optimal assignment S^* will be given by the largest set of associations which are compatible with $\{a_1, a_2\}$ within an η threshold. This is a consequence of the max operator in the definition of disp_m . The optimal assignment S^* can be then found by checking each pair of associations against all remaining ones and

Algorithm 1 RANSAC-based matching algorithm

Require: A set of plausible associations $S' \subset R \times Q$

- 1: $S_{\text{best}}, z_{\text{best}}, i, j \leftarrow \emptyset, 0, 0$
- 2: **while** $i \leq k \wedge j \leq w$ **do**
- 3: $a_1, a_2 \leftarrow \text{SamplePair}(S')$
- 4: $S_{\text{try}} \leftarrow \{a_1, a_2\}$
- 5: $j \leftarrow j + 1$
- 6: **if** $\text{disp}_m(S_{\text{try}}) > \eta$ **then**
- 7: **continue**
- 8: **end if**
- 9: **for each** $a_3 \in S' \wedge a_3 \neq a_1 \wedge a_3 \neq a_2$ **do**
- 10: **if** $[|\theta(a_3) - \theta(\{a_1, a_2\})| \varepsilon(\mathbf{t}(a_3), \mathbf{t}(\{a_1, a_2\}))]^T \leq \eta$ **then**
- 11: $S_{\text{try}} \leftarrow S_{\text{try}} \cup \{a_3\}$
- 12: **end if**
- 13: **end for**
- 14: $z_{\text{try}} \leftarrow \text{ObjectiveFunction}(S_{\text{try}})$
- 15: **if** $z_{\text{try}} > z_{\text{best}}$ **then**
- 16: $S_{\text{best}}, z_{\text{best}} \leftarrow S_{\text{try}}, z_{\text{try}}$
- 17: **end if**
- 18: $i \leftarrow i + 1$
- 19: **end while**
- 20: **return** Rototranslation(S_{best})

adding to S^* the pair that maximizes the objective function, along with all other compatible associations. If n is the number of possible associations to be considered, this results in a $\Theta(n^3)$ algorithm.

When dealing with a real scan matching problem, however, the assumptions made, in general, need to be relaxed and, while no optimality guarantee can be given, the results are expected to be close to the optimum. For instance, the best associated corners may provide a very noisy estimate of the rototranslation, and while constant weights ensure the monotonicity assumption, the weights defined in (3) do not. Nevertheless, we experimentally found that, in indoor environments, the rototranslation estimate does not change significantly for associations which are mutually compatible, therefore, the objective function roughly approximates a sum of constant weights.

An improvement in the computational effort can be achieved by adopting a RANSAC formulation: rather than evaluating all possible pairs of associations, we randomly samples just a subset of them. This is a reasonable choice because, often, indoor domains present many corners. Moreover, two sampled associations $a_1 = (i, j)$ and $a_2 = (i', j')$ may already be incompatible and can be safely discarded. This naturally leads to a double-bound RANSAC formulation: rather than sampling always k times, we run the procedure at least k times but up to w times. Then, $w - k$ would represent the “grace window” allowed for the discards; if less than $w - k$ discards happen, the double bound ensures that exactly k “meaningful” samplings are considered. This results in an algorithm with an improved time complexity of $\Omega(kn)$ and $\mathcal{O}(wn)$, which is outlined in Algorithm 1.

E. Computing the rototranslation

Up until now we have assumed that an oracle provides the best rototranslation for a particular assignment (Section III-B). In order to implement this feature in practice we adopt a technique similar to that of Zhang and Ghosh [20]. We consider the estimation of the rotation and translation separately. The first is obtained by solving a weighted least squares

problem, while the latter by minimizing the weighted sum of squared point-to-line distances, after having applied the rotation. Mathematically, given an assignment S , we seek to minimize, with respect to θ and \mathbf{t} , the cost functions:

$$J(\theta, S) = \sum_{(i,j) \in S} u_{ij} (\alpha_i - \alpha_j - \theta)^2 \quad (5)$$

$$J(\mathbf{t}, S) = \sum_{(i,j) \in S} v_{ij} \mathcal{D}^2(\mathbf{c}_j - \mathbf{t}, \mathbf{R}(\theta) \cdot \mathcal{S}_i) \quad (6)$$

