2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

On mutual information-based control of range sensing robots for mapping applications

Brian J. Julian[1,2], Sertac Karaman[3], and Daniela Rus[1]

Abstract— **In this paper we examine the correlation between the information content and the spatial realization of range measurements taken by a mapping robot. To do so, we consider the task of constructing an occupancy grid map with a binary Bayesian filter. Using a narrow beam-based sensor model (versus an additive white Gaussian noise model), we prove that any controller tasked to maximize a mutual information reward function is eventually attracted to unexplored space. This intuitive behavior is derived solely from the geometric dependencies of the occupancy grid mapping algorithm and the monotonic properties of mutual information. Since it is a function of both the robot's position and the uncertainty of the surrounding cells, mutual information encodes geometric relationships that are fundamental to robot control, thus yielding geometrically relevant reward surfaces on which the robot can navigate. Lastly, we present the results of two experiments employing an omnidirectional ground robot equipped with a laser rangefinder.**

I. Introduction

For occupancy grid mapping, the most successful robot control algorithms are grounded in geometric-based principles. Frontier detection, or the identification of boundaries between unknown and unoccupied space [1], drives many exploration algorithms, while path planning algorithms considering these frontiers often employ spatial metrics and heuristics [2]. This approach is not surprising given that the mapping algorithm's output is a representation of shapes, sizes, and positions of obstacles within the environment. However, underneath the mask of maximum likelihood, a map is fundamentally a field of binary random variables [3], and a sensor is a probabilistic channel that links robot motion in the physical world to information gain for the Bayesian inference. One should expect that geometric-based intuition agrees with information-based reasoning concerning robot control, however, this agreement up till now is at most a conjecture.

Thus, our goal in this paper is to rigorously characterize how the robot's motion in the physical world affects the relevance of its sensor measurements with respect to occupancy grid mapping. This goal goes beyond identifying which of the map's grid cells correspond to the largest amount of uncertainty, i.e., *entropy* [4]. We instead consider *mutual information* between the map and future sensor measurements to be the main reward for information-based control. Mutual information is an information-theoretic quantity [5] that, in occupancy grid mapping applications, predicts how much future measurements will decrease the robot's uncertainty associated with all grid cells [4]. Since it is a function of both the robot's position and the uncertainty of the surrounding cells, mutual information encodes geometric relationships that are fundamental to robot control, thus yielding geometrically relevant reward surfaces on which the robot can navigate, e.g., Figure 1.

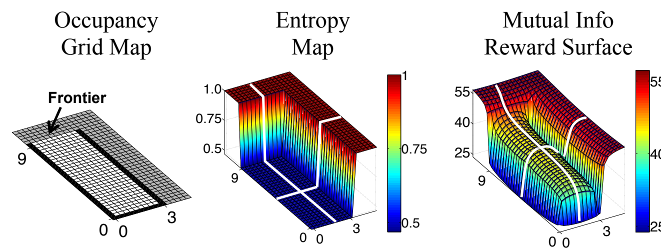

Fig. 1. A 3 m by 9 m hallway "dead end" with three walls and one open frontier. *Left:* The occupancy grid map with occupied (black), unoccupied (white), and unknown (gray) cells. Surrounding cells not shown are assumed to be unknown. *Middle:* A drawing of the the resulting entropy map where the open space is indistinguishable from the walls and, more importantly, the open frontier is indistinguishable from boundaries between occupied and unknown cells. *Right:* A more geometrically relevant mutual information reward surface that guides a robot away from the three walls and towards the open frontier at the top left. Note that the white curves highlighting the entropy map's and mutual information reward surface's geometries are for illustrative purposes, and that the two linear color scales represent different bit ranges.

To this effect, the main result of this paper provides a geometric interpretation for the mutual information-based reasoning of a robot using the occupancy grid mapping algorithm and a narrow beam-based range sensor. More specifically, we prove that any controller tasked to maximize a mutual information reward function is eventually[1] attracted to unexplored space.[2] Remarkably, this theoretical finding is independent of the implemented control algorithm. Sampling-based, gradient-based, and potential field methods all exhibit attractive behavior as long as they aim to

[1]Brian J. Julian and Daniela Rus are with the Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139, USA, bjulian@mit.edu and rus@csail.mit.edu

[2]Brian J. Julian is also with MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420, USA

[3]Sertac Karaman is with the Department of Aeronautics and Astronautics, MIT, Cambridge, MA 02139, USA, sertac@mit.edu

This work is sponsored by the Department of the Air Force under Air Force contract number FA8721-05-C-0002. The opinions, interpretations, recommendations, and conclusions are those of the authors and are not necessarily endorsed by the United States Government.

This work is also supported by the ONR-MURI grant N00014-09-1-1051, the ONR grant N00014-12-1-1000, the SMART Future of Urban Mobility project, and MIT Lincoln Laboratory.

[1]By *eventually* we mean that this attraction does not need to be initially true, but instead it will be true at and after some time in the future.

[2]By *attracted to unexplored space* we mean that unknown cells within the map yield a higher mutual information reward compared to the hypothetical situation where these cells are correctly known.

978-1-4673-6357-0/13/$31.00 ©2013 IEEE

maximize a mutual information reward function. To support these claims, we present the results of two experiments employing an omnidirectional ground robot equipped with a laser rangefinder.

A. Related Works

The seminal work of [3] introduced the concept of occupancy grid maps as a space representation model. Since then, much research has focused on constructing maps in higher dimensions [6], fusing measurements from multiple sensors/robots [7], [8], and simultaneously/concurrently localizing the robot [9], [10]. Unlike these cited works, the purpose of this paper is not to propose new occupancy grid mapping algorithms, but instead to conduct information-theoretic analyses using these preexisting algorithms for the underlying Bayesian inference framework. For example, we evaluate our theoretical results in Section IV-B using a Monte Carlo-based SLAM algorithm [11] that employs the occupancy grid mapping algorithm. We believe that our analyses can be applied to many other Bayesian derived mapping algorithms where sensor noise depends on the robot's position in the environment, such as those relying solely on object-based [12] or landmark-based [13] maps.

The task of maximizing a mutual information reward function has recently emerged as a powerful control objective to improve the quality of the occupancy grid map. The early work of [4] proposed employing mutual information between sensor observations and the occupancy grid map as a reward function for general information acquisition tasks. Specific to SLAM exploration, [14] combined this reward function with the mutual information between sensor observations and robot localization. Many following works focused on mutual information-based control for target and/or robot localization, see [15], [16], [17]. Specific to frontier identification, [18] used a gradient-based approach for selecting frontiers between explored and unexplored space. The resulting explorative behaviors are well described by the results in this paper, as are the behaviors from many other exploration algorithms, e.g., [19], [20], [21]. Not surprisingly, these (if not all) exploration approaches for occupancy grid mapping are benchmarked against the traditional geometric-based approach of frontier exploration first proposed in [1].

