Dense map inference with user-defined priors:
from priorlets to scan eigenvariations

Paloma de la Puente   Andrea Censi

Abstract—When mapping is formulated in a Bayesian framework, the need of specifying a prior for the environment arises naturally. However, so far, the use of a particular structure prior has been coupled to working with a particular representation. We describe a system that supports inference with multiple priors while keeping the same dense representation. The priors are rigorously described by the user in a domain-specific language. Even though we work very close to the measurement space, we are able to represent structure constraints with the same expressivity of methods based on geometric primitives.

I. INTRODUCTION

The Simultaneous Localization and Mapping (SLAM) problem is usually formulated in a Bayesian framework [1]. This paper concerns the use of prior distributions for the map: how to rigorously specify them and how to create an inference engine that works with multiple user-defined priors.

To see where exactly the prior comes in the problem, let us introduce some notation. Let q be the robot pose, let m be a variable representing the map, and let z be the measurements (including odometry and exteroceptive sensors), which follow the known sensor model $p(z|q, m)$. SLAM can be formulated as the problem of estimating $p(q, m|z)$, the joint distribution of pose and map conditioned to the measurements. We focus on the case of mapping with dense sensors and maps; if the map consists of landmarks, then most of the following remarks are not relevant. To make things concrete, we describe the formulation that uses the Rao-blackwellization technique [2], where one approximates the target distribution as $p(q, m|z) \simeq p(q|z)p(m|q, z)$, thereby factorizing SLAM in two subproblems: estimating the pose of the robot given the measurements ($p(q|z)$), and mapping with known poses ($p(m|q, z)$). Let us focus on the latter. Given the sensor model, we can compute the posterior using Bayes' theorem: $p(m|q, z) \propto p(z|q, m)p(m)$.

Therefore, if we want to compute the posterior distribution of the map m given the observations, we have to know the prior $p(m)$. We remark that, had we formulated SLAM as a maximum-likelihood problem (find m that maximizes $p(z|q, m)$), the knowledge of $p(m)$ would not be strictly necessary. That, however, would only work for finite-dimensional problems. In fact, if the underlying map is an arbitrary surface, the maximum-likelihood problem is ill posed, because the solution is any curve that perfectly interpolates the readings.

Paloma de la Puente is with the Intelligent Control Group, Universidad Politecnica de Madrid, 28006 Madrid, Spain. Comunidad de Madrid PhD grant. This work was partially funded by the Spanish Ministry of Science and Technology, DPI2007-66846-c02-01. E-mail: paloma.delapuente@upm.es

Andrea Censi is with the department of Control & Dynamical Systems, Division of Engineering and Applied Science, California Institute of Technology, Pasadena, MC 107-81, 91125, CA, USA. E-mail: andrea@cds.caltech.edu
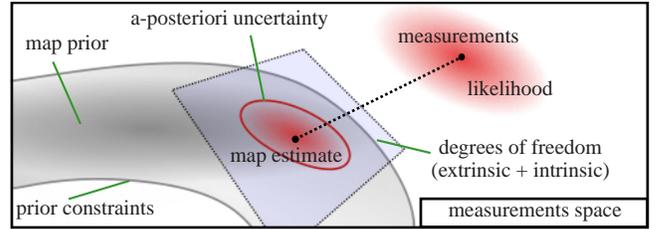
Figure 1. Computing the posterior distribution of the map in the measurement space has a clear geometric interpretation. The prior $p(m)$ defines a (thin) surface in the measurement space. The initial measurements define a thick ellipse of uncertainty that gets projected and constrained to the prior surface.

To obtain a more reasonable solution, we always need some kind of regularization, which is the prior. Therefore, we conclude that, to make the SLAM problem with dense sensors and maps well posed, we have to specify a prior $p(m)$.

Other than to make the mathematical formulation correct, the knowledge of the prior helps in reducing the uncertainty of the estimate. For example, constraints such as collinearity are very powerful in reducing the map uncertainty. In general, any assumption about the environment that the user can provide helps in making the filter more efficient. Yet, incorporating generic prior information in filters has never been done, and that can be attributed to the representation used.

For example, let us consider SLAM methods that represent maps using occupancy/evidence grids. Firstly, the grid resolution introduces some kind of spatial regularization, and makes it impossible to represent precise geometric primitives such as line segments. The other limitation is that each cell is assumed to be independent: this makes it impossible to effectively use the prior information because geometric constraints between different parts of the environment result in long-range correlation of cells occupancy.

The alternative to occupancy grids is using a map composed of geometric primitives (segments, circles, splines, etc.). In that case, the prior is implicit in the representation: representing a map by segments automatically gives non-segments map a zero prior. Using geometric primitives presents two major advantages: they provide explicit information about the geometrical nature of the environment, and the resulting maps are much more compact. With proper bookkeeping, the correlation between different parts of the environment can be precisely represented. However, they lack in flexibility. For example, in most realistic environments —except perhaps completely engineered factory floors— there will be parts of the environment that cannot be described by the prior. Moreover, often one wishes to impose "soft constraints": for example, rather than imposing that all walls are exact line segments, probably a better prior is that they are likely straight, or that they are of a bounded variation from straight; all these

details should be figured out by the user. This flexibility cannot be accommodated by existing features-based methods.

A. Contribution

We began this work by asking the question of whether it is possible to decouple the concept of prior from a particular representation. Instead of the prior being hidden in the representation, can it be made completely explicit and under the direct control of the user? Can we have an inference engine that works with multiple priors?

In Section II, we start by defining a new representation. A range-finder provides an array of numbers measuring the distance to the obstacles. We augment that by associating to each measurement the corresponding surface normal. This gives us the flexibility of occupancy grids with the precision of geometric primitives. Just like large environments can be represented by a collection of patches, one can represent any environment by a collection of augmented scans; still, in this paper, we focus on processing the data from a single scan.

In our system, the prior is entirely provided by the user, who describes the structure constraints and the model likelihood as a function of readings and normals, in a particular domain-specific language (Fig. 2 on the following page).

The inference engine, described in Section III, takes two inputs: the noisy raw distance readings from a laser sensor, and the user-specified prior. The output is the posterior distribution for the local map, represented as a Gaussian distribution on the space of readings and normals. This is obtained by a two-step process. In the first step, we solve a nonlinear optimization problem to obtain the mode of the distribution.

In the second step, described in Section IV, we use the knowledge of the structure constraints to shrink the measurements covariance, by projecting it onto the allowed submanifold. Fig 1 shows a geometric interpretation of the process.

