# Exploration using an Information-Based Reaction-Diffusion Process

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

University of Technology, Sydney (UTS), Australia

{maani.ghaffarijadidi, jaime.vallsmiro, gamini.dissanayake}@uts.edu.au

Örebro University, Sweden  rafael.valencia-carreno@oru.se

Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Spain  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.

## 1 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 [Yamauchi, 1997] 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) [Stocker, 2011]. In contrast with frontier based methods which extract frontiers from an occupancy grid map, a 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 2), in Section 3 we provide a brief introduction to Gaussian Processes. Next, in Section 4, we in-

troduce our proposed information maps learned and inferred from GPs, and in Section 5 we present the derivation of the frontier maps as a result of modeling the information distribution as a diffusion process. Sections 6 and 7 describe the process of extracting goals for exploration and map regeneration, respectively. Finally, sections 8 and 9 describe simulation results of the proposed approach on two well known mapping data-sets and conclusion, including possible extensions to this work.

## 2  Related Work

Juliá *et al.* [Juliá *et al.*, 2012] 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 [Yamauchi, 1998], cost-utility [González-Banos and Latombe, 2002], behaviour-based coordinated [Lau, 2003], coordinated [Burgard *et al.*, 2000], market-based coordinated [Zlot *et al.*, 2002], integrated [Makarenko *et al.*, 2002], and hybrid integrated [Juliá *et al.*, 2010]. 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.* [Valencia *et al.*, 2012] presented an active exploration strategy that tackles the limitation of the granularity of the occupancy 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 [Ila *et al.*, 2010; Valencia *et al.*, 2013], 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 [Valencia *et al.*, 2012] candidate goals were either frontiers [Yamauchi, 1997], or loop closing entropy minimizers. The paths to either frontiers or loop closure candidates were computed as shortest paths in a probabilistic roadmap [Kavraki *et al.*, 1996]. 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 sections. In [Prestes e Silva *et al.*, 2002] for instance, the gradient is computed 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 [Shade and Newman, 2011]

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, in [Shen *et al.*, 2012] is proposed 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. In [Vallve and Andrade-Cetto, 2013], 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.

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 [T O'Callaghan and Ramos, 2012]. 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 [Gan *et al.*, 2009].

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 without thresholding of the continuous map (see Fig. 6 in [T O'Callaghan and Ramos, 2012]). Moreover, it limits selection of an appropriate kernel and results in extrapolated obstacles or low quality free areas map. 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.

# 3 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 [Rasmussen and Williams, 2006]. A GP $f(x)$ is described by its mean, $m(x)$, and covariance (kernel) function, $k(x, x')$, as

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

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

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

where

$$\mu = E(f'|x, y, x') = k(x', x)[k(x, x) + \sigma_n^2 I]^{-1} y, \qquad (3)$$
$$\sigma^2 = k(x', x') - k(x', x)[k(x, x) + \sigma_n^2 I]^{-1} k(x, x'), \qquad (4)$$

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

# 4 Inferring Information Distribution Maps with GPs

This section presents our solution for the mapping problem. Through inference with GPs, we are able to estimate the full map posterior and avoid marginalization. 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 [Ila *et al.*, 2010].

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 [T O'Callaghan and Ramos, 2012]. 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. Note that in training with one GP the target value can be $-1$ and $1$ for occupied and free points, respectively.

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 well. To this end, we have chosen to use Matérn covariance functions [Stein, 1999],

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

where $\Gamma$ is the Gamma function, $K_\nu(\cdot)$ is the modified Bessel function of the second kind of order $\nu$, $\kappa$ is the characteristic length scale, and $\nu$ 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 $\mu_i$, and associated variances $\sigma_i^2$. 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, $\Lambda_o(x, t) : \mathbb{R}^3 \mapsto \mathbb{R}$, and a free area information map, $\Lambda_f(x, t) : \mathbb{R}^3 \mapsto \mathbb{R}$. 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 $l(x_i) = \log(|\Sigma_{o_i}|/|\Sigma_{f_i}|) + \|x_{o_i} - \mu_{o_i}\|_{\Sigma_{o_i}}^2 - \|x_{f_i} - \mu_{f_i}\|_{\Sigma_{f_i}}^2$. 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