B. Paper Organization

This paper is organized as follows. We introduce in Section II the well-known occupancy grid mapping algorithm, then provide a formal statement of our main theoretical result. In Section III, we prove the main result and provide numerous other results that are interesting in their own right. In Section IV, we focus on the process of generating mutual information reward surfaces that can drive information-based controllers, and in doing so discuss the experimental results from our two hardware experiments. We then summarize the contribution of the paper in Section V.

II. Eventual attraction to unexplored space

In this section we review the occupancy grid mapping algorithm then formally state the main theoretical result of this chapter. A detailed proof of this is given in Section III-C, and experimental results that support our finding are provided in Section IV-B.

A. The occupancy grid mapping algorithm

Consider the problem of constructing a map of an environment using a single robot equipped with a sensor that provides range measurements, i.e., the robot's distance to nearby obstacles. Suppose the robot employs the well-known occupancy grid mapping algorithm [3] for constructing the map from the noisy observations. For simplicity, we focus on the two-dimensional mapping problem, meaning that a robot of known position $c_k \in \mathbb{R}^2 \times \mathbb{S}$ at the discrete time step $k \in \mathbb{Z}_{\geq 0}$ moves in the two-dimensional Euclidean plane. The occupancy grid mapping algorithm models the static map of the robot's environment as an n_m-tuple random variable $M = (M^{[1]}, \ldots, M^{[n_m]})$, where each independent binary random variable $M^{[i]}$ takes the value 0 or 1 when an obstacle is absent or present, respectively, within a uniquely identified cell $i \in \{1, \ldots, n_m\}$. Inference of this map is enabled by observations the robot receives originating from its sensor. These observations are modeled as random variables Z_k that take values z_k, for all times k.

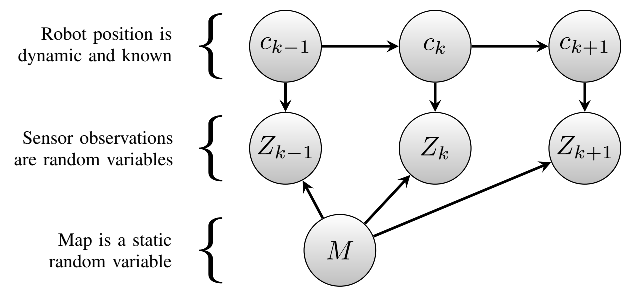

Fig. 2. Dynamic Bayesian network for occupancy grid mapping. Note that since the robot's position is known, we show in the top row the realized values c_k instead of corresponding random variables.

As illustrated in Figure 2, both the unknown map and the robot's known positions influence the noisy observations within this dynamic Bayesian network. Our goal of building a map is equivalent to estimating the posterior probability over all possible maps given the previous observations, which leads to the binary Bayesian filter

$$r^{[i]}(z_{1:k}) := \frac{\mathbb{P}(M^{[i]} = 1 | z_{1:k})}{\mathbb{P}(M^{[i]} = 0 | z_{1:k})} = \delta^{[i]}(z_k) r^{[i]}(z_{1:k-1}), \quad (1)$$

where $r^{[i]}$ is the odds ratio[3] of the posterior of cell i being occupied, $\delta^{[i]}$ is the odds ratio of the inverse sensor model, and $z_{1:k}$ is shorthand for $\bigcap_{k'=1}^{k} \{Z_{k'} = z_{k'}\}$. Note that for all cells $i \in \{1, \ldots, n_m\}$, we assume an uninformative map prior $\mathbb{P}(M^{[i]} = 0) = \mathbb{P}(M^{[i]} = 1) = 0.5$. We also use $r_k^{[i]}$ as shorthand for $r^{[i]}(z_{1:k})$ and refer to it as the robot's *belief* for cell i. When considering this belief as a random variable, we

[3]The reader may be more familiar with the log odds representation of the binary Bayesian filter (1). Although we use log odds ratios in the implementation for our experiments, the mathematical conclusions in this paper are more compactly expressed using odds ratios.

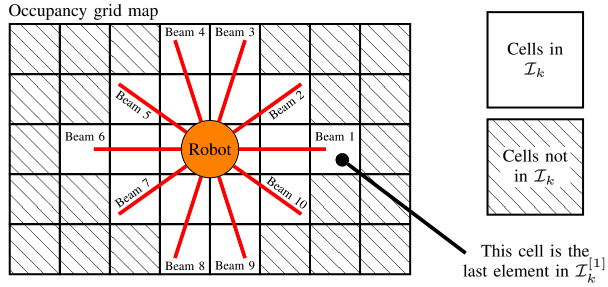

Fig. 3. The perceptual ranges for an example robot with a ten beam range sensor. Note that the assumption of measurement independence within a single observation may overlook the property that one cell is traversed by multiple measurements' beams.

use the notation $R_k^{[i]} := r^{[i]}(Z_{1:k})$, where $Z_{1:k}$ is shorthand for $\bigcap_{k'=1}^{k}\{Z_{k'}\}$.

We consider in this paper a beam-based range sensor that produces finite radial measurements of known bearing to nearby obstacles, even though our analysis is applicable to a wide range of sensor models. In addition, we assume that the beam is narrow, i.e., the width of the sensor's beam is negligible when compared to the size of our occupancy grid cells. This is the case for many laser-based sensors, e.g., laser rangefinders. Multiple measurements per time step are accommodated by modeling the observation Z_k as an n_z-tuple random variable $(Z_k^{[1]}, \ldots, Z_k^{[n_z]})$ whose elements take values $z_k^{[j]}$ from a common set of range measurements $\mathcal{Z} \subset \mathbb{R}_{\geq 0}$. The occupancy grid mapping algorithm assumes independence among the measurements $j \in \{1, \ldots, n_z\}$, and computes the odds ratio of the inverse sensor model by $\delta^{[i]}(z_k) = \prod_{j=1}^{n_z} \delta^{[i,j]}(z_k^{[j]})$, where

$$\delta^{[i,j]}(z_k^{[j]}) = \begin{cases} r_{occ}, & \text{if } z_k^{[j]} \text{ and } c_k \text{ imply a measurement location within cell } i, \\ r_{emp}, & \text{if } z_k^{[j]} \text{ and } c_k \text{ imply a sensor beam passed through cell } i, \\ 1, & \text{otherwise,} \end{cases}$$

with $r_{occ} \in (1, \infty)$ and $r_{emp} \in (0, 1)$ being constants. Note that for all functions, we do not include the robot's position c_k as an argument since it is known at all times k.

