

Exploration in Information Distribution Maps

Maani Ghaffari Jadidi*, Jaime Valls Miró*, Rafael Valencia[†], Juan Andrade-Cetto[‡] and Gamini Dissanayake*

*Faculty of Engineering and IT, University of Technology Sydney, NSW 2007 Australia

Email: maani.ghaffarijadidi, Jaime.VallsMiro, gamini.dissanayake@uts.edu.au

[†]Center for Applied Autonomous Sensor Systems, Örebro University, SE-701 82 Örebro, Sweden

Email: rafael.valencia-carreno@oru.se

[‡]Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Llorens Artigas 4-6, 08028 Barcelona, Spain

Email: cetto@iri.upc.edu

Abstract—In this paper, a novel solution for autonomous robotic exploration is proposed. We model the distribution of information in an unknown environment as an unsteady diffusion process, which can be an appropriate mathematical formulation and analogy for expanding, time-varying, and dynamic environments. This information distribution map is the solution of the diffusion process partial differential equation, and is regressed from sensor data as a Gaussian Process. Optimization of the process parameters leads to an optimal frontier map which describes regions of interest for further exploration. Since the presented approach considers a continuous model of the environment, it can be used to plan smooth exploration paths exploiting the structural dependencies of the environment whilst handling sparse sensor measurements. The performance of the approach is evaluated through simulation results in the well-known Freiburg and Cave maps.

I. INTRODUCTION

Exploration in unknown environments is among the major challenges for an autonomous robot. Autonomous exploration should be efficient with regards to one or more criteria (minimum energy consumption, shortest path, reduced map and localization uncertainties, reduced computational complexity), and be performed with limited perception capabilities due to sensor limitations.

Yamauchi [22] expressed the central question in exploration as: “Given what you know about the world, where should you move to gain as much new information as possible?” Thus, exploration entails building a reliable and accurate world model in which to evaluate such criteria. Currently, almost all available exploration methods rely on planar grid-based representations. Relying on grid-based maps, however, ignores the structural dependency in the environment due to the assumption of independence between cells. Moreover, the use of grid-maps at a fixed resolution have scalability issues both in the size of the area they can handle, and in the dimension of the space to explore, leading to an increase in computational complexity.

We tackle these two shortcomings of grid maps by devising a continuous map representation that takes into account structural dependencies by learning the environment information

distribution from sensor data. In this paper, information is defined as the inverse of the environment’s structural variance. This means that maximization of information (completion of a map) leads to reduction of uncertainty in the environment’s map. We model the distribution of information in the environment as an unsteady diffusion process, which is an appropriate mathematical formulation and analogy for expanding, time-varying, and dynamic environments. We use Gaussian Processes (GPs) as a regression tool in a Bayesian framework to learn the information distribution map. The resulting map is the solution of the diffusion process PDE (partial differential equation) [16]. In contrast with frontier based methods which extract frontiers from an occupancy grid map, an optimal continuous frontier map is computed by optimizing the process parameters. By clustering the frontier map, possible goals for further exploration are extracted.

After a review of the work related to our approach (Section II), in Section III we provide a brief introduction to Gaussian Processes. Next, in Section IV, we introduce our proposed information maps learned and inferred from GPs, and in Section V we present the derivation of the frontier maps as a result of modeling the information distribution as a diffusion process. Sections VI and VII describe the process of extracting goals for exploration and map regeneration, respectively. Finally, sections VIII and IX describe simulation results of the proposed approach on two well known mapping data-sets and conclusion, including possible extensions to this work.

II. RELATED WORK

Juliá et al. [6] studied and compared seven methods for autonomous exploration and mapping of unknown environments on the basis of exploration time and mapping error, both for single and cooperative mapping. The seven major approaches analyzed are nearest frontier [23], cost-utility [3], behaviour-based coordinated [8], coordinated [1], market-based coordinated [24], integrated [10], and hybrid integrated [5]. The comparison yielded conclusive results on the strategy to be used depending on the type of scenario and number of robots used. However, all of the techniques analyzed used occupancy grids for map representation.

Valencia et al. [21] presented an active exploration strategy that tackles the limitation of the granularity of the occupancy

J. Andrade acknowledges support from the Spanish Ministry of Economy and Competitiveness project PAU+ DPI 2011-27510 and the EU FP7 project ARCAS ICT-287617. R. Valencia acknowledges support from the EU FP7 project SPENCER ICT-600877.

grid by minimizing overall map and path entropies. Whereas the map entropy is computed on an occupancy grid at a very coarse resolution, the path entropy is the outcome of maintaining a Pose SLAM [4, 20], a variant of SLAM in which only the robot trajectory is estimated through the observation of relative constraints between robot poses. In the work presented here, we resort also to Pose SLAM for localization and navigation, but maintain instead a continuous information distribution map.

In [21] candidate goals were either frontiers [22], or loop closing entropy minimizers. The paths to either frontiers or loop closure candidates were computed as shortest paths in a probabilistic roadmap [7]. Another strategy to compute exploration paths is to treat the frontier between explored and unexplored areas as boundary conditions, and the explored area as a scalar potential field. The objective here is to find optimal paths to the unexplored. Prestes e Silva et al. [11] for instance, compute the gradient on this field solving Laplace’s PDE with Dirichlet boundary conditions (setting 1 for obstacles and 0 for free cells). This gradient descent direction indicates the path to the unexplored sections. Shade and Newman [13] suggest the use of both Dirichlet and Neumann boundary conditions, and extend the method to work with octomaps. Optimal paths to the boundaries of unexplored sections are computed using steepest descent on the associated gradient field.

To cope with varying resolutions for the explored and unexplored regions, Shen et al. [14] propose to map unexplored regions sparsely and to determine the regions for further exploration based on the evolution of a stochastic differential equation that simulates the expansion of a system of particles with Newtonian dynamics. The method is also presented for a 3D environment, although no considerations are made on how to compute the path to the goal.