We believe our method presents a novel approach for data segmentation and preprocessing in a flexible manner, being able to reduce the uncertainty of noisy measurements and providing information about the environment's geometrical nature. Section V includes some experiments showing how it works. In our opinion the proposed framework has a lot of potential for future research: it may also be very useful for scan-matching techniques with laser data and it is very promising for overall map optimization as well, helping us build better models of the physical world with a unique system in different situations.

B. Related work

So far, prior information about the environment has been used explicitly only in features-based SLAM methods. For example, Chong and Kleeman [3] employ collinearity constraints to enhance the state estimation with a Julier-Uhlman Kalman Filter. Rodríguez-Losada *et al.* [4] alleviate the inconsistency due to the linearization errors introduced by the Extended Kalman Filter by enforcing parallelism or orthogonality constraints. Nguyen *et al.* [5] apply orthogonality constraints to build accurate simplified plane based 3D maps. Beevers

and Huang [6] show that imposing a-priori known relative constraints also leads to consistency and efficiency improvements for particle filters. In all these works, a particular geometrical model for representation is used, and they only support equality constraints.

We know of no previous work using a dense representation and allowing the use of a different priors provided by the user. Modelling the scans as a Gaussian process [7] does allow to impose a prior distribution, corresponding to a smoothness constraint, but cannot capture structured priors such as polygonal environments.

C. Notation

Let $q = (t, \theta) \in \mathrm{SE}(2)$ be the robot pose. Assume, without loss of generality, that the range-sensor frame coincides with the robot frame. The sensor model for the range-sensor measurements $\tilde{\rho} = \{\tilde{\rho}_i\}_{i=1}^n$ is defined by $\tilde{\rho}_i = \rho_i + \epsilon_i$, where ρ_i is the true distance to the obstacle, and ϵ_i is additive Gaussian noise with covariance $\Sigma_{ij} = \mathrm{cov}\{\epsilon_i, \epsilon_j\}$, not necessarily diagonal. The true distance to the obstacle can be written as

$$\rho_i = r(\boldsymbol{m}, \langle \boldsymbol{t}, \theta + \phi_i \rangle), \tag{1}$$

where the angle ϕ_i is the direction of each reading in the scan, and the function $r : \mathcal{M} \times \mathrm{SE}(2) \to \mathbb{R}_+ \cup \{\infty\}$ is the "ray-tracing" function that returns the distance to the closest obstacle from a certain pose. The function r depends on the map $\boldsymbol{m} \in \mathcal{M}$. For now, we do not specify anything about $\boldsymbol{m}$, other than it represents the underlying map of the environment.

D. Problem statement

Formally, we divide the problem of approximating $p(\boldsymbol{m}|\boldsymbol{z}, \boldsymbol{q})$, where $\boldsymbol{z} = \tilde{\rho}$, in two sub-problems. First, we solve the maximum-a-posteriori problem to obtain the mode of the distribution. Because $p(\boldsymbol{m}|\boldsymbol{z}, \boldsymbol{q}) \propto p(\boldsymbol{z}|\boldsymbol{m}, \boldsymbol{q})p(\boldsymbol{m})$, this can be posed as the following.

Problem 1: Find $\boldsymbol{m}$ that maximizes

$$\log p(\boldsymbol{z}|\boldsymbol{m}, \boldsymbol{q}) + \log p(\boldsymbol{m}).$$

The first term is simply the measurements likelihood; the second term is the map prior. After we have found the mode of the distribution, we obtain a Gaussian approximation to $p(\boldsymbol{m}|\boldsymbol{z}, \boldsymbol{q})$ by projecting the initial covariance onto the prior constraints (Fig. 1).

This process is conducted in a representation very close to the measurement space, as described in the next section.

II. DEFINING MAP PRIORS WITH PRIORLETS

The environment prior is specified by the user in a domain-specific language; a representative set of user-supplied prior definition files is shown in Fig. 2. Providing a flexible way to parametrize environment priors posed two challenges. The first mathematical challenge is choosing a unified representation that allows for the description of a multitude of priors. The second challenge is that this representation must also be user-friendly.

```
name: Polygonal prior                                             1
order: 2                                                          2
max_curvature: 0                                                  3
p_1 = [cos( φ_1 ); sin( φ_1 )] * ρ_1; # define cartesian coords   4
p_2 = [cos( φ_2 ); sin( φ_2 )] * ρ_2; # as shortcuts             5
priorlet same_region:                                            6
    α_1 == α_2                                                    7
    (p_2 - p_1)' * [cos( α_1 ); sin( α_1 )] == 0                 8
```

(a) User-supplied definition for polygonal prior

```
name: Rectangular prior                                          1
specializes: Polygonal prior                                     2
priorlet different_region:                                       3
   ( α_2 == α_1 - pi/2 ) || ( α_2 == α_1 + pi/2 )                4
```

(b) User-supplied definition for rectangular prior

```
name: Rectangular prior (relaxed)                                1
specializes: Polygonal prior                                     2
priorlet different_region:                                       3
    tolerance = 3; # 3deg tolerance                              4
    cos( α_2 - α_1 ) <= cos(deg2rad(90+tolerance))               5
    -cos( α_2 - α_1 ) <= -cos(deg2rad(90-tolerance))             6
```

(c) User-supplied definition for relaxed rectangular prior

```
name: Rectangular prior (relaxed - alternative)                  1
specializes: Polygonal prior                                     2
priorlet different_region:                                       3
    model_likelihood cos( α_2 - α_1 )^2                          4
```

(d) User-supplied definition for alternative relaxed rectangular prior

```
name: Circular prior                                             1
order: 3                                                         2
max_curvature: 10 # min radius = 0.1 m                           3
# two oriented points define a circle. This is the radius.       4
r12 = sin(( α_2 - α_1 )/2) / norm(p_1 - p_2);                    5
r23 = sin(( α_3 - α_2 )/2) / norm(p_3 - p_2);                    6
r13 = sin(( α_3 - α_1 )/2) / norm(p_3 - p_1);                    7
priorlet same_region:                                            8
    r12 == r23 # the three oriented points                       9
    r23 == r13 # lie on the same circle                         10
```

(e) User-supplied definition for circular prior

```
name: Circular prior (with prior on radius)                      1
specializes: Circular prior                                      2
priorlet same_region: # it is likely that the radius is around 2.0   3
    model_likelihood (r13 - 2.0)^2                               4
```

(f) Circular prior, with prior information for the radius

```
name: Splines prior                                              1
order: 2                                                         2
max_curvature: 10                                                3
priorlet same_region:                                            4
    model_likelihood ( α_2 - α_1 )^2                             5
```

(g) User-supplied definition for spline prior

Figure 2. The environment prior is specified by the user with a domain-specific language. These are examples of actual source code interpreted by the inference engine (apart from some omissions in the interest of clarity). Using UNICODE, the special variables alpha_i, rho_i, phi_i can also be typed with Greek letters; this was inspired by Sun's Fortress language.

