

“© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Exploration on Continuous Gaussian Process Frontier Maps

Maani Ghaffari Jadidi, Jaime Valls Miró, Rafael Valencia and Juan Andrade-Cetto

Abstract—An information-driven autonomous robotic exploration method on a continuous representation of unknown environments is proposed in this paper. The approach conveniently handles sparse sensor measurements to build a continuous model of the environment that exploits structural dependencies without the need to resort to a fixed resolution grid map. A gradient field of occupancy probability distribution is regressed from sensor data as a Gaussian process providing frontier boundaries for further exploration. The resulting continuous global frontier surface completely describes unexplored regions and, inherently, provides an automatic stop criterion for a desired sensitivity. The performance of the proposed approach is evaluated through simulation results in the well-known Freiburg and Cave maps.

I. INTRODUCTION

The standard approach to autonomous robotic exploration in mobile robotics relies on the use of grid-based representations [1]. Relying on grid-based maps however, ignores the structural dependency in the environment due to the assumption of independence between cells. In this paper we suggest to compute frontiers through continuous map building and dynamic updating of the environment. To this end, we introduce the C-frontier map, a continuous model of the boundaries of known-free and unknown areas that accounts for structural dependencies, from which we are able to select goals for further exploration. The C-frontier map is in turn extracted from another continuous representation that models the occupancy probability distribution from sensor data. We use Gaussian processes, as a powerful regression tool in a Bayesian framework, to learn such continuous occupancy map.

II. RELATED WORK

Traditional autonomous exploration strategies have been devised to use occupancy grid maps to represent free, occupied and unknown regions [2]. Despite their popularity, occupancy grid maps exhibit shortcomings due to the assumption of independence between cells, fixed resolution, and the memory intensive requirements to model three-dimensional environments.

This work was supported by Centre for Autonomous Systems, University of Technology Sydney, and by the Spanish Ministry of Innovation and Science project PAU+ (DPI2011-27510).

M. Ghaffari Jadidi and J. Valls Miró are with the Faculty of Engineering and IT, University of Technology Sydney (UTS), NSW 2007, Australia maani.ghaffarijadidi@uts.edu.au, Jaime.VallsMiro@uts.edu.au

J. Andrade-Cetto is with the Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, 08028 Spain. chetto@iri.upc.edu

R. Valencia is with the Center for Applied Autonomous Sensor Systems, Örebro University, SE-701 82 Örebro, Sweden. rafael.valencia-carreno@oru.se




Fig. 1. Cave map entropy evolution

In an effort to tackle the limitation of the granularity in occupancy grid maps in [3], an occupancy grid is maintained for the sole purpose of evaluating spatial entropy when comparing candidate actions. Thus, it can be computed at a very coarse resolution since it is not used to maintain neither the robot localization estimate, nor the structure of the environment. The technique in [3] evaluates exploratory and place revisiting paths, which are selected based on entropy reduction estimates of both the map and the path. Whilst the map entropy is computed on an occupancy grid at coarse resolution, path entropy is the outcome of Pose SLAM [4, 5], a variant of SLAM in which only the robot trajectory is estimated through the observation of relative constraints between robot poses. The evolution of the map entropy for this method v.s. our Gaussian process continuous occupancy map is shown in Fig. 1.

Other strategies have been proposed to move beyond simple grid structures for efficient exploration. In [6] the use of octomaps is suggested, treating the frontier between explored and unexplored areas as boundary conditions, and the explored area as a scalar field. Optimal paths to the boundaries of unexplored sections are computed using steepest descent on the associated gradient field. Another effort to cope with varying resolutions for the explored and unexplored regions in grid maps is presented in [7], where registered sensor data is used to populate an occupied voxel cell, but a sparse free space representation is generated by a particle set. The evolution of a stochastic differential equation that simulates the expansion of the system of particles in free space with Newtonian dynamics determines the sparse unexplored regions. In [8], an information potential field by taking into account the joint entropy of map and path is defined. The robot trajectory is computed from descending on the gradient of the potential field, still, an accurate computation of occupancy maps is preliminary to compute the potential field.