Where u_{ij} and v_{ij} are particular chosen weights, $\mathbf{R}(\theta)$ is a rotation by θ around the origin, $\mathcal{D}(\mathbf{p}, \mathcal{S})$ is the point-to-line distance of point \mathbf{p} from line \mathcal{S} , and α_i and \mathbf{c}_i are the angular orientation and centroid, respectively, of line segment \mathcal{S}_i .

The minima of (5) and (6) are given by the values (proof is rather straightforward and thus omitted):

$$\hat{\theta} = - \left(\sum_{(i,j) \in S} u_{ij} \right)^{-1} \sum_{(i,j) \in S} u_{ij} (\alpha_i - \alpha_j) \quad (7)$$

$$\hat{\mathbf{t}} = - \left(\sum_{(i,j) \in S} v_{ij} \mathbf{n}_i \mathbf{n}_i^T \right)^+ \sum_{(i,j) \in S} v_{ij} \mathbf{n}_i \mathbf{n}_i^T (\mathbf{p}_i - \mathbf{c}_j) \quad (8)$$

Where \mathbf{A}^+ denotes the Moore-Penrose pseudoinverse of \mathbf{A} , \mathbf{n}_i is the unit vector normal to $\mathbf{R}(\theta) \cdot \mathcal{S}_i$, and \mathbf{p}_i is one of its end points. Notice that when S consists of a single association (i, j) , $\hat{\mathbf{t}}$ correctly simplifies to a translation by the point-to-line distance along the normal vector \mathbf{n}_i :

$$\hat{\mathbf{t}} = \mathbf{n}_i \mathbf{n}_i^T (\mathbf{p}_i - \mathbf{c}_j) = \mathbf{n}_i \mathcal{D}(\mathbf{c}_j, \mathbf{R}(\theta) \cdot \mathcal{S}_i) \quad (9)$$

The values u_{ij} and v_{ij} provide an effective way to weight the importance of associations, as spurious line segments tend to corrupt the estimate. While it is possible to provide a maximum likelihood formulation with particular values of u_{ij} and v_{ij} , we have found that this choice often gives results inferior to even a simple uniform weighting approach. This outcome is possibly explained by the inadequate estimation of the covariance of a line segment, when the interpolated points do not follow a consistent linear model (e.g., clutter).

Beyond a maximum likelihood or uniform approach, we propose adopting the following weights for cost functions (5) and (6) (ℓ_i is the length of \mathcal{S}_i):

$$u_{ij} = v_{ij} = (\ell_i^{-1} + \ell_j^{-1})^{-1} \quad (10)$$

This formula takes into account the fact that the length of a line segment roughly correlates with a measure of its information. Therefore, if we assume the presence of an abstract variance σ_S^2 , we can argue that $\ell^{-1} \propto \sigma_S^2$, in which case u_{ij} and v_{ij} would represent a value proportional to the combined information of the two line segments.

As a final note, if the rotation $\mathbf{R}(\theta)$ in (6) is carried out around the initial guess (rather than around the origin of the global frame), then it is possible to ensure that in a corridor-like environment no translation will occur along the axis tangent to the walls. This comes as a favorable consequence to the use of the Moore-Penrose pseudoinverse. Rotation along the initial guess should be also employed in

the computation of the dispersion function (2), as rotating closer to the line segments helps reduce the effect of the lever arm when their orientation is subject to non-negligible errors.

IV. EXPERIMENTAL RESULTS

In order to evaluate quantitatively the performance of the proposed algorithms we assume the presence of ground truth data in the form of an exact map and of a set of scans which are known to have been recorded at certain poses relative to that map. Under this framework, the quality of a scan matcher can be then tested by picking a scan, corrupting its pose with Gaussian noise, and evaluating the accuracy, namely the frequency of how often the rototranslation computed by the scan matcher is in a neighborhood of the real recorded pose.