A. Representation: distances ρ, normals α, topology $\mathcal{T}$

For what concerns the representation, our solution is parametrizing $p(m)$ by three finite-dimensional quantities.

True distance to the obstacle ρ: The quantities $\{\rho_i\}_{i=1}^n$ were already defined as part of the sensor model in equation (1). They represent a zeroth-order approximation of the environment shape.

Surface normals α: The surface normals represent a first-order approximation of the environment shape, and will play an important role in defining the priors. The surface normal α_i can be written similarly to ρ_i as a function of the derivative of the ray-tracing function[1]. We define $x \triangleq (\rho, \alpha)$ and we write compactly:

$$x = (\rho, \alpha) = r(m, q). \qquad (2)$$

Environment topology $\mathcal{T}$: We assume that the environment is partitioned into *surfaces,* and each *surface* is partitioned into one or more *regions.* Thus, for each two consecutive readings in the scan, there are three possible topology cases:

1) They belong to the same *surface* and the same *region.*
2) They belong to the same *surface,* but different *regions.*
3) They belong to different *surfaces.*

Having this fine distinction allows to precisely define the prior's constraints. To keep track of the topology information, we define a variable $\mathcal{T} = \{\mathcal{T}_k\}_{k=1}^n$, where each $\mathcal{T}_k \in \{\texttt{sameRegion}, \texttt{differentRegion}, \texttt{differentSurface}\}$ describes the relation between a pair of consecutive points.

B. Expressing priors as functions of $\rho, \alpha, \mathcal{T}$

We can express the prior as a function of the readings ρ, normals α, and topology $\mathcal{T}$ instead of as a function of the infinite-dimensional map m. Assuming that it is possible, we rewrite Problem 1 as follows.

Problem 2: Find $\rho, \alpha, \mathcal{T}$ that maximize

$$\log p(\tilde{\rho}|\rho) + \log p(\rho, \alpha, \mathcal{T}).$$

Now we are dealing with a finite-dimensional optimization problem: the infinite-dimensional map "m" has disappeared from the formalization. The limitation is that we can only define shape priors by their 0th (ρ) and 1st order (α) Taylor expansions. In the same spirit, we could use successive derivatives (curvature, and so on); nevertheless, we found that this parametrization has good expressivity. This does not mean that we are limited to piece-wise linear shapes; in fact, we can define shapes such as circles (Fig. 2e) and splines (Fig. 2g).

C. Expressing $p(\rho, \alpha, \mathcal{T})$ with local constraints and energies

Now we have fixed the representation, but we still have to solve the challenge of allowing the user to specify a prior in an intuitive way. It is clear that we can express almost any shape using a function $p(\rho, \alpha, \mathcal{T})$. In theory, we could ask the user to provide a symbolic expression for $p(\rho, \alpha, \mathcal{T})$. This, however, would be burdensome: assuming, for example, that there are 180 readings in a scan, the user would need to

---

[1]An explicit expression for the normal α_i as a function of the ray-tracing function r is $\alpha_i = \pi/2 + \arctan\left(\frac{\partial}{\partial \phi_i} r\left(m, \langle t, \theta + \phi_i \rangle\right)\right)$, but we are not going to need it in this paper.

provide a symbolic expression with 540 variables. Moreover, that expression would have to be changed if the number of readings changed.

Our observation was that one can define interesting priors by describing *local* constraints between consecutive points. For example, if the environment prior is polygonal, we want to impose that nearby points have the same normals if they belong to the same region: $\alpha_1 = \alpha_2 = \cdots = \alpha_n$. This can be expressed compactly by saying that $\alpha_i = \alpha_{i+1}$ if points i and $i+1$ belong to the same region (compare Fig. 2a, line 7). In addition to these, we need constraints on ρ_i to ensure that the points are aligned (Fig. 2a, line 8).

In the case of a rectangular prior, we have the additional constraint that $(\alpha_i - \alpha_{i+1}) = k\frac{\pi}{2}$ if the two points do *not* belong to the same region (compare Fig. 2b, line 4). Similarly, one can defined different relaxations for a rectangular prior (Fig. 2c-2d). We will not describe in detail the interpretation of all the expressions in Fig. 2, but they all correspond to simple geometric constraints.

Certain priors cannot be specified by considering only two successive points. For example, it takes three consecutive points to describe a circular prior (compare Fig. 2e), because it takes three points to define a circle. The *order* of a prior is the number of consecutive points needed for describing it.

D. Formal definitions of priorlets

We call *priorlet* a set of local constraints plus energies imposed on n consecutive points in the environment.

Definition 1: A *priorlet* of order n is a tuple $\langle F, G, H \rangle$ described by three sets of functions $F = \{f_k\}$, $G = \{g_k\}$, $H = \{h_k\}$. The arguments of all these functions are n couples of (distance, normal angle) and they all return a scalar. The functions $\{f_k\}$ represent equality constraints, the functions $\{g_k\}$ represent inequality constraints, and the functions $\{h_k\}$ represent "energies" (negative log-likelihoods).

The semantics of a priorlet is the specification of a small part of a larger optimization problem:

$$\min_{\rho_{1:n}, \alpha_{1:n}} \quad \cdots + \sum_k h_k((\rho_1, \alpha_1), \cdots, (\rho_n, \alpha_n)) + \cdots ,$$
$$\text{subject to} \quad f_k((\rho_1, \alpha_1), \cdots, (\rho_n, \alpha_n)) = 0,$$
$$g_k((\rho_1, \alpha_1), \cdots, (\rho_n, \alpha_n)) \leq 0.$$

The philosophy is very close to that of *factor graphs* [8]; the formalization, however, does not match perfectly because usually factor graphs do not include constraints.

Definition 2: A *user-defined environment prior* is a collection of three priorlets: a "`same_region`" priorlet, a "`different_region`" priorlet, and a "`different_surface`" priorlet, describing the constraints/energies for neighbouring points for the three topology cases.

Recall that the variable $\mathcal{T}$ specifies the environment partition in regions and surfaces. Given a particular choice of $\mathcal{T}$, we know which priorlet to apply to each couple (or triplet) of consecutive points. Therefore, we can define three functions $h_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha})$, $f_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha})$, $g_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha})$. These represent, respectively, the cumulative effect of all the energies, and the stacked equalities and inequalities given by the application of the priorlets to each neighbourhood of points (we do not write

them explicitly to avoid drowning in a sea of indices). We can rewrite Problem 2 as follows.