Although the above methods attempt to ease the short-




Fig. 2. A regressed continuous occupancy map of the Cave environment, with size of $20m \times 20m$. (a) Pose SLAM map; red dots indicate robot poses, green lines depict loop-closures. (b) The obstacle probability map Λ_o . (c) The free space probability map Λ_f . (d) The overall fused COM Λ .

comings of using grid maps to explore, they still fall short of accounting for structural correlations in the environment. Recent developments in Bayesian regression and classification methods, particularly from the machine learning community, are providing strong mathematical tools for continuous learning and inference in complex data sets. Non-parametric kernel models, such as Gaussian processes (GP), have proven particularly powerful to represent the affinity of spatially correlated data, overcoming the assumption of independency between cells [9]. The GP associated variance surface equates to a continuous representation of uncertainty in the environment, which can be used to highlight unexplored regions and optimize a robot's search plan. The continuity property of the GP map can improve the flexibility of the planner by inferring directly on collected sensor data without being limited by the resolution of the grid cell [10].

In this paper, we propose an approach to detect frontiers from a continuous representation of the environment that does not suffer from the issues associated to grid maps. It computes exploration goals on the GP associated variance surface. However, training a unique GP for both occupied and free areas results in a mixed variance surface and it is not possible to disambiguate between boundaries of occupied-unknown and free-unknown space without thresholding of the continuous map (see Fig. 6 in [9]). Moreover, it limits selection of an appropriate kernel and results in extrapolated obstacles or low quality free areas. To address this problem we propose training two separate GPs, one for free areas and one for obstacles, and fuse them to come up with a unique continuous occupancy map. Given its associated gradient field over consecutive maps, areas of higher variation would effectively represent incomplete parts of the map, hence equating to an information-driven frontier map for exploration, our so-called C-frontier map.

III. GAUSSIAN PROCESSES

In GPs, statistical inference is employed to learn dependencies between points in a data set [11]. A GP $f(x)$ is described by its mean, $m(x)$, and covariance (kernel) function, $k(x, x')$, as

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \quad (1)$$

where x and x' are the training and test (query) input vectors, respectively. By assuming that the target data, y , is jointly Gaussian, it follows

$$f(x') \sim \mathcal{N}(\mu, \sigma^2) \quad (2)$$

$$\mu = E(f'|x, y, x') = k(x', x)[k(x, x) + \sigma_n^2 I]^{-1} y, \quad (3)$$

$$\sigma^2 = k(x', x') - k(x', x)[k(x, x) + \sigma_n^2 I]^{-1} k(x, x'), \quad (4)$$

where σ_n^2 is the variance of the Gaussian observation noise and f' represents the output values at the test locations.

IV. INFERRING CONTINUOUS OCCUPANCY MAPS WITH GPs

Through inference with GPs, we are able to estimate the full map posterior and avoid the common marginalization in building occupancy grid maps. In this work it is assumed that the robot is equipped with a laser range finder and that local sensor measurements are mapped into a global reference frame, e.g. with a Pose SLAM approach [4].

Training points for the free area GP map are sampled along the laser beam between the robot and the sensed obstacles as in [9]. Computing the obstacles GP map is more straight forward, as it is possible to use the global measured points directly. In both instances the target value can be simply set to one depending on the nature of the map, i.e. a binary classification problem with static state: obstacle or free. Note that in training with one GP the target value can be $y_+ = 1$ and $y_- = -1$ for occupied and free points, respectively.

Since environments are constructed from sudden changes from free areas to obstacles, we are interested in covariance and mean functions which produce as sharp a distribution as possible. However, sharp kernels are inappropriate for covering large free areas, or for learning structural dependencies such as walls.

