

Data-Association-Free Landmark-based SLAM

Yihao Zhang¹, Odin A. Severinsen¹, John J. Leonard¹, Luca Carlone¹, Kasra Khosoussi²

Abstract—We study landmark-based SLAM with unknown data association: our robot navigates in a completely unknown environment and has to simultaneously reason over its own trajectory, the positions of an unknown number of landmarks in the environment, and potential data associations between measurements and landmarks. This setup is interesting since: (i) it arises when recovering from data association failures or from SLAM with information-poor sensors, (ii) it sheds light on fundamental limits (and hardness) of landmark-based SLAM problems irrespective of the front-end data association method, and (iii) it generalizes existing approaches where data association is assumed to be known or partially known. We approach the problem by splitting it into an inner problem of estimating the trajectory, landmark positions and data associations and an outer problem of estimating the number of landmarks. Our approach creates useful and novel connections with existing techniques from discrete-continuous optimization (e.g., k-means clustering), which has the potential to trigger novel research. We demonstrate the proposed approaches in extensive simulations and on real datasets and show that the proposed techniques outperform typical data association baselines and are even competitive against an “oracle” baseline which has access to the number of landmarks and an initial guess for each landmark.

I. INTRODUCTION

This paper studies landmark-based Simultaneous Localization and Mapping (SLAM) and investigates the challenging case where data associations between measurements and landmarks are not given. Consider a robot moving in an unknown environment and trying to estimate its own trajectory using odometry and detections of unknown landmarks. When data associations are given, this leads to a well-studied and now easy-to-solve setup, *i.e.*, standard landmark-based SLAM [6]. However, the problem becomes extremely challenging when data associations are unknown, *i.e.*, when the robot does not have prior knowledge about data association and does not even know how many unique landmarks it has observed. This case arises in practice when the detections do not provide rich information (e.g., appearance, semantics) to perform data association, or after failures resulting from incorrect data associations (where one needs to recover an estimate by reasoning over potentially wrong associations).

The problem of data association in SLAM has been extensively studied [4], see Section II for a review. However, most

papers consider an *incremental* setup where data associations are greedily established at each time step (e.g., one has to associate newly acquired measurements to the existing set of landmarks). The setup where all associations have to be established at the same time (*i.e.*, a *batch* optimization setup) has not received the same amount of attention. The batch setup is relevant for several reasons. First, the incremental setup makes hard data-association decisions at each step, but data associations that look plausible at a given time may turn out to be wrong after more measurements are collected; on the other hand, a batch setup simultaneously reasons over the entire history of measurements and looks for an optimal association that is consistent across all measurements. Second, typical incremental data association techniques are based on probabilistic tests that rely on the measurement covariances; however, the covariances are often hand-tuned and may not be reliable, hence calling for an alternative framework that is less sensitive to the covariance tuning. Finally, incremental data association may fail after an incorrect data association (or may accumulate large drift if correct data associations are discarded); therefore, it is desirable to have a batch approach that revisits all data associations and corrects mistakes, in order to recover after SLAM failures.

Contribution. This paper investigates batch-SLAM with unknown data associations. Our first contribution (Section III) is to provide a formulation for the inner problem, which simultaneously optimizes for the robot poses, landmark positions, and data associations assuming the number of landmarks is given. We then present a formulation to estimate the number of landmarks.

Our second contribution (Section IV) presents algorithms to solve the optimization problems arising in our formulations. Our algorithm for the inner problem is based on an alternation scheme, which alternates the estimation of the robot poses and landmark positions with the choice of data associations. The algorithm is not guaranteed to compute optimal solutions due to the hardness and combinatorial nature of the optimization problems it attempts to solve, but it works well in practice.

Finally, we demonstrate the proposed approaches in extensive simulations and on real datasets (Section V), and show that the proposed techniques outperform typical data association baselines and are even competitive against an “oracle” baseline which has access to the number of landmarks and an initial guess for each landmark.

II. RELATED WORK

The data association problem of associating measurements to states of interest has been extensively studied in the target tracking literature [4] and later computer vision and robotics.

*This work was partially funded by ARL DCIST CRA W911NF-17-2-0181, Lincoln Laboratory’s Resilient Perception in Degraded Environments program, ONR Neuroautonomy MURI grant N00014-19-1-2571, and ARPA-E award DE-AR0001218 under the DIFFERENTIATE Program.

¹Y. Zhang, O. Severinsen, J. Leonard, and L. Carlone are with the Massachusetts Institute of Technology, United States, {yihaozh, odinase, jleonard, lcarlone}@mit.edu

²K. Khosoussi is with the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia, Kasra.Khosoussi@data61.csiro.au

Data associations can be made in the front-end via feature descriptors [24] and recently via deep learning [17]. However, these front-end methods are typically sensor-dependent and require additional information (e.g., landmark appearance) from the sensor other than pure positional measurements of the landmarks. We are interested in the case where such information is not available and data associations must be inferred from the positional measurements. Standard robust optimization methods based on M-estimators used in SLAM (e.g., [29], [1], [31]) are not directly applicable to such scenarios in an efficient manner (i.e., they would require adding robust losses for all possible associations). Therefore, in the following we focus our attention on techniques designed for data association in landmark-based SLAM.