We consider as ground truth the pre-aligned versions of two datasets, the Rawseeds Bicocca 25b dataset² and MIT CSAIL dataset³, which represent a considerable challenge to line segment methods, due to the large amount of clutter and the many non-linear obstacles. Each dataset is composed of a sequence of 2D scans acquired by laser range scanners. We approximate the points in a scan with line segments: we segment the points with a split and merge approach and interpolate the resulting point clusters with a Theil-Sen estimator [16]. We then perform a realistic incremental matching test that requires algorithms to match a scan \mathcal{V}_m acquired at time m against a ground truth map composed by the union of all preceding scans \mathcal{V}_l $l < m$. This process is repeated 500 times for each scan, each time with a different corrupted pose.

We evaluated the following algorithms: Li and Griffiths ICL [10], ICL without filtering, point-to-line ICP [4] (computed on the original point-based map), Filtered ICL (Section III-C), and RANSAC matcher (Algorithm 1). In the results we will shorten these names to: L&G, Standard, PLICP, Filtered, and RANSAC, respectively.

The Standard, Filtered, and RANSAC algorithms determine the potential set of associations with a modified Elseberg et al. distance [5] that penalizes line segments which are tangentially distant, while the rototranslation is calculated with the method proposed in Section III-E, for both uniform (U) and length-based (L) weighting. We set the η dispersion threshold to 4.5° in rotation and 8 cm in translation, while the λ vector parameter for the Filtered ICL to $[0.3 \ 0.7]^T$. The RANSAC iterations range from $k = 4000$ to $w = 4500$ for maximum accuracy (generally much less iterations are required, e.g., reducing them to 200-250 affects the accuracy by less than 3%).

Table I reports the accuracy results on the two datasets. We considered as acceptability threshold 5° in rotation and 10 cm in translation, although it should be noted that the results remain roughly consistent by either reducing

TABLE I
ACCURACY RESULTS FOR DIFFERENT SCAN MATCHERS.

MIT CSAIL Dataset				
Error	L&G	Standard/U	Filtered/U	RANSAC/U
Small	38.3%	38.6%	88.8%	96.1%
Medium	31.4%	34.6%	83.9%	90.6%
Large	16.8%	26.0%	66.8%	72.7%
Error	PLICP	Standard/L	Filtered/L	RANSAC/L
Small	89.3%	58.3%	91.6%	96.4%
Medium	79.1%	55.2%	87.3%	91.1%
Large	54.6%	45.4%	69.2%	72.7%

Rawseeds Bicocca 25b Dataset				
Error	L&G	Standard/U	Filtered/U	RANSAC/U
Small	27.7%	58.1%	91.3%	93.9%
Medium	31.4%	50.2%	80.1%	80.5%
Large	13.2%	39.1%	64.4%	64.8%
Error	PLICP	Standard/L	Filtered/L	RANSAC/L
Small	93.0%	72.3%	92.4%	94.0%
Medium	80.5%	63.7%	81.4%	80.5%
Large	54.2%	51.6%	65.4%	64.6%

or enlarging these thresholds. The three error profiles reported in Table I refer to the covariance of the Gaussian noise by which the pose was corrupted, with $\Sigma_s = \text{diag}(0.002, 0.002, (\pi/120)^2)$ for the small profile, $10 \Sigma_s$ for the medium profile, and $100 \Sigma_s$ for the large profile. The errors are expressed in meters, i.e., a variance of 0.002 results in a standard deviation of approximately 4.5 cm.

The results show that the proposed methods provide a valid alternative to point-based scan matchers even in heavily cluttered environments. Although there is no substantial gain in accuracy when dealing with smaller errors, for larger errors the proposed methods provide, with confidence greater than 95%, statistically significant improvements with respect to PLICP. It should also be noted that, not only Filtered ICL improves the accuracy of the Standard ICL, but also reduces the average number of iterations required for convergence. Depending on the error profile, in the MIT CSAIL dataset we achieved a reduction in the number of iterations between 25.2% and 28.2%, while in the Bicocca 25b dataset between 12.3% and 16%.