It is useful to consider the ordered set of cell indices $\mathcal{I}_k^{[j]}$ located along the jth measurement's beam in the absence of obstacles. In other words, the jth measurement's *perceptual range* contains all cells that align with this beam and are within the maximum range of the sensor. As illustrated for the example system in Figure 3, by 'ordered' we mean that set operators will consider the cell indices sequentially based on proximity to the robot, with the nearest cells being considered first. Also note that for all cells ℓ not in the observation's perceptual range $\mathcal{I}_k := \bigcup_{j=1}^{n_z} \mathcal{I}_k^{[j]}$, any realized observation z_k results in an odds ratio of the inverse sensor model of $\delta^{[\ell]}(z_k) = 1$. This implies that the belief $r_k^{[\ell]}$ of this cell will not change over time if the robot remains stationary. On the other hand, for a cell i in the observation's perceptual range $\mathcal{I}_k$, any realized observation z_k with $\delta^{[i]}(z_k) \neq 1$ will change the corresponding belief, i.e., $r_k^{[i]} \neq r_{k-1}^{[i]}$. We call such an observation that influences the belief an *informative observation* for cell i. Clearly an informative observation $z_k = (z_k^{[1]}, \ldots, z_k^{[n_z]})$ must be composed of at least one *informative measurement* $z_k^{[j]}$ for cell i. For simplicity, we sometimes say the jth measurement is informative for cell i when $\delta^{[i,j]}(z_k^{[j]}) \neq 1$.

Throughout the paper, we assume that the sensor model is unbiased and the measurement's prior is clamped. We now formally define these two assumptions.

Assumption 1 (Unbiased sensor model) *For all times $k \in \mathbb{Z}_{\geq 0}$, given that the next observation Z_{k+1} is informative $\delta^{[i]}(Z_{k+1}) \neq 1$ for a cell $i \in \mathcal{I}_{k+1}$, the expected value of $\delta^{[i]}(Z_{k+1})$ is less than 1 if the cell is unoccupied and greater than 1 if it is occupied.*

Assumption 2 (Clamped measurement's prior) *For all times $k \in \mathbb{Z}_{\geq 0}$, cells $i \in \mathcal{I}_{k+1}^{[j]}$, and measurements $j \in \{1, \ldots, n_z\}$, there exists an $\epsilon_0 > 0$ such that if $r_k^{[i]} < \epsilon_0$, then*

$$\mathbb{P}(Z_{k+1}^{[j]} | z_{1:k}) = \mathbb{P}(Z_{k+1}^{[j]} | z_{1:k}, M^{[i]} = 0).$$

Similarly, there exists an $\epsilon_1 < \infty$ such that if $r_k^{[i]} > \epsilon_1$, then

$$\mathbb{P}(Z_{k+1}^{[j]} | M) = \mathbb{P}(Z_{k+1}^{[j]} | z_{1:k}, M^{[i]} = 1).$$

B. Identification of mutual information-based behaviors

Given the previous observations $z_{1:k}$, we are interested in minimizing the expected uncertainty (i.e., conditional entropy[4]) of the map M after receiving the robot's next observation Z_{k+1}. For mutual information-based control, this objective is equivalent to maximizing the mutual information $I(M; Z_{k+1}|z_{1:k})$ between the map and the next observation. In addition, since both M and Z_{k+1} are tuple random variables with independence among elements, we express total mutual information as the summation of $I(M^{[i]}; Z_{k+1}^{[j]}|z_{1:k})$ between $M^{[i]}$ and $Z_{k+1}^{[j]}$ over all cells $i \in \{1, \ldots, n_m\}$ and measurements $j \in \{1, \ldots, n_z\}$. We further simplify this expression to

$$I(M; Z_{k+1}|z_{1:k}) = \sum_{j=1}^{n_z} \sum_{i \in \mathcal{I}_{k+1}^{[j]}} I(M^{[i]}; Z_{k+1}^{[j]}|z_{1:k}) \qquad (2)$$

since the mutual information between a measurement and any cell not in the measurement's perceptual range is equal to 0, i.e., $I(M^{[i]}; Z_{k+1}^{[j]}|z_{1:k}) = 0$ for all $j \in \{1, \ldots, n_z\}$ with $i \notin \mathcal{I}_{k+1}^{[j]}$ (see Lemma 5 for a rigorous treatment). We refer to $I(M^{[i]}; Z_{k+1}^{[j]}|z_{1:k})$ as the *mutual information contribution* of cell i, but note that it is also specific to the particular measurement j.

We refer to expression given in (2) as the *mutual information reward function*, a function that drives many information-based control approaches that are cast in terms

[4]The *conditional* entropy given the next observation should not be confused with the current entropy of the map M. The latter is independent of the next robot position c_{k+1} and the next observation Z_{k+1}.

of an optimization problem. In this paper, we exploit certain algebraic properties of this reward function to show that maximizing (2) implies eventual attraction to unknown space. Intuitively speaking, we show that unknown cells (i.e., those cells i for which the belief $r_k^{[i]}$ is equal to 1) at and after some time are expected to yield a higher mutual information reward compared to the hypothetical situation where these cells are correctly known (i.e., the belief $r_k^{[i]}$ is less than or greater than 1 depending on if cell i was realized to be unoccupied or occupied, respectively).

To formalize the simultaneous existence of an actual situation and a hypothetical one, we construct a novel analytical approach that we call *two independent theoretical experiments*.

Definition 3 (Two independent theoretical experiments) *Let E and $\tilde{E}$ denote two independent theoretical experiments with identical robots, paths, sensor models, and inverse sensors models, i.e., all deterministic variables and functions are the same. Suppose the first experiment's map M is of equal value to the second experiment's map $\tilde{M}$, although the value itself is unknown to the robots. In addition, suppose the observations $Z_{1:k} := \bigcap_{k'=1}^{k}\{Z_{k'}\}$ and $\tilde{Z}_{1:k} := \bigcap_{k'=1}^{k}\{\tilde{Z}_{k'}\}$ are independently realized in E and $\tilde{E}$, respectively. At some time $K \leq k$ in the past, the robots have known beliefs $R_K^{[i]}$ and $\tilde{R}_K^{[i]} := r^{[i]}(\tilde{Z}_{1:K})$ for all cells, which can be used to condition their expected beliefs $\mathbb{E}(R_k^{[i]})$ and $\mathbb{E}(\tilde{R}_k^{[i]})$, respectively, at time k.*

The intuition behind the experiments E and $\tilde{E}$ comes from thinking of time k as being in the future versus thinking of time K as being in the past. Doing so immediately raises the following question: *Given information about the map at time K, can we reason about the expected behavior of mutual information-based control at time k?* Even for arbitrary maps, the expected beliefs $\mathbb{E}\big(R_k^{[i]}\big)$ and $\mathbb{E}\big(\tilde{R}_k^{[i]}\big)$ give insight into how past observations can affect future uncertainty - note that future beliefs are indeed random variables. It is this insight that provides a basis for qualitatively evaluating the behaviors of controllers employing the mutual information reward function.

C. A theorem for attractive behavior

The main theoretical result implies that any controller tasked to maximize the mutual information reward function is eventually attracted to unexplored regions of the map. Below we formally state this result with the use of the analytical approach described in the previous section.