Maximum likelihood data association [4] is a classical approach where one first rejects unlikely associations for which the Mahalanobis distance between measurements and their predictions is greater than a statistical threshold. This results in a number of *individually compatible* association candidates for each measurement, among which one picks the one with the maximum likelihood. In its most basic form, this approach processes measurements in a sequential manner and makes a decision for each measurement independently. Alternatively, one can solve a linear assignment problem between a batch of measurements and landmarks in which the cost of each measurement-landmark assignment is set to negative log-likelihood if they are *individually compatible*, and to infinity otherwise [4].

Joint compatibility branch and bound (JCBB) [25] is a popular technique for solving the data association problem *simultaneously* for a batch of measurements. This method leverages correlations between the estimated positions of landmarks and seeks the largest set of *jointly compatible* associations. JCBB uses the branch and bound technique to prune the exponentially large search space of the problem.

The combined constraint data association (CCDA) method [3] forms the *correspondence graph* whose nodes correspond to individually compatible associations and undirected edges correspond to a pair of mutually compatible associations. CCDA then aims to find the largest set of individually compatible associations that are also *pairwise compatible* by equivalently searching for maximum clique in the correspondence graph. A similar idea is proposed in [18] in which the authors propose to form an edge-weighted version of the correspondence graph and search for the densest clique.

Another approach is CLEAR [11], which identifies landmarks by clustering raw measurements using a spectral graph clustering-type method. This is done by forming an association graph whose nodes correspond to measurements and whose edges specify potential matches based on a compatibility criterion. Then, CLEAR uses spectral graph-theoretic methods to first estimate the number of distinct landmarks K , and then to rectify the association graph to obtain a disjoint union of K cliques, each of which corresponds to a unique landmark.

Unlike our approach, the abovementioned works aim to find an explicit set of associations *prior* to estimating robot trajectory and the map in each time step. An alternative

approach is proposed in [7] where the authors treat associations as latent variables and propose a scheme based on expectation-maximization (EM) for estimating poses and 3D points. This is done by maximizing the expected value of the log-likelihood function where the expectation is taken over all possible associations. Similar approaches are proposed in [5], [21] where, in addition to geometric information, semantic information is also incorporated. Unlike [7], [5], [21], in this work we aim to find the *most likely* estimate for robot poses, landmark positions, and *associations*, as well as the number of landmarks given the collected measurements.

In a separate line of work [22], the map is modeled as a random finite set (a finite-set-valued random variable), inferring the number of landmarks from the collected measurements and circumventing the need for conventional data association.

Our formulation for the inner problem of simultaneously estimating the robot poses, landmark positions and data associations has a strong connection to the max-mixture formulation [26] in the sense that landmark choices for a measurement can be modeled as different modes in the mixture. A max-mixture-type method for data association for semantic SLAM is proposed in [8] where additional semantic label and confidence information of landmarks is available to assist data association. Our inner problem is a special case of the general formulation recently presented in the discrete-continuous smoothing and mapping [9]. In Section V, we compare our method with the alternating minimization approach proposed in [9] (see Oracle baseline).

III. PROBLEM FORMULATIONS:

DATA-ASSOCIATION-FREE SLAM (DAF-SLAM)

We break the Data-Association-Free SLAM (DAF-SLAM) problem into an inner problem and an outer problem. The inner problem solves for the robot poses, landmark positions and data associations assuming the number of landmarks is known. The outer problem estimates the number of landmarks assuming one can solve the inner problem. We first introduce our notations. Let us denote the robot trajectory as $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \text{SE}(d)^n$, where $\mathbf{x}_i \in \text{SE}(d)$ is the (unknown) i -th pose of the robot and $\text{SE}(d)$ is the group of poses in dimension $d = 2$ or 3 ; in particular, $\mathbf{x}_i = (\mathbf{R}_i, \mathbf{t}_i)$, and \mathbf{R}_i and \mathbf{t}_i are the rotation and translation components of the pose. We denote the (unknown) position of the j -th landmark as $\mathbf{y}_j \in \mathbb{R}^d$; we then store all landmark position estimates as columns of a matrix \mathbf{y} , the *landmark matrix*. Finally, we denote the (given) measurement of a landmark position as $\bar{\mathbf{z}}_k \in \mathbb{R}^d$, for $k \in [m] \triangleq \{1, \dots, m\}$, and use the standard generative model for the measurements:

$$\bar{\mathbf{z}}_k = \mathbf{R}_{i_k}^\top \mathbf{y}_{j_k} - \mathbf{t}_{i_k} + \epsilon_k \quad (1)$$

where a landmark \mathbf{y}_{j_k} is observed in the local frame of robot pose i_k , up to zero-mean Gaussian measurement noise ϵ_k . Since we know the timestamp at which a measurement is taken, we can associate the measurement with the corresponding robot pose index i_k (e.g., the k -th measurement taken from the second robot pose has $i_k = 2$). However we do not know what is the index j_k of the landmark measured

by \bar{z}_k (*i.e.*, we have unknown data associations). With this notation, we can now introduce our problem formulations.

