

# Functional Path Optimisation for Exploration in Continuous Occupancy Maps

Gilad Francis, Lionel Ott and Fabio Ramos

**Abstract** Autonomous exploration is a complex task where the robot moves through an unknown environment with the goal of mapping it. The desired output of such a process is a sequence of paths that efficiently and safely minimise the uncertainty of the resulting map. However, optimising over the entire space of possible paths is computationally intractable. Therefore, most exploration methods relax the general problem by optimising a simpler one, for example finding the single next best view. In this work, we formulate exploration as a variational problem which allows us to directly optimise in the space of trajectories using functional gradient methods, searching for the *Next Best Path* (NBP). We take advantage of the recently introduced Hilbert maps to devise an information-based functional that can be computed in closed-form. The resulting trajectories are continuous and maximise safety as well as mutual information. In experiments we verify the ability of the proposed method to find smooth and safe paths and compare these results with other exploration methods.

**Key words:** Robotic Exploration, Functional Gradient, Path Planning

## 1 Introduction

The objective of autonomous exploration is to produce a consistent representation of the environment. It involves complex decision-making, selecting the trajectories a robot should take in order to minimise the overall uncertainty in the model. Essentially, exploration is a path optimisation procedure to find trajectories that efficiently learn the environment. The difficulty lies in the dimensionality and shape of the search space which prohibits a closed-form solution to the general exploration problem, making autonomous exploration an active field of research. The plethora of exploration methods in the literature offer different strategies for relaxing this

---

Gilad Francis (e-mail: gfra8070@uni.sydney.edu.au), Lionel Ott and Fabio Ramos  
School of Information Technologies, The University of Sydney, Australia

intractable problem, most commonly by reducing the search space dimensionality, for example by discretising the path.

In this work, we formalise exploration as a variational problem. We present a novel approach based on *functional gradient descent* (FGD) to efficiently optimise exploratory paths over continuous occupancy maps. We use stochastic FGD to overcome the limitations of standard FGD methods in order to ensure convergence. This process enables optimisation over the entire path, resulting in continuous smooth paths that maximise the overall map quality while keeping the robot safe from collisions. Our contributions are:

1. A *Next Best Path* method. An information-driven variational framework for safe autonomous exploration in continuous occupancy maps. Path optimisation is performed directly in the space of trajectories using a combined objective; which considers safety, efficiency and information. The method is invariant to the choice of path representation as it uses stochastic functional gradient descent to optimise the objective along the entire path.
2. Developing a *mutual information* (MI) variational objective for continuous occupancy maps. This replaces the common and expensive approach of computing MI explicitly over the entire map for each evaluated path. Instead, our method modifies the path using MI functional gradients without the need to compute MI explicitly. These gradients are obtained from local perturbations on the map model which are derived in closed-form.

The remainder of this paper is organised as follows. Literature on autonomous exploration is surveyed in Section 2. The basic building blocks, such as Hilbert maps and FGD are reviewed in Section 3. Section 4 describes in detail the functional exploration algorithm. Experimental results and analysis are presented in Section 5. Finally, Section 6 draws conclusions on the proposed method.

## 2 Related Work

The goal of autonomous exploration is to produce a consistent environment model by minimising any uncertainties. In a mapping context, exploration is the process of producing high-fidelity maps [21]. This is a complex problem mainly due the dimensionality of the solution space. Most exploration methods use occupancy grid maps in their planning [22], rather than continuous occupancy maps. Regardless of the type of occupancy map, exploration methods take one of two forms; frontier-driven or information-theoretic. Juliá et al. [12] provide a quantitative comparison between these exploration methods.

Frontier-based exploration methods drive the robot toward the borders of the known space [25]. In a grid map, frontiers are clusters of free cells neighbouring unknown cells. Once frontiers are identified, a separate path planner finds a safe path toward a selected frontier. Various utility functions can be used when choosing the most desirable frontier. The simplest form considers only the travelling costs [25]. González-Baños and Latombe [8] choose a goal point based on a score of expected coverage penalised by the travelling costs. A generalised approach for goal point selection given several criteria were suggested by Holz et al. [9].

Information-theoretic exploration methods optimise a utility function associated with the uncertainty of the map. Early work optimised the selection of a discrete goal point rather than optimising an entire path. Elfes [4] suggested MI as an information metric for exploration, while Whaite and Ferrie [24] proposed a *next best view* (NBV) approach using the entire map entropy. Vallvé and Andrade-Cetto [23] use a potential field computed over the entire configuration space to find exploration candidates. However, this method assumes discrete steps, disregarding the reduction of entropy between consecutive robot poses. Charrow et al. [1] combined the frontier and information-based methods. While they optimised the information heuristic over a continuous control input space, in effect the path consisted of a fixed number of time steps. Lauri and Ritala [14] formulate the exploration problem as *partially observable Markov decision process* (POMDP) and used sample-based approach to solve the POMDP. Similarly to the work of Charrow et al., the action space is continuous, however, the path consists of a finite set of time steps, which is in contrast to the proposed method where optimisation is performed in the space of trajectories.