For the first experiment E, suppose that a subset of cells $\mathbb{I} \subset \{1,\ldots,n_m\}$ is composed of *unknown* cells at time K, meaning that the beliefs $R_K^{[i]}$ are all equal to 1 for all $i \in \mathbb{I}$. In the second experiment $\tilde{E}$, however, suppose that the same subset of cells are *correctly known* at time K. By correctly known we mean that for any cell $i \in \mathbb{I}$ at time K, the belief $\tilde{R}_K^{[i]}$ is less than 1 if the cell is unoccupied, otherwise, it is greater than 1. We again note that the maps M and $\tilde{M}$ are of equal value yet unknown to the robots. Lastly, suppose that for all other cells not in the subset $\mathbb{I}$, the beliefs between experiments E and $\tilde{E}$ are equal but otherwise arbitrary in value. More formally, we have that $R_K^{[\ell]} = \tilde{R}_K^{[\ell]}$ for all cells $\ell \in \mathbb{I}^c = \{1,\ldots,n_m\}/\mathbb{I}$.

Let the event $\mathcal{A}_K$ represent the intersection of all aforementioned events between the two experiments at time K. When reasoning about the mutual information contributions $I(M^{[i]}; Z_{k+1}^{[j]}|z_{1:k})$ for all measurements $j \in \{1,\ldots n_z\}$ and cells $i \in \{1,\ldots n_m\}$ at time $k \geq K$, we wish to evaluate these values based on the expected beliefs $\mathbb{E}(R_k^{[i]})$ and $\mathbb{E}(\tilde{R}_k^{[i]})$ given the event $\mathcal{A}_K$. In other words, we are evaluating the resulting mutual information reward function $I(M; Z_{k+1}|z_{1:k})$ employing the expected beliefs for the robots in both experiments. This is formally described by

$$\mathcal{B}_{K,k} := \bigcap_{i=1}^{n_m} \Big\{\big\{R_k^{[i]} = \mathbb{E}(R_k^{[i]}|\mathcal{A}_K)\big\} \cap \big\{\tilde{R}_k^{[i]} = \mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)\big\}\Big\}.$$

Note that this event does not explicitly include the event $\mathcal{A}_K$ since this information (e.g., the state of the map) is unknown to the robots. We now state the main theoretical result.

Theorem 4 (Attraction to unexplored space) *Consider the two independent experiments E and $\tilde{E}$. Given enough number of informative observations with $\delta^{[i]}(z_k) \neq 1$ for the subset of cells $i \in \mathbb{I}$, there exists a time $\varkappa \geq K$ such that for all times $k \geq \varkappa$, we have that*

$$I\big(M; Z_{k+1}|\mathcal{B}_{K,k}\big) \ > \ I\big(\tilde{M}; \tilde{Z}_{k+1}|\mathcal{B}_{K,k}\big) \tag{3}$$

if any cell $i \in \mathbb{I}$ is in the observation's perceptual range.

Note that although an unknown cell always has the maximum amount of entropy at time K, more certain beliefs (i.e., $r_k^{[i]} \neq 1$) of that cell may result in larger mutual information contributions at time K. Conditioning on the expected beliefs for all cells, Theorem 4 states that, in expectation, there is a future time at an after which the mutual information reward function is greater if these cells are unknown at time K than if they were correctly known. This statement is stronger than one stating that the cell's mutual information contribution will be greater for a particular measurement. This statement also holds for any observation that perceives at least one of these cells, regardless of the sensor model. Thus, we conclude that mutual information-based controllers are eventually attracted to unexplored space.

III. On informative behaviors

In this section, we prove the main theoretical result stated in Theorem 4. Along the way, we provide a number of other results that may be interesting in their own right.

A. Algebraic properties of mutual information

Concerning an arbitrary cell $i \in \{1,\ldots,n_m\}$ and measurement $j \in \{1,\ldots,n_z\}$, the following lemma expresses the corresponding mutual information contribution as an integral of two functions, one of which has useful algebraic properties.

Lemma 5 (Mutual information contribution) *The mutual information at time k between a single measurement j and a single cell i given the observations $z_{1:k}$ can be written as*

$$I(M^{[i]}; Z_{k+1}^{[j]} | z_{1:k}) = \int_{z \in \mathcal{Z}} \mathbb{P}(Z_{k+1}^{[j]} = z | z_{1:k}) \; f\left(\delta^{[i,j]}(z), r_k^{[i]}\right) \; dz \quad (4)$$

where the nonnegative function $f(\delta, r)$ is endowed with the following two algebraic properties:

1) *For $\delta = 1$ we have that $f(\delta, r) = 0$ for any $r \in (0, \infty)$;*
2) *For any value of $\delta \in (0, \infty)/\{1\}$, there exists some $\tau(\delta) \in (0, \infty)$ such that $f(\delta, r)$ is monotonically increasing in r for all $r < \tau(\delta)$ and monotonically decreasing in r for all $r > \tau(\delta)$.*

Proof: [Lemma 5] From information theory [5], the mutual information contribution is defined as

$$I(M^{[i]}; Z_{k+1}^{[j]} | z_{1:k}) = \int_{z \in \mathcal{Z}} \sum_{m^{[i]} \in \{0,1\}} \mathbb{P}(M^{[i]} = m^{[i]} | Z_{k+1}^{[j]} = z, z_{1:k}) \times \mathbb{P}(Z_{k+1}^{[j]} = z | z_{1:k}) \log\left(\frac{\mathbb{P}(M^{[i]} = m^{[i]} | Z_{k+1}^{[j]} = z, z_{1:k})}{\mathbb{P}(M^{[i]} = m^{[i]} | z_{1:k})}\right) dz. \quad (5)$$

Since the complement of $\mathbb{P}(M^{[i]} | z_{1:k})$ and $\mathbb{P}(M^{[i]} | Z_{k+1}^{[j]}, z_{1:k})$ are equal to $1 - \mathbb{P}(M^{[i]} | z_{1:k})$ and $1 - \mathbb{P}(M^{[i]} | Z_{k+1}^{[j]}, z_{1:k})$, respectively, we have from (1) that $\mathbb{P}(M^{[i]} = m^{[i]} | z_{1:k}) = ((r_k^{[i]} - 1)m^{[i]} + 1)/(r_k^{[i]} + 1)$ and $\mathbb{P}(M^{[i]} = m^{[i]} | Z_{k+1}^{[j]} = z, z_{1:k}) = \frac{(r_k^{[i]} \delta^{[i,j]}(z) - 1)m^{[i]} + 1}{r_k^{[i]} \delta^{[i,j]}(z) + 1}$. Substituting these two expressions into (5) gives (4), where

$$f(\delta, r) := \log\left(\frac{r+1}{r + \delta^{-1}}\right) - \frac{\log(\delta)}{r\delta + 1}. \quad (6)$$

Immediately, the first algebraic property of $f(\delta, r)$ follows directly from (6) with $\delta = 1$ from the inverse sensor model.

