Mutual Information-based Exploration on Continuous Occupancy Maps

Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Dissanayake

Abstract— The problem of active perception with an autonomous robot is studied in this paper. It is proposed that the exploratory behavior of the robot be controlled using mutual information (MI) surfaces between the current map and a one-step look ahead measurements. MI surfaces highlight informative areas for exploration. A novel method for computing these surfaces is described. An approach that exploits structural dependencies of the environment and handles sparse sensor measurements to build a continuous model of the environment, that can then be used to generate MI surfaces is also proposed. A gradient field of occupancy probability distribution is regressed from sensor data as a Gaussian Process and provide frontier boundaries for further exploration. The continuous global frontier surface completely describes unexplored regions and, inherently, provides an automatic termination criterion for a desired sensitivity. The results from publicly available datasets confirm an average improvement of the proposed methodology over comparable standard and state-of-the-art exploratory methods available in the literature by more than 20% and 13% in travel distance and map entropy reduction rate, respectively.

I. INTRODUCTION

Active exploration implies expanding the explored regions whilst maximizing map accuracy. Information gain-based methods [1]–[3] tend to minimize entropy-based cost functions to achieve "accuracy". However, the fact that the entropy of a random variable is a function of its distribution [4] gives rise to a major objection to the commonly employed grid-based occupancy mapping techniques [5, 6] which infer the map posterior through marginalization and ignore the structural dependencies in the environment due to the assumption of independence between cells.

In contrast, Gaussian Processes (GP) [7] learn distributions at query locations using correlation in the training set (observations) hence capturing the environment's structural dependencies. We adopt GP occupancy maps [8]–[10] to infer the map in its original high dimensional space. In this paper, a novel method for active exploration based on computation of MI surfaces is presented. The GP-based mapping and exploration technique proposed in [10] is also exploited here. However, instead of using mean values to define the frontier map, a logistic regression classifier outputs a probabilistic representation of frontiers based on the inferred map uncertainty. Since the proposed method utilizes a high-dimensional regression technique to define posterior distributions over map points and frontiers, it can alleviate the shortcoming of the occupancy grid-based exploration

M. Ghaffari Jadidi, J. Valls Miro, and G. Dissanayake are with the Centre for Autonomous Systems, Faculty of Engineering and IT, University of Technology Sydney (UTS), NSW 2007, Australia {maani.ghaffarijadidi, jaime.vallsmiro, gamini.dissanayake}@uts.edu.au
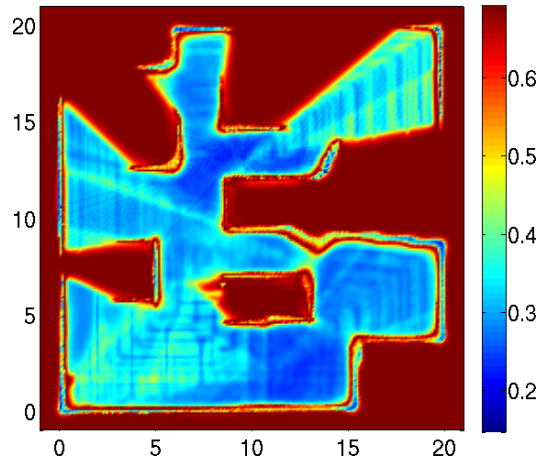
Fig. 1. An illustrative exemplar of the MI surface from the Cave dataset. The map areas beyond the current perception field of the robot preserve their initial entropy values and the higher values demonstrate regions with greater information gain. The map dimensions are in meters and the MI values in NATS. The associated continuous occupancy and probabilistic frontier maps are shown in Figure 3.

methods through capturing the uncertainty from the estimated underlying distribution. Formally speaking, entropy minimization through MI is equivalent to maximizing the Kullback-Leibler divergence of the marginal distribution of the random variables and their joint posterior which results in maximizing the information gain. As can be seen in the MI surface sample depicted by Figure 1, this representation conveniently encapsulates a cost map appropriate for information-driven robot guidance. The map shows the difference between a prior map entropy derived from an initial robot pose and map, and the conditional map entropy from future measurements. In short, the main contributions of this paper are:

1) Development of an exploration method using a probabilistic frontier representation and continuous occupancy maps able to handle sparse observations.
2) A novel method to predict the one-step look-ahead MI surface to control the exploratory behavior of the robot using a continuous map representation.

II. RELATED WORK

Frontier-oriented approaches are based on the geometry of the map and the extraction of frontiers as unknown regions in the neighborhood of known free areas [11]. Attempting to develop uncertainty minimizer algorithms, in context of the extended Kalman filter (EKF) simultaneous localization and mapping (SLAM) [12], a combined information utility

suggested in [13] using the proposed information-based cost function in [14] and an occupancy grid (OG) map MI. A local optimization was implemented for one step look-ahead optimal control action. The reported results indicated that the utility for mapping attracts the robot to unknown areas whilst the localization utility competes to maintain the robot well localized relative to known features in the map. In order to enhance the map quality of the EKF-based SLAM, in [15], an a-optimal criterion for active exploration was used. To simplify the objective function, the (feature) map covariance matrix was approximated through ignoring the correlation among the features. Furthermore, the search space of the trajectories was limited by pruning a breadth-first search algorithm to find global trajectories, leading to a tractable method.