Only a handful of algorithms tackle exploration in continuous occupancy maps. Yang et al. [26] employed *rapidly-exploring random tree* (RRT) to sample a set of feasible path candidates for a *Gaussian Processes* (GPs) maps. An adaptation of the frontier method for continuous occupancy maps was introduced by Jadidi et al. [10], where a discretised frontier map was built from the continuous map. More recently, Jadidi et al. [11] employed MI to rank frontiers. *Bayesian optimisation* (BO) has also been used for exploration. Marchant and Ramos [15] utilised BO to optimise the selection of continuous informative paths over a continuous environmental model. Francis et al. [5] used constrained BO for safe exploration to learn an MI objective. While BO optimises continuous paths, in practice it uses only a limited parameterisation such as quadratic or cubic splines.

In summary, the exploration method proposed in this paper uses an information-based utility to optimise path selection. Employing calculus of variations, the optimisation procedure is invariant to the path representation. This is a major difference from existing methods that typically depend on a finite path parametrisation, such as finite sets of time steps or waypoints, or employ simple representations such as quadratic or cubic splines. Instead, our method can utilise a highly expressive path representation; such as non-parametric [6] or approximate kernel paths [7]. As our method uses continuous occupancy maps, the MI utility can be derived directly from the map model in closed form, which simplifies computations.

### 3 Preliminaries

In this section, we review the basic building blocks of the functional exploration method; Hilbert maps and the functional gradient path planning method. In section 4, we extend these components to support an information-driven safe exploration algorithm.

### 3.1 Hilbert Maps

Hilbert maps [19] are continuous discriminative models that predict occupancy based on sensors observations. Unlike grid maps [3] that discretise space into a set of independent cells, Hilbert maps maintain neighbourhood information. Other continuous mapping methods based on *Gaussian Processes* (GPs) [10, 18] hold similar properties. However, these methods are limited by the scalability of the associated GPs. Hilbert maps, on the other hand, use an approximate kernel logistic regression model, which renders updating and querying the map independent of the dataset size. In addition, *stochastic gradient descent* (SGD) is used to enable real-time performance.

Formally, a Hilbert map is a *logistic regression* (LR) classifier model that predicts occupancy anywhere in the map. It uses nonlinear projection into a *reproducing kernel Hilbert space* (RKHS) in order to represent a real environment. Given a set of weights  $\mathbf{w}$ , the predictive occupancy posterior transforms into:

$$p(y = +1|\mathbf{x}, \mathbf{w}) = \frac{1}{1 + \exp(-\mathbf{w}^T \hat{\Phi}(\mathbf{x}))} = \sigma(\mathbf{w}^T \hat{\Phi}(\mathbf{x})). \quad (1)$$

In an occupancy map context  $\mathbf{x} \in \mathbb{R}^D$  represents either a  $2D$  or  $3D$  location, while  $y \in \{-1, +1\}$  denotes two possible binary outputs, unoccupied and occupied. We define  $\sigma$  as the logistic sigmoid function.  $\hat{\Phi}(\cdot)$  is a set of features that, in expectation, approximate the inner product defined by the kernel function  $k(\cdot, \cdot)$ . This approximation enables fast training of a map model, suitable for use in a robotics setting. There are several methods to approximate the kernel matrix [19]. Given the desired approximation, the weights vector  $\mathbf{w}$  can be then trained by minimising the regularised *negative log-likelihood* (NLL) as commonly performed in logistic regression methods [19].

### 3.2 Functional Gradient Descent

*Functional gradient descent* (FGD) is a variational framework to optimise nonlinear models. It has been successfully applied to motion planning problem in recent years with the main objective of producing safe, collision-free paths. It was recently suggested as an alternative approach to sampling-based methods for path planning using occupancy maps [6]. In this section, the general method is discussed, before the extension for autonomous exploration is described in section 4.

We first introduce notation. A path,  $\xi : [0, 1] \rightarrow C \in \mathbb{R}^D$ , is a function that maps a time-like parameter  $t \in [0, 1]$  into configuration space  $C$ . The objective functional  $U(\xi) : \mathcal{E} \rightarrow \mathbb{R}$  returns a real number for each path  $\xi \in \mathcal{E}$ , corresponding to a cost or loss associated with  $\xi$ . FGD path planners capture path features such as smoothness and safety in  $U(\xi)$ . The goal of the optimisation process is to find a path that minimises the overall costs:

$$\xi_{optimal} = \arg \min_{\xi} U(\xi) \quad (2)$$

Finding the optimal path is performed by following the functional gradient of the objective. This is an iterative process where the functional gradient update rule is derived from a linear approximation of the cost functional around the current trajectory,  $\xi_n$ :

$$U(\xi) \approx U(\xi_n) + \nabla_{\xi} U(\xi_n)(\xi - \xi_n). \quad (3)$$

We enforce small updates by adding a regularisation term based on the norm of the update:

$$\xi_{n+1} = \arg \min_{\xi} U(\xi_n) + (\xi - \xi_n)^T \nabla_{\xi} U(\xi_n) + \frac{1}{2\eta_n} \|\xi - \xi_n\|_{\Lambda}^2. \quad (4)$$

The regularisation term  $\|\xi - \xi_n\|_{\Lambda}^2 = (\xi - \xi_n)^T \Lambda (\xi - \xi_n)$  is the squared norm with respect to a metric tensor  $\Lambda$  and  $\eta_n$  is a user-defined learning rate. By differentiating the right hand side of (4) with respect to  $\xi$ , we obtain the iterative update rule:

$$\xi_{n+1}(\cdot) = \xi_n(\cdot) - \eta_n \Lambda^{-1} \nabla_{\xi} U(\xi_n). \quad (5)$$

We note that (5) forms a general update rule, regardless of the choice of the objective function or the path representation. The only requirements are that  $\Lambda$  is invertible and the gradient  $\nabla_{\xi} U(\xi_n)$  exists.