We repeated the experiments of Table I assuming to know *all* the ground truth map (and not only the scans acquired before the corrupted one). In these localization-only tests, the Standard ICL algorithm behaves comparably close to the proposed methods, even in the absence of a mutual compatibility constraints. This is because any query line segment (almost) always has a corresponding map line segment, which is not the case for incremental matching.

Table II provides a comparison of the single-threaded average execution times for the two datasets with small errors, as computed on an Intel® Core™ i7-3770K 3,5 GHz processor. The execution times of our algorithms are currently larger than those of both PLICP and L&G, although this is mostly due to the time required for the lookup of plausible associations. The lookup was implemented with a trivial $\Theta(kn)$ algorithm, where k and n are the number of line segments in the scan and in a reduced locality of the map, respectively, which provides significant slow-downs due to the complexity of the Elseberg et al. distance [5].

²http://www.rawseeds.org/rs/capture_sessions/view/5

³<http://kaspar.informatik.uni-freiburg.de/~slamEvaluation/datasets.php>

TABLE II
AVERAGE EXECUTION TIMES FOR DIFFERENT SCAN MATCHERS.

Scan matcher	MIT CSAIL	Bicocca 25b
L&G	5.0 ms	1.0 ms
Standard	88.8 ms	11.8 ms
Filtered	63.2 ms	10.3 ms
RANSAC	52.1 ms	45.1 ms
PLICP	7.9 ms	3.9 ms

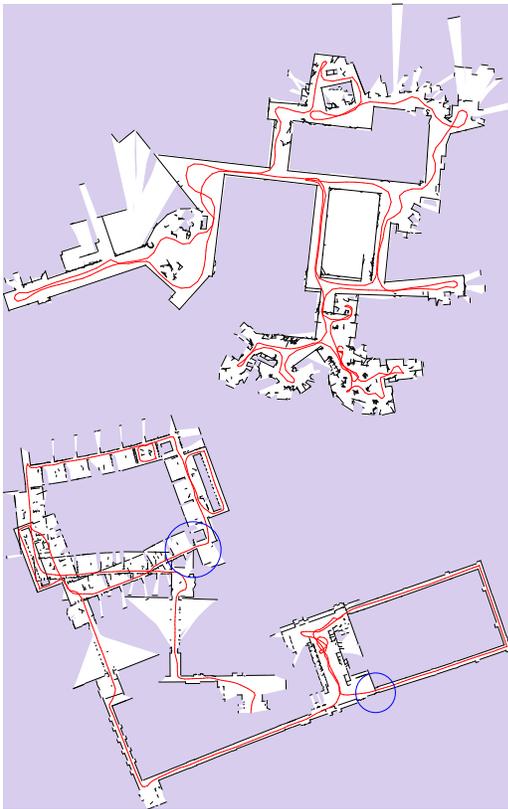


Fig. 3. Incremental RANSAC matching on the MIT CSAIL (top) and Bicocca 25b (bottom) datasets. Due to the absence of a SLAM framework it is impossible to close the highlighted loops in the Bicocca 25b dataset.

Our PLICP implementation, on the other hand, adopts k -d trees for association lookup, which achieve a much lower complexity of $\Theta(k \log n)$, while L&G relies on a much less complex distance. While this aspect could be definitely improved, the proposed RANSAC scan matcher is still fast enough to be executed in real time.

Finally, a qualitative evaluation of the RANSAC matcher is provided in Fig. 3, where the MIT CSAIL and Bicocca 25b datasets are matched incrementally, with no SLAM framework. Clearly, exact convergence is not possible due to incremental errors, nevertheless, the MIT CSAIL dataset results in only minor imperfections when closing the loop, while in the Bicocca 25b dataset two failed loop closures can be witnessed. This overall favorable performance can be explained by the ability of our approach to exploit the structure of indoor environments.