Problem 3: Find $\mathcal{T}, \boldsymbol{\rho}, \boldsymbol{\alpha}$ as the solution of the problem:

$$\max_{\mathcal{T}, \boldsymbol{\rho}, \boldsymbol{\alpha}} \quad \log p(\tilde{\boldsymbol{\rho}}|\boldsymbol{\rho}) + h_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha}),$$
$$\text{subject to} \quad f_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha}) = \mathbf{0},$$
$$g_{\mathcal{T}}(\boldsymbol{\rho}, \boldsymbol{\alpha}) \leq \mathbf{0}.$$

E. A domain-specific language for priorlets

We have given a formal description of priorlets that might appear overly complicated. In practice, the process of specifying a prior is intuitive, using a domain-specific language whose syntax we believe easy to understand even without a formal definition.

The user must minimally specify a `name`, and the `order` of the prior. Then she specifies the three priorlets (a `same_region`, `different_region`, `different_surface`), by specifying equalities (`==`) and inequalities (`<=`) over the predefined variables `rho_i`, `alpha_i`, `phi_i`, with $1 \leq i \leq$ order. Using UNICODE input, the variables can also be indicated with Greek letters. At any point in the file, other variables can be introduced using "`=`" (Fig. 2a, line 4). The syntax for the expressions is the one used by MATLAB/Octave. A "`||`" operand is supported for specifying a logical or (Fig. 2b, line 4). The model likelihood (the h function) is introduced by the keyword `model_likelihood`. The user can subclass existing priors using the `specializes` keyword; for example, the rectangular prior specializes the polygonal prior (Fig. 2b, line 2). Finally, we let the user specify an explicit `max_curvature` parameter that is used in the inference process.

III. INFERENCE WITH GENERIC PRIORS

Our goal has been to build an inference engine that works for arbitrary user-specified priors. Of course, we are doomed to be less efficient than an optimization method designed for a particular prior; however, we believe there is value in showing a completely general approach. In this section, we briefly recall the standard constrained-optimization methods that we use, we show how additional constraints can be added to the problem, and finally we discuss the two-level optimization procedure.

A. Homotopy methods

The idea of homotopy methods [9] consists in solving the constrained optimization problem by solving a *sequence* of *unconstrained* optimization problems. The *penalty function* method is useful for dealing with equalities or inequalities. Suppose the minimization problem to solve is

$$\min_{\boldsymbol{x}} h(\boldsymbol{x}), \qquad \text{subject to } f(\boldsymbol{x}) = 0,$$

and assume that we do not know a feasible point. We then consider a sequence of unconstrained minimization problems, where we add to the objective function a penalty function representing the distance from the feasible set:

$$\min_{\boldsymbol{x}} h(\boldsymbol{x}) + \lambda f(\boldsymbol{x})^2.$$

Similarly, the penalty function for an inequality $g(\boldsymbol{x}) \leq 0$ would be $\lambda \max\{0, g(\boldsymbol{x})\}^2$. As $\lambda \to \infty$, the solution of the

unconstrained problem tends to the solution of the constrained one. Therefore, we can solve the constrained problem by solving a sequence of unconstrained optimization problems, starting from $\lambda = 0$ and progressively raising it. One can prove proper convergence under appropriate conditions [9].

The *log-barrier* method is useful for dealing with inequalities. Suppose we have to solve the problem

$$\min_{\boldsymbol{x}} \ h(\boldsymbol{x}), \quad \text{subject to } \boldsymbol{x} \leq \overline{\boldsymbol{x}},$$

and assume that we start from a feasible point $\boldsymbol{x}_0 \leq \overline{\boldsymbol{x}}$. Then we solve the sequence of unconstrained optimization problems

$$\min_{\boldsymbol{x}} \ h(\boldsymbol{x}) - \frac{1}{\mu} \sum_i \log \left((\overline{x}_i - x_i) \right).$$

The $\log$ term represents a "barrier" that goes to infinity near the bounds. As $\mu \to \infty$, the solution of the unconstrained problem tends to the solution of the constrained one.

B. Additional details

Outliers: We expect that the prior supplied by the user describes most of the environment, but there will always be points that are clearly outside the prior, caused, for example, by clutter in the environment. Therefore, we define another optimization variable, the set INMODEL of points that do respect the prior. Suppose that the likelihood of a point being described by the prior model is $\beta \in (0, 1]$, and, for simplicity, that each point is independent. Then the log-likelihood component $\log(p(\text{INMODEL}))$ can be represented in the cost function by a term $\gamma|\text{INMODEL}|$, with $\gamma = \log(\beta/(1-\beta))$ and $|\text{INMODEL}|$ indicating the number of points.

Upper and lower bounds on $\boldsymbol{\rho}$, $\boldsymbol{\alpha}$: It is possible to derive upper and lower bounds for the variables $\boldsymbol{\alpha}$ and $\boldsymbol{\rho}$. For $\boldsymbol{\rho}$, bounds are obtained by using the initial covariance information. During the optimization, each ρ_i is allowed to vary at most $4\sigma_i$ from the initial estimate $\tilde{\rho}_i$. Because of that, outliers and clutter produce constraints that are impossible to satisfy, and eventually those points are removed from the INMODEL set. As for the normals, it is possible to derive bounds for α_i based on the allowed variation of $\rho_{i-1}, \rho_i, \rho_{i+1}$ and the knowledge of the maximum curvature in the environment.

C. Optimization overview

We rewrite again the form of the optimization problem, with the new variable INMODEL and the bounds on the state.

Problem 4: Find $\boldsymbol{\mathcal{T}}$, INMODEL, $\boldsymbol{x}$ as the solution of:

$$\max_{\boldsymbol{\mathcal{T}}, \boldsymbol{x} \in \text{INMODEL}} \log p(\boldsymbol{x}|\boldsymbol{\rho}) + h_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}) + \gamma|\text{INMODEL}|,$$

$$\text{subject to} \quad f_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}) = 0, \tag{3}$$

$$g_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}) \leq 0, \tag{4}$$

$$\underline{\boldsymbol{x}} \leq \boldsymbol{x} \leq \overline{\boldsymbol{x}}. \tag{5}$$

We have to optimize over discrete and continuous variables. The discrete variables are the set INMODEL and the topology $\boldsymbol{\mathcal{T}}$. The continuous variable is $\boldsymbol{x} = (\boldsymbol{\rho}, \boldsymbol{\alpha})$. We solve the problem using two nested levels: the outer level (Algorithm 1) optimizes over INMODEL and $\boldsymbol{\mathcal{T}}$, while the inner level (Algorithm 2) optimizes over $\boldsymbol{x}$, given a particular choice of INMODEL and $\boldsymbol{\mathcal{T}}$. We describe the inner level first.