The general rule for computing the objective functional gradient  $\nabla_{\xi} U$  stems from the calculus of variations. As a variational method, the objective functional must take the form of an integral or sum. Generally, the objective functional  $F$  takes the form  $F(\xi) = \int_a^b v(t, \xi, \xi') dt$ , therefore the functional gradient can be computed using the Euler-Lagrange equation [27]:

$$\nabla F_{\xi}(\xi) = \frac{\partial v}{\partial \xi} - \frac{d}{dt} \frac{\partial v}{\partial \xi'}. \quad (6)$$

These gradients are then used to compute the iterative update rule of (5). Constraints are incorporated using KKT conditions, similar to [16], and are an explicit part of the path representation.

## 4 Functional Exploration

The following section introduces our proposed functional exploration method. The use of functional gradient descent on an information-based objective results in trajectories which are highly expressive and safe while optimising the amount of information gained along the path.

### 4.1 Notation

We first introduce the notation used throughout the following sections. The workspace of the robot  $W \in \mathbb{R}^3$  defines the space where obstacles lie and the map is queried. In addition, to account for the robot's finite size or its pose uncertainty, a set of body points,  $B \in \mathbb{R}^3$  are defined. As the trajectory  $\xi$  lies in configuration space  $C$ , we set a

forward kinematic transform  $g$  that maps a robot configuration,  $\xi(t)$  and body point  $b \in B$  to a point in the workspace  $g : C \times B \rightarrow W$ . To simplify notation, we define  $x_t = g(\xi(t), b)$  as the workspace location for the pair  $(t, b)$ . We assume that the robot is equipped with a sensor, such as laser range finder, with a maximum range  $R_{max}$  and an angular field of view  $\Omega$ .

For a given function  $\xi$ , a functional returns a single value  $U : \mathbb{R}^D \rightarrow \mathbb{R}$ . Functionals are usually represented by an integral. However, whenever we compute the safety or MI functionals, the cost is computed over  $W$ . As a result, the functional must be approximated by a *reduce* operator, e.g., average or maximum, that aggregates the cost along  $\xi(\cdot)$ . Given a workspace cost function  $c(g(\xi(t), b)) : \mathbb{R}^3 \rightarrow \mathbb{R}$ , we can approximate the functional by a sum over a finite set  $\mathcal{T}(\xi) = \{t, b\}_i$  of time and body points:

$$U_{obs,MI}(\xi) \approx \sum_{(t,b) \in \mathcal{T}(\xi)} c_{obs,MI}(g(\xi(t), b)) \equiv \sum_{(t,b) \in \mathcal{T}(\xi)} c_{obs,MI}(x_t) \quad (7)$$

In addition, we note the difference between gradient operators. We define  $\nabla$  as a gradient with respect to  $\xi$ , while  $\nabla_x$  is the workspace gradient.

#### 4.2 Exploration Functional Objective

The goal of autonomous exploration is to safely reduce the uncertainty of the environment model. Some exploration methods compute a finite set of go-to points that locally maximise information gain. Other methods optimise an information-based objective over the entire path, however, these are highly dependent on the path parameterisation. Formulating exploration as a variational problem and solving it using functional gradient descent, provides a general optimisation framework which is invariant to the choice of path parameterisation and can even take a non-parametric form as shown in [6].

The approach taken in our work relies on maximising mutual information along the entire path while keeping the trajectory safe. This is attained by constructing an objective functional  $U$  which is a weighted sum of three components:


- $U_{obs}$  which maintains path safety by penalising proximity to obstacles,
- $U_{dyn}$  which penalises based on the shape of the trajectory, keeping path smooth and short and,
- $U_{MI}$  which rewards the mutual information gained along the path.

The overall objective takes a form of a weighted sum, as is also shown schematically in Figure 1;

$$U(\xi) = \beta_{obs}U_{obs}(\xi) + \beta_{dyn}U_{dyn}(\xi) + \beta_{MI}U_{MI}(\xi). \quad (8)$$

Here  $\beta_{obs}, \beta_{dyn}, \beta_{MI}$  are user-defined coefficients. In the following sections, we will introduce the different components of the functional objective;  $U_{obs}(\xi)$ ,  $U_{dyn}(\xi)$  and  $U_{MI}(\xi)$ . For each component we will derive its functional gradient assuming a Hilbert map as the environment model.

**Fig. 1** Given a sample  $(t, b)$ , the functional  $U$  is the weighted sum of the various objectives; obstacle, dynamic and MI.  $M$  and  $\xi$  are global variables used to compute these objectives.