To this end, for the obstacles GP map we have chosen the Matérn covariance function [12], given by

$$k_{Matern}(x, x') = \frac{1}{\Gamma(\nu)2^{\nu-1}} \left[\frac{\sqrt{2\nu}|x-x'|}{\kappa} \right]^\nu K_\nu \left(\frac{\sqrt{2\nu}|x-x'|}{\kappa} \right), \quad (5)$$

where Γ is the Gamma function, $K_\nu(\cdot)$ is the modified Bessel function of the second kind of order ν , κ is the

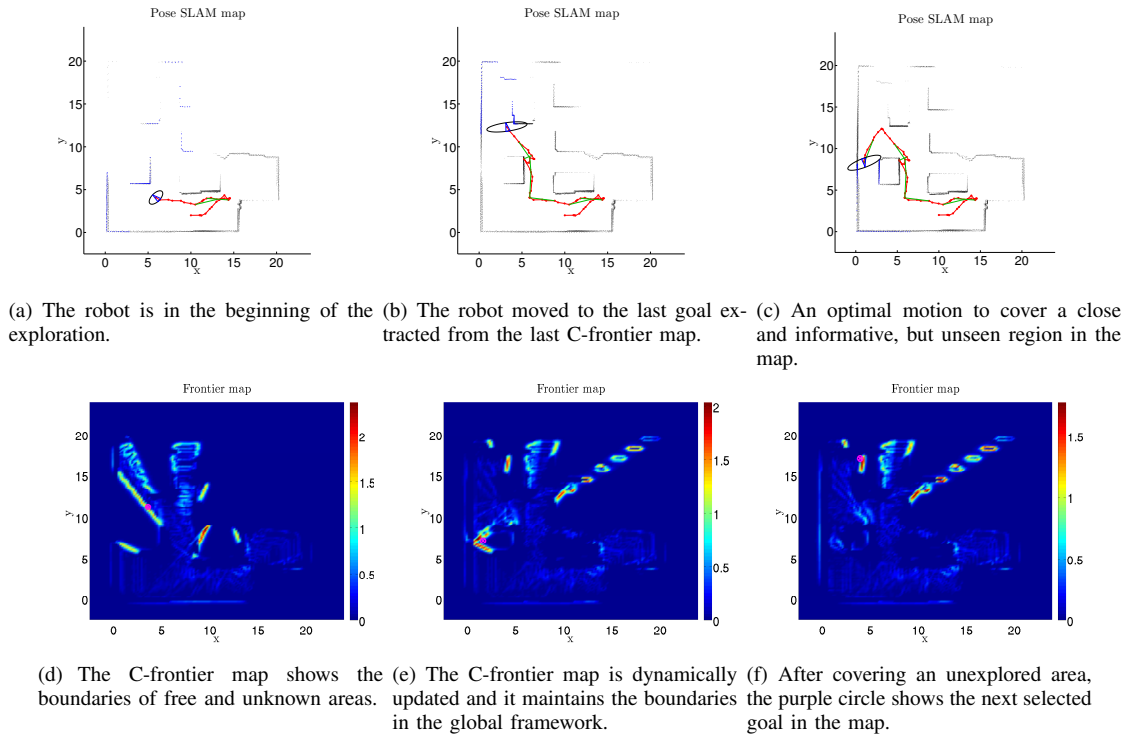


Fig. 3. Evolution of the frontier surface in the Cave map. The C-frontier map is global and facilitates optimal exploration goal extractions. The map is non-dimensional and illustrates a generalized idea for continuous frontiers.

characteristic length scale, and ν is a parameter used to control the smoothness of the covariance.

In order to cover larger spaces with less measurements, the GP map of the free area is defined with a compound covariance function given by the product of a squared exponential covariance function [11], expressed by

$$k_{SE}(x, x') = \exp\left(-\frac{|x - x'|^2}{2\kappa}\right), \quad (6)$$

and a Matérn covariance function.