The information gain-based exploration methods remained appealing by evolution of the SLAM algorithms. In [1], Rao-Blackwellized particle filters are employed to compute the map and robot pose posteriors. The proposed uncertainty reduction approach was based on the joint entropy minimization of the SLAM posterior. The information gain was approximated using ray casting for a given action. In practice, due to the computational complexity a particle was drawn from the particle set with a probability proportional to its weight. In a similar framework in [2], the problem of active SLAM and exploration and, specifically, the inconsistency in the filter due to the information loss for a given policy using the relative entropy concept was addressed. In [3], with the assumption that all random variables are normally distributed and employing the relative entropy metric, the combined traveling cost and expected information gain was developed.

The techniques in [16, 17] evaluate exploratory and place revisiting paths, which are selected based on entropy reduction estimates of both map and path. Whilst the map entropy is computed on an occupancy grid at coarse resolution, path entropy is the outcome of Pose SLAM [18, 19], a delayed-state SLAM algorithm from the pose graph family. Given the inherent complexity in the formulation to calculate joint entropy between robot pose and map, conditional independence is assumed. In [20], the MI surface between a map and future measurements was computed numerically. The work assumes known robot poses, and relies on an OG map representation and measurements from a laser range-finder. The algorithm integrates over an information gain function with an inverse sensor model at its core. It was formally proven that any controller tasked to maximize an MI reward function is eventually attracted to unexplored areas, thus providing a strong basis for the MI-based study of autonomous robotic exploration hereby proposed. The novel solution is distinctly flexible in being able to deal with sparse measurements and predicated on inferring a map posterior using Bayesian updates with a forward sensor model, and a continuous representation to define informative priors in the estimation of conditional entropy (MI) maps. Moreover, the work is thoroughly demonstrated with examples under realistic conditions.
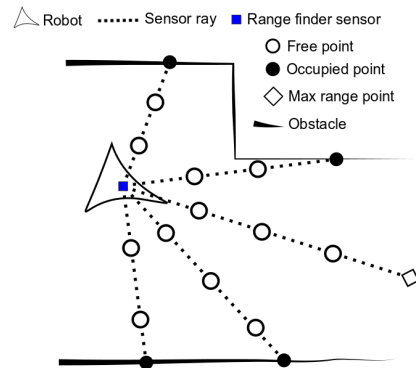

Fig. 2. Environment's setup and the concept of generating training data from range-finder beams can be seen in this picture. The negative information from each beam provide required training data for free areas which is captured through line segmentation.

III. GAUSSIAN PROCESSES MAPPING AND EXPLORATION

We employ continuous occupancy mapping technique which benefits from GPs to learn correlation between map points and infer the map in its original high dimensional space. For a given query point in the map GP predicts a mean, μ, and an associated variance, σ. The target value can be $y_+ = 1$ and $y_- = -1$ for occupied and free points, respectively, depending on the nature of the map, i.e. a binary classification problem. Optimization of the hyper-parameters is performed once at the beginning of each experiment by minimization of the negative log of the marginal likelihood function. For the prevailing case of multiple runs in the same environment, the optimized values can then be loaded off-line. In general, there is no guarantee that predicted GPs mean values be constrained to the target values. A logistic function squashes the predicted outputs into the range $(0, 1)$ and guarantees a valid probabilistic interpretation. By enquiring over a uniformly sampled range of points, we fuse the obstacle, m_o, and the free area, m_f, maps to achieve the desired continuous occupancy map (COM), m, as thoroughly explained in [10]. In order to incorporate new information incrementally, map updates are performed using the Bayesian Committee Machine (BCM) technique [21]. Figure 2 shows the environment's setup and training points generation concept. Note that query points can have any desired distributions. However building the map over a grid, as done in Section IV, facilitates comparison with standard occupancy grid-based methods, i.e. at similar map resolutions.

In information gain-based exploration the commonly used utility function is defined to maximize the MI between the state and new measurements. The expectation over new sets of measurements and actions gives a path and a goal which is considered as the optimal behavior. The underlying process involves moving the robot towards a candidate goal, collect a set of measurements, and compute their likelihood. The widely-employed approach to approximate the expected information gain is using an inverse sensor model through ray

Algorithm 1 build_frontier_map($m, m_o, \sigma, \beta, \gamma$)

1: $dm, dm_o \leftarrow$ Compute $\|\nabla m\|_1$ and $\|\nabla m_o\|_1$ numerically
2: $\sigma_{\min} \leftarrow \min(\sigma)$
3: $f \leftarrow zeros(size(m))$
4: **for all** $i \in \mathcal{M}$ **do**
5: $\quad \bar{f}^{[i]} \leftarrow dm^{[i]} - \beta(dm_o^{[i]} + m_o^{[i]} - 0.5)$
6: $\quad w^{[i]} \leftarrow \gamma\sqrt{\sigma_{\min}/\sigma^{[i]}}$
7: $\quad f^{[i]} \leftarrow (1 + \exp(-w^{[i]}\bar{f}^{[i]}))^{-1}$
8: **end for**
9: **return** f