The above mentioned methods can consider uncertainty in the occupancy cells but not take into account structural correlations in the environment. The path to the boundary computed on the scalar field assume a uniformly discretized occupancy map. 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, have proven particularly powerful to represent the affinity of spatially correlated data, hence overcoming the traditional assumption of independence between cells, characteristic of the occupancy grid method for mapping environments [17]. 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 [2].

In this paper, we propose an exploration method that 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 differentiate between boundaries of occupied-unknown and free-unknown space (see Fig. 6 in [17]). To address this problem we propose training two separate GPs, one for free areas and one for obstacles, and compute the difference between them to come up with a unique information distribution map for exploration.

III. GAUSSIAN PROCESSES

GPs are a non-parametric Bayesian regression technique in the sense that they do not explicitly define a functional relationship between inputs and outputs. Instead, statistical inference is employed to learn dependencies between points in a data set [12]. 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

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

where

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

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

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

IV. INFERRING INFORMATION DISTRIBUTION MAPS WITH GPs

We maintain two GPs, one for free areas and another one for obstacles. To this end, we assume that the robot is equipped with a laser range finder and that local sensor measurements are mapped into a global reference frame with Pose SLAM [4].

To compute the free area GP map, the training points are sampled along the laser beam between the robot and the sensed obstacles as in [17]. Computing the obstacle GP map is more straight forward, as it is possible to use the measured points in the global reference frame 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.

The selection of an appropriate kernel lies at the heart of GP regression. The covariance function places a prior likelihood on the possible functions used to evaluate the dependencies between the observations. 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. This imposes a trade-off between two competing objectives, smoothness to cover large areas and structure, and sharpness to model discontinuities

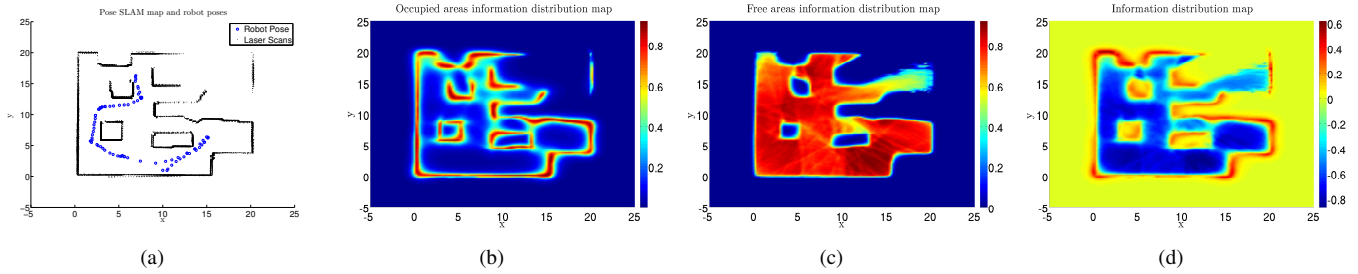


Fig. 1. A regressed information distribution map for the Cave environment, with size of $20m \times 20m$. (a) Map of explored area, where the blue circles are robot poses and the black points are relevant laser scans. (b) The obstacle distribution map Λ_o . (c) The free space distribution map Λ_f . (d) The overall information distribution map Λ . For Λ_o , Λ_f , data is locally normalized between 0 and 1 at each iteration.

well. To this end, we have chosen to use Matérn covariance functions [15],

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

where Γ is the Gamma function, $K_\nu(\cdot)$ is the modified Bessel function of the second kind of order ν , κ is the characteristic length scale, and ν is a parameter used to control the smoothness of the covariance.

The nice feature of this kernel is that with $\nu = 5/2$ and a linear mean function, the resulting distribution maps are twice mean square differentiable. This continuity of the distribution map guarantees smooth planning on the information surface, a feature particularly beneficial for planning in higher dimensions.

Regressing our two GP maps with the sensor data during navigation induces the most likely probability of occupancy and free space. 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 . By inverting these variances, we can compute the information associated to that location, $\lambda_i = 1/\sigma_i^2$. Querying over a uniformly sampled range of robot locations, we assemble both an obstacle information map, $\Lambda_o(x, t) : \mathbb{R}^2 \mapsto \mathbb{R}$, and a free area information map, $\Lambda_f(x, t) : \mathbb{R}^2 \mapsto \mathbb{R}$. With x representing a point in the map and t expressing the time.

In the classical occupancy mapping scenario, the state of a cell can be seen as a binary state (obstacle or free), and one usually computes the log odds $l(x_i) = \log p(x_i) - \log(1 - p(x_i))$, i.e., obstacle over free, which for Normal distributions reduces to something of the form $l(x_i) = \log(|\Sigma_{o_i}|/|\Sigma_{f_i}|) + \|x_{o_i} - \mu_{o_i}\|_{\Sigma_{o_i}}^2 - \|x_{f_i} - \mu_{f_i}\|_{\Sigma_{f_i}}^2$. To simplify the expression, we assume equal mean values for free and obstacle states, i.e., no a priori knowledge. This simplification further reduces the log-odds expression to the difference between the inverse covariances. With this simile, we define our desired information distribution map as

$$\Lambda = \Lambda_o - \Lambda_f \quad (6)$$

Figure 1 illustrates an example of a regressed information distribution map.

A. Selection of Training and Query Sets

Defining a reasonable training set is a key factor in GPs. In our case, 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.

During open loop, each new pose in the map gives an associated scan which is fed to the GP. These updates are local, and seldom overlap previous learned regions in the domain of the GP. At loop closure however, the whole map is recomputed and a new training sample is produced for all scans at their newly computed pose means. The map is updated globally and so is the GP. In this way, we guarantee that flow of information into the GP is equivalent to that in Pose SLAM.

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. However, a more efficient alternative is to compute the query set locally over a moving window of fixed size centered at the robot's current pose and covering the current sensor range.

B. Updating the Global Information Map

After inferring a local information map, we need to fuse it with the global information map. The Bayesian committee machine (BCM) [19], 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

$$E(f'|D) = C^{-1} \sum_{i=1}^p cov(f'|D_i)^{-1} E(f'|D_i) \quad (7)$$

with

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

where p is the total number of GPs. Note that in this paper, $E(f'|D_i)$ represents information values from locally normalized information distribution map.