D. Inner loop: optimizing $\boldsymbol{x}$ given INMODEL, $\boldsymbol{\mathcal{T}}$

Algorithm 2 solves Problem 4 assuming that INMODEL, $\boldsymbol{\mathcal{T}}$ have been fixed. We apply a double homotopy transformation to find $\boldsymbol{x}$. We use a penalty function for constraints (3)-(4) and a log-barrier method for constraint (5). Using the log-barrier for the bounds ensures that those are always satisfied during each iteration. Instead, the constraints on the prior are satisfied only in the limit: we start from the measurements and eventually arrive to the surface defined by the prior (Fig. 1).

At each iteration, we take a Newton step with backtracking. All the necessary gradients and Hessians are computed in closed form using symbolic derivations from the user-specified constraints. Moreover, we "convexify" the Hessian if it is not positive-definite by setting negative eigenvalues to a small positive value (Algorithm 2, line 12); this turns the Newton method into gradient descent in the non-convex parts of the state space.

Note that Algorithm 2 might fail to return a feasible point; this will be interpreted by the outer level as a sign that the topology $\boldsymbol{\mathcal{T}}$ is wrong and must be relaxed.

Algorithm 1 Discrete Optimization of INMODEL, $\boldsymbol{\mathcal{T}}$

function $[\boldsymbol{x}, \boldsymbol{\mathcal{T}}]$ = map_optimization($\tilde{\boldsymbol{\rho}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\rho}}}$,prior):	1
% initialize by using all points, and the strictest topology	2
INMODEL = all; $\mathcal{T}_k$ = sameRegion;	3
while True:	4
$[\underline{\boldsymbol{x}}, \overline{\boldsymbol{x}}]$ = geometric_bounds($\boldsymbol{\mathcal{T}}, \tilde{\boldsymbol{\rho}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\rho}}}$)	5
% Estimate surface normals	6
$[\boldsymbol{\alpha}_0$, covalpha] = estimate_initial_alpha($\tilde{\boldsymbol{\rho}}, \boldsymbol{\Sigma}_{\tilde{\boldsymbol{\rho}}}, \boldsymbol{\mathcal{T}}$)	7
% Restrict optimization to the INMODEL set	8
$\boldsymbol{x}_0 = \{(\tilde{\boldsymbol{\rho}}, \alpha_0)\}$ **for** $i \in$ INMODEL	9
[feasible, $\boldsymbol{x}$, link_penalties] =	10
inner_optimization(prior,$\boldsymbol{x}_0$,cov0,$\boldsymbol{\mathcal{T}}$,$[\underline{\boldsymbol{x}}, \overline{\boldsymbol{x}}]$)	11
if feasible: **break**	12
% If not feasible, break the topology based on the penalties	13
$\boldsymbol{\mathcal{T}}$ = break_greedily($\boldsymbol{\mathcal{T}}$,link_penalties)	14
% Remove points in small regions from the INMODEL set	15
[INMODEL,$\boldsymbol{\mathcal{T}}$] = remove_lonely_points(INMODEL,$\boldsymbol{\mathcal{T}}$)	16
return $[\boldsymbol{x}, \boldsymbol{\mathcal{T}}]$	17

Algorithm 2 Continuous optimization of $\boldsymbol{x}$

[feasible,$\boldsymbol{x}$,link_penalties] = inner_optimization(prior,$\boldsymbol{x}_0$,$\boldsymbol{\Sigma}_{\boldsymbol{x}_0}$,$\boldsymbol{\mathcal{T}}$,$[\underline{\boldsymbol{x}}, \overline{\boldsymbol{x}}]$)l	
% Obtain functions from prior and topology	2
$f_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}), g_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}), h_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x})$ = prior_to_constraints(prior, $\boldsymbol{\mathcal{T}}$)	3
for $\lambda = \lambda_0$; $\lambda \leq \lambda_{\max}$; $\lambda = \lambda_{\text{mult}}\lambda$:	4
for $\mu = \mu_0$; $\mu \leq \mu_{\max}$; $\mu = \mu_{\text{mult}}\mu$:	5
% return if the point is feasible	6
if $f_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}) < \epsilon$: **return** [true, $\boldsymbol{x}$]	7
% compute gradient and Hessian of objective + penalties	8
$J(\boldsymbol{x}) = p(\boldsymbol{x}|\boldsymbol{x}_0, \boldsymbol{\Sigma}_{\boldsymbol{x}_0}) + h_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x})$	9
$+$ log_barrier($[\underline{\boldsymbol{x}}, \overline{\boldsymbol{x}}], \mu, \boldsymbol{x}$) $+ \lambda$ penalty($f_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x}), g_{\boldsymbol{\mathcal{T}}}(\boldsymbol{x})$)	10
% convexify the Hessian (do gradient descent if nonconvex)	11
H = convexify($\nabla^2_{\boldsymbol{x}} J(\boldsymbol{x})$)	12
newton_direction = $-$inv(H) $* \nabla_{\boldsymbol{x}} J(\boldsymbol{x})$	13
$\boldsymbol{x}$ = back_tracking($\boldsymbol{x}$, newton_direction)	14
% the problem is infeasible: compute the penalty for each link	15
link_penalties = compute_link_penalties($\boldsymbol{\mathcal{T}}, \boldsymbol{x}$)	16
return [false, null, link_penalties]	17

E. Outer loop: optimizing INMODEL, $\mathcal{T}$

Algorithm 1 optimizes over the set INMODEL and the topology $\mathcal{T}$. Solving this problem exactly has combinatorial complexity, as we would have to try each possible grouping of points into surfaces and regions. To obtain an approximate solution, we use a heuristic approach based on relaxation.

We initialize INMODEL to contain all the points, and $\mathcal{T}$ to result in the strictest set of constraints ($\mathcal{T}_i = $ sameRegion). Iteratively, we call the inner level to find a corresponding x. If Algorithm 2 finds a feasible x, we are done. Otherwise, we try to relax the problem. If the problem is infeasible, some of the prior constraints ($g_{\mathcal{T}}(x) \leq 0, f_{\mathcal{T}}(x) = 0$) are not respected and the corresponding penalty functions are non-zero. We check which couple of nearby points gave the most contribution to the penalty function, and we relax the topology (line 14). If the corresponding $\mathcal{T}_k$ was sameRegion, we set it to differentRegion; if it was differentRegion, we set it to differentSurface. We observed that this simple algorithm was effective in finding region and surface boundaries.