Algorithm 2 build_MI_map(m)

Require: Current robot pose, sensor model;
1: $\bar{m} \leftarrow m$
2: $I \leftarrow -(m\log(m) + (1-m)\log(1-m))$
3: **for all** k **do**
4: $\quad$ Compute $\hat{z}_{t+1}^{[k]}$ and $\mathcal{I}_{t+1}^{[k]}$ using ray casting in m
5: $\quad$ **for** $i \in \mathcal{I}_{t+1}^{[k]}$ **do**
6: $\quad\quad \bar{h} \leftarrow 0$
7: $\quad\quad$ **for all** $z \leq \hat{z}_{t+1}^{[k]}$ **do**
8: $\quad\quad\quad p_1 \leftarrow p(\hat{z}_{t+1}^{[k]}|M=0)$
9: $\quad\quad\quad p_2 \leftarrow 0$
10: $\quad\quad\quad$ **for** $j \in \mathcal{I}_{t+1}^{[k]}$ **do**
11: $\quad\quad\quad\quad p_1 \leftarrow p_1 \times (1 - m^{[j]})$
12: $\quad\quad\quad\quad p_2 \leftarrow p_2 + p(\hat{z}_{t+1}^{[k]}|M=m^{[j]})m^{[j]}\prod_{l<j}(1-m^{[l]})$
13: $\quad\quad\quad$ **end for**
14: $\quad\quad\quad p_z \leftarrow p_1 + p_2$
15: $\quad\quad\quad \bar{m}^{[i]} \leftarrow p_z^{-1}p(\hat{z}_{t+1}^{[k]}|M=m_i)m^{[i]}\prod_{l<i}(1-m^{[l]})$
16: $\quad\quad\quad \bar{h} \leftarrow \bar{h} + pz[\bar{m}^{[i]}\log(\bar{m}^{[i]}) + (1-\bar{m}^{[i]})\log(1-\bar{m}^{[i]})]$
17: $\quad\quad$ **end for**
18: $\quad\quad I^{[i]} \leftarrow I^{[i]} + \bar{h}N_z^{-1}$
19: $\quad$ **end for**
20: **end for**
21: **return** I

---

casting operation in OG maps [1, 16, 20, 22, 23]. A number of policies define the robot path to simulate the action and acquire new poses and measurements over possible informative areas. Instead, we suggest to compute the MI of the map and future measurements over all areas in the current robot's perception field using a continuous representation of the map. We estimate the map posterior through the Bayes update formula and compute the map conditional entropy numerically. The COM technique infers a joint predictive distribution over the map, therefore, the computed entropy is more descriptive of the real map uncertainty (see fig. 1 in [20] for an illustration of map entropy vs. MI in OGs).

A. Frontier map

Constructing a frontier map is the fundamental ingredient of any geometry based exploration approach. Since it reveals the boundaries between known-free and unknown areas which are potentially informative regions, it is also applicable in information gain-based exploration strategies. From [10], given any location i in the map, the value of a frontier can be computed as

$$\bar{f}^{[i]} = \|\nabla m^{[i]}\|_1 - \beta(\|\nabla m_o^{[i]}\|_1 + m_o^{[i]} - 0.5), \quad (1)$$

where ∇ denotes the gradient operator, and β is a factor that controls the effect of obstacle boundaries. $\|\nabla m^{[i]}\|_1$ indicates
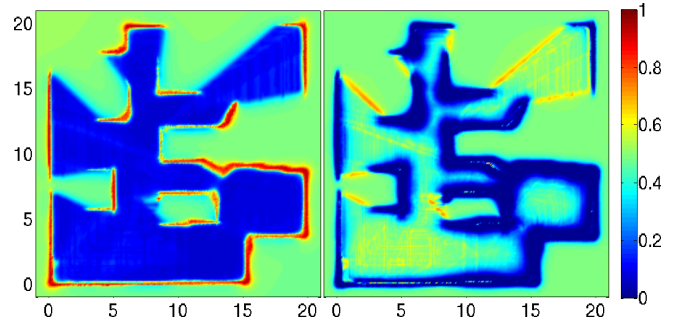

Fig. 3. Probabilistic frontier map (right) computed from COM (left). The map highlights the informative regions for further exploration by defining distributions over frontier points. The lower probabilities show the obstacles and walls whilst the values greater than the no discrimination probability, 0.5, can be considered as frontiers.

all boundaries whilst $\|\nabla m_o^{[i]}\|_1$ defines obstacle outlines. The subtracted constant is to remove the biased probability for unknown areas in the obstacles probability map.