For a given query point in the map x'_i , our GPs predict mean values for its occupancy and free states μ_i , and associated variances σ_i^2 . In general there is no guarantee that predicted GPs mean values be constrained to the target values or in case of one GP lie in the interval $[y_-, y_+]$, thus we are interested to interpret outputs as distributions by benefiting from GPs mean and variance values concurrently. Following the idea in binary discriminative classification (see sections 3.1 and 3.2 in [11] and chapter 8 in [13]), a logistic function (response function) can be used to squash the GPs outputs into the range $[0, 1]$ and guarantee a valid probabilistic interpretation. By inverting and bounding GPs variances, we can compute the bounded information associated to each location, $\lambda_i = \sigma_{min}^2 / \sigma_i^2$. GPs variances represent the amount of occupancy uncertainty at each point on the map, hence the square root of these information values are proportional to the required weights for logistic regression model, $w_i \propto \lambda_i^{1/2}$. Therefore

$$p(y_{\pm, i} | x_i) = \frac{1}{1 + \exp(-w_i \mu_i)} \quad (7)$$

where $w_i = y_{\pm, i} \gamma \lambda_i^{1/2}$ and $\gamma > 0$ is a constant.

An alternative solution is discussed in [9] through employing a probabilistic least-square classifier (see section 6.5 in [11]). However, it requires a unique covariance matrix inversion for each point which can be a bottleneck for long-term exploration experiments.

Querying over a uniformly sampled range of points, we assemble both an obstacle probability map, $\Lambda_o(x) : \mathbb{R}^2 \mapsto \mathbb{R}$, and a free area probability map, $\Lambda_f(x) : \mathbb{R}^2 \mapsto \mathbb{R}$. The fusion of the two (see section IV-B) is our desired continuous occupancy map (COM). Figure 2 illustrates an example of a regressed COM.

A. Selection of training and query sets

The training set consists of temporally annotated measurement points in the global reference frame. Given that each Pose SLAM pose is annotated with its corresponding laser scan, and that such a map contains only such poses that introduce a relevant amount of information change into the map, the obvious choice is to use such sensor data to train the GPs.

A plausible query set could be a dense uniform distribution over the entire GP domain. That is, sampling all areas covered by the robot sensor range. This is also a computationally intractable choice. Instead, at each iteration, the query set is computed locally over a moving window of fixed size centered at the robot's current pose and covering the current perception field.

Algorithm 1 explorationCOM($p_{robot}, p_{local}, reset$)

```
1:  $p_{global} \leftarrow \text{transform2global}(p_{robot}, p_{local})$ 
2:  $p_{free} \leftarrow \text{lineSegmentation}(p_{robot}, p_{local})$ 
3: if  $firstFrame$  then
4:    $\Lambda = \Lambda_o = \Lambda_f = \emptyset$ 
5:   optimize GP hyperparameters  $\theta_o, \theta_f$ 
6: else if  $reset == true$  then
7:    $\Lambda = \Lambda_o = \Lambda_f = \emptyset$ 
8: end if
9:  $x'_o, x'_f \leftarrow \text{testData}(p_{robot})$ 
10:  $x_o, x_f \leftarrow \text{trainingData}(p_{global}, p_{free})$ 
11:  $\sigma_o^2 \leftarrow \text{GP}(\theta_o, x_o, x'_o), \sigma_f^2 \leftarrow \text{GP}(\theta_f, x_f, x'_f)$ 
12:  $\lambda_o \leftarrow \sigma_{min,o}^2 / \sigma_o^2, \lambda_f \leftarrow \sigma_{min,f}^2 / \sigma_f^2$ 
13:  $p(y_o | x_o) \leftarrow \text{logisticReg}(\mu_o, \lambda_o)$ 
14:  $p(y_f | x_f) \leftarrow \text{logisticReg}(\mu_f, \lambda_f)$ 
15:  $\Lambda_o \leftarrow \text{fusionBCM\_local}(p(y_o | x_o), \Lambda_o)$ 
16:  $\Lambda_f \leftarrow \text{fusionBCM\_local}(p(y_f | x_f), \Lambda_f)$ 
17:  $\Lambda \leftarrow \text{fusionBCM}(\Lambda_o, \Lambda_f)$ 
18:  $\mathcal{F} \leftarrow \|\nabla\Lambda\|_1 - \beta(\|\nabla\Lambda_o\|_1 + \Lambda_o - 0.5)$ 
19:  $goal \leftarrow \text{goalExtraction}(\mathcal{F}, p_{robot})$ 
20: return  $goal$ 
```