A. DAF-SLAM Inner Problem

We provide a problem formulation for the inner problem of DAF-SLAM to estimate the robot poses and landmark positions as well as the data associations given the number of landmarks.

Problem 1 (DAF-SLAM inner problem).

$$\min_{\substack{\mathbf{x} \in \text{SE}(d)^n \\ \mathbf{y} \in \mathbb{R}^{d \times K}}} f_{\text{odom}}(\mathbf{x}) + \sum_{k=1}^m \min_{j_k \in [K]} \|\mathbf{R}_{i_k}^\top \mathbf{y}_{j_k} - \mathbf{t}_{i_k} - \bar{z}_k\|_{\Sigma}^2 \quad (2)$$

where K is the number of landmarks (*i.e.*, the number of columns of the matrix \mathbf{y}). The first term $f_{\text{odom}}(\mathbf{x})$ in (2) is the standard odometric cost (which is only a function of the robot trajectory)¹; the second term is similar to the one arising in landmark-based SLAM, but also includes an optimization over the unknown data association j_k . Each summand in this term can be viewed as a max-mixture by considering each association choice as a mode in the mixture [26].

Proposition 1 (Monotonicity²). *Let us define the function $f_{\text{slam}}(\mathbf{x}, \mathbf{y}, K) = f_{\text{odom}}(\mathbf{x}) + \sum_{k=1}^m \min_{j_k \in [K]} \|\mathbf{R}_{i_k}^\top \mathbf{y}_{j_k} - \mathbf{t}_{i_k} - \bar{z}_k\|_{\Sigma}^2$, which is the objective function in (2). After minimizing over \mathbf{x} and \mathbf{y} , it remains a function of K , namely $f_{\text{slam}}^*(K)$. Then, the function $f_{\text{slam}}^*(K)$ is monotonically decreasing in K and attains a minimum $f_{\text{slam}}^*(K) = 0$ when $K = m$.*

Remark 2 (Hardness). *Problem 1 is NP-hard even when the optimal poses are known: in the next section we show that it indeed is equivalent to K -means clustering when the optimal poses and the number of landmarks are known, which is known to be NP-hard (see Proposition 5).*

B. DAF-SLAM Outer Problem

We now present our formulation to the outer problem of estimating the number of landmarks K in Problem 2.

Problem 2 (DAF-SLAM as simplest explanation). *Using the definitions in Proposition 1, we write an outer minimization over K to estimate the number of landmarks (K):*

$$\min_{K \in \mathbb{N}} \left(\overbrace{\min_{\mathbf{x} \in \text{SE}(d)^n, \mathbf{y} \in \mathbb{R}^{d \times K}} f_{\text{slam}}(\mathbf{x}, \mathbf{y}, K)}^{f_{\text{slam}}^*(K)} \right) + \beta K \quad (3)$$

Note that $f_{\text{slam}}^*(K)$ decreases as we allocate more landmarks (*i.e.*, as K increases) according to Proposition 1. On the contrary, βK increases as more landmarks are allocated. The formulation in (3) thus minimizes a cost that trades-off the SLAM error with the cost β ($\beta > 0$) of using more landmarks. It essentially follows Occam’s razor: we try to explain the measurements with the minimal number of landmarks possible. Problem 2 together with Problem 1 simultaneously decides the SLAM estimates \mathbf{x} and \mathbf{y} , while

¹We refer the reader to standard references [6], [28] for the expression of f_{odom} , which is irrelevant for the rest of the discussion in this paper.

²The proofs are provided in the supplementary [32].

also deciding the number of landmarks (*i.e.*, the number of columns in \mathbf{y}) and the data associations (minimization over j_k in (2)). The following remarks give some insights on the shape of the cost function in (3) and on how to tune β .

Remark 3 (Shape of “U”). *Proposition 1 states that the objective function in (3) is the sum of a monotonically decreasing function of K (*i.e.*, $f_{\text{slam}}^*(K)$) and a monotonically increasing linear function of K (*i.e.*, βK). This suggests that for sufficiently large values of β , the overall cost function $f(K) = f_{\text{slam}}^*(K) + \beta K$, roughly has a “U” shape, which starts at a large value for $K = 1$ (this is the case where all measurements are associated to a single landmark, in which case the cost $f_{\text{slam}}^*(K)$ is typically large), then reaches a minimum for a suitable value of K , and then increases again until reaching a value $f(m) = \beta m$ for $K = m$. This insight will be useful to design our solver in Section IV-A, whose outer loop searches over potential values of K .*

Remark 4 (Interpretation of β). *The parameter β describes the cost of using one more landmark to explain the measurements. Choosing a large β leads to underestimating the number of landmarks, while a small β leads to overestimating it. More specifically, consider having a set of measurements \mathcal{M}_j for landmark j . The corresponding residual errors for these measurements will be $r_j \doteq \sum_{k \in \mathcal{M}_j} \|\mathbf{R}_{i_k}^\top \mathbf{y}_j - \mathbf{t}_{i_k} - \bar{z}_k\|_{\Sigma}^2$. Intuitively, when r_j exceeds β , the optimization may prefer “breaking” the measurements across two or more landmarks (*i.e.*, increasing K). We explain how we set β in Section V-B.*

IV. ALGORITHMS FOR DAF-SLAM

A. Algorithms for Problem 1

This section proposes a heuristic algorithm to solve Problem 1. Before introducing the algorithm, we provide fundamental insights that connect Problem 1 with Euclidean clustering and standard landmark-based SLAM.