#### 4.2.1 Obstacle Functional $U_{obs}(\xi)$

Following our previous work [6], we define the workspace cost function  $c_{obs}$  as the map occupancy of (1), i.e.  $c_{obs}(x_t) = p(y = +1|x_t)$ . Given (7), the obstacle functional can be approximated by  $U_{obs}(\xi) \approx \sum_{x_t \in \mathcal{T}(\xi)} c_{obs}(x_t)$ . The functional gradient can then be computed as

$$\nabla U_{obs} = \sum_{x_t \in \mathcal{T}(\xi)} \mathbf{J}(x_t) \nabla_x c_{obs}(x_t). \quad (9)$$

Here,  $\mathbf{J}(x_t)$  is the workspace Jacobian. Since the map model is continuous and at least twice differentiable [19], the spatial gradient of occupancy can be computed in closed-form from (1) as:

$$\nabla_x c_{obs}(x_t) = \nabla_x p(y = +1|x_t) = \sigma(\mathbf{w}^T \hat{\Phi}(x_t)) \left(1 - \sigma(\mathbf{w}^T \hat{\Phi}(x_t))\right) \mathbf{w}^T \nabla_x \hat{\Phi}(x_t). \quad (10)$$

#### 4.2.2 Path Dynamics Functional $U_{dyn}(\xi)$

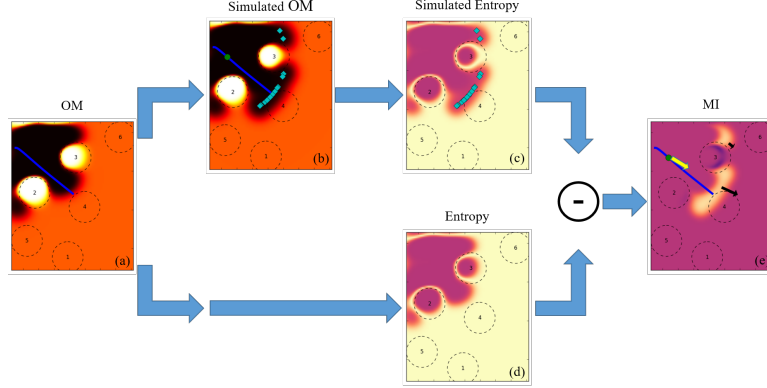
$U_{dyn}(\xi)$  penalises kinematic costs associated with  $\xi$ . The straightforward approach is to regularise on the trajectory length, which can be attained by optimising the integral over the squared velocity norm:  $U_{dyn}(\xi) = \frac{1}{2} \int_0^1 \left\| \frac{d}{dt} \xi(t) \right\|^2 dt$ . Following the Euler-Lagrange equation (6), the functional gradient of  $U_{dyn}$  is

$$\nabla_{\xi} U_{dyn}(\xi(t)) = -\frac{d^2}{dt^2} \xi(t). \quad (11)$$

#### 4.2.3 Mutual Information Functional $U_{MI}(\xi)$

*Mutual information* (MI) is used in many autonomous exploration algorithms as an information-based objective function [13]. In this context, MI is defined as the reduction in entropy conditioned on expected observations. Given an occupancy map  $M$  and a set of expected observations  $\hat{\mathbf{z}}$ , we can define MI as:

$$MI(M; \hat{\mathbf{z}}) = H(M) - H(M|\hat{\mathbf{z}}) \quad (12)$$



**Fig. 2** MI Functional gradient. The MI gradient is computed for a time sample  $t$  given a continuous occupancy map  $OM$  (a). The current path estimate  $\xi$  is depicted in blue, while the expected pose at time  $t$  is shown in green. The modified OM (b) is generated using hypothetical observations, shown as cyan diamonds, based on a robot's configuration at time  $t$ . Only observations at the edge of sensor range are considered for  $t$ . The entropy of the OM (c) and the modified OM (d) can be computed from the occupancy values (high entropy shown in white), which produces MI (e). The MI gradient is estimated by samples around observations (shown as black arrows). Accumulating gradient samples results in the overall MI gradient for  $t$ , shown by the yellow arrow on the robot. Note that the images of occupancy, entropy and MI are only given here for presentation purposes and full maps are never computed. The planner only accesses occupancy, entropy and gradient values through stochastic samples.

where  $H$  denotes Shannon's entropy. The main computational challenge of (12) is resolving the expected observations  $\hat{\mathbf{z}}$ . These are emulated observations that are produced by ray casting based on the sensor model. Determining  $\hat{\mathbf{z}}$  and MI over an entire path is computationally intensive, leading most exploration methods to solve a relaxed MI optimisation problem, by either discretising or parameterising the paths.

The approach taken in this work uses the MI functional  $U_{MI}$  and its gradient to maximise MI efficiently over the entire trajectory. To compute  $U_{MI}$ , an MI reward function  $c_{MI}(x_t) \approx MI(M; \hat{\mathbf{z}}|x_t)$  is computed over the robot's workspace  $W$ , in a similar fashion to the computation of  $U_{obs}$ . However, computation of the conditional entropy  $H(M|\hat{\mathbf{z}})$  entails changes to the Hilbert map model.

In the following section, we will describe the stages involved in computing the MI functional gradient. The three stages executed when computing  $\nabla U_{MI}$  are:

- Simulating expected observations
- Creating a perturbed Hilbert map model  $M|\hat{\mathbf{z}}$
- Obtaining MI functional gradient  $\nabla U_{MI}$

Fig. 2 shows the steps required for MI gradient computations at a given time  $t$ .


#### Simulating expected observations

Similarly to the obstacle functional, the MI functional is approximated by a sum over a finite set of points  $x_t$ ;  $U_{MI}(\xi) \approx \sum_{x_t \in \mathcal{T}(\xi)} c_{MI}(x_t)$ .  $c_{MI}$  is chosen in a way that will estimate the infinitesimal change in MI at  $x_t$ . To compute  $c_{MI}$  we emulate observations by ray casting, as done in other information-driven exploration techniques



[4, 11]. However, while other methods evaluate the MI reward over the entire field of view of the sensor, our method is only interested in the expected observations at the sensor's limits (maximum range). The rationale behind this approach is that new information about the environment will be mainly obtained at the sensor's sensing limits  $R_{max}$ . Fig. 3 illustrates the difference between the two approach to compute MI.

