# *An Efficient and Continuous Approach to Information-Theoretic Exploration*

# An Efficient and Continuous Approach to Information-Theoretic Exploration

Theia Henderson, Vivienne Sze, Sertac Karaman

***Abstract*— Exploration of unknown environments is embedded and essential in many robotics applications. Traditional algorithms, that decide where to explore by computing the expected information gain of an incomplete map from future sensor measurements, are limited to very powerful computational platforms. In this paper, we describe a novel approach for computing this expected information gain efficiently, as principally derived via mutual information. The key idea behind the proposed approach is a continuous occupancy map framework and the recursive structure it reveals. This structure makes it possible to compute the expected information gain of sensor measurements across an entire map much faster than computing each measurements' expected gain independently. Specifically, for an occupancy map composed of $|M|$ cells and a range sensor that emits $|\Theta|$ measurement beams, the algorithm (titled FCMI) computes the information gain corresponding to measurements made at each cell in $O(|\Theta||M|)$ steps. To the best of our knowledge, this complexity bound is better than all existing methods for computing information gain. In our experiments, we observe that this novel, continuous approach is two orders of magnitude faster than the state-of-the-art FSMI algorithm.**

**Supplementary video: `youtu.be/j_O1vOCrUME`**

## I. Introduction

Robot exploration is an autonomous navigation task with the goal of charting a map of an unknown environment. Typical formulations of the problem aim to minimize relevant metrics, such as, time, travel distance and energy, while constructing the map.

In this paper, we focus on problem instances that seek a completed *occupancy map*, in which the environment is divided into regions that are each labeled as either *free* or *occupied*. Arguably, the most common realization of such a map is an occupancy grid [1], where each cell of a regular grid spanning the environment is labeled with an occupancy probability. In this framework, the map is inferred from *range sensor* measurements, which radially measure distances from the sensor to the nearest occupied points in space. In practice, such measurements can be collected by lidar sensors.

In this setting, a useful operation is to determine how informative each potential measurement location is, since choosing to measure at a more informative location will help to minimize the total time, travel distance or energy required to construct the map. In fact, most existing robot exploration strategies rely on such a metric. For instance, greedy approaches find and navigate to the location that maximizes information gain in an iterative manner [2] and receding horizon approaches maximize information gain along paths within a look-ahead window [3]. Oftentimes, the information gain is combined with path cost [4], localization uncertainty [5] or measurement dependency [6].

This work was partially funded by the National Science Foundation, Cyber-Physical Systems (CPS) program, through grant no. 1837212. The authors are with the Massachusetts Institute of Technology. Emails: {tfh,sze,sertac}@mit.edu

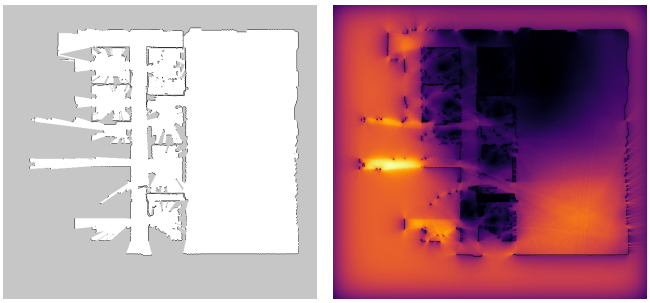

(a) An occupancy grid (b) A mutual information surface

Fig. 1. Mutual information in an incomplete map of MIT's building 31. In the occupancy grid, black pixels are occupied, white pixels are free and gray pixels are unknown. In the mutual information surface the brightness of a pixel indicates how much information is expected to be gained from a 360° range sensor measurement made at that pixel.

Despite the widespread reference to an "expected information gain" function, however, it has no universally accepted definition. For instance, expected information gain has been quantified as the solid angle of unknown space visible from the sensor [2], [7] as well as the volume of unknown space that is visible to the sensor [3], [4], [8]. The former heuristic measures the visible *frontier* [9], and greedily maximizing it is a slightly stronger variant of *frontier exploration*, which states to simply move to the nearest unknown space. Suppose the range sensor is comprised of $|\Theta|$ radial measurement beams each with a maximum range of $l$ cells. Then, computing each information heuristic at $|M|$ potential measurement locations takes time $O(l|\Theta||M|)$. Even in two dimensions, this complexity can be limiting; authors report downsampling potential measurement locations [4] and using adaptive data structures [10] to make the computation manageable.

A more principled approach to quantify expected information gain was introduced by Bourgault et al. [11]. If the information content of a map $M$ is measured by its entropy $H$, then the expected information to be gained from a sensor measurement $Z$ is the *mutual information* [12] between $M$ and $Z$, which is denoted by $I(M;Z)$ and defined as follows:

$$I(M;Z) = H(M) - H(M|Z). \tag{1}$$

Bourgault et al. [11] computed mutual information in a Monte Carlo fashion, by averaging the resulting map entropy over random simulated sensor measurements.