A probabilistic representation of the frontier map is highly favorable. It facilitates interpretation of the frontier concept and provides a meaningful way to set an automatic termination criterion for exploration experiments. Thus, the mean frontier surface proposed in [10] is converted to a probability frontier map incorporating the map uncertainty. In order to squash mean and variance values into the interval $(0, 1)$, a logistic regression classifier with inputs from $\bar{f}^{[i]}$ and the map uncertainty $\sigma^{[i]}$ is applied to data which yields

$$p(f^{[i]}|m^{[i]}, w^{[i]}) = \frac{1}{1 + \exp(-w^{[i]}\bar{f}^{[i]})} \quad (2)$$

where $w^{[i]} = \gamma\sqrt{\lambda^{[i]}}$ denotes the required weights, $\lambda^{[i]} = \sigma_{\min}/\sigma^{[i]}$ is the bounded information associated to location i, and $\gamma > 0$ is a constant to control the sigmoid shape. The details of the frontier map computations are presented in Algorithm 1. Figure 3 depicts an instance of the frontier map from an exploration experiment in the Cave environment. The aforementioned representation suggests a compatible interface with other modules of our Bayesian approach. In contrast to classical binary representation, defining frontiers in probabilistic terms as distributions over informative areas using map uncertainty appears more fitting to compute expected behaviors.

B. Mutual information map

The MI is defined as the reduction in uncertainty due to the another random variable [4]. In other words, assuming robot poses are known, for an exploratory action given a measurement $Z = z$ from $\mathcal{Z} \subset \mathbb{R}_{\geq 0}$ we are interested to know what will be the reduction in the map $M = m$ uncertainty. Therefore, the expected MI is applicable in active exploration where given a future measurement the utility function tries to maximize the uncertainty reduction. The MI between the map and the future measurement $Z_{t+1} = \hat{z}$ is

$$I(M; Z_{t+1}|z_{1:t})$$
$$= \int_{\hat{z} \in \mathcal{Z}} \sum_{m \in \mathcal{M}} p(m, \hat{z}|z_{1:t}) \log \frac{p(m, \hat{z}|z_{1:t})}{p(m|z_{1:t})p(\hat{z}|z_{1:t})} d\hat{z}$$
$$= H(M|z_{1:t}) - \overline{H}(M|Z_{t+1}, z_{1:t}), \quad (3)$$

TABLE I

COMPARISON OF THE DIFFERENT EXPLORATION STRATEGIES IN THE INDOOR DATASETS (AVERAGED OVER 30 EXPERIMENTS, MEAN $\pm$ STANDARD ERROR)

Cave environment 20m $\times$ 20m; map resolution: 0.2m; frontier map threshold: 0.59				
	NF	APS	GPNF	GPMI
Travel distance (m)	109.12 $\pm$ 4.72	317.83 $\pm$ 0.40	108.60 $\pm$ 7.64	**100.54 $\pm$ 3.46**
Exploration time (min)	14.50 $\pm$ 0.48	18.58 $\pm$ 0.28	**5.40 $\pm$ 0.43**	11.25 $\pm$ 0.38
Map ent. rate (NATS/step)	-14.1576 $\pm$ 0.4683	-10.2131 $\pm$ 0.1500	-15.1586 $\pm$ 0.4250	**-16.5754 $\pm$ 0.4905**
Localization MSE (m)	0.3392 $\pm$ 0.0629	0.1234 $\pm$ 0.0348	0.0930 $\pm$ 0.0189	**0.0784 $\pm$ 0.0135**
Number of closed loops	17.10 $\pm$ 2.17	**157.30 $\pm$ 1.28**	31.30 $\pm$ 5.08	25.60 $\pm$ 2.11
Freiburg environment 40m $\times$ 15m; map resolution: 0.2m; frontier map threshold: 0.57				
	NF	APS	GPNF	GPMI
Travel distance (m)	265.26 $\pm$ 10.59	n/a	170.37 $\pm$ 6.20	**154.46 $\pm$ 5.39**
Exploration time (min)	57.29 $\pm$ 1.25	n/a	18.23 $\pm$ 0.87	**15.13 $\pm$ 0.90**
Map ent. rate (NATS/step)	-8.2347 $\pm$ 0.1582	n/a	-12.6333 $\pm$ 0.3509	**-13.0932 $\pm$ 0.3100**
Localization MSE (m)	0.6135 $\pm$ 0.1432	n/a	0.3257 $\pm$ 0.0470	**0.2035 $\pm$ 0.0256**
Number of closed loops	**17.00 $\pm$ 1.11**	n/a	9.07 $\pm$ 0.64	7.87 $\pm$ 0.60

where $H(M|z_{1:t})$ and $\overline{H}(M|Z_{t+1}, z_{1:t})$ are map and map conditional entropy respectively, which by definition are

$$H(M|z_{1:t}) = - \sum_{m \in \mathcal{M}} p(m|z_{1:t}) \log p(m|z_{1:t}) \qquad (4)$$

$$\overline{H}(M|Z_{t+1}, z_{1:t}) = \int_{\hat{z} \in \mathcal{Z}} p(\hat{z}|z_{1:t}) H(M|Z_{t+1} = \hat{z}, z_{1:t}) d\hat{z} \qquad (5)$$

The exact solution requires computation of the conditional predictive distribution for each map point. This involves N_q (number of query points) times inversion of the corresponding $(N_q - 1) \times (N_q - 1)$ covariance matrix. A more tractable alternative is using the marginal predictive distribution through marginalizing the GP's latent function (joint predictive distribution). Although this simplification equates to independence between map points for prediction purpose, it should be noted that the current estimate is computed by considering full correlations between map points.