In addition, we check whether some regions are too small, and we remove the corresponding points from INMODEL (line 16). This is useful to deal with outliers.

IV. RECOVERING THE DEGREES OF FREEDOM

We have shown how to define generic priors (Section II) and how to perform inference with them (Section III). We have decoupled the environment prior from the environment representation: while the priors are most general, the representation is always the same. This approach certainly has its advantages in terms of generality and flexibility. However, we lose something with respect to a features-based approach. If we fit a circle to the environment, we implicitly state that 1) the points are constrained to lie on a circle *(constraints)*; and 2) the circle can change in radius and position (degrees of freedom). In this section, we show how we can perform a similar analysis even using augmented scans as the representation.

A. The geometric structure of the map space

In equation (1), we let the sensor model depend on the underlying true map "m", interpreted as an abstract infinite-dimensional quantity belonging to a certain set $\mathcal{M}$. In order to derive well-grounded results, we have to formalize some intuitive ideas about $\mathcal{M}$ (some of these are commented in more detail elsewhere [10]).

It is intuitive that, for each map $m \in \mathcal{M}$, there will be other elements in $\mathcal{M}$ that have the same *shape* but are rotated/translated to different *poses*. Thus, we can assume that all reasonable sets $\mathcal{M}$ are isomorphic to the product $\mathcal{S} \times \text{SE}(2)$, where $\mathcal{S}$ is called the *shape space*. Given this factorization, we can write an element $m \in \mathcal{M}$ as a couple $\langle S, p \rangle \in \mathcal{S} \times \text{SE}(2)$. This factorization is the basis of many works in the shape-space analysis [11], [12]. Based on that, we introduce a technical condition on the user-defined prior.

Definition 3: A prior $p(m)$ is *pose-independent* if it only depends on the map shape S but not on the map pose p:

$$p(m) = p(\langle S, p \rangle) = p(S).$$

Intuitively, this means that, if the prior allows a certain shape, then it must allow the same shape, rotated, with equal probability; or, equivalently, that observing the environment does not give any information on the robot pose in an external frame. We also state a simple lemma on the ray-tracing function.

Lemma 1: The observations do not change if robot and map are jointly roto translated: $r(\langle S, p \rangle, q) = r(\langle S, \delta \oplus p \rangle, \delta \oplus q)$.

B. Analyzing the degrees of freedom

Assume we have found a feasible solution x. By analyzing the constraints given by the prior, we can recover the degrees of freedom in the solution. More formally, we consider infinitesimal variations δx and we examine which ones are allowed by the prior. Recall that x contains both scan readings and surface normals, therefore δx belongs to $\mathbb{R}^{2n}$, where n is the number of readings. We first give the mathematical results and then we comment on the derivation.

Proposition 1: Suppose the prior is pose-independent (Definition 3). Then the space of the allowed variations δx to the solution can be factorized as ("$\sqcup$" indicates disjoint union):

$$\mathbb{R}^{2n} = \text{Constr} \sqcup \text{Free} = \text{Constr} \sqcup (\text{Intr} \sqcup \text{Extr}),$$

where the subspaces are defined (and computed) as follows:

$$\text{Free} \triangleq \ker \nabla_x f_{\mathcal{T}}(x), \quad (6)$$
$$\text{Constr} \triangleq \mathbb{R}^{2n} - \text{Free}, \quad (7)$$
$$\text{Extr} \triangleq \text{span}\{\nabla_q r\}, \quad (8)$$
$$\text{Intr} \triangleq \text{Free} - \text{Extr}. \quad (9)$$

The subspaces Free are the directions corresponding to the map variations allowed by the prior. The subspace Free is further divided in *intrinsic* (Intr) degrees of freedom, due to the uncertainty in map shape; and *extrinsic* (Extr) degrees of freedom, due to the uncertainty in map pose.

To explain the first division in the subspaces Constr and Free, we just need to consider the equality constraints in the prior, which are represented by the equation $f_{\mathcal{T}}(x) = 0$. This equation defines a hyper-surface inside $\mathbb{R}^{2n}$ where x is constrained to lie. The tangent plane to this surface is given by directions orthogonal to the gradient $\nabla_x f_{\mathcal{T}}$, and corresponds to the (infinitesimal) directions that are allowed by the prior. The subspace Constr is simply the complement of Free.

The further division of Free in intrinsic (Intr) and extrinsic (Extr) degrees of freedom is a more delicate topic. We have seen that the map m can be represented as a couple shape-pose $\langle S, p \rangle$. The subspace Extr identifies the variation in the readings due to the uncertainty in the pose p; or, more precisely, due to the uncertain pose between map and sensor. We can state the following result.

Proposition 2: If the prior is pose-independent, the subspace Extr $\triangleq \text{span}\{\nabla_q r\}$ is contained in Free.

Proof: Using (2), we write $x = r(m, q) = r(\langle S, p \rangle, q)$. If the prior does not depend on p, then $f_{\mathcal{T}}(x) = 0$ implies

$$f_{\mathcal{T}}(r(\langle S, \delta \oplus p \rangle, q)) = 0, \qquad \text{for all } \delta \in \text{SE}(2). \quad (10)$$

Given Lemma 1, we obtain that $f_{\mathcal{T}}(r(\langle S, p \rangle, \ominus \delta \oplus q)) = 0$, for all $\delta \in \text{SE}(2)$. This means that

$$f_{\mathcal{T}}(r(\langle S, p \rangle, q)) = 0, \qquad \text{for all } q \in \text{SE}(2).$$
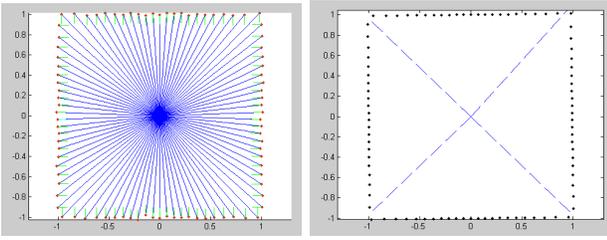
Figure 3. (a): Noisy simulated scan. (b): Corrected measurements and topology $\mathcal{T}$ (blue edges for different regions)

Intuitively, this says that the updated readings still respect the prior no matter where the robot is placed in the environment. If a function ($f_{\mathcal{T}}$) is constant with respect to an argument (q), the derivative with respect to that argument is 0. In our case, using the chain rule, we obtain:

$$\nabla_q f_{\mathcal{T}} = \nabla_x f_{\mathcal{T}} \cdot \nabla_q r = \mathbf{0}.$$

Therefore, $\nabla_q r$ is always orthogonal to $\nabla_x f_{\mathcal{T}}$; that is, span $\{\nabla_q r\} \subset \ker \nabla_x f_{\mathcal{T}} = $ Free. ∎

Extr $= $ span $\{\nabla_q r\}$ are the possible variations in the measurements due to the sensor movement. By contrast, the directions in Intr are due to the variation in the map shape S, and correspond to the intuitive notion of the degrees of freedom in the structure. To compute them, we use equations (6), (8), (9). Note that we used the concept of map factorization in $m = \langle S, p \rangle$ only as a theoretical device in deriving the results. In practice, we do not need to know anything about such abstract representation; the only quantities we have to compute are $\nabla_q r$ and $\nabla_x f_{\mathcal{T}}$, which lie in the very concrete measurement space. The procedure is completely automatic and allows to recover the degrees of freedom for any prior.