Julian et al. [13] derived a deterministic expression for the computation of mutual information $I(M;Z)$. Their algo-

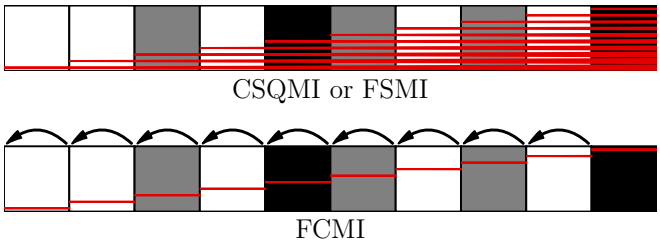


Fig. 2. Using FSMI or CSQMI to compute the expected information gain of every map cell above requires scanning over the remaining length of the array for each cell. Our algorithm, FCMI, computes the expected information gain at every cell in a single pass by reusing calculations from already scanned cells.

rithm runs in time $O(\lambda_z l^2|\Theta||M|)$, where $\lambda_z$ is the resolution of a numerical integral. Charrow et al. [6] improved upon this in their Cauchy-Schwarz Quadratic Mutual Information (CSQMI) algorithm which runs in time $O(l|\Theta||M|)$. However, the mutual information that CSQMI computes is defined according to a Cauchy-Schwarz cost function [14] rather than the self-information function originally proposed by Shannon [15], which is the only local, proper and smooth cost function for an alphabet of size at least three [16], [17].

More recently, Zhang et al. [18] proposed the Fast Shannon Mutual Information (FSMI) algorithm, which computes mutual information as classically defined by Shannon in time $O(l|\Theta||M|)$, matching both CSQMI and the information heuristics.

In this paper, we aim to compute a principled mutual information metric even faster. The main contribution of this paper is a novel algorithm that computes the mutual information between a map and sensor measurements made at all $|M|$ cells in the map in just $O(|\Theta||M|)$ time, improving upon the state-of-the-art by a factor of $O(l)$. The key to this improvement is a recursive structure, as shown in Figure 2, that results from a novel, continuous occupancy map framework. The algorithm is therefore titled **F**ast **C**ontinuous **M**utual **I**nformation (FCMI). In realistic experiments, we observe that this speedup makes FCMI roughly two orders of magnitude faster than the state-of-the-art. Specifically, we observe up to a 300 times speed up when compared to the FSMI algorithm, which was reported to be roughly twice as fast as CSQMI in experiments [18]. This speedup makes it practical to exhaustively compute the mutual information between a map and all possible measurement states as shown in Figure 1b. These dense information surfaces are applicable to a zoo of exploration strategies that involve maximizing expected information gain. We also provide the results of experiments that verify the correctness of FCMI.

This paper is organized as follows. Section II discusses certain technical assumptions, introduces our notation and presents the continuous occupancy map framework. Section III derives the FCMI algorithm for computing the mutual information between an occupancy map and sensor measurements made at all cells within that map. Section IV presents the results of experiments. Section V concludes.

## II. Assumptions, Definitions and Notation

### A. Technical Assumptions

In this section, we summarize the various technical assumptions of our work, for transparency. Notably, we have relaxed two assumptions that were necessary in [6], [13], [18]: we do not assume an inverse sensor model nor that the intersections between a ray and a grid are of a constant size.

*1) Sensor Certainty:* We assume that range sensors measurements are distorted by an arbitrarily small but nonzero amount of additive noise. This contrasts with other rigorously derived information metrics which consider the sensor noise to be additive and Gaussian [6], [13], [18]. Our reasoning is that many modern lidar sensors in use, like the Velodyne Puck or Hokuyo UST-10, have a very small amount of additive noise relative to their maximum range. Any perceivable noise is more likely to be from localization error or object reflections and won't necessarily follow a Gaussian model. Moreover, the authors in [18] observe that changing the sensor noise distribution from Gaussian to uniform makes an imperceptible difference in results, which suggests that accounting for this type of sensor noise may not contribute much to exploration.

*2) Sensor Range:* We assume that range sensors have no maximum range which is necessary to prove the recursive formula presented in Lemma 1. This assumption is clearly unrealistic and differs from most previous work. However, we argue that in a cluttered environment the difference is irrelevant. We also believe that this assumption can be relaxed, albeit with some tedious bookkeeping.

*3) Occupancy Independence:* We assume that the occupancy probabilities of nonoverlapping map regions are independent. This is equivalent to the independence assumption used in occupancy grid mapping [1], and it is also necessary for [6], [13], [18]. In structured, predictable environments this assumption is weak, however in cases where robotic exploration is useful, the dissociation becomes reasonable. Additionally, if the sensor noise is negligible, no dependencies will be introduced by the measurement process itself.

*4) Occupancy Cells:* We assume that the occupancy map is discrete and that the cells are small enough that rays passing through a cell approximate rays emanating from a single point.

*5) Localization:* Finally, we assume that localization is perfect. There are several exploration strategies that consider localization uncertainty in tandem with information gain which is helpful in Simultaneous Localization and Mapping (SLAM) applications [19], [20]. While it is important to consider whether the ideas developed here can be extended to include localization error, it is out of scope of this paper.

### B. Occupancy Maps

The following definitions describe a continuous occupancy map, generalizing occupancy grids [1].

**Definition 1.** An **occupancy map** $M$ is a binary-valued random field indexed by points $\mathbf{x} \in \mathbb{R}^n$. The outcome of each binary random variable $M_\mathbf{x}$ is either "free" or "occupied".