V. INFORMATION DISTRIBUTION MODEL

As the robot explores the environment, the concentration of information varies with time and depends on the location of the sensed data. Therefore, the concentration of information is a time-dependent scalar field over the explored areas and the movement of the robot spreads the scalar field out as sensory data fill unknown areas with information. This is an analogy with reaction-diffusion processes and we can describe the distribution of information similar to a dynamic system with parabolic PDEs. Inherently, this mathematical form takes into account time and space evolution of the information and even its internal reaction combining newly detected and previously gained data. The unsteady diffusion equation takes the form

$$\frac{\partial \Lambda}{\partial t} = a^2 \nabla^2 \Lambda - g \Lambda \quad (9)$$

where $\Lambda(x, t)$ is the total concentration of information at pose x and at time t , and a^2 and g are the diffusion constant and the information disintegration rate, respectively.

1) *Transient State*: In contrast to the regular procedure with diffusion processes, we are not interested in solving this equation. We are interested instead in computing, at each iteration, the residual

$$\mathcal{R} = \frac{\partial \Lambda}{\partial t} - a^2 \nabla^2 \Lambda + g \Lambda . \quad (10)$$

This residual shows potential areas for further exploration. Exploring these areas result in both, satisfying the information distribution model and compensating for the difference between ideal and regressed data.

To compute the optimal residual, we minimize the squared norm of \mathcal{R}

$$\min_{a, g \in \mathbb{R}} \|\text{vec}(\mathcal{R})\|_2^2 . \quad (11)$$

This optimization leads to optimal values for a^* and g^* in each iteration and, accordingly, an optimal residual, \mathcal{R}^* . We call the negative valued part of the optimal residual a frontier map and compute it with

$$\mathcal{R}^* = \frac{\partial \Lambda}{\partial t} - a^{*2} \nabla^2 \Lambda + g^* \Lambda \quad (12)$$

$$\mathcal{F} = \begin{cases} \mathcal{R}^* & \mathcal{R}^* < 0 \\ 0 & \text{otherwise} \end{cases} . \quad (13)$$

The imposed constraint to select only negative values in the frontier map is a result of the construction of the information map from obstacles and free areas. Only negative values correspond to variations of information between known and unknown areas, whereas positive values represent information variation around obstacles.

2) *Steady State*: When $\frac{\partial \Lambda}{\partial t} = 0$ and $g = 0$, the reaction-diffusion process PDE enters a steady state phase where the behaviour is similar to Laplace's equation. Hence, if after optimizing the PDE parameters, the solution parameters are close to zero, the PDE equation is not a valid proposition anymore, and the frontier map can be obtained solving the Laplace equation. In this case, the frontier map can be defined as

$$\mathcal{F} = \{-\nabla^2 \Lambda : \mathcal{F} < 0\} . \quad (14)$$

On the other hand, for any scalar field it holds that the curl of the gradient is always zero, therefore $|\nabla \times \nabla \Lambda| = 0$, and to exploit both equations concurrently in the frontier map we can write

$$\mathcal{F} = \{-\nabla^2 \Lambda |\nabla \times \nabla \Lambda| : \mathcal{F} < 0\} . \quad (15)$$

As a result of this separation between transient and steady phases, the method considers both time and spatial evolution of the environment. In the transient state, time evolution of the map plays a key role and the most recent informative goals are more relevant, whereas the steady state response offers the possibility to cover all potential regions that have not yet been explored. Fig. 2 shows the frontier surface evolution in transient state during an exploration scenario in the Cave map.

VI. GOAL EXTRACTION

In this section, extraction of exploration goals together with summarization of the proposed algorithm is presented. Once the frontier map is computed, we can generate an arbitrary number of targets on the map. The available continuous frontier map shows the most recent informative regions on the map, and we are free to choose the best target based on desired requirements and priorities. Greedy techniques maximize the difference between the information gain and the cost of an action [18]. We propose a clustering method using the k-means algorithm. Initially, a thresholded frontier map which contains top 30% values (can be variable) is created. Afterwards, a target which balances information gain and traverse distances can be selected

$$f = \max\{f_i = \alpha \bar{m}_i \frac{n_i}{N} - d_i^{1/2}, i = 1, 2, \dots, M\} \quad (16)$$

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 total number of points, M is total number of clusters, and α is a factor to relate information gain to the cost of motion. Algorithm 1 shows the overall procedure 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

In the proposed approach the traditional occupancy grid map has been substituted by the information distribution map. To show the viability of this map for planning, a simple