To compute the map conditional entropy, the predicted map posterior given the new measurement $Z_{t+1} = \hat{z}_{t+1}$ is required. The Bayesian inference finds the posterior probability for each map point $m^{[i]}$ and k-th beam of the range-finder as

$$p(m^{[i]}|\hat{z}_{t+1}^{[k]}, z_{1:t}) = \frac{p(\hat{z}_{t+1}^{[k]}|m^{[i]}) p(m^{[i]}|z_{1:t})}{p(\hat{z}_{t+1}^{[k]}|z_{1:t})} \qquad (6)$$

$$p(\hat{z}_{t+1}^{[k]}|z_{1:t}) = \sum_{m^{[i]} \in \mathcal{M}} p(\hat{z}_{t+1}^{[k]}|m^{[i]}) p(m^{[i]}|z_{1:t}) \qquad (7)$$

The likelihood function $p(\hat{z}_{t+1}^{[k]}|M = m^{[i]})$ is a beam-based mixture measurement model [24], where the term $p(\hat{z}_{t+1}^{[k]}|M = 0)$ can be interpreted as the likelihood of not observing the map point at location i, i.e. uniform distribution, and the term $p(\hat{z}_{t+1}^{[k]}|z_{1:t})$ is the marginal distribution over measurements which is calculated in line 14 of Algorithm 2 and denoted by p_z. By numerically integrating over a desired beam range, we can compute the predicted map posterior entropy in Equation 5. For all points in the k-th sensor beam's perception field at the current robot location, $\mathcal{I}_{t+1}^{[k]}$, the MI can be written as

$$I^{[i]} = h(m^{[i]}) - \overline{h}(m^{[i]}) \qquad (8)$$

where $h(m^{[i]})$ is the current entropy of the map point $m^{[i]}$ and $\overline{h}(m^{[i]})$ is the estimated map conditional entropy. In Figure 1, an estimated MI map during an exploration experiments in Cave environment is depicted. In practice, at each time step, the map is initialized with the current map entropy, $H(M|z_{1:t})$, and for all map points inside the current perception field the estimated map conditional entropy values are subtracted from the corresponding initial values. In Algorithm 2, the implementation of the MI map is given where N_z denotes the resolution of numerical integration.

C. Decision making

The resulting MI map shows the expectation for uncertainty reduction in the map at each place. In order to define a utility function, the frontier map is initially thresholded and, through k-means, clusters of geometric frontiers are extracted. Let each geometric frontier be regarded as an action from the exploration point of view. The action space can thus be defined as $\mathcal{A} = \{a^{[j]}\}_{j=1}^n$. We define a utility function as the difference between information gain reward and traversal distance. The optimal action will therefore be the one that maximizes the utility function. Therefore,

$$a^\star = \underset{a \in \mathcal{A}}{\operatorname{argmax}} \{\alpha \overline{I}(a) - d(a)^{1/2}\} \qquad (9)$$

where $d(a)$ is the direct distance from the current robot pose to the exploration action (squared root to prevent steep variations), $\overline{I}(a)$ is the mean MI for an exploration action, and α is a factor to relate information gain to the cost of motion.

IV. RESULTS AND DISCUSSION

In this section the exploration results from extensive scenarios are presented. The indoor experiments are focusing on comparison of the presented method with comparable exploratory methods available in the literature, whereas the outdoor experiment demonstrates the scalability. An average improvement over the standard and state-of-the-art exploratory methods by more than 20% and 13% in travel distance and map entropy reduction rate, respectively, while
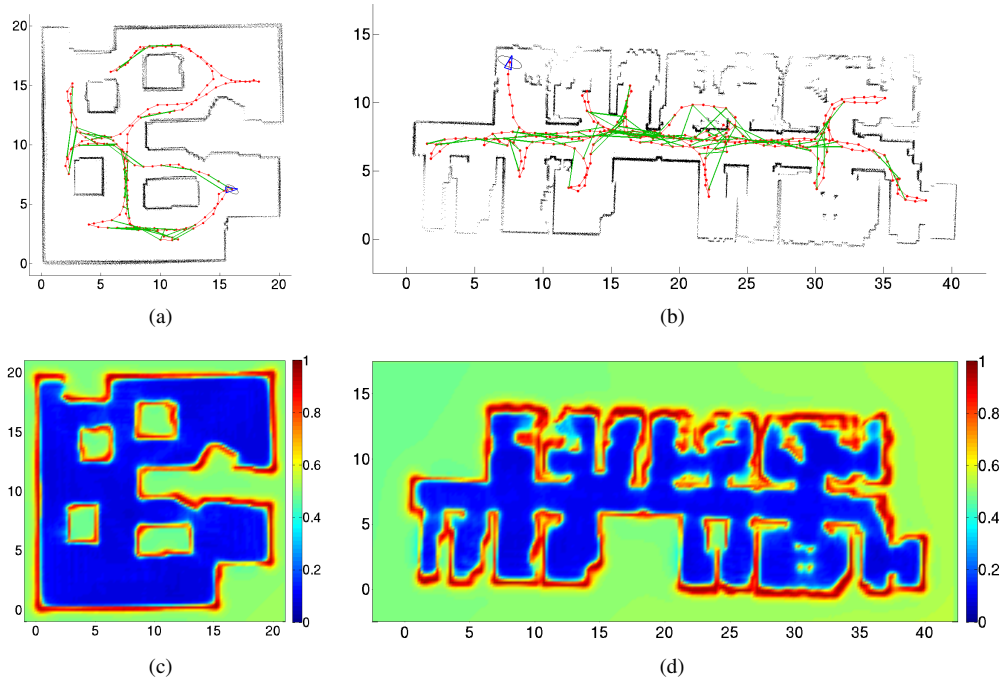
Fig. 4. MI-based exploration in Cave and Freiburg environments. (a) and (b) Pose SLAM maps, (c) and (d) continuous occupancy maps. In (a) and (b), red curves are the robot path and green lines indicate loop-closures. The continuous occupancy maps show the occupancy probability at each location where mid-probability value of 0.5 represents unknown points. Map dimensions are in meters.