**Definition 2.** Each point $\mathbf{x} \in \mathbb{R}^n$ in occupancy map $M$ has an **occupancy** $o_M(\mathbf{x}) \in [0,1]$ and a **vacancy** $v_M(\mathbf{x}) = 1 - o_M(\mathbf{x})$. For any piecewise smooth curve $\mathcal{C} \subset \mathbb{R}^n$ the probability that $M_\mathcal{C}$ is free is equal to the continuous product of vacancies along $\mathcal{C}$,

$$P(M_\mathcal{C} \text{ is free}) = \prod_{\mathbf{x}\in\mathcal{C}} v_M(\mathbf{x})^{ds} \tag{2}$$

$$= \exp\left(\int_\mathcal{C} \log v_M(\mathbf{x})\, ds\right). \tag{3}$$

**Remark 1.** In occupancy grid mapping, works [6], [13], [18] intuitively consider the probability that a light beam passes through cells with occupancy probabilities $o_1 \dots o_m$ to be

$$\prod_{i=1}^{m}(1 - o_i). \tag{4}$$

Definition 2 generalizes this property to continuous space.

In this generalization however, one cannot equate the occupancy $o_M(\mathbf{x})$ to the probability that $M_\mathbf{x}$ is occupied. To shed some insight on the relation between occupancy as we've defined it and occupancy probability [1], consider a space with constant vacancy $v$. By Definition 2,

$$P(M_\mathcal{C} \text{ is free}) = v^{\int_\mathcal{C} ds}. \tag{5}$$

$\int_\mathcal{C} ds$ is simply the arc length of $\mathcal{C}$. If $\mathcal{C}$ has unit length then the probability that $M_\mathcal{C}$ is free is exactly equal to the vacancy. As $\mathcal{C}$ shrinks, the probability that $M_\mathcal{C}$ is free grows exponentially and vice versa. A particularly useful consequence is that occupancy and vacancy, unlike dimensionless occupancy and vacancy probabilities, are invariant to the map resolution; dividing a region in two decreases the probability that the region is occupied, but does not effect its occupancy.

### C. Range Measurements

Similarly, the following definition generalizes the returns of many sensors including sonar, radar, and lidar.

**Definition 3.** $R$ is a random field consisting of **range measurements** $R_{\mathbf{x},\Theta}$. Each $R_{\mathbf{x},\Theta}$ evaluates to a set of distances that measure from source point $\mathbf{x} \in \mathbb{R}^n$ to the nearest occupied points of $M$ within field of view $\Theta \subset \mathbb{R}^{n-1}$. Each angular coordinate $\boldsymbol{\theta} = (\theta_1, \dots, \theta_{n-1}) \in \Theta$ denotes a *beam*.

In this paper, we make a symbolic distinction between ground truth range measurements $R_{\mathbf{x},\Theta}$ and distorted range measurements $Z_{\mathbf{x},\Theta}$. Although we assume that the noise is arbitrarily small, it must be nonzero, otherwise mutual information is undefined in a continuous setting [12].

## III. Computing Mutual Information

This section derives the FCMI algorithm that computes the mutual information $I(M; Z_{\mathbf{x},\Theta})$ between a map $M$ and range measurements $Z_{\mathbf{x},\Theta}$ for all cell centers $\mathbf{x}$ in $M$. The derivation is organized as follows. Section III-A establishes a general formula for $I(M; Z_{\mathbf{x},\Theta})$. Section III-B completes the formula by deriving an expression for the probability density function of a range measurement. With the goal of making the formula computationally tractable, Section III-C investigates how a piecewise map simplifies the probability density function. This analysis reveals that components of $I(M; Z_{\mathbf{x},\Theta})$ can be recursively defined in terms of components of $I(M; Z_{\mathbf{y},\Theta})$ for points $\mathbf{y}$ adjacent to $\mathbf{x}$. This result is key to giving the FCMI algorithm state-of-the-art complexity in Section III-D.

Due to space limitations, we only sketch the key points of each proof.

### A. Formulating Mutual Information in a Continuous Setting

**Theorem 1.** *Suppose measurements $Z_{\mathbf{x},\Theta}$ are range measurements $R_{\mathbf{x},\Theta}$ with distances distorted by additive, independent, identically distributed noise $N$. Then, the mutual information between map $M$ and the measurements $Z_{\mathbf{x},\Theta}$ is the sum of the mutual information between $M$ and individual measurement distances $Z_{\mathbf{x},\boldsymbol{\theta}}$ for beams $\boldsymbol{\theta} \in \Theta$.*

$$I(M; Z_{\mathbf{x},\Theta}) = \int_{\boldsymbol{\theta}\in\Theta} I(M; Z_{\mathbf{x},\boldsymbol{\theta}}). \tag{6}$$

*Additionally, $I(M; Z_{\mathbf{x},\boldsymbol{\theta}})$ is equal to the difference of differential entropies,*

$$I(M; Z_{\mathbf{x},\boldsymbol{\theta}}) = h(Z_{\mathbf{x},\boldsymbol{\theta}}) - h(Z_{\mathbf{x},\boldsymbol{\theta}}|M), \tag{7}$$