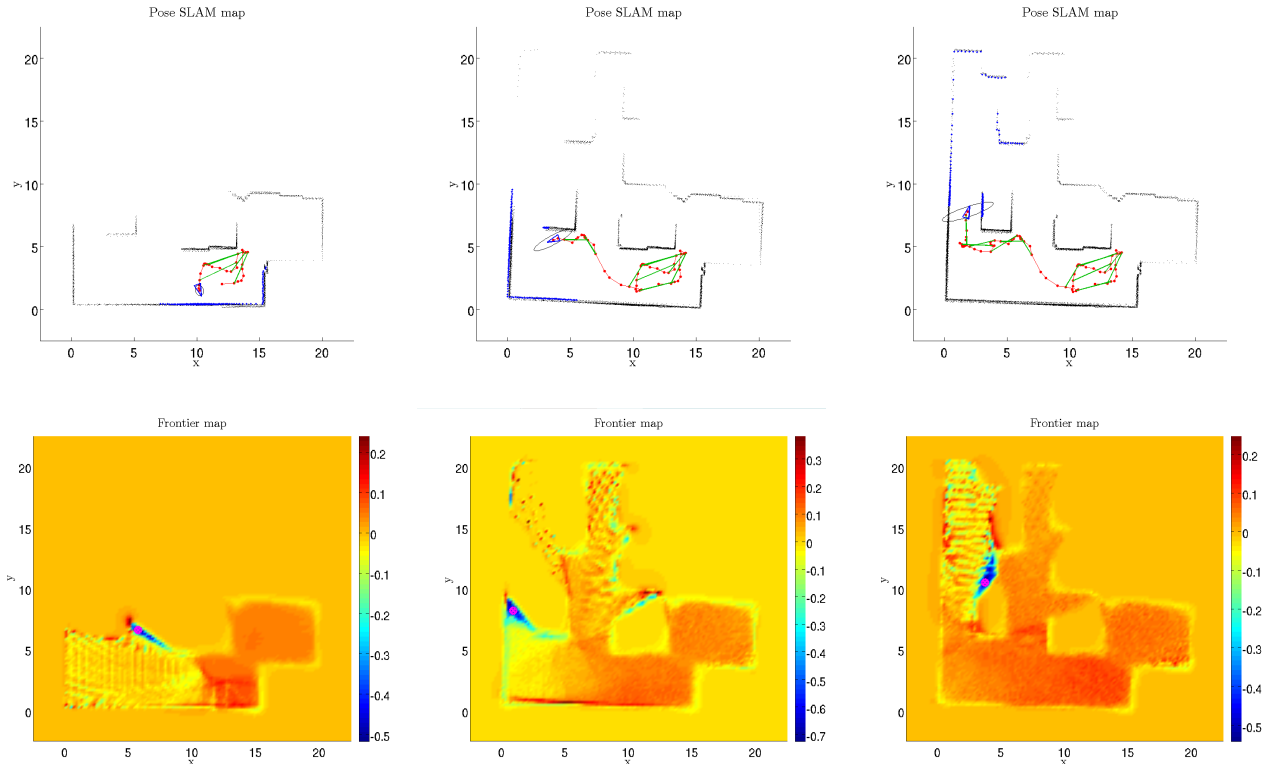


Fig. 2. Evolution of the frontier surface in the Cave map. The frontier map facilitates exploration goal extractions. The map is non-dimensional.

thresholded information map containing the desired unoccupied spaces for safe traversal has been used.

In practice, Pose SLAM may change the map significantly due to loop closure. Large changes in the map can cause notable shifts in robot pose, hence resulting in a low quality information map. To solve this issue, one solution is to reset and learn the information map with all the available data again. A possible measure to detect such drift in the information map can be efficiently calculated with the bounded measure provided by Jensen-Shannon divergence [9]. 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

$$JS_{\pi}(p_1, p_2, \dots, p_n) = H\left(\sum_{i=1}^n \pi_i p_i\right) - \sum_{i=1}^n \pi_i H(p_i) \quad (17)$$

where H is the Shannon entropy function.

$$H(p) = - \sum_{i=1}^n p(x_i) \log(p(x_i)) \quad (18)$$

and $p(x_i)$ is the probability associated to variable x_i .

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

Fusion of local and global information maps as a mixture of GPs inherently provides a smooth recovery of the map. Alternatively, the proposed approach by Valencia et al. [21] for replanning based on map entropy can also be considered to improve the exploration strategy and avoid planning on significantly altered maps.

VIII. RESULTS AND DISCUSSION

A. Maps with varying levels of detail

This section demonstrates the proposed approach through exploration simulation in the Cave¹ and Freiburg² maps. These maps represent varying difficulty of the environments which test the capabilities of the method. The Cave map is small with connected and clear walls. The Freiburg map is relatively large in relation to Cave, and contains many points and disconnected obstacles which make it a challenging environment to explore.

Figure 3 illustrates the fully explored Cave map. As it can be seen in frame (a), by exploiting the information map data, the robot was able to complete the map with a reduced set of movements. Despite the sparse measurements in some regions, the GPs are able to learn the information map, depicted in Fig. 3(b), for goal extraction and planning. Note that calculation of the information map as the difference between obstacles and free areas map not only makes differentiation between occupied-unknown and free-unknown parts possible,

¹Radish: The robotics dataset repository.

²www.slam.org.

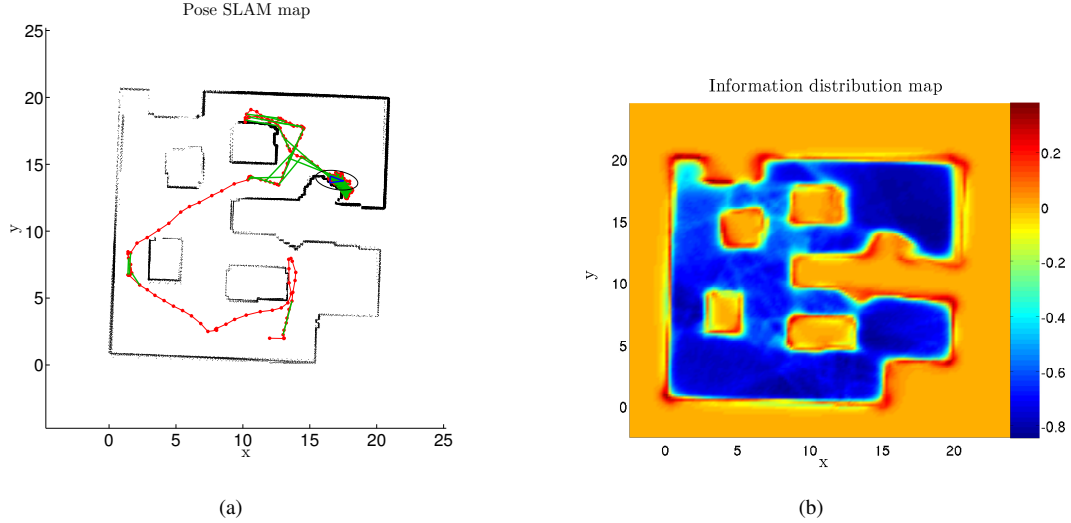


Fig. 3. (a) Exploration in Cave map, with size of $20m \times 20m$. The red lines and points show the planned path and the robot poses, respectively. The green lines are detected loop-closure which caused small rotations in the map in this experiment. (b) Final information distribution map; In spite of the sparse measurements in some locations, at the end of the experiment the whole map is completely explored based on the information map.