Proposition 5 (Clustering²). *Given K and for fixed \mathbf{x} , Problem 1 becomes an Euclidean clustering problem assuming the measurement noise covariances Σ are isotropic. In this clustering problem, the landmarks \mathbf{y} are the cluster centers and the data associations j_k are equivalent to the cluster assignments.*

Proposition 6 (SLAM²). *Given K and the data associations j_k for each measurement $\bar{z}_k \in \mathbb{R}^d$, $k = 1, \dots, m$, Problem 1 reduces to standard landmark-based SLAM.*

Propositions 5-6 suggest a simple alternation scheme, where we alternate solving for the data associations $j_k, \forall k = 1, \dots, m$ (given an initial guess for \mathbf{x}), with solving \mathbf{x}, \mathbf{y} (given the data associations). The former can be solved using standard clustering approaches (*e.g.*, Lloyd’s algorithm for k-means clustering with k-means++ initialization [2]) according to Proposition 5; the latter can use standard landmark-based SLAM solvers according to Proposition 6. We name our algorithm K-SLAM due to the combined use of k-means and SLAM solvers.

The pseudo-code of K-SLAM (inner problem solver) is in Algorithm 1. Algorithm 1 performs a fixed number of iterations, where each iteration alternates between two steps:

- **Data association update** (lines 2-4) computes the data associations j_k and landmark positions \mathbf{y} given \mathbf{x} ; this reduces to clustering per Proposition 5, which we solve (locally) using k-means and k-means++ [2];
- **Variable update** (line 5) computes the robot poses \mathbf{x} and landmark positions \mathbf{y} given the data associations j_k and the result of \mathbf{y} from the first step as initial guess; this reduces to standard landmark-based SLAM per Proposition 6, hence it can be readily solved with both local [13], [10] and global solvers [28].

Algorithm 1: K-SLAM inner solver.

Input: Odometry and corresponding initial guess \mathbf{x} ;
landmark measurements $\tilde{\mathbf{z}}_k, k = 1, \dots, m$; number of landmarks K ; maxIterations

Output: Estimate of robot poses \mathbf{x} ; landmark positions \mathbf{y} ;
SLAM objective value f_{slam} .

```

1 for iter = 1, 2, ..., maxIterations do
    /* Project each measurement to the world frame
       based on current estimate of  $\mathbf{x}$  */
2      $\hat{\mathbf{z}}_k = \mathbf{R}_{i_k} \tilde{\mathbf{z}}_k + \mathbf{t}_{i_k}, \text{ for } k \in [m];$ 
3     initialize  $\mathbf{y}$  by k-means++( $\{\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_m\}, K$ );
4      $\{\mathbf{y}, j_1, \dots, j_m\} = \text{k-means}(\{\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_m\}, K)$ ;
5      $\{\mathbf{x}, \mathbf{y}\} = \text{SLAM}(\mathbf{x}, \mathbf{y}, \{j_1, \dots, j_m\})$ ;
6 end
7 return  $\{\mathbf{x}, \mathbf{y}, f_{\text{slam}}(\mathbf{x}, \mathbf{y}, K)\}$ .
```

B. Algorithms for Problem 2

In consistency with Remark 3, we also empirically observe $f(K) = f_{\text{slam}}^*(K) + \beta K$ has roughly a U-shape. Therefore, we simply use a gridding method to solve for K^* . Naive gridding, which probes the objective function at every K ($K = 1, \dots, m$), requires $m \cdot \text{maxIterations}$ invocations of the landmark-based SLAM solver. For practical problems where the number of measurements m is in the thousands and considering that in our experiments $\text{maxIterations} = 15$, this is too expensive. We thus introduce a multi-resolution gridding technique, whose pseudo-code is in Algorithm 2.

In Algorithm 2, we iteratively grid the interval where the optimal K could lie into smaller intervals until we find the exact value of K in a resolution-one grid. Specifically, we start with the interval $[1, m]$ (line 1) and uniformly evaluate $f_{\text{slam}}^*(K) + \beta K$ at N_K values (N_K is set to 11), K_1, K_2, \dots, K_{N_K} , within this interval and including both end points (lines 3-5). We find the optimal K_{n^*} out of these N_k values (line 6), shrink the interval down to $[K_{n^*-1}, K_{n^*+1}]$ (line 10), and repeat the process until the resolution is one (line 2). This multi-resolution strategy makes approximately $N_k \log_{\frac{N_k}{2}}(m)$ evaluations of $f_{\text{slam}}^*(K)$ compared to m evaluations in the naive gridding. For $m = 1000$ and $N_K = 11$, this is 45 versus 1000 evaluations.