*which can be defined in terms of probability density functions $f_{Z_{\mathbf{x},\boldsymbol{\theta}}}(r)$ and $f_{R_{\mathbf{x},\boldsymbol{\theta}}}(r)$ of the distorted and ground truth distances as follows:*

$$h(Z_{\mathbf{x},\boldsymbol{\theta}}) = -\int_{r\in\mathbb{R}} f_{Z_{\mathbf{x},\boldsymbol{\theta}}}(r) \log f_{Z_{\mathbf{x},\boldsymbol{\theta}}}(r)\, dV \tag{8}$$

$$h(Z_{\mathbf{x},\boldsymbol{\theta}}|M) = -\int_{\rho\in\mathbb{R}} f_{R_{\mathbf{x},\boldsymbol{\theta}}}(\rho) \int_{r\in\mathbb{R}} f_N(r-\rho) \log f_N(r-\rho)\, dV d\rho. \tag{9}$$

Equation (6) results from the independence of measurement distances $Z_{\mathbf{x},\boldsymbol{\theta}}$, which is implied by Definition 2 and the noise assumptions. The definition of differential entropy [12] provides the rest of the equations, with the additivity of the noise needed in Equation (9). Note that the volume element is $dV = r^{n-1} dr d\boldsymbol{\theta}$.

### B. Density Functions of Range Measurements

Evaluating Equations (8) and (9) from Theorem 1 requires the knowledge of the probability density functions $f_{Z_{\mathbf{x},\boldsymbol{\theta}}}$, $f_{R_{\mathbf{x},\boldsymbol{\theta}}}$, and $f_N$. Since the noise is additive and independent, the distorted distribution is a convolution of the other two, *i.e.*, $f_{Z_{\mathbf{x},\boldsymbol{\theta}}} = f_{R_{\mathbf{x},\boldsymbol{\theta}}} \star f_N$. The noise distribution will be given explicitly in Section III-D which leaves the distribution of the range measurement distance $f_{R_{\mathbf{x},\boldsymbol{\theta}}}$ as the only unknown.

In this section, we develop an expression for $f_{R_{\mathbf{x},\boldsymbol{\theta}}}$. Without loss of generality, we consider probability distributions of measurement made along a line. With a slight abuse of notation, we index the map and measurement fields by real numbers with the intent of indicating points along that line defined by arbitrary angle $\boldsymbol{\theta}$ and translation $\mathbf{t}$, *i.e.*,

$$v_M(x) = v_M(x\hat{\mathbf{u}}(\boldsymbol{\theta}) + \mathbf{t}) \quad \text{and} \quad R_x = R_{x\hat{\mathbf{u}}(\boldsymbol{\theta})+\mathbf{t},\boldsymbol{\theta}}. \tag{10}$$

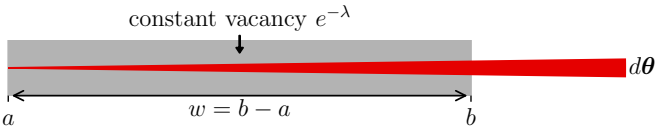


Fig. 3. A visualization of the notation presented in Lemma 1

A range measurement, shown in red, is at made from point $a$ and passes through point $b$. The region in between $a$ and $b$ shown in grey has width $w$ and constant vacancy $e^{-\lambda}$.

where $\hat{\mathbf{u}}(\boldsymbol{\theta})$ is a unit vector pointing in the $\boldsymbol{\theta}$ direction.

**Theorem 2.** *The probability density function of the range measurement distance $R_x$ is the following for $r \geq 0$:*

$$f_{R_x}(r) = -\exp\left(\int_x^{x+r} \log v_M(s)\, ds\right) \log v_M(x+r). \tag{11}$$

The proof follows directly from Definition 2.

### C. A Piecewise Vacancy Assumption

Applying the result of Theorem 2 to the equations in Theorem 1 forms a complete expression for the mutual information $I(M; Z_{\mathbf{x},\Theta})$. However, it is by no means practical to perform a *quadruple* integration numerically. In this section, we apply the assumption that the map is discrete, like an occupancy grid [1], so that vacancies within the map are piecewise constant. This gives rise to the foreshadowed recursive structure. In the rest of our derivation, we use the notation introduced in Lemma 1 to describe a piecewise region and its vacancy. This notation is visualized in Figure 3.

**Lemma 1.** *Suppose that $v_M(x) = e^{-\lambda}$, a constant for $a \leq x < b$ with $\lambda \geq 0$. Let $w = b - a$ be the width of this constant region. $f_{R_a}(r)$ has the following recursive form.*

$$f_{R_a}(r) = \begin{cases} 0 & r < 0 \\ \lambda e^{-\lambda r} & 0 \leq r < w \\ e^{-\lambda w} f_{R_b}(r-w) & r \geq w \end{cases} \tag{12}$$

Proof of Lemma 1 follows from Theorem 2. Informally, Lemma 1 shows that near its source, each measurement has an exponential distribution. Otherwise, each is defined recursively in terms of further measurement distributions.