The output of the ray casting process is a set of unoccupied expected observations  $\hat{\mathbf{z}}_f$  for any time and body point  $x_t$ .  $\hat{\mathbf{z}}_f$  differs from the expected observations of (12), as we are only interested in unoccupied (no obstacle) observation at the sensor maximum range. Fig. 2b depicts  $\hat{\mathbf{z}}_f$  as cyan diamonds.



**Fig. 3** Difference in MI calculation.  $x_t$  is the robot's pose for which MI is computed, the black lines depict the sensor field of view and the blue area is the region where MI is computed. (a) MI is computed over the entire field of view of sensor, producing  $\hat{\mathbf{z}}$ . (b) MI is computed only at the sensor's range limits, producing  $\hat{\mathbf{z}}_f$

#### Creating a perturbed Hilbert map model $M|\hat{\mathbf{z}}$

In this step, we generate  $M|\hat{\mathbf{z}}_f$ , the modified Hilbert map conditioned on the expected observations. The straightforward approach is to train a new map model based on  $\hat{\mathbf{z}}_f$ . This approach is commonly used in exploration methods for occupancy grid maps. However the computational costs of such an approach are high, as new maps must be generated along the entire trajectory during optimisation. Instead, we propose the use of a perturbed Hilbert map model.

The perturbed Hilbert map model modifies the predictive map model  $\sigma(\mathbf{w}^T \hat{\Phi}(\mathbf{x}))$  (1) with a perturbed model  $\sigma(\Psi(\mathbf{x}, \hat{\mathbf{z}}_f))$ . This model uses the expected observations as a dataset  $\hat{\mathbf{z}}_f = \{x_i, r_i\}_{i=1}^N$ , where  $x_i \in \mathbb{R}^D$  is a point in 2D or 3D space and  $r_i$  is the log odds of the desired predictive occupancy posterior at  $x_i$ , to fit a *Gaussian process* (GP):

$$\Psi(\mathbf{x}, \hat{\mathbf{z}}_f) \sim GP(\mathbf{w}^T \hat{\Phi}(\mathbf{x}), k(x_i, x_i)). \quad (13)$$

We note that  $\mathbf{w}^T \hat{\Phi}(\mathbf{x})$  of the current occupancy map is the mean function of the GP. The kernel function  $k$  is the same function approximated by the Hilbert map features  $\hat{\Phi}(\cdot)$ . The predictive probability of the perturbed map  $\hat{p}$  is given by:

$$\hat{p}(\mathbf{x}, \hat{\mathbf{z}}_f) = \sigma(\Psi(\mathbf{x}, \hat{\mathbf{z}}_f)) \quad (14)$$

Fig. 2b depicts the resulting occupancy map following the embedding of the expected observations in an existing map.

The computational cost of the perturbed Hilbert map is cubic in the number of expected observations  $|\hat{\mathbf{z}}_f|$ . As we are only modifying the map in a small region, a small set of observations is required to generate perturbation, keeping the computation load low.

#### Obtaining MI functional gradient $\nabla U_{MI}$

The workspace MI cost function for  $x_t$  is defined as the MI summed over the entire map, i.e.  $c_{MI}(x_t) = \int_M MI(M(m); \hat{\mathbf{z}}_f(x_t)) dm$ . However, as we are only interested in the change at the limits of the sensor range,  $c_{MI}$  can be replaced by;  $c_{MI}(x_t) = \int_{m \in m_{max}(x_t)} MI(M(m); \hat{\mathbf{z}}_f(x_t)) dm$ . Where  $m_{max}(x_t)$  denotes workspace locations which lie on the arc given by the maximum sensing range  $R_{max}$  and the sensor's field of view  $\Omega$ , as shown in Fig. 3.  $c_{MI}$  can be approximated with a sum using either a deterministic or a Monte-Carlo schedule:

$$c_{MI}(x_t) \approx \sum_{m \in m_{max}(x_t)} MI(M(m); \hat{\mathbf{z}}_f(x_t)). \quad (15)$$

Given the approximations (7) and (15), the MI functional can be represented as

$$U_{MI}(\xi) \approx \sum_{x_t \in \mathcal{T}(\xi)} \sum_{m \in m_{max}(x_t)} MI(M(m); \hat{\mathbf{z}}_f(x_t)). \quad (16)$$

The MI functional gradient follows the same form as (9):

$$\nabla U_{MI}(x_t) = \sum_{x_t \in \mathcal{T}(\xi)} \sum_{m \in m_{max}(x_t)} \mathbf{J}(x_t) \nabla_x MI(M(m); \hat{\mathbf{z}}_f(x_t)). \quad (17)$$