B. Fusion and updating the occupancy probability maps

After learning a local COM by using the local query set, we need to fuse it with the global COM. The problem is similar to the mixture of Gaussian processes. However, here we fuse a local part with the current global map. The Bayesian committee machine (BCM) [14], suggests an approach to combine estimators which were trained on different data sets. Assuming a Gaussian prior with zero mean and covariance Σ and each GP with mean $E(f'|D_i)$ and covariance $cov(f'|D_i)$, where $D_i = \{(x, y)_i\}$ is the dataset of observations used for each process, it follows that

$$E(f'|D) = C^{-1} \sum_{i=1}^p cov(f'|D_i)^{-1} E(f'|D_i) \quad (8)$$

with

$$C = cov(f'|D)^{-1} = -(p-1)(\Sigma)^{-1} + \sum_{i=1}^p cov(f'|D_i)^{-1} \quad (9)$$

where p is total number of processes. Note that in this paper, $E(f'|D_i)$ represents probability values from a COM.

V. FRONTIER MAPS

We are in the quest for the extraction of boundaries between known-free and unknown areas directly from a continuous representation of the occupancy probability distribution in the surroundings, without having to resort to a grid map. To this end, assuming the COM has been inferred, we propose to look at its associated gradient field, which effectively represents variations around the boundaries, or incomplete parts of the map. The gradient field is a vector field and its dimension is double the COM in 2D environments. As the C-frontier map is a scalar field and has the same dimensions as the COM, by computing the L_1 -norm (Manhattan distance) of the gradient field, we can reduce the dimensions back to the COM, and also obtain sharp and desirable boundaries. However, boundaries that are related to obstacles need to be accounted for as these are not real frontiers of interest.



Fig. 4. The final occupancy grid maps; a) exploration with the grid-based method and b) exploration with COM. Inferring the COM with GP regression leads to a partially incomplete equivalent occupancy grid map. See Fig. 5 for a completed experiment with GP regression.

To fulfill the above, we define the C-frontier map \mathcal{F} as

$$\mathcal{F} = \|\nabla\Lambda\|_1 - \beta(\|\nabla\Lambda_o\|_1 + \Lambda_o - 0.5), \quad (10)$$

where ∇ is the gradient operator, and β is a factor that controls the effect of obstacle boundaries. $\|\nabla\Lambda\|_1$ defines all boundaries, whilst the second and third terms define obstacle outlines and obstacles, respectively. The subtracted constant from the obstacles is to remove the biased probability for unknown areas in the obstacles probability map. The resulting frontier map \mathcal{F} contains only known-free and unknown boundaries and by clustering points in this map we can select goals for further exploration. Frames (d) to (f) in Fig. 3 show three instances of the C-frontier map for the Cave scenario.

VI. GOAL EXTRACTION

To obtain goal candidates for further exploration, the frontier map is thresholded in order to contain the most informative boundaries, which are in turn clustered with a k-means method. Subsequently, we select the goal with the best balance of information gain and traverse distance, according to

$$f = \max\{\alpha \bar{m}_i \frac{n_i}{N} - d_i^{1/2}, i = 1, 2, \dots, M\}, \quad (11)$$

where d_i is the distance from the current robot pose to the i -th cluster centroid (squared root to prevent steep variations), \bar{m}_i is the mean value of the valid frontier points in the cluster, n_i is the number of points in the i -th cluster, N is the total number of points, M is the total number of clusters, and α is a factor to relate information gain to the cost of motion. Alg. 1 describes the overall procedure to C-frontier map construction and goal extraction for exploration in a continuous space, taking as inputs the local sensor measurements p_{local} and the robot pose p_{robot} in the global reference frame.

VII. MAP REGENERATION

Loop closure during SLAM may change the map significantly. To account for such changes, we propose to reset and learn the COM with all the available data again. To efficiently detect such a drift in the COM we measure the




Fig. 5. Exploration in Freiburg map (a and c) and Cave map (b and d). Map dimensions in meters.