V. EXPERIMENTS

This section shows that the proposed K-SLAM algorithm computes accurate estimates across multiple simulated and real datasets, and it is competitive against two baselines, one of which has access to the true number of landmarks. We first

Algorithm 2: K-SLAM solver.

Input: Odometry and corresponding initial guess \mathbf{x} ;
landmark measurements $\tilde{\mathbf{z}}_k, k = 1, \dots, m$;
maxIterations; parameter β ; number of K 's to probe at each level N_K .

Output: Estimate of robot poses \mathbf{x} ; landmark number (K) and positions \mathbf{y} .

```

1  $\{K_1, K_2, \dots, K_{N_K}\} = \text{UniformDiv}([1, m], N_K)$ ;
2 while  $|K_1 - K_2| > 1$  do
3     for  $K = K_1, K_2, \dots, K_{N_K}$  do
4          $\{\mathbf{x}^{(K)}, \mathbf{y}^{(K)}, f_{\text{slam}}^{(K)}\} =$ 
           Alg. 1( $\mathbf{x}, \{\tilde{\mathbf{z}}_1, \dots, \tilde{\mathbf{z}}_m\}, K, \text{maxIterations}$ );
5          $f^{(K)} = f_{\text{slam}}^{(K)} + \beta K$ ;
           /* Save best result */
6         if  $f^{(K)} \leq f^*$  then
7              $f^* \leftarrow f^{(K)}$ ;
8              $\{\mathbf{x}^*, \mathbf{y}^*\} \leftarrow \{\mathbf{x}^{(K)}, \mathbf{y}^{(K)}\}$ ;
9              $K_{n^*} \leftarrow K$ ;
10     $\{K_1, \dots, K_{N_K}\} = \text{UniformDiv}([K_{n^*-1}, K_{n^*+1}], N_K)$ ;
11 return  $\{\mathbf{x}^*, \mathbf{y}^*, K_{n^*}\}$ .
```

introduce the experimental setup and datasets (Section V-A) and the baseline techniques (Section V-B). We then evaluate the techniques in terms of accuracy (Section V-C). An ablation study is provided in the supplementary [32].

A. Datasets and Setup

Datasets. We evaluate K-SLAM along with baseline techniques on (i) two fully synthetic “grid” datasets (2D and 3D), (ii) two semi-real datasets (Intel 2D and Garage 3D) where the trajectories and odometry measurements are real, but the landmarks are synthetic, and (iii) a real dataset (Indoor). See the supplementary [32] for visualizations of the datasets.

For the two synthetic datasets, we first generate a grid-like ground-truth trajectory and uniformly sample ground-truth landmark positions around the trajectory. We then add Gaussian noise to create odometry and landmark measurements following eq. (1). Each landmark is observed by a fixed number of poses that are closest to it in the ground-truth positions.

For the two semi-real datasets (Intel and Garage), we use two real datasets for pose graph optimization [14], [15], which include both odometry measurements and loop closures. We obtain a proxy for the ground-truth trajectory by optimizing the original pose graph with loop closures. Then we remove the loop closures from the dataset, and add landmarks following the same procedure used for the synthetic datasets.

The real dataset (Indoor) is a benchmarking dataset available online [20]. In this dataset, a robot navigates in a real office environment and observes a number of AprilTags [27] as landmarks along the trajectory (we only use the ID of the AprilTags for evaluating different methods).

Setup. In our synthetic and semi-real experiments, we sweep through different parameter values that define the map (e.g., odometry noise level, number of landmarks, and landmark measurement noise level). Twenty Monte Carlo simulations are performed for each set of parameters where

TABLE I: Dataset default parameters. Values in bold are varied in our experiments. n/a stands for a dense covariance matrix not representable by a single value.

Dataset	2D Grid	3D Grid	Intel	Garage	Indoor
# of poses	500	216	942	1661	571
# of landmarks	100	43	94	166	14
# of lm. measurements	1000	430	1880	3320	233
odometry translation std.	0.05	0.05	0.045	1	0.1
odometry rotation std.	0.005	0.005	0.014	n/a	0.1
lm. measurement std.	0.05	0.05	0.05	0.05	0.1

we vary noise realizations and landmark placement. The default map parameter values can be found in Table I.

B. Compared Techniques and Baselines

Proposed Technique. K-SLAM uses the Levenberg-Marquardt optimizer in GTSAM [10] for landmark-based SLAM. We set $\text{maxIterations} = 15$. A rule of thumb is to set β to the maximum expected residuals for the measurements of a single landmark. Assuming that a landmark j is observed $n = |\mathcal{M}_j|$ times and that the measurements are affected by Gaussian noise, r_j (defined in Remark 4) at the ground truth is distributed according to a χ^2 distribution, hence $\mathbb{P}(r_j \leq \gamma^2) = p$, where γ^2 is the quantile of the $\chi^2(dn)$ distribution with $d \times n$ degrees of freedom and lower tail probability equal to p . As a heuristic, we set $\beta = \gamma^2$ where γ can be computed for a given choice of tail probability (in our experiments $p = 0.997$) and for an average number of landmark observations n . Assuming having knowledge of n , the values of β used are 41.72 (Grid 2D), 55.64 (Grid 3D), 68.94 (Intel), and 94.47 (Garage). In practice, a rough estimate of K can be obtained by clustering the projected measurements (line 2 in Algorithm 1) and applying the elbow method [30] to estimate the number of clusters (*i.e.*, K). An estimate of n can be then derived from this estimate of K given the total number of landmark measurements which is known. The accuracy of K-SLAM is not sensitive to the exact value of β as seen in our sensitivity analysis in Table II and in the supplementary [32].