The spatial gradient of MI,  $\nabla_x MI(M(m))$ , can be expressed in closed-form using the Hilbert maps' continuous model. Using the MI definition from (12), the gradient is defined as  $\nabla_x MI(M; \hat{\mathbf{z}}) = \nabla_x H(M) - \nabla_x H(M|\hat{\mathbf{z}})$ , where  $\nabla_x H$  is the spatial gradient of the entropy. Using the chain rule we rewrite the spatial gradient of  $H$  around a query point  $x$  as:

$$\nabla_x H(x) = \frac{dH}{dp} \nabla_x p(x), \quad (18)$$

where  $p$  is the probability of occupancy at  $x$  given by (1). As  $p$  is a Bernoulli random variable,  $\frac{dH}{dp}$  is simply;  $\frac{dH}{dp} = \log_2 \frac{1-p}{p}$ . The occupancy gradient  $\nabla_x p(x)$  depends on the occupancy map used. For the unperturbed map,  $\nabla_x p(x)$  is given by (10). The spatial gradient of the perturbed map (14) can be computed similarly to the unperturbed map;

$$\begin{aligned} \nabla_x \hat{p}(\mathbf{x}, \hat{\mathbf{z}}_f) &= \nabla_x \sigma(\Psi(\mathbf{x}, \hat{\mathbf{z}}_f)) \\ &= \sigma(\Psi(\mathbf{x}, \hat{\mathbf{z}}_f)) \left(1 - \sigma(\Psi(\mathbf{x}, \hat{\mathbf{z}}_f))\right) \nabla_x \Psi(\mathbf{x}, \hat{\mathbf{z}}_f) \end{aligned} \quad (19)$$

As  $\Psi$  is a GP, its gradient can be computed in closed form.

Fig. 2e shows the computation of MI gradient by sampling along the arc defined by the sensor range. Each sample produces an MI gradient, schematically shown by the black arrows. The overall MI gradient, which pushes optimisation toward

exploratory trajectories is computed from the sum of all samples and is shown in Fig. 2e as a yellow arrow.

### 4.3 Functional Exploration Algorithm

In this section, we describe the functional exploration algorithm, which aims to find a safe path that maximises MI over its entire course.

Functional exploration is a general optimisation framework, meaning it is invariant to the choice of path representation. Functional optimisation methods have used waypoint parameterisation [27], Gaussian process representation [17, 2], or defined trajectories as RKHS [16]. However, in all methods the objective is sampled via a deterministic schedule. Such approaches have proved unsatisfactory for planning using occupancy maps [6]. To ensure convergence to a safe solution, the path is stochastically sampled in the entire  $t \in [0, 1]$  domain and any uninformative gradient updates, such as those coming from unsafe parts of the map, are rejected.

The path representation used in this work is based on kernel matrix approximations [7]. The path is essentially a weighted sum of nonlinear features, similar to a nonlinear regression. This path model is highly expressive, yet concise. Most importantly, the path can be optimised by SGD [20]. Given a set of weights  $\mathbf{w}^\xi$ , the path is defined as:

$$\xi(t) = (\mathbf{w}^\xi)^T \hat{\mathbf{Y}}(t) + \xi_o(t) + \xi_b(t). \quad (20)$$

$\xi_o$  is an initial path, which can be randomly chosen or can be computed by a simple and fast planner.  $\xi_b$  is a term used to adjust boundary conditions, such as the start and goal pose.  $\hat{\mathbf{Y}}(\cdot)$  are nonlinear features that approximate the inner product defined by a kernel function  $k_p$  by the following dot product [7];

$$k_p(t, t') = \langle \mathbf{Y}(t), \mathbf{Y}(t') \rangle \approx \hat{\mathbf{Y}}(t)^T \hat{\mathbf{Y}}(t'). \quad (21)$$

The kernel  $k_p(t, t')$  maintains the correlation between various time points.

Using (20), the general functional optimisation given in (2) is transformed into an optimisation of the weights,  $\mathbf{w}_{optimal}^\xi = \arg \min_{\mathbf{w}^\xi} U(\xi)$ . The iterative update rule of (5) takes the following form:

$$\mathbf{w}_{n+1}^\xi = \mathbf{w}_n^\xi - \eta_n \Lambda^{-1} \hat{\mathbf{Y}}(t_i)^T \hat{\mathbf{Y}}(t_i) \nabla_{\xi} U(\xi_n)(t_i), \quad (22)$$

where  $t_i \in \mathcal{T}$ .

The algorithm for functional exploration is given in Algorithm 1. The essential inputs are the initial robot state and the occupancy map. Boundary conditions or an initial solution are optional inputs. In each iteration, a mini-batch is drawn. The safety of each sample is checked, and if it is below  $P_{safe}$ , the sample is used to update the path. The gradient of the various components (obs, dyn, MI) of the objective functional are computed, and summed according to (8). Once the overall functional gradient is computed, the weights  $\mathbf{w}_{n+1}^\xi$  are updated according to (22).

Figure 4 provides an insight into the optimisation process of a single planning iteration. Each iteration starts with an initial guess, which is depicted in blue. The

**Algorithm 1:** Stochastic Functional Exploration

---


**Input:**  $M$ : Occupancy Map.  
 $\xi(0)$ : Start state.  
 $P_{safe}$ : No obstacle threshold.  
Optional: boundary conditions  $\xi_b$ , initial guess  $\xi_o$ .