Jensen-Shannon divergence [15]. The generalized Jensen-Shannon divergence for n probability, p_1, p_2, \dots, p_n , with weights $\pi_1, \pi_2, \dots, \pi_n$ is

$$JS_{\pi}(p_1, p_2, \dots, p_n) = H\left(\sum_{i=1}^n \pi_i p_i\right) - \sum_{i=1}^n \pi_i H(p_i) \quad (12)$$

where H is the Shannon entropy function.

$$H(p) = -\sum_{i=1}^n p(x_i) \log(p(x_i)) \quad (13)$$

and $p(x_i)$ is the probability associated to variable x_i . All the weights are set to one as all points are equal.

Alternatively, cumulative relative entropy by summing the computed Jensen-Shannon entropy in each iteration shows map drifts over a period of time and contains the history of map variations. Consequently, the method is less sensitive to small sudden changes.

VIII. RESULTS AND DISCUSSION

The proposed approach is demonstrated with exploratory simulations in two standard mapping environments, the Cave and Freiburg maps [16]. The Cave map represents a simple hand-drawn environment with a few rough obstacles, whereas the Freiburg map is computed from a real LMS-laser log taken with a Pioneer2 robot at the University of Freiburg (building 079 AIS-Lab), and contains many rooms and disconnected obstacles which make for a challenging environment to explore. GP computations have been implemented with the Open Source GP library in [11].

TABLE I

COMPARISON OF THE EXPLORATION METHODS IN THE CAVE MAP (AVERAGED OVER 10 EXPERIMENTS, MEAN \pm STANDARD ERROR)

	Grid-based	GP COM
Travel distance (m)	331	112 \pm 4
Exploration time (min)	33.38	31.75 \pm 5.45
Map entropy rate (NATS/step)	-0.4438	-2.1480 \pm 0.1991
Final map entropy (NATS)	143.8696	146.9389 \pm 0.6757

The proposed approach is compared with a state-of-the-art grid-based method [3] and results are collected in table I. Data are averaged over ten experiments for the listed criteria. The presented method outperforms the grid-based method in all metrics except for the final map entropy.

The final map entropy for both methods is computed from the final occupancy grid map, Fig. 4. In our approach, the occupancy grid map is computed once at the end of the experiment as the COM is inherently different. The final map entropy in our approach is higher than the grid-based method, however, the fact that GP regression covers regions with sparse measurements and computes full map posterior by considering cell dependencies results in a partially incomplete equivalent occupancy grid map. This in part improves travel distance and map entropy rate as it can be seen in table I.

A. Maps with varying levels of detail

The fully explored results of the maps are illustrated in Fig. 5. The complete COM, shown in frame (c), nicely handles sparse measurements as a result of the continuous regression, which is significantly effective for large environ-




Fig. 6. Three points in time during a frontier-based process using a grid map. Beyond the discretization effects, the effect of the independence assumption between cells can be seen as small frontiers appearing since information of near free and occupied cells is not propagated to the rest of the cells.

ments.

The exploration process was set to continue until the C-frontier map was cleared, i.e. no more significant frontiers appear. Thresholding the C-frontier map to derive the most informative areas prior to a k-means clustering process allowed the algorithm to explore the whole area and set the termination condition. This tuning allows the algorithm to disregard frontiers which appear not sufficiently informative to warrant further exploration, reducing the exploration strategy's computational time.

B. Comparison with grid maps

Figure 6 shows three points in time during a traditional frontier-based exploration with grid maps, on the Cave map with a cell size of $0.2m \times 0.2m$. The robot is always driven to the nearest frontiers, with size larger than 9 cells. Besides the loss of information due to the discretization, such a sequence makes evident the effect of the independence assumption between cells. In Fig. 6 small frontiers appear because the information of near free and occupied cells is not propagated to the rest of the cells. Hence, when the robot has eventually explored the larger frontiers, it might be driven to such artifacts instead of more informative regions. In a more realistic model, the occupancy in each place is not randomly distributed over the world as implicitly assumed by grid structures. Instead, a spatial correlation between points in the map should exist given the structured spatial nature of the world around us, and this is exactly what is achieved with the proposed GP maps. Training values for free and occupied space are appropriately propagated over the environment, thus generating well defined frontiers as shown in Fig. 3.