Algorithm 1 $\text{explorationIDM}(p_{robot}, p_{local}, reset)$

```

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 GPs 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:  $\lambda_o \leftarrow 1/\sigma_o^2, \lambda_f \leftarrow 1/\sigma_f^2$ 
13: normalize  $\lambda_o, \lambda_f$ 
14: update  $\Lambda_o, \Lambda_f$ 
15:  $\Lambda \leftarrow \Lambda_o - \Lambda_f$ 
16: find  $a^*, g^*$ 
17: if  $\frac{\partial \Lambda}{\partial t} = 0$  and  $g = 0$  then
18:    $\mathcal{F} \leftarrow \{-\nabla^2 \Lambda | \nabla \times \nabla \Lambda | : \mathcal{F} < 0\}$ 
19: else
20:    $\mathcal{F} \leftarrow \{\mathcal{R}^* : \mathcal{R}^* < 0\}$ 
21: end if
22:  $goal \leftarrow \text{goalExtraction}(\mathcal{F}, p_{robot})$ 
23: return  $goal$ 

```

it results in a sharper information map around an obstacle's corner which itself improves the performance of the planner and obstacle avoidance technique.

Fig. 4 shows exploration results in the Freiburg map ($40m \times 15m$). In the illustrated example, the bottom frame shows the computed transient frontier map and the extracted goal in which valid regions for clustering and goal extraction have negative values. As an advantage of planning in the

information map, when the robot reached the target, concurrent map completion and increase in information or reduction of uncertainty occurs. This inherent active link between exploration and mapping is highly demanded and desirable for any autonomous robotic scenario. However, structural complexity of Freiburg map and its size cause a high computational load and requires special computing facilities. Fig. 5 shows the capabilities of the method to accommodate to varying distributions of obstacles in a exploration run over the same environment. In the region in the left part of the map, with little obstacle density, the robot is commanded with larger motion commands than in the right part of the map, in which there is significantly more clutter, and the robot is commanded with slow, short movements.

Stop criterion depends on application, but a common choice in map completion is exploring till no frontier remains. Thus, the stop criterion is $\mathcal{F} = \emptyset$ or in a more practical way $\|vec(\mathcal{F})\|_\infty < \delta$. In this research, all the presented results are achieved with $\delta = 0.09$. Another useful measure is the predefined relative entropy. No significant change of a map's relative entropy over a period of time means there is no variation in the map and the robot is in a fully explored area. In addition, cumulative relative entropy provides a measure to control the desired quality of the information map. Therefore, in trivial environments by selection of a higher threshold, it can reduce computational time of the exploration strategy.

GP computations have been implemented with the Open Source GP library in [12].

B. Comparison with gridmaps

Fig. 6 shows three points in time during a traditional frontier-based exploration with grid maps, using the same environment and specifications from the results shown in Fig. 2, 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.

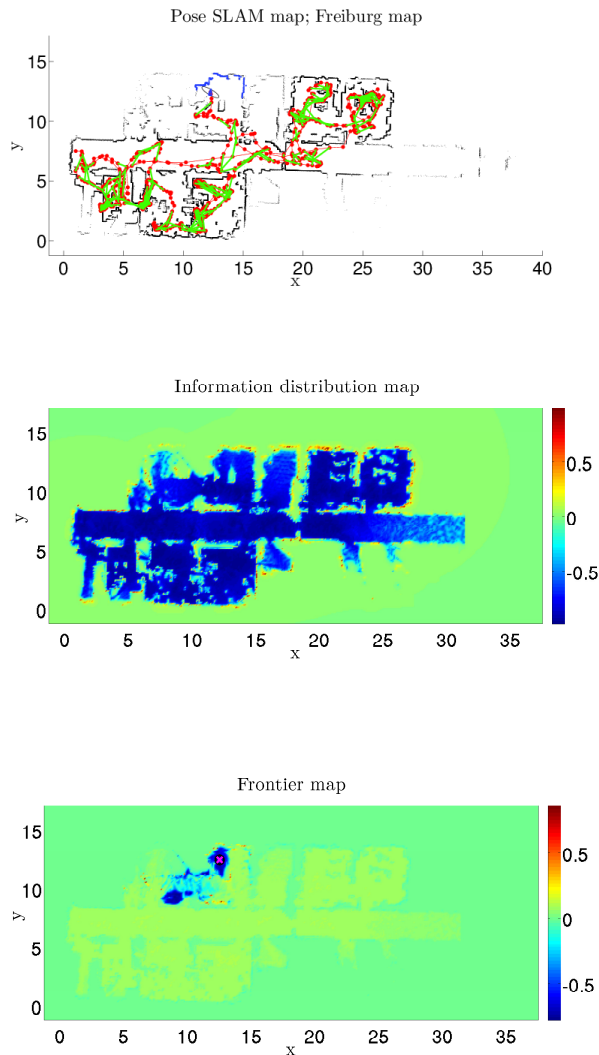


Fig. 4. Exploration in Freiburg map which is a challenging environment due to the many discontinuities, obstacles, and rooms. In the information distribution map, fully and partially explored regions can be easily identified. The information variation over boundaries of unexplored areas shows potential regions for further exploration. The bottom frame shows the transient frontier map, in which lower and negative values indicate regions of interest for goal extraction. The extracted goal is shown with a purple circle and, as it can be seen in the Pose SLAM map, the robot successfully moved toward the goal.

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 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.