**Output:**  $w_{optimal}$


$n = 0$   
**while** *solution not converged* **do**  
  //Stochastic sampling:  
   $S \sim U[0, 1]$  Draw mini-batch  
  **foreach**  $x_t \in S$  **do**  
     $P_{occ} \leftarrow M(x_t)$  Eq. (1)  
    **if**  $P_{occ} \leq P_{safe}$  **then**  
       $\hat{z}_f(x_t) \leftarrow$  simulated observations  
       $\nabla U_{obs} \leftarrow \frac{\partial}{\partial \xi} x_t \left( \sigma(w^T \hat{\Phi}(x_t)) \left( 1 - \sigma(w^T \hat{\Phi}(x_t)) \right) w^T \nabla_x \hat{\Phi}(x_t) \right)$  Eq.(9)  
       $\nabla U_{dyn} \leftarrow -\frac{d^2}{dt^2} \xi(t)$  Eq.(11)  
       $\nabla U_{MI} \leftarrow \sum_{m \in m_{max}(x_t)} \frac{\partial}{\partial \xi} x_t \nabla_x MI(M(m); \hat{z}_f(x_t))$  Eq.(17)  
       $\nabla U \leftarrow$  Combine using Eq.(8)  
       $w_{n+1} \leftarrow w_n^\xi - \eta_n \Lambda^{-1} \hat{Y}(t_i)^T \hat{Y}(t_i) \nabla_\xi U(\xi_n)(t_i)$  update rule Eq. 22  
    **end**  
  **end**  
   $n = n + 1$   
**end**

---

obstacle functional gradient repels the path from obstacles and unknown space. The MI objective pulls the path toward unexplored space. The intermediate path updates, following the functional gradient, are shown in grey. The final optimal path is shown in green.



**Fig. 4** Functional planning iteration. Functional gradients deform the initial solution in blue. The obstacles gradient repels path from obstacles and unknown space. The MI gradient pulls path toward unexplored space. The resulting path, in green, is the optimised solution.



**Fig. 5** Comparison of exploration methods - 2 repetitions; the proposed functional exploration method (red), RRT-based exploration (black) and frontier-based exploration (green). RRT and the propose method converge in a similar fashion as both optimise MI over entire path. The goal-based approach of frontier converges slower as it does not explicitly optimise on MI.

## 5 Experimental Results

In this section, we evaluate the performance of the stochastic functional gradient path planner and compare it to other exploration methods for continuous occupancy maps. As a benchmark, we chose to compare our method to the RRT-based exploration method of Yang et al. which we modified to use Hilbert maps. This method optimises MI during path selection, and while its bottleneck lies in computational complexity, it takes advantage of the fact the RRTs are probabilistically complete. Another method used for comparison is based on frontier exploration [25], which takes a grid approximation of the continuous map in a similar approach to [10]. The path is then constructed by a smooth RRT planner which reasons only about the path safety. All methods, including Hilbert maps, are implemented in Python and tested on an Intel i5-6200U with 8GB RAM.


Figure 6 shows a qualitative comparison at various planning iterations. While all exploration methods successfully build the map, there are some clear differences. The frontier-based method produces jerky paths and tends to move towards the edges and corners of the map. Path optimisation methods, on the other hand, stay closer to obstacles. As the RRT-based method maximises MI only, its path moves closer to the edges of obstacles, while the proposed method keeps a bigger distance, as its objective includes obstacle safety explicitly. We note that increasing the safety margin by applying a blurring filter, as done with grid maps, is not applicable in continuous maps, as it requires expensive discretisation of the map.

Quantitative comparisons are shown in Fig. 5 and in Table 1. Fig. 5 depicts the reduction in the entropy of the map. The rate of reduction is similar for both our method and the RRT-based exploration, albeit slightly better for the latter, mainly due the fact that the paths generated by the RRT move closer to obstacles. As a result, the RRT planner covers the map faster, however with a higher probability of collision, as shown in Table 1 by the maximum occupancy values exceeding the 50% occupancy threshold. The convergence of the frontier planner is slower, as a result of the over-estimation of the path utility by the choice of a single goal at each iteration. In addition, paths are jerky, leading to longer time to cover the same area. As the frontier planner does not explicitly minimise collision risk, the maximum occupancy over the path is high. In contrast, the maximum occupancy of the proposed method along the path is significantly below the 50% occupied threshold.

Comparing the median runtime results shown in Table 1 reveals that the RRT planner is significantly slower than our method, mainly because MI is computed

**Table 1** Performance comparison (40 planning iterations)

	Proposed method	RRT [26]	Frontier[25]
Mean occupancy	1.2	5.1	3.0
Max. occupancy	26.3	75.2	47.9
Median Iter. Plan Time [s]	65.4	200.6	13.4
Mean Iter. Plan Time [s]	73.7	196.6	84.3
Max. Iter. Plan Time [s]	174.4	239.6	1345.2



**Fig. 6** Comparison of exploration methods at various planning iterations using continuous occupancy maps. The hexagonal markers are the planning poses. Triangle and star represent start and end points, respectively. The grey lines are either candidate paths (RRT) or intermediate solutions (functional exploration). The pink line depicts the optimised path and the green line is the traversed path. The path is assessed during execution, and a re-plan step is invoked if a path is no longer safe.

over the entire map. Since frontier only queries the map for occupancy, its runtime is significantly smaller. However, its average runtime is similar to our method as occasionally resolving longer path is required. It is worth noting that the runtime results are limited by the Python implementation of the Hilbert map. C++ implementation of the Hilbert map proved to be two to three orders of magnitude faster, which makes it suitable for online applications.