Using Lemma 1 we conclude Corollaries 1 through 4, which together provide solutions to expected values

$$\mathbb{E}[R_a^k] = \int_0^\infty f_{R_a}(r) r^k\, dr \quad \text{and} \tag{13}$$

$$\mathbb{E}[R_a^k I(R_a)] = -\int_0^\infty f_{R_a}(r) r^k \log f_{R_a}(r)\, dr \tag{14}$$

for any $k \in \mathbb{Z}$. These values will be used to compute the mutual information in Theorem 1; note the similarity between Equations (14) and (8) and Equations (13) and (9) (recall that $r^{n-1}$ is hidden in the volume element, $dV$). Properly defining the sensor noise distribution in Section III-D will resolve the remaining differences.

Corollaries 1 and 2 have a closed form, while Corollaries 3 and 4 are defined recursively, which makes it possible to compute $\mathbb{E}[R_a^k]$ and $\mathbb{E}[R_a^k I(R_a)]$ and therefore $I(M; Z_a)$ in constant time if $\mathbb{E}[R_b^k]$ and $\mathbb{E}[R_b^k I(R_b)]$ are known.

**Corollary 1.** *The $R_a < w$ contribution to $\mathbb{E}[R_a^k]$ is*

$$\int_0^w f_{R_a}(r) r^k\, dr = \lambda^{-k} \gamma(k+1, \lambda w). \tag{15}$$

**Corollary 2.** *The $R_a < w$ contribution to $\mathbb{E}[R_a^k I(R_a)]$ is*

$$-\int_0^w f_{R_a}(r) r^k \log f_{R_a}(r)\, dr = \lambda^{-k}(\gamma(k+2, \lambda w) - \gamma(k+1, \lambda w) \log \lambda). \tag{16}$$

**Corollary 3.** *The $R_a > w$ contribution to $\mathbb{E}[R_a^k]$ is*

$$\int_w^\infty f_{R_a}(r) r^k\, dr = e^{-\lambda w} \sum_{i=0}^k \binom{k}{i} w^{k-i} \mathbb{E}[R_b^i]. \tag{17}$$

**Corollary 4.** *The $R_a > w$ contribution to $\mathbb{E}[R_a^k I(R_a)]$ is*

$$-\int_w^\infty f_{R_a}(r) r^k \log f_{R_a}(r)\, dr = e^{-\lambda w} \sum_{i=0}^k \binom{k}{i} w^{k-i} (\mathbb{E}[R_b^i I(R_b)] + \lambda w \mathbb{E}[R_b^i]). \tag{18}$$

Corollaries 1 and 2 reference the following function:

**Definition 4.** The **lower incomplete gamma function** is defined as [21]:

$$\gamma(k+1, x) = \int_0^x t^k e^{-t}\, dt \tag{19}$$

The recurrence relation $\gamma(k+1, x) = k\gamma_a(k) - x^k e^{-x}$ gives the function a closed-form expression for any $k \in \mathbb{Z}_{\geq 0}$, with base case $\gamma(1, x) = 1 - e^{-x}$.

### D. For Barely Distorted Range Measurements

Given Corollaries 1 through 4, it appears that we have the resources to compute the mutual information between the map and distortionless range measurements $R_{\mathbf{x},\boldsymbol{\theta}}$. Unfortunately, in the absence of any noise $h(R_{\mathbf{x},\boldsymbol{\theta}}|M) = -\infty$, a bizarre caveat of information in a continuous setting [12].

To solve this, consider range measurements that are *barely distorted* as follows.

**Lemma 2.** *Suppose the measurement noise $N$ has probability density function $f_N(r) = \Lambda e^{-\Lambda r}$ for $r \geq 0$ and $f_N(r) = 0$ otherwise. Then, if $\Lambda \gg \lambda$ then $f_{Z_a}(r)$ approximates $f_{R_a}(r)$ for $r \leq w$ and if $\lambda \gg \Lambda$ then $f_{Z_a}(r)$ approximates $f_N(r)$ for $r \leq w$.*

The barely distorted measurement $f_{Z_a}$ can be approximated by clipping values of $\lambda$ larger than $\Lambda$ to $\Lambda$ and then applying the formula for $f_{R_a}$ given by Lemma 1. This ensures the that both limits in Lemma 2 hold. As $\Lambda$ grows, the range of vacancies $e^{-\lambda}$ that are not well approximated shrinks exponentially.

A large value of $\Lambda$ also lets one simplify Equation (9) from Theorem 1.

**Lemma 3.** *If $Z_{\mathbf{x},\boldsymbol{\theta}}$ is barely distorted with sufficiently large $\Lambda$ then the following approximation holds.*

$$h(Z_{\mathbf{x},\boldsymbol{\theta}}|M) \approx (1 - \log \Lambda) \mathbb{E}[R_{\mathbf{x},\boldsymbol{\theta}}^{n-1}] d\boldsymbol{\theta}. \tag{20}$$

With that, the main result:

**Theorem 3.** *Suppose map $M$ contains $|M|$ cells that each have constant vacancy and all other space occupied. Then, Algorithm 1 (FCMI) approximates the mutual information $I(M; Z_{\mathbf{x},\Theta})$ between the map and barely distorted range measurements $Z_{\mathbf{x},\Theta}$ made from every cell center $\mathbf{x}$. The field of view $\Theta$ is constant and quantized into $|\Theta|$ beams with width $d\boldsymbol{\theta}$. The algorithm runs in time $O(|\Theta||M|)$.*