To prove the second algebraic property of $f(\delta, r)$, consider a measurement value $z \in \mathcal{Z}$ that results in $\delta \neq 1$. Taking the partial derivative of $f(\delta, r)$ with respect to r, we get that

$$\frac{\partial f(\delta, r)}{\partial r} = \frac{1}{r+1} - \frac{1}{r + \delta^{-1}} + \frac{\delta \log(\delta)}{(\delta r + 1)^2}. \quad (7)$$

Now define the belief threshold $\tau(\delta)$ as

$$\tau(\delta) := \Big(\delta \log(\delta) - \delta + 1\Big) \Big/ \Big(\delta(\delta - \log(\delta) - 1)\Big), \quad (8)$$

and note that $\tau(\delta) \in (0, \infty)$ for all $r \in (0, \infty)$. Substituting $(1 + \epsilon)\tau(\delta)$ for r in (7) gives

$$\left.\frac{\partial f(\delta, r)}{\partial r}\right|_{r = (1+\epsilon)\tau(\delta)} = \frac{-\epsilon \delta \xi \beta^3}{(\epsilon\xi + (\delta - 1)^2)(\epsilon\xi + \gamma)^2}, \quad (9)$$

where ξ, β, and γ are defined as $(\delta \log(\delta) - \delta + 1)$, $(\delta - \log(\delta) - 1)$, and $(\log(\delta) - \delta \log(\delta))$, respectively. By applying the natural logarithm inequalities of $(\delta - 1) \le \delta \log(\delta) \le \delta(\delta - 1)$, we have that ξ, β, and γ are all positive. This implies that (9) is finite and positive for all $\epsilon \in (0, \infty)$, which gives the desired result for negative monotonicity.

For positive monotonicity, we again apply the previously mentioned natural logarithm inequalities to have that ξ is less than or equal to both $(\delta - 1)^2$ and γ. This implies that both $(\epsilon\xi + (\delta - 1)^2)$ and $(\epsilon\xi + \gamma)$ are finite and positive for $\epsilon \in (-1, 0)$, hence (9) is finite and negative. ■

Remark 6 (Belief thresholds) *For an informative measurement, we have that δ is equal to either r_{occ} or r_{emp}, implying that the belief threshold $\tau(\delta)$ is either less than 1 or greater than 1, respectively. Suppose cell i is unknown, i.e. $r_k^{[i]} = 1$. Then measurement locations within cell i result in $\tau(r_{occ}) < 1$, implying that $f(\delta, r_k^{[i]})$ is monotonically decreasing in the belief $r_k^{[i]}$. For measurements of greater value, $\tau(r_{emp}) > 1$ and $f(\delta, r_k^{[i]})$ is instead monotonically increasing.*

One can interpret $f(\delta, r)$ in Lemma 5 as an information gain function for potential measurement values $z \in \mathcal{Z}$. For example, the second algebraic property states that no information about a cell is received from a measurement that does not intersect that cell, whether because the measurement "stops short" of the cell or simply because the cell is not in the measurement's perceptual range. This geometric dependency relates back to the construction of the index sets $\mathcal{I}_k$ in Section II-A, and likewise allows for the mutual information reward function to be calculated over a reduced domain size. In other words, no matter how the jth measurement's prior $\mathbb{P}(Z_{k+1}^{[j]} | z_{1:k})$ behaves, the integral of (4) does not need to evaluate measurement values $z \in \mathcal{Z}$ that imply $\delta^{[i,j]}(z) = 1$ for any cell i in the measurement's perceptual range $\mathcal{I}_{k+1}^{[j]}$.

B. Expected values of future beliefs

For the main theoretical result in Theorem 4, we are interested in the geometric interpretations of the monotonic properties of the information gain function $f(\delta, r)$, i.e., the second algebraic property of Lemma 5. For example, consider an unoccupied cell i with $M^{[i]} = 0$ that is unknown (i.e., $r_k^{[i]} = 1$) to a stationary robot. Suppose this robot receives a sequence of measurements all extending past the aforementioned cell. Lemma 5 says that the information gain function $f(\delta, r_k^{[i]})$ itself provably *decreases* with each observation, implying that the potential to receive informative observations was highest when the cell was originally unknown. This behavior suggests that unknown cells have stronger attractive properties for measurements implying the absence of an obstacle.

We would like to have the same result for measurement locations falling within a cell. This would imply that for all measurements large enough to "reach" the unoccupied cell, the information gain function provably *decreases* as the robot becomes more certain of the cell's unoccupied state. From Remark 6, this is not the case for $r_k^{[i]} = 1$. Fortunately, $\tau(r_{occ})$ in Lemma 5 gives the necessary belief threshold below which the monotonicity result holds for any informative measurement. Given the expectation to infer the correct state for observable cells, we have that all unoccupied cells would eventually exhibit the desired behavior, i.e., the expected information gain decreases with increased certainty of the state of a cell. The same interpretation holds true for occupied cells, which motivates the following lemma and corollaries that explore the robot's beliefs and measurements' priors in expectation.

Lemma 7 (Relationship between expected beliefs) *Suppose a numerical relationship (i.e., less than, greater than, or equal to) is given at time $K \geq 0$ between the beliefs $R_K^{[i]}$ and $\tilde{R}_K^{[i]}$, for any cell i. Then for all times $k \geq K$, the numerical relationship between the expected beliefs $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K^{[i]})$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K^{[i]})$ will remain the same.*

Proof: [Lemma 7] The proof proceeds by induction on k. If $k = K$, then the beliefs are equal to their realized values. Thus, this lemma trivially holds for $k = K$. Let us now assume that the result is true for some $k \geq K$. Note that if $i \notin \mathcal{I}_{k+1}$, then the lemma trivially holds for $k+1$. Now suppose $i \in \mathcal{I}_{k+1}$, and define the conditioning event $\mathcal{S}_K := \{M = m\} \cap \{\tilde{M} = m\} \cap \{R_K^{[i]} < \tilde{R}_K^{[i]}\}$ for some $m \in \mathcal{M}$. For the first experiment E, we have from the property of independence between observations and between measurements that

$$\mathbb{E}(R_{k+1}^{[i]}|\mathcal{S}_K) = \prod_{j=1}^{n_z} \mathbb{E}(\delta^{[i,j]}(Z_{k+1}^{[j]})|\mathcal{S}_K)\mathbb{E}(R_k^{[i]}|\mathcal{S}_K), \quad (10)$$

In addition, (10) is identical to the expression for the second experiment $\tilde{E}$ except with $R_k^{[i]}$, Z_k, and $Z_k^{[j]}$ being replaced by $\tilde{R}_k^{[i]}$, $\tilde{Z}_k$, and $\tilde{Z}_k^{[j]}$, respectively, for all times k. From the inverse sensor model and the law of total probability, we have for the first experiment E that

$$\mathbb{E}(\delta^{[i,j]}(Z_{k+1})|\mathcal{S}_K) = \int_{z\in\mathcal{Z}} \mathbb{P}(Z_{k+1}^{[j]} = z|\mathcal{S}_K)\delta^{[i,j]}(z)dz.$$