Oracle Baseline. The Oracle baseline is an alternation scheme that has access to the true number of landmarks and an initial guess for each landmark. The baseline alternates two steps. In the first step, for each measurement \tilde{z}_k it finds the landmark having a position estimate \mathbf{y}_j that minimizes $\|\mathbf{R}_{i_k}^\top \mathbf{y}_j - \mathbf{t}_{i_k} - \tilde{z}_k\|$. After making associations for all the measurements, it solves the SLAM problem with the estimated associations. This process repeats until convergence or reaching a user-specified number of iterations. This method is similar to the technique recently proposed in [9]. Note that this strong baseline assumes extra knowledge (*i.e.*, number of landmarks, initial guess for the landmarks) that is not available to the other techniques and in real scenarios. The initial guess for the landmarks is computed by projecting the first measurement of each landmark (*i.e.*, measurement from the first pose that observes the landmark) into the world frame (in the same way as line 2 in Algorithm 1).

ML Baseline. The second baseline is the popular Maximum Likelihood data association used in an incremental formulation of SLAM, where odometry is propagated se-

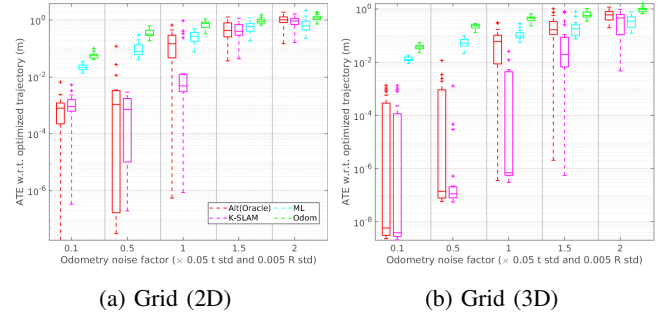


Fig. 2: Absolute trajectory error (ATE) for increasing odometry noise level on the synthetic grid datasets.

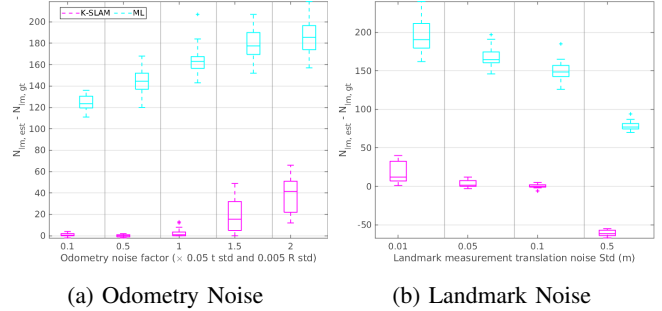


Fig. 3: Difference between the estimated and the ground truth number of landmarks for K-SLAM and the ML baseline on the 2D grid dataset with different noise levels.

quentially and newly acquired measurements are incrementally associated to the most likely landmark estimate; see [4]. At each time step, for each measurement, it selects a set of admissible landmarks by performing a χ^2 -test between the landmark estimates \mathbf{y}_j and the measurement \mathbf{z}_k . The measurement is associated to the most likely landmark, while each landmark is only associated to one of the newly acquired measurements. The measurements that are not associated to any landmark are initialized as new landmarks. After data associations are established, batch optimization with the Levenberg-Marquardt optimizer [16], [19] is performed to compute the SLAM estimate at each time step. Additional implementation details are in the supplementary [32].

Odom Baseline. This baseline only computes an odometric estimate of the robot trajectory by concatenating odometry measurements without using landmark measurements. We report the accuracy of this baseline as a worst-case reference.

C. Results: Accuracy

We measure the accuracy of the techniques by computing the Absolute Trajectory Error (ATE) of the trajectory estimated by each technique with respect to an optimized trajectory obtained via standard landmark-based SLAM with ground-truth data associations. Therefore, if the algorithms estimate all data associations correctly, the ATE will be zero. Unlike the normal ATE with respect to the ground-truth trajectory, this metric aims to capture data association correctness. In the following, we evaluate the ATE for varying levels of odometry noise, number of landmarks, and landmark measurement noise. Then we provide summary statistics for the real dataset (Indoor) in Table II.