maintaining the localization mean square error (MSE) 36% lower, can be clearly seen in Tables I and II. The localization MSE was computed at the end of each experiment by the difference between the robot traveled path, i.e. estimated poses and ground truth poses. The required parameters for the beam-based mixture measurement model, frontier maps, and MI maps computations are listed in Table III.

The experiments include exhaustive comparison among the original nearest frontier (NF) [11], active Pose SLAM (APS) [16], COM-based nearest frontier (GPNF) which is considered as the immediate outcome in [10], and the developed MI-based (GPMI) exploration approaches. In order to maintain NF and APS methods in their original frameworks, the results are computed using OG maps while in COM-based methods the developed COM in [10] and the probabilistic frontier map in this paper are employed. In all of the presented results, Pose SLAM [18] is included as the backbone to provide localization data together with number of closed loops. Additionally, regardless of the exploration method, the same set of Pose SLAM parameters was used in each environment. The simulated robot and its sensors, i.e. the odometric and laser range-finder, provide required sensory inputs for Pose SLAM and the environment is built by loading a binary image. The odometric and laser range-finder sensors noise covariances were set to $\Sigma_u = \text{diag}(0.1\text{m}, 0.1\text{m}, 0.0026\text{rad})^2$ and $\Sigma_y = \text{diag}(0.05\text{m}, 0.05\text{m}, 0.0017\text{rad})^2$, respectively. The robot started the experiments with an initial pose uncertainty of $\Sigma_0 = \text{diag}(0.1\text{m}, 0.1\text{m}, 0.09\text{rad})^2$ and laser beams were simulated through ray-casting operation over the ground truth map using the true robot pose.

The proposed approach is demonstrated with exploratory simulations in three mapping environments. The indoor datasets include the Cave and Freiburg maps [25] and the outdoor dataset is a parking area as an experiment in a larger scale (16 times on average in area) to demonstrate the scalability of the presented approach. The Cave map represents a simple hand-drawn environment with a few rough obstacles, whereas the Freiburg map is computed from a real LMS-laser log taken with a Pioneer2 robot at the University of Freiburg (building 079 AIS-Lab), and contains many rooms and disconnected obstacles which make for a challenging environment to explore. The implementation has been developed in MATLAB (2014a) and GP computations have been implemented using the open source GP library in [7].

A. Indoor mapping and exploration

Table I summarizes the exploration results in the indoor environments. The figures are averaged over 30 repetitions of the experiments in order to demonstrate repeatability and reliability of each method. To achieve the best outcome for all of the compared methods, in the nearest frontier methods (NF and GPNF), frontiers closer than 2m to the current robot pose are ignored. Furthermore, only frontiers with the size larger than 10 cells are considered valid in NF and APS. For GPNF and GPMI, a threshold was set for the frontier map as the termination condition of an experiment. These conditions lead to avoiding excessive search and fluctuation in a small area and exploring the whole map faster.

Despite the common belief in the literature, classic NF performs sensible in terms of minimum requirements for an
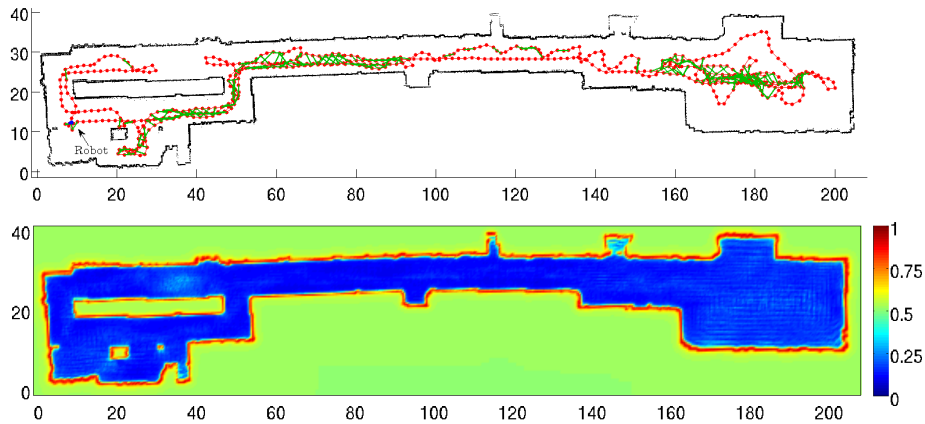
Fig. 5. MI-based exploration in an outdoor parking area. This map, on average, is 16 times larger than the indoor maps in area and demonstrates the scalability of the presented method. The lack of distinctive features in many region makes closing informative loops extremely challenging which can easily lead to the localization inconsistency in such a large environment. Map dimensions are in meters.