Figure 1 shows the evolution of the map entropy (in nats) for the frontier-based grid exploration and the GP COM for the same number of steps on the Cave map. For the case of the GP continuous map, the map inferred at the same grid size ($0.2m$) to make it a comparable metric. The plateau of entropy values for the grid map case can be explained by the fact that, at each step, the information gained by each cell is not propagated over its vicinity. The opposite happens in a GP map, where at each step the correlation between points in the map helps in reducing the map entropy at a faster rate.

IX. CONCLUSIONS

The novel solution presented in this paper suggests a continuous parametrization of frontiers for autonomous robotic

exploration. The method benefits from GPs to iteratively estimate structural variances or correlations of map points, hence, an inferred COM represents occupied and free regions. In addition, developing the notion of frontiers beyond occupancy grid maps in the form of a continuous non-dimensional frontier surface, where regions for further explorations are the natural boundaries between free-known and unknown areas, justifies the relevance of the proposed COM for exploration. Further efforts are currently being devoted to exploiting the COM in the planning process, and a closer integration with uncertainties from the SLAM process.

REFERENCES

- [1] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Int. Conf. Autonomous Agents*, 1998, pp. 47–53.
- [2] M. Juliá, A. Gil, and O. Reinoso, "A comparison of path planning strategies for autonomous exploration and mapping of unknown environments," *Auton. Robot.*, pp. 1–18, 2012.
- [3] R. Valencia, J. Valls Miro, G. Dissanayake, and J. Andrade-Cetto, "Active Pose SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 1885–1891.
- [4] V. Ila, J. Porta, and J. Andrade-Cetto, "Information-based compact Pose SLAM," *IEEE Trans. Robot.*, vol. 26, no. 1, pp. 78–93, 2010.
- [5] R. Valencia, M. Morta, J. Andrade-Cetto, and J. Porta, "Planning reliable paths with Pose SLAM," *Robotics, IEEE Transactions on*, vol. 29, no. 4, pp. 1050–1059, 2013.
- [6] R. Shade and P. Newman, "Choosing where to go: Complete 3D exploration with stereo," in *Proc. IEEE Int. Conf. Robot Automat.*, 2011, pp. 2806–2811.
- [7] S. Shen, N. Michael, and V. Kumar, "Stochastic differential equation-based exploration algorithm for autonomous indoor 3d exploration with a micro-aerial vehicle," *The Int. J. Robot. Res.*, vol. 31, no. 12, pp. 1431–1444, 2012.
- [8] J. Vallve and J. Andrade-Cetto, "Mobile robot exploration with potential information fields," in *6th European Conference on Mobile Robots*, 2013, pp. 222–227.
- [9] S. T O'Callaghan and F. Ramos, "Gaussian process occupancy maps," *The Int. J. Robot. Res.*, vol. 31, no. 1, pp. 42–62, 2012.
- [10] K. Yang, S. Keat Gan, and S. Sukkarieh, "A gaussian process-based rrt planner for the exploration of an unknown and cluttered environment with a uav," *Advanced Robotics*, vol. 27, no. 6, pp. 431–443, 2013.
- [11] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. MIT press, 2006, vol. 1.
- [12] M. Stein, *Interpolation of spatial data: some theory for kriging*. Springer Verlag, 1999.
- [13] K. P. Murphy, *Machine learning: a probabilistic perspective*. The MIT Press, 2012.
- [14] V. Tresp, "A Bayesian committee machine," *Neural Computation*, vol. 12, no. 11, pp. 2719–2741, 2000.
- [15] J. Lin, "Divergence measures based on the Shannon entropy," *IEEE Trans. Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [16] A. Howard and N. Roy, "The robotics data set repository (Radish)," <http://radish.sourceforge.net>, 2003.