Since this expression also contains the condition $\tilde{M} = m$ for the second experiment $\tilde{E}$, and both experiments have identical sensor and inverse sensor models, we have that $\mathbb{E}\big(\delta^{[i,j]}(\tilde{Z}_{k+1}^{[j]})|\mathcal{S}_K\big)$ is equal to $\mathbb{E}\big(\delta^{[i,j]}(Z_{k+1}^{[j]})|\mathcal{S}_K\big)$. Thus from (10), we have that

$$\mathbb{E}(R_{k+1}^{[i]}|\mathcal{S}_k)/\mathbb{E}(R_k^{[i]}|\mathcal{S}_k) = \mathbb{E}(\tilde{R}_{k+1}^{[i]}|\mathcal{S}_k)/\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{S}_k), \quad (11)$$

and thus this proof holds for $R_K^{[i]}$ being less than $\tilde{R}_K^{[i]}$ given that M and $\tilde{M}$ are both equal to the map value m. To complete the proof, we note that the derivation of (11) is not specific to the map value $m \in \mathcal{M}$ nor the numerical relationship between $R_K^{[i]}$ and $\tilde{R}_K^{[i]}$. ∎

Corollary 8 (Convergence of expected beliefs) *The expected beliefs $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K^{[i]})$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K^{[i]})$ of any unoccupied (occupied) cell i monotonically converge to 0 (∞, respectively) as the number of informative observations for this cell tends to ∞.*

Proof: [Corollary 8] Considering (10) in Lemma 7, the proof follows from Assumption 1. ∎

Corollary 9 (Equality of measurements' priors) *Consider the set $\mathbb{I}$ from Section II-C as a set consisting of a single cell i. Given enough number of informative observations for this cell, there exists a time $\varkappa \geq K$ such that for all times $k \geq \varkappa$, we have that the jth measurements' priors $\mathbb{P}(Z_{k+1}^{[j]}|\mathcal{B}_{K,k}^{[i]})$ and $\mathbb{P}(\tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k}^{[i]})$ are equal for all measurements j.*

Proof: [Corollary 9] From Corollary 8, there exists a time $\varkappa \geq K$ such that for all times $k \geq \varkappa$, the two expected beliefs $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)$ are both smaller than some $\epsilon_0 > 0$ or larger than some $\epsilon_1 < \infty$, depending on whether $M^{[i]}$ and $\tilde{M}^{[i]}$ are both of value 0 or 1, respectively. From Assumption 2, this implies that

$$\begin{aligned}\mathbb{P}(Z_{k+1}^{[j]}|\mathcal{B}_{K,k}) &= \mathbb{P}\Big(Z_{k+1}^{[j]}|\mathcal{B}_{K,k} \cap \bigcap_{i\in\mathbb{I}}\{M^{[i]} = \tilde{M}^{[i]}\}\Big) \\ &= \mathbb{P}\Big(\tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k} \cap \bigcap_{i\in\mathbb{I}}\{\tilde{M}^{[i]} = M^{[i]}\}\Big) = \mathbb{P}(\tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k}),\end{aligned}$$

where the third equality is due to all other expected beliefs $\mathbb{E}(R_k^{[\ell]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[\ell]}|\mathcal{A}_K)$ being equal between experiments for any cell $\ell \in \mathbb{I}^c$. ∎

C. A proof for attractive behavior

We now give the proof for the main theoretical result in Theorem 4.

Proof: [Theorem 4] Suppose the subset $\mathbb{I} \subset \{1, \ldots, n_m\}$ for the two experiments E and $\tilde{E}$ consists of a single cell i of arbitrary state. Denote this subset as $\mathbb{I}_1$. We first show that (3) holds for this subset regardless of the state of cell i.

Suppose this cell is realized to be unoccupied, i.e., $M^{[i]} = 0$. From Corollary 8, the two expected beliefs $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)$ monotonically converge to 0 as the number of informative observations for cell i tends to ∞. Thus, there exists a time $\varkappa_1^{[1]} \geq K$ such that for all $k \geq \varkappa_1^{[1]}$, both these expected beliefs are smaller than the belief thresholds $\tau(r_{occ})$ and $\tau(r_{emp})$. In addition, the second algebraic property of Lemma 5 guarantees that for all times $k \geq \varkappa_1^{[1]}$, the information gain function $f(\delta, r)$ employing $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)$ is monotonically decreasing in r for both $\delta = r_{occ}$ and $\delta = r_{emp}$. Since at time K cell i was unknown (i.e., $R_K^{[i]} = 1$) in the first experiment E and correctly known (i.e., $\tilde{R}_K^{[i]} < 1$) in the second experiment $\tilde{E}$, we have from Lemma 7 that $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K)$ is larger than $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)$, hence the corresponding $f(\delta, r)$ is larger. Also due to Lemma 7, we have for any time k that the expected beliefs $\mathbb{E}(R_k^{[\ell]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[\ell]}|\mathcal{A}_K)$ of all other cells $\ell \in \mathbb{I}_1^c$ are equal. This implies that the information gain function $f(\delta, r)$ for these expected beliefs are also equal between experiments.

Now consider the expected behavior of the measurements' priors $\mathbb{P}(Z_{k+1}^{[j]}|\mathcal{B}_{K,k})$ and $\mathbb{P}(\tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k})$ for the two experiments E and $\tilde{E}$, respectively. From Corollary 9, there exists another time $\varkappa_1^{[2]} \geq K$ such that for all times $k \geq \varkappa_1^{[2]}$, these measurements' priors are equal. From Lemma 5, we have for all time $k \geq \max\{\varkappa_1^{[1]}, \varkappa_1^{[2]}\}$ that $I(M^{[i]}; Z_{k+1}^{[j]}|\mathcal{B}_{K,k})$ is greater than $I(\tilde{M}^{[i]}; \tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k})$ for any measurement j with $i \in \mathcal{I}_{k+1}^{[j]}$. For all other measurements or cells in $\mathbb{I}_1^c$, these mutual information contributions are equal. This proves (3) holds for the subset $\mathbb{I}_1$ consisting of a single unoccupied cell i with a threshold time of $\max\{\varkappa_1^{[1]}, \varkappa_1^{[2]}\}$.

The proof that (3) holds for the subset $\mathbb{I}_1$ consisting of a single occupied cell i follows verbatim; however, this proof results in two subtle difference. Firstly, there exists a time $\varkappa_1^{[3]} \geq K$ such that for all times $k \geq \varkappa_1^{[3]}$, the two expected beliefs $\mathbb{E}(R_k^{[i]}|\mathcal{A}_K)$ and $\mathbb{E}(\tilde{R}_k^{[i]}|\mathcal{A}_K)$ are greater than both belief thresholds $\tau(r_{occ})$ and $\tau(r_{emp})$. Secondly from Corollary 9, there exists a time $\varkappa_1^{[4]} \geq K$ such that for all times $k \geq \varkappa_1^{[4]}$, the measurements' priors $\mathbb{P}(Z_{k+1}^{[j]}|\mathcal{B}_{K,k})$ and $\mathbb{P}(\tilde{Z}_{k+1}^{[j]}|\mathcal{B}_{K,k})$ are equal. Thus, (3) holds for the subset $\mathbb{I}_1$ consisting of a single cell i of arbitrary state with a threshold time of $\varkappa_1 := \max\{\varkappa_1^{[1]}, \varkappa_1^{[2]}, \varkappa_1^{[3]}, \varkappa_1^{[4]}\}$.