$$\Lambda = \Lambda_o - \Lambda_f \ . \qquad (6)$$

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

## 4.1 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 anno-
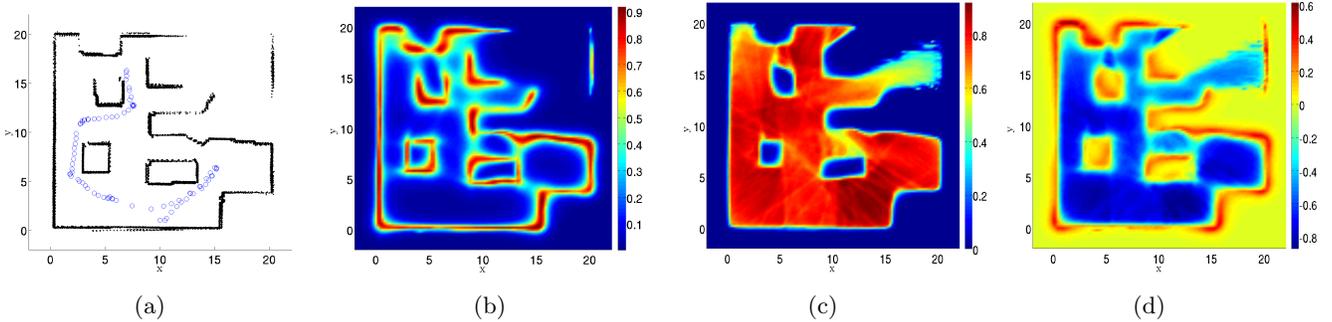
Figure 1: A regressed information distribution map for the Cave environment, size of $20m \times 20m$. *a)* Map of explored area, where the blue circles are robot poses and the black points are relevant laser scans. *b)* The obstacle distribution map $\Lambda_o$. *c)* The free space distribution map $\Lambda_f$. *d)* The overall information distribution map $\Lambda$. For $\Lambda_o$, $\Lambda_f$, data is locally normalized between 0 and 1 at each iteration.

tated 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.

## 4.2 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) [Tresp, 2000], suggests an approach to combine estimators which were trained on different data sets. Assuming a Gaussian prior with zero mean and covariance $\Sigma$ 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

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

with

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

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.

## 5 Information Distribution Model

As the robot explores the environment, the concentration of information (uncertainty) 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

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

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.

### 5.1 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

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

The residual contains the deviation between the proposed model and the real process and, therefore, 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 $\mathcal{R}$

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

This optimization leads to optimal values for $a^*$ and $g^*$ in each iteration and, accordingly, an optimal residual, $\mathcal{R}^*$. The negative parts of the information distribution map define free areas and frontiers are boundaries between free and unkown areas, hence, we call the negative valued part of the optimal residual a frontier map and compute it with

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

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

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 variations around obstacles.

## 5.2 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

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

On the other hand, for any scalar field it holds that the curl of the gradient is always zero, therefore

$$\lim_{t \to \infty} |\nabla \times \nabla \Lambda| = 0,$$

and to exploit both equations concurrently in the frontier map we can write

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

---

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

1: $p_{global} \leftarrow$ **transform2global**$(p_{robot}, p_{local})$
2: $p_{free} \leftarrow$ **lineSegmentation**$(p_{robot}, p_{local})$
3: **if** $firstFrame$ **then**
4: $\quad \Lambda = \Lambda_o = \Lambda_f = \emptyset$
5: $\quad$ optimize GPs hyperparameters $\theta_o, \theta_f$
6: **else if** $reset == true$ **then**
7: $\quad \Lambda = \Lambda_o = \Lambda_f = \emptyset$
8: **end if**
9: $x'_o, x'_f \leftarrow$ **testData**$(p_{robot})$
10: $x_o, x_f \leftarrow$ **trainingData**$(p_{global}, p_{free})$
11: $\sigma_o^2 \leftarrow$ **GP**$(\theta_o, x_o, x'_o)$, $\sigma_f^2 \leftarrow$ **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: $\quad \mathcal{F} \leftarrow \{-\nabla^2 \Lambda | \nabla \times \nabla \Lambda| : \ \mathcal{F} < 0\}$
19: **else**
20: $\quad \mathcal{F} \leftarrow \{\mathcal{R}^* : \ \mathcal{R}^* < 0\}$
21: **end if**
22: $goal \leftarrow$ **goalExtraction**$(\mathcal{F}, p_{robot})$
23: **return** $goal$