In Algorithm 1, $I_{\mathbf{x}}$ approximates $I(M; Z_{\mathbf{x},\Theta})$. $\alpha_k$ and $\beta_k$ approximate $\mathbb{E}[Z^k_{\mathbf{x},\boldsymbol{\theta}} I(Z_{\mathbf{x},\boldsymbol{\theta}})]$ and $\mathbb{E}[Z^k_{\mathbf{x},\boldsymbol{\theta}}]$ respectively. If the noise is small enough then according to Theorem 1, $h(Z_{\mathbf{x},\boldsymbol{\theta}}) \approx \alpha_{n-1} d\boldsymbol{\theta}$. According to Lemma 2, $h(Z_{\mathbf{x},\boldsymbol{\theta}}|M) \approx (1 - \log\Lambda)\beta_{n-1} d\boldsymbol{\theta}$. Therefore, according to Theorem 1, Lines 23 and 24 accumulate the mutual information contribution of beam $\boldsymbol{\theta}$ to a measurement at cell $\mathbf{x}$.

To compute $\alpha_k$ and $\beta_k$, Corollaries 1 through 4 are applied (Lines 17-22). This constant-time update step makes it possible to compute the mutual information contribution of beam $\boldsymbol{\theta}$ to all cells a ray intersects in a single pass. By casting rays that touch all $|M|$ cells in all $|\Theta|$ orientations, the algorithm takes $O(|\Theta||M|)$.

**Remark 2.** $n$ is the dimension of the of the space and for physical exploration $n \leq 3$, so it is omitted from FCMI's runtime.

**Remark 3.** FCMI is approximate for two reasons: the barely distorted approximation and the finite cell size. The parameter $\Lambda$ can be made arbitrarily large without effecting computation time. In practice we set $\Lambda = 10^{100}$ which makes this approximation numerically imperceptible. More subtly, the beams used to compute the mutual information at each cell center $\mathbf{x}$ are not truly radial. The beams contributing to mutual information at cell center $\mathbf{x}$ start at the boundary of $\mathbf{x}$'s cell and point through it. This approximates a truly radial scan if the cells are small as discussed in Section II-A.

**Remark 4.** FCMI as presented computes $I(M; Z_{\mathbf{x},\Theta})$ with a constant field of view $\Theta$. However in many cases it is valuable to have a variable $\Theta$ to account for different orientations of the sensor. In some cases FCMI can be augmented to this end with no change to complexity.

For example, consider a planar sensor that has field of view $\Theta = [0, d\theta \ldots \pi]$ in one particular orientation. FCMI can compute $I(M; Z_{\mathbf{x},\Theta})$ at all cells from this fixed orientation in $O(|\Theta||M|)$. Then to compute $I(M; Z_{\mathbf{x},\Theta+d\theta})$, the mutual information between the map and measurements with a slightly rotated field of view, one can add $I(M; Z_{\mathbf{x},\pi+d\theta})$ and subtract $I(M; Z_{\mathbf{x},0})$ from all cells. This can be done by in $O(|M|)$ via two iterations of Loop 5-26. A sweep of all quantized orientations still takes $O(|\Theta||M|)$.

## IV. Experimental Evaluation

In this section, we investigate the proposed FCMI algorithm empirically. Section IV-A measures the speed of FCMI and compares it to FSMI [18]. Section IV-B verifies the accuracy of the algorithm, both in terms of the integration it evaluates and as a measure of actual expected information gain. Finally, Section IV-C studies a hidden parameter of the algorithm, occupancy probability initialization.

**Algorithm 1** **F**ast **C**ontinuous **M**utual **I**nformation (FCMI)

**Require:** $n$-dimensional occupancy map $M$ containing $|M|$ constant vacancy cells and all other space occupied; field of view $\Theta$; large distortion parameter $\Lambda$

*Initialize the mutual information to zero*
1: **for** each cell center $\mathbf{x}$ in map $M$ **do**
2: $I_{\mathbf{x}} \leftarrow 0$
3: **end for**

4: **for** $\boldsymbol{\theta} \in \Theta$ **do**
5: Trace rays at angle $\boldsymbol{\theta}$ that start and end in occupied
6: space such that each cell in $M$ is hit exactly once.
7: **for** each ray $T$ **do**

*Initialize the expected values*
8: **for** $k = 0$ to $n - 1$ **do**
9: $\alpha_k \leftarrow \Lambda^{-k}((k+1)! - k! \log\Lambda)$
10: $\beta_k \leftarrow \Lambda^{-k} k!$
11: **end for**

12: **for** each cell that $T$ intersects in reverse **do**
13: $\mathbf{x} \leftarrow$ the center of the cell
14: $w \leftarrow$ the width of $T$ passing through the cell
15: $e^{-\lambda} \leftarrow$ the constant vacancy of the cell

*Clip to approximate small distortion*
16: $\lambda \leftarrow \min(\lambda, \Lambda)$