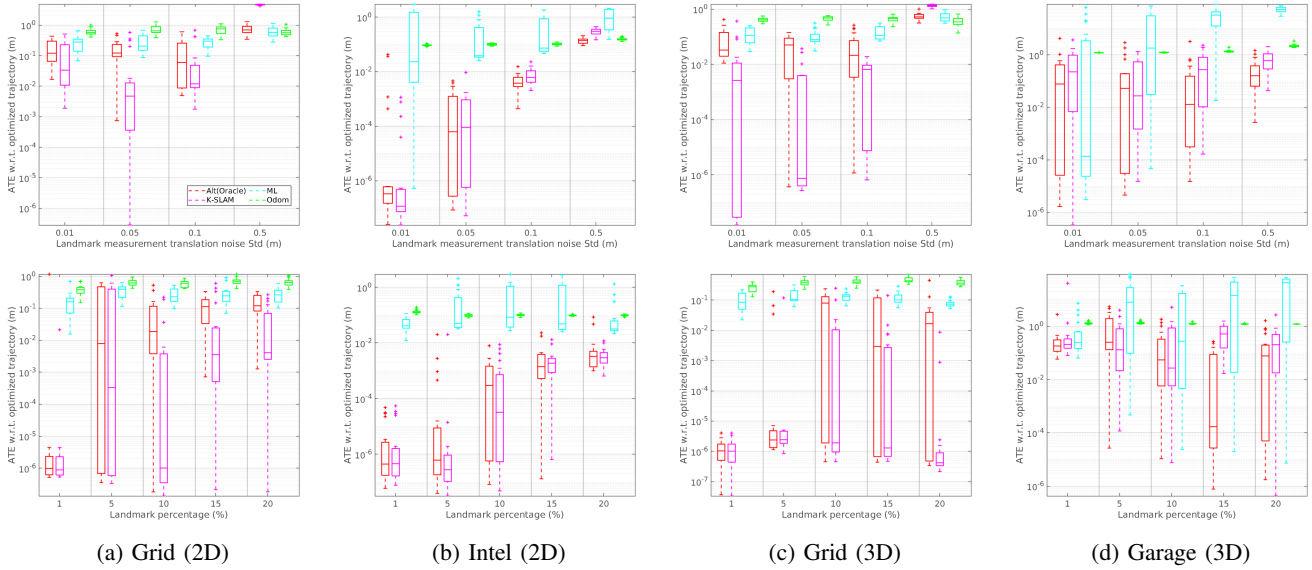


Fig. 4: Absolute trajectory error (ATE) for increasing landmark measurement noise and increasing number of landmarks (shown as a percentage of the number of poses).

Influence of odometry noise. Fig. 2 shows the ATE for increasing odometry noise in our synthetic datasets. For low and moderate odometry noise levels, K-SLAM generally outperforms other techniques. When the odometry noise is high, ML starts to outperform Oracle and K-SLAM. The reason is that the poor odometry causes the initial projected measurements of a landmark (line 2 in Algorithm 1) to be too far to cluster and associate. As an incremental method, ML reduces the odometry drift incrementally, alleviating this problem. Fig. 3a shows the estimated number of landmarks for K-SLAM and ML (Oracle is given the ground-truth number). We observe that ML tends to largely overestimate the number of landmarks. In contrast, K-SLAM provides fairly accurate estimation of the number at low to medium odometry noise. It only overestimates the number of landmarks at high odometry noise. We note that underestimating the number of landmarks in SLAM can lead to catastrophic failures due to incorrect associations. Overestimation, however, generally has less severe consequences, causing the method to miss some of the correct associations and thus lose loop-closing information.

Influence of number of landmarks and landmark noise. Fig. 4 shows the ATE across the four datasets for increasing landmark measurement noise (top row) and increasing number of landmarks shown as a percentage of the number of poses (bottom row). We observe that our K-SLAM still outperforms all other techniques almost across the spectrum; it is only worse at high landmark noise where the raw odometry becomes the best among all the techniques (with the exception of the garage dataset). In this case, two factors come into play for K-SLAM. First, as mentioned earlier, our approach can underestimate the number of landmarks under high noise (see Fig. 3b). Second, high landmark noise makes solving the inner problem (Problem 1) more challenging in K-SLAM by reducing the inter-cluster distances of the projected measurements. Solving the inner problem sub-

TABLE II: ATE (m) and runtime (sec) on the Indoor dataset.

Metric	Odom	ML	Oracle	K-SLAM		
				$\beta = 10$	$\beta = 80$	$\beta = 240$
ATE	0.415	1.34	0.076	0.054	0.046	0.148
Runtime	n/a	7.57	1.25	18.80	18.75	18.79

optimally further affects the outer estimation of the number of landmarks, causing more performance loss. The Garage dataset is particularly challenging as its odometry translation covariance is specified to be large despite the odometry being accurate. The ML baseline is the most vulnerable to such mismatched covariances due to its reliance on them. Number of landmarks is a relatively weak factor for data association. There is a slight trend of having less accurate data association as more landmarks are in the map, which is expected.

Performance on the real dataset. Table II reports the results on the real dataset. We observe that although being slow, K-SLAM still dominates the others. Additionally, the performance is not sensitive to β . The highest accuracy is obtained over a wide range of values for $\beta \in [20, 140]$. Discussions about the runtime are in the supplementary [32].