---

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. Figure 2 shows the frontier surface evolution in transient state during an exploration scenario in the Cave map.

## 6 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 [Thrun et al., 2005]. 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

$$f = max\{f_i = \alpha \bar{m}_i \frac{n_i}{N} - d_i^{1/2}, \ i = 1, 2, ..., M\} \tag{16}$$
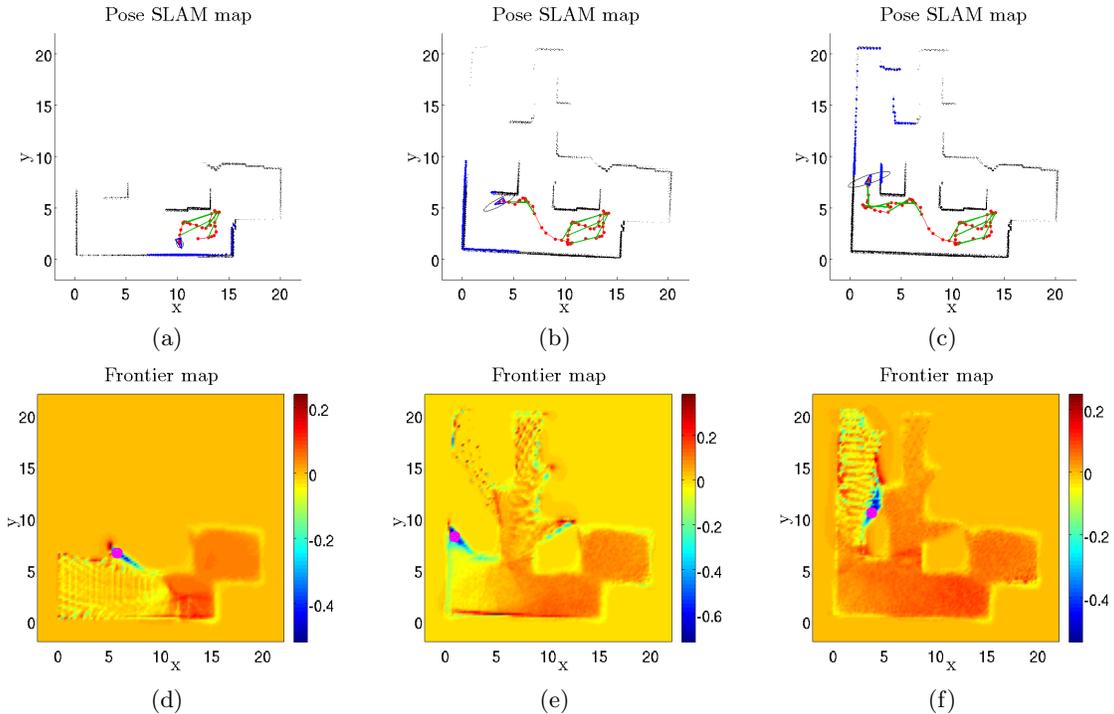
Figure 2: Evolution of the frontier surface in the Cave map during an exploration experiment. $a - c$) In pose SLAM maps red lines and points show the planned path and the robot poses, respectively. Green lines indicate detected loop-closure. $d - f$) The frontier map facilitates exploration goal extractions. Negative values show information variations between known and unknown areas, whereas positive values represent information variations around obstacles. Maps are non-dimensional.

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 $\alpha$ 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.