*Update the expected values*
17: **for** $k = n - 1$ to $0$ **do**
18: $\alpha_k \leftarrow e^{-\lambda w} \sum_{i=0}^{k} \binom{k}{i} w^{k-i} (\alpha_i + \lambda w \beta_i)$
19: $\alpha_k \leftarrow \alpha_k + \lambda^{-k}(\gamma(k+2, \lambda w) - \gamma(k+1, \lambda w)\log\lambda)$
20: $\beta_k \leftarrow e^{-\lambda w} \sum_{i=0}^{k} \binom{k}{i} w^{k-i} \beta_i$
21: $\beta_k \leftarrow \beta_k + \lambda^{-k}\gamma(k+1, \lambda w)$
22: **end for**

*Accumulate mutual information*
23: $I_{\mathbf{x}} \leftarrow I_{\mathbf{x}} + \alpha_{n-1} d\boldsymbol{\theta}$
24: $I_{\mathbf{x}} \leftarrow I_{\mathbf{x}} - (1 - \log\Lambda)\beta_{n-1} d\boldsymbol{\theta}$
25: **end for**
26: **end for**
27: **end for**

The software is available open-source at `https://github.com/sportdeath/range_mi`.

### A. The Speed of FCMI

In our first timing experiment we use a single threaded 1.9GHz i7-8650U CPU, comparable with the 2.10GHz Intel Xeon E5-2695 CPU used in the FSMI experiments [18]. Our implementation of FCMI computes the mutual information contribution of a single beam pointing in the positive direction at every cell in a random, 1-dimensional, 100-cell map in 2.3 microseconds on average. For comparison, the fastest reported variant of the FSMI algorithm (Uniform FSMI) computes the mutual information contribution at a *single* cell with a beam length of 100 cells in 10 microseconds [18]. By extrapolating, computing the mutual information at every cell takes FSMI 505 microseconds, making the FCMI 219 times faster than FSMI, in this case.

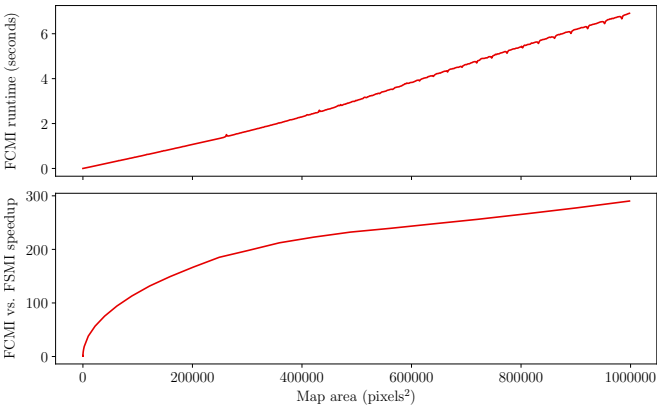


Fig. 4. Timing on square maps with $|\Theta| = 200$ beam angles

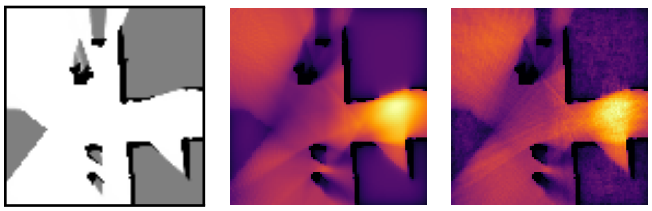

(a) An incomplete map (b) FCMI output (c) Simulated info gain

Fig. 5. Comparison of FCMI to average simulated information gain

In our second experiment we look at the time it takes to compute 2D mutual information surfaces as a function of map area on a 3.7GHz i7-8700K CPU. Figure 4 shows that FCMI computes mutual information surfaces of sizes up to $1000 \times 1000$ in seconds with computation time scaling linearly with map area. Compared to our optimized implementation of the FSMI algorithm, Figure 4 shows that FCMI is hundreds of times faster with a speedup that scales with the root of the map area.

### B. The Accuracy of FCMI

We performed two experiments to verify that the computation from FCMI actually produces the correct result.

In the first, we compute the mutual information numerically in random, 1-dimensional, 100-cell maps according the to the integrals provided in Theorems 1 and 2. When the integration step size is $10^{-6}$, the difference is on average $< 0.02\%$ of FCMI's output. Decreasing the step size reduces the error further. Similar results hold for dimensions $1 < n \leq 5$ (what 5-dimensional exploration could is used for is another question). The experiment also verifies that the mutual information is always non-negative as required axiomatically [12].

In the second experiment, we compare the output of FCMI to average simulated information gain. To compute the average simulated information gain of a single measurement location, we randomly generate "ground truth" values for an incomplete occupancy map based on the map's occupancy probabilities. Then, we use those values to simulate a measurement from the desired location and update the map accordingly. Averaging the change in entropy of the map (computed as the sum of binary entropies) over many possible map assignments gives the claimed result.

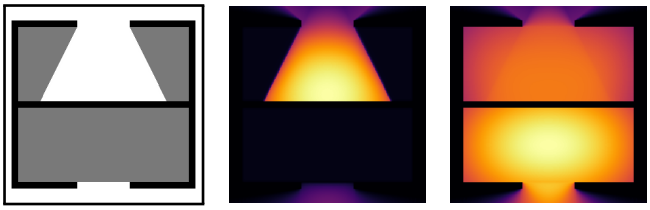

(a) Two cubicles (b) 50% occupied (c) 1% occupied

Fig. 6. The effect of occupancy initialization on mutual information

Figure 5 shows how the average simulated information gain converges to values that are practically indistinguishable from the output of the FCMI algorithm. Some noise remains in the simulated surface because of its Monte Carlo computation and long runtime.