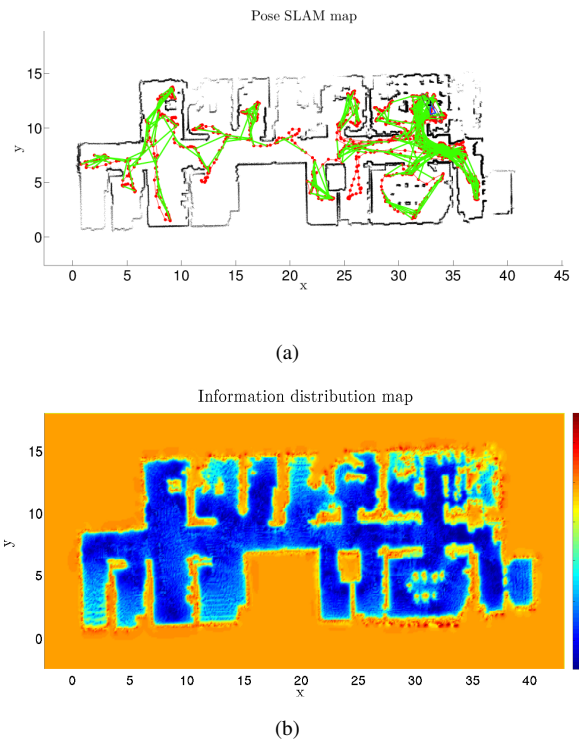


Fig. 5. Long-term exploration in the Freiburg map. In this run, the robot has been successfully exploring the map with respect to the stop criteria for a significantly larger period of time. note how in the eastern part of the map there are discontinuities in obstacles, and it is more challenging for the robot to explore. b) Final information distribution map.

Training values for free and occupied space are appropriately propagated over the environment, thus generating well defined frontiers as shown in Fig. 2.

IX. CONCLUSION

A solution to the robot exploration problem is proposed in this work with the introduction of a novel information distribution map. The main improvement that this solution brings in relation to traditional mapping approaches is avoiding the need to create an occupancy grid map, thus explicitly avoiding the strong assumption of independence between cells, which in turn facilitate accounting for structural dependencies in the environment. Moreover, the solution is also a better fit to efficient exploration of larger environments given the well-known scalability issues of grid-maps at fixed resolutions, and the increase in computational complexity that factor leads to.

It is also plausible to extract goals directly from the global information distribution map. Although this map varies with time, working with this map to extract potential exploration goals is a time-independent calculation, and only the currently available global information map is relevant at any given instant. In addition, the multivariate and continuous nature of the approach together with the possibility of iterative updating appears promising for 3D applications.




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.

REFERENCES

- [1] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proc. IEEE Int. Conf. Robot Automat.*, volume 1, pages 476–481, 2000.
- [2] S.K. Gan, K.J. Yang, and S. Sukkarieh. 3D online path planning in a continuous Gaussian process occupancy map. In *Proc. Australian Conf. Field Robotics*, 2009.
- [3] H.H. González-Banos and J.C. Latombe. Navigation strategies for exploring indoor environments. *The Int. J. Robot. Res.*, 21(10-11):829–848, 2002.
- [4] V. Ila, J.M. Porta, and J. Andrade-Cetto. Information-based compact Pose SLAM. *IEEE Trans. Robot.*, 26(1):78–93, 2010.
- [5] M. Juliá, Ó. Reinoso, A. Gil, M. Ballesta, and L. Payá. A hybrid solution to the multi-robot integrated exploration problem. *Eng. Appl. Artif. Intell.*, 23(4):473–486, 2010.
- [6] M. Juliá, A. Gil, and O. Reinoso. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Auton. Robot.*, pages 1–18, 2012.
- [7] L. Kavraki, P. Svestkaand, J. C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. Robot.*, 12:566–580, 1996.
- [8] H. Lau. Behavioural approach for multi-robot exploration. In *Australasian Conf. Robot. Automat.*, 2003.
- [9] J. Lin. Divergence measures based on the Shannon entropy. *IEEE Trans. Information Theory*, 37(1):145–151, 1991.
- [10] A.A. Makarenko, S.B. Williams, F. Bourgault, and H.F. Durrant-Whyte. An experiment in integrated exploration. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, volume 1, pages 534–539, 2002.
- [11] E. Prestes e Silva, P.M. Engel, M. Trevisan, and M.A.P. Idiart. Exploration method using harmonic functions. *Robot. Auton. Syst.*, 40(1):25–42, 2002.
- [12] C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*, volume 1. MIT press, 2006.
- [13] R. Shade and P. Newman. Choosing where to go: Complete 3D exploration with stereo. In *Proc. IEEE Int. Conf. Robot Automat.*, pages 2806–2811, 2011.
- [14] 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.*, 31(12):1431–1444, 2012.
- [15] M.L. Stein. *Interpolation of spatial data: some theory for kriging*. Springer Verlag, 1999.
- [16] T. Stocker. *Introduction to climate modelling*. Springer, 2011.
- [17] S. T O’Callaghan and F.T. Ramos. Gaussian process occupancy maps. *The Int. J. Robot. Res.*, 31(1):42–62, 2012.
- [18] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*, volume 1. MIT press, 2005.
- [19] V. Tresp. A Bayesian committee machine. *Neural Computation*, 12(11):2719–2741, 2000.
- [20] R. Valencia, J. Andrade-Cetto, and J.M. Porta. Path planning in belief space with Pose SLAM. In *Proc. IEEE Int. Conf. Robot Automat.*, pages 78–83, 2011.
- [21] R. Valencia, J. Valls Miro, G. Dissanayake, and J. Andrade-Cetto. Active Pose SLAM. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pages 1885–1891, 2012.
- [22] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Int. Sym. Comput. Intell. Robot. Automat.*, pages 146–151, 1997.
- [23] B. Yamauchi. Frontier-based exploration using multiple robots. In *Int. Conf. Autonomous Agents*, pages 47–53, 1998.
- [24] R. Zlot, A. Stentz, M.B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proc. IEEE Int. Conf. Robot Automat.*, volume 3, pages 3016–3023, 2002.