## 7 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 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 cal-

culated with the bounded measure provided by Jensen-Shannon divergence [Lin, 1991]. The generalized Jensen-Shannon divergence for $n$ probability, $p_1, p_2, ..., p_n$, with weights $\pi_1, \pi_2, ..., \pi_n$ is

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

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. Furthermore, the proposed approach by [Valencia et al., 2012] for replanning based on map entropy can also be considered to improve the exploration strategy and avoid planning on significantly altered maps.

(a)
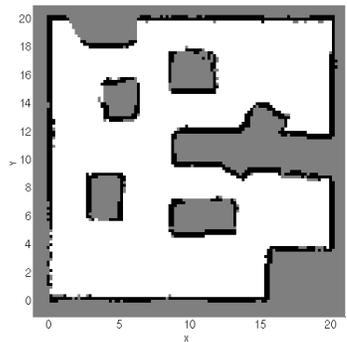


(b)

Figure 3: The Final occupancy grid maps; *a)* exploration with the grid-based method and *b)* exploration in information distribution map. Inferring the information distribution map with GP regression leads to a partially incomplete equivalent occupancy grid map. See Fig. 4 for a completed experiment with GP regression.

# 8 Results and Discussion

This section demonstrates the proposed approach through exploration simulation in the Cave[1] and Freiburg[2] maps. These maps represent varying difficulty of the environments which test the capabilities of the method.

The proposed approach is compared with a state-of-the-art grid-based method, [Valencia *et al.*, 2012], in table 1. Data are averaged over ten experiments for travel distance, exploration time, average map entropy rate, and final map relative entropy. The presented method outperforms the compared grid-based method in travel distance and map entropy rate. In spite of higher exploration time, larger absolute value for map entropy rate means fewer steps are required to complete an experi-

---
[1]Radish: The robotics dataset repository.
[2]www.slam.org.

Table 1: Comparison of the exploration methods in Cave map (averaged over 10 experiments, mean ± std)

|  | Grid-based | GP-based |
|---|---|---|
| Travel dist. (m) | 331 | 221 ± 22 |
| Expl. time (min) | 33.38 | 35.53 ± 9.39 |
| Map ent. rate (NATS/step) | -0.4438 | -0.9795 ± 0.1384 |
| Final map ent. (NATS) | 143.8696 | 148.3280 ± 1.5910 |

ment. The map entropy rate reduces 54% more than the grid-based approach, therefore, expected overall exploration time remains close.

Final map entropy for both methods is computed from the final occupancy grid map, Fig.3. For our approach occupancy grid map is computed once at the end of the experiment as the information distribution map is inherently different. The final map entropy value in our approach is higher than the compared grid-based method, however, the fact that GP regression covers regions with sparse measurements and computes full posterior of a map 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 1.

## 8.1 Maps with varying levels of detail

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 4 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. 4(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, it results in a sharper information map around an obstacle's corner which itself improves the performance of the planner and obstacle avoidance technique.

Figure 5 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 com-
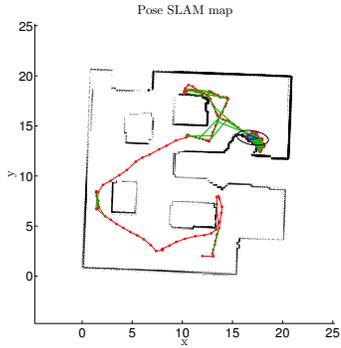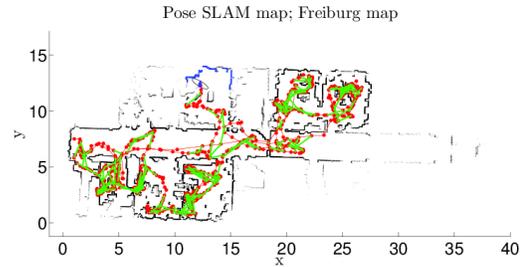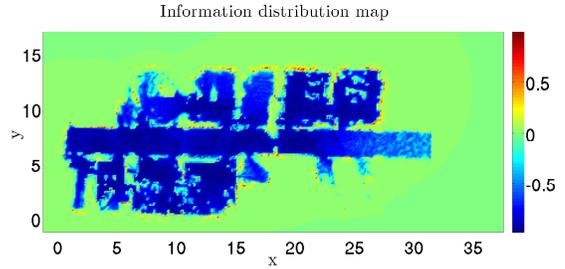
(a)
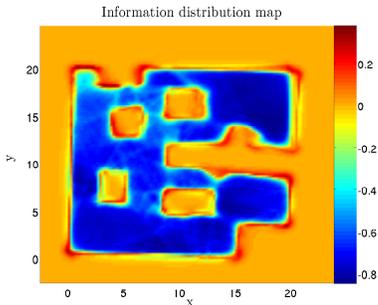