### C. The Effect of the Unknown Prior on Mutual Information

It is common practice to initialize occupancy probabilities to $1/2$ which authors claim is a state of zero information [1], [6], [13], [18]. However it is apparent from definition 2 or the discrete equivalent mentioned in remark 1 that this practice is synonymous to the assumption that a measurement beam is expected to pass through $\approx 2$ cells before hitting an obstacle. 2 cells is an entirely arbitrary number and in a map with a fine resolution, a very small distance.

Now, treating this initialization value as a free parameter, we study how it effects mutual information.

Figure 6a features two cubicles, the top one partially explored and the bottom one unexplored. If the initialized vacancy is relatively low, the mutual information is high in the top cubicle. This is because measurement beams are not expected to travel very far into unknown space and so mutual information effectively measures the perimeter of unknown space visible to the sensor. Conversely, if the vacancy is high, the most information is gained by viewing the largest volume of unknown space which is found in the bottom cubicle. Effectively, the initialization parameter interpolates between the two classic information gain heuristics discussed in the introduction — no wonder they work!

## V. Conclusion

In this work we derived the FCMI algorithm for computing the mutual information between an occupancy map and a range sensor measurement that has a lower complexity and is hundreds of times faster than the state of the art. This makes it possible to exhaustively compute the measurement state that maximizes expected information gain which is useful for a host of exploration strategies.

It would be interesting for future work to consider making the information metric account for various types of noise like localization error or sensor distortion while maintaining the same complexity. The algorithm presented is also ripe for parallelism and we believe with a GPU implementation, it would be possible to evaluate mutual information surfaces at sensor rate.

## REFERENCES


[1] Larry Matthies and Alberto Elfes. Integration of sonar and stereo range data using a grid-based representation. In *Proceedings of the 1988 IEEE International Conference on Robotics and Automation*, pages 727–733. IEEE, 1988.

[2] Christopher I. Connolly. The determination of next best views. In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, volume 2, pages 432–435. IEEE, 1985.

[3] Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon "next-best-view" planner for 3D exploration. In *Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1462–1468. IEEE, 2016.

[4] Héctor H. González-Banos and Jean-Claude Latombe. Navigation strategies for exploring indoor environments. *The International Journal of Robotics Research*, 21(10-11):829–848, 2002.

[5] Francesco Amigoni and Alessandro Gallo. A multi-objective exploration strategy for mobile robots. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 3850–3855. IEEE, 2005.

[6] Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. Information-theoretic mapping using Cauchy-Schwarz quadratic mutual information. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4791–4798. IEEE, 2015.

[7] Richard Pito. A solution to the next best view problem for automated surface acquisition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(10):1016–1030, 1999.

[8] S. Y. Chen and Y. F. Li. Vision sensor planning for 3-D model acquisition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 35(5):894–904, 2005.

[9] Brian Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA'97)*.

[10] Kok-Lim Low and Anselmo Lastra. Efficient constraint evaluation algorithms for hierarchical next-best-view planning. In *Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT'06)*, pages 830–837. IEEE, 2006.

[11] Frederic Bourgault, Alexei A. Makarenko, Stefan B. Williams, Ben Grocholsky, and Hugh F. Durrant-Whyte. Information based adaptive robotic exploration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 540–545. IEEE, 2002.

[12] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.

[13] Brian J. Julian, Sertac Karaman, and Daniela Rus. On mutual information-based control of range sensing robots for mapping applications. *The International Journal of Robotics Research*, 33(10):1375–1392, 2014.

[14] Deniz Erdogmus. *Information Theoretic Learning: Renyi's Entropy and Its Applications to Adaptive System Training*. PhD thesis, University of Florida Gainesville, Florida, 2002.

[15] Claude Elwood Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[16] Flávio du Pin Calmon and Nadia Fawaz. Privacy against statistical inference. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1401–1408. IEEE, 2012.

[17] Neri Merhav and Meir Feder. Universal prediction. *IEEE Transactions on Information Theory*, 44(6):2124–2147, 1998.

[18] Zhengdong Zhang, Theia Henderson, Sertac Karaman, and Vivienne Sze. FSMI: Fast computation of Shannon mutual information for information-theoretic mapping. In *Proceedings of the 2019 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6912–6918. IEEE, 2019.

[19] Luca Carlone, Jingjing Du, Miguel Kaouk Ng, Basilio Bona, and Marina Indri. Active SLAM and exploration with particle filters using Kullback-Leibler divergence. *Journal of Intelligent & Robotic Systems*, 75(2):291–311, 2014.

[20] Henry Carrillo, Philip Dames, Vijay Kumar, and José A. Castellanos. Autonomous robotic exploration using occupancy grid maps and graph SLAM based on Shannon and Rényi entropy. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 487–494. IEEE, 2015.

[21] Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables*, volume 55. Courier Corporation, 1965.