VI. CONCLUSION

In this paper we investigated a challenging batch SLAM problem where the goal is to simultaneously estimate the poses of the robot, the locations of unknown landmarks, and the data associations between measurements and landmarks. We proposed problem formulations and algorithms for the resulting data-association-free SLAM problem. The algorithms were tested on a mix of synthetic and real datasets and perform better or on par with strong baselines that have access to privileged information. While these initial findings are encouraging, our algorithm does show limitations such as degraded performance at high noise and slow runtime. It would be desirable to better understand the behavior of $f_{\text{slam}}^*(K)$ for K estimation and improve the robustness of association and clustering at high noise.

REFERENCES

- [1] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard. Robust map optimization using dynamic covariance scaling. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013.
- [2] D. Arthur and S. Vassilvitskii. K-Means++: The advantages of careful seeding. pages 1027–1035, 2007.
- [3] Tim Bailey. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, University of Sydney, 2002.
- [4] Yaakov Bar-Shalom, Thomas E Fortmann, and Peter G Cable. Tracking and data association, 1990.
- [5] S.L. Bowman, N. Atanasov, K. Daniilidis, and G.J. Pappas. Probabilistic data association for semantic SLAM. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1722–1729, 2017.
- [6] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016.
- [7] F. Dellaert, S.M. Seitz, C.E. Thorpe, and S. Thrun. Structure from motion without correspondence. In *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, volume 2, pages 557–564, 2000.
- [8] Kevin J Doherty, David P Baxter, Edward Schneeweiss, and John J Leonard. Probabilistic data association via mixture models for robust semantic SLAM. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1098–1104. IEEE, 2020.
- [9] Kevin J Doherty, Ziqi Lu, Kurran Singh, and John J Leonard. Discrete-continuous smoothing and mapping. *arXiv preprint arXiv:2204.11936*, 2022.
- [10] F. Dellaert et al. Georgia Tech Smoothing And Mapping (GTSAM). <https://gtsam.org/>, 2019.
- [11] Kaveh Fathian, Kasra Khosoussi, Yulun Tian, Parker Lusk, and Jonathan P How. Clear: A consistent lifting, embedding, and alignment rectification algorithm for multiview data association. *IEEE Transactions on Robotics*, 36(6):1686–1703, 2020.
- [12] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- [13] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2011.
- [14] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. g 2 o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*, pages 3607–3613. IEEE, 2011.
- [15] Rainer Kümmerle, Giorgio Grisetti, Hauke Strasdat, Kurt Konolige, and Wolfram Burgard. G2o datasets. https://github.com/OpenSLAM-org/openslam_g2o/tree/master/data, 2011.
- [16] Kenneth Levenberg. A method for the solution of certain nonlinear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.
- [17] Kejie Li, Daniel DeTone, Yu Fan Steven Chen, Minh Vo, Ian Reid, Hamid Rezaatfighi, Chris Sweeney, Julian Straub, and Richard Newcombe. ODAM: Object detection, association, and mapping using posed RGB video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5998–6008, 2021.
- [18] Parker C Lusk, Kaveh Fathian, and Jonathan P How. CLIPPER: A Graph-Theoretic Framework for Robust Data Association. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2021.
- [19] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [20] MathWorks. MathWorks indoor dataset. <https://www.mathworks.com/help/nav/ug/landmark-slam-using-apriltag-markers.html>.
- [21] Elad Michael, Tyler Summers, Tony A Wood, Chris Manzie, and Iman Shames. Probabilistic data association for semantic SLAM at scale. *arXiv preprint arXiv:2202.12802*, 2022.
- [22] John Mullane, Ba-Ngu Vo, Martin D Adams, and Ba-Tuong Vo. A random-finite-set approach to Bayesian SLAM. *IEEE transactions on robotics*, 27(2):268–282, 2011.
- [23] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- [24] Raul Mur-Artal and Juan D Tardós. ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras. *IEEE transactions on robotics*, 33(5):1255–1262, 2017.
- [25] José Neira and Juan D Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on robotics and automation*, 17(6):890–897, 2001.
- [26] E. Olson and P. Agarwal. Inference on networks of mixtures for robust robot mapping. In *Robotics: Science and Systems (RSS)*, July 2012.
- [27] Edwin Olson. Apriltag: A robust and flexible visual fiducial system. In *2011 IEEE International Conference on Robotics and Automation*, pages 3400–3407, 2011.
- [28] David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. SE-Sync: A certifiably correct algorithm for synchronization over the special Euclidean group. *The International Journal of Robotics Research*, 38(2-3):95–125, 2019.
- [29] N. Sünderhauf and P. Protzel. Switchable constraints for robust pose graph SLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [30] Robert Thorndike. Who belongs in the family? *Psychometrika*, 18(4):267–276, 1953.
- [31] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone. Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robotics and Automation Letters*, 5(2):1127–1134, 2020. [arXiv preprint arXiv:1909.08605](https://arxiv.org/abs/1909.08605) (with supplemental material), (pdf).
- [32] Yihao Zhang, Odin A. Severinsen, John J. Leonard, Luca Carlone, and Kasra Khosoussi. Data-association-free landmark-based SLAM. <https://arxiv.org/abs/2302.13264>.