We conclude by showing that (3) holds for subsets $\mathbb{I}_s$ of arbitrary size s. The proof proceeds by induction on s, with $s = 1$ already proved. Let us assume that the result is true for some $s \geq 1$. For $s+1$ with the subset $\mathbb{I}_{s+1}$, define a one cell subset $\mathbb{J} := \mathbb{I}_{s+1}/\mathbb{I}_s$ containing the one cell not originally in $\mathbb{I}_s$, then let $\mathbb{J}^c := \{1, \ldots, n_m\}/\mathbb{J}$. The proof to show that the mutual information reward function $I\big(M; Z_{k+1}|\mathcal{B}_{K,k}\big)$ in the first experiment E is strictly greater that $I\big(\tilde{M}; \tilde{Z}_{k+1}|\mathcal{B}_{K,k}\big)$ in the second experiment $\tilde{E}$ follows verbatim as the proof for $\mathbb{I}_1$, with the exception that the analysis yields a threshold time of $\varkappa_s$. Thus, Theorem 4 holds in general for an arbitrary subset $\mathbb{I}_s$ with a threshold time of $\varkappa = \max \bigcup_s \{\varkappa_s\}$. ∎

IV. Experiments in mapping

We discuss the technical aspects of our software and hardware platform, then present the resulting mutual information reward surfaces for two separate experiments.

A. Experimental setup

For the two hardware experiments, we implemented a surprisingly compact form of the mutual information reward function $I(M; Z_{k+1}|z_{1:k})$ at time k. For the sake of brevity, we skip the derivation of the expression and instead refer the reader to [22]. We first state the assumption that all obstacles in the map are opaque, which is inherent to the occupancy grid mapping algorithm.

Assumption 10 (Sensor obscuration) *For any cell ℓ appearing after an occupied cell i in the index set $\mathcal{I}_{k+1}^{[j]}$, we have that*

$$\mathbb{P}\big(Z_{k+1}^{[j]}|M^{[i]} = 1, M^{[\ell]} = 0\big) = \mathbb{P}\big(Z_{k+1}^{[j]}|M^{[i]} = 1, M^{[\ell]} = 1\big)$$

for any time k and measurement j.

With this assumption, we can use the property that the index set $\mathcal{I}_k^{[j]}$ is ordered to write $I(M; Z_{k+1}|z_{1:k})$ as

$$\sum_{j=1}^{n_z} \sum_{i \in \mathcal{I}_{k+1}^{[j]}} \int_{z \in \mathcal{Z}} f\big(\delta^{[i,j]}(z), r_k^{[i]}\big) \Bigg(\sum_{\ell \in \mathcal{I}_{k+1}^{[j]}} \frac{r_k^{[\ell]}}{\prod_{\zeta \leq \ell}(r_k^{[\zeta]} + 1)}$$
$$\times\, \mathbb{P}(Z_{k+1}^{[j]} = z|M = \mathbf{e}_\ell) + \frac{1}{\prod\limits_{\zeta \in \mathcal{I}_{k+1}^{[j]}} (r_k^{[\zeta]} + 1)} \mathbb{P}(Z_{k+1}^{[j]} = z|M = \mathbf{0}) \Bigg) dz.$$

where $\ell \leq i$ denotes the sequence of indices in $\mathcal{I}_{k+1}^{[j]}$ appearing before or equal to i. Note that $\mathbf{0}$ and $\mathbf{e}_i$ denote the zero tuple and ith Euclidean basis tuple, respectively.

Using the Robot Operating System [23], we employed an omnidirectional ground robot custom built at MIT Lincoln Laboratory (Figure 4). We equipped this robot with a Hokuyo scanning laser rangefinder (LIDAR) module that produces $n_z = 1024$ horizontal measurements over 360 degrees at a rate of 10 Hz. For both hardware experiments, we selected measurement range values of $\mathcal{Z} = (0.0, 5.6)$ m, mapping algorithm parameters of $(r_{occ}, r_{emp}) = (1.5, 0.66)$, and a beam-based proximity mixture sensor model [24] with parameters $z_{hit} = 0.7$, $z_{short} = 0.1$, $z_{max} = 0.1$, $z_{rand} = 0.1$, $\sigma_{hit} = \max\{0.03 \text{ m}, 0.03 z_{exp}\}$, and $\lambda_{short} = 5.0 \text{ m}^{-1}$.

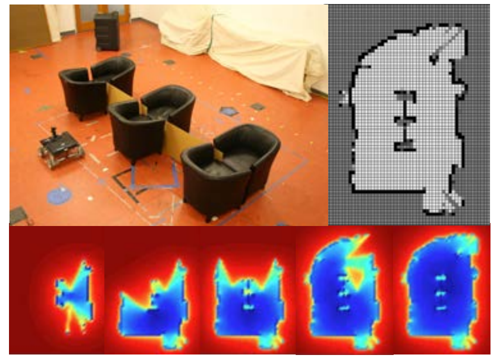

Fig. 4. *Top left:* An omnidirectional ground robot with 360 degree LIDAR is driven clockwise around a center opaque barrier for the occupancy grid mapping experiment. *Top right:* The resulting occupancy grid map after the 60 s run. *Bottom, left to right:* The corresponding reward surfaces at times of 0, 15, 30, 45 and 60 s. Note that the linear color scale from blue to red represents $[4, 68]$ bit for the mutual information reward surfaces.

B. Mutual information reward surfaces

We now discuss the resulting mutual information reward surfaces describing the overlap between map uncertainty and information gain from the robot's laser rangefinder. To simplify these surfaces, we prohibited the omnidirectional robot from rotating, i.e., for a particular translational position, the robot is always assumed to be facing in one direction. This constraint is not overly restrictive due to the sensor's rotational symmetry, yet it allows us to ignore the, in general, nontrivial effect rotational motion has on mutual information.

Figure 4 shows the setup for the occupancy grid mapping experiment. The environment to be mapped was set up with a center barrier and several scattered obstacles. A motion capture system tracked the position of the robot in realtime. In order to build a mostly complete map using a simplistic path, we manually drove the robot around the center barrier in a single loop. A time series of occupancy grid maps of 20 cm cell size was generated in realtime for this 60 second run. Figure 4 also shows five time instances for which we calculated the resulting mutual information reward surfaces. Due to the strong propagating frontiers (red regions) leading the robot, the evolution of these surfaces suggest that mutual information-based controllers would have drawn the robot down a similar path for exploration. In addition, we noted that the center barrier and exterior walls exhibited repulsive tendencies (i.e., low mutual information reward) due to the likelihood of sensor obscuration. These conclusions are not only consisted with the theoretical findings of this chapter, but are consisted among 10 additional runs of varying time durations that we conducted.