exploration method. However, NF does not contemplate any form of uncertainty reduction which explains the highest localization error in Table I. APS tries to minimize the approximated joint entropy of the map and robot pose which results in completion of the experiment with a lower localization error. However, by aiming for the most informative goal together with the replanning policy, according to the utility function, the robot may fluctuate among several areas of the map. The increased travel distance and significantly higher number of closed loops imply over-visiting of some already known areas. The results show that grid-based mapping does not give sufficient support for information-gain based exploration as the robot may travel a long distance to approach an informative area and meanwhile collect sparse measurement along its way. This is where importance of GP mapping becomes highlighted through handling sparse measurements and fast map entropy reduction. In other words, in many cases, the robot does not need to revisit partially observed areas and the correlation between points in the map helps in reducing the map entropy at a faster rate.

GPNF and GPMI exploit COM for mapping, exploration, and planning providing a common ground for both travel distance and uncertainty reductions. GP-based methods handle sparse sensor measurements and by learning the environment's structural dependencies define distribution over the map points occupancy and also frontier points. The significant increase in the map entropy reduction rate while maintaining the localization error low is due to this fact and can be concluded from the data in Table I. GPMI through the developed utility function tries to maximize the MI between map and future measurements. It demonstrates a better long term decision making, since it has the fastest map entropy reduction rate as the MI map suggests. The developed MI surface is a suitable cost map for planners such as A^* as demonstrated in the results. Although the map is one-step look-ahead it is assumed to be fixed during the planning time. Furthermore, traveling through uncertain traversable regions facilitates closing informative loops in the SLAM process by taking overlapping observations from diverse orientations.

Outdoor parking environment 210m × 42m; map resolution: 0.4m frontier map threshold: 0.53		
	GPNF	GPMI
Travel distance (m)	634.44 ± 36.54	**594.94 ± 53.53**
Exploration time (min)	**20.51 ± 1.00**	38.65 ± 2.17
Map ent. rate (NATS/step)	-72.8809 ± 4.0572	**-74.1553 ± 3.8043**
Localization MSE (m)	2.5288 ± 0.4777	**1.7357 ± 0.4664**
Number of closed loops	**34.50 ± 4.83**	24.10 ± 8.89

Figure 4 illustrates the indoor exploration results with GPMI method in Cave and Freiburg datasets.

B. Outdoor dataset

The outdoor scenario consists of a large parking area, approximately 16 times larger than the earlier indoor setting. Figure 5 shows the exploration results with the GPMI method. The environment is not challenging only due to its size, but also for the long sections with few distinguishing features. Closing informative loops is non-trivial and moving towards frontiers as they emerge along the long straight section in the middle of the map significantly increases robot pose uncertainty. Table II presents the comparative results of running GPNF and GPMI in this dataset. It can be seen how the latter performs slightly better in regards to travel distance, map entropy rate, and localization error, while computational cost is slightly worse. A longer run experiment such as this reveals how the lack of an explicit utility term for loop-closure actions results in statistically closer figures. In exploratory problems future state and measurements are unknown. As such, a method such as GPMI with an integrated outlook on uncertainty minimization can be a natural companion to a decision making framework that more closely incorporates SLAM outcomes within the planning process. This is not considered in this paper and has been left for future work.

TABLE III

PARAMETERS IN THE PRESENTED ALGORITHMS FOR FRONTIER AND MI
MAPS COMPUTATIONS.

Parameter	Symbol	Value
1) Beam-based mixture measurement model:		
Hit std	σ_{hit}	0.05 m
Short decay	λ_{short}	0.2 m
Max range	r_{max}	10.00 m
Hit weight	z_{hit}	0.7
Short weight	z_{short}	0.1
Max weight	z_{max}	0.1
Random weight	z_{rand}	0.1
2) Frontier map:		
Occupied boundaries factor	β	3.0
Logistic regression weight	γ	10.0
3) MI map and utility function:		
Number of sensor beams over 360 deg	B_z	266
Numerical integration resolution	N_z	10 m^{-1}
Information gain factor	α	15
Occupied probability threshold	p_o	0.65
Unoccupied probability threshold	p_f	0.35

C. Computational complexity

The computational cost of GPs due to the need to invert a matrix of the size of training data N_t is $O(N_t^3)$. For MI surface, the time complexity is at worst quadratic in the number of map points in the current perception field, N_p, and linear in the number of sensor beams, B_z, and numerical integration's resolution, N_z, respectively, resulting in $O(N_p^2 B_z N_z)$.