Home
RSS X
Content
Authors
RSS IX
Content
Authors
RSS VIII
Content
Authors
RSS VII
Content
Authors
RSS VI
Content
Authors
RSS V
Content
Authors
RSS IV
Content
Authors
RSS III
Content
Authors
RSS II
Content
Authors
RSS I
Content
Authors
Ethics and Malpractice
Statement
Contact
Webmaster

ROBOTICS

SCIENCE AND SYSTEMS

ONLINE PROCEEDINGS

Robotics: Science and Systems XI

July 13 - July 17, 2015
 Sapienza University of Rome
 Rome, Italy
[conference web site](#)

Robotics: Science and Systems X

July 12 - July 16, 2014
 University of California
 Berkeley, USA
[conference web site](#)
 Edited by Dieter Fox, Lydia E. Kavraki and Hanna Kurniawati
 ISBN 978-0-9923747-0-9

Robotics: Science and Systems IX


June 24 - June 28, 2013
 Technische Universität Berlin
 Berlin, Germany
[conference web site](#)
 Edited by Paul Newman, Dieter Fox and David Hsu
 ISBN 978-981-07-3937-9
 ISSN 2330-765X Vol.9 No. 1 2013

Robotics: Science and Systems VIII

July 09 - July 13, 2012
 University of Sydney
 Sydney, NSW, Australia
[conference web site](#)
 MIT Press, Cambridge, Massachusetts, 2013
 Edited by Nicholas Roy, Paul Newman and Siddhartha Srinivasa
[table of contents](#)
[publisher's web site](#)
 ISBN-13: 978-0-262-51968-7

Robotics: Science and Systems VII

June 27 - June 30, 2011
 University of Southern California
 Los Angeles, USA
[conference web site](#)
 MIT Press, Cambridge, Massachusetts, 2012
 Edited by Hugh Durrant-Whyte, Nicholas Roy and Pieter Abbeel



We are using google analytics for statistics!

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-51779-5

ISBN-13: 978-0-262-51779-9

Robotics: Science and Systems VI

June 27 - June 30, 2010

Universidad de Zaragoza

Zaragoza, Spain

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2011

Edited by Yoky Matsuoka, Hugh Durrant-Whyte and José Neira

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-51681-0

ISBN-13: 978-0-262-51681-5

Robotics: Science and Systems V

June 28 - July 1, 2009

University of Washington

Seattle, USA

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2010

Edited by Jeff Trinkle, Yoky Matsuoka and Jose A. Castellanos

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-51463-X

ISBN-13: 978-0-262-51463-7

Robotics: Science and Systems IV

June 25-28, 2008

Eidgenössische Technische Hochschule Zürich (ETHZ)

Zurich, Switzerland

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2009

Edited by Oliver Brock, Jeff Trinkle and Fabio Ramos

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-51309-9

ISBN-13: 978-0-262-51309-8

Robotics: Science and Systems III

June 27-30, 2007

Georgia Institute of Technology

Atlanta, Georgia

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2008

Edited by Wolfram Burgard, Oliver Brock and Cyrill Stachniss

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-52484-8

ISBN-13: 978-0-262-52484-1

Robotics: Science and Systems II

August 16-19, 2006

University of Pennsylvania

Philadelphia, Pennsylvania

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2007

Edited by Gaurav S. Sukhatme, Stefan Schaal, Wolfram Burgard and

Dieter Fox

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-69348-8

ISBN-13: 978-0-262-69348-6

Robotics: Science and Systems I

June 8-11, 2005

Massachusetts Institute of Technology

Cambridge, Massachusetts

[conference web site](#)

MIT Press, Cambridge, Massachusetts, 2005

Edited by Sebastian Thrun, Gaurav S. Sukhatme and Stefan Schaal

[table of contents](#)

[publisher's web site](#)

ISBN-10: 0-262-70114-6

ISBN-13: 978-0-262-70114-3

The International Standard Serial Number (ISSN) of the Robotics:
Science and Systems Proceedings is ISSN: 2330-765X.

RSS '13 Workshop

On Robotic Exploration, Monitoring, and Information Collection: Nonparametric Modeling, Information-based Control, and Planning under Uncertainty

[HOME](#) [INVITED SPEAKERS](#) [INVITED TALK ABSTRACTS](#) **[CONTRIBUTED PAPERS](#)** [SCHEDULE](#)

[ORGANIZERS](#)

Contributed Papers

Resource Constrained Exploration in Reinforcement Learning

- Jen Jen Chung (The University of Sydney)
- Nicholas Lawrance (The University of Sydney)
- Salah Sukkarieh (The University of Sydney)

Abstract

This paper examines temporal difference reinforcement learning (RL) with adaptive and directed exploration for resource-limited missions. The scenario considered is for an energy-limited agent which must explore an unknown region to find new energy sources. The presented algorithm uses a Gaussian Process (GP) regression model to estimate the value function in an RL framework. However, to avoid myopic exploration we developed a resource-weighted objective function which combines an estimate of the future information gain using an action rollout with the estimated value function to generate directed explorative action sequences. The results show that under this objective function, the learning agent is able to continue exploring for better state-action trajectories when platform energy is high and follow conservative energy gaining trajectories when platform energy is low.

[Download PDF](#)

Autonomous Methods for Environmental Modelling and Exploration

- Asher Bender (ACFR)

- Stefan Williams (ACFR)
- Oscar Pizarro (ACFR)

Abstract

The problem of efficiently gathering information from large-scale, natural environments applies to a wide variety of applications from planetary exploration to remote sensing and environmental monitoring. This paper addresses two issues which arise when using robotic platforms to explore large-scale, natural environments. The first issue is efficient processing of observed data, often from multiple sources, into a useful model of the environment. In this paper, a benthic habitat model is created using in-situ optical imagery collected by an autonomous underwater vehicle (AUV) and a digital elevation model of the seafloor collected using ship-borne multibeam echo-sounder data. To efficiently process large numbers of in-situ observations a variational Dirichlet process model is used to categorize the AUV images. A broader model of the environment is established using these observations and a Gaussian process classifier. Both models are Bayesian and non-parametric. This data driven approach, frees human experts from making narrow modelling assumptions and the burden of processing large volumes of data. Once a habitat model is available, the second issue that arises is to make effective use of the model to plan additional informative surveys. In this paper a mutual information objective function is described and used to select the optimal location to place a pre-specified survey template. The modelling and exploration methods are demonstrated on real marine data including two AUV deployments designed to test the proposed planning objective function. The results show that the methods are successful at modelling the environment with little human intervention and are able to recommend informative surveys.