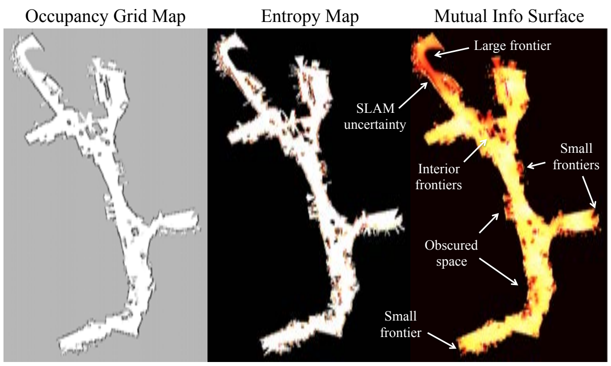

Fig. 5. The occupancy grid map, entropy map, and mutual information reward surface generated after a 50 minute SLAM experiment performed on the first floor of the MIT Stata Center. *Left:* The resulting occupancy grid map. *Middle:* The entropy map showing cells of high uncertainty. *Right:* A more geometrically relevant mutual information reward surface that highlights locations of frontiers, obscured space, and potential SLAM convergence issues. Note that the linear color scales from white to black represent $[0, 1]$ bit for the entropy map and $[9, 48]$ bit for the mutual information reward surface.

For the second experiment, we employed the GMapping simultaneous localization and mapping (SLAM) package. We again manually drove the robot, however, the environment to be mapped was much larger ($1500+$ m^2). Over 50 minutes of data was collected, after which SLAM was performed to build an occupancy grid map using a 30 cm square cell size. Figure 5 compares the resulting entropy map with the mutual information reward surface. In the entropy map, we see many cells of high uncertainty, however, they tend to be in isolated clusters near objects. On the other hand, the mutual information surface gives strong visual indicators throughout the open areas of the map. In particular, locations of interest to a mapping robot (e.g., frontiers, obscured space, and potential SLAM issues) are easily identifiable by either human or algorithm. Thus, we conclude that mutual information can both enable and monitor the SLAM algorithm. Lastly, we note that even for this large map, the mutual information reward surface over all known cells can be calculated from scratch in less than ten seconds on our desktop computer (3.4 Ghz quad-core CPU, GTX 690 GPU, MATLAB R2012a).

V. Conclusion

In this paper we prove that any controller tasked to maximize a mutual information reward function is eventually attracted to unexplored space. More importantly, the insight obtained in deriving this proof is exemplified by the results of our experiments. This solidifies the intuition that inherently guides our logic when designing motion planning algorithms. We believe that much more can be proven about the behavior of mutual information-based controllers given the geometric dependencies of the occupancy grid mapping algorithm and the monotonic properties of mutual information. In particular, we leave as a future work proving that obstacles eventually repel these controllers due to sensor obscuration (see Figure 1). This would result in trajectories that naturally avoid obstacles – a property that could give insight into the safety of these controllers in unstructured environments.

References

[1] B. Yamauchi. Frontier-based exploration using multiple robots. In *Proc. Int. Conference on Autonomous Agents*, St. Paul, MN, USA, May 1998.

[2] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006.

[3] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conference on Robotics and Automation*, St. Louis, MO, USA, Mar 1985.

[4] A. Elfes. Robot navigation: integrating perception, environmental constraints and task execution within a probabilistic framework. In *Reasoning with Uncertainty in Robotics*, Amsterdam, The Netherlands, Dec 1995.

[5] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley, New York, 1991.

[6] A. E. Johnson and S. B. Kang. Registration and integration of textured 3D data. *Image and Vision Computing*, 17(2):135–147, Feb 1999.

[7] S. Thrun, A. Bücken, W. Burgard, D. Fox, T. Fröhlinghaus, D. Hennig, T. Hofmann, M. Krell, and T. Schimdt. Map learning and high-speed navigation in rhino. In *AI-based Mobile Robots: Case studies of successful robot systems*. MIT Press, 1998.

[8] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proc. IEEE Int. Conference on Robotics and Automation*, San Francisco, CA, USA, Apr 2000.

[9] J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proc. IEEE/RSJ Int. Conference on Intelligent Robots and Syst.*, Osaka, Japan, Nov 1991.

[10] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Autonomous Robots*, 5(3-4):253–271, Jul 1998.

[11] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: a factored solution to the simultaneous localization and mapping problem. In *Proc. Nat. Conference on Artificial Intell.*, Edmonton, Alberta, Canada, Jul 2002.

[12] R. Chatila and J. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proc. IEEE Int. Conference on Robotics and Automation*, St. Louis, MO, USA, Mar 1985.

[13] R. C. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *The Int. J. of Robotics Research*, 5(4):56–68, Dec 1986.

[14] F. Bourgault, A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte. Information based adaptive robotic exploration. In *Proc. IEEE/RSJ Int. Conference on Intelligent Robots and Syst.*, Lausanne, Switzerland, Sep 2002.

[15] B. J. A. Kröse, N. Vlassis, R. Bunschoten, and Y. Motomura. A probabilistic model for appearance-based robot localization. *Image and Vision Computing*, 19(6):381 – 391, Apr 2001.

[16] B. Grocholsky, A. Makarenko, and H. F. Durrant-Whyte. Information-theoretic control of multiple sensor platforms. In *Proc. IEEE Int. Conference on Robotics and Automation*, Taipei, Taiwan, Sep 2003.

[17] E. Stump, V. Kumar, B. Grocholsky, and P. M. Shiroma. Control for localization of targets using range-only sensors. *The Int. J. of Robotics Research*, 28(6):764–757, May 2009.

[18] R. Rocha, J. Dias, and A. Carvalho. Cooperative multi-robot systems: a study of vision-based 3D mapping using information theory. *Robotics and Autonomous Syst.*, 53(3-4):282–311, Nov 2005.

[19] C. Stachniss. *Exploration and mapping with mobile robots*. PhD thesis, University of Freiburg, 2006.

[20] F. Amigoni and V. Caglioti. An information-based exploration strategy for environment mapping with mobile robots. *Robotics and Autonomous Syst.*, 58(5):684–699, May 2010.

[21] G. M. Hoffmann and C. J. Tomlin. Mobile sensor network control using mutual information methods and particle filters. *IEEE Trans. on Automatic Control*, 55(1):32–47, Jan 2010.

[22] B. J. Julian, S. Karaman, and D. Rus. On mutual information-based control of range sensing robots for mapping applications. *The Int. J. of Robotics Research*, 2013. Submitted.

[23] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. ROS: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, Kobe, Japan, May 2009.

[24] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents Series. The MIT Press, 2005.