V. CONCLUSION

A novel solution to the active perception problem with a mobile robot has been presented in this paper. The exploratory behavior of the robot is controlled using MI surfaces between the current map and a one-step look ahead measurements. It is based on learning spatial correlation of map points with iterative GP-based regression from sparse range measurements, and computing MI surfaces estimation from a one-step ahead map posterior and conditional entropy. The proposed exploration scheme exploits these Bayesian updates to compute the most likely frontiers for further inspection. The strategy has been contrasted against comparable exploratory methods available in the literature with improvements of more than 20% in travel distance and 13% in map entropy reduction rate. The methodology has thus been proven capable of dense uncertainty minimization, also in large-scale settings, while maintaining mean square localization errors along the traversed path on average 36% lower than comparable methods. GPs predictive distributions over map points and frontiers provide necessary inputs for MI computation, which truly depends on the distribution of the defined random variables, resulting in fast entropy reduction.

REFERENCES

[1] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using rao-blackwellized particle filters." in *Robotics: Science and Systems*, vol. 2, 2005.

[2] L. Carlone, J. Du, M. K. Ng, B. Bona, and M. Indri, "An application of kullback-leibler divergence to active SLAM and exploration with particle filters," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2010, pp. 287–293.

[3] F. Amigoni and V. Caglioti, "An information-based exploration strategy for environment mapping with mobile robots," *Robot. Auton. Syst.*, vol. 58, no. 5, pp. 684–699, 2010.

[4] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 1991.

[5] H. P. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *Proc. IEEE Int. Conf. Robot Automat.*, vol. 2. IEEE, 1985, pp. 116–121.

[6] A. Elfes, "Sonar-based real-world mapping and navigation," *Robot. and Automat., IEEE Journal of*, vol. 3, no. 3, pp. 249–265, 1987.

[7] C. Rasmussen and C. Williams, *Gaussian processes for machine learning*. MIT press, 2006, vol. 1.

[8] S. T O'Callaghan and F. Ramos, "Gaussian process occupancy maps," *The Int. J. Robot. Res.*, vol. 31, no. 1, pp. 42–62, 2012.

[9] S. Kim and J. Kim, "Continuous occupancy maps using overlapping local Gaussian processes," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2013, pp. 4709–4714.

[10] M. Ghaffari Jadidi, J. V. Miró, R. Valencia, and J. Andrade-Cetto, "Exploration on continuous Gaussian process frontier maps," in *Proc. IEEE Int. Conf. Robot Automat.*, 2014, pp. 6077–6082.

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

[12] M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem," *Robot. and Automat., IEEE Transactions on*, vol. 17, no. 3, pp. 229–241, 2001.

[13] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information based adaptive robotic exploration," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, vol. 1. IEEE, 2002, pp. 540–545.

[14] H. J. S. Feder, J. J. Leonard, and C. M. Smith, "Adaptive mobile robot navigation and mapping," *The Int. J. Robot. Res.*, vol. 18, no. 7, pp. 650–668, 1999.

[15] R. Sim and N. Roy, "Global a-optimal robot exploration in SLAM," in *Proc. IEEE Int. Conf. Robot Automat.* IEEE, 2005, pp. 661–666.

[16] R. Valencia, J. Valls Miro, G. Dissanayake, and J. Andrade-Cetto, "Active Pose SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 1885–1891.

[17] J. Vallve and J. Andrade-Cetto, "Dense entropy decrease estimation for mobile robot exploration," in *Proc. IEEE Int. Conf. Robot Automat.*, 2014, pp. 6083–6089.

[18] V. Ila, J. Porta, and J. Andrade-Cetto, "Information-based compact Pose SLAM," *IEEE Trans. Robot.*, vol. 26, no. 1, pp. 78–93, 2010.

[19] R. Valencia, M. Morta, J. Andrade-Cetto, and J. Porta, "Planning reliable paths with Pose SLAM," *IEEE Trans. Robot.*, vol. 29, no. 4, pp. 1050–1059, 2013.

[20] B. J. Julian, S. Karaman, and D. Rus, "On mutual information-based control of range sensing robots for mapping applications," *The Int. J. Robot. Res.*, vol. 33, no. 10, pp. 1375–1392, 2014.

[21] V. Tresp, "A Bayesian committee machine," *Neural Computation*, vol. 12, no. 11, pp. 2719–2741, 2000.

[22] A. Makarenko, S. Williams, F. Bourgault, and H. Durrant-Whyte, "An experiment in integrated exploration," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, vol. 1, 2002, pp. 534–539.

[23] R. Sim and J. J. Little, "Autonomous vision-based exploration and mapping using hybrid maps and rao-blackwellised particle filters," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2006, pp. 2082–2089.

[24] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005, vol. 1.

[25] A. Howard and N. Roy, "The robotics data set repository (Radish)," http://radish.sourceforge.net, 2003.