(b)

Figure 4: *a*) Exploration in Cave map, size of $20m \times 20m$. In this experiment, loop closures detected during Pose SLAM caused small rotations in the map. *b*) Final information distribution map; in information form, these sharp loop-closure changes are better absorbed in its final representation. In spite of the sparse measurements at some locations, at the end of the experiment the whole map is completely explored based on the information map.



(a)



(b)



(c)

Figure 5: 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.

putational load and requires special computing facilities. Figure 6 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.
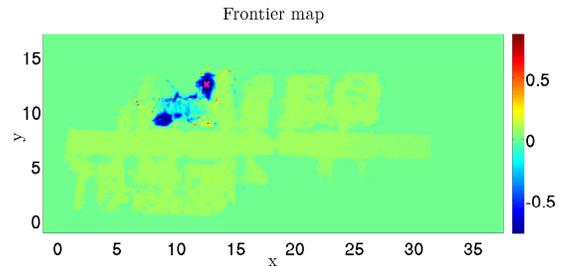
Stop criterion depends on application, but a common choice in map completion is exploring till no frontier remains. Thus, the stop criterion is $\mathcal{F} = \emptyset$ or in a more practical way $\|vec(\mathcal{F})\|_\infty < \delta$. In this research, all the presented results are achieved with $\delta = 0.09$. 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, cu-

mulative 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
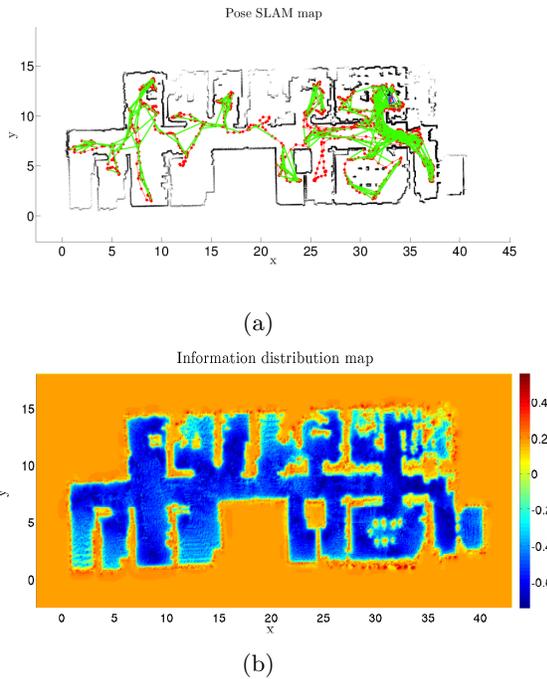
(a)



(b)

Figure 6: 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.

strategy.

GP computations have been implemented with the Open Source GP library in [Rasmussen and Williams, 2006].

## 8.2   Comparison with gridmaps

Figure 7 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. Besides the loss of information due to the discretization, such a sequence makes evident the effect of the independence assumption between cells. In Fig. 7 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. Training values for free and occupied space are appropriately propagated over the environment, thus generating well defined frontiers as shown in Fig. 2.

## 9   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.

## References

[Burgard *et al.*, 2000] 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.

[Gan *et al.*, 2009] 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.

[González-Banos and Latombe, 2002] 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.

[Ila *et al.*, 2010] V. Ila, J.M. Porta, and J. Andrade-Cetto. Information-based compact Pose SLAM. *IEEE Trans. Robot.*, 26(1):78–93, 2010.