C. Covariance shrinking

Other than for visualization purposes, we can use the degrees of freedom knowledge for computing the posterior uncertainty of the estimate. Assume that the covariance of the initial estimate x_0 was Σ_{x_0}. If the prior has only constraints and not energies (i.e., there is no term $h_{\mathcal{T}}(x)$), we can obtain the posterior covariance Σ_x simply by projecting Σ_{x_0} onto the subspace Free. Let the matrix P_{Free} be a projector onto Free. Then the posterior covariance estimate is $\Sigma_x = P_{\text{Free}} \Sigma_{x_0} P_{\text{Free}}^T$. If there is a term $h_{\mathcal{T}}(x)$, we have to account for the further reduction of uncertainty. Treating it as an additional observation, we obtain that $\Sigma_x = P_{\text{Free}}(\Sigma_{x_0}^{-1} + \nabla^2 h_{\mathcal{T}}^{-1})^{-1} P_{\text{Free}}^T$. Similarly, one can recover the contribution to Σ_x due to extrinsic or intrinsic uncertainty by projecting onto Extr or Intr. This process has a solid geometric intuition; see also, for example, Chapter 3 of Paul Newman's thesis [13]. The reader should note that this linearized analysis has the usual limitations [14].

For priors with many constraints, the rank of Σ_x is very low. Thus, it is better to represent it by its non-null eigensystem, which can be interpreted as the allowed scan eigenvariations $\{\langle v^m, \sigma_m \rangle\}_{m=1}^{\dim(\text{Free})}$, each representing a direction $v^m = \langle \rho^m, \alpha^m \rangle$ and corresponding uncertainty σ_m in that direction.
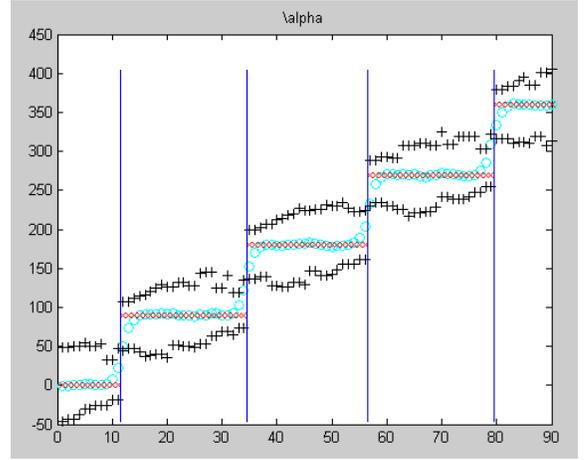

Figure 4. Corrected orientation angles (red), ground truth (green), initial values (cyan) and extracted topology $\mathcal{T}$ (blue edges for different regions)

V. EXPERIMENTS

We have conducted experiments with both synthetic and real data so as to validate our approach.

Fig. 3 shows an example test case, with a simulated scan from a square environment, using the rectangular prior. The original noisy simulated data is depicted in Fig 3(a), whereas Fig 3(b) presents the corrected measurements and the proper division of the scan into different regions. Fig 4 shows in red the orientation angles for all the readings corrected by applying our method, with the ground-truth represented in green and the initial estimates in cyan. The black x-marks indicate the bounds. The solution gets so close to the ground-truth that they can hardly be distinguished in the plot, with an average error of $0.26°$ for several tests and the walls being well aligned. We obtain similar results with a variety of other simulated environments.

The whole process described in Section III can be seen in action with real data from a Hokuyo laser sensor in Fig. 5. Even if most of the environment is polygonal, the regions of the polygonal surfaces are interrupted by random clutter and outliers (Fig. 5a). The first part of the relaxation introduces several breaks around outliers (Fig. 5b) producing a very fragmented topology. Then we remove the clutter from INMODEL and we can return to the simple correct topology for the rest of the points (Fig.5c).

Regarding the degrees of freedom extraction, the system recognizes well, for instance, that a circular environment has only one intrinsic degree of freedom (the radius of the circle) (Fig. 6a). A rectangular environment has two degrees of freedom (Fig. 6b) using the rectangular prior (Fig. 2b), but 5 if we use the polygonal prior (Fig. 2a), because the walls orientation is not constrained. Our system recognizes well the degrees of freedom in more complicated situations. For example, an environment with two circles has three degrees of freedom (the radii and the distance between the centers); however, they quickly become hard to visualize. Moreover, in the figures we plot only the variation of the readings because the variations of the normals are hard to visualize as well.

Fig. 6 also shows an example of covariance shrinkage with a

(a) Topology $\mathcal{T}$ after the first iteration   (b) Intermediate iteration   (c) Final topology
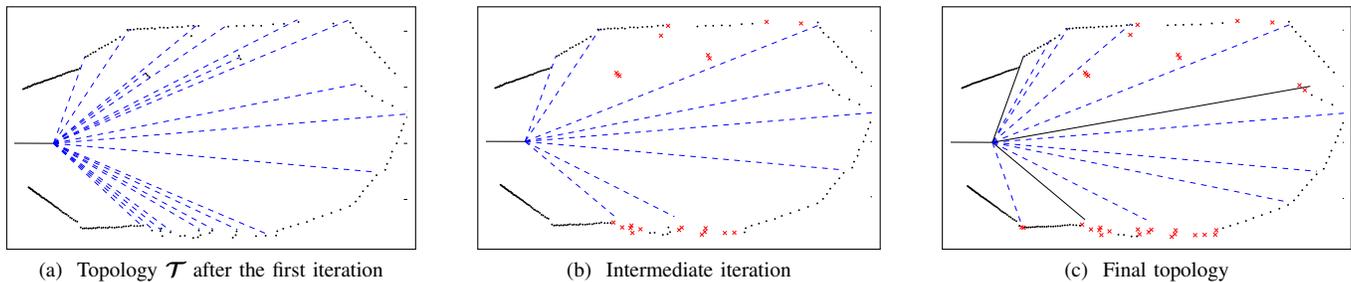
Figure 5. The outer level optimization (Algorithm 1) works on the discrete variables, deciding which sensor readings can be described by the prior (variable INMODEL) and the division in regions/surfaces (variable $\mathcal{T}$). The pictures show the evolution of the topology. Black lines indicate boundaries between surfaces; blue dashed lines indicate borders between regions. Red crosses indicate readings outside the INMODEL set (outliers). (a): Some decisions on the division in regions/surfaces can be taken based on the geometric constraints and the knowledge of the maximum curvature, specified in the prior. (b): The rest of the algorithm guesses where the region/surfaces boundaries are based on a greedy relaxation algorithm. Clutter and outliers tend to be isolated in small regions that are later removed. (c): The final result is feasible according to the prior; the outliers have been removed from the set INMODEL.