[Download PDF](#)

Exploration in Information Distribution Maps

- Maani Ghaffari Jadidi (Univ. of Tech., Sydney)
- Jaime Valls Miro (Univ. of Tech., Sydney)
- Juan Andrade-Cetto (Univ. of Tech., Sydney)
- Rafael Valencia (Institut de Robotica i Informatica Industrial)
- Gamini Dissanayake (Univ. of Tech., Sydney)

Abstract

In this paper, a novel solution for autonomous robotic exploration is

proposed. The distribution of information in an unknown environment is modeled as an unsteady diffusion process, which can be an appropriate mathematical formulation and analogy for expanding, time-varying, and dynamic environments. This information distribution map is the solution of the diffusion process partial differential equation, and is regressed from sensor data as a Gaussian Process. Optimization of the process parameters leads to an optimal frontier map which describes regions of interest for further exploration. Since the presented approach considers a continuous model of the environment, it can be used to plan smooth exploration paths exploiting the structural dependencies of the environment whilst handling sparse sensors measurements. The performance of the proposed approach is evaluated through simulation results in the well-known Freiburg and Cave maps.

[Download PDF](#)

Bayesian Optimisation for Active Perception and Smooth Navigation

- Jefferson Souza (University of Sao Paulo)
- Roman Marchant (University of Sydney)
- Lionel Ott (University of Sydney)
- Denis Wolf (Universidade de São Paulo)
- Fabio Ramos (University of Sydney)

Abstract

A key challenge for long-term autonomy is to enable a robot to automatically model properties of the environment while actively searching for better decisions to accomplish its task. This amounts to the problem of exploration-exploitation in the context of active perception. This paper addresses active perception and presents a technique to incrementally model the roughness of the terrain a robot navigates on while actively searching for paths that reduce the overall vibration experienced during travel. The approach employs Gaussian processes in conjunction with Bayesian optimisation for decision making. The algorithms are executed in real-time on the robot while it explores the environment. We present experiments with an outdoor vehicle navigating over several types of terrains demonstrating the properties and effectiveness of the approach.

[Download PDF](#)

Sampling-based Planning for Sensor Scheduling in Constrained Environments

- Marin Kobilarov (Johns Hopkins University)
- Gaurav Sukhatme (University of Southern California)
- Jerrold Marsden (California Institute of Technology)

Abstract

This paper considers the optimal estimation of the state of a dynamic observable using a mobile sensor. The main goal is to compute a sensor trajectory which minimizes the estimation error over a given time horizon taking into account the uncertainty in the observable dynamics and its sensing and respecting the constraints of the workspace. The main contribution is a methodology for handling nonlinear stochastic dynamics and environment constraints in a global optimization framework. It is based on sequential Monte Carlo methods and sampling-based motion planning. Three heuristic variance reduction techniques—utility sampling, shuffling, and pruning—based on importance sampling, are proposed to speed-up convergence. The developed framework is applied to two typical scenarios: a simple vehicle operating in a planar polygonal obstacle environment; a simulated helicopter searching for a moving target in a 3-D terrain.

[Download PDF](#)

Efficient modeling of non-homogeneous field dynamics for distributed robotic sampling

- Young-Ho Kim (Texas A&M University)
- Dylan Shell (Texas A&M University)

Abstract

Environmental monitoring is an important application for robotics. In a marine setting this could involve estimating attributes such as water temperature or pH. Selecting measurement locations intelligently for adaptive sampling requires estimates of uncertainty across the region of interest and, recently, several researchers have used linear least squares regression techniques to estimate the values of such attributes and uncertainty. Dynamic fields continue to pose a challenge for robotic sampling despite the fact that many of the measured attributes change naturally over time. This is partly because an evolving environment may have non-homogeneous spatial variability and require continual re-sampling of the field or re-modeling of the field

dynamics. This paper proposes an efficient modeling technique for treating changing fields by augmenting Ordinary Kriging to enable representation of the field's (potentially nonhomogeneous) evolution through Bayes filtering of parameters that characterize the underlying dynamics. The new modeling approach not only enables adaptive path planning in the field but also leads to a straight-forward formulation of the optimal workload distribution through modification of an approximate graph partitioning algorithm. Using a simulated multi-robot sampling scenario, we demonstrate and validate the approach with simulated data. The experiments show good performance in terms of cross-validation using real values and illustrate how dynamic fields with hotspots are identified and modeled, in turn affecting the division of labor.

[Download PDF](#)

Fast Marching Trees: a Fast Marching Sampling-Based Method for Optimal Motion Planning in Many Dimensions

- Lucas Janson (Stanford University)
- Marco Pavone (Stanford University)

Abstract

In this paper we present a novel probabilistic sampling-based motion planning algorithm called the Fast Marching Tree algorithm (FMT*). The algorithm is specifically aimed at solving complex motion planning problems in high-dimensional configuration spaces. This algorithm is proven to be asymptotically optimal and is shown to converge to an optimal solution faster than its state-of-the-art counterparts, chiefly, PRM* and RRT*. An additional advantage of FMT* is that it builds and maintains paths in a tree-like structure (especially useful for planning under differential constraints). The FMT* algorithm essentially performs a “lazy” dynamic programming recursion on a set of probabilistically-drawn samples to grow a tree of paths, which moves steadily outward in cost-to-come space. As such, this algorithm combines features of both single-query algorithms (chiefly RRT) and multiple-query algorithms (chiefly PRM), and is conceptually related to the Fast Marching Method for the solution of eikonal equations. As a departure from previous analysis approaches that are based on the notion of almost sure convergence, the FMT* algorithm is analyzed under the notion of convergence in probability: the extra mathematical flexibility of this approach allows for significant algorithmic advantages and is of independent interest. Numerical experiments over a range of dimensions and obstacle

configurations confirm our theoretical and heuristic arguments by showing that FMT returns substantially better solutions than either PRM* or RRT*.

[Download PDF](#)

[Proudly powered by WordPress](#)