[Juliá *et al.*, 2010] 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.

[Juliá *et al.*, 2012] M. Juliá, A. Gil, and O. Reinoso. A comparison of path planning strategies for au-
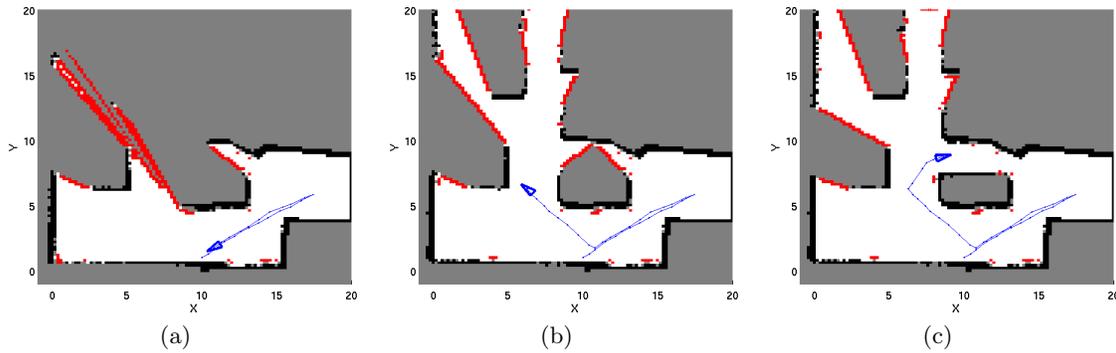
Figure 7: Three points in time during a frontier-based process using a grid map. Beyond the discretization effects, the effect of the independence assumption between cells can be seen as small frontiers appearing since information of near free and occupied cells is not propagated to the rest of the cells. Cell colours indicate free (white), unknown (gray), occupied (black), and frontier (red).

tonomous exploration and mapping of unknown environments. *Auton. Robot.*, pages 1–18, 2012.

[Kavraki *et al.*, 1996] 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.

[Lau, 2003] H. Lau. Behavioural approach for multi-robot exploration. In *Australasian Conf. Robot. Automat.*, 2003.

[Lin, 1991] J. Lin. Divergence measures based on the Shannon entropy. *IEEE Trans. Information Theory*, 37(1):145–151, 1991.

[Makarenko *et al.*, 2002] 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.

[Prestes e Silva *et al.*, 2002] 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.

[Rasmussen and Williams, 2006] C.E. Rasmussen and C.K.I. Williams. *Gaussian processes for machine learning*, volume 1. MIT press, 2006.

[Shade and Newman, 2011] 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.

[Shen *et al.*, 2012] 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.

[Stein, 1999] M.L. Stein. *Interpolation of spatial data: some theory for kriging*. Springer Verlag, 1999.

[Stocker, 2011] T. Stocker. *Introduction to climate modelling*. Springer, 2011.

[T O'Callaghan and Ramos, 2012] S. T O'Callaghan and F.T. Ramos. Gaussian process occupancy maps. *The Int. J. Robot. Res.*, 31(1):42–62, 2012.

[Thrun *et al.*, 2005] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*, volume 1. MIT press, 2005.

[Tresp, 2000] V. Tresp. A Bayesian committee machine. *Neural Computation*, 12(11):2719–2741, 2000.

[Valencia *et al.*, 2012] 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.

[Valencia *et al.*, 2013] R. Valencia, M. Morta, J. Andrade-Cetto, and J.M. Porta. Planning reliable paths with pose slam. *Robotics, IEEE Transactions on*, 29(4):1050–1059, 2013.

[Vallve and Andrade-Cetto, 2013] J. Vallve and J. Andrade-Cetto. Mobile robot exploration with potential information fields. In *6th European Conference on Mobile Robots*, 2013.

[Yamauchi, 1997] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Int. Sym. Comput. Intell. Robot. Automat.*, pages 146–151, 1997.

[Yamauchi, 1998] B. Yamauchi. Frontier-based exploration using multiple robots. In *Int. Conf. Autonomous Agents*, pages 47–53, 1998.

[Zlot *et al.*, 2002] 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.