(a) Eigenvariation for circular environment with circular prior (Fig. 2e)   (b) Eigenvariations for rectangular environment with rectangular prior (Fig. 2b)   (c) A-posteriori readings covariance   (d) Readings standard deviations
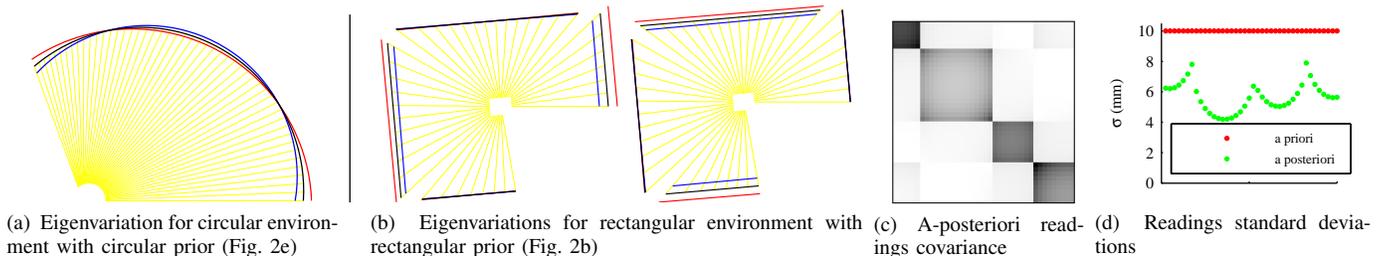
Figure 6. After we have found the solution to the optimization problem, the inference engine uses the knowledge of the prior for extracting the intrinsic degrees of freedom (*scan eigenvariations*) and for *shrinking* the covariance by projecting it onto the Free subspace. (a): For example, the inference engine can recognize that a circular environment has one allowed scan eigenvariation. (b): In the case of a rectangular environment and rectangular prior (Fig. 2b), we find 2 allowed scan eigenvariations. These can be interpreted as the variations of width and height of the environment. (c): We can shrink the a-priori readings covariance by projecting it onto the constraints. In this case, we assume that the a-priori covariance (not shown) has slight correlation between consecutive readings. The a-posteriori covariance has very low rank, and distant readings become correlated because of the structure. (d): The shrinking can be visualized by plotting the diagonal elements of a-priori and a-posteriori covariance.

rectangular environment. We assume that the initial covariance of $\tilde{\rho}$ is band-diagonal with slight correlation across neighbours. After the projection, the posterior covariance (Fig. 6c) correlates readings corresponding to the same surface or region. As shown in Fig. 6d, there is a dramatic uncertainty reduction.

VI. CONCLUSIONS AND FUTURE WORK

We have shown how to build an inference engine that can use different priors with the same representation. The priors are defined by the user in a domain-specific language. The problem of approximating the map posterior is turned into a constrained optimization problem, and a covariance projection over the unconstrained directions. We consider the main merit of this approach in the fact that it is possible to reason about structured priors (such as polygonal, rectangular, circular, etc.) using a unified representation. Still, one can recover the "structure" information under the form of scan eigenvariations, using the degrees-of-freedom analysis. Therefore, this approach brings together the feature space and the measurement space, providing the synthesis of the former, with the precision and flexibility of the latter.

As part of future work, we plan to improve the greedy Algorithm 1 by introducing backtracking. We are also interested in testing how preprocessing different sensor data with our method may help scan matching standard techniques. Finally, we are working on the integration of this algorithm into complete SLAM methods, by using the reduced degrees of freedom for global map optimization.

REFERENCES

[1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.

[2] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with Rao-Blackwellized particle filters," *Trans. of Robotics*, vol. 23, pp. 34–46, Feb. 2007.

[3] K. Chong and L.Kleeman, "Sonar based map building for a mobile robot," in *IEEE Conf. on Robotics and Automation*, 1997.

[4] D. Rodríguez-Losada, F. Matía, A. Jiménez, and R. Galán., "Consistency improvement for SLAM-EKF for indoor environments," in *IEEE Conf. on Robotics and Automation*, 2006.

[5] V. Nguyen, A. Harati, and R. Siegwart, "A lightweight SLAM algorithm using orthogonal planes for indoor mobile robotics," in *IEEE/RSJ Int. Conf. on Intelligent Robots and System*, 2007.

[6] K. Beevers and W. Huang1, *Algorithmic Foundation of Robotics VII*, ch. Inferring and Enforcing Relative Constraints in SLAM. Springer Berlin/Heidelberg, 2008.

[7] C. Plagemann, K. Kersting, P. Pfaff, and W. Burgard, "Gaussian beam processes: A nonparametric bayesian measurement model for range finders," in *Robotics: Science and Systems (RSS)*, 2007.

[8] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.

[9] A. P. Ruszczynski, *Nonlinear Optimization*. Princeton U. Press, 2006.

[10] A. Censi, "On achievable accuracy for pose tracking," in *IEEE Conf. on Robotics and Automation*, 2009.

[11] H. Le and D. G. Kendall, "The Riemannian structure of Euclidean shape spaces: A novel environment for statistics," *Annals of Statistics*, vol. 21, no. 3, pp. 1225–1271, 1993.

[12] P. W. Michor and D. Mumford, "Riemannian geometries on spaces of plane curves," *J. of the European Math. Soc.*, vol. 8, pp. 1–48, 2004.

[13] P. Newman, *On the Structure and Solution of the Simultaneous Localisation and Map Building Problem*. PhD thesis, U. of Sydney, 1999.

[14] S. Julier and J. LaViola, "On Kalman filtering with nonlinear equality constraints," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, 2007.