## 6 Conclusions

This paper introduces a novel method for exploration over continuous occupancy maps using stochastic functional gradient descent. This approach formalises exploration as a variational problem, where optimisation is performed directly in the space of trajectories. The functional objective of the proposed method explicitly optimises both safety and information collection over the entire path, finding the *Next Best Path*. While this approach can be used with any type of occupancy map, it is highly effective with Hilbert maps, where the introduced MI objective and its gradient can be computed from a perturbed model of the map. Our proposed approach eliminates the need for computing MI over the entire map as done in other exploration tech-

niques. Rather, it computes variations to the path based on functional gradient of MI which are efficiently derived in closed-form from the map model.

Comparisons with other exploration methods show that the proposed method improves on both safety and MI. Point exploration methods, such as frontier, which do not optimise the path selection, exhibit slower exploration rate. On the other hand, sampling-based exploration methods, such as [26], do not include safety in their objective, hence the resulting paths tends to move closer to obstacles. Moreover, these methods are computationally expensive due to the need to repeatedly sample the MI objective over the entire path. In comparison our proposed method achieves similar exploration rates to [26] while taking less time to compute and still maximising safety.

## References

1. Charrow, B., Kahn, G., Patil, S., Liu, S., Goldberg, K., Abbeel, P., Michael, N., and Kumar, V. Information-Theoretic Planning with Trajectory Optimization for Dense 3D Mapping. In *Proceeding of Robotics: Science and Systems*, 2015.
2. Dong, J., Mukadam, M., Dellaert, F., and Boots, B. Motion Planning as Probabilistic Inference using Gaussian Processes and Factor Graphs. In *Proceeding of Robotics: Science and Systems*, 2016.
3. Elfes, A. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, 1989.
4. Elfes, A. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In *Reasoning with Uncertainty in Robotics*, 1996.
5. Francis, G., Ott, L., Marchant, R., and Ramos, F. Occupancy Map Building through Bayesian Exploration. *preprint arXiv:1703.00227*, 2017.
6. Francis, G., Ott, L., and Ramos, F. Stochastic Functional Gradient for Motion Planning in Continuous Occupancy Maps. In *Proceeding of the IEEE International Conference on Robotics and Automation*, 2017.
7. Francis, G., Ott, L., and Ramos, F. Stochastic Functional Gradient Path Planning in Occupancy Maps. *preprint arXiv:1705.05987*, may 2017.
8. González-Baños, H. H. and Latombe, J.-C. Navigation Strategies for Exploring Indoor Environments. *The International Journal of Robotics Research*, 21(10-11):829–848, 2002.
9. Holz, D., Basilico, N., Amigoni, F., and Behnke, S. A Comparative Evaluation of Exploration Strategies and Heuristics to Improve Them. In *Proceedings of the European Conference on Mobile Robots*, 2011.
10. Jadidi, M. G., Miro, J. V., Valencia, R., and Andrade-Cetto, J. Exploration on continuous Gaussian process frontier maps. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2014.
11. Jadidi, M. G., Miro, J. V., and Dissanayake, G. Mutual information-based exploration on continuous occupancy maps. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.

12. Juliá, M., Gil, A., and Reinoso, O. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Autonomous Robots*, 33(4):427–444, 2012.
13. Julian, B. J., Karaman, S., and Rus, D. On mutual information-based control of range sensing robots for mapping applications. *The International Journal of Robotics Research*, 33(10):1375–1392, 2014.
14. Lauri, M. and Ritala, R. Planning for robotic exploration based on forward simulation. *preprint arXiv:1502.02474*, 2015.
15. Marchant, R. and Ramos, F. Bayesian Optimisation for informative continuous path planning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2014.
16. Marinho, Z., Boots, B., Dragan, A., Byravan, A., Gordon, G. J., and Srinivasa, S. Functional Gradient Motion Planning in Reproducing Kernel Hilbert Spaces. In *Proceedings of Robotics: Science and Systems*, 2016.
17. Mukadam, M., Yan, X., and Boots, B. Gaussian process motion planning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2016.
18. O’Callaghan, S. T. and Ramos, F. T. Gaussian process occupancy maps. *The International Journal of Robotics Research*, 31(1):42–62, 2012.
19. Ramos, F. and Ott, L. Hilbert maps: scalable continuous occupancy mapping with stochastic gradient descent. In *Proceedings of Robotics: Science and Systems*, 2015.
20. Robbins, H. and Monro, S. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400–407, 1951.
21. Stachniss, C. *Robotic Mapping and Exploration*. Springer, 2009.
22. Thrun, S., Burgard, W., and Fox, D. *Probabilistic Robotics*. MIT Press, 2005.
23. Vallvé, J. and Andrade-Cetto, J. Potential information fields for mobile robot exploration. *Robotics and Autonomous Systems*, 69:68–79, 2015.
24. Whaite, P. and Ferrie, F. P. Autonomous exploration: driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.
25. Yamauchi, B. A frontier-based approach for autonomous exploration. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 1997.
26. Yang, K., Keat Gan, S., and Sukkarieh, S. A Gaussian process-based RRT planner for the exploration of an unknown and cluttered environment with a UAV. *Advanced Robotics*, 27(6):431–443, 2013.
27. Zucker, M., Ratliff, N., Dragan, A. D., Pivtoraiko, M., Klingensmith, M., Dellin, C. M., Bagnell, J. A., and Srinivasa, S. S. CHOMP: Covariant Hamiltonian optimization for motion planning. *The International Journal of Robotics Research*, 32(9-10):1164–1193, 2013.