V. CONCLUSIONS

In this paper we introduced an approach for matching scans composed of line segments with mutual compatibility

between associations. Mutual compatibility allows to better exploit the structure of indoor environments while triggering the development of scan matching algorithms that go beyond the traditional iterative model of ICL. Experimental results are encouraging, showing that our algorithms are significantly more accurate (although slower) than those usually employed in building line segment maps. Moreover, by pure scan matching, we can reconstruct the maps of rather complex indoor environments with reasonable quality.

Future works will address the further enhancement of our approach, especially by extracting more reliable line segments from laser scans. Moreover, the use of our scan matchers in a complete SLAM framework, such as [8], and their extension to 3D will contribute to further assess their significance and utility.

REFERENCES

- [1] M. Alshawa. ICL: Iterative closest line. A novel point cloud registration algorithm based on linear features. *Ekscentar*, (10):53–59, 2007.
- [2] F. Amigoni, G. Fontana, and F. Garigiola. A method for building small-size segment-based maps. In *Proc. DARS*, pages 11–20, 2006.
- [3] J. Castellanos, J. Montiel, J. Neira, and J. Tardós. The spmap: A probabilistic framework for simultaneous localization and map building. *IEEE T Robot*, 15(5):948–952, 1999.
- [4] A. Censi. An ICP variant using a point-to-line metric. In *Proc. ICRA*, pages 19–25, 2008.
- [5] J. Elseberg, R. Creed, and R. Lakaemper. A line segment based system for 2D global mapping. In *Proc. ICRA*, pages 3924–3931, 2010.
- [6] O. Faugeras and M. Hebert. A 3-D recognition and positioning algorithm using geometrical matching between primitive surfaces. In *Proc. IJCAI*, pages 996–1002, 1983.
- [7] A. Garulli, A. Giannitrapani, A. Rossi, and A. Vicino. Simultaneous localization and map building using linear features. In *Proc. EECMR*, pages 44–49, 2005.
- [8] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. ICRA*, pages 3607–3613, 2011.
- [9] R. Lakaemper. Simultaneous multi-line-segment merging for robot mapping using mean shift clustering. In *Proc. IROS*, pages 1654–1660, 2009.
- [10] Q. Li and J. Griffiths. Iterative closest geometric objects registration. *Comput Math Appl*, 40(10):1171–1188, 2000.
- [11] M. Lopez-Sanchez, F. Esteva, R. de Mantaras, C. Sierra, and J. Amat. Map generation by cooperative low-cost robots in structured unknown environments. *Auton Robot*, 5(1):53–61, 1998.
- [12] J. Neira and J. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE T Robot*, 17(6):890–897, 2001.
- [13] V. Nguyen, S. Gächter, A. Martinelli, N. Tomatis, and R. Siegwart. A comparison of line extraction algorithms using 2D range data for indoor mobile robotics. *Auton Robot*, 23(2):97–111, 2007.
- [14] V. Nguyen, A. Harati, A. Martinelli, R. Siegwart, and N. Tomatis. Orthogonal SLAM: A step toward lightweight indoor autonomous navigation. In *Proc. IROS*, pages 5007–5012, 2006.
- [15] S. Pfister. *Algorithms for Mobile Robot Localization and Mapping, Incorporating Detailed Noise Modeling and Multi-Scale Feature Extraction*. PhD thesis, Caltech, 2006.
- [16] P. Sen. Estimates of the regression coefficient based on Kendall’s tau. *J Am Stat Assoc*, 63(324):1379–1389, 1968.
- [17] J. Tardós, J. Neira, P. Newman, and J. Leonard. Robust mapping and localization in indoor environments using sonar data. *Int J Robot Res*, 21(4):311–330, 2002.
- [18] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press, 2005.
- [19] J. Vanderpe, H. Van Brussel, and H. Xu. Exact dynamic map building for a mobile robot using geometrical primitives produced by a 2D range finder. In *Proc. ICRA*, pages 901–908, 1996.
- [20] L. Zhang and B. Ghosh. Line segment based map building and localization using 2D laser rangefinder. In *Proc. ICRA*, pages 2538–2543, 2000.